

invited talk

A Model-driven Approach for Time-energy Performance of Heterogeneous Systems

Yong Meng TEO, Lavanya RAMAPANTULU, Dumitrel LOGHIN
Department of Computer Science
National University of Singapore
Email: teoym@comp.nus.edu.sg
url: www.comp.nus.edu.sg/~teoym

The advent of processors with diverse performance-to-power ratios, such as high-performance (brawny) processors, low-power (wimpy) servers, accelerators, introduces new opportunities and complexities to efficiently execute parallel applications on heterogeneous systems. This invited talk is divided into *three* main parts. First, we review system heterogeneities including inter-node, intra-node, inter-chip and intra-chip, and the challenges faced by users to achieve efficient execution. Secondly, we discuss a new model-driven approach to determine time-energy efficient configurations for executing parallel applications on inter-node heterogeneous clusters. By modeling the overlap of workload service demands on resources such as processor cores, memory and I/O devices, we obtain energy-efficient mixes of nodes that execute an application with minimum energy usage for a given execution time deadline, or vice-versa. Our analysis shows that these configurations (or “sweet spots”) form the energy-deadline Pareto frontier where the energy usage reduces as the deadline is relaxed. Lastly, we highlight new opportunities in advancing the performance impact of different system heterogeneities.

Biography



TEO Yong Meng is an Associate Professor of Computer Science at the National University of Singapore (NUS) and an Affiliate Professor at the NUS Business Analytics Centre. He was a Visiting Professor at various institutes in the Chinese Academy of Science, China from 2010-2014. He received his PhD in Computer Science from the University of Manchester. His research interest is on parallel and distributed systems and applications. In the last five years, he focused on the performance of parallel systems, cloud computing, and emergent properties in complex systems. The paper on strategy-proof dynamic pricing of cloud computing resources won the **Best Paper Award** at the 10th International Conference on Algorithms and Architectures for Parallel Processing in 2010. Another paper, co-authored with his PhD student, on time-based semantic validation won the **ACM SIGSIM Best PhD Student Paper Award** in 2009. He leads the Computer Systems Research Group. He also served as Advisor (Director's Office) on Large-Scale Computing Systems, Asia-Pacific Science and Technology Centre, Sun Microsystems Inc. from 2007-2008, and External Grant Evaluator, European Research Council (Ideas Specific Program) from 2008-2013. He has received numerous external research grants including European Commission, Fujitsu Computers (Singapore) Pte Ltd, Fujitsu Laboratories Ltd (Japan), Sun Microsystems/Oracle (USA), Nvidia, and PSA Corporation (Singapore) among other institutions.