



Coordination of Heavy-Duty Vehicle Platooning

SEBASTIAN VAN DE HOEF

Doctoral Thesis
Stockholm, Sweden 2018

KTH Royal Institute of Technology
School of Electrical Engineering and Computer Science
Department of Automatic Control
SE-100 44 Stockholm
SWEDEN

TRITA-EECS-AVL-2018:22
ISBN 978-91-7729-722-2

Akademisk avhandling som med tillstånd av Kungliga Tekniska högskolan fram-
lägges till offentlig granskning för avläggande av teknologie doktorsexamen i Elektro-
och systemteknik fredagen den 27 april 2018 klockan 10.00 i F3, Kungliga Tekniska
högskolan, Lindstedtsvägen 26, Stockholm, Sverige.

© Sebastian van de Hoef, April 2018. All rights reserved.

Tryck: Universitetservice US AB

Abstract

A network-wide coordination system for heavy-duty vehicle platooning with the purpose of reducing fuel consumption is developed. Road freight is by far the dominating mode for overland transport with over 60% modal share in the OECD countries and is thus critically important for the economy. Overcoming its strong dependency on fossil fuels and manual labor as well as handling rising congestion levels are therefore important societal challenges. Heavy-duty vehicle platooning is a promising near-term automated-driving technology. It combines vehicle-to-vehicle communication and on-board automation to slipstream in a safe manner, which can reduce fuel consumption by more than 10%. However, in order to realize these benefits in practice, a strategy is needed to form platoons in an operational context. We propose a platoon coordination system that supports the process of automatically forming platoons over large geographic areas.

We develop an architecture in which fleet management systems send start locations, destinations, and arrival deadlines to a platoon coordinator. By computing desirable speed profiles and by letting the vehicles' on-board systems track them, vehicles can meet en route and form platoons. Matching vehicles into platoons and deriving suitable speed profiles is treated as an optimization problem with the objective of maximizing the overall fuel savings under the constraint that vehicles arrive in time at their destinations. By updating the speed profiles and the platoon configurations based on real-time measurements of vehicle position and platoon state, the system can accommodate new vehicles joining on the fly. Using real-time measurements also makes the system resilient to disturbances and changing operating conditions. This thesis seeks to develop the theoretical foundations of such a system and evaluate its potential to improve transport efficiency.

We first explore the coordination of vehicle pairs. Fuel-optimal speed profiles are derived. The uncertainty arising from traffic is taken into account by modeling travel time distributions and considering the probability of two vehicles successfully merging. Building on this coordination algorithm for vehicle pairs, we derive algorithms for larger platoons and vehicle fleets. This results in an NP-hard combinatorial optimization problem. The problem is formulated as an integer program and results on the solution structure are derived. In order to handle realistic fleet sizes with thousands of vehicles and continental sized geographical areas under real-time operation, heuristic algorithms are developed. The speed profiles resulting from the combinatorial optimization are further improved using convex optimization. Moreover, we derive efficient algorithms to identify all pairs of vehicles that can platoon. Simulations demonstrate that the proposed algorithm is able to compute plans for thousands of vehicles. Coordinating approximately a tenth of Germany's heavy-duty vehicle traffic, platooning rates over 65% can be achieved and fuel consumption can be reduced by over 5%. The proposed system was implemented in a demonstrator system. This demonstrator system has been used in experiments on public roads that show the technical feasibility of en route platoon coordination.

Sammanfattning

Ett nätverksomfattande system för att koordinera körning i lastbilskolonner i syfte att minska bränsleförbrukningen utvecklas i denna avhandling. Vägtransport är med marginal det mest dominerande sättet för landtransport med en andel över 60 % i OECD-länderna och är således avgörande för samhällsekonomin. Dess starka beroende av fossila bränslen och arbetskraft samt ökande trängsel är därför viktiga utmaningar för samhället. Vägtransport är på väg att genomgå grundläggande förändringar genom elektrifiering, kommunikation och automation. Lastbilskolonnkörning är en teknologi för att automatisera fordon, som är redo att lanseras inom en snar framtid. Kolonnkörning kan ge bränslebesparingar på över 10 % tack vare att automatisering och kommunikation mellan fordon möjliggör så kallad slipstreaming på ett säkert sätt. Det behövs dock en strategi för att sätta ihop lastbilskolonner för att kunna dra nytta av denna teknologi. Vi föreslår därför ett automatiserat koordineringssystem som stödjer processen att sätta ihop lastbilskolonner inom stora geografiska områden.

Vi utvecklar en arkitektur där fleet management system skickar start- och målpunkter samt senaste ankomsttid till koordineringssystemet. Genom att beräkna passande hastighetsprofiler och genom att låta lastbilarnas färdatorer följa dem kan lastbilar mötas på vägen och bygga ihop kolonner. Problemet att matcha fordon till kolonner och att komma fram till passande hastighetsprofiler hanteras som ett optimeringsproblem med målet att maximera den totala bränslebesparingen. Genom att uppdatera lösningen baserade på realtidsmätningar av fordonspositioner och kolonntillstånd kan systemet hantera tillkommande lastbilar och kan vara robust gentemot störningar och förändrade operativa förutsättningar. Denna avhandling strävar efter att utveckla teoretiska förutsättningar för ett sådant system och att utvärdera dess potential att öka transportsystemets effektivitet.

Vi börjar med att undersöka koordinering av fordonspar och härleder bränsleoptimala hastighetsprofiler. Den osäkerhet som beror på trafiken tas hänsyn till genom att modellera hastighetsfördelningar och genom att betrakta sannolikheten att två fordon träffas och kan påbörja kolonnkörning. Vi bygger på denna algoritm för fordonspar för att härleda algoritmer som kan hantera större kolonner och större antal lastbilar. Det resulterar i ett NP-svårt kombinatoriskt optimeringsproblem. Problemet formuleras som ett heltalsoptimeringsproblem och vi kommer fram till resultat som visar strukturen av lösningen. Resulterande hastighetsprofiler förbättras ytterligare med hjälp av konvex optimering. Förutom detta utvecklas effektiva algoritmer för att identifiera alla fordonspar som kan köra i kolonn. Simuleringar demonstrerar att den föreslagna algoritmen klarar att beräkna planer för tusentals lastbilar. Vi visar att en samordning av c:a 10 % av Tysklands lastbilstrafik kan tillåta att mer än 65 % av den totala körsträckan körs i kolonn och att bränsleåtgången kan minskas med mer än 5 %. Det presenterade systemet implementerades i ett demonstratorsystem. Experiment med demonstratorn på allmän väg har visat att koordinering av kolonnkörning på vägen är tekniskt genomförbart.

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my supervisors Dimos Dimarogonas and Karl Henrik Johansson. Their continuous inspiration and feedback have made this thesis possible, and it has been an excellent experience to work with them. I would like to thank Jonas Mårtensson for many interesting discussions and support, working together on projects, courses, and joint papers. I would like to thank the people that I have been working with in course of this thesis. In particular, I would like to mention my co-authors Håkan Terelius, Bart Besselink, Valerio Turri, Kuo-Yun Liang, Assad Alam, and Panagiotis Tsiotras. Thanks to all participants in the COMPANION project for a very enjoyable collaboration and enabling testing coordinated en route platoon formation in reality. I am grateful to Matteo Vanin, Ari Hauksson, David Spulak, and Christian Lindeborg for their commitment in implementing and maintaining our part of the demonstrator system. I would like to express my gratitude to Jeroen Ploeg and the team from the TNO Integrated Vehicle Safety department for making an exiting research visit possible. I would like to thank Mario Romero Vega for our collaboration on visualizing platoon planning.

I would like to thank Meng Guo, Antonio Adaldo, Demia Della Penda, David Umsonst, Christos Verginis, and Alexandros Nikou for the great collaboration in developing and running the hybrid and embedded control course; Pedro Lima and Valerio Turri for together supervising a really nice project course.

I would like to extend my gratitude to Mladen Čičić, Manne Held, Alexander Johansson, Ehsan Nekouei, Meng Guo, Pedro Lima, and Riccardo Risuleo for proofreading parts of this thesis.

I am thankful to present and past colleagues in the automatic control department for together creating such an enjoyable working environment. All the discussions over lunch, joint trips to conferences, after works, presentations, research discussions, reading groups, joint course work and so on, has made the time at the department really special for me. I want to thank in particular those who put time and effort in organizing workshops, reading groups, social events, and taking care of the office space. I would like to thank the administrators at the department for keeping the lab running.

The research in this thesis has been financed by the School of Electrical Engineering through the program of excellence, by the European Union through the COMPANION project, by the Swedish Research Council, and the Knut and Alice Wallenberg Foundation. I am grateful for their support.

Last but not least, I want to thank my family and friends for their support and company, in particular Anni for her dedication in proofreading my texts and both her and my parents for morally supporting me.

Sebastian van de Hoef
Stockholm, April 2018

Contents

Acknowledgments	v
Contents	vii
1 Introduction	1
1.1 Motivation	2
1.2 Platoon Coordination Example	8
1.3 Problem Formulation	10
1.4 Thesis Outline and Contributions	14
2 Background	21
2.1 Freight Transport	21
2.2 Information and Communication Technology in Transportation Systems	26
2.3 Vehicle Platooning	31
2.4 Platoon Coordination	36
2.5 Summary	42
3 Coordinated Platooning System	45
3.1 System Architecture	45
3.2 Platoon Coordinator Architecture	49
3.3 Summary	52
4 Coordination Algorithms for Pairs of Vehicles	53
4.1 The Structure of Adapted Platoon Plans	54
4.2 The Optimal Rendezvous Speed	55
4.3 Adapted Platoon Plans for a Fixed Speed Range	59
4.4 Adapted Platoon Plans for Position-Dependent Speeds	62
4.5 Summary	66
5 Coordination Algorithms under Uncertain Travel Times	69
5.1 Modeling	70
5.2 Optimal Speed Control	72

5.3	Modeling Controlled Travel-Time Distributions	77
5.4	Simulations	78
5.5	Summary	82
6	Coordination Algorithms for Many Vehicles	85
6.1	Combining Pairwise Plans by Selecting Coordination Leaders	85
6.2	Exact Solution to the Coordination Leader Selection Problem	89
6.3	Heuristic and Approximate Solutions to the Coordination Leader Selection Problem	97
6.4	Joint Platoon-Plan Optimization	102
6.5	Summary	105
7	Computational Efficiency Improvements	107
7.1	Candidate Platoon Pairs	108
7.2	Culling Candidate Platoon Pairs	111
7.3	Features and Classifiers for Culling Platoon Pairs	113
7.4	Simulations	118
7.5	Summary	121
8	Evaluation	123
8.1	Simulations on an Artificial Road Network	123
8.2	Simulations on a Real Road Network	129
8.3	Dynamic Simulations and Sensitivity Analysis	136
8.4	Experimental Evaluation	142
8.5	Summary	151
9	Conclusions and Future Work	153
9.1	Conclusions	153
9.2	Directions for Future Work	155
	Bibliography	157

Chapter 1

Introduction

TRANSPORTATION is a key aspect in any human civilization. Already in ancient times, trade over land and water was indispensable and without the ability to transport large quantities of goods more efficiently, the industrial revolution would have been impossible. Today's society and economy more than ever depend on the ability to move freight and people fast, efficiently, and reliably. In order to meet these requirements in the future and to satisfy the growing transport demand in a sustainable way, major developments in how we organize transport have to be implemented. Those developments are driven by automation, connectivity, and electrification, facilitated by a range of technological advances, most notably in the field of information and communication technology.

Heavy-duty vehicle platooning, also referred to as truck platooning, is one of the technologies that can help create a more sustainable transport system in the near future. Platooning has been shown to decrease fuel consumption and greenhouse gas emissions, to increase road-network capacity, and to improve safety. It is a first step towards the automation of road vehicles. However, for heavy-duty vehicle platooning to have a significant impact, an approach for its large-scale coordination is needed. This thesis introduces a framework for such coordination and demonstrates its ability to leverage the potential of heavy-duty vehicle platooning in practice.

