DEPARTMENT OF ECONOMICS
AND BUSINESS ECONOMICS
AARHUS UNIVERSITY

Review

of the doctoral thesis entitled

"Scheduling under energy consumption limits"

submitted by István Módos to Czech Technical University in Prague, Faculty of Electrical Engineering, Department of Control Engineering

Department of Economics
and Business Economics

Simon Emde
Associate Professor

Date: 8 May 2021

Direct Tel.: +45 871 52378
E-mail: siem@econ.au.dk
Web: au.dk/en/siem@econ

Sender's CVR no 31119103

Topic and structure of the thesis

Mr. Módos's thesis deals with integrating electric energy consumption into classic scheduling problems. Specifically, the thesis considers the following baseline problem: Given is a set of machines, each of which must process a given set of machine-specific jobs without preemption. The objective is to find a conflict-free schedule for the parallel machines such that some regular objective (such as the makespan or the sum of completion times) is minimal. This is a very general parallel machine scheduling problem, which has a plethora of applications in diverse fields such as industrial production and task scheduling on processors.

These types of parallel machine scheduling problems have been widely studied for decades; however, the energy consumption of concurrently processed jobs has only comparatively rarely been considered. This is surprising insofar as energy consumption limits play an important role in practice: in industrial production – especially in energy-intensive industries like steelmaking – the total as well as the peak power consumption can be a critical concern. This thesis therefore enriches the classic scheduling problem by subdividing the planning horizon into a number of metering intervals. Each job is associated with a certain power consumption. Jobs that are processed during the same metering interval add to the total power draw in that interval, which must be below a certain maximum value. This extension is practically motivated by steel production, where peak power consumption inside a metering interval (e.g., 15 minutes) must remain below a contractual maximum, otherwise the steel plant owes a penalty fee to the electricity provider.

Besides this deterministic problem statement, the thesis also considers the robustness of the ensuing schedule. Specifically, if one or multiple jobs are delayed such that (part of) their processing interval moves into a later metering interval than originally planned, the solution should still be feasible, i.e., the energy consumption should still be below the maximum level in all intervals. This is can be achieved by injecting buffers of idle time in-between some jobs.

The author investigates this problem in multiple settings and from multiple angles (e.g., deterministic, stochastic, using diverse modelling and solution approaches). Parts of the author's research have already been published in well-regarded international peer-reviewed journals (*Computers and Operations Research*, *Computers and Industrial Engineering*, and *Constraints*) and in conference proceedings.

Evaluation of each chapter

The thesis is written in English. The main body of text, which is divided into six chapters, consists of 76 pages. The thesis also contains a comprehensive seven-page bibliography and the usual tables (of contents etc.).

Chapter 1 (9 pages) introduces the topic of the thesis and illuminates its relevance in the context of energy pricing. The chapter also outlines the structure of the remainder of the thesis and states the main research questions to be addressed.

The chapter does a good job of motivating the subject and outlining the research gap, although some of the information could use more qualification (e.g., on p. 10, "uncertainty scenarios [...] are used when the probability distribution of the uncertain events is either not known or is uniform").

Chapter 2 (6 pages) gives a formal description of the baseline problem that is analyzed throughout the rest of the thesis. The problem is defined in a concise, mathematical, yet readable fashion, and very nicely explained via a comprehensive illustrative example.

Chapter 3 (5 pages) lays the groundwork of the modelling and solution techniques that are later used to formalize the scheduling problem. Specifically, this chapter presents (mixed-)integer linear programming, constraint programming, and local search frameworks.

DEPARTMENT OF ECONOMICS
AND BUSINESS ECONOMICS
AARHUS UNIVERSITY

The brief introduction mentions the major aspects of the relevant techniques, but does not go into great detail. Given that this chapter only deals with basic modelling techniques, this is understandable, although the selection of subjects could have been motivated better. For instance, why is local search explained, but not (logic-based) Benders decomposition, which is also used in a later chapter?

Chapter 4 (29 pages) constitutes the main part of the thesis along with Chapter 5. It investigates the scheduling problem with maximum power consumption limits in a deterministic environment, i.e., when all parameters are known with certainty, and robustness of the solutions is immaterial. The chapter starts with a complexity analysis of several special cases. The results are important, not least of all because they can give a hint as to whether decomposition approaches may be viable.

The author then goes on to present several alternative modelling formulations, namely a constraint programming model, a disjunctive mixed-integer linear programming model (MILP), a time-indexed MILP, and iterative MILP, and an implicit MILP. The latter is particularly noteworthy, because it improves upon a previous formulation from the literature. Since default solvers working on these models can struggle with solving large instances in acceptable time, the author also develops a local search heuristic.

The approaches are meticulously tested on a diverse set of randomly generated problem instances. The results are quite interesting, because the author manages to analyze in great detail which modelling and solution approaches work well in what instance groups, also giving some additional insights into, e.g., convergence speed (Fig. 4.5). Especially the local search heuristic is shown to work quite efficiently for the large instances. However, this begs the question: Could these models and methods be adapted to work on instances where the power consumption of a job is not uniform during its processing interval (e.g., if the energy requirement is high during the initial heat-up phase but then trails off)?

Chapter 5 (24 pages) extends the model from the previous section by also considering stochasticity in the start times, i.e., random delays may be injected into the schedule during execution. This makes it so that an ostensibly feasible and optimal schedule may suddenly become infeasible because an energy-hungry job is pushed into the next metering interval, where it cannot be accommodated.

The goal in this chapter is to find a fat solution, i.e., a solution is considered feasible only if it is feasible in every single delay scenario. While this is a relatively simple way of dealing with stochasticity in a problem, it allows the author to develop a pseudo-polynomial algorithm that returns an optimal schedule for a given job sequence. This pseudo-polynomial scheme is then used in a decomposition approach based on logic-based Benders decomposition (LBBD). The idea behind this LBBD is to determine a baseline schedule (without considering random delays) in the master model, which is solved by an off-the-shelf solver. The subproblem's task is then to check whether the proposed schedule is feasible in all scenarios. A similar decomposition is also embedded in a branch-and-bound scheme, and, finally, a local search heuristic is implemented.

The solution methods are again tested in a comprehensive computational study, which shows that the local search heuristic works quite well across the board, whereas the exact methods are only suitable for small instances.

The decomposition approaches proposed in this chapter are quite cleverly designed, because they play to the strength of the relatively efficient algorithm for solving the subproblem. The eventual performance, however, is not stellar, maxing out at about 10 jobs, which is not a lot. The heuristic is again shown to perform well. Why do the exact approaches not perform better? Are there any ways to improve them?

Chapter 6 (3 pages) concludes the thesis and gives some pointers for future research.

Conclusion and recommendation

This doctoral thesis addresses an interesting and relevant research gap. While the modelling and solution techniques have partially been adapted from the literature, they are cleverly designed, analyzed, implemented, and methodically tested. The generated insights can be of interest for further academic study as well as implementation in practice. In terms of style, the thesis is well written and pleasant to read. The high quality of this work is also supported by the author's good publication record (three published papers in international peer-reviewed journals at the time of submission).

Overall, the thesis constitutes an independent and creative effort, which contributes to the state of knowledge and fully satisfies the requirements for a doctoral thesis. I recommend that the Faculty of Electrical Engineering accept this thesis.

Aarhus, May 8, 2021,

Simon Emde
Associate Professor