

# Experiences from the Sigcomm 2005 European Shadow PC Experiment

Anja Feldmann  
Technische Universität München, Germany  
anja@in.tum.de

## ABSTRACT

This note is an informal report about the Shadow Program Committee experiment that I organized for Sigcomm 2005. It (i) discusses the motivation for a Shadow Program Committee (PC), (ii) presents the expected benefit to the community, (iii) explains the process by which it was organized, (iv) gives some statistics on the program that the Shadow PC selected, how much, and why it differed from the real PC's program, and (v) discusses some of the insights from the Shadow PC process.

## Categories and Subject Descriptors

A.m [General Literature]: Miscellaneous

## General Terms

Experimentation, Management, Documentation

## Keywords

Shadow Program Committee, Experience report

## 1. MOTIVATION

Ever since learning about the concept of a Shadow PC and the experiences from NSDI with Shadow PCs, I have been intrigued by the idea. It seems like an ideal concept for educating the participants, usually junior researchers, about the review and paper selection process of highly competitive conferences and exposing the criteria that determine if a paper is perceived to be a good one. A Shadow PC is a PC that will not have *any impact* on which papers are going to be accepted at the conference. Yet the Shadow PC will be run just as the real PC. (Shadow PC members are assigned some number of papers for review and are expected to attend a PC meeting at which highly rated papers are discussed and a program is selected.)

My motivation for proposing to organize a European Shadow PC experiment for Sigcomm originated from several observations: First, some members of the community perceive Sigcomm as highly US-biased in the sense that most Sigcomm authors and PC members are either at research institutions in the US or have strong ties to US research labs such as Sprint or AT&T. Second, some people consider Sigcomm as cliquish, with the chance of having a paper rejected being proportional to the author's distance to the in-crowd. The impression is that even the double blind submission and review process does not prevent this. The Shadow PC experience should allow people outside the "inner" core to experience a Sigcomm review process (the Shadow PC) first hand. Third, it seems

that it is always the same subset of Europeans that serve on the PCs indicating that there is a lack of potential reviewers and PC members from outside the US. Yet the last straw for going ahead and proposing to organize the Shadow PC experiment came while attending a PC meeting for some other European conference. I was underwhelmed by its professionalism with regards to handling of conflicts-of-interest and unwillingness of PC-members to stand up to their reviews. The first two observations highlight the value of participating in PCs, which is almost impossible for junior researchers, while the second two focus on why a Shadow PC experiment is rather valuable for Europe.

## 2. BENEFITS

In terms of discussing the expected benefits from the Shadow PC experiment I want to distinguish three different groups: the authors, the Shadow PC members, and the larger Sigcomm community.

The authors receive additional reviews for their papers that will reflect the differences of opinions on papers from people with varied backgrounds: more experienced reviewers (actual PC) vs. reviewers that are actively working in the area (Shadow PC) and people at the core of a conference (actual PC) vs. reviewers that are new entrants to the Sigcomm community. Eventually this might lead to a more transparent review process.

The Shadow PC members get a first hand experience on how a PC is organized and how it operates. Accordingly, they experience the process Sigcomm uses to select papers, have a chance to read top notch papers in their area of expertise, get to know the other PC members, and experience the double blind reviewing system first hand. But most importantly they get to see good quality papers as well as bad quality papers (helpful for calibrating expectations) and experience how papers are viewed by others. The latter is probably the most important aspect and should provide a good foundation for exploring what ingredients make a paper a good one and what parts have to be addressed in order to succeed in getting a paper accepted at a highly selective conference such as Sigcomm.

The larger Sigcomm community can benefit by making the Sigcomm review process more transparent. Additionally, the Shadow PC experiment is an outreach activity and it should make Sigcomm more accessible to non-US based people.

## 3. PROCESS

The process of running the Shadow PC consisted of several steps: the proposal, choosing participants, the organization, and the actual experience.

### 3.1 Proposal

The process of organizing the Shadow PC got started when I took the initiative and contacted the Sigcomm chair, this year's Sigcomm PC chairs, and some Sigcomm Technical Advisory Committee (TAC) members regarding the possibility of organizing a "European" Shadow PC for this year's Sigcomm.

