

## **Message from the ISCA 2020 Program Chair**

It is my great pleasure and honor to welcome you to the program of the 47th IEEE/ACM International Symposium on Computer Architecture (ISCA), the flagship conference in our community.

### **Program Overview**

We received a record number of 428 submissions (17% higher than ISCA 2019), out of which the program committee selected 77 papers for inclusion in the program – an acceptance rate of 18%. Seven papers were withdrawn at various stages during the review process. In addition to the 77 papers, the program also features three keynotes and an industry track. The keynotes are delivered by Prof. Margaret Martonosi (NSF CISE and Princeton University), Prof. Valeria Bertacco (University of Michigan) and Prof. Li-Shiuan Peh (National University of Singapore). The industry track is a new initiative as of this year and was led by Prof. David Patterson (Google and UC Berkeley). The industry track papers were selected through a separate review process which David Patterson describes in his foreword.

### **Program Committee and External Review Committee**

I had the honor to recruit and work with 62 Program Committee (PC) members who are experts on the broad range of topics within the field of computer architecture. Each PC member reviewed 20 or 21 papers. I also recruited 148 External Review Committee (ERC) members who each reviewed about half a dozen papers. When composing the PC and ERC, I paid special attention to balance domain expertise, gender, employment (academia versus industry), geography, and seniority. Finally, I also solicited additional reviews from 53 external reviewers who reviewed one or two papers.

### **Timeline: 14 weeks from Paper Submission to Author Notification**

The timeline for ISCA 2020 was as follows. Abstracts were due on Tuesday November 19, and full submissions were due on Tuesday November 26, 2019 – two days before the US Thanksgiving. Authors could not register a paper without having registered an abstract first. Reviewers were assigned by Friday December 13, 2019. The reviews were due Friday January 31, 2020. The authors were then given the opportunity to rebut the reviews and/or submit a revised manuscript between February 5 – 14. The PC, ERC and the external reviewers started the online discussion right after the author response period, up until the PC meeting. The PC met in person for 1.5 days to make the final paper acceptance decisions on February 26 – 27 at the University of California, San Diego, CA. The authors were notified with the paper decision outcomes on Wednesday March 4, 2020.

### **Single-Round Reviewing to Maximize Fairness and Balance Work/Life**

The review process consisted of a single round. A total of more than 2,150 reviews were submitted – all papers received at least 5 reviews; 44 papers received 6 reviews. All reviewers were assigned manually by me – I read all paper abstracts to match papers with reviewer expertise. Prof. Daniel Sorin (Duke University) assigned the papers I was conflicted with. Each paper got assigned three

PC reviews and two (or three) ERC or external reviewers. The reason for a single-round review cycle is twofold. First, it maximizes fairness among all submitted papers. Although a two-round or multiround review period may reduce the overall review workload to some extent, it also adds noise to the review process which may bias paper outcomes: good papers may be rejected early in the process based on a few reviews, and papers that pass on to the next round may have an unfair bias when assigned to new reviewers who know that the paper has already passed one or more review rounds. Second, a single review deadline allows reviewers to budget their time and balance work and life. More specifically, organizing multiple review rounds for ISCA would imply a review deadline during or shortly after the Christmas and New Year's break. Putting a review deadline towards the end of January enabled the reviewers to commit to the ISCA review load while at the same time being able to enjoy family time over the holidays.

### **Authors Were Given the Opportunity to Respond**

Authors were able to see the reviews (and the review scores) between February 5 – 14 and were invited to submit a revised manuscript and/or a rebuttal. The authors were given three options: (i) to not submit a rebuttal nor a revision, (ii) to submit an 800-word rebuttal but no revision, or (iii) to submit a revised paper and a rebuttal. The revised paper had to comply with the original submission guidelines (no extra pages were allowed); authors could submit a diff to highlight changes between the original and the revised paper. The goal of the revision option was to give the authors the opportunity to address (relatively) small issues that could be fixed within the response period, with the hope to accept as many excellent papers as possible.

A dozen papers each received an additional review after the author response phase. There were two major reasons for soliciting additional reviews. Either, the paper had less than five reviews before the rebuttal/revision phase – adding a fifth review provided fairness for all papers. Or, the PC/ERC reviewers, while discussing the paper online, realized they needed more expertise to reach a comprehensive decision, and hence asked for an additional expert review – adding an additional expert review increased the overall level of expertise to make a comprehensive decision during the online discussion and/or at the PC meeting.

