# Scheduling to Minimize Energy Usage and Heat

*Kirk R. Pruhs*

Computer Science Department, University of Pittsburgh, kirk@cs.pitt.edu

Energy usage and heat are quickly becoming limiting factors in the design of computational devices as the power usage of these devices increase. The energy in batteries in mobile devices (such as your laptop) is growing at a much slower rate than processor speeds and power usage, thus making conserving energy ever more important. Processors are requiring ever big fans and more elaborate schemes to cool them. There are many different techniques being developed to attack these problems. One of these techniques is dynamic voltage scaling, that is, the processor may be run at different speeds. Most laptops now come with multi-speed processors. As power is approximated by the processor speed cubed, slowing down the processor when possible can greatly extend the life of the battery, as well as allowing less expensive cooling methods.

In a traditional scheduling problem, the scheduler has to determine which job to run at each point in time. In the setting of a multi-speed processor, the scheduler also has to determine at what speed/power to run the processor at each time. This leads to dual criteria scheduling problems. That is, one wants to maximize the Quality of Service (QoS) while minimizing the heat/energy.

There is relatively little literature on these problems to date. I'll survey what is known, and will highlight open problems that I believe are interesting. The easiest QoS objective function to handle is deadline feasibility, that is, jobs have deadlines and the goal is to finish all jobs before their deadlines. The minimum energy deadline feasible schedule can be computed in polynomial time[4]. Further [4] gives an online algorithm that uses at most a constant factor more energy than the optimal algorithm. The paper [2] considers a generalization where the processor can be put into a lower-power sleep state, and that bringing the processor back to the on state requires a fixed amount of energy. The paper [2] gives a polynomial-time offline algorithm that uses at most three times the minimum possible energy, and an online algorithm that uses at most a constant times the minimum possible energy. The paper [3] consider the problem of minimizing the flow time subject to an energy constraint, and give a polynomial time offline algorithm for the case that all the jobs have the same work. In [1] results are given for the deadline feasibility problem with a constraint on the maximum temperature.

**Acknowledgments:** Supported in part by NSF grant CCR-0098752, NSF grant ANIR-0123705, NSF grant ANI-0325353, and a grant from the United States Air Force.

## References

[1] N. Bansal and K. Pruhs, "Temperature Aware Scheduling", in preparation.

[2] S. Irani, S. Shukla, and R. Gupta. "Algorithms for power savings", Symposium on Discrete Algorithms (SODA), 2003.

[3] K. Pruhs, P. Uthaisombut, and G. Woeginger, "Getting the Best Response for Your Erg", submitted for publication.

[4] F. Yao, A Demers, and S. Shenker. “A scheduling model for reduced cpu energy”, IEEE Symposium on Foundations of Computer Science (FOCS), 374–382, 1995.