

Program Chair's Message

Thank you for the privilege and burden of being the Program Committee Chair for the *2005 International Symposium on Computer Architecture (ISCA 2005)*. I first describe the process we used to select the ISCA 2005 technical program to both aid those who will design processes for future conferences and inform those who participated in the current process. I then thank those who made it happen. After this message, I print *Guidelines for SIGARCH Sponsored Conferences*.

We selected ISCA 2005's technical program with a process that largely followed recent practice, in accordance with SIGARCH guidelines, but with a few modest adjustments that I will touch upon. In consultation with others, I selected a *23-member program committee (PC)*. I tried to balance areas, institutions, and experience. The PC was somewhat smaller than recent years to facilitate PC meeting discussion dynamics at a cost of each PC member having to review more papers (about 26 reviews per PC member). At least some PC members felt, however, that reviewing 26 papers was too much a burden and encouraged a return to 30-member PC committees. Others thought that the smaller committee helped discussions. I restricted PC members to, at most, two submissions each to follow SIGARCH guidelines and reduce the likelihood of too many accepted PC papers making the conference (appear) inbred.

We received 194 submissions. Author names and institutions were not explicitly provided on manuscripts to facilitate "blind" review. I assigned most papers to be reviewed by three non-conflicting (defined below) PC members and two non-conflicting external reviewers. I did not follow past practice of having PC members assign external reviewers, because, in my experience, this practice often leads to reviews from junior students, reviews correlated with the assigning PC member, and puts extra burden on the PC whom I would rather have reading papers. We did not reveal author names to PC or external reviewers. I did not personally review any papers, because I knew author names. Joel Emer assigned all reviewers for most papers for which I had a conflict. Susan Eggers assigned papers that had conflicts with Emer and me. In all cases, I sent all email to PC members and external reviewers, so that they would not learn whether papers had conflicts.

I defined that authors, reviewers, and PC members have a conflict of interest with:

- Your Ph.D. advisor and Ph.D. students forever.
- Family relations by blood or marriage forever (if they might be potential reviewers).
- People with whom you collaborated in the last five years. Collaborators include co-authors on an accepted/rejected/pending research paper, co-PIs on an accepted/rejected/pending grant, those who fund of your research, and researchers who you fund. You may exclude "service" collaborations like writing CSTB report or serving on a program committee together.
- People who shared your primary institution in the last five years.
- Others with whom you believe a conflict of interest exists.


I, for example, added all former and current members of the Wisconsin Wind Tunnel and Multifacet Projects to my permanent conflict set.

With some nagging, I received all solicited external reviews before the start of the rebuttal period. Twenty-one PC members completed their reviews before the rebuttal period. We marked all later reviews with something like "RECEIVED AFTER THE START OF THE REBUTTAL PERIOD" so that the PC would understand that those reviews may not be rebutted. Most authors submitted review responses (i.e. rebuttals) during the 48-hour rebuttal period.

We did more work between the rebuttal period and the PC meeting. For some papers with varying views, I solicited late reviews from both PC members and external reviewers. I also encouraged PC members to reflect on other reviews and rebuttals, to add comments to the web site and optionally adjust their recommendations. I reviewed reviews and rebuttals and sent out dozens of email questions to prime the process. In the end, I received many comments and almost all reviews.

The PC met at the Chicago O'Hare Hilton for a one-day meeting from 8:00 AM to 6:30 PM. We endeavored to discuss a paper by having the reviewer with the most-favorable ranking identify the authors, summarize the paper, give its strengths, and then discuss weaknesses, external reviewer comments, and rebuttal. Next, other PC members who reviewed the paper added comments. Finally, the whole PC discussed the paper. In this manner, we discussed all of the papers ranked in the “top” 50, as well as more than 30 additional papers that some PC members indicated were worthy of discussion with credible chance of being accepted. Most papers were accepted or rejected in this way. A few papers were deferred as “conditional accepts” (good depending on opportunity cost) or “middle of the road.” We resolved these at the close of the meeting. To encourage papers in new areas, I assigned a paper’s primary rank with $\frac{2}{3} * (\text{mean recommendation}) + \frac{1}{3} * (\text{mean potential})$. To mitigate the effect of a harsh review, I assigned a secondary rank as mean recommendation after removing the single lowest recommendation score. We noted when the two ranks differed significantly. We handled PC papers mid-way through the above process by having PC authors leave the room (in random order) when their papers were discussed. I only asked a PC member to leave the room if they had a submission ranked in the top 100. This is painful, but, in my view, much better than the practice of disallowing PC submissions. We accepted 6 of 23 PC submissions.