The reason why I wanted to restrict the Shadow PC to Europe is that I felt that the Shadow PC is an outreach activity. As I am now based in Europe I felt that Europe is a good starting point. In addition, limiting the participation to researchers working in Europe reduces the cost as well as time for travel.

The proposed process was to select a PC of roughly 40-50 PC members. This is about 2-2.5 times larger than the actual PC which reduces the workload while ensuring that each member has to read a significant number of papers. The PC members should either be senior graduate students or junior faculty members. For the purpose of choosing PC members I wanted to be as open as possible yet reach as many different organizations as possible. To seed the process of PC member selection I proposed to use private contact lists, publication records, participants in European projects, etc. Each prospective PC member would be asked to name other groups or persons that should receive an invitation to the Shadow PC.

In terms of the PC meeting I envisioned a two day meeting. The first day should provide time for a Sigcomm style PC meeting based on the reviews of the Shadow PC. On the second day each Shadow PC member would be able to access the reviews of the actual PC and compare them to the reviews by the Shadow PC. Next we would discuss the differences between the reviews and the differences between the accepted papers by the real PC vs. the Shadow PC. The goal of such a two part PC meeting is to allow participants to understand why some of the decisions are different and what qualities the actual Sigcomm reviewers are in fact looking for.

Finally the reviews by the Shadow PC should be returned to the authors, clearly marked as Shadow PC reviews.

After quite some discussion about the advantages and disadvantages, the contacted people agreed on the upside of the proposal and the decision was made to enable it with support from the PC-chairs as a onetime experiment.

### 3.2 Participation

The participation in the Shadow PC experiment was absolutely voluntary for both the authors of submitted papers as well as the actual PC members. Each author and each PC member was given the choice of not participating in the sense of not making their paper available to the Shadow PC or by not sharing their reviews. To help authors and PC members with their decision I put together a Web page<sup>1</sup> outlining the motivation, the benefits, and the criteria for participating in the experiment as well as the list of Shadow PC members. The authors of 41 out of a total of 255 papers declined to make their paper available and 2 out of 20 PC members declined to share their anonymized reviews. Thanks to all who made the experiment possible! In order to avoid conflicts of interest I furthermore decided to remove 5 papers I had conflicts with. (A reviewer has a conflict-of-interest with a paper if (i) the reviewer is an advisor/student of one of the paper authors (ii) the reviewer and one of the paper authors have recently worked on a common project, e.g.,

<sup>1</sup>[http://www.net.in.tum.de/~anja/sigcomm2005/shadow\\_pc.html](http://www.net.in.tum.de/~anja/sigcomm2005/shadow_pc.html)

published a paper, (iii) the reviewer and one of the paper authors work at the same institution.)

After the go-ahead I put together a seed list of research groups all around Europe and started to contact each one regarding participation. I decided to mainly contact the heads of the research groups and ask them to nominate participants as they are most familiar with the work and the standing of their students and junior faculty members. More explicitly, I started soliciting participants from more than 60 groups from December through January stating: "I would appreciate if you could help me identify an interested candidate within your group or institution. Also if you can think about someone else that I should contact please let me know."

This resulted in 37 Shadow PC members. Another person was added, after he/she enquired about the Shadow PC on receiving the email to the authors outlining the Shadow PC experiment. After a discussion with some PC members about the Shadow PC it was strongly suggested that individuals from Asia and Latin America should also be enabled to participate. As a result we constructed another seed list and I contacted more than 10 groups. This resulted in the addition of four more Shadow PC members from Asia. Overall the 42 Shadow PC members represent more than 34 different research groups from 15 different countries.

### 3.3 Organization

For the organization of the Shadow PC we used the same software as the actual PC did. This means that we set up a separate database for the Shadow PC that was seeded with the 209 papers available to the Shadow PC. The major difference between this database and the actual one was that it only contained anonymized author names. In order to ensure that the double blind review process was not jeopardized, I as PC member should have no way to access the author names. Unfortunately it made the identification of conflicts more difficult as we could not use the standard procedure. This procedure identifies conflicts based on the name of the collaborators of the PC member and his/her institution. In the end we relied on the Shadow PC members to look through the titles and abstracts and identify if they had a conflict with a paper. We identified a conflict if the Shadow PC member suspected that someone they collaborated with in the last year is a co-author on the paper or if they suspected that someone that is at their institution at the same location is a co-author on the paper. While this is not a perfect solution it seemed sufficient. Furthermore, since the Shadow PC members are junior researchers from outside the typical Sigcomm community the total number of conflicts we identified was significantly smaller than in the actual PC.

