New challenges in scheduling theory

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

Published online: 4 July 2018

© Springer Science+Business Media, LLC, part of Springer Nature 2018

Machine scheduling traditionally is the study of the sequencing of tasks on a single or several parallel or dedicated machines, with possibly different characteristics, subject to a set of constraints. Constraints commonly capture the limited availability of resources, reflect precedence relations between tasks or generally express some restrictions of processing tasks over time. Beyond satisfaction of the constraints, there is generally the goal to optimize an objective as a criterion of the quality of the solution delivered by some search method. There is a vast literature on this area of study, and it has led to numerous sophisticated heuristics, approximate and exact algorithms. Many of these problems are computationally intractable, and problems have been classified by their varying degrees of complexity.

The motivation of many problems has come from applications in staff rostering and personnel planning, scheduling in parallel and distributed systems and production plan-

ning (see Blazewicz et al. (2018)). Over recent years, researchers have started to study scheduling problems that are derived from new applications and settings. These applications include scheduling in decentralized systems and selfish organizations, in sea ports and 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 on real industrial datasets. This listing is only a small sample of the many new applications and scheduling problems that researchers are studying.

The selected papers not only offer valuable insights on different facets of this current trend, but they also pave the way for future developments. In addition to the challenges related to changing industrial requirements and market conditions, future research studies will need to address challenges posed by emerging technologies. Industry 4.0, also called the Internet of Things, is closely tied with the digitalization of industrial processes and equipment, cyber-physical systems and the capability of real-time big-data processing.

Between March 29 and April 2, 2016, the 12th 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 this workshop. We received many high-quality submissions. After the normal thorough and rigorous refereeing process of the Journal of Scheduling, only 6 papers were chosen for inclusion in this special issue. The papers are listed below alphabetically by the first author. The length of each summary does not reflect any relevancy or importance.

In the paper on shared processor scheduling by D. Dereniowski and W. Kubiak, a set of jobs is to be processed by agents. Alternatively a smaller part of each job can also be processed by a subcontractor. As a consequence, the agent’s processing time will shrink but there is a payment to the subcontractor, who cannot process jobs in parallel, which

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

Jacek Blazewicz
jblazewicz@cs.put.poznan.pl

Benjamin Moseley
moseleyb@andrew.cmu.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 Tepper School of Business, Carnegie Mellon University, Pittsburgh, USA

⁴ Faculty III, University of Siegen, Siegen, Germany

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

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

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

is proportional to its weighted processed share of the job. The subcontractor likes to maximize the total payoff while the agents' goal is to minimize their completion times. The authors highlight that in order to find an optimal schedule it is enough to consider synchronized schedules. In the case that for any pair of jobs 1 and 2 the weight of 1 is smaller than that of 2 and the reverse is true for the processing times, then they provide a fast optimization algorithm. For the general case, they describe a $1/2$ approximation of the same speed.

D. Kress, M. Barketau and E. Pesch address single-machine batch scheduling to minimize the total setup cost in the presence of deadlines. A set of jobs that can be partitioned in subsets of job-families is to be processed on a single machine. Whenever a job of a new family is started then there is a setup cost which is constant for all setups. There are no setup times, and each job has a deadline. The objective is to minimize the number of setups. The authors prove that the problem is strongly NP-hard followed by an approximation algorithm that approximates an optimal schedule by a factor equal to the number of families.

"Applying recoverable robustness to tackle disturbances in single machine scheduling problems," the paper by M. van Akker, H. Hoogeveen and J. Stoef, deals with the minimization of the number of tardy jobs on a single machine. In contrast to the well-known problem, there is a small probability that the processing times are a little longer, and therefore, jobs might be rejected which would not have happened before. It leads to recoverable robustness which is the idea of repairing an infeasible solution to get back to feasibility if processing times were slightly disturbed. Disturbances are given as a list of scenarios. Besides some complexity results, the authors prove some dominance criteria and evaluate exact methods, dynamic programming, branch and bound as well as branch and price.

The paper "Scheduling Fully Parallel Jobs" by K. Wang, V. Chau and M. Li deals with the well-studied problem of scheduling a set of independent jobs (called also tasks) on parallel identical machines targeting the minimization of the sum of weighted completion times. They focus on the variant in which the jobs are fully parallel (the jobs can be scheduled in parallel at the same moment). They provide new results for this variant, including the design of a PTAS for fixed m . They also consider the special case where the weights correspond exactly to the processing times, which is polynomially solvable with fixed m (using a dynamic programming scheme). Finally, they show that the greedy algorithm Largest-Ratio-First has a bounded approximation ratio.

In the paper "Speed Scaling Problems with Memory/Cache Consideration," by W. Wu, M. Li, K. Wang, H. Huang, and E. Chen, a combination of energy-aware speed scaling and caching is studied. In the model, a set of jobs needs to be scheduled by their deadline. Each job requires computing and memory resources. The processor used for computation can be dynamically speed-scaled such that the processor uses more energy to process the jobs faster. The memory resource is a cache, and each job requires time to access the cache that cannot be sped up. The goal is to schedule jobs by their deadline using the minimum energy possible. Optimizing a speed scalable processor and memory constraints have been studied separately, but this is one of the first papers to combine these constraints that are commonly faced in practice. The paper gives a series of results to understand the problem both online and offline.

In the paper "Online Scheduling of Moldable Parallel Tasks," D. Ye, D. Chen and G. Zhang study the online scheduling problem with moldable parallel tasks on m identical processors where the tasks arrive one by one. In this model, the number of processors to execute a task is determined by the algorithm and cannot change during the execution. The objective is to minimize the makespan (longest completion time over all tasks). The main result is to provide a constant competitive ratio for this problem based on a general result that transforms a competitive algorithm of ratio ρ to an algorithm with ratio 4ρ . As there exists an old bound with $\rho = 6.6623$, the corresponding ratio is 26.65, which is quite large. Finally, the authors developed an improved bound of 16.74 by smart adaptations of the baseline algorithm.

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.

References

- Blazewicz, J., Ecker, K., Pesch, E., Schmidt, G., Sterna, M., & Weglarz, J. (2018). *Handbook on scheduling*. Berlin: Springer.