

# Advice for Program Chairs

Alex Aiken

October 3, 2010

## 1 Introduction

The title says it all: This document is written for people who are, will be, or think they might want to be the program chair for a technical conference. I am a former program chair for POPL (the ACM Symposium on Principles of Programming Languages) and PLDI (the ACM Conference on Programming Language Design and Implementation), both CS programming languages conferences. Given my limited experience, my comments may be most useful to people chairing these particular conferences and next to chairs of other computer science conferences. Program chairs for conferences outside of computer science should understand that computer scientists treat conferences very differently from other disciplines, and that may reduce the relevance of at least some of the points discussed here.

Currently there is little written on the general issues facing a program chair. There are a few (very good) summaries with recommendations [Hil05, McK05, Eis07], and also some longer articles covering specific aspects of the process [FC04, Wad06, Fel06].

A program chair has a lot to do, and I will not attempt to cover everything. For example, you will need to produce a Call for Papers, but unless radical changes to the conference's focus are contemplated, the time-tested method of updating last year's call will suffice. The intention here is to discuss less obvious aspects of being a program chair.

## 2 Goals

A program chair manages a conference's program committee, which chooses the program for the conference. A program committee has two goals:

- choose the best papers for the conference, and

- give critical feedback (reviews) to the authors.

It is important to understand that these goals are different and, in a process where resources are always limited, sometimes in conflict. Giving feedback to authors is about doing the best job of understanding and critiquing a particular paper; selecting the best program is about comparing different papers to decide which papers to accept. The effort required to understand papers well enough to decide whether they should be in the program is generally less (and sometimes much less) than the effort required to produce good reviews. Long ago the emphasis was primarily on choosing the best papers and authors frequently received only cursory feedback. As the role of conferences in computer science has become more important, the trend has been to have more reviews and to emphasize thorough reviewing.

At many points in deciding a conference's program the program chair makes decisions about how to allocate resources among the submitted papers:

- Has a paper been reviewed well enough to decide whether to include it in the program?
- Do the reviews for a paper give adequate feedback to the authors? Are there enough reviews?
- Has a paper been discussed enough? More time spent on one paper means less time spent on other papers.

Too many "no" answers to any of these questions in the final, time-constrained phases of the program committee process leads to serious problems. Reviews take time and expert people and are very difficult or impossible to do at the last minute. Trying to compensate for inadequate reviewing of more than a very small number of papers while simultaneously deciding on the program detracts from the time available to discuss the bulk of the papers. In the worst case the committee simply runs out of time before all the papers are adequately discussed and the final decisions are made with some papers receiving very little attention.

It is the program chair's responsibility to avoid such a disaster. Time management during the program committee meeting is important, but decisions made earlier in the process are probably even more important. This document is mostly about what can be done at each step of the process to minimize the problems that arise when the crunch comes at the end.

### 3 General Issues

There are a few points that are probably self-evident but worth stating explicitly:

- *The General Chair.* Conferences usually have separate general and program chairs. The general chair is responsible for the whole conference: the hotel, the food, excursions, affiliated workshops, etc. The program chair is responsible for the technical program. A good working relationship with the general chair is very important, because you will be working together a lot.
- *Conference Software.* There are several free and commercial software packages for managing conference submissions and reviewing; use one. I've used HotCRP [Koh06] and was happy with it. (The HotCRP web site, given in the bibliography, includes links to many other conference management systems.) Whatever system you use, make sure the database of submissions and reviews is backed up regularly, and double-check your work when using the notification system—it can be surprisingly easy to send messages to the wrong set of people, possibly revealing information they were not supposed to know.<sup>1</sup>
- *Slack.* You will be dealing with a lot of people, perhaps dozens of program committee members and external reviewers and hundreds of authors. With so many people involved, it is a certainty that there will be some problems, no matter how careful you are. Reviewers will get sick or forget to do their reviews, investigations will have to be conducted to decide whether authors bent or broke the submission rules, and the computer hosting the submission/review site will crash at inopportune times. Some slack needs to be included in both the program committee's schedule and your personal schedule to deal with such issues. It is probably not a good idea to be a program chair unless you can make it your highest professional priority, at least during the critical times around paper submission and the weeks leading up to the program committee meeting.

---

<sup>1</sup>For SIGPLAN conference chairs, Phil Wadler has posted some of the key mass notices he used for POPL on the POPL wiki <http://wiki.acm.org/sigplan.popl> (login required).

