

Real-time Co-planning in Synchromodal Transport Networks using Model Predictive Control

Larsen, R.B.

DOI

[10.4233/uuid:6b36a004-0b37-417c-a77a-ab423a19cbb9](https://doi.org/10.4233/uuid:6b36a004-0b37-417c-a77a-ab423a19cbb9)

Publication date

2022

Document Version

Final published version

Citation (APA)

Larsen, R. B. (2022). *Real-time Co-planning in Synchromodal Transport Networks using Model Predictive Control*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:6b36a004-0b37-417c-a77a-ab423a19cbb9>

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.



Summary

Under the synchromodal transport paradigm, transport providers decide how freight is transported. Thereby, real-time information on the transport system can be used to integrate the routing decisions of both freight and vehicles to utilize the transport capacity well between multiple stakeholders. This dissertation proposes co-planning, where consciously chosen information is exchanged between cooperating partners that plan individually towards shared goals. In the dissertation multiple routing methods based on model predictive control are presented. The conclusions illustrate that co-planning can contribute to make freight transport more efficient and thereby alleviate the environmental impacts.

About the Author

Rie B. Larsen has conducted her PhD research at the Department of Maritime and Transport Technology at Delft University of Technology in the Netherlands. Her aim is to use advanced control and learning techniques to develop realistic cooperation methods.

TRAIL Research School ISBN 978-90-5584-312-1




TRAIL THESIS SERIES T2022/9

Rie B. Larsen Real-time Co-planning in Synchromodal Transport Networks using Model Predictive Control

Real-time Co-planning in Synchromodal Transport Networks using Model Predictive Control

Rie B. Larsen



Real-time Co-planning in Sychromodal Transport Networks using Model Predictive Control

Rie B. Larsen



The research leading to this dissertation has received funding from the Netherlands Organisation for Scientific Research (NWO) under the project "Complexity Methods for Predictive Synchronodality"(project 439.16.120).

Real-time Co-planning in Sychromodal Transport Networks using Model Predictive Control

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof.dr.ir. T.H.J.J. van den Hagen,
chair of the Board for Doctorates
to be defended publicly on Friday 16 September 2022 at 10:00 o'clock

by

Rie Brammer LARSEN,
Master of Science in Engineering, Mechatronics
University of Southern Denmark, Denmark,
born in Aarhus, Denmark.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus

Prof.dr. R. R. Negenborn

Dr. B. Atasoy

chairperson

Technische Universiteit Delft, promotor

Technische Universiteit Delft, copromotor

Independent members:

Prof.dr.ir. T. Keviczky

Prof.dr.ir. L. A. Tavasszy

Prof.dr. C. Sys

Prof.dr. M. A. Botto

Prof.dr. S. Sacone

Technische Universiteit Delft

Technische Universiteit Delft

University of Antwerp, Belgium

University of Lisbon, Portugal

University of Genoa, Italy



TRAIL Thesis Series no. T2022/9, the Netherlands Research School TRAIL

Published and distributed by: Rie B. Larsen

ISBN 978-90-5584-312-1

Keywords: synchromodal transport, co-planning, model predictive control, vehicle and container routing.

Copyright © 2022 by Rie B. Larsen

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

Contents

1	Introduction	1
1.1	Transport developments towards synchromodality	2
1.1.1	Traditional transport	3
1.1.2	Multi-modal transport	4
1.1.3	Synchromodal transport	4
1.2	Model predictive control	5
1.3	Cooperation	6
1.4	Problem statement and research questions	7
1.5	Dissertation outline	8
2	Context and Related Literature	11
2.1	Synchromodal transport	11
2.1.1	Quantitative methods for synchromodal transport	12
2.1.2	Operational method types	13
2.1.3	Problems addressed with operational methods	18
2.2	Model predictive control	19
2.2.1	Introduction to MPC	19
2.2.2	MPC feature discussion	21
2.2.3	MPC for transport	26
2.3	Cooperation in the transport domain	29
2.3.1	Research perspectives	29
2.3.2	Information sharing and responsibility transfer	32
2.4	Conclusions	35
3	Integrating Routing Decisions for Containers and Vehicles	37
3.1	Introduction	38
3.2	Model description	41
3.3	Proposed control method	46

3.4	Simulation experiments	49
3.4.1	MPC parameters	49
3.4.2	Benchmark method	51
3.4.3	Simulation scenarios	52
3.5	Results and discussion	56
3.6	Conclusions	61
4	Real-time Co-planning for Efficient Container Transport	63
4.1	Introduction	64
4.2	From intermodal planning to synchromodal co-planning	67
4.2.1	Comparison of different transport paradigms	70
4.2.2	Impacts of planning under different paradigms	73
4.3	Co-planning between LSP and FSO	76
4.3.1	LSP optimization model	77
4.3.2	FSO optimization model	79
4.3.3	Communication	82
4.4	Numerical Experiments	84
4.4.1	Improved computation time	84
4.4.2	Experimental setup	85
4.4.3	Illustrative case study	88
4.4.4	Paradigm comparison	89
4.5	Conclusions	92
5	Learning Discrete Actions over Time	95
5.1	Introduction	96
5.2	System and problem definition	98
5.2.1	Problem definition	99
5.2.2	Stability and feasibility	101
5.2.3	Sequence selection	104
5.3	Implementation example	106
5.3.1	Benchmark controllers	106
5.3.2	System	107
5.3.3	Results	108
5.4	Conclusions	109
6	Co-planning with Learning	111
6.1	Introduction	112
6.2	Problem statement	115

6.3	Real-time centralized decision making	116
6.3.1	System dynamics	118
6.3.2	Centralized method	120
6.4	Departure learning	123
6.4.1	Learning good departure times	125
6.4.2	Evaluating the performance	128
6.5	Simulation experiments	129
6.5.1	Benchmark methods	129
6.5.2	Scenarios	130
6.5.3	Impact of the tunable parameters	132
6.5.4	Prediction horizon - T_p	133
6.5.5	Exchanged schedules - n	134
6.5.6	Learning parameters - α and β	135
6.5.7	Comparison between departure learning and the bench- mark methods	139
6.6	Conclusions	142
7	Conclusions and Future Research	145
7.1	Conclusions	145
7.2	Contributions	149
7.3	Future research	150
	Bibliography	153
	Abbreviations	171
	Curriculum vitae	173
	Summary	177
	Samenvatting	179
	TRAIL Thesis Series publications	181

Chapter 1

Introduction

Using the available resources optimally should always be the goal. This is true in most domains, and specifically in transport. Transportation accounts for around one-fifth of global carbon dioxide (CO₂) emissions with freight transport on roads responsible for 30% thereof [135]. Globally it is currently discussed how humanity can decrease both our CO₂ emissions and our negative impact on Earth. It requires overarching political visions, e.g. the European Union's *Green Deal* [25], and specific solutions.

In this dissertation, we propose methods to take operational decisions in synchromodal transport networks which emphasises the new possibilities this freight transport paradigm provides over traditional paradigms. Under the synchromodal transport paradigm, decisions on mode and route are transferred from the shipper to the transport provider. It ties in with ideas such as physical internet, smart transport and (freight) transport-as-a-service, as it increases the transport providers' flexibility. The focus in this dissertation is on how to obtain the benefits that can be gained from automated transport decisions that are integrated across stakeholders and are real-time. In this first chapter of the dissertation, we first motivate the dissertation's contributions in relation to freight transport at large, synchromodal transport specifically, and model predictive control. Then we state the dissertation's main focus and research questions together with the research method and a few core assumptions. Finally, the relations between the dissertation's chapters is discussed and the outline of the dissertation is stated.

1.1 Transport developments towards synchronomodality

Traditionally, plans in the transport sector are made hierarchically. Strategic plans and decisions such as investment in new infrastructure and tactical decisions like hiring new staff have impact on all later decisions and plans. Having a hierarchical framework for these kind of decisions can thus help model and understand their consequences. However, decisions and plans at an operational level are in the academic literature also often thought of as hierarchical. This is mentioned explicitly by few (e.g., [10]), but is mostly apparent from the boundaries of the considered problems. When containers are being routed through the transport network, the barge schedule is assumed in place (e.g., [55]). When barge departure times are decided, the available truck capacity is considered sufficient at any point of time (e.g., [33]). Figure 1.1 shows the decision hierarchy that is often assumed in transport literature. Only considering the routing of vehicles in an inland container transport network, the top of the hierarchy is the train schedule. This schedule is planned for a long period of time and is hard to change later, as freight trains typically are assigned slots on the rails in-between passenger trains, which cannot be rescheduled without causing inconvenience to many people ([49]). In the academic literature on container transport, the next level is the routing of barges. Sometimes the routing is planned as a full schedule with

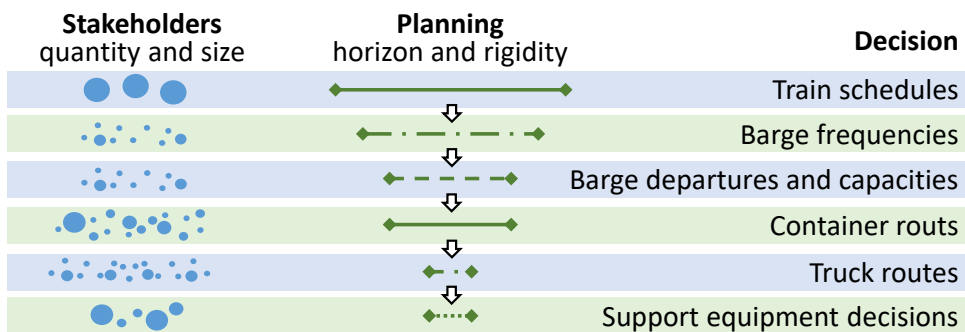


FIGURE 1.1: *The decision hierarchy that is often assumed in transport research. For each decision type, indications of how many stakeholders of which size are given together with a visual description of how long a time horizon the decision typically is fixed for, and how rigid it is (with solid lines being most rigid and dotted most flexible).*

departure times, sometimes with a service frequency. In the literature on multi-modal transport, barges and trains are often treated as the same type of mode, only varying in cost, emission and capacity and their routing decisions are thought of as a service network design problem ([152]). If only routes and departure frequencies are decided upon, planning exact departure times is the next step. This is assumed done after reliable information about which containers are to be routed, i.e. the transport requests, is available (e.g., [130]). After the transport requests are fully known, trucks are routed with specification of their departure times (e.g., [170]). At the bottom of the transport planning hierarchy are the decisions regarding cranes, stackers and other terminal equipment (e.g, [23]). This rough outline regards the primary vehicles holding the containers. Several other plans exist for support material, especially in ports. These plans are often created after the primary plans, even if rough estimates are used as limitations when creating plans at a higher level (e.g, [171]).

1.1.1 Traditional transport

When plans are made hierarchically, it is simpler to coordinate between different stakeholders than if they are integrated. However, if a plan that has been made earlier has to be changed because of external factors, reorganizing plans from the lower layers of the hierarchy becomes harder as there is no system in place for influencing plans across the levels. This complication reflects in the way shippers choose to transport their containers. In interviews and surveys ([124], [71]), it is often indicated that shippers prefer the certainty and flexibility of road transport compared to the cost savings and possible decreased environmental impact of the water and rail modes. Often, using water and rail modes require additionally the use of a road vehicle as first and/or last mile transport. When a container is transported by multiple modes, it is planned according to the different schedules, often assuming trucks and handling equipment are available as needed (e.g.,[88]). A further complication is that each mode typically is operated by a different entity, which limits the information known by each stakeholder. Giusti et al. highlights both collaboration and information flow as critical, but hard to achieve, enablers for synchronodal transport in [47]. To account for delays, buffer time can be added. This may lead to additional costs, as found in [53] where Gumuskaya et al. study the impact of the realization of stochastic transport requests on the actual cost of transport when the barge departure

times have been optimized in advance waiting the cost of all possible realizations according to their probability of occurrence. If the additional buffer-time is added to each container individually, it makes the transport slower, increases the space needed for stacking containers in multi-mode terminals and increases the use of handling equipment to reorganize stacks. Under the commonly assumed hierarchical model of planning of transport operations, there are thus many incentives for shippers towards using one mode of transport.

1.1.2 Multi-modal transport

From a resource perspective, it is often favourable to use multiple modes. Barges and short-sea ships have the advantages that very few people can move many containers using less fuel over a mode that is seldom congested. There is debate whether the environmental impact of barges is better than that of trucks, as the barge fleet generally is older and being replaced slower ([72]). It is, however, a common assumption that barge is the most sustainable inland transport vehicle for container transport, and as its societal impact is better ([159]) we adopt this assumption here. Barges and short-sea ships are limited to move on appropriate waters and are generally very slow modes. Trains move faster, overland and have many of the same advantages as barges, as few people can move many containers without congesting roads. However, trains often share tracks with passenger trains on parts of their journeys, and are as such very hard to reschedule. Maintaining a rail line is furthermore complex. Trains can be electrically driven, and as such, theoretically, able to drive on greener energy. Trucks are nearly always the fastest and most flexible mean of transport. They are however also more labour intensive and impact the traffic flow and safety in the road network negatively.

1.1.3 Synchromodal transport

To help transport providers to better utilize their resources and to simplify transport planing from a shippers perspective, synchromodal freight transport has been proposed earlier. The concept has rather emerged than been invented, and is as such not well defined. A good overview of the early literature can be found in [133], [10] and [160]. It is commonly agreed upon that for a freight transport network to be synchromodal, it must have

multiple modes, the transport requests must be without mode specifications (a-modal), and the concluded contracts must allow last minute changes. A shipper concluding a transport contract with a transport provider will thus not know in advance what modes will be used. This flexibility enables the transport provider to plan more containers to be moved by barge and train, even when there is a tight deadline, because it is possible to change the plan and transport any container by truck if need be. Synchronomodal transport takes away the assumption that operational plans can be made hierarchically, and thus emphasises that new, more flexible ways of taking decisions are needed. Synchronomodal transport needs methods for real-time planning to optimally use the available vehicles. This dissertation addresses the need for such methods, as the research on the topic is still limited and mainly consider the container-routing problem.

1.2 Model predictive control

Control theory regards all dynamic system where an input to the system produces an output from the system, and is concerned with how to choose the input so that the described output is achieved over time. It is a very mature field with application in a wide variety of fields, from manufacturing of chemicals, over manoeuvring of humanoids to information spreading. In large scale transportation outside the warehouse floor, control is often discussed at the level of control of individual vehicles (e.g., [59]), platoons of vehicles (e.g., [182] or groups of infrastructure (e.g., [154])). A few researchers have applied control to planning of freight transport over multi-modal networks ([87] among others). None of the works so far have used control methods to integrate the plans of containers with those of the vehicles that carry them.

Model predictive control is a method that combines the advantages of future predictions with those of real-time updates based on feedback ([102]). Usually, the dynamics of the system to be controlled is described by a discrete-time, state-space model for which there is a quantifiable goal and restrictions on either the state or the controllable input. From the model, the objective and the constraints, an optimization problem can be formulated to find the finite time-series of inputs that will fulfil the goal as well as possible. Typically, only the first set of inputs from this time-series is implemented before the output of the system is again measured to re-evaluate what inputs are

optimal. This way, the system is controlled towards long term goals with the ability to frequently make corrections. The objective can be described over the infinite horizon, if the difference between the truncated cost and the infinite cost depends on states in the finite horizon that is being optimized over. This is not the case in a transport system, as the future transport requests are unknown perturbations of the system that cannot be neglected and e.g., train schedules create irregular time-varying constraints. To capture the long term goals of the transport plan, predictions further into the future are needed [11]. However, large optimization problems takes a long time to solve, which in turn decreases the frequency with which decisions can be taken. It is thus not trivial to apply model predictive control to transport, as each container and each vehicle have unique features.

1.3 Cooperation

Cooperation between stakeholders is necessary in many systems in general and specifically in almost all transport systems. Some cooperation is part of a formalized hierarchy, e.g., the custom authorities dictates when an import container has passed their check and is ready to be picked up, while other is based on trust and negotiations, e.g., barge and terminal operators discuss when the barge can be assigned quay time. In the transport literature, cooperation is mostly discuss as either auctions or distributed optimization. Some also study what happens when nobody cooperate in fully decentralized systems. Auctions usually concerns cooperation between homogeneous competitors, e.g., [96]. Distributed optimization is additionally applied between complimentary organisations, e.g. in [88] organisations that routes containers through geographically non-overlapping intermodal transport networks cooperate to fulfil demand that goes across their individual networks.

Distributed optimization is also used in the literature on model predictive control to improve the joint performance of multi-agent systems in general. Here the information flow is usually considered symmetric, such that pre-defined updates of local copies of shared variables steers the negotiation towards the common best solution [168]. There is as such a need for methods that are tailored to specific stakeholders and use the information they realistically would be willing to share. In this dissertation we provide such methods for two cases. When an logistics service provider and a service

operator discuss alternatives for transport over flexible services, and when a barge operator and a truck operator cooperates on creating an efficient joint network.

1.4 Problem statement and research questions

The research presented in this dissertation aims at solving operational planning problems for synchromodal transport networks using the mindset and theory from model predictive control. The main question addressed in this dissertation is:

How can container transport realistically be planned in real-time when several different stakeholders own the vehicles?

On the way to the answer the main research question, the following sub-questions were studied:

- Q1 What is the impact of integrating decisions across the planning-hierarchy layers that concerns container and vehicle routing?
- Q2 How can operational planning under synchromodal transport take advantages of the opportunity for real-time mode-changes?
- Q3 What is the impact of stakeholders planning cooperatively at the operational level?
- Q4 How can containers and vehicles be routed cooperatively through a synchromodal network, if only traditional transport requests and their expected fulfilment are communicated?
- Q5 How can Bayesian optimization help solve a model predictive controller's mixed integer optimization problem?
- Q6 How can we bridge the information gap that comes from low communication frequency?

To answer the research questions, we first explore how to integrate container and vehicle routing in real-time assuming one entity has all information and authority to take all decisions. We hereafter consider two cases of cooperation between entities. One focuses on the value of replanning and

information, while the other provides a method to overcome the information gap. To extract more information out of limited communication, ideas inspired by Bayesian optimization are used. To ensure these ideas are compatible with model predictive control, a method to control general, switched linear systems is researched. Figure 1.2 shows how these four main bodies of work relates in the dissertation. As research is exploring new directions, the goal for each method is to get a proof-of-concept for the core ideas in the method. It is therefore assumed that infrastructure, legal framework, business models, etc., for synchromodal transport is in place, leaving the research with the operational planning problems. It is however emphasised that it should be imaginable that the involved stakeholders/companies would be interested in using the methods after they have been extended to meet all requirements for a real world application. Another fundamental assumption in the presented methods is that all events within the time horizon we plan for are accurately known and that no information about later events is available. This is chosen to highlight that model predictive control can rely on feedback to react to disturbances as an integrated part of the planning method. If stochastic methods were used as part of the model predictive control methods, it would be harder to differentiate what part of the performance under uncertainties could be attributed to the real-time aspect.

To answer question Q3, Q4, and Q6 we introduce the concept of co-planning. Co-planning is the process of two or more autonomous entities that create their individual plans with limited communication between them sharing carefully selected information while striving towards a common goal. The decentralized nature of the transport sector has to be taken into account when presenting distributed planning methods. Otherwise, the methods will never create a sector-wide impact and will only benefit large stakeholders.

1.5 Dissertation outline

This dissertation is organized as follows (see Figure 1.2). Additional information about the numerical experiments and digital versions of the presented figures from which exact data-points can be extracted are publicly available at 4TU.ReserchData. The research presented in this dissertation is in Chapter 2 positioned in the current academic literature. A comprehensive review of the state of the art of synchromodal transport is first presented. Hereafter,

the core assumptions which are used in the dissertation are discussed in relation to the literature on model predictive control. This includes a summary of the commonly used way to apply model predictive control to multi-agent systems. Finally, the chapter relates co-planning to trends in cooperative planning in the transport field.

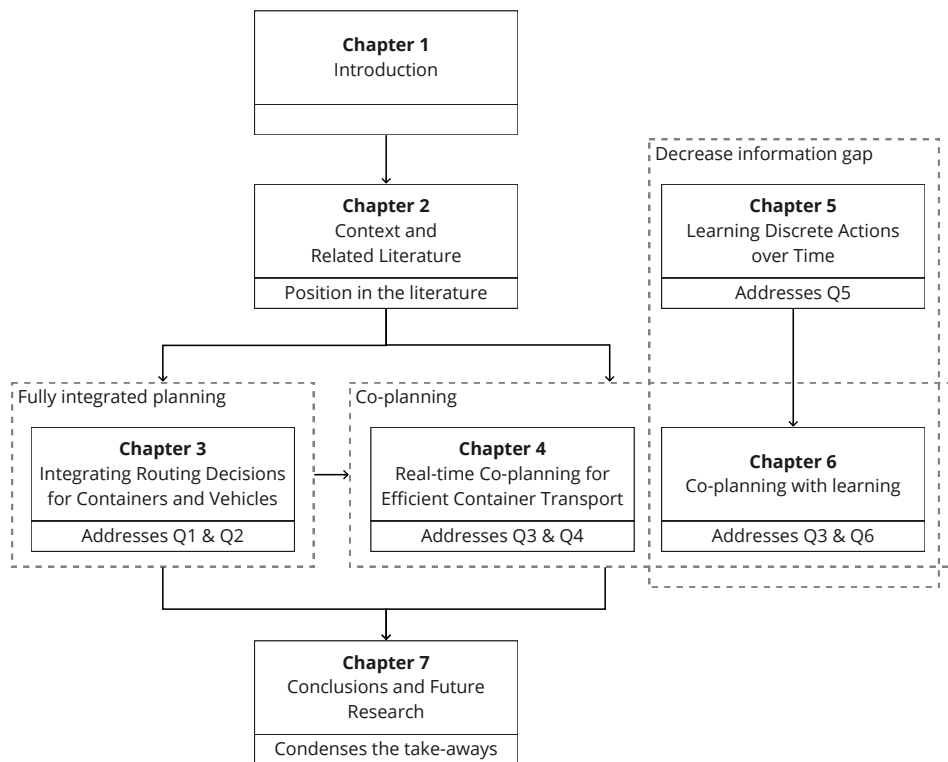


FIGURE 1.2: *Relation between the chapters of the dissertation and with the research questions.*

In Chapter 3, a real-time method for routing of containers and vehicles simultaneously is presented. The method uses model predictive control directly, and in the results it is shown how it reacts to travel time disturbances. Variations of the presented method are used in both Chapters 4 and 6.

The decentralized nature of the transport sector is introduced in Chapter 4. Here a method for co-planning between a logistics service provider and a vehicle fleet operator is presented.

Exchanging information is not always practically feasible, neither is it always desirable as competitors may be able to interpret what is considered sensitive information from the communicated data. In Chapter 5 we explore the use of ideas from Bayesian optimization to learn discrete actions over time instead of optimizing them directly as part of the model predictive control method.

The general method is in Chapter 6 modified to suit the case of co-planning between a barge operator and a truck operator. The goal is to use very limited information exchange to decrease the total cost of transport over the joint network by changing the departure plan of the barge to allow as many containers as possible to use this mode.

In the final chapter, Chapter 7, the overarching conclusions are drawn, the research question and its sub questions are answered explicitly and the key contributions of this dissertation are outlined. The chapter ends with an overview of directions for further research that are relevant in light of this dissertation.

The bibliography can be found on page 153 and a list of the used abbreviations is on page 171. Hereafter follows information about the author, a short summary of the dissertation in English and Dutch, and a list of other recently published titles in the TRAIL Thesis Series.

Chapter 2

Context and Related Literature

This chapter has three main parts. First, we present a comprehensive overview of the academic literature on synchronomodal transport. The concept is discussed and the research presented in the dissertation is positioned in the full body of decision methods developed for synchronomodal transport. Second, an introduction to model predictive control (MPC) is given and the core assumptions behind the methods presented in Chapter 3 to 6 are discussed from a control theoretical point of view. Furthermore, applications of MPC to transport problems are discussed. Third, we provide an overview of the research on cooperation in the transport and logistics domain. This chapter ends with the main conclusions and a clear positioning of this dissertation into the existing literature.

2.1 Synchronomodal transport

Synchronomodal concept is a fairly recent reconsideration of how freight transport should work. A formal definition of synchronomodal transport has never been agreed upon, but concept descriptions usually emphasise that the transport contract of the freight is without mode specification (called *a-modal*) and allows for changes in both the transport route and mode choice during the transport [133]. Beside this fundamental feature, definitions often include integration (synchronization) of plans [10] and improved service for the shippers [160]. Many early discussions of synchronomodal transport were held in business environment, were published in Dutch and/or used other names (e.g., [156], [169], and [97]).




FIGURE 2.1: *Sustainability in the synchromodal literature (indexed by Scopus by 1st December 2021). Shows how many papers mention any word starting with “sustainab” per publication year and the total number of new citations.*

Most research on synchromodal transport have an economical focus, while a few includes environmental sustainability in their definition of the concept. The review article by Giusti et al. [47] define synchromodal transport as *the provision of efficient, reliable, flexible, and sustainable services through the coordination and cooperation of stakeholders and the synchronization of operations within one or more supply chains driven by information and communication technologies*. Sustainability is, however, often seen as an implicit part or consequence of a synchromodal transport even if it is seldom part of the used definitions. This is visible in literature reviews (e.g., [147]) and metadata on published articles as seen in Figure 2.1. In this dissertation, it is defined as a-modal freight transport with real-time planning of the transport operations that seek to achieve the system-wide best solution. We leave the definition of what the best solution is open, as different definitions are suitable for different cases.

2.1.1 Quantitative methods for synchromodal transport

To fully use the flexibility that synchromodal transport offers via real-time mode changes, the applied methods should be different than those applied to e.g. intermodal transport. Therefore, we here review of the methods published specifically for synchromodal transport. On the 1st of December 2021, 99 papers on synchromodal freight transport were indexed by Scopus. Of these 58 were ICT tools, decision methods or quantitative methods to study the effects of synchromodal transport: Nine papers are in the ICT domain, proposing either communication platform structures ([12], [20],

[148]), software structures ([46], [63]) or unified data formats ([4], [64],[149], [158]). ICT considerations are largely related to bringing a concept to implementation, so we will not describe them further. Eight papers either focus on details such as how to create a network map from geographical, maritime positions [35] or discuss methods quantitatively [90]. There is a big overlap between papers that analyse consequences and those that propose decision models, as real-time mode changes require a decision model, so in the remainder of the section we focus on the 41 papers that model sychromodal transport mathematically.

The vast majority of the published decision methods are at the operational level. This is expected, as the direct difference to intermodal transport is at the operational level. The two papers that are at the strategic level both consider the location of new hubs. In [26] the hub capacity is furthermore decided upon. A binary, single commodity flow model with stochastic travel times is used to find the location and capacity of hubs that over time will be most cost beneficial when considering both the investment cost and the operation cost. In [48] similar costs are considered, but the profit from and handling capacity needed for each future piece of freight is uncertain instead of the travel time. The model here is a two stage stochastic model where the hub location decision comprises the first stage and the commodity flow during operation the second stage, whose stochastic nature is captured by scenario sampling.

Three papers consider tactical decisions. In [166] and [167] van Riessen and co-authors propose firstly to create fare classes as known from the passenger aviation industry on well established routes on sychromodal transport corridors and secondly propose a decision model for computing the price of each class under uncertain future demand. The first paper considers a the direct transfer between two locations, while the later describes the idea for a network. The perspective is from a logistics service provider with control over and access to multiple modes. The objective when establishing the fare classes is purely economical. Another paper that fits at the tactical level is [40], where no decision method is proposed, but the impact of tactical decisions on the energy consumption over a supply chain is analysed.

2.1.2 Operational method types

An overview of the 36 papers that propose a sychromodal, operational decision method is presented in Table 2.1 and 2.2. Some of them are developed

with focus on the insights their application to a use case can bring, while others are the main contribution of the research. The references marked with * are of the author of this dissertation, and the content of reference [80] marked with ** is presented in Chapter 3 of this dissertation. In Table 2.1 the main characteristics of the methods are summarised together with their goals. All methods describe the system they consider using a mathematical model. Three main types were identified: Binary, flow and state space. With binary models, we denounce models where decisions are described with binary variables, e.g., in [58] batches of containers are as a group matched with a service if the corresponding variable is 1. Next to the binary variables there are often continuous variables describing time limitations. Flow models often describe time in a similar way, but instead of binary decisions for fixed units, integer or continuous variable flows describe how much is being moved, allowing ,e.g., batches of containers to be split [130]. Finally, the work that forms part of the PhD research leading to this dissertation and a few others use state space models, where in all cases, the dynamics of the quantity of the subjects at different locations is described using a discrete-time difference equation for a finite time period. The decisions are modelled as flows between the different locations. All methods based on state-space models use a rolling horizon and are as such MPCs. It is worth to notice that the authors of [88] also have MPC based methods for intermodal transport, e.g. [87].

The operational methods published for synchronodal transport are mainly assuming one decision maker has access to all information and authority to take all decisions. In [5] and [30], there is no communication between the different agents in the systems. In the former, decisions are taken completely independently, and the resulting system behaviour is studied under synchronodal and intermodal transport paradigms. In the later, each shipper is an agent and decisions are taken sequentially using up the available transport capacity. More integrated approaches are presented in [88], where distributed optimization is used to align the operations of three homogeneous transport providers that operate in non-overlapping geographical areas. In [16] distributed optimization is achieved in hindsight by assigning extra fees to the local optima of each agent to make the sum of local optimum coincide with the global optima for the system. The work on co-planning takes outset in the restricted information which can be communicated between the agents and a clear limits to their responsibilities. It is discussed further in Chapter 6.

TABLE 2.1: *Operational methods for synchromodal transport*

Paper	Method			Objective			
	Model	Structure	Disturbance mitigation	Time	Cost	Lateness	Other
[1]	Binary			x	x		
[5]	Binary	Decentralized	replan		x		
[9]	Flow			x			
[10]	Flow			x	x		
[16]	Flow	Distributed			x		
[30]	Binary	Sequential		x			
[31]	Binary		Scenario		x		
[36]	Flow				x		Back-order
[56]	Binary		Rolling		x	x	Emmision
[58]	Binary		1)		x	x	Emmision
[57]	Binary		2)		x	x	
[70]	Flow				x		
[78]*	State space		Rolling		x	x	
[79]*	State space	Co-planning	Rolling		x	x	
[80]**	State space		Rolling		x	x	
[81]*	State space	Co-planning	Rolling		x	x	
[82]*	Binary				x	x	
[85]	Flow		Switch policy				
[88]	State space	Distributed	Rolling	x	x	x	
[92]	Binary						Quality
[105]	Binary				x	x	Emmision
[112]	State space		Rolling		x		Emmision
[110]	State space		Rolling		x		Emmision
[111]	State space		Rolling		x		Mode split
[121]	Flow		Quantify risk		x		Preference
[137]	Binary		Rolling		x		
[126]	Flow		Rolling		x		
[136]	Binary		Replan		x		
[139]	Flow				x		
[130]	Flow		Replan		x	x	
[134]	Binary		Replan	x	x		Emmision
[165]	Flow				x	x	
[164]	Flow		Replan		x	x	
[175]	Flow				x		Profit
[178]	Binary		3)		x	x	
[184]	Binary+flow		4)		x		

TABLE 2.2: *Operational methods for synchromodal transport (continued)*

Paper	Decisions						
	Quantity	Container routing	Mode choice	Cancel services	Departure times	Vehicle routing	
[1]			x		x		
[5]		x			x		Unclear
[9]		x		x			Unclear
[10]			x		x		
[16]		x					
[30]		x					
[31]			x		x		
[36]	x		x				
[56]		x					
[58]		x					
[57]	reject	x					
[70]		x		x			
[78]*		x					Truck
[79]*		x			x		Truck
[80]**		x					Truck
[81]*		x			x		Truck
[82]*		x					Truck
[85]	x		x				
[88]		x					
[92]		x					
[105]		x					
[112]		x	5)	x			
[110]		x		x			
[111]			x				
[121]		x					
[137]			x				Barge
[126]			x		x		
[136]							Truck
[139]		x					
[130]		x			x		
[134]		x					
[165]		x					
[164]		x					
[175]				x			
[178]		x					
[184]		x					Terminal

		Replanning intervals	
		Periodic	Irregular
Decisions	All	Reconsider [80]	Change [164]
	Infeasible	Update [56]	Handle exception [130]

FIGURE 2.2: *Structured replanning can be categorized regarding the regularity of the replanning and how large a part of the decisions are taken into account. The references are examples of the four categories. Each category contains a spectra.*

The real-time aspects of synchromodal transport is in many of the published methods used to allow for replanning. This can be either periodically at regular intervals (denounced rolling in this section) or when needed, e.g., because new demand needs planning for or disturbances render an existing plan infeasible. In both cases, the re-optimization of the plan can be for all decisions that are not yet started or for only those decisions that are no longer feasible. Figure 2.2 illustrates this and provides references to methods. A few papers use robust optimization methods to create plans that will remain feasible for more potential scenarios, which means less replanning will be needed. In [121] the risk of a plan becoming infeasible is quantified and used to describe the different preferences shippers may have. A policy, rather than a set of decisions is optimized in [85], which provides a way of changing actions up until departure without replanning. A few methods mitigate disturbances using multiple methods, which are marked with a number in Table 2.1, corresponding to: 1) Rolling horizon and scenario based optimization, 2) Rolling horizon, chance constrained and scenario based optimization, 3) Replan when needed and scenario based optimization, 4) Two-stage optimization with scenario based optimization in the second stage. In conclusion, replanning is the most commonly used method to address dynamic demand and react to unforeseen disturbances in operational synchromodal transport methods.

2.1.3 Problems addressed with operational methods

The goals of the published synchromodal methods are in nearly all cases to reduce the operational costs, as can be seen in Table 2.1. What this cost consists of varies greatly, both depending on the scope of the problem and how detailed the model is. In [36] synchromodal transport is considered as part of a supply chain, so here stock is part of the operational costs, while most papers focus on the transport costs, e.g., [178]. In e.g., [137] only the transit cost is considered, while e.g., [56] also consider the cost of changing modes and storing containers. Most methods consider time as part of the objective, either directly as travel time, i.e. aims at the fastest possible delivery, or as a penalty for containers that arrive too late. Some have a penalty for arriving too early, that is usually much smaller and in nature closer to an operational cost. Lateness penalties can be interpreted as a quantification of customer dissatisfaction with arrivals after the agreed due date [164]. Some methods consider additional components in their objective functions. Among those, reduced emissions is the most common goal. In most methods, the different objective components are quantified in monetary terms and summed, but a few investigate the Pareto frontier between the components (e.g., [134]). In [105] the ideal plan for achieving each component is made and the objective of the final plan is constrained to be within a margin from these solutions.

The definition of synchromodal transport covers all types of freight transport, but almost all methods discuss container transport. As seen in Table 2.2, all papers except one ([175]) consider the route of each (batch of) container(s) through a transport network or a simplified mode choice for networks where that defines the full path the container(s) will follow. The mode choice is in one paper a constraint of the optimization problem, not a decision, this is marked with 5). In the papers that regard broader logistics systems, it is also decided how many containers are to be sent. One paper from a transport provider's perspective includes the possibility to reject new transport requests ([57]). In more than half of the papers, the decisions regarding containers are integrated with decisions regarding the transporting vehicles. Mostly it is done for barges, where the route is fixed. Here some methods decide whether a given service is to be operated or cancelled while others finalize departure times. One paper integrates the container routing with the assignment of time-slots at terminals to the barges that potentially can transport them ([184]). The routing of the vehicles is only done in a few methods,

of which five are part of the work leading to this dissertation. When the vehicles are routed, their movements are considered both when they are loaded and when they are empty. Only by routing the vehicles, continuity between the services they can offer can be ensured. In methods where services can be cancelled it is not clear how the vehicles are repositioned. Integrating both container and vehicle routes increases the size of the involved optimization problems significantly, both directly by the additional (combinations of) decisions, but also indirectly as vehicle routing only adds value if containers are likely to move in both directions of each arc in the network. In the academic research on container transport often the flow only in one direction is considered (e.g., [130]). The increased computational complexity seems to be the main reason why integration of routes is not more common.

2.2 Model predictive control

The research presented in this dissertation uses model predictive control (MPC) as the core of the suggested methods. In this section, we first introduce the basics of MPC and then we position the dissertation in the MPC literature by discussing the relation between the existing knowledge and the core assumptions that apply in Chapter 3 to 6.

2.2.1 Introduction to MPC

Model predictive control is a method to use information about a system actively for feedback control [145]. The (approximate) dynamics of the system one wants to control is often known, which makes it possible to predict how the system will react to future inputs. To find the best time sequence of inputs, an optimization problem can be constructed. The objective function is a quantification of the deviation from the system's desired behaviour and the optimization variables are the controllable inputs to the system. The optimization is constrained by the model of the system's dynamics, limitations of the inputs (e.g., physical properties of the actuators) and considerations on what system behaviour is acceptable. Furthermore, it is a requirement that the starting point of the predictions in the optimization problem matches the current reality. If the whole series of optimal inputs are applied after this problem is optimized, the system is controlled using an (advanced) open-loop controller [74].




FIGURE 2.3: An MPC controller regularly gathers information and recomputes the optimal sequence of inputs for the prediction horizon. Only the first part of the plan is implemented.

The disadvantage of open-loop control is that there is no possibility for correcting the inputs, hence new information cannot be used which leads to a lower performance of the system [114]. This new information is typically measurements from the system that shows it does not behave as expected or predictions of upcoming (external) uncontrollable events. Feedback control uses this type of information to constantly align the inputs to the expectations for the system. Several different types of controllers using feedback loops exist. One of the simplest lets the input be proportional to the system measurements at any time. MPC is an advanced feedback control technique that combines the advantages of updating the inputs based on relevant information with the refined control that optimization provides.

An MPC controller solves an optimization problem as described above and starts applying the optimal sequence of inputs. After a time interval, the MPC controller receives new measurements of the system and new information about external events. This information is used to solve the optimization problem again resulting in a new sequence of optimal inputs. The MPC controller starts applying the new sequence of inputs immediately, i.e. before the end of the old sequence is implemented in full. Figure 2.3 shows this visually. The three bars represent the optimal control sequences obtained at different times with the upper one being the earliest. When researchers evaluate a controller, they typically make a simulation experiment with a finite time duration. It is important to be aware that at the end of such a simulation, parts of the optimal inputs will not yet have been applied (e.g., [110]) unless the researcher stops the MPC controller before the end of the simulation time (e.g., [94]). In reality, this is not an issue, as ending a system's

processes often is a separate task handled manually or by a dedicated controller.

A typical MPC controller optimizes the input sequence for a given time horizon at regular time intervals and uses a discrete time, state space model of the system and applies only piece-wise constant inputs. If the full state of the system is directly available (2.1)-(2.6) is the optimization problem solved by the MPC controller. The notation used here, is very common in the literature on MPC. The notation used in Chapter 3, 5, and 6 follows these conventions closely. The state that describes the system right now is x_{measured} , if it is not directly available several different methods exists to either estimate it (e.g.,[176]) or otherwise ensure the properties of the controller [107]. The state of the system is predicted over a prediction horizon that is divided into T_p timesteps and the estimated state k timesteps from now is denounced $x(k)$. $u(k)$ is the input which is applied k timesteps from now and $d(k)$ is the uncontrollable disturbances that influence the system at the same timestep. The difference equation that describes the dynamics is denoted $f(\cdot)$. The set \mathcal{X} describes the states the system is allowed to be in and \mathcal{U} contains the possible inputs. Often the state at the end of the prediction horizon $x(T_p)$ is additionally restricted to be within the set $\mathcal{X}_f \in \mathcal{X}$. The divination of the last state is furthermore quantified differently in the objective function, denoted $V(x(T_p))$. At all other timesteps, the objective usually has the same structure and depends on both the state and the inputs. For an extended introduction, we refer to [132].

$$\min_{u(0)\dots u(T_p)} \sum_{k=0}^{T_p-1} l(x(k), u(x)) + V(x(T_p)) \quad (2.1)$$

$$\text{s.t.} \quad x(0) = x_{\text{measured}} \quad (2.2)$$

$$x(k+1) = f(x(k), u(k), d(k)) \forall k \in [0 \dots T_p - 1] \quad (2.3)$$

$$x(k) \in \mathcal{X} \quad \forall k \in [0 \dots T_p] \quad (2.4)$$

$$u(k) \in \mathcal{U} \quad \forall k \in [0 \dots T_p - 1] \quad (2.5)$$

$$x(T_p) \in \mathcal{X}_f \quad (2.6)$$

2.2.2 MPC feature discussion

To apply a MPC controller to a specific system, several choices have to be made. In the following sections, we discuss the different options that are

relevant to the synchromodal transport planning problem. The choices we cover are:

1. Controller structure
2. External uncontrollable disturbances
3. System model
4. Objective function
5. End of prediction horizon

Controller structure

If one decision maker knows all available information about the system, e.g. measurements, dynamics and constraints, and has authority to take all decisions, a centralized controller can be applied. Such a controller enables full integration of the considered decisions and a system-wide performance objective. A centralized MPC will typically achieve the goal better than MPC controllers with other structures as it works with the global optimum. It is therefore a sensible choice when possible or to represent a 'best possible' benchmark controller for systems without a single decision maker.

In some cases, decisions and/or information is private to a part of the system. In these cases and if the system is too large to be practically controlled by one decision maker, multiple decision makers have to cooperate. Such cooperating controllers can either be structured in a hierarchy or work on a peer to peer basis. It is in some systems also necessary to control each sub-system individually without explicit coordination and treat the interaction with the remainder of the system as disturbances [77]. In systems with multiple decision makers, it is important to consider what can be communicated between the entities and with what frequency [99]. A commonly used structure is to have a central entity that coordinates high-level goals and leave the fine control of each subsystem to the local controller (e.g., [60]). In flat structures with peer-to-peer communication, various distributed optimization techniques can be used depending on the structure of the objective and coupling between the subsystems. Nedić and Liu provide in [117] an overview for systems where the global objective is the sum of the local objective. This is an assumption we also adopt here. The core of these methods is typically that subsystems that are connected, exchange local information about the variables that couples them. This information is either used to update their local copy of the coupling variables or to penalize the difference between the copies. This process is repeated until consensus is reached

within some margin. In a transport routing planning setting, this could be interpreted as either negotiation of quantities or updates to the transport fees (see e.g., [88]). In the distributed scenarios considered in Chapter 4 and 6 the communication-limits prevents direct application of the well-established distributed optimization methods.

External uncontrollable disturbances

MPC is a control method that updates its plans regularly. Depending on the complexity of the involved optimization problem and on the demands of the system to be controlled the update frequency can be measured in milliseconds (e.g., [120]) or days (e.g., [110]). Every time the plan is reoptimized, all available information is taken into account. This way, the MPC reacts to external disturbances after they happen besides planning for the future. In systems where some (probabilistic) information about the disturbances are available, it can be included into the optimization problem. In (2.1)-(2.6) deterministic information about the disturbance enters as the time varying parameter $d(k)$. Probabilistic information about the disturbance can be used to ensure the likelihood of an acceptable performance is high by creating more conservative plans. Different robust optimization techniques can be applied, e.g. chance constrained MPC [52], tube MPC [76], and, scenario based MPC [143]. Common for these approaches is that while the resulting implementation is more likely to remain feasible, the implementation cost is suboptimal for the deterministic system. When the disturbances are expected to be large, the methods may also limit the actions that can be taken so much that the performed task is significantly slower or even infeasible. For these reasons, it is important to consider if constraint-violations are detrimental for the system or can be alleviated by other means before blindly choosing a robust approach. In freight transport systems, it is usually possible to do nothing without endangering the system. The use of robust optimization approaches will thus mainly change the obtained cost/profit. In this dissertation, deterministic MPC is used, to highlight the real-time aspect of synchromodal transport and how it differs from earlier transport paradigms.

