

Solving Large-Scale Dynamic Collaborative Vehicle Routing Problems An Auction-Based Multi-Agent Approach

Los, J.

DOI

[10.4233/uuid:db6851f3-0fce-4874-949b-e4fe0ec1cadd](https://doi.org/10.4233/uuid:db6851f3-0fce-4874-949b-e4fe0ec1cadd)

Publication date

2021

Document Version

Final published version

Citation (APA)

Los, J. (2021). *Solving Large-Scale Dynamic Collaborative Vehicle Routing Problems: An Auction-Based Multi-Agent Approach*. [Dissertation (TU Delft), Delft University of Technology].
<https://doi.org/10.4233/uuid:db6851f3-0fce-4874-949b-e4fe0ec1cadd>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Solving Large-Scale Dynamic Collaborative Vehicle Routing Problems

An Auction-Based Multi-Agent Approach

Johan Los

Solving Large-Scale Dynamic Collaborative Vehicle Routing Problems

An Auction-Based Multi-Agent Approach

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
maandag 1 november 2021 om 12:30 uur

door

Johannes LOS,
doctorandus in de kunstmatige intelligentie, Rijksuniversiteit Groningen,
geboren te Groningen.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof. dr. R.R. Negenborn	Technische Universiteit Delft, <i>promotor</i>
Dr. M.T.J. Spaan	Technische Universiteit Delft, <i>promotor</i>
Dr. F. Schulte	Technische Universiteit Delft, <i>copromotor</i>

Onafhankelijke leden:

Prof. dr. M. Gansterer	Universiteit van Klagenfurt, Oostenrijk
Prof. dr. R.D. van der Mei	Vrije Universiteit Amsterdam
Prof. dr. F.M. Brazier	Technische Universiteit Delft
Prof. dr. ir. L.A. Tavasszy	Technische Universiteit Delft
Prof. dr. ir. P.H.A.J.M van Gelder	Technische Universiteit Delft, <i>reservelid</i>



This thesis is supported by the project “Dynamic Fleet Management (P14-18 – project 3)” (project 14894) of the research programme i-CAVE, partly financed by the Netherlands Organization for Scientific Research (NWO), domain Applied and Engineering Sciences (TTW).

TRAIL Thesis Series no. T2021/26, the Netherlands Research School TRAIL
P.O. Box 5017, 2600 GA Delft, The Netherlands
E-mail: info@rsTRAIL.nl

Cover illustration: Johan Los
E-mail: johan_los@hotmail.com

ISBN 978-90-5584-301-5

Keywords: Collaborative vehicle routing, collaborative transportation, platform-based transportation, multi-agent system, logistics, dynamic pickup and delivery problem, dynamic fleet management.

Copyright © 2021 by Johan Los

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

Printed in the Netherlands

*Ze maakten dankbaar gebruik van alle computertechnieken die God hun schonk,
maar ze wisten dat de vooruitgang van het weten nooit de omvang van het
onbekende zou verminderen.*

Geert Mak
(Hoe God verdween uit Jorwerd, p. 130)

Voorwoord

Minstens zo belangrijk als het doel is de weg ernaartoe. Dat platgetreden pad moet men zoveel mogelijk vermijden, maar in het licht van dit proefschrift krijgt het tóch een nieuwe lading. Het vinden van de juiste weg en het bereiken van het doel hangen hier zozeer samen, op verschillende niveaus, dat de weg niet zomaar weggelaten kan worden.

Er zijn legio parallellen te trekken tussen het onderwerp van dit proefschrift en het voltooien ervan. Er is een doel, er is een uitgangspunt, er zijn, binnen zekere grenzen, oneindig veel routes mogelijk. Er wordt een plan gemaakt, een route gekozen, men gaat op weg, en haast voordat men vertrokken is wordt het plan alweer bijgesteld. Op willekeurige momenten is er nieuwe informatie beschikbaar, met soms kleine, soms grote gevolgen voor de te volgen route. Halverwege kan er zomaar een nieuwe weg ingeslagen worden. Het project groeit, er moet samengewerkt worden, maar de verschillende partijen hebben verschillende doelen, verschillende niveaus van communicatie en verschillende mates van autonomie. Zo nu en dan worden er uit strategisch oogpunt dingen verzwegen of juist mooier voorgesteld dan ze in werkelijkheid zijn. Soms is men de weg even helemaal kwijt. Maar toch tracht men, rekening houdend met ieders voorkeuren, uit de verschillende alternatieven steeds de beste route te kiezen.

Of er nu een proefschrift of een bestelling afgeleverd moet worden, uiteindelijk bereikt men zijn bestemming wel. Met wat omwegen, uiteraard, maar die blijken meestal niet nutteloos. Hoewel de optimale route vaak het onderspit delft, leiden alle wegen naar Rome.

En zo komt er een eind aan het fietsen tussen de twee steden die in de vorige zin zo mooi in elkaars buurt gepositioneerd zijn. Tijd om aan beide einden van de Vliet dank te betuigen.

Eerst zuidwaarts.

Rudy, dank voor je uitnodiging om in te stappen in i-CAVE en voor je vertrouwen om mij aan het stuur van dit deelproject te zetten. Dank voor je stimulatie van het gebruik van heuristieken in het gehele proces: eerst maar eens een eind

de goede richting op zonder te weten waarin het ging uitmonden werkte in deze dynamische context beter dan 4 à 5 jaar plannen zonder daadwerkelijk vooruit te gaan. Je hebt me er ook succesvol van weerhouden om in de open probleemvariant verzeild te raken: in plaats van alleen maar verder te zoeken zijn we nu steeds netjes teruggegaan om verslag uit te brengen.

Matthijs, dank voor je hulp bij het herhaaldelijk zoeken van het juiste landweggetje en het aanzwengelen van de motor als die weer eens was afgeslagen. Als we Rudy's begeleiding vergelijken met de planning van het vervoer tussen grote knooppunten, was jouw expertise juist het *last-mile* transport. Dank ook voor je mooie verhalen, het sporadisch gezamenlijk fietsen en het kletsen over de kinderen.

Frederik, dank voor het wijzen op verscheidene nog onbekende maar veelbelovende richtingen. Dankzij jou is samenwerking niet alleen een hoofdonderwerp van dit proefschrift geworden, je was ook de drijvende kracht achter de samenwerking met onze Oostenrijkse collega's. Je hebt me niet alleen op het belang van het delen van informatie gewezen, maar me ook regelmatig een relevant artikel en bovendien een goede dataset gestuurd. Je hebt me niet alleen in de richting van het bestuderen van strategisch gedrag gestuurd, je hebt zelf ook herhaaldelijk mooie staaltjes van constructief strategisch gedrag laten zien, waarbij je niet alleen het voordeel voor jezelf, maar ook juist voor alle andere betrokkenen zocht.

Speciale dank aan Margaretha Gansterer en Richard Hartl voor de uitnodiging om in Wenen onze onderzoeksplannen te bespreken en voor de daaruit voortgekomen samenwerking. Het was interessant om onze verschillende richtingen met elkaar te vergelijken.

Dank aan de promotiecommissie voor het instappen aan het eind van de rit om het gehele proefschrift door te nemen en aan de gebruikerscommissie van het project *Dynamic Fleet Management* en de overige leden van i-CAVE voor de verscheidene discussies.

Dank aan alle collega's van de afdeling *Maritime and Transport Technology* voor de hilariteit, het gemopper en het overige algemene kantoorleven. Dank aan de huidige en voormalige collega's van de sectie *Transport Engineering and Logistics* voor de inhoudelijke besprekingen en het gezamenlijke nadenken. In het bijzonder dank aan Breno, Marc en Lindert. Breno, het was fijn om met jou op hetzelfde project te werken, zo nu en dan eens een goed gesprek daarover te voeren (maar vaker nog over belangrijker dingen), samen naar bijeenkomsten af te reizen en daar je enorme eetlust te aanschouwen. Marc, ik heb het gezamenlijk theedrinken in de koffiehoeck zeer gewaardeerd. Ik bewonder je gematigdheid, mildheid, en je vermogen om te accepteren dat dingen moeilijker of complexer zijn dan menigeen wil toegeven. Lindert, het was fijn om met jou

in een kantoor te zitten en regelmatig opborrelende willekeurige malligheden of serieuze aangelegenheden met je te bespreken. Je kennis van en inzicht in de grotemensenwereld is memorabel.

En dan gelukkig weer noordwaarts.

Dank aan mijn familie voor alle praktische hulp en alle briljantheid, en in het bijzonder dank aan mijn ouders. Papa en mama, dank voor al jullie steun, maar meer nog voor jullie levenshouding, rust en nuchterheid.

Matthias en Jakob, dank voor jullie geschater en geknuffel en dank voor het verbeteren van mijn kwaliteiten in *dynamic preemptive scheduling*. Matthias, dankjewel voor je bijna nooit afwezige vrolijkheid en je geduld. Jakob, dankjewel voor al je goede vragen en hypothesen en je prachtig geformuleerde volzinnen. En dankjewel dat je onszelf altijd een goede reis wenst.

Tjitske, dankjewel voor het meefietsen tot aan de A12 op de eerste van november in 2016, het samen zoeken van onze weg in Leiden en het belachen van de talloze postbusjes in de straat. Dankjewel voor al je goede inzichten, je bemoedigingen en je ondersteuning. Dank voor het aanhoren van talrijke klagen en het meedenken over velerlei zaken. Je bent zo lief! Dankjewel dat je met mij op weg bent gegaan en ons ook steeds weer de juiste kant op stuurt. Ik zie ernaar uit om samen met jou de weg die ook de waarheid en het leven is te blijven volgen en zo onze bestemming te bereiken.

Johan Los

Dalfsen, augustus–oktober 2021

Contents

Voorwoord	vii
1 Introduction	1
1.1 Cooperative automated fleet management	2
1.2 Collaboration approaches	4
1.2.1 Centralized approaches	4
1.2.2 Decentralized approaches with central auctions	6
1.2.3 Decentralized approaches with local auctions	6
1.3 Multi-agent systems for collaborative routing	7
1.4 Problem statement and research questions	8
1.5 Thesis outline	9
2 Background	11
2.1 Dynamic collaborative pickup and delivery problems	12
2.1.1 Notation and objectives	12
2.1.2 Constraints and dynamics	15
2.1.3 Example problem and solution	18
2.2 Multi-agent auction approach	20
2.3 Conclusions	23
3 Information Sharing	25
3.1 Introduction	26
3.2 Related work	27
3.2.1 Full information sharing with a central coordinator	29
3.2.2 Partial information sharing with a central coordinator	30
3.2.3 Decentralized partial information sharing	31
3.2.4 Comparing information sharing policies	33
3.3 Information types	34
3.3.1 Vehicle route information	34
3.3.2 Marginal cost information	35
3.4 Auction approach	35

3.4.1	Order agent	37
3.4.2	Vehicle agent	39
3.5	Computational study	42
3.5.1	Problem instances	42
3.5.2	Experimental settings	43
3.5.3	Information sharing scenarios	43
3.5.4	Number of reauctions	47
3.5.5	Emission or congestion penalties	50
3.5.6	Fleet capacity and order urgency	52
3.5.7	Mixed information sharing attitudes	55
3.6	Implications	58
3.7	Conclusions	61
4	Large-Scale Collaboration	65
4.1	Introduction	66
4.2	Related work	67
4.3	Auction approaches	69
4.3.1	Local auction procedure	69
4.3.2	Bundling	73
4.3.3	Marginal costs and route improvements	74
4.3.4	Reference approach: central combinatorial auction	75
4.4	Computational study	77
4.4.1	Cooperation gains	77
4.4.2	Bundling benefits	82
4.4.3	Comparing central and local combinatorial auctions	83
4.5	Implications	89
4.6	Conclusions	90
5	Strategic Behaviour	93
5.1	Introduction	94
5.2	Related work	95
5.3	Auction approaches	96
5.3.1	First-price auctions	96
5.3.2	Second-price auctions	98
5.4	Computational study	100
5.4.1	Strategic behaviour in a first-price auction system	101
5.4.2	Strategic behaviour in a second-price auction system	104
5.5	Implications	109
5.6	Conclusions	110

6	User Preferences	113
6.1	Introduction	114
6.2	Related work	114
6.3	Generalized pickup and delivery problems with preferences	116
6.4	Auction approach	119
6.4.1	Order agent	121
6.4.2	Vehicle agent	122
6.4.3	Example	123
6.5	Computational study	123
6.5.1	Improvements through alternative locations	125
6.5.2	Comparing the MAS with ALNS	127
6.6	Implications	129
6.7	Conclusions	130
7	Conclusions and Future Research	131
7.1	Conclusions	132
7.2	Contributions	135
7.3	Future research	137
A	Marginal Costs	141
B	Comparing the Multi-Agent System with Centralized Heuristics	145
C	Results for Constrained Instance Sets	151
D	Comparing the Multi-Agent System with Combinatorial Auctions	161
	Bibliography	165
	Nomenclature	173
	Summary	177
	Samenvatting	181
	Curriculum Vitae	185
	TRAIL Thesis Series Publications	189

Chapter 1

Introduction

Freight transportation is necessary in supply chains but a large contributor to air pollution. Daily, numerous unloaded trips are performed, leaving room for improvement. Although individual carriers might not have a large potential to increase the efficiency of their operations, economies of scale could be obtained if different carriers collaborate. Besides the difficult problem of finding efficient routes, carriers' attitudes towards cooperation are crucial: if they do not trust the cooperation system or need to sacrifice too much of their autonomy, they likely decline to join a coalition. Consequently, big challenges for collaborative transportation arise in the areas of information sharing, incentives to participate, potential fraud, and individual preferences of the stakeholders.

Whereas this thesis will focus on these key problems of carrier cooperation in the next chapters, this chapter first sets the scene, presents the challenges in collaborative vehicle routing, and introduces the approaches that we will adopt. First, in Section 1.1, we sketch the major problems related to freight transportation; then we present carrier cooperation as a possible solution but identify some challenges of large-scale cooperative fleet management in a dynamic world. In Section 1.2, we review current approaches to realize collaborations between carriers and argue that a decentralized approach with local auctions might be most suitable for dynamic large-scale problems. Then, in Section 1.3, we set out the properties that such an approach ideally has and state the related research challenges. Subsequently, the research questions that this thesis addresses are formulated in Section 1.4, and an outline of the thesis is given in Section 1.5.

1.1 Cooperative automated fleet management

Every city dweller regularly experiences the results of increased direct-to-customer consumerism (Savelsbergh and Van Woensel, 2016): delivery vans from several companies are continuously blocking the streets, since they all deliver parcels in the same area around the same time. The resulting road congestion and air pollution are problems encountered by many cities. The road transport sector is responsible for about 30% of nitrogen oxides and 10% of particulate matter (European Environment Agency, 2019). Furthermore, approximately 12% of greenhouse gas emissions can be attributed to road transportation, of which 40% is caused by freight transportation (Ritchie and Roser, 2020).

Carrier cooperation is an important way to reduce these negative effects of transportation. Trucks are often only partly loaded, and the number of empty backhauls is still estimated to be roughly 10–30% (Terrazas, 2019). This capacity could often be used for other loads. If similar or complementary tasks from different carriers are combined into the same routes, the total vehicle mileage can be reduced (see Figure 1.1). Various studies have shown that this results in overall savings of 20–30% (Gansterer and Hartl, 2018b), not only in terms of costs, but also in terms of emissions and congestion.

Achieving a successful collaboration between carriers, however, is a complicated problem since different actors exist that have their own objectives. Traditional fleet management systems (see, e.g., Zeimpekis et al., 2007) assist in the assignment of transportation orders to a fleet of vehicles belonging to one single carrier. The goal is then to find optimal routes for the vehicles transporting the cargo according to certain criteria. The objective for this Vehicle Routing Problem (VRP) is often formulated as minimization of the total travel costs. Within collaborative routing, however, different carriers each have their own fleets of vehicles. Furthermore, each carrier may have a set of preassigned transportation orders, but customers or shippers can also ask the collective to transport a load instead of committing to a fixed carrier beforehand. This leads to a more complicated set of objectives when assigning or transferring orders to carriers and planning the vehicle routes: from a high level perspective, the total travel costs must be minimized, but also the service level must be maximized, that is, from the requests without an initial contract with one of the carriers, as few as possible should be rejected. Moreover, from their own perspective, each of the separate carriers will try to maximize its own profits from the collaboration. These goals will often conflict with one another.

Various trends in the field of transportation put even more challenges on carrier collaboration. First, there is an increasing demand for fast dispatching. Customers expect orders to be delivered the next day, the same day, or even, in certain application areas, within several minutes. The high degree of dynamics

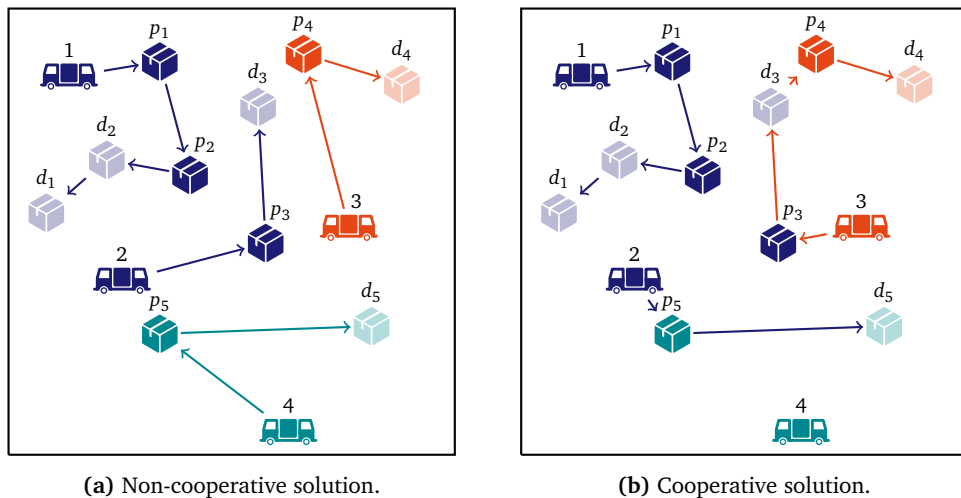


Figure 1.1: Non-cooperative and cooperative solution for a pickup and delivery problem with 3 carriers. In the non-cooperative case, each carrier serves its own orders. In the cooperative case, travel costs can be decreased by taking over orders from other carriers.

due to continuous changes in demand (and real-time changes and disturbances in supply as well) complicates the transportation planning process, and this difficulty might increase if coordination between different carriers is needed. The availability of large amounts of data (both real-time vehicle data and historical demand data) could facilitate real-time decisions, but its application potential is not straightforward. Furthermore, different data sources are owned by different stakeholders, and it is not likely that they are willing to share all information.

Second, with the trends of vehicle automation and the sharing economy, it is possible that privately owned autonomous vehicles will be made available to serve other tasks at idle times (Beirigo, 2021). Although not so relevant for full truckloads, this could be an interesting opportunity for parcel delivery. Hence, flexible coordination between a lot of (possibly small-scale, possibly temporarily available) vehicle owners becomes necessary, instead of traditional cooperation between only a few carriers. Those private individuals likely want to put out the operational planning to an automated platform that matches supply and demand, while larger carriers might want to stay at least partly autonomous, but could benefit a lot from participating in larger collaborations. Thus, we envision cooperation of thousands of heterogeneous transportation service providers, henceforth just called carriers. The recent rise of transportation platform companies (e.g., Quicargo and UberFreight) enables cooperation at a large scale, but the scientific knowledge on suitable large-scale approaches is limited: typical studies consider three collaborating carriers, and a few studies investigate

collaboration of some dozens of carriers. For future applications, however, it is necessary to explore methods to coordinate the operations of a number of carriers that is some orders of magnitude larger.

This thesis leaves the picture of human planners of different firms exchanging orders by numerous phone calls behind and turns to the vision of a transportation platform that automatically and efficiently matches millions of transportation requests that continuously arrive with the dynamic fleet of (autonomous) vehicles that is collectively brought together by thousands of carriers.

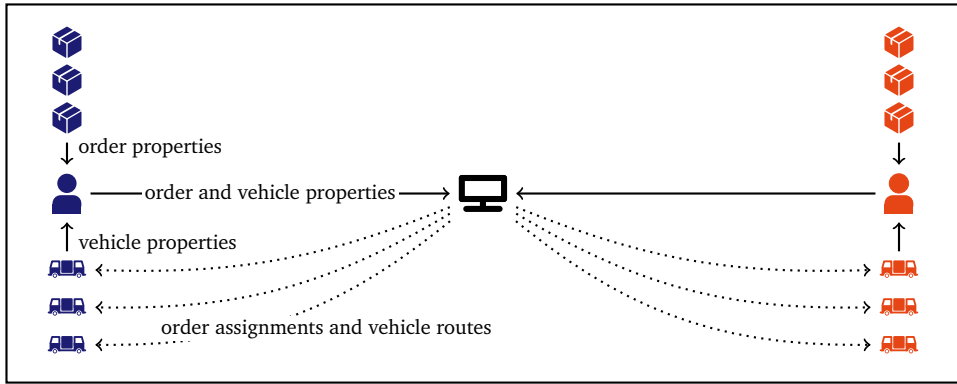
1.2 Collaboration approaches

Cooperation of carriers is studied within the field of collaborative vehicle routing, which in its broad definition encompasses all kinds of cooperations that are aimed at increasing the efficiency of vehicle fleet operations (Gansterer and Hartl, 2018b; Pan et al., 2019). Within the literature on collaborative vehicle routing, two main research areas can be distinguished: centralized collaboration and decentralized collaboration, where we make an additional distinction between decentralized approaches with and without a central auctioneer.¹

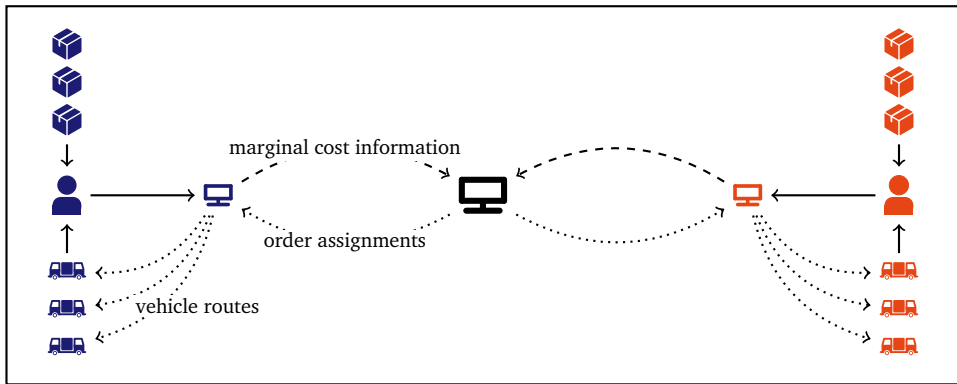
1.2.1 Centralized approaches

Centralized collaboration approaches typically assume that the vehicle and order information of all carriers is known by a single central decision maker having complete control over the resources of the different participants (see Figure 1.2a). Centralized approaches usually assume a set of orders for each carrier and compute what gains could theoretically be obtained if orders are exchanged. Approximation algorithms are used to compare the solution where each carrier performs only its own orders and the solution where (part of) the orders can be exchanged (Fernández et al., 2018; Molenbruch et al., 2017; Schulte et al., 2017; Montoya-Torres et al., 2016). Although the theoretical savings by cooperation could be computed this way, the practical applicability is limited. First, since the VRP is NP-hard, a central optimal solution cannot be computed if the

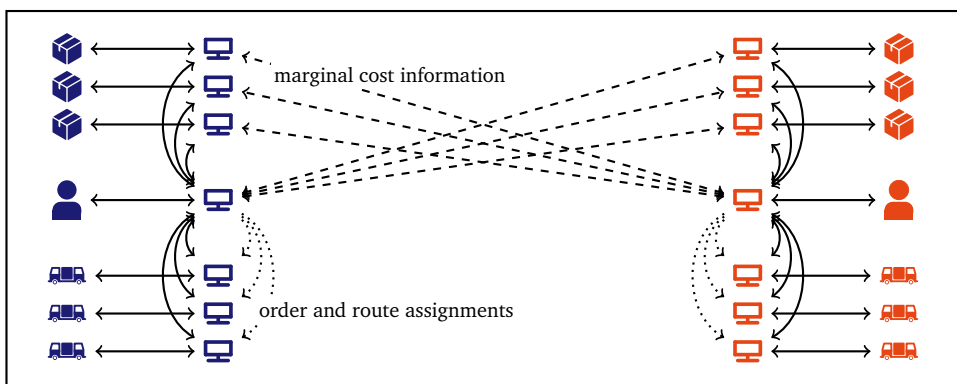
¹The terms *decentralized* and *distributed* are often confused (see, e.g., Vergne, 2020, Footnote 1). Although *distributed* (and its opposite *concentrated*) generally refer to decision making or control, and *decentralized* (together with *centralized*) to communications or the access to data, these notions often overlap in the field of collaborative vehicle routing: Gansterer and Hartl (2018b) use *centralized collaborative planning* for scenarios where “collaborative decisions are made by a central authority having full information” and *decentralized planning* if there is no central planner with full information. In line with the literature on collaborative vehicle routing, we use the terms *centralized* and *decentralized* in conjunction with, for instance, *planning*, *approach*, and *collaboration* throughout this thesis to refer to these notions.



(a) Centralized approach.



(b) Decentralized approach with a central auction.



(c) Decentralized approach with local auctions.

—→ full information - - - → limited information ·····→ instructions

Figure 1.2: Collaboration approaches with different control structures.

problem size increases. The number of cooperating carriers in most available computational studies ranges from 2–4. Only Schulte et al. (2017) use larger instances of up to 50 carriers, while in future practice thousands of carriers must cooperate. This scale problem is even more problematic if the available time is limited in a dynamic world. Second, carriers tend to be reluctant to share all their confidential information due to competition or legislation. Third, carriers likely want to stay autonomous, rather than having an external party prescribing them what to do.

1.2.2 Decentralized approaches with central auctions

The last two problems can be avoided in decentralized approaches, since there is no central omniscient controller (see Figure 1.2b). Typical decentralized approaches consider a single central auctioneer that interacts with all carriers but does not have complete information (Berger and Bierwirth, 2010; Gansterer and Hartl, 2018a). Each carrier can submit orders to the auctioneer, who offers them to all other carriers. Based on the received bids, the auctioneer computes an optimal assignment. Hence, carriers themselves are in control when it comes to deciding whether they submit orders for exchange, whether they share some information by placing a bid, and how they plan their individual routes. This approach, however, still suffers from scalability issues: due to the central role of the auctioneer, available computational studies are limited to static problems with small numbers of carriers and orders.

1.2.3 Decentralized approaches with local auctions

To be able to solve VRPs of a larger size, another decentralized approach, based on the concept of multi-agent systems (MASs), is used. A MAS consists of several intelligent autonomous computational entities that can act and interact in a certain environment to obtain their goals (Wooldridge, 2013), usually without a central main actor. If vehicles and orders are each represented by individual agents, an auction-based MAS can be used for allocation of orders to vehicles (see Figure 1.2c). There is no central auctioneer involved, but each individual can act as auctioneer and offer an order, while all carriers can bid for it (Máhr et al., 2010; Mes et al., 2013). Such market-based approaches are of increasing interest: quick adjustments based on real-time data (for example new carriers, changing orders, or schedule disturbances) are possible within a large system, since each modification of a plan only affects a few actors and does not disturb the planning of the others. The drawback of an auction-based MAS is that no guarantees on optimality can be given, but its application value lies in scalability and flexibility.

Auction-based MAS approaches in the field of vehicle routing are mostly applied for non-cooperative problem variants. There they serve as a heuristic for large-scale dynamic problems when traditional centralized heuristics cannot be applied: all vehicles of a single carrier are represented by autonomous agents that are fully cooperative in finding the best allocation of orders to vehicles. A natural and potentially beneficial extension is to apply this approach in an environment where different stakeholders exist, such that individual parties correspond with separate agents. Hence, this thesis explores the applicability of an auction-based MAS in a dynamic environment where multiple competitive carriers exist that want to benefit from cooperation.

1.3 Multi-agent systems for collaborative routing

Whether an auction-based MAS is suitable for real-world carrier cooperation highly depends on two aspects. First, an efficient routing solution needs to be found, that is, the total traveled distance must be as short as possible. Second, all participants need to have incentives to participate in the system in the desired way.

Whether a MAS can provide efficient solutions has been investigated in multiple studies. Van Lon and Holvoet (2017), for example, compared a MAS to centralized heuristics and found that a MAS approach performs better on large-scale problems with dynamic and urgent tasks. The current thesis develops a different MAS approach and compares it to other established methods to investigate the quality of the MAS. In line with earlier research, it is shown that the MAS is a competitive approach. The main focus of this thesis, however, is on the participation incentives for the carriers.

For a cooperative system to be applicable in practice, it is important that carriers want to participate. Obviously, their own profits should increase through cooperation. But other properties are important as well. A potential hurdle might be that carriers are requested to share certain information with the system (Cleophas et al., 2019). Although this problem is largest in a centralized approach where full vehicle and order information is requested, participants also might be hesitant to share partial information in a decentralized auction. Hence, it is important to investigate how much information is necessary in a decentralized collaboration system to find efficient solutions. This thesis contributes to answering this question: it investigates whether carriers that are hesitant in sharing their private information still have incentives to participate.

Another incentive to participate in a transportation collaboration pertains to fairness and trust. Carriers that cooperate should not be put at a disadvantage by other collaborators. Thus, carriers should not be able to increase their own

profits at the cost of the others. This directly relates to the concept of incentive compatibility, that is, the property that each participant obtains the best individual outcome when it acts according to its real preferences. In other words, strategic behaviour or cheating should not pay off. This thesis considers to what extent strategic behaviour of individual carriers can be prevented within a MAS approach.

1.4 Problem statement and research questions

As introduced in the previous sections, it is an open question how to organize large-scale carrier collaboration efficiently in a dynamic world, taking into account the individual interests and restrictions of the carriers. An auction-based MAS seems a possible approach, but it is not clear how the algorithm design and the system's properties are connected.

Thus, the overarching aim of this thesis is to answer the following question:

To what extent can an auction-based multi-agent system be applied to solve dynamic large-scale collaborative vehicle routing problems?

To answer this main question, five research questions are evaluated. First, it is important to know which *features* such a system must comprise to deal with collaborative problems:

1. How can an auction-based multi-agent system be applied to collaborative vehicle routing problems?

Within such a system, different amounts and types of information can be exchanged by the participants. Thus, it is evaluated what is *expedient* for the system in terms of available information:

2. What is the value of information sharing within this system?

Then, it is important to investigate whether the proposed approach indeed leads to efficient collaborations among a large number of carriers. Thus, the potential *result* of the system is assessed:

3. What gains can be obtained by large-scale carrier cooperation?

In a system that enables efficient collaborations, participants should not have incentives to diminish the overall efficiency for their own benefits. Thus, the *robustness*, in terms of potential misuse of the system, is examined:

4. To what extent can participants benefit from strategic behaviour in the system?

Finally, to be applicable in various real-world settings, a cooperation system should take the different wishes and attitudes of individual users into account. Thus, lastly, the practical *appropriateness* of such a system is considered:

5. How can the system assist in meeting specific user preferences?

1.5 Thesis outline

The remainder of this thesis is organized as follows (see Figure 1.3). First, Chapter 2 presents the background material for this thesis. The Dynamic Collaborative Pickup and Delivery Problem, that will serve as problem model in all chapters, is formally defined. Subsequently, a MAS approach that is suitable for collaborative problem variants is proposed. This answers Research Question 1 and provides the initial version of the method that is used in the subsequent chapters.

In Chapter 3, we vary the amounts and types of information that are exchanged by the participants within the proposed MAS. To answer Research Question 2, the effect on solution quality is measured, and the implications of different information sharing strategies are discussed.

Next, Chapter 4 extends the MAS with auctions of bundles of orders to improve solution quality. Then, a computational study based on real-world data is carried out to examine the possible gains of large-scale collaboration. This answers Research Question 3.

In Chapter 5, we discuss the problem of strategic behaviour and investigate which incentives carriers and shippers have to fairly participate in the system. To answer Research Question 4, we increase or decrease the bid prices for selected amounts of participants to learn whether they can increase their individual profits this way.

In Chapter 6, we look at the different wishes and preferences of individual users of the system. We propose a problem variant where users can specify different pickup and delivery alternatives and indicate which options they prefer. We show how the MAS can be used to solve these problems, answering Research Question 5.

Finally, Chapter 7 concludes this thesis by discussing to what extent the developed MAS can be applied in real-world cases and provides directions for future research.

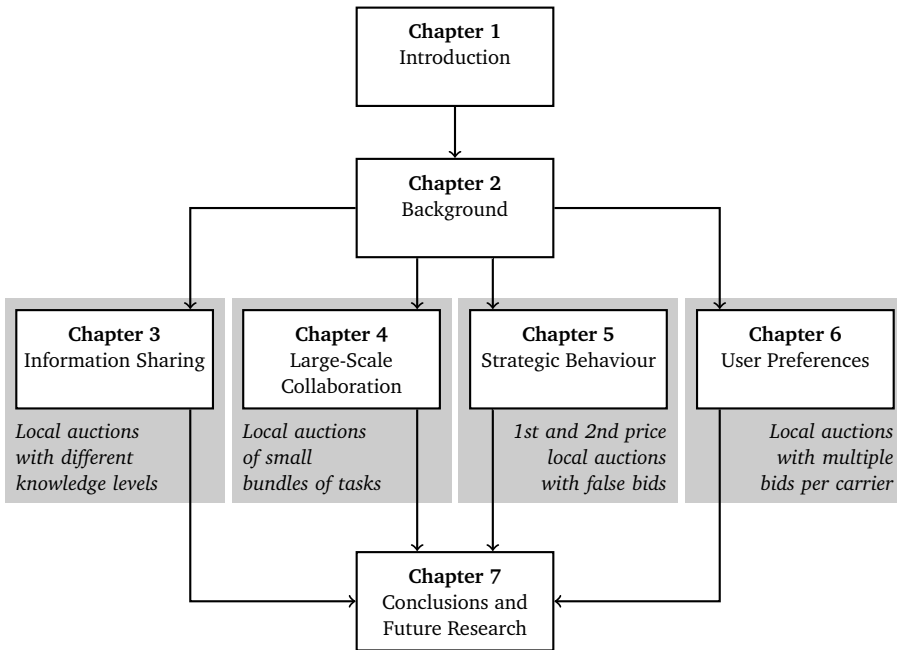


Figure 1.3: Structure of the thesis with solution method details.

Chapter 2

Background

In Chapter 1, we introduced the topic of this thesis, proposed to develop a MAS approach for large-scale dynamic collaborative problems, and stated the research questions. This chapter presents the background material that is necessary to understand the proposed MAS approach and serves as a basis for the subsequent chapters.

In this chapter, we introduce two subjects. First, in Section 2.1, we give a formal definition of the Dynamic Collaborative Pickup and Delivery Problem that will be used in the following chapters. Second, in Section 2.2, we describe the key elements of a multi-agent auction approach to solve large instances of this problem. Basically, this will answer the question how an auction-based multi-agent system can be applied to collaborative problem variants (Research Question 1). The elementary approach will be extended in several ways in the subsequent chapters. Finally, Section 2.3 concludes the chapter.

2.1 Dynamic collaborative pickup and delivery problems

We consider a transportation platform (see, e.g., Guo, 2020; Pourmohammad-Zia et al., 2020) to which thousands of carriers and millions of shippers may connect. Companies and individuals can make their (autonomous) vehicles available at any moment, to make some money by contributing to the transportation service. Professional shippers and private customers with specific requests can ask for transportation at any time and are partly autonomous in selecting a service. Hence, we assume a platform that connects large amounts of transport supply and demand in a complex, dynamic world.

The problem builds on the well-known Pickup and Delivery Problem (PDP; see Parragh et al., 2008), in which each order consists of a load of a certain size to be collected at a pickup location within a given time window and to be dropped at a delivery location within another time window. In addition, we assume the following:

- The problem is dynamic. Orders can be added, changed, or canceled at any moment during the operational time span. Furthermore, carriers can join or leave the platform at any moment or announce changes with respect to the availability of their vehicles. The system has to deal with these dynamics in real time.
- The number of orders might be very large. Hence, central reoptimization after arrival of an order will not be possible in reasonable time due to the complexity of the problem.
- The problem is inherently decentralized. Availability time windows, route plans, and carriers' operational cost structures might not be known to the platform. Classical centralized approaches for (dynamic) PDPs are hence not appropriate.

2.1.1 Notation and objectives

In the following, we formally define the Dynamic Collaborative Pickup and Delivery Problem (DCPDP). A problem instance (see Table 2.1) consists of a set of shippers S , a set of carriers C , a set of orders O , and a set of vehicles V . The orders can be divided into two categories: orders that are directly submitted to the platform by shippers themselves (O_S), and orders that already have been assigned to a specific carrier (due to a long-term contract between shipper and carrier) but are available for reassignment (O_C). In the later case, the contracted carrier can be seen as the owner – the original shipper is then irrelevant. For each shipper $s \in S$, its orders are denoted by O_s , and equivalently, O_c denotes

Table 2.1: Nomenclature for instances and solutions of the DCPDP

General instance properties	S	Set of shippers
	C	Set of carriers
	O	Set of orders
	V	Set of vehicles
	τ	Time horizon
Properties of order $o \in O$	r_o	Release time
	p_o	Pickup location
	d_o	Delivery location
	q_o	Quantity
	f_o	Price
Properties of vehicle $v \in V$	r^v	Release time
	e^v	Earliest availability time
	l^v	Latest availability time
	α^v	Start location
	ω^v	End location
Sets of locations	P	All pickup locations
	D	All delivery locations
	A	All vehicle start locations
	Ω	All vehicle end locations
	I	Possible interim locations
Properties of locations $i, j \in P \cup D \cup A \cup \Omega \cup I$	e_i	Earliest service time
	l_i	Latest service time
	s_i	Service duration
	c_{ij}	Travel cost
	t_{ij}	Travel time
Properties of route plan for vehicle v at time t	n^{vt}	Number of stops
	ρ_h^{vt}	Location of h -th stop
	$a(\rho_h^{vt})$	Arrival time at h -th stop
	$d(\rho_h^{vt})$	Departure time at h -th stop
	$s(\rho_h^{vt})$	Service start time at h -th stop
	$l(\rho_h^{vt})$	Load after service at h -th stop
Properties of final solution R^τ	O_A	Set of accepted orders
	O_R	Set of rejected orders

the orders initially assigned to carrier $c \in C$, such that $O_S = \bigcup_{s \in S} O_s$ is the total set of unassigned orders, $O_C = \bigcup_{c \in C} O_c$ is the total set of assigned orders, and $O = O_S \cup O_C$ is the total set of orders. Likewise, V_c denotes the set of vehicles of carrier $c \in C$, such that $V = \bigcup_{c \in C} V_c$ is the total set of vehicles.

Each order $o \in O$ has a release time r_o (i.e., the time the order becomes known to the system), a pickup location p_o , a delivery location d_o , a load quantity q_o , and a reservation price f_o (i.e., the maximum amount that the shipper is willing to pay for transportation). Each vehicle $v \in V$ has a release time r^v , a capacity k^v , an availability time window $[e^v, l^v]$, a start location α^v , and an end location ω^v . The release time specifies when the vehicle's availability becomes known to the corresponding carrier, whereas the time window $[e^v, l^v]$ specifies its actual availability. Hence, $r^v \leq e^v$. Carriers are active from their release time r_c until all their vehicles have become unavailable. Hence, for each time t , we can define the set of carriers that is known and active by $C^t = \{c \in C \mid r_c \leq t \wedge \exists v \in V_c \ t < l^v\}$. (Generally, the release time r^v for a vehicle $v \in V_c$ will equal the release time r_c of the owning carrier, but it is possible to model later vehicle release times, for example, when a carrier outsources its vehicles elsewhere without knowing the exact return details.)

Each location i has an earliest service start time e_i and a latest service start time l_i , that together determine the time window $[e_i, l_i]$ in which the service can start. Furthermore, s_i is the service duration, determining the time that is needed for the service at i . (For each vehicle $v \in V$, we set $e_{\alpha^v} = e_{\omega^v} = e^v$, $l_{\alpha^v} = l_{\omega^v} = l^v$, and $s_{\alpha^v} = s_{\omega^v} = 0$ to model that it only operates during its availability time window.) For each pair of locations (i, j) , we denote the travel time and travel costs from i to j by t_{ij} and c_{ij} , respectively. All times are assumed to be before the time horizon τ .

A (temporary) solution at time t for a problem instance is defined as a set of routes $R^t = \{\langle \rho^{1t} \rangle, \dots, \langle \rho^{|V|^t} \rangle\}$, where each route (plan) $\langle \rho_h^{vt} \rangle_{h=1}^{n^{vt}}$ is a sequence of n^{vt} locations representing the (partially completed) path of vehicle v at time t . Due to the dynamic nature of the problem, the (partially completed) route plans can change during operational time. Only orders $o \in O$ for which p_o and d_o occur in the definitive route $\langle \rho^{v\tau} \rangle$ are served by vehicle v . Let O_A^τ and O_R^τ denote the sets of served and rejected orders in the final solution R^τ :

$$O_A^\tau = \{o \in O \mid \exists v \in V \ \exists h \ \rho_h^{v\tau} = p_o\}; \text{ and} \quad (2.1)$$

$$O_R^\tau = O \setminus O_A^\tau. \quad (2.2)$$

The overall goal is to obtain a final solution R^τ with several objectives:

- maximize the service level

$$SL(R^\tau) = \frac{|O_A^\tau|}{|O|}; \quad (2.3)$$

- minimize the total travel costs

$$\text{TC}(R^\tau) = \sum_{v \in V} \sum_{h=2}^{n^{v\tau}} c_{ij}, \text{ where } i = \rho_{h-1}^{v\tau} \text{ and } j = \rho_h^{v\tau}; \text{ and} \quad (2.4)$$

- maximize the total profit

$$\text{PR}(R^\tau) = \sum_{o \in O_A^\tau} f_o - \text{TC}(R^\tau) - \gamma |O_R^\tau|, \quad (2.5)$$

where γ is a parameter representing the fine per rejected order.

The total profit might be used as an aggregate objective, in which the relative importance of service level and travel costs can be controlled by the values of γ and c_{ij} . In the definition of profit above, it is assumed that shippers always pay their reservation prices for transportation of their orders. In Chapters 4 and 5, however, we consider the more natural situation in which shippers can pay lower prices, dependent on the current market situation.

Each stakeholder might have a different focus: shippers or customers might mainly be concerned about the service level, while a platform owner or authority might also be concerned about environmental aspects, as represented by travel costs. Individual carriers, however, might have different goals. Without an effective profit sharing mechanism, they likely want to maximize their own profit at the expense of the global system. Define the equivalents of the global objectives on carrier level as follows:

$$\text{TC}_c(R^\tau) = \sum_{v \in V_c} \sum_{h=2}^{n^{v\tau}} c_{ij}, \text{ where } i = \rho_{h-1}^{v\tau} \text{ and } j = \rho_h^{v\tau}; \text{ and} \quad (2.6)$$

$$\text{PR}_c(R^\tau) = \sum_{v \in V_c} \sum_{o \in O_A^{v\tau}} f_o - \text{TC}_c(R^\tau), \text{ where } O_A^{v\tau} = \{o \in O \mid \exists h \rho_h^{v\tau} = p_o\}. \quad (2.7)$$

Carriers might completely focus on their local profit by maximizing PR_c , or might be concerned with the global system, for example, if all carriers are charged equally for the term $-\gamma |O_R^\tau|$ in the global profit function. In even more cooperative situations, carriers do not want to optimize their local profit, but focus completely on the global profit PR , and rely on a fair profit allocation system.

2.1.2 Constraints and dynamics

In a dynamic setting, we assume that vehicles might change direction at any time t , except during service time (i.e., loading or unloading needs to be fully completed). Hence, we define the DCPDP on the Euclidean plane, with travel

times and travel costs proportional to Euclidean distances between two locations. By doing so, the location of a vehicle can be determined at every time t if we keep track of departure and arrival times. Also, travel times and travel costs can be defined for each pair of locations.

Let $P = \bigcup_{o \in O} \{p_o\}$ and $D = \bigcup_{o \in O} \{d_o\}$ be the sets of all pickup and all delivery locations, respectively. Also, let $A = \bigcup_{v \in V} \{\alpha^v\}$ and $\Omega = \bigcup_{v \in V} \{\omega^v\}$ be the sets of all start and end positions of the vehicles. Furthermore, by I we denote a set of additional interim locations at which the vehicles might change direction. We require P , D , A , Ω , and I to be disjoint sets, but they can contain duplicates of the same physical locations. For $i \in I$, we always have $s_i = 0$.

A route (plan) $\langle \rho_h^{vt} \rangle_{h=1}^{n^{vt}}$ for a vehicle $v \in V$ at time $t \geq r^v$ consists of at least two locations ($n^{vt} \geq 2$), is associated with four functions $a : \bigcup_{h=2}^{n^{vt}} \{\rho_h^{vt}\} \rightarrow \mathbb{R}$, $d : \bigcup_{h=1}^{n^{vt}-1} \{\rho_h^{vt}\} \rightarrow \mathbb{R}$ (representing arrival and departure times, respectively), $s : \bigcup_{h=1}^{n^{vt}} \{\rho_h^{vt}\} \rightarrow \mathbb{R}$ (representing service start time), and $l : \bigcup_{h=1}^{n^{vt}} \{\rho_h^{vt}\} \rightarrow \mathbb{R}$ (representing vehicle load after service), and has the following constraints:

- the vehicle starts at its start location and stops at its end location:

$$\rho_1^{vt} = \alpha^v \text{ and } \rho_{n^{vt}}^{vt} = \omega^v; \quad (2.8)$$

- all time constraints are respected:

$$e_{\rho_h^{vt}} \leq s(\rho_h^{vt}) \leq l_{\rho_h^{vt}} \quad \forall h \in \{1, \dots, n^{vt}\}; \quad (2.9)$$

$$a(\rho_h^{vt}) \leq s(\rho_h^{vt}) \quad \forall h \in \{2, \dots, n^{vt}\}; \quad (2.10)$$

$$s(\rho_h^{vt}) + s_{\rho_h^{vt}} \leq d(\rho_h^{vt}) \quad \forall h \in \{1, \dots, n^{vt} - 1\}; \quad (2.11)$$

$$d(\rho_{h-1}^{vt}) + t_{\rho_{h-1}^{vt}, \rho_h^{vt}} = a(\rho_h^{vt}) \quad \forall h \in \{2, \dots, n^{vt}\}; \quad (2.12)$$

- all capacity constraints are respected:

$$l(\rho_1^{vt}) = l(\rho_{n^{vt}}^{vt}) = 0; \quad (2.13)$$

$$\rho_h^{vt} = p_o \Rightarrow l(\rho_h^{vt}) = l(\rho_{h-1}^{vt}) + q_o \quad \forall h \in \{2, \dots, n^{vt} - 1\}, o \in O; \quad (2.14)$$

$$\rho_h^{vt} = d_o \Rightarrow l(\rho_h^{vt}) = l(\rho_{h-1}^{vt}) - q_o \quad \forall h \in \{2, \dots, n^{vt} - 1\}, o \in O; \quad (2.15)$$

$$l(\rho_h^{vt}) \leq k^v \quad \forall h \in \{2, \dots, n^{vt} - 1\}; \quad (2.16)$$

- pickups and deliveries are paired and respect precedence constraints:

$$\rho_h^{vt} = p_o \Rightarrow \exists \hat{h} > h \rho_{\hat{h}}^{vt} = d_o \quad \forall h \in \{2, \dots, n^{vt} - 1\}, o \in O; \quad (2.17)$$

$$\rho_h^{vt} = d_o \Rightarrow \exists \hat{h} < h \rho_{\hat{h}}^{vt} = p_o \quad \forall h \in \{2, \dots, n^{vt} - 1\}, o \in O; \text{ and } \quad (2.18)$$

- start and stop locations occur only at the beginning and end of a route:

$$\rho_h^{vt} \in P \cup D \cup I \quad \forall h \in \{2, \dots, n^{vt} - 1\}. \quad (2.19)$$

For each vehicle $v \in V$ and each time $t < r^v$, we have an empty route $\langle \rho^{vt} \rangle = \langle \rangle$.

A (temporary) solution $R^t = \{\langle \rho^{1t} \rangle, \dots, \langle \rho^{|V|t} \rangle\}$ for a problem instance has the following restrictions:

- the pickup and delivery corresponding to each order $o \in O$ occur within the route of at most one vehicle:

$$\begin{aligned} \rho_h^{vt} = p_o \Rightarrow (\rho_{\hat{h}}^{\hat{v}t} \neq p_o \vee (\hat{v} = v \wedge \hat{h} = h)) \\ \forall o \in O, v, \hat{v} \in V, h \in \{2, \dots, n^{vt} - 1\}, \hat{h} \in \{2, \dots, n^{\hat{v}t} - 1\}; \text{ and} \end{aligned} \quad (2.20)$$

- no elements unknown at time t appear:

$$\begin{aligned} (\rho_h^{vt} = p_o \vee \rho_h^{vt} = d_o) \Rightarrow r_o \leq t \\ \forall o \in O, v \in V, h \in \{2, \dots, n^{vt} - 1\}. \end{aligned} \quad (2.21)$$

Due to the dynamic nature of the problem, the solution needs to be iteratively updated during run time. For each vehicle $v \in V$, an initial route $\langle \rho^{vr^v} \rangle = \langle \alpha^v, \omega^v \rangle$ is given at time r^v . Given a vehicle v , a previous time u and the current time t (with $r^v \leq u < t$), the route $\langle \rho^{vu} \rangle$ might be changed into a route $\langle \rho^{vt} \rangle$ only if the part until time t remains unaffected. Formally we have the following requirements for each ρ_h^{vu} within $\langle \rho^{vu} \rangle$ (except for $h = 1$, since $\rho_1^{vt} = \rho_1^{vu}$):

- if the service at ρ_h^{vu} has already started at time t (i.e., $s(\rho_h^{vu}) \leq t$), then the location is fixed and arrival and service time cannot be changed anymore:

$$\rho_h^{vt} = \rho_h^{vu}, a(\rho_h^{vt}) = a(\rho_h^{vu}), \text{ and } s(\rho_h^{vt}) = s(\rho_h^{vu}); \quad (2.22)$$

the departure time, however, may be changed according to $d(\rho_h^{vt}) \geq t$ if $d(\rho_h^{vu}) > t$;

- if the vehicle has arrived at ρ_h^{vu} at time t but has not yet started the service (i.e., $a(\rho_h^{vu}) \leq t < s(\rho_h^{vu})$), there are two options:

- the vehicle can still process the planned pickup or delivery:

$$\rho_h^{vt} = \rho_h^{vu}; \text{ or} \quad (2.23)$$

- the vehicle may withdraw from service and leave the location by changing it into an interim location:

$$\rho_h^{vt} = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_h^{vu} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = t; \quad (2.24)$$

in either case, $a(\rho_h^{vt}) = a(\rho_h^{vu})$, but the service start time and departure time may be changed according to $d(\rho_h^{vt}) \geq s(\rho_h^{vt}) \geq t$;

- if the vehicle is driving towards ρ_h^{vu} at time t (i.e., $d(\rho_{h-1}^{vu}) \leq t < a(\rho_h^{vu})$), there are again two options:
 - the vehicle may stick to the original plan:

$$\rho_h^{vt} = \rho_h^{vu}; \quad (2.25)$$

in this case, the arrival time stays the same, but service and departure times may be changed, if possible:

$$a(\rho_h^{vt}) = a(\rho_h^{vu}) \text{ and } d(\rho_h^{vt}) \geq s(\rho_h^{vt}) \geq a(\rho_h^{vu}); \text{ or} \quad (2.26)$$

- the vehicle may change direction immediately at time t :

$$\rho_h^{vt} = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ the actual location of the vehicle at time } t$$

$$\text{with } e_{\hat{\rho}} = l_{\hat{\rho}} = t, \quad (2.27)$$

and arrival, service, and departure times are set accordingly:

$$a(\rho_h^{vt}) = s(\rho_h^{vt}) = e_{\hat{\rho}} \text{ and } d(\rho_h^{vt}) \geq e_{\hat{\rho}}; \quad (2.28)$$

- if none of the tree above conditions hold, there are no restrictions for ρ_h^{vu} , that is, ρ_h^{vu} may be changed or removed from the route (subject to the definitions of a route and a solution, of course).

