

Message from the Program Chair

It has been an honor for me to serve as the Program Chair of MICRO 2015. I am delighted to present to you the technical program, which consists of 61 papers and three keynote speeches.

With the goal of improving the quality of the paper evaluation process, this year we experimented with a revision-based review model, which allowed the authors of top-ranked (approximately 80) submissions to revise their papers during the response period. For facilitating revision, the response period was extended to three weeks. I had presented the preliminary proposal of revision at the business meeting at MICRO-47 to seek suggestions from the community. The feedback was incorporated into a formal proposal, and this proposal was submitted to the steering committee. The steering committee provided valuable suggestions and recommendations and approved the new process. An online document describing the new review process in detail was made available to the community three months before the paper submission deadline. Here, I describe the review process of MICRO 2015 and some of the innovations we tried this year.

Overall Statistics: The 61 papers you will be exposed to were selected from 283 submissions. The Program Committee (PC), consisting of 52 distinguished experts, made this selection, closely aided by an External Review Committee (ERC) of 135 additional experts active in our field. An additional 16 external reviewers aided the evaluation process by providing one review each. We obtained a total of 1512 reviews, with an average of 5.34 reviews per paper. All submissions, except one, received at least 5 reviews. Pradip Bose handled the process for the 28 submissions for which I had a conflict-of-interest.

Submission Format: One of the goals in deciding the submission format for MICRO 2015 was to make the submission format identical to the camera-ready format, with the idea of avoiding the unnecessary time to port the paper from one version to another. To aid this, the submission format was derived from ACM-recommended templates. We found that changing the font-type from the default one could give some authors 10%-15% extra space. Therefore, to ensure fairness for all submissions, it was required that the authors use the default font, and not override it with another space-efficient font. I am glad that the program chairs of future conferences (HPCA 2016 and ISCA 2016) improved and adopted this template.

Format Check for Fairness: Some submissions can grossly violate the format to get more space, which would be unfair to submissions that adhered to the prescribed format. The day after the submission deadline, we did a word count of all the submitted papers (the mean was 8800 and the standard deviation was 1200). All 35 papers that had a word count of more than 10,000 words were manually checked for format violations. We found that 15 submissions had gross format violations, including one paper that had in excess of 14,000 words. All 15 submissions were rejected immediately. However, the authors were given an option to resubmit their paper within 2 days if they agreed to have a full compliance of the required format. The authors were allowed to remove anything from the paper but not add/edit the existing content. Any further violation was ground for rejection without an option for resubmitting. All 15 papers resubmitted a version that complied with the formatting requirements.

Review Assignment: While I was aware of the research expertise of all PC members and most ERC members, I spent the week before the submission deadline to go over the recent work of the PC/ERC members and gathered data about their expertise and recent work. I assigned the reviews myself after carefully examining each submitted paper, reading the abstract and related work, and matching it to the PC/ERC members based on their expertise. Overall, the assignment was done such that each paper received at least 3 PC reviewers, and 2 ERC reviewers. As the robustness of the entire review process relies strongly on proper review assignments, the review assignment was done carefully over seven days.

Obtaining Reviews: We required that reviewers do the reviews themselves without consulting colleagues or graduate students. If consulting an expert was deemed necessary, then the expert had to be approved by the Program Chair, with the input from the expert deemed advisory and the reviewer still responsible for writing the review themselves. I kept an eye on the reviews as they came in, and in few (uncommon) cases requested the reviewers to update their review if the review was short or did not meet the quality expected for MICRO. With only a few reminders, 70% of the requested reviews were submitted by the deadline.

(Friday) and 90% were submitted within two days (Sunday). To aid with emergency reviews, one PC member (Milos) was held in reserve, with no reviews assigned during the original review assignment but agreed to do a dozen reviews during the week after the review deadline. The reserved PC member helped in handling reviewers who were unable to do their reviews, or if an extra review was required for any paper. All submissions, except one, had at least five reviews before the reviews were released to the authors.

Online-Discussion (Before Reviews are Released): This year, we also introduced a one-week online-discussion period between the reviewers so that they can read the reviews of other co-reviewers, discuss in case there is a disagreement between reviewers, and provide a concrete set of concerns/questions to be responded in the author response. We found that the online discussion was quite active, with more than 350 comments posted. The review scores of several papers got changed because of the online discussion. To help steer the discussion, a lead was assigned for each paper that was ranked within the Top 100. The External reviewers were fully involved in the entire review and discussion process.