System model

MPC usually relies on a state-space model of the system as described Section 2.2.1. Transport systems, on the other hand, are often modelled as a series of binary choices with associated continuous time variables (e.g., [57]).

Such models can also be applied in a receding horizon fashion, but the feasibility and stability proofs that come with MPC do not necessarily apply in that case. As those proofs are of less importance in freight transport system where abstaining from action is usually a possibility, it is important to consciously chose model type. The advantage of the traditional binary models is the direct representation of the unique attributes of the involved subjects (e.g. pick-up time windows for containers) and the ability to plan for an unrestricted time horizon. The disadvantage is a very high computation time that often results in methods that are useful for few, large batches of containers that must follow the exact same trajectory (e.g., [33]). For intercontinental shipping, this is a reasonable assumption, but for hinterland transport network, it may be questionable. State-space models use discrete-time difference equations to describe the dynamics of a system. This limits how precise time-durations can be described to predefined intervals. If the timesteps in a transport planning MPC is 10 minutes long, a transport time of 31 minutes will appear the same as one that is 40 minutes. Using smaller intervals between timesteps increase the precision but also the computational complexity. The tradeoff between modelled time and computational complexity extends to the prediction horizon. The longer a prediction horizon, the more things are planned in advance and the longer a process can take and still be described in full. However, a long prediction horizon also increases the computation complexity. State-space models suffers as binary models from slow computation times if all the unique features of a transport system is modelled. To achieve real-time control, aggregation is needed. A state-space method is however more suitable for splitting batches of containers, which makes the resulting plans able to utilize vehicles/modes with low capacity better.

Objective function

The objective of an optimization problem can be a reflection of the true cost of implementing the computed solution. It is however sometimes favourable to also use the objective function to guide the solution. One example of this is the use of soft constraints, where deviations from the desired values (i.e. constraint violations) are allowed at a penalty. If violation of a specific constraint is not detrimental to the system, but may cause the optimization problem to become infeasible it is usually possible to formulate the constraint as a violation penalty in the objective function instead of a strict constraint. In

a transport system, it is often sensible to use soft constraints for the timely delivery of freight [165].

If transport cost or profit is used as objective function, the solution found is often not unique as the structure often is linear. When the optima is not unique, an MPC may change plans more frequently, because it finds a different optimal solution at each timestep. To decrease the number of optima, small variations in the costs can be used to promote the preferred optima. In a transport system, it is often an advantage to act earlier. This can be promoted by adding a very small additional cost on actions that is proportional to how far into the planning horizon the action is to take place.

The objective function, furthermore, influences how easy it is to solve the problem to optimality. For continuous variable problems, the fastest is to have a convex search space (typically form having only linear constraints) and a convex or linear objective function. This type of problems can be solved efficiently with off-the-shelves solvers. A good introduction to methods is provided in [14]. If there are integer or binary variables in the optimization problem, the computation quickly becomes intractable [86]. Convex relaxations and other heuristics may be necessary to use integer and binary variables as part of the MPC optimization problem. Sufficiently frequent updates and careful rounding of the results can limit the sub-optimality that may occur from using a convex relaxation [141].

End of prediction horizon

MPC implements actions that are found to be part of the optimal sequence of actions over a finite horizon. The optimization problem that describes what action sequence is best is highly effected by how long this horizon is and what conditions and costs apply at the end. In MPC for general systems, it is very common to ensure the recursive feasibility of the optimization problem by ensuring the state of the system at the end of the prediction horizon is within a set of states for which a feasible (typically very sub-optimal) solution always exist. This means that a control-law that ensures feasibility exist for all states in the set \mathcal{X}_f in (2.6). If this control-law stabilizes the system, the stability of the system under MPC feedback control is also ensured [102]. Since the planning horizon is truncated, the objective of the optimization problem is also truncated. To mimic the infinite horizon cost that a sequence of actions will cause, the term $V(x(T_p))$ is often used

in the objective function (2.1). This value can also play a role in stabilizing the system and ensuring the recursive feasibility of the MPC [11]. In a transport system, the resulting cost from the implemented actions is often a very important performance parameter, while recursive feasibility can be ensured by model choice. However, choosing the function $V(\cdot)$ is nontrivial in transport systems where different modes are available, some of which are scheduled. There exist methods for time-varying constraints that are periodic (e.g, [50]), but they are often impractical for transport systems as the combination of schedules seldom is periodic over a reasonable horizon. Some applications of MPC in transport systems disregard the possibility for transport by different modes and use the trucking cost as an estimate of the infinite horizon cost [87]. This pushes containers in the system towards their destination, but cannot be meaningfully applied to vehicles that has no final destination. Whether or not special conditions and cost apply to the last state in the prediction horizon, the consequences of an action must be described in the optimization problem in order for the solution to truly consider that action. If a barge journey takes two days and the system is modelled as a simple delay system, the prediction horizon must be longer than two days to consider the barge option truly.

2.2.3 MPC for transport

With suitable choices for the MPC, it can be a powerful methodology to develop real-time, operational transport planning methods. Since most other freight transport paradigms do not allow for real-time mode changes, control is rarely applied to operational planning problems outside synchromodal transport. Noticeable exceptions are the paradigms of physical internet [157] and smart logistics [41], but neither paradigm seem to have caught the attention of control researchers.

The main challenges when applying MPC to centralized operational transport planning problems are the discrete nature of the problem, the simultaneous need for detailed time discretization and long prediction horizons, and the uniqueness of each container/vehicle/control subject. Therefore, MPC is more commonly used to control, e.g., vehicle movements, as the control variables are continuous, affects the system relatively quickly and come from a few unique actuators. To understand what kind of problems MPC is applied to in freight transport, we categorized the relevant papers indexed on Scopus the 15th of March 2022. 303 papers were found that in their title,

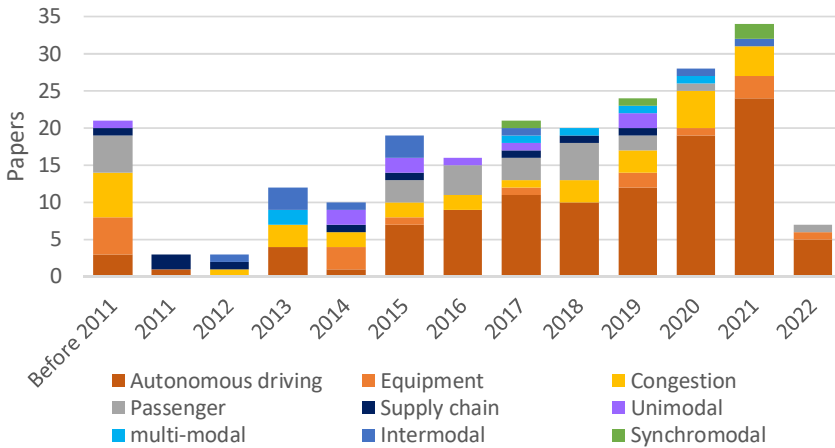


FIGURE 2.4: *Topics covered by the papers indexed in Scopus 15th March 2022 on MPC and freight transport*

abstract or keywords contained at least one word from each of the following three groups:

1. Model predictive control, MPC, predictive control
2. *modal, container, freight, planning
3. transport, transportation

Of the 303 papers, 218 were on relevant topics. Among the irrelevant topics were handling of radiative waste and tunnel ventilation systems. In Figure 2.4 the relevant papers were categorized based on the type of problems they describe. From the graph it is clear that there is an increasing interest in applying MPC to transport problems, but the majority of the papers (in total 104) and the main growth is in the category of *Autonomous driving*. All methods that decide the speed and/or steering angle of a vehicle falls within this category. Many regard autonomous, augmented or guided driving of cars, either individually [83] or cooperating [119], but e.g., trains are also considered [61]. A related category is *Equipment*, where the movement and maintenance of freight transport equipment is considered. Most methods in this category maximize the number of containers a crane can move without swaying dangerously [118]. Maintenance of railways [153] and optimal energy storage and use in hybrid vehicles [140] are also common topics.

Another growing category for application of MPC is that of traffic control, usually aimed at decreasing congestion on highways [123] or in urban

areas [150]. In these problems, the vehicles in the system are usually aggregated into homogeneous flows, which helps overcome the challenges of discreteness and uniqueness. MPC is sometimes used in *Passenger transport*, but no trends over time can be observed. The main difference between passenger and freight transport is the ability of passengers to independently change their plans and move. Research on passenger transport is thus more about providing possibilities than deciding movements [89]. MPC is applied in a wide spectra of passenger transport applications, from on-demand taxi services [131], over evacuation planning [173] to metro-scheduling[89].

The application of MPC to freight transport problems does also not show any trend over time. It has been used in *supply and distribution chains* [142]) to control flows of material and in *unimodal transport* such as railways and pipe systems to e.g., decide fleet size [106] and multi-product transport flows [45]. In these problems, the control subjects are again often aggregated into a few homogeneous flows and the discretization of time matches the necessary planning horizon well. The applications of MPC to *Multi-modal transport* mainly focus on transshipments between modes [21] and terminal decisions, such as berth allocation [18]. How MPC is applied varies, but generally the discretization of time again matches the necessary planning horizon well. When MPC is applied to *intermodal transport* and *sychromodal transport*, similar problems are considered.

Sychromodal transport is a relatively new term where the main difference to intermodal transport is the ability to change mode after the transport has started. This means intermodal transport researchers with a control perspective often consider sychromodal transport systems without naming them as such. It is often part of the MPC setting to reconsider the mode choice for a container on the next leg in its journey up until departure. Another difference between intermodal and sychromodal transport is that the freight must be containerized in intermodal transport. However, MPC has not been applied to other types of sychromodal transport than container transport. The methods for intermodal transport considers sometimes the terminal operations in greater details than the ones for sychromodal transport.




FIGURE 2.5: Overview of the published literature on cooperation in freight transport logistics. The numbers indicates how many papers indicates they consider the corresponding keywords.

2.3 Cooperation in the transport domain

There are usually multiple organizations involved in transporting a container or piece of freight. For unimodal, domestic transport, they are usually a supplier, a buyer and a transport provider. The shipper can be either the supplier, the buyer or with responsibility transfer during transport. If the transport is international or make use of multiple modes, the number of organizations that must cooperate to make the transport successfully increases. The efficiency of a container transport system is thus highly depended on how well the involved actors cooperate.

2.3.1 Research perspectives

The cooperation in transport systems has been studied from many different perspectives. Figure 2.5 shows how many published papers take the perspective of ICT solutions, cooperation platforms, that uses qualitative, optimization or other quantitative methods. The numbers in the Venn diagram shows how many papers mention words related to the five perspectives in

the title, abstract or keywords. Since only metadata was used, there may be inaccuracies in the results, but as the number of processed papers is large, the findings give a good indication of what is being researched in relation to cooperation in freight transport. In total the metadata from 1775 papers that were indexed in Scopus 17th of March 2022 were used. The papers were selected as they contained at least one of the terms in each of the following two categories:

1. cooperat*, multi-agent, multi agent, distributed, decentralized, collaborative, platform
2. *modal transport*, container transport*, freight transport*, pick up and delivery, pick-up and delivery, freight logistics

This search resulted in 2144 papers written in English. Of these 293 were irrelevant. Irrelevant topics included handling of radiative materials, transporter proteins, propulsion systems and papers specific to passenger transport. 76 papers were reviews and thus also excluded from the results. Of the analysed papers, 177 were case studies and 161 did not report any of the considered perspectives.

The search terms favour papers with a quantitative perspective on cooperation in freight transport, and this is reflected in the results, where very few papers do not mention terms related to optimization or other quantitative methods. Some papers focus on facilitating cooperation and thus lie in the overlap between ICT solutions and platforms. The majority of papers with the ICT perspective use optimization or other quantitative methods, which indicates that there in papers with a stronger focus on practical implementation is a firm connection to the theoretical sides of the field. Around the same number of papers that indicate a quantitative method also indicate that an optimization method is used. This is in line with the detailed analysis of the research on synchronomodal transport, where most papers that propose quantitative methods are optimization-based (see Section 2.1.1).

Looking at the papers that take a quantitative and optimization perspective on cooperation in freight transport, the distribution between methods and cooperation structure can be seen in Figure 2.6. As for Figure 2.5, the information is based on related terms found in title, abstract or key-words. In total, 1250 paper's metadata was analysed. Most papers use optimization methods or do not indicate what method they use. Game theory is applied more often than learning techniques, and they are hardly ever combined. Transport and logistics are traditional operations research fields, so

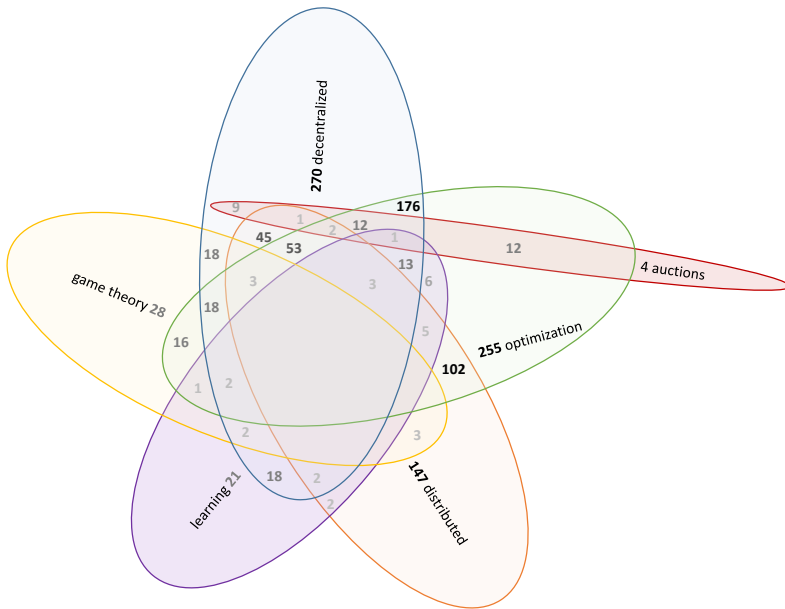


FIGURE 2.6: Overview of the published literature that propose cooperation methods or uses quantitative methods to model cooperation in freight transport and logistics. The numbers indicates how many papers consider the corresponding keywords.

the strong focus on optimization falls in line with the tradition of the field. Using learning and game theory methods to describe the cooperation between actors in freight transport may bring new insights and is as such an open gap in the literature.

In a few papers, the cooperation between partners are governed by bidding mechanisms in auctions. These are typically considering competing transport providers that can minimize their costs by distributing their orders between them such that each competitor gets to deliver freight that creates economically cheaper routes than what can be achieved if each competitor only transports the orders they themselves have received from shippers. In [96], information sharing between the actors in such systems is discussed and the value hereof is analysed. Information in such systems is mostly shared with a central platform, but sometimes directly between competitors. Commonly, the pick-up and delivery location and time of (a subset of) shipping orders are shared and bids in the form of a price of insertion to the competitors common routes are used to distribute the shipping orders. The transfer of responsibility is often clear in the presented methods and follow

the orders that are being auctioned (e.g., [28]).

In papers considering distributed and decentralized transport systems, a large variety of approaches exist, and a general idea cannot be sketched. The definitions on distributed and decentralized vary a little between sub-fields, but generally in decentralized systems there is no cooperation between the participating actors. In the strict definition, there is also no information exchanged and each actor is responsible only for their own operations. Decentralized systems are thus often used to model the aggregated behaviour of a multi-agent system (e.g., [2]).

In distributed systems, there is always information exchanged and the responsibility is shared, partially or fully. With enough information sharing and cooperation, a distributed methods can create solutions similar to centralized methods [19]. However, in a competitive environment, the actors will often be reluctant to join collaborations where they have to share sensitive information and hand over responsibility and decision-power over core decision.

2.3.2 Information sharing and responsibility transfer

The variety in methodology used and problems researched in the field of cooperation in freight transport and logistic is evident from the ten most cited papers which provides a quantitative method or model. Table 2.3 provides an overview. Half of these papers analyses a logistics or transport situation, while the other half propose strategic or operational decision methods. As such, most of the papers concern strategic problems. This trend is supported by the findings of other reviews on freight transport (e.g., [152]) and is in contrast to the trend in the literature specific to synchromodal transport. As argued in Section 2.1.1, the focus on operational methods for synchromodal transport is well in line with the real-time flexibilities of that transport paradigm.

The most cited papers on operational cooperation in freight transport discuss the possibilities for cooperation and provides three very different cooperation structures. In [42] a very realistic approach is described in the text, where multiple companies are independent agents and within each company, the headquarter and each truck driver are independent, but fully cooperating agents. However, in the explicit mathematical model and the experiments, only one company is considered and the negotiation between competitors

TABLE 2.3: *Summary of the most cited papers that provide quantitative models or methods for cooperation in transport*

Citations	Actors	Level	Method Structure	Information Ex-change	Responsibility Transfer	Actor Objective	Decisions
[113] 371	factories, retailers, markets	analysis (strategic)	central	non, noncooperative	[116]	-	manufacturing & purchase quantities
[75] 218	carriers	strategic	central	request info	full to platform	profit sharing, transport cost	truck routes
[155] 176	truck companies, platform	strategic	two-level	1) hub details 2) traffic flows	non	transport cost	1) hubs size & location 2) truck routes
[73] 146	1 supplier, 1 buyer	analysis (strategic)	central	complete	non	transport & inventory cost	quantities & frequency
[179] 140	shippers, platform	analysis (strategic)	central platform form	volume of order	orders to platform	profit sharing, efficient transport	truck routing
[42] 138	shipping companies, trucks	operational	text: multi-agent; exp.: auction	orders & free capacity	order to auction winner	freight' and trucks' route	text: announce orders & free capacity; exp.: match freight to trucks & truck routes
[66] 134	vehicles	analysis (case study)	multi-agent Monte Carlo	non	non	-	task & vehicle type
[51] 129	retailers, 1 warehouse	analysis (strategic)	central	complete	full	transport & inventory cost, customer waiting time	stock level for order and emergency order
[15] 128	factories, warehouses, retailers	operational	semi-decentralized	different degrees of forecasts	full to entity-type's decision maker	transport cost, stock level	manufacturing & purchase quantities
[6] 124	ad-hoc & back-up drivers	operational	central platform form	available tasks, ad-hoc drivers' route & time	full to platform form	transport distance	match tasks to drivers

is thus omitted. The information shared is limited to the that of specific orders that should be distributed and that of periods where the trucks run empty. The responsibility follows the transport orders, however, in the explicit mathematical model it is on the headquarter agent as the driver agents must accept rewarded orders, even if it lowers their individual profit.

In [15] a logistics system with factories, warehouses and retailers is considered. Here full cooperation between all actors of a specific type is assumed and the the method becomes distributed between three different layers. This simplifies the communication and leads to better coordinated plans, but may in may cases be an unrealistic setting. The paper researches the impact of sharing different kind of information in the form of forecasts between the layers and finds that forecasts on consumer demand or retailer replenishment leads to the most economical production and distribution. The final paper providing an operational method, [6], gives all decision power to a central platform. Ad-hock drivers can submit periods where they are available to the platform, leaving all responsibility within the periods to the platform.

Assuming a central platform, or analysing the system from a central perspective is a very common approach to cooperation also at the strategic level. All relevant information and a complete transfer of responsibilities happens in [75] and [51]. In [179] the shippers can decide what order to transfer to the platform, that once transferred take over all information and responsibility, and in [73] the two actors keep their responsibilities, but can take decisions based on complete information from the other party. No information is communicated between the actors in [113] and they remain independent, competing entities. The behaviour of each actor is assumed to follow the dynamics given by Nash in [116] and the system is analysed centrally. The decisions in [155] are described from a central perspective with a two level model. The information flow between the two levels is very limited and well-established and each level keeps responsible for the decisions corresponding to the level. The method could thus be easily modified to a co-planning method between independent entities. Finally, [66] analyses the aggregated traffic flows when each agent takes independent decisions drawn from probability functions that depend on the agent types. There is thus no communication nor responsibility transferred.

2.4 Conclusions

For synchromodal transport to be efficient, operational level, real-time planning methods that can integrate plans of several different stakeholders are needed. That is agreed upon in the academic literature. Some operational methods with a real-time aspect have been published and a few regards the cooperation between different stakeholders, however there is still several open questions and the methods presented in this dissertation comprise a substantial part of the research on real-time planning methods that integrate different types of decisions.

In the broader scope, cooperation in freight transport systems has been researched from many different perspectives. In the field there is a strong tradition for using optimization, both when analysing a transport system and as part of e.g., decision support methods. Even-though, cooperation between stakeholders is frequently researched, the proposed methods and models often rely on a central, trustworthy platform that has access to all relevant information and has the authority to take all decision. It is often not highlighted in the published literature what assumptions are taken regarding information sharing and allocation of responsibilities. Awareness of these aspects of distributed systems will help making the applicability to practise of the proposed methods and results more apparent and it will lead to insights of higher relevance. In this dissertation, we propose the term *co-planning* to describe planning methods that seek to achieve efficient transport by integrating plans based on realistic information exchange and where each participating party maintain their responsibilities and thus authorities as much as possible.

Model predictive control (MPC) has earlier been applied to a variety of transport problems, but seldom for freight transport. It is likely related to the origin of the method in the control of individual systems and the strong theoretical background present for systems with continuous variables. The main challenges when applying MPC to operational transport planning problems are the discrete nature of the problem, the simultaneous need for detailed time discretization and long prediction horizons, and the uniqueness of each container/vehicle/control subject. With careful consideration of the application, MPC can provide a strong foundation for real-time planning methods for synchromodal transport and is thus the base of the methods proposed in this dissertation.

Chapter 3

Integrating Routing Decisions for Containers and Vehicles

Among the most important pre-requisites of synchromodal transport are the integration of decisions and smart, real-time decision making, as outlined in Chapter 2, Section 2.1. In this chapter, we show how decisions on vehicle and container routes can be integrated in a real-time planning method. The proposed method is an application of model predictive control (MPC), that uses the assumptions discussed in Chapter 2, Section 2.2. To address Research Question Q1, the integrated vehicle and container routing method is used to demonstrate the impact of integrating vehicle and container routing on the usage of the vehicle fleet. Mode-choice is in the method re-evaluated at every decision moment, and the method thus provides an answer to Research Question Q2.

The core of this chapter is published in *European Journal of Control*¹. The scope of the problem is introduced and positioned in the literature in Section 3.1. Then, in Section 3.2, the synchromodal transport system is described as state space dynamics and in Section 3.3 an MPC method is proposed to perform real-time planning. The numerical experiments are presented in Section 3.4 and their results are discussed in Section 3.5. The conclusions are drawn up in Section 3.6.

¹Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. Model predictive control for simultaneous planning of container and vehicle routes. *European Journal of Control*, volume 57, pages 273–283, 2021

3.1 Introduction

The longer it takes from the moment a plan is made until it is implemented, the larger is the risk that something unexpected will happen. In container transport this unexpected event could be extreme weather delaying a barge, or an extra control check by customs delaying a container. Traditionally, such events are handled manually, hence making direct truck transport the easiest mode to use. Truck transport is however often the least environmentally friendly and the most man-hour consuming mode of transport. From an environmental, societal, and economical perspective it is therefore desirable to use other modes of transport such as rail and water instead. Multi-modal, intermodal and synchromodal transport, as well as the physical internet, supply chain logistics, etc. are all concepts that enable such a shift away from simplistic solutions and towards overall efficient solutions.

The shift towards an overall efficient approach creates new challenges on both the strategic, network design level, the tactical, flow scheduling level and on the operational, specific movements level. For synchromodal transport it can be argued that the time-horizon of decisions taken on the tactical level becomes closer to the time-horizon of decisions on the operational level [133], when the flows and services can be re-planned based on online information. A key enabler for this change is the concept of a-modal bookings where the service of transport is bought instead of a slot on a specific connection. This lets the transport supplier decide which modes and which vehicles are used to fulfil a specific transport order, and allows the supplier to change this decision during the execution of the transport.

It is however not enough to change decisions in real-time, it is also necessary to take good decisions. Smart planning, disruption handling, dynamic switching, and demand aggregation are in [147] identified to be the four categories of necessary actions to obtain synchromodality. Real-time switching and integrated planning are also in the literature review [54] found to be among the 8 most important properties of synchromodality. It is thus agreed upon that the success of synchromodal transport is closely linked to the ability to switch plans when disturbances occur and the ability to plan container moves and equipment use simultaneously.

This chapter presents such a framework which chooses modality and routes for containers simultaneously with routes and loading/unloading actions for trucks in real-time. The framework uses model predictive control (MPC) to take decisions based on the latest available information with a

conscious trade-off between the cost of transport for the containers and the utilization rates of the vehicles.

In the current literature on transport planning under uncertainty, transport suppliers create vehicle routes based on estimations of the demand. In [175], a static plan that accommodates uncertain future events is created by optimizing over different scenarios, while they in [137] are accommodated by using the probabilistic knowledge of the future events in an approximate dynamic programming method. Another approach is to plan truck flows and barge and train schedules ahead of time based on an assumed demand and handle undercapacity during implementation with expensive ad hoc alternatives (e.g., [161] and [10]).

In the literature there are very few attempts that directly plan container and vehicle routes simultaneously at the operational level. The authors of [160] state that “the flexibility in transportation routes may be used in conjunction with the operational fleet deployment problem. This creates new and more complex optimisation challenges”, but the statement is not explored further. A planning model that besides container routes also decides if a specific service is operated or not is presented in [175]. The services are however not routed, which for a scenario with more import than export will lead to overcapacity of empty vehicles on the import side. In other words, the need for vehicles performing round-trips is not considered. In the container route planning model presented in [134], trucks are likewise modelled as links between locations which for a given time can be used or not. It is here taken into account that trucks may not always be available, but the model does not route the trucks. In [130], import containers, trucks, trains and barges are scheduled simultaneously by solving a mixed integer optimization problem. However, all vehicles, including trucks, have pre-determined routes and thus only the departure times are decided. In contrast, this chapter routes the trucks and handles both import and export containers.

Both container and vehicle planning problems have separately been studied extensively in the literature for several different transport systems. In [152], a comprehensive overview of the Operations Research planning models used in multi-modal, intermodal, and synchromodal transport can be found. To route containers through a synchromodal network, [105] finds the k shortest paths through a network where barges and trains depart according to a schedule. This framework does not reconsider decisions on future actions automatically, but the ability to do so when disruptions occur is discussed. In [87], last minute decisions are used to route commodity flows online over

a network with scheduled barge and train services, assuming truck capacity is infinite and instantly available. In [139], a similar problem is addressed by learning a preferred policy with Approximate Dynamic Programming. To obtain higher utilization rates of vehicles, the literature on dynamic vehicle routing problems combine pre-defined pick-up and delivery appointments in the most efficient way ([129]). Most papers in this category do not relate themselves to intermodal or synchromodal transport. Some accommodate transshipments in their models (e.g., [38] and [17]) and cover thereby some of the challenges of intermodal transport planning.

The ability to change decisions during transport without confirmations from shippers as well as the increasing volumes to be transported motivate the use of control methods in container transport problems. Model predictive control (MPC) has already been used to address the container routing problem, but has not yet been used to integrate the planning of container and truck routes. In aforementioned [87], receding horizon control is used to plan the container flows in a hinterland network, but in contrast to this chapter, they only consider import and assume trucks are available when needed. In [88], that model is extended to the distributed case, where the geographical network is divided into non-overlapping regions served by different cooperating stakeholders. They consider commodity flows between multiple origins and destinations, but still assume trucks to be instantly available when needed. The container routing problem is furthermore solved distributed in [34] in an MPC-like framework. Trucks are here considered instantly available and mainly used for last-mile transport.

MPC has also been used for planning and execution of related problems. It has been used to coordinate supply to demand in different supply chains (e.g. [125], [172], [98] and [62]). These models generally treat transport as a known input delay, without considering modes and timetables. [3] and [111] employed MPC to improve efficiency inside container terminals. The former considers equipment as queues, and is only suitable for small geographical areas, as it does not consider the advantages of handling containers based on their geographical location. The latter considers trucks to be instantaneously available.




FIGURE 3.1: *Example network. Circles 1-9 and icons 10-12 are nodes of the system. Green and yellow lines are long and short distance truck networks, dashed lines indicate time dependent connections (connections to scheduled services) and red lines show the connections between network nodes and their adjacent virtual destination nodes.*

3.2 Model description

The transport network is modelled as a continuous state, discrete time, state-space commodity flow model of a hinterland network. The network is described by an undirected graph, like depicted in Figure 3.1, where the nodes represent locations where containers are transferred between modes, locations where containers or trucks are stored or parked for longer periods of time, or scheduled services with high capacity. The arcs represent truck routes between physical locations or (un)loading actions for scheduled services. Vehicles and containers are modelled on separate networks that are coupled by the constraint that containers can only flow on a directed arc if there is at least the same number of trucks flowing on the same arc. If one of the nodes is a train or barge node (a scheduled service), no trucks are required. The main features of the model are:

- Demand is modelled as containers available to the network and needed from the network. Unsatisfied demand is penalized. At all timesteps the demand is fully known over the planning horizon.

- Commodity flows are considered to be continuous variables. This simplifies the model and can capture the desired level of accuracy, see [110].
- Unscheduled vehicles, with trucks as example, are also modelled as continuous variable flows. This again allows for balance between model complexity and accuracy.
- Each scheduled service is modelled separately. Two trains serving the same route are modelled as two nodes.
- A limited number of containers can be (un)loaded to trucks at a given node and a limited number can be loaded to and from the scheduled services at any given time.
- Trucks can wait at a node to be unloaded at a later time or drive through with its load.
- Travel times and capacity limits are known for the planning horizon at all timesteps.
- Terminal operating hours, truck drivers resting hours, and predictable travel time delays due to peak hours are not considered.

The model is an extension of the model presented in [78]. The current model ensures recursive feasibility of the MPC even when truck travel-times are uncertain. Trucks can here wait at nodes or drive through nodes without unloading and loading containers. This ensures that the capacity for unloading and loading trucks is not exceeded if delays cause multiple trucks to arrive at the same time. It furthermore brings the model closer to reality.

The model supports multi-commodity flows for both import and export. The demand profiles at the destinations are created based on time windows for each single container, but as commodity flows are considered, one container of a certain commodity can replace another (similar to the assumption in [87]). In [110] and [93] it is shown how this classification can be used to keep track on due dates and expiration dates. Trucks are modelled in the same fashion as containers, allowing to distinguish different kinds of vehicles. Each truck network includes a free parking node that represents the trucks that are not being used in the network but are available to the network.

The travel time and capacity limits can vary over time as time dependent parameters. This way, e.g., expected congestions can be modelled as time dependent increased travel times, and lower stacking height on barges due to high water levels can be modelled as time dependent decreased capacity. To simplify the notation they are used without a time indication in the model. It is assumed that when a travel is started, it is also fulfilled. In other words,

no decisions can be taken when a truck or container is on an arc. The scheduled services (barge and train) are modelled as nodes with time dependent arc capacities that correspond to the timetable of the respective connection. When the scheduled service is at a terminal, it has a predetermined time slot to unload and hereafter a predetermined time slot to load before it departs according to schedule.

The mathematical description of the transport network is kept general, while the specifications of the network used as example can be found in Section 3.4 and Figure 3.1.

The state x_i of each node $i \in \mathcal{N}$ in the system at every time step k is given by:

$$x_i(k) = \begin{bmatrix} x_i^c(k) \\ x_i^v(k) \\ u_{i,m_1}^h(k) \\ \vdots \\ u_{i,m_{n_v}}^h(k) \\ v_i^h(k) \end{bmatrix}, \quad (3.1)$$

where $x_i^c(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the quantity in number of containers of each of the n_c different commodities stacked at node i and $x_i^v(k) \in \mathbb{R}_{\geq 0}^{n_v}$ is the quantity of each of the n_v different truck types parked at node i . In this chapter superscripts are used to distinguish variables with similar functions, while subscripts are used for indexing the variables. Notice that most variables are vectors such that different commodities are represented by different elements in the vector. It is for simplicity assumed that all containers are of the same size and that all truck types can transport one container. However, these assumptions can be overcome by introducing additional commodities for containers of different sizes and vehicle capacities different than 1.

The vector $u_{i,m}^h(k)$ is the amount in containers of each commodity that are on the way to node i by a truck of type m at time step k . It is necessary to keep a record of the containers that are on the way to node i but have not yet arrived, since each arc in the truck network is associated with a travel time τ_{ji} that acts as a delay. Formally, $u_{i,m_1}^h(k) = [u_{j_i,m_1}(k-1)^T \dots u_{j_i,m_1}(k - \tau_{ji})^T \dots u_{j_i,m_1}(k-1)^T \dots u_{j_i,m_1}(k - \tau_{j_i})^T]^T$, $\{j \dots j'\} = \mathcal{I}_i$, $\{m_1 \dots m_{n_v}\} = [1, n_v]$, where $u_{j_i,m}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the volume in containers of each commodity that leave node j at time step k on the arc to node i using truck type m . The

set \mathcal{T}_i contains all nodes with a truck connection to node i . Likewise, $v_i^h(k) = [v_{j_i}(k-1)^T \dots v_{j_i}(k-\tau_{j_i})^T \dots v_{j'_i}(k-1)^T \dots v_{j'_i}(k-\tau_{j'_i})^T]^T$, $\{j \dots j'\} = \mathcal{T}_i$ is the amount of trucks of the different types that are on the way to node i at time step k . Here, $v_{ji}(k) \in \mathbb{R}_{\geq 0}^{n_v}$ is the amount of trucks that leave node j towards i at time step k .

The demand is modelled on virtual destination nodes $d \in \mathcal{D}$ that are adjacent to network nodes. The virtual destination nodes are copies of the network nodes, which instead of modelling the container flows model the satisfaction and accumulation of new demand. It is thus possible for a container to arrive at the network node corresponding to its destination before it is used to satisfy the demand at the virtual destination node. The arc between a virtual destination node and its adjacent network node has unlimited capacity and zero travel time, letting demand being satisfied unrestricted as soon as containers arrive at the network node. The unsatisfied demand (both available and needed containers) at the virtual destination nodes is penalized, while containers stacked at the network node waiting for demand to satisfy are only accumulating storage costs and taking up stack space. We say that node i has outgoing demand when i is the origin of the commodity and that node i has incoming demand when i is the destination. The virtual destination nodes have different dynamics than the nodes in the network, namely

$$x_i^d(k+1) = x_i^d(k) - u_{di}(k) - u_{id}(k) + d_i(k), \quad (3.2)$$

where $x_i^d(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the amount of incoming and outgoing demand in containers of each commodity at time step k . Both incoming and outgoing demand are modelled as positive values, since the commodities are defined based on destination. The variable $u_{id}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the containers that were available at network node i that are used to satisfy the incoming demand at time step k , and likewise, $u_{di}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the containers used to satisfy the outgoing demand. Demand satisfaction can be postponed (hence the integral dynamics), and the new demands $d_i(k) \in \mathbb{R}_{\geq 0}^{n_c}$, that can be satisfied from time step k , act as disturbances to the system and are thus not controllable.

The remaining nodes in the network are described as in (3.1) and have the same dynamics. For describing the dynamics three sets are defined for each node i : \mathcal{T}_i as introduced earlier, \mathcal{S}_i and \mathcal{D}_i . The set \mathcal{S}_i contains all nodes to which i is linked via a time-dependent arc connection. If node i is a scheduled service, \mathcal{S}_i contains the terminals it serves, and if node i is a terminal, \mathcal{S}_i contains the scheduled services that depart from here. Notice

that if i is a scheduled service $\mathcal{T}_i = \emptyset$. Likewise \mathcal{D}_i contains the adjacent destination node for node i . This set contains maximum one element. The dynamics of $x_i^c(k)$ is

$$\begin{aligned} x_i^c(k+1) = & x_i^c(k) + \sum_{m \in [1, n_v]} \sum_{j \in \mathcal{T}_i} (u_{ji,m}(k - \tau_{ji}) - u_{ij,m}(k)) \\ & + \sum_{s \in \mathcal{S}_i} (u_{si}(k) - u_{is}(k)) + \sum_{d \in \mathcal{D}_i} (u_{di}(k) - u_{id}(k)), \end{aligned} \quad (3.3)$$

where the control action $u_{is}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the containers moved from node i over a time-dependent connection to node s . If node i is a barge, $u_{is}(k)$ is unloading containers at terminal s . $u_{si}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the reverse movement.

As there are no scheduled services nor demand in the truck network the dynamics hereof is given by:

$$x_i^v(k+1) = x_i^v(k) + \sum_{j \in \mathcal{T}_i} (v_{ji}(k - \tau_{ji}) - v_{ij}(k)). \quad (3.4)$$

The two networks are connected by the constraint that containers cannot be moved without a truck if they are transported on a truck-arc.

$$\sum_{m \in [1, n_v]} \mathbf{1}_{n_c} u_{ij,m}(k) \leq \mathbf{1}_{n_c} v_{ij}(k) \quad \forall j \in \mathcal{T}_i. \quad (3.5)$$

The bold $\mathbf{1}_a = \{1\}^a$ is a row vector of size a with all ones.

The network is furthermore constrained by capacities:

$$\mathbf{1}_{n_c} x_i^c(k) \leq c_i^c \quad (3.6)$$

$$x_i^v(k) \leq c_i^v \quad (3.7)$$

$$-c_i^m \leq \sum_{m \in [1, n_v]} \mathbf{1}_{n_c} \sum_{j \in \mathcal{T}_i} \text{abs}(u_{ji,m}(k - \tau_{ji}) - u_{ij}(k)) \leq c_i^m \quad (3.8)$$

$$\mathbf{1}_{n_c} u_{si}(k) \leq c_{si}(k), \quad s \in \mathcal{S}_i \quad (3.9)$$

$$\mathbf{1}_{n_c} u_{is}(k) \leq c_{is}(k), \quad s \in \mathcal{S}_i, \quad (3.10)$$

where, at location i , the scalar c_i^c is the maximum number of containers that can be stored, $c_i^v \in \mathbb{R}_{\geq 0}^{n_v}$ is the maximum number of vehicles of each kind that can be parked ((3.7) is to be satisfied element wise). The notation $\text{abs}(\cdot)$ is the element-wise absolute value of a vector. The scalar c_i^m is the maximum number of containers that can be moved to and from trucks within one time

step at location i . Notice that this constraint is the crane capacity and thus does not effect containers that remain on the same truck. Trucks can thus drive through nodes without limitations. This is different from the model presented in [78]. The schedules of the barge and train connections are implemented by the time varying crane speeds $c_{si}(k)$ and $c_{is}(k)$. To illustrate, assume i is a barge and s is a terminal. When the barge is at the terminal and can be unloaded $c_{si}(k) = 0$ and $c_{is}(k) \neq 0$, and when the barge can be loaded $c_{si}(k) \neq 0$ and $c_{is}(k) = 0$, otherwise $c_{si}(k) = 0$ and $c_{is}(k) = 0$.

3.3 Proposed control method

To achieve an efficient execution of container transport and truck routing that can adapt to delays online, a convex MPC is proposed. The control variables are, for all $i \in \mathcal{N}$, the amount of departing trucks and the containers they bring, $v_{ij}(k), \forall j \in \mathcal{T}_i$ and $u_{ij,m}(k), \forall j \in \mathcal{T}_i$, the quantity to load and unload for scheduled services, $u_{ij}(k), \forall j \in \mathcal{S}_i$, and the amount of demand to satisfy $u_{di}(k)$ and $u_{id}(k), d \in \mathcal{D}_i$.

The proposed control model is based on [78] and extended, such that trucks can arrive to a node and continue driving without unloading the container it carries. Trucks with containers are furthermore able to wait at a node until the crane is available to unload them, this is modelled as a road leading back to the same node the next timestep, hence $\tau_{ii} = 1 \forall i \in \mathcal{N}$. This is a more realistic assumption than what was used in [78].

It is assumed that the controller has an accurate model for the dynamics of the transport system, and access to accurate information of the state of the global system every ΔT minute. Furthermore, a prediction of the future demand is assumed available to the controller. At each time $t = i\Delta T, i \in \mathbb{N}$ the controller gets up to date information and uses it to find the sequence of decisions that will minimize a cost function over a prediction horizon T_p . Only the decisions that require an action at this timestep $t = i\Delta T$ are implemented, and when $t = (i + 1)\Delta T$, the process starts over.

The dynamics presented in Section 3.2 is known by the controller, but since only trucks that load or unload containers require crane movements, the decision variable $z_{i,m}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is introduced to represent the containers departing at timestep k from node i on the same vehicle which they arrived with and have not been unloaded from. This way (3.8) can be formulated as a convex constraint. Only crane movements are restricted and bare a cost.

$$\begin{aligned} \min_U \sum_{k=0}^{T_p} \left(\sum_{i \in \mathcal{N}} \left(M_i^c x_i^c(k) + M_i^v x_i^v(k) + \sum_{j \in \mathcal{T}_i} M_{ij}^t v_{ij}(k) \right. \right. \\ \left. \left. + \sum_{m \in [1, n_v]} M_i^l \left(\sum_{j \in \mathcal{T}_i} (u_{ij,m}(k) + u_{ji,m}(k - \tau_{ji})) - 2z_{i,m}(k) \right) \right. \right. \\ \left. \left. + \sum_{s \in \mathcal{S}_i \cap \mathcal{N}_i} (M_i^s (u_{si}(k) + u_{is}(k))) \right) + \sum_{i \in \mathcal{D}} (x_i^d(k))^T M_i^d x_i^d(k) \right) \quad (3.11) \end{aligned}$$

$$\text{s.t. (3.2) - (3.7), (3.9), (3.10)} \quad \forall i \in \mathcal{N}, \quad \forall k \in [0, T_p - 1] \quad (3.12)$$

$$z_{i,m}(k) \leq \sum_{j \in \mathcal{T}_i} u_{ij,m}(k) \quad \forall i \in \mathcal{N}, \quad \forall m \in [1, n_v], \quad \forall k \in [0, T_p - 1] \quad (3.13)$$

$$z_{i,m}(k) \leq \sum_{j \in \mathcal{T}_i} u_{ji,m}(k - \tau_{ji}) \quad \forall i \in \mathcal{N}, \quad \forall m \in [1, n_v], \quad \forall k \in [0, T_p - 1] \quad (3.14)$$

$$\sum_{m \in [1, n_v]} \sum_{j \in \mathcal{T}_i} (u_{ij,m}(k) + u_{ji,m}(k - \tau_{ji})) - 2z_{i,m}(k) \leq c_i^m \quad \forall i \in \mathcal{N} \quad (3.15)$$

$$v_{ij}(k) = 0 \quad \forall i \in \mathcal{N}, \quad \forall j \in \mathcal{T}_i, \quad \forall k > T_p - \tau_{ij} \quad (3.16)$$

$$x_i(k=0) = \tilde{x}_i(t) \quad \forall i \in \mathcal{N} \quad (3.17)$$

Containers arriving and leaving on the same trucks do not. The subscript $m \in [1, n_v]$ denotes the different vehicle types.