## 4 The Program Committee

Choosing the program committee is one of the first responsibilities of a program chair. Typically, the program chair proposes a program committee which is approved by the general chair and the conference's steering committee (if the conference has a steering committee). Program committees are rarely if ever rubber-stamped. General chairs and steering committees can and do object to people or combinations of people and suggest additional members. Often the conference's governing organization has specific criteria for program committees such as desired gender, seniority, and geographical distributions. Other typical considerations are the overall size of the committee, the proposed committee's expertise in conference topics, and whether specific people have previously served well, poorly, or too frequently on other committees. The general chair/steering committee will be some of the most senior and experienced people in a conference's field and they usually know what they are talking about; it is important to listen to their advice.

Thus, the program committee is selected in consultation with others, and you can expect this consultation to be active and to take some time. Issuing invitations and getting responses from all of the invited program committee members will take even more time. So, whatever you do, start early! The earlier you ask someone to serve on the program committee, the less likely they are to have a conflict and the more likely they are to say "yes". A year in advance of the program committee meeting is not too early to begin working on the program committee.

A few of the finer points of program committee selection are listed below (in no particular order):

- *Covering Topics.* It is important to cover all of the topics on which submissions can be expected. The papers published in the conference in the previous 2-3 years are a good guide, but because conference interests evolve fairly quickly using more history than that can be misleading. If your conference favors small program committees you will need to look for at least some people who are qualified to review papers in more than one topic.
- *Covering Small Topics.* Watch out for specialties that always have submissions and also have a relatively small number of researchers who often collaborate among themselves. The chances are high that any program committee member covering that area will have a conflict

with one or more of the submissions; you will need at least two program committee members who can reliably cover that topic.

- *Backups.* Some people (hopefully very few!) will decline to serve on the program committee, so you must be ready to issue an additional invitation when someone says “no”. Save valuable time by having your list of backups pre-approved by the general chair and steering committee.
- *The Invitation.* Your invitation to the program committee should include the schedule (the dates of major deadlines, see below), the expected workload (in number of papers to review), and whether program committee members are expected to attend a program committee meeting. It is important that program committee members know what they are signing up for. Having clear expectations from the beginning also reduces the chances for problems with the committee.
- *External Reviewers.* Many conferences use external reviewers for one or more reviews of each paper. The idea is to have at least one expert outsider review each paper, to get a broader perspective than the committee can provide and to provide a calibration check on the committee process. Consider inviting the external reviewers early, too, with an invitation similar to the program committee invite but asking only for a commitment to review a few (say 4-5 at most) papers. If you are going to use external reviewers, it is much better to line them up early.<sup>2</sup> The alternative is to spend precious time after the papers are submitted trying to find people who can do external reviews right away.
- *Should you review papers?* There are two basic styles of program chair. There is the chair who is a super-sized program committee member. This chair reviews papers just like other committee members (and perhaps reviews many more papers than the rest of the committee) and also does the other jobs (picking and inviting the committee, managing the review process, etc.). Then there is the program-chair-as-manager, who manages the program committee but does not review papers.

I prefer the manager-style program chair. There are some critical tasks that only the program chair can do, and it is important that the program chair have both the time and focus to do them well. If you

---

<sup>2</sup>PLDI, for example, has formalized the use of external reviewers in the “External Review Committee”, which is recruited at the same time as the program committee.

are worried that there won't be enough reviewing horsepower without you, add one more person to the committee to read your share—they will likely do a more thorough job than you would be able to do.

It must be noted that program chairs have varying opinions about how much reviewing is appropriate or necessary for the program chair to do, and the two extremes described above are not the only possibilities. Some review all of the controversial papers, some review all of the papers that might be accepted or discussed. And even if you try to be a manager-style chair, as the reviewer of last resort you should expect to end up reading at least a few papers.

For any kind of program chair, the program committee should make all or almost all of the decisions about what papers to accept. Let them; you pick the committee, but the committee's job is to pick to the program.

## 5 Program Committee Submissions

Conferences struggle with how program committee submissions should be handled. Some conferences ban submissions by the program committee, which simplifies life for the program chair, but can make it more difficult to recruit the program committee. Allowing program committee submissions, however, creates the potential for serious conflicts of interest in deciding whether to accept those submissions.

