Energy-Efficient Algorithms

Susanne Albers[*1]

1 **Department of Computer Science**
Humboldt-Universität zu Berlin, Germany
`albers@informatik.hu-berlin.de`

Abstract

This presentation surveys algorithmic techniques for energy savings. We address power-down as well as dynamic speed scaling mechanisms.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases energy efficiency, power-down mechanisms, dynamic speed scaling, offline algorithm, online algorithm

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2011.1

Summary

Algorithmic techniques for energy savings have received considerable research interest over the past years. Power dissipation is critical in portable, battery operated devices where the amount of available energy is severely limited. Moreover, energy consumption is a concern in desktop computers and servers as electricity costs impose a substantial strain on the budget of computing and data centers. Last but not least, power dissipation causes thermal problems. Most of the consumed energy is eventually converted into heat, resulting in wear and reduced reliability of hardware components.

Algorithmic work on energy conservation focuses mostly on the system and device level: How can we save energy in a given computational device? The following two mechanisms have been studied extensively.

- *Power-down mechanisms*: When a system is idle, it can be transitioned into low-power standby or sleep states. This technique is well-known and widely used. The goal is to determine good shutdown times subject to the constraint that a transition back to the active mode requires extra energy.
- *Dynamic speed scaling*: Microprocessors currently sold by chip makers such as AMD and Intel are able to operate at variable speed. The higher the speed, the higher the power consumption is. The goal is to save energy by applying the full speed/frequency spectrum of a processor and utilizing low speeds whenever possible.

In this talk we will review the above techniques along with some fundamental results. Most of the presentation will be devoted to dynamic speed scaling. We concentrate on a basic processor scheduling problem introduced in a seminal paper by Yao, Demers and Shenker [3]. Here a set of jobs, each specified by a release time, a deadline and a processing volume, has to be scheduled on a variable-speed processor so as to minimize energy consumption. Yao et al. devised a polynomial time offline algorithm as well as two online algorithms. We study settings with parallel processors and present efficient offline and online algorithms developed in [2]. Moreover, we investigate a setting where a variable-speed processor is equipped with

* Work supported by the German Research Foundation.

31st Int'l Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011).
Editors: Supratik Chakraborty, Amit Kumar; pp. 1–2
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

a sleep state. We show how to integrate power-down and speed scaling mechanisms and give efficient offline approximation algorithms [1].

References

1 S. Albers and A. Antoniadis. Race to idle: New algorithms for speed scaling with a sleep state. *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms,* 2012.

2 S. Albers, A. Antoniadis and G. Greiner. On multi-processor speed scaling with migration. *Proc. 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures,* 2011.

3 F.F. Yao. A.J. Demers and S. Shenker. A scheduling model for reduced CPU energy. *Proc. 36th IEEE Symposium on Foundations of Computer Science*, 374–382, 1995.