



PC Chair Report

Robby Findler
Northwestern University

This region of the pdf contains comments
that (roughly) correspond to what was said
during the talk.

2012 SIGPLAN Awards

Matthias Felleisen

Achievement Award

Shriram Kirshnamurthi

Robin Milner Young Researcher Award

Jikes RVM

Software System Award

Jens Palsberg

Distinguished Service Award


Dan Marino


Student Dissertation Award


Lets start by taking a moment to recognize the winners of this year's SIGPLAN awards. I encourage you to check out the work represented here, but also to keep in mind that these awards exist and to consider nominating worthy awardees for the coming years.

The following slides are the obligatory information from a chair's report showing basic stats from this year and previous years about attendance submission and acceptance rate.

More on these specific numbers later.









This shows the geographic range of the submitting authors. With the exception of .com, .org, and .net email addresses, I used the domain of the email address to determine the country. Within those domains, organizations and companies that are restricted to just one country were all given that country and otherwise I inspected the affiliation of the submitter directly. Duplicate email addresses were discarded. I was unable to determine where one of the submitters is from; that person is not represented on this slide (or the next).


The Australasia and West Asia regions contain one country each: Australia and India respectively




Most papers had three reviewers (the three PC members assigned the paper, generally speaking), but a substantial number had more than three.



Every paper had at least one expert.





The 7 papers who had only two PC member reviews had external reviewers who were selected specifically by a PC member who was originally assigned the paper. In each case, the PC member read the paper, agreed with what the external reviewer said (often consulting them in person), and was held responsible for being able to articulate their opinion about the paper and discuss it intelligently in the PC meeting. All were able to do this.

[illegible]

3

This year I insisted that there be no submissions by PC members. I was quite happy to avoid that particular distraction, but this was the cost: 3 people turned down my request to serve on the PC because of this rule. (I got reasons from all who turned me down.)

Amusingly, none of these three ended up submitting and one who turned me down citing too much work ended up submitting.

Overall, I think it this number is nearly meaningless; specifically it is unlikely to tell us if a no-PC-submission rule makes it hard to effectively recruit a PC.

87

This is the number of submissions and it is lower than I had hoped it would be. I do not know why, but the record attendance gives me hope that this is not an indication of disaster.

Let me pause a moment here to point out that there is a meme going around that FP is useful. People are recognizing it works better with recent(ish) developments in software process management, e.g., test-driven design. People are realizing the benefits for parallelism with a functional-style of programming, even in decidedly imperative PLs. We have a chance to connect to the larger world and share the wonder and joy we all experienced with FP.

Lets not miss this chance. FP will probably take over the world. With or without us.

Instead of the conventional A/B/C/D
“identify the champion” approach to paper
ranking, I consulted with some economists
and algorithm mechanism design people to
try to understand if there was a better set
of rules that I might use for paper
assignments.



Paper ranking as an economic game

Idea: trade off information quality to combat bias

Mechanics:

- All reviewers had a pool of \$50 to spend on papers
- More money \Rightarrow better paper, in the reviewers opinion
- All money must be spent (the system enforced this)

Thus, reviewers all have the same average score

The Bias: The 😊 / 😞 reviewer

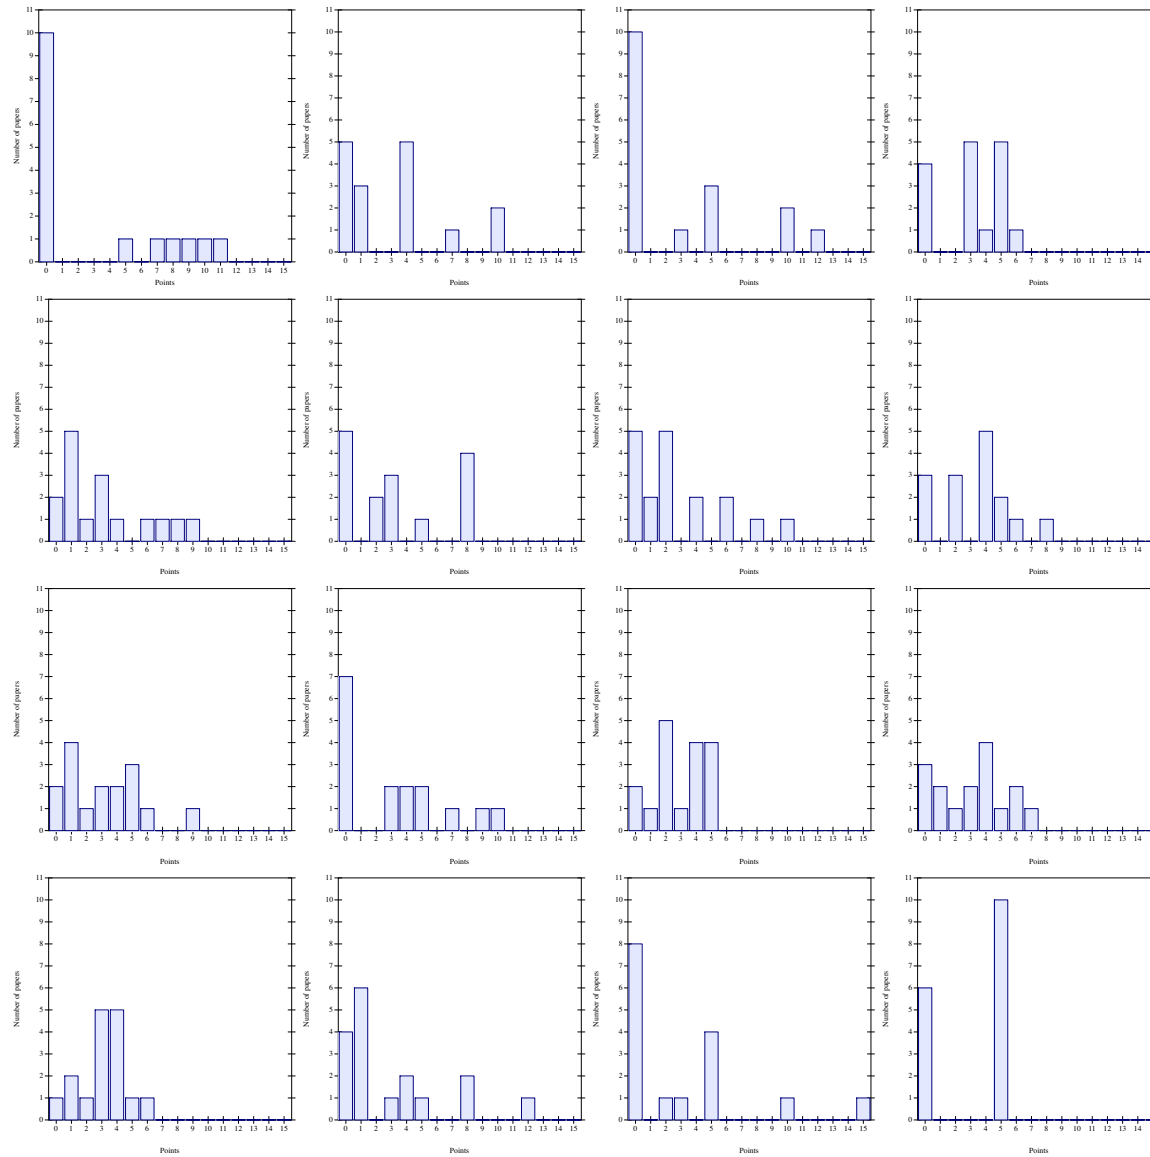
The FAQ: But what if I get a
extraodinarily good/bad
pile of papers?

