

10071 Abstracts Collection

Scheduling

— Dagstuhl Seminar —

Susanne Albers¹, Sanjoy K. Baruah², Rolf H. Möhring³ and Kirk Pruhs⁴

¹ HU Berlin, DE

`albers@informatik.hu-berlin.de`

² University of North Carolina - Chapel Hill, US

`baruah@cs.unc.edu`

³ TU Berlin, DE

⁴ University of Pittsburgh, US

`kirk@cs.pitt.edu`

Abstract. From Feb. 14, 2010 to Feb. 19, 2010, the Dagstuhl Seminar 10071 “Scheduling” was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Scheduling, real-time, complexity, approximation algorithms

10071 Executive Summary – Scheduling

The primary objectives of this seminar were to bring together leading researchers working on scheduling problems in three different research communities – operations research, theoretical computer science, and real-time systems – to expose each community to the important problems addressed by the other communities; to enable and encourage cooperation among the researchers; and to facilitate a transfer of solution techniques from each community to the others.

Keywords: Scheduling, real-time, complexity, approximation algorithms

Joint work of: Albers, Susanne; Baruah, Sanjoy K.; Möhring, Rolf H.; Pruhs, Kirk

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2541>

10071 Open Problems – Scheduling

Collection of the open problems presented at the scheduling seminar.

Keywords: Open problems, scheduling

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2536>

A Stochastic Framework for Multiprocessor Soft Real-Time Scheduling

James Anderson (University of North Carolina - Chapel Hill, US)

Prior work has shown that the global earliest-deadline-first (GEDF) scheduling algorithm ensures bounded deadline tardiness on multiprocessors with no utilization loss; therefore, GEDF may be a good candidate scheduling algorithm for soft real-time workloads. However, such workloads are often implemented assuming an average-case provisioning, and in prior tardiness-bound derivations for GEDF, worst-case execution costs are assumed. As worst-case costs can be orders of magnitude higher than average-case costs, using a worst-case provisioning may result in significant wasted processing capacity. In this paper, prior tardiness-bound derivations for GEDF are generalized so that execution times are probabilistic, and a bound on expected (mean) tardiness is derived. It is shown that, as long as the total expected utilization is strictly less than the number of available processors, the expected tardiness of every task is bounded under GEDF. The result also implies that any quantile of the tardiness distribution is also bounded.

The uploaded paper is from the upcoming RTAS. I would like to hear suggestions about how to ease the assumption of independent execution times in this analysis.

Keywords: GEDF, multiprocessor, tardiness

Joint work of: Anderson, James; Mills, Alex

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2537>

Hardness of approximation for Scheduling Problems

Nikhil Bansal (IBM TJ Watson Research Center, US)

We give a survey type talk on how to show hardness of approximation based on the Unique Games Conjecture. We illustrate this with two cases studies: A 0.878 hardness for the Max-Cut Problem, and a 2 hardness for the problem of minimizing weighted completion time with precedence constraints.

Keywords: Hardness of approximation, scheduling, Unique Games Conjecture

Hierarchical Real-Time Scheduling

Enrico Bini (Scuola Superiore Sant'Anna - Pisa, IT)

Issues on Hierarchical Real-Time Scheduling arising in compositional analysis

Keywords: Hierarchical Real-Time Scheduling

The Parallel Supply Function Abstraction for a Virtual Multiprocessor

Enrico Bini (Scuola Superiore Sant'Anna - Pisa, IT)

A new abstraction — the Parallel Supply Function (PSF) — is proposed for representing the computing capabilities offered by virtual platforms implemented atop identical multiprocessors. It is shown that this abstraction is strictly more powerful than previously-proposed ones, from the perspective of more accurately representing the inherent parallelism of the provided computing capabilities. Sufficient tests are derived for determining whether a given real-time task system, represented as a collection of sporadic tasks, is guaranteed to always meet all deadlines when scheduled upon a specified virtual platform using the global EDF scheduling algorithm.

Keywords: Virtual multiprocessor

Joint work of: Bini, Enrico; Marko, Bertogna; Baruah, Sanjoy K.

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2542>

Feasibility analysis of recurrent multiprocessor task systems

Vincenzo Bonifaci (MPI für Informatik - Saarbrücken, DE)

After a brief introduction to recurrent real-time scheduling problems, we review some recent results, including a new condition for Earliest Deadline-schedulability and Deadline Monotonic-schedulability that leads to approximate tests of feasibility for periodic and sporadic multiprocessor systems. We also highlight some outstanding open problems in this area.

