Improving Sigcomm: A few straw proposals

Abstract

This report discusses several proposals for improving the Sigcomm conference. It was prepared by Sigcomm's Technical Advisory Committee and approved by the Sigcomm Executive Committee.[1]

1 Introduction

The Sigcomm[2] conference is widely seen as the foremost academic conference in the area of computer networking and Internet architecture. The quality and depth of the research performed by the Sigcomm community is quite impressive. However, as in all things, striving for improvement is the only way to maintain excellence. The task before us is to explore ways in which the annual Sigcomm conference could be improved.

Perhaps the most pervasive complaint about Sigcomm is that it could do a better job of selecting papers for acceptance. Dissatisfaction with the paper selection process is the subject of much conversation in the Sigcomm community. No conference with a high rejection rate will ever be free of complaints, but Sigcomm can certainly improve in this regard. We address this issue in Section 2.

While the selection of papers attracts much of our attention, we should keep in mind that there are larger issues at stake. Sigcomm, both as a conference and as an organization, has a chance to influence the research agenda of the wider community and, in turn, the evolution of the Internet. We believe that Sigcomm can play a more active role in fostering wide-ranging discussion and debate of various topics. This would benefit the research community in general, and improve the quality of submissions to future Sigcomms. Section 3 presents a few proposals in this direction.

Sigcomm traditionally allows PC chairs wide latitude in carrying out their task, and this document is not intended to overly constrain their choices. The philosophy behind the proposals is more fundamental than how they are carried out, and most of the proposals here are fairly general in nature and leave substantial room for different implementations. While we strongly believe that PC chairs should continue to have substantial discretion, we also strongly believe that this discretion should be exercised in consultation with the Sigcomm Technical Advisory Committee (TAC). The TAC can provide a degree of continuity and long-term perspective that cannot be supplied by a progression of extremely talented and dedicated but inherently ephemeral PC chairs. We discuss the role of the TAC in Section 4.

We make these proposals with the knowledge that any (or all) of them may fail to achieve their desired aim and that some of them are controversial. We encourage Sigcomm to consider the proposals here to be *work-in-progress* and continue monitoring their effectiveness. We hope that Sigcomm as an organization achieves consensus on the changes that appear to be working and abandons those that appear to have failed. Some of the more debatable proposals we present here are explicitly labeled as experiments; we recognize that they are gambles, and hope that the experience gained from these experiments will help guide the TAC and future PC chairs.

[1] The writing and much of the work on this document was done by a subcommittee of the TAC consisting of Tom Anderson, Jim Kurose, and Scott Shenker.

[2] In what follows, the term Sigcomm will refer to the Sigcomm conference. The term Sigcomm organization will be used to refer to the Sigcomm SIG.

2 Selection of Papers

Sigcomm is an extremely selective conference, with acceptance rates that have consistently been less than 15% in the past five years. The selection of papers plays a crucial role in determining the quality of the Sigcomm program. In this section we make several proposals for improving the Sigcomm selection process.

2.1 General Philosophy

We propose that Sigcomm adopt the following philosophy:

> Sigcomm looks favorably on papers that open up new areas, present new ideas, and/or serve as a foundation for new work. Such papers are often unable to present complete and convincing cases for their ideas and, because the territory they cover is relatively new, the execution of the paper may be somewhat less than ideal. When the two are in conflict, Sigcomm values interesting and novel ideas and results over complete and flawless execution.

This philosophy involves taking risks; we recognize that the Program Committee will occasionally accept papers that are later found to be uninteresting or technically flawed. Conferences are *caveat emptor* and it should be generally understood that acceptance at Sigcomm (or indeed any conference) does not imply an iron-clad guarantee of the correctness of the results or of the wisdom of the design. Instead, the PC is choosing papers that it feels would be valuable for the community to read, discuss, and debate; all is not lost if, at times, the outcome of those discussions is the refutation of the original paper.[3]

None of the above should be taken to mean that execution is unimportant or irrelevant. Flaws in execution should *always* be taken seriously. However, the severity of the flaws should be weighed against the novelty or depth of the ideas; when the ideas are interesting, flaws in execution should not always be considered fatal.[4]

Balancing the tradeoff between innovation and execution involves a degree of judgment that cannot be captured in explicit guidelines. However, one can structure the way in which papers are considered to reinforce the philosophy being proposed here. For instance, one way to embody this general philosophy is to ask each reviewer to answer three simple questions:

1. **Does this paper address an interesting issue?** To what extent is the *topic* of the paper important and interesting? If the issue or problem addressed were (or now is) completely understood or solved, how important would that be, in terms of either fundamental concepts, or increased understanding, or practical relevance?