2.1.3 Example problem and solution

In Figure 2.1, we show how a solution for a DCPDP could develop over time. Initially, two vehicles and two orders are known to the system. Although vehicle 1 is only available from $t = 4$ on, it can already plan its route at $t = 0$. Since the capacity of vehicle 1 is not sufficient for order 4, the best plan is to assign order 4 to vehicle 2 and order 5 to vehicle 1. Vehicle 2 immediately starts driving towards p_4 . At $t = 4$, vehicle 1 becomes available and starts driving towards p_5 .

Then, at $t = 7$, vehicle 3 connects to the platform and announces its availability from $t = 8.5$ on. The total driven distance can be decreased if vehicle 3 serves order 4 and vehicle 2 serves order 5, since vehicle 2 is already quite close to p_5 at $t = 7$. Hence, the plan is updated at $t = 7$ as follows: vehicle 1 stops driving towards p_5 and stays idle, vehicle 2 starts driving towards p_5 , and vehicle 3 plans to go to p_4 immediately when it becomes available at $t = 8.5$. At $t = 9.5$, vehicle 2 reaches p_5 , but cannot yet start the pickup, since the pickup

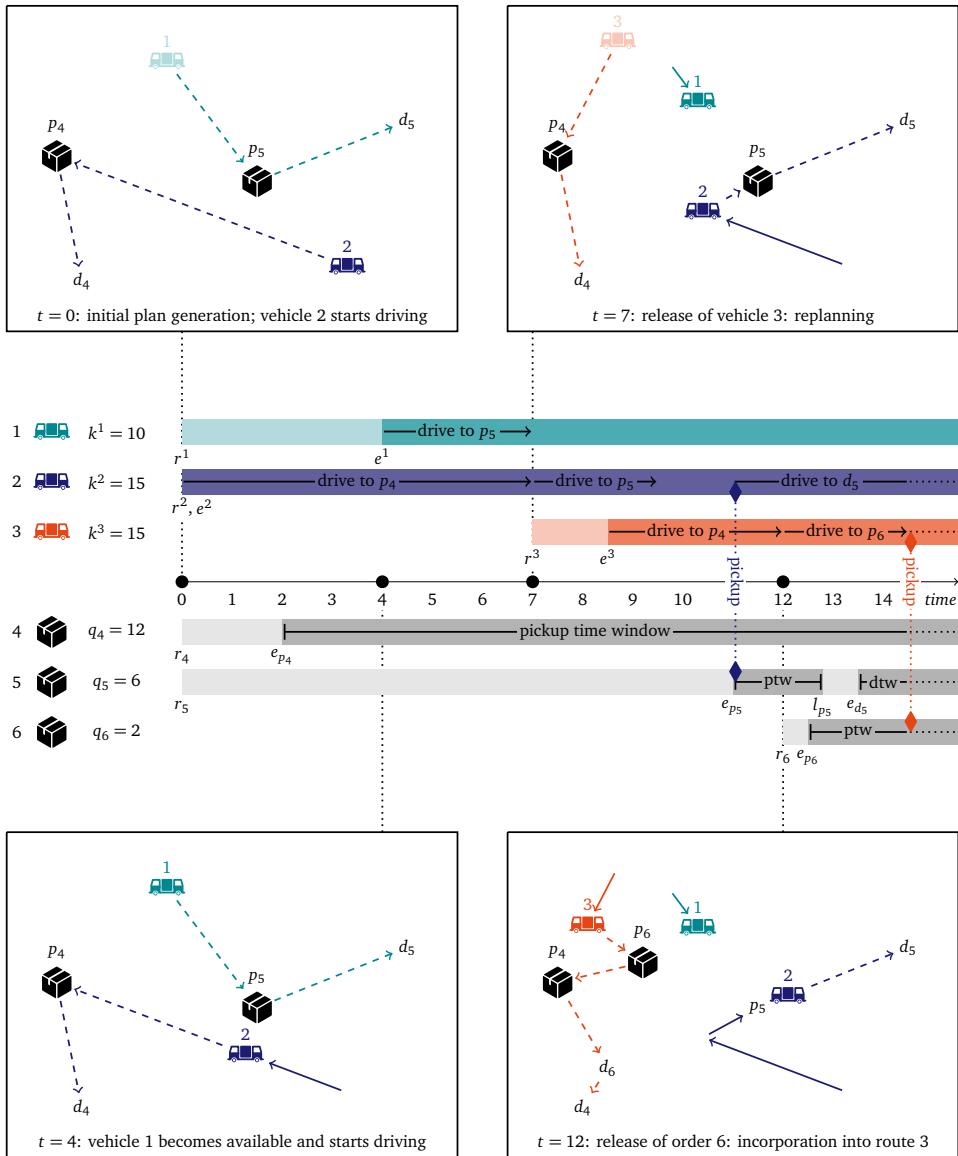


Figure 2.1: First part of the development of a solution for a DCPDP instance with 3 orders and 3 vehicles. The time line shows the vehicle and order properties, as well as the vehicle actions. The four panels show the positions, planned routes (dashed lines), and completed routes (solid lines) of the vehicles at $t = 0$, $t = 4$, $t = 7$, and $t = 12$.

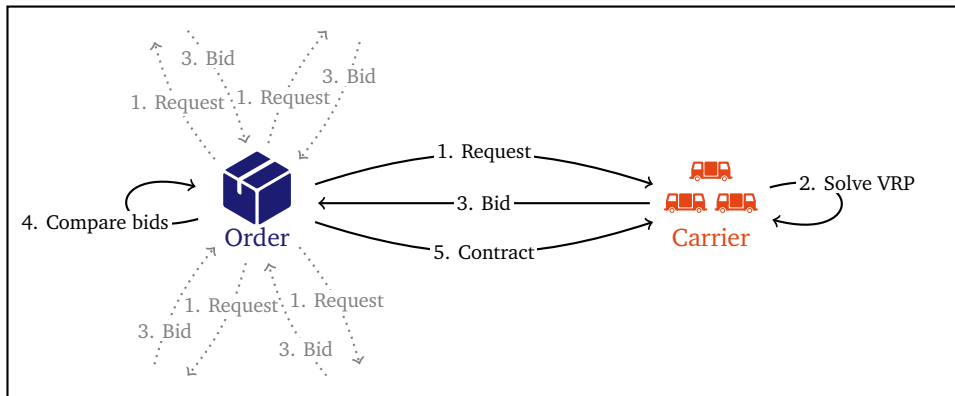


Figure 2.2: Standard MAS approach for large-scale dynamic collaborative VRPs.

time window of order 5 has not yet opened. Hence, vehicle 2 waits until $t = 11$, performs the pickup, and next drives towards d_5 .

At $t = 12$, a new order is submitted to the platform. Although its pickup location is close to vehicle 1, a better option is to assign it to vehicle 3 since vehicle 3 has enough capacity for carrying order 4 and 6 together and d_6 is close to d_4 . Hence, at $t = 12$, vehicle 3 includes order 6 into its route, and changes direction towards p_6 .

Note that the final solution is not optimal: with hindsight it would have been better if either vehicle 1 or vehicle 2 had stayed completely idle and the other had served order 5.

2.2 Multi-agent auction approach

We argued in Chapter 1 that a MAS approach would be suitable for solving large instances of the problem defined in the previous section, whereas approaches with a single central coordinator or auctioneer suffer from privacy, autonomy, or scalability issues. Thus, throughout this thesis, we develop a multi-agent auction approach where orders are iteratively offered in reverse auctions (see Figure 2.2). The carriers, acting as sellers of service, can bid for the orders, and the carrier with the lowest bid wins the auction: it becomes responsible for filling the order. In the current section, we present the basic MAS approach, based on Máhr et al. (2010) and Mes et al. (2013). Whereas these previous approaches consider individual vehicles as autonomous entities (either all vehicles belong to the same carrier or each vehicle corresponds to a separate carrier), we generalize the approach such that carriers are considered as autonomous actors. This gives the opportunity to model realistic collaborative scenarios where carriers

can have multiple vehicles, although computing the marginal costs is more complex with multiple vehicles than with one vehicle: an instance of the VRP must be solved instead of a Traveling Salesman Problem (TSP) instance.

When an order $o \in O$ becomes available at r_o , an auctioneer for order o (acting on behalf of shipper s if $o \in O_s$ or acting on behalf of carrier c if $o \in O_c$, but operated by the platform) is initialized and becomes active. When active, the auctioneer repeatedly organizes auctions. Given a maximum number of auctions m per auctioneer and its activation time r_o , the time between subsequent auctions is set to $(l_{p_o} - r_o)/m$. The auction at time t then consists of the following steps (see Figure 2.2):

1. **Requesting transportation:** The auctioneer sends a request for transporting order o (including the details of the order in terms of locations, time windows and the load quantity) to all known and active carriers $c \in C^t$.
2. **Computing marginal costs:** Each carrier $c \in C^t$ computes its marginal costs $MC_c^t(o)$ for order o at time t , that is, the minimal extra travel costs for inserting o into one of its routes, according to the constraints and given the situation at time t . A formal description is given in Appendix A. For a quick computation of the marginal costs, however, we often use a greedy insertion heuristic that respects the current sequences of stops in the routes (see Figure 2.3) instead of solving the insertion problem exactly. If transporting o is infeasible for c , $MC_c^t(o)$ is set to ∞ .
3. **Bidding:** The carriers submit a bid with value $MC_c^t(o)$ to the auctioneer (i.e., they indicate that they can transport the order at this price).
4. **Comparing:** The auctioneer compares the received bids; let b_0 be the lowest bid provided by carrier c_0 . Furthermore, the auctioneer examines the current costs $CC^t(o)$ for order o at time t , given by the actual marginal costs for the order if it is incorporated in a carrier's route plan, or by the order's reservation price otherwise.
5. **Updating contracts:** If $b_0 < CC^t(o)$, the bid is accepted. (In a more cooperative scenario, if the order is still unassigned, the bid also could be accepted independently of the order's reservation price.) The platform informs the winning carrier, who updates its route plan by incorporating the order into it, and the shipper of the order, who updates its contract. If a contract already existed for the order, the contracted carrier is informed about the cancellation. Which payments between the different actors will be made depends on the scenario.

When transportation of o starts or l_{p_o} has passed without a contract for the order, the auctioneer stops starting auctions and becomes inactive.

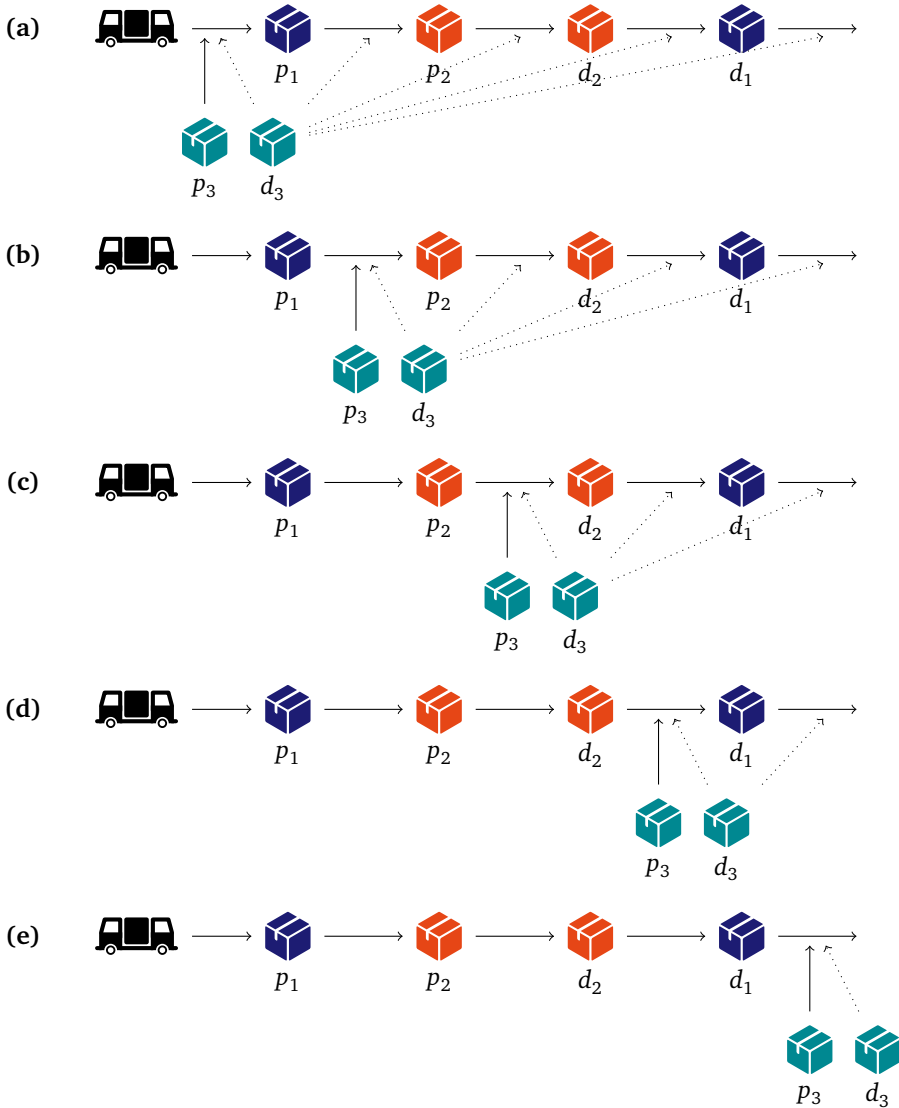


Figure 2.3: Different insertion possibilities for a new pickup (p_3) and delivery (d_3) into a vehicle's route consisting of two orders. The greedy heuristic keeps the sequence of the current route ($\langle p_1, p_2, d_2, d_1 \rangle$). The number of possible routes to check is quadratic in the number of pickups and deliveries.

The described MAS will be extended in the next chapters. In Chapter 3, the number of carriers that will be sent a request per auction will be limited to decrease the computational load, and the cost information given within a bid will sometimes be hidden for privacy reasons. In Chapter 4, auctioneers can offer bundles of orders in addition to single orders to benefit from interaction effects. Also, we define ways to distribute the gains of each auction among the participants. Then, in Chapter 5, we will compare first-price and second-price auctions; this requires ways to make sure that enough money is available to pay the value of the lowest or the second-lowest bid. Finally, in Chapter 6, we extend the approach to a variant where carriers can submit multiple alternative bids and auctioneers have other criteria for comparing and weighing them.

2.3 Conclusions

In this chapter, we have defined and explained the Dynamic Collaborative Pickup and Delivery Problem to model real-world transportation problems with multiple carriers involved. The constraints and dynamics have been formally described. Furthermore, we have proposed a multi-agent auction approach that is suitable for problems with multiple carriers. Instead of considering vehicle controllers as individual agents, we consider carriers as the main decision makers. This gives the opportunity to apply the MAS to collaborative vehicle routing problems, which answers Research Question 1. The basic version of the approach that has been presented in this chapter will be extended in the subsequent chapters.

Chapter 3

Information Sharing

In Chapter 2, we introduced the Dynamic Collaborative Pickup and Delivery Problem, and proposed a MAS approach using a platform for solving it. For such a transportation platform to function, carriers and shippers need incentives to participate and cooperate on it. A limitation of the types and amounts of operational information to be shared might be one of these incentives, since this information might be confidential. Although carriers are often requested to reveal (part of) their private information, not all information is always necessary. The exact relation between available carrier information and solution quality, however, is not known.

In this chapter, we investigate the value of sharing different classes and amounts of carrier information for solution quality in a decentralized routing approach to answer Research Question 2. First, we explain the problem of information sharing in Section 3.1. Then, in Section 3.2, we review the literature on collaborative vehicle routing with emphasis on the information sharing assumptions made. In Section 3.3, we further elaborate on two types of information that will be of importance in this chapter. Next, in Section 3.4, we develop a MAS in which the amount of exchanged information can be limited. Then, in Section 3.5, we perform a series of simulations with the MAS, showing trade-offs between carriers' tendency to protect information and system performance. Finally, in Section 3.6, we discuss the design of incentives to share the relevant information, and Section 3.7 concludes the chapter.

3.1 Introduction

The increasing demand for transportation requires efficient and robust vehicle routing. Providing a reliable and quick transportation service while minimizing costs, energy consumption, and pollution is important for carriers, shippers, and society in general. A significant gain can be obtained when different carriers cooperate: combining deliveries can reduce travel costs, pollution, and congestion. Various studies have shown improvements of about 20–30% in comparison with non-collaborative situations (Gansterer and Hartl, 2018b). For individual carriers, even profit improvements of up to 800% may be achieved if they cooperate in larger coalitions (Schulte et al., 2019).

With the rise of platform companies (e.g., Quicargo and UberFreight), cooperative planning is becoming even more important. Such platforms act as intermediary between carriers and shippers, to connect transport supply and demand dynamically without having direct control over these actors. Hence, incentives for cooperation need to be provided to both shippers and carriers.

A limitation of the types and amounts of operational information that needs to be shared with the platform or other participants might be one of these incentives. Carriers can be requested to reveal part of their private information, such as vehicle availability, intended routes, or costs for satisfying a request to enable cooperation on a platform. However, they might be reluctant to share all confidential information due to competition or legislation – a problem referred to as *coopetition* (Crujssen et al., 2007; Cleophas et al., 2019). On the other hand, not all carrier information is always needed. In highly dynamic scenarios, there is generally no time for computing optimal solutions. Reasonable approximations based on partial information may be of great value in these situations.

The precise effect of available carrier information on solution quality, however, is not known. Current research on collaborative transportation mainly applies centralized models or combinatorial auctions (Gansterer and Hartl, 2018b), in which a central coordinator collects information and defines a solution. These approaches (which are also applied in platforms) require certain fixed levels of information sharing with a central authority. Multi-agent approaches, on the other hand, allow to model different information sharing levels in a flexible manner, even without an active central authority (shippers and carriers might need the platform mainly for getting to know each other). Nonetheless, current multi-agent approaches have considered either fully cooperative carriers, that is, they accept orders that are unprofitable for themselves if the global system is better off (Máhr et al., 2010; Gath, 2016), or fully competitive carriers, that is, they just try to maximize their own profit (Figliozzi et al., 2004, 2005). In a realistic system, however, carriers might want to find a balance between cooperation and competition by sharing a limited amount of information.

In this chapter, we examine the value of sharing different types and amounts of carrier information for the quality of the routing solution in a dynamic collaborative transportation setting. To this end, we extend the MAS of Chapter 2 with different information sharing options. In a series of simulations, we focus on two types of information: we vary what is known about a vehicle's position and route plan, and what cost information a carrier includes in its bid. The results show trade-offs between carriers' tendency to protect information and system performance based on service level, travel costs, and profits. Although the approach abstracts from direct incentives and explicit gain sharing, a system-wide improvement in solution quality is observed with increased information sharing. This leaves room for the design of incentives to share the relevant information.

3.2 Related work

Horizontal collaboration between carriers has extensively been considered in the literature (see, e.g., the reviews by Cruijssen et al., 2007; Verdonck et al., 2013; Gansterer and Hartl, 2018b; Cleophas et al., 2019; Pan et al., 2019). Although several collaboration aspects (e.g., types of collaboration, request exchange mechanisms, gain sharing) have been investigated in depth, the issue of information sharing has hardly gotten any attention. Most authors mention that information sharing might be difficult in practice due to competition and privacy regulations, but disregard the implications for their approach because they focus on other aspects. Some authors emphasize that their approach does not require full information sharing, but do not state explicitly what information is being exchanged in their approach and what can be inferred from this information.

We give an overview of the information sharing properties of several recent articles on collaborative vehicle routing and compare them in Table 3.1. The table is divided into three parts, each having a different information sharing paradigm.

The first part contains articles assuming full information sharing with a central coordinator. Generally, these studies assume that centralized optimization is possible, and compare a centrally computed optimal solution for the cooperative scenario with the corresponding non-cooperative solution.

The second part contains approaches assuming partial information exchange between carriers and a central coordinator. Although not all information is shared, the coordinator can contribute to a solution by, for instance, proposing order exchanges or iteratively updating prices. Hence, the central coordinator supports the process of cooperation but has no full control. In general, this category corresponds to the category of decentralized approaches with central auctions as described in Section 1.2.

Table 3.1: Overview of information sharing properties in collaborative transportation. Explicit comparisons of different levels of information sharing are given in boldface.

Category	Reference	Problem characteristics						Dynamism		Information sharing				
		R	T	P	L	#Ord	#Carr	#Veh	O	V	A	N	I	B
Centralized optimization	Dahl and Derigs (2011)	✓	✓	✓	✓	⊗	~ 50	8905	✓		F	C	⊕	full information; system suggests exchanges
	Molenbruch et al. (2017)	✓	✓	✓	✓	400	4	32			F	C	⊕	full information; system computes solution
	Montoya-Torres et al. (2016)	✓				⊕ 61	3	3			F	C	⊕	full information; system computes solution
	Schulte et al. (2017)	✓	✓	✓		10–75	4–50	4–50			F	C	⊕	full information; system computes solution
Centralized support	Berger and Bierwirth (2010)	✓		⊕	< 100	3	3			P	C	✓	marginal profit for (bundles of) orders	
	Dai et al. (2014)	✓	✓	✓	✓	15–24	3	3–30		P	C	✓	yes-no bids + variable: profit	
	Gansterer and Hartl (2018a)	✓		✓	✓	30–210	3	3		P	C	✓	marginal profit for bundles	
	Gansterer et al. (2020a)	✓		✓	✓	30–90	3–6	9–18		P	C	✓	marginal profit for bundles + variable: aggregate geographical data, marginal profit for requests	
	Lai et al. (2017)	✓		✓		30–245	3–24	∞		P	C	✓	request portfolio, marginal profit for order	
	Li et al. (2015)	✓		✓	✓	9–15	3	6		P	C	✓	marginal profits for outsourcing and sourcing orders	
	Lyu et al. (2019)	✓	✓	✓	✓	9–45	3	9		P	C	✓	outsourcing prices, yes-no bids + variable: profit	
	Wang and Kopfer (2014)	✓	✓	✓	✓	104–266	2–5	19–61		P	C	✓	number of vehicles, costs of candidate routes	
	Wang and Kopfer (2015)	✓	✓	✓		~ 1767	⊗	⊗	✓	P	C	✓	number of vehicles, costs of candidate routes	
	Decentralized development	Dai and Chen (2011)	✓	✓	✓	✓	9	3	3-30	✓		P	D	✓
Figliozzi et al. (2004)			✓	✓		⊗	4	8	✓		P	D		marginal cost for order
Figliozzi et al. (2005)			✓	✓		⊗	⊗	4	✓		P	D		adapted marginal cost for order (also broadcasted)
Gath (2016)			✓		✓	100–200	⊕	3–29			P	D	✓	marginal cost for order
Máhr et al. (2010)			✓	✓		65	⊕	40	✓	✓	P	D	✓	marginal cost for order
Mes et al. (2013)			✓	✓		⊗	10	10	✓		P	D	✓	marginal cost for order (including opportunity costs)
This chapter			✓	✓	✓	1000	75–150	75–150	✓	✓	P	D	✓	variable: marginal cost for order, location information

R: Reallocation of orders (i.e., there is an initial assignment of orders to carriers); **T:** Problems with time windows; **P:** Problems with pickups and deliveries; **L:** Less than truckload problems; **#Ord:** Number of orders; **#Carr:** Number of carriers; **#Veh:** Number of vehicles; **O:** Dynamic orders; **V:** Dynamic vehicles; **A:** Amount of information sharing (**F:** full information sharing; **P:** partial information sharing); **N:** Nature of information sharing (**C:** centralized information sharing, i.e., with a central coordinator; **D:** decentralized information sharing, e.g., with other carriers or with shippers); **I:** Iterative information sharing (multiple rounds); **B:** Sharing of information regarding bundles of orders (combinatorial approach); **⊗:** Not available; **⊕:** Not applicable.

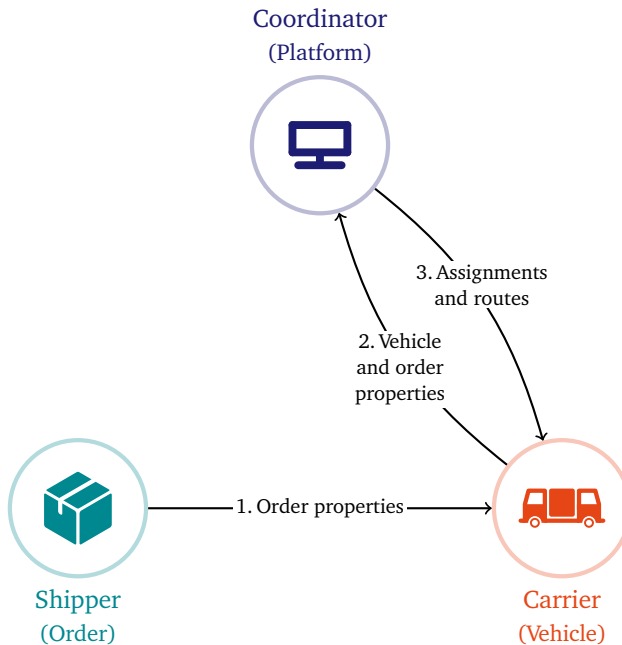


Figure 3.1: Communication between shipper, carrier, and coordinator in a centralized optimization scenario.

The third part considers partial information exchange on a local level, for example, between carriers among each other or between carriers and shippers. Generally, solutions are developed in a decentralized way, while a platform is mainly needed for connecting the different participants and for providing a reliable communication framework. Still, the platform can be in charge of minor but important decisions. This category mainly corresponds to the category of decentralized approaches with local auctions as described in Section 1.2.

Throughout the three categories, we furthermore distinguish between methods with iterative information sharing and methods with a single round of information exchange, and between methods where aggregate information is being exchanged (information regarding bundles of requests) and methods with information exchange on the level of a single request.

3.2.1 Full information sharing with a central coordinator

In approaches with full information sharing, a central coordinator is aware of all information about orders and vehicles (see Figure 3.1). While this coordinator computes and imposes a global solution in most approaches (Montoya-Torres et al., 2016; Schulte et al., 2017; Molenbruch et al., 2017), it suggests order

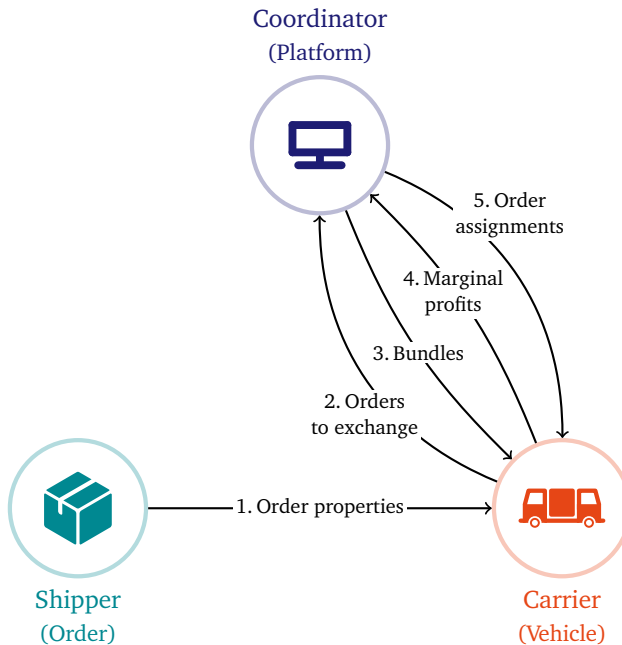


Figure 3.2: Communication between shipper, carrier, and coordinator in a combinatorial auction, a typical example of centralized support.

exchanges in the approach of Dahl and Derigs (2011). Individual carriers might accept or deny these exchanges, and hence stay autonomous. (The approach of Dahl and Derigs (2011) thus tends to the category of centralized support.) The majority of studies in this area focuses on the difference in profit between a scenario without cooperation (i.e., each carrier serves its own orders), and a fully cooperative scenario for which full information sharing is necessary. Although the initial assignment of orders to carriers is fixed in general, Molenbruch et al. (2017) compare different initial assignments for dial-a-ride problems.

3.2.2 Partial information sharing with a central coordinator

Several methods apply partial information sharing with a central coordinator (see Figure 3.2). The coordinator does not have full control, but has a high-level overview over (part of) the problem since it collects information from several carriers. It supports the carriers by proposing or establishing profitable collaborations.

In combinatorial auctions with a central auctioneer, each carrier can submit some requests to a request pool, and subsequently bid its marginal profit for the bundles of requests that the auctioneer generated (Berger and Bierwirth, 2010;

Gansterer and Hartl, 2018a; Gansterer et al., 2020a). Hence, aggregated profit information is exchanged only once with a central coordinator.

Li et al. (2015) also consider a central coordinator, but focus on single request exchange. Carriers submit their request with lowest marginal profit to a pool, and subsequently bid their marginal profit for all requests in the pool. In each iteration, the coordinator selects the request with highest profit gain for exchange. Hence, carriers iteratively need to share marginal profit information with the coordinator.

Lai et al. (2017) reverse the order: carriers may bid their marginal profit on single requests that they want to obtain, and based on the demand and price, every carrier decides which bundle of requests it is willing to sell. The coordinator then tries to find the best match. A drawback of this approach is that all carriers have to disclose their full request portfolio.

Dai et al. (2014) consider a central auctioneer that proposes outsourcing prices for requests, and iteratively updates the prices based on the bidding behaviour of the carriers. The auctioneer may adapt the prices with or without knowing profit information from the carriers.

Lyu et al. (2019) propose a similar system: carriers determine outsourcing prices for bundles of requests, and may bid on all bundles in the request pool. An auctioneer then solves the winner determination problem, either based on profit information from the carriers, or without this information.

In the approach of Wang and Kopfer (2014, 2015), carriers submit requests that they want to exchange to a pool, and then iteratively propose prices for bundles of requests (to be served in one route) that they want to obtain. The information that is being exchanged is comparable to that in combinatorial auctions, but carriers themselves can propose the bundles now. However, they also need to reveal their number of vehicles to the central coordinator, since the assigned number of bundles per carrier may not exceed the number of available vehicles.

3.2.3 Decentralized partial information sharing

Whereas all communications go via a central coordinator in the previously described approaches, a third category of information sharing is fundamentally different: interaction takes place directly between the cooperating participants, without the intervention of a central coordinator (see Figure 3.3). The main role of the coordinator is to connect the participants to each other, while the participants themselves locally develop a solution. However, the coordinator can have a crucial influence in the selection of participants in large-scale scenarios.

In the approach of Dai and Chen (2011), carriers offer their requests at an outsourcing price to all other carriers. Based on the bids of other carriers, the

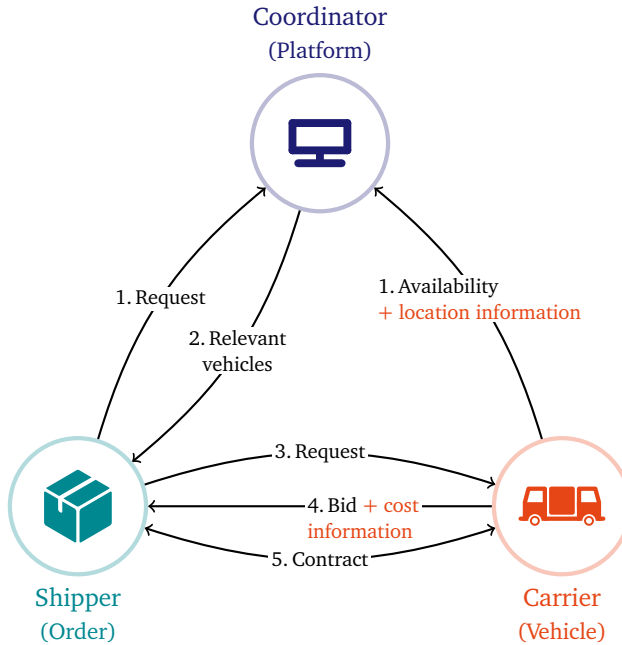


Figure 3.3: Communication between shipper, carrier, and coordinator in our multi-agent approach, a typical example of decentralized solution development. Location information sharing and cost information sharing are optional within our approach.

outsourcing price is adapted, or the request is exchanged. Hence, carriers do not include profit information in their bids, but they reveal private information by iteratively communicating outsourcing prices.

Although most approaches assume initial assignment of requests to carriers, Máhr et al. (2010) and Gath (2016) consider cooperative approaches where orders are not initially assigned. In their approaches, orders act as individual agents that try to find the most appropriate vehicle by organizing auctions, and carriers share their marginal costs only with the order agents. Although no direct cost information is being exchanged between carriers, they might infer information of other carriers from iterative auctions. Furthermore, Máhr et al. (2010) apply an improvement step in which carriers directly share uncompleted parts of their routes.

Mes et al. (2013) consider a comparable mechanism, but order agents are allowed to delay commitments if the best bid is still higher than a reservation price. Furthermore, carriers can reject already accepted orders, and consider the possible impact of accepting an order on future revenues. Hence, opportunity costs are included within their bids. Figliozzi et al. (2004, 2005) also consider a marketplace without initial assignment of orders, but focus on competitive carrier

behaviour. Carriers do not need to reveal their real costs, but may bid strategically, based on experience from earlier auctions or even based on all previous bids of all other carriers (Figliozzi et al., 2005).

3.2.4 Comparing information sharing policies

In all considered approaches, the necessary information is simply assumed to be shared by the carriers. Although this information is slightly different from approach to approach, there is hardly any assessment of the value of information sharing, as considered important by Gansterer and Hartl (2018b). To the best of our knowledge, an explicit assessment of the value of information sharing has only been done by Gansterer et al. (2020a), Lyu et al. (2019), and Dai et al. (2014).

Gansterer et al. (2020a) consider a combinatorial auction and show that it is beneficial for a carrier to share aggregate geographical information about its requests, since other carriers might consider to submit requests to the pool that are close to these requests. Furthermore, it is beneficial to share the marginal profit of individual requests within the pool with the auctioneer, since the auctioneer can generate more adequate bundles based on this information. The experiments, however, were performed on rather small static instances.

Lyu et al. (2019) compare a setting in which carriers do not reveal any profit information to the auctioneer with a setting in which carriers share their marginal profit for both outsourced and demanded bundles of requests. In the first case, the winner determination problem maximizes the number of bundles being exchanged. In the second case, the total profit increase is maximized. Lyu et al. (2019) do not find a significant difference between the two scenarios, but they test on very small instances of 3 carriers with 5 requests each. In fact, both scenarios obtain the same results as a centralized planning approach in most cases.

The different information sharing assumptions in the approach of Dai et al. (2014) are similar: within the conflict resolution step of the winner determination problem, the auctioneer favours carriers with the largest profits if profits are known. Otherwise, it selects carriers randomly. Dai et al. (2014) conclude that the scenario with profit information sharing yields better results than the scenario without this information. However, their instances are rather small (3 carriers, 5 or 8 requests each), and in more than half of the tests, both methods find the centrally computed optimum.

In this section, we have classified the collaborative vehicle routing literature based on information sharing properties. We identified full information sharing with a central coordinator (generally used for central optimization), partial information sharing with a central coordinator (generally used for centralized

support), and decentralized partial information sharing (generally used for decentralized solution development). The literature review has shown that there is a lack of systematical assessments of types and levels of information sharing within a collaborative routing approach for complex, large-scale, dynamic problems. This chapter contributes to filling this gap by explicitly proposing different levels of location and cost information sharing, and analyzing their influence on solution quality.

3.3 Information types

If carriers are concerned with their privacy, they might not want to share (part of) their properties with the platform. Therefore, we propose to construct vehicle routes on the local level in a MAS approach: transportation contracts are made through local interactions between carriers and shippers (see Figure 3.3). To determine the influence of different levels of information sharing in such an approach, we define several sharing policies for carriers. We distinguish between information about locations and route plans of vehicles (Section 3.3.1), and information about travel costs of the vehicles (Section 3.3.2). Although the information on locations and travel costs is not necessary for the mechanism to work, we expect that more information will improve solution quality.

3.3.1 Vehicle route information

Information about the current and future locations of vehicles might be of importance for shippers, to be able to select the most appropriate vehicles. We consider the following three (incrementing) levels of location information that carriers might share with the platform (see Figure 3.3), and discuss the potential for shippers when this information is known.

- **No position sharing (NPS):** If carriers do not share any information about their vehicle positions, shippers cannot infer which vehicles will be most appropriate.
- **Current position sharing (CPS):** If carriers share their actual vehicle positions, shippers can select the vehicles that are closest to their pickup locations for interaction. Although there is no guarantee that the closest vehicle will transport the order in an optimal solution, nearer vehicles are in general more likely to have lower detour costs for the order.
- **Full plan sharing (FPS):** If carriers share their complete route plans (i.e., give access to $\langle \rho^{vt} \rangle$), shippers might construct even better estimates of

the appropriateness of a vehicle. The spatiotemporal distance of the vehicle route to either the pickup or the delivery location can be used as a heuristic: if the vehicle has already planned to be in the neighborhood of the pickup or delivery location of an order, just at a time that is within or close to the pickup or delivery time window, the order might be included into the route at relatively low extra costs. Formally, the spatiotemporal distance of a pickup or delivery location i to a vehicle route $\langle \rho^{vt} \rangle$ is given by $\min_{u \in [t, t^v], \hat{u} \in [e_i, l_i]} \theta |u - \hat{u}| + c_i \hat{\rho}^{vu}$, where $\hat{\rho}^{vu}$ is the location of vehicle v at time u , and θ is a parameter controlling the relative importance of time.

3.3.2 Marginal cost information

Information about the marginal travel costs of a vehicle for an individual order is relevant for the global solution: if an order can be served by another vehicle with lower marginal travel costs, this will decrease the total travel costs. Although there are not always direct advantages of sharing marginal travel costs for a carrier, system-wide costs may be reduced on the longer term. If gains are redistributed, this can be beneficial for individuals as well (see Section 2.1.1). In other scenarios, however, carriers might want to keep their cost information confidential, or want to provide their marginal costs only if the order will contribute to an increase of the local carrier profit PR_c . We consider the following three (decreasing) levels of cost sharing.

- **Full cost sharing (FCS):** A carrier is always willing to share cost information, even if this does not improve its local profit PR_c directly.
- **Partial cost sharing (PCS):** A carrier is only willing to share cost information if this can contribute to an increase of the local carrier profit PR_c .
- **No cost sharing (NCS):** A carrier is never willing to share cost information.

Note that cost information is not shared with the platform, but only with the individual shippers (see Figure 3.3), within an auction system that will be described in the next section.

3.4 Auction approach

To investigate the value of different types and amounts of (locally) shared information, we introduce a MAS for solving the DCPDP in a decentralized manner, in accordance with the assumption that vehicles and orders can independently attach to the platform without a central authority that is aware of all information. For each order, an order agent is introduced that needs to make a contract for

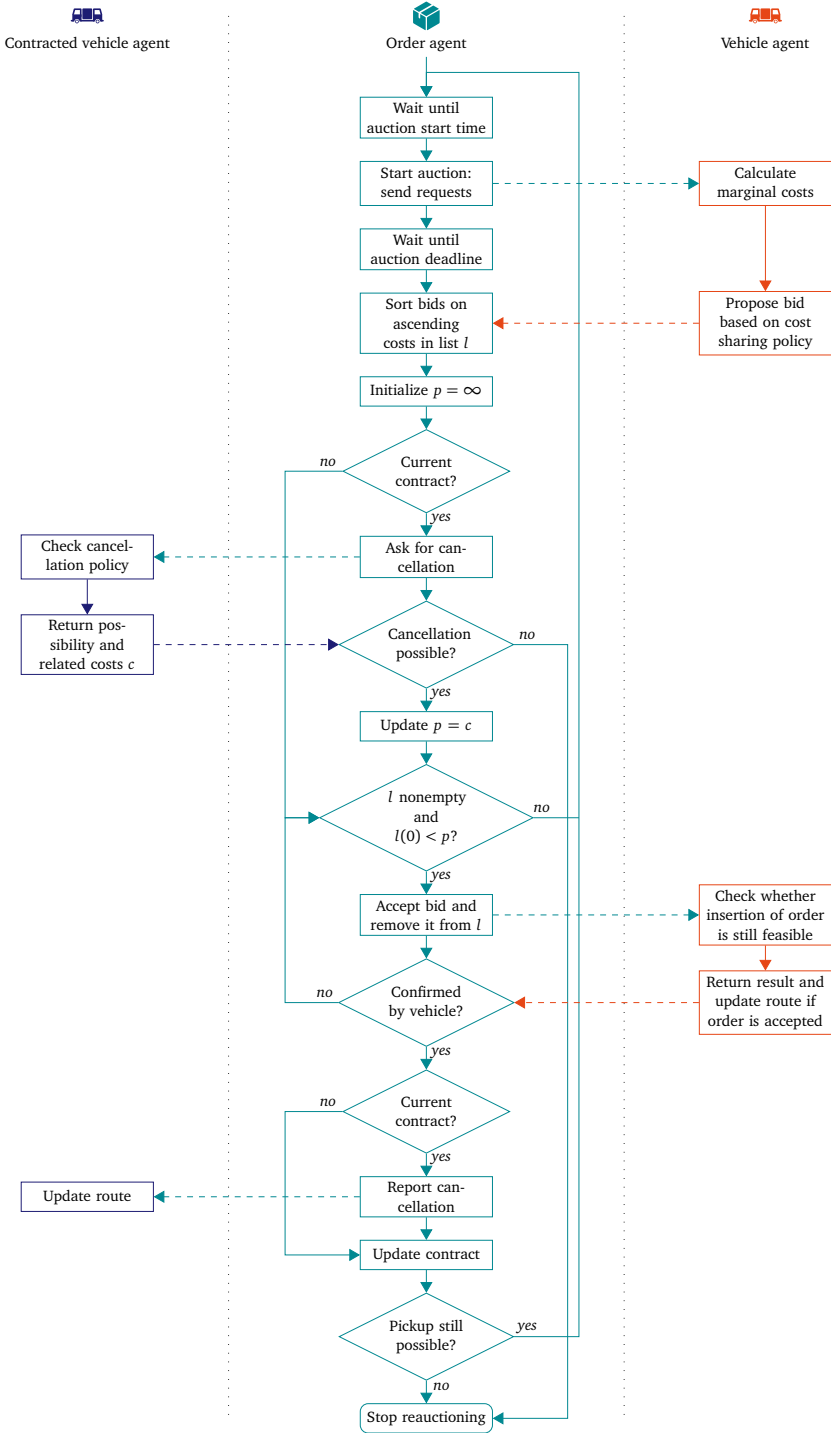


Figure 3.4: Flowchart of the iterative auction procedure within the multi-agent system. Dashed arrows represent exchange of information between different actors.

transportation. Each carrier, in this chapter limited to having only one vehicle, is represented by a vehicle agent, that has the goal of making some profit by efficiently serving orders. Order agents and vehicle agents try to make contracts by interacting in multi-agent auctions: order agents act as auctioneers by offering a transportation task; vehicle agents respond by proposing a bid (see Figure 3.4)¹. This approach allows fast responses within a dynamic supply and demand market. In addition to the MAS approach introduced in Section 2.2, we now vary the information that is being exchanged per auction, and limit the number of carriers that are consulted in each auction. Although there is no guarantee on optimal solutions, our MAS is competitive with standard centralized heuristics for the Dynamic Pickup and Delivery Problem (DPDP): we found improvements of up to 4% on the benchmark instances of Mitrović-Minić et al. (2004) (see Appendix B).


3.4.1 Order agent

An order agent for an order $o \in O$ is responsible for making and, if possible, improving a transportation contract with a vehicle agent. At time r_o , it sends a request, containing p_o and d_o (with the corresponding time windows and service durations), the quantity q_o , the price f_o , and an auction end time, to a well-selected group of vehicle agents. While in current MAS approaches a request is sent to all vehicle agents (Máhr et al., 2010), a lot of computational effort might be saved if the order agent interacts only with a restricted group of most promising vehicles, especially if the number of vehicles is large. Therefore, the platform selects a certain percentage of the currently available vehicles for the order agent, based on the available vehicle location information (see Section 3.3.1):


- In case of NPS, the platform cannot do better than selecting vehicles randomly (see Figure 3.5a).
- With CPS, the platform selects the vehicles that are closest to p_o at the time of the request (see Figure 3.5b).
- For FPS, the platform selects the vehicles with lowest spatiotemporal distance of the vehicle route to either p_o or d_o (see Figure 3.5c).

After communicating the request to the selected vehicles, an order agent waits for bids from the vehicle agents. At the auction end time, the received

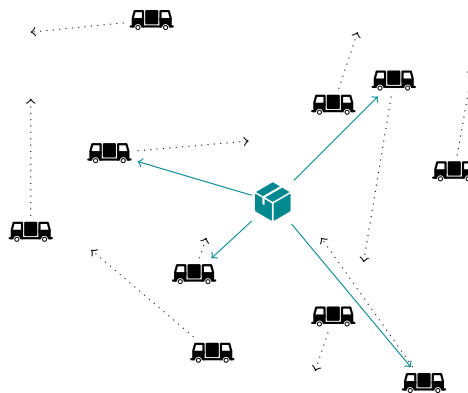
¹In line with other literature on MASs for logistics (e.g., Van Lon and Holvoet, 2017), we use the auction terminology, although we do not focus on a real competitive setting yet where we aim for incentive compatibility.



(a) No position sharing.



(b) Current position sharing.



(c) Full plan sharing.

Figure 3.5: Selection of 40% of the available vehicles under different location information sharing policies. A solid arrow represents selection of a vehicle; a dotted arrow represents the route plan of a vehicle.

bid with lowest value is chosen (this decision might be random if the bids do not contain values; see Section 3.4.2). The order agent asks the vehicle agent that submitted this bid to incorporate the task into its route. Since the route plan of the vehicle has possibly been changed between proposing the bid and receiving the acceptance message, the vehicle agent needs to check if the bid has not become outdated. If insertion is still feasible, they establish a contract. Otherwise, the order agent chooses the vehicle agent of the second-best bid, and so on, until a contract is made or no feasible bid remains.

After each auction, the order agent schedules to start a new auction after some time, since a more suitable vehicle might appear in the dynamic world. If an order agent already has a contract, it asks the contracted vehicle agent at the end of the auction if cancellation of the contract is still possible, and what the actual costs are. Only if cancellation is still possible and the costs of the best new bid are lower than these actual costs, the contract will be replaced. If contract cancellation is no longer possible (since the pickup has already taken place or it is too expensive for the vehicle agent) or the order agent did not succeed in making a contract at time l_{p_o} , the order agent stops scheduling new auctions.


3.4.2 Vehicle agent

A vehicle agent for a vehicle $v \in V$ (representing a carrier $c \in C$) is responsible for constructing a feasible route, while maximizing profit by accepting or rejecting orders. As discussed in Section 2.1.1, carriers can merely focus on their own profit, or be more cooperative and focus on the global profit. This difference is expressed in the bidding strategy of the vehicle agent.


When the vehicle agent of carrier c receives a request with auction end time t from the order agent for order o , it computes the marginal cost $MC_c^t(o)$ for this order with respect to its current route plan at time t (i.e., the additional travel costs for including o into its route after time t).² Furthermore, the marginal profit $MP_c^t(o) = f_o - MC_c^t(o)$ is computed. If including the order is feasible, the vehicle agent returns a bid to the order agent, based on the specific cost sharing policy (see Section 3.3.2):

- With FCS, the vehicle agent always submits a bid with $MC_c^t(o)$, even when the marginal profit $MP_c^t(o)$ for including o is negative (see Figure 3.6a). The vehicle agent is fully cooperative in the sense that it allows the order agent to make a contract with the vehicle agent with lowest marginal cost.

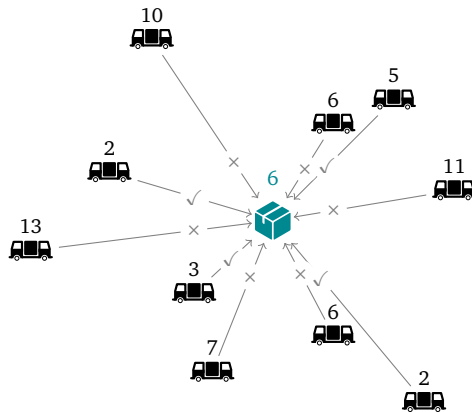
²Although an exact or metaheuristic approach for computing the marginal cost might be used in real-world applications (where each vehicle agent has its own computational resources), we approximate it by the elementary insertion heuristic presented in Section 2.2 throughout our simulations.



(a) Full cost sharing.



(b) Partial cost sharing.



(c) No cost sharing.

Figure 3.6: Vehicle bids given their marginal costs and the order price under different cost sharing policies. An arrow with a number, checkmark or cross represents a bid with marginal cost, a bid without any value, or no bid at all, respectively.

- With PCS, the vehicle agent only submits a bid with $MC_c^t(o)$ if $MP_c^t(o) > 0$ (see Figure 3.6b). If the marginal profit is negative, no response is given. Hence, the vehicle agent focuses more on PR_c than with FCS, but is still cooperative in the sense that it allows vehicle agents with lower marginal cost to make a contract with the order agent.
- For NCS, the vehicle agent submits a bid if $MP_c^t(o) > 0$, but does not include the marginal cost within the bid (see Figure 3.6c). This means that the order agent cannot compare the different bids and needs to select one at random.

When the vehicle agent of carrier c has a contract with the order agent of order o and receives a cancellation request of this order agent for time t , it checks whether the order has not yet been picked up at time t and whether its cancellation policy allows cancellation at time t . Only in this case, the contract may be terminated. Cancellation policies depend on the cost sharing policy as follows:

- With FCS, the vehicle agent returns the actual marginal cost $MC_c^t(o)$ and allows cancellation if the marginal cost of another vehicle for the order is lower.
- With PCS, the vehicle agent computes $MC_c^t(o) - \phi \cdot f_o$, for $\phi \in [0, 1]$. If this value is negative, the vehicle agent will not allow cancellation. Otherwise, this value is returned as actual cost. For low values of ϕ , the vehicle agent is rather cooperative, but probably shares too much information. For high values of ϕ , the vehicle agent is more reticent regarding cancellation and shares in general less information. Only if the new vehicle has a significantly lower marginal cost, a cancellation is allowed. If $\phi = 1$, cancellation is never possible, since the vehicle agent would not have submitted a bid if the marginal cost is higher than f_o .
- With NCS, cancellation is only allowed if $MC_c^t(o) - \psi \cdot f_o > 0$, for $\psi \in [0, 1]$. The vehicle agent does not share any value. Since a new vehicle is selected randomly, the value of ψ needs to be not too low: this would result in multiple contract cancellations, although removing an order is likely to decrease the quality of a route. On the other hand, if $\psi = 1$, cancellations are never allowed, while some successful reauctions might improve the global solution quality.

