

08071 Abstracts Collection
Scheduling
— **Dagstuhl Seminar** —

Jane W. S. Liu¹, Rolf H. Möhring² and Kirk Pruhs³

¹ Academia Sinica - Taipei, TW

janeliu@mx.iis.sinica.edu.tw

² TU Berlin, D

³ Univ. of Pittsburgh, USA

kirk@cs.pitt.edu

Abstract. From 10.02. to 15.02., the Dagstuhl Seminar 08071 “Scheduling” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. 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, supply chain

08071 Executive Summary – Scheduling

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.

Keywords: Scheduling, real-time, supply chain

Joint work of: Liu, Jane W. S.; Möhring, Rolf H.; Pruhs, Kirk

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2008/1487>

Tardiness Bounds for Global Multiprocessor Scheduling Algorithms

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

We consider the issue of deadline tardiness under global multiprocessor scheduling algorithms.

We present a general tardiness-bound derivation that is applicable to a wide variety of such algorithms, including global earliest-deadline-first (EDF), first-in, first-out (FIFO), EDF until zero laxity (EDZL), and least-laxity-first (LLF). For these tardiness bounds to hold, constraints on overall utilization are not required (other than requiring that the processing platform is not over-utilized). Our derivation is very general: job priorities may change rather arbitrarily at runtime, arbitrary non-preemptive regions are allowed, and capacity restrictions may exist on certain processors. Our results show that, with the exception of static-priority algorithms, most global algorithms considered previously have bounded tardiness. In addition, our results provide a simple means for checking whether tardiness is bounded under newly-developed algorithms. These results were obtained as part of a broader effort to understand how to best schedule real-time workloads on multicore platforms, particularly workloads with mostly soft real-time constraints.

An introduction to real-time scheduling

Theodore P. Baker (Florida State University, USA)

The defining characteristic of real-time computer systems is the existence of constraints on the timing of certain operations.

Timing constraints typically arise from the need for a computer system to maintain some degree of synchrony with external processes. Loss of synchronization may have potentially severe consequences, so the primary goal of scheduling in a real-time system is to guarantee that the timing constraints will always be satisfied. This goal can be difficult to achieve because of several complicating factors, such as unpredictable timing of external events, consideration of multiple resources, concurrent operation of multiple processing units, software and hardware complexity, and the inherent computational complexity of the scheduling problems introduced by these factors.

A great deal of research has been one on real-time scheduling.

This has largely focussed on scheduling CPU usage to meet timing constraints, but also on other resources and other constraints, such as power and energy. A variety of models have been considered for timing constraints, workloads, processors, and run-time environments. A corresponding variety of scheduling algorithms has been proposed, and analyzed. All of this work has been characterized by a trade-off between generality and tractability of analysis. That is, strong timing guarantees have been derived for some very restrictive models, and comparatively weak results have been obtained for more general models.

In this talk we will survey some of the application requirements that lead to real-time scheduling constraints, and some of the well known models for timing constraints, workloads, processors and run-time environments. We will explain the difficulty of execution time prediction and some techniques for coping with it.

We will briefly introduce some examples of models which have been analyzed with the most success, including the periodic and sporadic task models for single and multiple processors, and two important approaches to the analysis of such systems: Critical Zone analysis and demand analysis. We will conclude by mentioning some of the more complex scheduling problems, which present continuing scheduling problems.

Keywords: Real-time scheduling

Brute-force determination of multiprocessor schedulability for sets of sporadic hard-deadline tasks

Theodore P. Baker (Florida State University, USA)

This report describes a necessary and sufficient test for the schedulability of a set of sporadic hard-deadline tasks on a multiprocessor platform, using any of a variety of scheduling policies including global fixed task-priority and earliest-deadline-first (EDF). The contribution is to establish an upper bound on the computational complexity of this problem, for which no algorithm has yet been described. The compute time and storage complexity of the algorithm, which performs an exhaustive search of a very large state space, make it practical only for tasks sets with very small integer periods. However, as a research tool, it can provide a clearer picture than has been previously available of the real success rates of global preemptive priority scheduling policies and low-complexity sufficient tests of schedulability.

Keywords: Sporadic task, schedulability, feasibility, real time, hard deadline

See also: T.P. Baker and M. Cirinei, "Brute-force determination of multiprocessor schedulability for sets of sporadic hard-deadline tasks", Principles of Distributed Systems, Springer Lecture Notes in Computer Science No. 4878; the Proceedings of OPODIS 2007, the 10th International Conference on Principles of Distributed Systems", Guadeloupe (December 2007) 62-75.

The primal dual framework for online problems

Nikhil Bansal (IBM TJ Watson Research Center, USA)

While many general techniques have been developed for the design and analysis of approximation algorithms, this is not the case for online algorithms. Typically analyzing an online algorithm involves studying some potential function, which often seems to be pulled out of thin air.