2. **Does this paper present interesting results?** Do the results provide worthwhile insight into the topic addressed? Are the results likely to be widely used by others? Does the work open up new areas, present new ideas, and/or serve as a foundation for new work?

3. **Is this paper sufficiently well executed?** Are there flaws (*e.g.*, technical mistakes, important uncited related work, poor assumptions, insufficient scope of evaluation, unsubstantiated conclusions, poor writing) in the paper? Are the flaws fundamental or superficial? That is, are the results likely to be true despite the flaws, or do the flaws fundamentally impact the results in the paper?

[3] This point raises the issue of to what extent Sigcomm should encourage papers, or other presentations, *debunking* previous papers. That is, if we are going to take risks, we should also provide a mechanism whereby corrective contributions can also be published. This will be touched upon briefly in Section 3 but the general point probably deserves more consideration than presented here.

[4] For instance, papers that present superficially interesting but deeply flawed ideas or whose execution is so flawed that the ideas are essentially unsupported should be rejected. But papers whose ideas are worth considering, even if the case is not completely persuasive, should be given the benefit of the doubt.

The ordering of the questions is crucial: the quality of the ideas and the results is of more importance than flawlessness of execution. Execution should be evaluated by whether it was *sufficient* to support the ideas, not whether it was flawless.

In these questions, the term *interesting* should be interpreted broadly to include all papers that will (or should) be widely read and used by the research community. There is no single metric of being interesting; among many other reasons, papers can be worth reading because they add to the conceptual foundations independent of any practical application, or because they propose a novel idea for improving the operation of the Internet, or because they present measurements that deepen our understanding or will foster research by others in the field.

2.2 Program Committee

Composition: With very few exceptions, all Program Committee (PC) members should have three qualifications: (i) demonstrated ability to write high quality research papers (not necessarily at Sigcomm), (ii) judgment (particularly the ability to find what is good in a paper, not only what is bad), and (iii) sufficient time to devote to the process. Within these constraints, the PC should be diverse in terms of age and experience. As we state later, we encourage PC chairs to consult the TAC and other sources prior to selecting the PC committee, as research reputation does not always correlate with quality performance as a PC member.

In the past five years there has been an explicit goal of including a high fraction (roughly 33%) of new PC members (researchers who have never served on the PC) each year. This practice was useful in bringing diversity to the Sigcomm PC and, to a large extent, past concerns about "cliquishness" of the Sigcomm conference seem to have been allayed. There is now a large pool of researchers who have served on the PC and there should no longer be a quota on the number of new PC members each year. Moreover, while geographic diversity may be a desirable goal it is of lower priority than technical excellence and balance.

Size *(experiment)*: The size of the PC is a difficult issue. There is a tension between keeping the work load manageable (which calls for a larger PC) and achieving more coherent acceptance decisions (which calls for a smaller PC). This tradeoff is a difficult one. In recent years Sigcomm has tended towards larger PCs. To gain experience with the other end of the spectrum (and to test the hypothesis that a smaller PC would produce more coherent acceptance decisions) we propose that over the next few years Sigcomm experiment with smaller PCs – certainly no larger than 25 but perhaps smaller than that. The TAC should revisit this issue every year until it appears that Sigcomm has found a long-term solution.

Smaller PCs will require changes in past Sigcomm practice to keep the workload manageable. This will include making more use of outside reviews than in the past few years, and modifying the review procedures.

2.3 PC Process Issues

We have all been on PC's that, despite everyone's best intentions, have yielded less than the best possible result. We next discuss several process issues that will help PCs choose the most interesting papers. Of course, there are a myriad of detailed PC procedures that we do not address here.