If a vehicle agent receives an acceptance message from an order agent, it checks whether the order still can be placed into the route. If so, a contract is made, and the vehicle agent updates the route plan for the vehicle accordingly.

3.5 Computational study

In this section, we compare different information sharing scenarios. First, we explain how our problem data set was generated and which parameter settings were used. Then, we compare the different position and cost information sharing scenarios and examine how various parameters influence them. Finally, we consider a scenario with mixed information sharing policies.

3.5.1 Problem instances

To compare the information sharing policies, we generated a set of 10 DCPDP instances with time windows, quantities, and order prices.³ The instances were defined on an approximately 1000×1000 area with a 10h time horizon, and contain each 1000 orders and 150 vehicles. Recall that each vehicle represents a different carrier. Half of the vehicles is available during the complete time horizon; the other half has randomly selected r^v , e^v , and l^v . All vehicle capacities were set to 100.

Order quantities were sampled from $\mathcal{N}(20, 5^2)$, service durations (in seconds) for pickups and deliveries from $\mathcal{N}(120, 30^2)$, and time window lengths (in seconds) for pickups and deliveries from $\mathcal{N}(9000, 1800^2)$. The x and y coordinates of all pickup and delivery locations and vehicle start locations were sampled from $\mathcal{N}(500, 75^2)$ with probability 0.5, from $\mathcal{N}(200, 75^2)$ with probability 0.25, or from $\mathcal{N}(800, 75^2)$ with probability 0.25, resulting in 9 clusters of locations. All instances represent an open DCPDP, that is, end locations for vehicles are not prescribed.

Travel times t_{ij} equal the Euclidean distance between locations. Furthermore, $c_{ij} = 0.011t_{ij}$ for all locations i and j , and $f_o = 0.014t_{p_o d_o}$ for all orders $o \in O$.

We refer to the instance set described above as the *base set*. In Section 3.5.6, we check whether results are generalizable to other scenarios. For this, we created three related problem instance sets. First, we created two sets with lower transport capacity by reducing the number of vehicles. These *low capacity set* and *medium capacity set* are copies of the base set in which only 75 or 100 of the vehicles per instance are kept, respectively. The base set acts as a high capacity set where all orders could easily be served.

Second, we created a set in which the requests are more urgent. This *urgent set* is a copy of the base set in which the latest service times for pickups and deliveries have been changed according to a sampling from $\mathcal{N}(900, 300^2)$ instead of $\mathcal{N}(9000, 1800^2)$ for the time window lengths. Earliest service times are kept

³Available at <http://doi.org/10.4121/12763868>.

the same. The base set acts as corresponding non-urgent set with more time between release times and deadlines of orders.

3.5.2 Experimental settings

Throughout our experiments, we used the following default settings. The time between the start of an auction and the auction deadline was set to 10s; vehicle agents compute results for 1s after the actual auction deadline to make sure that the bid is still feasible after auction processing and acceptance communication. If an order agent tries to make a contract with a vehicle agent but the bid is no longer feasible, the new auction deadline is set to 1ms after the current auction deadline. The time between an auction deadline and the start time of a reauction for an order o is $(l_{p_o} - r_o)/m$, where m denotes a maximum number of auctions per order, which was set to 10.

The standard number of vehicles that an order agent interacts with is 10% of the currently known vehicles. For computational purposes, we aggregate over regions of 50×50 distance units and a time of 1800s instead of computing the exact spatiotemporal distances when selecting vehicles.

If a vehicle agent determines a new route plan and has some temporal flexibility in its schedule, it waits for 20% of the available time at its current position. By this, orders that appear in the neighborhood of the vehicle soon after plan determination might still be included into the route, but most of the available flexibility can be used to change plans when the vehicle is already driving.

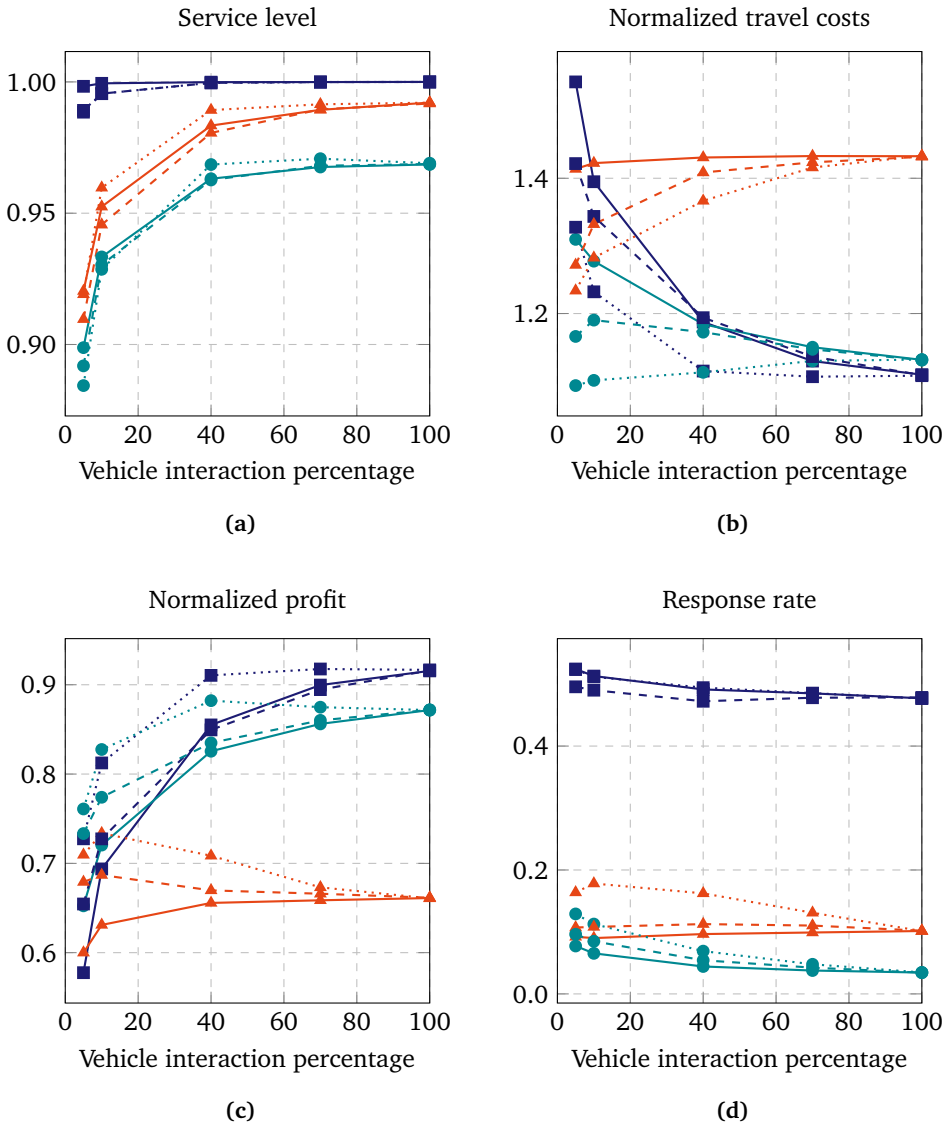
The values of ϕ and ψ were set to 0.2 and 0.6, respectively.

Instead of processing the auctions in real time, we keep track of the times and give all processes enough time to fully complete. In general, the total time needed for the experiments with this approach is much less than the real total time (one run with FCS and the settings described above takes about 25s, one run with interaction with 100% of the vehicles takes about 180s, and computation times for PCS and NCS are even lower).⁴ Furthermore, when running in real time, vehicle agents do not always have enough time to place a bid, since they all need to compute it at the same moment but share a single processor. In real-world applications, this will not be a problem, since computations can be distributed over multiple processors.

3.5.3 Information sharing scenarios

To compare the influence of the different information sharing policies, we combine each of the three vehicle position sharing policies (NPS, CPS, and FPS) with

⁴The experiments were performed single-threaded on a Linux pc with a 3.30 GHz Intel i5-4590 CPU and 8 GB of RAM.



<i>Cost sharing</i>		<i>Position sharing</i>	
■	Full cost sharing (FCS)	—	No position sharing (NPS)
●	Partial cost sharing (PCS)	- - -	Current position sharing (CPS)
▲	No cost sharing (NCS)	⋯	Full plan sharing (FPS)

Figure 3.7: Mean results on varying VIP for the three cost sharing and the three position sharing methods. In Figure 3.7c, there is no fine for rejected orders ($\gamma = 0$).

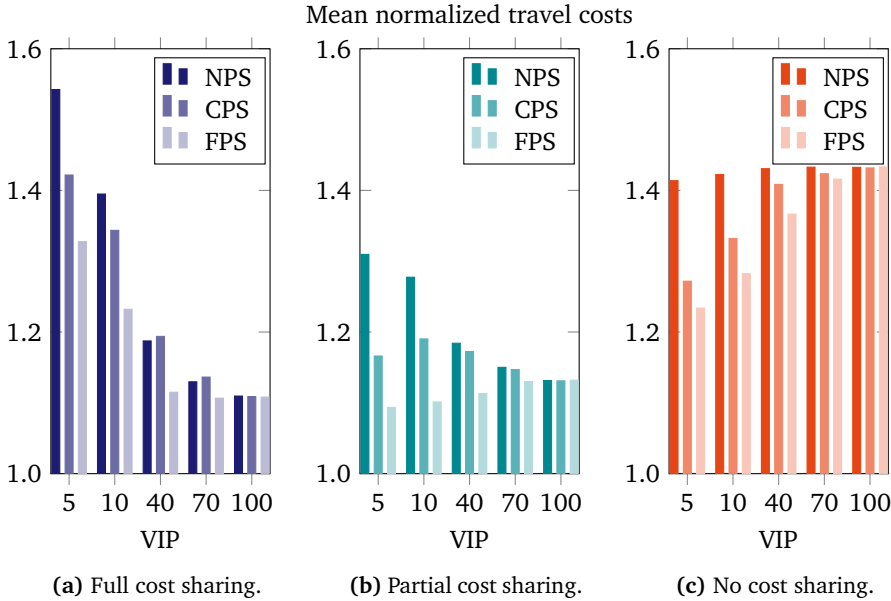


Figure 3.8: Mean normalized travel costs on varying vehicle interaction percentage (VIP) for the three cost sharing and the three position sharing methods.

each of the three cost sharing policies (FCS, PCS, and NCS), resulting in 9 scenarios. Furthermore, we want to compare the influence of the number of vehicles that each order interacts with, and hence run each scenario with different vehicle interaction percentages of the total known fleet.

To be able to compute average results over the different instances, we normalize the travel costs and profit as follows. For each instance, we compute 10 solutions by running the MAS with parameter settings that likely give a high quality solution. We use FCS, 100% vehicle interaction, and a fast reauctioning time of 60s between the deadline of an auction and the start time of a reauction for the order. Per instance, the solution with highest service level (and, in case of a tie, lowest travel costs) is used as best known solution (BKS). Let R be the solution for a run of the MAS on a problem instance, then normalized travel costs and normalized profit for R are defined by $TC(R)/TC(BKS)$ and $PR(R)/PR(BKS)$, respectively.

The results in terms of service level, normalized travel costs, and normalized profit for the different information sharing scenarios are shown in Figure 3.7, where we vary the percentage of (currently known) vehicles that each order interacts with (VIP). All results are averages over 10 instances and 10 runs per instance. The response rate in Figure 3.7d is the average number of times that a vehicle agent sends a response (a bid) upon a single request of an order agent.

Figure 3.7a shows that the service level is almost always 1 if costs are always shared (FCS). For a VIP of 5% or 10%, the service level is lower if positions are shared (CPS or FPS) than if positions are not shared (NPS). This could be caused by a myopic vehicle selection heuristic: for CPS and FPS, the order agent selects only a limited group of closest vehicles within each reaction, while this group might not differ that much from auction to auction. With NPS, however, the group of contacted vehicles could be much larger since it is randomly selected each auction.

The policies in which costs are not fully shared, PCS and NCS, result in significantly lower service levels than FCS, especially for the lower VIPs, as expected: the bid submission rate is much lower than for FCS (see Figure 3.7d), and it is likely that some orders do not receive any feasible bid. Interestingly, the service level is higher for NCS than for PCS. An explanation can be that some advantageous vehicles remain unused and still are available for subsequent auction rounds with NCS (since order agents choose vehicles randomly), while order agents immediately select the best vehicles with PCS. This greediness can cause a local shortage of feasible vehicles on the long term. This explanation corresponds with the higher number of bids per request for NCS than for PCS (see Figure 3.7d).

Figure 3.7b and Figure 3.8 show decreasing travel costs with increasing VIP for the policies where costs are shared (FCS and PCS). Since the service level stays almost equal for FCS, we can conclude that a higher VIP results in better solutions. The communicational and computational load, however, is also higher. Without cost sharing (NCS), the travel costs increase with VIP, but the service level also does. Only for NCS with NPS, we can observe that a higher VIP is helpful, since the service level increases while the route costs are constant.

The same figures show a clear advantage of route information sharing: the route costs are in general lower for FPS than for CPS, and lower for CPS than for NPS (given the same cost sharing policy). Figure 3.7c shows the same pattern for the profit: FPS gives higher profit than CPS, and CPS results in higher profit than NPS, indicating that position sharing is profitable. This difference is the largest for lower VIPs. (Indeed, with a VIP of 100, position sharing is not relevant.) However, there is one exception: for FCS, NPS performs slightly better than CPS at a VIP of 40% and 70%. Again, this might be due to the larger range of vehicles that NPS selects within the subsequent reactions. Certainly, the received number of bids per request is a bit lower for CPS than for NPS with FCS (see Figure 3.7d), while this is not the case with PCS and NCS.

In general, the policy in which vehicle agents always share their costs (FCS) yields the best results. Since the profit is highest for FCS with the higher VIP values, vehicle agents have an incentive to use this cost sharing policy. Further-

more, sharing complete route plans (FPS) is beneficial: for a VIP of 40% with FPS, service level and profit are about the same as for a VIP of 100%, while the communicational and computational load is about 40% of it. Only with lower VIPs, using PCS is beneficial for vehicle agents compared to FCS.

3.5.4 Number of reauctions

A larger number of reauctions per order can increase the solution quality: due to the dynamic nature of the problem, combinatorial advantages might arise after some time. In Figure 3.9, we show the results of varying the maximum number of auctions per order (m) for different cost sharing methods and different fines for rejected orders (γ). We used a vehicle interaction percentage of 10% and no position sharing.

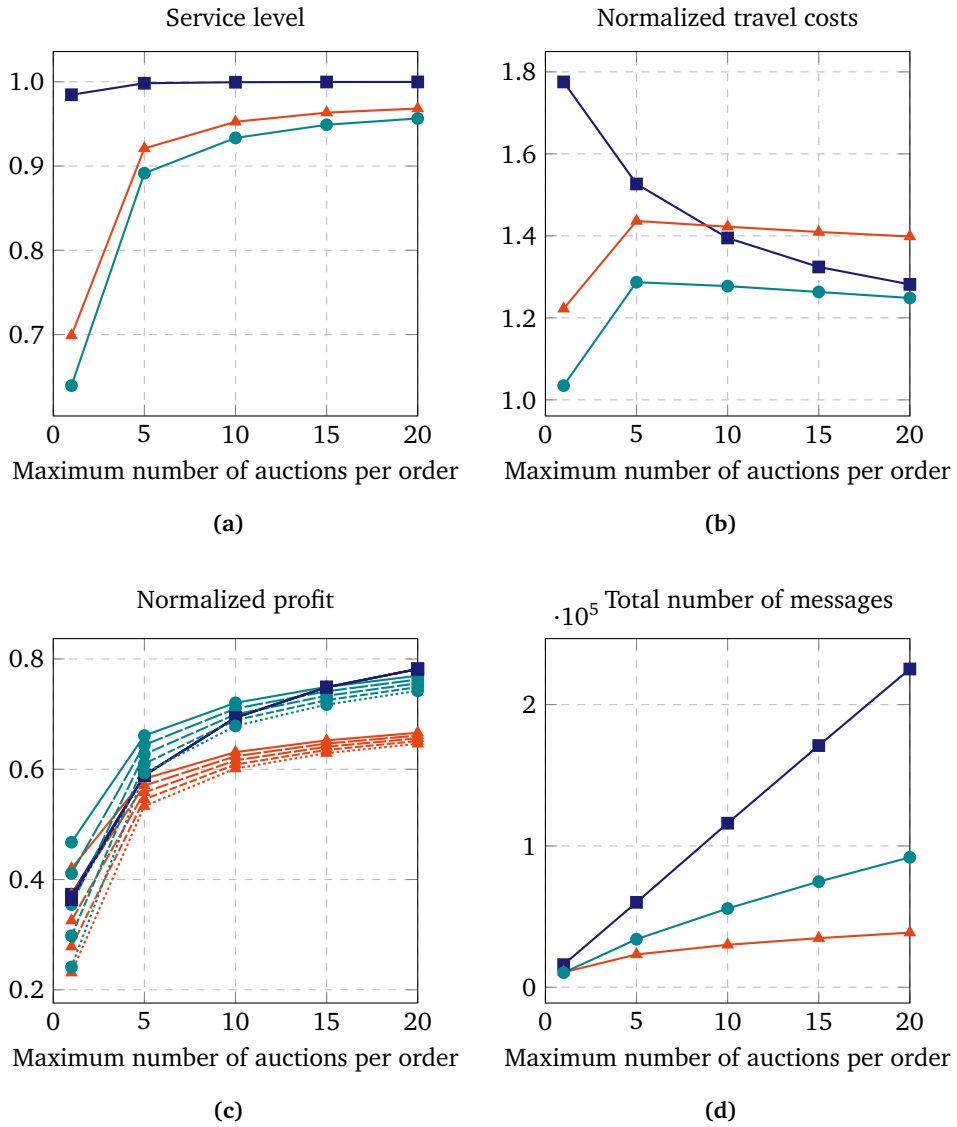
Omitting reauctions seriously deteriorates the service level if costs are not always shared: Figure 3.9a shows a steep increase in service level for PCS and NCS when m increases from 1 to 5. After 10 auctions, this increase is much smaller. When costs are always shared (FCS), however, the service level is always almost 1, even when reauctions are omitted.

For the travel costs, a similar pattern is visible (see Figure 3.9b): when m increases from 1 to 5, the travel costs increase a lot for PCS and NCS (since extra orders will be served). For $m > 5$, the travel costs decrease a little for PCS and NCS, while the service level still increases. Hence, the higher number of auctions results in slightly better solutions. For FCS, the costs decrease when increasing m , while the service level is almost the same. Hence, solutions are better with higher m , but the effect is largest for low values of m .

The relevance of a high number of auctions is also reflected in the profit (see Figure 3.9c): for all cost sharing policies, a higher m value yields higher profits. While PCS always outperforms NCS in terms of profit, it depends on the fine per rejected order (the value of γ) and m whether FCS outperforms PCS. In general, with a higher m value, a lower value of γ is needed to let FCS be superior. For $m = 5$, PCS performs better than FCS, even for $\gamma = 2$. For $m = 10$, however, PCS performs better than FCS for $\gamma \leq 1$, but worse for $\gamma \geq 1.5$. For $m = 15$, PCS is only better than FCS for $\gamma = 0$, and for $m = 20$, FCS outperforms PCS irrespective of the value of γ .


Hence, by enabling a larger number of auctions or setting a high fine for rejected orders, individual carriers might prefer the FCS policy over the PCS policy, resulting in a higher service level. A drawback, however, is that FCS in combination with a large number of auctions comes with the highest communicational load (see Figure 3.9d).

Furthermore, an increase in allowed auctions might result in an increase of sequential route changes for vehicles. This generally is no problem with au-

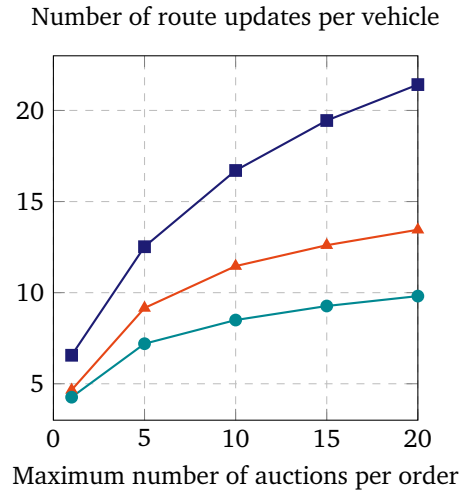


Cost sharing		γ			
■	Full cost sharing (FCS)	—	0.0	----	1.5
●	Partial cost sharing (PCS)	- - -	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		


Figure 3.9: Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ).




(a)



(b)



(c)



(d)

<i>Cost sharing</i>		<i>Position sharing</i>	
■	Full cost sharing (FCS)	—	No position sharing (NPS)
●	Partial cost sharing (PCS)	- - -	Current position sharing (CPS)
▲	No cost sharing (NCS)	⋯⋯⋯	Full plan sharing (FPS)

Figure 3.10: Mean numbers of plan updates per vehicle. Figures 3.10a and 3.10b show the total average number of plan updates (i.e., the number of times a vehicle agent inserts or removes a task), whereas Figures 3.10c and 3.10d show the average number of plan updates resulting in an immediate change in driving direction.

tonomous vehicles but might be uncomfortable for drivers. In Figure 3.10, we show the average number of route plan changes per vehicle under different circumstances. From Figure 3.10a, we observe that vehicles on average have more route updates with FCS than with NCS, and even less with PCS. Figure 3.10b shows that the number of route updates does increase with a larger maximum number of auctions, but not linearly. In Figure 3.10c, the number of route updates that require an immediate change of direction is shown. This number is very low compared to the total number of route updates. Interestingly, it is generally highest for CPS, followed by NPS and FPS, and decreases with increasing VIP. In Figure 3.10d, we observe that the number of turnoffs is very low when only one auction is allowed ($m=1$), and slightly higher (but not strictly increasing) for larger values of m .

3.5.5 Emission or congestion penalties

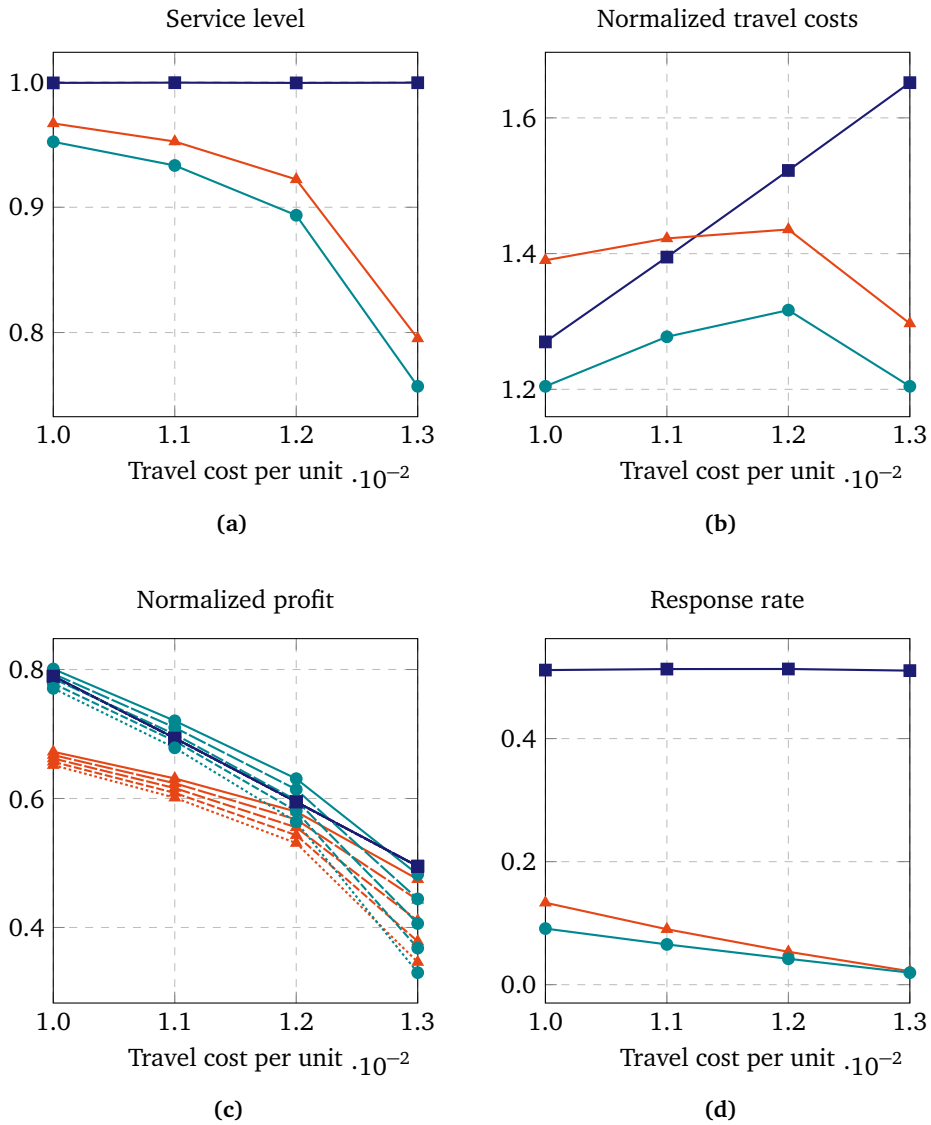
Since the cost information sharing policy (FCS) provides higher service levels than the policies where costs are not (always) shared (PCS and NCS), a platform owner or government might want to stimulate vehicle agents to share their costs. An extra charge on driven distance (next to the regular travel costs, for example, by applying taxes or specific emission or congestion penalties) can be expected to be helpful, since the bidding behaviour with FCS is not affected by travel costs, while the bidding behaviour with PCS and NCS is highly affected by travel costs.

In Figure 3.11, we show the results for varying values of the travel costs (including penalties) per unit (TCU) for different cost sharing methods and different fines for rejected orders (γ). Again, we used a vehicle interaction percentage of 10% and no position sharing.

Figure 3.11a shows that the service level decreases for PCS and NCS if TCU increases, as expected: the marginal profit for an order decreases if the marginal costs increase, and hence, offering a bid is attractive for less vehicles. Figure 3.11d confirms this observation: the number of bids that an order receives upon a request is lower for higher TCU values. The service level with FCS, on the other hand, is not dependent on the TCU, as expected. Also, the travel costs for FCS increase linearly with the TCU (see Figure 3.11b), and the profit decreases accordingly (see Figure 3.11c).

The profits for PCS and NCS decrease especially when TCU increases from 0.012 to 0.013 (see Figure 3.11c). This is as expected, because TCU is then close to 0.014, the price per unit that orders pay for transportation.

An interesting trade-off between FCS and PCS is mainly visible for these higher TCU values. While PCS results in higher profit than FCS if $TCU \leq 0.012$ and $\gamma \leq 0.5$ or even $\gamma \leq 1.0$, FCS results in highest profits when $TCU = 0.013$, irrespective of the value of γ .



<i>Cost sharing</i>		<i>γ</i>			
■	Full cost sharing (FCS)	—	0.0	-----	1.5
●	Partial cost sharing (PCS)	- - -	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure 3.11: Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ).

Setting TCU and γ wisely may result in an increased incentive for vehicle agents to fully share their costs, and hence improve the service level, without putting an extra communicational load upon the system (as it was the case with increasing the maximum number of auctions per order in Section 3.5.4).

3.5.6 Fleet capacity and order urgency

To analyze whether the results hold in other scenarios, we repeat the experiments on the low capacity set, the medium capacity set, and the urgent set (see Sect 3.5.1). Here, we describe the main findings; detailed results are given in Appendix C.

In Figure 3.12, we show the results on the high, medium, and low capacity set for different cost sharing policies. The position sharing policy is fixed as NPS, but the results for other position sharing policies are comparable.

The service level decreases if the available number of vehicles decreases, as expected. In all instance sets, FCS results in a higher service level than PCS, and PCS performs better than NCS.

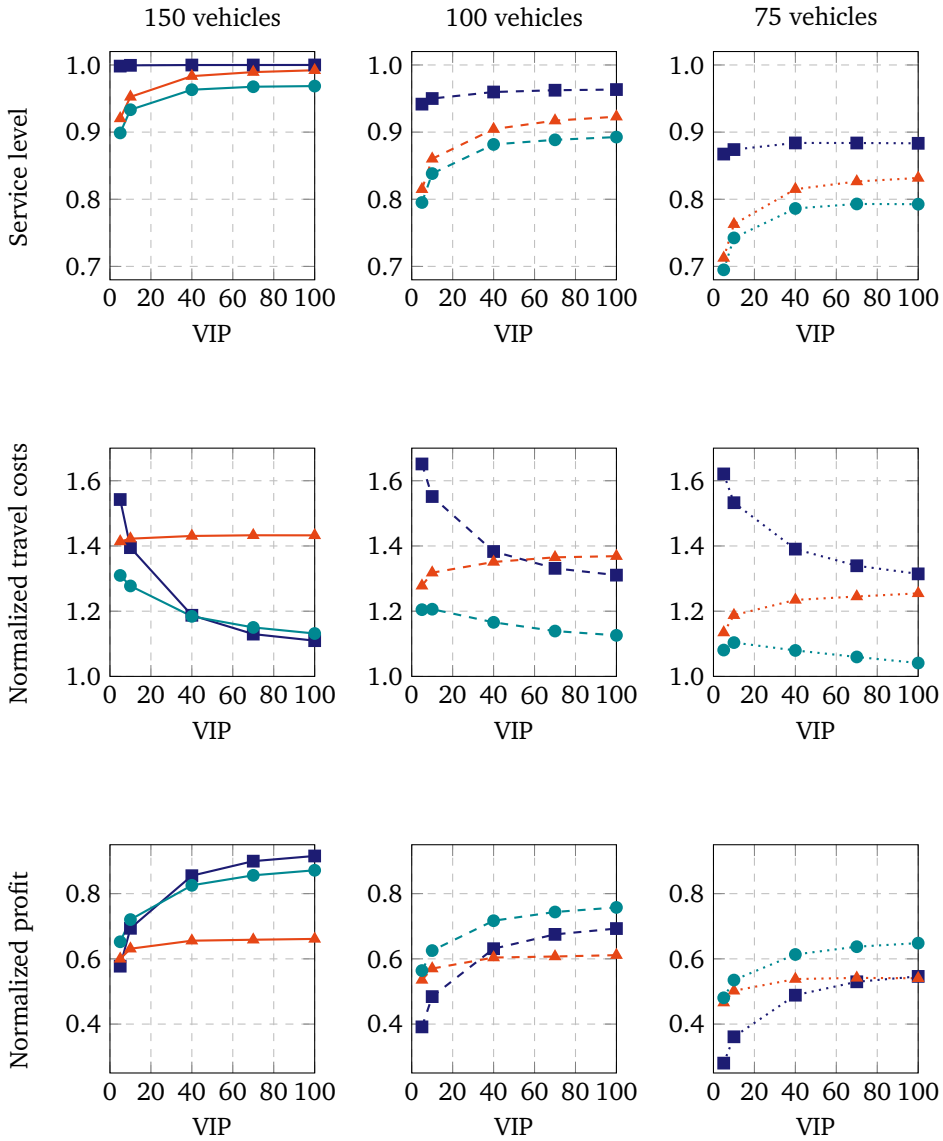
With respect to the travel costs, FCS performs worse if the vehicle capacity decreases: where FCS obtains lowest travel costs for a VIP of 70% or 100% in the high capacity set, it ends up between PCS and NCS for the medium capacity set, and even results in highest travel costs with the low capacity set. Note that the relative travel costs of PCS and NCS are not dependent on the ratio of supply and demand: PCS always results in lower travel costs than NCS.

A similar pattern is visible for the profit: with the high capacity set, FCS results in highest profits, whereas PCS outperforms FCS with the medium capacity set. Even NCS results in higher profits than FCS in most cases with the low capacity set.

For the urgent instances, it is expected that current position information is more important than for non-urgent instances. For the later ones, there might be enough time for efficient planning and replanning, while urgent orders need a vehicle relatively soon to be picked up before the latest service time. Only vehicles that are close enough to the pickup location are feasible ones.

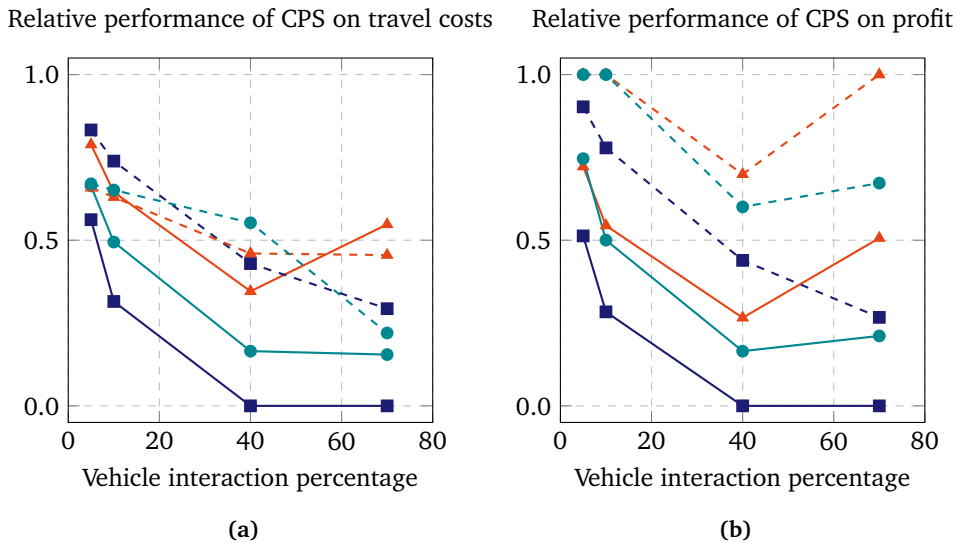
In Figure 3.13, we show the relative performance of the CPS policy on travel costs and profit with respect to the values obtained with the other position sharing policies. Per cost sharing policy and per interaction percentage, the relative performance of CPS on travel costs is calculated by $(TC_{\max} - TC_{\text{CPS}})/(TC_{\max} - TC_{\min})$, where TC_{\max} denotes the maximum of TC_{FPS} , TC_{CPS} , and TC_{NPS} , and TC_{\min} denotes the minimum of these three values. Equivalently, the relative performance of CPS on profit is given by $(PR_{\text{CPS}} - PR_{\min})/(PR_{\max} - PR_{\min})$.

From Figure 3.13a, we observe that CPS obtains relatively lower travel costs for the urgent set than for the non-urgent set. Only for NCS and a VIP of 5% or



<i>Cost sharing</i>		<i>Instance set</i>	
■	Full cost sharing (FCS)	—	High capacity set (150 vehicles)
●	Partial cost sharing (PCS)	- - -	Medium capacity set (100 vehicles)
▲	No cost sharing (NCS)	⋯	Low capacity set (75 vehicles)

Figure 3.12: Mean results on varying VIP for the three cost sharing methods on the different instance sets. All results are normalized with respect to the BKSs of the high capacity set. There is no fine for rejected orders ($\gamma = 0$) and NPS is used everywhere.



<i>Cost sharing</i>		<i>Instance set</i>	
■	Full cost sharing (FCS)	—	Non-urgent set (large time windows)
●	Partial cost sharing (PCS)	- - -	Urgent set (small time windows)
▲	No cost sharing (NCS)		

Figure 3.13: Performance of CPS on travel costs and profit relative to all three position sharing methods (FPS, CPS, and NPS) on varying VIP for the urgent and the non-urgent instance set. A performance value of 1 indicates that CPS performs best (has lowest travel costs or highest profit) among FPS, CPS, and NPS, whereas a performance value of 0 indicates that CPS performs worst (has highest travel costs or lowest profit) among FPS, CPS, and NPS. (For a VIP of 100%, all three position sharing policies perform equally.)

70%, CPS performs better on the non-urgent set than on the urgent set. Note also that CPS performs worst (and has even higher travel costs than NPS, see Figure 3.7b) with FCS and higher VIPs for the non-urgent set, while this is not the case for the urgent set.

The effect is even larger on profits, as shown in Figure 3.13b. CPS always performs better on the urgent set than on the non-urgent set. Furthermore, CPS reaches maximal performance on profit for a VIP of 5% or 10% with PCS or NCS on the urgent set. (Indeed, CPS outperforms FPS in these cases; see Figure C.7c).

Besides the differences in the value of position sharing, the results for the urgent instances are similar to the results on the instances with lower capacity: the service level is lower than for the non-urgent set, and the FCS policy results in relatively higher travel costs and lower profits (see Figure C.7).


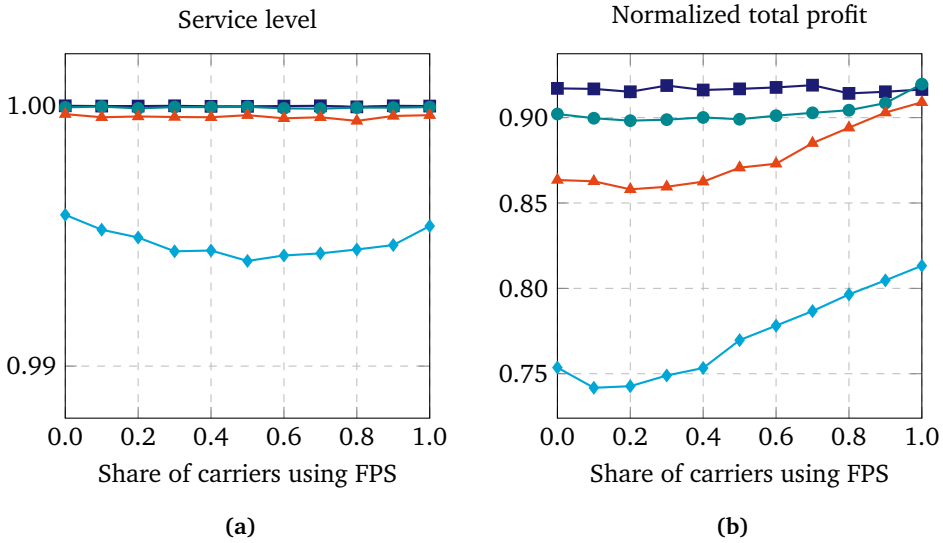
3.5.7 Mixed information sharing attitudes

In the previous experiments, we have assumed that all carriers use the same information sharing policies. In reality, however, different carriers might have different attitudes towards cooperation. Not all carriers are equally prepared to share confidential information due to competitive reasons. Larger companies, for example, might be more concerned with their privacy, while smaller firms could have larger collaboration advantages and are more open to information sharing.

To complete the experiments of this chapter, we consider a hybrid information sharing scenario: we experiment with different information sharing policies by different carriers to investigate the value of information sharing dependent on what others are willing to share. We limit ourselves to only FPS and CPS as position sharing policies, and FCS and PCS as cost sharing policies, since meaningful comparisons with NPS or with NCS cannot be made.

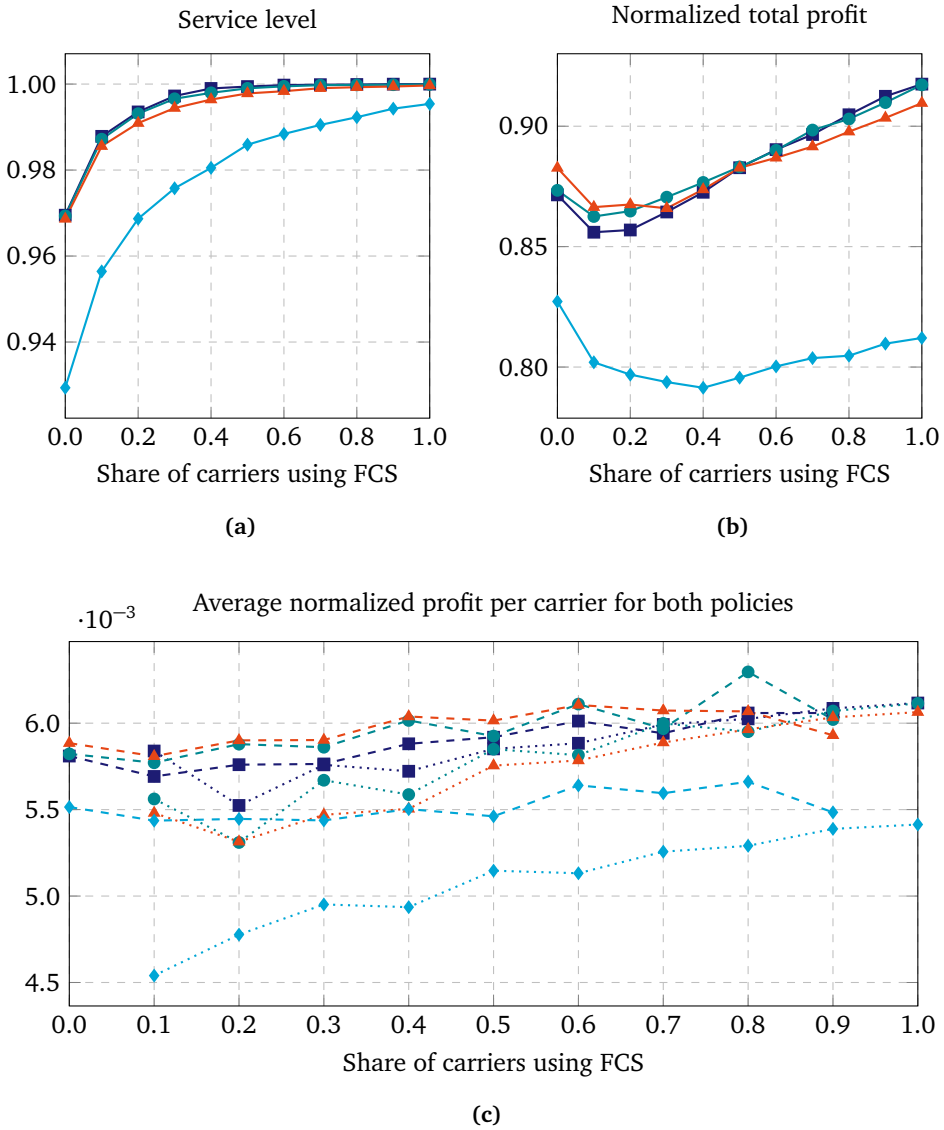
First, we vary the position sharing policies of the vehicles. In Figure 3.14, we compare results for an increasing share of the carriers using FPS. We let order agents interact with 10, 40, 70 or 100% of the available vehicles, and marginal costs are always shared by the vehicles.

Figure 3.14a shows that the service level is almost always 1, that is, hardly any orders are rejected, indifferent of the position sharing policy. In Figure 3.14b, we observe that the normalized total profit is higher for higher VIPs and generally increases with increasing share of FPS agents for lower VIPs (10% and 40%). However, a share of at least 50% of FPS agents is then needed to obtain a higher total profit than a fleet consisting of only CPS agents obtains. For higher VIPs (70% and 100%), there is no influence of position sharing policy. In Figure 3.14c, we show the average normalized profit per carrier, broken down to position sharing attitude. FPS agents obtain higher average profits than CPS agents, and this



Vehicle interaction percentage		Position sharing policy	
♦	10%	●	70%
▲	40%	■	100%
		Full plan sharing (FPS)
		---	Current position sharing (CPS)

Figure 3.14: Results for varying proportions of full plan sharing (FPS) carriers and current position sharing (CPS) carriers.



Vehicle interaction percentage		Cost sharing policy	
♦	10%	●	70%
▲	40%	■	100%
		Full cost sharing (FCS)
		---	Partial cost sharing (PCS)

Figure 3.15: Results for varying proportions of full cost sharing (FCS) carriers and partial cost sharing (PCS) carriers.

difference is larger for lower VIPs. The individual profits generally decrease a bit if more vehicle agents share full route plans.

Second, we vary the cost sharing policies of the vehicles. Figure 3.15 shows the results for an increasing share of the carriers having an FCS policy. Again, order agents interact with 10, 40, 70 or 100% of the available vehicles, and route plans are always shared by the vehicles.

Figure 3.15a shows an increasing service level when a larger share of FCS agents is present. Furthermore, there is a large gap between a VIP of 10 on the one hand, and a VIP of 40 or more, on the other hand. This gap is visible as well for the normalized total profit in Figure 3.15b. While the total profit increases with increasing share of FCS agents from 0.4 on for a VIP of 10%, it does so from 0.1 on for higher VIPs. Figure 3.15c shows the average normalized profit per carrier, broken down to cost sharing attitude. The average profit for PCS agents is in general higher than the average profit for FCS agents, but this difference becomes smaller when the share of FCS agents is larger and when higher VIPs are used. Average profits slightly increase for both FCS agents and PCS agents if the share of FCS agents increases.

3.6 Implications

We have compared different levels of information sharing that can be used in platform-based collaborative vehicle routing and have demonstrated the influence of the different information sharing policies on service levels, total travel costs, and total profits. We have also discussed interaction effects of the different information sharing policies with vehicle interaction percentages, number of auctions, and penalties on driven distance. Furthermore, we have investigated the effects of urgency and the ratio of supply and demand on the value of information. Finally, we considered scenarios with mixed information sharing attitudes. Our findings contain various insights for platform providers.

- **Impact of information sharing:** Sharing cost information and route information improves solutions in terms of all observed criteria, but the effect of cost information is generally the dominant factor. Thus, if carriers are hesitant to share route information with the central platform, platform providers may want to encourage them to share at least their cost information. Since this is only shared on a local level (between individual shippers and carriers) it may more easily be accepted. Moreover, partial information sharing leads to significant improvements in all conducted experiments and should therefore be considered when full information sharing is not accepted.

- **Route information sharing:** Sharing route information becomes more important if only low vehicle interaction percentages are possible. Significant improvements in travel costs and profits are then obtained. In this situation, it is important for a platform provider to convince carriers of the need of route information sharing. The increased profits resulting from sharing position information may be used by platform providers to create specific incentives for sharing this kind of information.

If only limited interaction is possible, platform providers may want at least 50% of the carriers to share their plans, since the total system profit (and related, the driven distance) is then equal to the scenario in which only positions are shared, and increases when more carriers share their complete plans. For carriers themselves, on the other hand, there is a strong monetary incentive to share full position information: individual profits can increase with up to 134% if future plan information is shared instead of only actual positions. In particular, if carriers know that the number of participants in auctions is limited, sharing full plans is beneficial.

With limited interaction, individual profits become lower if more carriers share their full plans. The global system profit, however, still increases, since carriers that share only their positions obtain even lower profits. Hence, although individual carriers benefit most from sharing position information if all competitors are hesitant to share it, it is better for society if all plans are shared.

- **Cost information sharing:** Sharing cost information is mainly important for obtaining high service levels (i.e., a high customer acceptance rate). In most cases, no orders are rejected if costs are shared, although no or partial cost sharing results in several order rejections.

The experiments on mixed information sharing attitudes, however, showed that it is unprofitable for individual carriers to share their costs always: they might obtain a slightly higher profit if they hide their costs sometimes.

Hence, to offer a good service to customers, a platform provider may want to set strong (external) incentives for carriers to share cost information. If the platform does not only provide contact details of carriers to shippers, but actually acts as facilitator for local communication between individual carriers and shippers, it can be designed in such a way that costs need to be provided. Alternatively, a profit sharing mechanism by the platform needs to be developed.

- **Limited cost sharing:** It is not necessary for a platform provider to motivate all carriers to share their costs. If about 60% of the carriers share

cost information and more than 40% of the fleet is involved in each auction, the service level is almost maximal. Hence, platform providers may want to motivate at least 60% of the carriers to share full cost information. Nonetheless, total profits increase if more carriers are willing to share full cost information.

Travel costs (and the related congestion and emission penalties) are significantly larger when no costs are shared at all. If carriers do not want to share their costs every time, occasional sharing already results in a notable improvement. Again, this is also serving the self-interest of carriers since profits are significantly lower if they do not share costs at all.

- **Impact of carrier interaction:** It is not necessary that shippers interact with all available carriers. An interaction percentage of approximately 40% generally already leads to a significant share of the possible gains from information sharing, especially if costs are fully or partially shared and full route plans are shared. For these situations, service levels as well as profits do not increase if the vehicle interaction percentage is increased above 40%.
- **Impact of reauctions:** Allowing reauctions (i.e., iterative interactions between shippers and carriers) is important. With more reauctions, the service levels can significantly be improved (if costs are not or not always shared), or the travel costs can be decreased (if costs are shared). It is therefore important that a platform allows multiple interaction rounds for each shipper.

By allowing multiple reauctions, a platform provider could encourage carriers to always share costs, since profits are larger in this setting than with a partial cost sharing policy. The advantage for shippers (and indirectly for the platform) is that the service levels are higher with a full cost sharing policy.

A potential problem is the communicational load that increases for larger numbers of reauctions, especially when more information is shared. If many reauctions are not possible, full cost information sharing can be induced by charging the coalition of carriers a high fine for rejected orders.

- **Impact of emission and congestion penalties:** Imposing emission or congestion penalties on driven distance is another means to stimulate cost information sharing. Strategies in which cost information is not or only partially shared, are more significantly impacted by these penalties than a full cost sharing policy.

- **Impact of problem characteristics:** When transportation problems get more difficult (due to a limited number of vehicles or due to more urgent orders), cost information sharing is still relevant for obtaining high service levels. However, total profits might end up higher if cost information is not or not always shared.

Sharing information about actual vehicle positions becomes much more relevant in scenarios with urgent orders than in situations where longer replanning times are available. Under certain circumstances, sharing full route plans is unnecessary since the availability of current positions already leads to the same profits.

- **Cooperation incentives:** In conventional centralized approaches, it is generally assumed that carriers trust on a (fair) profit sharing mechanism by the platform. In reality, however, carriers often want to be given a certain leeway in their own decision-making rather than being dominated by a central decision maker. The approach proposed in this chapter takes this aspect into account and lets carriers make their individual decisions. Moreover, individual carriers may not be convinced by average profit gains due to collaboration. The proposed MAS approach enables alternative incentives, for example, active design of the bidding process such that carriers learn through experimentation that they benefit from information sharing.

3.7 Conclusions

In this chapter, we have investigated the value of sharing different types and amounts of carrier information for solving real-world platform-based collaborative vehicle routing problems. While the need for realistic information sharing is consistently emphasized in research on collaborative transportation, dedicated information sharing policies have hardly been investigated. Thus, we focused on the question what is expedient for a collaboration system in terms of available information (Research Question 2). To analyze various information sharing policies without a central coordinator, we adapted a multi-agent system in which order agents locally organize auctions and vehicle agents, representing carriers, respond with proposals for transportation. In our approach, carriers share two types of information. First, we let them share details about the locations of their vehicles by revealing their full route plans, only their positions, or no location information at all. Second, we consider three different bidding strategies that allow carriers to hide their marginal costs for transporting the auctioned order, to share their marginal cost information only if the order is profitable, or to share their marginal costs always. Thus, the approach explicitly models the

information exchange of different actors in the considered collaborative pickup and delivery problem.

Based on computational experiments with various instances, we found that the solutions generally improve in terms of service level, travel costs, and profit when information sharing is increased. Marginal cost information and position information both contribute to solution quality, but cost information has the largest impact and is crucial for high service levels. Generally, when full cost information is available, travel costs can be lower and profits about 5% higher than when only partial cost information is available. However, with rather restricted problems (when fleet capacity is limited or orders are urgent), full cost information is not necessary to obtain maximal profits.

