TOP PICKS

Yale N. Patt
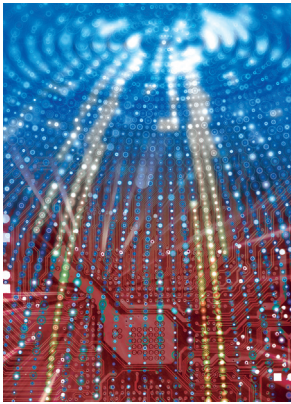University of Texas
at Austin

Onur Mutlu
Carnegie Mellon
University

• • • • • As co-chairs of the selection committee and guest editors of this issue, it is our pleasure to introduce you to this, the eighth Top Picks special issue of *IEEE Micro*. In the tradition started by then Editor in Chief Pradip Bose with the Nov./Dec. 2003 issue of *IEEE Micro*, we have attempted to share with you a sample of some of the best papers published in computer architecture conferences during the past year.

As should be expected, picking a dozen or so papers out of hundreds is difficult at best. The operative words in the above paragraph are "sample" and "some." In fact, while we know we did our best, we are equally certain that some will point to papers that were not selected that "should have made it," and to a few that were included that "absolutely shouldn't have been." In fact, the two of us are not entirely happy about every paper that was selected. Such is the nature of an attempt of this kind.

Our only response is that we treated the matter seriously and arrived at a majority decision within the selection committee on all papers accepted and rejected, regardless of how we, the co-chairs, felt. We describe the process in more detail below. Because all papers included were previously vetted by inclusion in one of our top conferences, we also provide some thoughts on how to improve the conference paper selection process. We feel strongly that work is needed in that area as well.

The process

A call for submissions was published in *IEEE Micro*, resulting in 87 papers submitted for consideration. Each submission consisted of a two-page summary of the paper, a one-page discussion of the paper's long-term impact, and a PDF of the published paper itself. In the case of Micro, which was held after the selection process, a PDF of the final camera-ready paper was submitted.

We appointed a committee of 29 professionals from our research community to evaluate the 87 papers. In keeping with the traditions of previous years, members of the selection committee, including the co-chairs, were invited to submit their published papers for inclusion in the special issue. Because the two co-chairs have a fundamental conflict with respect to papers done by each other (PhD advisor and student), any paper bearing either of our names, as well as any paper associated with our respective institutions or written by researchers we both had a conflict of interest with, was removed from the list of 87 papers and given to Professor Sudha Yalamanchili of Georgia Tech to handle independently. Nine papers fell into that category. Each paper was reviewed independently by three to five members of the selection committee plus an additional external reviewer when it became clear that a voice independent of the committee members was warranted. We would like to thank all members of the selection committee and the external reviewers for their reviews and valuable contributions.

An all-day meeting of the selection committee was held at the Marriott North close to DFW airport on October 15 to complete the review process. Of the 31 committee members, 26 were present at the meeting, and four of the remaining five were connected online via Skype. One member was sick and unable to participate.

The entire set of reviews (except for conflicts) was made available to all members of the selection committee four days before

the meeting. We did this to provide the opportunity for nonreviewers to read as many reviews as they wished to prior to the meeting. Because all papers had been previously published, committee members were for the most part seeing most of these papers for at least the second time. The co-chairs have for a long time felt that the usual practice of not making all reviews available to the entire committee (minus conflicts) until the day of the meeting prevents committee members from spending enough time on papers they have not reviewed to participate effectively in the discussion of those papers at the program committee (PC) meeting.

All 87 papers were available for discussion at the meeting. The selection committee chose the 11 papers included in this special issue as the 2010 Top Picks.

Observations regarding future conference reviewing

On the one hand, the *IEEE Micro* Top Picks special issue is outside the purview of selection criteria for our important conferences (ISCA, Micro, HPCA, ASPLOS, and so on). On the other hand, because these conferences supply the raw material for the Top Picks issue and because we have been involved with many program committees over the years, it is difficult to let this opportunity pass without putting an idea or two out there with respect to the conference review process. We feel that much needs to be done to improve the reviewing process, and our hope is that our comments will provide a stimulus to encourage discussion about how to fix it.