Recently, Buchbinder and Naor developed a very general technique based on the online primal dual framework. In addition to obtaining various new results, the technique also gives substantially simpler proofs of various previous results.

More significantly, the technique reveals the underlying combinatorial structure in seemingly different problems.

In this talk, I will give an overview of the technique and consider various applications including the recent poly-logarithmic competitive algorithms for paging problems with non-uniform weights and page sizes.

Joint work of: Bansal, Nikhil; Buchbinder, Niv; Naor, Joseph (Seffi)

Two open problems in multiprocessor real-time scheduling

Sanjoy K. Baruah (University of North Carolina- Chapel Hill, USA)

The recurrent task models commonly used to represent real-time workloads was discussed. This model differs from those typically studied in "traditional" scheduling theory in several ways. In particular, jobs are generated by recurrent processes called tasks, each of which generates a potentially infinite sequence of jobs; furthermore, the exact sequence of jobs generated by a given task is not known beforehand (and hence scheduling systems of such tasks is inherently an on-line problem).

Two open problems in multiprocessor real-time scheduling theory were described. 1. Determine feasibility for a given collection of recurrent real-time tasks upon a preemptive multiprocessor platform, and 2. design policies for arbitrating access to additional non-preemptive serially reusable resources, which may be accessed by individual jobs within critical sections guarded by semaphores.

Space of feasible parameters

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

The design of real-time systems can be view as an optimization problem. Hence it is important to find a characterization of the feasible parameters, which is well suited for optimizations.

Keywords: Real-time scheduling

Uniprocessor EDF Feasibility is an Integer Problem

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

The research on real-time scheduling has mostly focused on the development of algorithms that allows to test whether the constraints imposed on the task execution (often expressed by deadlines) are verified or not. However, in many design scenarios the task set is only partially known and these algorithms cannot be applied because they require the complete knowledge of all the parameters of

the task set. Moreover, very often the designer has the freedom to select some of the task set parameters in order to maximize the system performance, and an arbitrary selection of the free parameters can lead either to poor performance or to a constraint violation. It is then useful to describe the feasibility region of the task set parameters by equations instead of by algorithms, so that optimization algorithms can be applied to find the best assignment to the free variables.

In this paper we formulate the EDF schedulability on a single processor through a combination of linear constraints. We study the geometry of the feasibility region of task deadlines when computation times and periods are known.

Keywords: EDF schedulability condition, optimal deadline assignment

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1488>

Hierarchical Scheduling and Timebands

Alan Burns (University of York, GB)

Complex systems exhibit temporal behaviour at many different time scales from microseconds or less to years or more. At each of these levels scheduling activities take place. This short talk will introduce the notion of timebands that have been used to model the temporal structure of complex systems. This framework may be useful to describe hierarchical relationships between schedules. The talk will consider some useful distinctions between what is sometimes called scheduling and planning. The former delivers a policy; the latter a plan.

This contribution is more conceptual than technical!

Keywords: Scheduling hierarchical

Algorithms for Temperature-Aware Task Scheduling

Marek Chrobak (Univ. California - Riverside, USA)

We study scheduling problems motivated by recently developed techniques for microprocessor thermal management at the operating systems level.

The general scenario can be described as follows.

The microprocessor temperature is controlled by the hardware thermal management system that continuously senses the chip temperature and automatically reduces the processor's speed as soon as the thermal threshold is exceeded. Some tasks are more CPU-intensive than other and thus generate more heat during execution. The cooling system operates non-stop, reducing (at an exponential rate) the deviation of the processor's temperature from the ambient temperature.

As a result, the processor's temperature, and thus the performance as well, depends on the order of the task execution. Given a variety of possible underlying architectures, models for cooling and for hardware thermal management, as well

as types of tasks, this gives rise to a plethora of interesting and never studied scheduling problems.

We focus on scheduling real-time jobs in a simplified model for cooling and thermal management.

A collection of unit-length jobs is given, each job with a deadline and heat contribution. If, at some time step, the temperature of the system is t and the processor executes a job with heat contribution h , then the temperature at the next step is $(t + h)/2$. If the temperature exceeds the given thermal threshold τ , the processor stays idle. The objective is to maximize the throughput, that is, the number of tasks that meet their deadlines. We prove that in the offline case computing the optimum schedule is NP-hard. In the online case, we show a 2-competitive deterministic algorithm and a matching lower bound.