Position information is most important if only small groups of vehicles are involved within an auction. Profits can then significantly be improved (by 5–13%) if vehicles' positions are known, and even more (by up to 26%) if their full route plans are available. With urgent orders, however, information about current vehicle positions is sufficient – knowing full plans then does not improve the solutions.

Our results provide detailed insight into the prevailing trade-off between privacy and protection of competitive advantages, on the one hand, and service levels, total travel costs, carriers' profits, and communicational requirements, on the other hand. Sharing information plays a major role in improving routing solutions, although individual carriers may be hesitant to reveal it. Hence, depending on the specific application or desired outcomes, platform providers could stimulate sharing of certain kinds or parts of information. We showed that allowing multiple interaction rounds and imposing emission or congestion penalties to carriers can stimulate them to share more information.

Quite remarkably, an interaction with only about 40% of the carriers enables the largest part of possible improvements. When complete route plans are shared, solutions do not even improve anymore for larger interaction rates. Our experiments with mixed strategies showed that it is not necessary to have all carriers sharing complete cost information: the highest service level is already obtained if about 60% of the carriers share full cost information. Platform providers may thus not need to require full information sharing – as widely assumed in collaborative transportation research – from their customers and could possibly convince hesitant individuals to engage in limited information sharing. Moreover, the proposed multi-agent approach does not center the decision-making power and gives individual carriers a certain leeway to make decisions on their own which allows the platform providers to create respective incentives for each of these decisions. The results also demonstrate that an iterative interaction process, for instance, several auction rounds, may reduce costs by roughly 10–

30% as well as improve service levels and let carriers realize the benefits from information sharing.

In this chapter, we considered separate vehicles as autonomous decision makers. This approach can be extended by considering full information sharing and optimization within groups of vehicles belonging to the same carrier. We will consider such problems in Chapter 4 and investigate the possible gains of cooperation at a larger scale using a real-world data set.

Chapter 4

Large-Scale Collaboration

In Chapter 3, we investigated a multi-agent system for the Dynamic Collaborative Pickup and Delivery Problem where different levels and amounts of information could be shared by the participants. The decentralized nature of this method allows to solve large-scale problems. This gives us the opportunity to analyze the benefits from large-scale carrier cooperation, whereas previous collaborative vehicle routing approaches only examined what gains could be obtained by collaboration between limited numbers of carriers. The current chapter fills the gap of academic insight into possible large-scale collaboration gains (Research Question 3) by estimating the possible benefits of cooperation for large numbers of carriers based on a real-world data set of over 12000 orders.

This chapter is organized as follows. First, in Section 4.1, we introduce the problem of large-scale cooperation and enumerate the contributions of this chapter. Next, Section 4.2 discusses previous approaches for large-scale problems and explains the differences with our approach. In Section 4.3, we extend the multi-agent system that we proposed before by integrating combinatorial aspects. Then, in Section 4.4, we perform a computational study in which we investigate the benefits from cooperation for up to 1000 carriers, and analyze what the impact of bundling within the multi-agent system is. Moreover, we compare our approach with a central combinatorial auction to benchmark the quality of our approach and to determine its value for large instance sizes. Finally, we discuss the implications of the study in Section 4.5 and draw conclusions in Section 4.6.

4.1 Introduction

Horizontal collaboration is an effective approach to increase transportation efficiency (Verdonck et al., 2013; Gansterer and Hartl, 2018b; Pan et al., 2019) and has received increasing attention of governments, companies, and academia in the last years (Crujssen, 2020). While traditional collaborative vehicle routing focuses on exchange of orders between limited numbers of carriers, recent technological developments enable large-scale collaboration in real time. Different transportation platform companies already match orders with (partly) empty truck trips in practice, but there is a lack of academic insight into possible large-scale collaboration gains, optimization approaches, and participation incentives.

Centralized collaboration approaches have been studied to assess the possible gains of collaboration (Fernández et al., 2018; Molenbruch et al., 2017; Schulte et al., 2017), but these make the assumption of complete control and full information availability – which generally cannot be assumed in real-world applications due to the heterogeneity of carriers and their autonomy and privacy concerns. Decentralized approaches with a central auctioneer, and combinatorial auctions in particular (Berger and Bierwirth, 2010; Gansterer and Hartl, 2018a), overcome these problems, but available computational studies are limited to static problems with small numbers of carriers and orders. For order allocation in larger dynamic problems, MASs have been used, where orders are iteratively offered in auctions and carriers place bids for them (Máhr et al., 2010; Mes et al., 2013). Such market-based approaches are becoming more and more important: quick modifications to existing plans could be made based on real-time data, without having direct control over the cooperative (but nevertheless rational) participants. Although there is no guarantee on optimality, the scalability of the approach is a considerable advantage.

In this chapter, we again consider the auction-based MAS for solving large-scale dynamic collaborative pickup and delivery problems (see Section 2.1). Shippers can request transportation for their orders, and carriers can both source profitable jobs and outsource less profitable tasks. We examine various possible advantages and properties of this system:

- First, we investigate the possible gains of cooperation among a large number of carriers. Although MASs generally have been used for allocation of orders to vehicles, they are suitable for scenarios of mere reallocation as well. Hence, we are able to examine cooperation gains on large instances with up to 1000 carriers, while such gains have only been investigated for cooperation between a few carriers so far.
- Second, we compare the performance of the MAS (consisting of multiple small iterative auctions) with the performance of (single-round, large)

combinatorial auctions. Both approaches adopt limited information and decentralized control, but they differ in nature. The combinatorial auction theoretically gives the optimal solution if the auctioneer proposes all possible bundles and if all carriers give exact bids based on their individual optimal solutions. In practice, these conditions cannot be fulfilled, but good solutions can be found if the auctioneer offers a subset of well-selected bundles and the carriers use heuristics for generating bids (Gansterer and Hartl, 2018a; Gansterer et al., 2020b). The MAS does not give any guarantee on optimality, but since the individual auctions are relatively cheap to perform, several subsequent reactions might be used, which has a positive effect on solution quality (see Chapter 3). In this chapter, we compare both methods on instances of different size to see how they perform under different circumstances.

- Third, this chapter contains a methodological contribution: to improve solution quality, we propose the integration of combinatorial aspects within the MAS. Although single-order auctions are computationally beneficial, MASs have a limited ability to deal with interaction effects of orders. Offering bundles of orders can be necessary to avoid preventable rejections, as we illustrate with the following cases. Consider two orders that are relatively close to each other, but too far from any of the available carriers to make it profitable for them to accept an individual order. If the orders are offered sequentially, none of them will be accepted. The revenue for the two orders together, however, can for some carriers be higher than the combined transportation costs, and they might gladly accept both orders when offered in a bundle. Similarly, consider two orders that have already been assigned to different carriers that each do not have capacity for combining both orders into one route, or two orders that have been assigned to the same carrier but could be served more efficiently by another carrier. In both cases, offering the orders in a bundle could cause a reallocation, while offering the individual orders in sequence might not. Hence, we expect that offering bundles within a MAS can improve the results, while the extra effort for carriers to compute a bid on a bundle is limited if bundle sizes are kept very small.

4.2 Related work

In Section 1.2, we made the distinction between centralized and decentralized collaboration approaches. There, we argued that centralized approaches suffer from scalability issues.

A few studies succeed in solving large-scale single-carrier problems by centralized approaches: Bertsimas et al. (2019) solve dynamic large-scale problems by iteratively solving an updated mixed-integer program. They are able to do so in limited time by reducing the number of arcs in the flow graph such that only the potentially best arcs are kept. Arnold et al. (2019) use a knowledge-guided local search in which they drastically reduce neighborhood sizes to solve instances of up to 30000 orders. We are, however, not aware of collaborative centralized approaches for large-scale problems.

This chapter will develop a decentralized collaborative approach. Here, we recapitulate two types of decentralized collaboration (with central and with local auctions) since this chapter integrates aspects of both of them.

Decentralized collaboration with central auctions assumes that one central auctioneer interacts with all carriers but does not have complete information. An advantage is that the auctioneer can give some guarantees, for example, it can ensure that all orders are assigned by solving the winner determination problem. The complexity of such subproblems for the coordinator, however, restricts the size of instances that can be solved. In combinatorial auctions (Berger and Bierwirth, 2010; Gansterer and Hartl, 2018a; Gansterer et al., 2020a), each carrier submits unprofitable orders to the auctioneer. To reduce complexity, the auctioneer proposes only a limited subset of attractive bundles of these orders, and all carriers can bid on them. The auctioneer then computes the optimal assignment. Various iterative variants where bundles of orders are considered and the auctioneer finally determines a solution based on the information of different carriers have been studied by Dai et al. (2014), Lyu et al. (2019), and Wang and Kopfer (2014, 2015) (see Table 3.1). Other variants where bids are made only for single orders have been considered by Lai et al. (2017) and Li et al. (2015). Still, central auctions can only be applied to cooperative problem instances of limited size and are restricted to static problems. This hinders their applicability to the large-scale dynamic problem we focus on.

In decentralized collaboration with local auctions, no central auctioneer is considered. In contrast, any actor can act as auctioneer at any time by starting an auction on (part of) the order(s) that it is responsible for. Hence, local improvements can be made without guarantees on the feasibility of other orders and on global solution quality. Consequently, quick adjustments in dynamic large-scale problems are possible. Generally, this approach is used for allocation of orders to carriers (or even to separate vehicles of one carrier), but Dai and Chen (2011) apply it for reallocation as well (see Table 3.1). Máhr et al. (2010) and Van Lon and Holvoet (2017) consider MASs with local auctions to examine whether such a decentralized approach can outperform centralized approaches, without focusing on incentives for different carriers. Several carrier strategies

and learning mechanisms are considered by Figliozzi et al. (2004, 2005). Mes et al. (2013) investigate the interaction of several look-ahead policies for shippers and carriers, namely delaying commitments, breaking commitments, and valuation of opportunities with respect to future orders.

The present chapter investigates the interface between decentralized collaboration with central auctions and decentralized collaboration with local auctions: we compare both approaches and integrate them to benefit from their respective advantages.

4.3 Auction approaches

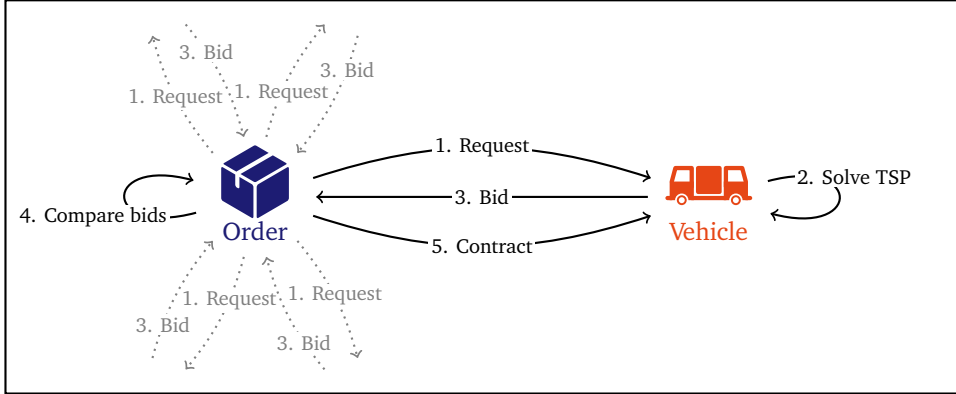
We propose a multi-agent approach where orders are iteratively offered in reverse auctions (see Figure 4.1). All available carriers (acting as sellers of service) can bid for them, and the carrier with lowest bid wins the auction: it receives the price of its bid, and becomes responsible for filling the order. In contrast to previous approaches (Máhr et al., 2010; Mes et al., 2013), we do not restrict an auctioneer to be a shipper or carrier offering a separate order: we introduce bundle auctioneers as well, that offer a group of orders $B \subseteq O$ (see Figure 4.1b). The orders within a bundle are not necessarily owned by the same shipper or carrier, since bundle auctioneers can be generated by the platform.

4.3.1 Local auction procedure

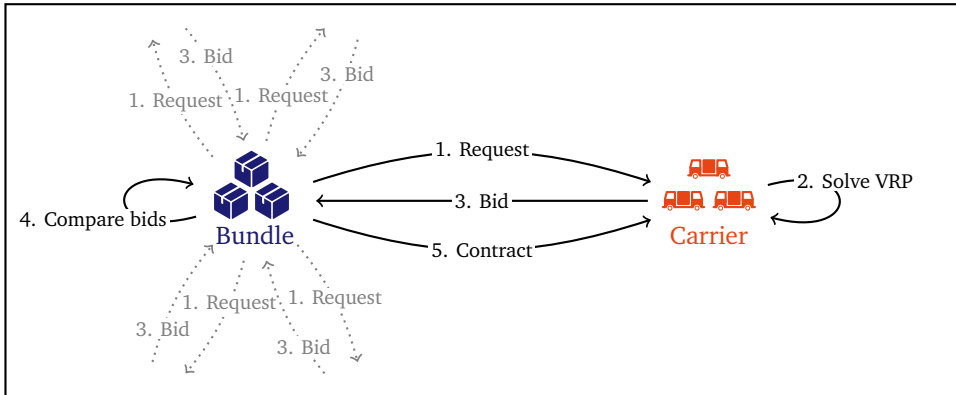
When order $o \in O$ becomes available at r_o , an auctioneer for order o (acting on behalf of shipper s if $o \in O_s$ or acting on behalf of carrier c if $o \in O_c$, but operated by the platform) is initialized and becomes active. Furthermore, the platform immediately generates, if possible, auctioneers for different bundles $B \subseteq O$, such that $o \in B$ and $|B| > 1$ (based on similarity of o and previously released orders that are known to the platform, as we will define in Section 4.3.2) and activates them shortly after the auctioneer for o has been activated.

When active, an auctioneer for a bundle B repeatedly organizes auctions. Given a maximum number of auctions m per auctioneer and its activation time r_B , the time between subsequent auctions is set to $(\min_{o \in B} l_{p_o} - r_B)/m$. The auction at time t then is as follows (see Figure 4.2; this is an extension of the auction of Section 2.2 with bundle auctions included):

1. **Requesting transportation:** The auctioneer sends a request for transporting bundle B to all known and active carriers $c \in C^t$.
2. **Computing marginal costs:** Each carrier $c \in C^t$ computes its individual marginal costs $MC_c^t(B)$ for bundle B at time t , that is, the extra travel



(a) Standard MAS approach for large-scale dynamic single-carrier VRPs.



(b) Extended MAS approach with bundling of orders for large-scale dynamic collaborative VRPs.

Figure 4.1: The standard MAS approach proposed by Máhr et al. (2010) (Figure 4.1a) has been extended in two ways (Figure 4.1b). First, it is used for assignment and exchange of orders between carriers rather than for assignment of orders to vehicles of a single carrier. Thus, instead of vehicle agents solving a TSP, carriers need to solve a VRP. Second, bundles of orders can be offered within each auction instead of single orders only.

costs for inserting all orders in B , according to their constraints, into its routes, given the situation at time t . If one or more of the orders in B have already been planned in the routes of the carrier, the marginal costs are computed as if these orders were not yet planned. If transporting bundle B is infeasible for carrier c , $MC_c^t(B)$ is set to ∞ .

3. **Bidding:** The carriers submit a bid with value $MC_c^t(B)$ to the auctioneer (i.e., they indicate that they can transport the orders if they receive at least that price).
4. **Comparing:** The auctioneer compares the received bids; let b_0 be the lowest bid provided by carrier c_0 . Furthermore, the auctioneer examines the current costs for the bundle by asking all involved carriers and shippers to report their marginal costs and reservation prices. Formally, the current costs $CC^t(B)$ for bundle B at time t are given by the sum of the marginal costs for assigned orders and the reservation prices for unassigned orders:

$$CC^t(B) = \sum_{c \in C} MC_c^t(B \cap O_c^t) + \sum_{o \in B \cap O_S^t} f_o, \quad (4.1)$$

where $O_c^t = \{o \in O \mid \exists v \in V_c \exists h \in \{1, \dots, n^{vt}\} \rho_h^{vt} = p_o\}$ is the total set of orders that carrier c has in its route plans at time t and $O_S^t = \{o \in O \mid \neg \exists v \in V \exists h \in \{1, \dots, n^{vt}\} \rho_h^{vt} = p_o\}$ is the set of unassigned orders at time t .

5. **Updating contracts:** If $b_0 < CC^t(B)$, the bid is accepted. The platform informs all involved shippers and carriers, who update their contracts and routing plans. Furthermore, the platform receives in total $CC^t(B)$ from the outsourcing shippers and carriers and pays b_0 to the winning carrier c_0 . The gain of $CC^t(B) - b_0$ is divided over the participants as incentive to cooperate, following some profit distribution function. Within this chapter, the gain is shared among the winning carrier, the (group of) currently contracted agent(s), and the platform, as defined by the following two parameters:

- **Winner gain share (WGS):** This parameter defines what fraction of the gain $CC^t(B) - b_0$ is paid by the platform to the carrier winning the auction.
- **Contracted gain share (CGS):** This parameter defines the total fraction of the gain $CC^t(B) - b_0$ that is paid by the platform to the currently contracted carrier(s) and/or shipper(s) for the orders within B . Each of them gets an equal amount.

If WGS and CGS do not add up to 1, the remaining gains are kept by the platform. If $b_0 \geq CC^t(B)$, no (re)allocations and no payments take place.

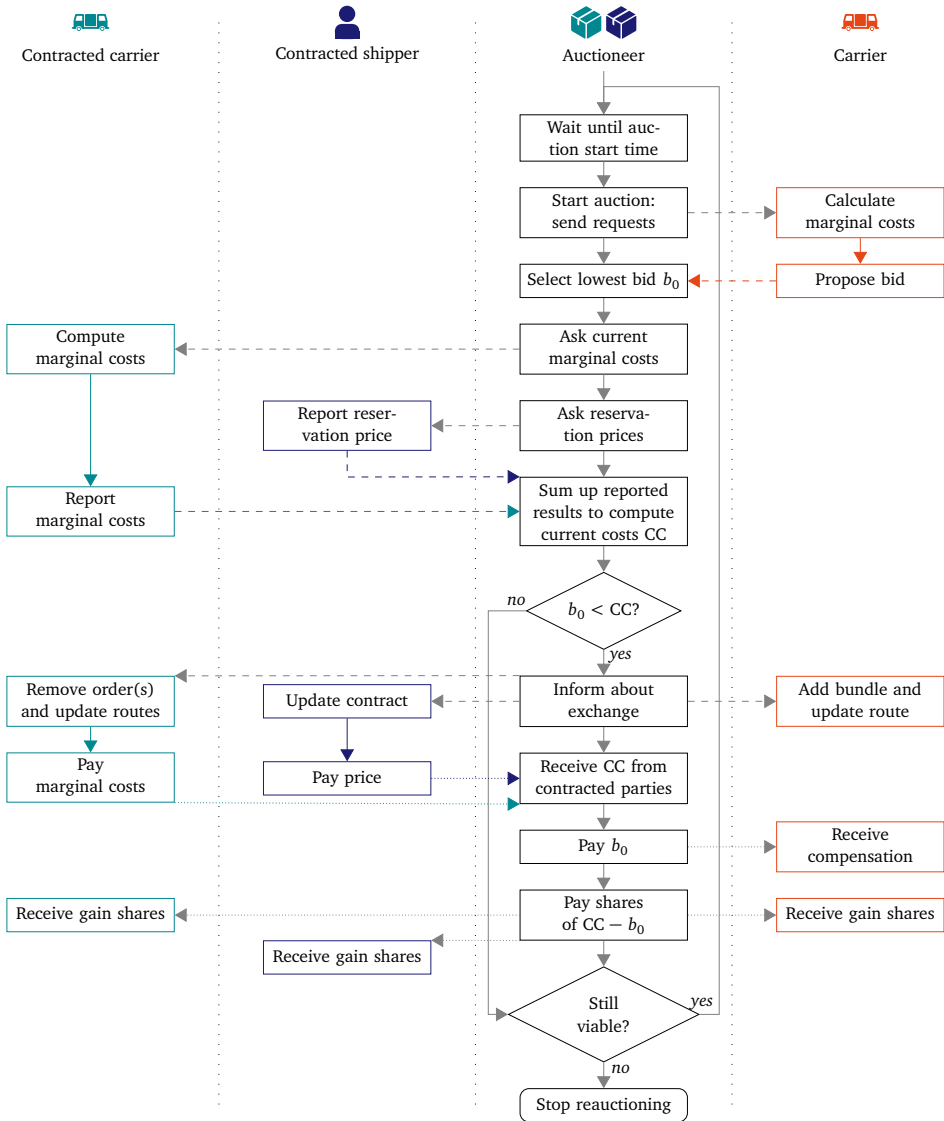


Figure 4.2: Flowchart of the iterative auction procedure within the Multi-Agent System. Dashed arrows represent exchange of information between different actors and dotted arrows represent cash flow.

When transportation of one of the orders in B starts or the latest pickup time of one of the orders has passed without a contract for that order, the auctioneer stops initializing auctions and becomes inactive.

The approach guarantees that no carrier is worse off per auction, since outsourcing carriers do not pay more than their current costs for the order(s), and the winning carrier gets at least its marginal costs for the order(s). They might, however, be worse off on the long term if they get dynamically revealed yet assigned tasks that produce bad interactions with the tasks they acquired before, or that would have had good interaction effects with the tasks that they just outsourced. Nevertheless, individual rationality is guaranteed if all assigned tasks are known by the carriers beforehand.

4.3.2 Bundling

Selling bundles of orders within a MAS is relevant if for (some of the) individual orders, the best bid is higher than the current costs, while the best bid for the bundle is below the current costs for the bundle. This is likely to happen if orders are close to each other (both in space and time) since they might be combined within the same vehicle route with lower marginal costs.

Relatedness of orders has been defined by Ropke and Pisinger (2006) for PDPs in the context of large neighborhood search (LNS). Since the goal there is to select orders from routes that can be reinserted at each other's places, both pickup locations and delivery locations need to be similar and actual visiting times are compared. For our application, it is already sufficient if one of the locations of one order is similar to one of the locations of the other order and the time windows are not too different. Gansterer and Hartl (2018a) have investigated bundle criteria based on isolation, density and tour length. Isolation, however, is not useful in our context (since we do not require partitions of the complete set of requests) and time windows are not considered in their approach. Hence, we propose a new relatedness measure and bundling procedure that can be applied in the MAS.

We define a relatedness measure $r(o, \hat{o})$ for two orders o and \hat{o} as follows:

$$r(o, \hat{o}) = \min(\text{sim}(p_o, d_{\hat{o}}), \text{sim}(d_o, p_{\hat{o}}), 0.5(\text{sim}(p_o, p_{\hat{o}}) + \text{sim}(d_o, d_{\hat{o}}))), \quad (4.2)$$

where the similarity of two pickup or delivery locations i and j is defined based on both travel time and time windows:

$$\text{sim}(i, j) = \zeta t_{ij} + w(i, j). \quad (4.3)$$

Here, ζ is a parameter (generally $\zeta > 1$) representing the cost of travel time relative to waiting time. In this chapter, we use $\zeta = 2$. Moreover, w represents the

minimal waiting time (due to time window restrictions) at one of the locations if a vehicle serves both locations immediately after each other. Formally,

$$w(i, j) = \max(0, \min(u(i, j), u(j, i))), \quad (4.4)$$

where

$$u(i, j) = \begin{cases} \infty & \text{if } e_i + s_i + t_{ij} > l_j; \\ \max(e_i + s_i + t_{ij}, e_j) - \min(l_i + s_i + t_{ij}, l_j) & \text{otherwise.} \end{cases} \quad (4.5)$$

In Equation 4.2, the minimum over three terms is taken. If the pickup of one of the orders is similar to the delivery of the other order, the orders might form a good match, irrespective of the other pickup and delivery locations and times. If, however, both pickup locations are similar, it does matter whether the delivery locations are similar. If they are at opposite directions, combining the orders might appear less useful than if they are similar as well. Hence, the third term in Equation 4.2 involves similarity of both pickup and delivery locations.

The platform dynamically generates bundles based on the relatedness measure r . Given a new order o at release time t and the pool of not yet being transported orders O^t , x bundles of size 2 and y bundles of size 3 are generated as follows:

- **Bundles of size 2:** The platform generates bundles $\{o, \hat{o}\}$ for $\hat{o} \in O^t$ and keeps the x bundles with minimal $r(o, \hat{o})$.
- **Bundles of size 3:** The platform generates bundles $\{o, \hat{o}, \check{o}\}$ for $\hat{o}, \check{o} \in O^t$ and keeps the y bundles for which $r(o, \hat{o}, \check{o})$ is minimal, where

$$r(o, \hat{o}, \check{o}) = \min(r(o, \hat{o}) + r(\hat{o}, \check{o}), r(o, \check{o}) + r(\check{o}, \hat{o}), r(o, \hat{o}) + r(o, \check{o})). \quad (4.6)$$

We have defined relatedness for three orders in such a way that not all three orders have to be highly related to each other to form an attractive bundle. Instead, each order in the bundle needs to be highly related to at least one other order in the bundle.

4.3.3 Marginal costs and route improvements

For a system dealing with dynamic reassignments, fast approximations of marginal costs are necessary. Throughout our experiments, all carriers use an elementary insertion heuristic that keeps the current sequence of orders, and inserts the new order(s) into this route at the best possible position. For bundles, the orders that can be inserted at least costs are inserted first. Thus, for a carrier

$c \in C$ approximating its marginal costs for a bundle B , there are $|B|$ main iterations in which the insertion costs for all resulting orders (at most $|B|$) at all routes ($|V_c|$ in total) are checked. Let n_c^{\max} denote the current maximum vehicle route length for carrier c . Then insertion of both the pickup and the delivery needs to be checked for each position in the route (which can be up to $n_c^{\max} + 2|B| - 2$ positions when the last order of the bundle must be inserted). Furthermore, even though we can maintain earliest and latest times along the route, a chain of time consistency updates might be necessary along the complete route in the worst case as well (Campbell and Savelsbergh, 2004). Hence, the insertion heuristic has a complexity of $\mathcal{O}(|B|^2 |V_c| (n_c^{\max} + |B|)^3)$. For single orders, this reduces to $\mathcal{O}(|V_c| (n_c^{\max})^3)$. In practice, a lot of options might be quickly pruned due to time, precedence and capacity constraints. Nevertheless, to keep computation times manageable, we limit ourselves to bundles of size 2 and 3.

To improve the quality of routes constituted by the insertion heuristic, we let carriers apply an LNS improvement phase (Pisinger and Ropke, 2019) after each insertion or deletion in one of their routes. Throughout our computational study, we use the following settings. Two destroy operators, worst removal and related removal, and four repair operators, k -regret for $k \in \{1, 2, 3, 4\}$, are used, as defined by Ropke and Pisinger (2006). Within each LNS iteration, a random neighborhood size below a given maximum is selected, and a random destroy and repair operator are applied. A simple hill-climbing acceptance criterion is used (i.e., no worse solutions are accepted).

To save computation time, we do not apply the LNS improvement phase for computation of the marginal costs, but only after a bid has been accepted or an order has been outsourced. The advantage is that bids can be submitted fast. Furthermore, carriers can improve their own routes, independent of other participants, only when it is assured that a bid is accepted or an order is outsourced. Hence, they do not need to make the computational effort for each bid, with the risk of delaying the auction too much. In real-world applications, however, carriers might apply different optimization techniques, depending on the available time and resources.

4.3.4 Reference approach: central combinatorial auction

To benchmark the quality of the solutions found by the MAS, we will compare it with the combinatorial auction as proposed by Gansterer et al. (2020b,a). In this approach, a central auctioneer creates various sets of attractive bundles on which the carriers can bid. In contrast to the MAS, only one auction round is applied after which the auctioneer reassigns tasks to carriers. The central combinatorial auction (CCA) generally consists of 5 steps (Berger and Bierwirth, 2010):

1. **Request selection:** Carriers can select part of their orders to submit for the auction, while they might keep other orders private.
2. **Bundling:** The auctioneer creates attractive bundles of the submitted orders and opens the auction.
3. **Bidding:** Carriers submit their bids, based on their marginal profits, for all bundles that they want to obtain.
4. **Winner determination:** The auctioneer solves the winner determination problem, such that the total profits are maximized and each carrier obtains at most one bundle.
5. **Profit sharing:** The obtained profits are shared among the participants.

In the first step, it is necessary to limit the number of submitted orders for reasons of complexity if the instance size increases. We follow the approach of Gansterer et al. (2020a), where orders that either have a low marginal profit for the carrier itself, or are expected to be attractive to other carriers are selected. To estimate potential attractiveness, all carriers provide aggregate information about the locations of their orders: a grid is superimposed upon the transport area, and each carrier provides the number of pickup and delivery locations that it has within each cell. Then, the total count by all other carriers for the two grid cells in which the pickup and delivery of an order o are located, indicates the attractiveness of this order to other carriers. A carrier computes for each order o a score, consisting of the rank of the attractiveness of o minus the rank of the marginal profits for o , and submits the orders with highest scores to the auctioneer.

Since proposing all possible bundles of submitted orders results in a too large computational load, the auctioneer applies a genetic algorithm to propose a smaller set of attractive bundles. Several partitions of the total request pool are generated. The appropriateness of a bundle is based on the distance to other bundles, the density of orders within the bundle, the minimum length of a tour visiting all the orders within the bundle, and the valuations of all carriers for the separate orders within the bundle. For details on this bundling process, we refer to Gansterer and Hartl (2018a) and Gansterer et al. (2020a).

Next, the carriers place bids consisting of their marginal costs for all offered bundles. As in Gansterer et al. (2020b), a variable neighborhood search metaheuristic is applied to build the routes for the carriers.

The fourth step consists of solving the winner determination problem as described by Gansterer and Hartl (2018a). The auctioneer uses an exact approach to maximize the total profits, while each carrier is assigned at most one bundle to make sure that the solution is still feasible.

Finally, the profits of the exchange of orders can be divided over the participants in several ways (Guajardo and Rönnqvist, 2016). Within the current chapter, however, we do not consider allocation of the profits to individuals, since we only focus on the total possible gains.

4.4 Computational study

For our computational study, we use a real-world data set of over 12000 orders from a Dutch transportation platform company. This company matches any submitted orders to the available load capacity of empty or partly empty trips of subscribed carriers. The data set contains locations and time windows for both pickup and delivery of each order, as well as order release times and load quantities. In contrast to the computational study of Chapter 3, we assume that carriers have multiple vehicles.

To investigate possible cooperation gains (Section 4.4.1) and the impact of bundling (Section 4.4.2), we define 6 instances of 2000 orders each, and impose different assignments of orders to various numbers of carriers. To be able to compare central and local combinatorial auctions (Section 4.4.3), we define smaller instances consisting of 50–200 orders. Additionally, we use the data set provided by Gansterer and Hartl (2016) as a benchmark.

4.4.1 Cooperation gains

For determining possible cooperation gains in large-scale problems, we generated 6 instances with the following properties. Each instance consists of 2000 orders with pickup and delivery locations (in and close to the Netherlands) and load quantities approximately as in the original data set. Original time windows have been kept, except for shifts of whole days, such that all orders fall within a time span of 10 days. Release times have been set to the start of the time span to avoid problems with initial assignments. Per instance, 1000 identical vehicles of capacity 13.6 (loading meters) are available during the complete time span, distributed over 50 randomly chosen depots (such that each depot accommodates 20 vehicles). All vehicles are assumed to have a constant speed of 72 km/h, and Euclidean distances between all locations are used. The open problem variant is used, that is, vehicles do not have to return to their depots after the last service. Travel costs are equal to traveled distances, and the reservation price for an order equals 1.5 times the distance between the pickup and delivery location.

Per instance, 5 carrier configurations (10, 50, 100, 500, or 1000 carriers) and 2 assignment configurations (close assignment or random assignment) are considered. With 10 carriers, each carrier owns 100 vehicles (i.e., precisely 5 de-

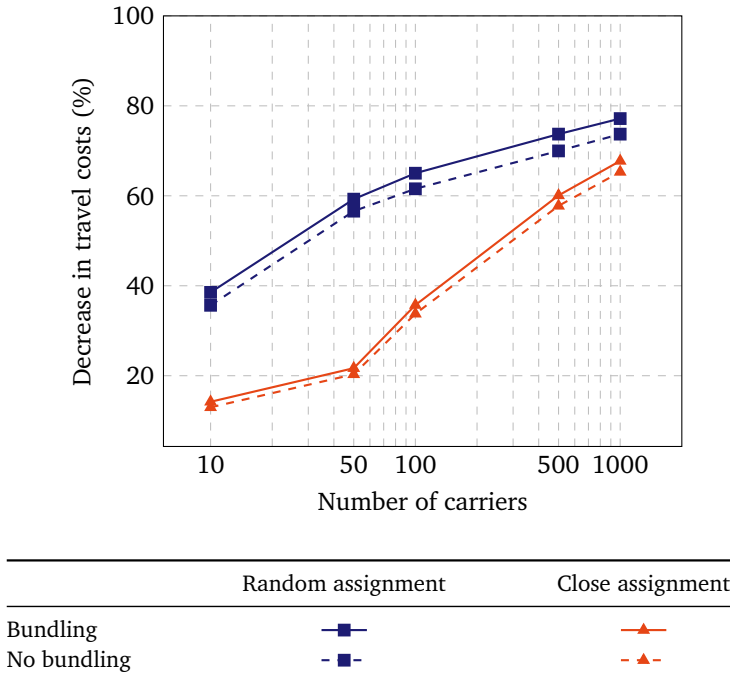


Figure 4.3: Decrease in travel costs for the cooperative scenarios (with and without bundling) compared to the non-cooperative scenario.

pot). With 50 carriers, each carrier has exactly 1 depot with 20 vehicles. With 100, 500, and 1000 carriers, each carrier owns 10 vehicles, 2 vehicles, or 1 vehicle, respectively (i.e., each depot contains the vehicles of 2, 10, or 20 carriers). Each order is assigned to a depot – the depot closest to its pickup location for the close assignment configuration and a randomly selected depot for random assignment – and then randomly to a carrier having vehicles in that depot. Hence, the theoretical optimum is dependent on the carrier and assignment configurations if cooperation is not considered, but not if cooperation is considered.


To obtain the cooperative solutions, we apply the MAS both with and without bundling three times on all instance configurations. In the runs without bundling, a maximum of 30 reactions per order is allowed. In the runs with bundling, single orders are reacted a maximum of 10 times. In addition, we select the three most promising bundles of size 2 for the order and the most promising bundle of size 3 for the order (see Section 4.3.2), and auction them a maximum of 5 times each. These parameters are selected in such a way that the total number of reactions for each order is equal with and without bundling. Note, however, that some orders might be offered more than 30 times if they appear in bundles of other orders as well. In both cases, each carrier applies a

small LNS improvement phase (100 iterations, at most 5 orders per iteration) only after an auction causes an insertion or deletion in one of its routes.


To obtain the solutions of the non-cooperative scenario, we use the following procedure for each carrier. Initially, the insertion heuristic is used to include all the tasks of the carrier into the routes of its vehicles, and afterwards an LNS improvement phase of 2500 iterations with a maximum of 100 orders per iteration is applied to improve this solution. Since we have to compute this non-cooperative solution only once for each carrier, we could use a much larger LNS improvement phase than the small LNS improvement phases that are iteratively performed after each auction in the cooperative scenario.

We show the average decrease in total travel costs for the cooperative scenarios compared to the non-cooperative scenario in Figure 4.3. As expected, cooperation gains increase with the number of participating carriers, but remarkably can be as large as 77% for 1000 carriers with random assignment. Although the non-cooperative solutions with close assignment are expected to be much better than their random assignment equivalents, cooperation can also drastically reduce the travel costs for the larger instances with close assignment: we observe savings of 68% for 1000 carriers. Note that the cooperative scenarios with bundling result in higher gains than the cooperative scenarios without bundling. We will explore this in depth in Section 4.4.2. Furthermore, note that all of the 2000 orders have been accepted in all cases, except for the non-cooperative scenarios with 1000 or 500 carriers (for 1000 carriers, 2 orders on average have been rejected with random assignment and 10 orders on average with close assignment; for 500 carriers, only 2 orders on average have been rejected with close assignment). Hence, cooperation may even improve the service level.


In Figure 4.4, we give an indication of profits for the platform and for the carrier collective as a fraction of the sum of all reservation prices (i.e., the total price the shippers have paid). Analogously to the gains in travel costs, the profits for both the carriers and the platform increase if cooperation is applied, and slightly more with bundling than without bundling. Furthermore, the profit increases are larger when more carriers participate. Note that the exact values highly depend on the WGS and CGS parameters for larger numbers of carriers, as well as on the prices that shippers pay for transportation. Under the current settings, shippers have paid 1.5 times the travel costs from pickup to delivery locations of the orders. With random assignment among 1000 carriers, this does not compensate the high travel costs if cooperation is not allowed. With low gain shares for the carriers and shippers (WGS=CGS=0.1; Figure 4.4a), carriers even make no profit after exchange of tasks (although the platform does). With higher WGS and CGS values, carriers do make profit when collaborating (Figures 4.4b and 4.4c).



(a) WGS=CGS=0.1.



(b) WGS=CGS=0.3.



	Random assignment			Close assignment		
	Non-coop	No bundling	Bundling	Non-coop	No bundling	Bundling
Platform's gain						
Carriers' gain						

Figure 4.4: Profits for the platform and the collective of carriers on instance 1 as a percentage of the system's revenue for different values of winner gain share (WGS) and contracted gain share (CGS). *Non-coop* denotes the non-cooperative scenario.




Figure 4.5: Decrease in travel costs for the bundling scenario compared to the non-bundling scenario.

4.4.2 Bundling benefits

We expected that applying small bundling within the MAS could improve solutions, which is further supported by the results from Figure 4.3. In the following, we consider both problems in which all tasks are initially assigned, as before, and problems in which part of the orders is initially unassigned (i.e., shippers connect to the platform to find a carrier).

First, we consider the same instances as in Section 4.4.1, but now we take the scenario without bundling as base case. In Figure 4.5, we show how much of the travel costs can be avoided by offering bundles. We observe that gains again increase with increasing numbers of carriers, up to 7% for 1000 carriers with close assignment and even to 13% for 1000 carriers with random assignment.

Second, we consider a more dynamic problem set in which part of the orders is not initially assigned to carriers. Again, we create 6 instances of 2000 orders each, of which only 1000 are initially assigned to carriers. We use 3 carrier configurations, namely 125, 250, or 500 carriers per instance. Each carrier has a single depot, in which it has 1–3 vehicles available. Each of the initially assigned orders is associated with a random carrier from the 10% closest carriers with respect to the pickup location. One third of the carriers have limited availability time windows, the other two third are available during the complete time span.

Original order release times have been kept, except for initially assigned orders. For these, the release times equal the corresponding carrier's release time.

We run the MAS on these instances with different numbers of carriers and various reservation price factors, both with and without bundling. The results are summarized in Table 4.1. The decrease in travel costs using bundling is generally between 0 and 1 percent, and there is a small positive influence from bundling on the service level. There is, however, no consistent pattern for increasing numbers of carriers or increasing price factors.

While bundling clearly outperforms no bundling on the instances with assigned orders, it does not on the instances where part of the orders is inassigned. To explain the difference, we again consider an instance of Section 4.4.1, but remove all initial assignments. We run the MAS both with and without bundling, and define a non-cooperative scenario as well. The latter one uses in this case only 1 auction per order to get an initial assignment, followed by an LNS improvement phase by the winning carrier. In Table 4.2, we compare the results of these experiments to the results of the instance with initial assignment. The travel costs of the non-cooperative solution for the instance without initial assignment are generally much lower than the travel costs of the non-cooperative solution for the instances with random or close assignment. Furthermore, the number of vehicles used in the solutions for the instance without initial assignment is much lower – it is actually quite close to the final number of vehicles used in the cooperative scenarios. Hence, the average route length is larger (see Figure 4.6).

This might explain the relative small difference between bundling and no bundling for the instances without initial assignment: first, the possible improvements are already lower than for instances with close or random assignment, and second, bundles of orders might be less easily accepted in longer routes, since these generally are more constrained. Note, however, that bundling still has a slight advantage on instances without initial assignment, not only in travel costs, but also in service level.

4.4.3 Comparing central and local combinatorial auctions

Now we have seen that large cooperation gains could be obtained if we apply the MAS on large-scale instances, we naturally come to the question what the quality of the MAS itself is. Since there are no optimality guarantees, both the results for the non-cooperative scenarios and the results for the cooperative scenarios might differ from the optimal solutions, leaving some space for lower or even higher possible cooperation gains. To get more grip on the quality of the MAS, we compare it with established methods, both on our own instances, and on a benchmark data set.


Table 4.1: Results for bundling on the partly assigned instance set where reservation prices are equal to the distance between pickup and delivery multiplied by a price factor.

Carriers		Price factor		
		1.25	1.5	2
125	Average decrease in travel costs (%)	0.98	0.26	0.78
	Rejected orders with bundling (avg, [min-max])	9.22[4-18]	3.44[1-7]	0.56[0-3]
	Rejected orders without bundling (avg, [min-max])	10.89[3-17]	3.83[1-7]	1.00[0-3]
250	Average decrease in travel costs (%)	0.79	0.49	0.05
	Rejected orders with bundling (avg, [min-max])	6.22[2-11]	1.50[0-4]	0.39[0-2]
	Rejected orders without bundling (avg, [min-max])	7.06[3-18]	1.61[0-4]	0.67[0-3]
500	Average decrease in travel costs (%)	0.24	0.88	0.70
	Rejected orders with bundling (avg, [min-max])	4.67[0-8]	1.22[0-4]	0.61[0-2]
	Rejected orders without bundling (avg, [min-max])	4.61[2-7]	1.67[0-4]	0.44[0-2]


Table 4.2: Average results (over 3 runs) on instance 1 in terms of travel costs, service level, and used vehicles for the three scenarios. (For instances without initial assignment, the non-cooperative scenario consists of only 1 auction per order, followed by an improvement phase by the winning carrier.)

Carriers	Assignment	ETC(%)			#RO			#V		
		NC	NB	B	NC	NB	B	NC	NB	B
10	Random	84.3	17.4	11.9	0.0	0.0	0.0	228	179	165
	Close	24.3	8.4	6.6	0.0	0.0	0.0	284	229	224
	No	23.0	6.2	5.8	33.0	6.7	3.3	132	125	126
100	Random	244.5	34.7	23.1	0.0	0.0	0.0	443	247	201
	Close	72.8	14.6	8.9	0.0	0.0	0.0	524	292	260
	No	26.9	6.3	6.2	27.7	5.7	4.7	138	125	125
1000	Random	438.9	44.5	26.6	5.0	0.0	0.0	862	341	246
	Close	247.8	23.1	15.2	16.3	0.0	0.0	591	235	202
	No	34.3	9.2	7.7	26.0	7.0	4.3	131	120	118

ETC(%): extra travel costs compared to a reference LNS solution where all vehicles belong to the same carrier; **#RO**: number of rejected orders (out of 2000); **#V**: number of vehicles used in the solution; **NC**: non-cooperative scenario; **NB**: cooperative scenario without bundling; **B**: cooperative scenario with bundling.



(a) Close assignment: 284 routes. Q_1 : 4 stops; Q_2 : 10 stops; Q_3 : 20 stops.



(b) No assignment: 131 routes. Q_1 : 10 stops; Q_2 : 22 stops; Q_3 : 42 stops.

Figure 4.6: Routes for the non-cooperative scenario on instance 1 with 10 carriers, both for close assignment and no assignment. Examples of routes for the three main quartiles of length (in terms of number of stops) are highlighted in green, purple, and orange.

Company-based instances (50–200 orders)

First, we compare the MAS with the central combinatorial auction as proposed by Gansterer et al. (2020b,a) (see Section 4.3.4) on instances of size varying from 50 to 200 orders. Larger problem sizes turn out to take too much time for the CCA, unless the number of bundles would be reduced drastically. We consider 5 or 10 carriers per instance, each having their own depot. The number of vehicles equals 10% of the number of orders, and time windows are omitted, but all orders need to be done within 24 hours. Other settings are equal to the settings of Section 4.4.1.

In Table 4.3, we show the increases in total profit by cooperation, both for the CCA and for the MAS with local combinatorial auctions, compared to a non-cooperative solution obtained by LNS. As expected, the CCA performs better on the smallest instances. The MAS, however, performs increasingly better when instance size increases. For the largest instances, the number of submitted orders and the total number of bundles generated within the CCA need already to be lowered to be able to solve the winner determination problem to optimality.

There is a notable difference between instances with random assignment and instances with close assignment. While the CCA finds comparable improvements for both assignments on the instances of size 100 and 200, the improvements for the MAS are much better on the instances with random assignment. An analysis of the profit values discloses that the cooperative solutions for random and close assignment instances are similar for the MAS, but different for the CCA. Hence, the CCA is much more dependent on the initial assignment than the MAS. Of course, this effect is dependent on the parameters used for the CCA, and in particular on the number of submitted orders per carrier. For the instance with 50 orders, where 5 carriers each submit at most 10 orders, the auctioneer has an almost complete view on the total set of orders, resulting in a larger improvement with random assignment.

In one case (200 orders, 5 carriers, close assignment), the MAS obtains a negative improvement. Although this appears counterintuitive, it is explainable since we did not use the non-cooperative LNS solution referred to in the table as starting point for the MAS; instead, we used the same fast LNS approximations as are used by carriers after an auction causes any change for them. These generally arrive at about 7% lower profits than the non-cooperative LNS solutions referred to in Table 4.3. Although the MAS compensates this in all other cases, it did not even obtain the non-cooperative solution under these specific settings. Hence, it largely depends on the parameters whether the MAS is competitive with the CCA, but in general, the MAS seems to be a reasonable alternative when the CCA suffers from scalability issues.

Table 4.3: Solution improvement due to cooperation (in terms of profit increase, relative to a non-cooperative LNS solution) for both the central combinatorial auction and the MAS with local combinatorial auctions on instances with 50–200 orders and 5–10 carriers. Each row comprises the average over 10 instances.

Instance properties			Central combinatorial auction		Local combinatorial auctions	
Orders	Carriers	Assignment	#SO	#B	Improvement (%)	Improvement (%)
50	5	Random	10	500	28.64	22.33
50	5	Close	10	500	12.62	7.53
100	5	Random	10	500	5.19	9.15
100	5	Close	10	500	5.47	2.98
200	5	Random	10	500	1.67	3.14
200	5	Close	10	500	1.76	-1.24
200	10	Random	5	100	2.65	12.80
200	10	Close	5	100	3.52	3.64

#SO: Maximum number of submitted orders per carrier; **#B:** Total number of bundles generated by the auctioneer.

Benchmark data set (30–45 orders)

Next, we apply our method on the static data set proposed by Gansterer and Hartl (2016). We benchmark against the best known solutions (BKSs) that have been found for those instances by any method, as described by Gansterer et al. (2020b,a). All instances consist of 3 carriers with depots located at 200 distance units from the others. Each carrier initially has 10 (set O_x_10) or 15 (set O_x_15) orders for which the pickup and delivery locations are in a radius of 150 (set O1_xx), 200 (set O2_xx) or 300 (set O3_xx) distance units around its depot. Thus, the area of overlap is smallest for sets O1_10 and O1_15 and largest for sets O3_10 and O3_15.

We ran the MAS under standard settings (see Section 4.4.1) on those instances (except for the fact that no maximum number of orders is specified for an LNS iteration). Furthermore, we calculated the solutions where the number of allowed auctions was increased by a factor 10. For all settings, we conducted 25 runs of the algorithm.

For each instance, the best result out of 25 runs was used to compute the improvement I with respect to the BKS, given by

$$I = (\text{PR}(\text{MAS}) - \text{PR}(\text{BKS})) / \text{PR}(\text{BKS}) \times 100\%, \quad (4.7)$$

where PR(MAS) and PR(BKS) denote the profit obtained by the MAS and the profit of the BKS, respectively. The average improvement per instance set is given in Table 4.4. For instance sets O2_10, O2_15, O3_10, and O3_15, our best

Table 4.4: Average maximum improvements in profit using the MAS with respect to the BKSs. The average results per instance set are calculated using the maximum profit value out of 25 runs of the MAS per instance.

Instance set	MNA=30		MNA=300	
	Bundling	No bundling	Bundling	No bundling
O1_10	-0.31	-0.37	-0.42	-0.43
O1_15	-0.20	-0.30	-0.28	-0.27
O2_10	2.42	1.32	3.18	1.61
O2_15	0.63	0.41	0.75	0.44
O3_10	6.25	3.35	6.69	2.23
O3_15	2.49	1.68	2.73	1.34

MNA: Maximum number of auctions per order.

Table 4.5: Average improvements in profit using the MAS with respect to the BKSs. The average results per instance set are calculated using the average profit value out of 25 runs of the MAS per instance.

Instance set	MNA=30		MNA=300	
	Bundling	No bundling	Bundling	No bundling
O1_10	-3.17	-3.27	-3.01	-3.12
O1_15	-2.51	-2.60	-2.49	-2.55
O2_10	-1.04	-1.92	-0.80	-1.92
O2_15	-2.03	-2.36	-2.05	-2.26
O3_10	-0.36	-1.82	-0.18	-1.82
O3_15	-1.21	-1.85	-1.20	-1.78

MNA: Maximum number of auctions per order.

solutions outperform the BKSs, with up to 6% on average for set O3_10. It should be noted, however, that the number of order exchanges might have been limited in the approaches to find the BKSs, while this was not the case with our MAS. For instance sets O1_10 and O1_15, our best solutions are slightly lower than the BKSs. We observe that allowing bundling generally results in better solutions, while allowing more auctions has a much lower impact. The detailed results provided in Tables D.1, D.2, and D.3 in Appendix D show that our MAS finds improvements of up to 15% on individual instances. Although we have used the best results out of 25 runs of the algorithm here, the average profits among the 25 runs are not much lower than the profits of the BKSs, as can be observed from Table 4.5. Thus, the MAS is competitive with the other approaches used to solve the benchmark data set, especially for the instances with large areas of customer overlap.

4.5 Implications

We developed a local auction system for large-scale dynamic collaborative pickup and delivery problems, and ran various experiments to investigate the advantages and disadvantages of this system. Here, we discuss the implications and limitations of the computational study.

- **Cooperation gains:** Other studies generally underestimate the possible gains of cooperation due to a very small number of cooperating carriers. We found significant improvements of about 77% for 1000 collaborating carriers, but have to note that the exact savings are highly dependent on several parameters. First, the initial assignment is of importance. While the difference in improvement between the two assignments that we examined is not very large for 1000 carriers, it is more significant for lower numbers of carriers. For 50 carriers, for example, the savings with random assignment are even about three times as high as with close assignment. Second, since we compared instances with exactly the same order set, our experiments with large numbers of carriers suffer from a low number of orders per carrier. Individual routes might hence be very inefficient. Third, our approximation of the non-cooperative solutions can be too conservative. We already have seen that it performs about 7% worse than a more extensive LNS on the small instances. For larger numbers of carriers with short individual routes, however, the fast LNS approach might give a good approximation. Thus, in real-world scenarios, the benefits of cooperation can highly depend on the number of orders and on the acceptance criteria that the different carriers have for them, as well as on their individual routing approaches. Certainly, the population of carriers in the real world is much more heterogeneous than in our experiments, which might have interesting consequences for cooperation.