In terms of reviewing assignments and process we followed the same scheme as the actual PC, including using a logarithmic scale for rating papers, i.e., from top 5% over top 10%, top 25%, top 50% to bottom 50%. This meant that I could usually turn around most of the emails I received from the PC chairs after query/replace the Chair's names with mine and "PC" with "Shadow PC":-)

However, before the start of the review process I emailed each Shadow PC member instructions regarding the Shadow PC process. This included the following comments:

There are *obligations* attached with reviewing:

- the papers and their content are confidential!
- you are not allowed to plagiarize material gleaned from the reviewed papers!
- try not to find out who the authors are!
- not to review a paper that you have a conflict with

- not to participate in the discussion of a paper that you have a conflict with.

Also, keep these recommendations in mind while reviewing:

- be open minded
- look for papers of broader interest
- do not look just for a neat technical trick
- if there are flaws are they easy to fix or fatal
- step back before judging a paper. Try to consider the bigger picture.

In addition, I put together a set of recommendations derived from a document that describes how a Sigcomm PC works: “Improving Sigcomm: A few straw proposals”<sup>2</sup>. This document was prepared by Sigcomm’s Technical Advisory Committee and approved by the Sigcomm Executive Committee.

Each Shadow PC member got access to the list of paper titles and abstracts and was asked to submit two lists: one with papers they preferred to review (in order of preference) and one with papers they did not want to review. Based on these lists and the research area of each Shadow PC member (obtained either via personal knowledge or via their Web pages), I assigned each paper one primary reviewer and two secondary reviewers from the Shadow PC<sup>3</sup>. The difference between a primary review and a secondary one is that the primary reviewer is asked to file a full review. The secondary reviewer are only asked to file a full review if the paper either is in the top 50% part of his/her papers or is rated by one of the other reviewers as being in the top 10%. This resulted in assigning between 14 and 17 papers to each reviewer. The actual PC members got up to 39 papers in this round of assignments.

After filing the reviews, the Shadow PC entered the discussion phase. At this point all reviewers who had filed their reviews got access to the other reviews and a discussion about the merits of the paper was initiated. This resulted in a grading of most papers as to whether they should be discussed at the PC meeting or not. In addition, it allowed the PC members to recalibrate their ratings. At the same time, two Shadow PC members who helped me with the administration went through the reviews to identify papers that did not have enough reviews or that had reviews with widely differing ratings. They either assigned an additional reviewer to the paper or asked reviewers who had rated their papers in the bottom 50% without filing a detailed review to provide one. I delegated this task to ensure that I myself kept unbiased until after the actual PC meeting.

Since the actual PC meeting happened before the Shadow PC meeting I also asked every member to be cautious in the following sense: “There is the natural curiosity to find out which papers have been accepted by the PC! Yet in order to make our experiment work please try to keep a certain distance and do not inquire if any of the papers that you reviewed made it...”

For the Shadow PC meeting we selected the following set of papers for discussion:

- those that had at least one rating in the top 10%
- those that had at least one grade stating it should be discussed
- those that had a high average rating

After sending out this list to the Shadow PC we removed papers where the primary reviewer explicitly stated that the paper should

<sup>2</sup><http://www.acm.org/sigs/sigcomm/admin/July2001RepFinal.pdf>

<sup>3</sup>I myself did not review any paper for the Shadow PC.

not be discussed and added those that were suggested for discussion. This resulted in about 65 papers. In addition, it was possible to add papers to the discussion list during the PC meeting. This added 10 additional papers to the discussion list.

### 3.4 Shadow PC experience: Day 1

The Shadow PC meeting on April 26th 2005 started in Munich at 9:15am and ended at 7:15pm with 34 participants. Unfortunately, the Shadow PC members from Asia could not secure funding for travel to Germany, two excused themselves due to family emergencies, and one due to being a reviewer in an EU meeting that happened at the same time.

