

Energy Efficient Data Access and Storage through HW/SW Co-design

Minyi Guo

Shanghai Jiao Tong University, Shanghai, China

guo-my@cs.sjtu.edu.cn

ABSTRACT

Massive energy consumption has become a major factor for the design and implementation of datacenters. This has led to numerous academic and industrial efforts to improve the energy efficiency of datacenter infrastructures. As a result, in state-of-the-art datacenter facilities, over 80% of power is now consumed by servers themselves. Historically, the processor has dominated energy consumption in the server. However, as processors have become more energy efficient, their contribution has been decreasing. On the contrary, energy consumed by data accesses and storage is growing, since multi- and many-core servers are requiring increased main memory bandwidth/capacity, large register file and large-scale storage system. Accordingly, energy consumed by data accesses and storage approaching or even surpassing that consumed by processors in many servers. For example, it has been reported that main memory contributes to as much as 40–46% of total energy consumption in server applications. In this talk, we present our continuing efforts to improve the energy efficiency of data accesses and storage. We study on a series of approaches with hardware-software cooperation to save energy consumption of on-chip memory, register file, main memory and storage devices for embedded systems, multi- and many-core servers, respectively. Experiments with a large set of workloads show the accuracy of our analytical models and the effectiveness of our optimizations.

Categories and Subject Descriptors

C.5.m [Computer system implementation]: *Miscellaneous*

Keywords

Energy consumption; Embedded systems; Main memory; Storage devices; Register file;

Bio

Minyi Guo is now a chair professor of Shanghai Jiao Tong University, China. He is also the head of the Department of Computer Science and Engineering at Shanghai Jiao Tong University, China. Before 2009, he was a professor at the University of Aizu, Japan. Dr. Guo is serving as an associate editor of IEEE Transactions on Computers and IEEE Transactions on Parallel and Distributed Systems. His research interests include automatic parallelization and data-parallel languages, bioinformatics, compiler optimization, high-performance computing, and pervasive computing. He has published more than 200 papers in major conferences and journals in these areas. He is a senior member of the IEEE, a member of the ACM, IEICE and IPSJ.



Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

LCIES'14, June 12–13, 2014, Edinburgh, UK.

ACM 978-1-4503-2877-7/14/06.

<http://dx.doi.org/10.1145/2597809.2602569>