More focus on insights

We have seen a continually growing number of papers accepted on the basis of substantial quantitative improvements, often without sufficient insights into what is causing the benefit. Authors have picked up on this trend, resulting in the "formula" paper: k experiments, n bar graphs, x percent improvement. Two things are wrong with this picture. First, submissions focus on the amount of the improvement, rather than on insights. This leads to incremental tweaking rather than big ideas. Second, the thrust

of the paper, the actual quantitative results, without sufficient explanation and insights, are almost impossible for the reviewer to validate. Is the improvement the result of a good idea—or the result of a carefully massaged experiment, a weak baseline, or a bug in the experimental apparatus? One could ask the author to release the entire experimental apparatus to allow the reviewer to replicate the results. This would be unfair to diligent researchers who have spent years developing their infrastructure; they should not have to weigh publication against giving away their jewels.

We suggest that the focus of our acceptance criteria needs to change. Insights, backed up by reasonable experimentation, should trump quantitative results, and quantitative results should only be considered if accompanied by sufficient supporting material so that the reviewer does not have to accept the author's results at face value. Admittedly, a 10-page paper often does not provide enough space for sufficient documentation. We suggest that submissions should be accompanied by sufficient supporting material—not necessarily sufficient for someone to appropriate the entire experimental apparatus, but enough for a careful reviewer to evaluate the work. If the paper is subsequently accepted, the documentation is put in the public domain, perhaps as an internal technical report, accessible from the author's website. By emphasizing insights and holding data accountable, we submit that research in computer architecture will be better served.

Identity of reviewers

Reviewing papers is a thankless job, so it is no wonder that too many reviews are too superficial. We do not see how to legitimately break the cloak of blind negative reviews of a paper, since to do so invites too many opportunities for abuse. However, we do not see similar problems with accepted papers. Ergo, we suggest that accepted papers should carry loud and clear the identities of reviewers who recommend acceptance. Our hope is that at least reviews that recommend acceptance will be based on strong positive reviews, resulting from critical reading of the paper. Putting the name of the reviewer

Top Picks of 2010

This year's 11 selected papers come from four major computer architecture conferences: ISCA (five papers), Micro (three papers), ASPLOS (two papers), and HPCA (one paper). The overwhelming majority of submissions to this special issue came from these four conferences.

The selected articles, listed below, represent a diversity of topics, including reliability, on-chip interconnects, multicore architectures, main memory and caches, memory consistency, and phase change memory.

The first two papers deal with reliable design of microprocessors and memory. ''Virtualized ECC: Flexible Reliability in Main Memory,'' by Doe Hyun Yoon and Mattan Erez, presents a mechanism that virtualizes main memory error correction mechanisms such that dedicated hardware for storing error correction information is not required. ''Voltage Noise in Production Processors'' by Vijay Janapa Reddi et al. characterizes voltage variation in an existing processor and proposes techniques to mitigate the negative effects of such variations.

In the next article, ''Aérgia: A Network-on-Chip Exploiting Packet Latency Slack,'' Reetuparna Das et al. describe a packet scheduling policy that takes into account the importance of packets on application performance to improve the performance provided by on-chip interconnects.

The next three articles deal with various aspects of multicore architecture and microarchitecture. In ''Cohesion: An Adaptive Hybrid Memory Model for Accelerators,'' John Kelm et al. describe a hybrid memory model that supports both hardware-implemented and software-implemented cache coherence, with the goal of obtaining the advantages of each while minimizing their disadvantages. In ''Data Marshaling for Multicore Systems,'' M. Aater Suleman et al. describe a hardware/software cooperative method to reduce the performance overhead of remotely executing a code segment. In ''ReMAP: A Reconfigurable Architecture for Chip Multiprocessors,'' Matthew Watkins and David Albonesi describe uses and benefits of reconfigurable fabric that is shared among multiple cores to improve communication and synchronization between threads.