Joint work of: Chrobak, Marek; Durr, Christoph; Hurand, Mathilde; Robert, Julien

Scheduling Orders to Minimize Sum of Weigthed Completion Times

José R. Correa (Universidad Adolfo Ibáñez - Santiago, RCH)

We consider a general scheduling problem in which orders, consisting of several jobs, have to be processed in unrelated machines. The goal is to find a schedule minimizing the sum of weighted completion time of orders. We give a constant factor approximation algorithm for this problem, improving and extending a recent result by J. Leung, H. Li, M. Pinedo and J. Zhang. We also study the parallel machine environment and design polynomial time approximation schemes for special cases of the problem.

Joint work of: Correa, José R.; Skutella, Martin; Verschae, José

Validating schedulability analyses and some other open problems in stochastic real-time scheduling

Liliana Cucu (INRIA - Nancy, F)

Validating feasibility and schedulability analyses in stochastic real-time scheduling can be done either by Monte-Carlo simulations, either analytically. Unfortunately, simulation is not well suited to estimate rare events (e.g. less frequent than 10^4) because of the size of the sample that is needed to achieve reasonable error bounds. Central Limit Theorem tell us that the convergence rate is of order $n^{1/2}$, where n is the number of random draws, which means that adding one significant digit requires increasing n by a factor 100. Therefore we believe that an analytical validation method must be always preferred when one wants to validate feasibility and schedulability analyses in stochastic real-time scheduling.

Keywords: Real-time, stochastic, schedulability and feasibility analyses

One (short) open problem in real-time multiprocessor scheduling + One (short) answer on what predictability is for real-time multiprocessor scheduling

Liliana Cucu (INRIA - Nancy, F)

For the problem of scheduling parallel tasks on identical processors we had provided a polynomial algorithm using a 2D solution. The problem on unrelated processors is still open and one should try a LP formulation.

When one wants to provide feasibility intervals based on periodicity results for schedules of periodic tasks on unrelated procesors, then she/he has to take into account scheduling anomalies regardless to the execution times. By proving that a scheduling algorithm is predictable we ensure the existence of a feasibility interval.

Keywords: Multiprocessor, real-time, scheduling, predictability, LP formulation

Full Paper:

http://www.loria.fr/~cuculili/rtns07_cucu.pdf

Polynomial Time Algorithms for Minimum Energy Scheduling

Christoph Durr (Ecole Polytechnique - Palaiseau, F)

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: Dynamical programming, minimizing energy

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

See also: Proc. of the 15th Annual European Symposium on Algorithms (ESA), 136-150, 2007

Broadcast Scheduling: Algorithms and Complexity

Thomas Erlebach (University of Leicester, GB)

Broadcast Scheduling is a popular method for disseminating information in response to client requests. There is a set of pages of information, and clients request pages at different times. However, multiple clients can have their requests satisfied by a single broadcast of the requested page. We consider several related broadcast scheduling problems. One central problem asks to minimize the maximum response time (over all requests). Another related problem is the version in which every request has a release time and a deadline, and the goal is to maximize the number of requests that meet their deadlines. While approximation algorithms for both these problems were proposed several years back, it was not known if they were NP-complete. One of our main results is that both these problems are NP-complete. Furthermore, we give a proof that FIFO is a 2-competitive online algorithm for minimizing the maximum response time and that there is no better deterministic online algorithm (these results had been claimed earlier, but without complete proofs). We also give a lower bound on the integrality gap of the natural LP formulation of the problem of maximizing the number of requests that meet their deadlines.

Keywords: NP-completeness, competitive algorithm, lower bound, FIFO

Joint work of: Chang, Jessica; Erlebach, Thomas; Gailis, Renars; Khuller, Samir

See also: In Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, California, January 20 - 22, 2008). Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics, Philadelphia, PA, 473-482.

Making good rosters for a 24/7 environment

Han Hoogeveen (Utrecht University, NL)

Our problem is to find rosters for the personnel of some department at UMC. The department has 35 workers and works 24/7. There are three shifts: Early, Late, and Night; the required number of persons varies per shift. Except for a number of part-timers, each worker has a contract for 36 hours per week, but gets approximately 34 hours per week on average to compensate beforehand for extra work that has to be carried out to replace people that are ill. Next to the ordinary working shifts, there are the so-called stand-by shifts. It is allowed to deviate from the desired number of workers per shift. Night shifts and weekend shifts

should not be over-booked; a shortage is covered by hiring outsiders. For the early and late shifts during week-days, we can have more workers than necessary. This is even encouraged during the Wednesday morning shift, since the superfluous workers can then follow a training session.

Currently, the department works with non-personalized, cyclic rosters. Our task is to produce a set of better rosters that take into account the personal preferences of the personnel; each roster covers a period of one year. These rosters have to satisfy several regulations. For instance, each roster should consist of a number of early shifts, followed by a number of late shifts, and finally by a set of night shifts, with a sufficient number of days off in between. The rosters have to be personalized, such that they reflect the personal preferences as much as possible; we wish to maximize the total satisfaction, with the side-constraint that the unluckiest person is not extremely unlucky. The satisfaction level is mostly determined by issues like having a fixed day off each week and the number of days off following a night shift. We further want to find a roster in which the total weighted deviation from the desired number of personnel is minimum; thereto, we minimize the weighted sum of both components.