### **Online Discussion to Enable an Efficient In-Person Meeting**

Papers were heavily discussed online by the PC, ERC and the external reviewers. The ERC and the external reviewers actively participated in the discussions. Each paper was assigned a discussion lead to initiate and monitor the discussion – discussion leads were selected based on expertise (primarily) and positivity and balance (secondarily). More than 4,400 comments were posted on the HotCRP paper submission site during the review process – on average, this is more than 10 comments per paper. The goal of the online discussion was threefold. First, reviewers were asked to provide a post-response overall merit score after having read the rebuttal, the revised submission (if available), and the other reviews. Second, papers were classified in three categories: #online-accept, #meeting-discuss and #online-reject. Papers were categorized by consensus among all the reviewers: 33 papers were accepted online, 119 papers were categorized as meeting-discuss, and the remaining 269 papers were rejected online. The third goal of the online discussion was to prepare a succinct paper summary, including the contributions, strengths, weaknesses and

discussion points, for the in-person PC meeting. This enabled efficient and effective discussions – it was the discussion lead’s responsibility to write up the summary in preparation for the PC meeting.

### **PC Meeting Co-Located with HPCA to Reduce Carbon Footprint**

The PC meeting was co-located with HPCA 2020 in San Diego, CA. The reason for co-locating the ISCA PC meeting with HPCA was to reduce travel time and cost and to reduce our carbon footprint. In addition, as the PC meeting took place earlier than usual, it gave authors of (a) rejected paper(s) ~6 weeks of time to thoroughly revise their work for resubmission to another conference (e.g., MICRO 2020 with a submission deadline mid April). The PC meeting was a single-track event and lasted for 1.5 days on Feb 26 (1pm – 6pm) and Feb 27 (8:30am – 6pm). Most PC members were able to attend the PC meeting in person. Nine people could not make it because of medical emergencies or personal reasons. The PC members who were unable to attend in person participated on the phone or had posted their final discussion comments online.

### **PC Meeting Process**

I emphasized at the beginning of the PC meeting that the goal is to select a high-quality program with papers that advance the field. The program should comprise of a good mix of (i) new-idea papers with high novelty, that are thought-provoking, open up a new direction for research, and have potentially high impact, and (ii) well-executed papers in more established areas, providing novel insight, making a significant step forward, and/or providing compelling results. I also reminded the PC that ISCA has a broad scope and the PC members should be positive and look for reasons to accept papers in the program. I clearly noted that there is no target acceptance rate.

The papers were ranked for discussion based on the post-response overall merit score. The discussion order for papers with the same overall merit score was randomized to limit the bias introduced by the discussion order. The process below was followed by all papers discussed at the PC meeting. The discussion lead introduced the paper by providing a summary of the paper’s contributions, its strengths, weaknesses and discussion points raised during the online discussion. The other PC reviewers then chimed in and provided their perspectives on the paper. Other PC members were allowed to ask questions. The initial goal of the discussion was to reach consensus among the PC reviewers. This means that a paper is accepted if all PC reviewers agreed to accept the paper after the discussion. Shepherding was provided as an option to reach consensus – shepherding was allowed to fix writing and clarification issues, not to provide new experiments and results. If all PC reviewers agreed to not accept the paper, the paper was rejected. In case the PC reviewers were unable to reach consensus, the whole PC was asked to vote. Before doing so, the PC was invited to ask additional clarification questions, and eventually one PC reviewer was asked to summarize why the paper should not be accepted whereas another PC reviewer would summarize why the paper should be included in the program. The outcome of the paper was then determined by the majority upon a PC-wide vote. In case of a tie, the paper was considered accepted. Voting was done electronically using HotCRP and was anonymous. The reason for keeping the voting anonymous was to not create a side-channel through which to leak votes. At the end of the PC meeting, the PC had accepted 77 submissions, out of which 13 with shepherding.

The discussion leads provided a discussion summary for all the papers that had submitted a rebuttal and/or revision, and were discussed online or at the PC meeting.

PC members were able to view all submitted papers, reviews and comments during the PC meeting, until the author notifications were sent out. The PC members could no longer see the list of the submitted papers, reviews and comments once the authors were notified – this is to close another potential side-channel about the review process. ERC members and external reviewers were able to see only the papers, reviews and comments for which they were a reviewer, i.e., ERC members and external reviewers could not ever see the list of submitted papers, reviews and comments – this was to close yet another potential side-channel for leaking review information. Author names of rejected papers were never revealed throughout the review process. Author names of accepted papers were revealed to the reviewers only after the final decision was reached.

### **Review Ethics Agreed Upon by All Reviewers**

All reviewers were asked to abide to and agree with the following review ethics upon entering a review into the HotCRP paper submission system. These review ethics were published on the ISCA 2020 website as well.