Initial Organizational Communication: Program Chairs should not assume that everyone on the PC shares the same perspective on the tradeoff between execution and innovation. The best time to increase agreement on this balance is *before* paper reviewing begins, rather than afterwards; in the past, these issues have often been addressed only at the PC meeting itself, by which time reviewers have hardened opinions on specific papers and the contentious atmosphere and time pressure makes agreement difficult.

We believe it would be useful to have, at the very beginning of the reviewing process, a PC teleconference where such issues are discussed. A teleconference is strongly preferred to an email outlining such issues, as dialog and

a common understanding among the PC members is important. Past experience shows that the busy lives of PC members results in email being read at different times, and even a multiple-email-exchange discussion invariably only actively engages a small subset of the PC. But whatever the form of this discussion, it must occur at the beginning of the reviewing process not at the end.

Technical Reviews: All papers under serious consideration should be subject to some number of *technical* reviews. These are in-depth reviews in which the reviewers pay careful attention not only to the ideas in the paper, but also to the technical details. Papers in the upper tier (perhaps 2-3 times as many as there is room to accept, or roughly a third of all submissions given recent patterns) should be subject to additional technical reviews.

Outside Reviews: We propose that the use of outside reviews be encouraged. While there are often problems calibrating these outside reviews, outside reviews can play an invaluable role in providing coverage where the PC is weak and in ensuring that a wide variety of perspectives are incorporated in the reviewing process. In addition, use of external reviews can significantly reduce PC workload. Both of these advantages of outside reviews become increasingly important as we experiment with smaller PCs. External reviewers should be aware of the evaluation criteria and paper-acceptance-philosophy of the conference (see Section 2.1).

The Sigcomm Technical Advisory Committee (TAC), as we comment below, can help PC chairs find good outside reviewers.

Light Reviews: We believe it is extremely valuable for as much of the PC as possible (at least a half, preferably more) to have read each paper being considered at the PC meeting. A major complaint has been inconsistency in quality across the program; having more of the PC read each paper can help reduce unevenness and chance in the reviewing process. These readings have a distinctly different purpose than the technical reviews – their goal is to understand the basic ideas presented in the paper and to understand the comments of the technical reviewers. These *light* reviews should be able to answer the questions of whether the paper addressed an interesting issue and whether the results are of interest, but may miss flaws in execution. Whether or not these light reviewers file a report (or submit a score) is open to debate; the main point of these light reviews is to make sure the PC discussion has a high number of informed discussants.

Pre-meeting Discussions: In the past five years, the set of PC members who have reviewed a particular paper have held online discussions. There has been an increasing trend towards reaching an acceptance/rejection decision in these small groups, even though this has *not* been a stated goal of these discussions in the past. In some cases these early decisions led to worthy papers not even being considered at the PC meeting.

While these pre-meeting discussions appear necessary, we caution that the goals of these discussions should be clear. We propose that these discussions should be geared towards two goals:

1. *Understand the opinion of other reviewers.*

   The reviewers need not come to agreement, but they should understand the differences in opinions between the various reviews.

2. *Calling for additional technical reviews if required.*

   Additional technical reviews will be needed if, for example, there are technical questions that need to be addressed by additional referees, or if the paper touches on a topic outside of the reviewers' expertise, or if there is substantial technical disagreement among the reviewers about the correctness of the paper.

To re-emphasize an important point, it is explicitly *not* a goal of these pre-meeting discussions to reach consensus (on anything!). In particular, the reviewers are not being asked to reach consensus on the value of the paper or on accept/reject decisions.

Papers Considered at the PC Meeting: Because of time constraints, only a small subset of the submitted papers can be considered in any depth at the PC meeting. In general the number of considered papers should be roughly 2-3 times the number of expected acceptances. An aggressive triage process is necessary to focus the PCs attention on this top tier of papers. We propose that the decision of whether or not to consider a paper at the PC meeting be based on a mechanical criterion involving the reviewer scores (*e.g.*, any paper rated by any reviewer as being in the top quartile). In addition, any PC member should be able to nominate a paper to be considered at the PC meeting.[5]

PC Discussions: Typically PC discussions start with the reviewers of the paper stating their views on the paper. If there are conflicts among the reviewers, a discussion among the reviewers ensues.