We have solved this problem by applying column generation in combination with a rolling horizon approach, because of the size of the problem. In this way we have found a large pool of good, feasible rosters for each person, from which we pick one per worker, such that the combination is optimal. This program took about 4 hours and found an almost optimal solution. Our approach can be applied to larger problem instances without increasing the running time, and the results even get better.

Joint work of: Hoogeveen, Han; Penninx, Eelko

Approximation Algorithms for 2-dimensional packing problems

Klaus Jansen (Universität Kiel, D)

During the last years there has been a lot of interest in two dimensional packing problems, like 2D strip packing, 2D bin packing and 2D knapsack. These problems play an important role in diverse applications like cutting stock, VLSI layout, internet advertisement, and multiprocessor scheduling with parallel tasks.

The 2D strip packing problem can be described as follows: Let R be a set of n rectangles with width and height at most 1. The problem is to pack all the rectangles R into a strip of unit width and minimum length. We allow only orthogonal packings; this means that the rectangles must not overlap and their sides must be parallel to the sides of the strip. In the first of the talk we describe shortly classical algorithms for the strip packing problem including the asymptotic fully polynomial time approximation scheme (AFPTAS) by Kenyon and Remila. The AFPTAS generates for each error guarantee $\epsilon > 0$ a packing with height at most $(1 + \epsilon)OPT + O(1/\epsilon^2)$ in time polynomial in n and $1/\epsilon$. In the

second part we describe a new developed approximation algorithm (in joint work with Roberto Solis-Oba) where we improve the additive constant from $O(1/\epsilon^2)$ to 1. The algorithm was derived from the study of the 2D knapsack problem. Given a set R of rectangles with positive profits, the goal is to pack a subset of R into a unit size square $[0, 1] \times [0, 1]$ so that the total profit of the selected rectangles is maximized. We present an algorithm that for any $\epsilon > 0$ finds a subset $R' \subset R$ of rectangles with total profit at least $(1 - \epsilon)OPT$, where OPT is the profit of an optimum solution, and packs them into the augmented bin $[0, 1] \times [0, 1 + \epsilon]$.

Finally we discuss several important open question and describe our further new results.

Keywords: Scheduling with parallel jobs, 2D strip packing, 2D knapsack

Average-Case and Smoothed Competitive Analysis of the Multi Level Feedback problem

Alberto Marchetti-Spaccamela (University of Rome "La Sapienza", I)

In this paper, we introduce the notion of smoothed competitive analysis of online algorithms. Smoothed analysis has been proposed by Spielman and Teng to explain the behavior of algorithms that work well in practice while performing very poorly from a worst-case analysis point of view. We apply this notion to analyze the multilevel feedback algorithm (MLF) to minimize the total flow time on a sequence of jobs released over time when the processing time of a job is only known at time of completion.

It is known that MLF has a worst case competitive ratio that is $\Omega(p_{\max})$ where p_{\max} denotes the longest processing time. We use a partial bit randomization model, i.e., the initial processing times are smoothed by changing the k least significant bits under a quite general class of probability distributions. We show that MLF admits a smoothed competitive ratio of $O((2^k/\sigma)^3 + (2^k/\sigma)^2 p_{\max}/2^k)$ where σ denotes the standard deviation of the distribution. A direct consequence of our result is the first average-case analysis of MLF. We show a constant expected ratio of the total flow time of MLF to the optimum under several distributions including the uniform one. We also prove an $\Omega(p_{\max}/2^k)$ lower bound for any deterministic algorithm that is run on processing times smoothed according to the partial bit randomization model.

For various other smoothing models, including the additive symmetric smoothing one, which is a variant of the model used by Spielman and Teng, we give a higher lower bound of $\Omega(p_{\max})$.

Keywords: Smooth analysis, non clairvoyant scheduling

Full Paper:

<http://www.personeel.unimaas.nl/t.vredeveld/PUB/MLF-MOR.pdf>

Single machine precedence constrained scheduling is a vertex cover problem

Monaldo Mastrolilli (IDSIA - Lugano, CH)

We consider the single machine precedence constrained scheduling problem of minimizing the sum of weighted completion time ($1|\text{prec}|\sum w_j C_j$). We settle an open problem first raised by Chudak & Hochbaum and whose answer was subsequently conjectured by Correa & Schulz. Namely we prove that the addressed scheduling problem reduces to solving a vertex cover problem on the well-known graph of incomparable pairs, which is induced by the precedence constraints. As a consequence of our result, previous results for the scheduling problem can be explained, and in some cases improved. In particular, by combining techniques from vertex cover and dimension theory of partial orders it is possible to obtain the best known approximation ratios for all classes of precedence constraints considered so far.