The cost to be optimized by the MPC is the total cost of transporting the containers. It is assumed that the transport provider has pre-approved all incoming orders, which means that the deadline and payment from the shipper for each container is fixed. The planning tool should thus minimize the cost the transport provider needs to pay to fulfil the accepted orders, namely storing of containers, (un)loading of vehicles, slots on scheduled services, movement of trucks and parking of trucks. It is assumed that there is a central planner that can decide which plan will be followed. To evaluate what the best sequence of decisions is, the MPC controller solves the optimization problem (3.11)-(3.17), where the measured state (3.1) for node i at time t is denoted by $\tilde{x}_i(t)$. The decision vector U contains all inputs $u_{ij,m}(k)$, $v_{ij}(k)$, $u_{id}(k)$ and $u_{di}(k)$ for all $i \in \mathcal{N}$ and $k \in [0, T_p - 1]$. The time-invariant weight M_i^c is the cost of storing a container at node i , while M_i^v is the cost of parking

a truck. M_{ij}^t is the cost of a truck journey from i to j and M_i^l is the cost associated with moving a container from a stack to a truck or vice versa. Moving a container to or from a scheduled service has the cost M_i^s , which is only paid at the terminals. Transport by scheduled service is paid per container per time step as the container storage cost M_i^c . The cost of unsatisfied demand is a quadratic term scaled by M_d^i , which lets less delays be significantly cheaper than more delays.

The optimal decisions are constrained by the dynamics of the transport system by (3.12). (3.13)-(3.14) ensure that for each type of truck, only containers that are not unloaded or loaded are counted as such. The total (un)loading actions cannot exceed the available crane capacity by (3.15). By (3.17), the initial conditions of the optimization problem corresponds to the current, measured state of the transport system.

Typically, MPC ensures recursive feasibility of the optimization problem and stability of the controlled system by special constraints and costs at the end of the prediction horizon [101]. The synchromodal transport system described in this chapter is inherently marginally stable and recursively feasible, but as the actions taken within the prediction horizon will effect the state of the system in the future and thus the long-term (infinity) cost, considerations regarding the two concepts are important. The methods to address these challenges often impose conservatism that will cause underutilization of the scheduled services in the synchromodal transport problem, see ,e.g., [24]. A way to address the long-term cost of the MPC problems, when no formulation of the expected infinity costs and constraints exist, is to use a long prediction horizon, see, e.g., [11] or [29]. The current literature on this assumes different symmetric cost functions around a reference point (here the global zeros-state) that lies in the interior of the feasible set. If the transport cost is formulated based on absolute numbers and the reference point is set to be a vector of very small positive numbers instead the origin, then the assumptions hold and only the time-varying constraints prevent a calculation of the necessary length of the prediction horizon. To ensure the controller sees the consequences of its decisions, only trucks that will arrive within the prediction horizon are allowed to depart (3.16).

3.4 Simulation experiments

To evaluate the potential benefits of simultaneous routing of containers and trucks, simulation experiments of hinterland transport scenarios have been carried out. The experiments are performed both with the planning method presented in Section 3.3 that determines container and truck routes simultaneously and with a benchmark method that considers truck capacity to be infinite and instantly available. To focus on the added value of simultaneous routing, the benchmark method is an MPC-based method that has the same parameters and constraints as the proposed method except for the cost structure and assumptions related to the movement of empty trucks. In each experiment, the applied control method decides the routing over 600 timesteps. The simulations were performed in Matlab with Yalmip [95] and Gurobi.

In this section, first the parameters of the MPC are discussed, then the benchmark method is introduced followed by descriptions of the hinterland transport scenarios.

3.4.1 MPC parameters

The choice of costs and prediction horizon has significant impact on the MPC's resulting control since the MPC's prediction horizon is finite and without estimations of the infinity cost. For the presented results, the proportional costs shown in Table 3.1 are used. They are chosen to reflect the expenses from a system-wide perspective. To encourage movement and capture the cost of unnecessary crane-moves at small stacks, the costs of stacking containers and parking trucks are fairly high except at the central stack (node 6) and the parking lots (node 8 and 9), respectively. In the literature this cost is often either disregarded, e.g. in [130] or very low, e.g. [87]

TABLE 3.1: *Costs parameters*

$M_i^v = 1 \cdot \mathbf{1}_{n_v} \forall i \in [1, 7]$	$M_i^v = 0 \cdot \mathbf{1}_{n_v} \forall i \in [8, 9]$
$M_i^c = 1.2 \cdot \mathbf{1}_{n_c} \forall i \in [1, 7] \setminus \{6\}$	$M_6^c = 0.12 \cdot \mathbf{1}_{n_c}$
$M_{11}^c = 1.2 \cdot \mathbf{1}_{n_c}$	$M_{12}^c = 1.6 \cdot \mathbf{1}_{n_c}$
$M_i^l = 3 \cdot \mathbf{1}_{n_c} \forall i \in [1, 7]$	$M_i^s = 3 \cdot \mathbf{1}_{n_c} \forall i \in [1, 5]$
$M_{ij}^l = \tau_{ij} \cdot 3 \cdot \mathbf{1}_{n_c} \forall i, j \in [1, 9]$	$M_3^d = 30$
$M_5^d = 30$	$M_{10}^d = 30$

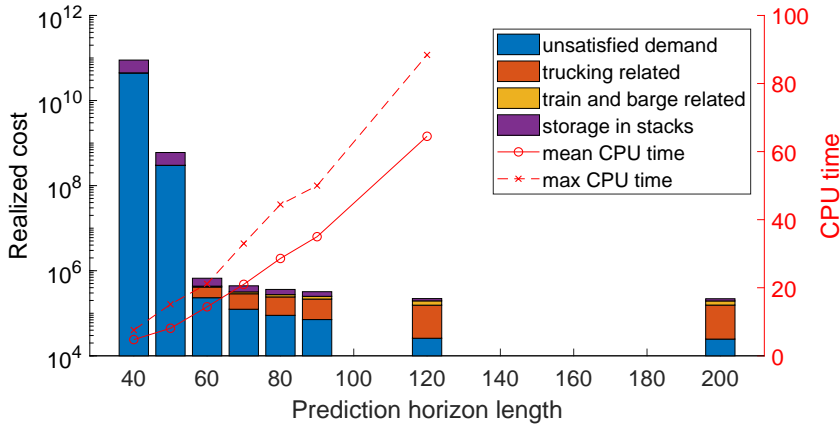


FIGURE 3.2: Comparison of realized cost and computation time in seconds per timestep for different prediction horizon lengths for scenario P1 with unbalanced demand.

where it is less than 0.1% of the hourly transport cost by barge. The results presented in Section 3.5 are simulated based on an update rate of $\Delta T = 15$ minutes and a prediction horizon $T_p = 80$ timesteps. In the remainder of this chapter, time is measured in timesteps, not minutes, as the model does not consider opening hours.

Different prediction horizon lengths allows the MPC to take different events into consideration. The longer the horizon is, the better overview over available connections the MPC will have. However, as increasing the prediction horizon length leads to increased computation time, a trade-off has to be established. It is generally advisable to choose prediction horizons long enough that the MPC can foresee both departure and arrival of the scheduled services at all times and truck roundtrips from parking node over container origin to container destination. For one of the scenarios that will be introduced in Section 3.4.3 (PIU with deterministic travel time) the realized costs when using simultaneous planning is shown for different prediction horizon lengths in Figure 3.2. The bar diagram shows that for smaller prediction horizons ($T_p = 40$ and $T_p = 50$) the MPC cannot foresee the benefit of sending an empty truck from its initial parking spot at node 9 to pick up an import container in node 7, as the delivery of that container in the inland terminals will lie outside the predicted future. When $T_p = 50$ the MPC can however predict the delivery of an import container in virtual demand node 5, if an empty truck is send from node 5 to pick up said container. For longer

$$\min_U \sum_{k=0}^{T_p} \left(\sum_{i \in \mathcal{N}_c} \left(M_i^c x_i^c(k) + \sum_{j \in \mathcal{T}_i} \left(\sum_{m \in [1, n_c]} M_{ij}^t u_{ij,m}(k) + M_i^l v_{ij}(k) \right) \right) \right. \\ \left. + \sum_{s \in \mathcal{S}_i \cap \mathcal{N}_i^s} (M_i^s (u_{si}(k) + u_{is}(k))) \right) + \sum_{i \in \mathcal{D}} (x_i^d(k))^T M_i^d x_i^d(k) \quad (3.18)$$

s.t. (3.12), (3.13), (3.14), (3.15), (3.16), (3.17) (3.19)

prediction horizons, where the MPC can foresee truck roundtrips to all destinations, the cost still decreases when the prediction horizon gets longer. The portion of the cost that is used on transport compared to the cost of unsatisfied demand also increases with increased prediction horizons. To reduce computation time, upcoming simulation experiments are limited to $T_p = 80$.

3.4.2 Benchmark method

To illustrate the impact of performing simultaneous planning, the proposed method is compared to a benchmark method that assumes trucks are instantaneously available and only optimizes the container routes in the hinterland network. The benchmark method is an MPC controller with the same update rate and prediction horizon as the proposed method. This ensures that the differences in the results obtained by the two methods only show the impact of considering container and truck routes simultaneously compared to assuming trucks instantaneously available. For the same reason, the constraints of the proposed method are used in the MPC problem of the benchmark method. Instantly available trucks are implemented in the benchmark method by ensuring sufficient capacity is available at all times in all nodes and by assigning the travel cost M_{ij}^t to the transported container instead of the truck. In the literature, it is common to assign travel costs this way (e.g. [130]). Furthermore, the handling cost M_i^l in the benchmark model is charged per departing truck to discourage the movement of empty trucks. The benchmark MPC solves thus the optimization problem (3.18)-(3.19).

TABLE 3.2: *Travel times on truck networks in time steps*

	End node								
	1	2	3	4	5	6	7	8	9
1	1	15	27	1	25	1	1	1	30
2	15	1	12	25	20	14	16	-	20
3	27	12	1	23	10	28	28	-	5
4	1	15	25	1	23	1	1	1	30
5	25	20	10	23	1	23	23	-	5
6	1	14	28	1	23	1	1	1	30
7	1	16	28	1	23	1	1	1	30
8	1	-	-	1	-	1	1	1	-
9	30	20	5	30	5	30	30	-	1

3.4.3 Simulation scenarios

The hinterland transport network used for the simulation experiments can be seen in Figure 3.1. It consists of three virtual destinations: one adjacent to ship connections, and two adjacent to inland terminals, from where last-mile delivery and pick-up are assumed to be arranged. The ships arrive and depart according to a predetermined schedule. The network has a barge and a train connection with fixed schedules. In the port area (between node 1,4,6,7,8) port vehicles transport the containers (yellow network), while long distance trucks are responsible for the remaining routes (green network). In this example the two truck networks are not overlapping, but the proposed planning method is able to address overlaps as well. The travel times τ_{ij} for both networks can be seen in Table 3.2.

The initial container state and the initial states of arriving containers and vehicles are zero in all scenarios, $\tilde{x}_i^c(t=0) = \mathbf{0}_{n_c}$, $\tilde{u}_i^h(t=0) = \mathbf{0}_{n_c}$ and $\tilde{v}_i^h(t=0) = \mathbf{0}_{n_v} \forall i \in \mathcal{N}$.

The demand profiles for the virtual demand nodes 3, 5, and 10 were generated based on individual transport orders with an allowable lead time of minimum 40 time steps. Two different demand profiles were used, as seen in Figure 3.3. One where significantly more containers are imported (destination 3 and 5) than exported (destination 10/ship), and one where the import and export are proportional, which are referred to as unbalanced and balanced demand, respectively. When empty and full containers are considered as disconnected problems (e.g. [65]) the unbalanced demand is more common in Europe. Balanced demand is on the other hand more representative if

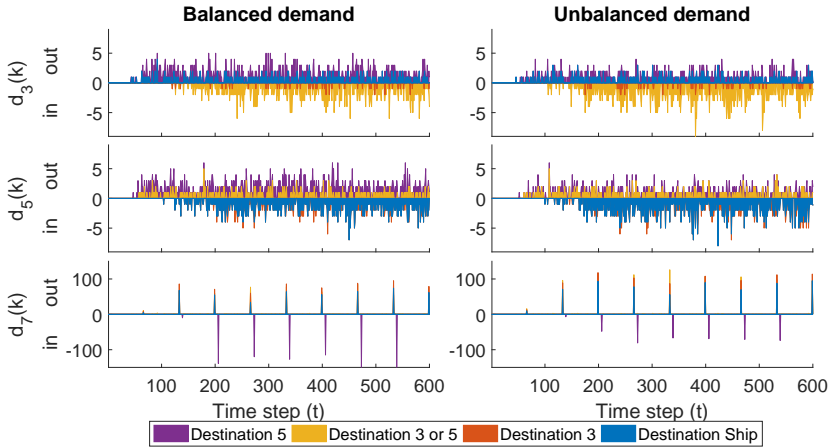


FIGURE 3.3: *Demand profile at the three virtual destination nodes for both balanced and unbalanced demand. The quantity of new demand $d_i(k)$ is shown over time steps. Outgoing demand is shown as positive and the incoming demand as negative.*

movements of all containers are considered in one problem (e.g. [144]). It is assumed that the controller at all timesteps has access to an accurate demand prediction for the prediction horizon.

Three scenarios with different levels of available resources have been used in the experiments. In the first scenario the constraints on container storage, truck parking, truck (un)loading and ship (un)loading are sufficiently loose that they do not become active during the simulation experiments. These constraints are all tightened to restrictive levels for the second scenario. The third scenario combines the constraints from the second scenario with limitations on the number of trucks available. Since the third scenario requires control of the total number of trucks in the system, the benchmark method is, like most methods in synchromodal transport literature, not applicable and only simulation results for the proposed method are presented.

A summary of the differences between the available resources in the three scenarios can be found in Table 3.3 together with the non-zero initial states. The train and barge schedules, capacity and maximum (un)loading rates are the same for all three scenarios with the following values: $c_{11}^c = 80$, $c_{11i} = c_{i11} \in \{0, 50\} \forall i \in \{4, 5\}$ and $c_{12}^c = 45$, $c_{12i} = c_{i12} \in \{0, 30\} \forall i \in \{1, 2, 3\}$ for the barge and train respectively.

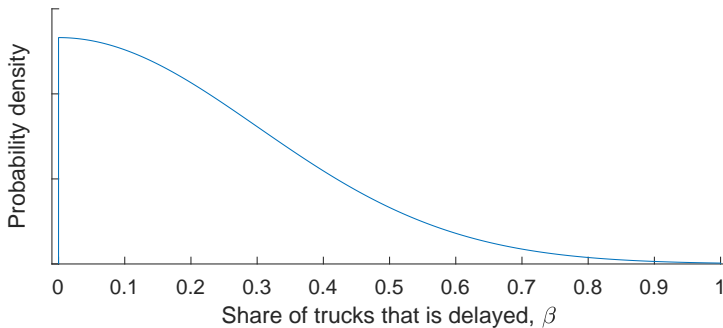


FIGURE 3.4: *The probability that a share of the trucks on an arc is delayed by 15 minutes. If no trucks are delayed $\beta = 0$ and if all trucks are delayed $\beta = 1$.*

If a truck is delayed, then not only the container it currently transports is affected, but also the containers it was scheduled to transport in the future. The MPC can react to delays and reschedule, such that other trucks transport the most urgent containers. This, however, requires that other trucks are available. The simulations from this chapter are therefore performed for both the case without delays (nominal case) and the case where the trucks may be delayed (uncertain case). The delays are not predicted by the MPC, but are added to the system by changing the distribution of the incoming container and truck flows, $u_i^h(k)$ and $v_i^h(k)$. At each timestep, a percentage of the trucks that are almost arriving is delayed exactly one timestep. One truck can be delayed several times, this corresponds to the trucks not informing the central planner in advance if they foresee a longer delay. It is assumed that the probability that an empty truck is delayed is the same as the probability that a loaded truck is delayed, hence the differences between the predicted state at time k computed at $k-1$ and the measured state are $\tilde{u}_{ji,m}(t - \tau_{ji}) = \beta u_{ji,m}(-\tau_{ji}|t-1)$, $\tilde{u}_{ji,m}(1 - \tau_{ji}) = (1 - \beta)u_{ji,m}(-\tau_{ji}|t-1) + u_{ji,m}(1 - \tau_{ji}|t-1)$, $\tilde{v}_{ji}(t - \tau_{ji}) = \beta v_{ji}(-\tau_{ji}|t-1)$, and $\tilde{v}_{ji}(1 - \tau_{ji}) = (1 - \beta)v_{ji}(-\tau_{ji}|t-1) + v_{ji}(1 - \tau_{ji}|t-1)$ where β is the share of the trucks that are delayed and the notation $a(k|t)$ indicates the prediction of $a(k)$ computed at time t . The parameter β is drawn per arc in the network per timestep from the truncated normal distribution seen in Figure 3.4.

TABLE 3.3: *The difference in parameters and initial conditions between the performed simulations. The five scenarios are evaluated using both the balanced and the unbalanced demand profiles.*

METHOD	RESOURCES		
	Unrestricting	Tight	Limited trucks
Proposed $c_8^v = [1000 \ 0]^T$ $c_9^v = [0 \ 1000]^T$ $\bar{x}_i^v(0) = [0 \ 0]^T$ $\forall i \in [1, 7]$	Scenario P1 $c_i^c = 1000 \ \forall i \in [1, 7]$ $c_i^m = 300 \ \forall i \in [1, 7]$ $c_{12,7}, c_{7,12} \in \{0, 200\}$	Scenario P2 $c_i^c = 75 \ \forall i \in [1, 5]$ $c_i^c = 250 \ \forall i \in \{6, 7\}$ $c_i^m = 20 \ \forall i \in [1, 5]$ $c_i^m = 50 \ \forall i \in \{6, 7\}$ $c_{12,7}, c_{7,12} \in \{0, 40\}$	Scenario P3 $c_i^c = 75 \ \forall i \in [1, 5]$ $c_i^c = 250 \ \forall i \in \{6, 7\}$ $c_i^m = 20 \ \forall i \in [1, 5]$ $c_i^m = 50 \ \forall i \in \{6, 7\}$ $c_{12,7}, c_{7,12} \in \{0, 40\}$
	Benchmark $c_8^v = [0 \ 0]^T$ $c_9^v = [0 \ 0]^T$ $\bar{x}_8^v(0) = [0 \ 0]^T$ $\bar{x}_9^v(0) = [0 \ 0]^T$	Scenario B1 $\forall i \in \{1, 4, 6, 7\}$: $c_i^v = [1000 \ 1000]^T$ $\bar{x}_i^v(0) = [500 \ 500]^T$ $\forall i \in \{2, 3, 5\}$: $c_i^v = [0 \ 1000]^T$ $\bar{x}_i^v(0) = [1000 \ 0]^T$ $\bar{x}_9^v(0) = [0 \ 1000]^T$	Scenario B2 $\forall i \in \{1, 4, 6, 7\}$: $c_i^v = [1000 \ 1000]^T$ $\bar{x}_i^v(0) = [500 \ 500]^T$ $\forall i \in \{2, 3, 5\}$: $c_i^v = [0 \ 1000]^T$ $\bar{x}_i^v(0) = [0 \ 1000]^T$

3.5 Results and discussion

This section presents the results of the comparison between the proposed method that considers the movements of trucks and containers simultaneously and the benchmark method that assumes infinite and instant truck capacity. Both methods compute what actions to take fast due to the optimization problems' convex nature. The average and maximum time the MPCs needed to compute the inputs at any timestep t are shown in Table 3.4 for all scenarios. The benchmark method is faster than the proposed method, but both are significantly faster than the chosen update rate of $\Delta T = 15$ minutes, i.e., the real-time performance is guaranteed. In a real-world implementation, one could thus choose to increase the prediction horizon or decrease the update rate. The trends from the results presented in this section are expected to hold in such cases too.

The results show very significant differences in the utilization of the transport modes for the proposed and the benchmark methods. The utilization do however not differ much between the nominal scenarios and their counterparts with uncertain travel time. Hence, in the following only the vehicle utilization in the nominal scenarios will be discussed.

In Figure 3.5 the results for the nominal scenarios with unbalanced and balanced demand are presented. The solid color blocks show the results for the proposed method, while the lines show the results produced by the benchmark method. The dark blue area indicates how many hinterland trucks were transporting containers at a given time, while the translucent blue indicates

	Balanced demand					Unbalanced demand				
	P1	P2	P3	B1	B2	P1	P2	P3	B1	B2
Average	31	29	26	23	25	29	27	26	22	22
CPU time (s)	27	30	29	22	25	26	26	29	22	23
Maximum	42	41	34	30	38	44	36	39	31	31
CPU time (s)	44	53	37	29	42	43	35	47	30	30

TABLE 3.4: *The average and maximum CPU time for the MPC to compute outputs at one time step t . Black: nominal case. Grey: case with uncertain travel time.*

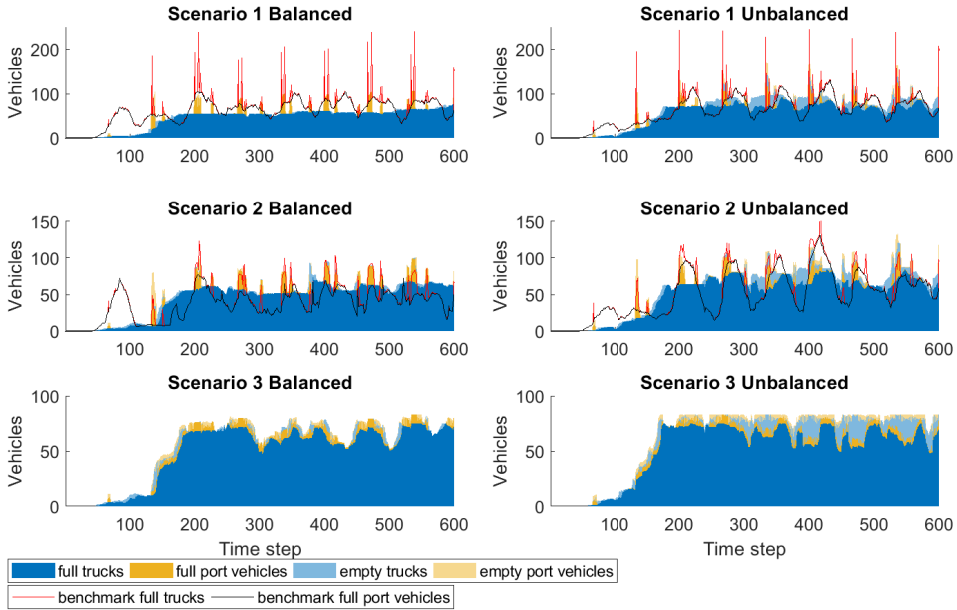


FIGURE 3.5: *Number of driving vehicles $\sum_{i \in \mathcal{N}, j \in \mathcal{N}} \sum_{l=0}^{\tau_{ij}-1} v_{ij}(k-l)$, i.e., for each time step the y-axis shows the number of vehicles that has departed a node but has not yet reached the next node. The shaded area is the number of vehicles for the proposed method with the light part being the portion of empty vehicles. The lines are the full vehicles when using the benchmark method. The information is stacked for each method.*

how many empty hinterland trucks were driving in the network. The yellow and translucent yellow show the same for the port vehicles. In scenario 3 with unbalanced port demand profile all 75 hinterland trucks and 8 port vehicles are driving either empty or full at nearly all timesteps $k > 200$. For the benchmark method, only information about the vehicles that transport containers exists due to the assumption of instant and infinite truck capacity. Thus for the benchmark method no information on empty vehicles is available to be shown. The full hinterland trucks are shown in black and the full port vehicles are shown in red. The information is stacked in the same manner as the information for the proposed method, i.e. for scenario 1 with unbalanced demand the benchmark method used 72 hinterland vehicles and 173 port vehicles to transport containers at timestep $k = 400$. Notice that all experiments start with an empty system and slowly increasing demand.

In all scenarios, less containers are moved by port vehicles when containers and trucks are routed simultaneously. The number of port vehicles needed in peaks is also significantly lower. In scenario P1 with unbalanced demand, port vehicles were driving 3925 timesteps, while only transporting containers 2041 of those timesteps. The maximum amount of containers that were transported by port vehicles at any timestep was 74 in this scenario. In the corresponding scenario B1, 174 containers were transported at one timestep. This is likely to incur even higher costs for driving empty. The only scenario where the port vehicles are almost continuously in use is P3, where only a very limited number is available. The ability to take decision based on the actual number of trucks available thus have a large impact on the realistic viability of the operational decisions. The methods presented in this chapter are intended for the operational level, but if they are used on the tactical level to dimension a truck fleet, the results show that it is important to consider container and truck routes simultaneously, since the peaks here give a realistic indication of the necessary fleet size.

The number of vehicles driving empty is as expected higher in the scenarios with unbalanced demand since the trucks have to be replaced to the port before they can transport new containers. The benchmark method always has trucks available, and do thus not need to wait. The fluctuation in the truck usage is therefore much higher for the benchmark method. The same trends can be seen for the scenarios with balanced demand, however with smaller peaks. The number of vehicles driving empty with the proposed method is very low with empty to full ratios of less than 3% in all scenarios with balanced demand. This indicates that considering containers and trucks simultaneously provides benefits such that trucks may wait for a new transport demand before departing from a node.

The utilization of the barge and train is shown in Figure 3.6. It shows the time the barge and train operate at different utilization rates. The results are shown in separate bars for import and export. Hence, a barge-import-utilization rate of (40%-80%] is achieved when (40%-80%] of the barge capacity is filled with import containers. Only results for timesteps $t \geq 250$ are used to generate the data, as the network starts empty and only around this timestep is fully developed according to Figure 3.3 and Figure 3.5. The results are furthermore shown in percentage of this time-interval rather than timesteps. Hence in scenario 1 with balanced demand and the controller considering simultaneous routing (P1B), import containers take up (1%-40%] of the barge capacity 29% of the time, (40%-80%] of the barge capacity 18%

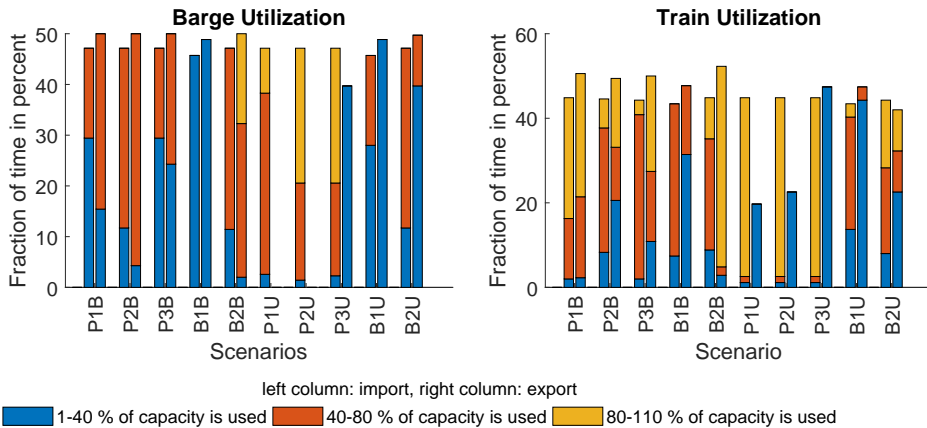


FIGURE 3.6: *The occupation rate at the barge and the train for each simulation. Only results for $k \in [250, 600]$ are used as the initial zero-states impacts the earlier results. Example of how to read figure: In scenario P3 with balanced demand profile (P3B) the import takes up (1-40)% of the barge capacity 29% of the time. In the same scenario the export takes up (40-80)% of the capacity 50-24=26% of the time.*

of the time and never takes up more than 80% of the barge capacity. In the formulation of the model, it is not specified when a container should be unloaded from a scheduled service at the time it is loaded to the service, it is thus possible for both import and export containers to stay on the barge or train while they travel in both directions. This do however only occur in the beginning of the simulations and never while $t > 250$.

The number of containers that are not delivered in time varies largely between the different scenarios, and from the nominal cases to the cases with uncertain travel time. Since the model considers commodity flows, the deadline of each container is not considered, instead the satisfaction of demand for a given commodity at a given destination is discussed. In Table 3.5 the unsatisfied demand at the three destinations are shown over all scenarios for $t > 200$. The numbers count each timestep a demand was not satisfied, so if a container of a given commodity was needed at time $t = 300$ but is only satisfied at $t = 305$, it adds 5 to the count. The demand at the ship (node 1) follows the capacity at which the ship can be (un)loaded. In scenario 1 all demand is thus to be satisfied at the first timestep possible, while the demand is spread over more timesteps in scenario 2 and 3.

node	Balanced demand					Unbalanced demand				
	P1	P2	P3	B1	B2	P1	P2	P3	B1	B2
3	26	26	29	19	25	38	37	1435	17	24
	126	110	96	105	65	180	144	1924	176	107
5	25	28	28	19	26	39	37	1163	18	23
	97	80	65	120	50	165	120	1546	152	86
Ship	3	2	5	2	2	3	1	22	3	1
	8	15	44	6	5	6	8	33	4	2

TABLE 3.5: *Unsatisfied demand where each delayed container is counted each timestep. Black numbers are from nominal scenarios, while grey numbers are from scenarios with travel time disturbances.*

When both container and truck routes are considered, more demand is left unsatisfied in most cases. This is especially pronounced in the case of unbalanced demand, where there is more import than export. In order to reduce the cost of empty trucks driving from the hinterland to the port, the proposed method prefers the use of scheduled services as seen in Figure 3.6. This results in longer travel times and thus more unsatisfied demand. The scheduled services are priced per slot, so the MPC does not consider the cost of sailing/driving back empty for those services. All scenarios prioritize satisfying the more costly ship demand over the inland demand and therefore the unsatisfied demand is much lower at the ship destination.

When the travel times of trucks become uncertain, all scenarios with sufficient numbers of trucks available adjust the plans and obtain results with expected increases in the unsatisfied demand. Scenario P3 with unbalanced demand does not have enough vehicles available to transport the containers quick enough. This is also visible in Figure 3.5 where all vehicles are driving either empty or full at all times. The results from this scenario does thus not represent an implementation we recommend, but serves to ensure that even in cases with under capacity the MPC does not violate the constraints or render infeasible. Eventually the MPC would accumulate very high costs which may give numerical issues, but in a real-world implementation high rates of unsatisfied demand will raise awareness and foster changes in either fleet size or accepted demand.

When the constraints are unrestricting more demand is left unsatisfied with the proposed method under deterministic travel times. However when the truck travel time becomes uncertain, the performance is very similar to the performance of the benchmark method. This is due to the benchmark method's constant availability of trucks in combination with the cheap stacking of containers in node 6. For the benchmark method in scenario 1 the cheapest option is often to truck the containers just in time to satisfy the demand and otherwise keep them at node 6. When trucks and containers are considered simultaneously, it becomes cheaper for the MPC to truck containers ahead of demand if there is otherwise an empty truck driving in that direction. The stacks at node 3 and 5 do thus contain containers with destination at their adjacent virtual destinations more often for the proposed method or when the capacity in node 6 is a limiting factor. More local storage increases the chance that a delay in the truck travel time will not cause a delay in demand satisfaction.

3.6 Conclusions

The method presented in this chapter, is an answer to Research Question Q2 *How can operational planning under synchromodal transport take advantages of the opportunity for real-time mode-changes?* The presented results addresses additionally Q1 *What is the impact of integrating decisions across the planning-hierarchy layers that concerns container and vehicle routing?* The method assumes a single decision maker has the authority to take all decisions regarding containers and trucks (i.e. vehicles with low capacity and no pre-determined schedule). A core assumption in the presented method is that commodity flows accurately represent container transport. This may be the case for large quantity flows of non-perishable goods, but may not hold when the deadline of each specific container is important.

The proposed method is based on MPC, which previously has been used for container transport planning in some instances, but not for planning containers and trucks simultaneously. It routes containers and trucks simultaneously in real-time and thereby change the mode the containers are transported by depending on the availability of trucks. At the same time, the truck routes changes to serve the known demand best possible. It is shown that the MPC can adjust the plan when the travel times are subject to disturbances. When containers and trucks are planned simultaneously, the MPC

is encouraged to transport containers when an empty truck is available at the right location and not only based on deadlines, which makes the system less sensitive to travel time delays.

It was found that the often used assumption that trucks are instantly available at any location in the synchromodal network significantly changes what the optimal actions are. A plan under this assumption is thus likely to perform worse in reality where only a finite number of trucks are available. The proposed method routes trucks and containers simultaneously and successfully smooths out peaks in the needed number of trucks, even when it has a large number of trucks available. This creates better plans for companies that hire third-party trucks for excess capacity, as the company's basic fleet will be utilized better between peaks and less third-party trucks will be needed to serve the peaks.

The proposed method can plan for a finite number of trucks available in the system. When the plan is optimized for the actually available number of trucks, infeasible plans are avoided and the utilization rates of the available trucks are improved, i.e. the number of empty trucks driving in the network is reduced to the benefit of the environment, society and economy. The proposed method furthermore gives information on how many trucks are to be relocated between specified locations. This information can be used in hindsight or already at the planning stage to indicate beneficial volume-changes to the transport company's sales department.

This chapter shows that integration of operational decisions across the traditional hierarchy boundaries makes the resulting plans more realistic and uses the vehicle fleet better. However, to integrate decisions, multiple stakeholders need to cooperate. In the next chapter, Chapter 4, the co-planning of container and truck routing is considered.

Chapter 4

Real-time Co-planning for Efficient Container Transport

Integration of operational decisions in a synchromodal transport network improves the utilization rate of the vehicle fleet. Chapter 3 shows this for container and vehicle routing decisions in a network where one centralized entity has full authority over both container and vehicle movements. Often, this decision power is split between a logistics service provider and a transport service operator. In this chapter, Research Question Q3 is addressed by comparing the proposed co-planning method to a traditional synchromodal method, where the logistics service provider re-plans the container movements based on real-time location information from terminals, instead of the feedback from the transport service operator. Furthermore, Research Question Q4 is answered as the presented co-planning method relies on communication of traditional transport requests and their expected fulfilments.

The core of this chapter will soon be submitted to a journal. The chapter is organized as follows: the problem is described in Section 4.2 where the differences between synchromodal transport with co-planning and earlier transport paradigms are also discussed. In Section 4.3 the mathematical formulation of the proposed method is detailed and in Section 6.5 numerical experiments that show the performance of co-planning in comparison to intermodal and synchromodal planning without cooperation is shown. Conclusions and outlooks on future research are drawn in Section 4.5.

4.1 Introduction

Increasing the flexibility and responsiveness of transport networks is expected to increase transport sector's efficiency and lower its environmental impact. Flexibility requires a high degree of cooperation and responsiveness can only be achieved if plans adapt to real-time information. In the literature, transport models at the operational level (e.g., [8, 22, 33, 108, 162]) are often developed for static scenarios and consider the decisions of only one stakeholder.

Synchromodal transport is a paradigm that seeks to achieve flexible and responsive transport systems, mainly by allowing transport providers to decide which mode a given container is transported by, even after departure, without waiting for response from the shippers [133]. This way, the transport provider can use the available capacity better and adapt plans if disruptions or disturbances occur. Typically, slower transport modes such as train, barge and short sea shipping are economically and environmentally better options than truck transport [159], but since their reach is limited to the extent of the rails and waterways, transport by these modes often requires road transport for first and last mile. Small delays can easily occur in transport systems, which creates uncertainties in transfers of containers from one mode of transport to another. This risk can be mitigated by making more robust plans [100], but in practise, most shippers instead choose for uni-modal road transport [36]. Surveys show that if a synchromodal transport network can deliver a reliable service, shippers are mostly willing to hand over the mode choice [71].

The real-time aspects of synchromodal transport clearly sets it apart from intermodal transport at the operational level. In [5] multi-agent simulations show that dynamic synchromodal solutions perform better than intermodal ones when the transport system is subject to disturbance. New methods that can provide up-to-date decisions are thus needed. Delbart et al. [32] give insights into how uncertainty is handled in parts of the literature. Some works focus on robust plans that are likely to perform well under several different scenarios (e.g., [33] and [37]), while others rely on replanning (e.g., [56] and [80]). Many only replan when a disturbance makes the current plan infeasible, a few systematically replan at time intervals. Replanning when a plan becomes infeasible [164] and possibly only replan for parts of the transport system [130] can reduce computational complexity significantly. This will create faster responses. On the other hand, replanning all decisions

at regular time intervals ensures the best possible plans for all parts of the system. The advantage of using regular intervals is that the time available for computing the plan is known. This is especially an advantage if new information is available frequently and irregularly. In this paper, we call the regular replanning times as *decision moments*. To replan regularly, model predictive control (MPC) has earlier been used in e.g., [80], [87], and [110]. This method optimizes the modelled decisions for a finite time horizon, but implements only the actions corresponding to time period until the next decision moment in a receding horizon fashion.

Real-time mode changes performed by the transport provider cannot only improve efficiency by making rail and water modes more attractive, but can also improve the performance of the transport system by changing modes and routes depending on the available capacity of each vehicle [163]. Often container routing decisions in the literature on intermodal and synchromodal transport are taken assuming that the capacity of non-scheduled vehicles are infinite [152]. This can be a reasonable assumption in busy port areas, but even here, it can, according to practitioners, easily take two hours before a vacant truck is available on site. If decline in truck driver professionals continue [67], availability of ad-hoc trucks will decrease further. In a fully synchromodal transport system that uses real-time information, integration of container and equipment plans is expected to improve the efficiency significantly. In [80] it is shown that the integration of container and truck routes improves the use of a truck fleet. The same was found in [82] that furthermore showed the truck utilization (distance driven with a container compared to total distance) improves with integration. Qu et al. [130] provide a method to align the schedules of barges and trains with the routing of containers in an import network with limited truck capacity. Resat et al. includes the vehicle speed decisions in [134] while [181] consider the qualitatively formulated preferences of different shippers.

The previously mentioned works all consider a centralised decision maker. In contrast, transport systems often consist of several independent organisations. Very little research has been made on cooperation methods in synchromodal transport systems. On the 1st of December 2021, only five papers in the synchromodal transport field were indexed in Scopus which provide planning methods that does not assume a central decision maker, either in the form of an organization or platform. Of these, [16] focus on computing fairness fees that push the decisions of decentralized decision

makers towards the global optimum of the static problem. In [30] new shipping orders appear dynamically, but the updated information does not lead to re-planning, as all decisions are taken sequentially. Three methods are based in MPC and replan as such all decisions within a receding horizon at regular decision moments. [88] use distributed optimization methods to facilitate cooperation on the routing of containers between multiple transport operators. Each operator has a unique network that does not overlap with the others. The information exchanged between the operators is symmetric and the responsibilities of each organisation is similar. In contrast, [79] provide a co-planning method to let a transport provider with only trucks in its vehicle fleet influence the departure times of a barge that belongs to a different operator. The responsibility of each organisation is realistic and clear, and the information exchange is limited to potential schedules and aggregated costs. In [81] the method is extended to multiple transport providers and facilitates complete privacy of information.

In this paper, we present a real-time co-planning method for the cooperation between a logistics transport provider and a transport service operator whose services have flexible departure times and thus operate ‘on-demand’. We define *co-planning* as a cooperation method that aim at improving the common operation of a transport system by exchanging consciously chosen, realistic information and keep the responsibilities within each autonomous, cooperating party as much as possible. Co-planning is a necessary concept as it distinguishes research that emphasises realistic assumptions on how entities cooperate.

In the broader literature on freight transport and logistics, cooperation is treated from many different perspectives, from fully decentralised multi-agent systems [2] to centralized platforms [6]. Among the methods to distribute shipping orders among competing, unimodal transport providers, auction based methods have many of the qualities that are sought for in co-planning. The responsibility of each transport provider is typically realistic and only transferred for individual shipping orders when a provider choose to do so (e.g., [28]). The information sharing in auction based methods and the value of different information is discussed in [96]. Commonly, only order information, such as pick-up and delivery time and locations, is shared. Sometimes the communication is peer to peer, sometimes to a central platform.

When heterogeneous freight transport partners cooperate, each partner has different focus and priorities, and is thus likely to use different planning

method than the other parties. Very little research has been conducted in the area of freight transport with interfaces between different method-types. [138] integrate the long-haul with the drayage operations and study the interface between the two. They consider both decisions taken by the same entity, but the long-haul is modelled as a Markov Decision Process model and the drayage as a Mixed-Integer Linear Programming model. Interfaces between different methods will influence the performance of any system and must therefore be taken into account.

In this paper, we propose a co-planning method that entails communication between a logistics service provider (LSP) and a transport provider (FSO) in a synchronomodal container transport network with additional stakeholders. The cooperation relies on automated communication of order information similar to what is being communicated manually in current practice. The responsibilities of each organization are clearly defined and they use two different planning methods that each suits their decision-types better. The LSP's real-time optimization is an extension of the method published in [56] where an LSP books flexible and scheduled services for dynamic demand. The method is in this paper extended to reconsider plans periodically and incorporate expected arrival times communicated by the FSO. The FSO's method has its ground in [80] where MPC is used to route containers and trucks in a centrally controlled, synchronomodal transport network. The main differences are that in this paper, the method is extended to distinguish containers from different transport requests and earlier actions are promoted. Co-planning has earlier been introduced in [79] in the context of planning barge departure times together with container and truck routes between a barge and an FSO with focus on using learning to bridge the information gap. In this paper, the focus is keeping the communication between the LSP and the FSO realistic and close to current practise. The communication between such different organizations with different planning methods to achieve a more efficient transport system is our core contribution.

4.2 Problem description: from intermodal planning to synchronomodal co-planning

The scarcity of resources is important to take into account when planning any transport operation. When a container's departure is decided immediately before it is to be loaded to a vehicle, knowing the availability of that




FIGURE 4.1: *Schematic of communication (dashed lines) between and actions (solid lines) of the LSP and the FSO in a traditional intermodal transport system.*

vehicle becomes crucial for the feasibility and efficiency of the transport plans. In transport networks, different stakeholders are in control of different parts of the infrastructure and resources needed to move goods making access to information limited. We study the problem of co-planning between two decision-makers in container transport. The aim of the method is to maintain the responsibilities of the different stakeholders as they are now and add only realistic communication. The current situation, transport under the intermodal paradigm, is depicted in Figure 4.1. The first decision maker is a Logistics Service Provider (LSP) that receives transport requests from shippers. The LSP is responsible for organizing the transports, i.e. book spots on scheduled services, like trains, and book flexible transport services, like trucks. The second decision-maker operates transport services with flexible departure times, henceforth denounced Flexible Service Operator (FSO). The FSO receives transport requests from the LSP with an earliest pick up time and a latest drop off time. The FSO routes the vehicles

in their fleet and decides where to load and unload containers, they furthermore decide how many containers from each transport request to transport. The FSO may store containers at intermediate stacks if it is beneficial. If the FSO fleet contains multiple types of vehicles, the FSO is responsible for mode-selection within their network. Flexible services are typically trucks and barges, but could also be, e.g., automated ground vehicles in a port area.

On one hand, the LSP has contact with the shippers, and thus better predictions of upcoming demand. On the other hand, the FSO knows the location and occupation-status of each vehicle in their fleet. To transport the containers as cost-efficient as possible, as well as utilize the flexible vehicle fleet optimally, the LSP and the FSO can communicate.

The transport network consists besides the LSP and the FSO of multiple passive stakeholders whose decisions we do not model. We assume there is no competitors to the LSP nor the FSO. For each stakeholder, the assumptions are as follows:

Shipper All shipping requests must be fulfilled eventually; delays are fined. All containers are standard 40ft. All transport requests are synchronomodal (lets LSP change mode during transport) and can be split. A shipping request consists of a request ID, announce time, release time and origin, due time and destination, and the container quantity.

Logistics Service Provider (LSP) [Decision maker] Decides what mode and route each container in a shipping request will follow. Decides (un)loading of scheduled services. Formulates transport requests to the FSO consisting of a request ID, a pick-up time and location, a drop off time and location, and the container quantity.

Flexible Service Operator (FSO) [Decision maker] Decides vehicle routing and (un)loading operations. The fleet size is limited. Delays on delivery of transport requests are fined.