Our proposal here is that the goal of the discussions should *not* be to resolve the disagreements among the reviewers, but to educate the rest of the PC about the key issues so the PC as a whole can reach an informed decision. This might entail referring to the three questions that embody the philosophy of acceptance (Section 2.1) and appear on the review form.

In addition, we propose that after the main reviewers have spoken that the rest of the PC have a chance to contribute to the discussion. Moreover, when the accept/reject decisions are finally made, they should reflect the sense of the entire PC not just of the main reviewers of the paper. This will become more natural if a large fraction of the PC has read each paper, but even PC members who have not read the paper should be able to contribute their sense of whether, based on what they've heard, they would favor acceptance or rejection.

2.4 Reviewing

General Goals: The primary purpose of the reviewing process is to select a set of papers. The reviewing form and process should be geared towards that end. While it is valuable to provide authors with useful feedback on their papers, that goal becomes secondary when it comes in conflict with either the quality of the paper selection decisions or the overall PC workload.

In particular, we propose that the main goal of a review should be to help other PC members understand the reviewer's answers to the three basic questions relating to the quality of the ideas, the quality of the results, and the sufficiency of the execution. The main audience for the reviewer's comments should be, contrary to present practice, the PC. While it is important that authors see the reasons the PC members gave for recommending for or against acceptance, providing authors detailed feedback on how to improve the paper is at best a secondary goal.

The review form should be modified to reflect this change in emphasis. We leave that for a more detailed discussion below.

In passing we note that long and detailed reviews directed to the authors often tend to focus more on the execution of the paper rather on the quality of the basic ideas. De-emphasizing the length and depth of written reviews, and reminding reviewers that the other PC members are the audience for the reviews, may have an ancillary value of causing reviewers to focus more on the quality of the ideas.

Review Form *(experiment)*: We don't intend to delve into the all the details of the review form here but there are some basic changes that might be helpful, and we discuss a few of them here.

The current review form consists of a single numerical rating plus comments to the author(s) and comments to the PC (not to be seen by the author(s)). Currently, the content of most reviews follows the implicit priorities of the form, and address the author(s), not the PC. The numerical ratings in recent years have become sharply clustered around 3.

We propose that the review form should be changed in three basic ways.

[5] Of course, the PC chairs may have to fine tune the mechanical criterion and the nomination process to ensure that the number of papers considered at the PC meeting is manageable.

- First, the quantitative rating should be a rough percentile ranking. The reviewer's estimation of the percentile ranking is, of course, only a rough approximation but we think using percentiles has an advantage over the current 1-5 ratings. The hope is that the use of percentiles will mitigate the current tendency for ratings to cluster around the average (3.0) since all reviewers, over time, should have ratings appropriately distributed among the percentiles. In addition, this rating system reduces the stigma of the lowest score and the unqualified adoration of the highest score, so they will be applied more often.

- Second, we propose that the two comment sections should be labeled "Public comments to the PC" and "Private comments to the PC" making the point clear that the main audience for the review comments are the PC members.

- Third, the public comments to the PC should be organized about the three questions mentioned in section 2.1.

An example review form embodying these ideas is in Appendix A. This is only a proof-of-concept example; PC chairs are free to implement the review form differently.

2.5 PC Gadflies *(experiment)*

PC chairs are typically swamped with administrative duties and rarely have the time to provide much technical input into the reviewing process. For this reason, we propose that PC chairs appoint one or two *gadflies.* These are PC members whose duty is to lightly read each paper (at least the abstract, introduction, and conclusion) and read the technical reviews. Gadflies can help identify situations where further reviews are needed and cases where the reviews are not consistent with the general philosophy discussed in Section 2.1 (particularly "idea" papers that are being inappropriately hammered for imperfect execution). They can either review these papers themselves or suggest to the PC chairs that other reviewers may be needed. In addition, like any other PC member, they can nominate papers for consideration at the PC meeting. Often the last few frantic moments of the PC meeting are devoted to "resurrecting" papers that were rejected earlier. One can think of gadflies as being ombudspersons who are resurrecting papers *before* the PC meeting, not at the very last minute.