We showed that the profits of both the platform and the carrier collective increase with cooperation if certain percentages of the gains per transaction are given to the carriers. This may act as an incentive to participate. If the profit increases for the carriers are too low, however, they might not consider it worth the effort to cooperate. The platform is rather powerful in its decision what amount will be given to the carriers. Even if a certain share is promised, the carriers cannot verify it. Furthermore, certain individuals can significantly contribute to a better solution without receiving a significant compensation, due to the disconnected local auctions. Incentives to participate and fair profit allocations need more study, although this might be rather difficult in large-scale dynamic settings.

- **Central auctions versus local auctions:** A comparison between a central combinatorial auction and an approach with local combinatorial auctions showed that the local approach is competitive with the central one for larger instance sizes. We need to emphasize, however, that both methods depend on various parameters settings. A notable finding is that the local approach is less dependent on initial assignment. What is preferable and feasible in a real-world scenario might highly depend on the specific problem properties, computational resources, and time available.
- **Bundling:** Although allowing bundling within a system with unconnected local auctions can improve results by up to 13%, it is again dependent on the problem properties whether it will be useful or not. The benefits seem to be much larger when all orders have been initially assigned to carriers than in open systems where shippers still are looking for a contract. In general, however, it is advisable to use bundling, since the extra computational efforts are limited. Also, individuals may simply approximate a bid value or refuse to bid on too complex bundles if that is not feasible timewise, for instance, in a highly dynamic environment where bids need to be submitted in less than a second.

4.6 Conclusions

Carrier cooperation is commonly seen as a promising approach to reduce transportation costs and emissions, but existing studies only show which gains can be obtained on relatively small instances. The current chapter investigated the potential of large-scale transportation collaboration to answer Research Question 3. Based on real-world problems, we have shown that gains of 77% can be obtained with 1000 cooperating carriers. The societal advantages in terms of emissions and traffic density are directly related. Hence, both policy makers and platform operators should provide incentives for carriers to cooperate on a larger scale.

We compared a platform-based multi-agent auction approach to a central combinatorial auction mechanism and observed that the local auction approach is competitive with the central auction if instance size increases – and even outperforms it for the largest instances. Also, the local approach is less impeded by the structure of the initial assignment. For a small-scale benchmark instance set, the local auction approach on average approximates the best known solutions, and often finds better solutions with profit improvements of up to 15%. To combine the advantages of both approaches, we integrated them by allowing auctions of small bundles of orders within the multi-agent system. Although the

extra computational effort is limited, bundling improves the results with up to 13%.

The approach within this chapter assumed fully cooperative agents. In practice, however, carriers might behave competitively. Hence, to make real-world large-scale collaboration possible, it is important to ensure that carriers cannot benefit themselves. We will investigate the possibilities of strategic behaviour within the developed system in Chapter 5.

Chapter 5

Strategic Behaviour

In Chapters 2–4, we have investigated an auction-based multi-agent system that provides good solutions for dynamic large-scale collaborative vehicle routing problems. We assumed throughout the experiments that carriers and shippers always provide their real valuations of orders. In reality, however, individual actors might try to exploit the system by behaving strategically. In this chapter, we empirically investigate to what extent strategic bidding can increase the profits of carriers, and how this can be prevented (Research Question 4). Besides a first-price auction approach, as proposed in the earlier chapters, we also consider a second-price auction.

This chapter is organized as follows. In Section 5.1, we explain the possible problem of strategic behaviour and describe our goals for this chapter. Then, in Section 5.2, we review the literature on preventing strategic behaviour in transportation markets. Section 5.3 theoretically analyzes when strategic bidding could pay off for carriers and shippers within our approach and formally describes the modeling of this behaviour, both in a first-price and in a second-price auction approach. Next, in Section 5.4, we perform a series of experiments that show under which circumstances strategic bidding pays off. The implications for transportation platform providers are discussed in Section 5.5, and Section 5.6 concludes the chapter.

5.1 Introduction

Cooperation of carriers in transportation markets can reduce the number of driven kilometers. Whereas centralized collaboration approaches have been applied for relatively small instances, multi-agent auction approaches have successfully been developed for larger, dynamic problems (Máhr et al., 2010; Mes et al., 2013). In the previous chapters, we have investigated such a decentralized auction system to analyze the value of information, possible cooperation gains, and the impact of offering bundles of orders within a local auction approach. In our experiments, we always assumed that (estimates of the) real marginal costs are reported throughout the auctions. In practice, however, carriers and shippers might bid strategically and try to increase their individual profits at the cost of the others. Nonetheless, strategic behaviour is not straightforward: Gansterer and Hartl (2018a) present a small computational example with a central combinatorial auction in which the cheating carrier always incurs a loss compared to truthful bidding. They emphasize that no general conclusions can be drawn from the example, but suggest that it might be rather difficult in practice to find a profitable cheating strategy.

It is important to know whether the MAS that we have developed within this thesis is incentive compatible, that is, can withstand strategic behaviour. If the system is robust and cannot be misused by individuals, its practical applicability value might be rather high. If, on the other hand, successful cheating is possible, we need to find ways to prevent it. Hence, the goal of this chapter is twofold. First, we want to get insight into the possible benefits of strategic bidding within the MAS developed in this thesis. Second, we look for ways to prevent this strategic behaviour.

Throughout the chapter, we will observe that carriers can sometimes successfully outplay other carriers by asking for a lower value than their true marginal costs for a bundle of orders. Although they will incur a small loss by doing so, they will often be compensated: they can either directly get a share of the auction profits, or they might be compensated later on through future events. A possible solution to the problem of direct compensation lies in applying second-price auctions, where the lowest bidding carrier is compensated with the amount of the second best bid. Under certain conditions, participants in second-price auctions do not have any incentives to deviate from their true values. Although this incentive compatibility property holds for auctions with a single indivisible good, it is not guaranteed for our scenario where we have multiple dependent auctions. Carriers can still be compensated indirectly, either by reselling the orders, or by obtaining other orders that have positive interaction effects with the orders already in their routes. Still, to find a strategic policy seems more difficult with second-price auctions than with first-price auctions because the direct

compensation is only dependent on the price of the second best bid. Hence, we extend our approach to second-price auctions and examine how they perform under strategic behaviour.

5.2 Related work

A mechanism for exchanging orders between carriers should not only be robust with respect to strategic behaviour, but requires some other qualities as well. Ideally, it has the following four properties from standard auction theory:

- **Efficiency:** The mechanism leads to a routing solution that cannot be further improved.
- **Individual rationality:** For each carrier, participating in the collaboration does not result in worse results than not participating.
- **Incentive compatibility:** Carriers do not have incentives to report other values than their true valuations.
- **Budget balance:** No extra money from outside the system is needed.

It is, however, not possible to obtain all four properties simultaneously in standard environments (Myerson and Satterthwaite, 1983). A couple of studies investigate trade-offs of these properties in static carrier collaboration situations: Xu et al. (2017) propose a bundle double auction for a problem where each carrier can exchange only one full truckload, and show that their method realizes budget balance, incentive compatibility and individual rationality, but only asymptotical efficiency. They extend the model to the exchange of multiple truckloads and propose two extended mechanisms that either are not incentive compatible for outsourcing carriers or not asymptotically efficient anymore. Gansterer et al. (2019) analyze combinatorial auctions where carriers can act as buyers and sellers at the same time. The marginal costs for insourcing an order then do not only depend on their current orders, but also on which orders they will outsource, making the problem more complex. They compare a Vickrey-Clarke-Groves mechanism and a team bidder approach: both are incentive compatible and efficient, but the properties of individual rationality and budget balance are violated. In an experimental study, they show the trade-offs of both approaches.

In addition to the possible interactions between insourced and outsourced orders, our case is even more complicated because we consider a dynamic environment where future orders might influence the value of current orders. As far as we know, only Figliozzi (2006) studies incentive compatible mechanisms

for dynamic carrier collaboration. He uses a second-price auction scheme for each newly arrived order and claims that the approach is incentive compatible, individually rational, and budget balanced, but not fully efficient. Efficiency is hindered by possible future orders (as is common for dynamic systems) and also by the fact that no reassignment is made if the current costs for the owner of the order are lower than the value of the second bid, but higher than the value of the first bid. (In that case, the owner would make a loss by paying the second price, but a better allocation could be made.)

The claim for incentive compatibility, however, can be opposed. It is argued that a carrier will not place a bid lower than its true marginal costs for transporting the order, since it will make a loss if also the second price is below its true marginal costs. This indeed holds under the assumptions that the the marginal costs “include all relevant costs (including opportunity costs) associated with servicing (or not servicing) an additional shipment or shipments” and that “all participating carriers compute these costs accurately” (Figliozzi, 2006, p. 35). However, these assumptions are too strong: it is impossible to give a certainly accurate prediction of opportunity costs in dynamic systems, simply because it is not known what orders might appear later on, and even more because it is not known whether these can be lucratively obtained or outsourced via the auction system. (Furthermore, in large-scale systems, an exact computation of the insertion costs may take too much time to be practical.) Hence, carriers may strategically bid lower values to obtain orders at a loss if they expect that advantageous interaction effects can occur later on.

In this chapter, we experimentally investigate to what extent our earlier developed first-price auction system and a second-price auction system (which is comparable to that of Figliozzi (2006)) are incentive compatible in the dynamic context that this thesis considers.

5.3 Auction approaches

We build upon the auction approach developed in Chapter 4. First, we analyze the different options of strategic behaviour and show how they can be modeled in the first-price approach. Next, we describe how the approach can be transformed into a second-price auction system.

5.3.1 First-price auctions

We consider possibilities for strategic behaviour in the first-price auction system developed in Sections 4.3.1–4.3.3. Direct strategic behaviour can take place at two points within the auction procedure of Section 4.3.1: in the bidding be-

haviour of the carriers (step 3) or in the cost reporting behaviour of carriers and shippers (step 4).

First, instead of submitting a bid with their true marginal costs as value, carriers can bid another value in step 3. For a carrier c placing a bid to acquire a bundle B , we can reason as follows (assuming that the carrier has the lowest bid), where $MC_c^t(B)$ denotes the carrier's marginal costs, b_0 denotes the carrier's bid, and g denotes the profit that the carrier makes, that is, g is a fraction of $CC^t(B) - b_0$, dependent on the used profit distribution function (see step 5 of the procedure in Section 4.3.1).

- The carrier will not bid a value $b_0 < MC_c^t(B)$ if g is expected to be relatively small, since the compensation $b_0 + g$ will not cover the marginal costs $MC_c^t(B)$.
- The carrier might place a bid $b_0 < MC_c^t(B)$ if g is expected to be relatively high. If $b_0 + g > MC_c^t(B)$, lowering the bid is a good strategy to outbid another carrier with a bid between b_0 and $MC_c^t(B)$.
- The carrier might speculate on getting a high gain from reselling the bundle later on, or foresee good interaction effects with orders that will appear later on, and hence might place a bid $b_0 < MC_c^t(B)$.
- The carrier might bid a value $b_0 > MC_c^t(B)$ to get a higher compensation, but this comes at the risk of not winning the auction anymore.

Formally, the bid value for a strategic carrier c bidding for bundle B at time t will be

$$\sigma_c MC_c^t(B), \quad (5.1)$$

where σ_c represents the degree of strategic bidding for carrier c .

Second, instead of reporting the current costs of a bundle that they own in step 4, carriers can perturb this value. Similarly, shippers can misreport their reservation prices. For such carriers or shippers mentioning their marginal costs or reservation prices for outsourcing orders, we make the following assumptions.

- Carriers and shippers do not report a value above their true value, since they need to pay this value.
- Carriers and shippers might report a lower value, but this comes with the risk that the lowest bid b_0 is not lower than $CC^t(B)$, hindering the trade. Indeed, they might report lower values and slightly increase them in next auction rounds, but due to the dynamic environment, there is no guarantee on success.

Formally, the perturbed current costs in step 4 of the auction in Section 4.3.1 are given by

$$CC^t(B) = \sum_{c \in C} \lambda_c MC_c^t(B \cap O_c^t) + \sum_{s \in S} \sum_{o \in B \cap O_s^t} \lambda_s f_o, \quad (5.2)$$

where $O_s^t = \{o \in O_s \mid \neg \exists v \in V \exists h \in \{1, \dots, n^{vt}\} \rho_h^{vt} = p_o\}$ is the set of unassigned orders of shipper s at time t , and λ_c and λ_s represent the degree of false reporting for carrier c and shipper s , respectively.

5.3.2 Second-price auctions

We extend the auction approach of Section 4.3.1 to a second-price auction. Instead of getting the value b_0 (and possibly an extra gain g , dependent on the profit distribution function), the winning carrier gets the amount of the second-lowest bid b_1 . In second-price auctions with a single, indivisible item, the carriers do not have any incentive to deviate from their true valuation: the winning carrier either would have won the auction anyhow (if its true value is below the second price), or makes a loss (if its true value is above the second price). This property does not hold anymore in our case, since our environment is dynamic and the value of a bundle depends on later events as well, but we still want to test whether second-price auctions perform better than first-price auctions in terms of preventing strategic behaviour.

The use of second price auctions raises a new problem. The amount of the second price needs to be paid by someone. In a budget balanced setting, still the shippers or already contracted carriers must pay this price. The second price, however, is more likely to be higher than their current costs than the first price is. This might result in less (re)allocations, and hence a worse final solution than with first-price auctions. A solution could lie in the bundling approach proposed in Chapter 4. If bundles of orders from different owners are considered, the interaction advantages of the orders might result in lower bids of the carriers, while the separate current costs are not influenced by interaction effects. Hence, paying the second price could be less problematic if the platform generates bundles of orders from different owners.

Still, the risk that an auction does not succeed due to false current costs is higher than with first-price auctions. To prevent the current owners of the orders from reporting too low current costs, we let the auctioneer ask them a certain amount such that the second price can be paid to the winning carrier. The current owners only have to accept or refuse the proposed price from the auctioneer.

We describe the complete second-price auction procedure below. As before, we have bundle auctioneers that repeatedly organize auctions, until transportation of one of the orders in the bundle starts or the latest pickup time of one of

the orders has passed without a contract. The second-price auction procedure followed by an auctioneer for a bundle of orders B at time t is as follows:

1. **Requesting transportation:** The auctioneer sends a request for transporting bundle B to all known and active carriers $c \in C^t$.
2. **Computing marginal costs:** Each carrier $c \in C^t$ computes its marginal costs $MC_c^t(B)$ for bundle B at time t , that is, the extra travel costs for inserting all orders in B , according to their constraints, into its routes, given the situation at time t . If transporting B is infeasible for c , $MC_c^t(B)$ is set to ∞ .
3. **Bidding:** The carriers submit a bid with value $\sigma_c MC_c^t(B)$ to the auctioneer, (i.e., they indicate that they can transport the orders if they receive at least that price), where σ_c represents the degree of strategic bidding.
4. **Comparing:** The auctioneer compares the received bids; let b_0 be the lowest bid provided by carrier c_0 and let b_1 the second lowest bid.
5. **Proposing prices:** The auctioneer needs to pay b_1 to c_0 for a (re)allocation, and hence must make sure to get at least b_1 from the current owner(s) of the orders in B . If a lower amount is gathered, the auctioneer will make a loss, and has no incentive to make a (re)allocation. All amounts above b_1 can be kept as profit for the auctioneer. Thus, the auctioneer proposes a price a_c for all carriers $c \in C_B^t$, and a price a_s for all shippers $s \in S_B^t$ such that $\sum_{c \in C_B^t} a_c + \sum_{s \in S_B^t} a_s \geq b_1$, where C_B^t represents the set of all carriers contracted at time t for at least one order in B and S_B^t represents the set of shippers having an order in B that is yet unassigned at time t . Prices could be determined in different ways. Here, we use a straightforward approach that divides b_1 proportionally to the distance between pickup and delivery locations of the orders in B , and adds a small profit factor to it, defined as follows:
 - **Platform gain share (PGS):** This parameter defines what fraction of the second bid b_1 is additionally requested from the current owners of the orders as a gain for the platform.

The requested prices are then given by

$$a_c = \frac{\sum_{o \in B \cap O_c^t} t_{p_o} d_o}{\sum_{o \in B} t_{p_o} d_o} (1 + \text{PGS}) b_1 \quad \forall c \in C_B^t \quad (5.3)$$

and

$$a_s = \frac{\sum_{o \in B \cap O_s^t} t_{p_o d_o}}{\sum_{o \in B} t_{p_o d_o}} (1 + \text{PGS}) b_1 \quad \forall s \in S_B^t \quad (5.4)$$

such that $\sum_{c \in C_B^t} a_c + \sum_{s \in S_B^t} a_s = (1 + \text{PGS}) b_1$.

When the auctioneer has proposed the prices, the current owners of the orders can check whether the requested prices are less than or equal to their current costs or reservation prices. If so, they will accept the proposed prices.


6. **Updating contracts:** If all current owners accept the proposed prices, the bid is accepted. The platform informs all involved shippers and carriers, who update their contracts and routing plans. The auctioneer receives the payments from the outsourcing shippers and carriers as proposed (i.e., $(1 + \text{PGS}) b_1$ in total), and pays b_1 to the winning carrier c_0 . The remaining gain of $\text{PGS} \cdot b_1$ is kept by the auctioneer. If one of the current owners does not accept to outsource its orders at the proposed price, no (re)allocations and no payments take place.

The approach guarantees that the second price is paid to the winning carrier if a (re)allocation takes place, and that the auctioneer does not incur a loss (if $\text{PGS} \geq 0$). The drawback, however, is that current owners need to accept the prices that are proposed by the auctioneer to have a successful (re)allocation. This becomes less likely with higher PGS values.

5.4 Computational study

In this section, we empirically test what the influence of strategic bidding and strategic value reporting is within the MAS approach. First, we investigate the possible advantages of strategic behaviour in the first-price auction system as proposed in earlier chapters. Second, we consider strategic bidding in a second-price system.

Throughout the computational study, we again use the real-world data set from a Dutch transportation platform company introduced in Chapter 4, and generate instances of 2000 orders each. To prevent any bias from unprofitable initial contracts, we use problem instances without initial assignment. Per instance, 250 carriers with 1–3 vehicles each are considered, of which one third have restricted availability time windows. Further instance characteristics are as described in Section 4.4.1.



	Truthfully bidding carriers	Strategically bidding carriers
10% strategically bidding carriers	—▲—	—■—
20% strategically bidding carriers	- -▲- -	- -■- -
30% strategically bidding carriers	...▲...	...■...

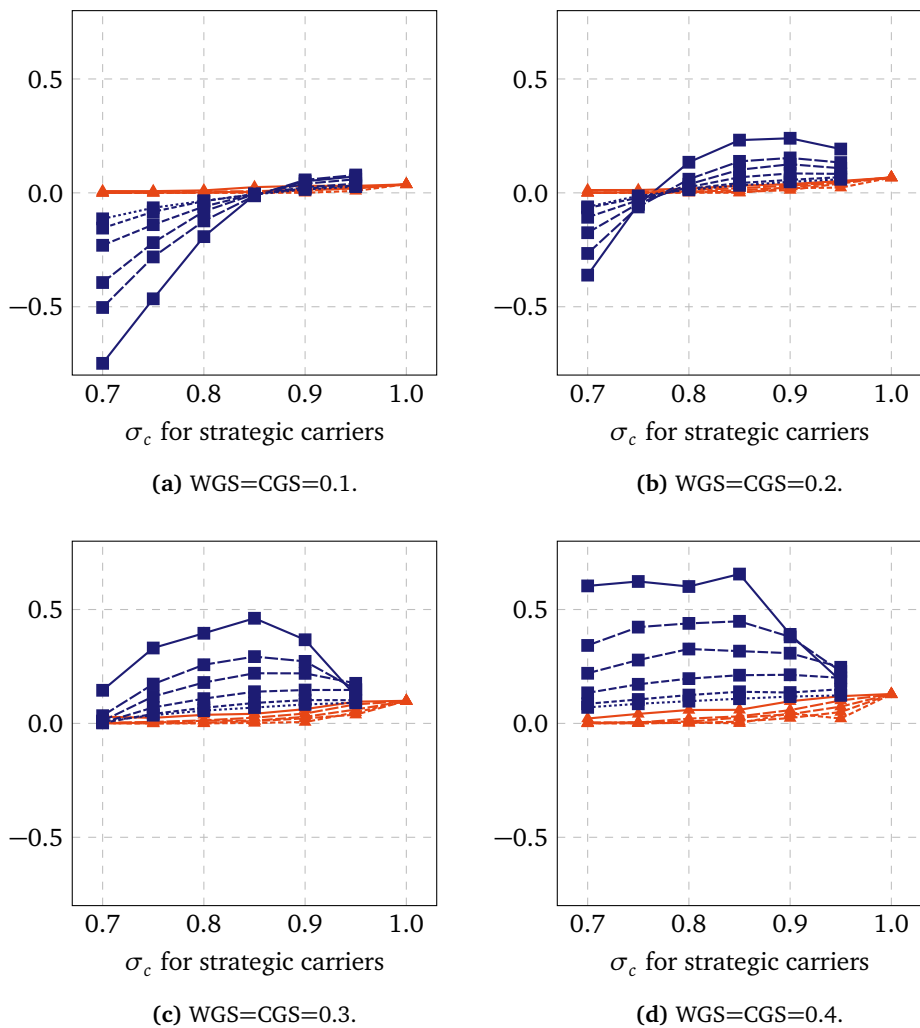
Figure 5.1: Average carrier profits if part of the carriers bid a fraction of their real (estimated) insertion costs in the first-price auction system.

5.4.1 Strategic behaviour in a first-price auction system

First, we analyze whether carriers can benefit from placing false bids in a first-price auction system. We run the MAS with different percentages of carriers (10%, 20%, or 30%) placing strategic bids ($\sigma_c \in \{0.8, 0.85, 0.9, 0.95, 1.05\}$). We test three configurations for winner gain share and contracted gain share (WGS = 0.1, CGS = 0; WGS = CGS = 0.1; and WGS = 0, CGS = 0.1). In the last configuration, winning carriers do not take any of the profit generated by a successful auction (they even lose some profit if their bid is lower than their real costs), but they might obtain a gain if they resell the order later on.

In Figure 5.1, we give the average profit per carrier, both for the carriers that bid truthfully and for the carriers that bid strategically. (As a reference, we also show the average profit if all carriers bid their true values at $\sigma_c = 1$.) We observe that strategic bidding pays off for $\sigma_c = 0.9$ or $\sigma_c = 0.95$ if WGS = 0.1, but not for other values of σ_c . The truthful carriers are worse off if the strategic carriers bid lower than their true prices, even if the strategic carriers themselves also do not gain any extra profit. With WGS = 0 and CGS = 0.1, there is no incentive

Average profit per carrier (% of total reservation prices)



	Truthfully bidding carriers	Strategically bidding carriers
10% strategically bidding carriers	—▲—	—■—
20% strategically bidding carriers	-▲-	-■-
30% strategically bidding carriers	-▲-	-■-
50% strategically bidding carriers	-▲-	-■-
80% strategically bidding carriers	-▲-	-■-
100% strategically bidding carriers	-▲-	-■-

Figure 5.2: Average carrier profits if part of the carriers bid a fraction of their real (estimated) insertion costs in the first-price auction system, for increasing values of WGS and CGS on instance 1.

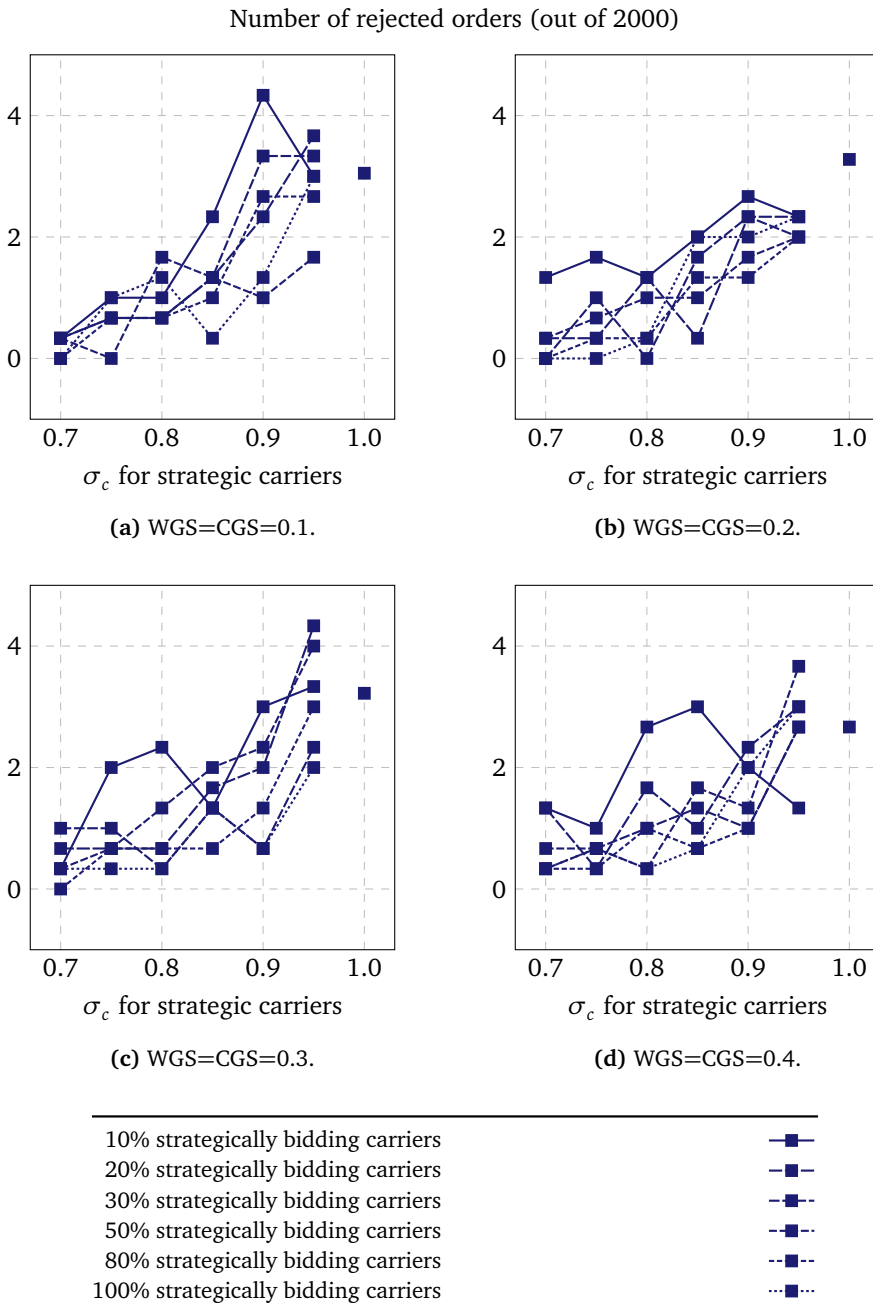


Figure 5.3: Average number of rejected orders if part of the carriers bid a fraction of their real (estimated) insertion costs in the first-price auction system, for different values of WGS and CGS.

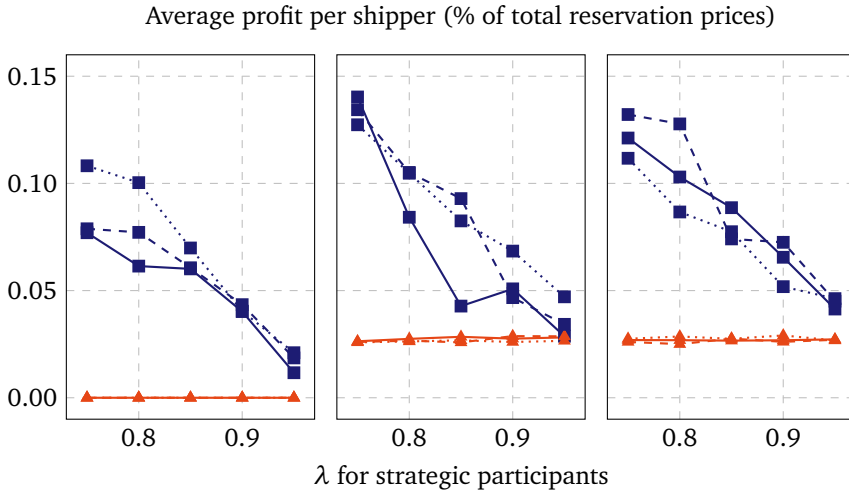
to bid another value than the true value. Note that the highest profits can be obtained if only low numbers of carriers bid strategically. Similarly, the losses that strategic carriers can obtain will be largest with low numbers of strategic carriers, that is, if most other carriers just report their true costs. These losses can be already very large with slightly lower values for σ_c . Hence, finding a beneficial strategic bidding value could be a critical process.

With higher values for WGS, lower values of σ_c are expected to be beneficial: if the system assigns large shares of the gains to the winning carriers, cheating might appear too easy. We tested this hypothesis on a single instance, even for larger shares of strategic carriers, and show the results in Figure 5.2. Indeed, for $WGS = CGS = 0.2$, the turning point below which strategic bidding does not pay off decreases to $\sigma_c = 0.8$. For $WGS = CGS = 0.3$ and $WGS = CGS = 0.4$, even the lowest tested value of $\sigma_c = 0.7$ is still beneficial for strategic carriers. Thus, the higher the value of WGS, the easier it is for carriers to find a beneficial strategic bidding policy. From Figure 5.3, we observe that the number of rejected orders is rather low in general, and is lowest for lower values of σ_c .


Next, we analyze how much shippers and carriers can benefit from communicating false (lower) reservation prices or current costs. In Figure 5.4, we show average obtained profits per shipper and per carrier when 10–30% of the participants communicates current costs or reservation prices equaling 75–95% of their true values. Strategic shippers can obtain considerably higher profits if they lower their communicated reservation prices. This can be explained by the large difference between reservation price and insertion costs for a carrier that already had planned a route in which the order fits quite well. The shipper then might easily outsource its order at a low price. Likewise, carriers can obtain extra profits by outsourcing orders for a lower price than their actual costs, but the differences are smaller. The drawback of using lower reservation prices, however, is that less orders will be served, as can be observed from Figure 5.5.

5.4.2 Strategic behaviour in a second-price auction system

We now investigate whether a second-price auction system can reduce the motivation to bid strategically in a dynamic world. To remove any interference between the bid value that a carrier c submits for a bundle B and the expected price a_c that the platform proposes to this carrier if c already owns any order $o \in B$, we restrict our experiments in such a way that bids are only made for bundles that do not contain any currently owned orders. We run the MAS with different percentages of carriers (10%, 20%, 30%, 50%, 80%, or 100%) that place strategic bids ($\sigma_c \in \{0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$), and test four different values for platform gain share ($PGS \in \{0, 0.01, 0.1, 1\}$). For $PGS = 0$, the auctioneer asks in each auction round exactly b_1 in total from the current own-



(a) WGS=0.1, CGS=0. (b) WGS=0.1, CGS=0.1. (c) WGS=0, CGS=0.1.



(d) WGS=0.1, CGS=0. (e) WGS=0.1, CGS=0.1. (f) WGS=0, CGS=0.1.

	Truthful participants	Strategic participants
10% strategically reporting shippers and carriers	—▲—	—■—
20% strategically reporting shippers and carriers	- -▲- -	- -■- -
30% strategically reporting shippers and carriers	...▲...	...■...

Figure 5.4: Average shipper profits (Figures 5.4a–5.4c) and carrier profits (Figures 5.4d–5.4f) if part of the shippers and carriers mention lower reservation prices and lower current costs than their true ones in the first-price auction system.




Figure 5.5: Average number of rejected orders if part of the shippers and carriers mention lower reservation prices and lower current costs than their true ones in the first-price auction system.

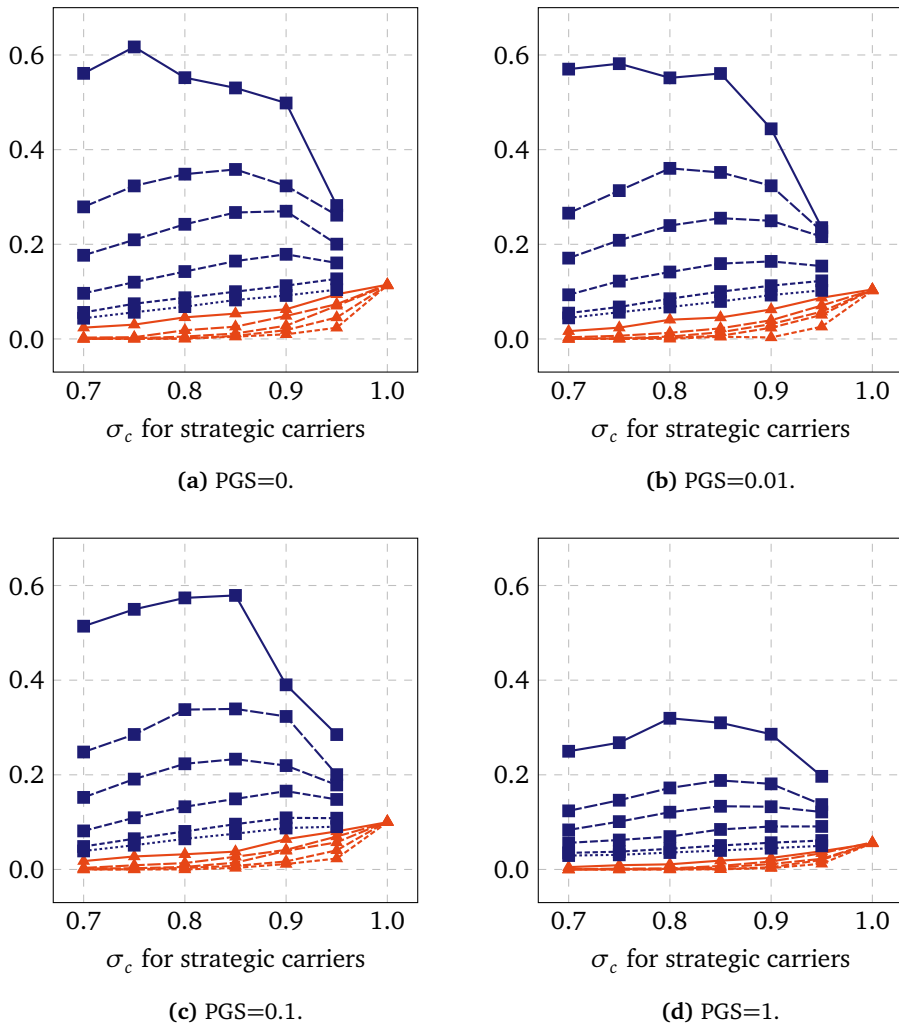
ers, and hence, makes no profit itself. In the extreme case of $PGS = 1$, on the other hand, the auctioneer asks $2b_1$ from the current owners, and tries to make a profit of b_1 itself each auction.

The average profits (as a percentage of the sum of the reservation prices for the transported orders) for the carriers that bid strategically and for the carriers that bid truthfully are given in Figure 5.6. (Again, as a reference, the average profit with only truthful carriers is given at $\sigma_c = 1$.)

Strikingly, strategic bidding always results in higher profits than true bidding, regardless of the number of strategic carriers, the value of σ_c , or the value of PGS . The profits of strategic carriers, however, highly depend on the total number of strategic carriers within the system. If 80% or 100% of the carriers act strategically, their profits are easily becoming lower than the profits in a scenario with only truthful carriers, leading to a kind of prisoner’s dilemma: irrespective of what the others do, strategic bidding results in higher individual profits than truthful bidding, but carriers are better off when they all bid truthfully than when they all bid strategically.

In Figure 5.7, we show the corresponding average numbers of rejected orders. With more strategically bidding carriers and lower σ_c values, the number

Average profit per carrier (% of total reservation prices)



	Truthfully bidding carriers	Strategically bidding carriers
10% strategically bidding carriers	—▲—	—■—
20% strategically bidding carriers	-▲-	-■-
30% strategically bidding carriers	...▲...	...■...
50% strategically bidding carriers	-▲-	-■-
80% strategically bidding carriers	-▲-	-■-
100% strategically bidding carriers	—▲—	—■—

Figure 5.6: Average carrier profits if part of the carriers bid a fraction of their real (estimated) insertion costs in the second-price auction system, for different values of PGS.

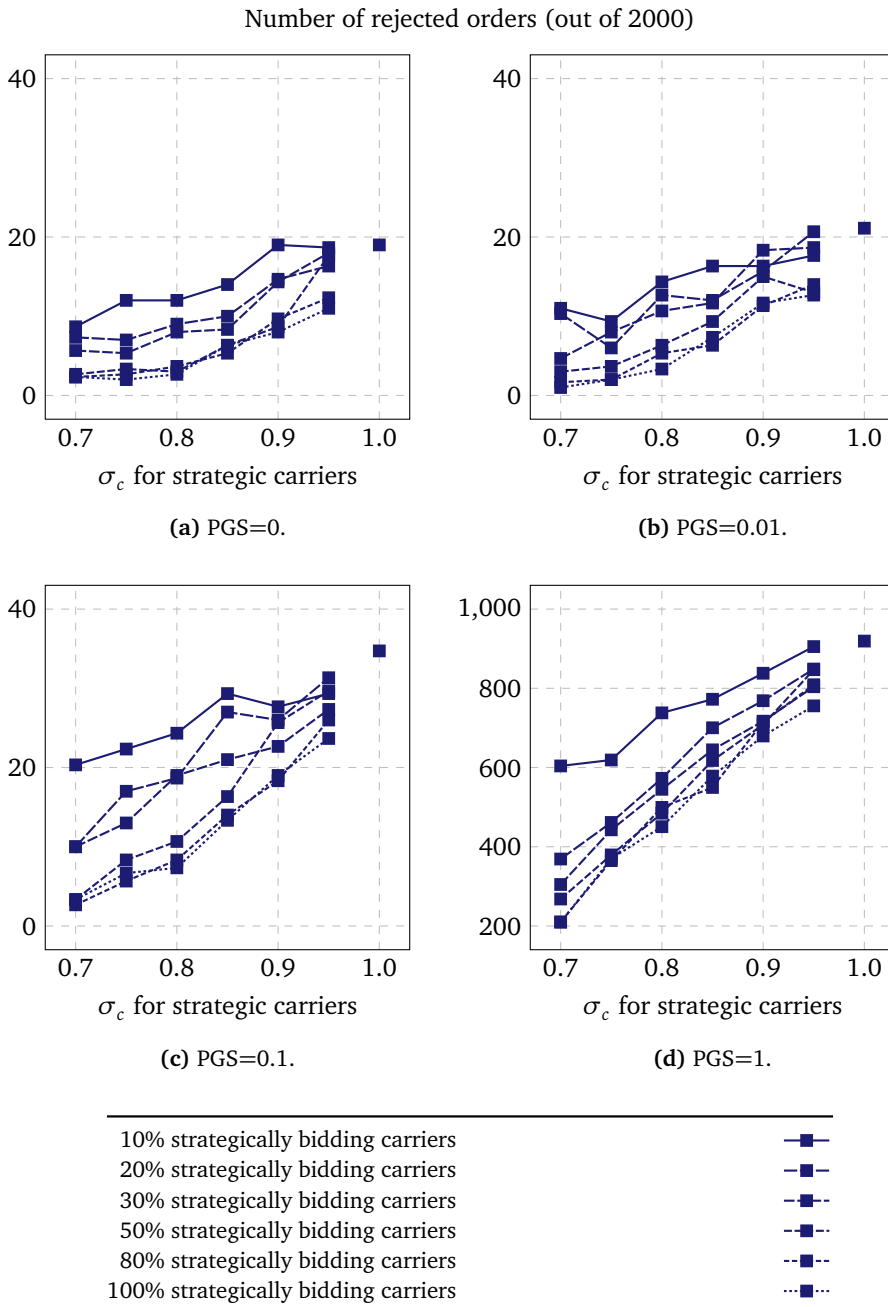


Figure 5.7: Average number of rejected orders if part of the carriers bid a fraction of their real (estimated) insertion costs in the second-price auction system, for different values of PGS. (Note that the scale of the y-axis is different for Figure 5.7d.)

of rejected orders decreases, as expected, since the value of b_1 is likely to get lower. For increasing values of PGS, the number of rejected orders increases, with almost half of the orders rejected under some conditions for $\text{PGS} = 1$. This can be explained by the (too) high prices that the auctioneer asks from the current owners of the orders. A comparison with the number of rejected orders within the first-price auction system (see Figure 5.3) reveals that the number of rejected orders is generally lower in the first-price auction system than in the second-price auction system. Furthermore, the rejection rate is not dependent on the WGS and CGS parameters in the first-price auction system, while it is heavily dependent on the PGS parameter in the second-price auction system.

5.5 Implications

We analyzed the potential of strategic behaviour within a local auction system for large-scale dynamic collaborative pickup and delivery problems. We investigated both a first-price auction system and a second-price auction system. The following implications for platform providers can be derived.

- **Strategic bidding in a first-price auction:** We have shown that it is not straightforward to bid strategically in a first-price auction. At the same time, it is also not impossible. Whether strategic bidding can pay off highly depends on the gains per auction that are attributed to the winning carrier:
 - If this share is relatively low, carriers can benefit if they slightly lower their bids. If they make their bids too low, however, they easily will make a loss. The exact value of the turning point will not be clear beforehand, making it difficult for carriers to cheat. The drawback of a system with a low gain share for carriers is that carriers have little incentive to participate in the system.
 - If the gain share for a winning carrier is relatively high, carriers might be interested in participating in the cooperation system. The problem is that it then will be easy for them to cheat the system: they can bid lower prices to get more orders and will be compensated for their too low bids. At the same time, the total routing solution will become worse, since the orders will often not be assigned to the (truthful) carriers that can perform them at least costs.

It might thus be possible for a platform provider to use a first-price auction system, but the procedure details must be selected carefully to prevent strategic bidding. With too high gain shares for carriers, strategic bidding

easily pays off for them, while they might have no incentive to participate with too low gain shares.

- **Strategic value reporting in a first-price auction:** Reporting of reservation prices or current costs of an order in a route is prone to strategic behaviour in a first-price auction system. The owners of the orders (shippers or carriers) have incentives to lower their prices, since their profit then increases. Shippers generally will get larger profits with lower reservation prices. For carriers, the amount by which they lower their prices is less important.

The number of rejected orders within the system, however, will increase with lower reported prices. Hence, for individual shippers there is a risk that their orders will not be accepted anymore if they lower their prices too much. For a platform provider, it is important to know what risks shippers are willing to take to increase their profits.

- **Strategic bidding in a second-price auction:** Within a second-price auction system, it pays off for carriers to report false bid values. Their profits will increase at the cost of the carriers that bid truthfully. However, a prisoner's dilemma appears: if a large share of the carriers act strategically, the average profit is lower than the average profit when all carriers act truthfully. Since also the number of rejected orders will be larger than in the first-price auction system, the currently investigated second-price auction system will not be relevant for a platform provider.

5.6 Conclusions

To verify whether the multi-agent auction approach proposed in this thesis will be feasible in practice, we analyzed whether it is possible for individual participants to benefit from strategic behaviour (Research Question 4). Asking lower prices than the real costs for serving an order turned out to be advantageous for carriers under certain circumstances and can yield profits of up to 5 times the profits in a truthful setting. It is, however, highly dependent on gain shares (and hence not evident) by what amount bids can be lowered without becoming disadvantageous. On the other side, shippers or carriers outsourcing orders have an incentive to report lower marginal costs or reservation prices than their true ones. The drawback, however, is that a larger number of orders will not be (re)assigned. This hinders the improvement of the total system, but also might harm individual shippers if their orders will not be accepted.

Motivated by the strategy-proofness of second-price auctions for single indivisible items, we experimentally tested whether second-price auctions could be

applied successfully in our multi-item dynamic context. The hypotheses were that a second-price auction could reduce the profit of strategic bidding and that auctioning bundles of orders might solve the budget balance problem as well in this case. It turned out, however, that strategic bidding always pays off for carriers within the proposed system: average profits for strategic carriers are higher than the related profits for truthful carriers. Apparently, the long-term advantages of having a larger set of orders outweigh the lower compensations when acquiring them. Carriers will only have a disadvantage if too many carriers cheat: when 80–100% of the carriers bid strategically, the average profits are lower than in a completely truthful setting.

Whereas this chapter compared the influence of true and false bids from a carrier, we will consider different true bids from the same carrier for the same order in Chapter 6: we will analyze a setting with multiple pickup and delivery alternatives per order, for which carriers can indicate their availability.

Chapter 6

User Preferences

In Chapters 2–5, we considered transportation problems with predetermined pickup and delivery locations and times per order. If these are determined by an operator, it can be rather inconvenient for customers and might lead to a high number of missed deliveries. If users, on the other hand, are allowed to specify the pickup or delivery details, a very inefficient routing is expected. In this chapter, we propose a solution for these problems by introducing the concept of multiple time-location alternatives. Through this, carriers are given more routing flexibility and the service for customers can be improved in delivery processes. Furthermore, we introduce preference indications for each option, and aim for finding solutions that balance minimizing total travel costs and customer or operator dissatisfaction. To answer Research Question 5, we show how a multi-agent approach can assist in reaching this goal.

This chapter is organized as follows. We introduce the concept of multiple pickup and delivery alternatives in Section 6.1. In Section 6.2, we discuss similar approaches and argue where they need to be extended. Then, in Section 6.3, we formally define the Generalized Pickup and Delivery Problem with Preferences. Subsequently, in Section 6.4, we extend the MAS developed in earlier chapters to solve the problem in a decentralized way. A computational study in Section 6.5 compares this approach with a centralized approach. In Section 6.6, we discuss the implications, after which the conclusions follow in Section 6.7.

6.1 Introduction

In home delivery processes, a frequently encountered problem is the absence of the customer at the moment a parcel is delivered. A solution could be to serve each customer in a customer-defined time window, but this can highly impair the efficiency of vehicle routes, and hence result in high delivery costs and a negative environmental impact.

In this chapter, we propose another solution, by allowing customers or operators to specify multiple time-location combinations for delivery. Next to the (still preferred) option of home delivery between 10:00 and 12:00, for example, a (less appreciated but still acceptable) delivery at a locker box station two streets away might be allowed, with the advantage of a larger time window. One might also think of delivery at one's work location, or the specification of several time slots during a week for home delivery, with the earlier time slots as preferred options. Similarly, multiple pickup alternatives can be considered, for instance, in the case of distribution from companies with different production or storage locations. Another scenario might be a mobility service where the customer is willing to conduct part of the trip by foot or other modality.

For such scenarios with multiple alternatives, a higher delivery success rate can be achieved, and the transport operator has more flexibility in designing efficient routes. From the alternative pickup and delivery locations and times, some might be preferred over others by customers, shippers, carriers, or other stakeholders. Taking these preferences into account could result in a higher service level.

6.2 Related work

To deal with multiple pickup or delivery options, several extensions of the classic VRP have been proposed (see Table 6.1). The Generalized VRP (GVRP) considers clusters of possible customer locations (Golden, 1978; Ghiani and Improta, 2000). For each cluster, a delivery needs to be performed at one location. To our knowledge, only Moccia et al. (2012) study the GVRP with Time Windows (GVRPTW). On the other hand, the Multi-Depot VRP (MDVRP) has single delivery locations for each customer, but allows pickup at one of several depots (Renaud et al., 1996; Cordeau et al., 1997; Vidal et al., 2012; Montoya-Torres et al., 2015). The MDVRP has its applications in cases where a company with multiple production facilities or warehouses has to supply customers. The GVRP in contrast, is generally applied in situations where it is too costly or time-consuming to serve all customers. Instead, one central point in each cluster is chosen and all customers in the area will be served from that point (Baldacci et al., 2010; Bektaş

Table 6.1: Overview of transportation approaches with multiple alternative locations.

Category	Reference	VRP	PDP	PA	DA	TW	PV	#Ord	#Veh
GVRP	Bektaş et al. (2011)	✓			✓			6–131	2–12
	Ghiani and Improta (2000)	✓			✓			24	4
	Moccia et al. (2012)	✓			✓	✓		6–131	2–12
MDVRP	Cordeau et al. (1997)	✓		✓				48–360	2–∞
	Renaud et al. (1996)	✓		✓				50–360	∞
	Vidal et al. (2012)	✓		✓				50–417	∞
VRPRDL	Ozbaygin et al. (2017)	✓			✓	✓		15–120	∞
	Reyes et al. (2017)	✓			✓	✓		15–120	∞
VRPDO	Dumez et al. (2021)	✓			✓	✓	✓	50–400	2–41
	Tilk et al. (2021)	✓			✓	✓	✓	25–50	∞
GPDPP	This chapter		✓	✓	✓	✓	✓	500–2000	100–400

GVRP: Generalized VRP; **MDVRP:** Multi-Depot VRP; **VRPRDL:** VRP with Roaming Delivery Locations; **VRPDO:** VRP with Delivery Options; **GPDPP:** Generalized PDP with Preferences; **VRP:** Vehicle Routing Problem; **PDP:** Pickup and Delivery Problem; **PA:** Pickup alternatives; **DA:** Delivery alternatives; **TW:** Time windows; **PV:** Preference values; **#Ord:** Number of orders; **#Veh:** Number of vehicles.

et al., 2011). More recently, Reyes et al. (2017) and Ozbaygin et al. (2017) proposed a trunk delivery application: instead of only at one home location, parcels can be delivered to the car of a customer, that can be at different locations during a day. The resulting VRP with Roaming Delivery Locations (VRPRDL) differs from the GVRPTW in the fact that time windows cannot overlap: the itinerary of a car is respected.

When the material for this chapter was prepared, no application of service levels or preferences had been proposed for scenarios with multiple locations. Besides soft time windows, linear preference functions had been proposed for VRPs (Ghannadpour et al., 2014; Lin et al., 2014), but these were limited to a single time window per customer. Later on, customer preferences have been proposed in the VRP with Delivery Options (VRPDO), by Tilk et al. (2021) and Dumez et al. (2021).

In all models described above, all requests have the same pickup location(s). In contrast, we consider problems where each request has its own set of possible pickup and delivery locations. The transport operator must choose one pickup location and one delivery location from these sets to actually visit. In addition to a time window for each location, we expand the model with a preference value for each location, resulting in the Generalized Pickup and Delivery Problem with Preferences (GPDPP). With this framework, we provide in finding a balance between route costs and satisfaction of customers or other stakeholders.

6.3 Generalized pickup and delivery problems with preferences

The GPDPP shares most of its properties with the DCPDP defined in Section 2.1, but instead of a single pickup location p_o and a single delivery location d_o , each order $o \in O$ has a set of possible pickup locations P_o and a set of possible delivery locations D_o . We give a complete description and an integer linear program (ILP) formulation below, in which we abstract from individual carriers, individual shippers, and the dynamic aspect.

An instance is described by the structure given in Table 6.2, constrained to the restrictions of Table 6.3. Given are a set of transport orders, a set of locations, travel times and travel costs associated with each pair of them, and a fleet of vehicles for conducting the transport requests. Instead of single pickup and delivery locations, as is the case in the PDP, each order $o \in O$ has a set of possible pickup locations P_o and a set of possible delivery locations D_o . Multiple separated time windows for the same location can be modeled by introducing location duplicates. This contrasts with current definitions, where only one time window per location is considered, even in case of soft time windows.

A load size, modeled as integer number, is associated with each request. For ILP modeling purposes, this load quantity is related to all possible pickup locations of the request and the negative load is related to all its delivery locations. Time windows in which loading or unloading can start are given by $[e_i, l_i]$ for location $i \in P \cup D$; note that these can differ per location alternative for the same order. The same holds for the service times (loading or unloading duration). For each order, a preference value π_i of 1 is assigned to the most preferred delivery location i . The alternative delivery locations for that request (if they exist) will be assigned a number in $(0, 1]$ to describe the satisfaction if that location is served, relative to the most preferred delivery location. The same holds for pickup locations: at least one pickup location is fully preferred and the others are measured relative to it.