In Section 1.1, we further motivate why improving goods transportation is important for society. Therein, we argue that road freight transport is in many aspects superior to other transport systems. However, road freight transport poses a number of challenges to remain sustainable. Heavy-duty vehicle platooning can help overcome some of these challenges. We argue, that a coordination system is needed to benefit from platooning on large scale. The positive effects and feasibility of such a coordination system are motivated by an example in Section 1.2. Section 1.3 formulates the problem of developing an integrated system to coordinate heavy-duty vehicle platooning. The remainder of this thesis presents a framework for such a system and highlights different design aspects as well as its potential and its limitations. Section 1.4 gives an overview of the thesis and summarizes its contributions

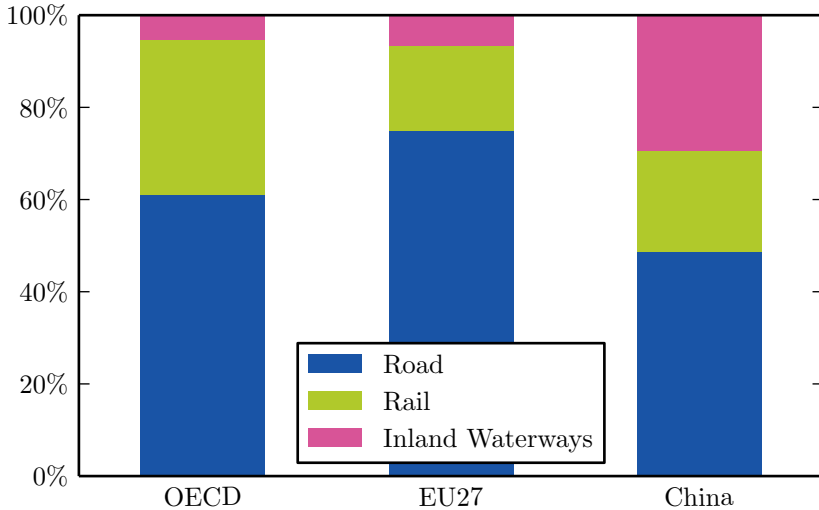


Figure 1.1: The modal split of surface freight transport volumes in the OECD countries, in the EU27 countries, and in China in the year 2014. In all three regions road transport is the dominating form of transport. (Source: OECD/ITF [169])

chapter by chapter. Furthermore, we indicate where the presented material has been published.

1.1 Motivation

Goods transport is critical for the economy, and transport volumes are tightly coupled with economic prosperity [201, 169, 76]. For example, in the European Union, the entire transport sector accounts for 5% of the gross domestic product [2]. Developments in transportation systems are key enablers for industrial development. Without improved ships and the invention of railways, the industrial revolution could not have taken place. Over the last decades, road freight transport has become the dominant mode of surface freight transport [4, 169]. In the European Union, about three quarters of all surface freight is transported by road, see Figure 1.1. This is due to a number of advantages road freight transport with trucks has compared to alternative means of transportation such as rail or inland waterways.

Trucks are a very flexible way of transporting goods. They can reach virtually every location that goods need to be transported to or from. The organizational

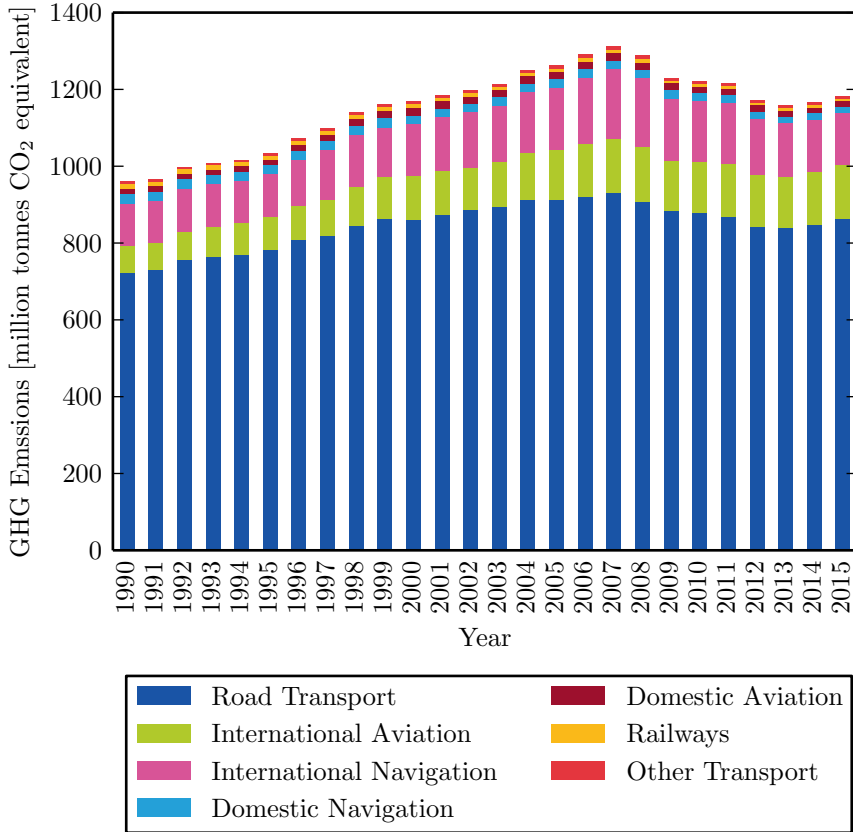


Figure 1.2: Greenhouse gas emissions of different transport modes in the EU28 measured in million tonnes of CO₂ equivalent. Road transport is consistently the largest source of greenhouse gas emissions in transportation. (Source: Eurostat [7])

overhead of trucks is low and many operators are small companies [63]. This enables quick adaptation to changing demands, and competition keeps prices low [25]. Since a truckload is relatively small, it is often possible to transport goods directly from source to destination with little overhead for combining different transports in order to fill the vehicle [43]. Road freight transport has, on the other hand, a number of challenges to overcome. Caused by the decentralized nature of road infrastructure, transport times can vary due to traffic and they are difficult to predict [239]. Costs

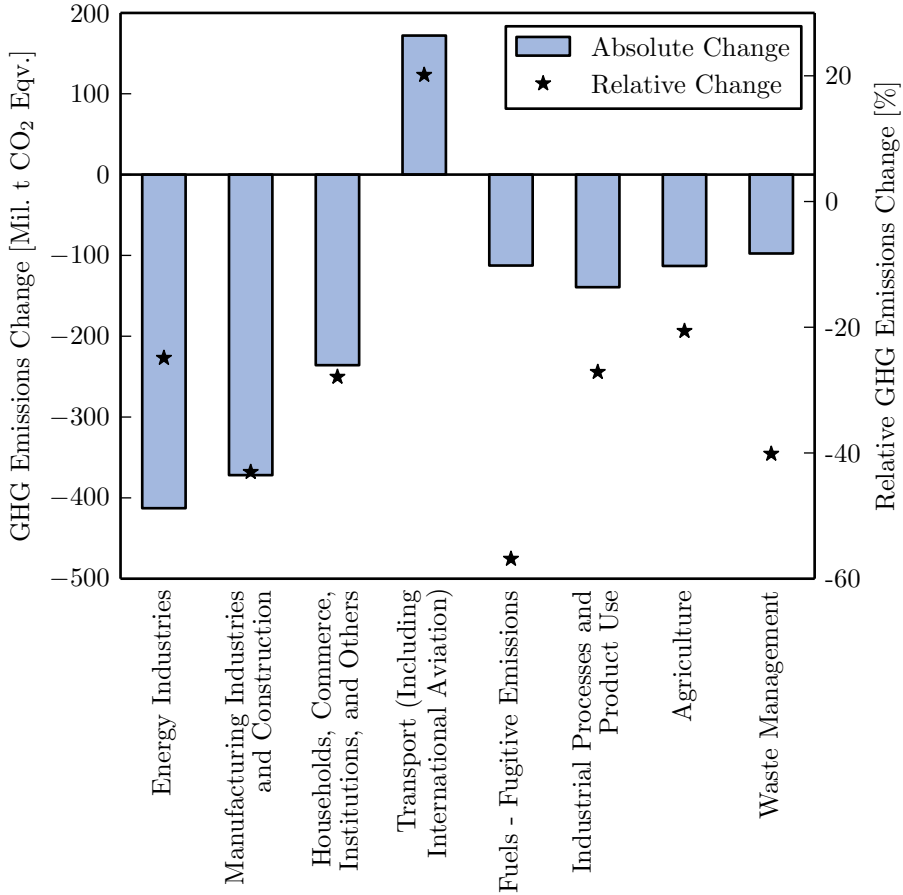


Figure 1.3: Absolute and relative change of greenhouse gas (GHG) emissions by different sectors in the EU28 from 1990 until 2015. The absolute change is measured in million tonnes of CO₂ equivalent. Transport is the only sector where greenhouse gas emissions are still rising. (Source: Eurostat [7])

due to congestion are high [216]. Furthermore, every truck needs a driver, and vehicles are limited in size compared to other transport modes. This leads to high labor costs [10, 110], accounting to roughly a third of the total operation cost of a heavy-duty freight truck in long-haulage operation. Driving a truck over long distances on highways can be at times a monotonous task. Nevertheless, the driver's full attention is required, since even short moments during which a driver

is not attentive can lead to fatal accidents [19]. Another problem is that the great majority of trucks is powered by fossil fuels, and despite various research efforts such as electric highways and alternative fuels, this is not likely to change soon, in particular in the domain of long haulage transport [54, 169]. The use of fossil fuels leads to problematic emissions, most prominently carbon-dioxide [169]. In 2014, the transport sector accounted for 20 % of the greenhouse gas emissions in the European Union, of which 72 % were due to road transport (Figure 1.2). Roughly one third of the emissions from road transport is caused by freight vehicles [82]. The demand for road freight transport is expected to continue growing [9]. Opposed to other sectors, emissions from transport are still growing [169] as shown in Figure 1.3.

Heavy-duty vehicle platooning is a technology that can help solve some of these problems. In platooning a group of vehicles forms a road train without any physical coupling between the vehicles. A short inter-vehicle distance is maintained by automatic control and vehicle-to-vehicle communication. Figure 1.4 shows two demonstrations of vehicle platooning. The small inter-vehicle spacing leads to an improved road throughput and the automatic control of the trailing vehicles improves safety [199, 88]. Similar to what racing cyclists exploit, the follower vehicles and, to a lesser degree, the lead vehicles experience a reduction in air drag, which translates into reduced fuel consumption [17, 18, 42, 134, 217, 247, 251]. Reduced fuel consumption, in turn, implies decreased greenhouse gas emissions. Since fuel accounts for roughly a quarter to a third of a heavy truck's operation costs in long haulage transport [10, 110], there is an economic incentive to introduce platooning technology. Platooning can also be a way of fully automating the trailing vehicles so that the driver can rest. This in turn can lead to increased productivity in case this time can count as rest time as required by the hours of service regulations [118]. Even scenarios in which no driver in the trailing vehicles is required are possible, leading to even higher economic gains.

The business case for platooning can be established considering fuel consumption reduction alone. The marginal cost per truck will be small in near future, since most hardware needed will be standard as it is used in other advanced driver assistance systems [36, 103, 124]. As a result, even small gains in fuel economy are sufficient to offset these costs. Productivity gained from increased driving times or even driverless vehicles in the platoon makes the business case even more interesting. Platooning is seen as one of the first commercially viable applications of heavy-duty vehicle automation [124, 90].

Advances in wireless communication, satellite-based positioning, available computing power, and driver support systems in general have made the deployment of platooning systems technically feasible and platooning systems are under active development by most major vehicle manufacturers and widespread adoption is expected in the near future [118, 75]. The European Automobile Manufacturers' Association expects market introduction in Europe as soon as the year 2022 [8].

Integrating platooning into the road freight transport system leads to a challenging coordination problem. While there have been promising demonstrations of intra-platoon control systems [102, 111, 130, 217, 220, 75], the question remains



Figure 1.4: Platooning demonstrations in the scope of the European Truck Platooning Challenge 2016 (Source: European Truck Platooning Challenge [6])

open how to benefit from heavy-duty vehicle platooning at large scale and in a regular operational context. In some special cases, trucks have the entire or the first part of their journey in common, for instance, when leaving from a distribution center. However, such special cases account only for a small fraction of road freight transport.