It should be noted that gadflies have no more authority than anyone else on the PC; any PC member can alert the PC chairs that a paper requires additional reviews and any PC member can nominate a paper to be considered at the PC meeting. However, because of their different reviewing load, gadflies bring a special perspective to the process; they will have a broad overview of the submissions rather than a deep understanding of a few papers. We think that having at least one PC member with this broad perspective will make it easier to spot mistakes-in-the-making and will ensure that all interesting papers get full consideration.

2.6 Ongoing Evaluation *(experiment)*

Put in theoretical terms, the selection of papers is a distributed algorithm among energy-constrained nodes with limited communication faced with hard deadlines. Diagnosing any malfunctions in such a distributed process is extremely difficult. We don't pretend to be sure how our proposed changes in the process will effect the final results. For that reason, we recommend that the TAC be charged with the responsibility for ongoing evaluation of the selection process. We suggest that at least one member of the TAC be on the PC each year, and this person should serve as a *monitor* of the process. The monitor should be in conversation with many of the PC members and should take notes during the process noting what seemed to work, what didn't seem to work, and other various aspects of the process. These notes will provide crucial input for the TAC discussions (described below in Section 4). The point is that someone should be designated to collect this input *during the process* rather than after the details have faded from memory.

Note that the monitor is *not* checking up on the PC chairs or on individual PC members. Paper selection is a complicated distributed process in which well-meaning individuals sometimes collectively make imperfect decisions. The goal here is to understand how the collective algorithm worked, and how it might be improved.

The roles of the gadfly and the monitor, while quite distinct, are very synergistic. Thus, they could well be served by the same person.

3 Fostering Better Submissions

Improving the selection of papers is important to the continued health of Sigcomm as a conference. However, Sigcomm as an organization should have a more expansive view of its mission. In particular, the Sigcomm organization should seek to improve research in the community as a whole (which, as a fortunate byproduct, will lead to better submissions). The Internet architecture is largely a product of this research community, and the research community continues to play a special role in the Internet's evolution. The Internet's future vitality depends on the fundamental insight and understanding that can only come from the research community.

While the selection of papers is a clean well-bounded problem, the generation of better submissions is not. The proposals we suggest here address only a very few aspects of the problem, focusing only on areas where we think Sigcomm (the conference) could have some impact. Because the issues here are so amorphous, we start with a short discussion of the problems we are addressing.

3.1 Problems

We address several interconnected problems.

Narrowly Technical: In recent years, Sigcomm has become *narrowly technical* in that most submitted papers provide in-depth technical answers to purely technical problems. Only a very few submissions address general architectural principles or high-level design decisions where judgment and wisdom are more relevant than simulations and measurements. However, these architectural and high-level design issues are often much more critical to the future of networking in general and the Internet in particular than purely technical results.

Encouraging presentation or discussion of material of this nature would enliven the conference itself and improve the submissions for future conferences. Thus, as we outline below, we propose that Sigcomm include such material in its conference program as either papers, invited talks, or panel discussions.

Technically Narrow: In addition to being narrowly technical, many people contend that Sigcomm has become *technically narrow* in that the range of topics being considered is quite constrained. While the selection process may be partially responsible,[6] it is clearly true that the vast majority of submissions focus on a very few topics (*e.g.*, congestion control, QoS, multicast). While these topics are, and will remain, important, they probably receive more than their share of attention because of the self-reinforcing nature of academic disciplines. We assume that there are many other open issues of significant importance that are being neglected.[7]

As we discuss below, we propose that Sigcomm seek out presentations, either invited or on panels, describing important but understudied problems. Soliciting input from the real world is critical here, to prevent Sigcomm from becoming detached from reality.

Having a "home" at Sigcomm: Concerns have been voiced that certain technical topics do not have a "home" at Sigcomm, in the sense that papers within a certain topic area have difficulty being accepted into the conference, or that there is minimal PC technical expertise in the area. Hardware-oriented (*e.g.*, router design) papers, ATM

[6] Papers on new topics often have a hard time because there is little agreement on the proper set of assumptions or on the importance of the problem. This particular difficulty will hopefully be lessened, given the reviewing philosophy espoused in Section 2.1

[7] The web page maintained by Sally Floyd on open problems in networking could be of some use here. Sigcomm should find a way to make better use of this site. One suggestion would be to publish the list of problems each year in CCR. Another would be to have a series of "open problem" short papers (perhaps as short as a page) in CCR.

networking, mobility, and most recently measurement-oriented papers have been anecdotally noted as areas without a firm home in Sigcomm. We believe that it is important that Sigcomm remain a "broad-spectrum" conference that provides an umbrella under which papers from many different areas of networking can be published.

