Approximation Algorithms in Allocation, Scheduling and Pricing

Citation for published version (APA):

Oosterwijk, T. (2018). *Approximation Algorithms in Allocation, Scheduling and Pricing*. [Doctoral Thesis, Maastricht University]. Universitaire Pers Maastricht. https://doi.org/10.26481/dis.20180119to

Document status and date:
Published: 01/01/2018

DOI:
10.26481/dis.20180119to

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.

Download date: 04 Feb. 2023

VALORISATION

This chapter discusses the contribution of the research of this thesis to society. The first part outlines the contributions of the general fields, and the second part focuses more on the actual contents of the chapters.

This thesis deals with different optimization problems, in the fields of congestion models, high multiplicity scheduling, vector scheduling, and optimal stopping theory and posted price mechanisms. In every optimization problem, the goal is to minimize costs, maximize profit or to solve something as quickly as possible. The optimization of timetables for public transport, vehicle routing problems, scheduling problems and other applications has provided major efficiency improvements. Timetables are optimized to reduce travelling time; vehicle routes are optimized to reduce time, fuel costs and the impact on the environment; scheduling jobs on machines in factories or tasks in processes is optimized to improve efficiency and so on. The influence of mathematical optimization to society has been huge.

Although most optimization problems occurring in practice are intrinsically hard to solve to optimality, the impact of implementing non-optimal solutions is still huge. Approximation algorithms compute solutions whose value are not too far from the optimum. Though suboptimal, they can still provide a huge decrease in costs, time requirements, or fuel costs, having a positive impact on the users and the environment.

Chapter 2 deals with congestion models. Since there is a lot of congestion on road and data networks, this is a field with relevant applications to practice. By understanding where congestion comes from and the dynamics underlying it, congestion can be reduced to save costs, time and the negative impact on the environment. For example, drivers arrive sooner and cheaper at their destination, thereby using less fuel which is good for the environment. And by reducing congestion in internet networks, data packets arrive faster.

Chapter 3 and Chapter 4 are problems in the field of scheduling. This field is concerned with scheduling jobs on machines in general, but its applications are more varied and range from scheduling exams to scheduling tasks on assembly lines. Companies use techniques from this field to optimize their processes, educational institutes use scheduling methods to find timetables, timetables for public transport are constructed using scheduling, and so on. In all applications, time or costs are saved by optimizing the efficiency.

Finally, Chapter 5 considers problems in optimal stopping theory and posted price mechanisms, in which the objective is to maximize the expected reward or profit. Both fields are concerned with optimizing expected values given a random arriving sequence of events, such as applicants for a vacant position or customers for an item someone is selling. In particular, posted price mechanisms can be used in auctions and by companies who want to set personalized prices for their customers, in order to maximize their profit.

The remainder of the valorisation will outline more details about the contributions of the results in the different chapters.

Chapter 2. Polymatroid Congestion Models

The problem considered in Chapter 2 is to minimize a non-decreasing separable function over a polymatroid. A polymatroid is an abstract structure that incorporates many different structures like singletons, spanning trees in graphs and matroids. The results in this chapter show that efficiently finding an optimal solution in such a system is hard, i.e. there is an intrinsic hardness in the problem that hinders researchers to find an optimal solution. To complement this hardness result, the chapter provides an algorithm that finds the best possible solution, up to a constant factor, one can find in an efficient amount of time. Researchers in discrete optimization can use the method in this chapter whenever they need to minimize a function over a polymatroid to find a satisfactory solution.

Minimizing a function over such a general structure subject to common constraints has a wide variety of applications in different theoretical topics. This of course includes the perspective from which the chapter is written, namely polymatroid congestion models. Here, the strategic choices of a set of players induce a strategy profile whose social costs we are minimizing. Researchers working in congestion models can draw inspiration from the methods and analysis leading to the result of this chapter. Moreover, minimum cost solutions can serve as building blocks for other cost-efficient solutions in related problems.

In practice, the results of this chapter are mainly relevant in two different scenarios. The first one is in cooperative games, where players can collaborate to minimize their costs. Collaborations between different parties is happening more and more. Whenever the strategy space of every player can be modelled as a polymatroid, the result of this chapter can be applied. This is for example true when players need to connect (a subset of) different objects in a network, e.g. computers, servers, or physical locations. Using the algorithm of this chapter results in a lower total cost, saving e.g. money, time or fuel, which has a positive impact on the players and the environment.