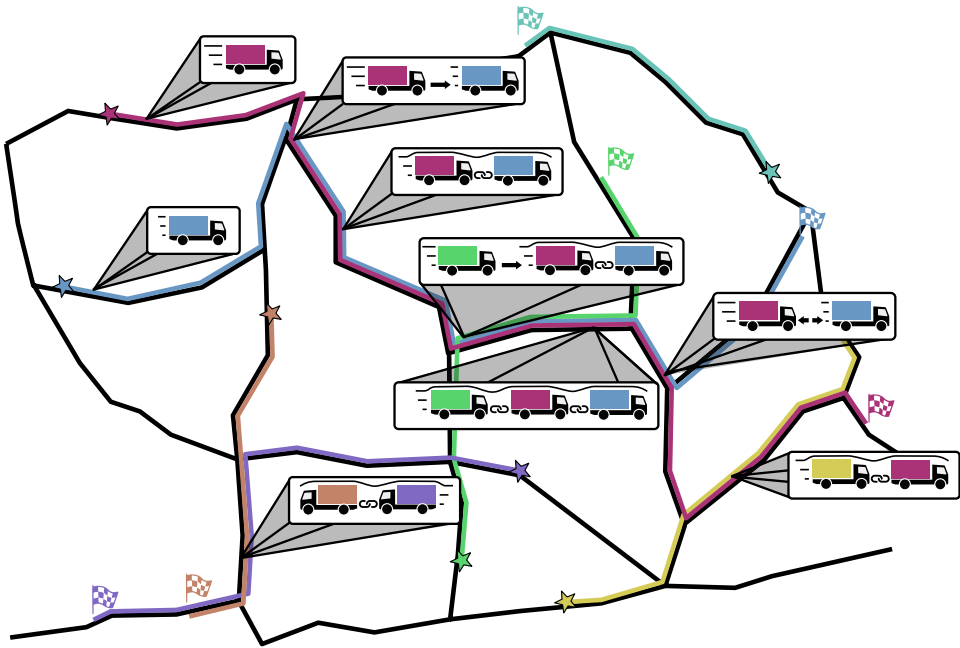


Figure 1.5: Vehicles travel from start to destination in a road network. By small adjustments of speed, platoons can be formed where routes overlap and fuel can be saved.

One approach to benefit from platooning technology is to control the vehicles' reference speeds in order to form platoons en route as illustrated in Figure 1.5. Small adjustments of the reference speeds can be sufficient to let vehicles get into each others' vicinity on the road, platoon for some time, and then split up and continue driving individually. By leveraging such situations systematically, it is possible to maximize the gains from platooning, specifically reduced fuel consumption, while retaining the advantages trucks have over other transportation systems, such as flexibility and independence.

In order to facilitate en route platoon formation across vehicles from different manufacturers and fleet owners, an integrated platoon coordination system as illustrated in Figure 1.6 is needed. This thesis seeks to develop the theoretical foundations of such a system and to evaluate its potential to improve transport efficiency. In order to retain the flexibility of road freight transport and to be robust to disturbances, such a system needs to automatically adapt based on real-time information streams from a number of sources. This makes it different from many other transport planning problems, where dynamic plan updates are infrequent and often require human intervention. Algorithms need to be efficient enough to




Figure 1.6: Illustration of an integrated platoon coordination system. Trucks communicate via vehicle-to-infrastructure communication with the centralized platoon coordinator that coordinates the dynamic formation of platoons and integrates with logistic operators.

run real-time, they must be able to incorporate new information, and they must guarantee correct system operation.

1.2 Platoon Coordination Example

In this section, we discuss the example illustrated in Figure 1.7 that motivates the potential of en route platoon coordination. At time 7:15 the blue vehicle departs at the sea port of Bremerhavn and starts its trip to Kassel. At 7:43 the green vehicle starts a trip to southern Germany at the port of Hamburg. By selecting an average speed of 83 km/h, it gains 2 minutes and 42 seconds compared to driving 80 km/h and forms a platoon with the blue vehicle at 8:55 where both routes merge on the same highway. At 9:15 the purple vehicle starts in Hanover. By selecting an average speed of 75 km/h, it joins the platoon of the other two vehicles at 9:45, 22 km after the point where the routes meet. By traveling at a reduced speed of 75 km/h, it loses 2 minutes and 6 seconds. The three vehicles drive in a platoon for 1 hour and 18 minutes until they arrive at 11:03 at a point where the route of the purple vehicle forces it to split up from the platoon. It continues its way by itself to Eisenach with an average speed of 84 km/h to compensate for the time lost at the beginning of the trip. The blue and the green vehicle continue to platoon until the blue vehicle arrives at 11:32 at its destination in Kassel. The green vehicle continues its trip to a rest stop at 11:53 after 4 hours and 10 minutes of driving.

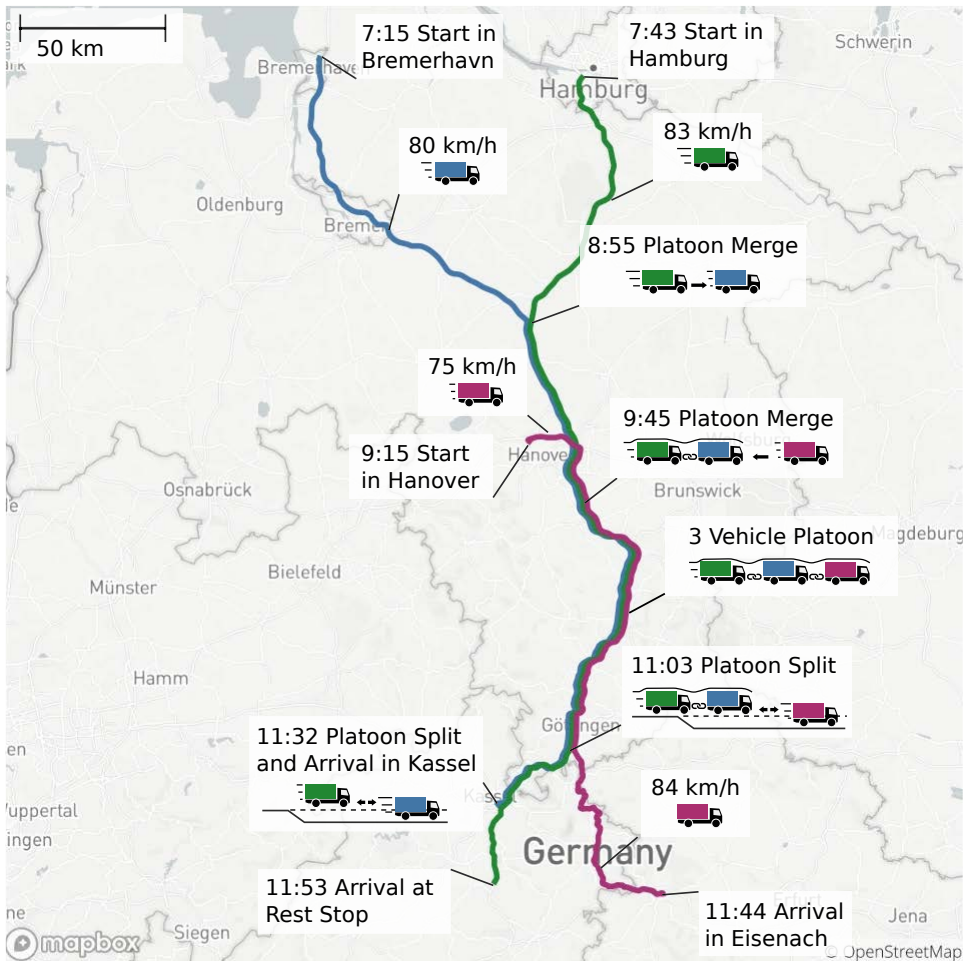


Figure 1.7: Three vehicles with different origin-destination pairs form a platoon en route by small speed adaptations.

In this scenario, the blue vehicle is a platoon follower, i.e., a trailing vehicle in the platoon, for 104 km and the green vehicle for 209 km. In total 34.4% of the 908 km traveled by the three vehicles is traveled as a platoon follower. Assuming an air-drag reduction of 40% when being a platoon follower and using the fuel consumption model from [35], we get 14% of fuel consumption reduction when being a platoon follower assuming the vehicle is fully loaded with a total mass of 40 tons. In absolute values the fuel consumption in liters of diesel is lowered by 4.5 liters per 100 km. In this scenario, 14.1 liters of diesel are saved due to platooning.

The green vehicle selecting a higher speed at the beginning of its trip and the purple vehicle at the beginning and the end of its trip, leads to an additional fuel consumption of 0.86 liters of diesel reducing the overall savings to 13.2 liters diesel or 5.59 % of the fuel consumption without platooning. The total carbon-dioxide reduction is 34.9 kg. This might not seem like a lot at first glance. However, road freight transport with trip length at least 150 km in the EU28 countries was in the year 2016 estimated to be 110 551 million km [12]. Making a conservative assumption that 10 % of that distance can be traveled as a platoon follower instead of the vehicle driving by itself, 497 million liters of diesel would be saved annually corresponding to 1.31 million tonnes CO₂ alone. We show that with a platoon coordination system, high platooning rates can be achieved, with over 40 % of the distance traveled as a platoon follower. Due to higher permitted speed in some countries, for instance the US, even higher potential benefits from platooning can be expected. Recall also that fuel consumption reduction is only one of the benefits from platooning technology. Under future legislation, it might be possible that the 2 hours and 37 minutes the green vehicle spends as platoon follower count as a rest time. This would allow the driver to continue the trip until 16:02 without the 45 minutes of rest mandatory in the European Union after 4.5 hours of driving.

1.3 Problem Formulation

The problem considered in this thesis is to design a coordinated platooning system that supports platooning-enabled vehicles to form platoons en route. As shown in Figure 1.8, it is composed of fleet management systems, a platoon coordination system, and a platoon vehicle system. The platoon vehicle system comprises of platooning-enabled vehicles with index set \mathcal{N} traveling in a road network. The focus is on the development of the platoon coordination system while the fleet management systems and the platoon vehicle system are considered as given.

The road network is represented as a directed graph $\mathcal{G}_r = (\mathcal{N}_r, \mathcal{E}_r)$ with nodes \mathcal{N}_r and edges $\mathcal{E}_r \in \mathcal{N}_r \times \mathcal{N}_r$. Nodes correspond to intersections or endpoints in the road network, and edges correspond to road segments connecting these intersections as shown in Figure 1.10. The function $L : \mathcal{E}_r \rightarrow \mathbb{R}^+$ maps each edge in \mathcal{E}_r to the length of the corresponding road segment. A position in the road network is a pair $(e, x) \in \mathcal{E}_r \times [0, L(e)]$ where e indicates the current road segment and x how far the vehicle has traveled along that segment.

The platoon coordination system is designed as a feedback control system shown in Figure 1.9. The platoon coordination system receives positions $(e_n, x_n) \in \mathcal{E}_r \times \mathbb{R}$ and platoon configuration $\pi_n \in 2^{\mathcal{N}}$ of each vehicle $n \in \mathcal{N}$ and suggests next road segment $e_n^+ \in \mathcal{E}_r$, reference speed $v_n^{\text{ref}} \in \mathbb{R}$ and reference platoon configuration $\pi_n^{\text{ref}} \in 2^{\mathcal{N}}$ in order to fulfill the objective that each vehicle reaches its respective destination \mathcal{P}_n^D before its deadline t_n^D . This should happen with minimal overall fuel consumption leveraging the ability of vehicles to platoon in order to lower fuel consumption. For notational convenience we define $\mathcal{P}_{\mathcal{N}}^D = \{\mathcal{P}_n^D : n \in \mathcal{N}\}$,

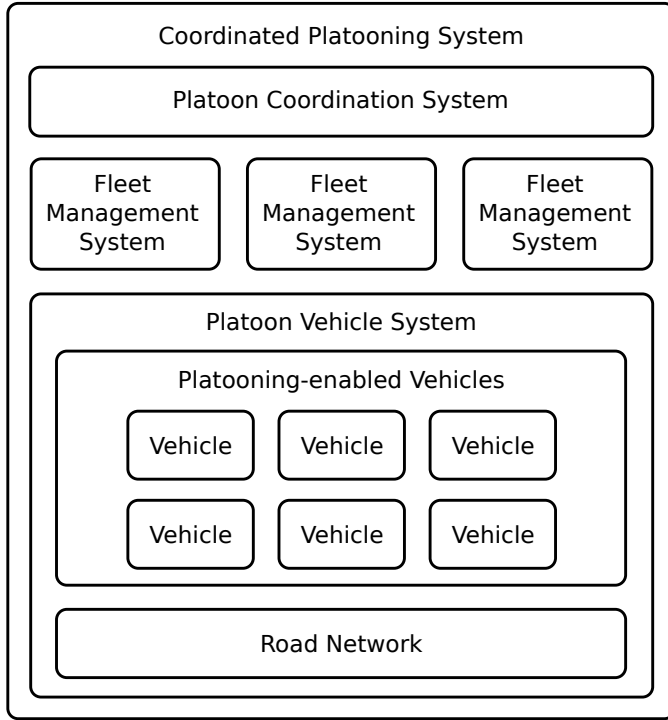


Figure 1.8: The platoon coordination system, the fleet management systems, and the platoon vehicle system form together the coordinated platooning system. The platoon vehicle system is composed of platooning-enabled vehicles traveling in a road network.

$$t_{\mathcal{N}}^D = \{t_n^D : n \in \mathcal{N}\}, e_{\mathcal{N}} = \{e_n : n \in \mathcal{N}\}, x_{\mathcal{N}} = \{x_n : n \in \mathcal{N}\}, \pi_{\mathcal{N}} = \{\pi_n : n \in \mathcal{N}\}, \\ e_{\mathcal{N}}^+ = \{e_n^+ : n \in \mathcal{N}\}, v_{\mathcal{N}}^{\text{ref}} = \{v_n^{\text{ref}} : n \in \mathcal{N}\}, \pi_{\mathcal{N}}^{\text{ref}} = \{\pi_n^{\text{ref}} : n \in \mathcal{N}\}.$$

The evolution of the continuous part of the position x is according to the reference speed v_n^{ref} with the relation $x_n(t) = \int_{t_0}^t v_n^{\text{ref}}(\tau) d\tau$, where t_0 is the time when the vehicle starts traversing the road segment, i.e., $x_n(t_0) = 0$. We assume $v_n^{\text{ref}} \geq 0$.

