

New challenges in scheduling theory

Jacek Blazewicz^{1,2} · Benjamin Moseley³ · Erwin Pesch^{4,5} ·
Denis Trystram⁶ · Guochuan Zhang⁷

Published online: 30 November 2016
© Springer Science+Business Media New York 2016

Traditionally, machine scheduling is the study of the sequencing of tasks on a single or several parallel machines, possibly with different characteristics, subject to a set of constraints. The constraints express the limited amount of resources available, and the goal is to optimize an objective by capturing the quality of service delivered by the system. There is vast literature in this field, leading to numerous sophisticated heuristics and approximate and exact algorithms. Many of these problems are computationally intractable, and problems have been classified by their varying degrees of complexity.

The motivation for many problems has come from applications in production planning, staff rostering and personnel planning, scheduling in parallel and distributed systems, etc. [Blazewicz et al. \(2007\)](#). Over recent years, researchers have begun to study scheduling problems derived from new appli-

cations and settings. These applications include scheduling in decentralized systems and selfish organizations, in seaports or automotive production plants. Energy-efficient processing, fast data processing, and online scheduling are challenging applications. Some of these scheduling problems reflect real-life situations by including results from real industrial datasets. This list is only a small sample of the many new applications and scheduling problems that researchers are now studying.

Between March 31 and April 4, 2014, the 11th scheduling workshop, entitled “New Challenges in Scheduling Theory,” was held at the Centre CNRS Paul-Langevin, Aussois, France. (The participants really enjoyed the wonderful atmosphere and excellent organization.) The objective of the workshop was to explore new areas in scheduling theory and applications that have emerged in recent years.

The papers in this special issue highlight some of the results of that workshop. We received many high-quality submissions. After the normal thorough and rigorous refereeing process, only nine papers were chosen for inclusion in this special issue. The papers are listed below alphabetically by first author. The length of each summary does not reflect relevance or importance.

In the paper “Throughput maximization for speed-scaling with agreeable deadlines,” E. Angel, E. Bampis, V. Chau, and D. Letsios study the energy-efficient scheduling problem where a set of jobs has to be scheduled on a single machine, the speed of which can be changed dynamically. The objective is to finish as many jobs on time as possible, subject

Editing this issue would not have been possible without the help of many referees. We greatly appreciate their critical and encouraging evaluation of the submissions.

✉ Erwin Pesch
erwin.pesch@uni-siegen.de

Jacek Blazewicz
jblazewicz@cs.put.poznan.pl

Benjamin Moseley
bmoseley@wustl.edu

Denis Trystram
trystram@imag.fr

Guochuan Zhang
zgc@zju.edu.cn

¹ Institute of Computing Science, Poznan University of Technology, Poznan, Poland

² Institute of Bioorganic Chemistry, Polish Academy of Science, Poznan, Poland

³ Department of Computer Science and Engineering, Washington University, St. Louis, MO, USA

⁴ Faculty III, University of Siegen, Siegen, Germany

⁵ Center for Advanced Studies in Management, HHL Leipzig, Leipzig, Germany

⁶ Laboratoire d’Informatique de Grenoble, University Grenoble-Alpes, Grenoble, France

⁷ College of Computer Science, Zhejiang University, Hangzhou, China

to a limited energy budget that is not to be exceeded and some additional job-related constraints. The authors provide dynamic programming algorithms that solve certain problems with specific parameter settings optimally.

G. Aupy, M. Shantharam, A. Benoit, Y. Robert, and P. Raghavan are the authors of “Co-scheduling algorithms for high-throughput workload execution,” where they consider computing applications whose run time can be significantly reduced when a large number of processors are used in parallel. The idea is to partition the application into packages that are to be assigned to the processors such that the sum of the execution times of the packages is minimized.