Note that the stakeholder having the preferences is not made explicit, and can differ depending on the problem domain. For example, in multi-depot grocery distribution applications, the company can have preferences for pickup locations (based on stock and work force available in the depots), where the customer can have preferences for delivery locations (based on, for instance, distance from home or parcel weight).

The fleet of vehicles, each with its own capacity, is dispersed over multiple locations. Vehicles may start at the same location, but this is not necessary. Start and terminal locations of the vehicles are independent, that is, it is not required that a vehicle ends at the point where it started its route. In the ILP below,

Table 6.2: Problem structure.

O	A set of orders
V	A set of vehicles
N	A set of locations
$N = P \cup D \cup A \cup \Omega$	A partition of N
$P = P_1 \cup \dots \cup P_{ O }$	A partition of P , where P_o is the set of possible pickup locations for order $o \in O$
$D = D_1 \cup \dots \cup D_{ O }$	A partition of D , where D_o is the set of possible delivery locations for order $o \in O$
$A = \{\alpha^1, \dots, \alpha^{ V }\}$	The set of start locations of the vehicles, such that vehicle $v \in V$ will start at location α^v
$\Omega = \{\omega^1, \dots, \omega^{ V }\}$	The set of end locations of the vehicles, such that vehicle $v \in V$ will end at location ω^v
$G = \langle N, E \rangle$	A complete directed graph over N
$\tau \in \mathbb{N}^+$	The time horizon of the problem
$q_i \in \mathbb{Z}$	The load quantity corresponding to location $i \in P \cup D$
$e_i \in \{0, \dots, \tau\}$	The earliest service start time at location $i \in N$
$l_i \in \{0, \dots, \tau\}$	The latest service start time at location $i \in N$
$s_i \in \mathbb{N}$	The service duration at location $i \in N$
$t_{ij} \in \mathbb{N}$	The travel time from i to j for each edge $\langle i, j \rangle \in E$
$c_{ij} \in \mathbb{R}^{\geq 0}$	The travel cost from i to j for each edge $\langle i, j \rangle \in E$
$\pi_i \in (0, 1]$	The customer's or operator's preference value for each location $i \in P \cup D$
k^v	The maximum load capacity of vehicle $v \in V$
β	The weight of dissatisfaction relative to travel cost

Table 6.3: Problem constraints.

$\forall i \in P \quad q_i \geq 0$	Pickup quantities are non-negative
$\forall o \in O \quad \forall i \in P_o \quad \forall j \in D_o \quad q_j = -q_i$	Delivery quantities are the opposite of corresponding pickup quantities
$\forall i \in A \cup \Omega \quad (e_i = 0 \wedge l_i = \tau \wedge s_i = 0)$	Start and end locations are represented as requests with specific properties
$\forall o \in O \quad \exists i \in P_o \quad \exists j \in D_o \quad \pi_i = \pi_j = 1$	At least one of the pickup and delivery locations of each request is fully appreciated
$\forall \theta \in \Theta(G) \quad \exists \langle i, j \rangle \in \theta \quad s_i + t_{ij} > 0$	Each cycle of pickup and delivery locations in G takes at least some travel time

$\Theta(G)$: the set of cycles in G for which all nodes are element of $P \cup D$.

Table 6.4: Decision variables.

$x_{ij}^v \in \{0, 1\}$	Assigned the value 1 when vehicle v travels along the arc $\langle i, j \rangle$
$y_i^v \in \{0, \dots, k^v\}$	The load of vehicle v after serving location $i \in N$
$z_i \in \{0, \dots, \tau\}$	The service start time at location $i \in N$

we limit ourselves to the scenario where all vehicles are available during the complete time span.

A solution to a GPDPP instance consists of a set of routes, such that all transport requests are handled: for each request, a vehicle needs to pick up the load at one of the possible pickup locations in the corresponding time frame; the same vehicle needs to deliver the load later on in its route at one of the possible delivery locations, respecting the time window. Vehicle capacities need to be respected during the route, and all vehicles need to satisfy their start and end location requirements. The goal is to find a solution with minimal cost. Cost is defined as the sum of travel costs for all vehicles added to the sum of realized dissatisfaction values multiplied by a weight β , as is formalized below by the mathematical program objective 6.1. The larger the value of β , the more important it is to satisfy the stakeholders. Note that all dissatisfaction values are counted equally in this model, but that any individual differences can be captured by setting the π_i values wisely. In more advanced models, different weight factors could be introduced for different stakeholders, or the product of all dissatisfaction values could be used to make sure that dissatisfaction is evenly distributed among the different stakeholders.

Given the problem structure and constraints in Tables 6.2 and 6.3 and the decision variables in Table 6.4, the mathematical programming formulation of the problem is as follows. The ILP can easily be obtained by applying standard linearization techniques on Constraints 6.4, 6.14, and 6.15.

$$\min \sum_{v \in V} \sum_{(i,j) \in E} c_{ij} x_{ij}^v + \beta \sum_{v \in V} \sum_{i \in N} \sum_{j \in PUD} (1 - \pi_j) x_{ij}^v \quad (6.1)$$

subject to

$$\sum_{v \in V} \sum_{i \in N} \sum_{j \in D_o} x_{ij}^v = 1 \quad \forall o \in O \quad (6.2)$$

$$\sum_{i \in D_o} \sum_{j \in N} x_{ij}^v - \sum_{i \in P_o} \sum_{j \in N} x_{ij}^v = 0 \quad \forall o \in O, v \in V \quad (6.3)$$

$$\sum_{i \in D_o} \sum_{j \in N} z_i x_{ij}^v - \sum_{i \in P_o} \sum_{j \in N} z_i x_{ij}^v \geq 0 \quad \forall o \in O, v \in V \quad (6.4)$$

$$\sum_{j \in N} x_{ij}^v - \sum_{j \in N} x_{ji}^v = 0 \quad \forall i \in PUD, v \in V \quad (6.5)$$

$$\sum_{j \in N} x_{\alpha^v j}^v = 1 \quad \forall v \in V \quad (6.6)$$

$$\sum_{v \in V} \sum_{j \in N} x_{ij}^v = 1 \quad \forall i \in A \quad (6.7)$$

$$\sum_{i \in N} x_{i\omega^v}^v = 1 \quad \forall v \in V \quad (6.8)$$

$$\sum_{v \in V} \sum_{i \in N} x_{ij}^v = 1 \quad \forall j \in \Omega \quad (6.9)$$

$$\sum_{v \in V} \sum_{i \in N} \sum_{j \in A} x_{ij}^v = 0 \quad (6.10)$$

$$\sum_{v \in V} \sum_{i \in \Omega} \sum_{j \in N} x_{ij}^v = 0 \quad (6.11)$$

$$\sum_{v \in V} \sum_{i \in N} x_{ii}^v = 0 \quad (6.12)$$

$$e_i \leq z_i \leq l_i \quad \forall i \in N \quad (6.13)$$

$$(z_i + s_i + t_{ij})x_{ij}^v - z_j x_{ij}^v \leq 0 \quad \forall i, j \in N, v \in V \quad (6.14)$$

$$(y_i^v + q_j)x_{ij}^v - y_j^v x_{ij}^v = 0 \quad \forall i \in N, j \in P \cup D, v \in V \quad (6.15)$$

$$y_{a^v}^v = 0 \quad \forall v \in V \quad (6.16)$$

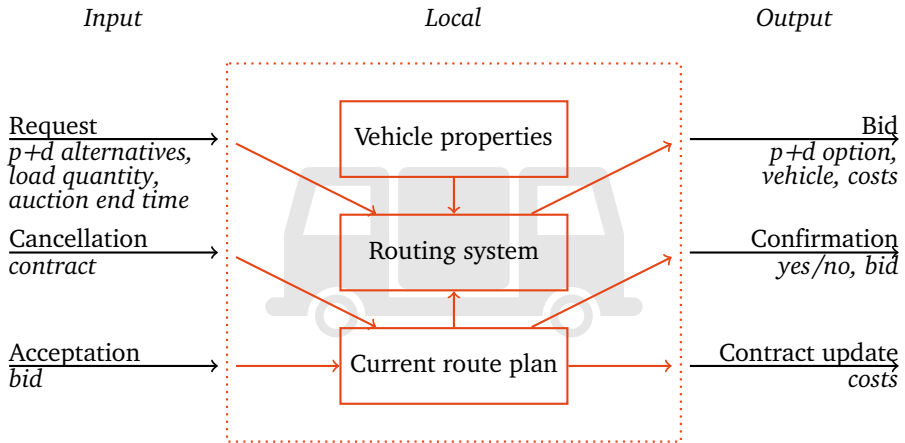
Constraints 6.2 state that all requests are handled exactly once. Constraints 6.3 and 6.4 enforce that related pickup and delivery tasks are coupled into the same vehicle, and occur in the right order. A consistent vehicle flow is guaranteed by Constraints 6.5–6.12, together with the subtour elimination properties given by Constraints 6.14 and the last constraint of Table 6.3. Temporal requirements are represented by Constraints 6.13 and 6.14 and capacity constraints by Constraints 6.15 and 6.16.

6.4 Auction approach

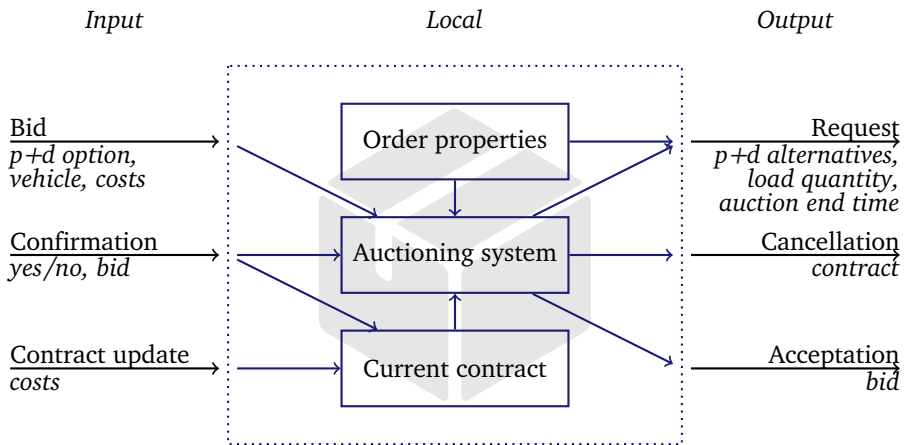
We extend the MAS approach proposed in Chapter 3 to solve the GPDPP in a decentralized manner, in accordance with the assumption that vehicles and orders can independently attach to a platform. Again, we have two types of agents that represent the main stakeholders of the problem: order agents, each responsible for getting one of the orders transported, and vehicle agents, each representing one vehicle.

Order agents try to make a contract with a vehicle agent for a pickup and a delivery alternative with high preference values, but are cooperative in the sense that they take vehicle routing costs into account and accept lower preference values if this decreases the routing costs enough. Vehicle agents are responsible for making contracts with order agents, but have the local goal of minimizing the sum of travel costs while keeping a feasible route.

When they enter the system, order agents send a request for transportation to a well-selected set of vehicle agents. These compute the marginal costs for inserting the order into their current route (by solving multiple subproblems – one for each combination of an alternative pickup and delivery location of the



(a)



(b)

Figure 6.1: Information input and output, as well as local information flows, for a vehicle agent (Figure 6.1a) and an order agent (Figure 6.1b).

new order, together with the orders already included in their route) and propose different bids – one for each option. Order agents evaluate the different bids from the vehicle agents and choose the one that is best, not only based on the routing costs of the bid, but also based on their own preferences. If no changes have occurred meanwhile in the route of the chosen vehicle agent, the order will be inserted into its route. Orders auction themselves again after some time to check if there are better options due to the dynamics within vehicle routes.

The MAS approach differs at two points from the approaches earlier developed in this thesis. First, vehicle agents can send multiple bids to an order agent, resulting from different combinations of pickup and delivery alternatives. Second, preference costs for each of these bids are locally added by the order agents themselves. Thus, preferences are not reported to the vehicle agents that process the request. This limits the private information that is exchanged between the different actors (see Figure 6.1). Furthermore, instead of considering auctions to run instantaneously, we again make the more natural assumption that they take some time, as we did in Chapter 3. Hence, the situation may have changed after a bid has been accepted. The consequence is that a vehicle agent still has the right to withdraw when an order agent has accepted its bid.

6.4.1 Order agent

An order agent keeps track of the contract of the order, consisting of a transporting vehicle, one pickup and one delivery location that are agreed on, as well as the costs for transportation. Initially, there is no contract; the order agent organizes auctions for obtaining and improving a contract.

An order agent starts its first auction in the system immediately after its release time. First, it selects a set of vehicle agents to send a request for transportation. Similarly to the full plan sharing approach in Chapter 3, we select the vehicles based on the spatiotemporal distance of the order to the planned routes of the vehicles, but now, we take the different pickup and delivery alternatives into account. The order agent opens the auction by sending all its possible pickup and delivery locations, the corresponding time windows and service durations, the load quantity, and the time at which the auction will end to all selected vehicles.

If an order agent receives a bid (consisting of a pickup location i , a delivery location j , and the marginal travel costs) from a vehicle agent, it adds its dissatisfaction costs ($\beta(2 - \pi_i - \pi_j)$, see Section 6.3) for the specific alternatives to the route costs to obtain the total costs for the bid. Subsequently, it stores the bid in a sorted list with increasing bid costs. When the auction time has ended, the order agent selects the first bid of its list, compares the costs of that bid to the costs of its current contract, if possible, and acts appropriately:

- If the costs of the selected bid are lower than that of the current contract, or there is no current contract, the order agent asks the vehicle agent that proposed the bid to insert the order into its route. If a positive response follows, the order agent updates its current contract, cleans up its bid list and schedules to start a new auction after some time. Furthermore, a message is sent to the vehicle agent of the previous contract (if applicable) to inform this agent that the order can be removed from its route. In case of a negative response of the vehicle agent, the bid has become outdated. In this case, the order agent possibly includes a new bid of the vehicle agent into its bid list, selects the next bid of its bid list and repeats the procedure.
- If the costs of the selected bid are not lower than the costs of the current contract, the current contract is still the best option. The agent cleans up its bid list and schedules to start a new auction after some time.
- If there is no bid selected (i.e., the bid list was empty) and there is no current contract, the order agent immediately starts a new auction. If vehicle routes have been changed in the meantime, probably it will obtain some bid from the new auction. This is urgent since there is no contract yet.

6.4.2 Vehicle agent

A vehicle agent keeps track of the planned route of the vehicle, along with earliest and latest possible times for each location, and the used vehicle capacity at each trajectory. Initially, the route only consists of the vehicle's start and end locations.

When a vehicle agent receives a request from an order agent, it checks if the auction has not yet ended. If there is still time, it computes the marginal travel costs for inserting each combination of alternatives into its current route, that is, it solves the single-vehicle PDP multiple times: once for each possible combination of a pickup and a delivery alternative of the new order. If an insertion is possible, a bid consisting of the marginal travel costs, the pickup location, and the delivery location is sent to the order agent. Hence, a vehicle agent can return multiple bids based on one request.

For quick vehicle computations, we again use a fast greedy insertion heuristic instead of solving the local vehicle problem in an exact way. The current sequence of the route will be kept, and feasibility (of time windows and capacities) will be checked for insertion of the new pickup and delivery at all possible positions (see Figure 2.3).

If a vehicle agent receives a message of acceptance of a bid from an order agent, it checks whether including the corresponding pickup and delivery loca-

tions into its route is still possible for the same (or less) costs. If this can be done, the vehicle agent updates its route accordingly and confirms this to the order agent. Otherwise, it sends a negative response to the order agent, together with a new bid for the same pickup and delivery locations, if possible. The rationale is that the vehicle agent still might have a better offer than other vehicle agents, although the costs might be higher than in the initial bid.

Each time a vehicle agent changes its route plans (after insertion or removal of an order), it informs all order agents that are affected by the changes about their new travel costs: for all order agents that have a pickup or delivery directly before or after an inserted or deleted location in the route, the vehicle agent computes what it would gain by removing the pickup and delivery of that order. These actual marginal costs will be sent to the corresponding order agents; they update the costs of their contracts, which is useful when order agents compare bids to their contract in a new auction.

6.4.3 Example

An example of a auction round is given in Figure 6.2, where we abstract from time windows and preferences. We consider an order agent having one pickup and two delivery alternatives, and two vehicles with current routes $\langle \alpha^1, 1, 2, \omega^1 \rangle$ and $\langle \alpha^2, 3, 4, \omega^2 \rangle$. The order agent sends a request with its pickup option 5 and delivery alternatives 6 and 7 to both vehicle agents. The vehicle agents each consider the two options, one with delivery alternative 6 and one with delivery alternative 7. They insert the new stops into their current routes and compare the costs (as defined by the graph of Figure 6.3) of the different new routes to the cost of their current routes. A bid with the least increase in costs is sent back to the order agent. Note that delivery alternative 6 is not feasible for vehicle agent 2; hence, only one bid is sent back. The order agent selects the best bid (consisting of pickup option 5 and delivery alternative 6 with a cost of 2), and notifies vehicle agent 1 of its acceptance.

6.5 Computational study

To get insight in the complexity of the newly defined problem, we perform two computational experiments. First, we define a set of small-scale scenarios, both with and without alternative locations, and compare the solutions to learn what gains can be obtained by allowing pickup and delivery alternatives. Next, we apply the MAS developed in Section 6.4 on large-scale problem instances, and compare its performance with an adaptive large neighborhood search (ALNS) approach.

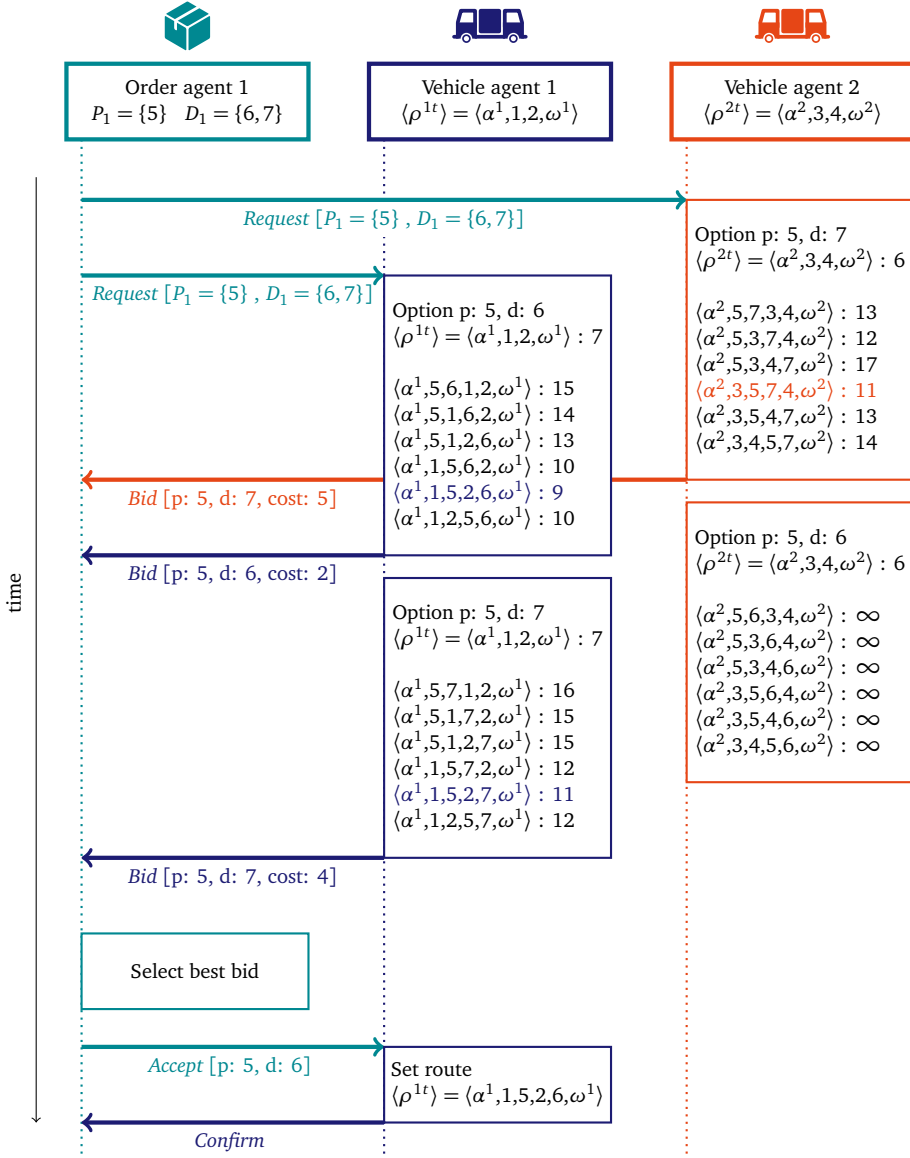


Figure 6.2: Schematic overview of an auction round.




Figure 6.3: Graph with initial vehicle routes and an order with two possible delivery locations corresponding to the auction process overview of Figure 6.2. Edges represent travel costs between locations.

6.5.1 Improvements through alternative locations

To investigate the possible improvements through alternative locations, we construct a data set with small-scale instances.¹ We consider a scenario in which 5 cooperative retail chains each have 2 branches on a 100×100 area. The retail chains have vehicles available, which they share in a fleet to serve all customers. All vehicles, with capacity 250 or 500, have one branch location as start and end point. We consider $|O|$ customers that place an order. Each customer has 4 randomly generated possible delivery locations, and the two branches of a random chain are the possible pickup locations. We generate 40 instances, 10 of each for $|O| \in \{2, 5, 10, 20\}$. To compare the scenario with multiple locations with the single location scenario, we make a duplicate of each instance in which we keep only one pickup and one delivery location with preference value 1. All requests have a random load between 1 and 100. For large loads (> 50), delivery preference values for all locations except one are set to 0.1, 0.2, or 0.3 to represent that the customers prefer home delivery. For small loads (< 10), they are set to 0.8, 0.9, or 1. For other load quantities, delivery preference values are assigned randomly. All pickup preference values equal 1, that is, it does not matter from which branch the goods are picked up. Travel times are equivalent to the Euclidean distance between two locations (rounded up to integers) and

¹Available at <http://doi.org/10.4121/16638283>.

Table 6.5: Instance parameter specifications.

General values	Values for $i \in P$	Values for $i \in D$
$\tau = 960$	$s_i = 3$	$3 \leq s_i \leq 15$
$t_{ij} = \lceil \sqrt{ x_i - x_j ^2 + y_i - y_j ^2} \rceil$	$e_i = 0$	$30 \leq l_i - e_i \leq 960$
$c_{ij} = 0.5t_{ij} + \sigma, \sigma \in [-4, 4]$	$l_i = 960$	$-100 \leq q_i \leq -1$
$\beta = 20$	$\pi_i = 1$	
$k^v \in \{250, 500\}$		

travel costs are about half of the travel time to represent an operation cost of about 30 euros per hour per vehicle. Furthermore, $\beta = 20$ to represent that a dissatisfied (but still served) customer can be compared to an extra cost of 20 euros. Hence, about 40 minutes extra driving time is allowed to fully satisfy a customer. Further instance details can be found in Table 6.5.

We solve the problem both with an ALNS algorithm and with the exact solver Gurobi. The first method has shown promising results in solving other routing and scheduling problems (Ropke and Pisinger, 2006; Pisinger and Ropke, 2007). Based on an initial solution, part of the orders is removed from the routes, and subsequently reinserted into the impaired solution. By iteratively applying this procedure, we aim for convergence to a low-cost solution. ALNS combines different heuristics for exploring the neighborhood; in each iteration, the destroy and reinsert heuristic are selected based on their performance in previous iterations. We mainly follow the procedure described by Ropke and Pisinger (2006), adapted to the GPDPP where necessary.

We applied ALNS 10 times on each instance and stopped each run if there was no improvement in the last 2000 iterations. The best result of these 10 runs was provided as initial solution to the ILP solver. We restricted the time for finding an exact solution to 9000 seconds. All experiments were performed on a 64-bit machine running Linux with Intel i5-4590 CPU at 3.30GHz and 8 GB of RAM.

Table 6.6 presents the computational results for both the scenario with 1 pickup and 1 delivery option and the scenario with multiple alternatives in terms of the gap between the solver's best value and its best bound and the computation times for both methods. Furthermore, the average improvement in objective value for the scenario with multiple locations compared to the scenario with single locations is given for each problem size.

The solution found by the ILP solver is never better than the ALNS solution. Hence, the metaheuristic finds optimal solutions for all cases where optimality is proven by the ILP solver. Note that the scenario without alternatives for 5 orders is always solved to optimality by the ILP solver, whereas the scenario

Table 6.6: Average results for solutions by the exact solver and by ALNS, based on 10 instances per problem size.

#Orders	Configuration	Gap (%)	Exact time (s)	ALNS time (s)	Improvement (%)
2	1 pickup, 1 delivery	0.00	0.71	0.45	
	2 pickups, 4 deliveries	0.00	11.25	0.66	40.02
5	1 pickup, 1 delivery	0.00	324.77	2.14	
	2 pickups, 4 deliveries*	0.00	3765.23	2.31	38.10
10	1 pickup, 1 delivery	32.45	⊕	5.55	
	2 pickups, 4 deliveries	43.28	⊕	15.86	31.88
20	1 pickup, 1 delivery	50.62	⊕	27.20	
	2 pickups, 4 deliveries	⊗	⊗	132.33	28.67

⊕: Time limit reached; ⊗: Memory limit reached.

* The results for the exact solver are the average of 4 out of 10 instances; the other 6 instances could not be solved within the available time and resulted in a gap of 29.88% on average.

with alternatives for 5 orders is not always solved to optimality. For 10 orders, both scenarios are already too complex to solve exactly in 9000 seconds, whereas ALNS uses about 16 seconds. For the scenario with alternatives for 20 orders, the exact solver runs out of memory, whereas the metaheuristic still quickly comes up with a solution.

When comparing the scenario with alternatives to the scenario without alternatives, we see that improvements of about 30% can be made. The duration of the computation, however, has increased by a factor 5.

6.5.2 Comparing the MAS with ALNS

To gain insight into the performance of the MAS, we compare the results of the MAS for different auction sizes with the results of a standard ALNS approach on instances of moderate size.

We generated a problem set with instances of 500 to 2000 orders.² All orders have a random load between 0 and 100, 2 pickup locations, and 4 delivery locations on a 100×100 area. Travel times and travel costs correspond to Euclidean distances between the locations. Time windows are randomly generated with length at least 30, and for each instance, $\tau = 960$ and $\beta = 20$. The number of vehicles equals 20% of the number of orders, and vehicles have capacity 250 or 500.

To compare the performance of the MAS and ALNS and the influence of the number of vehicles in an auction, we consider the following methods:

²Available at <http://doi.org/10.4121/16638319>.

Table 6.7: Normalized costs for the different methods relative to the baseline solution (mean \pm standard deviation of 5 runs per instance).

O	T	I	Baseline	MAS-25%	MAS-50%	MAS-100%	ALNS	
500	5	1	1.00 \pm 0.04	1.14 \pm 0.02	1.15 \pm 0.02	1.18 \pm 0.02	0.97 \pm 0.01	
		2	1.00 \pm 0.03	1.08 \pm 0.01	1.13 \pm 0.00	1.12 \pm 0.03	0.95 \pm 0.02	
		3	1.00 \pm 0.06	1.15 \pm 0.01	1.15 \pm 0.03	1.22 \pm 0.04	0.98 \pm 0.02	
		4	1.00 \pm 0.05	1.18 \pm 0.01	1.20 \pm 0.03	1.25 \pm 0.03	1.06 \pm 0.04	
		5	1.00 \pm 0.03	1.20 \pm 0.02	1.20 \pm 0.02	1.23 \pm 0.01	1.09 \pm 0.06	
		Avg.	1.00 \pm 0.04	1.15 \pm 0.01	1.17 \pm 0.02	1.20 \pm 0.03	1.01 \pm 0.03	
	10	1	1.00 \pm 0.03	1.19 \pm 0.02	1.21 \pm 0.02	1.23 \pm 0.04	0.96 \pm 0.03	
		2	1.00 \pm 0.06	1.17 \pm 0.00	1.19 \pm 0.02	1.20 \pm 0.02	0.95 \pm 0.01	
		3	1.00 \pm 0.03	1.25 \pm 0.02	1.24 \pm 0.02	1.30 \pm 0.02	1.02 \pm 0.02	
		4	1.00 \pm 0.02	1.26 \pm 0.02	1.26 \pm 0.02	1.29 \pm 0.03	1.04 \pm 0.04	
		5	1.00 \pm 0.05	1.22 \pm 0.02	1.25 \pm 0.02	1.26 \pm 0.02	1.02 \pm 0.02	
		Avg.	1.00 \pm 0.04	1.22 \pm 0.02	1.23 \pm 0.02	1.25 \pm 0.02	1.00 \pm 0.02	
	1000	5	1	1.00 \pm 0.03	1.00 \pm 0.03	1.07 \pm 0.04	⊗	⊗
			2	1.00 \pm 0.03	1.03 \pm 0.02	1.12 \pm 0.02	⊗	1.14 \pm 0.01
			3	1.00 \pm 0.02	1.05 \pm 0.03	1.17 \pm 0.01	⊗	⊗
			4	1.00 \pm 0.02	1.06 \pm 0.04	1.17 \pm 0.02	⊗	⊗
			5	1.00 \pm 0.03	0.98 \pm 0.02	1.06 \pm 0.02	⊗	⊗
			Avg.	1.00 \pm 0.03	1.02 \pm 0.03	1.12 \pm 0.02	⊗	1.14 \pm 0.01
10		1	1.00 \pm 0.02	1.09 \pm 0.04	1.14 \pm 0.02	1.27 \pm 0.01	1.10 \pm 0.05	
		2	1.00 \pm 0.02	1.14 \pm 0.04	1.16 \pm 0.02	1.31 \pm 0.02	1.11 \pm 0.01	
		3	1.00 \pm 0.04	1.10 \pm 0.02	1.15 \pm 0.02	1.28 \pm 0.02	1.16 \pm 0.04	
		4	1.00 \pm 0.02	1.07 \pm 0.03	1.14 \pm 0.02	1.28 \pm 0.02	1.11 \pm 0.04	
		5	1.00 \pm 0.03	1.07 \pm 0.03	1.13 \pm 0.01	1.24 \pm 0.01	1.05 \pm 0.03	
		Avg.	1.00 \pm 0.02	1.09 \pm 0.03	1.14 \pm 0.02	1.28 \pm 0.02	1.10 \pm 0.03	
2000		5	1	1.00 \pm 0.03	⊗	⊗	⊗	⊗
			2	1.00 \pm 0.04	⊗	⊗	⊗	⊗
			3	1.00 \pm 0.02	⊗	⊗	⊗	⊗
			4	1.00 \pm 0.03	⊗	⊗	⊗	⊗
			5	1.00 \pm 0.01	⊗	⊗	⊗	⊗
			Avg.	1.00 \pm 0.03	⊗	⊗	⊗	⊗
	10	1	1.00 \pm 0.01	1.02 \pm 0.14	⊗	⊗	⊗	
		2	1.00 \pm 0.03	0.95 \pm 0.04	⊗	⊗	⊗	
		3	1.00 \pm 0.04	1.04 \pm 0.13	⊗	⊗	⊗	
		4	1.00 \pm 0.02	1.10 \pm 0.16	⊗	⊗	⊗	
		5	1.00 \pm 0.02	1.08 \pm 0.10	⊗	⊗	⊗	
		Avg.	1.00 \pm 0.02	1.04 \pm 0.11	⊗	⊗	⊗	

O: Number of orders per instance; T: Computational time in minutes; I: Instance number; ⊗: No solution found within the given time.

- **Baseline:** The centralized reference baseline first computes an initial solution with a greedy heuristic and thereafter gets the full computation time to improve this solution by applying ALNS. Hence, it can be seen as ALNS where an initial solution is already given.
- **MAS-25%, MAS-50%, MAS-100%:** The three MAS methods differ in vehicle interaction percentage: order agents send in each auction a request to the most promising 25%, 50%, or 100% of the available vehicles, respectively, as described in Sections 6.4.1 and 3.4.1.
- **ALNS:** ALNS starts from scratch and uses the complete available computation time to both build an initial greedy solution and improve this solution by applying ALNS.

We run the different methods 5 times on 5 problem instances of 500, 1000, and 2000 orders. Since solutions need to be provided quickly in highly dynamic real-world cases, we limit the computation time of our experiments to 5 and 10 minutes on an Intel i5-4590 CPU at 3.30GHz. Table 6.7 shows the normalized costs relative to the baseline result, along with averages per group. Some experiments did not result in feasible solutions since some orders were not assigned to a vehicle at all, due to limited time.

The MAS solutions get closer to the baseline solutions for smaller computation times, for lower vehicle interaction percentages, and for larger problem instances. Although ALNS produces better results than the MAS for the smaller instances, there is an opposite result for the larger instances: for 1000 order instances and a time limitation of 10 minutes, the MAS-25% method outperforms ALNS, and for the 5 minutes case, both the MAS-25% and the MAS-50% methods outperform ALNS. Furthermore, for the 2000 order instances, the MAS-25% method is still able to find a solution in 10 minutes while ALNS is not. In addition, the MAS-25% method is even highly competitive with the baseline for 1000 orders in 5 minutes and for 2000 orders in 10 minutes. Note that there are two problem instances (no. 5 of the 1000 order series and no. 2 of the 2000 order series) for which the mean MAS-25% result even outperforms the baseline result.

6.6 Implications

The preliminary computational studies in this chapter have shown that significant gains can be obtained if alternative locations are offered. Furthermore, the MAS approach can successfully be applied on this type of problem. For platform providers, this gives some opportunities.

First, they can stimulate shippers to provide as many pickup or delivery alternatives as possible – not only in terms of locations, but also in terms of different time slots. Although some shippers might still have specific needs and cannot provide alternatives, others are more flexible and might be willing to give up some of their control, for instance, in exchange for lower prices.

Second, platforms offering flexible requests might be more attractive to carriers: with more alternatives, it is more likely that carriers can find an option that is efficient for them and hence can reduce the costs. Moreover, carriers might apply dynamic pricing strategies to steer the request alternatives they get and hence affect the properties of their routing plans according to their own preferences or future opportunities.

6.7 Conclusions

We have introduced the GPDPP to be able to model logistical problems with multiple alternative service locations where preferences are taken into account. As expected, adding multiple pickup and delivery locations to a problem instance is highly beneficial in terms of total costs: our scenario allowed for objective value improvements of about 30% when 1 other pickup alternative and 3 other delivery alternatives were added. Although the amount of improvement may highly depend on the specific problem characteristics, these results indicate that a significant cost reduction can be obtained when multiple locations are taken into account.

When the number of orders in a GPDPP instance increases, the developed MAS approach can be of value. Multiple bids can be submitted per request, but only limited information is being exchanged, and preference values can be kept private. The performance of the MAS approach on instances with alternatives and preferences increases for large-scale, time-limited instances in comparison to a centralized approach. For instances of 2000 orders, the decentralized MAS is able to find solutions while a centralized ALNS approach suffers from time limitations. Thus, the multi-agent approach developed in this thesis can successfully be applied to practically relevant problems with specific user preferences, which answers Research Question 5.

Chapter 7

Conclusions and Future Research

This thesis considered the potential of carrier cooperation in dynamic environments where a lot of transportation requests must be served. It is well-known that large reductions in terms of operational costs and emissions can be obtained if carriers exchange part of their orders, but previous centralized approaches are limited in size and assumptions. Therefore, this thesis explored decentralized auction-based approaches that benefit from scalability and flexibility. We have investigated the applicability of such a system in terms of incentives for carriers to participate. More specifically, we have explored what private information needs to be exchanged, whether strategic behaviour can pay off, and how individual user preferences can be incorporated. Furthermore, we have analyzed what cost reductions can be realized through large-scale cooperation.

We summarize our conclusions and contributions in Sections 7.1 and 7.2, respectively, and give recommendations for future research in Section 7.3.

7.1 Conclusions

Carrier cooperation is recognized as an important way to increase the efficiency of fleet operations and reduce the negative effects of transportation. There is, however, a lack of approaches for large-scale collaboration in a dynamic world. Decentralized auction-based approaches seem promising in this respect, since they can benefit from scalability and flexibility. Therefore, this thesis centered at answering the following main question:

To what extent can an auction-based multi-agent system be applied to solve dynamic large-scale collaborative vehicle routing problems?

We formulated five research questions in Chapter 1 to cover different aspects of the requested applicability, and addressed them in Chapters 2–6. Here, we answer these research questions and integrate the respective conclusions to provide an answer to the main research question.

1. How can an auction-based multi-agent system be applied to collaborative vehicle routing problems?

In Chapter 2, we have introduced the Dynamic Collaborative Pickup and Delivery Problem. To be able to solve it, we generalized an existing auction-based multi-agent system to a multi-carrier variant. Current multi-agent approaches for vehicle routing problems generally assume that all vehicles belong to one carrier and represent the individual vehicles by autonomous agents that can acquire transportation orders being offered in auctions. To deal with a situation in which multiple carriers collaborate, we have proposed an extended approach in which carriers (each potentially having multiple vehicles) can bid for the orders and internally assign them to their own fleet of vehicles. Although computing the marginal costs for an order usually takes more time for a carrier with multiple vehicles than for a single vehicle agent, quick approximations of the marginal costs can be sufficient. This gives the opportunity to solve collaborative problem variants.

2. What is the value of information sharing within this system?

Multi-agent auction approaches for vehicle routing normally assume that real marginal costs are communicated as bid prices and that no other information is available. We have investigated in Chapter 3 what cost information is expedient to share within a system for a good solution, and what gains can be obtained if also information about vehicles' locations is available to the system. We found that more carrier information generally improves solution quality, but is not always necessary. Sharing cost information contributes to higher service levels,

but partial cost information can be sufficient in problems with limited fleet capacity or urgent orders. If information about vehicles' positions or route plans is available, it is no longer necessary to involve all carriers in each auction: a 40% interaction rate is sufficient then for obtaining a significant share of the possible profits. Hence, a reduction in communicational and computational load can be obtained.

3. What gains can be obtained by large-scale carrier cooperation?

In Chapter 4, we investigated the potential of large-scale carrier cooperation. Based on a real-world data set, We modeled situations with up to 1000 cooperating carriers and found that cost reductions of up to 77% can be obtained, whereas the current literature on collaborative transportation reports reductions of 20–30% for small-scale cooperation. Even when we assumed that carriers initially only possess tasks in their own neighborhood, still cost reductions of up to 68% through collaboration are observed. Both the cooperating carriers and the transportation platform benefit from the gains in our approach, but the distribution of the profits among them highly depends on the system's settings.

4. To what extent can participants benefit from strategic behaviour in the system?

To check whether the multi-agent auction approach developed in this thesis is robust in terms of misuse, we analyzed in Chapter 5 when strategic behaviour pays off for individual carriers and shippers. Asking lower prices than the real marginal costs can be advantageous for individual carriers. Although it will hurt the total efficiency, individual profits of up to 5 times the profits in a truthful setting can be obtained. It is, however, not clear beforehand by what amount bids can safely be lowered before it turns disadvantageous for the carrier. This highly depends on the gains that are attributed to the winning carrier in each auction. For shippers or carriers that want to outsource orders, reporting lower prices than they are willing to pay is a good strategy – their profit generally increases. The drawback, however, is that it might hinder successful transactions, and consequently, decrease the service level.

An expected advantage of a second-price auction scheme does not apply to our context: carriers that bid strategically always obtain higher profits than their competitors that bid truthfully. Since the service level is also lower than with a first-price auction scheme, the proposed second-price auction system is not favourable in the assumed setting.

5. How can the system assist in meeting specific user preferences?

In Chapter 6, we introduced a problem variant where modeling of multiple alternative service locations is supported and where preferences for the different

options can be defined. We extended the multi-agent auction in such a way that carriers can place multiple bids – one for each alternative. The auctioneer then weighs the bid values and the indicated preferences for the alternatives, and accepts one of the bids. This way, specific user preferences can be met, while preference values can be kept private.

Now we have discussed the different aspects of an approach for large-scale collaborative transportation raised by the respective research questions, we are able to answer the main question:

To what extent can an auction-based multi-agent system be applied to solve dynamic large-scale collaborative vehicle routing problems?

In Chapter 2, we extended an auction-based multi-agent system to a multi-carrier variant, which has then been used in Chapters 3–6 to solve dynamic large-scale collaborative pickup and delivery problems. In Chapter 4, we have shown that the approach is suitable for large-scale carrier cooperation and that the potential benefits of involving numerous carriers into a cooperation are huge: cost reductions of up to 77% have been observed. As shown in Chapter 3, the approach is rather flexible in terms of required information: the best solutions could generally be found if full information is provided, but carriers hesitant to reveal some private information could still participate in a collaboration – often it is not necessary that all information is available. Furthermore, the approach leaves room for flexibility in the transportation requests: richer problems with alternative locations and user preferences could be solved without explicitly disclosing the preferences, as shown in Chapter 6. When it comes to fairness and trust, applying the approach is more delicate. In Chapter 5, we investigated the possible advantages of strategic bidding, and observed that cheating can pay off for individuals. However, this is not straightforward and may largely depend on the problem and settings at hand. The distribution of realized cooperation gains among the different participants – carriers, shippers, and the platform – plays a key role in this matter.

In brief, the developed auction-based multi-agent system can be an important mechanism to enable and stimulate cooperation between carriers, and hence reduce driven kilometers – leading to less costs, less emissions, and less congestion. It can be applied in various large-scale dynamic situations with specific transportation requests and provides flexibility in its information assumptions. This way, it is likely attractive for carriers to cooperate. Nevertheless, a proper way of profit distribution and incentives that prevent cheating must be designed and implemented for specific applications.

7.2 Contributions

This thesis addressed various challenges in the field of collaborative vehicle routing. We provide an overview of the different approaches that we have used and their characteristics in Table 7.1 and summarize our contributions in Table 7.2.

Table 7.1: Characteristics of the approaches in the different chapters.

		Chapter 3 – Information Sharing Local auctions with different knowledge levels	Chapter 4 – Large-Scale Collaboration Local auctions of small bundles of orders	Chapter 5 – Strategic Behaviour 1st and 2nd price local auctions with false bids	Chapter 6 – User Preferences Local auctions with multiple bids per carrier
Instance size	Number of orders	1000	2000	2000	500–2000
	Number of carriers	75–150	10–1000	250	100–400
	Number of vehicles	75–150	1000	250–750	100–400
Order assignment	Initial		✓		
	Later on	✓	(✓)	✓	✓
Instance type	Real-world based		✓	✓	
	Artificial	✓			✓
Objective	Travel costs minimization	✓	✓	(✓)	✓
	Service level maximization	✓	(✓)	(✓)	
	Preference fulfillment				✓
	Profit maximization	✓	✓	✓	
Cost information sharing	No	✓			
	Limited	✓			
	Full	✓	✓	✓	✓
	False			✓	
Position information sharing	No	✓	✓	✓	
	Limited	✓			
	Full	✓			✓

(✓): Used only in parts of the computational study or of minor importance.

Table 7.2: Thesis contributions with references to the related publications.

		[1]	[2]	[3]	[4]	[5]	[6]	[7]
Chapter 2	■ We introduced the Dynamic Collaborative Pickup and Delivery Problem (DCPDP) to study large-scale transportation problems with multiple interacting carriers in a changing world and gave a formal description.			✓	✓			✓
	■ Furthermore, we developed a multi-agent auction approach suitable for the collaborative problem variant.			✓	✓			✓
Chapter 3	■ We investigated what the value of information sharing is in the multi-agent transportation system and provided trade-offs between participants' privacy and the performance of the system.							✓
	■ In addition, we studied scenarios where individual carriers have different policies with respect to the amount of information they share.							✓
Chapter 4	■ We investigated what transportation cost reductions can be obtained if numerous carriers exchange their requests in a large-scale collaborative approach.			✓	✓			
	■ We extended the system with the possibility of locally auctioning bundles of similar orders to improve solution quality.			✓	✓			
	■ Moreover, we compared the system with local bundling to a central combinatorial auction approach on instances of various sizes.			✓				
Chapter 5	■ We experimentally explored whether the proposed auction approach is incentive compatible.			✓	✓			
	■ We also examined possible payoffs of strategic behaviour in a system with a second-price auction scheme.			✓				
Chapter 6	■ We defined the Generalized Pickup and Delivery Problem with Preferences (GPDPP), a problem variant in which alternative pickup and delivery locations and user preferences could be modeled.							✓
	■ Finally, we extended the multi-agent auction approach in such a way that carriers could place separate bids for the alternative options and the auctioneers find a trade-off between costs and preferences.			✓				

[1]: Los et al. (2022a). [2]: Los et al. (2022b); [3]: Los et al. (2021); [4]: Los et al. (2020a); [5]: Los et al. (2020b); [6]: Los et al. (2020c); [7]: Los et al. (2018);

7.3 Future research

The research described in this thesis calls attention to several topics that need further investigation. In this section, we provide an overview of recommended directions for future research.

- **Predictive analytics and learning:** Throughout this thesis, we assumed that carriers are willing to accept orders if they can make at least some profit. The investments in resources and time, however, are not considered. In real-world situations, the potential profit of using the fleet capacity for other purposes will be important as well: carriers might be able to make much larger profits if they wait for other requests. Hence, opportunity costs (Mes et al., 2013) should be incorporated into the approach in some way.

Furthermore, participants might learn from previous bids. Whereas we assumed that carriers base their bids only on their marginal costs, information of previous bids of other participants can be incorporated to improve the bidding strategy (Figliozzi et al., 2005; Mes et al., 2013; Van Heeswijk, 2020).

More generally, stochastic information about demand, supply, and traffic conditions could be used to make more accurate predictions about the real value of an order for participants. Besides this, a topic for further investigation is how incentives can be designed such that platform users can learn about cooperation benefits through experimentation.

- **Profit distribution:** The approach that we proposed in Chapter 4 provides a way to allocate the gains of each small improvement to the associated carriers and shippers, which may act as incentive to participate. However, a certain percentage can be kept by the platform without the carriers being able to check the amount. A certain level of trust is thus necessary, but the platform is also not invulnerable: if the profits for carriers are repeatedly too low, they might withdraw. Experimental research on real-world data will be interesting in this respect.

Furthermore, the profit shares are fair on a local level within our approach, but the separate auctions are disconnected from each other, even though they are interrelated. Hence, profits could be highly imbalanced from a higher level perspective. Although more sophisticated gain sharing methods do exist (Guajardo and Rönnqvist, 2016), they are intractable in large-scale settings, and are not directly applicable in decentralized approaches (Gansterer et al., 2020b). Thus, research on fair and acceptable profit

allocation methods in dynamic, large-scale contexts is necessary for the acceptance of auction-based request exchange mechanisms in practical applications.

- **Strategic behaviour:** Strategic behaviour in the developed system is sometimes advantageous for individual carriers, but harms the others. As shown in Chapter 5, the way profits are distributed plays a role here. Thus, besides contributing to a fair profit allocation, individual gains per transaction should prevent successful cheating.

An alternative to the current approach – although it ignores the fairness property to even greater extent – could be to experiment with an additional predetermined financial reward (e.g., proportional to the scale of a request) for cooperating participants, instead of giving them part of the auction gains. This will remove the direct incentives for reporting lower bids. The indirect, long-term ones, however, might still exist, but at the same time, higher bids could generate more income as well.

Within our second-price auction approach, the auctioneer always offered a fixed proportional price to outsourcing carriers and shippers. Different ways to get the required amount from the current owners might improve the results. Another promising idea is to let the auctioneer make a loss in the auctions where it is difficult to gather the amount of the second price, if this can be compensated in other auction rounds.

- **Autonomy and individual differences:** A final topic for further investigation is how individual carriers' attitudes can be modeled appropriately and how they will influence a cooperation system. We already discussed different attitudes with respect to information sharing. In addition, different levels of autonomy can be of interest, where shippers and carriers either can be in charge of outsourcing orders themselves, or subcontract this process to the platform. In this respect, the risk of being left with unassigned orders could deliberately be carried by the shipper itself, or could be transferred to the platform, that may adapt its cost structure accordingly.

It is also relevant to conduct realistic computational studies that focus in depth on individual carriers and their behaviours rather than on group averages. Individuals might, for example, obtain severe losses even if the average profit increases with false bidding strategies. Carriers with different risk management strategies should be modeled to examine the interaction effects of their strategic behaviours.

Finally, future research should consist of experiments on real-world problems with specific user preferences to fully investigate the potential of our

approach. Besides preferences for alternative locations, different pricing approaches and service characteristics could be relevant to model.

In brief, incentives to cooperate should be developed for various types of carriers to engage them all. Individual trust and autonomy may be key conditions to fully exploit the benefits of cooperation in practice. The above-mentioned directions will contribute to an improvement in the way fleets of (autonomous) vehicles belonging to different carriers can be merged to jointly provide an efficient, sustainable, and reliable transportation service.

Appendix A

Marginal Costs

Here, we formally describe the marginal costs $MC_c^t(\hat{\delta})$ for an order $\hat{\delta}$ to be included into one of the routes of a carrier c at time t . We define the following:

- for all vehicles $v \in V_c$, let $\hat{\alpha}^v$ be the location of v at time t :

- if $t < d(\alpha^v)$, then $\hat{\alpha}^v = \alpha^v$;
- if $\exists h a(\rho_h^{vt}) \leq t < d(\rho_h^{vt})$, then

$$\hat{\alpha}^v = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_h^{vt} \\ \text{with } e_{\hat{\rho}} = l_{\hat{\rho}} = \max(t, s(\rho_h^{vt}) + s_{\rho_h^{vt}}); \quad (\text{A.1})$$

- if $\exists h d(\rho_{h-1}^{vt}) \leq t < a(\rho_h^{vt})$, then

$$\hat{\alpha}^v = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ the actual location of the vehicle at time } t, \\ \text{with } e_{\hat{\rho}} = l_{\hat{\rho}} = t; \quad (\text{A.2})$$

- $\hat{O}_c = \{o \in O \mid \exists v \in V_c \exists h \in \{1, \dots, n^{vt}\} (\rho_h^{vt} = p_o \wedge s(\rho_h^{vt}) > t)\} \cup \{\hat{\delta}\}$ is the set of current orders of c for which the service has not yet started at t , together with the new order;
- for all vehicles $v \in V_c$, $\hat{O}_c^v = \{o \in O \mid \exists h, \hat{h} \in \{1, \dots, n^{vt}\} (\rho_h^{vt} = d_o \wedge s(\rho_h^{vt}) > t \wedge \rho_{\hat{h}}^{vt} = p_o \wedge s(\rho_{\hat{h}}^{vt}) \leq t)\}$ is the set of orders of v for which the delivery has not yet started, but the pickup has already started;
- $\hat{P} = \{p_o \mid o \in \hat{O}_c\}$ is the set of relevant pickup locations;
- $\hat{D} = \{d_o \mid o \in \hat{O}_c\} \cup \{d_o \mid o \in \hat{O}_c^v\}$ is the set of relevant delivery locations;
- $\hat{A} = \{\hat{\alpha}^v \mid v \in V_c\}$ is the set of relevant start locations;

- $\hat{\Omega} = \{\omega^v \mid v \in V_c\}$ is the set of relevant end locations;
- $N = \hat{A} \cup \hat{\Omega} \cup \hat{P} \cup \hat{D}$ is the set of all relevant locations;
- $E = \{\langle i, j \rangle \mid i, j \in N, i \notin \hat{\Omega}, j \notin \hat{A}, i \neq j\}$ is the set of all relevant arcs; and
- $\hat{q}_i = \begin{cases} q_o & \text{if } i = p_o \in \hat{P} \\ -q_o & \text{if } i = d_o \in \hat{D} \\ 0 & \text{if } i \in \hat{\Omega} \\ \sum_{o \in \hat{O}_c^v} q_o & \text{if } i = \hat{\alpha}^v \in \hat{A} \end{cases}$ represents the quantity to pick up at location i .