The vehicle transitions to the next road segment e_n^+ when it reaches the end of a road segment, i.e., when $x_n = L(e_n)$. It starts traversing the next segment e_n^+ at $x_n = 0$. Hereby, the vehicle has to follow a path in \mathcal{G}_r , i.e., if it transitions from $e_n = (n_1, n_2)$, then it has to transition to road segment $e_n^+ \in \{(n_2, n_3) : (n_2, n_3) \in \mathcal{E}_r\}$ for some appropriate n_3 .

The platoon state corresponds to a logical state in the platoon manager indicating that the vehicle-to-vehicle communication is established and the vehicles

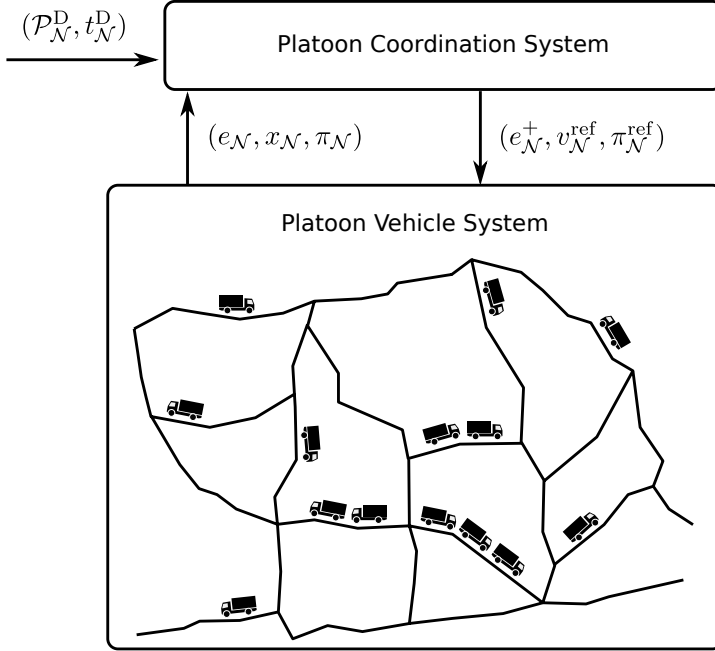


Figure 1.9: The objective of the platoon coordination system is to set reference routing decisions e_n^+ , reference speeds v_n^{ref} , and reference platoon partners π_n^{ref} such that vehicles arrive at their destinations \mathcal{P}_n^D before time t_n^D with minimum overall fuel consumption.

platoon. The platoon manager is a system on board the vehicles that manages the platoon. The platoon state is modeled as a set π_n of vehicle indices in \mathcal{N} . When vehicles platoon, it implies that their inter-vehicle distances are small. We approximate this by requiring that vehicles have the same position when they platoon at time t , i.e.,

$$(e_n(t), x_n(t)) = (e_m(t), x_m(t)), \quad \forall n, m \in \pi_n(t). \quad (1.1)$$

Platoons are formed for those vehicles in the reference platoon configuration π_n^{ref} for which (1.1) is fulfilled, i.e., $\pi_n(t) = \{m \in \pi_n^{\text{ref}}(t) : (e_n(t), x_n(t)) = (e_m(t), x_m(t))\}$.

The instantaneous combined fuel consumption of all vehicles $F_c(\pi_{\mathcal{N}}(t), v_{\mathcal{N}}^{\text{ref}}(t)) \in \mathbb{R}$ at time t depends on the platoon configurations $\pi_{\mathcal{N}}(t)$ and the reference speeds $v_{\mathcal{N}}^{\text{ref}}(t)$. The total fuel consumption for all vehicles at time \underline{t}^h to reach their destinations is

$$F = \int_{\underline{t}^h}^{\bar{t}^h} F_c(\pi_{\mathcal{N}}(\tau), v_{\mathcal{N}}^{\text{ref}}(\tau)) d\tau,$$




Figure 1.10: The road network is modeled as a directed graph where edges correspond to road segments and nodes to connections between road segments. The position (e, x) of a vehicle is represented by and edge e in the road network graph and how far the vehicle has traveled along the segment, x . The platoon configuration π is a set of vehicle identifiers.

where $\bar{t}^h = \max_{n \in \mathcal{N}}(t_n^D)$. Platooning is used as a measure to reduce fuel consumption. Once a vehicle has reached its destination, it stops and does not use any fuel.

The problem considered in this thesis is how to design a platoon coordination system. The coordination system computes control inputs $e_{\mathcal{N}}^+(t), v_{\mathcal{N}}^{\text{ref}}(t), \pi_{\mathcal{N}}^{\text{ref}}(t)$ as a function of the state $e_{\mathcal{N}}(t), x_{\mathcal{N}}(t), \pi_{\mathcal{N}}(t)$. The control objective is that each vehicle reaches its destination $e_n(t_n^A), x_n(t_n^A) = \mathcal{P}_n^D$ before deadline $t_n^A \leq t_n^D$ with minimal fuel consumption F from initial time \underline{t}^h to time \bar{t}^h with initial state $e_{\mathcal{N}}(\underline{t}^h), x_{\mathcal{N}}(\underline{t}^h), \pi_{\mathcal{N}}(\underline{t}^h)$ and with the system dynamics as described above.

Remark 1. The control inputs $e_n^+, v_n^{\text{ref}}, \pi_n^{\text{ref}}$ are processed by the platoon manager and the vehicle controller residing on board the vehicles. The vehicle controller of vehicle n tracks either a reference speed v_n^{ref} or controls the relative position in the platoon. It also ensures that the vehicle speed is always safe and within the legal limits. The platoon manager organizes platoon formation, splitting, and ordering according to the reference platoon configuration π_n^{ref} . The vehicle controller and the platoon manager work in unison to both follow the average speed v_n^{ref} suggested by the platoon coordination system and the reference platoon configuration π_n^{ref} . Refer to [245] for further details on the vehicle controller and the platoon manager.

Remark 2. Disturbances acting on the control inputs e_n^+ , v_n^{ref} , π_n^{ref} are compensated for by feedback from the state measurements e_n, x_n, π_n . Traffic, weather conditions, power limitations of the vehicles, and many other factors can lead to such disturbances. For instance, if the traffic situation is such that the suggested reference speed cannot be implemented, the vehicle controller senses the surrounding vehicles and lowers the speed to maintain a safe following distance. Also the driver can choose to override the reference speed in such situations. The difference between the actual speed and the reference v_n^{ref} is from the perspective of the platoon coordination system an input disturbance.

Remark 3. For the feedback mechanism to function, the state of the platoon vehicle system needs to be measured. The vehicle position $(e_n(t), x_n(t))$ at time t can be measured using global navigation satellite systems, odometry and digital maps. The platoon state $\pi_n(t)$ at time t is provided by the platoon manager on the vehicle.

Remark 4. Vehicles are not connected to the platoon coordination system all the time. The connection periods are determined by the fleet management system the vehicle is associated with and cannot be influenced by the platoon coordination system. When a vehicle is disconnected neither can its state be measured nor can it be controlled. Typically, one period of activity corresponds to reaching a destination from some start position without intermediate stops. Such a destination can also be an intermediate stop for the driver to take a rest on a longer trip. The vehicle might also be disconnected when traveling on roads outside the highway network. We consider in large parts of the thesis that the time and the location a vehicle becomes active in as well as the destination of a vehicle can be predicted some time in advance. Multi-stop vehicle routing problems, rest time planning, staff scheduling etc. are performed by the fleet management systems and they are out of the scope of the platoon coordination system.

1.4 Thesis Outline and Contributions

This section provides an overview of the thesis. It describes each chapter's content and contribution. We also indicate publications in which material used in this thesis has been or is going to be published.

Chapter 2: Background Chapter 2 provides relevant background for the topics considered in the thesis. Section 2.1 gives a brief overview over research on freight transport systems. Road freight transport is one of several alternative transport modes. Each transport mode faces its unique challenges but they have in common that they are complex and large-scale systems that are crucial for the economy. This results in complex decision problems. Heavy-duty vehicle platooning is a relatively new technology that has been made possible by information and communications technology (ICT). ICT is transforming the transportation systems is a multitude

of ways. We review some of these developments in Section 2.2. The topic of this thesis is the coordination of heavy-duty vehicle platooning and it takes the ability of vehicles to be part of a platoon, once formed, for granted. Section 2.3 summarizes work on the control of heavy-duty vehicle platooning, enabling technology, and the effect of platooning on fuel consumption. In Section 2.4, work on the coordination of platooning is reviewed. We structure the work based on several features and discuss the individual contributions.

Chapter 3: Coordinated Platooning System In Chapter 3, we introduce the coordinated platooning system as a feedback system consisting of the platoon coordination system and the platoon vehicle system. In Section 3.1, we introduce the architecture of such a system consisting of four hierarchical layers operating at different time scales. The focus of the thesis is hereby on the platoon coordinator, that interfaces with fleet management systems and data providers in order to accomplish platoon coordination. In Section 3.2, the architecture of the platoon coordinator is discussed. It comprises of several modules, which together compute platoon plans maximizing the predicted fuel consumption reduction from platooning.

The chapter is primarily based on the publication:

- S. van de Hoef, J. Mårtensson, D. V. Dimarogonas, and K. H. Johansson. A Predictive Framework for Dynamic Heavy-Duty Vehicle Platoon Coordination. *ACM Transactions on Cyber-Physical Systems*, 2017. submitted

Chapter 4: Coordination Algorithms for Pairs of Vehicles Chapter 4 considers the computation of fuel-efficient platoon plans for two vehicles. Section 4.1 introduces the structure of such plans. One of the two vehicles adapts its speed at the beginning of the trip in order to meet the other vehicle en route. The two vehicles platoon for some distance until they split up. The speed at the end of the journey is adjusted in a way that the vehicle meet their arrival deadlines. Fuel is saved during the platooning phase.

Section 4.2 considers an affine fuel model taking into account that fuel consumption depends on speed and whether or not the vehicle platoons. The speed is assumed to be freely chosen in a fixed interval. It derives the fuel optimal rendezvous speed to meet the other vehicle and platoon for the rest of the trip. Section 4.3 builds on this result and derives fuel optimal adapted plans according to the structure introduced in Section 4.1. In Section 4.4, we introduce a way to compute adapted plans when the feasible speed range cannot be approximated as constant along the route. Instead, the fastest possible speed profile considering vehicle limitations, legal constraints, and traffic is computed for both vehicles. From that speed profile, the adapted plan is derived.

The chapter is primarily based on the publications:

- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Coordinating Truck Platooning by Clustering Pairwise Fuel-Optimal Plans. In *18th IEEE International Conference on Intelligent Transportation Systems*, pages 408–415, Sept. 2015
- S. van de Hoef, J. Mårtensson, D. V. Dimarogonas, and K. H. Johansson. A Predictive Framework for Dynamic Heavy-Duty Vehicle Platoon Coordination. *ACM Transactions on Cyber-Physical Systems*, 2017. submitted

Chapter 5: Coordination Algorithms under Uncertain Travel Times

Chapter 5 introduces a systematic way to explicitly consider disturbances in traffic in the planning. Section 5.1 models the problem. Like Chapter 4, it considers two vehicles, one adapting to the other. It is assumed that travel times over a road segment are stochastic with a known distribution. This distribution is conditioned on the selected reference speed. The objective is to determine the probability of both vehicles successfully merging into a platoon taking into account that the adapting vehicle adjusts its reference speed to maximize the merge probability after each segment. In Section 5.2, we derive the optimal selection using dynamic programming. An extension to the case that travel time distributions of adjacent segments are correlated is given. Furthermore, we show how computations can be made more efficient in the case that a small error can be taken into account. Section 5.3 discusses how the effect of the reference speed can be modeled in order to obtain travel time distributions from data of vehicles traveling at maximum speed. A simulation study presented in Section 5.4 demonstrates the method.