Operators of Scheduled Services A predetermined spot-market priced capacity on each departure is available for last minute bookings by the LSP. This capacity and the service schedules are known by the LSP.

Terminal Operators Sufficient (un)loading capacity is available to satisfy the decisions. Opening hours are disregarded. (Un)loads containers to/from flexible and scheduled services.

Flexible Service Vehicles Travel the arcs decided by the FSO without delays. Drivers' working hours are disregarded. Cost of driving is the same when loaded and unloaded.

4.2.1 Comparison of different transport paradigms

The main difference between intermodal and synchromodal transport is the ability to adapt transport decisions at any point in time. Since the LSP gets new shipping requests of varying priority at different times (often referred as dynamic demand), changing previous plans can ensure a better overall transport performance. Re-evaluating plans and creating new ones for new requests can happen when it becomes relevant or at regular intervals. In this paper, regular intervals are used and denounced *decision moments* and indexed by k . What decisions that are re-evaluated at decision moments and the information that is available to do so differs between the different transport paradigms. Table 4.1 gives an overview of intermodal, active synchromodal, passive synchromodal and synchromodal with co-planning. In the following, the joint planning problem of an LSP and an FSO is described under each paradigm.

TABLE 4.1: *Overview of the differences between the compared transport paradigms in replanning of decisions and communicated information.*

	Decisions			Communication	
	LSP	FSO		LSP→FSO	LSP←FSO
	Container mode choice	Container departure time	Vehicle and container decisions	Transport requests	Container locations
Intermodal	fixed	adjust if arrive late	reconsider periodically	once per container	upon late arrival
Passive synchromodal	adjust if arrive late	adjust if arrive late	reconsider periodically	as necessary	upon late arrival
Active synchromodal	reconsider periodically	reconsider periodically	reconsider periodically	periodically	upon departure & arrival
Co-planning	reconsider periodically	reconsider periodically	reconsider periodically	periodically	expected & upon arrival

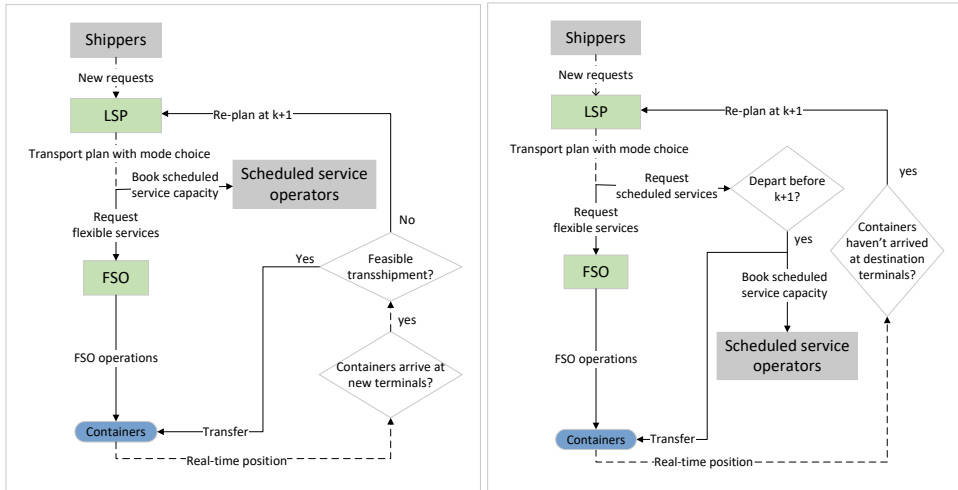


FIGURE 4.2: Schematics of communication between and actions of the LSP and the FSO in passive (left) and active (right) synchronomodal transport systems.

The cooperation between the LSP and the FSO under a traditional intermodal paradigm is described in Figure 4.1. At each decision moment, the LSP collects all new shipping requests and creates transport plans for them based on the current knowledge of capacity on the scheduled services. The operators of the services used in the plan are then notified of the containers to be transported by them and fulfill the wishes as possible. If the plan for a container becomes infeasible (e.g., there is no available truck or the container arrive after the departure of the scheduled service it was planned for), the operator lets that container wait for the next available service of the mode specified by the LSP. As the FSO’s vehicles can depart at any time, the FSO is even under the intermodal transport paradigm reevaluating their internal plans for all vehicles and containers at every decision moment.

Under synchronomodal transport, plans can be re-evaluated when it becomes clear they cannot be executed as intended (*passive synchronomodal*) or at every decision moment (*active synchronomodal*). Both concepts are represented in Figure 4.2. The FSO’s internal operations remain the same. In a passive synchronomodal transport system, the LSP is notified when a container arrives at a transshipment terminal too late to continue its journey on the intended scheduled service. The LSP will then cancel all future bookings for this container and include it in the planning made at the next decision moment. The transport operators used in the new plan will be notified about

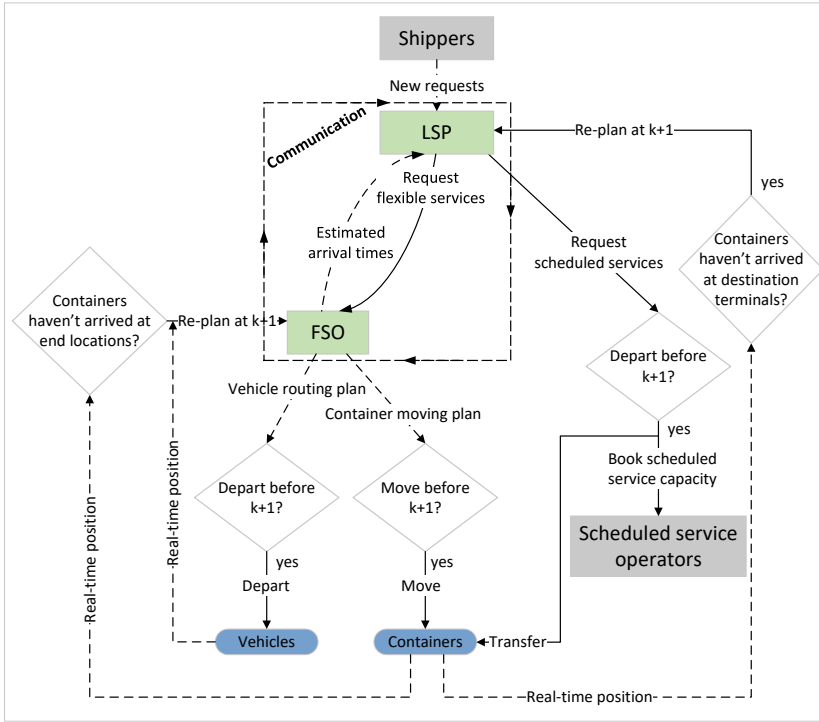


FIGURE 4.3: Schematic of communication between and actions of the LSP and the FSO in an active synchromodal transport systems with co-planning.

the new needed services. In active synchromodal transport systems, the LSP is not only re-planned if the initial plan becomes infeasible. Every time a container is at a transshipment terminal, their further route is reconsidered based on the available information on other shipping requests. Notice more information is available at this point of time, than when the LSP initially received the container’s shipping request. The LSP informs under active synchromodality the transport service operators about upcoming demand at every decision moment, but only commits to transports that will prepare to depart before the next decision moment.

When co-planning is added to active synchromodality, the LSP additionally gets feedback on the plans of the FSO. To co-plan, both stakeholders must work towards a common good by communicating in a realistic manner and without compromising each other’s authority. In this case, the common good is transporting all the containers at the lowest total operation cost and the responsibilities are clearly divided as outlined in the beginning of the section. The communication is kept as it is in current practice, the LSP only

communicates transport requests which they want the FSO to carry out and the FSO only responds to these requests. The communication and authority assumptions prevent a global optimization, so the achieved operation cost may not be the lowest possible, as the FSO also cannot suggest container transports. The communications and actions of the proposed co-planning framework are shown in Figure 4.3. The FSO's internal operation and the actions and communication related to the operator of scheduled services do not differ from active synchronomodality. At each decision moment, the LSP requests transports from the FSO who will provide feedback on if the containers can reach their drop-off location in the time the LSP expects. If the containers will be late, the FSO will also inform the LSP on when all containers in the requests is expected to arrive and how many containers that will be picked up before the next decision moment and the expected arrival time of this portion. If there are multiple rounds of communication, the LSP uses this information to replan all shipping requests and the process repeats. The earliest expected arrival time of the containers that will be late is used to compute a lower bound for when that portion of the shipping request can be picked-up by the corresponding service. After the last communication round, the feedback from the FSO does not influence the actions taken at the current decision moment, but it will at the next decision moment.

4.2.2 Impacts of planning under different paradigms

When plans are adapted differently under the four transport paradigms, the resulting actions in the transport system will also differ. Figure 4.4 shows the realised movements of containers and trucks in a small, illustrative case study with two shipping requests and a unimodal FSO operating trucks. The details of the case study are discussed together with further results in Section 4.4.2 and Section 4.4.3. Each of the nine locations in the transport network appear on the horizontal axes of the figure without respecting distances and network structure. The time dimension appears on the vertical axes. The black lines show the position of the FSO's trucks. The width of the black lines indicates how many trucks are at a given location or in transit. For example, under the intermodal paradigm, all the trucks drive from Node 9 to 1, while under co-planning, only some drive from Node 9 to 5 and the remaining stays parked at Node 9. For scheduled barge and train services, the width of the lines indicates how many containers are transported by the service. If a scheduled service is not used, it appears as a dotted line. The

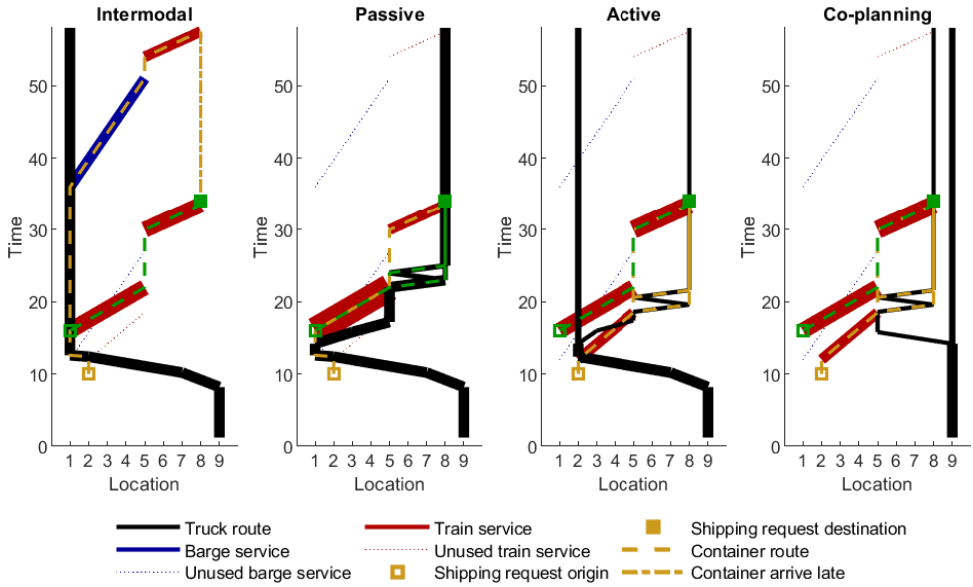


FIGURE 4.4: *The routes of containers and vehicles under each of the four transport paradigms.*

two routes of the containers from the two shipping requests are marked with green and yellow respectively.

In the case study, the LSP expects that the containers from the yellow shipment can be picked up by a truck and transported to Node 1 in time to take the barge which departs from there at time 12. However, as all trucks initially are far away (at Node 9) and the shipping is requested at time 8, the containers cannot be picked up before time 12.4, making them unable to reach the barge in time. Under intermodal transport, the containers stays at Node 1 until the next barge becomes available. This leads to the containers arriving at their destination late (arrival: 57.5, due date: 34) and a very inefficient use of trucks. All trucks drive the long way from Node 9 to 2, while a single truck would be able to transfer all containers from Node 2 to 1 in the available time before the second barge departs. The green shipping request is announced after it becomes clear that the yellow request will be late, and it is therefore known that the train from Node 5 to 8 leaving at time 30 has enough capacity to transport all containers in the green request.

Under passive synchronomodal transport, the late arrival of the containers to Node 1 causes the LSP to consider the yellow order as part of the plan made at the next decision moment. This results in the mode being changed

from barge to train travelling from Node 1 to 5. The cheapest route continues with a train service from Node 5 to 8. This train has limited capacity, so when the green shipping request is announced, there is not enough capacity left to transport these containers. The green containers are thus transported by the more expensive truck mode instead. This increases the total cost of transport as there are more containers in the green shipping request than in the yellow. Under passive synchronomodal transport all containers arrive at their destination in time, but again, more trucks are used than what is needed resulting in an unnecessary low truck utilization (distance driven full over total distance). Compared with the intermodal paradigm, the plan under passive synchronomodal paradigm benefits from mode changes when containers arrive late at transfer nodes.

In contrast to the passive synchronomodal paradigm, the LSP's plans are reconsidered periodically under the active synchronomodal paradigm. At the decision moment at time 12, the LSP knows that the containers from the yellow shipping request still are at Node 2, and thus cannot reach the first barge from Node 1. The route of the containers is changed to go by the train departing immediately from Node 2. After the green shipping request is announced, both requests are part of the planning problem. Since there is more containers in the green shipping request, they are assigned to the train from Node 5 to 8 leaving at time 30 and the yellow are to be transported by truck. The FSO receives a transport request from Node 5 to 8 with sufficiently long time between pick up and delivery, that a few trucks are able to transport all containers within the FSO's planning horizon. Therefore, the FSO only sends a few trucks from Node 2 to 5. The decision moment after some of the yellow containers have been picked up at Node 5, the LSP replans based on the information that some of the containers still are at location 5. This results in a new request to the FSO for transporting the remaining yellow containers. The FSO has (as they predicted earlier) enough time to use the trucks that have dropped off the first batch of yellow containers at Node 8 to pick up some of the remaining containers. The distance from Node 8 to 5 is smaller than that from Node 2 to 5, so this is the most economical decision. After the FSO has picked up the second batch of yellow containers, only a few containers are left in Node 5. When the LSP now does the periodic replanning, these remaining containers can fit on the train leaving Node 5 at time 30. The periodic reconsideration of the LSP's plan ensures a cheaper transport of the containers as they can use barge and train for more of their journeys. It also improves the truck utilization rate, as less containers are to

be transported from Node 5 to 8, allowing the FSO to let some trucks remain parked at Node 2.

When the LSP and the FSO communicate under co-planning, the containers follow the same route as under the active synchronodal paradigm. However, the trucks usage is significantly improved. The FSO lets the LSP know that the trucks cannot pick up the containers from the yellow shipment in time to reach the barge leaving from Node 1 at time 12. Since the LSP replans based on this knowledge before committing to any actions, the containers from the yellow shipment are immediately rescheduled to take the train instead of the barge. The FSO is thus never asked to transport anything from Node 2, which prevents the unnecessary relocation of trucks to Node 2 as in the case of active synchronodal paradigm. In this way, the empty truck travel distance is significantly reduced, in turn, the total operation cost is lower than the active synchronodal paradigm.

The case study demonstrates that when the LSP's decisions can be re-considered frequently, the containers are more likely to arrive at their destinations before their due date and at a lower operational cost. The cost of operation is furthermore improved when co-planning is used and the flexible vehicle fleet is better utilized.

4.3 Co-planning between LSP and FSO

The proposed co-planning framework, outlined in Figure 4.3, is detailed in this section. First, the receding horizon aspects of the method is presented. Hereafter, the LSP's and the FSO's individual planning problems are detailed. In Section 4.3.3, the proposed co-planning framework is presented. The communication is then described and an algorithm shows all details of the presented co-planning.

To facilitate periodic replanning, co-planning uses a receding horizon. At every Δt hour, the LSP and the FSO plan the movements of containers and vehicles. The LSP's plan spans over an infinite time horizon, while the FSO plans for the next $T_p \delta t$ hours using a state-space model discretized in time by δt . From these plans, only the decisions that are realized before a new plan is available are implemented and all other decisions are reconsidered, as illustrated by Figure 4.5. Decisions that are implemented before a given




FIGURE 4.5: *Illustration of the receding horizon.*

decision moment cannot be changed, but their consequences are known in form of future arrivals of containers and vehicles at the end of the arcs they are currently travelling. Shipping requests that are announced before the decision moment can be incorporated in the plan, while requests that are announced later remain unknown, regardless of release time. It is assumed that containers that are planned for scheduled services at decision moment k can be transported by services departing at time $k\Delta t + \delta t$ as travel time includes (un)loading of multiple containers and planning takes less than δt hours.

4.3.1 LSP optimization model

The LSP receives shipping orders for transports of high quantity, spanning long periods of time. Therefore, the LSP uses a matching model to optimize the routes of the containers through the synchronomodal transport network. We state the LSP's optimization problem without indicating the time the decision is taken for the ease of notation. This means that at any time k , e.g., R is the set of shipping requests on which the LSP can take decisions at this moment k and the variable $x_{rs} = 1$ if in this plan the containers from request r is transported by service s from the set of all services S . Scheduled services are only available according to their schedules while flexible services are considered always available. The later assumption will be tightened in Section 4.3.3. The LSP does not split requests as part of the planning problem, but before each decision moment, the LSP is notified of departures and arrivals of containers. If the containers of a shipping request r are at different places, e.g. some at a node and some being transported by the FSO, the LSP splits the request into bundles that appear as separate requests r' and r'' in set R . These bundles can be split again later on but never merged.

$$\min \sum_{r \in R} \sum_{s \in S} x_{rs} q_r c_s + \sum_{r \in R} t_r^{\text{delay}} q_r c_r^{\text{delay}} \quad (4.1)$$

subject to

$$\sum_{s \in S_{o_r}^+} x_{rs} = 1, \quad \forall r \in R, \quad (4.2)$$

$$\sum_{s \in S_{d_r}^-} x_{rs} = 1, \quad \forall r \in R, \quad (4.3)$$

$$\sum_{s \in S_i^+} x_{rs} = \sum_{s \in S_i^-} x_{rs}, \quad \forall r \in R, i \in N \setminus \{o_r, d_r\}, \quad (4.4)$$

$$\sum_{r \in R} x_{rs} q_r \leq Q_s, \quad \forall s \in S^{\text{scheduled}}, \quad (4.5)$$

$$t_{r o_r} = t_r^{\text{release}}, \quad \forall r \in R, \quad (4.6)$$

$$t_{ri} \leq TA_s + M(1 - x_{rs}), \quad \forall r \in R, i \in N \setminus \{o_r\}, s \in S_{i^-}^{\text{scheduled}}, \quad (4.7)$$

$$t_{ri} \geq TA_s + M(x_{rs} - 1), \quad \forall r \in R, i \in N \setminus \{o_r\}, s \in S_{i^-}^{\text{scheduled}}, \quad (4.8)$$

$$t_{ri} \leq t_{rs}^{\text{depart}} + t_s + M(1 - x_{rs}), \quad \forall r \in R, i \in N \setminus \{o_r\}, s \in S_{i^-}^{\text{flexible}}, \quad (4.9)$$

$$t_{ri} \geq t_{rs}^{\text{depart}} + t_s + 2M(x_{rs} - 1), \quad \forall r \in R, i \in N \setminus \{o_r\}, s \in S_{i^-}^{\text{flexible}}, \quad (4.10)$$

$$t_{ri} \leq TD_s + M(1 - x_{rs}), \quad \forall r \in R, i \in N \setminus \{d_r\}, s \in S_{i^+}^{\text{scheduled}}, \quad (4.11)$$

$$t_{ri} \leq t_{rs}^{\text{depart}} + M(1 - x_{rs}), \quad \forall r \in R, i \in N \setminus \{d_r\}, s \in S_{i^+}^{\text{flexible}}, \quad (4.12)$$

$$t_r^{\text{delay}} \geq t_{r d_r} - t_r^{\text{due}}, \quad \forall r \in R. \quad (4.13)$$

The LSP optimizes the transport cost and the fees for late delivery (4.1). q_r is the quantity of containers in request r . The costs c_s is the cost of transporting one container with service s and c_r^{delay} is the fee per hour for delivering one container from request r late. t_r^{delay} is the lateness of bundle r . The variable t_{ri} is the arrival time of request $r \in R$ at terminal $i \in N$ and the parameter t_{rs}^{depart} is the departure time of request r with the flexible service $s \in S^{\text{flexible}}$. From each node i in the set of terminals N , a set of services depart S_i^+ and arrive S_i^- . o_r is the origin of shipping request r until the request is released. After release, o_r is either the location of request r if it is at a node, or the destination of the scheduled service if r is being transported by a scheduled service, or the drop-off location of the transport request if the FSO is transporting request r . d_r is the destination of shipping request r .

(4.2)-(4.4) ensures all containers are transported from their origin to destination. The LSP restricts with (4.5) how many containers are planned for a scheduled service based on their exact knowledge of the service capacity Q_s .

Shipping request r 's release time t_r^{release} reflects its next earliest pick up time. Before the release of request r , t_r^{release} is the release time at its origin o_r ; if the request is stacked at a node, it is the current time; if the request is being transported by a scheduled service, it is the arrival time of that service; if request r is being transported by the FSO, r is the expected drop-of time of the transport request. In case the containers do not arrive in time, the LSP misses the arrival notification and will thereafter expect the containers to arrive δt hours after the decision moment. The connection between the first place and time of request r in the plan is ensured by (4.6). The arrival time t_{ri} of request r at location i is ensured to match the arrival time TA_s of the scheduled services that transports it by (4.7)-(4.8), while (4.9)-(4.10) match it with arrival times of the relevant flexible services. The departure times are matched by (4.11)-(4.12), where TD_s is the departure time of the scheduled service s . (4.13) count how long after its due time t_r^{due} , the containers from shipping request r arrives.

4.3.2 FSO optimization model

The FSO not only has to route the containers through their network, they also have to route the full and empty vehicles. To do so efficiently the FSO uses a state-space model that describes the dynamics of the system over a prediction horizon T_p that is discretized into what we denote *planning-timesteps* each spanning δt hour, i.e. at every decision moment, the FSO plans for the next $T_p \delta t$ hours and implements actions corresponding to the next $\frac{\Delta t}{\delta t}$ planning-timesteps. We use κ to indicate the planning-timesteps in the FSO's plan. We state the FSO's optimization problem without indicating the time the plan is made to ease the notation. For a plan made at decision moment k , $\xi_{rij}(\kappa)$ is thus the number of containers from transport request r that are planned to be loaded on a vehicle that at time $k\Delta t + \kappa\delta t$ will depart from node i to j . We assume (un)loading time to be included in the transport time. When necessary, the decision moment is indicated with $[k]$ after the general notation. The set of transport requests from the LSP to the FSO that are currently being transported or that are new requests at decision moment k is thus denoted \mathcal{R} when the decision moment is clear and $\mathcal{R}[k]$ otherwise. The

parameters and variables that are exclusive to the FSO's optimization problem are denoted with Greek letters to increase clarity. Indices are marked with superscripts differentiating similar variables/parameters, e.g., different kind of costs. Only flexible services with the same travel-time and capacity can operate on an arc. Vehicles with different characteristics transporting between the same nodes can be modelled by duplicating the nodes.

$$\begin{aligned} \min \sum_{i \in H} \left(\sum_{\kappa=1}^{T_p-1} \left(\sum_{j \in H_i} (1 + \omega^e \frac{\kappa}{T_p-1}) \omega_{ij}^v v_{ij}(\kappa) + \right. \right. \\ \left. \left. \sum_{r \in \mathcal{R} \setminus \{r|d_r=i\}} \omega^s \chi_{ri}(\kappa+1) + \sum_{r \in \mathcal{R}} (\hat{c}_r + \omega^d) \psi_{ri}(\kappa+1) \right) \right) \\ + \sum_{r \in \mathcal{R}} (1 + \omega^e) \omega_{id_r}^{direct} \left(\chi_{ri}(T_p) + \sum_{j \in H_i} \sum_{l \in [1, \dots, \tau_{id_r}]} \xi_{rji}(T_p - l) \right) \end{aligned} \quad (4.14)$$

subject to

$$\begin{aligned} \chi_{ri}(1) = \tilde{\chi}_{ri}(k), \quad \rho_i(1) = \tilde{\rho}_i(k), \quad \xi_{rij}(-l) = \tilde{\xi}_{rij}(k - l\delta t), \\ v_{ij}(-l) = \tilde{v}_{ij}(k - l\delta t) \\ \forall l \in [0, \dots, \tau_{ij}], \quad \forall i \in H, \quad \forall j \in H_i, \quad \forall r \in \mathcal{R}[k] \cup \mathcal{R}[k-1] \end{aligned} \quad (4.15)$$

$$\begin{aligned} \chi_{ri}(1) = \phi_{ri}^+(1), \quad \xi_{rij}(-l) = 0 \\ \forall l \in [0, \dots, \tau_{ij}], \quad \forall i \in H, \quad \forall j \in H_i, \quad \forall r \in \mathcal{R}[k] \cap \mathcal{R}[k-1] \end{aligned} \quad (4.16)$$

$$\psi_{ri}(1) = \phi_{ri}^-(1), \quad \forall i \in H, \quad \forall r \in \mathcal{R}[k] \quad (4.17)$$

$$\begin{aligned} \forall i \in H, \quad \forall \kappa \in [1, \dots, T_p - 1] : \\ \chi_{ri}(\kappa+1) = \chi_{ri}(\kappa) - \sum_{j \in H_i} \xi_{rij}(\kappa) + \sum_{j \in H_i} \xi_{rji}(\kappa - \tau_{ji}) + \phi_{ri}^+(\kappa+1) - \\ \varepsilon_{ri}(\kappa) \quad \forall r \in \mathcal{R}, \end{aligned} \quad (4.18)$$

$$\rho_i(\kappa+1) = \rho_i(\kappa) - \sum_{j \in H_i} v_{ij}(\kappa) + \sum_{j \in H_i} v_{ji}(\kappa - \tau_{ji}), \quad (4.19)$$

$$\psi_{ri}(\kappa+1) = \psi_{ri}(\kappa) - \varepsilon_{ri}(\kappa) + \phi_{ri}^-(\kappa+1) \quad \forall r \in \mathcal{R}, \quad (4.20)$$

$$\sum_{r \in \mathcal{R}} \xi_{rij}(\kappa) \leq \gamma_v v_{ij}(\kappa) \quad \forall j \in H_i. \quad (4.21)$$

The FSO optimizes, as given in (4.14), the predicted cost of operating their vehicles and late delivery. The FSO furthermore promotes early transport of containers by adding a fee for transporting later. This is added as a small percentage increase, ω^e , to the travel cost ω_{ij}^v , for all $i \in H$ and $j \in H_i$. H denotes hubs and terminals in the flexible mode's network, $N \subseteq H$; H_i denotes set of nodes reachable from node i by flexible service; ω_{ij}^v denotes the cost of operating a vehicle, regardless of its load, from node $i \in H$ to a node j in the set of nodes connected to i by the FSO's transport network H_i . At each location except the destination of the request, a stacking and (un)loading cost, ω^s , is applied to each container each timestep. This cost is an average estimate based on historical data. Furthermore, the stacking cost promotes early departures. d_r is the destination and \hat{c}_r is the fee for late arrival of request $r \in \mathcal{R}$ received from the LSP. In addition to the fee for late arrival, an internal cost of not satisfying demand, ω^d , drives the system.

The last term in the objective function estimates the cost the containers will cause after the end of the plan, this makes plans with shorter planning horizons, T_p , more accurate. It is a well-known technique from MPC [103] that has been adapted to transport contexts, e.g., in [87]. A long planning horizon increases significantly the computational complexity of the optimization problem. The used estimate is optimistic in the sense that it penalizes the containers that are not at their destination at the end of the plan with the minimum remaining travel cost ω_{ij}^{direct} from node i to j . If no direct arc exists, it is the cost of the shortest path between the nodes. It does as such not reflect the true operation cost as that includes repositioning of empty trucks, but it is computationally tractable.

$v_{ij}(\kappa)$ is the number of vehicles leaving node i at timestep κ in the plan towards node j . $\chi_{ri}(\kappa)$ is the number of containers from transport request r that is planned to be stacked at node $i \in H$ at time $k\Delta t + \kappa\delta t$. Loading fees are assumed to be paid by the LSP as part of the service cost. The number of containers from request r that should have been dropped off at node i before time $k\Delta t + \kappa\delta t$ but has not yet arrived is denoted $\psi_{ri}(\kappa)$. τ_{ij} is the travel-time from node $i \in H$ to node $j \in H_i$ including potential (un)loading.

The FSO plans according to the current state of the transport system. Therefore we need to distinguish between planned values and realized values. The latter we denote with a tilde, such that, e.g. $\tilde{\chi}_{ri}(k)$ is the number of containers from trucking request r that is stacked at node i at time $k\Delta t$. The planned number of trucks parked at node i at planning-timestep κ is $\rho_i(\kappa)$, while the realized number parked at decision moment k is $\tilde{\rho}_i(k)$. $\tilde{\psi}_{ri}(k)$ is

the number of containers from request r that is still missing to be delivered after their drop off time at node i at decision moment k . The initial state of the transport system is in (4.15) set to be the current state of the containers the FSO has previously committed to transport but not dropped off and of the vehicles, and in (4.16) to the new transport requests. (4.17) initiates the demand for both new and already picked up requests. The dynamics of the number of containers, vehicles and unsatisfied demand in each node are described by (4.18)-(4.20). $\varepsilon_{ri}(\kappa)$ is the number of containers from request r that arrive at their drop off location, $d_r = i$, at time $k\Delta t + \kappa\delta t$. The number of containers from transport request r that requested to be picked up at node i at time $k\Delta t + \kappa\delta t$ is denoted $\phi_{ri}^+(\kappa)$ and the corresponding requested drop offs are $\phi_{ri}^-(\kappa)$. (4.21) states that containers can only travel over an arc if there is sufficient vehicles to transport them. γ_v is the capacity of the vehicles that operate the arc from node i to j .

4.3.3 Communication

Communication between the LSP and the FSO gives the LSP new possibilities to foresee when containers are dropped off before committing to their transport by the FSO. Instead of splitting shipping requests into bundles based on the containers current locations, the LSP can create bundles based on the FSO's plan. For each transport request, the containers that the FSO expects to drop off in time, constitute one bundle which will keep the request's properties. The containers that are expected to arrive late will be a similar bundle, but with a pick up time limit. This time limit LB_{rs} regards only the transport of containers from the new request bundle r with the flexible service $s \in S^{flexible}$. The LSP then recomputes the plan using (4.1)-(4.13) with the additional constraint (4.22). This process is repeated until the pre-decided number of communication rounds are elapsed and the decisions are implemented.

$$t_{rs}^{depart} \geq LB_{rs} + M(x_{rs} - 1), \quad \forall r \in R, s \in S^{flexible}. \quad (4.22)$$

To provide an example, a shipping request r in the LSP's plan made at decision moment k is scheduled to take a scheduled service from its origin o_r to node i and thereafter a flexible service from Node i to its destination d_r with a pick-up time, $\hat{t}_r^{pick-up}$, and drop-off time, $\hat{t}_r^{drop-off}$. The FSO receives

the transport request from Node i to d_r and plans vehicle and container moments. The LSP is notified that the plan show all containers will arrive late at time t_r^{late} . The LSP subtracts the traveltime they expect for the service (transport from Node i to d_r) and attach the departure time lower bound LB_{rs} to the combination of request bundle r and service s . In the next iteration of the communication round, the LSP's plan may still assign request r to the

Algorithm 4.1 Co-planning

while simulation is running **do**

$k=k+1$

 LSP recieves new orders

 LSP updates active requests $R[k] = R[k-1] \cup R^{new}[k] \setminus \{r | o_r = d_r\}$

 LSP updates schedules $S^{scheduled} = S^{scheduled} \setminus \{s | TD_s < k\Delta t + \delta t\}$

for each communication round **do**

 LSP optimizes container routes (4.1)-(4.13), (4.22)

 LSP communicates corresponding transport requests $\hat{\mathcal{R}}_j^{new}$ to FSO

 where $\forall r \in \hat{\mathcal{R}}_j^{new}$, $\hat{q}_r = q_r$, $\hat{o}_r = \{i | s \in S_{i+}^{flexible}\}$, $\hat{t}_r^{pick-up} = TD_s$,

$\hat{d}_r = \{i | s \in S_{i-}^{flexible}\}$, $\hat{t}_r^{drop-off} = TA_s$, $s = \{s | x_{rs} = 1\}$

 FSO updates transport request information

 FSO optimizes container and vehicle routes (4.14)-(4.21)

 FSO communicates feedback to the LSP

for all $r \in \hat{\mathcal{R}}_j^{new}$ **do**

 LSP splits the request

 FSO updates volume and drop-of time of the request

end for

end for

for All $\kappa = 1, \dots, \frac{1}{\delta t}$ the FSO implements the decisions $\forall r \in \mathcal{R}$: **do**

$\tilde{\xi}_{rij}(k\Delta t + \kappa\delta t) = \xi_{rij}(\kappa)$, $\tilde{v}_{ij}(k\Delta t + \kappa\delta t) = v_{ij}(\kappa) \forall i \in H, j \in H_i$,

end for

 FSO informs LSP on arrivals of containers

 LSP assumes late containers arrive at time $k\Delta t + \delta t$ and splits requests

for departing requests, $r \in \{r \in R | \exists s | x_{rs} = 1 \text{ and } TD_s \leq (k+1)\Delta t\}$

do

$S_r = \{s | x_{rs} = 1 \text{ and } TD_s \leq (k+1)\Delta t\}$

 LSP updates information on expected drop off for each r, s pair

end for

end while

flexible service, just with a later pick-up and drop-off time, or it may show taking another route through the network is better.

After the communication rounds are over, the FSO implements the actions and adjusts the volumes in the new requests to the number of containers that departs before the next plan is made. The drop-of time is also updated to the expected arrival time that is communicated to the LSP. The actions and communication flow of the proposed co-planning method are outlined in Algorithm 4.1.

4.4 Numerical Experiments

To demonstrate the impact of co-planning between an LSP and an FSO, results achieved with the proposed method are compared in numerical experiments to results under the intermodal, passive synchronomodal and active synchronomodal transport paradigms. All experiments are carried out in Matlab using CPLEX to solve the LSP's optimization problem and using Yalmip [95] with Gurobi to solve the FSO's optimization problem. In this section, first a heuristic algorithm that decreases computation time is described. Thereafter, the experimental setup is presented. The details of the case study and additional results are discussed next. Finally, the results of the comparison of the paradigms are discussed under different problem sizes with varying shipping requests.

To compare co-planning with existing transport paradigms, the proposed co-planning method is compared to the intermodal, passive synchronomodal and active synchronomodal methods introduced in Section 4.2.1. In all four methods, the FSO reconsiders all decisions at every decision moment by solving (4.14)-(4.21) and informs the LSP about the expected arrival times upon departures of containers. Under intermodal, passive synchronomodal and active synchronomodal the LSP optimizes (4.1)-(4.13) to update the routing of containers and splits shipping requests into bundles based on the containers' current locations when their routes are to be reconsidered.

4.4.1 Improved computation time

The faster the LSP and the FSO can optimize the routing problems, the more frequent the decision moments can be kept and the more rounds of communication are possible for each decision moment. The heuristic presented by




FIGURE 4.6: *Transport network used in all experiments. The road network is only known by the FSO. The travel times stated in grey are the traveltimes used in the LSP’s optimization problem.*

Guo et al. in [56] is used to optimize the LSP’s routing problem. It pre-processes the data to establish a set of feasible combinations of services for each request. The maximum number of services one request can use is furthermore limited to three, as sensitivity analysis on this parameter in [56] shows three is a good trade-off between optimality and computation time for a similar geographical network.

4.4.2 Experimental setup

In this section, the baseline setup is described which is used unless otherwise stated. The decision moments occur every hour, i.e. $\Delta t = 1$, for all methods and with two rounds of communication for co-planning. With two rounds of communication, the LSP adjusts their plan according to the feedback from the FSO once before the immediate actions are implemented.

Both the case study and the large scale experiments are performed on the geographical network shown in Figure 4.6 using realistic costs and travel times. The top network displays the locations of interest to the LSP, the scheduled services operating between them and the origin-destination pairs of shipping requests. To ease the presentation of results, the FSO provides

TABLE 4.2: *Travel times and costs for truck transport between nodes.*

	1	2	3	4	5	6	7	8	9		1	2	3	4	5	6	7	8	9
1	0	0.2		2.6	3.2		3.5	4		LSP travel time [h]	1	0	17		140	171		187	213
2	0.2	0		2.8	3.3		3.6	4.2			2	52	0		151	176		192	223
3	1.57	1.71	0								3	41	44	0					
4	1.95	2.08	0.93	0							4	50	54	24	0				
5			1.24	0.87	0	0.8		0.9	1.5		5		32	22	0	47		53	84
6				1.28	0.8	0		0.5	0.9		6			33	21	0		32	53
7	2.07	2.19		0.33			0				7	53	56		9		0		
8					0.9	0.5		0			8				23	13		0	
9					1.5	0.9	1.95		0		9				39	23	50		0
	FSO direct travel time [h]										FSO direct travel cost [€]								

unimodal truck transport where one vehicle can carry one container, i.e. $\gamma_v = 1$. The lower network shows the roads with indications of the travel times used in the FSO's planning problem. The road network has additional locations where trucks can park and containers can be stacked.

The travel times, costs and other network parameters are adopted from [56] and based on existing practices. The cost and travel time for road transport are seen in Table 4.2. In the upper triangle of the left matrix are the travel times used in the LSP's planning problem. In the lower triangle are the actual travel times known by the FSO. If no time is indicated, the LSP, respectively the FSO, do not use a direct connection between the indicated nodes in their planning problem. The upper triangle of the right matrix shows the cost the LSP assumes for road transport of a container between the two indicated nodes, while the lower triangle shows the actual cost for transversing the arc with a vehicle. The LSP's estimated travel cost includes (un)loading costs and assumes the FSO's vehicles drive empty 50% of the distance. Transversing an arc takes the same time and costs in both directions. In the beginning of the simulation, i.e. before the first decision moment, no shipping requests are known and all trucks are parked at one node. In the case study, 10 trucks are parked in Node 9, Dortmund, and in the large scale experiments, 20 trucks are parked in Node 1, Maasvlakte I.

The dynamics of FSO's transport system is in the experiments discretized using $\delta t = 0.2$ h and the FSO's optimization problem considers $T_p = 50$ planning timesteps. The baseline delay cost is $\omega^d = 1000$, the storage cost

TABLE 4.3: *Scheduled services available to the LSP. Two vehicles operates each route, departing from either end simultaneously.*

Mode	Routes	Capacity	First departure	Frequency	Travel time [h]	Cost per container [€]
Barge	5-1, 1-5	100	12	24	15	48.11
	6-1, 1-6	100	10	24	18	50.91
	8-6, 6-8	50	52	48	4.5	38.33
Train	5-1, 1-5	100	16	24	6	77.59
	6-1, 1-6	100	38	24	8	98.39
	9-1, 1-9	100	34	48	9	108.78
	5-2, 2-5	100	12	24	6.5	82.79
	6-2, 2-6	100	10	24	8.5	103.58
	8-5, 5-8	50	52	48	3.5	51.6
	9-5, 5-9	50	52	48	4.5	61.99
	9-6, 6-9	50	60	48	3.5	51.6

$\omega^s = 0.2$ and the penalty used to promote earlier actions in the FSO's optimization problem is $\omega^e = 0.001$. The delay cost known to the FSO is the true delay cost of the request $\hat{c}_r = c_r^{delay}$.

In the case study, the shipping request marked with yellow in Figure 4.4 comprises of 10 containers with delay cost $c_r^{delay} = 5$, while the green one comprises of 15 containers with $c_r^{delay} = 10$. The scheduled services between Node 1, 2 and 5 have capacity for 30 containers, while the train between 5 and 8 can transport 20 containers.

In all the large scale experiments, the same collection of 100 shipping requests are used. When fewer requests are used, the subset is chosen based on announce time. All requests are released immediately after they are announced. The attributes of each request are randomly drawn. The probability of a request originating in Rotterdam is 75% with equal split between Node 1 and 2. These requests have even probability for being due at Node 5, 6, 8 and 9. Reversely, there is 25% probability the request origins in either Node 5, 6, 8 or 9 with an even chance between them and have destination in Node 1 or 2, also with an even chance between them. The number of containers in the request q_r is drawn from a uniform distribution between 1 and 19. The difference between the announce time of a request and the following request is drawn from a Poisson distribution with a mean of 45 min. The release

TABLE 4.4: *Additional results for the illustrative case study.*

Paradigm	Actual cost	Change	Estimated cost	Truck utilization	Truck distance	Change
Intermodal	5722€		4761€	2.31%	3250 km	
Passive	5401€	-6%	3894€	17.32%	6712 km	107%
Active	4568€	-21%	3289€	11.26%	4795 km	48%
Co-planning	3351€	-42%	3289€	42.86%	1260 km	-61%

time, $t_r^{release}$, is the decision moment following the announce time and the due time t_r^{due} is 12, 18 or 24 hours thereafter with an even probability. Requests with 12 hour lead time have delay cost $c_r^{delay} = 15$, 18 hours have $c_r^{delay} = 10$, and 24 hours have $c_r^{delay} = 5$ following the assumption that short lead times imply urgency and priority.

The details on the scheduled services used in the large scale experiments are shown in Table 4.3. Each route consists of two locations, between which two identical vehicles travel. The vehicles depart at the same time from either location. Barge services are cheaper than train services, but have longer travel times.

4.4.3 Illustrative case study

The decisions that are implemented in the transport system depend strongly on which paradigm the planning method adhere to. In Section 4.2.2 the paradigm's impacts on the routes of containers and trucks are discussed. In this section we supplement with a discussion of the additional results shown in Table 4.4.

With more flexibility the total cost of operating the transport system decreases 20% from the intermodal paradigm to active synchronomodality. Adding co-planning decreases the cost additionally by 27% of the cost under active synchronomodal transport. In the case study, communication (using co-planning) is as important for decreasing the operation cost as planning flexibility (active synchronomodal instead of intermodal paradigm). The total costs include the spot prices paid for transport of containers on scheduled services, loading and unloading to both flexible and scheduled services, driving full and empty trucks, intermediate stacking of containers in the truck network, and penalties paid by the LSP for late delivery of containers.

The LSP expects a fixed cost of c_s , $s \in S^{\text{flexible}}$ per container to use a flexible service between the pick-up and drop-off location. The true operation cost of the service is more complex, as the FSO may have to drive empty to the pick up location. There is therefore a significant gap between the cost the LSP expects and the actual cost. In intermodal transport, the mode choice does not change, and the gap is a moderate 178% while it for both active and passive synchromodal transport is 30%. Co-planning reduces the gap significantly to only 2% and shows the benefit of communication clearly.

Co-planning also improves the truck utilization significantly, i.e. trucks more often drive full compared to the total distance they drive. The total distance is, furthermore, significantly lower than the other paradigms (as discussed in Section 4.2.2). It is noteworthy that the utilization rate under passive synchromodality is higher than under active at the same time the distance driven is also higher, so the distance the trucks drive empty is higher under passive than active synchromodal transport. This is due to that only shipping requests that cannot make their next planned service are replanned under passive synchromodality, while all shipping requests are replanned at every decision moment under active synchromodality. So while it looks like the efficiency of the truck fleet is better under passive synchromodal transport, the cost of road transport and its environmental impact is better under the active synchromodal paradigm. In the case study, the flexibility and coordination in co-planning outperforms the other paradigms.