Keywords: Real-time scheduling, feasibility analysis, approximation, periodic task system, sporadic task system

Algorithms and Complexity for Periodic Real-Time Scheduling

Vincenzo Bonifaci (MPI für Informatik - Saarbrücken, DE)

We investigate the preemptive scheduling of periodic tasks with hard deadlines. We show that, even in the uniprocessor case, no polynomial time algorithm can test the feasibility of a task system within a constant speedup bound, unless $P = NP$. This result contrasts with recent results for sporadic task systems. For two special cases, synchronous task systems and systems with a constant number of different task types, we provide the first polynomial time constant-speedup feasibility tests for multiprocessor platforms. Furthermore, we show that the problem of testing feasibility is coNP-hard for synchronous multiprocessor tasks systems. The complexity of some of these problems has been open for a long time.

We also propose a profit maximization variant of the feasibility problem, where every task has a nonnegative profit, and the goal is to find a subset of tasks that can be scheduled feasibly with maximum profit. We give the first constant-speed, constant-approximation algorithm for the case of synchronous task systems, together with related hardness results.

Keywords: Real-time scheduling, periodic task system, feasibility test, Earliest Deadline First, approximation algorithm

Joint work of: Bonifaci, Vincenzo; Chan, Ho-Leung; Marchetti-Spaccamela, Alberto; Megow, Nicole

Full Paper:

http://www.siam.org/proceedings/soda/2010/SODA10_109_bonifaciv.pdf

See also: V. Bonifaci, H.-L. Chan, A. Marchetti-Spaccamela, N. Megow. Algorithms and Complexity for Periodic Real-Time Scheduling. In Proc. 21st ACM-SIAM Symposium on Discrete Algorithms. SIAM Press, 2010.

Minimum Energy Scheduling

Marek Chrobak (University of California - Riverside, US)

With any schedule we can associate its energy usage in the following way: When the processor is on, it uses 1 unit of energy per step. During idle times (gaps), the processor can be on or off. However, turning the processor on requires L units of energy. Thus we can assume that the processor is on in gaps of length at most L and it's off in gaps longer than L . Given n jobs, each with a release time, deadline, and processing time, we want to compute their preemptive schedule (if it exists) with minimum energy.

In this talk we will start with the fundamental algorithm by Philippe Baptiste that computes a schedule for unit-length jobs that has the minimum number of

gaps (this corresponds to the case $L < 1$). Then we will review other results in this area, and we will discuss techniques for speeding up Baptiste's algorithm and for extending it to jobs with arbitrary processing times. If time suffices, we will also cover the multiprocessor case and the case of jobs with agreeable deadlines.

Keywords: Scheduling, algorithm, dynamic programming, energy

Polynomial Time Algorithms for Minimum Energy Scheduling

Marek Chrobak (University of California - Riverside, US)

The aim of power management policies is to reduce the amount of energy consumed by computer systems while maintaining satisfactory level of performance. One common method for saving energy is to simply suspend the system during the idle times. No energy is consumed in the suspend mode. However, the process of waking up the system itself requires a certain fixed amount of energy, and thus suspending the system is beneficial only if the idle time is long enough to compensate for this additional energy expenditure. In the specific problem studied in the paper, we have a set of jobs with release times and deadlines that need to be executed on a single processor. Preemptions are allowed. The processor requires energy L to be woken up and, when it is on, it uses the energy at a rate of R units per unit of time. It has been an open problem whether a schedule minimizing the overall energy consumption can be computed in polynomial time. We solve this problem in positive, by providing an $O(n^5)$ -time algorithm. In addition we provide an $O(n^4)$ -time algorithm for computing the minimum energy schedule when all jobs have unit length.

Keywords: Scheduling, algorithm, dynamic programming, energy

Joint work of: Chrobak, Marek; Baptiste, Philippe; Dürr, Christoph

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2535>

Optimal fixed-priority scheduling algorithm for probabilistic real-time with variable execution times

Liliana Cucu-Grosjean (INRIA - Nancy, FR)

We deal with periodic tasks with release times, deadlines and (probabilistic) variable execution times.

The problem is to determine if there is an optimal fixed-priority algorithm that can schedule these tasks with arbitrary preemption on one processor. To the best of my knowledge this open problem arises from the paper:

D. Maxim and L. Cucu-Grosjean, Towards optimal priority assignment for probabilistic real-time systems with variable execution times, Proceedings of the 3rd Junior Researcher Workshop on Real-Time Computing (JRWRTC 2009)

Related Results and Comments: Optimality for fixed-priority algorithms: such algorithm is optimal in the sense that if there is (at least) a priority assignment that satisfies the constraints then the algorithm will find it. No (general) optimal algorithm is known for this problem.