Bonus: reviewers are more likely to revisit papers

These are two sides of the same coin. What do you see when you sit on these committees? Do you see people that are overly negative/positive, or do you see that your submission piles are more uniformly great/horrible?

My feeling is that the former is a more common occurrence than the latter, but I have not conducted a randomized control trial (and I think it would be unwise to attempt one).


That question aside, I believe this mechanism had an additional positive impact because it forced reviewers to revisit papers (to reallocate funds) and, in my experience, a second look at a paper is typically far more accurate than the first look.



These are the histograms of the scores for the individual reviewers. Each graph represents a single PC member. Each column in a graph represents a score, and the height of the column shows the number of times that reviewer gave that score.


As you can see, people generally played along with the game, using their money in interesting ways.

Two histograms stand out (to me), that I've put in the upper left and bottom right corners of this slide. In the bottom right, you see a reviewer that just gave binary decisions to the papers. In the upper right, this is someone that actually ordered all of the papers that they were willing to accept. These are the least and most informative reviewers (in the economic game sense, not in the content of their reviews!). Most people did something in between.



Each bubble on this slide corresponds to a submission. Reading from left to right and then top to bottom, the order is the discussion order, and the color of the bubbles indicates if the paper was ultimately accepted or not.

As you can see, I randomized the order of the bubbles. Rather than going in order from high- to low-ranked papers, which would encourage more and more negativity as the meeting progressed, each paper was more likely to be considered on its own merits, not on its score, (altho we quite naturally seemed to use the score to calibrate expectations more and more as time went on).



These are the same bubbles, but now arranged in score order. Each column is a different score, as labelled at the bottom. The bubble in the bottom right is the paper that I am conflicted with; I don't know its score, only that it is rejected.

Interestingly, despite the random ordering, the score ended up being a pretty good predictor of acceptance. Also quite interestingly, the left-most green bubble and the right-most red bubble were the papers that the PC debated the most, although we did not know, at the time, that these bubbles would have been in those positions.

(I don't want to say more about those papers, for fear of revealing too much, but I believe that the PC made the right decision in both cases, given the information we had about the papers.)

Why are we here?

Alright, one last thought.

I doubt anyone will be surprised at my answer: we are here to improve the quality of functional programming research and, through that, to have some positive benefit to society.

That said, I don't hear us reminding each other of this root motivation much. At all, really. I even get contradictory advice, e.g. when reviewing "Just focus on the good papers in your pile."

I beat the PC members about the head and shoulders with these ideas. I told them when I invited them that their primary job is not to evaluate the submissions, but to provide the experience and wisdom to the community in the form of reviews. Deciding which papers appear here is important, but secondary. I also sent back many reviews, telling reviewers in no uncertain terms that their reviews were not good enough. (I know that some less-than-optimal reviews still went out, but be sure I was not happy about that.)

So, please do keep that in mind. For authors, when you hear your paper is accepted to ICFP, don't think "Good! Now I have three papers at top conferences and so I can defend my dissertation!" That's wrong on many levels. Instead think "Oh no! I'm now representing FP to the world! I better not screw up the final version of my paper!"

This kind of attitude, at least for me, leads to great satisfaction with my work. So let me leave you with a final thought as a quote (next slide).

**“There is nothing better
for people to do than
to eat, drink, and find
satisfaction in their
work.”**

Ecclesiastes 2:24