4.4.4 Paradigm comparison

To understand the value of co-planning in more realistic scenarios, numerical studies on instances with varying number of shipping requests have been carried out. The results in Table 4.5 show that with increased flexibility, the operational cost of the transport system decreases when the number of shipping requests starts filling up the network capacity. When the number of shipping requests is low (e.g., five requests), there is no difference between intermodal, passive and active synchromodal transport. In these cases, co-planning achieves improvements. When more shipping requests are being considered, the LSP's estimated cost is higher than the actual operation cost, while they underestimate the cost when few requests are planned for. The gaps between the estimated and actual costs are on average the lowest under the intermodal paradigm, i.e. the correlation between operation cost and transport prices are more stable under this paradigm. Co-planning improves

TABLE 4.5: Comparison of transport under the four different paradigms for varying numbers of shipping requests.

Paradigm	Number of Requests	Actual cost per container [€]	Overestimate by LSP	Delay time per container [h]	Mean delaycost per lateness [€]	Barge share	Train share	Truck share	Truck distance [km]	Truck utilization	CPU [s]
Intermodal	100	1450	2.2%	32	10.14	21%	56%	23%	105756	47%	1413
	75	1425	3.2%	30	10.28	21%	56%	23%	75940	52%	1064
	50	1535	-1.1%	37	10.58	24%	60%	16%	44668	41%	740
	25	1324	-1.6%	21	10.93	32%	47%	21%	28050	47%	513
	10	1116	-6.6%	19	10.01	68%	18%	14%	9355	37%	105
	5	1429	-11.6%	25	11.38	54%	33%	13%	6030	30%	161
	Mean	1380		27	10.54	37%	45%	18%	44966	42%	
Passive	100	1264	6.3%	18	9.48	21%	51%	28%	112537	53%	1995
	75	1231	6.8%	15	9.28	21%	52%	27%	79297	57%	1597
	50	1271	1.2%	16	9.65	24%	55%	21%	46871	49%	1181
	25	1213	-0.6%	14	9.31	32%	46%	21%	28227	48%	549
	10	1116	-6.6%	19	10.01	68%	18%	14%	9355	37%	107
	5	1429	-11.6%	25	11.38	54%	33%	13%	6030	30%	162
	Mean	1254		18	9.99	37%	43%	21%	47052	46%	
Active	100	1221	7.8%	20	10.07	19%	55%	25%	98509	55%	1000
	75	1204	7.3%	16	9.97	20%	56%	25%	74189	59%	702
	50	1224	1.6%	16	10.31	22%	61%	17%	42135	51%	383
	25	1158	7.2%	15	9.56	32%	41%	27%	28395	57%	274
	10	1116	-6.6%	19	10.01	68%	18%	14%	9355	37%	173
	5	1429	-11.6%	25	11.38	54%	33%	13%	6030	30%	155
	Mean	1225		19	10.31	36%	44%	20%	43102	48%	
Co-planning	100	1228	6.5%	20	10.12	19%	57%	24%	97420	44%	2427
	75	1204	7.2%	16	9.91	20%	56%	24%	73775	49%	1506
	50	1220	1.5%	16	10.32	22%	61%	17%	41705	45%	629
	25	1164	7.4%	15	9.67	32%	41%	27%	28300	47%	297
	10	1110	-6.1%	19	10.03	68%	18%	14%	9170	31%	196
	5	1416	-10.9%	25	11.38	54%	33%	13%	5845	21%	154
	Mean	1224		19	10.33	36%	44%	20%	42703	39%	

the LSP's estimations compared to active synchronomodal transport without any cooperation. The overestimation by the LSP depends on the actual truck utilization and what average utilization is used to compute the transport prices that appear in the LSP's optimization problem.

Surprisingly, the average time delay per container and the total delay time are the least under the passive synchronomodal paradigm. Here, the average delay cost per hour of the containers that were late is similar to the average of all shipping requests assigned delay cost. The mean delay cost indicates that the high-priority requests with short lead time and high delay cost also arrive late. The good performance of passive synchronomodality on this indicator may be because the scheduled services get fully booked by requests with long lead times before the urgent requests are announced as the release of containers follows immediately after the announcement. The distance travelled by truck compared to the total transport distance (the modal share for truck) is the highest for passive synchronomodal transport, hence a larger part of the urgent containers, than in under the other paradigms, may be initially assigned to direct truck transport and thus not miss connections.

The total operating costs of the active synchronomodal transport network and that with co-planning are similar in all instances. The need for repositioning of empty trucks is higher in the beginning of the simulations, causing the truck utilization to be lower for instances with less requests as the concentration of shipping requests per time is comparable for the periods with active requests in all instances. It is expected that co-planning will have a larger impact when trucks need to be repositioned more frequently. The scenario used for the experiments was not designed to encourage this. In all the instances, the distance driven by trucks is lower when co-planning is used, even though the truck utilization in many cases is lower. This indicates that, as expected, the limited truck fleet is used wiser when the LSP gets basic feedback. The communication increases the computation time of co-planning significantly, especially in instances with more shipping requests. Surprisingly, the computation time is lowest under the active synchronomodal paradigm. This is unexpected, as more shipping requests are considered at each time step than under intermodal and passive synchronomodal paradigm where only infeasible requests are reconsidered. However, the results show that the splitting of bundles increases the number of transport requests that the FSO needs to reconsider more for intermodal and passive synchronomodal than for active synchronomodal and co-planning. The complexity of the LSP's

optimization problem increases with increased flexibility, but the size of the FSO's optimization problem decreases.

4.5 Conclusions

Cooperation under different transport paradigms sets different requirements to the coordination and cooperation between the stakeholders in transport systems. In this chapter, we present a method where a logistics service provider (LSP) and an operator of a flexible vehicle fleet (FSO) co-plan in real-time and compare its impact on a transport system with that of planning methods without communication under intermodal and synchronomodal transport paradigms. The results show that when the LSP's decisions can be reconsidered as part of a synchronomodal paradigm, significant economical savings and reduced lateness can be achieved compared to the less flexible intermodal paradigm.

The chapter addresses Research Question Q3 *What is the impact of stakeholders planning cooperatively at the operational level?* The case study shows that in some instances, the operational cost and overall efficiency of a transport system can improve significantly when the LSP uses feedback on expected arrival times to reconsider container routes. In large scale experiments, the economical benefits were limited as less repositioning of empty trucks were necessary to satisfy the demand. In all cases co-planning improves the LSP's ability to predict the operation costs of the transport network and decreases the distance driven by truck significantly without changing the modal share for trucks significantly. The FSO's operational costs decrease when the LSP adjusts the routes and departure times of containers that overburdens the FSO's service network.

The chapter, furthermore, answers Research Question Q4 *How can containers and vehicles be routed cooperatively through a synchronomodal network, if only traditional transport requests and their expected fulfilment are communicated?*. The proposed method relies on automation of the traditional communication between an LSP and an FSO. The LSP sends transport requests with container quantity, pick up and drop off information and the LSP responds to these requests with quantity and expected arrival times for the part of the transport request that can be fulfilled as requested and the part that will be late. Co-planning using this communication scheme is very realistic for implementation in practice.

In this chapter, we have considered co-planning between two stakeholders that both takes a multitude of decisions. In Chapter 6, we will address co-planning between a truck operator that plans container and truck moves similarly to what is done in Chapter 3 and a barge operator. On one hand, barge operator's decision changes the dynamics of the transport system drastically. On the other hand, these changes (departures of a barge) do not happen often. Therefore, the transport system can in such cases be modelled as a switched linear system. In the next chapter, Chapter 5, we lay the theoretical foundation for using Bayesian optimization as a parallelization heuristic to improve the computation time for MPC on switched linear systems.

Chapter 5

Learning Discrete Actions over Time

In the previous two chapters, we improve the computation time of the methods by approximating integer variables with their continuous counterparts. Controlling systems with both continuous and discrete actuators using model predictive control is often impractical, since mixed-integer optimization problems are too complex to solve sufficiently fast. When the integer actuators have a limited number of possibilities, the system can be modelled as a switched system. This chapter addresses Research Question Q5 by proposing and analysing a parallelizable method to control both the continuous input and the discrete switching signal for linear switched systems. The method uses ideas from Bayesian optimization to limit the computation to a predefined number of convex optimization problems.

The core of this chapter has been presented at *IFAC World Congress*¹. It is organized as follows: in Section 5.2, the considered system and problem is specified. In Section 5.2.1, MPC with MDS is presented together with an analysis of stability and recursive feasibility. In Section 5.3, MPC with MDS is assessed using a simulation and compared with three other controllers. Finally, Section 5.4 draws the main conclusions.

¹Rie B. Larsen, Bilge Atasoy, Rudy R. Negenborn. Model Predictive Control with Memory-based Discrete Search for Switched Linear Systems. *IFAC-PapersOnLine for IFAC World Congress*, volume 53, pages 6769-6774, 2020.

5.1 Introduction

Many systems are controlled by a combination of discrete and continuous actuators. A class of them can be described by switched linear dynamics without dwell-time where both the continuous input and the discrete switching signals are decisions taken by the applied controller. For very small systems, model predictive control (MPC) can be applied directly to improve the performance of these systems under constraints. However, since the system must be described using both continuous and discrete variables, the MPC controller must solve a mixed-integer program every time the control sequence is updated. For larger systems, the computation of this mixed-integer program may be unacceptably long compared to the system specifications.

Switched linear systems can be categorized as either internally or externally forced. An overview of stability analysis for different classes of switched linear systems can be found in [91], while [183] reviews the literature on MPC for switched systems. The literature on MPC for internally forced systems focuses mainly on the case where the switching signal is state dependent (piecewise affine systems) as in, e.g., [109], or on systems under uncontrolled switching, as in, e.g., [180]. This chapter considers externally forced systems, hence systems where the applied MPC must decide on both the continuous inputs and the discrete switching signals.

To improve the computation time of MPC for externally forced switched linear systems, most frameworks focus on solving each iteration faster. Often the proposed integer solvers are iteration based, e.g., [7] or [115], which limits how fast the computation can be and increases the computation time uncertainty. This complicates the implementation of MPC, since a typical MPC re-computes the input at predefined time intervals. One way to avoid performing iterative computations of the discrete variables is using convex relaxations of the mixed-integer problem. In [141] it is shown that continuous variables can be used to approximate binary variables in the MPC optimization problem if the timesteps are sufficiently small. In [44] this result is used to control a switched linear system without other input than the switching signal. In [128] the switching signals are modelled as switching times resulting in a non-linear optimization problem with solely continuous variables. [104] decrease the number of iterations needed to solve the mixed-integer problem by distributing the computation between different agents coupled by auxiliary variables.

Another stream of research, namely that on suboptimal MPC, acknowledge that it is not always possible to obtain the optimal solution. This research establishes bounds and criteria on the suboptimality and the properties of the implemented control law that ensures stability and feasibility of the system. Stability and recursive feasibility of the suboptimal MPC is in, e.g., [84] and [122] established by imposing specific improvements in the objective function. In contrast, [146] ensure stability and recursive feasibility by improving a known solution which satisfy stability and feasibility criteria.

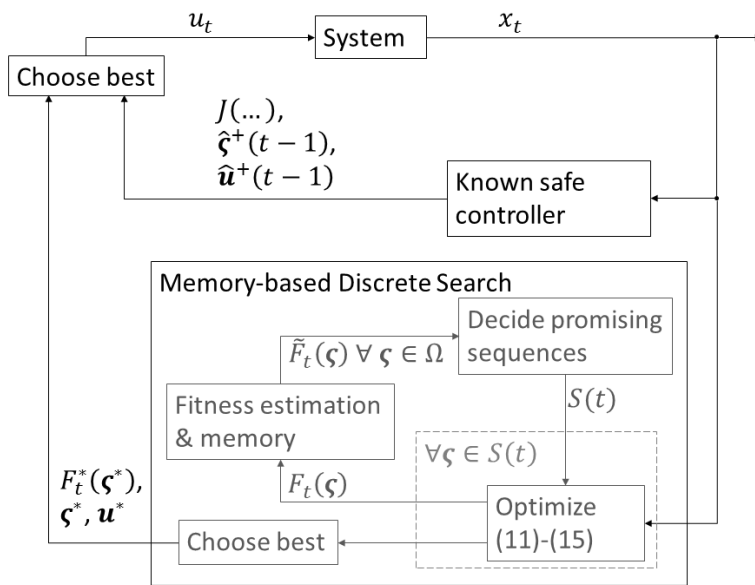


FIGURE 5.1: *Schematics of MPC with MDS.*

In this chapter, the proposed method finds better solutions by remembering the performance at previous timesteps, and is thus called MPC with Memory-based Discrete search (MDS). The method is outlined in Figure 5.1 with notation as introduced in the remainder of this chapter. MPC with MDS solves the part of the usual MPC problem, which is continuous in the variables, for a set of switching signal sequences. It then implements the first elements of the sequence with the best performance, if this performance is better than that of a stabilizing input sequence computed based on the previous MPC prediction and a known conservative control law. If the known stabilizing input sequence performs better, its solution corresponding to the

current timestep is implemented and the process starts over.

It is assumed that the performance of a switching signal sequence at a given time is comparable to the performance of related sequences at the previous timestep. This assumption is used to choose the switching signal sequences for which MPC with MDS computes the convex optimization problem. The ideas are analogous to Bayesian Optimization where the current knowledge of a function is used to assess which function evaluation to perform at the next iteration. See e.g. [68] or [151] for an introduction. MPC with MDS utilizes the rolling horizon from MPC to obtain information about not only current solutions, but also future solutions. In this fashion, the optimization problems in MPC with MDS has low complexity and can be solved in parallel.

5.2 System and problem definition

We consider a discrete time, switched, linear system without dwell-time with the following dynamics:

$$x_{t+1} = A_{\sigma_t}x_t + B_{\sigma_t}u_t, \quad (5.1)$$

where $x_t \in \mathbb{R}^n$ and $u_t \in \mathbb{R}^{m_{\sigma_t}}$ are the state of and the input applied to the system at discrete time t . The switching signal $\sigma_t \in \mathbb{S} = \mathbb{I}_{[1:l]}$ is the input at time t , that determines from finite sets, the dynamic matrices $A_{\sigma_t} \in \mathcal{A} = \{A_1, A_2, \dots, A_l\}$ and $B_{\sigma_t} \in \mathcal{B} = \{B_1, B_2, \dots, B_l\}$. $\mathbb{I}_{[a;b]}$ denotes the set of integers $\{x \in \mathbb{I} | a \leq x \leq b\}$. The system is controllable for one or more switching signals.

Definition 5.1 Basic switching signal and dynamics

The basic switching signal is denoted by γ . It is one switching signal $\gamma \in \mathbb{S}$ for which system (5.1) is controllable. The corresponding dynamics $x_{t+1} = A_{\gamma}x_t + B_{\gamma}u_t$ will be referred to as the basic dynamics. \square

The system's state and inputs are constrained by

$$x_t \in \mathbb{X} \subseteq \mathbb{R}^n \quad \forall t \quad (5.2)$$

$$u_t \in \mathbb{U}_{\sigma_t} \subseteq \mathbb{R}^{m_{\sigma_t}} \forall t. \quad (5.3)$$

The set \mathbb{X} is convex, closed and contains the origin in its interior, while each set \mathbb{U}_{σ_t} is convex and compact. \mathbb{U}_{γ} contains the origin.

It is assumed that more information about the basic dynamics is known.

Assumption 5.1 *A control law $u = \kappa(x)$, called the safe control law, and a set \mathcal{X}_f are known for which the following is true:*

1. *The basic dynamics under the safe control law, i.e. $x_{t+1} = A_\gamma x_t + B_\gamma \kappa(x_t)$ is asymptotically stable.*
2. *$\mathcal{X}_f \subseteq \mathbb{X}$, $0 \in \mathcal{X}_f$ and \mathcal{X}_f is closed.*
3. *$\kappa(x) \in \mathbb{U}_\gamma \forall x \in \mathcal{X}_f$.*
4. *$A_\gamma x + B_\gamma \kappa(x) \in \mathcal{X}_f \forall x \in \mathcal{X}_f$.*

5.2.1 Problem definition

We consider stability to the equilibrium at the origin, but the results can be generalized to set point stability.

When system (5.1) is controlled by an MPC controller, the controller solves at timestep t the following mixed-integer program:

$$\begin{aligned} \min \quad & J(\boldsymbol{\zeta}, \mathbf{u}, \mathbf{x}) & (5.4) \\ & u_0, \dots, u_{N-1} \\ & \zeta_0, \dots, \zeta_{N-1} \end{aligned}$$

$$J(\boldsymbol{\zeta}, \mathbf{u}, \mathbf{x}) = \sum_{k=0}^{N-1} (x_k^T Q x_k + u_k^T R_{\zeta_k} u_k) + x_N^T Q_N x_N \quad (5.5)$$

$$\text{s.t.} \quad x_0 = x_t \quad (5.6)$$

$$x_{k+1} = A_{\zeta_k} x_k + B_{\zeta_k} u_k \quad \forall k \in \mathbb{I}_{[0;N-1]} \quad (5.7)$$

$$x_k \in \mathbb{X}, u_k \in \mathbb{U}_{\zeta_k}, \forall k \in \mathbb{I}_{[0;N-1]} \quad (5.8)$$

$$x_N \in \mathcal{X}_f, \quad (5.9)$$

where variables $x_k \in \mathbb{R}^n$, $\zeta_k \in \mathbb{I}_{[1;l]}$ and $u_k \in \mathbb{R}^{m_{\zeta_k}}$ represent the controller's model of the state, switching signal and input to the system at prediction time $t + k$, respectively. N is the prediction horizon. The cost matrices R and Q_N are symmetric and positive definite, while Q is symmetric and at least positive semi-definite. Bold symbols denote ordered time sequences, e.g., $\mathbf{a}_{[0;N]} = \langle a_0, a_1, \dots, a_N \rangle$, where each element a_k is a vector related to timestep k . When appropriate, the interval subscript is omitted for a simpler notation.

Assumption 5.2 *The stage cost for the basic dynamics $l(x, u, \gamma) = x_k^T Q x_k + u_k^T R_\gamma u_k$ and the final cost $V_f(x) = x^T Q_N x$ are such that $V_f(A_\gamma x + B_\gamma \kappa(x)) - V_f(x) + l(x, \kappa(x), \gamma) \leq 0 \forall x \in \mathcal{X}_f$.*

Notice that Assumption 1 and 2 can be satisfied by the design of $\kappa(x)$, \mathcal{X}_f and $V_f(x)$.

The time it takes to solve (5.4)-(5.9) often limits which switched linear systems can realistically be controlled with MPC. Most solvers are iteration based, and thus not suitable for parallel implementations. Furthermore, current efforts typically do not take the recursive nature of MPC into account.

Instead of solving a slow mixed-integer problem at each timestep, we suggest to search over the possible sequences of switching signals based on their previous performance and only solve the MPC problem for the chosen sequences. We call the method MPC with Memory-based Discrete Search (MDS). For given switching signal sequences, the optimization problem corresponding to each sequence is continuous in the variables, and hence solvable by fast, convex optimization solvers. When the set of switching signals, for which the MPC problem should be solved, depends on the method's memory of previous performances, the optimizations can be performed in parallel to decrease the computation time further.

MPC with MDS searches for good switching signal sequences based on previous performance. Since the controlled system's state changes dynamically, the expected performance of a given sequence will depend on the previous performance of sequences that are time-shifted and similar, not identical to that sequence. MPC with MDS is thus very suitable for systems that require long prediction horizons, as more information about a sequence's potential performance can be gathered before the dynamic information becomes obsolete.

To choose for which set of switching signal sequences to solve the MPC problem, MPC with MDS uses ideas from Bayesian optimization to ensure the set contains both sequences that are expected to perform well, and sequences that will provide new information. To measure the performance and information value of a switching signal sequence, the sequence's fitness $F_t(\boldsymbol{\varsigma}_{[0;N-1]})$, expected fitness $\tilde{F}_t(\boldsymbol{\varsigma}_{[0;N-1]})$ and uncertainty value $Y_t(\boldsymbol{\varsigma}_{[0;N-1]})$ are introduced.

Definition 5.2 Fitness

The fitness of switching signal sequence $\boldsymbol{\varsigma}_{[0;N-1]} = \{\varsigma_0, \dots, \varsigma_{N-1}\}$ is

$$F_t(\boldsymbol{\varsigma}_{[0;N-1]}) = \min_{u_1, \dots, u_{N-1}} J(\boldsymbol{\varsigma}_{[0;N-1]}, \mathbf{u}, \mathbf{x}) \quad (5.10)$$

$$\text{s.t.} \quad x_0 = x_t \quad (5.11)$$

$$x_{k+1} = A_{\varsigma_k} x_k + B_{\varsigma_k} u_k \quad \forall k \in \mathbb{I}_{[0;N-1]} \quad (5.12)$$

$$x_k \in \mathbb{X}, u_k \in \mathbb{U}_{\varsigma_k}, \forall k \in \mathbb{I}_{[0;N-1]} \quad (5.13)$$

$$x_N \in \mathcal{X}_f \quad (5.14)$$

where the notation corresponds to that of (5.4)-(5.9). \square

Definition 5.3 Expected fitness

The expected fitness for a sequence $\boldsymbol{\varsigma}$ at time t is denoted by $\tilde{F}_t(\boldsymbol{\varsigma})$ and its uncertainty value is denoted by $\tilde{Y}_t(\boldsymbol{\varsigma})$. \square

Figure 5.1 presents a schematics of MPC with MDS. The detailed notation will be provided in the remainder of this section. At each timestep t , fitness estimates of all switching signal sequences are used to determine a set of promising sequences. For each promising sequence, the fitness is evaluated for the current system state. The best performance of the promising sequences is then compared to the performance of a known controller. This controller knows a switching signal sequence and a corresponding continuous input sequence which satisfy all constraints and stabilizes the system. The first switching signal and continuous input corresponding to the sequence that performs the best are implemented on the system and the process starts over at the next timestep.

How the promising sequences are chosen does not affect the stability of a system controlled by MPC with MDS. Therefore, the stability and recursive feasibility of the proposed method are proven in Section 5.2.2 before the selection method is presented in Section 5.2.3.

5.2.2 Stability and feasibility

To guarantee stability and recursive feasibility of the MPC problem, MPC with MDS compares the performance of the potential sequences with the performance of a control law, which is known to satisfy all constraints. This section outlines how this is done, and proves that the implemented inputs will stabilize the system.

The algorithm knows a feasible solution at time t , but needs to access the performance of the potential sequences at time $t + 1$. It is thus necessary to define how the solution at time t is shifted in time and to prove that the shifted solution is indeed feasible.

Definition 5.4 Shifted sequences

The sequences $\boldsymbol{\zeta}_{[0;N-1]}^+ = \{\boldsymbol{\zeta}_{[1;N-1]}, \gamma\}$, $\mathbf{u}_{[0;N-1]}^+ = \{\mathbf{u}_{[1;N-1]}, \kappa(x_N)\}$ and $\mathbf{x}_{[0;N]}^+ = \{\mathbf{x}_{[1;N]}, A_\gamma x_N + B_\gamma \kappa(x_N)\}$ are said to be the shifted sequences of $\boldsymbol{\zeta}_{[0;N-1]}$, $\mathbf{u}_{[0;N-1]}$ and $\mathbf{x}_{[0;N]}$. \square

The outer framework of MPC with MDS can be seen in Algorithm 5.1. The method is initialized with a feasible solution to the mixed-integer problem (5.4)-(5.9) for the initial state. In the algorithm the best solution that is known to the controller at a given time is marked with a hat. The first

Algorithm 5.1 MPC with MDS

- 1: Let $\hat{\boldsymbol{\zeta}}(t)$, $\hat{\mathbf{u}}(t)$ and $\hat{\mathbf{x}}(t)$ be a feasible solution to (5.4)-(5.9)
 - 2: **while** system is to be controlled **do**
 - 3: Implement $\boldsymbol{\sigma}_t = \hat{\boldsymbol{\zeta}}_0$, $u_t = \hat{u}_0$
 - 4: $t = t + 1$
 - 5: Measure x_t
 - 6: Compute $\mathcal{S}(t)$ using Algorithm 5.2
 - 7: **for** all $\boldsymbol{\zeta} \in \mathcal{S}(t)$ **do**
 - 8: **if** (5.10)-(5.14) is feasible **then**
 - 9: $\tilde{F}_t(\boldsymbol{\zeta}) = F_t(\boldsymbol{\zeta})$
 - 10: **else**
 - 11: $\tilde{F}_t(\boldsymbol{\zeta}) = M$
 - 12: **end if**
 - 13: $\tilde{Y}_t(\boldsymbol{\zeta}) = \beta J(\hat{\boldsymbol{\zeta}}^+(t-1), \mathbf{u}^+(t-1), \mathbf{x}^+(t-1))$
 - 14: **end for**
 - 15: Find $F_t^*(\boldsymbol{\zeta}^*) = \min_{\boldsymbol{\zeta} \in \mathcal{S}(t)} F_t(\boldsymbol{\zeta})$
 - 16: **if** $F_t^*(\boldsymbol{\zeta}^*) \leq J(\hat{\boldsymbol{\zeta}}^+(t-1), \mathbf{u}^+(t-1), \mathbf{x}^+(t-1))$ **then**
 - 17: $\hat{\boldsymbol{\zeta}}(t) = \boldsymbol{\zeta}^*$, $\hat{\mathbf{u}}(t) = \mathbf{u}^*$, $\hat{\mathbf{x}}(t) = \mathbf{x}^*$
 - 18: **else**
 - 19: $\hat{\boldsymbol{\zeta}}(t) = \hat{\boldsymbol{\zeta}}^+(t-1)$, $\hat{\mathbf{u}}(t) = \hat{\mathbf{u}}^+(t-1)$, $\hat{\mathbf{x}}(t) = \hat{\mathbf{x}}^+(t-1)$
 - 20: **end if**
 - 21: **end while**
-

elements of the best known switching signal sequence and continuous input sequence are implemented as is usually done with MPC. Hereafter MPC with MDS decides on the set of potential sequences. How the method computes $\mathcal{S}(t)$ in line 6 will be detailed in Section 5.2.3. For feasible sequences, the expected fitness is updated to be the actual fitness, while infeasible sequences' expected fitnesses are penalized with a large number M . The uncertainty value is in both cases updated to be a factor, $\beta > 0$, times the objective function value of the best known solution. This ensures that the uncertainty values are of reasonable size compared to the fitness values. If the best of the potential sequences is better than the best known solution, that potential sequence becomes the new best solution known by the controller.

Lemma 5.1 *For any switching signal sequence $\hat{\boldsymbol{\zeta}}_{[0;N-1]}$ that has a feasible solution to (5.10)-(5.14) at time t , the shifted sequence $\hat{\boldsymbol{\zeta}}_{[0;N-1]}^+$ will have a feasible solution to (5.10)-(5.14) at time $t + 1$. \square*

Proof: Lemma 5.1 follows directly from (5.14) and Assumption 5.1. \square

With the outer framework of MPC with MDS fully defined, we now analyse the stability of a system under Algorithm 5.1 using standard techniques, see, e.g., [102].

Theorem 5.2 *If a feasible solution to (5.4)-(5.9) exists for System (5.1) at time $k = k_0$, then the system under Algorithm 5.1 converges to the origin.*

Proof: To ensure stability, first recursive constraint satisfaction must be established.

Recursive Feasibility: If (5.10)-(5.14) is infeasible at any time t for all $\boldsymbol{\zeta} \in \mathcal{S}(t)$ the sequences from $t - 1$ are shifted and applied by Algorithm 5.1. The recursive feasibility of the shifted sequences is given by Lemma 5.1.

Stability: Sufficient criteria for stability of System (5.1) subject to Algorithm 5.1 are: A) $J(\hat{\boldsymbol{\zeta}}(t + 1), \hat{\boldsymbol{u}}(t + 1), \hat{\boldsymbol{x}}(t + 1)) < J(\hat{\boldsymbol{\zeta}}(t), \hat{\boldsymbol{u}}(t), \hat{\boldsymbol{x}}(t)) \forall t \in \{t | J(\hat{\boldsymbol{\zeta}}(t), \hat{\boldsymbol{u}}(t), \hat{\boldsymbol{x}}(t)) > 0\}$ and B) $J(\hat{\boldsymbol{\zeta}}(t + 1), \hat{\boldsymbol{u}}(t + 1), \hat{\boldsymbol{x}}(t + 1)) = 0 \forall t \in \{t | J(\hat{\boldsymbol{\zeta}}(t), \hat{\boldsymbol{u}}(t), \hat{\boldsymbol{x}}(t)) = 0\}$, because $J(\hat{\boldsymbol{\zeta}}(t), \hat{\boldsymbol{u}}(t), \hat{\boldsymbol{x}}(t)) = 0$ only at the origin and $\boldsymbol{\sigma}_t = \hat{\boldsymbol{\zeta}}_0$, $\boldsymbol{u}_t = \hat{\boldsymbol{u}}_0 \forall t$. Due to Line 16 in Algorithm 5.1 and Assumption 5.1, $J(\hat{\boldsymbol{\zeta}}(t + 1), \hat{\boldsymbol{u}}(t + 1), \hat{\boldsymbol{x}}(t + 1)) \leq J(\hat{\boldsymbol{\zeta}}^+(t), \hat{\boldsymbol{u}}^+(t), \hat{\boldsymbol{x}}^+(t)) \forall t$. Using the notation of Assumption 5.2, criterion A) thus corresponds to $V_f(A_\gamma \hat{\boldsymbol{x}}_N(t) + B_\gamma \boldsymbol{\kappa}(\hat{\boldsymbol{x}}_N(t))) - V_f(\hat{\boldsymbol{x}}_N(t)) +$

$l(\gamma, \kappa(\hat{x}_N(t)), \hat{x}_N(t)) - l(\hat{\zeta}_0(t), \hat{u}_0(t), \hat{x}_0(t)) < 0$. The time-argument is omitted for the rest of the proof. Due to convexity $l(\hat{\zeta}_0, \hat{u}_0, \hat{x}_0) > 0$ except at the origin, where $l(\hat{\zeta}_0, \hat{u}_0, \hat{x}_0) = 0$. Both criteria A) and B) are thus fulfilled if $V_f(A_\gamma \hat{x}_N + B_\gamma \kappa(\hat{x}_N)) - V_f(\hat{x}_N) + l(\gamma, \kappa(\hat{x}_N), \hat{x}_N) \leq 0 \forall t$ which is ensured by Assumption 5.2. \square

Notice that the strategy used to update $\mathcal{S}(t)$ has no influence on the feasibility and stability properties of Algorithm 5.1. It only affects the quality of the solution and the computation time. Notice furthermore that infeasible solutions and suboptimal solutions to (5.10)-(5.14) do not impact the stability and feasibility of MPC with MDS.

5.2.3 Sequence selection

Several methods can be used to select the set of potential sequences. MPC with MDS estimates, like in Bayesian optimization, whether a sequence is likely to be either the best solution (exploitation) or bring new information (exploration). The set of potential sequences is assembled to reflect a trade-off between exploitation and exploration. In the following it is assumed that N and size l of \mathbb{S} are sufficiently small, that saving and sorting information on each l^N switching signal is realistic. If l and N are large, sampling methods and parallelization can be used to represent the switching signals.

When the fitness and uncertainty value of a sequence is estimated, only information from the previous timestep is used. It is thus necessary to establish what switching signal sequences from the previous time $t - 1$ will have influence at a given sequence at time t . For this, we define neighbour sequences and implemented ancestors.

Definition 5.5 Neighbours

Two switching signal sequences ζ^a and ζ^b are neighbours, if they differ at only one timestep except for the last timestep, that is $\zeta^a \in \{\zeta | \zeta_i = \zeta_i^b \forall i \in \mathbb{I}_{[0;N-2]} \setminus \{j\}, \text{ where } j \in \mathbb{I}_{[0;N-2]}\}$. Switching signal sequence ζ^a has o_{ζ^a} neighbours. \square

Definition 5.6 Ancestors

A switching sequence ζ 's implemented ancestor $\mathbf{a}(\zeta)$ at time t is a switching sequence whose first element is applied to the system at time $t - 1$, and whose remaining elements are identical with the first $N - 2$ elements in the sequence in question, i.e. switching signal sequence ζ has implemented ancestor $\mathbf{a}(\zeta) = \{\Psi | \sigma_{t-1} = \Psi_0, \Psi_{[1;N-1]} = \zeta_{[0;N-2]}\}$. \square

Algorithm 5.2 Deciding $S(t)$

```

1: for all  $\boldsymbol{\varsigma} \in \Omega$  do
2:    $\tilde{F}_t(\boldsymbol{\varsigma}) = \alpha \tilde{F}_{t-1}(\mathbf{a}(\boldsymbol{\varsigma})) + \frac{1-\alpha}{o_{\boldsymbol{\varsigma}}} \sum_{\boldsymbol{\psi} \in \mathcal{N}_{\boldsymbol{\varsigma}}} \tilde{F}_{t-1}(\boldsymbol{\psi})$ 
3:    $\tilde{Y}_t(\boldsymbol{\varsigma}) = \alpha \tilde{Y}_{t-1}(\mathbf{a}(\boldsymbol{\varsigma})) + \frac{1-\alpha}{o_{\boldsymbol{\varsigma}}} \sum_{\boldsymbol{\psi} \in \mathcal{N}_{\boldsymbol{\varsigma}}} \tilde{Y}_{t-1}(\boldsymbol{\psi})$ 
4: end for
5:  $S(t) = \emptyset$ 
6: while  $S(t)$  is not full do
7:    $S(t) = S(t) \cup \arg \min_{\boldsymbol{\varsigma} \in \Omega \setminus S(t)} \tilde{F}_t(\boldsymbol{\varsigma})$ 
8:    $S(t) = S(t) \cup \arg \min_{\boldsymbol{\varsigma} \in \Omega \setminus S(t)} \tilde{F}_t(\boldsymbol{\varsigma}) - \tilde{Y}_t(\boldsymbol{\varsigma})$ 
9:    $S(t) = S(t) \cup \arg \max_{\boldsymbol{\varsigma} \in \Omega \setminus S(t)} \tilde{Y}_t(\boldsymbol{\varsigma})$ 
10:   $S(t) = S(t) \cup \{\boldsymbol{\psi}\}$  where  $\boldsymbol{\psi} \in \Omega \setminus S(t)$ 
11: end while
12: return  $S(t)$ 

```

A switching sequence's neighbours can be precomputed and does not vary over time, while the implemented ancestor will vary depending on the switching signal implemented at the previous timestep. However, a set of potential ancestors of a switching signal sequence can be precomputed and the correct implemented ancestor can be found online. This reduces computation time.

Algorithm 5.2 shows how MPC with MDS computes expected fitness and uncertainty value for the switching signal sequences and uses simple optimization over these values to decide the set of potential sequences. Before Algorithm 5.1 is started, $\tilde{F}_t(\boldsymbol{\varsigma})$ and $\tilde{Y}_t(\boldsymbol{\varsigma})$ must be given initial values. We recommend a positive factor times the optimal function value of the initially known solution, to ensure correct scaling.

The expected fitness $\tilde{F}_t(\boldsymbol{\varsigma})$ and uncertainty value $\tilde{Y}_t(\boldsymbol{\varsigma})$ are updated for all sequences in Ω at all times t . Ω contains all possible switching signal sequences, even those for which (5.10)-(5.14) may be infeasible at time t . When the values are updated, the impact of the implemented ancestor is weighted to the impact of the average of the neighbour sequences by $0 < \alpha < 1$. Notice that the expected fitness and uncertainty value could be easily computed in parallel. The selection of the set of potential sequences

could also be parallelized, if a centralized step ensures the uniqueness of each potential sequence in the set by replacing duplicates with random sequences. This will lead to increased exploration.

5.3 Implementation example

To illustrate the performance of the proposed method, simulation experiments are conducted where MPC with MBS has been implemented on a small switched linear system and the performance of MPC with MBS is compared to three benchmark controllers. In the following, first the benchmark algorithms are introduced, then the system is given, and finally the results are shown.

5.3.1 Benchmark controllers

The performance of MPC with MDS is compared to three other controllers, named Optimal, Random and Safe. All controllers are implemented in serial, since the authors do not have access to parallel computing yet. MPC with MDS and the Random controller are however prepared for parallel implementation as mentioned in the end of Section 5.2.1. The experiments are implemented in Matlab using Yalmip, [95], and Gurobi.

MPC with MDS: follows Algorithm 5.1 and 5.2 with the selection of the potential sequences prepared for parallelization as mentioned in the end of Section 5.2.3. (5.10)-(5.14) is solved using Gurobi's barrier method limited to 150 iterations. The experiment is repeated 10 times to illustrate the effect of the random selections used in Algorithm 5.2.

Optimal controller: solves the optimal mixed-integer MPC problem (5.4)-(5.9) at each time t . To obtain a convex continuous approximation of the mixed-integer program, the switching signals are represented by binary variables.

Safe controller: solves (5.4)-(5.9) for the system's initial state and uses the solution as switching signal and continuous input the first 9 timesteps. Hereafter the basic switching signal and the safe control law is implemented, i.e. $\sigma_t = \gamma$ and $u_t = \kappa(x_t) \forall t \geq 10$.

Random controller: selects the set of potential sequences randomly instead of using Algorithm 5.2 but follows otherwise Algorithm 5.1. Problem (5.10)-(5.14) is solved using Gurobi's barrier method with a limitation of

150 iterations. The experiment is repeated 10 times.

5.3.2 System

The system used to illustrate the performance has initial state $x_0 = [-7 \ -2]^T$ and dynamics

$$x_{t+1} = A_{\sigma}x_t + B_{\sigma}u_t \quad \text{where } \sigma \in [0, 1, 2, 3],$$

$$A_0 = A_2 = \begin{bmatrix} 0.9 & 0.1 \\ 0 & 1.1 \end{bmatrix}, B_0 = B_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1.1 \end{bmatrix}, B_1 = \begin{bmatrix} 0.1 \\ 0 \end{bmatrix},$$

$$A_3 = \begin{bmatrix} 1 & 0.3 \\ 0 & 1 \end{bmatrix} \text{ and } B_3 = \begin{bmatrix} 0.1 & 0 \\ 0 & 0.1 \end{bmatrix}.$$

It is known that the unconstrained system is stable when $\sigma_t = 0$ and $u_t = [-0.1 \ -0.2]x_t$ for all t .

The system is subject to state constraint $H_s x_t \leq h_s \ \forall t$, and input constraints $H_{\sigma} x_t \leq h_{\sigma} \ \forall t$, with

$$H_s = \begin{bmatrix} 1 & 0 \\ -1 & 0 \\ 0 & 1 \\ 0 & -1 \\ 0.1 & 0.2 \end{bmatrix}, h_s = \begin{bmatrix} 10 \\ 10 \\ 5 \\ 5 \\ 1 \end{bmatrix}, H_3 = \begin{bmatrix} 1 & 1 \\ -1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & -1 \end{bmatrix}, h_3 = \begin{bmatrix} 3 \\ -1 \\ 3 \\ -1 \\ 3 \end{bmatrix}, H_0 = \begin{bmatrix} 1 \\ -1 \end{bmatrix},$$

$$h_0 = \begin{bmatrix} 0.25 \\ 0.25 \end{bmatrix}, H_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, h_1 = \begin{bmatrix} 2 \\ 2 \end{bmatrix}, H_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, h_2 = \begin{bmatrix} 5 \\ -0.05 \end{bmatrix}.$$

Notice that the continuous input under switching signal 2 is bounded away from zero and that the continuous input under switching signal 4 has a different dimension.

Between timestep $t = 50$ and $t = 100$ the system is subject to a constant disturbance

$$x_{t+1} = A_{\sigma}x_t + B_{\sigma}\kappa(x_t) - \begin{bmatrix} 0.1 \\ 0 \end{bmatrix} \quad \forall 50 \leq t \leq 98. \quad (5.15)$$

The controllers have no information about this disturbance, besides the measured state, and thus do not guarantee state constraint satisfaction. The disturbance is however so small that the Optimal controller remains feasible.

Both MPC with MDS, the Optimal and the Random controller are implemented with prediction horizon $N = 10$ and the following cost

$$J(\boldsymbol{\varsigma}, \mathbf{u}, \mathbf{x}) = \sum_{k=1}^{N-1} (x_k^T Q x_k + u_k^T R_{\varsigma_k} u_k) + x_N^T Q_{N \times N}, \quad (5.16)$$

where $R_0 = 1$, $R_1 = R_2 = 0.001$, $R_3 = \begin{bmatrix} 0.1 & 0 \\ 0 & 1 \end{bmatrix}$, $Q = \begin{bmatrix} 0.1 & 0 \\ 0 & 0.1 \end{bmatrix}$ and $Q_N = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. The terminal set is $\mathcal{X}_f = \{x | x_N^T Q_{N \times N} \leq 0.2\}$. Satisfaction of Assumption 5.1 part 1) and 2) is immediately clear, and satisfaction of Assumption 5.1 part 3) and 4) and Assumption 5.2 can be shown using the S procedure.

5.3.3 Results

The four controllers performed as expected in the conducted experiments. In Figure 5.2 the realized cost, $x_t^T Q x_t + u_t^T R_{\varsigma_t} u_t$, is shown accumulated over time. As expected the Optimal controller is over time the cheapest controller, while the Safe controller is the most expensive. Furthermore, the Optimal controller only accumulates a little more cost when the system is disturbed between $t = 50$ and $t = 100$ while the Safe controller is costly. MPC with MDS performs well compared to both Safe and Random controllers. MPC with MDS performs significantly better than the Safe controller when the system is disturbed. The safe input was chosen by MPC with MDS mainly when the system was close to the reference value. During the disturbance ($50 \leq t \leq 98$), a better solution was found at 94,4% of the timesteps.

In the experiment, the maximum computation time for one iteration of the Optimal controller was 0.6032 s. MPC with MDS was implemented in serial, with a maximum computation time of 1.6575 s. However, the strength of MPC with MDS is its parallelizability. Computing the parts of the algorithm that must be serial took maximum 0.0189 s and the longest time it took to solve (5.10)-(5.14) was 0.3669 s. Computing Algorithm 5.2 in serial is time consuming due to the large amount of data. Updating the estimated




FIGURE 5.2: *The accumulated cost of the applied inputs and realized states under the different controllers.*

fitness and uncertainty value took 0.4389 s in serial. The computation for each switching signal sequence is independent, and therefore easily parallelizable. It is expected that a parallel implementation of MPC with MDS reaches computation times comparable to that of the Optimal controller on this small system.

5.4 Conclusions

This chapter addresses Research Question Q5 *How can Bayesian optimization help solve a model predictive controller’s mixed integer optimization problem?* A method is proposed that uses ideas from Bayesian optimization to parallelize the computations performed by MPC on switched linear systems. The presented method, MPC with Memory-based Discrete search (MDS) guarantees recursive feasibility and stability of the system. It finds ‘good enough’ inputs fast. The solution is not optimal, but the performance improvements compared to implementing a known safe input makes it an interesting method for many applications.

In MPC with MDS, the computational complexity is decreased by reducing the mixed-integer MPC problem to a prespecified number of convex optimization problems. The results indicate that this decreases computation times, if the algorithm is implemented in parallel. It is expected that, with sufficient parallelization, the computation time will increase less than that of

a mixed-integer MPC when the system size increases.

In the next chapter, Chapter 6, MDS is applied to a synchromodal transport network. Here the learning capabilities of MDS is not only used to improve computation time, but also to decrease the need for communication between two agents.

