

Which Way Microarchitecture?

ERIK R. ALTMAN

caltman@us.ibm.com

..... I was recently looking over session topics of Micro conference papers since 2000. The top five topics in descending order of frequency were memory; microarchitecture and superscalar; caches; energy and power; and parallelization, concurrency, and multi-core. All of these topics had 10 to 16 sessions devoted to them over the last dozen years, and in total accounted for almost 50 percent of all sessions. Obviously, these are all important topics and are core to our field. However, with the exception of energy and power, all focus mostly on performance.

Aside from the 7.5 percent of sessions devoted to power and energy, only about 20 percent of our sessions focus on nonperformance topics—for example, reliability, security, debugging, simulation and testing, programming models, reconfigurability, virtualization, and big data and data centers. These other topics address important system capabilities. Indeed, many of them have conferences and journals devoted specifically to them. What is our view in the microarchitecture community—that microarchitecture can play a limited role for many of these things? That the role for microarchitecture is well understood in most dimensions? That these things are best handled in software? That these areas are sufficiently specialized that they should be covered largely as separate disciplines? Or something else?

Aside from aggregate numbers over a dozen years, some topics, of course, are increasing in popularity, others decreasing, and many are staying relatively constant. Most of the increases are probably not surprising: data centers and big data, network-on-chip, virtualization, and parallelization. Areas in at least temporary decline include branch prediction and speculation, compilers, and microarchitecture and superscalar. Interestingly, security had a mini-boom from 2003 to 2009, but it seems to have lapsed since.

The topics we study are an important issue for the community and also for us at *IEEE Micro*. Looking at the themes for our issues (<https://sites.google.com/site/ieemicro/issue-history>), we also tend to focus on performance-related issues, although our previous issue was on energy-aware computing, and we have issues in the pipeline for the next 14 months on reliability, dark silicon, and reconfigurable computing. (And there is still time to submit papers for consideration for all of these issues!) Nevertheless, we continually examine whether our *IEEE Micro* mix of topics, and especially performance and nonperformance topics, is appropriate. We also appreciate that the community really drives these topics. If you are doing research in traditional disciplines or emerging areas, we welcome your submissions.


In that vein, I think that this issue features a good and eclectic mix of topics.

As with last year's November/December issue, we have articles on Cool Chips, guest edited by Fumio Arakawa and Makoto Ikeda. Arakawa and Ikeda have selected and shepherded two excellent articles from the 2012 Cool Chips conference that provide insights into the microarchitectures of two application-oriented and highly power-constrained systems: controlling humanoid robots and object recognition in a mobile vision system. I hope that you will find them interesting and useful.

We also have several other articles, and I particularly want to call attention to our cover article, "Redefining the Role of the CPU in the Era of CPU-GPU Integration," by Manish Arora, Siddhartha Nath, Subhra Mazumdar, Scott B. Baden, and Dean Tullsen. This article explores what computation will be mapped to GPUs, and what work will remain on CPUs. The remaining CPU work has significantly different characteristics than historical CPU workloads, and the authors explore the implications for CPU microarchitecture. Fascinating stuff.

Happy reading!

Erik R. Altman
Editor in Chief
IEEE Micro

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.