The satisfaction of the constraints is verified using results provided in J.L Diaz, D.F. Garcia, K. Kim, C.G. Lee, L.L. Bello, J.M. Lopez and O. Mirabella, Stochastic Analysis of Periodic Real-Time Systems, Proceedings of 23rd of the IEEE Real-Time Systems Symposium (RTSS 2002)

Keywords: Periodic, probabilistic execution times, real-time

In Multiprocessor Global Fixed Priority Pre-emptive Scheduling, What is the arrival pattern that leads to the worst-case response time?

Rob Davis (University of York, GB)

This is a short presentation about a research problem in multiprocessor systems using global fixed priority pre-emptive scheduling (gFPPS).

For the non-concrete periodic / sporadic task models, the pattern of task arrivals that leads to the worst-case response time for a particular task is unknown. Obtaining this arrival pattern would provide an exact schedulability test for gFPPS. Any properties that can be derived for the worst-case arrival pattern would provide an opportunity to improve upon current sufficient schedulability tests.

Keywords: Multiprocessor global fixed priority pre-emptive scheduling

Resource Sharing in Global Fixed-Priority Preemptive Multiprocessor Scheduling

Arvind Easwaran (ISEP-IPP - Porto, PT)

In this paper we consider global fixed-priority preemptive multiprocessor scheduling of constrained-deadline sporadic tasks that share resources in a non-nested manner. We develop a novel resource-sharing protocol and a corresponding schedulability test for this system. We also develop the first schedulability analysis of priority inheritance protocol for the aforementioned system. Finally, we show that these protocols are efficient (based on the developed schedulability tests) for a class of priority-assignments called *reasonable* priority-assignments.

Keywords: Scheduling

Joint work of: Easwaran, Arvind; Andersson, Björn

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2538>

Every Deterministic Nonclairvoyant Scheduler has a Suboptimal Load Threshold

Jeff Edmonds (York University - Toronto, CA)

Problem Statement: The goal is to prove a surprising lower bound for resource augmented nonclairvoyant algorithms for scheduling jobs with sublinear nondecreasing speed-up curves on multiple processors with the objective of average response time. Edmonds and Pruhs in SODA09 prove that for every $\epsilon > 0$, there is an algorithm Alg_ϵ that is $(1+\epsilon)$ -speed $O(\frac{1}{\epsilon^2})$ -competitive. A problem, however, is that this algorithm Alg_ϵ depends on ϵ . The goal is to prove that every fixed deterministic nonclairvoyant algorithm has a suboptimal speed threshold, namely for every (graceful) algorithm Alg , there is a threshold $1+\beta_{\text{Alg}}$ that is $\beta_{\text{Alg}} > 0$ away from being optimal such that the algorithm is $\Omega(\frac{1}{\epsilon\beta_{\text{Alg}}})$ competitive with speed $(1+\beta_{\text{Alg}})+\epsilon$ and is $\omega(1)$ competitive with speed $1+\beta_{\text{Alg}}$.

I have worked very hard on it and have felt that I was close. The proof technique is to use Brouwer's fixed point theorem to break the cycle of needing to know which input will be given before one can know what the algorithm will do and needing to know what the algorithm will do before one can know which input to give. Every thing I have can be found at

Keywords: Scheduling

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2544>

Full Paper:

<http://www.cse.yorku.ca/~jeff/research/schedule/lowerbound.pdf>

Full Paper:

<http://www.cse.yorku.ca/~jeff/research/kirk/laps/laps.ppt>

Scalably Scheduling Processes with Arbitrary Speedup Curves

Jeff Edmonds (York University - Toronto, CA)

We give a scalable $((1 + \epsilon)$ -speed $O(1)$ -competitive) nonclairvoyant algorithm for scheduling jobs with sublinear nondecreasing speed-up curves on multiple processors with the objective of average response time.

Keywords: Scheduling

Joint work of: Edmonds, Jeff; Pruhs, Kirk

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2546>

Power-Aware Real-Time Scheduling: Models, Open Problems, and Practical Considerations

Nathan Fisher (Wayne State University, US)

Power-related issues have received considerable research attention from the real-time community in the past decade. In our talk, we introduce a recent model and set of assumptions made in the recent real-time literature on energy and thermal issues; suggest two high-level open problems for power-aware real-time scheduling: *peak-temperature minimization* and *energy-minimization with temperature as a constraint*; and discuss practical considerations that should be considered in proposed solutions.