Chapter 6

Co-planning with Learning

Cooperation between container transport service providers can increase efficiency in the logistics sector significantly. This is established in Chapter 3 and 4. However, cooperation between competitors requires co-planning methods that not only give the cooperating partners an advantage towards external competition but also protect the partners from losing information, clients and autonomy to one another. In this chapter, we apply the real-time method proposed in the previous chapter, Chapter 5, to the cooperation between a barge and a truck operator. Since the synchromodal transport system is marginally stable when capacity constraints on parking and stacks are omitted, the recursive feasibility of the method is guaranteed without an explicit safe control law. The proposed co-planning method called *departure learning* lets a barge operator consider the joint cost of themselves and a truck operator when deciding barge departures. To answer Research Question Q3, departure learning is compared to, among others, integrated container and truck routing with traditional, pre-scheduled departure plans. To address Research Question Q6, the communication between the two operators is minimal. The barge operator's decisions are taken based on feedback from the truck operator on only the total cost invoked by different barge departure plans.

The core of this chapter has been submitted to a journal. The initial work has been presented at *the International Conference on Computational Logistics*¹. This chapter is organized as follows: Section 6.2 details the problem

¹Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. Learning-based co-planning for improved container, barge and truck routing. In Eduardo Lalla-Ruiz, Martijn Mes, and Stefan Voß, editors, In *Proc. of the International Conference on Computational Logistics*,

of co-planning barge departures. In Section 6.3, an MPC planning barge, truck and container moved in a network with a single decision maker is presented. Hereafter, departure learning is introduced in Section 6.4. The overall framework is presented first, followed by the barge operator's learning strategy and the truck operator's feedback. In Section 6.5, simulation experiments are used to show the impact of the parameters of departure learning and how departure learning performs compared to three benchmark methods. Section 6.6 concludes the chapter.

6.1 Introduction

Better co-planning between stakeholders in transport systems for planning barge schedules, truck and container routes in real-time will help utilizing the transport capacity better. A more efficient transport system will help alleviating the negative impacts on the environment, since less resources will be used and less pollution will be emitted on each transport. The transport sector is a large contributor of CO₂ emissions and has a low efficiency, with, e.g., trucks being empty 26% of the kilometres they drive in the Netherlands [39]. CO₂ emission is however not the only negative impact of freight transport. The report [159] estimates the external costs of transport, such as the cost of accidents, climate impact, and noise nuisance. Here it is concluded that maritime transport induces the lowest external cost, followed by rail, inland waterway and road transport in this order. It is therefore desirable not only to improve the vehicle utilization, and hence efficiency, of truck transport, but also the utilization across transport modes.

Synchromodal transport uses a-modal bookings and change acceptance to enable transport providers to optimize plans in accordance with the realisation of uncertainties [47, 160]. The concept gives transport providers more flexibility than the previous concepts, intermodal and multi-modal transport, since mode decisions in these concepts were fixed, at latest, before departure. In the traditional transport literature, decisions are divided into strategic, tactical and operational levels [152]. Strategic decisions have long lasting impact and usually high impact on revenue. Tactical decisions have impact over a tangible time horizons and are typically based on estimates of future events. Plans are often made on the tactical level and corrected at the operational level. Operational decisions regard what to do right now with

the realised events. With synchromodal transport, decisions from the tactical and the operational levels are intertwined: uncertain long term plans for operational decisions can be formulated without commitment, and tactical decisions can be changed during operation. This intertwining requires additional research to utilize the potential of synchromodality. Model Predictive Control (MPC) provides a framework for combining predictions of future events with real-time decision making. MPC has previously been used to route containers in several cases, e.g., [110], [87] and [80].

Barge schedules are typically decided on at the tactical level based on estimated demand [33]. When plans are made in advance, the realised demand is often different and external factors, like weather, cause unforeseen limitations. Some methods plan in accordance with these uncertainties [161], others adjust predefined departure times after the demand realization [10] or cancel unprofitable departures [175]. Truck routing is typically decided on at the operational level based on pick up and delivery locations and times of the goods [129]. In [82] we demonstrated the negative impact of planning first container routes and then truck routes compared to planning them simultaneously in a synchromodal network. The results of [130] show the same on a network with only one origin of the demand.

Barges and trucks are often operated by different stakeholders, so simultaneous planning requires cooperation. Cooperation can involve both information sharing and loss of autonomy. Many companies are interested in the benefits of cooperation [27], but participate reluctantly due to these implications. Many cooperation schemes in the transport literature are constructed such that missing information or sudden changes in the willingness to follow the scheme can damage the other participating parties. Traditional cooperation schemes vary from auctions [174] to distributed optimization [34, 88]. [43] provides an overview that classifies the existing research into (1) centralized methods, which requires a neutral party, (2) decentralized optimization methods, and (3) auction-based methods.

We use the term co-planning to describe the act of cooperating to achieve the vehicle and container transport plans that are best for the group of cooperating stakeholders without sharing sensitive information or being vulnerable to defiance of the other parties. Very little research on co-planning methods exist. [13] proposes a co-planning method that optimizes the joint profit cycle-time imposes on a port and a shipping liner. In their method in each round of communication, each operator optimizes its costs for a given cycle-time. These costs are communicated and the expected joint profit is

computed. If the expected profit has improved, a higher cycle-time is used in the next round of communication. Since the joint problem under certain conditions is concave, the joint profit and the corresponding cycle-time can be optimized without exchanging detailed information. For more complex problems, their method may find the lowest cycle-time corresponding to a local maxima, not the global maximum. In the literature on transport contract negotiation, planning at a tactical level is researched under privacy assumptions similar to those of co-planning. In [177] a carrier offers two contracts to a shipper. Since it is a tactical problem, the carrier only has probabilistic information about the future demand when deciding what contracts to propose. For operational co-planning, the decision frequency is higher, and there is thus a potential to learn from the other stakeholders' previous actions.

In this chapter, we show the impact of co-planning and describe the details of a co-planning method, called *departure learning*, for real-time co-planning between a barge and a truck operator. The co-planned actions are departures of a barge and the remaining actions considered are loading and unloading of containers to trucks and barge, and departures of the trucks. An initial version of departure learning was presented in [79] together with preliminary results and a variation of this method for one barge and multiple truck operators was presented in [81], again with limited experimental work. The method is based on Model Predictive Control (MPC) and uses ideas from Bayesian optimization to learn good departure times through continuous communication. Departure learning requires communication of a number of schedules and indications of the corresponding performances between the barge operator and the truck operator. The initial departure learning algorithm is now enhanced with better initial guesses for the performance of schedules and the impact of the method's learning parameters is presented. Furthermore, we show the impact of actively learning what schedules are expected to perform well by comparing departure learning to a similar method that uses randomly chosen schedules. It is assumed that no party seeks to exploit the framework, but if one party acts autonomously the other party is not damaged. The main advantage of departure learning is the ability to co-plan barge departures, i.e. enable a barge operator to depart when it improves the operation cost of the transport network, without losing control over own operation, transferring responsibilities between stakeholders or communicating detailed information.




FIGURE 6.1: *Small Dutch transport network used as example in this chapter.*

6.2 Problem statement

When each transport operator can change planned action up until the time the action is carried out, co-planning between transport operators must happen in real-time. One planning problem in a synchromodal transport network is co-planning between truck operators responsible for routing trucks and delivering containers in time, and a barge operator responsible for barge departures. Figure 6.1 shows the synchromodal transport network used as example in this chapter. Changing the departure time of a barge impacts the other operator in the network significantly. If the barge departure fits well with the release and due date of container demand, it is often attractive for the truck operator to send the containers by barge instead of trucking them. They can do so as the containers are booked a-modal under synchromodal transport. Good departure times thus often benefit the truck operator by decreasing transport cost and the barge operator by increasing the transported volume.

The barge operator does not know which barge departure times will decrease the joint cost of the transport the most. In a highly dynamic environment with many small shippers that do not request transport of containers regularly, optimizing barge departures based on past transport flows can give very ill-fitting schedules for upcoming periods. In those cases, assuming the barge operator has no knowledge about demand but what he receives in real-time from the other operators is thus reasonable. On the other hand, the other operators do not have authority to decide on the barge schedule as it is a core

business decision for a barge operator. In case the barge operate in a system with multiple other entities, no one operator will have full knowledge of the system. We consider a transport system that consists of a barge operator who decides on the departure times of one barge between two different terminals and one other operator, namely the truck operator, who decides on the routing of trucks and the mode and route containers are transported by. In [81] the case of one barge operator and several truck operators is presented without a thorough investigation of the impact of the method's parameters on the achieved performance.

The truck operator wants to keep all information about a given container's transport-order (such as release time, due date and quantity) private until the container is committed to take the barge. In that case, the truck operator is willing to share the necessary practical information. For simplicity, we assume all containers are standard 40ft containers. The truck operator is willing to share the total cost they will occur over a certain time horizon if a given barge schedule is implemented. To prevent that the barge operator can infer information, as well as to decrease the computational burden, there is, however, a limit on how many times they are willing to share this information per timestep. We, furthermore, assume in the numerical experiments that the truck operator shares the cost truly and that the barge operator indeed sends back the best schedule honestly. It is, however, worth noticing that each operator can incorporate other expenses into the costs used in the method and has full authority over actions they are responsible for.

We assume the barge and the truck operators have an agreement outlining the distribution of the economical gains and burdens. The goal of the co-planning method is thus to obtain the cheapest possible operation of the total transport system.

In short the presented method, called *departure learning*, answers the question 'How can a barge operator learn which barge departures make the total operation cost of a synchromodal transport network lower under limited real-time information about the total cost?'.

6.3 Real-time centralized decision making

Before presenting the proposed co-planning method *departure learning*, we here introduce a real-time, centralized method, the operator of a synchro-modal transport network can use to achieve receding horizon-optimal control

of the network if one entity has full knowledge and authority. The method is a model predictive control (MPC) method that uses the known dynamics of the transport network. The assumptions and notation of the centralized method is similar to those of departure learning.

MPC addresses uncertainties by adjusting future plans based on feedback from the system [102]. Every Δt timeunits there is a new decision moment where actions to be taken until the next decision moment are fixed based on a predicted plan for the next $T_p \Delta t$ timeunits. We denote with t the running time and with k the count of decision moments. We call the latter timesteps and define $t = k \Delta t$.

Using MPC for problems that require frequent updates, i.e. low Δt , and a long prediction horizon T_p necessitates fast optimization of the model. In synchromodal transport problems, a long prediction horizon is needed because of the long travel times of barges and the need to describe at least one departure from each terminal. We therefore formulate the truck and container routing problem as flows with continuous variables. This decreases the computational complexity sufficiently such that no heuristics, such as the ones described in [69], are needed. Frequent updates can ensure sufficient precision when the continuous optimal decisions are rounded to integer variables [141]. We separate the containers into commodity flows based on their destinations. If a finer granularity is needed, due date and container type can also be included in the definition of a commodity. This has earlier been considered in, e.g., [110], which differentiates between containers with different destination and due time combinations, and [127], which only distinguishes full and empty containers. A finer granularity will increase the computational complexity and thus the time period, Δt , between decision moments. It is assumed local decision makers disaggregate the commodity flow decisions into actions for separate containers. The barge capacity is much larger than that of trucks and the vessel movements are thus described by binary variables.

The key assumptions of the synchromodal transport system dynamics are:

- Any node in the network can be the origin and destination of transport demand, if it is defined as such, hence both import and export are considered.
- Demand is modelled as containers available to the network and needed from the network. Unsatisfied demand is penalized. The demand is fully known over the prediction horizon.

- Containers are modelled as continuous variable, commodity flows.
- Trucks are also modelled as continuous variable flows.
- The number of trucks is finite and each truck can transport one container.
- The barge has invariant, finite capacity.
- Quay capacity, crane capacity, etc., and (un)loading rates are considered sufficient.
- Terminal operating hours, drivers resting hours, etc., are not considered.

Each geographical location in the synchromodal transport network is represented by a node in a graph with multiple directed arcs connecting the nodes. A location can be a terminal where trucks and/or the barge can be loaded and unloaded; a waypoint where trucks can change their intended route; or a hub where trucks can pick up or deliver containers and possibly park. The set of nodes is denoted \mathcal{N} and the set of road-arcs is denoted \mathcal{R} . There is one barge in the network that sails between node 1 and node 2. The two directional arcs describing this waterway comprise the set \mathcal{W} . Nodes 1 and 2 can also be connected by road. The operators' decisions can only be changed when the vehicles and containers are at the nodes. It is, e.g., not possible to make the barge return to its departure terminal if a delay occurs. Furthermore, it is assumed that only the truck operators have contact to clients and therefore the barge operator receives the demand only through truck operators.

6.3.1 System dynamics

In the following we describe the realized dynamics of the transport system. Departure learning and the centralized method are both based on predictions of the consequences of future actions. These predictions are made using the same dynamics. To distinguish between predicted and realized actions and states, we use a bar over the notation for the realized case. The method used to describe this simultaneous truck and container planning problem is a simplification of the one presented in [80]. The following assumptions allow the use of a simplified method: 1) the travel time is deterministic, 2) all trucks are of the same kind, 3) there is no scheduled services in the transport network, and 4) there is unlimited capacity for loading and unloading. The full method can be used with departure learning, but as it adds complexity to the description, the simplified model is used here.

We define virtual demand nodes adjacent to the graph nodes where containers can have origin or destination. They act as a reminder of unfulfilled bookings. The set of virtual nodes is denoted by \mathcal{D} . The dynamics of the virtual demand nodes are given as

$$\bar{z}_i^d(k+1) = \bar{z}_i^d(k) - \bar{u}_{di}(k) - \bar{u}_{id}(k) + \bar{d}_i(k) \quad \forall i \in \mathcal{D}, \forall k, \quad (6.1)$$

where $\bar{d}_i(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the newly realized demand of each commodity at time k . Notice that all values are positive, so whether the demand indicates containers that are ready to be transported or that are due depends on the commodity, i.e. the element in the vector. n_c is the number of commodities, i.e. the number of virtual demand nodes. The mappings $p_i^r \in \{0, 1\}^{1 \times n_c}$ and $p_i^d \in \{0, 1\}^{1 \times n_c}$ are defined such that $p_i^r \bar{d}_i(k)$ is the sum of containers that are ready to be picked up at node i at time k and $p_i^d \bar{d}_i(k)$ is the sum of containers that are due at node i at time k . The variable $\bar{z}_i^d(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the unsatisfied demand at node i at time k and $\bar{u}_{id}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the containers of each commodity from terminal node i that are used to satisfy due dates at the virtual demand node at time k . $\bar{u}_{di}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is the opposite. To guide the direction of the demand satisfaction, the following must be true:

$$p_i^r \bar{u}_{id}(k) = 0 \quad \forall i \in \mathcal{D}, \forall k \quad (6.2)$$

$$p_i^d \bar{u}_{di}(k) = 0 \quad \forall i \in \mathcal{D}, \forall k \quad (6.3)$$

Each node in the network can be connected with three kinds of other nodes: \mathcal{D}_i , \mathcal{W}_i and \mathcal{R}_i . \mathcal{D}_i contains node i 's adjacent virtual demand nodes and \mathcal{W}_i the node to which i is connected by waterways. These sets are either empty or have one element. The set \mathcal{R}_i contains all nodes that are connected to node i by road. Based on these sets, the dynamics of the stacks of containers are

$$\begin{aligned} \bar{z}_i^c(k+1) = & \bar{z}_i^c(k) + \sum_{j \in \mathcal{D}_i} (\bar{u}_{di}(k) - \bar{u}_{id}(k)) + \sum_{j \in \mathcal{W}_i} (\bar{u}_{ji}^b(k - \tau_{ji}^b) - \bar{u}_{ij}^b(k)) \\ & + \sum_{j \in \mathcal{R}_i} (\bar{u}_{ji}^r(k - \tau_{ji}^r) - \bar{u}_{ij}^r(k)) \quad \forall i \in \mathcal{N}, \forall k. \end{aligned} \quad (6.4)$$

The variable $\bar{z}_i^c(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is a vector of how many containers of each commodity that are stacked at node i at time k . $u_{ij}(k) \in \mathbb{R}_{\geq 0}^{n_c}$ has the same structure and is for the containers transported from i to j by road at time k . The road travel between i and j takes τ_{ij}^r timesteps and the waterway travel takes

τ_{ij}^b . $u_{ij}^b(k) \in \mathbb{R}_{\geq 0}^{n_c}$ is a vector with the number of containers of each commodity that is transported from i to j by barge departing at time k .

Two binary variables $y_1(k)$ and $y_2(k)$ are used to describe the departures of the barge at time step k from node 1 and 2 respectively. The travel time from node 1 to 2, τ_{12}^b , and the return, τ_{21}^b , include loading, travel time, mooring and unloading. Containers that arrive at the terminal after loading has started will not be accepted on the barge and containers can only be picked up after the barge has finished unloading all containers. The dynamics of the barges is described as

$$\bar{z}_i^b(k+1) = \bar{z}_i^b(k) - \bar{y}_i(k) + \bar{y}_j(k - \tau_{ji}^b) \quad i, j \in \{1, 2\}, i \neq j, \forall k \quad (6.5)$$

where $\bar{z}_i^b(k) \in \{0, 1\}$ is the number of barges at the quay of node i at time k .

The barge has a capacity of c^b and only carries containers that were ready for loading at the departure time. Hence

$$\mathbf{1}_{n_c} \bar{u}_{ij}^b(k) \leq c^b \bar{y}_i(k) \quad \forall \langle i, j \rangle \in \mathcal{W}, \forall k, \quad (6.6)$$

where $\mathbf{1}_{n_c} \in \mathbb{R}^{1 \times n_c}$ is a vector of ones. The variable $\bar{z}_i^v(k) \in \mathbb{R}$ is the number of trucks parked at node i at time k , and has the dynamics

$$\bar{z}_i^v(k+1) = \bar{z}_i^v(k) + \sum_{j \in \mathcal{R}_i} \bar{v}_{ji}(k - \tau_{ji}^r) - \bar{v}_{ij}(k) \quad \forall i \in \mathcal{N}, \forall k, \quad (6.7)$$

where $\bar{v}_{ij}(k) \in \mathbb{R}$ is the number of trucks departing from i on the road to j at time k . To ensure containers only travel by roads if they are loaded on trucks, the sum of containers departing node i at time k on the road to node j must not exceed the number of trucks departing on the same road at the same time. Trucks are on the other hand allowed to drive empty. Both are modelled by

$$\mathbf{1}_{n_c} \bar{u}_{ij}(k) \leq \bar{v}_{ij}(k) \quad \forall j \in \mathcal{R}_i, \forall i \in \mathcal{N}, \forall k. \quad (6.8)$$

6.3.2 Centralized method

At each discrete timestep, the centralized method optimizes the predicted cost of operating the synchronodal transport system over the time horizon k to $k + T_p$. It is assumed that it is cheaper but slower to use the barge than to only use the road mode if we look isolated at sending one container

between the waterway terminals on a barge with a realistic utilization and do not consider the cost of driving a truck empty. The actual cost will depend on the vehicles' utilization and the container's origin and destination. We consider four kind of operation costs:

$w_i^b \in \mathbb{R}$ Total operation cost associated with an empty barge that departs from node i

$w_{ij}^l \in \mathbb{R}_{\geq 0}^{1 \times n_c}$ Cost of transporting one additional container with the barge from i to j . Can vary based on commodity

$w_{ij}^v \in \mathbb{R}$ Cost of driving a truck from i to j , regardless of load

$w_d \in \mathbb{R}_{\geq 0}^{1 \times n_c}$ Cost per timestep delay per container. Can vary based on commodity. Demand satisfaction is formulated as a soft constraint

The base cost of sailing the barge is defined as:

$$J^b(k) = w_1^b y_1(k) + w_2^b y_2(k) \quad \forall k. \quad (6.9)$$

It is assigned to the timestep where the barge departs, i.e. the total travel-cost is incurred at departure and not during the travel. This reflects the assumption that plans for a specific vehicle can only be changed when that vehicle is at a node. Running costs like owning the equipment and hiring people are disregarded, as they are out of scope of the real-time, operational problem.

The remaining cost of operating the synchromodal network is, as the barge cost, assigned to the timestep a truck or a container departs. The soft constraint penalty for late delivery of containers at their destinations is added per container, per timestep. The remaining cost is thus:

$$J^t(k) = \sum_{\langle i,j \rangle \in \mathcal{R}} w_{ij}^v v_{ij}(k) + \sum_{\langle i,j \rangle \in \mathcal{W}} w_{ij}^l u_{ij}^b(k) + \sum_{i \in \mathcal{D}} w_d z_i^d(k+1), \quad (6.10)$$

At each timestep, the centralized method gathers the number of containers, trucks and barge that are located at each node and the quantities due to arrive in the future as a consequence of previous decisions. The information

related to the barge location is:

$$\bar{x}^b(k) = \begin{bmatrix} \bar{z}_1^b(k) \\ \bar{y}_2(k-1) \\ \vdots \\ \bar{y}_2(k - \tau_{12}^b) \\ \bar{z}_2^b(k) \\ \bar{y}_1(k-1) \\ \vdots \\ \bar{y}_1(k - \tau_{21}^b) \end{bmatrix}, \quad (6.11)$$

and the remaining information is:

$$\bar{x}^d(k) = \begin{bmatrix} \left[\bar{z}_1^d(k) \cdots \bar{z}_{|\mathcal{D}|}^d(k) \right]^T \\ \left[\bar{u}_{j_1}^b(k - \tau_{j_1}^b), j \in \mathcal{W}_1, \dots, \bar{u}_{j_{|\mathcal{N}|}}^b(k - \tau_{j_{|\mathcal{N}|}}^b), j \in \mathcal{W}_{|\mathcal{N}|} \right]^T \\ \left[\bar{z}_1^c(k), \bar{u}_{j_1}(k - \tau_{j_1}^r) \forall j \in \mathcal{R}_1, \dots, \bar{z}_{|\mathcal{N}|}^c(k), \bar{u}_{j_{|\mathcal{N}|}}(k - \tau_{j_{|\mathcal{N}|}}^r) \forall j \in \mathcal{R}_{|\mathcal{N}|} \right]^T \\ \left[\bar{z}_1^v(k), \bar{v}_{j_1}(k - \tau_{j_1}^r) \forall j \in \mathcal{R}_1, \dots, \bar{z}_{|\mathcal{N}|}^v(k), \bar{v}_{j_{|\mathcal{N}|}}(k - \tau_{j_{|\mathcal{N}|}}^r) \forall j \in \mathcal{R}_{|\mathcal{N}|} \right]^T \end{bmatrix} \quad (6.12)$$

This information forms the initial constraints of optimization problem (6.13)-(6.25), which is used to decide what actions to implement at the current timestep k . At the next timestep, the process is repeated such that at any timestep k the actions and their consequences are optimized for k to $k + T_p$, but only the actions corresponding to k are implemented.

$$\min \sum_{\kappa=k}^{k+T_p-1} J^t(\kappa) + J^b(\kappa) \quad (6.13)$$

$$\text{s.t. } x^t(k) = \bar{x}^t(k), \quad (6.14)$$

$$x^b(k) = \bar{x}^b(k), \quad (6.15)$$

$$\text{and } \forall \kappa \in \{k, \dots, k+T_p-1\} : \quad (6.16)$$

$$z_i^d(\kappa+1) = z_i^d(\kappa) - u_{di}(\kappa) - u_{id}(\kappa) + d_i(\kappa) \quad \forall i \in \mathcal{D}, \quad (6.17)$$

$$z_i^c(\mathbf{\kappa} + 1) = z_i^c(\mathbf{\kappa}) + \sum_{j \in \mathcal{D}_i} (u_{di}(\mathbf{\kappa}) - u_{id}(\mathbf{\kappa})) + \sum_{j \in \mathcal{W}_i} \left(u_{ji}^b(\mathbf{\kappa} - \tau_{ji}^b) - u_{ij}^b(\mathbf{\kappa}) \right) + \sum_{j \in \mathcal{R}_i} \left(u_{ji}(\mathbf{\kappa} - \tau_{ji}^r) - u_{ij}(\mathbf{\kappa}) \right) \quad \forall i \in \mathcal{N}, \quad (6.18)$$

$$z_i^b(\mathbf{\kappa} + 1) = z_i^b(\mathbf{\kappa}) - y_i(\mathbf{\kappa}) + y_j(\mathbf{\kappa} - \tau_{ji}^b) \quad i, j \in \{1, 2\}, i \neq j, \quad (6.19)$$

$$z_i^v(\mathbf{\kappa} + 1) = z_i^v(\mathbf{\kappa}) + \sum_{j \in \mathcal{R}_i} v_{ji}(\mathbf{\kappa} - \tau_{ji}^r) - v_{ij}(\mathbf{\kappa}) \quad \forall i \in \mathcal{N}, \quad (6.20)$$

$$\mathbf{1}_{n_c} u_{ij}(\mathbf{\kappa}) \leq v_{ij}(\mathbf{\kappa}) \quad \forall j \in \mathcal{R}_i, \forall i \in \mathcal{N}, \quad (6.21)$$

$$p_i^r \bar{u}_{id}(\mathbf{\kappa}) = 0 \quad \forall i \in \mathcal{D}, \quad (6.22)$$

$$p_i^d \bar{u}_{di}(\mathbf{\kappa}) = 0 \quad \forall i \in \mathcal{D}, \quad (6.23)$$

$$\mathbf{1}_{n_c} u_{ij}^b(\mathbf{\kappa}) \leq c^b y_i(\mathbf{\kappa}) \quad \forall \langle i, j \rangle \in \mathcal{W}, \quad (6.24)$$

$$y_1(\mathbf{\kappa}) \in \{0, 1\}, y_2(\mathbf{\kappa}) \in \{0, 1\}. \quad (6.25)$$

6.4 Departure learning

Centralized planning is only possible when one entity has all information and authority to take all decisions. Traditional distributed optimization requires several rounds of communications and introduces artificial fees to shift the local optima. It is often not realistic to assume transport operators will commit to such a scheme. We therefore propose the novel co-planning method departure learning, where the global optimum is the goal of the planning, but only a pre-specified number of potential schedules and the truck operator's expected costs are communicated at each timestep. The method builds on the same assumptions and system dynamics as presented in Section 6.3.

At each timestep, the barge operator sends a set $I(k)$ of barge schedules to the truck operator. The truck operator hereafter computes the transport cost over the prediction horizon for each of the schedules. The costs are send back to the barge operator who adds the costs related only to the barge. The barge operator compares the total costs with the estimated costs of all other feasible schedules and communicates the best schedule to the truck operator. The actions corresponding to the current timestep in the schedule with the best performance are implemented by the barge operator and truck operator separately, and the process is repeated at the next timestep.

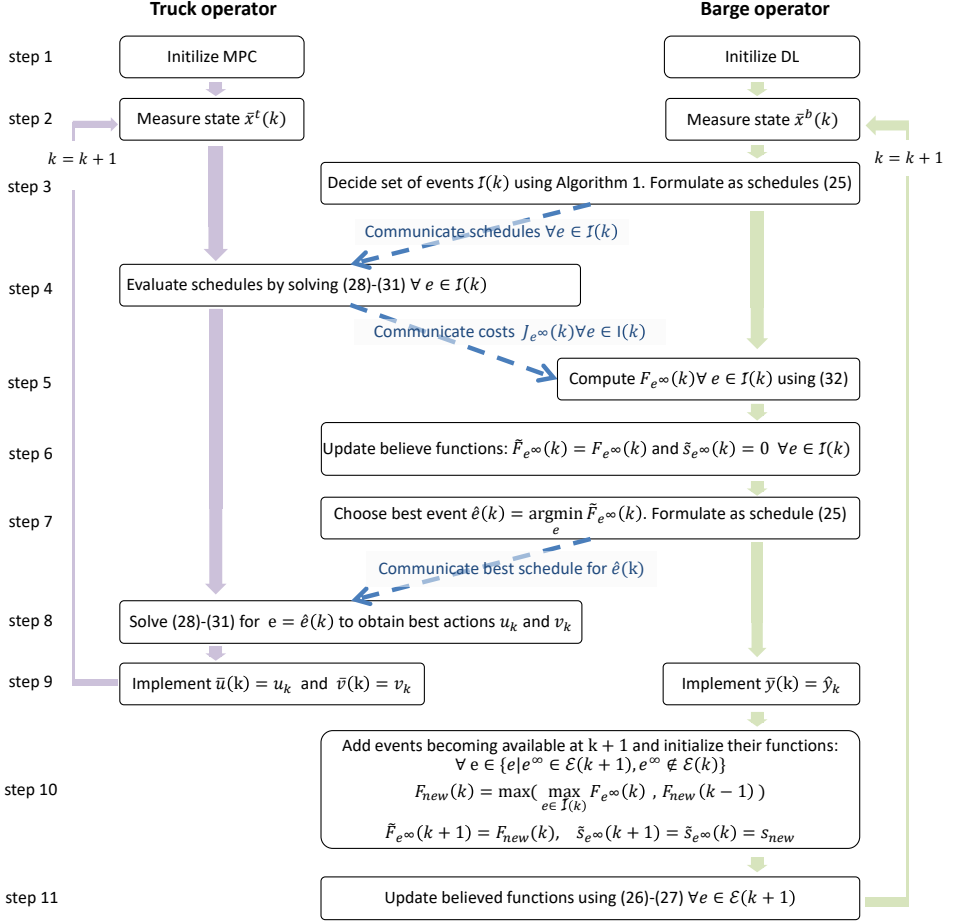


FIGURE 6.2: *Departure learning. Truck operator actions shown as purple flow to the left and barge operator actions as green flow to the right. Blue, dashed arrows indicates the necessary communication.*

To estimate which schedules will perform better, the barge operator uses the performances indicated by the truck operator at previous timesteps to estimate the performance at the current timestep. It is ensured that the set of potential schedules includes both schedules that will perform well and schedules that helps identifying good schedules in the future by using selection strategies that focus on both exploitation and exploration. The overview of the departure learning is shown in Figure 6.2. In the following, it is described how the barge operator learns good schedules, and how the truck operator evaluates the cost of a schedule.

6.4.1 Learning good departure times

To estimate what the performance of all schedules are, all schedules must be identified. However, the first departure in a schedule must be from the terminal where the barge currently is, or to which it is travelling. It is thus possible to describe the performance of all feasible schedules if only half the schedules are identified as long as the location of the barge is known. This reduces the number of binary options per timestep to one (to depart or not). Such a reduced schedule is called an event e . Figure 6.3 shows an example of a schedule and its corresponding event. Events can be decoded into schedules using the known location of the barge at time k as it determines the first departure terminal and all departures hereafter alternate terminals, mathematically:

$$y_i(k + \gamma) \leq z_i^b(k) + \sum_{\kappa=\tau_{ji}^b-\gamma}^{\tau_{ji}^b} y_j(k - \kappa) \quad \forall \gamma \leq \tau_{ji}^b, \forall \langle i, j \rangle \in \mathcal{W}. \quad (6.26)$$

$[y_1(k) \cdots y_1(k + T_p - 1)]$	$=$	$[001000]$
$[y_2(k) \cdots y_2(k + T_p - 1)]$	$=$	$[100010]$
e	$=$	$[101010]$

FIGURE 6.3: A schedule consists of two vectors of binary variables describing the departure times from the two end terminals. The corresponding event combines the two.

$e_{a1} \in \mathcal{E}(k)$	=	[0 1 0 0 1 0]
$e_{a2} \in \mathcal{E}(k+1)$	=	[1 0 0 1 0 0]
$e_{a1}^\infty = e_{a2}^\infty = e_a^\infty$	=	0 0 0 0 0 1 0 0 1 0 0 0 0 0 0 0
$e_b \in \mathcal{N}_{e_a^\infty}(k)$	=	[1 0 0 0 1 0]
$e_{c1} \in \mathcal{N}_{e_a^\infty}(k)$	=	[0 0 1 0 1 0]
$e_{c2} \in \mathcal{N}_{e_a^\infty}(k+1)$	=	[0 1 0 1 0 0]

FIGURE 6.4: *Illustration of e_∞ and $\mathcal{N}_{e_\infty}(k)$. Note that the set of neighbours varies over time.*

Each element of the event is a binary variable denoted by b_k . An event is thus $e = [b_k, \dots, b_{k+T_p-1}]$ where each element is a specific realization of $b_k \in \{0, 1\}, \dots, b_{k+T_p-1} \in \{0, 1\}$. It takes time for the barge to travel between the terminals, and therefore not all events are feasible at all timesteps. The set of events that are feasible at time k is denoted by $\mathcal{E}(k)$. Events at two different timesteps may correspond to the same sequence of events when viewed over an infinite timespan, and are as such identical. e^∞ denotes an event over the infinite timespan and is defined as $e^\infty = [\mathbf{0}_{1:k} \ e \ \mathbf{0}_{k+T_p:\infty}]$, where $\mathbf{0}_{a:b} = \{0\}^{b-a}$ is a zero-vector of suitable size. If two events are identical except for two subsequent elements, the events are said to be neighbours, i.e. for an event $e_1 = [b_k^1, \dots, b_{k+T_p-2}^1] \in \mathcal{E}(k)$, the set of neighbouring events is $\mathcal{N}_{e_1}(k) = \left\{ e = [b_k^2, \dots, b_{k+T_p-2}^2] \in \mathcal{E}(k) \mid b_i^2 = b_i^1 \forall i \setminus \{i = \{j, j+1\} \text{ for exactly one } j \in \{k, \dots, k+T_p-2\} \text{ for which } b_j^1 = 1 \text{ and } b_j^2 = b_{j+1}^1, b_{j+1}^2 = b_j^1\} \setminus \{e_1\} \right\}$. This corresponds to two barge schedules only differing in one departure time and for that departure only with one timestep. The set of neighbours are indexed with the event's e_∞ and time, since two events $e_1 \in \mathcal{E}(k)$ and $e_2 \in \mathcal{E}(k+1)$ with $e_1^\infty = e_2^\infty$ will have the same set of neighbours \mathcal{N}_{e_∞} for all k where $\mathcal{N}_{e_\infty} \in \mathcal{E}(k) \cap \mathcal{E}(k+1)$. Both e_∞ and $\mathcal{N}_{e_\infty}(k)$ are exemplified in Figure 6.4.

The barge operator holds an estimate of the total operation cost for the barge and the truck operator for each event. This estimate of an event's performance is called the event's expected fitness and is denoted by $\tilde{F}_{e_\infty}(k)$. If the barge operator has received the operation cost the truck operator will incur if the barge departs according to an event e , we say event e has been evaluated. To indicate how certain the expected fitness is, an uncertainty

Algorithm 6.1 The strategy used to assemble $I(k)$

```

1: input  $\tilde{F}_{e^\infty}(k), \tilde{s}_{e^\infty}(k), \mathcal{E}(k)$ 
2: return  $I(k)$  with  $n$  unique events
3:  $I(k) = \emptyset$ 
4: for  $i \leftarrow 1$  to  $\text{floor}(n/6)$  do
5:    $I(k) = I(k) \cup \arg \min_{e \in \mathcal{E}(k) \setminus I(k)} \tilde{F}_{e^\infty}(k) + \tilde{s}_{e^\infty}(k)$ 
6:    $I(k) = I(k) \cup \arg \min_{e \in \mathcal{E}(k) \setminus I(k)} \tilde{F}_{e^\infty}(k)$ 
7:    $I(k) = I(k) \cup \arg \max_{e \in \mathcal{E}(k) \setminus I(k)} \tilde{s}_{e^\infty}(k)$ 
8:    $I(k) = I(k) \cup \arg \min_{e \in \mathcal{E}(k) \setminus I(k)} \tilde{F}_{e^\infty}(k) - \tilde{s}_{e^\infty}(k)$ 
9:   for  $j \leftarrow 1$  to 2 do
10:     $I(k) = I(k) \cup \text{rand}(e \in \mathcal{E}(k) \setminus I(k))$ 
11:   end for
12: end for
13: for  $i \leftarrow \text{floor}(n/6)6$  to  $n$  do
14:    $I(k) = I(k) \cup \text{rand}(e \in \mathcal{E}(k) \setminus I(k))$ 
15: end for

```

function $\tilde{s}_{e^\infty}(k)$ is used. $\tilde{s}_{e^\infty}(k)$ decreases when an event corresponding to e^∞ or its neighbours are evaluated and increases slowly over k . It is expected that the performance indicator for events that share e^∞ evolve slowly over time, and that the performance indicators of neighbouring events are related. Like in Bayesian optimization, $\tilde{F}_{e^\infty}(k)$ and $\tilde{s}_{e^\infty}(k)$ are used to sample a number of candidate events that are expected to either correspond to good barge schedules or provide useful information for the future. Unlike most implementations of Bayesian optimization, the number of feasible events is finite in departure learning, and thus $\tilde{F}_{e^\infty}(k)$ and $\tilde{s}_{e^\infty}(k)$ can be computed for all events.

The set of candidate events $I(k)$ is sampled using strategies based on ranking of $\tilde{F}_{e^\infty}(k)$, $\tilde{s}_{e^\infty}(k)$ and functions of the two, together with random selection as outlined in Algorithm 6.1 for balanced exploitation and exploration. The cardinality of $I(k)$, denoted by n , is the number of schedules the truck operator must evaluate. Notice that the cost of each schedule is independent of the other schedules and the operator therefore can evaluate the schedules in parallel.

After the barge operator receives the cost of each evaluated event from the truck operator, the expected fitness of these events are updated and their uncertainty values are set to zero. Some events will be feasible at the next

time $k + 1$ which were not feasible at time k . These events are initialized with the maximum fitness evaluated at k and the uncertainty value s_{new} . Hereafter, all the fitness and uncertainty values of all events are updated as follows:

$$\tilde{F}_{e^\infty}(k+1) = \alpha \tilde{F}_{e^\infty}(k) + \frac{1-\alpha}{|\mathcal{N}_{e^\infty}(k) \cup \mathcal{N}_{e^\infty}(k+1)|} \sum_{i \in \mathcal{N}_{e^\infty}(k) \cup \mathcal{N}_{e^\infty}(k+1)} \tilde{F}_i(k) \quad (6.27)$$

$$\tilde{s}_{e^\infty}(k+1) = (\alpha + \beta) \tilde{s}_{e^\infty}(k) + \frac{1-\alpha}{|\mathcal{N}_{e^\infty}(k) \cup \mathcal{N}_{e^\infty}(k+1)|} \sum_{i \in \mathcal{N}_{e^\infty}(k) \cup \mathcal{N}_{e^\infty}(k+1)} \tilde{s}_i(k) \quad (6.28)$$

The learning parameter α balances the emphasis laid on each events' previous value and on neighbouring events' values and the factor β controls the speed at which information from previous timesteps become uncertain. To initialize departure learning prior knowledge can be used, otherwise it is recommended that $\tilde{F}_{e^\infty}(1) = \tilde{F}_{init} \forall e \in \mathcal{E}(1)$ where \tilde{F}_{init} is higher than the expected maximum fitness and $\tilde{s}_{e^\infty}(1) = s_{new} \forall e \in \mathcal{E}(1)$. s_{new} is the maximum uncertainty and is also used to update new feasible events in step 10 of the method overview in Figure 6.2.

6.4.2 Evaluating the performance

The truck operator evaluates the performance of the communicated schedules by planning container and truck routes simultaneously for each $e \in I(k)$. To do so he solves the following optimization problem, initiated from the current state for the given schedule:

$$J_{e^\infty}(k) = \min \sum_{\kappa=k}^{k+T_p-1} J^t(\kappa) \quad (6.29)$$

$$\text{s.t. } x^t(k) = \bar{x}^t(k) \quad (6.30)$$

$$\{ \langle y_1(\kappa), y_2(\kappa) \rangle \mid \kappa \in \{k, \dots, k+T_p-1\} \} = e \quad (6.31)$$

$$(6.14), (6.17), (6.18), \text{ and } (6.20) - (6.24). \quad (6.32)$$

After receiving the truck operator's cost for an event, the barge operator adds its private costs to compute the total predicted cost which serves as the

event's fitness, i.e.

$$F_{e^\infty}(k) = J_{e^\infty}(k) + \sum_{\kappa=k}^{k+T_p-1} J^b(\kappa) : \{ \langle y_1(\kappa), y_2(\kappa) \rangle \mid \kappa \in \{k, \dots, k+T_p-1\} \} = e. \quad (6.33)$$

6.5 Simulation experiments

When departure learning is used to co-plan barge departures, the learning rate of the method and the realised cost are dependent on departure learning's four tunable parameters: the prediction horizon, T_p ; the learning parameter, α ; the forgetfulness parameter, β ; and the number of communicated schedules, n . These dependencies were investigated numerically in simulated experiments. A well-tuned departure learning was hereafter compared to the performance of a method without cooperation, the centralized method presented in Section 6.3, and a co-planning method without learning. In all experiments it is assumed that decisions are taken every $\Delta t = 15$ min. In this section, the used benchmark methods and the scenarios are first described in detail. Second, the impacts of the tunable parameters are presented. Finally, the departure learning and the three benchmark methods are compared. All experiments are performed in Matlab formulated with Yalmip [95] and solved by Gurobi.

6.5.1 Benchmark methods

The performance of departure learning is benchmarked against three methods: (1) the *centralized method*, presented in Section 6.3, which requires full cooperation and unlimited information sharing, (2) a *fixed method* that does not require any cooperation and (3) an *uninformed co-planning method* without memory of previous plans.

The fixed method that requires no cooperation mimics the traditional division between decisions taken at the tactical and the operational level, while still assuming a-modal bookings. In this method, a pre-defined barge schedule is used and thus only trucks and containers can be routed in real-time. During the simulation, the truck operator solves the truck and container planning problem (6.29)-(6.32) every Δt min in an MPC fashion using this publicly available schedule.

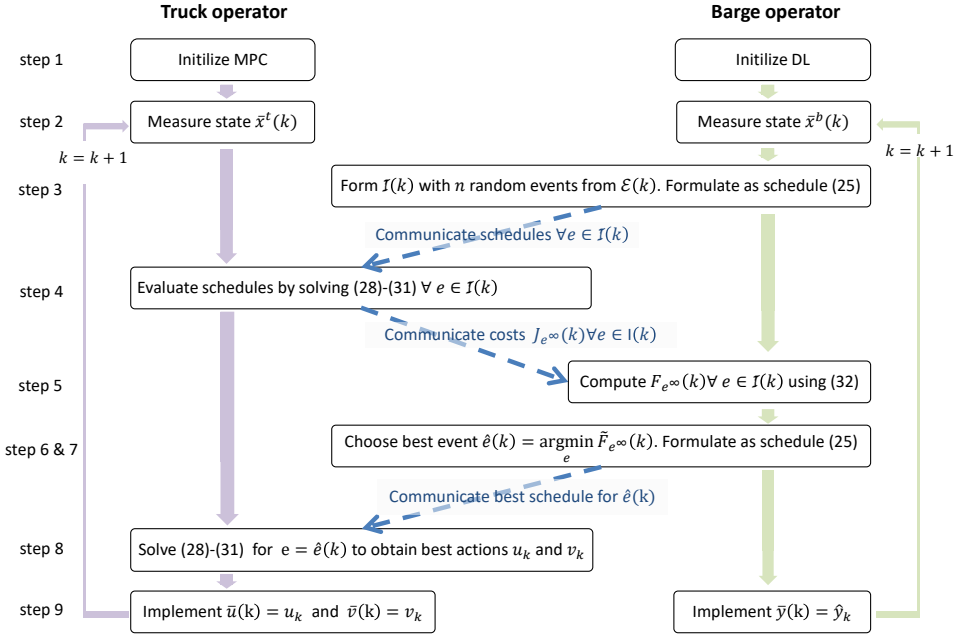


FIGURE 6.5: Actions of the uninformed method.