The second scenario is the situation in which a planner can implement a solution all players adhere to. This can occur e.g. in internet protocols or when autonomous cars make use of the road network. In both applications, there is a central planner that implements some protocol or software that tries to find a satisfactory solution for everyone. Minimizing costs can result in lower electricity usage, costs, time usage and fuel consumption.

Chapter 3. High Multiplicity Scheduling

Chapter 3 is dedicated to scheduling a set of products on a single machine that can produce one product at a time. Every product is associated with a maximum production rate, a demand rate and holding costs per time unit. Moreover, there are sequencing costs that need to be paid when the machine switches production. The objective is to find a cyclic schedule that minimizes the average costs.

The research in this chapter was inspired by a problem in practice. A multinational textile company wished to find the optimal cycle length for their production, of only three types of lycra in extremely large quantities on a single machine. The results in this chapter directly relate to problems like this, where a small number of products need to be scheduled on a machine. For a small number of products, the chapter gives the best possible schedule. For a general number of products, we show the problem is hard to solve efficiently and present an algorithm that finds a solution whose costs are not too far off from the costs of the optimal solution. All companies with similar problems can apply the methods in this chapter to find a solution to their problem.

This is the first research that tackles the problem in this form. By including both the high multiplicity encoding of the input and the sequencing costs, the problem considered in the chapter is closer to problems occurring in reality. Moreover, this research is a starting point for research in more complex problems and the ideas and results can help other researchers and practitioners to find methods to solve problems with additional practical constraints.

Chapter 4. Vector Scheduling

In Chapter 4 the Vector Scheduling problem is considered. There, a set of jobs with certain requirements need to be scheduled on a set of machines such that no machine is overloaded. One can think e.g. of applications on a computer that both require some amount of memory and have some CPU requirements, that need to be scheduled on a set of computers such that every computer can handle the workload.

The theoretical contribution is clear. We show almost matching upper and lower bounds on the running time required to find an approximate solution to the problem. Thus, it almost settles the theoretical complexity of the problem in terms of its approximation schemes.

In practice, the main contribution of this chapter is the knowledge that computing near-optimal solutions to this problem can take a rather long time. The lower bounds we show on the running time of these approximation schemes are high, indicating that no resources should be wasted on trying to find near-optimal solutions. The running time of the approximation scheme presented in the chapter, though nearly matching the provided lower bounds, is still too large for practical purposes.

There is a wide range of problems occurring in practice that can be modelled as an instance of the Vector Scheduling problem. The example above seems the most natural, in which the objective is to schedule a set of jobs on a set of computers, such that no computer is overloaded. For all these problems, and extensions of the problem, the results in this chapter show that finding near-optimal solutions to this problem is simply too hard to do efficiently.

Chapter 5: Optimal Stopping and Posted Prices

Finally, Chapter 5 concerns itself with two similar problems in different fields. In optimal stopping theory, a gambler faces a finite sequence of non-negative random variables whose values become apparent upon arrival. At every arrival, the gambler can claim the realized value and stop, or reject the value and continue, with the goal of maximizing his expected reward. In posted price mechanisms, a seller has a single item to sell to a set of potential customers who arrive uniform at random. The seller can set individual prices for the customers and makes a profit equal to the price he set for the first customer whose valuation exceeded the price. In both fields, the chapter provides approximation algorithms for a non-adaptive and an adaptive version of both problems.

The problem in optimal stopping theory has applications in human resources management. Suppose a company needs to hire a new employee and has interviews with some applicants. During the interviews, they can assess the value of the applicant for the position. The algorithms from this chapter provide a way to decide which of the applicants to hire, with the objective of maximizing the value of the new employee.

This is true for any practical problem in which true values are learned upon the arrivals, and the goal is to maximize the value of the first arrival that is accepted.

The applications following from the second field of posted price mechanisms are quite natural. Whenever some person or company needs to sell one item and wants to make the highest revenue, they can apply the algorithms in this chapter to do so. An example of the non-adaptive scenario is a retailer that sends a direct mail to all its potential customers, where the first customer to accept the price offered to her gets to buy the item. An example of the adaptive scenario is the situation in which at the gate of an airport, the employees sequentially offer travellers an upgrade to business class, where the extra revenue they make equals the price offered to the first accepting traveller.