Message from the Program Co-Chairs

Welcome to OSDI '10, the biggest OSDI yet, with 32 papers selected from an all-time high of 199 submissions. In approaching the task of chairing OSDI, we started with the explicit intention of accepting a larger set of papers, consistent with the growth in the field. Below we outline some of the rationale behind this goal, and the process we applied to achieve it.

Computer systems research is growing as a community. We believe that progress on computer systems research is limited by manpower, not by the limits of a finite domain for interesting research. By implication, as the number of systems researchers increases, the volume of interesting research likely goes up as well. Year after year, top research programs add faculty or research positions in the systems area, while at the same time new programs establish their presence in the field, including newfound growth outside the traditionally strong geographies. The expansion of our community is consistent with the robust scientific and commercial application of computer systems research, providing a strong economic basis for this growth. We believe a larger OSDI program is an appropriate reflection of this growth in the systems community.

We were also motivated by the challenge in making meaningful distinctions, under the pressure of program committee deadlines, between papers that are almost accepted and those almost rejected. The fragility of PC decision process has been documented and discussed elsewhere [A08]. Too often, rejections seem arbitrary in retrospect, hinging on the nuances of a PC discussion rather than clear merit. In accepting more papers we hope to incrementally improve on the fragility of these decisions, while also building a program that is more diverse and therefore of broader interest.

This goal of a larger program was a consideration throughout the review process. The PC was split into two groups: a "heavy" PC who participated in the first two rounds of reviewing, and a "heavier" PC who also reviewed papers in round three and attended a face-to-face meeting to decide final outcomes. In the first round, each paper received two reviews and approximately 35 papers were pruned. To reduce the risk of a premature pruning decision, we allowed reviewers to "rescue" a pruned paper by simply stating their support, with no discussion required. Each round-2 paper received three additional reviews. Another 80 or so papers were pruned after this round. This left us with a pool of 85 papers, each of which received two or three additional reviews in preparation for the PC meeting. After the second and third review rounds, borderline papers were discussed electronically by the reviewers and rejected by consensus of the reviewers.

In the single-day, face-to-face PC meeting each remaining paper was presented by a reviewer, generally an advocate, followed by a time-limited discussion. Based on the first discussion, we binned each paper into one of four categories: "accept," "acceptable," "questionable," and "reject." No rejects were allowed in the first part of the day, the goal of this rule being to avoid the problem of a negative start leading to rejecting good papers early. When all papers had been discussed once, we briefly considered and then accepted the "acceptable" papers as a group, then began the difficult work of reconsidering the "questionable" papers. At the end of the meeting about 30 papers had been accepted.

In the days following the PC meeting, a small set of additional papers were accepted based on an email vote by the heavier PC members. While unusual, we justified this process based on our goal to create a larger and more interesting program, and a sentiment shared by many PC members that the PC discussion had not given due consideration to several of the best liked but most controversial papers. In retrospect we believe these late accepts allowed us to create a stronger and more interesting program, and we would encourage future PC chairs to plan an appropriate process for thoughtful consideration of difficult papers after the bustle of the PC meeting has subsided. For example, even with a single-day PC meeting, it might make sense to put a small set of papers into an "overnight" category, allowing a broader collection of PC members to study them before a final decision the next day.

Apart from the review process, we took some additional measures to try and get more reviews and reviewers in a mindset to accept. We encouraged positivity, following Hill and McKinley's excellent advice [HM05]. We strictly applied conflict-of-interest rules, such that conflicted PC members were not given access to results for conflicted

papers until notifications had been sent to authors. We tried to lighten the PC load from papers that had no chance of acceptance, to leave more quality time for the remaining papers.

Before we close we'd like to briefly acknowledge a few individuals who made a difference in our bringing this program to you. The USENIX staff was fantastic throughout the entire process. We also thank Eddie Kohler for his continued support of HotCRP, a truly wonderful piece of software. We also would like to acknowledge the program committee for their tireless efforts and thoughtful reviews, and Haryadi Gunawi for his detailed note-taking during the PC meeting. Finally, we would like to thank our families and the families of PC members for supporting (and tolerating!) the long hours required to do this kind of work.

Thank you for attending OSDI '10, and have a great conference!

Remzi Arpaci-Dusseau, ***University of Wisconsin, Madison***
Brad Chen, ***Google***
OSDI '10 Program Co-Chairs

REFERENCES

[A08] "Towards a Model of Computer Systems Research," Thomas Anderson, University of Washington. WOWCS '08, April 2008.

[HM05] "Notes on Constructive and Positive Reviewing," Mark Hill, University of Wisconsin—Madison, and Kathryn S McKinley, University of Texas at Austin: http://userweb.cs.utexas.edu/users/mckinley/notes/reviewing.html, May 2005.