The uninformed co-planning method follows the steps outlined in Figure 6.5. The method deviates from departure learning in steps 3, 6, 7, 10 and 11 from Figure 6.2. Instead of using Algorithm 6.1 to actively choose which schedules to propose to the truck operator, the barge operator sends n randomly chosen schedules from the set of feasible schedules. Step 10 and 11 are omitted and step 6 and 7 are replaced by one step where the best of the schedules evaluated at the current timestep k is decided to be implemented.

6.5.2 Scenarios

The numerical experiments were performed on the Dutch network shown in Figure 6.1 where Rotterdam and Apeldoorn are origins and destinations and Nijmegen is a terminal for transshipments. The network thus accommodates $n_c = 2$ different commodities: *import* to be transported from Rotterdam to Apeldoorn, and *export* to be transported in reverse direction. It is assumed that trucks drive 90 km/h and (un)loading a truck in Rotterdam takes 20 min, while it is 10 min in Nijmegen and Apeldoorn. With these assumptions, the

TABLE 6.1: *Parameters and costs of the realistic scenario.*

Non-zero initial states	Network parameters	Costs	
$\bar{z}_2^b(0) = 1$	$\tau_{12}^b = \tau_{21}^b = 24$	$w_{12}^b = w_{21}^b = 210$	$w_d = 1000$
$\bar{z}_2^v(0) = 36$	$\tau_{13}^r = \tau_{31}^r = 9$	$w_{13}^v = w_{31}^v = 73.19$	
	$\tau_{23}^r = \tau_{32}^r = 4$	$w_{23}^v = w_{32}^v = 33.93$	
	$c^b = 100$	$w_{12}^l = w_{21}^l = 13.18 \mathbf{1}_{n_c}$	

140 km distance between Rotterdam and Apeldoorn corresponds to 123 min traveltime, and the 55 km distance between Nijmegen and Apeldoorn takes 56 min. The barge between Dordrecht and Nijmegen is by [130] reported to take 5 h including loading, so we assume the total travel time between Rotterdam and Nijmegen is 6 h. The capacity of the barge is assumed to be 100 containers. The truck operator is assumed to have 36 trucks, each of which can transport one container. They are all parked in Apeldoorn at the beginning of the simulation. The barge departure schedule used by the fixed method has its first barge departure at timestep 1 from Nijmegen and departs hereafter alternating between the two terminals every 6.5 hr corresponding to $\tau_{12}^b + 2 = 26$ timesteps.

The transport cost is in most of the literature on synchromodal transport computed primarily from the shippers perspective [10, 130, 134]. This does not capture the cost of repositioning empty trucks and under-full barges realistically. One strength of departure learning is the ability to track empty vehicles and thus the ability to assign cost to them, see [79]. We therefore use the vehicle-centered costs shown in Table 6.1.

The demand used in the experiments contains both import and export. Unless specified otherwise, we use the realistic demand profile with peaks for which the first 5 days can be seen in Figure 6.6. In this profile, a base-demand of 0 to 2 containers are released in Apeldoorn and 0 to 1 containers in Rotterdam every 15 min. On top of this, 80 to 100 containers are released at both locations at independent and irregular time intervals between 7 and 22.5 hr. The profiles were constructed such that each container was released at least 10 hr before they were due at the other location. The number of containers due at a virtual demand node is drawn at each timestep from a uniform distribution between zero and the number of containers that can have due date at this destination at this time. The demand was sampled once

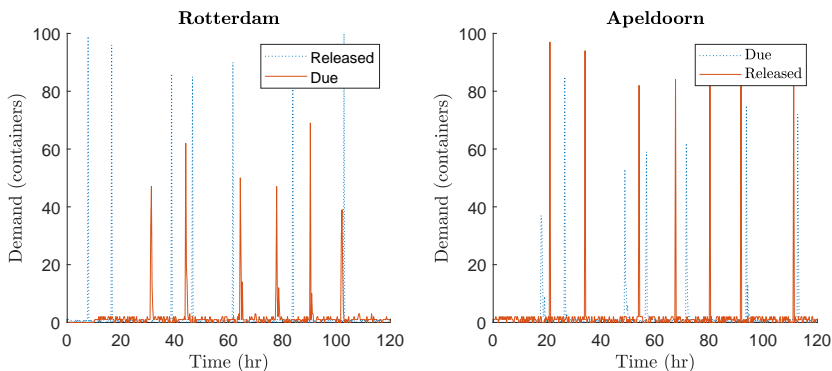


FIGURE 6.6: *The first 5 days of the demand profile with peaks used in the experiments.*

and used in all experiments. Within this period, a total of 878 export and 855 import containers were both released and due.

To investigate if the observed patterns vary with different demand profile characteristics, some of the experiments have been repeated with a demand profile without peaks. In this profile, between 0 and 1 containers are released at Rotterdam and between 0 and 3 are released at Apeldoorn every 15 min. The due demand is sampled in the same way as for the demand profile with peaks. The profile is sampled one time, which will be referred to as the *unbalanced* demand profile because of the surplus of export.

When appropriate, the experiments are a simulation of 5 days transport in the system where, initially, no containers are present and all trucks are parked in Apeldoorn. This simulation setup will be referred to as the *long simulation setup*. A *short simulation setup* has also been used. Experiments with this setup starts with the state of the system after the centralized method with $T_p = 80$ has been used for 31,75 hr (127 timesteps) on the long simulation setup. These experiments stops 101 timesteps after. This time period is chosen since it starts after the demand profile with peaks is fully established and covers a time period where the realized cost when using the optimal method is higher than average, which indicates that the problem is more complex during this period.

6.5.3 Impact of the tunable parameters

The impact of the tunable parameters has been investigated on a series of experiments, where conclusions made on earlier experiments impacted tuning

decisions on later ones. In the following sections, the impact of each tunable parameter will thus be presented after an introduction and a description of the experiments that lead to the insights. In all sections, time will be indicated as timesteps and to initialize new events $s_{new} = J_{new} = 10^7$ was used.

6.5.4 Prediction horizon - T_p

The prediction horizon impacts not only departure learning, but also the benchmark methods since they are all MPC-based. The longer the prediction horizon is, the more information each method will have available to optimize the cost. The optimization problem only sees the advantage of moving a vehicle or container, if the prediction horizon is longer than the travel time. We therefore have considered only prediction horizons longer than $T_p = \tau_{12}^b + \tau_{23}^r = 28$ timesteps. To ensure the random variables in departure learning and the uninformed method or the choice of schedule for the fixed method do not impact the results, the centralized method was used to show the impact of the prediction horizon. The long simulation setup was used in this experiment.

The results in Figure 6.7 show that as the prediction horizon increases,

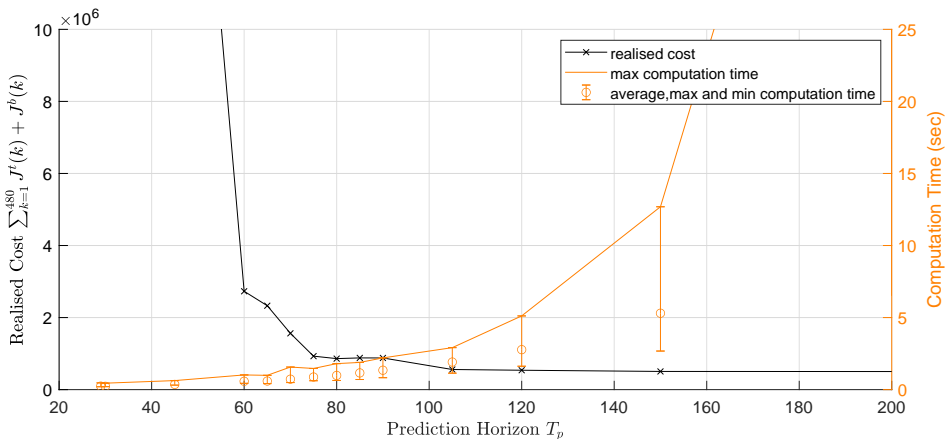


FIGURE 6.7: *The tradeoff between the total cost realised in the simulation and computation time for the optimal method. With increasing prediction horizon, the realised cost decreases and the time it takes to solve the optimization problem each timestep increases.*

the total cost of the realised actions decreases. It is noteworthy that the realised cost is significantly higher when the prediction horizon is too short to foresee the implications on the containers of a round-trip of the barge. The results, furthermore, show that longer prediction horizons increase the computation time significantly. Since the method is MPC-based, the optimization problem is solved at each timestep. In the figure, the shortest, longest and average computation times for solving the optimization problem one time are reported. The maximum computation time determines the speed of the method; ΔT must be higher than the slowest computation to ensure new decisions are available at all timesteps. This value rises quickly with increasing prediction horizon. For the remainder of the experiments in this chapter the prediction horizon is $T_p = 80$, as it is a reasonable tradeoff between the achievable realised cost and the computation time.

6.5.5 Exchanged schedules - n

The more schedules the barge operator gets feedback on from the truck operator, the more information is available to decide on departures and future communication. However, for each schedule communicated, the truck operator will have to optimize the planning problem (6.29)-(6.32). This can be

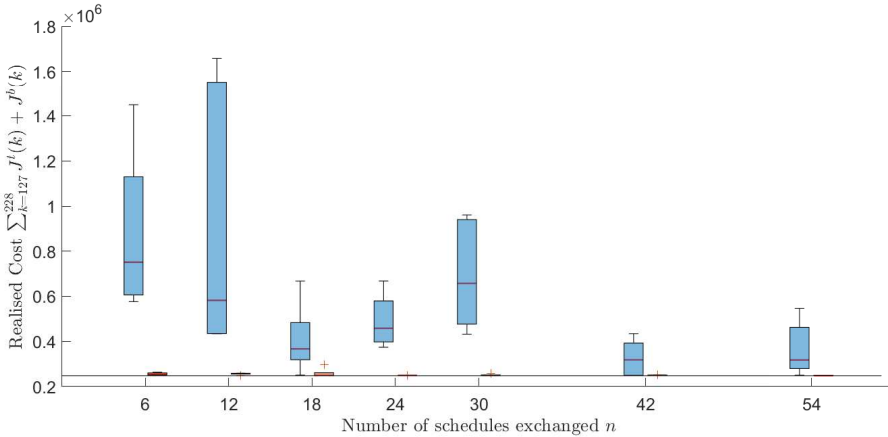


FIGURE 6.8: Blue boxplots show the min, max, 25th, 50th and 75th percentile of the realized cost when the uninformed method is exchanging different numbers of schedules. The small red boxplots show the same for departure learning with $\alpha = 0.7$ and $\beta = 0.1$. Red + indicates outliers. The black line is the realized cost when using the centralized method.

done in parallel to decrease the computation time if the truck operator has sufficient parallel computation capacity. With each schedule communicated, the barge operator gets a little more insight into the truck operator's cost structure and current demand profile since no natural noise from the shifting demand profiles is present. Therefore, it is desirable for the truck operator, both from a computation and an information perspective, to provide feedback on the lowest number of schedules that can ensure satisfactory barge departures.

The statistical information on the realised cost for five repetitions of using departure learning and the uninformed method with different numbers of exchanged schedules on the short simulation setup are in Figure 6.8 compared to the realized cost obtained by the centralized method. Especially the uninformed method benefit from exchanging more schedules, but the performance of departure learning does also improve. This is expected since the probability of randomly choosing schedules that results in a good realized performance increases when more schedules are exchanged. The results shows that the uninformed method is more sensitive to this effect than departure learning. For all considered n it is clear that departure learning deliver better and more consistent results than the uninformed method. The remaining experiments will be performed with $n = 6$ exchanged schedules.

6.5.6 Learning parameters - α and β

Departure learning's ability to learn good schedules is tightly linked to the update of the believed fitness and the uncertainty, (6.27) and (6.28). These updates are highly dependent on the learning parameters α and β . These parameters have no impact on the communication between the barge and truck operator, neither on the computation time. To investigate the impact, experiments using departure learning with different combinations of α and β -values were performed on the small simulation setup with $T_p = 80$ and $n = 6$. Each experiment was repeated five times.

The realized cost of each experiment is shown in Figure 6.9. Departure learning with all combinations of α and β perform better than the uninformed method, which, as seen in Figure 6.8, makes the total transport cost minimum €576,383. In four instances, the smallest realized cost obtained using departure learning was smaller than the cost of the centralized method, and in another four instances using $\alpha = 1$ even the 25th percentile was smaller. This happens because the departure learning's at some timesteps implement




FIGURE 6.9: *The statistics of the realized cost using departure learning with different combinations of α and β . Each group of boxplots shows the statistics for departure learning using α as indicated on the horizontal axis. Within a group, the color indicates the used β -value as indicated by the legend. The leftmost boxplot in a group is $\beta = 0$. The black line is the realized cost achieved by the centralized method.*

actions that at that timestep seem suboptimal, but over time open up for decisions which improves the realized cost.

It is clear from the results that departure learning with $\alpha = 1$ performs differently regardless of β . When $\alpha = 1$ the expected fitness of each event is only updated based on that event's earlier evaluations. Figure 6.10 shows the number of events that has an expected fitness different from the initialization of the event. In other words, it shows how many events departure learning has an expectation about, and thus implicit how wide the majority of the search is. We call this number the *active search space*. In the figure, the realized departure times for each repetition are also marked. All repetitions with all tunings of departure learning departs the barge at the simulation's first timestep $k = 127$ because of the implementation method using Matlab sort function and because this departure also for the centralized method is optimal. When departure learning implements a barge departure, the active search space collapses rapidly to a significantly smaller size. When the barge departs, all schedules with departure times within the travel time become infeasible. The large collapses just after a departure is realized thus indicate that departure learning had investigated several events with departure at the realized time or soon after. Slow decreases in the active search space occur




FIGURE 6.10: *Departure learning’s active search space over time for four different tunings of α and β . The borders of the coloured panels are min and max over the repetitions of the experiment. The realized departure times and corresponding search space is for each repetition indicated by a marker. Notice the gap in the vertical axis.*

when schedules with a departure at that timestep becomes obsolete since the barge did not depart.

Departure learning with $\alpha = 1$ has, in addition, rapid collapses in the active search space $\tau_{12}^b = 24$ timesteps after each departure, regardless of whether there is a new departure or not. For departure learning with $\alpha \neq 1$ this effect is also visible, but the decrease is less rapid. This indicates, that departure learning with $\alpha = 1$ to a higher extend focus on the same plan without investigating plans that has slightly different departure times, and thus becomes obsolete at different timesteps. This is furthermore supported by the very small active search space. It is thus likely that even though $\alpha = 1$ is a very good tuning for the investigated scenario, it may perform poorly in other cases or if different plans are found initially.

To ensure the conclusions regarding departure learning with $\alpha = 1$ holds for other scenarios, the experiments were repeated using the unbalanced demand profile. Figure 6.11 (a) shows that departure learning with this profile still performs better than the uninformed profile, but not as pronounced as with the demand profile with peaks. It furthermore shows that departure learning with $\alpha = 1$ has a very large variance. This can again be explained by the very narrow search space in combination with the relative performance of the uninformed method.

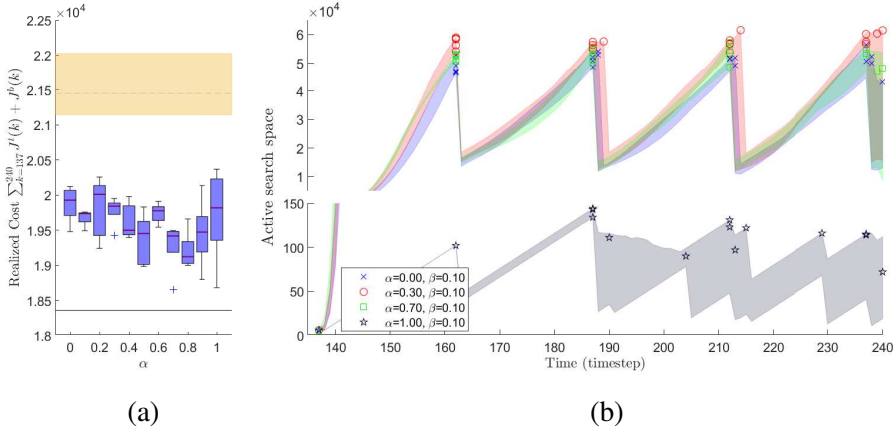


FIGURE 6.11: *The realized cost and search space for departure learning for experiments on the unbalanced demand profile. (a) corresponds to Figure 6.9 and (b) to Figure 6.10. In all experiments $\beta = 0.1$. The yellow area in (a) is the min, max and average realized cost obtained by the uninformed method.*

When there are peaks in the demand, the difference between choosing random and learned schedules is higher. This indicates that the time of departure is more important. For the unbalanced demand profile, the centralized method departs as frequently as possible. The first schedule departure learning learns thus remain a good schedule in the beginning of the simulation. However, this schedule only has three departures because of the prediction horizon length. When the simulation approaches the third departure, it is thus likely that one of the randomly chosen schedules in $I(k)$ will perform better, leading to a large diversity in the realized departure times. The centralized method only departs three times when used on the demand profile with peaks. These departures correspond to the departure found by departure learning with $\alpha = 1$. For this profile, it is not as important if the schedule that are planned around the third departure include other departures in the future or not, and it is thus likely that departure learning with $\alpha = 1$ can find this last “optimal” departure. When the unbalanced demand is used, the difference between the departure times has less impact and it is thus likely that departure learning with $\alpha = 1$ finds different good departures at each repetition of the experiment. In conclusion, it is not recommended to use departure learning with $\alpha = 1$ since the first chosen schedules will have a very large impact on the future realized cost.

Returning to Figure 6.9, a clear pattern in the impact of α and β is not visible. There is a tendency that higher α -values in combination with lower β -values gives better results. For $\alpha \leq 0.7$ and lower, departure learning with $\beta = 0$ performs very poorly. This is likely because when $\beta = 0$, equal confidence is put on evaluations performed in the recent and distant past. When new information becomes known by departure learning, earlier evaluations may become obsolete. Very high α likely compensates for this effect with the increased focus on previously evaluated schedules. For the final comparison between departure learning and the benchmark methods, $\alpha = 0.7$ and $\beta = 0.1$ are chosen.

6.5.7 Comparison between departure learning and the benchmark methods

The right tuning of departure learning can improve the methods performance, as seen in the previous sections. In this section departure learning with $\alpha = 0.7$ and $\beta = 0.1$ is compared to the three benchmark methods. The comparison is done on the results from the long simulation setup with $T_p = 80$ for all methods and $n = 6$ for departure learning and the uninformed method. The experiments are conducted on both the demand profile with peaks and the unbalanced profile. The experiments with departure learning and the uninformed method are repeated five times and average values over the repetitions are reported.

As seen in Table 6.2, the centralized method performs, as expected, best in terms of realized cost. However, among the methods that do not require full information sharing and loss of autonomy, departure learning performs very well. For the demand profile with peaks, departure learning outperforms both the fixed and the uninformed methods, while it performs nearly as well as the fixed method on the unbalanced demand profile. The performance of the fixed method is relative to the centralized performance better for the unbalanced demand profile because the cost savings obtained by having many barge departures in this profile is more significant than what can be gained by consolidating freight at specific times. In Figure 6.12 the departures realized by each method for the demand with peaks are seen. In the later half of the simulation, the fixed method's schedule is very misaligned with the "optimal" schedule provided by the centralized method, while departure learning remains similar. The uninformed method departs at relatively similar times in the different repetitions. A corresponding plot for the

unbalanced demand reveals no patterns. When the barge departure time has high correlation with the realized costs, departure learning is more likely to outperform both the fixed and the uninformed method.

Departure learning and the three benchmark methods all try to reduce the total realized cost. When doing so they achieve different utilizations of the barge and trucks. The average capacity utilization over the different barge departures is highest when the centralized method is used, followed by departure learning and thereafter the fixed method. Again the effect is more pronounced when the exact departure time is more important, i.e. when the demand has peaks. The random method utilize the barge capacity better than departure learning for the uninformed demand, but less well for the demand with peaks. For both demand profiles the uninformed method has fewer barge departures, which together with the correlation between departure times and costs may explain these results.

The truck utilization does not vary significantly between the different

TABLE 6.2: *Performance of departure learning (with $\alpha = 0.7$, $\beta = 0.1$) and the three benchmark methods on the long simulation setup.*

	Departure learning	Centralized	Fixed	Uninformed
Demand with peaks				
Realized costs (1.000)	2 068	860	7 936	8 587
Barge departures	16.6	17	19	12.8
Barge utilization	35.1%	42.8%	21.9%	28.1%
Truck departures	2 285	2 281	2 195	2 174
Truck utilization	83.6%	84.0%	85.8%	87.0%
Unsatisfied demand	1 922	720	7 787	8 439
Unbalanced demand				
Realized costs	85 132	81 825	84 066	91 919
Barge departures	14.8	17	19	9.4
Barge utilization	23.2%	25.1%	22.0%	24.6%
Truck departures	1 427	1 431	1 445	1 460
Truck utilization	66.9%	66.8%	66.7%	67.1%
Unsatisfied demand	0	0	0	0




FIGURE 6.12: *The departures realized by each of the four methods for the demand profile with peaks. Circular markers represent departures from Apeldoorn and squares from Rotterdam. For departure learning and the uninformed method, the results of the five repetitions are reported above each other. The departures, the central method considered but did not implement, are shown as smaller, transparent markers. The intensity of the color of these markers thus indicates how long the central method considered a departure beneficial.*

methods. Truck utilization is in this context the number of truck departures where the truck carries a container. For the demand profile with peaks the import and export was more or less balanced and all methods achieve a truck utilization above 80%. This high utilization is likely linked to the vehicle centred cost used by all methods. For the unbalanced demand, the truck utilization is much lower, around 67%. This is expected since more containers must be transported to Rotterdam than to Apeldoorn. Repositioning the needed empty trucks is, however, cheaper than to pay the fee for unsatisfied demand, and all methods manage to deliver all containers in time for the unbalanced demand. For the demand with peaks, more containers are transported in the system, and bottlenecks causes even the centralized method to have delayed deliveries. The reported unsatisfied demand is measured in containers times timesteps.

In conclusion, for the demand profile with peaks, departure learning is a very good method for systems where the barge and truck planning cannot be integrated due to stakeholder interests. For the unbalanced demand, high frequency of barge departures is more important than integration of plans which

leads departure learning to perform slightly worse than the fixed method. Systems with different characteristic will benefit from different approaches, but departure learning is a promising method if the system has higher correlation between barge departures and costs.

6.6 Conclusions

The efficiency in the logistics sector can improve significantly if the involved stakeholders cooperate in real-time. However, cooperation between competitors requires co-planning methods that not only give the cooperating partners an advantage towards external competition but also protect the partners from losing information, clients and autonomy to one another. In this chapter we address Research Question Q6 *How can we bridge the information gap that comes from low communication frequency?* We present a novel method, called *departure learning*, which facilitates real-time co-planning of barge schedules between a barge and a truck operator. Departure learning uses ideas from Bayesian optimization to determine what potential schedules a barge operator should propose to a truck operator based on the previous feedback from the truck operator. The more information there is available to each operator when planning, the cheaper the realization will be. The computation time does, however, increase significantly. Less feedback from the truck operator on the barge operator's departure plans decreases the possibilities for inferring information. It was found that even with feedback on only 6 schedules at each timestep, departure learning's performance was sufficient to achieve good performance.

We show that departure learning is a promising method for cooperation in practice and how the tunable parameters of departure learning affect the performance of the combined transport system. When the transport system is under pressure and consolidation of demand is only possible at specific barge departure times, departure learning outperforms the current practice where schedules are fixed ahead of time. Departure learning decides which schedules to receive feedback on using ideas from Bayesian optimization. The expected performance of all schedules is updated using two learning parameters. The results show that extreme parameters limit departure learning and that higher α values in combination with lower β values perform well. It was shown that, regardless of the parameters used, departure learning was

superior to receiving feedback on random schedules without remembering previous information.

Departure learning shows that it is possible to co-plan with very limited information sharing and no loss of autonomy. In contrast to the current, hierarchical, transport planning system, departure learning can adapt departure times to real-time information. With further research it can become a practical tool for transport operators to increase the utilization of their transport capacity and thus help alleviating the negative impacts on the environment.

In this and the previous chapters, we have explored different directions needed to answer the main Research Question : *How can container transport realistically be planned in real-time when several different stakeholders own the vehicles*. In the following chapter, Chapter 7, that question will finally be answered after examining the answers to the sub-questions.

Chapter 7

Conclusions and Future Research

Using the available resources optimally should always be the goal. This dissertation's introductory sentence is the key motivation behind the conducted research. In the transport sector, the synchromodal transport paradigm offers more flexibility than traditional paradigms, but how to convert flexibility into increased efficiency in a complex system with many stakeholders is not trivial. In this dissertation we have considered real-time, container transport problems in synchromodal networks and proposed co-planning methods to decrease the operational cost. In this chapter, firstly the conclusions of the dissertation and the Research Questions are presented in Section 7.1. Thereafter, the contributions to the academic fields that are presented in the dissertation are summarised in Section 7.2 and recommendations for additional research that can expand and strengthen the insights from this dissertation are discussed in Section 7.3.

7.1 Conclusions

The research presented in this dissertation is driven by one question:

How can container transport realistically be planned in real-time when several different stakeholders own the vehicles?

This is a very broad question that we have provided two specific answers for in Chapters 4 and 6. Several other, specific, answers exist. However, general conclusions that guide future research on this question and assist when practical applications are developed have been obtained during the research for this dissertation. Below are the most important ones:

- **Coordinate** container and vehicle routes. It results in economically better routes for both containers and vehicles and makes it easier to utilize the finite vehicle fleet better.
- **Replan** decisions periodically. Frequent reconsiderations ensure all available information is used to take informed decisions at all times. In any transport network, there are external factors that can change or delay processes. Periodic replanning are thus likely to improve the performance of the transport system. The main obstacle to frequent reconsiderations is the computational complexity of the optimization problems used to find the optimal plans. System boundaries and model type have large impact on computation time and the trade-off between these choices, optimality gap and reconsideration frequency must fit to the planning problem at hand.
- **Co-plan**. The efficiency of a transport system improves when the stakeholders actively coordinate. Both in research and in practical implementations it is important to consider what information it is viable to communicate and with what frequency. It is also necessary to clearly line out what responsibilities each organization has and how much autonomy they are willing to transfer to the co-planning method. If the co-planning organizations do not have the same goal in mind, profit distribution and participation intensives increases the complexity of cooperating.
- **Use learning** methods to decrease the communication burden. Learning can identify patterns in the other stakeholders' behaviours and thereby decrease the need for communicating information between co-planning parties.

The main conclusions were reached while answering the specific research questions introduced in Chapter 1. In the following, we summarize the insights gained into the research questions. The detailed insights can be found in Chapters 3 to 6.

Q1 What is the impact of integrating decisions across the planning-hierarchy layers that concerns container and vehicle routing?

Integrating decisions from the container routing and vehicle routing layers of the traditional planning-hierarchy leads to better operation of the vehicle

fleet. Chapter 3 shows that it smooths out the need for vehicles when containers' transport routes are planned together with the route of the vehicles that carry them, especially in scenarios with large differences between demand in peak and off-peak periods. This means both that a smaller vehicle fleet can satisfy the demand and that each vehicle is better utilized. The chapter furthermore demonstrates that the 'optimal' routing is very different depending on whether the containers' routes are planned with or without the vehicle routing in mind. It emphasises thus the importance of awareness of the planning-hierarchy when recommendations and conclusions are made based on academic research, especially when the subject is operational cost and environmental impacts.

Q2 How can operational planning under synchromodal transport take advantages of the opportunity for real-time mode-changes?

When the mode of the transport for each container is not fixed in advance, the container can follow the route through the transport network which is most advantageous at the time. A transport network often consists of locations like transshipment terminals and road/waterway/rail junctions that can serve as logical nodes where decisions regarding routing have to be taken. The mode and planned route of a container can this way be changed during transport. In Chapter 3, all departures from network nodes are reconsidered periodically using the model predictive control framework. The results show that when a truck is delayed, the proposed method changes the planned routes to perform better under the new situation.

Q3 What is the impact of stakeholders planning cooperatively at the operational level?

The advantages of co-planning lies mainly in the reduced operational cost and more efficient utilization of the transport vehicle fleet. In Chapters 4 and 6, transport under co-planning methods are compared to traditional transport. In both cases, one of the co-planning parties suggests plans and the other provides feedback. Chapter 4 highlights that the provided feedback improves the understanding and predictions of the operational costs for the party that suggests plans. The impact of co-planning is highest in transport networks where there is a large need for repositioning of empty vehicles. This is seen as better performance in instances with few shipping requests in Chapter 4 and instances with large differences between peak and off-peak demand in Chapter 6.

- Q4 How can containers and vehicles be routed cooperatively through a synchromodal network, if only traditional transport requests and their expected fulfilment are communicated?

Co-planning using automated communication of traditional transport requests and feasibility feedback is presented in Chapter 6. The method considers co-planning between a logistics service provider (LSP) and a operator of a vehicle fleet that is not bounded by schedules. The communication used in the proposed method is similar to the information exchanged in current practice with the communication frequency increased to match a synchromodal transport network's need for real-time decision making. The feedback on the feasibility and expected arrival times of the LSP's plans are used to split the shipping requests into suitable bundles and consider expected arrival times as part of the LSP's planning problem.

- Q5 How can Bayesian optimization help solve a model predictive controller's mixed integer optimization problem?

The core idea of Bayesian optimization is to improve an estimate of what is optimal by testing carefully chosen options. The options are chosen by combining an understanding of expected outcomes that is built on previous tests with a quantification of how certain this outcome is. It is a method that can find good solutions with relatively few tests. When a model predictive controller optimizes decisions that impact switched linear systems, the discrete/integer variables which describe the system switching signals slow down the computation time significantly. A method that uses the core idea of Bayesian optimization to enable parallel computation of the optimal decisions in a switched linear system is presented in Chapter 5. At each iteration, a number of tests can be performed in parallel and as the model predictive controller satisfies certain safety guarantees, the expected outcome can be projected forward in time. The results of the test can thus update the understanding of the expected outcome of different switching sequences.

- Q6 How can we bridge the information gap that comes from low communication frequency?

When transport system stakeholder co-plan, it will often be desirable to decrease the number of communication iterations per decision moment. This is both to decrease the amount of information the parties can infer to each other, but also for practical reasons, e.g., to decrease the computational burden.

When less information is communicated, the co-planning parties have less information available to base their decisions on. In Chapter 6 co-planning is improved by adjusting the method presented in Chapter 5 to the case of co-planning between a barge operator and a truck operator. The barge operator learns the truck operator's expected cost of different barge departure plans with the truck operator only providing feedback on a few of those at regular intervals. The barge operator can thus implement barge departures that fits the truck operator better (hence lowering the total operational cost) with a lower communication frequency.

7.2 Contributions

In addressing the research questions, the research leading to this dissertation has contributed to the academic understanding of especially real-time planning of synchromodal transport, but also learning-aided model predictive control. In the dissertation, model predictive control is applied to the problem of routing containers and vehicles, two novel and distinctively different co-planning methods are proposed and a theoretical proof that the ideas the learning-aided co-planning method can result in a stabilizing and recursively feasible MPC is presented. Thematically, the main contributions are summarised below:

- Integration of decisions from different layers in the traditional planning-hierarchy. As evident from Chapter 2, little research on the integration of container and vehicle routing was present when the research behind this dissertation began. The methods presented in Chapters 3, 4, and 6 all integrate decisions that traditionally are taken sequentially. Additionally, related methods are presented in [78], [79], and [81]. A static, binary planning method that integrate container and truck routing is proposed in [82].
- Real-time planning. Synchromodal transport enables more flexible planning and control of transport systems. As it is still a relatively small field, the real-time planning methods presented in Chapters 3, 4, and 6 as well as in [78], [79], and [81] contribute with inspiration into how real-time planning can be achieved (see also Research Question Q2) and what advantages it can bring.

- Using learning to parallelize MPC for switched linear systems. The performance of a controller improves if it can adjust the planned actions more frequently. In an MPC for switched linear systems, the computation time is usually slowed down by the integers used to describe the switching. In Chapter 5 a novel approach using ideas from Bayesian optimization to parallelize the MPC's computations is discussed theoretically for general switched linear systems, while it in Chapter 6, [79], and [81] is applied to different variations of the co-planning problem between barge and truck operators.
- The term co-planning. The literature on cooperation in the transport sector concerns, as described in Chapter 2, any kind of research where more than one stakeholder is considered. Everything from studies of the aggregated behaviour of completely independent, non-cooperating agents to planning methods that assume all stakeholders let a central organization take all decisions is covered. To differentiate realistic cooperation methods for operational planning we use the term co-planning. Co-planning is the process of two or more autonomous entities that create their individual plans with limited communication between them sharing carefully selected information while striving towards a common goal. Both Chapters 4 and 6, as well as [79], and [81] regard co-planning.

7.3 Future research

New research insights always bring additional questions and inspirations towards other relevant research directions. We have identified the following interesting directions during the research leading to this dissertation:

- Increase realism. When formulating optimization problems and performing numerical experiments, the level of realism and the system boundaries are always to be discussed. In this dissertation, opening hours, working hours and similar limitations were disregarded. As they increase the complexity of, especially, the vehicle routing problem we expect that the advantages of co-planning will increase if considered. Another interesting direction for gaining insights into the realism of the proposed methods is to study the impact of aggregating

the container and vehicle flows using different strategies. In Chapter 4, only containers belonging to the same transport request are aggregated, leading to slower computation times than what is achieved in Chapters 3 and 6, where only the destination is used to aggregate container flows. Exploring the trade-off between computation time and method performance is an interesting future research direction.

- Sensitivity to quality of predictions. In transport systems, there is often information available about future events (e.g., new transport requests or decreased barge capacity) which impact the plans that are made at any given moment. Using a model predictive controller, the newest available information is used every time plans are adjusted, but already executed decisions will always limit the decision space. It is thus interesting to study, both in case-studies and theoretically, the impact of inaccurate predictions of future events and quantify at what level of uncertainty, e.g., a robust MPC is favourable over the nominal MPCs used in this dissertation.
- Study the impact of transport paradigms on supply chain resilience. In this dissertation, the transport operations have been studied separately from the value chains they are part of. In many supply chains, the goods transported can be considered 'rolling stock' and for longer transports, tightening or relaxing the transport requests due dates can help balance stock levels to uncertain demand levels. Synchromodal transport's ability to change mode of transport can thus help create more resilient supply chains.
- Impact of co-planning between industrial stakeholders in a broader field. Real-time planning in systems where multiple stakeholders are involved in the value creation is not only needed in the transport sector. Valuable insights is thus expected from research into realistic co-planning methods for other sectors and for systems with a large number of homogeneous and heterogeneous stakeholders. An interesting angle is the identification of situations where co-planning improves the processes significantly over non-cooperative methods and situations where tighter cooperation is needed to advance efficiency and reduce operation cost. The environmental impacts of co-planning in different systems comprise another interesting research direction.

- Profit distribution and cooperation incentives. The co-planning methods presented in this dissertation assume the co-planning parties truthfully follow the method. However, organizations need tangible reasons to co-plan and reassurance that they are not exploited. The detection of misuse, the distribution of profit and the transparency of co-planning are thus necessary future research directions, both as separate topics and as interacting components of implementable co-planning methods.
- Active use of learning in cooperation and co-planning methods. The method presented in Chapter 6 uses learning to reduce communication and decrease the computational burden for a co-planning method. When the co-planning parties exchange information and take decisions frequently, learning can be used to identify patterns. The information inferred from these patterns can be used actively for by the participants to improve the joint performance. Identifying suitable learning techniques for various planning methods is an interesting direction for future research, both from a theoretical as well as an application point of view.

Bibliography

- [1] A.A. Agbo and Y. Zhang. Sustainable freight transport optimisation through synchromodal networks. *Cogent Engineering*, 4(1), 2017.
- [2] U. Ahmed and M. J. Roorda. Modeling freight vehicle type choice using machine learning and discrete choice methods. *Transportation Research Record*, 2676(2):541–552, 2022.
- [3] A. Alessandri, C. Cervellera, M. Cuneo, M. Gaggero, and G. Soncin. Management of logistics operations in intermodal terminals by using dynamic modelling and nonlinear programming. *Maritime Economics & Logistics*, 11(1):58–76, 2009.
- [4] T. Ambra and C. MacHaris. Agent-based digital twins (ABM-DT) in synchromodal transport and logistics: The fusion of virtual and physical spaces. In *Proc.of Winter Simulation Conference*, pages 159–169, Orlando, USA, 2020.
- [5] T. Ambra, A. Caris, and C. Macharis. Should I stay or should I Go? Assessing intermodal and synchromodal resilience from a decentralized perspective. *Sustainability*, 11(6), 2019.
- [6] A. M. Arslan, N. Agatz, L. Kroon, and R. Zuidwijk. Crowdsourced delivery—a dynamic pickup and delivery problem with ad hoc drivers. *Transportation Science*, 53(1):222–235, 2019.
- [7] D. Axehill, T. Besselmann, D. M. Raimondo, and M. Morari. A parametric branch and bound approach to suboptimal explicit hybrid MPC. *Automatica*, 50(1):240–246, 2014.
- [8] B. Ayar and H. Yaman. An intermodal multicommodity routing problem with scheduled services. *Computational optimization and applications*, 53(1):131–153, 2012.
- [9] N. Batarlienė and R. Šakalys. Mathematical model for cargo allocation problem in synchromodal transportation. *Symmetry*, 13(4), 2021.

- [10] B. Behdani, Y. Fan, B. Wiegman, and R. Zuidwijk. Multimodal schedule design for synchromodal freight transport systems. *European Journal of Transport and Infrastructure Research*, 16(3):424–444, 2016.
- [11] A. Boccia, L. Grüne, and K. Worthmann. Stability and feasibility of state constrained mpc without stabilizing terminal constraints. *Systems & Control Letters*, 72:14–21, 2014.
- [12] W. Bol Raap, M.-E. Jacob, M. van Sinderen, and S. Piest. An architecture and common data model for open data-based cargo-tracking in synchromodal logistics. In *Lecture Notes in Computer Science*, volume 10033, pages 327–343, 2016.
- [13] E. Boros, L. Lei, Y. Zhao, and H. Zhong. Scheduling vessels and container-yard operations with conflicting objectives. *Annals of Operations Research*, 161:149–170, 2008.
- [14] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [15] M. W. Braun, D. E. Rivera, M. E. Flores, W. M. Carlyle, and K. G. Kempf. A model predictive control framework for robust management of multi-product, multi-echelon demand networks. *Annual Reviews in Control*, 27 II(2):229–245, 2003.
- [16] L.A.M. Bruijns, F. Phillipson, and A. Sangers. Fair user equilibrium in a transportation space-time network. In *Lecture Notes in Computer Science*, volume 12433, pages 682–697, 2020.
- [17] P. Buijs, J. A. L. Alvarez, M. Veenstra, and K. J. Roodbergen. Improved collaborative transport planning at dutch logistics service provider fritom. *Interfaces*, 46(2):119–132, 2016.
- [18] R. T. Cahyono, E. J. Flonk, and B. Jayawardhana. Discrete-event systems modeling and the model predictive allocation algorithm for integrated berth and quay crane allocation. *IEEE Transactions on Intelligent Transportation Systems*, 21(3):1321–1331, 2020.
- [19] A. Camponogara and H. F. Scherer. Distributed optimization for model predictive control of linear dynamic networks with control-input and output constraints. *IEEE Transactions on Automation Science and Engineering*, 8(1):233–242, 2011.
- [20] A. Carvalho, P. Melo, M.A. Oliveira, and R. Barros. The 4-corner model as a synchromodal and digital twin enabler in the transportation sector. In *Proc. of IEEE International Conference on Engineering, Technology and Innovation*, Cardiff, UK, 2020.

- [21] A. A. Ceder. *New applications of ITS to real-time transit operations*, pages 19–79. *Modelling Intelligent Multi-Modal Transit Systems*. 2017.
- [22] T.-S. Chang. Best routes selection in international intermodal networks. *Computers & operations research*, 35(9):2877–2891, 2008.
- [23] L. Chen, A. Langevin, and Z. Lu. Integrated scheduling of crane handling and truck transportation in a maritime container terminal. *European Journal of Operational Research*, 225(1):142–152, 2013.
- [24] L. Chisci, J. A. Rossiter, and G. Zappa. Systems with persistent disturbances: predictive control with restricted constraints. *Automatica*, 37(7):1019–1028, 2001.
- [25] European Commission. European green deal. URL https://ec.europa.eu/clima/eu-action/european-green-deal_en.
- [26] T.G. Crainic, R. Giusti, D. Manerba, and R. Tadei. The synchronized location-transshipment problem. In *Transportation Research Procedia*, volume 52, pages 43–50, 2021.
- [27] F. Cruijssen, W. Dullaert, and H. Fleuren. Horizontal cooperation in transport and logistics: a literature review. *Transportation Journal*, 46(3):22–39, 2007.
- [28] B. Dai, H. Chen, and G. Yang. Price-setting based combinatorial auction approach for carrier collaboration with pickup and delivery requests. *Operational Research*, 14(3):361–386, 2014.
- [29] M. S. Darup and M. Cannon. Some observations on the activity of terminal constraints in linear mpc. In *Proc. of European Control Conference*, pages 770–775, 2016.
- [30] M.A.M. De Juncker, F. Phillipson, L.A.M. Bruijns, and A. Sangers. Optimising routing in an agent-centric synchromodal network with shared information. In *Lecture Notes in Computer Science*, volume 11184, pages 316–330, 2018.
- [31] W.J. de Koning, F. Phillipson, and I. Chiscop. Simulation approach for container assignment under uncertainty. In *Lecture Notes in Computer Science*, volume 12433, pages 616–630, 2020.
- [32] T. Delbart, Y. Molenbruch, K. Braekers, and A. Caris. Uncertainty in intermodal and synchromodal transport: Review and future research directions. *Sustainability*, 13(7), 2021.

- [33] E. Demir, W. Burgholzer, M. Hrušovský, E. Arıkan, W. Jammerneegg, and T. Van Woensel. A green intermodal service network design problem with travel time uncertainty. *Transportation Research Part B: Methodological*, 93:789–807, 2016.
- [34] A. Di Febbraro, N. Sacco, and M. Saeednia. An agent-based framework for cooperative planning of intermodal freight transport chains. *Transportation Research Part C: Emerging Technologies*, 64:72–85, 2016.
- [35] A. Dobrkovic, M.-E. Iacob, and J. van Hillegersberg. Maritime pattern extraction and route reconstruction from incomplete AIS data. *International journal of Data science and Analytics*, 5(2):111–136, 2018.
- [36] C. Dong, R. Boute, A. McKinnon, and M. Verelst. Investigating synchronomodality from a supply chain perspective. *Transportation Research Part D: Transport and Environment*, 61:42–57, 2018.
- [37] Jing-Xin Dong, Chung-Yee Lee, and Dong-Ping Song. Joint service capacity planning and dynamic container routing in shipping network with uncertain demands. *Transportation Research Part B: Methodological*, 78:404–421, 2015.
- [38] M. Drexl. Applications of the vehicle routing problem with trailers and transshipments. *European Journal of Operational Research*, 227(2):275–283, 2013.
- [39] Eurostat. Database: Annual road freight transport vehicle movements, loaded and empty, by reporting country, 2020. URL <https://data.europa.eu/euodp/en/data/dataset/oXV2zjVuHDQNM8vsw>.
- [40] N.Z. Farahani, J.S. Noble, C.M. Klein, and M. Enayati. A decision support tool for energy efficient synchronomodal supply chains. *Journal of Cleaner Production*, 186:682–702, 2018.
- [41] B. Feng and Q. Ye. Operations management of smart logistics: A literature review and future research. *Frontiers of Engineering Management*, 8(3):344–355, 2021.
- [42] K. Fischer, J. P. Müller, and M. Pischel. Cooperative transportation scheduling: an application domain for dai. *Applied Artificial Intelligence*, 10(1):1–34, 1996.
- [43] M. Gansterer and R. F. Hartl. Collaborative vehicle routing: A survey. *European Journal of Operational Research*, 268(1):1–12, 2018.