Losing People and Ideas: With its low acceptance rate and narrow set of topics, Sigcomm runs the risk of people tuning out and turning elsewhere. Whether the alternative outlet is another conference (Infocom or Mobicom) or just to stay home, Sigcomm will have lost a possible source of insight and lessened its ability to shape future research.

We should be clear that this is not a question of how best to compete with other conferences. The issue is one of keeping the intellectual base of Sigcomm broad and vibrant.

3.2 Organizational Responses

We now discuss several concrete steps that Sigcomm (the conference and the organization) can take to address these problems. Each of these steps address one or more of the issues cited above. We should also note that improving the selection process addresses these issues indirectly; evaluating papers on the quality of ideas may encourage more architectural papers and broaden the set of topics considered, and both of these may encourage people to continue contributing to Sigcomm.

Call for Papers: While Sigcomm regulars often ignore the formal Call for Papers (CfP), the CfP is the only formal description of Sigcomm's scope of interest and criteria for acceptance. As such, we should update the CfP to reflect the general philosophy described in Section 2.1. Also, the list of topics in the CfP should be revised to incorporate more recent developments.

Technical coverage of the PC: To ensure that papers from certain technical areas have a "home" in Sigcomm, it is important to have PC representation by visible, technically outstanding researchers from those areas. If a conference wants to proactively encourage submissions in a particular area, those PC members should be tasked with "beating the bushes" for top papers.

Panels and Invited Talks: Sigcomm should encourage the use of invited speakers and panel presentations to address important issues. These should help identify important but understudied problems, and should also provide an avenue for practitioners to talk about the reality they face. We propose that Sigcomm invite one outside speaker and have one panel. Example panel topics might include:

- The failed protocol society (a review of seemingly promising protocols that failed)
- An operator's view of the Internet
- Is peer-to-peer going to change everything?
- QoS: DiffServ, IntServ, or neither?

Example invited speakers would include those from outside the traditional Sigcomm research community, but with a broad perspective relevant to networking research.[8]

To make room for the invited speaker and the panel in the program, we make the following observations:

- The Outrageous Opinion session no longer serves any technical purpose and, if it is to be continued, should be moved to the banquet or some other purely social time.

[8] Panels and invited speakers are best arranged 6-12 months in advance of the conference, and thus should be started well before the PC meeting. In fact, these may be the first tasks of the incoming PC chairs.

- Since the Sigcomm Award is given based on past accomplishments and not on the basis of presentation quality, we recommend that the Award talk be shortened slightly (perhaps to 30 minutes).

- We should keep only one of the Work-in-Progress session or the Poster session; we don't need both. We advocate keeping and possibly expanding the poster session, as it allows a larger number of people to more actively participate in the conference and have their research visible at the conference. This helps address the "Losing People and Ideas" problem in Section 3.1. We note that the choice between poster sessions and work-in-progress sessions involves a tradeoff; compared to work-in-progress sessions, poster sessions are less efficient at presenting research results and ideas, but allow more researchers (students in particular) to participate. We imagine that over the years Sigcomm may switch back-and-forth between these two events, or will find a way to include both.

Position Papers *(experiment)*: In order to increase the level of discussion of architectural principles and other high-level design decisions, we recommend that Sigcomm's CfP include a special call for what we call *Position Papers*.[9] Position papers are supposed to address issues where wisdom and judgment are more important than detailed measurement or simulation. They may advocate (or repudiate) a certain design approach, or discuss a challenge facing the Internet, or argue that certain research agendas are misguided.

The implementation of how best to review these papers is largely up to the PC chairs, but for concreteness we present one possibility.

- Position papers are marked as such on submission.

- There will be no more than three such papers accepted in any year, but the PC can, in its wisdom, decide to accept fewer (or none at all).