For the purpose of classifying the papers we did not just use the two categories “Accept” and “Reject” but rather the following four:

- Accept
- Exciting but flawed
- Boringly correct
- Reject

The idea behind this classification is that just because a paper is well done and the results are well presented it should not necessarily be preferred to one that is exciting if it does not present a major step forward. Of course the papers in the categories “Exciting but flawed” and “Boringly correct” eventually have to be moved to either “Accept” or “Reject”. More initial categories allow for a more explicit decision if the excitement that a paper might create and the discussions that it may motivate supercede the flaws or if the flaws are so large that it is better to reject the paper in its current form. It also enables an explicit decision on whether those papers in the “Boringly correct” category present enough of a delta to be worth presenting at Sigcomm or whether they are YAP (Yet Another Paper) on some topic.

The process of selecting papers proceeded in four stages: In the first stage the primary reviewer was asked to briefly present a summary of the paper and then their standpoint on the paper. Then the other reviewers were asked for their standpoint. Then the PC as a whole reached a consensus regarding which category to put the paper in. As a side note, sometimes people jokingly suggested to put a paper in the category “Boringly flawed”. The discussion of the first set of papers took longer as the PC overall needed to be calibrated. In order to not focus too much on either the highly rated papers or the lower rated papers we used a discussion order that intermixed highly rated papers with lower rated ones.

After the first round the Shadow PC had 12 papers in the “Accept”, 11 in “Exciting but flawed”, 11 in “Boringly correct”, and 32 papers in “Reject”. This process led to many interesting discussions and arguments and helped the participants to realize what is required for a paper to be accepted outright.

The second round revisited the papers in the “Exciting but flawed” and the “Boringly correct” category. For each paper the reviewers were asked to briefly recap the paper and the standpoints and then after some discussion the PC decided about moving the paper to “Accept”, “Reject”, or whether to reconsider it once more. In addition, it was possible to revive papers by either adding additional papers to the discussion list or by reconsidering papers from the “Reject” category. This led to 7 revived papers, 10 accepted papers and 7 papers to be reconsidered. At this point it was already early evening and the PC was getting a bit tired, despite the exciting discussions.

After a short break the PC reconvened for some “scraping at the

bottom of the top”. Some opponents gave in or were convinced by the proponents and after some lengthy discussion another 4 papers were accepted. At this point the PC was tired but still very excited. After all 26 papers were accepted and only 3 papers were left in the non “Reject” categories.

Since I, as the Shadow PC chair, wanted to ensure that the Shadow PC selects a full program I did not want to stop at this point. After all it should be possible to identify 28 good papers among 209 considered papers. Therefore I decided to go for one more round with the goal of adding two more papers. The arguments regarding the last three papers were once more recapped and then the whole Shadow PC voted regarding acceptance or rejection of each of the papers. Not too surprisingly, this led to the selection of two more papers, one of them after a tie in the voting.

At this point the Shadow PC had selected its program. Therefore it was possible to reveal the outcome of the actual PC. This resulted in quite some ooh’s and aah’s as well as some comments. After this the PC rushed out for a well earned dinner at a Munich brew pub.

### 3.5 Shadow PC experience: Day 2

In the morning of the second day we emailed each Shadow PC member the “List of accepted papers by the Shadow PC”, the “List of accepted papers by the actual PC”, a high level comparison of the results of the Shadow PC vs. actual PC, and the anonymized reviews of the actual PC for all non-conflicting papers. In addition, I distributed copies of two papers for further discussion to the Shadow PC: one that was accepted by the actual PC but not the Shadow PC and one that was accepted by the Shadow PC but not the actual PC. Then I gave each Shadow PC member some time to look at the reviews and to read the papers.

At about 11am we started the meta discussion by reviewing the merits of the paper accepted by the Shadow PC but not the actual PC. Then we discussed possible reasons for why this paper might not have been accepted by the actual PC. This led to a rather lively discussion and in the end the Shadow PC understood the reasons for the decision of the actual PC. Then we did the same for the paper accepted by the actual PC but not the Shadow PC. We discussed and comprehended the reasons for the difference in the decision.

