

PhD Topic 10: Parallel Computing for Energy Efficiency

First Supervisor	Prof. Dr. Jesper Larsson Träff
Second Supervisor	Prof. Dr. Wolfgang Gawlik
Primary Field	Large-scale data- and compute-intensive systems
Secondary Field	Green Business
Interdisciplinary References	C2, C3, C4

Energy consumption and high performance are conflicting, major issues for computing devices ranging from mobile phones and embedded devices to dedicated high-performance computing systems. Parallel processing has traditionally been a means for improving the performance of applications, which is in itself difficult and not always possible. Parallel processing may contribute useful techniques to control the overall energy consumption of devices and applications. The energy consumption of a compute device depends on voltage and frequency, in such a way that lowering the frequency will lead to a more than linear decrease in energy consumption. A linear increase in processing elements could compensate the loss in performance (time), and at the same time have the potential to be more energy-efficient. This of course presupposes that efficient, linear speed-up parallel algorithms are available for the problems to be solved. However, the more than linear energy-decrease may leave room for some slackness by less than optimally efficient algorithms. At the same time, traditional parallel architectures based on extensions of cache-based sequential architectures are not likely to be the most energy-efficient solutions, as witnessed in the current interest in special-purpose architectures and accelerators like GPUs. Current multi-core architectures also invite some inefficient implementation techniques, like e.g. lock-spinning, and sometimes too little attention is paid to memory transfers; this problem can be considerable for offload-based solutions with GPUs.

This topic is concerned with trade-offs between time and energy-efficiency for parallel algorithms and architectures. It should be based on tractable and as far as possible realistic architecture models that can account for energy consumption (of instructions and data movements) and parallel performance. It will investigate new algorithms (for fundamental or application inspired problems) using these models, and should aim to verify results and predictions by concrete implementations, and model tradeoffs between performance and energy. It can investigate a more energy-efficient implementation technique that is avoiding spinning, locks, and tradeoffs between data-movement and re-computation. Use of special

purpose devices and heterogeneous systems, where possibly more energy- and time-efficient solutions have to be traded against programmability and data movement issues are likewise of high interest.

References:

- [1] Vijay Anand Korthikanti, Gul Agha, Mark R. Greenstreet: On the Energy Complexity of Parallel Algorithms. ICPP 2011:562-570
- [2] Ziliang Zong, Adam Manzanares, Xiaojun Ruan, Xiao Qin: EAD and PEBD: Two Energy-Aware Duplication Scheduling Algorithms for Parallel Tasks on Homogeneous Clusters. IEEE Trans. Computers (TC) 60(3):360-374 (2011)
- [3] Vijay Anand Korthikanti, Gul Agha: Analysis of Parallel Algorithms for Energy Conservation in Scalable Multicore Architectures. ICPP 2009:212-219
- [4] Lee Kee Goh, Bharadwaj Veeravalli, Sivakumar Viswanathan: Design of Fast and Efficient Energy-Aware Gradient-Based Scheduling Algorithms Heterogeneous Embedded Multiprocessor Systems. IEEE Trans. Parallel Distrib. Syst. (TPDS) 20(1):1-12 (2009)
- [5] Andreas Beckmann, Ulrich Meyer, Peter Sanders, Johannes Singler: Energy-efficient sorting using solid state disks. Green Computing Conference 2010: 191-202
- [6] Siegfried Benkner, Sabri Pllana, Jesper Larsson Träff, Philippas Tsigas, Uwe Dolinsky, Cédric Augonnet, Beverly Bachmayer, Christoph W. Kessler, David Moloney, Vitaly Osipov: PEPPER: Efficient and Productive Usage of Hybrid Computing Systems. IEEE Micro 31(5): 28-41 (2011)