The next three articles examine various problems in main memory, caches, and their intersection. ''Thread Cluster Memory Scheduling'' by Yoongu Kim et al. describes a memory scheduling algorithm for multicore and multithreaded systems that can sustain both high system performance and fairness by applying different scheduling policies for different types of threads. ''Coordinating DRAM and Last-Level-Cache Policies with the Virtual Write Queue'' by Jeffrey Stuecheli et al. shows that making the cache policies aware of the main memory layout and characteristics can improve overall system performance by reducing the impact of writes on the more critical read requests. In ''CHOP: Integrating DRAM Caches for CMP Server Platforms,'' Xiaowei Jiang et al. describe techniques to improve efficiency of DRAM-based caches by filtering out cold pages from the cache.

In ''Address Translation Aware Memory Consistency,'' Bogdan F. Romanescu et al. develop a framework for reasoning about the correctness of memory consistency models at the virtual address level as opposed to the physical address level that has been established in previous work, and propose a method to dynamically verify the correct implementation of address translation.

The last article, ''Security Refresh: Protecting Phase-Change Memory against Malicious Wear Out'' by Nak Hee Seong et al., tackles the problem of wearout caused by writing into phase-change memory. The authors develop a wear-leveling algorithm that can maximize lifetime in the presence of a malicious adversary aiming to destroy the phase-change memory.

Reliability and error tolerance

- ''Virtualized ECC: Flexible Reliability in Main Memory'' by Doe Hyun Yoon and Mattan Erez
- ''Voltage Noise in Production Processors'' by Vijay Janapa Reddi, Svilen Kanev, Wonyoung Kim, Simone Campanoni, Michael D. Smith, Gu-Yeon Wei, and David Brooks

On-chip interconnects

- ''Aérgia: A Network-on-Chip Exploiting Packet Latency Slack'' by Reetuparna Das, Onur Mutlu, Thomas Moscibroda, and Chita R. Das

Multicore architectures

- ''Cohesion: An Adaptive Hybrid Memory Model for Accelerators'' by John H. Kelm, Daniel R. Johnson, William Tuohy, Steven S. Lumetta, and Sanjay J. Patel
- ''Data Marshaling for Multicore Systems'' by M. Aater Suleman, Onur Mutlu, José A. Joao, Khubaib, and Yale N. Patt
- ''ReMAP: A Reconfigurable Architecture for Chip Multiprocessors'' by Matthew A. Watkins and David H. Albonesi

Main memory and caches

- ''Thread Cluster Memory Scheduling'' by Yoongu Kim, Michael Papamichael, Onur Mutlu, and Mor Harchol-Balter
- ''Coordinating DRAM and Last-Level-Cache Policies with the Virtual Write Queue'' by Jeffrey Stuecheli, Dimitris Kaseridis, Lizy K. John, David Daly, and Hillery C. Hunter
- ''CHOP: Integrating DRAM Caches for CMP Server Platforms'' by Xiaowei Jiang, Niti Madan, Li Zhao, Mike Upton, Ravi Iyer, Srihari Makineni, Donald Newell, Yan Solihin, and Rajeev Balasubramonian

Memory consistency

- ''Address Translation Aware Memory Consistency'' by Bogdan F. Romanescu, Alvin R. Lebeck, and Daniel J. Sorin

Phase-change memory

- ''Security Refresh: Protecting Phase-Change Memory against Malicious Wear Out'' by Nak Hee Seong, Dong Hyuk Woo, and Hsien-Hsin S. Lee

on the accepted paper for all to see could help that. In an ideal world, we would also recommend listing the names of those who opposed acceptance. Unfortunately, . . .