The traditional method for handling program committee submissions is to have only the subset of the committee that did not submit papers review and decide on the program committee papers. To avoid any appearance of favoritism, these papers are often held to a higher standard than the general pool of submissions. A typical rule is that there must be consensus that an accepted program committee paper is in the top half of the accepted papers; another (usually stricter) rule is that only clear accepts are taken (e.g., all reviews are at least mildly positive, and at least one reviewer is positive enough to champion the paper for acceptance). Of course, all such decisions are made separately from the regular program committee process, either with all submitting committee members out of the room or in a separate electronic discussion. I have served on committees that used this process and I think it works well, but it is not perfect. There are many objections that people raise, but the one that probably has the most merit is that bias can still creep in because the program committee is working closely together, and there is ample opportunity for committee members who submit papers

to discuss their current work with other committee members who did not submit, discussions that could easily cross the line into lobbying. Another issue is that with reviewers and submitters both on the program committee, maintaining anonymous reviews becomes more difficult for the program chair.

Another approach (which I used for PLDI'10) is to use only external reviewers to review program committee submissions. That is, no one on the program committee has anything to do with reviewing any program committee paper, and from the program committee member's point of view, it is as if they submitted to another conference. With no possibility for a flow of information between the reviewers and the authors, this approach has the advantage that no one can object that there is a potential conflict of interest. It does require a little more work from the program chair, mostly in recruiting sufficient outside reviewers to handle the program committee submissions, but if you are using external reviewers anyway, this is not a big difference. The external reviewers for program committee papers should mostly be experienced people who have served on similar committees in the past. I would also recommend that a program committee submission be accepted only if someone champions it and everyone else feels the paper is at least decent (say, not in the bottom 10% of accepted papers)—unlike regular papers where it may be worthwhile to take some risks, it is probably best not to accept controversial program committee submissions.

As a rule, program chairs and general chairs may not submit papers, and even if you are allowed to submit I would recommend that you do not. It is very awkward for the program committee to judge the program chair's paper.

## 6 Double-Blind Review

Double-blind review, where reviewers do not know the names and institutions of the authors of submitted papers, has become more popular in computer science conferences in recent years. The argument for double-blind review is that knowing who wrote the paper and where they are from prejudices a referee's opinion, if only subconsciously. By withholding this information, papers are judged more fairly. While it may do some good, the benefit of double-blind reviewing is probably at the margins (in my opinion), and it definitely imposes additional burdens on the reviewing process, and particularly on the program chair. That said, many people (including some reviewers) like double-blind reviewing. Of the arguments I've heard, the

one I like best is that not knowing who wrote the paper allows a reviewer to quickly reject a bad paper by a famous person without feeling guilty that they didn't spend more time trying to decipher it.

If your conference is double-blind, there are a number of additional issues that you will need to deal with:

- Issue clear rules for authors. Be precise about what information authors can and cannot include in the submission. Authors naturally do not want to violate the submission rules, and unclear rules leave them guessing about what is permitted and what is not, potentially wasting their time and damaging their submission. I recommend very simple rules: no authors and institutions on the title page, and all of the authors' prior work referred to in the third person.
- Conversely, don't let authors use double-blind reviewing as a shield. That is, don't set up a situation in which authors can withhold relevant information from the committee. For example, it is not acceptable to omit closely related work by the same author simply because it might cause a reviewer to guess the author's identity. An explicit statement in the call for papers that no information important to reviewing the paper should be withheld in the name of anonymity helps. If nothing else, this point makes clear that later claims (say in the author feedback) that the double-blind rules prevented the authors from making the best case for their paper will not be looked upon favorably.
- Reveal author identities after the reviews are in. Even many (but certainly not all!) advocates of double-blind reviewing believe it is necessary at some point in the process to let the reviewers know who wrote the papers. The author feedback and discussion of a paper will be much more straightforward if this information is out in the open. As a concrete example of why it is a good idea to disclose identities before final decisions are made, consider a paper that proposes a system X, argued to be much better than a previous system Y but with no head-to-head comparison. System Y was never made publicly available, so the lack of direct experimental comparison seems unavoidable. However, if the same authors wrote both X and Y, the standard is different! Thus, there are situations in which knowing who wrote a paper legitimately affects how that paper is evaluated. I believe the right goal for double-blind reviewing is not to hide the identities forever, but rather to hide them long enough that they don't affect a reviewer's first impression of a paper.



- Conflicts of interest are harder to sort out with double-blind reviewing. The usual procedure is for the authors to mark the reviewers with which they have a conflict. Note that this opens the door to authors blackballing reviewers that they simply don't want to review the submission (and this does happen). There will also be the usual problem that some authors simply don't fill out the conflict list or fill it out incorrectly. As a check, give the program committee and external reviewers a complete list of all the authors submitting to the conference and ask them to mark the ones they have conflicts with. You can then compare the two sets of conflicts and investigate any discrepancies.