Deciding Revision vs. Rebuttal: As part of the review process of MICRO 2015, we wanted to provide the authors with the option of revising their papers to better address the concerns of the reviewers. The authors were allowed to modify the paper and also include new results, but only the ones that were specifically asked by the reviewers. However, to reduce the workload of the reviewers, we wanted to limit the option of revision only to papers that are likely to be discussed at the PC meeting. A priori, we had decided that approximately 80 or so papers would be invited for revision. Given this was the first time the revision model was being tried, we wanted the decision for revision versus rebuttal to be a completely objective one, based entirely on paper rank alone, rather than being a subjective one. To rank the papers, we used the metric “*Average Without Lowest (AWOL)*” score, which computes the average by excluding one lowest score. The AWOL metric was chosen so that one outlier harsh review does not degrade the rank of the paper significantly. All papers with an AWOL score of 3.8 or higher were invited for revision. We ensured that all papers on the edge received an equal number (six) of reviews, so the only possible scores around the threshold of 3.8 were 3.6 and 4.0. A total of 82 papers were invited for revision. The remaining 201 papers got an option of responding only with a rebuttal, similar to the conventional model.

Guidelines During Revision Process: To provide the authors with sufficient time to do the revision, the author response period was extended to 3 weeks (July 20-Aug 10). The papers invited for revision were given strict guidelines to (a) not include any unsolicited additions – the revision was meant only for responding to the concerns of the reviewers, and (b) highlight the newly added content in the PDF so that the reviewers can easily identify the new content. The rebuttal space was repurposed both as change log as well as to respond to concerns that did not need revision to the paper (for example, the reviewer simply missed what was already there in the paper). We also communicated that an invitation for rebuttal was not a guarantee for acceptance. While revision was optional, all the 82 papers invited for revision opted for it.

Option of Early Reject: The 201 papers not invited for revision had the option of submitting a rebuttal. For these papers, we clarified that an invitation for revision was not a requirement for acceptance and that a reviewer was allowed to champion any assigned paper to be discussed at the PC meeting. Furthermore, if the rebuttal was effective at improving the scores, the paper still could be discussed at the PC meeting. However, given the ASPLOS deadline was within a few weeks (August 14th), the authors of several papers not invited for revision, requested that their paper be rejected immediately so that they could submit to ASPLOS without violating concurrent submission guidelines. To facilitate this, we introduced an option of “Requesting Early Reject”, whereby the authors can voluntarily request a reject for their paper and terminate the evaluation process immediately. A total of 49 papers opted for Early Reject or Withdrawal.

Reading Revision and Online-Discussion: There was a 10-day period (Aug 11- Aug 20) for the reviewers to read the revision and provide a “post-response” score. If after reading the revised version, the scores did not change, the reviewers were asked to update their reviews indicating how the revision falls short at addressing their concerns. For the papers that got an option for submitting only a rebuttal, the reviewers read the rebuttal and provided their “post-response” scores. They were also given an option of championing any assigned paper for discussion at the PC meeting. For deciding the papers to be discussed at the PC meeting we used the following criterion: (a) All papers with AWOL score based on “post-response” merit score higher than 3.6 (b) All papers invited for revision (regardless of their post-response

scores), and (c) Papers with a champion PC reviewer, regardless of the score. A total of 95 papers were selected for discussion at the PC meeting, including 13 papers that were not invited for revision.

PC Meeting: The Program Committee met at the Georgia Tech Hotel and Conference Center in Atlanta for 1.5 days on August 29-30. All but two of the PC members attended the PC meeting. I am very grateful for the hard work, professionalism, and fairness of the PC members during, before and after the PC meeting. A few important and noteworthy aspects from the PC meeting are described below:

- *Online Accept:* A total of 23 papers were tagged as online-accept, based on the paper having all review scores in the accept range and none of the reviewers objecting to the paper getting tagged as online-accept. These papers were still presented and vetted at the PC meeting, however their discussion was short. We also altered the discussion order to have the online-accept paper be discussed at the end of the first day when the energy level of PC members is low due to the hard work the entire day.
- *Decision Process:* The process for deciding an accept/reject decision for the papers was as follows. If all the PC reviewers are unanimous in their decision, then that decision was considered final, in order to give priority to PC members that read the paper. If there was a disagreement between PC members then the decision went to a PC-wide vote, where the outcome was based on a majority of the PC-wide. We required that at least half of the PC members must vote for the PC-vote to be valid, which helped in having an increased participation (and attention) from the PC members.
- *Introduction of Online Voting:* Manually counting the PC-wide votes by raising hands is time consuming, error-prone, and sometimes prone to bias (as the vote may depend on who else is voting, or the vote count for a “no” may get influenced by the vote counts for the “yes” vote). To avoid this, we introduced electronic voting within HotCRP, whereby PC members were allowed to provide their vote online in a period of ten seconds, after which the vote counts for Yes and No were displayed, along with the list of people who voted for each direction. We found that electronic voting reduced the time in counting, and was received overwhelmingly positively by all of the PC members.
- *Conflict Screen in Hallway:* We also introduced a conflict screen in the hallway that displayed the list of PC members conflicted with the current paper being discussed. This prevented the manual announcement to bring PC members in and to keep some PC members out depending on the changes in conflict list of consecutive papers. Manual announcement is both error prone and time consuming.