- [44] S. Gau, T. Leifeld, and P. Zhang. A novel and fast mpc based control strategy for switched linear systems including soft switching cost. In *Proc. of Conference on Decision and Control*, pages 6513–6518, Melbourne, Australia, 2017.
- [45] S. H. Ghenaati and S. Aghaei. Modeling and mpc-based method for planning transportation of multiple oil products in pipeline network. In *Proc. of Iranian Conference on Electrical Engineering*, pages 1145–1150, Yazd, Iran, 2019.
- [46] R. Giusti, C. Iorfida, Y. Li, D. Manerba, S. Musso, G. Perboli, R. Tadei, and S. Yuan. Sustainable and de-stressed international supply-chains through the synchro-net approach. *Sustainability*, 11(4), 2019.
- [47] R. Giusti, D. Manerba, G. Bruno, and R. Tadei. Synchromodal logistics: An overview of critical success factors, enabling technologies, and open research issues. *Transportation Research Part E: Logistics and Transportation Review*, 129:92–110, 2019.
- [48] R. Giusti, D. Manerba, and R. Tadei. Multiperiod transshipment location–allocation problem with flow synchronization under stochastic handling operations. *Networks*, 78(1):88–104, 2021.
- [49] T. Godwin, R. Gopalan, and T.T. Narendran. Locomotive assignment and freight train scheduling using genetic algorithms. *International Transactions in Operational Research*, 13(4):299–332, 2006.
- [50] R. Gondhalekar and C. N. Jones. Mpc of constrained discrete-time linear periodic systems — a framework for asynchronous control: Strong feasibility, stability and optimality via periodic invariance. *Automatica*, 47(2):326–333, 2011.
- [51] J. Grahovac and A. Chakravarty. Sharing and lateral transshipment of inventory in a supply chain with expensive low-demand items. *Management Science*, 47(4):579–594, 2001.
- [52] J. M. Grosso, C. Ocampo-Martínez, V. Puig, and B. Joseph. Chance-constrained model predictive control for drinking water networks. *Journal of Process Control*, 24(5):504–516, 2014.
- [53] V. Gumuskaya, W. van Jaarsveld, R. Dijkman, P. Grefen, and A. Veenstra. Dynamic barge planning with stochastic container arrivals. *Transportation Research Part E: Logistics and Transportation Review*, 144:102161, 2020.

- [54] W. Guo, W. B. van Blokland, and G. Lodewijks. Survey on characteristics and challenges of synchromodal transportation in global cold chains. In *Lecture Notes in Computer Science*, volume 10572, pages 420–434, 2017.
- [55] W. Guo, B. Atasoy, W. B. van Blokland, and R. R. Negenborn. Dynamic and stochastic shipment matching problem in multimodal transportation. *Transportation Research Record*, 2674(2):262–273, 2020.
- [56] W. Guo, B. Atasoy, W.B. van Blokland, and R.R. Negenborn. A dynamic shipment matching problem in hinterland synchromodal transportation. *Decision Support Systems*, 134, 2020.
- [57] W. Guo, B. Atasoy, W. B. van Blokland, and R. R. Negenborn. Global synchromodal transport with dynamic and stochastic shipment matching. *Transportation Research Part E: Logistics and Transportation Review*, 152:102404, 2021.
- [58] W. Guo, B. Atasoy, W.B. van Blokland, and R.R. Negenborn. Anticipatory approach for dynamic and stochastic shipment matching in hinterland synchromodal transportation. *Flexible Services and Manufacturing Journal*, 2021.
- [59] A. Haseltalab and R. R. Negenborn. Adaptive control for a class of partially unknown non-affine systems: Applied to autonomous surface vessels. *IFAC-PapersOnLine*, 50(1):4252–4257, 2017.
- [60] A. Haseltalab, F. Wani, and R. R. Negenborn. Multi-level model predictive control for all-electric ships with hybrid power generation. *International Journal of Electrical Power & Energy Systems*, 135: 107484, 2022.
- [61] D. He, L. Zhou, and Z. Sun. Energy-efficient receding horizon trajectory planning of high-speed trains using real-time traffic information. *Control Theory and Technology*, 18(2):204–216, 2020.
- [62] T. Hipólito, J. Nabais, R. B. C. Benítez, M. A. Botto, and R. R. Negenborn. Mpc approach for synchronized supply chains of perishable goods. In *Proc. of International Conference on Industrial Engineering and Systems Management*, Saarbrücken, Germany, 2017.
- [63] W. Hofman. Towards a federated infrastructure for the global data pipeline. In *Lecture Notes in Computer Science*, volume 9373, pages 479–490, 2015.

- [64] W. Hofman. Toward large-scale logistics interoperability based on an analysis of available open standards. In *Proc.of the I-ESA Conferences*, volume 9, pages 249–261, 2019.
- [65] A. Hosseini and T. Sahlin. An optimization model for management of empty containers in distribution network of a logistics company under uncertainty. *Journal of Industrial Engineering International*, 2018.
- [66] J. D. Hunt and K. J. Stefan. Tour-based microsimulation of urban commercial movements. *Transportation Research Part B: Methodological*, 41(9):981–1013, 2007.
- [67] IRU. New IRU survey shows driver shortages to soar in 2021. <https://www.iru.org/news-resources/newsroom/new-iru-survey-shows-driver-shortages-soar-2021>.
- [68] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- [69] A. A Juan, P. Keenan, R. Martí, S. McGarraghy, J. Panadero, P. Carroll, and D. Oliva. A review of the role of heuristics in stochastic optimisation: from metaheuristics to learnheuristics. *Annals of Operations Research*, pages 1–31, 2021.
- [70] K. Kalicharan, F. Phillipson, and A. Sangers. Cutting planes for solving logistic flow problems. In *Lecture Notes in Computer Science*, volume 12433, pages 569–583, 2020.
- [71] M. Khakdaman, J. Rezaei, and L.A. Tavasszy. Shippers’ willingness to delegate modal control in freight transportation. *Transportation Research Part E: Logistics and Transportation Review*, 141, 2020.
- [72] N. S. Kim and B. Van Wee. Toward a better methodology for assessing co2 emissions for intermodal and truck-only freight systems: A european case study. *International Journal of Sustainable Transportation*, 8(3):177–201, 2014.
- [73] S. . Kim and D. Ha. A jit lot-splitting model for supply chain management: Enhancing buyer-supplier linkage. *International Journal of Production Economics*, 86(1):1–10, 2003.
- [74] J. Kopp and F. Woittennek. Flatness based trajectory planning and open-loop control of shallow-water waves in a tube. *Automatica*, 122: 109251, 2020.

- [75] M. A. Krajewska, H. Kopfer, G. Laporte, S. Ropke, and G. Zaccour. Horizontal cooperation among freight carriers: Request allocation and profit sharing. *Journal of the Operational Research Society*, 59(11): 1483–1491, 2008.
- [76] W. Langson, I. Chrysochoos, S. V. Raković, and D. Q. Mayne. Robust model predictive control using tubes. *Automatica*, 40(1):125–133, 2004.
- [77] R. B. Larsen, A. Carron, and M. N. Zeilinger. Safe learning for distributed systems with bounded uncertainties. *IFAC-PapersOnLine*, 50(1):2536–2542, 2017.
- [78] R. B. Larsen, B. Atasoy, and R. R. Negenborn. Simultaneous planning of container and vehicle-routes using model predictive control. In *Proc. of European Control Conference*, pages 2177–2182, Naples, Italy, 2019.
- [79] R. B. Larsen, B. Atasoy, and R. R. Negenborn. Learning-based co-planning for improved container, barge and truck routing. In Eduardo Lalla-Ruiz, Martijn Mes, and Stefan Voß, editors, *Lecture Notes in Computer Science*, pages 476–491, Twente, The Netherlands, 2020.
- [80] R. B. Larsen, B. Atasoy, and R. R. Negenborn. Model predictive control for simultaneous planning of container and vehicle routes. *European Journal of Control*, 57:273–283, 2021.
- [81] R. B. Larsen, R. Baksteen, B. Atasoy, and R. R. Negenborn. Secure multi-party co-planning of barge departures. *IFAC-PapersOnLine*, 54(2):335–341, 2021.
- [82] R. B. Larsen, J. M. Sprokkereef, B. Atasoy, and R. R. Negenborn. Integrated mode choice and vehicle routing for container transport. In *Proc. of International Intelligent Transportation Systems Conference*, pages 3348–3353, Indianapolis, United States, 2021.
- [83] R. Lattarulo and J. P. Rastelli. A hybrid planning approach based on mpc and parametric curves for overtaking maneuvers. *Sensors*, 21(2): 1–19, 2021.
- [84] M. Lazar and W.P.M.H. Heemels. Predictive control of hybrid systems: Input-to-state stability results for sub-optimal solutions. *Automatica*, 45(1):180–185, 2009.
- [85] N. Lemmens, J. Gijbrecchts, and R. Boute. Sychromodality in the physical internet – dual sourcing and real-time switching between transport modes. *European Transport Research Review*, 11(1), 2019.

- [86] J.K. Lenstra and A.H.G. Rinnooy Kan. Computational complexity of discrete optimization problems. In P.L. Hammer, E.L. Johnson, and B.H. Korte, editors, *Discrete Optimization I*, volume 4 of *Annals of Discrete Mathematics*, pages 121–140. Elsevier, 1979.
- [87] L. Li, R. R. Negenborn, and B. De Schutter. Intermodal freight transport planning – a receding horizon control approach. *Transportation Research Part C: Emerging Technologies*, 60:77–95, 2015.
- [88] L. Li, R. R. Negenborn, and B. De Schutter. Distributed model predictive control for cooperative synchromodal freight transport. *Transportation Research Part E: Logistics and Transportation Review*, 105: 240–260, 2017.
- [89] S. Li, M. M. Dessouky, L. Yang, and Z. Gao. Joint optimal train regulation and passenger flow control strategy for high-frequency metro lines. *Transportation Research Part B: Methodological*, 99:113–137, 2017.
- [90] S. Li, R. R. Negenborn, and J. Liu. Stimulating inland waterway transport between seaports and the hinterland from a coordination perspective. In *Proc. of International Conference on Computational Logistics*, pages 67–85, Vietri sul Mare, Italy, 2018. Springer.
- [91] H. Lin and P. J. Antsaklis. Stability and stabilizability of switched linear systems: A survey of recent results. *Transactions on Automatic Control*, 54(2):308–322, 2009.
- [92] X. Lin, R.R. Negenborn, and G. Lodewijks. Towards quality-aware control of perishable goods in synchromodal transport networks. In *IFAC-PapersOnLine*, volume 49, pages 132–137, 2016.
- [93] X. Lin, R. R. Negenborn, M. B. Duinkerken, and G. Lodewijks. Quality-aware modeling and optimal scheduling for perishable good distribution networks: The case of banana logistics. In Tolga Bektaş, Stefano Coniglio, Antonio Martinez-Sykora, and Stefan Voß, editors, *Computational Logistics*, pages 483–497. Springer International Publishing, 2017.
- [94] X. Lin, R. R. Negenborn, and G. Lodewijks. Predictive quality-aware control for scheduling of potato starch production. *Computers and electronics in agriculture*, 150:266–278, 2018.
- [95] J. Löfberg. YALMIP : A toolbox for modeling and optimization in matlab. In *Proc. of International Symposium on Computer Aided Control System Design*, pages 284–289, 2004.

- [96] J. Los, F. Schulte, M. T.J. Spaan, and R. R. Negenborn. The value of information sharing for platform-based collaborative vehicle routing. *Transportation Research Part E: Logistics and Transportation Review*, 141:102011, 2020.
- [97] IMPJ Lucassen and T Dogger. Synchromodality pilot study: Identification of bottlenecks and possibilities for a network between rotterdam, moerdijk and tilburg. 2012.
- [98] J. M. Maestre, D. M. de la Pena, and E. F. Camacho. Distributed mpc: a supply chain case study. In *Proc. of the Conference on Decision and Control held jointly with Chinese Control Conference*, pages 7099–7104, Shanghai, China, 2009. IEEE.
- [99] J. M. Maestre, R. R. Negenborn, et al. *Distributed model predictive control made easy*, volume 69. Springer, 2014.
- [100] L. M. Maiyar and J. J. Thakkar. Robust optimisation of sustainable food grain transportation with uncertain supply and intentional disruptions. *International Journal of Production Research*, 58(18):5651–5675, 2020.
- [101] D. Mayne. An apologia for stabilising terminal conditions in model predictive control. *International Journal of Control*, 86(11):2090–2095, 2013.
- [102] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [103] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 6 2000. ISSN 0005-1098.
- [104] P. R.C. Mendes, J. M. Maestre, C. Bordons, and J. E. Normey-Rico. A practical approach for hybrid distributed mpc. *Journal of Process Control*, 55:30–41, 2017.
- [105] M. R. K. Mes and M. E. Iacob. Synchromodal transport planning at a logistics service provider. In *Logistics and Supply Chain Innovation*, pages 23–36. Springer, 2015.
- [106] M. S. Milenković, N. J. Bojović, L. Švadlenka, and V. Melichar. A stochastic model predictive control to heterogeneous rail freight car fleet sizing problem. *Transportation Research Part E: Logistics and Transportation Review*, 82:162–198, 2015.

- [107] H. Mo and G. Farid. Nonlinear and adaptive intelligent control techniques for quadrotor UAV—a survey. *Asian Journal of Control*, 21(2): 989–1008, 2019.
- [108] L. Moccia, J.-F. Cordeau, G. Laporte, S. Ropke, and M. P. Valentini. Modeling and solving a multimodal transportation problem with flexible-time and scheduled services. *Networks*, 57(1):53–68, 2011.
- [109] M. Morari and M. Barić. Recent developments in the control of constrained hybrid systems. *Computers & Chemical Engineering*, 30(10):1619–1631, 2006.
- [110] J. L. Nabais, R. R. Negenborn, and M. A. Botto. Model predictive control for a sustainable transport modal split at intermodal container hubs. In *Proc. of International Conference on Networking, Sensing and Control*, pages 591–596, Evry, France, 2013.
- [111] J. L. Nabais, R. R. Negenborn, R. Carmona-Benítez, and M. A. Botto. Cooperative relations among intermodal hubs and transport providers at freight networks using an mpc approach. In F. Corman, S. Voß, and R. R. Negenborn, editors, *Proc. of International Conference on Computational Logistics*, pages 478–494, Delft, the Netherlands, 2015. Springer.
- [112] J.L. Nabais, R.R. Negenborn, R.B.C. Benitez, and M.A. Botto. A constrained mpc heuristic to achieve a desired transport modal split at intermodal hubs. In *Proc. of IEEE Conference on Intelligent Transportation Systems*, pages 714–719, The Hague, the Netherlands, 2013.
- [113] A. Nagurney, J. Dong, and D. Zhang. A supply chain network equilibrium model. *Transportation Research Part E: Logistics and Transportation Review*, 38(5):281–303, 2002.
- [114] Z. K. Nagy and R. D. Braatz. Open-loop and closed-loop robust optimal control of batch processes using distributional and worst-case analysis. *Journal of Process Control*, 14(4):411–422, 2004.
- [115] V. V. Naik and A. Bemporad. Embedded mixed-integer quadratic optimization using accelerated dual gradient projection. *IFAC-PapersOnLine*, 50(1):10723–10728, 2017.
- [116] John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.
- [117] A. Nedić and J. Liu. Distributed optimization for control. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):77–103, 2018.

- [118] J. Neupert, E. Arnold, K. Schneider, and O. Sawodny. Tracking and anti-sway control for boom cranes. *Control Engineering Practice*, 18(1):31–44, 2010.
- [119] I. A. Ntousakis, I. K. Nikolos, and M. Papageorgiou. Optimal vehicle trajectory planning in the context of cooperative merging on highways. *Transportation Research Part C: Emerging Technologies*, 71:464–488, 2016.
- [120] J. Nubert, J. Köhler, V. Berenz, F. Allgöwer, and S. Trimpe. Safe and fast tracking on a robot manipulator: Robust mpc and neural network control. *IEEE Robotics and Automation Letters*, 5(2):3050–3057, 2020.
- [121] M.R. Ortega del Vecchyo, F. Phillipson, and A. Sangers. Alternative performance indicators for optimizing container assignment in a synchromodal transportation network. In *Lecture Notes in Computer Science*, volume 11184, pages 222–235, 2018.
- [122] G. Pannocchia, J. B. Rawlings, and S. J. Wright. Conditions under which suboptimal nonlinear mpc is inherently robust. *Systems & Control Letters*, 60(9):747–755, 2011.
- [123] I. Papamichail, A. Kotsialos, I. Margonis, and M. Papageorgiou. Coordinated ramp metering for freeway networks - a model-predictive hierarchical control approach. *Transportation Research Part C: Emerging Technologies*, 18(3):311–331, 2010.
- [124] G. Perboli, S. Musso, M. Rosano, R. Tadei, and M. Godel. Synchromodality and slow steaming: New business perspectives in freight transportation. *Sustainability*, 9(10), 2017.
- [125] E. Perea-Lopez, B. E. Ydstie, and I. E. Grossmann. A model predictive control strategy for supply chain optimization. *Computers & Chemical Engineering*, 27(8-9):1201–1218, 2003.
- [126] A.E. Pérez Rivera and M.R.K. Mes. Integrated scheduling of drayage and long-haul operations in synchromodal transport. *Flexible Services and Manufacturing Journal*, 31(3):763–806, 2019.
- [127] M. C.-P. Poo and T. L. Yip. An optimization model for container inventory management. *Annals of Operations Research*, 273(1):433–453, 2019.
- [128] C.de Prada, I. Grossmann, D. Sarabia, and S. Cristea. A strategy for predictive control of a mixed continuous batch process. *Journal of Process Control*, 19(1):123–137, 2009.

- [129] H. N. Psaraftis, M. Wen, and C. A. Kontovas. Dynamic vehicle routing problems: Three decades and counting. *Networks*, 67(1):3–31, 2016.
- [130] W. Qu, J. Rezaei, Y. Maknoon, and L. Tavasszy. Hinterland freight transportation replanning model under the framework of synchronomodality. *Transportation Research Part E: Logistics and Transportation Review*, 131:308–328, 2019.
- [131] M. Ramezani and M. Nourinejad. Dynamic modeling and control of taxi services in large-scale urban networks: A macroscopic approach. *Transportation Research Part C: Emerging Technologies*, 94: 203–219, 2018.
- [132] J. B. Rawlings. Tutorial overview of model predictive control. *IEEE Control Systems Magazine*, 20(3):38–52, 2000.
- [133] V. Reis. Should we keep on renaming a +35-year-old baby? *Journal of Transport Geography*, 46:173–179, 2015.
- [134] H Giray Resat and Metin Turkyay. A discrete-continuous optimization approach for the design and operation of synchronomodal transportation networks. *Computers & Industrial Engineering*, 130:512–525, 2019.
- [135] Hannah Ritchie. Cars, planes, trains: where do co2 emissions from transport come from?, 2020. URL <https://ourworldindata.org/co2-emissions-from-transport>.
- [136] A. E. P. Rivera and M. R. K. Mes. Scheduling drayage operations in synchronomodal transport. In *Lecture Notes in Computer Science*, volume 10572, pages 404–419, 2017.
- [137] A. E. P. Rivera and M. R. K. Mes. Anticipatory freight selection in intermodal long-haul round-trips. *Transportation Research Part E: Logistics and Transportation Review*, 105:176–194, 2017.
- [138] A. E. P. Rivera and M. R. K. Mes. Integrated scheduling of drayage and long-haul operations in synchronomodal transport. *Flexible services and manufacturing journal*, 31(3):763–806, 2019.
- [139] A. P. Rivera and M. Mes. Service and transfer selection for freights in a synchronomodal network. In *Lecture Notes in Computer Science*, pages 227–242. Springer, 2016.
- [140] T. Rudolf, T. Schurmann, S. Schwab, and S. Hohmann. Toward holistic energy management strategies for fuel cell hybrid electric vehicles in heavy-duty applications. *Proceedings of the IEEE*, 109(6):1094–1114, 2021.

- [141] S. Sager, H. G. Bock, and M. Diehl. The integer approximation error in mixed-integer optimal control. *Mathematical Programming*, 133(1):1–23, 2012.
- [142] H. Sarimveis, P. Patrinos, C. D. Tarantilis, and C. T. Kiranoudis. Dynamic modeling and control of supply chain systems: A review. *Computers and Operations Research*, 35(11):3530–3561, 2008.
- [143] G. Schildebach and M. Morari. Scenario-based model predictive control for multi-echelon supply chain management. *European Journal of Operational Research*, 252(2):540–549, 2016.
- [144] F. Schulte, R. G. González, and S. Voß. Reducing port-related truck emissions: Coordinated truck appointments to reduce empty truck trips. In Francesco Corman, Stefan Voß, and Rudy R. Negenborn, editors, *Computational Logistics*, pages 495–509. Springer International Publishing, 2015.
- [145] M. Schwenzer, M. Ay, T. Bergs, and D. Abel. Review on model predictive control: an engineering perspective. *The International Journal of Advanced Manufacturing Technology*, 117(5):1327–1349, 2021.
- [146] P. O. M. Scokaert, D. Q. Mayne, and J. B. Rawlings. Suboptimal model predictive control (feasibility implies stability). *Transactions on Automatic Control*, 44(3):648–654, 1999.
- [147] P. Singh, M. van Sinderen, and R. Wieringa. Synchromodal transport: pre-requisites, activities and effects. In *Proc. of International Conference on Information Systems Logistics and Supply Chain*, pages 1–4, Bordeaux, France, 2016.
- [148] P.M. Singh. Developing a service oriented it platform for synchromodal transportation. In *Lecture Notes in Computer Science*, volume 8842, pages 30–36, 2014.
- [149] P.M. Singh and M. Van Sinderen. Interoperability challenges for context-aware logistic services - the case of synchromodal logistics. In *Proc. of CEUR Workshop*, volume 1414, Nimes, France, 2015.
- [150] I. I. Sirmatel and N. Geroliminis. Economic model predictive control of large-scale urban road networks via perimeter control and regional route guidance. *IEEE Transactions on Intelligent Transportation Systems*, 19(4):1112–1121, 2018.
- [151] J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.

- [152] M. Steadieseifi, N.P. Dellaert, W. Nuijten, T. Van Woensel, and R. Raoufi. Multimodal freight transportation planning: A literature review. *European Journal of Operational Research*, 233(1):1–15, 2014.
- [153] Z. Su, A. Jamshidi, A. Núñez, S. Baldi, and B. De Schutter. Multi-level condition-based maintenance planning for railway infrastructures – a scenario-based chance-constrained approach. *Transportation Research Part C: Emerging Technologies*, 84:92–123, 2017.
- [154] S. Soroush T. Aval and N. Eghbal. Feedback-based cooperative ramp metering for highway traffic flow control: A model predictive sliding mode control approach. *International Journal of Robust and Nonlinear Control*, 30(18):8259–8277, 2020.
- [155] E. Taniguchi, M. Noritake, T. Yamada, and T. Izumitani. Optimal size and location planning of public logistics terminals. *Transportation Research Part E: Logistics and Transportation Review*, 35(3):207–222, 1999.
- [156] L. Tavasszy, R. Janssen, L. Van der Lugt, and L. Hagdorn. Verkenning synchromodaal transportsysteem. *Delft, The Netherlands*, 2010.
- [157] H. Treiblmaier, K. Mirkovski, P. B. Lowry, and Z. G. Zacharia. The physical internet as a new supply chain paradigm: a systematic literature review and a comprehensive framework. *International Journal of Logistics Management*, 31(2):239–287, 2020.
- [158] A. Tsertou, A. Amditis, E. Latsa, I. Kanellopoulos, and M. Kotras. Dynamic and synchromodal container consolidation: The cloud computing enabler. In *Transportation Research Procedia*, volume 14, pages 2805–2813, 2016.
- [159] H. van Essen, L. van Wijngaarden, A. Schroten, S. de Bruyn, D. Sutter, C. Bieler, S. Maffii, M. Brambilla, D. Fiorello, F. Fermi, R. Parolin, and K. El Beyrouty. *Handbook on the External Costs of Transport*. CE Delft, 2019. URL www.cedelft.eu.
- [160] B. Van Riessen, R. R. Negenborn, and R. Dekker. Synchromodal container transportation: an overview of current topics and research opportunities. In F. Corman, S. Voß, and R. R. Negenborn, editors, *Proc. of International Conference on Computational Logistics*, pages 386–397, Delft, The Netherlands, 2015. Springer.

- [161] B. Van Riessen, R. R. Negenborn, R. Dekker, and G. Lodewijks. Service network design for an intermodal container network with flexible transit times and the possibility of using subcontracted transport. *International Journal of Shipping and Transport Logistics*, 7(4):457–478, 2015.
- [162] B. Van Riessen, R. R. Negenborn, G. Lodewijks, and R. Dekker. Impact and relevance of transit disturbances on planning in intermodal container networks using disturbance cost analysis. *Maritime Economics & Logistics*, 17(4):440–463, 2015.
- [163] B. van Riessen, R.R. Negenborn, and R. Dekker. Synchronodal container transportation: An overview of current topics and research opportunities. In *Lecture Notes in Computer Science*, volume 9335, pages 386–397, 2015.
- [164] B. Van Riessen, R.R. Negenborn, G. Lodewijks, and R. Dekker. Impact and relevance of transit disturbances on planning in intermodal container networks using disturbance cost analysis. *Maritime Economics and Logistics*, 17(4):440–463, 2015.
- [165] B. van Riessen, R.R. Negenborn, and R. Dekker. Real-time container transport planning with decision trees based on offline obtained optimal solutions. *Decision Support Systems*, 89:1–16, 2016.
- [166] B. van Riessen, R.R. Negenborn, and R. Dekker. The cargo fare class mix problem for an intermodal corridor: revenue management in synchronodal container transportation. *Flexible Services and Manufacturing Journal*, 29(3-4):634–658, 2017.
- [167] B. Van Riessen, J. Mulder, R.R. Negenborn, and R. Dekker. Revenue management with two fare classes in synchronodal container transportation. *Flexible Services and Manufacturing Journal*, 33(3): 623–662, 2021.
- [168] P. Velarde, J. M. Maestre, H. Ishii, and R. R. Negenborn. Vulnerabilities in lagrange-based distributed model predictive control. *Optimal Control Applications and Methods*, 39(2):601–621, 2018.
- [169] Kees Verweij. Synchronic modalities—critical success factors. *Logistics yearbook edition*, pages 75–88, 2011.
- [170] M. Vidović, G. Radivojević, and B. Raković. Vehicle routing in containers pickup up and delivery processes. *Procedia - Social and Behavioral Sciences*, 20:335–343, 2011.

- [171] S. Wang, M. Zhu, I. Kaku, G. Chen, and M. Wang. An improved discrete pso for tugboat assignment problem under a hybrid scheduling rule in container terminal. *Mathematical Problems in Engineering*, 2014, 2014.
- [172] W. Wang and D. E Rivera. Model predictive control for tactical decision-making in semiconductor manufacturing supply chain management. *Transactions on Control Systems Technology*, 16(5):841–855, 2008.
- [173] Z. Y. Wu and Y. J. Liu. Research on real time scheduling of emergency transportation based on mpc in post-disaster dynamic environment. In *Proc.of International Symposium on Project Management*, pages 852–861, Wuhan, China, 2017.
- [174] S. X. Xu, G. Q. Huang, and M. Cheng. Truthful, budget-balanced bundle double auctions for carrier collaboration. *Transportation Science*, 51(4):1365–1386, 2017.
- [175] Y. Xu, C. Cao, B. Jia, and G. Zang. Model and algorithm for container allocation problem with random freight demands in synchromodal transportation. *Mathematical Problems in Engineering*, 2015, 2015.
- [176] J. Yan and R. R. Bitmead. Incorporating state estimation into model predictive control and its application to network traffic control. *Automatica*, 41(4):595–604, 2005.
- [177] R. Yang, C.-Y. Lee, Q. Liu, and S. Zheng. A carrier–shipper contract under asymmetric information in the ocean transport industry. *Annals of Operations Research*, 273(1-2):377–408, 2019.
- [178] H. Yee, J. Gijbrecchts, and R. Boute. Synchromodal transportation planning using travel time information. *Computers in Industry*, 125, 2021.
- [179] O. O. Özener and O. Ergun. Allocating costs in a collaborative transportation procurement network. *Transportation Science*, 42(2):146–165, 2008.
- [180] L. Zhang, S. Zhuang, and R. D. Braatz. Switched model predictive control of switched linear systems: Feasibility, stability and robustness. *Automatica*, 67:8–21, 2016.
- [181] Y. Zhang, B. Atasoy, and R. R. Negenborn. Preference-based multi-objective optimization for synchromodal transport using adaptive large neighborhood search. volume 0, pages 71–87, 2021.

-
- [182] H. Zheng, J. Wu, W. Wu, and R. R. Negenborn. Cooperative distributed predictive control for collision-free vehicle platoons. *Intelligent Transport Systems*, 13(5):816–824, 2019.
 - [183] F. Zhu and P. J. Antsaklis. Optimal control of hybrid switched systems: A brief survey. *Discrete Event Dynamic Systems*, 25(3):345–364, 2015.
 - [184] B.G. Zweers, S. Bhulai, and R.D. van der Mei. Planning hinterland container transportation in congested deep-sea terminals. *Flexible Services and Manufacturing Journal*, 33(3):583–622, 2021.

Abbreviations

MPC	Model predictive control
CO2	Carbon dioxide
ICT	Information communications technology, the term is used in its broadest sense (Chapter 2)
LSP	Logistics Service Provider (Chapter 4)
FSO	Flexible Service Operator (Chapter 4)
MDS	Memory-based discrete search (Chapter 5)

Curriculum vitae

Rie Brammer Larsen was born the first of October 1988 near Aarhus, Denmark. She finished her bachelor in Mechanical Engineering in 2013 from University of Southern Denmark, Odense, Denmark. Parts of her studies took place at Konkuk University, Seoul, South Korea. In 2015, she finished her Master of Engineering in Mechatronics with specialization in mathematical modelling and control from University of Southern Denmark, Soenderborg, Denmark. From August 2015 until January 2018 she worked in research and consultancy at University of Southern Denmark, Soenderborg, Denmark; ETH, Zurich, Switzerland; and Jacobs Engineering Group, the Hague, the Netherlands. In February 2018 she started a Ph.D. project in the Section of Transport Engineering & Logistics, Department of Maritime & Transport Technology, Delft University of Technology, the Netherlands, under the supervision of Prof. Rudy R. Negenborn and Dr. Bilge Atasoy. The project concerns real-time methods to facilitate (cooperative) planning in synchromodal transport networks and the results are presented in this dissertation. Her research interests lie in the intersection of real-time control, operational planning and artificial intelligence with a strong interest in multi-agent systems. She is fascinated by the challenges arising from co-planning between heterogeneous stakeholders with an uneven power-balance.

Publications

1. Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. A learning-based co-planning method with truck and container routing for improved barge departure times. Submitted to a journal.

2. Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. Model predictive control for simultaneous planning of container and vehicle routes. *European Journal of Control*, volume 57, pages 273–283, 2021.
3. Rie B. Larsen, Ruud Baksteen, Bilge Atasoy, and Rudy R. Negenborn. Secure multiparty co-planning of barge departures. *IFAC-PapersOnLine for IFAC Symposium on Control in Transportation Systems*. Volume 54(2), pages 335–341, 2021.
4. Rie B. Larsen, Jasper M. Sprokkereef, Bilge Atasoy, and Rudy R. Negenborn. Integrated mode choice and vehicle routing for container transport. In *Proc. of the International Intelligent Transportation Systems Conference*, pages 3348–3353, 2021.
5. Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. Model Predictive Control with Memory-based Discrete Search for Switched Linear Systems. *IFAC-PapersOnLine for IFAC World Congress*, volume 53, pages 6769–6774, 2020.
6. Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. Learning-based co-planning for improved container, barge and truck routing. In Eduardo Lalla-Ruiz, Martijn Mes, and Stefan Voß, editors, In *Proc. of the International Conference on Computational Logistics*, pages 476–491, 2020.
7. Rie B. Larsen, Bilge Atasoy, and Rudy R. Negenborn. Simultaneous planning of container and vehicle-routes using model predictive control. In *Proc. of European Control Conference*, pages 2177–2182, 2019.
8. Rie B. Larsen, Andrea Carron, and Melanie N. Zeilinger. Safe Learning for Distributed Systems with Bounded Uncertainties. *IFAC-PapersOnLine for IFAC World Congress*, volume 50, pages 2536–2542, 2017.
9. Rie B. Larsen, Jerome Jouffroy, and Benny Lassen. On the premature convergence of particle swarm optimization. In *Proc. of European Control Conference*, pages 1922–1927, 2016.

Abstracts and invited talks

1. Abstract presented at the TRAIL PhD Congress, Utrecht, the Netherlands, 2018. Title: Towards Predictive Sychromodality using Model Predictive Control
2. Abstract presented at the Modeling, Analysis, and Control of Complex Networks and Cyber-Physical Systems – workshop, Ischia, Italy, 2019. Title: Applying MPC to Container Transport Planning
3. Abstract accepted for the TRAIL PhD Congress, Utrecht, the Netherlands, 2019. Title: Real-Time Sychromodal Transport Planning at Operational Level
4. Nominated for submission and thereupon approved for participation in the ELA Doctorate Workshop, Lappeenranta, Finland, 2020. Did not participate due to COVID-19. Title: Predictive Sychromodality for more Efficient Container Transport
5. Abstract accepted for the Benelux Meeting on Systems and Control, Elspeet, The Netherlands, 2020. Title: Model Predictive Control for Integrated Sychromodal Transport
6. Invited to talk at the physical event which, due to COVID-19, was replaced with a shortened online edition. Netherlands OML Conference, online, 2020. Title: Real-time sychromodal transport planning

Grants

1. Erasmus+ staff mobility for teaching and training. Used for two weeks exchange stay with prof.dr. José María Maestre Torreblanca at Systems and Automation Engineering Department, University of Seville, Spain. 2019

Summary

Container transport is an essential part of the well-functioning, highly specialized, and global production chains society currently relies on. To improve the utilization of resources, it is important to ensure all processes are as efficient as possible. Synchronodal transport is a recent transport paradigm which seeks to increase the efficiency of freight transport by letting transport providers change the mode of transport of goods in real-time. This new flexibility alleviates some of the obstacles to using sustainable transport modes, e.g., barges and trains, as it simplifies the process of changing transport plans if something unpredicted happens, such as delays, cancellations or if shipping requests that were announced later makes different routing smarter. Furthermore, synchronodal transport can improve the utilization of the transport vehicles, as the freight can be routed using up-to-date information about vehicle availability.

Usually, multiple stakeholders are involved when goods are transported. Ensuring good coordination between stakeholders improves the efficiency of transport and often reduces costs. When coordinating different stakeholders' plans, the main challenge is that each organization is, typically, only willing to share limited information and needs to maintain its autonomy.

New methods are needed to transform the flexibility of synchronodal transport into efficiency improvements and costs reductions in transport networks with multiple stakeholders. The research presented in this dissertation provides insights into how container transport can realistically be planned in real-time when several different stakeholders own the vehicles.

Model predictive control (MPC) is a method that combines optimization of future plans with feedback control. The method is suitable for real-time control of complex systems and, therefore, forms the basis of the operational planning methods for synchronodal transport, which are presented in this dissertation. One centralized method which integrates the routing of

containers and vehicles has been presented as well as two novel and distinctively different co-planning methods and theoretical proof that the ideas behind learning-aided co-planning method can result in a stabilizing and recursively feasible MPC.

The research shows that the integration of container and vehicle routes indeed improves the efficiency of transport systems. Numerical experiments were carried out on realistic models of the transport networks that connect Rotterdam port, the Netherlands, with its hinterland. Centralized, real-time planning methods that integrate container and vehicle routes enable the transports to be carried out satisfactorily with a smaller vehicle fleet when compared with real-time planning that separate the two routing problems. This is especially pronounced in transport systems where more empty vehicles need repositioning, e.g., due to peaks in the transport demand, or irregular availability of scheduled, spot-priced transport services. The frequent reconsideration of transport plans provided by the MPC based methods decreases the operational cost of the transport systems. The studied transport systems perform better when all decisions are reconsidered in an integrated fashion compared to only routing newly announced shipping requests and adjusting plans that render infeasible. In transport systems with multiple stakeholders, the performance improves significantly when synchromodal co-planning methods are used. If co-planning is not used, the resulting sequential decision-making causes inefficient use of transport vehicles. If information exchange is limited, learning can help infer some of the missing information which makes co-planning more efficient.

Samenvatting

Containervervoer is een essentieel onderdeel van de goed functionerende, zeer gespecialiseerde, wereldwijde productieketens waarop de samenleving momenteel vertrouwt. Om het gebruik van resources te verbeteren, is het belangrijk om ervoor te zorgen dat alle processen zo efficiënt mogelijk zijn. Synchronodaal transport is een recente transportparadigma met als doel de efficiëntie van vrachtvervoer te verhogen door vervoerders de vervoermodaliteit van goederen in realtime te laten veranderen. Deze nieuwe flexibiliteit reduceert de belemmeringen bij het gebruik van duurzame vervoermodaliteiten, zoals binnenvaartschepen en treinen, omdat het proces bij wijzingen van vervoersplannen vereenvoudigt wanneer er iets onberekenbaars gebeurt, zoals vertragingen, annuleringen of als later aangekondigde verzendverzoeken een andere route maken slimmer. Bovendien kan synchronodaal transport het gebruik van de capaciteit van de transportvoertuigen verbeteren, omdat de vracht kan worden gerouteerd met behulp van actuele informatie over de beschikbaarheid van voertuigen.

Bij het transporteren van goederen zijn doorgaans meerdere stakeholders betrokken. Een goede coördinatie tussen de belanghebbenden zorgt er voor dat de efficiëntie van het vervoer verbeterd en de kosten vaak worden verlaagd. De grootste uitdaging bij het coördineren van de verschillende plannen van de belanghebbenden is dat elke organisatie meestal slechts beperkte informatie willen delen en haar autonomie moet behouden.

Nieuwe methoden zijn nodig om de flexibiliteit van synchronodaal transport om te zetten in efficiëntieverbeteringen en kostenreducties in transportnetwerken met meerdere stakeholders. Het onderzoek dat in dit proefschrift wordt gepresenteerd, geeft inzicht in hoe containervervoer realistisch in realtime kan worden gepland wanneer verschillende belanghebbenden eigenaar zijn van de voertuigen.

Model Predictive Control (MPC) is een methode die optimalisatie van toekomstige plannen combineert met feedbackcontrole. De methode is geschikt voor real-time besturing van complexe systemen en vormt daarom de basis van de operationele planningsmethoden voor synchrodaal transport die in dit proefschrift worden gepresenteerd. Een gecentraliseerde methode wordt gepresenteerd die de routing van containers en voertuigen integreert, evenals twee nieuwe en duidelijk verschillende co-planningsmethoden en een theoretisch bewijs dat de ideeën achter de learning gestuurde co-planningsmethode kunnen resulteren in een stabiliserende en herhaalbare MPC.

Uit het onderzoek blijkt dat integratie van container- en voertuigroutes inderdaad de efficiëntie van transportsystemen verbetert. Er zijn numerieke experimenten uitgevoerd op realistische modellen van de transportnetwerken die de Rotterdamse haven met het achterland verbinden. Gecentraliseerde, realtime planningsmethoden die container- en voertuigroutes integreren, kunnen de transporten voldoende uitvoeren met een kleiner wagenpark in vergelijking met realtime planning die de twee soort routing scheidt. Dit is vooral duidelijk in transportsystemen waar meer lege voertuigen moeten worden verplaatst, bijvoorbeeld vanwege pieken in de transportvraag of onregelmatige beschikbaarheid van geplande, spotgeprijsde transportdiensten. De herhaaldelijke heroverweging van transportplannen die door de op MPC gebaseerde methoden plaatsvinden, verlaagt de operationele kosten van de transportsystemen. De onderzochte transportsystemen presteren beter wanneer alle beslissingen op een geïntegreerde manier worden heroverwogen in vergelijking met het alleen beschouwen de routing van nieuw aangekondigde verzendverzoeken en het aanpassen van transportplannen die onhaalbaar zijn geworden. In transportsystemen met meerdere belanghebbenden verbeteren de prestaties aanzienlijk wanneer synchrodale co-planningsmethoden worden gebruikt. Als er geen gebruik wordt gemaakt van co-planning, leidt de daaruit voortvloeiende sequentiële besluitvorming tot inefficiënt gebruik van transportvoertuigen. Als de informatie-uitwisseling beperkt is, kan kunstmatige intelligentie methoden gebruikt worden om een deel van de ontbrekende informatie af te leiden, wat co-planning efficiënter maakt.

TRAIL Thesis Series

The following list contains the most recent dissertations in the TRAIL Thesis Series. For a complete overview of more than 200 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.

Larsen, R.B., *Real-time Co-planning in Sychromodal Transport Networks using Model Predictive Control*, T2022/9, September 2022, TRAIL Thesis Series, the Netherlands

Zeinaly, Y., *Model-based Control of Large-scale Baggage Handling Systems: Leveraging the theory of linear positive systems for robust scalable control design*, T2022/19, June 2022, TRAIL Thesis Series, the Netherlands

Fahim, P.B.M., *The Future of Ports in the Physical Internet*, T2022/7, May 2022, TRAIL Thesis Series, the Netherlands

Huang, B., *Assessing Reference Dependence in Travel Choice Behaviour*, T2022/6, May 2022, TRAIL Thesis Series, the Netherlands

Reggiani, G., *A Multiscale View on Bikeability of Urban Networks*, T2022/5, May 2022, TRAIL Thesis Series, the Netherlands

Paul, J., *Online Grocery Operations in Omni-channel Retailing: opportunities and challenges*, T2022/4, March 2022, TRAIL Thesis Series, the Netherlands

Liu, M., *Cooperative Urban Driving Strategies at Signalized Intersections*, T2022/3, January 2022, TRAIL Thesis Series, the Netherlands

Feng, Y., *Pedestrian Wayfinding and Evacuation in Virtual Reality*, T2022/2, January 2022, TRAIL Thesis Series, the Netherlands

Scheepmaker, G.M., *Energy-efficient Train Timetabling*, T2022/1, January 2022, TRAIL Thesis Series, the Netherlands

Bhoopalam, A., *Truck Platooning: planning and behaviour*, T2021/32, December 2021, TRAIL Thesis Series, the Netherlands

Hartleb, J., *Public Transport and Passengers: optimization models that consider travel demand*, T2021/31, TRAIL Thesis Series, the Netherlands

Azadeh, K., *Robotized Warehouses: design and performance analysis*, T2021/30, TRAIL Thesis Series, the Netherlands

Chen, N., *Coordination Strategies of Connected and Automated Vehicles near On-ramp Bottlenecks on Motorways*, T2021/29, December 2021, TRAIL Thesis Series, the Netherlands

Onstein, A.T.C., *Factors influencing Physical Distribution Structure Design*, T2021/219, December 2021, TRAIL Thesis Series, the Netherlands

Olde Kalter, M.-J. T., *Dynamics in Mode Choice Behaviour*, T2021/27, November 2021, TRAIL Thesis Series, the Netherlands

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

Nguyen, 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

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