## 7 The Schedule

In developing the reviewing schedule for a conference, perhaps the major concern is avoiding a time crunch. No matter how conscientious you and the program committee are, no matter how carefully considered the reviewing and decision-making procedures, it will all end in a train wreck if there is insufficient time. On the other hand, the whole point of a conference is to showcase the latest research results, and so it is desirable to keep the time between paper submission and the actual conference as short as possible. Here is one typical schedule (the one that I used for PLDI'10), phrased in terms of number of days for each stage:

- Day 0 Abstracts due.
- Day 7 Papers and program committee bids due.
- Day 12 Paper assignments given to PC members and external reviewers; reviewing begins.
- Day 59 Reviews due.
- Day 61-63 Author feedback period.
- Day 64 Start of electronic discussion.
- Day 76-77 Program committee meeting.
- Day 80 Author notification.
- Day 130 Final papers due.

The following subsections briefly discuss these phases.

## 7.1 Abstracts and Program Committee Bids

The idea behind having separate abstract and paper submissions is that in the week before papers are due the program committee can already have a look at the abstracts and express preferences (“bids”) for the papers they would like to referee. When the submission deadline finally arrives the program chair is then in a good position to quickly make paper assignments and get the reviewing process underway.

That is the theory. In practice, the abstracts often represent authors’ hopes and dreams rather than actual papers they will submit. Typically more than 1/4 of the abstracts don’t turn into papers, which means that some significant fraction of the work in bidding on abstracts is wasted and the bids on the papers that are submitted may be skewed. As a result, it is not entirely clear whether having a separate abstract deadline is worthwhile.

## 7.2 Paper Assignments

After the selection of the program committee, the program chair’s next major task, is assigning reviewers to papers. All of the conference management tools come with a feature to do paper assignments. Authors list categories describing their papers, program committee members and external reviewers express preferences for general categories and/or specific papers, the program chair presses a button, and out pops a solution to the problem of matching papers to referees.

Unfortunately, automatic paper assignment does not seem to work well and I recommend that you not use it. Despite the work involved, an assignment done by hand will be much, *much* better (in the sense of doing a better job of matching reviewer expertise to papers) and save time in the end. How can this be? There is nothing wrong with the matching algorithms themselves; they work as advertised. The problem is the data, the author and program committee preferences, that is the input to the matching algorithms. Some authors give detailed descriptions of their paper in whatever format is requested, others give no input at all, and still others give clearly incorrect information (presumably mistakes or user interface issues). Similarly, some program committee members may make detailed bids, some just enter very general blanket preferences, some enter nothing, and some actively bid against papers in their area of expertise (because they are tired of the subject) or bid for papers in areas that they know nothing about (because they want to learn something new). Enough data is missing or wrong that the automatic paper assignments are often very strange. Personally, I

have found automatic assignments unusable (and recall that in general I am a fan of the conference management tools).

To assign papers to referees by hand is a significant job; it takes a few whole days for a conference with about 200 submissions (notice the time in the schedule between the submission deadline and when paper assignments are given out). The authors' descriptions and reviewers' expressed preferences (when they exist) are still useful guidelines but in some cases may be contrary to obtaining the most expert and deepest reviews. The advantage of doing assignments by hand, then, is much better matching: the reviewers' expertise is well-matched to the papers they are reviewing and results in higher confidence, relatively expert reviews. In practice, if a good assignment of papers is done initially, relatively few papers will need additional reviews later. Minimizing the number of urgent extra reviews required in later stages is crucial, as there just isn't much time left after the reviewing deadline passes and the discussion about which papers to accept begins.

There are further simple steps you can take to try to ensure papers are well-reviewed the first time. Once the reviewing assignments are made, you can ask the reviewers to tell you within a short period of time whether they will be able to provide an expert review for the paper (note the question is not whether they want to review the paper, but whether if they review it they think they will rate the review as "expert"). At this stage, before any reviews have actually been done, a paper with fewer than two expected expert reviews is a potential problem. Such papers can either be reassigned, or you can recruit more external reviewers.

