

# Executive Summary: Dagstuhl Seminar 08071 on Scheduling February 10-15, 2008

Jane Liu and Rolf Mohring and Kirk Pruhs

April 16, 2008

Scheduling is a form of decision making that involves allocating scarce resources to achieve some objective. The study of scheduling dates back to at least the 1950's when operations research researchers studied problems of managing activities in a workshop. Computer systems researchers started studying scheduling in the 1960's in the development of operating systems and time-critical applications.

Several of these scheduling problems proved to be closed linked to fundamental concepts in theoretical computer science (such as approximability, intractability, and NP-completeness), thereby attracting the attention of theoretical computer scientists as well. Today, scheduling is widely studied as parts of the disciplines of mathematical programming and operations research; algorithmics and theoretical computer science; and computer systems, particularly real-time and embedded systems. The specific scarce resources, as well as the objectives to be optimized, differ in the different disciplines; nevertheless, there are remarkable similarities (as well as significant differences) in the general framework adopted by researchers in scheduling theory in these disparate disciplines.

The primary objectives of the seminar were to bring together leading and promising young researchers in the different communities to discuss scheduling problems that arise in current and future technology; to expose each community to the important problems addressed by the other communities; and to facilitate a transfer of solution techniques from each community to the others.

There were approximately fifty participants at the seminar, approximately evenly split between the three areas. The first morning consisted of three introductory talks by Ted Baker on real-time scheduling, Andreas

Schulz on mathematical programming techniques and by Nikhil Bansal on online scheduling. During the the first afternoon each participant gave a short talk about their research interests. The participants were then broken up into about ten clusters, of six people each, by the organizers based on research interests. Some of the clusters were:

- Game Theoretic Scheduling: This research line consider situations each user is motivated to try to optimize the performance for this users task. So for example, the user might not truthfully state the resources required by his/her task. The general goal is to study the effect of such selfish behavior and to design mechanisms to align players incentives with the global good.
- Stochastic Scheduling: This research line assumes that only probabilistic information is known about the resource requirements for a job. The goal is to try to find schedules that are good in expectation.
- Power Aware Scheduling: Modern processors can change their speed as a way of managing power. This research line focuses on designing speed changing policies that will be the most effective.
- Periodic Scheduling: This line of research is distinguished by the fact that jobs much be executed periodically over along period of time. Periodic scheduling occurs frequently in real-time applications. The resulting scheduling problems are notoriously difficult.

These clusters meet periodically through the week to discuss research problems, and to organize talks from within the cluster. The clusters organized 32 medium length talks on various recent research results. There was a plenary session on Thursday afternoon to discuss interesting directions for future research and future collaborations.

This seminar was essentially a first meeting of the real-time community with the operations research and theoretical computer science community. The general consensus was that both communities learned a lot about the other communities. And at least some collaborations between the communities were born from this seminar.