Keywords: Precedence constraints, vertex cover, dimension theory of partial orders, approximation algorithms

Joint work of: Ambuehl, Christoph; Mastrolilli, Monaldo; Mutsanas, Nikos; Svensson, Ola

Approximation Results for Preemptive Stochastic Online Scheduling

Nicole Megow (TU Berlin, D)

We present first constant performance guarantees for preemptive stochastic scheduling to minimize the sum of weighted completion times. For scheduling jobs with release dates on identical parallel machines we derive policies with a guaranteed performance ratio of 2 which matches the currently best known result for the corresponding deterministic online problem.

Our policies apply to the recently introduced stochastic online scheduling model in which jobs arrive online over time. In contrast to the previously considered nonpreemptive setting, our preemptive policies extensively utilize information on processing time distributions other than the first (and second) moments. In order to derive our results we introduce a new nontrivial lower bound on the expected value of an unknown optimal policy that we derive from an optimal policy for the basic problem on a single machine without release dates. This problem is known to be solved optimally by a Gittins index priority rule. This priority index also inspires the design of our policies.

Keywords: Stochastic online scheduling, sum of completion times

Joint work of: Megow, Nicole; Vredeveld, Tjark

Two-phase greedy algorithms for some classes of combinatorial linear programs

Britta Peis (TU Berlin, D)

We present greedy algorithms for some classes of combinatorial packing and cover problems within the general formal framework of Hoffman and Schwartz' lattice polyhedra. Our algorithms compute in a first phase Monge solutions for the associated dual cover and packing problems and then proceed to construct greedy solutions for the primal problems in a second phase. We show optimality of the algorithms under certain sub- and supermodular assumptions and monotone constraints. For supermodular lattice polyhedra with submodular constraints, our algorithms offer the farthest reaching generalization of Edmonds' polymatroid greedy algorithm currently known.

Keywords: Lattice polyhedra, submodular functions

Joint work of: Peis, Britta; Faigle, Ulrich

A Survey of Competitive Analysis of Online Scheduling Algorithms

Kirk Pruhs (University of Pittsburgh, USA)

Tentatively I plan to give a survey on competitive analysis of online scheduling algorithms.

Non-Clairvoyant Scheduling with Precedence Constraints

Julien Robert (ENS - Lyon, F)

We consider Edmonds's model (1999) extended by precedence constraints. In our setting, a scheduler has to schedule non-clairvoyantly jobs consisting in DAGs of tasks arriving over time, each task going through phases of different degrees of parallelism, unknown to the scheduler. As in the original model without precedence constraints, the scheduler is only informed of the arrival and the completion of each task, at the time of these events, and nothing more.

Furthermore, it is not aware of the DAG structure of each job beforehand neither of the precise characteristics of the phases of the tasks that compose each job.

We consider the preemptive strategy EquioEqui, that divides the processors evenly among the alive jobs and then divides the processing power allotted to each job evenly among its alive tasks. We show that whatever how complex the precedences are, EquioEqui is $(2 + \epsilon)$ -speed $O(k/\epsilon)$ -competitive for the flowtime metric, where k is the maximum number of independent tasks in

each job. That is to say, the flowtime of the schedule computed by EquioEqui is at a constant ratio of the optimal flowtime as soon as Equi is given slightly more than twice the resources as the optimum it is compared to. Interestingly, the extra speed needed to obtain a competitive algorithm, namely $(2+\epsilon)$, is the same in presence of precedence constraints, as in the original setting without precedences studied by Edmonds in 1999. This means that the maximum load that the system can handle without diverging, is the same with or without precedence constraints.

Furthermore, we propose a simple scheme to analyze a special class of schedulers, namely Equi-schedulers, which allows to obtain upper and lower bounds on particular precedences structures, such as independent chains, IN-trees, OUT-trees and Serial-parallel DAGs.

Keywords: Online scheduling, Precedences, Non-clairvoyant algorithm, Fairness, Equi-partition

Joint work of: Robert, Julien; Schabanel, Nicolas

A bi-criteria PTAS for Real-time Scheduling with fixed priorities

Thomas Rothvoss (Univ. Paderborn, D)

We present a bi-criteria polynomial time approximation scheme for the real-time scheduling problem with fixed priorities.

Thereby tasks T_1, \dots, T_n are given, such that task T_i generates a job with running time c_i each p_i time units, which has to be completed until the period p_i ends.

This scheduling problem was described by Liu and Layland more than 30 years ago and has received a lot of attention, especially in the real-time and embedded-systems community.