Let $x_{ij}^v \in \{0, 1\}$ be a decision variable representing whether vehicle v traverses arc $\langle i, j \rangle$, $y_i^v \in \mathbb{Z}$ a decision variable representing the load of vehicle v after serving location i , and $z_i \in [0, \tau]$ a decision variable representing the service start time at location i . Then, we define the following mathematical program $M^{\hat{\delta}}$, which is a slight adaptation of the standard PDP (see Parragh et al., 2008):

$$\min \sum_{v \in V_c} \sum_{\langle i, j \rangle \in E} c_{ij} x_{ij}^v \quad (\text{A.3})$$

subject to

$$\sum_{v \in V_c} \sum_{i \mid \langle i, j \rangle \in E} x_{ij}^v = 1 \quad \forall j \in \hat{P} \cup \hat{D} \cup \hat{\Omega} \quad (\text{A.4})$$

$$\sum_{j \mid \langle \hat{\alpha}^v, j \rangle \in E} x_{\hat{\alpha}^v j}^v = 1 \quad \forall v \in V_c \quad (\text{A.5})$$

$$\sum_{i \mid \langle i, \omega^v \rangle \in E} x_{i \omega^v}^v = 1 \quad \forall v \in V_c \quad (\text{A.6})$$

$$\sum_{i \mid \langle i, j \rangle \in E} x_{ij}^v - \sum_{i \mid \langle j, i \rangle \in E} x_{ji}^v = 0 \quad \forall j \in \hat{P} \cup \hat{D}, v \in V_c \quad (\text{A.7})$$

$$\sum_{j \mid \langle p_o, j \rangle \in E} x_{p_o j}^v - \sum_{j \mid \langle d_o, j \rangle \in E} x_{d_o j}^v = 0 \quad \forall o \in \hat{O}_c, v \in V_c \quad (\text{A.8})$$

$$\sum_{j \mid \langle d_o, j \rangle \in E} x_{d_o j}^v = 1 \quad \forall v \in V_c, o \in \hat{O}_c^v \quad (\text{A.9})$$

$$z_{p_o}^v \leq z_{d_o}^v \quad \forall o \in \hat{O}_c, v \in V_c \quad (\text{A.10})$$

$$x_{ij}^v (z_i + s_i + t_{ij}) \leq x_{ij}^v z_j \quad \forall \langle i, j \rangle \in E, v \in V_c \quad (\text{A.11})$$

$$z_i \geq e_i \quad \forall i \in N \quad (\text{A.12})$$

$$z_i \leq l_i \quad \forall i \in N \quad (\text{A.13})$$

$$x_{ij}^v (y_i^v + \hat{q}_j) \leq x_{ij}^v y_j^v \quad \forall \langle i, j \rangle \in E, v \in V_c \quad (\text{A.14})$$

$$y_i^v \geq \max(0, \hat{q}_i) \quad \forall i \in N, v \in V_c \quad (\text{A.15})$$

$$y_i^v \leq \min(k^v, k^v + \hat{q}_i) \quad \forall i \in N, v \in V_c \quad (\text{A.16})$$

In the above mathematical program, Constraints A.4 guarantee that all orders will be served and that each end location will be attended by exactly one vehicle. Together with Constraints A.5 and A.6, it is assured that vehicles start and end at their required positions. Next, Constraints A.7 guarantee a consistent vehicle flow, that is, vehicles entering a location also leave the location. Then, Constraints A.8 couple pickups and deliveries into the same vehicle, whereas Constraints A.9 make sure that remaining deliveries (for which the pickup already has taken place before time t) will be assigned to the appropriate vehicles. Furthermore, Constraints A.10 guarantee that pickups take place before the corresponding deliveries, Constraints A.11–A.13 let all services take place in the right time windows, and Constraints A.14–A.16 prevent exceeding vehicles' capacities. Constraints A.11 also eliminate subtours, given that travel time and/or service duration are always positive. To obtain a mixed-integer linear program, Constraints A.11 and A.14 can easily be linearized (see Cordeau, 2006).

Finally, to compute the marginal cost for order \hat{o} , also an equivalent mathematical program M where \hat{o} is omitted from \hat{O}_c is defined. The marginal cost $\text{MC}_c^t(\hat{o})$ for including order \hat{o} in one of the routes of carrier c at time t is then defined as the solution of $M^{\hat{o}}$ minus the solution of M .

Appendix B

Comparing the Multi-Agent System with Centralized Heuristics

To compare the MAS developed in Chapter 3 with standard centralized heuristics, we have run it on the benchmark DPDP instances of Mitrović-Minić et al. (2004), which were also considered by Holborn (2013). Mitrović-Minić et al. (2004) used cheapest insertion heuristics and tabu search improvement while they focus both on short-term and long-term goals. Holborn (2013) used more advanced insertion and improvement heuristics. Both use a central approach and process orders from certain intervals in batches.

The 90 instances of Mitrović-Minić et al. (2004), containing 100, 500, or 1000 dynamic pickup and delivery requests with time windows, were defined on a complete directed graph. Hence, we restricted I to contain only duplicates of locations within $P \cup D \cup \Omega$. Furthermore, Equation 2.27 was replaced by

$$\rho_h^{vt} = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_h^{vu} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = a(\rho_h^{vu}), \quad (\text{B.1})$$

and Equation A.2 was replaced by

$$\hat{\alpha}^v = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_h^{vt} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = a(\rho_h^{vt}). \quad (\text{B.2})$$

We applied our MAS using FCS, since prices of orders are not defined on the instances, and did not use position sharing since we always interacted with all vehicles. A maximum number of 100 auctions was allowed for each order. Durations for processing auctions were set to 0.6 times the standard durations.

We ran our MAS 10 times for all instances. Orders were never rejected. Tables B.1, B.2, and B.3 show the results of the MAS on the 100, 500 and 1000 order instances and compare them with previously published results.

Our MAS improved upon previously published results by about 4% on the 100 order instances, and about 0.5% on average for the 500 and 1000 order instances, with mean improvements of about 2% for the best MAS results.

Hence, we can conclude that the MAS performs as well as current approaches. Although the decentralized approach does not directly take into account combinatorial properties of orders, the orders' active requests immediately after their release time and the frequent reauctioning turn out to be competitive with current centralized approaches.

Table B.1: Comparison of results for the instances with 100 requests from Mitrović-Minić et al. (2004).

Instance	M-M	H	$I_H(\%)$	Min	$I_{Min}(\%)$	Avg \pm Std	$I_{Avg}(\%)$
0	2,656.41	2,642.97	0.51	2,432.29	7.97	2,465.86 \pm 11.80	6.70
1	2,700.60	2,605.27	3.53	2,302.93	11.60	2,302.93 \pm 0.00	11.60
2	2,774.64	2,797.31	-0.82	2,622.01	5.50	2,622.01 \pm 0.00	5.50
3	2,853.89	2,695.67	5.54	2,454.04	8.96	2,454.04 \pm 0.00	8.96
4	2,787.88	2,727.12	2.18	2,623.02	3.82	2,623.02 \pm 0.00	3.82
5	2,965.55	2,790.13	5.92	2,604.05	6.67	2,604.05 \pm 0.00	6.67
6	2,631.34	2,596.22	1.33	2,444.16	5.86	2,444.16 \pm 0.00	5.86
7	2,674.47	2,725.43	-1.91	2,516.61	5.90	2,516.61 \pm 0.00	5.90
8	2,888.39	2,726.56	5.60	2,745.16	-0.68	2,745.16 \pm 0.00	-0.68
9	2,978.87	2,778.56	6.72	2,671.48	3.85	2,671.48 \pm 0.00	3.85
10	2,576.58	2,523.94	2.04	2,412.21	4.43	2,412.21 \pm 0.00	4.43
11	2,812.76	2,638.94	6.18	2,584.79	2.05	2,584.79 \pm 0.00	2.05
12	2,677.90	2,658.15	0.74	2,597.43	2.28	2,603.71 \pm 2.21	2.05
13	2,703.01	2,594.72	4.01	2,587.62	0.27	2,610.13 \pm 11.86	-0.59
14	3,016.79	2,740.93	9.14	2,747.01	-0.22	2,747.01 \pm 0.00	-0.22
15	2,759.91	2,684.19	2.74	2,607.91	2.84	2,607.91 \pm 0.00	2.84
16	2,694.01	2,539.09	5.75	2,516.35	0.90	2,516.35 \pm 0.00	0.90
17	2,894.00	2,698.98	6.74	2,619.64	2.94	2,619.64 \pm 0.00	2.94
18	2,696.56	2,703.61	-0.26	2,613.42	3.08	2,613.42 \pm 0.00	3.08
19	2,537.65	2,508.39	1.15	2,414.55	3.74	2,414.55 \pm 0.00	3.74
20	2,819.49	2,561.16	9.16	2,527.43	1.32	2,527.43 \pm 0.00	1.32
21	2,704.22	2,693.24	0.41	2,655.24	1.41	2,655.24 \pm 0.00	1.41
22	2,860.85	2,848.82	0.42	2,608.83	8.42	2,608.83 \pm 0.00	8.42
23	2,479.15	2,388.92	3.64	2,473.74	-3.55	2,473.74 \pm 0.00	-3.55
24	2,894.79	2,816.63	2.70	2,616.85	7.09	2,616.85 \pm 0.00	7.09
25	2,543.57	2,396.13	5.80	2,398.44	-0.10	2,398.44 \pm 0.00	-0.10
26	2,889.89	2,788.19	3.52	2,651.95	4.89	2,651.95 \pm 0.00	4.89
27	2,780.39	2,629.16	5.44	2,482.36	5.58	2,482.36 \pm 0.00	5.58
28	2,653.85	2,454.65	7.51	2,280.93	7.08	2,282.77 \pm 1.27	7.00
29	2,763.60	2,690.50	2.65	2,377.00	11.65	2,377.00 \pm 0.00	11.65
Average	2,755.70	2,654.79	3.60	2,539.65	4.19	2,541.79	4.10

M-M: Travel costs found by Mitrović-Minić et al. (2004); **H:** Travel costs found by Holborn (2013); **$I_H(\%)$:** Improvement of H relative to $M-M$; **Min:** Minimum travel costs found by the MAS within 10 runs; **$I_{Min}(\%)$:** Improvement of Min relative to the best of $M-M$ and H ; **Avg \pm Std:** Average travel costs and standard deviation found by the MAS within 10 runs; **$I_{Avg}(\%)$:** Improvement of Avg relative to the best of $M-M$ and H .

Table B.2: Comparison of results for the instances with 500 requests from Mitrović-Minić et al. (2004).

Instance	M-M	H	$I_H(\%)$	Min	$I_{Min}(\%)$	Avg \pm Std	$I_{Avg}(\%)$
0	10,053.62	8,739.43	13.07	8,624.81	1.31	8,642.96 \pm 39.79	1.10
1	9,699.48	8,349.54	13.92	8,348.31	0.01	8,426.96 \pm 27.65	-0.93
2	9,608.40	8,202.41	14.63	8,053.33	1.82	8,078.43 \pm 22.75	1.51
3	9,807.06	8,350.71	14.85	8,253.04	1.17	8,253.35 \pm 0.40	1.17
4	10,176.05	8,832.41	13.20	8,564.35	3.03	8,692.37 \pm 85.38	1.59
5	10,133.55	8,733.24	13.82	8,535.39	2.27	8,548.53 \pm 5.05	2.12
6	10,045.82	8,485.01	15.54	8,452.11	0.39	8,541.94 \pm 83.63	-0.67
7	9,978.97	8,753.23	12.28	8,326.42	4.88	8,520.84 \pm 144.00	2.65
8	9,651.25	8,513.46	11.79	8,172.69	4.00	8,422.74 \pm 150.06	1.07
9	9,707.42	8,865.12	8.68	8,316.24	6.19	8,707.15 \pm 190.44	1.78
10	9,200.16	8,458.44	8.06	8,195.19	3.11	8,445.93 \pm 100.02	0.15
11	9,710.40	8,586.22	11.58	8,248.30	3.94	8,477.08 \pm 125.44	1.27
12	9,748.16	8,600.62	11.77	8,420.31	2.10	8,487.48 \pm 79.83	1.32
13	9,961.84	8,380.88	15.87	8,557.82	-2.11	8,642.67 \pm 35.88	-3.12
14	9,560.35	8,390.46	12.24	8,154.69	2.81	8,200.68 \pm 65.42	2.26
15	9,296.75	8,448.59	9.12	8,225.50	2.64	8,352.47 \pm 92.73	1.14
16	9,784.43	8,500.53	13.12	8,337.16	1.92	8,479.57 \pm 121.81	0.25
17	9,917.51	8,411.73	15.18	8,451.29	-0.47	8,534.18 \pm 60.07	-1.46
18	9,729.92	8,554.13	12.08	8,355.77	2.32	8,424.19 \pm 48.33	1.52
19	9,721.48	8,297.99	14.64	8,184.58	1.37	8,257.46 \pm 51.87	0.49
20	10,118.79	8,742.17	13.60	8,572.05	1.95	8,765.13 \pm 154.12	-0.26
21	9,458.99	8,742.40	7.58	8,220.11	5.97	8,323.56 \pm 132.64	4.79
22	10,126.10	8,739.42	13.69	8,695.30	0.50	8,821.53 \pm 85.22	-0.94
23	9,879.78	8,533.37	13.63	8,500.39	0.39	8,500.39 \pm 0.00	0.39
24	9,313.77	8,572.88	7.95	8,380.59	2.24	8,478.14 \pm 100.30	1.11
25	9,637.84	8,323.20	13.64	8,268.16	0.66	8,428.31 \pm 99.25	-1.26
26	10,349.09	8,684.06	16.09	8,296.99	4.46	8,399.11 \pm 53.83	3.28
27	9,925.99	8,411.79	15.25	8,143.39	3.19	8,338.56 \pm 194.02	0.87
28	9,823.70	8,572.82	12.73	8,549.65	0.27	8,687.32 \pm 97.74	-1.34
29	9,997.84	8,066.93	19.31	8,416.19	-4.33	8,527.18 \pm 174.76	-5.71
Average	9,804.15	8,528.11	12.96	8,360.67	1.93	8,480.21	0.54

M-M: Travel costs found by Mitrović-Minić et al. (2004); **H:** Travel costs found by Holborn (2013); **$I_H(\%)$:** Improvement of H relative to $M-M$; **Min:** Minimum travel costs found by the MAS within 10 runs; **$I_{Min}(\%)$:** Improvement of Min relative to the best of $M-M$ and H ; **Avg \pm Std:** Average travel costs and standard deviation found by the MAS within 10 runs; **$I_{Avg}(\%)$:** Improvement of Avg relative to the best of $M-M$ and H .

Table B.3: Comparison of results for the instances with 1000 requests from Mitrović-Minić et al. (2004). For these instances, only the average travel costs found by Mitrović-Minić et al. (2004) are available.

Instance	M-M	H	$I_H(\%)$	Min	$I_{Min}(\%)$	Avg \pm Std	$I_{Avg}(\%)$
0	⊗	13,909.70	⊗	13,478.84	3.10	13,938.26 \pm 225.96	-0.21
1	⊗	14,545.70	⊗	13,813.08	5.04	14,150.75 \pm 155.98	2.72
2	⊗	13,997.10	⊗	13,672.09	2.32	13,905.62 \pm 133.24	0.65
3	⊗	14,606.10	⊗	14,139.90	3.19	14,516.25 \pm 249.31	0.62
4	⊗	14,257.40	⊗	13,898.35	2.52	14,163.47 \pm 167.71	0.66
5	⊗	14,312.60	⊗	13,840.33	3.30	14,101.84 \pm 204.04	1.47
6	⊗	13,754.10	⊗	13,461.04	2.13	13,828.27 \pm 206.05	-0.54
7	⊗	14,202.80	⊗	13,842.95	2.53	14,077.49 \pm 144.50	0.88
8	⊗	14,003.10	⊗	13,762.18	1.72	14,121.62 \pm 227.92	-0.85
9	⊗	14,408.00	⊗	13,871.93	3.72	14,182.39 \pm 223.79	1.57
10	⊗	14,222.80	⊗	14,076.89	1.03	14,238.56 \pm 139.87	-0.11
11	⊗	14,496.50	⊗	14,004.43	3.39	14,225.29 \pm 158.21	1.87
12	⊗	14,324.20	⊗	13,741.95	4.06	14,176.29 \pm 214.21	1.03
13	⊗	14,079.10	⊗	13,587.32	3.49	13,923.32 \pm 148.97	1.11
14	⊗	13,923.80	⊗	13,649.60	1.97	14,073.83 \pm 257.04	-1.08
15	⊗	14,463.90	⊗	14,254.98	1.44	14,450.65 \pm 141.55	0.09
16	⊗	14,398.40	⊗	14,228.08	1.18	14,450.56 \pm 129.29	-0.36
17	⊗	14,626.00	⊗	13,897.35	4.98	14,202.78 \pm 163.18	2.89
18	⊗	14,288.40	⊗	13,586.74	4.91	13,747.76 \pm 123.50	3.78
19	⊗	13,366.20	⊗	13,206.87	1.19	13,503.14 \pm 212.28	-1.02
20	⊗	14,111.50	⊗	13,732.49	2.69	13,956.56 \pm 182.77	1.10
21	⊗	14,140.10	⊗	13,817.91	2.28	14,214.95 \pm 245.40	-0.53
22	⊗	14,010.20	⊗	13,732.70	1.98	14,085.22 \pm 198.89	-0.54
23	⊗	14,077.90	⊗	14,037.16	0.29	14,277.14 \pm 129.58	-1.42
24	⊗	14,203.80	⊗	13,936.68	1.88	14,156.15 \pm 180.83	0.34
25	⊗	13,709.20	⊗	13,270.70	3.20	13,626.73 \pm 214.01	0.60
26	⊗	14,080.90	⊗	13,545.05	3.81	13,972.07 \pm 215.28	0.77
27	⊗	13,907.80	⊗	13,950.05	-0.30	14,151.44 \pm 89.46	-1.75
28	⊗	14,655.20	⊗	14,307.43	2.37	14,463.46 \pm 93.16	1.31
29	⊗	14,562.10	⊗	13,950.44	4.20	14,182.80 \pm 139.83	2.60
Average	17,610.45	14,188.15	19.43	13,809.85	2.65	14,102.16	0.59

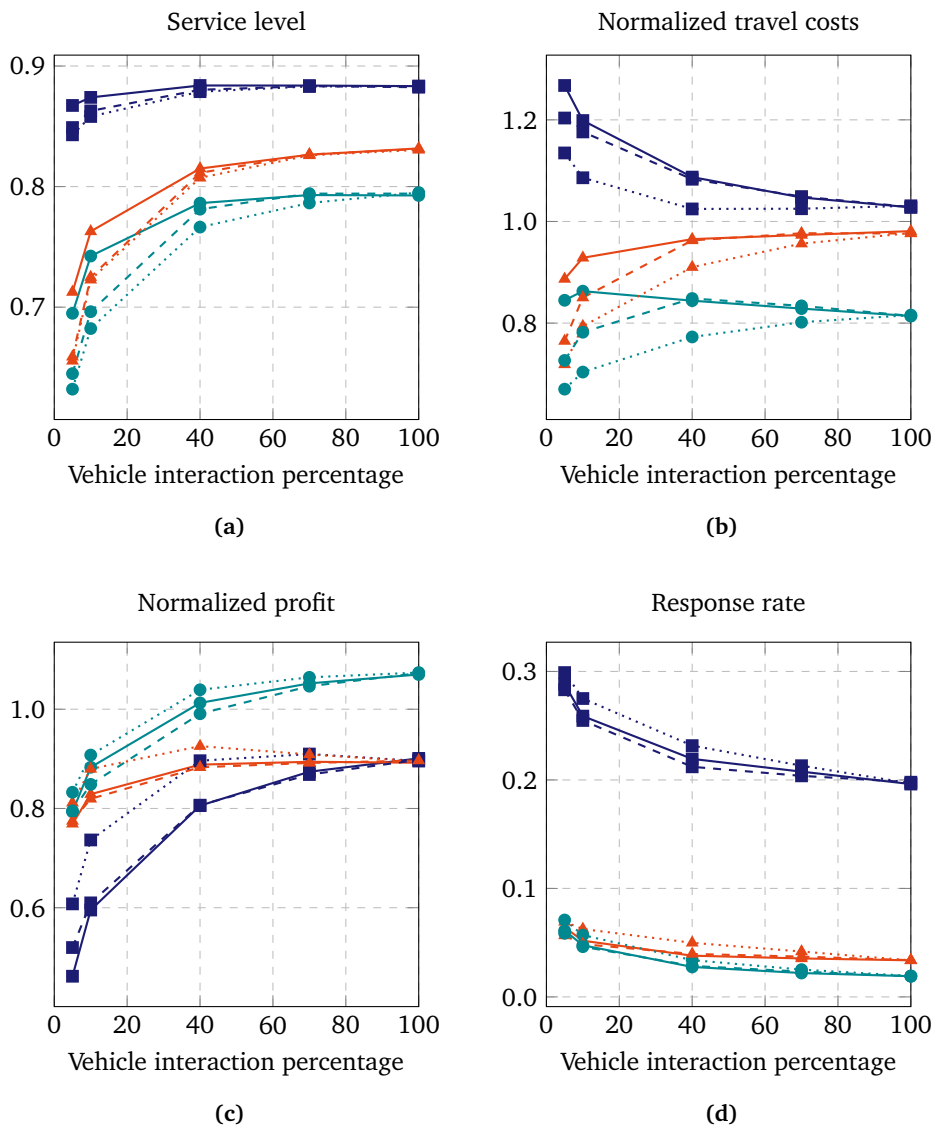
M-M: Travel costs found by Mitrović-Minić et al. (2004); **H:** Travel costs found by Holborn (2013); **$I_H(\%)$:** Improvement of H relative to $M-M$; **Min:** Minimum travel costs found by the MAS within 10 runs; **$I_{Min}(\%)$:** Improvement of Min relative to the best of $M-M$ and H ; **Avg \pm Std:** Average travel costs and standard deviation found by the MAS within 10 runs; **$I_{Avg}(\%)$:** Improvement of Avg relative to the best of $M-M$ and H ; **⊗:** Not available.

Appendix C

Results for Constrained Instance Sets

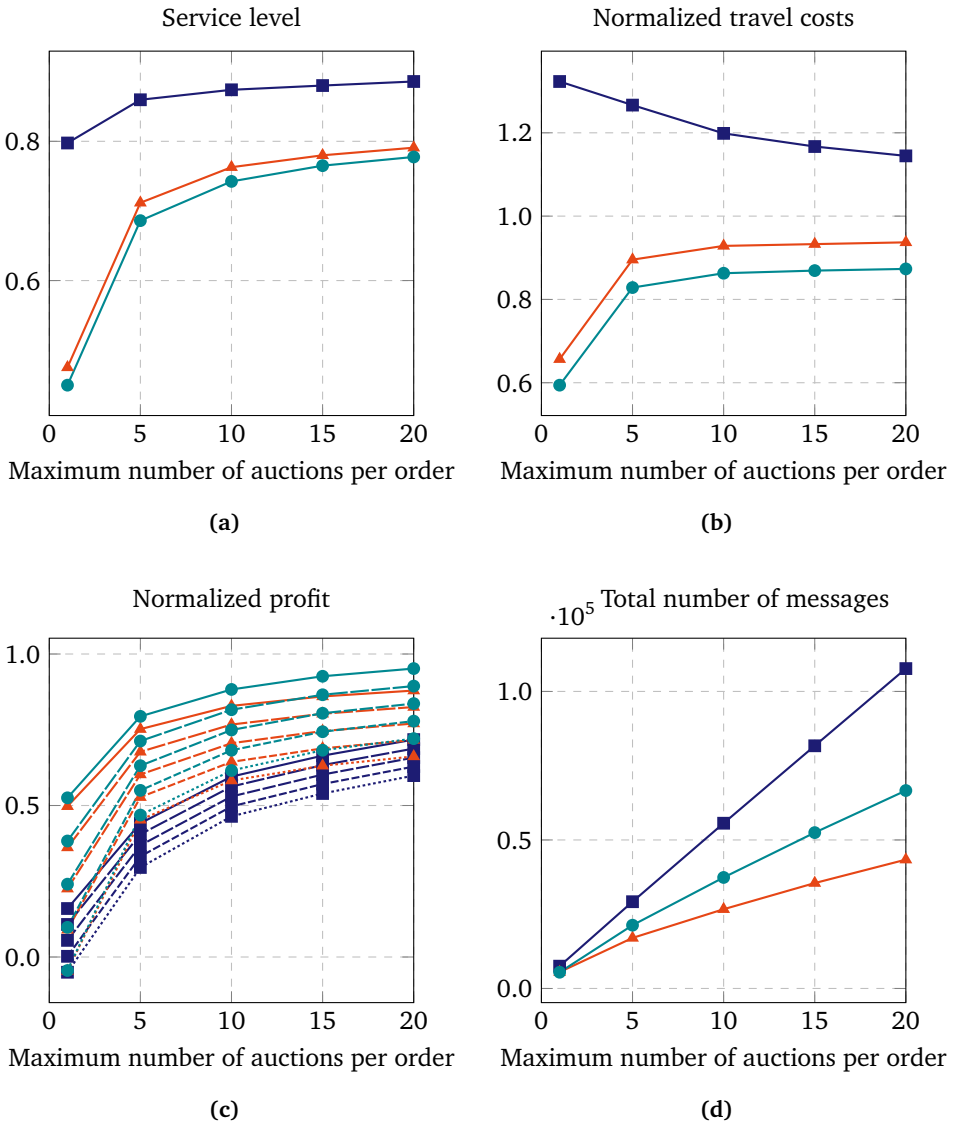
In this appendix, we provide the results of the experiments on the instance sets with lower fleet capacity and higher order urgency for different levels of information sharing, as described in Section 3.5.6.

The results for the low capacity set (instances with 75 vehicles) are given in Figures C.1–C.3 and the results for the medium capacity set (instances with 100 vehicles) are given in Figures C.4–C.6. Furthermore, the results for the urgent set (instances with small time windows) are given in Figures C.7–C.9.



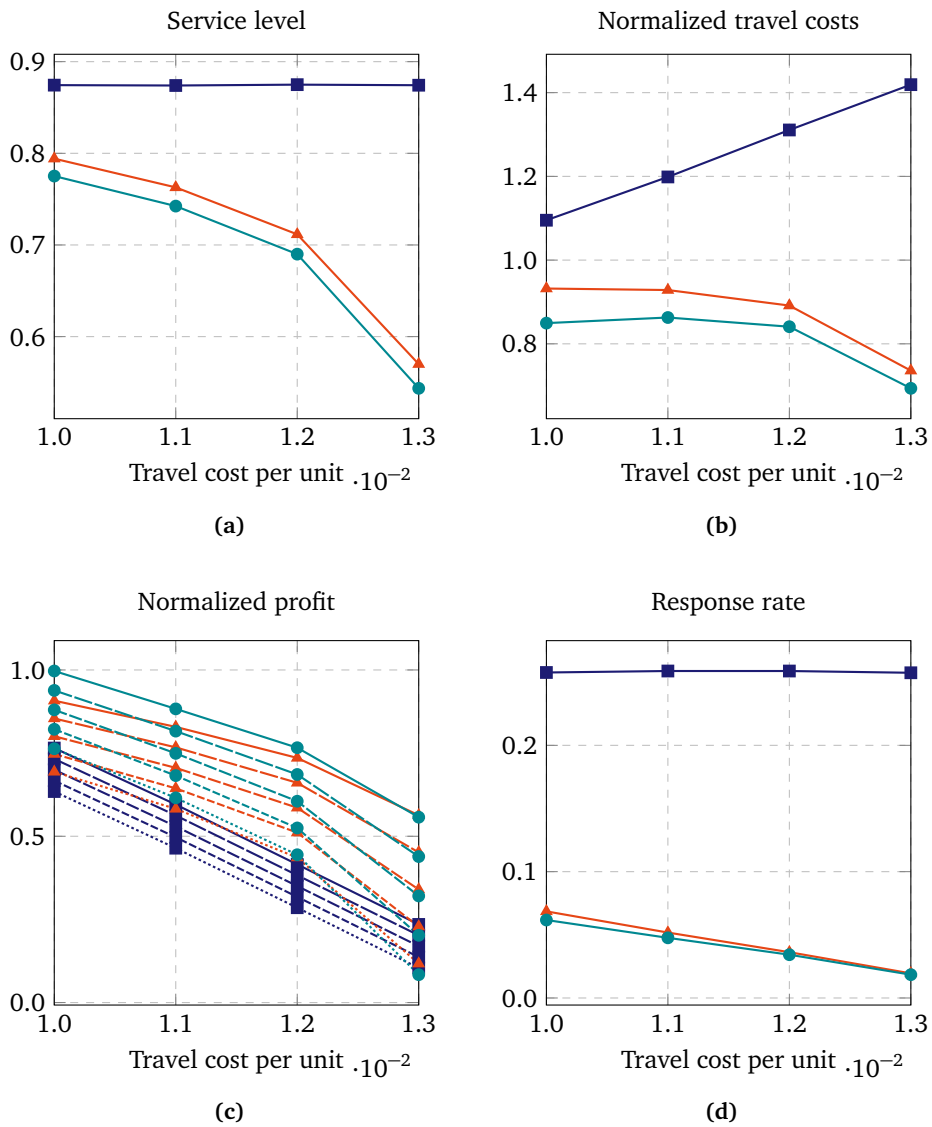
<i>Cost sharing</i>		<i>Position sharing</i>	
■	Full cost sharing (FCS)	—	No position sharing (NPS)
●	Partial cost sharing (PCS)	- - -	Current position sharing (CPS)
▲	No cost sharing (NCS)	⋯	Full plan sharing (FPS)

Figure C.1: Mean results on varying VIP for the three cost sharing and the three position sharing methods on the low capacity set. In Figure C.1c, there is no fine for rejected orders ($\gamma = 0$).



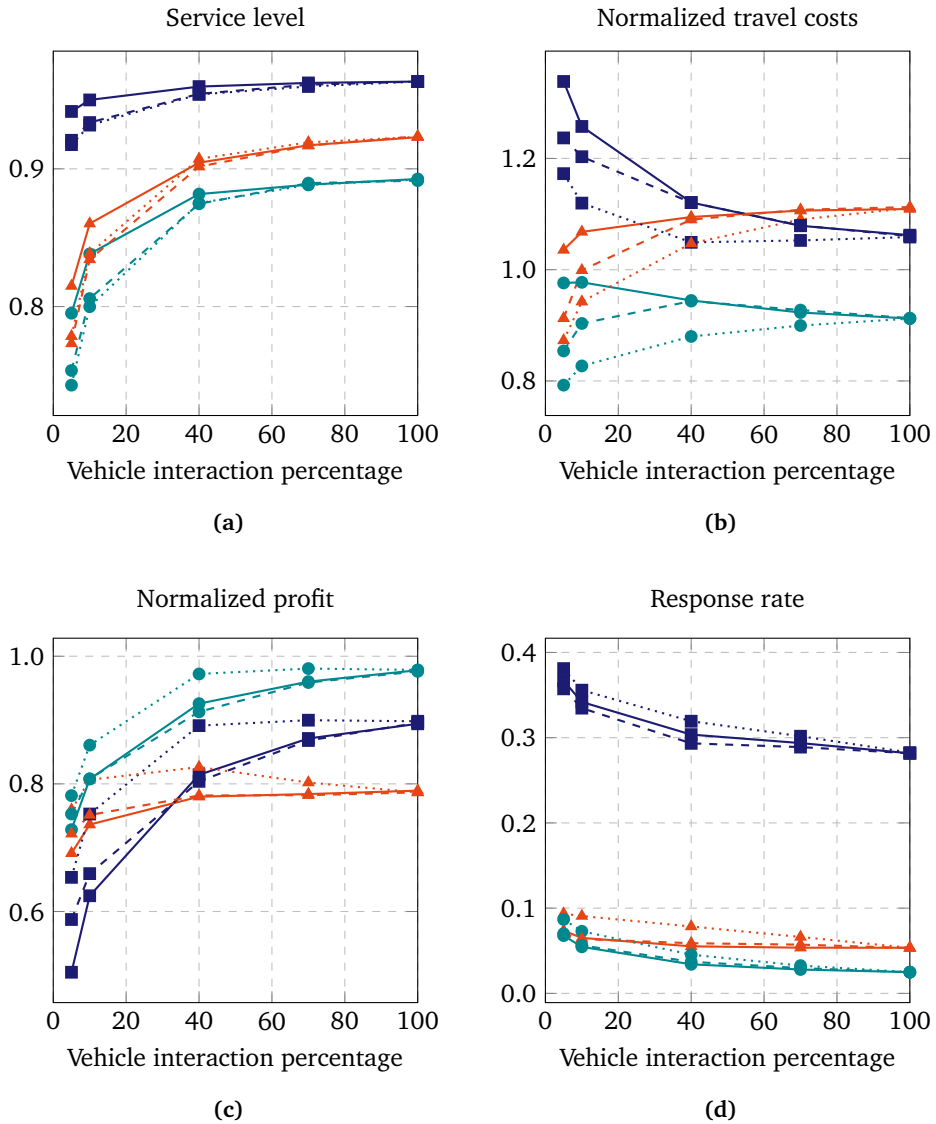
<i>Cost sharing</i>		γ			
■	Full cost sharing (FCS)	—	0.0	----	1.5
●	Partial cost sharing (PCS)	- · - ·	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure C.2: Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ) on the low capacity set.



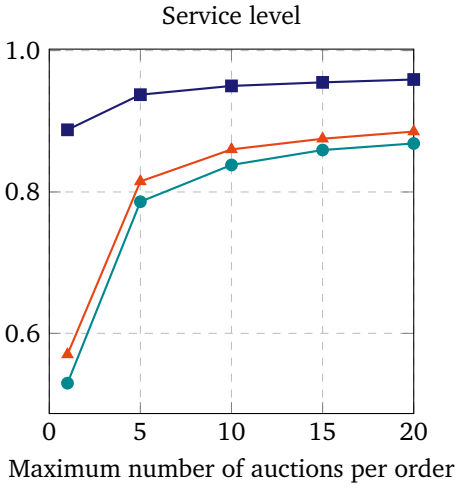
<i>Cost sharing</i>		<i>γ</i>			
■	Full cost sharing (FCS)	—	0.0	----	1.5
●	Partial cost sharing (PCS)	- - -	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure C.3: Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ) on the low capacity set.




<i>Cost sharing</i>		<i>Position sharing</i>	
■	Full cost sharing (FCS)	—	No position sharing (NPS)
●	Partial cost sharing (PCS)	- - -	Current position sharing (CPS)
▲	No cost sharing (NCS)	⋯	Full plan sharing (FPS)


Figure C.4: Mean results on varying VIP for the three cost sharing and the three position sharing methods on the medium capacity set. In Figure C.4c, there is no fine for rejected orders ($\gamma = 0$).




(a)



(b)



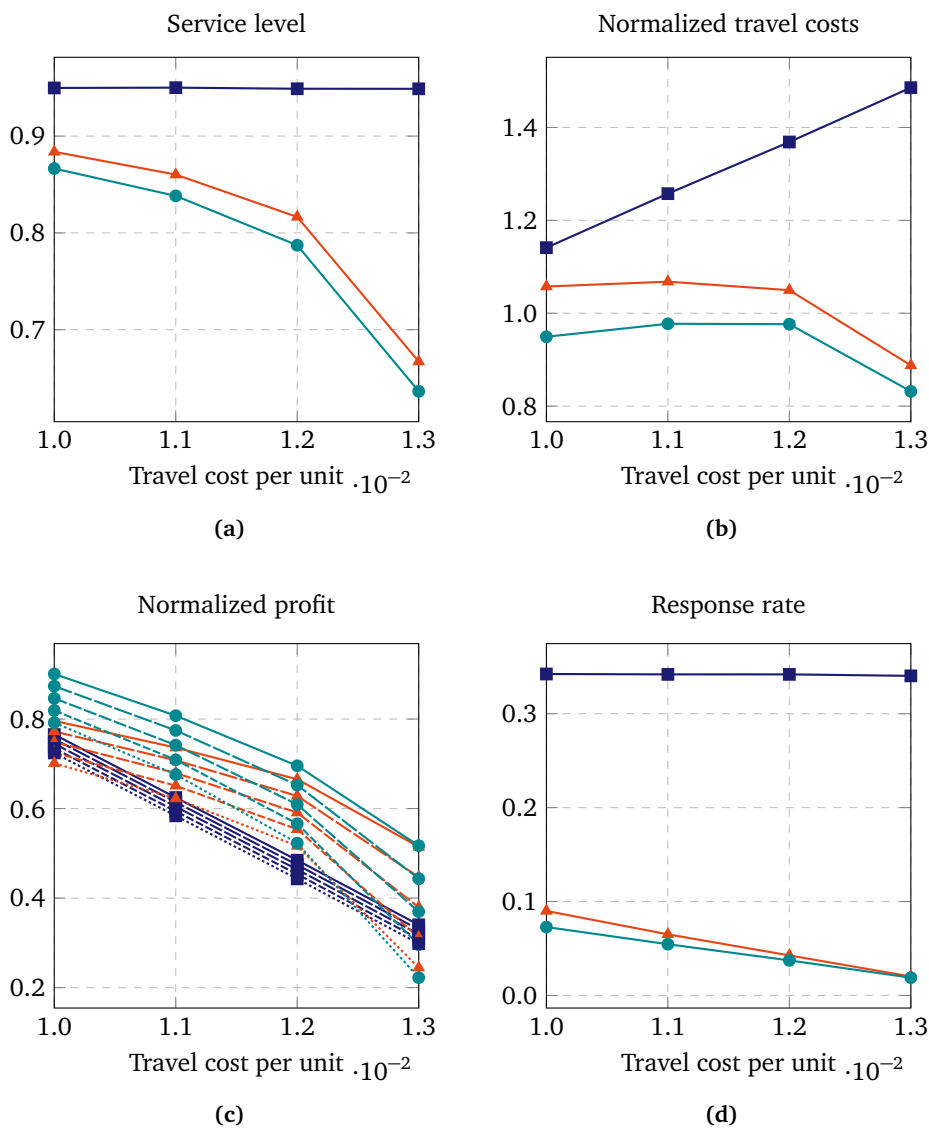
(c)



(d)

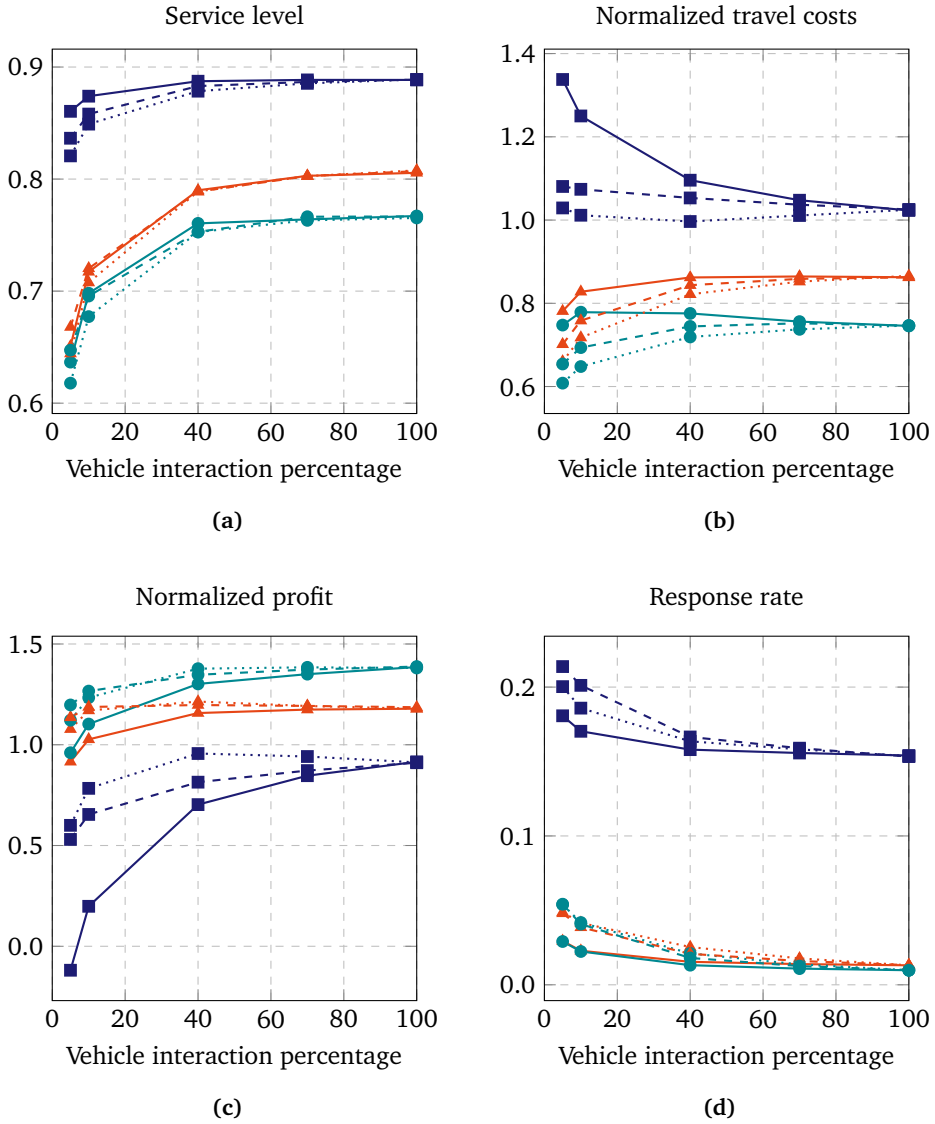
<i>Cost sharing</i>		γ			
■	Full cost sharing (FCS)	—	0.0	----	1.5
●	Partial cost sharing (PCS)	-·-	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure C.5: Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ) on the medium capacity set.



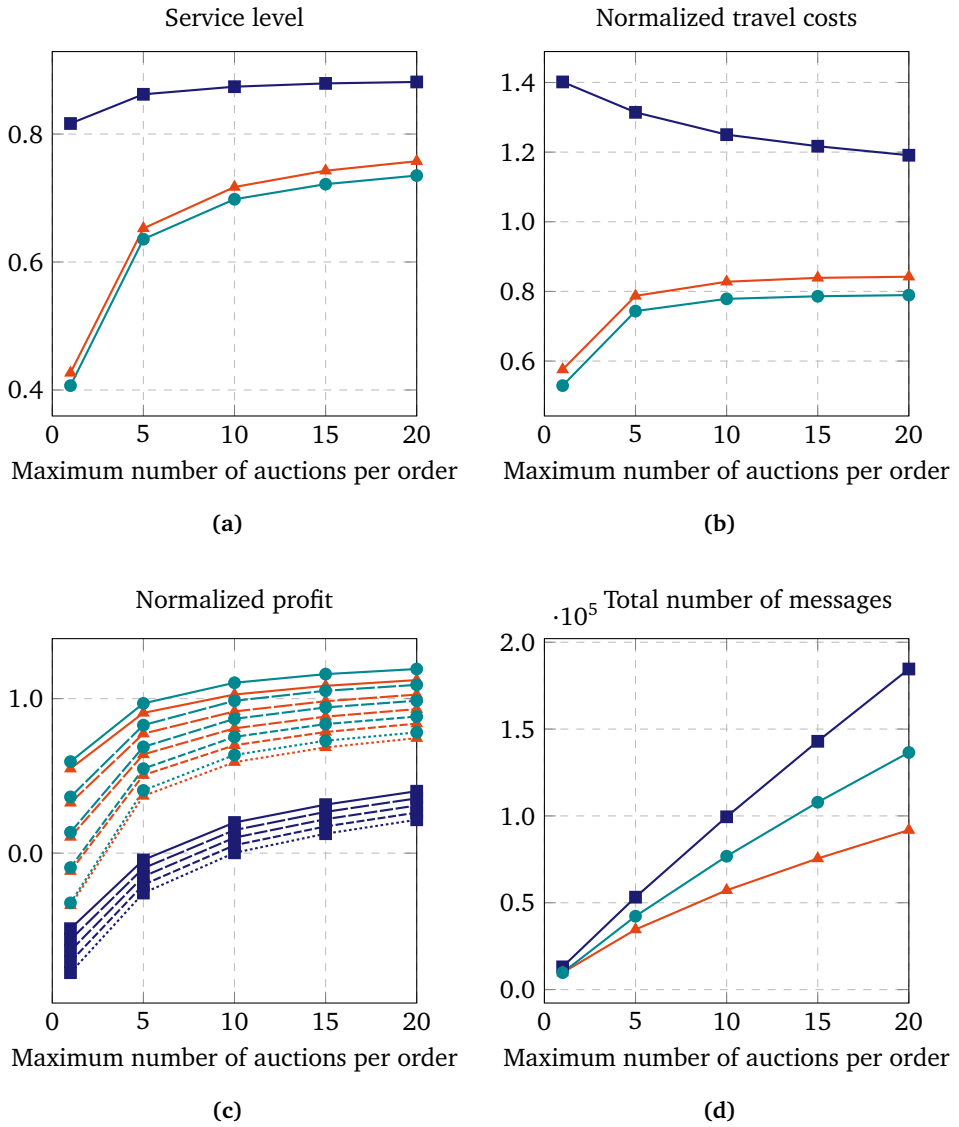
<i>Cost sharing</i>		γ			
■	Full cost sharing (FCS)	—	0.0	-----	1.5
●	Partial cost sharing (PCS)	- - -	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure C.6: Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ) on the medium capacity set.



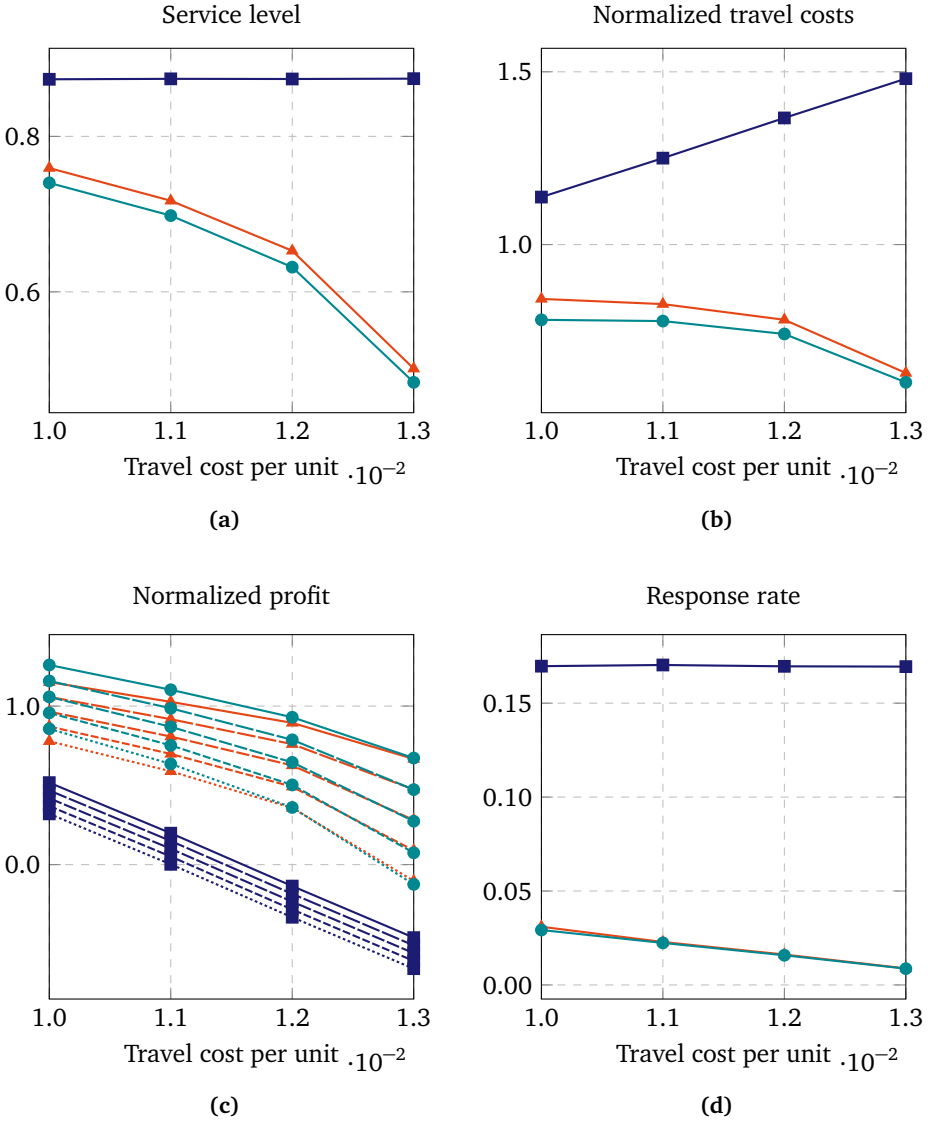
<i>Cost sharing</i>		<i>Position sharing</i>	
■	Full cost sharing (FCS)	—	No position sharing (NPS)
●	Partial cost sharing (PCS)	- - -	Current position sharing (CPS)
▲	No cost sharing (NCS)	⋯	Full plan sharing (FPS)

Figure C.7: Mean results on varying VIP for the three cost sharing and the three position sharing methods on the urgent set. In Figure C.7c, there is no fine for rejected orders ($\gamma = 0$).



<i>Cost sharing</i>		γ			
■	Full cost sharing (FCS)	—	0.0	----	1.5
●	Partial cost sharing (PCS)	- - -	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure C.8: Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ) on the urgent set.



<i>Cost sharing</i>		γ			
■	Full cost sharing (FCS)	—	0.0	----	1.5
●	Partial cost sharing (PCS)	- - -	0.5	2.0
▲	No cost sharing (NCS)	- - -	1.0		

Figure C.9: Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ) on the urgent set.

Appendix D

Comparing the Multi-Agent System with Combinatorial Auctions

Tables D.1, D.2, and D.3 show the results of the MAS proposed in Chapter 4 applied to instance sets O1, O2, and O3 proposed by Gansterer and Hartl (2016), and compare them with the BKSs as described by Gansterer et al. (2020a,b).¹ The profit in the MAS column is the best value among 100 runs in total (4 groups of 25 runs, with or without bundling, and with a maximum of 30 or 300 auctions per order).

¹Note that the results cannot be compared to the results reported by Lyu et al. (2019, Tables A14–A16), since other constraints or instance properties might have been used there. We thoroughly investigated where the differences in results might have come from, but the details of their solutions could not be obtained.

Table D.1: Best results for the MAS on benchmark set O1.