In the end, we selected **45 papers from 194 submissions (23%)**. We assigned “shepherds” to five of these papers to guide final versions. The histogram below shows how *accepted papers* over ISCA’s history [<http://www.cs.wisc.edu/~markhill/mp2001.html>]. Submission counts over the last twenty years have been: 194 papers (2005), 217 (2004), 184, 180, 163, 166, 135, 156, 147, 112, 180, 143, 208, 173, 197, 242, 170, 275, 126, and 127 (1986). For information on ASPLOS, HPCA, ISCA, and MICRO, please see: http://www.cs.wisc.edu/~markhill/AcceptanceRates_and_PC.xls.



Throughout the selection process, I encouraged the PC to consider additional papers and additional wilder papers, as I believe our community is doing more good work than is getting recognized due to the limited size of our conferences. My goal was to increase the number of acceptances toward 42. The program committee exceeded this goal.

Some have asked whether restricting the number of submissions per author would reduce the number of submissions we review. For ISCA 2005, the short answer is “no,” unless we restrict authors to an unreasonable single submission. For ISCA 2005, seven authors submitted more than three papers, 21 submitted three, and 48 submitted two. I had my assistant, Caitlin Scopel, select a random paper from an author of more than three papers, delete it for him/her and his/her co-authors, and repeat until paper counts were reduced. Numbers would be different if she selected different papers, but this gives the general trend:

Maximum Papers Per Author	Papers Eliminated	Percent Reduction	Papers Submitted
Unlimited	0	0%	194
3	8	4%	186
2	17	9%	177
1	40	21%	154

I want to thank the many people who made this process happen:

- Caitlin Scopel provided excellent administrative assistance.
- My Ph.D. student, Min Xu, was a tireless Web Master (who now also understands many practical database issues).
- Dirk Grunwald graciously provided the Conference Review Package (CRP) that we depended upon (<http://crp.sourceforge.net/>).
- I got sage advice from ISCA 2005 General Chair Guri Sohi and the ISCA 2005 Advisory Committee: Alan Berenbaum, Bill Dally, Michel Dubois, Margaret Martonosi, Per Stenström, and Mateo Valero.
- The ISCA 2005 Program Committee exhibited professionalism and wisdom: David Albonesi, Arvind, Krste Asanovic, David August, Luiz Barroso, Jose Duato, Susan Eggers, Joel Emer, Garth Gibson, Norman Jouppi, Steve Keckler, Christos Kozyrakis, Anders Landin, Mikko Lipasti, Kathryn McKinley, Sanjay Patel, David Patterson, Steven Reinhardt, Steve Scott, André Seznec, Kevin Skadron, Mateo Valero, and Uri Weiser.
- External reviewers (listed elsewhere) were also amazing, many committing to, and executing, four reviews.
- The hundreds of authors of ISCA 2005 submissions, both accepted and otherwise, are the ones we depend on to keep our field exciting and relevant.

I encourage all to reflect on ISCA 2005's paper section process and technical program to make suggestions to future program chairs.

Mark D. Hill
University of Wisconsin-Madison

Guidelines for SIGARCH Sponsored Conferences

David A. Patterson

Computer Science Division /EECS Dept.

University of California, Berkeley

April 20, 1994

Now appears at: http://www.acm.org/sigs/sigarch/conference_guidelines.html

At the business meeting on April 20, 1994 it was voted that SIGARCH adopt this set of guidelines to aid program chairs on matters of policy at SIGARCH sponsored conferences (e.g., ISCA, ASPLOS, and so on).

Amendments to the guidelines would be voted on at SIGARCH business meetings. There are several options in these guidelines, so the program chair will report which options were selected in the chair's message that appears in the proceedings. To ensure that the guidelines are part of a tradition that is not forgotten, ACM will include them as part of the conference proceedings of the International Symposium on Computer Architecture (and thereby appear in Computer Architecture News once a year).

1. Quality vs. Balance

Papers will be judged on their scientific merit and anticipated interest to conference attendees. Balance means that the paper should be accepted for some other reasons than quality as judged by reviewers. Examples of non-quality issues that have risen are:

- Too many papers from one author;
- Too many papers from one institution;
- Too many papers from one country;
- Too many papers on one topic.

The policy of SIGARCH is that quality alone should be the guide on the first three issues. (It is understood that papers in new areas are likely to contain less quantitative evaluations and comparisons than those in more established areas.)

The last issue is a very sensitive one, for the program committee must be extremely careful in limiting the number of high quality submissions even if all are in a single area. In particular, the program committee should keep in mind that impact and relevance are an important component of quality. SIGARCH has played a vital role in improving the architecture of the commercial computer industry, and it is important to the field for that role to continue. Balance at the cost of impact and relevance must be avoided. (Discussions on statistics of balance must be delayed until after papers have been evaluated, if such statistics are discussed at all.)