### 7.3 Reviewing and Re-reviewing

Too much analysis goes into the very fine points of how to score papers. There are many scoring systems that work just fine as long as everyone understands what the scores are supposed to mean. That said, I use the "ABCD" scoring method: "A" means a reviewer will champion the paper (would argue to see it accepted), "B" is a weak accept, "C" is a weak reject, and "D" means the reviewer will argue to reject the paper. This system has two advantages. First, there are few categories. Given that the ultimate outcome is a binary decision (accept or reject the paper), finer gradations don't add much. Second, there is no neutral, middle score: referees are forced to at least think about and declare a preference, if not a strong one.

One of the advantages of using conference management software is that the program chair gets early notice if someone is having trouble with their reviews. If someone is far behind or their reviews do not seem up to standard,

you should contact them to inquire how things are going while there is still time left before the review deadline. It is also a good idea to leave a couple of days of slack in the schedule at the end of the reviewing period for stragglers.

Once all the reviews are in for a paper (and if you encourage your committee to submit their reviews early, for many papers the reviews will be in before the reviewing deadline), the program chair can evaluate whether the paper needs further attention:

- If a paper has no expert reviews (whether those reviews are positive, negative, or indifferent) there is increased risk of making a bad decision. There is no good alternative but to try to get more, and expert, reviews on short notice.

Clearly, the sooner such papers are identified the better, and a good initial paper assignment is the only defense against having many papers in this category.

- An exception must be made for the rare paper that can never have expert reviews. If a paper's topic is truly new, or if a paper combines expertise from very different areas, expert reviewers may simply not exist.
- Typically electronic discussion of the papers begins when all of the reviews are in. Many disagreements about particular technical points are resolved here and the reviewers reach consensus, saving time at the program committee meeting. Conversely, discussions that end in strong disagreements signal a paper that may need another reviewer (see below). These benefits are lost if the discussion doesn't happen; as program chair, it is important to keep an eye on and encourage discussion of the papers that need it.
- A controversial paper (some reviews are quite positive and some are quite negative) may or may not require additional reviews. If the reviewers have understood the paper and at least some are expert, then it is unlikely that an additional review will bring up new issues. If the reviewers can't agree on what the paper is about or what the technical problems are, another review may help to get all the issues out.
- A few reviews will be badly done or possibly never done at all. If the missing review may make a difference to the outcome, another review is needed. If the outcome is clear without the review, the program chair

must decide whether the feedback in the other reviews is adequate or more reviews are still required.

- Watch out for program committee members who are overly reliant on subreviewers. A committee member who simply passes along a report from someone else and is unable to articulate an opinion of their own is not doing their job. Be sure your expectations on this point are clear to the committee from the beginning.

Notice that even if the program chair does not read papers, he or she will generally end up reading many of the reviews for the papers, particularly the under-reviewed, borderline and controversial papers. The 12 days of electronic discussion in the sample schedule allows time for the program chair and the program committee to identify controversial and inadequately reviewed papers, and to ask for and receive additional reviews prior to the program committee meeting. This is also the period in which the papers to discuss at the program committee meeting are identified (if there is a separate, physical committee meeting). Typically, papers to discuss include any paper with a champion (if using “ABCD” scoring, these are the papers with at least one “A” score), any paper someone on the committee wants to discuss, and some selection of the papers with no strong detractors (no “D” scores; for example, papers with all “B” scores, or more “B” scores than “C” scores).

## **7.4 The Program Committee Meeting**

Time is the critical resource at an in-person program committee meeting. Here are a few tips on making the most of the time:

- Do in advance what can be done in advance. Guest wireless access, projector, power strips, and snacks should all be arranged beforehand.
- Procedurally, using three categories works well: accepted, rejected, and tabled. If a paper discussion goes on for a while without consensus don’t be afraid to table it. Often decisions will be made very quickly when a tabled paper is revisited.
- A key question is: How many papers should be discussed at the meeting? There is not enough time to discuss all the papers, but discussing very few papers thoroughly means the program committee does not make very many decisions, or rather, a lot of the decisions have already been made before the meeting. I suggest discussing about twice

as many papers as you plan to accept; the meeting makes substantial decisions, and the pool is large enough that it is very unlikely any paper that could be accepted is excluded.