The chapter is primarily based on the publication:

- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Efficient Dynamic Programming Solution to a Platoon Coordination Merge Problem With Stochastic Travel Times. *20th IFAC World Congress*, 50(1):4228–4233, July 2017

Chapter 6: Coordination Algorithms for Many Vehicles

In chapter 6, we consider how to compute platoon plans that are fuel-efficient for a fleet of vehicles as a whole. This is done by first computing a default plan and a number of adapted plans similar to the ones derived in Chapter 4 for each vehicle. Then, a subset of these plans is selected and combined in order to maximize the total fuel consumption reduction. In Section 6.1, we formulate the problem of selecting and combining plans for two vehicles as a combinatorial optimization problem. In Section 6.2, methods to compute exact solutions to this optimization problem are presented. This includes the formulation as an integer programming problem, which allows the use of general purpose solvers. Results on the structure of the optimal solution are established that can help to reduce the search space. We also

prove that the problem is NP-hard, which means that it might take a lot of computational effort to compute an optimal solution. The result that the optimization problem is NP-hard motivates the development of heuristic algorithms presented in Section 6.3. The first algorithm is similar to heuristic algorithms used in clustering and community detection. It improves the fuel-efficiency of the combined plans in every step until it reaches a local maximum. The other algorithm iterates once over the set of coordinated vehicles. We show that optimization problem is submodular from which follows that the second algorithm guarantees an approximation ratio. Section 6.4 elaborates on how the selected combined plans for two vehicles can be jointly improved using convex optimization. Hereby, the pairs of vehicles that platoon remain fixed, but the speed profile that leads to this platooning is optimized considering all vehicles at once.

The chapter is primarily based on the publications:

- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Fuel-Efficient En Route Formation of Truck Platoons. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):102–112, Jan. 2018
- S. van de Hoef, J. Mårtensson, D. V. Dimarogonas, and K. H. Johansson. A Predictive Framework for Dynamic Heavy-Duty Vehicle Platoon Coordination. *ACM Transactions on Cyber-Physical Systems*, 2017. submitted

Chapter 7: Computational Efficiency Improvements Chapter 7 considers the computationally efficient identification of all vehicle pairs that can potentially platoon. To identify the platoon opportunities for the set of transport assignments is the first step in the computation of platoon plans, which can be computationally expensive. Section 7.1 introduces that pairs of vehicles have to share part of their routes and have to be able to meet on the overlapping segment in order to potentially platoon. All pairs of vehicles that can platoon according to this criterion can be identified by directly comparing route identifiers and time intervals for all pairs of vehicles. In Section 7.2, we propose an efficient approach based on extracting low dimensional features. These features can be used to efficiently dismiss a majority of the pairs that cannot platoon. The remaining pairs can then be processed using a computationally more expensive algorithm that compares the routes explicitly. Section 7.4 presents solutions that demonstrate the effectiveness of this approach using simulations on a real road network.

The chapter is primarily based on the publication:

- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Computing Feasible Vehicle Platooning Opportunities for Transport Assignments. *14th IFAC Symposium on Control in Transportation Systems*, 49(3):43–48, May 2016

Chapter 8: Evaluation Chapter 8 presents an evaluation of the proposed platoon coordination method by simulations and experiments. Section 8.1 presents simulations on an artificially generated road network focusing on systematic parameter variations. In Section 8.2, we study the coordination system on routes in a real road network generated from a population density map of mainland Sweden. We analyze potential fuel consumption reduction and size distribution of the created platoon for different numbers of coordinated vehicles. In Section 8.4, experimental results from a demonstrator system are shown. The experiment took place in September 2016 on public highways in Spain. Three heavy-duty trucks were successfully coordinated to form a platoon en route.

The chapter is primarily based on the publications:

- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Coordinating Truck Platooning by Clustering Pairwise Fuel-Optimal Plans. In *18th IEEE International Conference on Intelligent Transportation Systems*, pages 408–415, Sept. 2015
- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Fuel-Efficient En Route Formation of Truck Platoons. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):102–112, Jan. 2018
- S. van de Hoef, J. Mårtensson, D. V. Dimarogonas, and K. H. Johansson. A Predictive Framework for Dynamic Heavy-Duty Vehicle Platoon Coordination. *ACM Transactions on Cyber-Physical Systems*, 2017. submitted

Chapter 9: Conclusions and Future Work In Chapter 9 we conclude the thesis and discuss future work. Section 9.1 summarizes and discusses the obtained results. Section 9.2 outlines possible ways to continue this work.

Publications Work covered in this thesis has been presented in the following publications.

- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Fuel-Efficient En Route Formation of Truck Platoons. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):102–112, Jan. 2018
- S. van de Hoef, J. Mårtensson, D. V. Dimarogonas, and K. H. Johansson. A Predictive Framework for Dynamic Heavy-Duty Vehicle Platoon Coordination. *ACM Transactions on Cyber-Physical Systems*, 2017. submitted
- B. Besselink, V. Turri, S. van de Hoef, K.-Y. Liang, A. Alam, J. Mårtensson, and K. H. Johansson. Cyber-physical Control of Road Freight Transport. *Proceedings of the IEEE*, 104(5):1128–1141, Mar. 2016
- K.-Y. Liang, S. van de Hoef, H. Terelius, V. Turri, B. Besselink, J. Mårtensson, and K. H. Johansson. Networked Control Challenges in Collaborative Road Freight Transport. *European Journal of Control*, 30:2–14, July 2016
- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Fuel-Optimal Coordination of Truck Platooning Based on Shortest Paths. In *American Control Conference*, pages 3740–3745, Chicago, IL, July 2015
- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Coordinating Truck Platooning by Clustering Pairwise Fuel-Optimal Plans. In *18th IEEE International Conference on Intelligent Transportation Systems*, pages 408–415, Sept. 2015
- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Computing Feasible Vehicle Platooning Opportunities for Transport Assignments. *14th IFAC Symposium on Control in Transportation Systems*, 49(3):43–48, May 2016
- S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Efficient Dynamic Programming Solution to a Platoon Coordination Merge Problem With Stochastic Travel Times. *20th IFAC World Congress*, 50(1):4228–4233, July 2017
- S. van de Hoef. *Fuel-Efficient Centralized Coordination of Truck Platooning*. KTH Royal Institute of Technology, June 2016. Licentiate Thesis

The order of the author names in the publications listed above reflects their contribution.

Chapter 2

Background

THIS chapter provides background on the thesis. We begin in Section 2.1 by giving an overview of how freight transport is organized today. In order to develop a practical approach making use of heavy-duty vehicle platooning, it is important to understand the context in which freight trucks operate. Furthermore, work on planning and on related challenges in freight transportation is reviewed. Transportation systems are heavily influenced by information and communications technology (ICT). Platooning and its coordination is a technology enabled by ICT. Section 2.2 gives a brief overview of developments in this area. Section 2.3 introduces platooning technology and evidence for its potential to reduce fuel consumption. Platooning is classified into different levels of automation and the technologies needed for platooning are discussed. Finally, Section 2.4 reviews in detail existing work on platoon coordination and on the system level effects of platooning.

2.1 Freight Transport

Freight transport systems ensure the movement of goods between stages of the supply chain as illustrated in Figure 2.1. The figure depicts the main modes of transport using vehicles: transport by road, air, water, and rail. Other important elements in freight transport systems are ports, terminals, and warehouses. Ensuring that all these components efficiently work together is a non-trivial problem.

Freight transport is critically important for the world economy and tightly coupled to the GDP [201, 169, 76]. The total transport demand is steadily increasing [169], the reasons being expanding trade, global economic development, globalization, and the ability to handle complicated supply chains. A critical enabler for these developments is indeed ICT and transport systems themselves forming a positive feedback loop. The costs for transportation are significant. Transport accounts for roughly half of the logistic costs and is therefore the single most important cost factor in making goods, once produced, available to the customer [183, 184].

The environmental impact of the current transport system is highly problem-

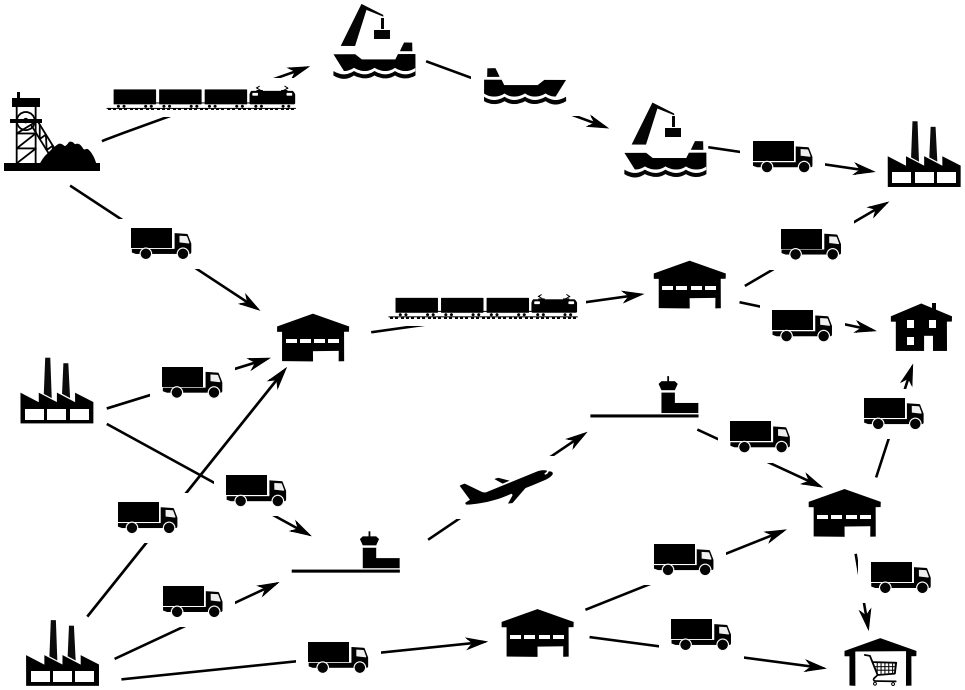


Figure 2.1: Freight transport systems move goods between stages of the supply chain. The main modes of transport are road, air, rail, and water. These modes can be linked into intermodal transport chains.

atic. For transport systems to continue growing in a sustainable way, massive improvements in the efficiency have to be implemented. This concerns handling of the sheer complexity of transport, more efficient use of resources, transition to alternative energy sources, and increased automation [4, 169, 54, 9].

Organization of Freight Transport Freight transport is organized as a complex interaction of interdependent actors [174, 68, 185, 211]. Figure 2.2 gives an overview of the most important roles in the freight transportation process. On the left side of the figure are transport operators. They operate the infrastructure and equipment required to physically move goods. Rail, road, air, barge, and sea carriers operate transport vehicles. Ports and terminal operators provide the necessary infrastructure to move goods from one vehicle to another, in many cases to another mode of transport, but also to reorganize and consolidate goods. In addition, they provide intermediate storage and services such as customs clearance. Warehouses are focused on the storage of goods. They also facilitate packaging and consolidating items, as well as loading them onto vehicles.