- Reviews of position papers ask the following questions:

  1. Does this paper address an interesting issue?
  2. Does the paper provide insight into how to think about the issue? Does it provide a new perspective on previous research efforts? Does it suggest a novel, broad research agenda?

- In some cases, it would be appropriate for the presentation of a position paper to be followed by a discussant who gets a few minutes (which would be taken from the presenter's time) to respond to the paper. A panel discussion of a coherent set of position papers is another option.

Thus, we are proposing that the Sigcomm technical program include the following: a normal slate of standard technical talks, perhaps as many as three position paper talks (taking regular talk slots), the Sigcomm Award talk (shortened), one additional invited talk, one panel, and a student poster session (or, alternatively, a work-in-progress session). To allay worries that this leads to an overcrowded schedule, we've laid out a possible program in Appendix B, showing that fitting in all these events is quite feasible.

Sponsoring Other Venues: Sigcomm should aggressively support other venues for networking research. Sigcomm has already sponsored the measurement workshop (scheduled for Fall of 2001), and there is at least one other proposal being floated (by David Wetherall and Larry Peterson). We view this as a very desirable development. Providing more outlets for networking research will lessen the pressure on Sigcomm as being one of the very few high-quality networking conferences, and will also foster more in-depth technical discussions on a wide range of topics. If these workshops include shorter papers on work that is in the formative stages, then they can also serve as places where early stages of work can receive exposure and feedback before being submitted to Sigcomm.

[9] This name is not optimal, and suggestions for a better name are hereby solicited.

CCR: Sigcomm currently has an arrangement with ToN to fast-track accepted papers. We propose augmenting this with an arrangement with CCR to fast-track some number of the rejected papers that were considered quite strong. The CCR editor could use the Sigcomm reviews to identify papers for publication in CCR. The CCR editor would invite the authors of those papers to submit with a known publication date. Of course, it is up to the authors to decide if they want to publish in CCR or submit elsewhere.

Using CCR as the publication vehicle for good papers Sigcomm was not able to fit into its program would provide authors with more options, and may improve the quality of CCR overall. CCR has uniquely fast turnaround and wide dissemination, and the entire research community benefits by getting ideas into circulation as quickly as possible. Rejected authors currently find themselves in a predicament where they must choose between waiting for Infocom (which entails another year's worth of delay) or sending it to a lesser conference or sending it straight to a ToN (which has its own long delay). Publication in CCR does not preclude future publication in ToN or other ACM journals, so CCR could be viewed as taking the place of a conference publication.

4 The Role of the TAC

There are two clear centers of authority when it comes to the Sigcomm conference: *(i)* the Sigcomm Executive Committee (ExCom), which has unambiguous and complete control over all matters of Sigcomm policy and *(ii)* the Sigcomm PC chairs who, following Sigcomm tradition, have a large degree of independence and freedom in their choice of PC membership and procedures. Nothing we recommend here is intended to infringe on either of these two authorities.

However, we feel that the TAC has a valuable *advisory* role to play in Sigcomm. We emphasize the term "advisory" in that both the ExCom and the PC chairs are free to ignore the TAC's advice. But the TAC can play a special role as it combines the ExCom's long-term perspective with the PC chairs' focus on the technical aspects of the Sigcomm conference (whereas the ExCom has a much broader, and more taxing, charter). The TAC is thus uniquely positioned to function as the institutional memory of the Sigcomm conference.

The TAC currently gives advice to the ExCom about future PC chairs. We recommend that this continue. We further recommend that if any of the proposals described here are adopted by the ExCom they be made clear to future PC chairs *before* they are selected. If future PC chairs are not sympathetic to these new directions then they will never be implemented.

The TAC can give PC chairs advice on who might make good PC members, which outside reviewers might be appropriate, and what PC procedures have worked in the past. We strongly urge that the ExCom *require* future PC chairs to consult with the TAC about all PC and procedural matters. The PC chairs are always free to ignore this input, but they should be required to at least hear the comments of the TAC.

In keeping with its role as the institutional memory of Sigcomm, the TAC should institutionalize the discussion of what worked and what didn't after every Sigcomm. As discussed earlier, we propose that a TAC *monitor* be charged with collecting notes up through the PC meeting itself. After the PC meeting, the TAC should meet to discuss how things went. A second discussion should be held after the conference itself.