Overall, this led to a long argument about the insights and take-aways from the previous day, including the review process, review quality, review depth, personality of the reviewer, as well as presenting and discussing a review, the possibility of steering a PC either by pushing for a paper or against it, what defines a good paper, including the presentation, the style of writing, etc. as well as how to pick good research problems to work on. For this discussion it would have been helpful if more senior people would have participated as some of my answers were obviously biased. After a late lunch break we got back together and examined among other questions how the double blind review process and how the Shadow PC meeting differed from other PC meetings that the PC members had participated in.

## 4. STATISTICS

In terms of pure numbers, of the 27 papers accepted by the actual PC the Shadow PC accepted 11, discussed but rejected 10, did not discuss 2, and could not consider 4 papers. (Of the 10 papers discussed by the Shadow PC but rejected, 4 were rejected in rounds three and four.) 10 papers accepted by the Shadow PC were not discussed by the actual PC and 2 papers accepted by the actual PC were not discussed by the Shadow PC. Note that the pool of papers

Topic	PC	Shadow PC
Architecture	3	2
Classification	4	1
Measurements	4	4
Overlays	3	3
Routing	4	6
Security	1	6
TCP	6	2
Wireless	2	4

Table 1: Difference in the selected papers by topics.

PC	ranking base	> 10				
		≤ 10	≤ 20	≤ 50	≤ 100	> 100
Shadow	Shadow	6	4	10	7	2
actual	Shadow	3	5	7	6	6
Shadow	actual	4	3	4	9	8
actual	actual	7	6	7	5	2

Table 2: Difference in the selected papers by rankings based on the reviews.

available to the Shadow PC was smaller than for the actual PC. Therefore it is not surprising that the number of accepted papers not discussed by the actual PC is larger for the Shadow PC than the reverse.

In terms of topic areas I classified the papers into the following eight areas: papers presenting an architectural idea (architecture), papers about packet classification and network processing (classification), papers presenting measurement methodologies or measurement results (measurements), papers regarding overlays (overlays), papers on improving routing (routing), papers proposing security mechanism or dealing with attacks (security), papers regarding improvements of TCP (TCP), and finally papers concerning wireless networks (wireless). Table 1 shows the results. It appears that the actual PC preferred papers about the evolution of TCP and packet classification while the Shadow PC had a preference for papers in the wireless and security domains. Put differently “the Shadow PC did not care as much about TCP details :) and was also unsympathetic to classification”.

At first the difference seems huge. Yet there are a few things to keep in mind. Which papers are accepted depends on the reviews that are available to each of the PC. Therefore we performed the following experiment: we ranked the papers according to the normalized review ratings and then counted how many papers the Shadow PC and the actual PC accepted out of the top 10, top 20 but not top 10, top 50, top 100 papers, see Table 2. Interestingly both PCs, using their reviews liked the highly ranked papers (those in the top 10) and did not accept many papers, 2 each, that were not ranked among the top 100. In terms of choosing from the other categories they behaved about the same. Therefore one can say that both PCs did their jobs well based on their respective inputs. If one tries to judge the performance of the PCs based on the ranking of the other PCs, both again behave about the same. The likelihood of accepting a paper out of the top 10 decreases and the likelihood of accepting a paper out of the top 50 increases substantially.

I note, that the enthusiasm of the Shadow PC members was amazing. Almost all of them filed their quite lengthy reviews on time in contrast to some of the actual PC members. Indeed, in terms of review length the median length for the actual PC reviews<sup>4</sup> are 316

<sup>4</sup>The reviews that were shared with the Shadow PC.

words vs. 199 for the Shadow PC. The length of the 5% longest reviews for the actual PC are 836 words vs. 808 for the Shadow PC. But the actual PC had one advantage: it gathered a huge number of outside reviews from experts in the area. This is the one aspect of a Sigcomm PC that I did not include for the Shadow PC. This means that while the Shadow PC had the same number of PC reviews, three to four reviews<sup>5</sup> the actual PC had at least three but often up to five PC reviews plus usually one to three outside reviews. A consequence is that the number of readers per discussed paper at the Shadow PC meeting was somewhat less. In addition assigning papers to reviewers was harder since I had not met all Shadow PC members and was not that familiar with their particular research interest. Accordingly, the level of knowledge of the Shadow PC member on a paper's topic was either significantly higher than that of the actual PC if the Shadow PC member is currently pursuing or has just finished a Ph.D. thesis in the area, or alternatively, might be lower. Therefore, more variance in the review depth and breadth can be expected from the Shadow PC than from the actual one. Our experience confirms this expectation.