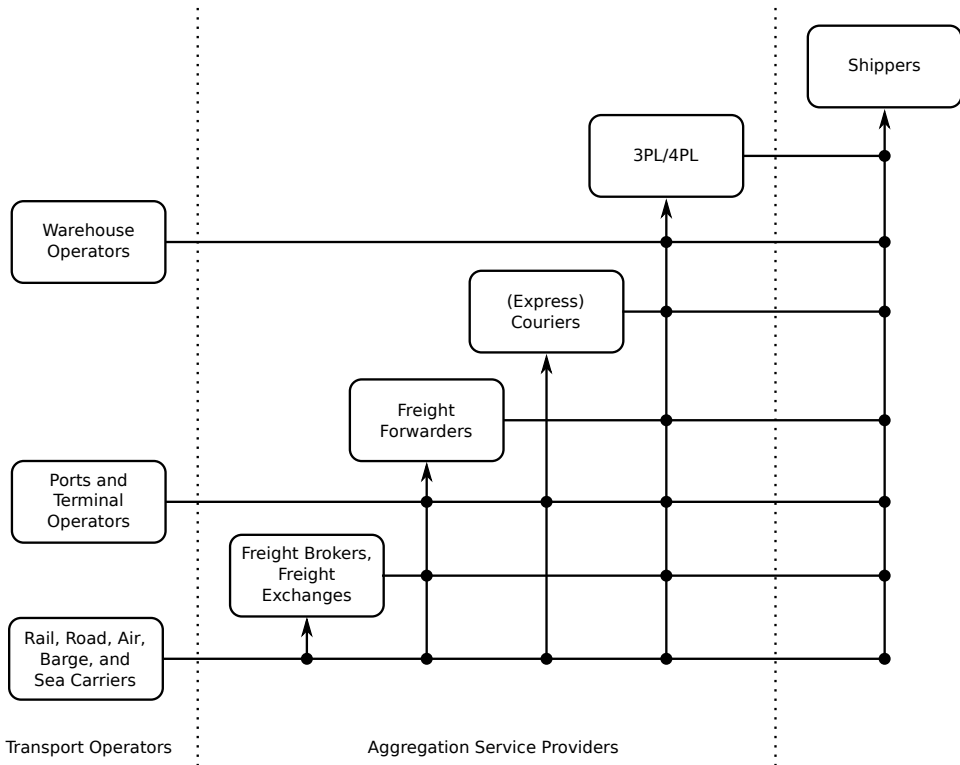


Figure 2.2: An overview of the main actors in the freight transport system

Between shippers and transport operators, there exist aggregation service providers. They aggregate transport services and organize the transport process [184, 154, 234]. In Figure 2.2, a black dot indicates that the actor to which the arrow points can make use of the service provided by the actors left of the black dot. Freight brokers find suitable carriers for a transport demand without taking any responsibility for the transport process as such. Freight exchanges are ICT systems that fill the same role and that gain considerable traction as online platforms [25]. Freight forwarders organize entire transportation chains, across different modes if needed. In contrast to brokers and exchanges, the customer of a freight forwarder does not get into direct contact with the carriers. They work with ports and terminals for inter-modal transport and they consolidate freight for higher transport efficiency. Couriers are similar to freight forwarders but they are focused on the door-to-door delivery of parcels using a standardized distribution network. Freight forwarders, on the other hand, are more focused on larger quantities, larger items and offer more customization. Third-party logistics providers (3PL) and fourth-

party logistics providers (4PL) organize the entire logistics system on behalf of their customers, of which transportation is a part. 4PLs offer in addition to 3PLs consultancy services on the supply-chain organization.

A black dot in Figure 2.2 only indicates that making use of service is common, but it is not mandatory. For example, some shippers organize all their transport demand by directly contracting carriers. A company that takes the role of a shipper could even act as a carrier transporting goods on own account. Other shippers source all their transport through freight brokers or hand only transport demand to brokers that cannot be fulfilled by directly contracted carriers.

Many companies combine several of the roles shown in Figure 2.1. It is, for instance, common that freight forwarders and couriers operate a fleet of vehicles themselves and subcontract to meet demand fluctuations. There can also be further levels of subcontracting, especially on carrier level [125]. Various other role definitions exist that can be partially overlapping with the roles listed here such as a multi-modal transport operator, a role that can be taken by some carriers, freight forwards, or third-party logistics providers.

We exclude transport by pipelines from the discussion since operational characteristics and challenges are quite different from transport by vehicles. We also do not discuss transport modes that play only a minor role in order to keep the presentation concise.

Planning and Control in Freight Transport Driven by its vital role in the economy, transport systems and freight transport are active fields of study. Planning for freight transport systems is divided into three levels: strategic level, tactical level, and operational level [68, 152]. The three levels are illustrated in Figure 2.3. On the strategic level, long term decisions are made. The construction of infrastructure such as depots and harbors fall in this category. In addition, the types of service offered by a transport operator are decided on this level. For example, the transportation of iron ore from a mine to a furnace poses different challenges than the just-in-time delivery of car components. On the tactical level decisions are made about how the transport should take place. Here, the transport mode is traditionally decided. On the operational level, the actual transport schedule is determined. Assignment of drivers falls into this category, too, which is a challenging task due to the uncertainties in transport times and strict hours of service regulations [156, 83].

Different transport modes have their own domain specific challenges. For instance in air traffic, airborne waiting times are very expensive and should be minimized while high safety standards have to be guaranteed [34]. In marine traffic, many of the challenges are related to the operations in ports [209, 38]. In railway systems, infrastructure disruptions are difficult to handle due to the infrastructure's high complexity [120]. There is an increasing interest to integrate different transport modes into multimodal or synchromodal transport systems and to handle their mutual dependencies in close to real-time [43, 141, 152].

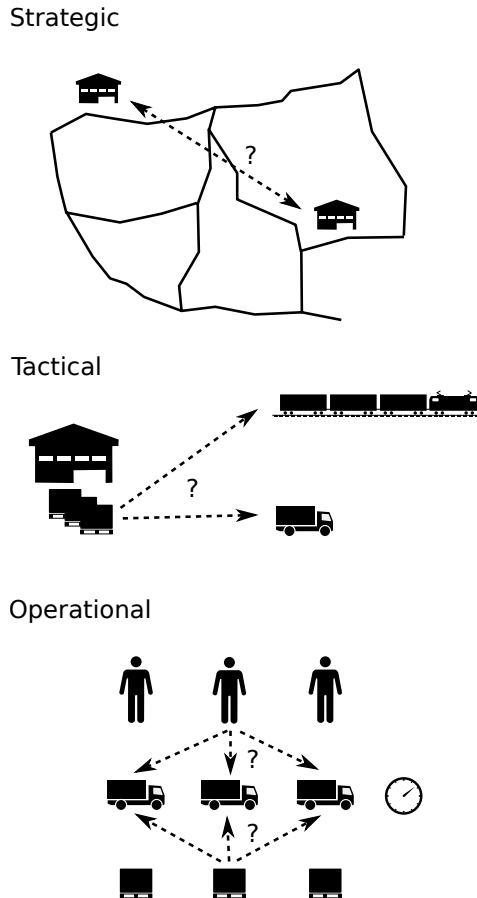


Figure 2.3: Planning for freight transport systems is typically divided into three levels.

Road transport is the dominating mode of surface freight transport [4, 169] with shares over 60 % in the OECD countries. The reason is very high accessibility through the road network and the high flexibility of trucks [25]. Road freight transport is characterized by a high degree of decentralization compared to other transport modes [125, 211]. There is almost no active control of vehicle movements by the infrastructure compared to, for instance, air or rail traffic and according to Eurostat [11], the average number of persons employed per enterprise in the road transport sector in the EU 28 countries in the year 2015 was just 5.4. Cost and life span of road vehicles are small and the training requirements to operate a heavy-duty vehicle are moderate. Compared to other modes of transport, the size of road

vehicles is small. This leads to fast adaptation of innovations on vehicle level and makes road freight transport dynamic and competitive on the one hand. On the other hand, the large number of parties involved and the relatively small resources per organization to invest in central ICT can make efforts of system wide prediction, optimization, and coordination such as decreasing empty movements (e.g., [52]) challenging [211, 106, 250]. Thus, developments in automation, electrification, and communication are challenging to introduce into the road freight system [29], but they have potentially high impact when ways are found to implement them.

2.2 Information and Communication Technology in Transportation Systems

ICT developments have major impact on transportation systems improving efficiency, safety, and reliability [25, 96, 174, 45, 127, 98, 67]. Coordinated platooning is one such ICT-enabled technology.

ICT is applied in multi-modal transport integrating the different transport modes [106]. Since transport is traditionally organized in a relatively decentralized way with many actors involved, multi-modal transport relies on efficient collaboration of various stakeholders. The focus of ICT in multi-modal transport is thus on the seamless integration and coordination of the different parties involved in the transport process.

ICT has applications in all individual modes of transportation. In the maritime transport sector, the application of ICT has a long history due to the need to communicate over long distances, due to its application for navigation purposes, and since large amounts of goods transported by ship travel internationally. GNSS and radar systems are commonplace in modern vessel operation. The automatic identification system (AIS) allows vessels to share their position data using radio communication. AIS is used for a range of applications such as collision avoidance, navigation, tracking etc [215]. Auto-pilot systems support the crew in maneuvering the ship and there are efforts in developing unmanned vessels for the use in transport [153, 51]. ICT is also heavily used in ports to automate and optimize the loading and unloading of vessels [57].

Similar to the maritime domain, ICT is an integral part of air transport systems with use of radars, radio communication, and radio navigation systems going a long way back. Fly-by-wire systems help reduce weight of the aircraft. Auto-pilots and more advanced flight management systems relieve the crew from manually controlling the aircraft in many situations and allow them to concentrate on exceptional situations [196]. Collision avoidance systems based on radar improve safety. Modern use of ICT include optimization of the traffic flow and air-port operations [34, 187]. GNSS allows for more economic and less polluting operation of the aircraft [41, 64].

In rail transport, ICT plays an increasingly important role. Modern signaling systems rely on ICT-technology. Communications-based train control uses radio communication to the trains instead of classic block-based signaling in order to

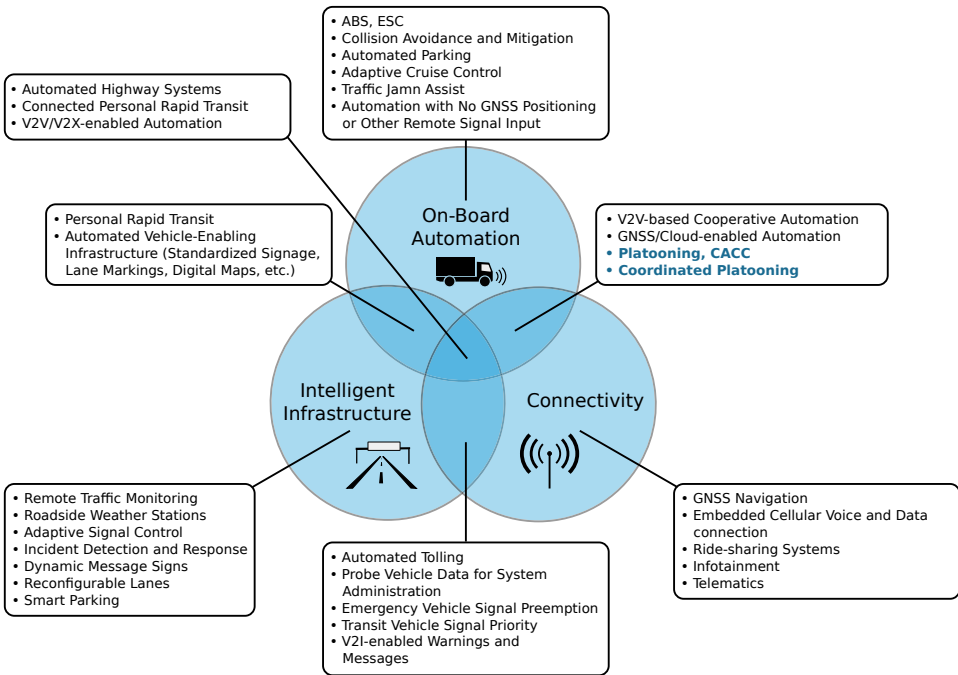


Figure 2.4: An overview of advanced transportation and vehicle technology adapted from [206]

improve efficiency and make the system more flexible [173, 164]. Technology permits to operate trains automatically with either the operator monitoring the system or being completely replaced by the control system [55]. It can allow for higher ride comfort and energy efficiency [249]. Gathering operational data in a central location opens up the opportunity to optimize the schedule and the vehicle utilization.

ICT in Road Transport Compared to other modes of transport, implementing ICT solutions into road transport is challenging. This is due to its decentralized nature, its fast dynamics, and a large number of relatively cheap vehicles. ICT applications in road transport are commonly referred to as intelligent transportation systems (ITS). Figure 2.4 shows some of the most important ITS technologies categorized by systems that rely on infrastructure, connectivity, and ICT situated on board the vehicle. The overview is adapted from [206].

The first ITS technologies focused on intelligent roadside infrastructure. Remote traffic monitoring allows to inspect the traffic state from a traffic control center. Similarly, roadside weather stations provide information on the road conditions. These information sources can be used to provide warnings to drivers,

help improve the design of the road network, and provide input to ITS components that can actively influence the traffic flow. One such component is adaptive signal control. It is used in an urban context to optimize the traffic flow based on traffic demand and in highway settings to do ramp metering, which allows to keep traffic density at the point of optimal flow rate [172, 26]. Incident detection and response algorithms process sensor data in order to quickly react to disruptions in the network. Dynamic message signs can display warnings, routing advice, and adjust speed limits dynamically. Reconfigurable lanes can make more efficient use of the infrastructure by adjusting lane usage according to traffic demand. Smart parking systems guide drivers to available parking spots thus reducing the up to 30% of traffic that arises in some cities just from vehicles searching for parking [200].