Shepherding: The PC decided that 7 papers would be accepted subject to getting an approval from the shepherd. The shepherding process ensured that some critical concerns were addressed in the paper before the papers are published. We are very pleased that all shepherded papers were accepted for publication.

Reflections on Revision-Based Model: The main objectives of the revision model was to (a) improve the quality of decisions in the paper selection process, so that good papers are not rejected because of easily fixable concerns (b) Allow the authors to better respond to the reviewer’s concerns by having a version of the paper where the concern was already addressed, and (c) to improve the quality of the program by having the papers go through a thorough round of revision. Overall, 59 out of the 82 papers invited for revision were accepted, and only 2 papers not invited for revision were accepted. Furthermore, out of the top 45 ranked papers (all of which were invited for revision), 44 were accepted, an indication that the revision model was more helpful in building consensus among the reviewers. The revision model also improved the scores of several papers – out of the 82 papers invited for revision, the score of 42 papers increased by at least 1 point, and the score of 23 papers increased by 2 or more points.

To assess how the community felt about the new process, I surveyed the authors of papers that got invited for revision (before the final accept/reject outcome was announced). Out of the 78 authors who responded, 85% felt that the revision helped them better address the concerns than the rebuttal, and that revision helped improve the quality of their papers. I also surveyed the PC members after the PC meeting. Out of the 36 PC members who responded, 85% felt that revision helped them in making a better decision either absolutely or somewhat. Furthermore, only 15% of the PC members felt that going over the revised version of the paper took significantly more time than rebuttal alone, with 85% responding that it took negligible more time or somewhat more time (but not too bad) compared to a rebuttal alone. Overall, the authors seem to find the revision-based model helpful, and the PC members feel it helps them do a better job and without significant effort. Based on the exercise, I would strongly recommend this model for future conferences.

Acknowledgements:

Coincidentally, I am writing this note over the Thanksgiving weekend in the US, and I have a large number of people I need to thank for their help and support throughout this process. This process, let alone the innovations, would not be possible without them. I am truly grateful for their support.

- I thank the steering committee of MICRO who not only gave me the honor of being the PC chair but also trusted me in trying out the new model. I would especially like to thank Rich Belgard, who helped me throughout the process, and provided valuable advice and guidance.
- I thank Milos Prvulovic, the general chair, for his help and mentorship throughout the PC review process. Milos was always available to give feedback on many of the things I wanted to try. He also agreed to do a dozen emergency reviews within a short period after the review deadline. It was a pleasure to be a part of the team with Milos – I learned a lot.
- I thank Yale Patt for discussion on several aspects of the MICRO review process and for his valuable advice and suggestions. Yale ensured that important decisions were properly vetted and discussed.
- I thank Onur Mutlu, Amer Jaleel, and Lieven Eeckhout for discussions and feedback. I discussed many aspects of the revision model and PC process with them and got useful feedback and encouragement.
- I thank Pradip Bose for his dedicated service in handling the 28 papers conflicted with me. Pradip was always available during the process and was a pleasure to work with throughout the process.
- I thank Vinson Young and Chia-Chen Chou for implementing the changes in HotCRP for supporting the revision-based model and ensuring a smooth process for submissions, review, and paper selection.
- I thank Prashant Nair and Chia-Chen for working tirelessly to help me throughout the review process. They also ensured a smooth PC meeting, and implemented the innovations such as Electronic Voting and Conflict Screen. They also kept track of the notes and decisions during the PC meeting.
- I thank Gururaj Saileshwar for helping with arrangements for the Program Committee meeting.
- I thank Ching-Kai Liang for promptly doing web updates and facilitating communication with authors.
- I thank Zeshan Chisthi for handling the publication chair duties diligently and professionally.
- I thank the three keynote speakers – Jan Rabaey, Todd Austin, and Thomas Pawlowski – for accepting my invitation, and sharing their vision and experience with the broader architecture community.
- I thank all the 52 PC members for their dedication, integrity, time and hard work; they have all spent many days and significant effort to provide constructive reviews for the 18+ assigned papers, attend the PC meeting in person, and in some cases follow-up with the authors during shepherding.
- I thank the 135 ERC members many of whom reviewed 4 or more papers and were actively involved in online discussions.
- Finally, I thank the authors who submitted their work, irrespective of whether their paper was accepted or not – they are the lifeblood of the conference.

The technical program presented here is the culmination of hard work by many people: the program committee, the external review committee, and the authors of the submitted papers. It is the creativity, innovation, technical expertise, and rigor of these people that continues to make MICRO a vibrant, thriving, and a top-notch conference in our field. It was an honor for me to serve as the Program Chair. I would very much welcome any feedback you may have on anything related to MICRO-48, especially on the new things we tried this year.

Moinuddin K. Qureshi
Program Chair of MICRO-48