One non-quality issue for a conference is papers likely to stimulate interesting and productive conversations at the conference. This is a perfectly acceptable reason to include a paper in the program.

2. Program Chair Selection

The program chair is key to the success of the technical papers, and hence the decision is critical to the success of the conference. Most SIGs separate the decision of the location/general chair of a conference (which needs to be made 3 or 4 years in advance) from the selection of the program chair (which needs to be selected no more than 2 years in advance). We recommend following this tradition, as career changes may make it difficult for program chair candidates to make or keep their commitments. (The decision should be made early enough so that the chair of the succeeding conference can be an unofficial member of the current program committee.)

3. Program Committee

The program chair has the dual goals of a high quality committee and finding representatives from several communities. Hence diversity is important in selection of the program committee (PC). This policy encourages submissions from many groups plus offers direct access to different groups of competent reviewers. The program chair should consider the following when selecting a PC:

- Personal qualities: judgment, reliability, ethics, standards;
- Coverage: technical areas, leading universities, leading companies;
- Balance: technical areas, geographic, academe vs. industry, youth vs. experience, and so on;
- Planning for the future: grooming future program chairs.

There should be significant turnover as well as some overlap between successive PC memberships.

Before picking a committee, the program chair should consult with prior program chairs for input along these four issues. It will be helpful if the program chair of the succeeding conference is invited as an ex-officio member of the program committee and participate at the PC meeting.

4. Reviewing Papers

At least two members of the PC must review each paper themselves in addition to external reviews. We encourage the recent tradition of blind reviews, asking the authors to blank out the names on self-identifying references in their submissions.

5. Program Committee Meeting

There must be a PC meeting where all PC members should be in attendance. Assurances should be made by the PC member that he or she will attend before agreeing to serve. If the makeup of the committee prevents the vast majority of the members being at a single meeting, then the recommendations of multiple meetings must be resolved at a single meeting containing many representatives of each meeting.

Members of the committee must not participate during the discussion of papers in which they have a conflict of interest. SIGARCH uses the National Science Foundation definition of conflict of interest, which will be distributed to the PC. (See NSF Policy on reviewer conflict of interest which recommends selecting reviewers who are not at the same institution, are not former PhD advisors/advisees or postdoc advisors/advisees, and are not known to be personal friends or antagonists of the authors.)

In consultation with the PC, the chair either:

- (1) Keeps the PC meeting blind, with papers unidentified; or
- (2) Identifies the authors so the PC meeting is not blind.

6. Papers by Program Committee Members

The first issue is the threshold of acceptance for program committee papers. The three policies are:

- (1) The quality of the PC papers must be at least as high as the threshold of the of acceptance of regular papers. (This encourages participation of active members of the field in the PC while being fair to the non-committee papers.)
- (2) The quality of the PC papers must be much higher than the threshold of the of acceptance of regular papers. (This reduces the chances of the appearance of bias.)
- (3) PC members cannot submit papers.

The program chair must pick the model and inform PC members before they are asked to join.

A second issue to address is the number of papers that a member of the PC can co-author. In return of the honor of being on the PC, the default policy should be no more than two papers submitted per PC member. (This addresses the appearance of impropriety of too many papers by a single PC author at the conference.)

The third issue is the method of acceptance PC papers. This is not a simple problem. The issues are being fair to PC papers, being fair (and looking fair) to papers that are not co-authored by the PC, and not having awkward situations or peer pressure at the PC meeting. The common points:

- The program chair and general chair are not authors or co-authors of any submissions.
- Submissions are divided into PC and non-PC papers.
- The whole committee decides the fate of the non-PC papers.

There are again three options:

- (1) "After regular papers" model: The PC authors are thanked and dismissed after the non-PC papers have been selected. The non-submitting members then decide the fate of the remaining (PC) papers.
- (2) "Before regular papers" model: The program chair sets up an informal committee, possibly chaired by a prior program chair or the general chair, and the fate of the PC papers is decided before the PC meeting. The meeting starts with the list of accepted PC papers. (This option makes it difficult to apply the non-PC threshold to PC papers.)
- (3) "Hot seat" model: Only the author is excused when a PC paper is evaluated by the PC. (This can be embarrassing, plus it is not blind.) A variation of this model is PC members are asked to leave the room whenever *any* paper from their institutions is discussed, which reduces embarrassment since it is unclear who the author. It also shields PC members from questions by colleagues about what happened to their paper.