On-board automation plays an increasingly important role [39, 170]. Anti-lock braking systems (ABS) and electronic stability control (ESC) are safety features that allow the driver to keep control over the vehicle in critical situations by sensing the vehicle movement and actuating the brakes individually. Collision avoidance and mitigation uses radar sensors in order to automatically brake when a collision is eminent. Automated parking actuates the steering wheel to maneuver the car into parking spots while the driver actuates throttle and brakes. Adaptive cruise control (ACC) is an extension to regular cruise control which senses the distance to preceding vehicles, typically by radar, and keeps a safe following distance. Traffic jam assist adds lateral control to the adaptive cruise control when driving with low speed on highways.

Communication technology becomes increasingly important in today's transport systems [244, 107]. Navigation with GNSS in combination with digital maps and fast routing algorithms [195] helps the driver to find the best routes and concentrate on traffic even in unknown environments. This is often complemented by historic and real-time information about traffic and different objectives such as travel time, fuel consumption, and arrival time reliability can be considered in the routing [44, 95, 94]. Cellular voice and data connections as well as infotainment systems aim to make the driving experience more enjoyable and productive. Telematics use wireless data transmission, typically over the cellular network, to provide information about vehicle location, transport conditions, driving times, etc. to the transport operators or clients. It is one of the key technologies to automate and optimize road freight transport and commercial passenger transport [115]. Ridesharing systems use real-time vehicle location, destination, and smartphones in order to dynamically match passengers to vehicles. As a result, they make more efficient use of the available transport capacity [16].

Areas in which on-board automation technology is combined with intelligent infrastructure are personal rapid transit and automated driving with automated vehicle-enabling infrastructure. Personal rapid transport systems use small automated vehicles running on dedicated infrastructure. Due to the smaller vehicles, trips with no or few stops and without transfer to another vehicle are possible decreasing journey times compared to conventional public transport. It is possible to simplify the automation of vehicles using on-board technology by adapting the

infrastructure [232]. An example is embedding magnetic markers into the pavement [160].

On-board automation is often combined with connectivity. Cooperative automation based on vehicle-to-vehicle (V2V) communication promises to improve the traffic flow significantly compared to automation based on vehicle sensors only [107, 81]. Platooning and cooperative adaptive cruise control [162, 165], where communication is used to safely reduce the inter-vehicle distance, belongs to this category. Another interesting area of development is automated intersection crossing [81, 255, 70]. Instead of using traffic lights, vehicles communicate and adapt their speed automatically prior to crossing the intersection so they do not collide but only slow down as much as needed. This can also be supported by infrastructure in the form of an intersection controller. Vehicle automation can also be supported by GNSS for localization and cloud computing. The idea is that the sensor data are shared via a communication network and possibly processed and augmented with other data sources off-board. Like this, vehicles can expand their environmental model beyond the scope of their own sensors. Coordinated platooning, as investigated in this thesis, uses this approach to reduce fuel consumption.

Increasingly, infrastructure connects directly with vehicles, known as vehicle-to-infrastructure (V2I) communication. One such application is automated tolling: roadside communication units connect to communication equipment on-board the vehicle to register it and bill the driver. Instead of expensive roadside infrastructure, increasingly probe vehicle data are used to collect information about the road network and the traffic situation [108]. Emergency vehicle signal preemption and transit vehicle signal priority are both technologies where vehicles communicate with traffic lights to get priority. Priority for emergency vehicles is granted with minimal delay, whereas priority for transit vehicles typically affects only the signal cycle length. V2I communication enabled warnings and messages are displayed in the vehicle so that expensive (variable) message signs are not needed.

On board automation, connectivity, and intelligent infrastructure is combined in automated highway systems (AHS) [233, 160]. AHS are envisioned as systems with dedicated infrastructure where regular passenger and freight vehicles are fully automated once they enter the system. A hierarchical control system enables higher throughput, less accidents, and less pollution than conventional highway systems. Connected personal rapid transit systems are similar to AHS with the difference that vehicles are dedicated for the system and cannot be used on regular roads. Finally, variants of AHS allow both automated and human operated vehicles trading some of the efficiency for a gradual introduction of the technology [21].

Applications of ICT specifically for road freight transport have been focusing on telematics and fleet management [238]. Telematic systems allow to locate the vehicle remotely using GNSS localization and mobile data connection. Digital tachographs facilitate monitoring vehicles' operation conditions and driver behavior. Digital tachographs can be integrated with a telematics system making the data centrally available in real-time. These data can be used in fleet management systems and help to optimize vehicle scheduling and routing, as well as maintenance [28].

Integration with systems on logistics level such as transport management systems, customer relationship management systems etc. leads to higher overall transport efficiency. Radio frequency identification (RFID) and other communication technologies enable so-called smart transportation management systems. These systems allow to track vehicles and goods along the transport chain in real-time enhancing efficiency and reliability of the transport process [210, 197]. As professional driver wages constitute a considerable share of the cost operating a road freight vehicle, there is a clear business case for implementing high levels of automation in freight vehicles [124]. Automation of freight vehicles not only promises higher productivity but could also lead to far reaching operational changes [90]. When there is no driver, constraints such as hour of service regulations are no longer present.

Travel Time Prediction An important ICT-enabled prerequisite for en route platoon formation is accurate travel time prediction. In order for two vehicles to arrive at the same position at the same time to form a platoon, the feasible travel times up to that point need to be estimated with good precision. Hereby, the primary source of variations in the travel time is due to traffic. Increasing availability of travel time data from connected probe vehicles open up for levels of accuracy required by platooning application.

Travel times can be predicted based on historic data and potentially also incorporating real-time information [133, 180, 61, 248, 202, 132]. Some methods estimate the state of a traffic model, which has been calibrated with historic data, and use it to predict travel times [237, 56, 212, 60]. A class of these models describe traffic flow as a system of partial differential equations that is then discretized in space and time for implementation on digital computers. By estimating the current state of the system from traffic measurements and forecasts on the travel demand, such a model can be used to predict traffic flow and thus travel times.

An accurate prediction of the travel times in a road network is a challenging task due to the large scale of the system and a large number of factors that influence travel times [128, 221]. Being able to model the distribution of the prediction error provides additional information useful for planning.

Various stochastic distributions have been explored to model travel time distributions. The input data for such models are typically direct or indirect travel time measurements but they can also be estimated from randomized simulations [62]. One motivation for modeling travel times distributions explicitly, is to use parameters of the distribution as a performance measurement. [59, 89, 128, 231, 239]. Another important application of travel time distribution models is planning and operations [109, 119, 243]. Such models have been successfully employed to find reliable routes rather than fastest routes [85, 194, 207]. In [23], the authors are concerned with modeling travel time distributions for stochastic routing. We use in the thesis such stochastic models to compute the probability that two vehicles will successfully meet and form a platoon. Is the probability of a successful merge too small, it can be preferable to not attempt the merge in the first place or attempt to form a platoon with another vehicle.



Figure 2.5: A three-vehicle platoon in the COMPANION project




Figure 2.6: Platooning leads to a slipstream effect that can reduce the fuel consumption of the follower vehicle.

2.3 Vehicle Platooning

Platooning means vehicles driving in groups with the gap between adjacent vehicles being controlled as shown in Figure 2.5. In its simplest form, this occurs naturally on busy roads [188]. Modern sensor and wireless communication technology enable automatic control of the inter-vehicle gaps. Such automatic control has a number of advantages over manual control by human drivers. Adaptive cruise control (ACC) [246], which has been in series production for some time [84], is both a convenience and a safety feature. It relieves the driver from the task of controlling the distance to the vehicle in front. Rear collision accidents due to insufficient gaps and inattentive drivers account for a significant number of accidents [19]. ACC can help to avoid such accidents. Automatic control of the inter-vehicle gaps, in particular cooperative adaptive cruise control (CACC) where vehicles communicate actively, opens

Table 2.1: Platooning technologies categorized by levels of automation according SAE International standard J3016 (Adopted from SAE International, [190])

SAE level	Name	Lon./Lat. Control	OEDR	DDT Fallback	ODD	Platooning Techn.
Driver performs part or all of the DDT						
0	No Driving Automation	Driver	Driver	Driver	n/a	n/a
1	Driver Assistance	Driver and system	Driver	Driver	Limited	CACC, DATP
2	Partial Driving Automation	System	Driver	Driver	Limited	DATP
Automated Driving System (“System”) performs the entire DDT (while engaged)						
3	Conditional Driving Automation	System	System	Fallback-ready user	Limited	DATP
4	High Driving Automation	System	System	System	Limited	FAP
5	Full Driving Automation	System	System	System	Unlimited	FAP

DDT: Dynamic Driving Task, OEDR: Object and Event Detection and Response, ODD: Operational Design Domain

up the possibility to reduce the inter-vehicle gaps compared to human-controlled gaps without compromising safety [114, 219, 162, 24]. This results in two other desirable effects. By reducing the inter-vehicle gaps, more vehicles can fit on the road [20, 101, 147, 199, 223, 72], leading to more efficient use of the infrastructure. Small inter-vehicle gaps lead to a slipstream effect (Figure 2.6), which reduces the air drag experienced by the trailing vehicles. Reduced air drag, in turn, leads to reduced fuel consumption. The slipstream effect is frequently exploited in bicycle races.

Different definitions of platooning exist in the literature. Some authors require automation of the lateral control in addition to the longitudinal control [123, 74, 32]. Another way to differentiate between CACC and platooning is that the vehicles in the platoon form a group and they are aware of each other, as opposed to CACC systems where only adjacent vehicles interact [236]. In this thesis, we use the latter definition of platooning. Another term used is driver assisted truck platooning (DATP) [40, 36]. DATP refers to platooning systems in which the driver constantly monitors and potentially steers the vehicle. Table 2.1 categorizes platooning technology according to the levels of automation defined by the Society

of Automobile Engineers (SAE International), [190]. CACC is classified as level 1 as it only requires the longitudinal control of the vehicle to be automated. DATP may range from SAE levels 1 to 3. DATP systems that only automate the longitudinal control are classified as level 1, while system that automate both steering and longitudinal control are on level 2 or 3, depending on the degree of which exceptional situations such as a passenger vehicle cutting in the platoon can safely be handled before the control is handed back to the human driver. We refer to platooning systems on level 4 or 5 as fully automated platooning (FAP). On these levels, the platooning function is a part of the capabilities of a self-driving vehicle. In the context of long haulage road freight transport, level 4 might become important. A scenario is possible in which vehicles can operate without the need for a driver on the highway network allowing the driver to rest and only drive manually on the final stretch to the destination.

Fuel Saving Potential of Platooning The main driving factor behind the early implementations of DATP is its potential to reduce fuel consumption [124]. There is strong evidence that platooning is an effective way to decrease fuel consumption. In a two-truck platoon, typical air-drag reduction lies in the range of 40 % for the trailing vehicle and around 10 % for the leading vehicle. This leads to nominal fuel savings of 10 % for the trailing vehicle and 4 % for the leading vehicle compared to the vehicle driving by itself [181]. These figures on fuel consumption reduction assume loaded heavy-duty vehicles traveling at regular motorway speeds. For longer platoons, most results suggest that the leading vehicle saves the same amount of fuel as in a two-truck platoon. The trailing vehicles save at least as much as the trailing vehicle in a two-truck platoon. To measure the exact fuel consumption reduction of platooning under realistic conditions is a difficult task since there are many factors that influence fuel consumption. This pertains both to the platooning technology used and the vehicle types as well as to external factors such as road topography, traffic, weather etc. Apart from the air-drag reduction, the platoon control itself has an effect on the fuel consumption. When road topography is, for instance, not taken into account [222, 53], platooning can even increase fuel consumption [18, 17] on hilly roads. On the other hand, there is a potential for improving fuel consumption further than the reduction caused by lowered air-drag. This is due to the ability of some platoon controllers to reject disturbances from the preceding vehicle [135, 252, 142, 208] and thus waste less energy in braking in order to maintain a safe inter-vehicle distance.

The potential of platooning to reduce air-drag and fuel-consumption has been shown in computational fluid dynamics (CFD) studies, wind tunnel tests, and field experiments. CFD studies [166, 80, 203, 112, 217, 242, 97, 69, 235, 14] allow to test a large number of configurations at relatively low incremental cost and to analyze the aerodynamic phenomena in more detail. It is however challenging to ensure accuracy of these simulations as the underlying dynamic equations are inherently difficult to solve numerically. They also might neglect important effects such as

cross-winds and disturbances from surrounding traffic.