To get a feeling for the success of the Shadow PC experiment I put together a anonymous review form<sup>6</sup>. All participants who filled out the form felt that participation in the Shadow PC experiment was "very much" a useful experience. More than 70% of the participants felt that the instructions were "very clear" while the others "kind of knew" what was expected. With regards to the kind of papers Sigcomm is looking for only one participant felt that it was unclear, while to most participants it was clear or they at least had a good sense about it. In terms of time-frames most (72%) felt that the review time-frame was sufficient. With regards to the one day PC meeting most felt it was sufficient but 20% did not deem it sufficient. Some of these suggested to remove some of the randomness in the decision process by getting additional reviews for the borderline papers.

In terms of the most useful and least useful part the opinion was split: 44% judged the participation in a PC meeting as the most useful aspect, 28% felt that the meta discussion about the process was most beneficial, 20% judged the reviewing of papers submitted to Sigcomm the most valuable aspect, while 8% felt that seeing the papers submitted to Sigcomm was crucial. While meeting other Shadow PC members was never mentioned as most useful it only was perceived as least useful by 16%. Overall judging from the comments it seems that folks had a hard time deciding which part to declare as the least useful one, i.e., one person stated: "really hard to choose if you were keen on the whole process. An additional option, e.g., did not dislike something, would help :)". Accordingly it is not that surprising that 40% picked seeing the papers submitted to Sigcomm, 28% picked the discussion about the process, 16% picked meeting the other Shadow PC members and 8% picked the other two options.

## 5. INSIGHTS

The most common reasons for papers to be rejected were:

- the reviewers did not understand the paper, either due to an unclear motivation, or because the relationship to related work

<sup>5</sup>In exceptional cases when one of the reviewer did not file the review in time there were only two reviews.

<sup>6</sup><http://www.net.in.tum.de/~anja/tmp/sigcomm-feedback.html>

was unclear, or due to questions regarding the practicality or the contributions, or due to a bad presentation.

- the novelty of the approach was not convincing.
- the reviewers identified unfixable flaws either in the experimental methodology (e.g., in their simulation setup or in the analysis), or in an unfair comparison.
- the feeling among the PC was that the paper was out of scope, e.g., too layer one/two oriented, much more of a security paper, etc.
- someone made a strong case that the paper should be rejected.

The most common reasons for papers to be accepted were:

- someone stood up for the paper and championed accepting it. A champion is convinced that the paper is good and usually either has a good knowledge of the area or can convince the rest of the PC of the benefit of the paper. The champion has to be able to convince the rest of the PC of his opinion.
- it appears that the paper can lead to interesting discussion
- a paper addresses a real problem in real networks with a feasible solution.
- the reviewers might not have understood the paper but are convinced by the argumentation and the presentation. No one identified a potential flaw.

Among the insights that some of the Shadow PC members identified about getting a paper accepted at a conference like Sigcomm are:

- A paper should be accessible to a non-expert reviewer. This requires concise presentation of the main ideas in order to not lose the reviewer. The writing should clearly identify why the problem is an important one and that the solution is a significant contribution. A paper should highlight the potential of the idea to motivate interest in it.
- A paper should make a good case in order to make it harder to reject it. It is good to discuss how the work relates to all relevant publications and describe why other publications are not relevant. Potential attacks by some reviewers can be prevented by addressing and resolving the potential issues in the paper itself.
- It is good to seek feedback in early stages of one's work. This helps in discovering related work, and gives one an impression of the work's relevance. In addition, this is a good dry run for writing the paper.
- Parts of the selection process are "semi-random". Borderline papers can go either way depending on the PC dynamics.
- It is a good idea to read a lot of papers and get inspirations on presentation style in general and especially on writing style.