For a fixed $\epsilon > 0$, our algorithm runs in polynomial time and finds an assignment using at most $(1 + \epsilon) * OPT + O(1)$ processors which is feasible if the processors have speed $1 + \epsilon$. The formerly best known algorithm gave a $7/4$ -approximation.

Joint work of: Rothvoss, Thomas; Eisenbrand, Fritz

Truthful Mechanisms for Scheduling Problems

Guido Schäfer (TU Berlin, D)

We survey recent advances that have been made in devising truthful cost sharing mechanisms for classical scheduling problems, such as makespan, completion time and flow time scheduling.

In the first part of the talk, we consider Moulin mechanisms that achieve a very strong notion of truthfulness. We present some negative results that show that these mechanisms inherently lead to poor approximation guarantees in the scheduling context.

In the second part of the talk, we introduce a novel class of cost sharing mechanisms, so-called acyclic mechanisms. These mechanisms achieve a slightly weaker notion of truthfulness than Moulin mechanisms, but provide additional means to improve upon the approximation guarantees.

Joint work with Janina Brenner, TU Berlin.

Graph Balancing: A Special Case of Scheduling Unrelated Parallel Machines

Jiri Sgall (Academy of Sciences - Prague, CZ)

We design a 1.75-approximation algorithm for a special case of scheduling parallel machines to minimize the makespan, namely the case where each job can be assigned to at most two machines, with the same processing time on either machine. (This is a special case of so-called restricted assignment, where the set of eligible machines can be arbitrary for each job.) We also show that even for this special case it is *NP*-hard to compute better than 1.5 approximation.

This is the first improvement of the approximation ratio 2 of Lenstra, Shmoys, and Tardos [Approximation algorithms for scheduling unrelated parallel machines, *Math. Program.*, 46:259–271, 1990], for any special case with unbounded number of machines. Our lower bound yields the same ratio as their lower bound which works for restricted assignment, and which is still the state-of-the-art lower bound even for the most general case.

Joint work of: Ebenlendr, Tomas; Krcal, Marek; Sgall, Jiri

Challenges for Multi-function SoC Scheduling and Assistive Living Services

Chi-Sheng Shih (National Taiwan University, TW)

Tasks on assistive living devices and multi-function SoC (MFSoc) have the requirements to meet certain timing constraint. Examples include the end-to-end throughput constraint on MFSoc and deadline constraint on assistive living devices. However, the workload model on the two application domains are different from conventional real-time workload models. The new workload model, hence, impose new scheduling challenges on real-time computing community.

In this short paper, we will discuss the new workload models and the challenging issues on scheduling such tasks.

Keywords: Real-time Task Scheduling

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1486>

Tradeoffs and Average-Case Equilibria in Selfish Scheduling

Alexander Souza (Universität Freiburg, D)

We consider the prices of anarchy (PoA) and stability (PoS) of selfish scheduling on uniformly related machines. PoA relates the value of a Nash equilibrium with highest cost to the optimum value. Similarly for PoS, where the Nash equilibrium with lowest cost is compared. Each of the selfish players controls a job and seeks to assign it to one of the machines as to minimize its experienced latency.

In the majority of related work, the cost of a solution is the maximum latency. In contrast, we value solutions by the sum of player-latencies. We show that PoA and PoS are $\Theta(n/t)$, where n is the number of players and t the sum of job-lengths. This situation also transfers to a stochastic model, where job lengths are random variables and we are interested in the expected PoA. As an extension, we introduce machine restrictions: Now, each player is allowed to assign its job to a subset of the machines only. For identical machines the PoA and PoS are $\Theta(n\sqrt{m}/t)$, where m is the number of machines.

These results can also be interpreted from the perspective of classical optimization:

The expected PoA can be seen as the average-case performance of a local search algorithm.

Latency problems with profits

Frits C.R. Spieksma (Katholieke Universiteit Leuven, B)

In the latency problem with profits a set of clients, each with a corresponding profit p_i , and a server are given. The server can collect profits at the clients and these profits are equal to $p_i - t$, where t denotes the time at which the server reaches client i . The goal is find a route for the server such that total collected profit is maximized. We focus on the variant where the clients are located on a line, and establish the complexity of different variants of this problem.

Keywords: Latency, complexity

Joint work of: Coene, S; Spieksma, Frits C.R.

Non-Preemptive Min-Sum Scheduling with Resource Augmentation

Clifford Stein (Columbia University, USA)

We give the first $O(1)$ -speed $O(1)$ -approximation polynomial-time algorithms for several nonpreemptive min-sum scheduling problems where jobs arrive over time and must be processed on one machine.

