How to Increase the Chances Your Paper is Accepted at ACM SIGCOMM

by Craig Partridge
BBN Technologies

This note is some informal and personal advice about ways authors can increase the chance that a paper they submit to ACM SIGCOMM will be accepted. My dual purpose in writing this note is to help authors submit better papers and help the SIGCOMM conference, by improving the quality of papers it receives.

My dubious qualifications to write this note are that I've served on the SIGCOMM Program Committee every year since 1989 and was Co-Program Chair in 1994 and that I've had two papers accepted, and several papers rejected, by ACM SIGCOMM.

I. THE PAPER EVALUATION PROCESS

It is useful to start by describing what happens to a paper after it is submitted to ACM SIGCOMM.

It is important to keep in mind that SIGCOMM runs a double-blind reviewing process. By double-blind we mean that authors don't know who their reviewers are, and reviewers don't know who the authors are. Personally, I'm a big fan of double-blind reviewing. Double-blind reviewing focuses attention on the papers, rather than who wrote the papers. And every year there are some surprises — SIGCOMM accepts an outstanding paper from previously unknown authors.

The first step in the evaluation process is that the program chairs circulate the title and abstract of each paper, and ask the program committee members to indicate which papers they are qualified to review. Based on the feedback from the program committee members, the program chairs then assign the papers to committee members for review. Observe that this process usually ensures that a paper is reviewed by someone who feels well-qualified in the paper's field.

In some years, each paper has initially been given to two reviewers, in other years each paper has been given to three reviewers. The program committee members are expected to read all their assigned papers themselves (though they may also solicit additional reviews from qualified colleagues).

Each reviewer rates the paper on a 1 to 5 scale, with a rating of 3 or above indicating the paper is acceptable. Each reviewer also writes a text review indicating what she or he sees as the strengths and weaknesses of the paper.

The program chairs collate all the ratings about two weeks before the program committee meeting. At this time, papers that have a wide range of ratings (such as a 1 and 5 from two reviewers) will be assigned to additional program committee members for review before the meeting. In years that only two reviewers initially review a paper, all papers that are candidates for acceptance also get a third review during this period.

The program committee meeting's purpose is to pick the best 24 to 30 papers for SIGCOMM from the 250-odd papers that are submitted. Inevitably there's a bit of randomness in the process (how a paper fares depends in part on which reviewers it is assigned to). But the process is intended to try to minimize the randomness. In particular, it is generally the case that every paper that received at least one rating of 3 or higher is discussed by the program committee. So the committee as a group is looking at a fairly large chunk of the papers (typically about half of them) to find the 10% of the papers it wants to accept.

The program committee meeting is an intense process, lasting all day and often into the night. Controversial papers are often given to additional people to review during the course of the meeting, and while the average paper's fate is decided in a minute or two, the last few decisions often take 20 or 30 minutes each.

It is also worth noting that program committees typically do not try to shape the program. That is, there's no attempt to ensure that a hot topic has papers accepted, or that there's a balance of theory and systems papers. Rather the goal is to take the best SIGCOMM-related papers that can be found, regardless of topic. Indeed, the program often accepts papers on topics known to be obscure or appealing to a small community on the grounds that the technical work is outstanding and deserves circulation to the broader community.

II. GENERAL ADVICE

This section gives general advice about how to improve your submission to SIGCOMM. The next section gives detailed advice about particular topics and types of papers.

II.1 Start Writing Early

One cannot stress enough the importance of making sure your paper is well-written. Because SIGCOMM is double-blind, the reviewers often have no idea who wrote the paper. So they cannot make allowances for your reputation and, for instance, say to themselves "Joe can easily fix this paper up." If the writing is sloppy or the presentation is bad, the paper probably will not be accepted. SIGCOMM does accept two or three papers each year on the condition that a program committee member oversees the revision of the paper. But these revisions are typically for technical content, not presentation.

I often tell prospective authors that they should aim to have a first draft of their paper done three to four weeks before the SIGCOMM submission deadline, solicit a few reviews from colleagues and friends, and then revise the paper before submitting it to the conference. This advice is doubly strong for authors for whom English is not their native language.