- Discuss papers in a random or at least arbitrary order. At the start of the meeting the committee will be reluctant to make decisions because they do not yet have a sense of where the quality threshold lies. By discussing a representative sample of papers at the start, the group will develop a sense of the overall quality of the pool much more quickly than if papers are discussed, say, ordered by initial scores. This saves time by helping the committee reach the point where it can make firm decisions sooner than it would otherwise.
- Allow the discussion some slack at the beginning of the meeting while people are still feeling their way. After the first 10-15 papers the pace of decision making should increase.
- Don't pay too much attention to scores. Scores help to decide which papers to discuss and who should lead the discussion, but scores are noisy and once a paper is being discussed it is the technical points, not the letter grades, that should determine the outcome.
- Making decisions is the purpose of the meeting. It is a mistake to discuss a paper further once the decision is obvious. There will be time for delving into the cool details of the accepted papers and the fatal flaws of the rejected papers later, after the meeting is over.
- People who have a conflict with a paper leave the room while that paper is discussed, and people who were out return if they have no conflict with the next paper. Notwithstanding the recommendation to randomize paper discussion order, when determining the order look for opportunities to group a few papers together with the same or similar conflicts. Minimizing the traffic in and out of the room saves time.
- Before the meeting, designate someone on the program committee to manage the discussion of any papers with which you have a conflict.
- The ultimate sin is to not have adequate time to discuss the papers; you definitely want to avoid this! Some program chairs enforce fairness by allocating a number of minutes to each paper. However, not all discussions are created equal. Some papers will take almost no time and others may require multiples of the average time to reach agreement.

I prefer not limiting the discussion time of papers in advance, but this requires being willing to table discussions that aren't going anywhere for the moment and ensuring that the average rate of progress through the papers is still sufficient.

## 7.5 The Electronic Meeting

Electronic meetings are similar to physical meetings in that there is still a limited amount of attention that must be divided among the papers in a way that ensures every paper that could potentially be accepted receives enough discussion. Thus, it is still useful in electronic meetings to focus on a relatively small subset of the submissions that is nevertheless considerably larger than the anticipated number of accepted papers. The program chair also should recognize when it is time to end a discussion and either accept or reject a paper so that program committee members can devote more effort to other papers.

The main difficulty with electronic meetings is ensuring that all the discussions that need to happen actually happen. It is up to the program chair to check up on each and every paper under discussion to see that it is actually being discussed, and to try to think of ways to help if a discussion becomes sidetracked or stuck. Most discussions will spontaneously take place and resolve themselves, but there are always some that languish unless the program chair notices and reminds those involved that a decision on the paper will be needed shortly.

In physical meetings, comparisons between papers and contributions from other members of the program committee listening to a discussion will happen naturally. In an electronic meeting, these interactions generally will not happen unless the program chair notices that it would be useful to have additional people involved in a discussion.

## 8 The Aftermath

Once the program committee meeting is over, everyone breathes a sigh of relief. The committee's job is done, but the program chair still has a long way to go. The crisis, however, has passed, and there should be plenty of time to deal with the remaining tasks: organizing the program into sessions, recruiting session chairs from the program committee, alerting authors to the final paper deadline and instructions, having plaques or certificates made for any award papers, working with the general chair on the final details of

the conference, and so on. You will find the list is long, but not really difficult.

You can expect at least a little complaining from authors or committee members at some point in the process. Try not to take it personally; people usually complain because they care about the conference and the papers selected, and the process can be stressful.

Eventually the big day finally comes and the conference actually begins! You will likely be busy with last minute details, but try to take a moment to enjoy the conference you helped put together.

## 9 Acknowledgments

Thanks go to Phil Wadler for suggesting that I write this document, and to fellow program chairs Kathleen Fisher and Mooly Sagiv for their comments. David Gay suggested the idea of having reviewers identify the authors with which they have a conflict in double-blind review conferences.

## References

- [Eis07] J. Eisner. How to serve as program chair of a conference. <http://www.cs.jhu.edu/~jason/advice/how-to-chair-a-conference.html>, 2007.
- [FC04] K. Fisher and C. Chambers. Experiences with author response at pldi and icfp 2004. *SIGPLAN Notices*, 39(12):1–6, December 2004.
- [Fel06] M. Felleisen. Time is money: A report from the POPL 2007 program chair. *SIGPLAN Notices*, 41(12):2–4, December 2006.
- [Hil05] M. Hill. Some advice for program committee chairs based on my isca 2005 experience. <http://pages.cs.wisc.edu/~markhill/pcadvice2005.html>, 2005.
- [Koh06] E. Kohler. Hotcrp conference management software. <http://www.cs.ucla.edu/~kohler/hotcrp/>, 2006.
- [McK05] K. McKinley. Notes on chairing program committees. <http://www.cs.utexas.edu/users/mckinley/notes/PC.html>, 2005.
- [Wad06] P. Wadler. Conferences with improved management and style. *SIGPLAN Notices*, 41(2):1–3, February 2006.