Keywords: Real-time scheduling, power-aware scheduling, sporadic tasks

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2010/2539>

Energy Efficient Scheduling via Partial Shutdown

Samir Khuller (University of Maryland, US)

We define a collection of new problems referred to as “machine activation” problems. The central framework we introduce considers a collection of M machines (unrelated or related), with machine i having an activation cost of a_i .

There is also a collection of N jobs that need to be performed, and p_{ij} is the processing time of job j on machine i . Standard scheduling models assume that the set of machines is fixed and all machines are available. We assume that there is an activation cost budget of A – we would like to select a subset S of the machines to activate with total cost $a(S) \leq A$ and find a schedule for the jobs on the machines in S minimizing the makespan. In this work we develop bi-criteria approximation algorithms for this problem based on both LP rounding and a greedy approach.

Keywords: Unrelated parallel machine scheduling, approximation algorithms

Joint work of: Khuller, Samir; Li, Jian; Saha, Barna

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2543>

Full Paper:
<http://www.cs.umd.edu/~samir/grant/soda10.pdf>

See also: Appeared in ACM-SIAM Symp. on Discrete Algorithms (2010)

Scheduling with varying Degrees of Parallelizability

Ravishankar Krishnaswamy (Carnegie Mellon University - Pittsburgh, US)

In this talk, we introduce, motivate and show how to work with the following class of problems: consider a multiprocessor system where jobs arrive online, and the scheduler is "non-clairvoyant" about the nature of the jobs – it only knows when a job arrives and when a job completes. However, between these times, any job may have several stages, and each stage may be parallelizable to a different extent (i.e one stage could be highly sequential where running the job on more than one machine may be useless, and another may be highly parallelizable, the rate at which the job gets processed could scale (almost) linearly with the number of machines it is scheduled on); but the scheduler is not aware of these parameters.

In this talk, we go over the algorithm and the proof techniques (using a non-local potential function based argument) introduced by Edmonds (1999) and Edmonds and Pruhs (2009), and talk about a recent generalization (joint work with Anupam Gupta, Sungjin Im, Benjamin Moseley and Kirk Pruhs) to the objective function of minimizing the L_k norm of flowtimes.

Keywords: Online algorithms, non-clairvoyant algorithms, scheduling, parallelizability

Joint work of: Gupta, Anupam; Im, Sungjin; Krishnaswamy, Ravishankar; Moseley, Benjamin; Pruhs, Kirk

See also: Anupam Gupta, Sungjin Im, Ravishankar Krishnaswamy, Benjamin Moseley, and Kirk Pruhs. Scheduling jobs with varying parallelizability to reduce variance. In *SPAA*, 2010.

A Primal Dual Approach to Online Optimization Problems

Seffi Naor (Technion - Haifa, IL)

A unified approach, based on the primal-dual method, is discussed for a wide range of online covering and packing problems, having various objective functions.

This approach has led to a simple alternative view and analysis of many previously suggested algorithms, as well as new results. In particular, a randomized $O(\log k)$ -competitive online algorithm for the weighted paging problem, where there is a (page dependent) weight for fetching each page into the cache, and k is the cache size.

This is the first randomized $o(k)$ -competitive algorithm for the problem and its competitiveness matches the known lower bound on the problem. Weighted paging is a special case (weighted star metric) of the well known k -server problem for which it is a major open question whether randomization can be useful in obtaining sublinear competitive bounds. The focus of the talk will be on developing the general methodology.

Based on papers with Nikhil Bansal, Niv Buchbinder, and Kamal Jain.

Keywords: Online algorithms, competitive analysis, primal-dual analysis

Scheduling periodic tasks in a hard real-time environment

Martin Niemeier (EPFL - Lausanne, CH)

We consider a real-time scheduling problem that occurs in the design of software-based aircraft control. The goal is to distribute tasks $\tau_i = (c_i, p_i)$ on a minimum number of identical machines and to compute offsets a_i for the tasks such that no collision occurs. A task τ_i releases a job of running time c_i at each time $a_i + k \cdot p_i$, $k \in \mathbb{N}_0$ and a collision occurs if two jobs are simultaneously active on the same machine.

We shed some light on the complexity and approximability landscape of this problem.

Although the problem cannot be approximated within a factor of $n^{1-\varepsilon}$ for any $\varepsilon > 0$, an interesting restriction is much more tractable: If the periods are dividing (for each i, j one has $p_i | p_j$ or $p_j | p_i$), the problem allows for a better structured representation of solutions, which leads to a 2-approximation. This result is tight, even asymptotically.