II.2 Properly Cite and Summarize Prior Work

A surprising number of papers fail to cite relevant prior work, or worse disparage or misrepresent the prior work to make the paper's contribution look better. Reviewers who are trying to assess a paper's contribution tend to get annoyed when that contribution is exaggerated or, worse, is a reinvention of prior work.

A corollary to this comment, is that the authors should make the readers aware of the limitations of the work in the paper. Doing so typically strengthens rather than weakens the paper. The reviewers generally are quite good at assessing the work's limitations, and if these points are not conceded/discussed, then the reviewers become concerned that (1) the authors fail to see the big picture, (2) claims made in the paper may not have been assessed objectively, and (3) if the paper is accepted, others reading it will also miss these limitations. Because the program committee usually does not have any way to ensure that the authors will revise their paper to fairly assess the shortcomings of the work, they will typically err on the side of caution and reject submissions that fail to do so.

II.3 Keep Within the Page Limits

Reviewers have to do a lot of reading. Overlength papers do not win friends.

II.4 But Be Complete

Sometimes authors, struggling to stay within page limits, leave out crucial details. A classic example is the paper that omits some proofs for brevity.

Don't leave out a critical detail. Find a way to fit it into the paper.

If there's a useful proof that won't fit in the main paper, or the paper relies on a prior work of yours that you cannot cite due to anonymity, consider attaching the proof or relevant piece of the (anonymized) paper in an appendix and referencing the appendix in the main body of the paper. This approach has the double advantage of keeping the reviewer informed and showing that the author knows how to edit his or her work down to size. A short appendix of this sort will not count against the page limits.

II.5 Write a Good Abstract

Reviewers choose what papers to read based on the abstract. So make sure that the abstract indicates what type of paper the reviewer may expect. Each year there's usually a paper whose abstract makes it look like a systems paper, but is actually a theory paper — causing systems people to work through the math. To their credit, the systems folks usually work through the math carefully (and vice-versa, the theory folks will read a systems paper with care) but it still means a review with less confidence behind it. Reviews with less confidence are a concern because, to be accepted, a paper often needs someone on the program committee to argue strongly for acceptance. If none of the paper's reviewers are confident of their reviews, they are less likely to argue strenuously in favor.

One long-time program committee member suggests writing your abstract to be of interest to the committee members you believe are best qualified to review your paper.

II.6 Avoid Needless Buzzwords

Every year seems to have a new buzzword or hot term. Sometimes authors think that having the hot buzzwords in their paper increases the chance of acceptance. Generally buzzwords don't help; and sometimes hurt. At more than one SIGCOMM PC meeting, certain buzzwords have rapidly become swear words in response to their overuse.

III. SPECIFIC ADVICE

SIGCOMM deals with certain types of papers very frequently. In this section, I give some guidance to authors of certain popular types of papers.

III.1 TCP Performance Papers

TCP performance is a well-trod ground and so the standards for getting a TCP paper accepted are now quite high.

To be accepted at SIGCOMM, a TCP performance paper should demonstrate that the proposed performance improvements have been thoroughly tested. For instance, any changes to TCP flow control should be tested over heavily loaded multi-hop topologies with cross traffic. Furthermore, the analysis should show not only that the enhanced TCP performs better, but also show the effects of the enhanced TCP on non-enhanced traffic. Note that TCP performance papers are often measurement papers and so the discussion of measurement papers in the next section applies.

III.2 Measurement Papers

Writing good measurement papers is very hard. It requires careful network monitoring, using good statistical techniques. A good monitoring paper should explain how the data was taken, why the data is believable (i.e., what statistical measures were taken to ensure the data was sound), and then analyze the data carefully with good charts and graphs and discussion that indicates the data was thoroughly analyzed and contemplated. Probably more than any other type of paper, measurement papers benefit from extra time for the author to refine the paper. Start writing early!

III.3 Systems Papers

SIGCOMM very much wants to publish systems papers. At the same time, SIGCOMM is consciously struggling with some difficulties in handling systems papers. (These problems are not unique to SIGCOMM; several journals also seem to have similar challenges). This subsection is a guide to the problems and pitfalls that may trap an unwary author.