Set	Instance	BKS	MAS	Improvement (%)
O1_10	run=0+dist=200+rad=150+n=10	4827.75	4827.75	0.00
	run=1+dist=200+rad=150+n=10	3036.25	3033.79	-0.08
	run=2+dist=200+rad=150+n=10	3744.46	3744.46	0.00
	run=3+dist=200+rad=150+n=10	3468.53	3464.23	-0.12
	run=4+dist=200+rad=150+n=10	4286.33	4286.32	0.00
	run=5+dist=200+rad=150+n=10	3686.73	3686.73	0.00
	run=6+dist=200+rad=150+n=10	4599.17	4583.83	-0.33
	run=7+dist=200+rad=150+n=10	4286.41	4286.41	0.00
	run=8+dist=200+rad=150+n=10	3450.11	3489.57	1.14
	run=9+dist=200+rad=150+n=10	4070.29	3969.57	-2.47
	run=10+dist=200+rad=150+n=10	3625.60	3630.59	0.14
	run=11+dist=200+rad=150+n=10	3828.01	3828.01	0.00
	run=12+dist=200+rad=150+n=10	3337.18	3337.18	0.00
	run=13+dist=200+rad=150+n=10	3048.79	3048.79	0.00
	run=14+dist=200+rad=150+n=10	3512.80	3512.80	0.00
	run=15+dist=200+rad=150+n=10	4020.28	4016.96	-0.08
	run=16+dist=200+rad=150+n=10	3564.18	3564.18	0.00
	run=17+dist=200+rad=150+n=10	4172.44	4115.91	-1.35
	run=18+dist=200+rad=150+n=10	3945.07	4016.95	1.82
	run=19+dist=200+rad=150+n=10	2885.98	2855.45	-1.06
	Average			-0.12
O1_15	run=0+dist=200+rad=150+n=15	6053.44	6038.02	-0.25
	run=1+dist=200+rad=150+n=15	6805.56	6805.56	0.00
	run=2+dist=200+rad=150+n=15	6885.41	6960.50	1.09
	run=3+dist=200+rad=150+n=15	7778.42	7770.54	-0.10
	run=4+dist=200+rad=150+n=15	6079.67	6110.73	0.51
	run=5+dist=200+rad=150+n=15	7594.15	7557.20	-0.49
	run=6+dist=200+rad=150+n=15	6251.17	6323.67	1.16
	run=7+dist=200+rad=150+n=15	7425.76	7408.46	-0.23
	run=8+dist=200+rad=150+n=15	6527.29	6491.17	-0.55
	run=9+dist=200+rad=150+n=15	6135.65	6271.66	2.22
	run=10+dist=200+rad=150+n=15	6775.30	6843.83	1.01
	run=11+dist=200+rad=150+n=15	7040.55	7028.88	-0.17
	run=12+dist=200+rad=150+n=15	6389.52	6360.01	-0.46
	run=13+dist=200+rad=150+n=15	6037.21	6021.42	-0.26
	run=14+dist=200+rad=150+n=15	6095.78	6055.23	-0.67
	run=15+dist=200+rad=150+n=15	6899.75	6886.70	-0.19
	run=16+dist=200+rad=150+n=15	6525.64	6525.64	0.00
	run=17+dist=200+rad=150+n=15	6217.41	6201.34	-0.26
	run=18+dist=200+rad=150+n=15	6911.93	6911.32	-0.01
	run=19+dist=200+rad=150+n=15	5592.14	5586.17	-0.11
	Average			0.11

BKS: Best known solution (profit) found by Gansterer et al. (2020a,b); **MAS:** Profit found by the MAS; **Improvement (%)**: Improvement of MAS relative to BKS;

Table D.2: Best results for the MAS on benchmark set O2.

Set	Instance	BKS	MAS	Improvement (%)
O2_10	run=0+dist=200+rad=200+n=10	5681.90	6029.02	6.11
	run=1+dist=200+rad=200+n=10	4105.88	4105.88	0.00
	run=2+dist=200+rad=200+n=10	3323.96	3605.80	8.48
	run=3+dist=200+rad=200+n=10	5840.33	5896.11	0.96
	run=4+dist=200+rad=200+n=10	3923.62	4284.09	9.19
	run=5+dist=200+rad=200+n=10	5172.38	5590.90	8.09
	run=6+dist=200+rad=200+n=10	4306.61	4336.51	0.69
	run=7+dist=200+rad=200+n=10	5113.16	5310.22	3.85
	run=8+dist=200+rad=200+n=10	4415.43	4585.39	3.85
	run=9+dist=200+rad=200+n=10	5474.34	5518.35	0.80
	run=10+dist=200+rad=200+n=10	5314.20	5351.46	0.70
	run=11+dist=200+rad=200+n=10	5351.21	5470.72	2.23
	run=12+dist=200+rad=200+n=10	5778.25	5869.91	1.59
	run=13+dist=200+rad=200+n=10	4885.80	5131.42	5.03
	run=14+dist=200+rad=200+n=10	5217.46	5402.51	3.55
	run=15+dist=200+rad=200+n=10	5518.18	5618.84	1.82
	run=16+dist=200+rad=200+n=10	5431.95	5582.08	2.76
	run=17+dist=200+rad=200+n=10	4566.98	4547.06	-0.44
	run=18+dist=200+rad=200+n=10	5372.42	5624.97	4.70
	run=19+dist=200+rad=200+n=10	4710.49	4881.78	3.64
	Average			3.38
O2_15	run=0+dist=200+rad=200+n=15	8405.39	8459.64	0.65
	run=1+dist=200+rad=200+n=15	8931.09	9090.49	1.78
	run=2+dist=200+rad=200+n=15	9037.81	9054.71	0.19
	run=3+dist=200+rad=200+n=15	9722.75	9994.23	2.79
	run=4+dist=200+rad=200+n=15	8970.14	8948.72	-0.24
	run=5+dist=200+rad=200+n=15	8168.25	8243.93	0.93
	run=6+dist=200+rad=200+n=15	7493.37	7429.49	-0.85
	run=7+dist=200+rad=200+n=15	7512.26	7608.29	1.28
	run=8+dist=200+rad=200+n=15	6264.43	6369.51	1.68
	run=9+dist=200+rad=200+n=15	7937.83	7949.73	0.15
	run=10+dist=200+rad=200+n=15	7852.35	8028.10	2.24
	run=11+dist=200+rad=200+n=15	8836.52	8954.58	1.34
	run=12+dist=200+rad=200+n=15	9493.88	9590.01	1.01
	run=13+dist=200+rad=200+n=15	8967.75	8985.83	0.20
	run=14+dist=200+rad=200+n=15	7622.74	7745.15	1.61
	run=15+dist=200+rad=200+n=15	9468.63	9769.13	3.17
	run=16+dist=200+rad=200+n=15	8265.40	8475.67	2.54
	run=17+dist=200+rad=200+n=15	8299.04	8607.71	3.72
	run=18+dist=200+rad=200+n=15	9543.06	9613.84	0.74
	run=19+dist=200+rad=200+n=15	8837.45	8916.38	0.89
	Average			1.29

BKS: Best known solution (profit) found by Gansterer et al. (2020a,b); **MAS:** Profit found by the MAS; **Improvement (%)**: Improvement of MAS relative to BKS;

Table D.3: Best results for the MAS on benchmark set O3.

Set	Instance	BKS	MAS	Improvement (%)
O3_10	run=0+dist=200+rad=300+n=10	9032.72	10085.99	11.66
	run=1+dist=200+rad=300+n=10	8336.44	8915.17	6.94
	run=2+dist=200+rad=300+n=10	10150.10	10821.15	6.61
	run=3+dist=200+rad=300+n=10	9075.06	9643.72	6.27
	run=4+dist=200+rad=300+n=10	8650.27	9248.99	6.92
	run=5+dist=200+rad=300+n=10	8359.13	8935.33	6.89
	run=6+dist=200+rad=300+n=10	8555.03	9167.59	7.16
	run=7+dist=200+rad=300+n=10	8580.78	9283.23	8.19
	run=8+dist=200+rad=300+n=10	7960.82	8624.07	8.33
	run=9+dist=200+rad=300+n=10	7328.59	7845.96	7.06
	run=10+dist=200+rad=300+n=10	7491.47	8238.26	9.97
	run=11+dist=200+rad=300+n=10	7242.82	7599.63	4.93
	run=12+dist=200+rad=300+n=10	7223.90	8033.24	11.20
	run=13+dist=200+rad=300+n=10	8095.11	9320.57	15.14
	run=14+dist=200+rad=300+n=10	7750.75	8860.07	14.31
	run=15+dist=200+rad=300+n=10	8304.09	8653.14	4.20
	run=16+dist=200+rad=300+n=10	7780.18	8256.28	6.12
	run=17+dist=200+rad=300+n=10	7022.65	7553.64	7.56
	run=18+dist=200+rad=300+n=10	8462.22	8640.40	2.11
	run=19+dist=200+rad=300+n=10	7733.83	7815.14	1.05
	Average			7.63
O3_15	run=0+dist=200+rad=300+n=15	14156.90	14141.17	-0.11
	run=1+dist=200+rad=300+n=15	13443.10	14067.83	4.65
	run=2+dist=200+rad=300+n=15	12429.70	12839.42	3.30
	run=3+dist=200+rad=300+n=15	12345.00	12755.36	3.32
	run=4+dist=200+rad=300+n=15	14617.30	14635.17	0.12
	run=5+dist=200+rad=300+n=15	13880.40	14520.15	4.61
	run=6+dist=200+rad=300+n=15	15396.60	15709.87	2.03
	run=7+dist=200+rad=300+n=15	12246.60	12813.90	4.63
	run=8+dist=200+rad=300+n=15	15753.90	16234.34	3.05
	run=9+dist=200+rad=300+n=15	12041.30	12411.35	3.07
	run=10+dist=200+rad=300+n=15	13388.90	13960.70	4.27
	run=11+dist=200+rad=300+n=15	11300.50	11507.86	1.83
	run=12+dist=200+rad=300+n=15	14750.60	15332.83	3.95
	run=13+dist=200+rad=300+n=15	15133.00	15822.57	4.56
	run=14+dist=200+rad=300+n=15	15972.70	16632.43	4.13
	run=15+dist=200+rad=300+n=15	16064.00	16617.58	3.45
	run=16+dist=200+rad=300+n=15	11956.80	12309.97	2.95
	run=17+dist=200+rad=300+n=15	14508.20	14923.94	2.87
	run=18+dist=200+rad=300+n=15	12937.30	13505.86	4.39
	run=19+dist=200+rad=300+n=15	13384.50	14179.11	5.94
	Average			3.35

BKS: Best known solution (profit) found by Gansterer et al. (2020a,b); **MAS:** Profit found by the MAS; **Improvement (%)**: Improvement of MAS relative to BKS;

Bibliography

- Arnold, F., Gendreau, M., and Sörensen, K. (2019). Efficiently solving very large-scale routing problems. *Computers & Operations Research*, 107:32–42.
- Baldacci, R., Bartolini, E., and Laporte, G. (2010). Some applications of the generalized vehicle routing problem. *Journal of the Operational Research Society*, 61(7):1072–1077.
- Beirigo, B. A. (2021). *Dynamic Fleet Management for Autonomous Vehicles: Learning- and optimization-based strategies*. PhD thesis, Delft University of Technology, Delft, the Netherlands.
- Bektaş, T., Erdoğan, G., and Røpke, S. (2011). Formulations and branch-and-cut algorithms for the generalized vehicle routing problem. *Transportation Science*, 45(3):299–316.
- Berger, S. and Bierwirth, C. (2010). Solutions to the request reassignment problem in collaborative carrier networks. *Transportation Research Part E: Logistics and Transportation Review*, 46(5):627–638.
- Bertsimas, D., Jaillet, P., and Martin, S. (2019). Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 67(1):143–162.
- Campbell, A. M. and Savelsbergh, M. (2004). Efficient insertion heuristics for vehicle routing and scheduling problems. *Transportation Science*, 38(3):369–378.
- Cleophas, C., Cottrill, C., Ehmke, J. F., and Tierney, K. (2019). Collaborative urban transportation: Recent advances in theory and practice. *European Journal of Operational Research*, 273(3):801–816.
- Cordeau, J.-F. (2006). A branch-and-cut algorithm for the dial-a-ride problem. *Operations Research*, 54(3):573–586.

- Cordeau, J.-F., Gendreau, M., and Laporte, G. (1997). A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, 30(2):105–119.
- Crujijssen, F. (2020). *Cross-Chain Collaboration in Logistics: Looking Back and Ahead*, volume 297 of *International Series in Operations Research & Management Science*. Springer, Cham, Switzerland.
- Crujijssen, F., Dullaert, W., and Fleuren, H. (2007). Horizontal cooperation in transport and logistics: A literature review. *Transportation Journal*, 46(3):22–39.
- Dahl, S. and Derigs, U. (2011). Cooperative planning in express carrier networks – An empirical study on the effectiveness of a real-time Decision Support System. *Decision Support Systems*, 51(3):620–626.
- Dai, B. and Chen, H. (2011). A multi-agent and auction-based framework and approach for carrier collaboration. *Logistics Research*, 3:101–120.
- Dai, B., Chen, H., and Yang, G. (2014). Price-setting based combinatorial auction approach for carrier collaboration with pickup and delivery requests. *Operational Research*, 14:361–386.
- Dumez, D., Lehuédé, F., and Péton, O. (2021). A large neighborhood search approach to the vehicle routing problem with delivery options. *Transportation Research Part B: Methodological*, 144:103–132.
- European Environment Agency (2019). Emissions of air pollutants from transport. <https://www.eea.europa.eu/data-and-maps/indicators/transport-emissions-of-air-pollutants-8/transport-emissions-of-air-pollutants-8>.
- Fernández, E., Roca-Riu, M., and Speranza, M. G. (2018). The shared customer collaboration vehicle routing problem. *European Journal of Operational Research*, 265(3):1078–1093.
- Figliozzi, M. A. (2006). Analysis and evaluation of incentive-compatible dynamic mechanisms for carrier collaboration. *Transportation Research Record*, 1966(1):34–40.
- Figliozzi, M. A., Mahmassani, H. S., and Jaillet, P. (2004). Competitive performance assessment of dynamic vehicle routing technologies using sequential auctions. *Transportation Research Record*, 1882(1):10–18.
- Figliozzi, M. A., Mahmassani, H. S., and Jaillet, P. (2005). Impacts of auction settings on the performance of truckload transportation marketplaces. *Transportation Research Record*, 1906(1):89–96.

- Gansterer, M. and Hartl, R. F. (2016). Request evaluation strategies for carriers in auction-based collaborations. *OR spectrum*, 38:3–23.
- Gansterer, M. and Hartl, R. F. (2018a). Centralized bundle generation in auction-based collaborative transportation. *OR Spectrum*, 40:613–635.
- Gansterer, M. and Hartl, R. F. (2018b). Collaborative vehicle routing: A survey. *European Journal of Operational Research*, 268(1):1–12.
- Gansterer, M., Hartl, R. F., and Savelsbergh, M. (2020a). The value of information in auction-based carrier collaborations. *International Journal of Production Economics*, 221:107485.
- Gansterer, M., Hartl, R. F., and Sörensen, K. (2020b). Pushing frontiers in auction-based transport collaborations. *Omega*, 94:102042.
- Gansterer, M., Hartl, R. F., and Vetschera, R. (2019). The cost of incentive compatibility in auction-based mechanisms for carrier collaboration. *Networks*, 73(4):490–514.
- Gath, M. (2016). *Optimizing Transport Logistics Processes with Multiagent Planning and Control*. Springer, Wiesbaden, Germany.
- Ghannadpour, S. F., Noori, S., Tavakkoli-Moghaddam, R., and Ghoseiri, K. (2014). A multi-objective dynamic vehicle routing problem with fuzzy time windows: Model, solution and application. *Applied Soft Computing*, 14:504–527.
- Ghiani, G. and Improta, G. (2000). An efficient transformation of the generalized vehicle routing problem. *European Journal of Operational Research*, 122(1):11–17.
- Golden, B. L. (1978). Recent developments in vehicle routing. In *Proceedings of the Bicentennial Conference on Mathematical Programming*, pages 233–240, Gaithersburg, MD.
- Guajardo, M. and Rönnqvist, M. (2016). A review on cost allocation methods in collaborative transportation. *International Transactions in Operational Research*, 23(3):371–392.
- Guo, W. (2020). *Dynamic, Stochastic, and Coordinated Optimization for Synchronodal Matching Platforms*. PhD thesis, Delft University of Technology, Delft, the Netherlands.

- Van Heeswijk, W. (2020). Smart containers with bidding capacity: A policy gradient algorithm for semi-cooperative learning. In *Proceedings of the 11th International Conference on Computational Logistics (ICCL 2020)*, pages 52–67, Enschede, the Netherlands.
- Holborn, P. L. (2013). *Heuristics for Dynamic Vehicle Routing Problems with Pick-ups and Deliveries and Time Windows*. PhD thesis, Cardiff University, Cardiff, United Kingdom.
- Lai, M., Cai, X., and Hu, Q. (2017). An iterative auction for carrier collaboration in truckload pickup and delivery. *Transportation Research Part E: Logistics and Transportation Review*, 107:60–80.
- Li, J., Rong, G., and Feng, Y. (2015). Request selection and exchange approach for carrier collaboration based on auction of a single request. *Transportation Research Part E: Logistics and Transportation Review*, 84:23–39.
- Lin, C., Choy, K. L., Ho, G. T. S., Lam, H. Y., Pang, G. K. H., and Chin, K. S. (2014). A decision support system for optimizing dynamic courier routing operations. *Expert Systems with Applications*, 41(15):6917–6933.
- Van Lon, R. R. S. and Holvoet, T. (2017). When do agents outperform centralized algorithms? A systematic empirical evaluation in logistics. *Autonomous Agents and Multi-Agent Systems*, 31:1578–1609.
- Los, J., Schulte, F., Gansterer, M., Hartl, R. F., Spaan, M. T. J., and Negenborn, R. R. (2020a). Decentralized combinatorial auctions for dynamic and large-scale collaborative vehicle routing. In *Proceedings of the 11th International Conference on Computational Logistics (ICCL 2020)*, pages 215–230, Enschede, the Netherlands.
- Los, J., Schulte, F., Gansterer, M., Hartl, R. F., Spaan, M. T. J., and Negenborn, R. R. (2021). Large-scale collaborative vehicle routing. Under revision.
- Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2020b). Collaborative vehicle routing when agents have mixed information sharing attitudes. *Transportation Research Procedia*, 44:94–101.
- Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2020c). The value of information sharing for platform-based collaborative vehicle routing. *Transportation Research Part E: Logistics and Transportation Review*, 141:102011.
- Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2022a). An auction-based multi-agent system for the pickup and delivery problem with autonomous vehicles and alternative locations. Submitted.

- Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2022b). Strategic bidding in decentralized collaborative vehicle routing. Submitted.
- Los, J., Spaan, M. T. J., and Negenborn, R. R. (2018). Fleet management for pickup and delivery problems with multiple locations and preferences. In *Proceedings of the 6th International Conference on Dynamics in Logistics (LDIC 2018)*, pages 86–94, Bremen, Germany.
- Lyu, X., Chen, H., Wang, N., and Yang, Z. (2019). A multi-round exchange mechanism for carrier collaboration in less than truckload transportation. *Transportation Research Part E: Logistics and Transportation Review*, 129:38–59.
- Máhr, T., Srouf, J., de Weerd, M., and Zuidwijk, R. (2010). Can agents measure up? A comparative study of an agent-based and on-line optimization approach for a drayage problem with uncertainty. *Transportation Research Part C: Emerging Technologies*, 18(1):99–119.
- Mes, M., van der Heijden, M., and Schuur, P. (2013). Interaction between intelligent agent strategies for real-time transportation planning. *Central European Journal of Operations Research*, 21:337–358.
- Mitrović-Minić, S., Krishnamurti, R., and Laporte, G. (2004). Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B: Methodological*, 38(8):669–685.
- Moccia, L., Cordeau, J.-F., and Laporte, G. (2012). An incremental tabu search heuristic for the generalized vehicle routing problem with time windows. *Journal of the Operational Research Society*, 63(2):232–244.
- Molenbruch, Y., Braekers, K., and Caris, A. (2017). Benefits of horizontal cooperation in dial-a-ride services. *Transportation Research Part E: Logistics and Transportation Review*, 107:97–119.
- Montoya-Torres, J. R., López Franco, J., Nieto Isaza, S., Felizzola Jiménez, H., and Herazo-Padilla, N. (2015). A literature review on the vehicle routing problem with multiple depots. *Computers & Industrial Engineering*, 79:115–129.
- Montoya-Torres, J. R., Muñoz-Villamizar, A., and Vega-Mejía, C. A. (2016). On the impact of collaborative strategies for goods delivery in city logistics. *Production Planning & Control*, 27(6):443–455.
- Myerson, R. B. and Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281.

- Ozbaygin, G., Karasan, O. E., Savelsbergh, M., and Yaman, H. (2017). A branch-and-price algorithm for the vehicle routing problem with roaming delivery locations. *Transportation Research Part B: Methodological*, 100:115–137.
- Pan, S., Trentesaux, D., Ballot, E., and Huang, G. Q. (2019). Horizontal collaborative transport: Survey of solutions and practical implementation issues. *International Journal of Production Research*, 57:5340–5361.
- Parragh, S. N., Doerner, K. F., and Hartl, R. F. (2008). A survey on pickup and delivery problems: Part II: Transportation between pickup and delivery locations. *Journal für Betriebswirtschaft*, 58:81–117.
- Pisinger, D. and Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34(8):2403–2435.
- Pisinger, D. and Ropke, S. (2019). Large neighborhood search. In Gendreau, M. and Potvin, J.-Y., editors, *Handbook of Metaheuristics*, volume 272 of *International Series in Operations Research & Management Science*, chapter 4, pages 99–127. Springer, Cham, Switzerland.
- Pourmohammad-Zia, N., Schulte, F., and Negenborn, R. R. (2020). Platform-based platooning to connect two autonomous vehicle areas. In *Proceedings of the 23rd International Conference on Intelligent Transportation Systems (ITSC 2020)*, pages 1–6, Rhodes, Greece.
- Renaud, J., Laporte, G., and Boctor, F. F. (1996). A tabu search heuristic for the multi-depot vehicle routing problem. *Computers & Operations Research*, 23(3):229–235.
- Reyes, D., Savelsbergh, M., and Toriello, A. (2017). Vehicle routing with roaming delivery locations. *Transportation Research Part C: Emerging Technologies*, 80:71–91.
- Ritchie, H. and Roser, M. (2020). Emissions by sector. <https://ourworldindata.org/emissions-by-sector#energy-electricity-heat-and-transport-73-2>.
- Ropke, S. and Pisinger, D. (2006). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40(4):455–472.
- Savelsbergh, M. and Van Woensel, T. (2016). City logistics: Challenges and opportunities. *Transportation Science*, 50(2):579–590.
- Schulte, F., Lalla-Ruiz, E., González-Ramírez, R. G., and Voß, S. (2017). Reducing port-related empty truck emissions: A mathematical approach for truck

- appointments with collaboration. *Transportation Research Part E: Logistics and Transportation Review*, 105:195–212.
- Schulte, F., Lalla-Ruiz, E., Schwarze, S., González-Ramírez, R. G., and Voß, S. (2019). Scalable core and Shapley value allocation methods for collaborative transportation. Technical report, University of Hamburg, Hamburg, Germany.
- Terrazas, A. (2019). What you need to know about empty miles in trucking. <https://convoy.com/blog/empty-miles-in-trucking/>.
- Tilk, C., Olkis, K., and Irnich, S. (2021). The last-mile vehicle routing problem with delivery options. *OR Spectrum*. DOI: 10.1007/s00291-021-00633-0.
- Verdonck, L., Caris, A., Ramaekers, K., and Janssens, G. K. (2013). Collaborative logistics from the perspective of road transportation companies. *Transport Reviews*, 33(6):700–719.
- Vergne, J. P. (2020). Decentralized vs. distributed organization: Blockchain, machine learning and the future of the digital platform. *Organization Theory*, 1(4):1–26.
- Vidal, T., Crainic, T. G., Gendreau, M., Lahrichi, N., and Rei, W. (2012). A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, 60(3):611–624.
- Wang, X. and Kopfer, H. (2014). Collaborative transportation planning of less-than-truckload freight: A route-based request exchange mechanism. *OR Spectrum*, 36:357–380.
- Wang, X. and Kopfer, H. (2015). Rolling horizon planning for a dynamic collaborative routing problem with full-truckload pickup and delivery requests. *Flexible Services and Manufacturing Journal*, 27:509–533.
- Wooldridge, M. (2013). Intelligent agents. In Weiss, G., editor, *Multiagent Systems*, chapter 1, pages 3–50. MIT Press, Cambridge, MA.
- Xu, S. X., Huang, G. Q., and Cheng, M. (2017). Truthful, budget-balanced bundle double auctions for carrier collaboration. *Transportation Science*, 51(4):1365–1386.
- Zeimpekis, V., Tarantilis, C. D., Giaglis, G. M., and Minis, I. (2007). *Dynamic Fleet Management: Concepts, Systems, Algorithms & Case Studies*, volume 38 of *Operations Research/Computer Science Interfaces Series*. Springer, New York, NY.

Nomenclature

The following lists provide an overview of the symbols and abbreviations used throughout this thesis, along with references to the pages where they are introduced or defined.

Symbols

Latin minuscules

<i>a</i>	Price proposed by auctioneer	99
<i>b</i>	Bid by carrier	21
<i>c</i>	Travel cost	14
<i>d</i>	Delivery location	14
<i>e</i>	Earliest service time	14
<i>f</i>	Reservation price	14
<i>g</i>	Profit for winning carrier in auction	97
<i>k</i>	Capacity	14
<i>l</i>	Latest service time	14
<i>m</i>	Maximum number of auctions per order	21
<i>n</i>	Number of stops in route	14
<i>p</i>	Pickup location	14
<i>q</i>	Load quantity	14
<i>r</i>	Release time	14
<i>s</i>	Service duration	14
<i>t</i>	Travel time	14
<i>x</i>	Decision variable representing arc traverse	117, 142
<i>y</i>	Decision variable representing vehicle load	117, 142
<i>z</i>	Decision variable representing service start time	117, 142

Latin capitals

A	Set of vehicle start locations	16
B	Bundle, set of orders	69
C	Set of carriers	12
D	Set of delivery locations	16
E	Set of edges	117, 142
G	Graph	117
I	Set of additional interim locations	16
N	Set of all locations	117, 142
O	Set of orders	12
O_A	Set of accepted orders	14
O_C	Set of orders initially assigned to carriers	12
O_R	Set of rejected orders	14
O_S	Set of initially unassigned orders	12
P	Set of pickup locations	16
R	Solution, set of routes	14
S	Set of shippers	12
V	Set of vehicles	12

Greek minuscules

α	Start location	14
β	Dissatisfaction weight	118
γ	Fine per rejected order	15
ζ	Relative cost of travel time	73
θ	Relative importance of time	35
λ	Degree of false reservation price or current cost reporting	98
π	Preference value	116
ρ	Location of stop in route	14
σ	Degree of strategic bidding	97
τ	Time horizon	14
ϕ	PCS cancellation parameter	41
ψ	NCS cancellation parameter	41
ω	End location	14

Greek capitals

Ω	Set of vehicle end locations	16
----------	------------------------------	----

Subscripts

$[\cdot]_c$	For carrier $c \in C$	14
$[\cdot]_h$	For position or element h	14
$[\cdot]_i$	For location $i \in N$	14
$[\cdot]_{ij}$	For locations $i, j \in N$	14
$[\cdot]_o$	For order $o \in O$	14
$[\cdot]_s$	For shipper $s \in S$	12

Superscripts

$[\cdot]^t$	At time t	14
$[\cdot]^u$	At time u	17
$[\cdot]^v$	For vehicle $v \in V$	14

Single-letter functions

$a(\cdot)$	Arrival time	16
$d(\cdot)$	Departure time	16
$l(\cdot)$	Load after service	16
$r(\cdot, \cdot)$	Relatedness of two orders	73
$r(\cdot, \cdot, \cdot)$	Relatedness of three orders	74
$s(\cdot)$	Service start time	16
$u(\cdot, \cdot)$	Minimal waiting time (directed)	74
$w(\cdot, \cdot)$	Minimal waiting time (undirected)	73

Multi-letter functions

CC(\cdot)	Current costs	21, 71
MC(\cdot)	Marginal costs	21, 141
MP(\cdot)	Marginal profit	39
PR(\cdot)	Profit	15
SL(\cdot)	Service level	14
TC(\cdot)	Travel costs	15

Abbreviations

ALNS	Adaptive Large Neighborhood Search	123
BKS	Best Known Solution	45
CCA	Central Combinatorial Auction	75
CGS	Contracted Gain Share	71

CPS	Current Position Sharing	34
DCPDP	Dynamic Collaborative Pickup and Delivery Problem	12
DPDP	Dynamic Pickup and Delivery Problem	37
FCS	Full Cost Sharing	35
FPS	Full Plan Sharing	34
GPDP	Generalized Pickup and Delivery Problem with Preferences	115
GVRP	Generalized Vehicle Routing Problem	114
GVRPTW	Generalized Vehicle Routing Problem with Time Windows	114
ILP	Integer Linear Program	116
LNS	Large Neighborhood Search	73
MAS	Multi-Agent System	6
MDVRP	Multi-Depot Vehicle Routing Problem	114
NCS	No Cost Sharing	35
NPS	No Position Sharing	34
PCS	Partial Cost Sharing	35
PGS	Platform Gain Share	99
PDP	Pickup and Delivery Problem	12
TCU	Travel Costs per Unit	50
TSP	Traveling Salesman Problem	21
VIP	Vehicle Interaction Percentage	45
VRP	Vehicle Routing Problem	2
VRPDO	Vehicle Routing Problem with Delivery Options	115
VRPRDL	Vehicle Routing Problem with Roaming Delivery Locations	115
WGS	Winner Gain Share	71

Summary

The freight transportation sector is one of the major contributors to air pollution. An important way to reduce emissions consists of collective route planning. Although unloaded trips and inefficient routes could not always be prevented by individual carriers, more efficient operations could often be obtained if multiple carriers collaborate by exchanging part of their shipments. The resulting vehicle mileage reductions not only lower the costs for the cooperating carriers, but also reduce emissions and decrease the level of congestion.

Achieving a successful collaboration between carriers, however, is a difficult problem. On top of the NP-hardness of the vehicle routing problem, the collaborative variants suffer from different carriers each having their individual policies, objectives, and preferences. Whereas information is generally assumed to be available in fleet management problems for individual carriers, this is problematic in collaborative cases: carriers might be hesitant to share confidential information with each other or with a platform that coordinates the cooperation. Furthermore, carriers might be more interested in increasing their own profits than in reducing the overall costs. Hence, they might try to exploit a cooperative approach.

This thesis explores how the above problems can be approached in the context of dynamic large-scale collaborative pickup and delivery problems. Earlier, centralized collaboration approaches have been proposed, but these are only applicable to problems of limited size: computation times increase with the number of orders, and hence, quick adaptations in a dynamic world will be hindered. Furthermore, information is assumed to be always available in centralized approaches, and carriers need to give up their autonomy. To avoid the last two problems, decentralized approaches with central auctions have been used, but these still suffer from scalability issues due to the role of a central auctioneer. This thesis therefore proposes a decentralized approach with local auctions: carriers can bid on transportation orders offered by individual shippers or associate carriers. Thus, no central authority is involved. The main aim of this thesis is to investigate to what extent such an auction-based multi-agent system can be applied to dynamic large-scale collaborative vehicle routing problems.

First, we investigate the value of information sharing, that is, the quality of solutions that can be obtained when different types and amounts of carrier information are known. In a computational study, we vary whether carriers' routing plans or the positions of their vehicles are made available and also whether carriers share or hide information about their marginal costs for orders within each auction. The solutions generally improve in terms of service level, travel costs, and individual profits if more carrier information is available. Cost information is important to obtain high service levels, whereas position information is most useful if only a limited number of carriers is consulted for an order. In scenarios with a small fleet or urgent orders, limited information often suffices.

Next, we analyze the potential results of large-scale carrier cooperation. In a computational study based on a real-world data set consisting of over 12000 orders, we vary the number of carriers that collaborate. Reductions in travel costs of up to 77% can be obtained with 1000 cooperating carriers. Thus, whereas previous studies only report improvements of 20–30% for small collaborations, our local auction approach allows to solve large-scale problems and exceeds the reported cost reductions by a factor of three. Furthermore, small bundles of orders can be offered within our approach to benefit from interaction effects. Although the extra computational effort is limited, bundling can improve the results with up to 13% for 1000 cooperating carriers.

A third major contribution of this thesis is the investigation of the possible advantages of strategic behaviour. Instead of reporting (estimates of) their marginal costs, carriers might bid strategically and try to increase their individual profits at the cost of the others. We analyze that incurring small losses in an auction might be acceptable for carriers since they can be compensated either by a share of the cooperation gains or by future events. A computational study shows that it is highly dependent on the distribution of the cooperation gains whether strategic bidding pays off. Hence, cheating is possible but not straightforward. Strikingly, a second-price auction system does not help in preventing strategic behaviour: the possible benefits of cheating even increase.

Finally, we extend the developed auction-based multi-agent system such that it can be applied to problem variants where multiple pickup and delivery alternatives can be specified. By this, carriers have more flexibility in choosing the most efficient options. Furthermore, users may specify their preferences for the different options. The auction approach then assists in finding a balance between constructing efficient routes and meeting the user preferences as much as possible. A computational study shows that the approach outperforms centralized heuristics on large-scale instances of 2000 orders.

In short, the proposed multi-agent approach with local auctions can contribute to enabling and stimulating collaboration between many carriers in a

dynamic world and thereby drastically reduce the overall number of driven kilometers – implying less costs, less emissions, and less congestion. The approach is rather flexible in its assumptions on information availability, it can withstand strategic behaviour under some conditions, and can successfully be applied to practically relevant problems with specific user preferences. To fully exploit the benefits of cooperation in practice, some open challenges still must be addressed: incentives for carriers to participate must be carefully designed, among others through a fair distribution of obtained collaboration profits, stronger guarantees on truthful behaviour of collaborators, and high levels of autonomy for individual carriers.

Samenvatting

Een gedecentraliseerd veilingssysteem voor grootschalige, dynamische voertuigrouteringsproblemen met meerdere vervoerders

Een aanzienlijk deel van de luchtvervuiling is toe te schrijven aan wegtransport. Gezamenlijke routeplanning is een belangrijke manier om deze uitstoot te verminderen. Afzonderlijke vervoerders kunnen niet altijd voorkomen dat ze inefficiënte routes rijden of ritten zonder lading hebben, maar vaak kunnen meerdere vervoerders wel veel efficiënter werken door onderling taken uit te wisselen. Op deze manier kunnen ze het totale aantal voertuigkilometers verminderen en zo niet alleen besparen op hun eigen kosten, maar ook zorgen voor minder uitstoot en minder drukte op de weg.

Het is echter ingewikkeld om een succesvolle samenwerking tussen vervoerders van de grond te krijgen. Het voertuigrouteringsprobleem op zich is al NP-moeilijk, maar in de variant waarin verschillende vervoerders een rol spelen moet er ook nog rekening gehouden worden met hun individuele doelen, voorkeuren en tactieken. Voor het beheer van het wagenpark van een enkele eigenaar kan meestal worden aangenomen dat alle informatie beschikbaar is, maar dat is lastiger bij samenwerking tussen meerdere vervoerders: vertrouwelijke of gevoelige informatie wordt niet zomaar met andere vervoerders of met een platform gedeeld. Bovendien zijn vervoerders meestal meer gericht op hun eigen winst dan op de totale kostenreductie die door samenwerking behaald kan worden, dus zullen ze waarschijnlijk proberen een samenwerkingssysteem uit te buiten.

In dit proefschrift wordt onderzocht hoe deze obstakels voor samenwerking aangepakt kunnen worden bij grootschalige dynamische vervoersvraagstukken waarin elke vracht een individuele ophaallocatie en bestemming heeft. Voor dit probleem zijn al verscheidene gecentraliseerde methodes voorgesteld, maar die kunnen alleen op kleinschalige problemen worden toegepast: de rekestijden nemen snel toe als het aantal taken groter wordt, waardoor onmiddellijke aanpas-

singen in een voortdurend veranderende wereld niet mogelijk zijn. Daarnaast zijn er twee problematische aannames bij gecentraliseerde methodes, namelijk dat alle informatie beschikbaar is en dat vervoerders hun autonomie opgeven. Deze twee problemen zijn er niet bij gedecentraliseerde methodes met een centrale veiling. Hierbij blijft het probleem met schaalbaarheid echter nog steeds bestaan, omdat één enkele veilingmeester verschillende combinaties van alle taken aanbiedt. In dit proefschrift wordt daarom een derde categorie voorgesteld: een gedecentraliseerde methode met kleinschalige, lokale veilingen. Vervoerders kunnen dan bieden op transporttaken die aangeboden worden door individuele verzenders of door andere vervoerders. Er is dus geen centrale alwetende tussenpersoon nodig. Het doel van dit proefschrift is om te onderzoeken in welke mate een dergelijk veilingstelsel gebruikt kan worden in grootschalige dynamische routeringsproblemen waarin meerdere vervoerders betrokken zijn.

Eerst nemen we de waarde van informatie onder de loep: we bekijken de kwaliteit van de oplossingen die gevonden worden als verschillende hoeveelheden en verschillende soorten informatie van vervoerders bekend zijn. In onze experimenten variëren we of vervoerders hun geplande routes of de actuele posities van hun voertuigen beschikbaar stellen en ook of ze informatie over de marginale kosten voor een transporttaak al dan niet openbaar maken in een veiling. Als er meer informatie beschikbaar is wordt in het algemeen het percentage gerealiseerde orders groter, de totaal afgelegde afstand kleiner en de individuele winst hoger. Informatie over de marginale kosten blijkt essentieel voor een hoog percentage ten uitvoer gebrachte orders, terwijl informatie over voertuigposities het nuttigst is als er per veiling slechts een kleine groep vervoerders geraadpleegd wordt. Als er relatief weinig voertuigen beschikbaar zijn of als vrachten snel nadat ze aangemeld zijn al opgehaald moeten worden, is een deel van de informatie vaak al voldoende om goede oplossingen te vinden.

Vervolgens analyseren we wat voor winst er behaald kan worden door grootschalige samenwerking tussen vervoerders. In een studie gebaseerd op een echte dataset met meer dan 12000 transporten variëren we het aantal vervoerders dat samenwerkt. Het totale aantal voertuigkilometers blijkt met tot wel 77 procent gereduceerd te kunnen worden als 1000 vervoerders hun opdrachten onderling uitwisselen. Terwijl eerdere onderzoeken verbeteringen van 20 à 30 procent melden op basis van kleinschalige samenwerkingsverbanden, blijkt onze aanpak met lokale veilingen goed toepasbaar op grootschalige problemen en zijn de kostenreducties daarin een factor 3 groter. Bovendien kunnen in deze aanpak kleine clusters van taken samen geveild worden zodat de wisselwerking tussen verschillende orders benut wordt. Dat kan de resultaten voor 1000 samenwerkende vervoerders met tot wel 13 procent verbeteren, terwijl de benodigde extra rekestijd beperkt blijft.

Een andere grote bijdrage van dit proefschrift is dat we onderzoeken of het voor individuele vervoerders mogelijk is om het voorgestelde veilingsysteem te slim af te zijn. Vervoerders kunnen namelijk, in plaats van hun daadwerkelijke marginale kosten te noemen, strategisch bieden om zo een grotere winst te behalen ten koste van andere vervoerders. Een klein verlies bij het verkrijgen van een order kan voor een vervoerder in principe acceptabel zijn, als dat later tenminste weer gecompenseerd wordt, ofwel door een verdeling van de winst die door de samenwerking is ontstaan, ofwel door een gunstige combinatie met orders die in de toekomst beschikbaar komen. Of het voordelig is om onjuiste biedingen te doen blijkt in grote mate af te hangen van welk deel van de samenwerkingswinst een vervoerder precies krijgt. Het is dus mogelijk het systeem te slim af te zijn, maar het kan wel complex zijn om daarvoor de juiste tactiek te vinden. Ook een Vickreyveiling kan dergelijk strategisch bieden niet voorkomen: de voordelen van strategisch gedrag nemen daarin juist toe.

Tenslotte breiden we het veilingsysteem uit zodat het op problemen met meerdere ophaal- en afleveralternatieven kan worden toegepast. Hierdoor zijn vervoerders flexibeler in hun routekeuze. Tegelijkertijd kunnen klanten hun voorkeuren voor de verschillende opties kenbaar maken. Het veilingsysteem helpt dan in het vinden van een balans tussen enerzijds zoveel mogelijk tegemoetkomen aan de wensen van de gebruikers, en anderzijds het vinden van efficiënte routes. Een experiment laten zien dat onze methode op problemen met 2000 orders beter werkt dan gecentraliseerde heuristieken.

Kortom, de in dit proefschrift ontwikkelde methode met lokale veilingen maakt het voor grote aantallen vervoerders die doorlopend nieuwe transporttaken krijgen mogelijk om efficiënt samen te werken en zo het totale aantal voertuigkilometers drastisch te verlagen. Dit leidt tot minder kosten, minder uitstoot en minder verkeersdruk. De aanpak vereist niet dat alle vertrouwelijke informatie van vervoerders beschikbaar is, en is onder bepaalde omstandigheden bestand tegen strategisch gedrag. Bovendien kan de methode toegepast worden op nieuwe probleemvarianten met specifieke gebruikersvoorkeuren. Voordat dit systeem in de praktijk toegepast kan worden, moeten de stimulansen voor vervoerders om mee te doen echter nog beter uitgedacht worden. Onder andere een eerlijke verdeling van de winst die door samenwerking ontstaat, garanties dat andere vervoerders het systeem niet te slim af kunnen zijn en een grote mate van autonomie kunnen hieraan bijdragen.

Curriculum Vitae

Johan Los was born on March 10, 1992 in Groningen, the Netherlands. He grew up in Sauwerd, which, according to W.F. Hermans, does not produce smoke but only mooing. Nevertheless, Sauwerd can be considered as an important village from a transportational viewpoint, since it has the northernmost railway interchange station of the Netherlands. Los, however, was more impressed by the complex traffic infrastructure and public transport processes observable in Amsterdam and Rotterdam. Thus, after obtaining his bachelor's and master's degrees at the University of Groningen (Los is for 25% a mathematician, one quarter a linguist, and for the rest, he is rather intelligent, mainly artificially, but slightly naturally as well), he combined his interests in optimization and multi-agent systems within a PhD project in the field of transportation at Delft University of Technology. Aside from his professional interests, he is fascinated by literature and linguistics, plays the organ, and still prefers mooing to smoke. He lives with his wife and two sons in Daltsen.

Education

2004–2010	Secondary education Gomarus College, Groningen, the Netherlands	
2010–2013	BSc Artificial Intelligence University of Groningen, the Netherlands	<i>Summa cum laude</i>
	<i>Thesis:</i> Application of Centering Theory in Coreference Resolution	
2010–2013	BSc Mathematics University of Groningen, the Netherlands	<i>Cum laude</i>
	<i>Thesis:</i> Poncelet Figures over \mathbb{Q}	

- 2013–2016 BA Dutch Language and Culture *Cum laude*
University of Groningen, the Netherlands
Thesis: Bé Nijenhuis’s Novel *The Tornado* in the Context of the Reformed Churches
- 2013–2016 MSc Artificial Intelligence *Cum laude*
University of Groningen, the Netherlands
Specializations: Multi-agent systems; logics; natural language processing
Thesis: Preference-Based Improvements on Solutions of Multi-Agent Temporal Problems by Automated Negotiation
- 2016–2021 PhD Transport Engineering and Logistics
Delft University of Technology, the Netherlands
Dissertation: Solving Large-Scale Dynamic Collaborative Vehicle Routing Problems: An Auction-Based Multi-Agent Approach

Publications

- [1] Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2022a). An auction-based multi-agent system for the pickup and delivery problem with autonomous vehicles and alternative locations. Submitted.
- [2] Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2022b). Strategic bidding in decentralized collaborative vehicle routing. Submitted.
- [3] Los, J., Schulte, F., Gansterer, M., Hartl, R. F., Spaan, M. T. J., and Negenborn, R. R. (2021). Large-scale collaborative vehicle routing. Under revision.
- [4] Los, J., Schulte, F., Gansterer, M., Hartl, R. F., Spaan, M. T. J., and Negenborn, R. R. (2020a). Decentralized combinatorial auctions for dynamic and large-scale collaborative vehicle routing. In *Proceedings of the 11th International Conference on Computational Logistics (ICCL 2020)*, pages 215–230, Enschede, the Netherlands.
- [5] Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2020b). Collaborative vehicle routing when agents have mixed information sharing attitudes. *Transportation Research Procedia*, 44:94–101.

- [6] Los, J., Schulte, F., Spaan, M. T. J., and Negenborn, R. R. (2020c). The value of information sharing for platform-based collaborative vehicle routing. *Transportation Research Part E: Logistics and Transportation Review*, 141:102011.
- [7] Los, J., Spaan, M. T. J., and Negenborn, R. R. (2018). Fleet management for pickup and delivery problems with multiple locations and preferences. In *Proceedings of the 6th International Conference on Dynamics in Logistics (LDIC 2018)*, pages 86–94, Bremen, Germany.
- [8] Los, J., Mepschen, T., and Top, J. (2018). Rational Poncelet. *International Journal of Number Theory*, 14(10):2641–2655.

Extended abstracts

- [1] Los, J., Spaan, M. T. J., Schulte, F., and Negenborn, R. R. (2019). The value of information sharing for distributed vehicle routing. In *Proceedings of the TRAIL PhD Congress*, Utrecht, the Netherlands.
- [2] Los, J., Spaan, M. T. J., Schulte, F., and Negenborn, R. R. (2018). Autonomous vehicle routing: Optimizing individual pickups and deliveries in a multi-agent system with auctions. In *Proceedings of the TRAIL PhD Congress*, Utrecht, the Netherlands.
- [3] Los, J., Spaan, M. T. J., and Negenborn, R. R. (2017). Fleet management for pickup and delivery problems with multiple locations and preferences. In *Proceedings of the TRAIL PhD Congress*, Utrecht, the Netherlands.

TRAIL Thesis Series

The following list contains the most recent dissertations in the TRAIL Thesis Series. For a complete overview of more than 275 titles see the TRAIL website: www.rsTRAIL.nl.

The TRAIL Thesis Series is a series of the Netherlands TRAIL Research School on transport, infrastructure and logistics.

Los, J., *Solving Large-Scale Dynamic Collaborative Vehicle Routing Problems: An Auction-Based Multi-Agent Approach*, T2021/26, November 2021, TRAIL Thesis Series, the Netherlands

Khakdaman, M., *On the Demand for Flexible and Responsive Freight Transportation Services*, T2021/25, September 2021, TRAIL Thesis Series, the Netherlands

Wierbos, M.J., *Macroscopic Characteristics of Bicycle Traffic Flow: A bird's-eye view of cycling*, T2021/24, September 2021, TRAIL Thesis Series, the Netherlands

Qu, W., *Synchronization Control of Perturbed Passenger and Freight Operations*, T2021/23, July 2021, TRAIL Thesis Series, the Netherlands

Nguyễn, T.T., *Highway Traffic Congestion Patterns: Feature Extraction and Pattern Retrieval*, T2021/22, July 2021, TRAIL Thesis Series, the Netherlands

Pudāne, B., *Time Use and Travel Behaviour with Automated Vehicles*, T2021/21, July 2021, TRAIL Thesis Series, the Netherlands

Gent, P. van, *Your Car Knows Best*, T2021/20, July 2021, TRAIL Thesis Series, the Netherlands

Wang, Y., *Modeling Human Spatial Behavior through Big Mobility Data*, T2021/19, June 2021, TRAIL Thesis Series, the Netherlands

Coevering, P. van de, *The Interplay between Land Use, Travel Behaviour and Attitudes: A Quest for Causality*, T2021/18, June 2021, TRAIL Thesis Series, the Netherlands

- Landman, R.L., *Operational Control Solutions for Traffic Management on a Network Level*, T2021/17, June 2021, TRAIL Thesis Series, the Netherlands
- Zomer, L.-B., *Unravelling Urban Wayfinding: Studies on the development of spatial knowledge, activity patterns, and route dynamics of cyclists*, T2021/16, May 2021, TRAIL Thesis Series, the Netherlands
- Núñez Velasco, J.P., *Should I Stop or Should I Cross? Interactions between vulnerable road users and automated vehicles*, T2021/15, May 2021, TRAIL Thesis Series, the Netherlands
- Duivenvoorden, C.W.A.E., *Speed Up to Safe Interactions: The effects of intersection design and road users' behaviour on the interaction between cyclists and car drivers*, T2021/14, April 2021, TRAIL Thesis Series, the Netherlands
- Nagalur Subraveti, H.H.S., *Lane-Specific Traffic Flow Control*, T2021/13, March 2021, TRAIL Thesis Series, the Netherlands
- Beirigo, B.A., *Dynamic Fleet Management for Autonomous Vehicles: Learning- and optimization-based strategies*, T2021/12, March 2021, TRAIL Thesis Series, the Netherlands
- Zhang, B., *Taking Back the Wheel: Transition of Control from Automated Cars and Trucks to Manual Driving*, T2021/11, February 2021, TRAIL Thesis Series, the Netherlands
- Boelhouwer, A., *Exploring, Developing and Evaluating In-Car HMI to Support Appropriate Use of Automated Cars*, T2021/10, January 2021, TRAIL Thesis Series, the Netherlands
- Li, X., *Development of an Integrated Analytical Model to Predict the Wet Collapse Pressure of Flexible Risers*, T2021/9, February 2021, TRAIL Thesis Series, the Netherlands
- Li, Z., *Surface Crack Growth in Metallic Pipes Reinforced with Composite Repair System*, T2021/8, January 2021, TRAIL Thesis Series, the Netherlands
- Gavrilidou, A., *Cyclists in Motion: From data collection to behavioural models*, T2021/7, February 2021, TRAIL Thesis Series, the Netherlands
- Methorst, R., *Exploring the Pedestrians Realm: An overview of insights needed for developing a generative system approach to walkability*, T2021/6, February 2021, TRAIL Thesis Series, the Netherlands
- Walker, F., *To Trust or Not to Trust? Assessment and calibration of driver trust in automated vehicles*, T2021/5, February 2021, TRAIL Thesis Series, the Netherlands

- Schneider, F., *Spatial Activity-Travel Patterns of Cyclists*, T2021/4, February 2021, TRAIL Thesis Series, the Netherlands
- Madadi, B., *Design and Optimization of Road Networks for Automated Vehicles*, T2021/3, January 2021, TRAIL Thesis Series, the Netherlands
- Krabbenborg, L.D.M., *Tradable Credits for Congestion Management: support/reject?*, T2021/2, January 2021, TRAIL Thesis Series, the Netherlands
- Castelein, R.B., *Accommodating Cold Logistics Chains in Seaport Clusters: The development of the reefer container market and its implications for logistics and policy*, T2021/1, January 2021, TRAIL Thesis Series, the Netherlands
- Huang, B., *The Influence of Positive Interventions on Cycling*, T2020/20, December 2020, TRAIL Thesis Series, the Netherlands
- Xiao, L., *Cooperative Adaptive Cruise Control Vehicles on Highways: Modelling and Traffic Flow Characteristics*, T2020/19, December 2020, TRAIL Thesis Series, the Netherlands
- Polinder, G.J., *New Models and Applications for Railway Timetabling*, T2020/18, December 2020, TRAIL Thesis Series, the Netherlands