In terms of additional insights by the participants, I would simply like to quote some of them:

- "This Shadow PC was a good experience to show young researchers that bias in the acceptance of the papers in a conference does not come from clusters of people who bias the decision of the PC in favor of their paper, but that any PC leads to a lot of randomness with regard to the papers that are not either obvious rejects or accept."
- "... the papers I reviewed surprised me at the large diversity of quality. ... The outcome was not really that surprising once we realized that the number of straight accepts was

small. The differences between our and the real PC are visible but ok for me.”

- “The experience of the (uncontrollable) group dynamics. . . . The necessity of stepping back to judge a paper more objective. . . .”
- “. . . thank you for the awesome experience you offered us. . . . I had the occasion to meet valuable people both from the human and intellectual perspective, not to mention the inspiring atmosphere!”
- “. . . Also, the participants were very nice, I had a lot of good personal talks.”
- “. . . At the end of the first day there was definitely a sense that the process was becoming more and more random. If this happens in the real PC too, then fine - and perhaps it’s a learning experience no matter what. . . .”
- “. . . The main goal of a paper that is discussed in a PC meeting should be to convince PC members to speak up for it. . . .”
- “. . . Who should attend to Shadow PCs? . . . I believe that anyone who want to get to know a conference, what’s the philosophy behind the conference, who are the PC members and what do they think. Of course, not to mention the possibility to read (generally) high quality papers in advance and to get to know what are the current research trends.”
- “. . . Generally, my main lesson learned was that the average quality of SIGCOMM submissions is amazingly poor, which I hadn’t anticipated.”
- “. . . All in all, an eye-opening and worthwhile experience!”

## 6. SUMMARY

In effect the experiment highlights that there are a significant number of submissions that some PC might find exciting and worthwhile to present. Which subset is selected depends to some degree on the composition of the PC and their different visions and expertise.

The feedback of the participants is very positive and suggests that the benefits of participating in a Shadow PC experiment are indeed the ones outlined earlier. The most common feedback regarding future Shadow PCs is exemplified by: “Do it again!!!”, “Two words: DO IT!”, “do Shadows PC on a regular basis, and for most important conferences!”.

In terms of suggestions for improvements, it was suggested that each Shadow PC member should prepare a one page summary about him-/herself and their research to put a “name to a face”. Also the idea of accepting papers with a shepherd to fix presentation flaws and the possibility of accepting papers as “position papers” should have been pointed out in advance and not as the necessity arose.

Furthermore some would have liked to see a more specific goal for the second day, besides the meta discussion. One Shadow PC member’s proposal for this is to come up with documented answers to questions such as “what endears a paper to a PC”, or “what factors cause a discussed paper to be rejected at a PC”, or “what factors caused different decisions in the two PCs”. Also it might help if several senior people would participate and the outside reviews would be available to the Shadow PC. An alternative for enabling more discussion about the differences in reviews and to enable a more in-depth analysis might be to split the group into subgroups on the second day according to their sets of interests. Another interesting suggestion was to have both PC meetings at the same time and then have a joint discussion about the differences on the second day. Indeed this would ensure that the Shadow PC chair has no way to bias the Shadow PC meeting. I have to admit that at times it was difficult to stay unbiased. Indeed, maybe next time the Shadow PC should also address the question if the actual Sigcomm process is correct or could be improved.

Overall, despite the work involved in organizing the Shadow PC, I personally very much enjoyed the adventure and the experiment seems to have had the expected benefits for the three target groups. Therefore I highly recommend that Sigcomm enable future Shadow PC. This is not to say that there should be a Shadow PC every year. But having one every so often for different subgroups of the community seems sensible. In terms of choosing subgroups, imposing some scope, e.g., a geographic one, seems to work well as it reduces travel costs and helps in establishing a local community. Yet, I recommend to keep the process of choosing participants as open as possible but somewhat restricted to reach a reasonable number of different research groups.

## Acknowledgments

Many thanks go to the authors that submitted to Sigcomm for participating in the experiment, the actual PC members for sharing their reviews, to Ramesh Govindan and Greg Minshall for their support, Vinay Aggarwal, Petra Lorenz, and Olaf Maennel for the local organization and their support during the Shadow PC experiment, and last but not least to all Shadow PC members.