External reviewers

Having served on many program committees, we know too well the feeling of being inundated with far too many (often more than 20) papers to review. It is tough to do this and also have a day job. It is also very unusual to find a PC member expert across the spectrum of papers he or she is asked to review. Thus, the need for external reviewers (that is, reviewers not on the program committee) who are asked to review two or three papers for the conference. Two or three papers enhances the likelihood of the paper being thoroughly reviewed, and an external reviewer enhances the likelihood of the reviewer actually being an expert on the paper's subject matter. In fact, many conferences already use external reviewers. The problem arises after the external reviewer has submitted his or her review and is done with the process. We suggest that the external reviewer continue to be part of the process through the rebuttal period and the subsequent discussion period, all the way up to the PC meeting, and that

the PC be careful to pay attention to the external review. Unfortunately, in too many cases, the external reviewer is marginalized because he or she is not (by definition) at the PC meeting, and in fact, is not part of the process after he or she submits the external review.

Acknowledging good reviewers and bad

Reviewers who provide useful comments and thorough, fair reviews are an important mechanism for the viability of our conferences, and those who do the opposite both clog up the works and derail good papers. Certainly we would never want to deny bad reviewers the privilege of publishing in our conferences, but something needs to be done about the imbalance in the efforts put forward by some at the expense of others. Is there some reward structure that could encourage good reviews? Is there some mechanism that would correct sloppiness or even mean-spirited reviews?

Self restraint

More and more program committees are facing larger and larger numbers of submissions. The result: too many weak papers containing at best one LPU (least publishable unit) each. These papers must be evaluated,

which reduces the amount of time a PC member can spend on each paper. Not surprisingly, the effect is a set of less-thorough reviews and more randomness in the evaluation and selection process. One could limit the number of submissions from any one author, but that seems wrong-headed. Given a reward system that insists on paper count, we recognize the folly of suggesting graduate students and untenured faculty show restraint in submitting—even though they usually know when their paper is not ready. There is just too much at stake. Consequently, we urge senior faculty to engage the dialogue to change the reward structure. What if tenure cases were based on candidates submitting only their best (and not more than) four pieces of work for evaluation? What if faculty recruiting committees insisted that an application identified only three papers, and no additional papers would be welcomed? The volume of submissions could recede, but only if senior faculty step up.

We hope you find the articles in this special issue useful. Note that each article was shortened to fit in this issue. We encourage you to also read the original conference versions. Finally, we welcome any feedback you have on anything about the special issue. MICRO

Acknowledgments

This special issue would not have been possible without the contributions of many people. We would like to thank the members of the selection committee and the external reviewers for their efforts. José Joao, a PhD student in the HPS Research Group at the University of Texas at Austin, maintained the submission/reviewing website and helped us run a smooth selection committee meeting by providing technical support.

David Albonesi, the Editor in Chief of *IEEE Micro* last year, was helpful along the way with prompt and clear answers to our questions. Finally, we would like to thank the *IEEE Micro* and IEEE Computer Society staff.

Yale N. Patt is the Ernest Cockrell Jr. Centennial Chair in Engineering at the University of Texas at Austin. His research interests include harnessing the expected fruits of future process technology into more effective microarchitectures for future microprocessors. He has a PhD in electrical engineering from Stanford University. His honors include the 1996 IEEE/ACM Eckert-Mauchly Award and the 2000 ACM Karl V. Karlstrom Award. He coauthored *Introduction to Computing Systems: From Bits and Gates to C and Beyond* (McGraw-Hill, 2003). He's a Fellow of IEEE and the ACM.

Onur Mutlu is an assistant professor in the Department of Electrical and Computer Engineering at Carnegie Mellon University. His research interests include computer architecture and systems. He has a PhD in electrical and computer engineering from the University of Texas at Austin. Some of his coauthored works were chosen as best papers at ASPLOS and HPCA conferences and as Top Picks by *IEEE Micro*. Also, he is the recipient of an NSF CAREER Award. He's a member of IEEE and the ACM.

Direct questions and comments about this article to Yale N. Patt at pattyn@austin.utexas.edu or to Onur Mutlu at onur@cmu.edu.