“Optimizing memory allocation for multi stage scheduling including setup time,” the paper by A. Benoit, M. Coqblin, J.-M. Nicod, and V. Rehn-Sonigo, considers the execution of applications distributed in intervals of stages on a number of processors. There is a setup time between the execution of stages on the same processor. A buffer for storage of intermediate results can reduce the setup. The authors consider the problem where the buffer sizes are not given beforehand and have to be fixed before execution in order to maximize the throughput on each processor.

The paper “Necessary and sufficient optimality conditions for scheduling unit time jobs on identical parallel machines” by P. Brucker and N.V. Shakhlevich provides necessary and sufficient optimality conditions for scheduling unit time jobs on parallel machines. The optimality conditions are equivalent to finding negative cycles in a specific graph.

M. Burcea, W.-K. Hon, H.-H. Liu, P.W.H. Wong, and D.K.Y. Yau consider in their paper “Scheduling for electricity cost in smart grid” an offline situation where electricity consumers send a power request together with a flexible set of time slots during which their request could be served. The electricity cost is a convex function of the load in each slot. The problem is to schedule all requests with minimum total electricity cost. A polynomial-time algorithm finds an optimal solution for the case where requests have unit power requirement and unit duration, the jobs can be served in arbitrary time slots, and the cost function is a general convex function.

B. Chen, E. Coffman, D. Dereniowski, and W. Kubiak consider “Normal-form preemption sequences for an open problem in scheduling theory.” The authors focus on the structural properties of the problem $P2|pmtn; \text{in-tree}; r_j; p_j = 1|\sum C_j$, whose computational complexity is open. They introduce the concept of normal schedules where shifts, i.e., the sizes of intervals bounded by the times created by preemptions, job starts, or completions, decrease as a function of time and prove that there are optimal schedules that are normal for in-tree precedence constraints. They show that there is an optimal schedule with a finite number of events which are times when jobs start, finish, or are preempted.

They also detect that a job may be preempted neither at a start nor a completion of a job.

“Scheduling on parallel identical machines with late work criterion,” authored by X. Chen, M. Sterna, X. Han, and J. Blazewicz, is a paper on offline and online scheduling where the objective value is the amount of late work, i.e., the part that is scheduled after a due date. Only the fact that it is late matters, not the amount by which it is late. The authors show that the problem and the criterion of minimizing late work is NP-hard, and propose a dynamic programming approach for the case of two machines. For the online case, they propose an algorithm with a competitive ratio that is a function of the number of machines. For the case of two machines and where the total size of all jobs is known, they prove a tight bound and that the suggested algorithm is optimal.

In the paper “Scheduling with time dependent discrepancy times,” F. Jaehn and H. Sedding describe the discrepancy time minimization problem as a scheduling problem for a steadily moving assembly line where a worker is supposed to perform a number of jobs which require material. The material for each job is positioned on shelves and needs to be collected by the worker while the line is moving. Thus, the walking distances to the shelves are time dependent. As repositioning the shelves to keep the walking distance of the worker for each job minimal is not an option, the question is how to sequence the jobs such that the makespan is minimized. Complexity proofs and a discussion of the structural properties of optimal schedules are supplemented by dynamic programming and a branch-and-bound algorithm.

Finally, the paper “On scheduling with non-increasing time slot cost to minimize total weighted completion time” by Y. Zhao, X. Qi, and M. Li deals with the problem of scheduling a set of jobs on a single machine. The planning horizon consists of multiple time slots whose use incurs different costs. The objective is to minimize the sum of the total weighted completion time and the total machine time slot cost, assuming that the time slot costs are non-increasing. They consider time slot costs decreasing linearly at a constant rate, in which case the problem can be solved by a simple sorting. In the case of accelerative decreasing time slot costs, a dynamic programming algorithm is developed. For the case of decelerative decreasing time slot costs, they propose a polynomial-time repetitive merge algorithm. Finally, they prove that the problem is strongly NP-hard for arbitrary non-increasing time slot costs.

Reference

- Blazewicz, J., Ecker, K., Pesch, E., Schmidt, G., & Weglarz, J. (2007). *Handbook on Scheduling*. Berlin: Springer.