The classic systems paper presents an implementation or planned implementation. The implementation can be in software or hardware or both. The implementation's contribution is usually that it either achieves some new function, never before achieved, or it realizes an existing function more efficiently or effectively than previously (and not just as a result of applying Moore's Law!).

It is often hard to describe an entire system completely in ten single-spaced pages. Make a strong effort to be as complete as possible. A number of systems papers get rejected because reviewers feel key details of the system are missing, details that in most (but not all) cases the authors could have provided.

Another way to express this idea is that small systems papers often fare better than large systems paper. By small systems paper, I mean a paper that tackles a modest problem and usually has a contribution that can be described in one sentence. A big systems paper is trying to accomplish a larger set of goals, and often has three or four important contributions. It is much harder to completely describe a large system.

Make sure the relevance to networking is clear. For instance, in a paper describing a wonderful new protocol verification language, it is very easy to get wrapped up in details of language syntax. Remember that SIGCOMM is a networking conference, not a language conference, and ensure that how the language improves the state of networking is made very clear.

SIGCOMM program committees have varied from year to year in how willing they are to accept preliminary papers (papers on systems still being built vs. papers on systems that have been completed). There's a tension between presenting interesting new ideas to the community as soon as possible and the danger of presenting ideas that will prove duds after a bit more testing. There is no consensus on this issue and authors considering submitting work on incomplete systems are advised to talk with the program chairs to learn the current year's leanings.

Be extra careful to cite relevant work. A common mistake in systems papers is to start with a blank sheet when designing a system. There's lots of prior experience about how to solve a wide range of systems' problems. Make sure the paper describes how it takes us to a new part of the solution space, rather than just blindly repeating prior work.

III.4 Modelling and Queuing Theory Papers

SIGCOMM believes it is very open-minded about modelling and queuing theory papers, provided their relevance to networking is clear. The first self-similarity paper was published at SIGCOMM, when other conferences found it too startling to publish. Furthermore, SIGCOMM is often willing to publish innovative work that tries to tackle hard problems, even if the resulting solution is incomplete. That said, there are three problems that frequently afflict modelling and queuing theory papers.

First, the problem being modelled or analyzed is often so abstracted from the reality of how networks operate that the results have little or no real relevance to networking. As one expert put it, too many papers look for a way to redefine a networking problem into a solvable math problem rather than actually solving the networking problem. A lesser version of this behavior is that papers often present the problem so formally that they never state how they actually contribute to networking. Don't leave the application of your work to the imagination of a tired reviewer who has already read ten papers this week.

The second problem, which is related to the first problem, is that some theory papers are really mathematics papers, masquerading as networking papers. A classic sign of such as paper is an introduction, explaining that the paper's topic has relevance to some problem in networking and then no further mention of networks until the conclusion.

Third, the traffic models are often unrealistic. It is now widely recognized that network traffic is almost never i.i.d. Yet astonishing number of papers still use i.i.d. models to analyze performance (esp. for switch performance). Unless the result is a negative one (i.e., this switch allocation scheme doesn't work, even for i.i.d traffic), or this work is the first time anyone has even attempted to analyze a very hard problem, use a realistic traffic model.

III.5 Architectural Papers

SIGCOMM occasionally publishes architectural papers; papers that attempt to get us to think about networking in a new way. Examples of such papers are the 1990 ALF paper and the 1992 paper on Integrated

Services.

The best papers present both new architectural thinking and an example of how the new way of thinking enables us to do new things, or radically reconsider existing work. For example, the ALF paper had performance numbers for combining multi-layer functions in an application.

IV. Final Comments

ACM SIGCOMM now accepts about 1 paper in 10, making it three times more competitive than the average ACM or IEEE journal. Furthermore, unlike the journal process, where authors get a chance to improve their papers and have them re-reviewed, the SIGCOMM process has just one accept/reject cycle. So it is very important that a submitted paper be the best possible paper it can be.

Finally, there's no stigma to having a SIGCOMM paper rejected. Some of my best journal papers were first rejected by ACM SIGCOMM (usually because I didn't have enough time to fix one of the problems noted above). If you've got a good idea, please send it in. The purpose of this note is not to cause you to avoid submitting a paper, but rather to help you make your paper the best it can be. Good luck!