Virtual BSC RS: Divide and Conquer Frontend Bottleneck

Objectives

You can watch the seminar in this **link**.

Abstract: The frontend stalls caused by instruction and BTB misses are a significant source of performance degradation. Server processors commonly use prefetchers to mitigate the frontend bottleneck. However, next-line prefetchers, which are available in server processors, are incapable of eliminating many L1 instruction misses. Temporal prefetchers, on the other hand, eliminate most of the misses but impose significant area overhead. Finally, while BTB-directed prefetchers offer low area overhead, as they rely on the BTB content for prefetching, BTB misses stall the prefetcher, which likely leads to costly instruction misses. In this talk, I present a divide-and-conquer approach to address the frontend bottleneck. The proposal, named SN4L+Dis+BTB, imposes the same area overhead as the state-of-the-art BTB-directed prefetcher, and at the same time, outperforms it by 5% on average and up to 16%.

Short

Pejman Lotfi-Kamran is an associate professor of computer science, the head of the school of computer science, and the director of Turin Cloud Services at the Institute for Research in Fundamental Sciences (IPM). His research interests include computer architecture with an emphasis on memory systems. His recent work on scale-out server processor design lays the foundation for Cavium ThunderX. Lotfi-Kamran has a Ph.D. in computer science from EPFL. He received his MS and BS in computer engineering from the University of Tehran. He is a member of IEEE and the ACM. (http://cs.ipm.ac.ir/~plotfi/)

Speakers

Pejman Lotfi-Kamran is an associate professor of computer science, the head of the school of computer science, and the director of Turin Cloud Services at the Institute for Research in Fundamental Sciences (IPM).
Barcelona Supercomputing Center - Centro Nacional de Supercomputación

---