

Invited Speaker

Performance of Shared-Memory Programs on Multicore Systems

Yong Meng TEO

Department of Computer Science
National University of Singapore
email: teoym@comp.nus.edu.sg
url: www.comp.nus.edu.sg/~teoym

With the advent of multicore systems, parallel computing is becoming mainstream. However, with the growing number of programming models and languages, as well as the size and heterogeneity of multicore systems, understanding the performance of parallel programs is an increasingly difficult task.

In this talk, we discuss a practical approach for characterizing the performance of shared-memory programs. Based on trace of the operating system run-queue and hardware events counters, our analytical model derives both the speedup performance and the parallelism loss due to data dependency and memory overhead in both UMA and NUMA systems. Our approach is highly practical as it does not require the source or binary code, and is independent of language and thread implementation. Using extensive measurements on state-of-the-art multicore systems with 8, 24 and 48 cores, we establish the relationship between memory access patterns and the degree of off-chip contention. Our model is validated against measurements using NPB 3.3 dwarfs and the PARSEC benchmark. As applications of our model, we show how memory contention changes with problem sizes, and how to determine the optimal number of cores for a program.

Biography

TEO Yong Meng is an Associate Professor with the Department of Computer Science at the National University of Singapore. He currently heads the Information Technology Unit and the Computer Systems Research Group. He has been a Visiting Professor at the Shanghai Advanced Research Institute, Chinese Academy of Sciences since 2010. He was a Fellow of the Singapore-Massachusetts Institute of Technology Alliance from 2002-2006. He received his Master and PhD in Computer Science from the University of Manchester, UK, in 1987 and 1989. His main research interest is parallel & distributed computing. Current projects include performance analysis of large systems, cloud computing, composability theory and validation, and peer-to-peer file distribution. He has held various visiting positions at MIT(USA), KTH (Sweden), Hitachi Central Research Lab (Japan) among others.