- Reviewers must treat all submissions as strictly confidential and destroy all papers once the technical program has been finalized.
- Reviewers must review their paper assignments themselves and provide unbiased reviews.
- Reviewers must contact the PC chair if they feel there is an ethical violation of any sort (e.g., authors seeking support for a paper, authors seeking to identify who the reviewers are).
- Reviewers are not allowed to actively look for author identities. Reviewers should judge a paper solely on its merits.
- Reviewers should review the current submission, and not judge previous versions of the paper they may have reviewed for other venues.

### **ISCA 2020 PC/ERC members and other reviewers are not allowed to:**

- Solicit external reviews of submitted papers from third parties, including their own students.
- Reviewers who have any reason to request a third-party review should contact the PC chair for third-party assignment.
- Discuss the content of a submitted paper with anyone other than unconflicted reviewers of the paper, PC/ERC members, and the PC chair during the review period (from now until paper publication in any venue).
- Discuss the content of reviews, including the reviewers' identities, with anyone other than unconflicted reviewers of the paper, PC/ERC members, and PC chair during the review period (from now until paper publication in any venue).
- Reveal the name of paper authors in case reviewers happen to be aware of author identity. (Author names of accepted papers will be revealed after the PC meeting; author names of rejected papers will never be revealed.)

- Disclose the outcome of a paper until its authors are notified of its acceptance or rejection.
- Disclose the content of a paper until its publication.
- Disclose the content of reviews, including the reviewers' identities, or discussions about a paper until its publication. After publication, reviewers, PC/ERC members, and the PC chair are allowed to discuss their own opinions of a paper, but not the content of others' reviews or any PC/ERC discussion about the papers.

## Challenges

The increasing number of submissions poses a number of challenges for our conference review model. Let me enumerate a couple:

- Identifying conflicts is a tedious process. I spent a full week double-checking conflicts. The submission site (HotCRP.com) provides suggestions for possibly missed conflicts – I got in touch with the authors to verify those and added conflicts if the authors had indeed missed the conflicts. Furthermore, we used scripts provided by Prof. Reetu Das and Prof. Tor Aamodt (MICRO 2019 Program Co-Chairs) to identify possibly unjustified conflicts based on a DBLP crawl – I got in touch with the authors to verify those conflicts and I removed the conflicts in case they were unjustified. A conflict database that keeps track of conflicts would facilitate the process for both the authors and the program chair.
- Manually assigning reviewers to papers is a tedious and time-consuming process, especially with 428 submissions and a PC plus ERC of more than 200 members. But, in my opinion, there is no other option. Automatically assigning reviewers to papers based on topic keywords provided by the authors and the reviewer expertise by the reviewers doesn't work because of the disparity in how authors chose topic keywords for their submissions – many provided only a single keyword whereas others provide half a dozen keywords. I spent about 10 days to do the manual review assignments.
- Managing an in-person PC meeting with 62 experts is challenging but is invaluable to 'set the bar', to find consensus if possible, and, if there is no consensus, to clearly spell out disagreements among the reviewers before making an informed decision.

## Thanks

Finally, I'm grateful to a large group of people without whom it would not have been possible to come up with the exciting program for ISCA 2020. A big thank you to:

- all the authors for having submitted their best work to ISCA 2020;
- the 62 PC members who each reviewed 20 or 21 papers, and who attended the in-person PC meeting or over the phone;
- the 148 ERC members who reviewed half a dozen papers;
- the 53 external reviewers;
- Daniel Sorin for handling my conflict papers;
- our three keynote speakers, Margaret Martonosi, Valeria Bertacco and Li Shiuan Peh;
- David Patterson for initiating and leading the industry track program;

- José Martinez and Jose Duato, the general chairs, for their support and help organizing the logistics for the PC meeting;
- Carole-Jean Wu, the publications chair, for collecting the final paper versions and assembling the proceedings;
- Eddie Kohler for promptly responding to my questions and requests regarding HotCRP.com;
- Josep Torrellas and Babak Falsafi, and their respective executive committees TCCA and SIGARCH, for their confidence in me and for supporting me when needed;
- Ajeya Naithani and Kartik Lakshminarasimhan, my PhD students, who spent many hours helping me run a smooth review process and PC meeting;
- Dean Tullsen, Hadi Esmaeilzadeh and Jocelyn Bernardo for helping me organize the PC meeting at the CSE Department at UCSD;
- Hillery Hunter and Erik Altman (ISCA 2019 Program Chair and Vice-Chair) and Reetu Das and Tor Aamodt (MICRO 2019 Program Chairs) for responding to my questions and providing feedback regarding the review process;
- Candace Lopez at the IEEE Computer Society for helping me negotiate the hotel contract for the PC meeting.

I hope you will enjoy ISCA 2020's technical program!

Yours sincerely,  
Lieven Eeckhout, *Ghent University*  
**ISCA 2020 Program Chair**