More precisely, we give the first $O(1)$ -speed $O(1)$ -approximations for the non-preemptive scheduling problems

- $1 \mid r_j \mid \sum w_j F_j$ (weighted flow time), - $1 \mid r_j \mid \sum T_j$ (total tardiness), - the broadcast version of $1 \mid r_j \mid \sum w_j F_j$,
 an $O(1)$ -speed, 1-approximation for $1 \mid r_j \mid \sum \bar{U}_j$ (throughput maximization),
 and an $O(1)$ -machine, $O(1)$ -speed $O(1)$ -approximation for $1 \mid r_j \mid \sum w_j T_j$ (weighted tardiness).

Our main contribution is an integer programming formulation whose relaxation is sufficiently close to the integer optimum, and which can be transformed to a schedule on a faster machine.

Joint work of: Stein, Clifford; Bansal, Nikhil; Chan, Ho-Leung; Khandekar, Rohit; Pruhs, Kirk; Scheiber, Baruch

Optimal Mechanisms for Scheduling Jobs on a Single Machine

Marc Uetz (University of Twente, NL)

The design of optimal auctions is recognized as an intriguing issue in auction theory; first studied by Myerson (1981) for the case of single item auctions. In that setting, the goal is to maximize the seller's revenue. We study the design of optimal auctions (more precisely, mechanisms) in a setting where job-agents compete for being processed on a single machine that can only handle one job at a time. Each job has a service time and a weight representing the disutility for waiting. In this setting, it is folklore that the total disutility of the jobs is minimized by Smith's rule: schedule jobs in order of non-increasing ratios weight over service time.

While the jobs' service times are public information, we assume a job's weight is only known to the job itself. Given jobs' reports about their private information, a mechanism for this setting determines both an order in which jobs are served, and for each agent a payment that the agent receives from the provider. The payment can be seen as a reduction of service fees to compensate for waiting. Publicly known probability distributions represent common beliefs about other jobs' weights. We study two cases: discrete and continuous probability distributions.

We aim at finding Bayes-Nash incentive compatible mechanisms that are optimal, that is, minimize the total expected expenses of the system. By a graph theoretic interpretation of the incentive compatibility constraints -first used by Rochet (1987)- we are able to derive optimal mechanisms. For both discrete and continuous distributions of weights, we obtain closed formulae for §virtualĤ job weights, and show that serving the jobs in the order of non-increasing ratios of virtual weights over service times is optimal, as long as a certain regularity condition is fulfilled. It also turns out that the optimal mechanisms do not necessarily minimize total disutility, but they do so if e.g. all jobs are symmetric.

Joint work with Birgit Heydenreich, Debasis Mishra, and Rudolf Mueller

Keywords: Scheduling, Mechanisms Design, Optimal Auction, Bayes-Nash

SOS: Stochastic Online Scheduling

Tjark Vredeveld (Maastricht University, NL)

We consider a model for scheduling under uncertainty. In this model, we combine the main characteristics of online and stochastic scheduling in a simple and natural way. Job processing times are assumed to be stochastic, but in contrast to traditional stochastic scheduling models, we assume that jobs arrive online, and there is no knowledge about the jobs that will arrive in the future. The model incorporates both, stochastic scheduling and online scheduling as a special case. The particular setting we consider is non-preemptive parallel machine scheduling, with the objective to minimize the total weighted completion times of jobs. We analyze simple, combinatorial online scheduling policies for that model, and derive performance guarantees that match the currently best known performance guarantees for stochastic and online parallel machine scheduling. For processing times that follow NBUE distributions, we improve upon previously best known performance bounds from stochastic scheduling, even though we consider a more general setting.

Keywords: Stochastic, online, approximation,

Joint work of: Vredeveld, Tjark; Megow, Nicole; Uetz, Marc

Full Paper:

<http://mathor.highwire.org/cgi/content/abstract/31/3/513>

See also: Models and algorithms for stochastic online scheduling; Mathematics of Operations Research 31(3): 513-525 (2006)

Scheduling and Resource Allocation Issues in the Akamai Platform

Joel Wein (Polytechnic Univ. - New York, USA)

We'll discuss in an introductory fashion some real-world resource allocation questions that arise in the Akamai platform.

Online Scheduling with Reordering

Matthias Westermann (RWTH Aachen, D)

In the classic minimum makespan scheduling problem, we are given an input sequence of jobs with processing times.