Keywords: Real-Time Scheduling, Periodic scheduling problem, Periodic maintenance problem, Approximation hardness, Approximation algorithm

Joint work of: Eisenbrand, Friedrich; Hähnle, Nicolai; Niemeier, Martin; Skutella, Martin; Verschae, Jose; Wiese, Andreas

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2010/2534>

Full Paper:

http://disopt.epfl.ch/webdav/site/disopt/shared/PM_EHNSVW10_report.pdf

See also: F. Eisenbrand, N. Hähnle, M. Niemeier, M. Skutella, J. Verschae, and A. Wiese. Scheduling periodic tasks in a hard real-time environment. In Proceedings of ICALP 2010, Lecture Notes in Computer Science. Springer, 2010. To appear.

Recent Hardness Results for Periodic Uni-processor Scheduling

Thomas Rothvoss (EPFL - Lausanne, CH)

Consider a set of n periodic tasks τ_1, \dots, τ_n where τ_i is described by an execution time c_i , a (relative) deadline d_i and a period p_i .

We assume that jobs are released synchronously (i.e. at each multiple of p_i) and consider pre-emptive, uni-processor schedules.

We show that computing the response time of a task τ_n in a Rate-monotonic schedule i.e. computing

$$\min \left\{ r \geq c_n + \sum_{i=1}^{n-1} \left\lceil \frac{r}{p_i} \right\rceil c_i \leq r \right\}$$

is (weakly) **NP**-hard (where τ_n has the lowest priority and the deadlines are implicit, i.e. $d_i = p_i$).

Furthermore we obtain that verifying EDF-schedulability, i.e.

$$\forall Q \geq 0 : \sum_{i=1}^n \left(\left\lfloor \frac{Q - d_i}{p_i} \right\rfloor + 1 \right) \cdot c_i \leq Q$$

for constrained-deadline tasks ($d_i \leq p_i$) is weakly **coNP**-hard.

Keywords: Hardness, periodic scheduling, uni-processor scheduling

Joint work of: Eisenbrand, Friedrich; Rothvoss, Thomas

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2545>

Full Paper:

<http://infoscience.epfl.ch/record/141575/files/EDFhardness-cameraready-SODA2010.pdf?version=1>

See also: Friedrich Eisenbrand and Thomas Rothvoss. EDF-schedulability of synchronous periodic task systems is coNP-hard. In *ACM-SIAM Symposium on Discrete Algorithms (SODA10)*, 2010.

Approximation Algorithms for Stochastic Sequencing and Scheduling Problems

David Shmoys (Cornell University, US)

In the design and analysis of approximation algorithms, one is interested in proving performance guarantees on the quality of near-optimal solutions produced by an efficient (i.e., polynomial-time) algorithm.

In stochastic scheduling, the tradition has been to consider the design of policies that produce asymptotically optimal solutions for inputs in which the processing time of jobs is specified by independent probability distributions and its realization is learned by the policy solely by running the job to completion. Over the past 10 years, there has been an important trend of marrying the two approaches by instead designing and analyzing algorithms in a stochastic setting to obtain performance guarantees relative to the optimal non-anticipatory policy. Furthermore, this approach has been considered not just in the setting where the processing times are random, but also when the existence of a job is probabilistic. We will give a few examples of work in this area.

Keywords: Approximation algorithms, stochastic scheduling problems, online scheduling models

Optimal Mechanisms for Scheduling

Marc Uetz (University of Twente, NL)

We study the design of optimal mechanisms in a setting where a service provider needs to schedule a set of non-preemptive jobs, one job at a time. Jobs need to be compensated for waiting, and waiting cost is private information. In this setting, an optimal mechanism is one that induces jobs to report truthfully their waiting cost, while minimizing the total expected compensation cost of the service provider. Here, truthful refers to Bayes-Nash implementability, and assumes that private information is independently drawn from known distributions. We derive closed formulae for the optimal mechanism, and show that it is a modification of Smith's ratio rule. We also show that it can be implemented in dominant strategies. Our analysis relies on a graph-theoretic interpretation of the incentive compatibility constraints. It parallels the analysis known for auctions with single parameter agents, yet it exhibits some subtle differences.

We also consider the multi-dimensional case where also the service times of jobs are private information. We show that for this problem the optimal mechanism generally does not satisfy an independence condition known as IIA, and thus known approaches are doomed to fail.

Keywords: Optimal Mechanism Design, Scheduling, Job Agents, Smith's Rule

Joint work of: Heydenreich, Birgit; Mishra, Debasis; Müller, Rudolf, Uetz, Marc

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2540>