

Hiking in the scheduling landscape

Citation for published version (APA):

Buchem, M. (2022). *Hiking in the scheduling landscape: exact and approximation algorithms for parallel machines*. [Doctoral Thesis, Maastricht University]. Maastricht University.
<https://doi.org/10.26481/dis.20221129mb>

Document status and date:

Published: 01/01/2022

DOI:

[10.26481/dis.20221129mb](https://doi.org/10.26481/dis.20221129mb)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Summary

This dissertation entails the theoretical and empirical study of various scheduling problems from theory and practice. In Chapters 2-4, the focus lies on *deterministic scheduling problems* for which the input is known and available without any uncertainty. Chapters 5 and 6 study *scheduling problems with uncertainty*. In the following we summarize the findings of each individual chapter and outline the main conclusions.

In *Chapter 2*, we investigate a family of load balancing problems on identical parallel machines. As it is unlikely that efficient algorithms exist which solve these problems to optimality, we aim at polynomial time algorithms which find solutions of provably good quality. More precisely, we formalize the concept of additive approximation schemes. These approximation schemes offer a solution technique with a clear trade-off between the quality of the solution and the running time needed to find the solution. The challenge when developing additive approximation schemes is that standard techniques for multiplicative approximation schemes are not applicable. To overcome this challenge, we introduce a new relaxation of the considered load balancing problems. Our additive polynomial time approximation schemes are based on structural properties which allow us to solve this relaxation complement with a local search technique which finds a final schedule whose objective value is provably close to that of the solution to the relaxation.

The new techniques as well as the concept of additive approximation schemes offer promising directions for future research such as extending the techniques to other scheduling problems. Especially considering problems for which it is known that no multiplicative approximation scheme exists, the notion of additive approximation schemes may offer a pathway to non-trivial approximation algorithms.

Chapter 3 investigates a natural extension of the identical parallel machine environment by combining this problem with the notion of machine conflicts. These conflicts are relevant when machines are cooperating with a single server responsible for pre- and post-processing of jobs and a pair of conflicting machines may not access the server at the same time. We establish a close connection between this scheduling problem and the graph theoretical problem of finding a maximum induced bipartite subgraph when all jobs are identical and pre-processing, processing and post-processing of any job takes one time unit. This connection leads to a strong inapproximability result implying that for general conflict graphs we cannot achieve any non-trivial approximation guarantees, unless $\mathcal{P} = \mathcal{NP}$. Therefore, we restrict our attention to special graph classes which are relevant from an applied or theoretical point of view.

The results in this chapter open up various interesting directions for future research. An open question which remains is whether, in general, the special case with unit jobs can be solved in polynomial time whenever a specially structured maximum induced bipartite subgraph can be computed in polynomial time or whether there exist graph classes for which this is not possible. If the former is the case, the presented algorithm for bipartite graphs may be a first step towards a polynomial time algorithm for other graph classes. Furthermore, considering more general job parameters with simple conflict graph classes poses an interesting research direction to investigate the quality of classic scheduling algorithms in this context of machine conflicts. Moreover, investigating graph classes which capture machine conflicts caused by spatial proximity such as geometric intersection graphs poses an interesting avenue for future research.

In *Chapter 4*, we turn our attention towards the problem of minimizing the total (weighted) squared deviation from job specific due dates. We focus on the case where all jobs have the same processing time. In the unweighted setting we devised a polynomial time algorithm based on the insight that in an optimal schedule jobs are evenly distributed among machines. In the weighted setting, we developed a polyno-

mial time algorithm for a single common due date if the due date is sufficiently large. If the due date is small, we devised an algorithm which runs in time exponential in the number of machines. Furthermore, when jobs have distinct due dates we present an additive fully polynomial time approximation scheme if the number of distinct due dates is constant.

Multiple interesting questions for future research remain open in the weighted setting. First of all, settling the complexity for a single common restrictive due date is of interest. Furthermore, the setting of k distinct due dates in the weighted version remains interesting since for absolute deviation penalties this can be solved in polynomial time via an assignment problem. However, this uses the fact that optimal schedules use integral starting times when considering absolute deviation penalties, which is not true for squared deviation penalties. Hence, different techniques and insights are needed for this next step.

Chapter 5 entails theoretical and empirical analyses of stochastic online scheduling policies for the problem of minimizing total weighted expected completion time. In particular, we consider the case where machines are either fast or slow. We adapt stochastic online scheduling policies which were previously known for the identical and unrelated machine environment by taking into account the machine speeds. From a theoretical perspective, we derive a performance guarantee for each of the policies in the online-list as well as the online-time model. In the online-list model, we further prove that the policy is asymptotically optimal. From a practical point of view, we analyse the realized performance of the policies compared to the theoretical performance. This analysis shows that the true performance is better than expected based on the theoretical results. The most striking result of the computational study is that the two lower bounds used in the theoretical performance analysis of the policies do not yield better lower bounds than the trivial lower bound of the weighted sum of expected processing times divided by the speed of the fast machine for small instances.

For future research, considering more general speed models is of interest. Furthermore, it remains an open question whether one can show that special policies for the uniform parallel machine environment can yield performance guarantees strictly better than those implied by the more general setting of unrelated parallel machines even for the special case considered in this chapter. This, however, would require new theoretical lower bounds or a different analysis of the policies.

In *Chapter 6* we study a scheduling problem from the perspective of a single vessel occurring in inland waterway management. We define the problem of finding an optimal velocity policy for an individual vessel which faces uncertainty at the lock with the goal of minimizing the expected fuel consumption for traversing an inland waterway split into two segments by a single lock between them. To solve this problem, we first develop a dynamic program to find a near-optimal velocity policy. In addition, we introduce a simple class of policies based on the idea of a fixed arrival time at the lock and describe how to find the best policy among this class. We analyse these two techniques in a computational study by comparing the fuel consumption of the two considered policy to a naive policy ignoring the uncertainty at the lock. The computational study shows, that taking into account uncertainty is vital and significant fuel savings can be achieved by both policies. In fact, we show that the simple policy choosing a fixed arrival time already leads to a significant share of the possible fuel savings due to a (near-)optimal velocity policy.

Many interesting avenues for future research exist based on the problem and results in this chapter. First of all, analysing the theoretical performance of simple velocity policies such as the fixed arrival policy may be of interest to establish theoretical performance guarantees. Furthermore, extending the model to more general settings such as a sequence of multiple locks or multiple vessels approaching the lock is a next step towards a better understanding and implementation of inland waterway transportation.