A Sample Review Form

Current Sigcomm Review Form

```
Reviewer name:

Paper title:

Overall recommendation:
5  =  Strong accept: An extremely fine paper, should be considered for a
      best paper award.
4  =  Likely accept: A very good paper. Should be published even if it
      means extending the program
3  =  Accept if room: A paper that I would be happy to see published.
2  =  Likely reject: Not a bad paper, but not exciting enough to get
      over the threshold.
1  =  Definitely reject

Reviewer familiarity:
5. Expert   4. Strong   3. General   2. Marginal   1. Not acquainted

Comments on paper (to be returned to authors):
  Please comment on the originality, technical merit (value, correctness),
  relevance (to the field and to SIGCOMM), and readability. Specific
  comments addressing these issues, as well as potential improvements
  in the paper, are particularly valuable.




Comments to PC (not returned to the authors):
```

Possible Revised Sigcomm Review Form

Reviewer name:

Paper title:

Reviewer familiarity:

```
1  =  I know nothing or almost nothing about this topic.
2  =  I am somewhat familiar  with  this  topic,  but  I can't  claim
      that  I am an expert; or, my work in this area is far out of date.
3  =  I am well versed in this area,  but  it  isn't  my direct area of
      specialty.
4  =  This is my area.
```

OVERALL EVALUATION. In your estimation, how would you rank this paper with respect to other papers that have been submitted to Sigcomm. If you are unfamiliar with Sigcomm submissions, how would you rank this paper with respect to papers that are submitted to a broad range of networking conferences.

```
__ Top  5% of submitted papers.
__ Top 10% of submitted papers, but not in the top 5%.
__ Top 25% of submitted papers, but not in the top 10%.
__ Top 50% of submitted papers, but not in the top 25%.
__ Bottom 50% of submitted papers.
```

PUBLIC COMMENTS TO THE PC. Your comments below, which will be used by the PC in making paper selection decisions, will be returned to the authors. Please address the following questions in your public comments:

1. **Does this paper address an interesting issue?** To what extent is the *topic* of the paper important and interesting? If the issue or problem addressed were (or now is) completely understood or solved, how important would that be, in terms of either fundamental concepts, or increased understanding, or practical relevance?

2. **Does this paper present interesting results?** Do the results provide worthwhile insight into the topic addressed? Are the results likely to be widely used by others? Does the work open up new areas, present new ideas, and/or serve as a foundation for new work?

3. **Is this paper sufficiently well executed?** Are there flaws (*e.g.*, technical mistakes, important uncited related work, poor assumptions, insufficient scope of evaluation, unsubstantiated conclusions, poor writing) in the paper? Are the flaws fundamental or superficial? That is, are the results likely to be true despite the flaws, or do the flaws fundamentally impact the results in the paper?

PRIVATE COMMENTS TO THE PC. Your additional comments below will NOT be returned to the authors.

B Sample Conference Schedule

This schedule includes 27 technical talks (expandable to 30 by adding on another talk session on the 3rd day), a panel, an invited talk, a poster session, the Sigcomm Award talk, the Sigcomm business meeting, and the banquet. If the Outrageous Opinions session is included, it could follow the Sigcomm business meeting (as a social event) or be included in the banquet.

Day 1

9:00- 9:15 Welcome

9:15-10:00 Sigcomm Award Presentation and Talk

10:30- 12:00 3 Technical Talks

12:00- 1:30 Lunch (and poster set-up)

1:30- 3:00 3 Technical Talks

3:00- 4:00 Poster Session

4:00- 5:30 3 Technical Talks

5:30- 6:30 Sigcomm Business Meeting

Day 2

9:00- 9:45 Invited Speaker

10:15- 11:45 3 Technical Talks

11:45- 1:15 Lunch

1:15- 2:45 3 Technical Talks

3:15- 4:45 3 Technical Talks

5:00- 6:00 Panel

Evening Banquet

Day 3

9:00-10:30 3 Technical Talks

11:00- 12:30 3 Technical Talks

12:30- 2:00 Lunch

2:00- 3:30 3 Technical Talks