A scheduling algorithm has to assign the jobs to m machines. The objective is to minimize the makespan, which is the time it takes until all jobs are processed. In this talk, we consider online scheduling algorithms without preemption. However, we do not require that each arriving job has to be assigned immediately to one of the machines. A reordering buffer with limited storage capacity can be used to reorder the input sequence in a restricted fashion so as to schedule the jobs with a smaller makespan. As main result, we present, for m identical machines, a lower bound of r_m on the competitive ratio of this problem with a reordering buffer whose size does not depend on the input sequence and a fairly simple scheduling algorithm matching this lower bound with a reordering buffer of size $\lceil (1 + 2/r_m) \cdot m \rceil + 2$. For example, $r_2 = 4/3$ and $\lim_{m \rightarrow \infty} r_m \approx 1.4659$. For the special case of two identical machines, we present a r_2 -competitive scheduling algorithm using only a buffer of size 2 which is the smallest buffer size allowing reordering. Further, we show that the general lower bound can be improved for buffer sizes below $m/2$.

Finally, we give a simple scheduling algorithm for m uniformly related machines, i. e., for m machines with different speeds, achieving a competitive ratio of $2 + \varepsilon$ with a reordering buffer of size m .

Keywords: Minimum makespan scheduling, online algorithms, reordering buffer

Joint work of: Englert, Matthias; Özmen, Denis; Westermann, Matthias

On Alcuin's scheduling problem

Gerhard Woeginger (TU Eindhoven, NL)

We consider a planning problem that generalizes Alcuin's river crossing problem (also known as: The wolf, goat, and cabbage puzzle) to scenarios with arbitrary conflict graphs. We derive a variety of combinatorial, structural, algorithmical, and complexity theoretical results around this problem.

Keywords: Transportation problem, scheduling and planning, graph theory

Joint work of: Csorba, Peter; Hurkens, Cor; Woeginger, Gerhard

Non-migratory Multi-processor Scheduling for Flow Time and Energy

Prudence W. H. Wong (University of Liverpool, GB)

Energy usage has been an important concern in recent research on online job scheduling, where processors are allowed to vary the speed dynamically so as to save energy whenever possible.

Notice that providing good quality of service such as deadline feasibility / flow time and conserving energy are conflicting objectives. For online algorithms, jobs arrive at arbitrary time and the job information is only known when a job arrives.

No future knowledge is given in advance. We aim to develop competitive on-line algorithms whose performance is not too far from the optimal off-line algorithm which has all information in advance.

The past two years have witnessed significant progress in the single-processor setting when energy is taken into concern, and online algorithms with performance close to optimal have been obtained.

For deadline feasibility, we are given a set of jobs each with a we extend the work to the two processor case.

We proposed an algorithm called Slow-SR which is based on the single processor algorithm Slow-D and the traditional deadline scheduling algorithm (without energy concern) Safe-Risky. Slow-SR is 3-competitive on throughput and $O(1)$ -competitive in energy.

We also extend the study of optimizing the tradeoff between flow time and energy to the multi-processor setting with $m \geq 2$ machines. We devise and analyze a simple non-migratory online algorithm that makes use of the classified-round-robin (CRR) strategy to dispatch jobs. If we allow speed relaxation, i.e., the on-line algorithm uses $(1 + \epsilon)$ times the maximum speed, we can obtain an $O(1 + 1/\epsilon)$ -competitive algorithm. On the other hand, if all the jobs are of power-of-two size and we do not allow speed relaxation, we have an algorithm which is $O(\log P)$ competitive, where P is the ratio of the maximum job size to the minimum job size. The competitive ratio in both cases hold even when the comparison is made against the optimal non-migratory off-line algorithm.

Joint work of: Lam, Tak-Wah; Lee, Lap-Kei; To, Isaac K. K.; Wong, Prudence W. H.

The price of anarchy on uniformly related machines revisited

Rob van Stee (Universität Karlsruhe, D)

Recent interest in Nash equilibria led to a study of the *price of anarchy* (poa) and the *strong price of anarchy* (spoa) for scheduling problems. The two measures express the worst case ratio between the cost of an equilibrium (a pure Nash equilibrium, and a strong equilibrium, respectively) to the cost of a social optimum.

We consider scheduling on uniformly related machines.

Here the atomic players are the jobs, and the delay of a job is the completion time of the machine running it, also called the load of this machine. The social goal is to minimize the maximum delay of any job, while the selfish goal of each job is to minimize its own delay, that is, the delay of the machine running it.

While previous studies either consider identical speed machines or an arbitrary number of speeds, focusing on the number of machines as a parameter, we consider the situation in which the number of different speeds is small. We reveal a linear dependence between the number of speeds and the poa. For a set of machines of at most p speeds, the poa turns out to be exactly $p + 1$. The growth of the poa for large numbers of related machines is therefore a direct result of the large number of potential speeds.

We further consider a well known structure of processors, where all machines are of the same speed except for one possibly faster machine. We investigate the poa as a function of both the speed of the fastest machine and the number of slow machines, and give tight bounds for nearly all cases.

Keywords: Price of anarchy, related machines

Joint work of: Epstein, Leah; van Stee, Rob