Wind-tunnel experiments [177, 99, 171, 163, 251, 241, 48] are quite similar to CFD studies in their characteristics, except for that they do not rely on modeling assumptions for the airflow dynamics. They are often conducted with simplified geometries at scale, which is a potential source of error. According to [203], no wind-tunnel facilities exist that would allow testing a heavy-duty vehicle platoon in full scale. In terms of equipment needed, they are more expensive than CFD studies. Both CFD and wind-tunnel studies only make statements about the air-drag reduction, which is connected to but not the same as fuel consumption reduction.

Studies evaluating data from platooning vehicles [203, 112, 18, 217, 149, 186, 175, 104, 134, 69, 160, 42, 49, 157, 17, 1, 139] have the highest potential to indicate realistic fuel savings of platooning when taking most factors into account that can influence the fuel consumption in a platoon. However, they are very expensive to conduct and it is non-trivial to establish a consistent baseline due to the large number of factors that influence fuel consumption. To mitigate this problem, the studies are often conducted under idealized conditions on test tracks. This approach bears the risk of missing important effects on the fuel consumption in a platoon as they would occur in driving on public roads such as road topography, traffic disturbances, etc. Opposed to CFD studies and wind-tunnel experiments, the air-drag reduction obtained in experiments with real vehicles can only indirectly be measured through the fuel consumption or the motor torque. Therefore, it is not trivial to single out the different effects that have an influence on the fuel consumption in a platoon and to compare with CFD studies and wind-tunnel experiments.

The work in this thesis explicitly considers that fuel consumption depends on vehicle speed and that platooning affects the fuel consumption. The development of an accurate fuel consumption model is non-trivial, as fuel consumption depends on a large number of factors such as road, weather, vehicle, driver, speed, load, traffic, etc. [71, 44, 139]. However, to assess and optimize the effect of platooning, an absolute prediction of fuel consumption is not needed. The reduction of fuel consumption obtained is the relevant quantity.

Enabling Technologies for Platooning and Platoon Control Platooning is enabled by the combination of modern sensor, communication and embedded computing technology. In [218], an overview of the technology used in three major platooning projects is provided. Figure 2.7 shows some of the most important elements used in state-of-the-art platooning systems. Radar systems are used to measure inter-vehicle distance and also to monitor the environment, for instance, to detect intruding vehicles. Lidar sensors can also be used for that purpose but they are currently less common due to their higher price. Cameras are used to detect lane markings for lateral control. They can also complement the measurements from the radar. A number of low-cost sensors on vehicle level are used to estimate vehicle speed, position, brake capacity etc. A human machine interface provides the interaction with the driver [189, 92, 91, 93]. Absolute vehicle position-




Figure 2.7: An overview of the technologies that enable platooning

ing is primarily provided by global navigation satellite systems (GNSS). Absolute positioning is needed to retrieve information from digital maps, for instance, for geo-fencing where platooning is permitted and supports also the estimation of the relative vehicle position, in particular, during merging maneuvers. Also magnetic markers embedded in the pavement have been put to use for absolute positing [160]. They require costly modifications of the infrastructure, though. V2V communication is central to platooning technology. By exchanging state information, higher levels of control performance can be achieved compared to vehicle following based on sensors alone. It is also applied to negotiate platoon formation and splitting. For medium access, IEEE 802.11p is one of the commonly used standards. The higher protocol layers are in the process of standardization. For further details refer to [236]. Also infrared light communication can serve for V2V communication. V2I communication is primarily important for platoon coordination. It can also be rel-

evant for receiving real-time information such as weather data and updating digital maps. It is typically realized through existing cellphone data networks but the use of IEEE 802.11p might become important in the future when such infrastructure is deployed. Platooning relies on the automatic control of the drive line and the brake systems. Such functionality is fairly standard for heavy-duty vehicles due to its use in advanced driver assistance systems. Less common are steering actuators, which are needed in case lateral control is automated. Finally, in order to process all the data coming from sensors and communication, embedded control units (ECU) are needed.

The automatic control of inter-vehicle gaps is a challenging problem that has attracted significant research interest [58, 20, 159, 102, 105, 114, 130, 140, 165, 179, 213, 220, 233, 219]. Apart from the stringent requirements of safety, i.e., that vehicles do not collide under any circumstances, it is required that disturbances in one part of the platoon do not get amplified as they travel from vehicle to vehicle through the platoon. The phenomenon of spontaneously occurring traffic jams in heavy traffic is an example of a small disturbance being amplified as it travels upstream through the string of vehicles. To formalize this phenomenon for the sake of control design, the notion of string stability is used [214, 176]. Roughly speaking, a system is string stable when a disturbance on one subsystem is attenuated as it propagates along the string of subsystems. When wireless communication is used to transmit control information between the vehicles, interference can cause information to be lost. The controller needs to handle such loss of information, for instance, by relying on data from other sensors or increasing the gap between the vehicles. Surrounding traffic needs to be taken into account. For instance, other vehicles still have to be able to enter and exit the highway. Long platoons need to detect such vehicles and open gaps for the other vehicles when needed. When platooning is used as a measure to reduce fuel consumption, it is important that the control of the inter-vehicle gaps is performed in a way that the reduced air drag actually translates into reduced fuel consumption. If the vehicles brake and accelerate a lot in order to keep the gap at the desired value, they might consume more fuel compared to not platooning [18]. In particular in hilly terrain, a sophisticated fuel-efficient control strategy is a prerequisite to realizing the potential fuel consumption benefits from platooning [222].

2.4 Platoon Coordination

Platooning requires in most scenarios some level of coordination to form, split and potentially to reorder platoons. We define *platoon coordination* for the context of this thesis as an approach to organize the formation and splitting of platoons. That a platoon can easily be formed and split, even while driving, is the main advantage of platooning compared to road trains with physical coupling. So only with a coordination approach, platooning cannot realize its full potential.

Platoon coordination can be categorized into local, regional, and global level




Figure 2.8: Platoon coordination is categorized into three levels: local platoon coordination, regional platoon coordination, and global platoon coordination.

[168]. These levels are illustrated in Figure 2.8. Local platoon coordination takes effect when vehicles are close enough to establish V2V communication, and use it to coordinate the platoon formation and splitting with a distributed protocol. When this approach is used without any higher level of coordination, we refer to it as spontaneous platooning. Spontaneous platooning relies on having a relatively high share of platooning-enabled vehicles on the road in order to benefit from platooning. The next level is regional coordination. On this level, the coordination of

platooning is organized over a region that is smaller than the entire road network but too large to rely on V2V communication only. An example is coordination on a particular stretch of highway. The third level is global coordination. On this level, vehicles are coordinated across the entire road network using a suitable communication network. This is the most complex coordination but also the one with most potential. Global coordination is especially interesting, when there are only few platooning-enabled vehicles or only vehicles from the same fleet can platoon. Note that these approaches build to some extent on top of each other. For example, two vehicles that coordinated to meet at some point still need to use V2V communication to locally form the platoon.

An overview of existing work on platoon coordination with a focus on regional and global platoon coordination is provided by Table 2.2. In addition to the categories local, regional, and global coordination, Table 2.2 characterizes work by the following features.

- **Fuel savings potential** These publications focus on the system-wide potential to reduce fuel consumption under coordination rather than a specific coordination approach to realize the potential.
- **AHS** Work that focuses on platoon coordination specifically in the context of AHS is listed here.
- **HDV platooning** Work in this category is specific to platoon coordination for vehicles transporting goods.
- **Optimization-based** These contributions use mathematical optimization for coordination.
- **Effect of traffic** This category lists contributions that consider the interaction of platoon coordination with traffic.

In the following, we discuss the references listed in Table 2.2 in detail.

Platoon Coordination in AHS and Local Coordination Platoon coordination has first been studied in the context of AHS. In AHS, platoons are primarily considered as an organizational unit and a way to increase traffic throughput. Fuel economy is a secondary effect and all vehicles on the highway or at least on a special lane are platooning enabled. This gives rather different requirements for the coordination algorithms needed compared to the ones needed in the context of this thesis, where only a small subset of vehicles can platoon. In [111], the authors consider a hierarchical model for an AHS. On the upper level the network layer controls traffic flows and below that on link level reference speeds are set and platoon maneuvers are suggested. The network layer and the link layers are implemented as an off-board system. The on-board system of each vehicle consists of a coordination layer responsible for executing maneuvers such as merging with a platoon and a regulator that controls the continuous dynamics of the vehicle. The

Table 2.2: Characterization of work on platoon coordination

Feature	References
Local Coordination	[21, 81, 111, 149, 161, 198, 240]
Regional Coordination	[15, 27, 47, 86, 100, 101, 129, 136, 143, 144, 193, 191, 192, 253]
Global Coordination	[87, 131, 137, 138, 150, 158, 167, 205, 254]
Fuel Savings Potential	[36, 79, 118, 145]
AHS	[21, 27, 81, 101, 111, 161, 240]
HDV Platooning	[36, 87, 100, 118, 131, 143, 145, 149, 158, 192, 253, 254]
Optimization-based	[15, 27, 47, 86, 129, 136, 137, 138, 150, 167, 191, 205, 240, 253, 254]
Effect of Traffic	[27, 86, 143, 192, 198, 205, 253]

focus of [27] is on the route choice in an AHS leading to an intractable optimization problem. By approximating the problem with a mixed integer linear program, it can be solved sufficiently fast to be used as a model predictive feedback controller. In [101] a setting is considered, where platoons are formed at highway entrances by waiting in dedicated zones. Different approaches are developed to sort vehicles into platoons in order to maximize throughput of the highway. It proves beneficial for the throughput to create platoons that stay intact for a long distance. The Grand Cooperative Driving Challenge 2016 [81] was a competition in which teams participated by designing on-board systems to drive cooperatively supported by V2V communication. One of the scenarios was that two platoons merged before a lane closure. The protocol for this local platoon coordination was specified for all participants.

Further examples for local platoon coordination are [21, 149, 198, 240]. In [149], field tests are reported with a heavy-duty vehicle platooning system. In addition to steady state cruising, acceleration, deceleration, join and split maneuvers are evaluated. A protocol for a vehicle joining a platoon is proposed in [198]. The protocol is evaluated through simulations and the effect of communication errors is investigated. The authors of [21] introduce an extensive framework for local platoon coordination in a highway context. The protocol covers all relevant situations and scenarios for vehicle platooning on highways. It is evaluated using an integrated micro traffic and V2V communication simulator. In [240], two trajectory generation methods for joining a platoon and splitting vehicles from a platoon are introduced and evaluated with respect to their emissions.

Coordinated Heavy-Duty Vehicle Platooning The setting considered in this work is that platooning is primarily a measure to improve fuel consumption. We consider that only a subset of the vehicles on the road can platoon with each other either due to technical reasons or because only vehicles from the same organization are allowed to platoon. We refer to this as coordinated heavy-duty vehicle platooning. In this setting, the focus is on finding and creating platoons of vehicles that would not spontaneously meet during their journeys. In AHS, there is primarily a need to form platoons in order to increase capacity, which is only needed when there are many vehicles on the road. But then, finding platoon partners locally is not a problem. Coordinated heavy-duty vehicle platooning thus requires different coordination approaches than the ones developed for AHS. A literature review of recent contributions in the area of coordinated heavy-duty vehicle platooning is provided by [37]. It also discusses the relation of coordinated heavy-duty vehicle platooning with other cooperative transport solutions such as freight consolidation and ridesharing.

In [79, 145, 36, 118], the potential of platooning is evaluated without much focus on the actual coordination scheme that realizes the potential. An artificial road network is generated in [79] and vehicles traveling on that network are simulated. The generated dataset is then analyzed in order to answer the question how much platooning is possible when vehicles wait up to a certain time at nodes of the network for platoon partners. The authors of [145] analyze location data from a fleet management system. They investigate how many vehicles spontaneously platoon and to what extent the platooning rate can be improved when adjustments to the schedule are made. In [36], data from fleet management systems is analyzed as well. The objective is to find out what the routes are on which vehicles could have platooned allowing for adaptation of the vehicle speeds. The speed adaptation is determined using an unspecified optimization algorithm and different types of adaptations are tested. Also the effect of different braking capabilities leading to differences in the inter-vehicle gaps is investigated. In [118], platooning potential for three road carriers is analyzed based on their operation patterns. The effect of fuel reduction and two different scenarios of labor cost reduction is considered.

Work that considers the formation of platoons when it is already decided which vehicles should form a platoon and which consider fixed routes are [144, 193, 191, 253, 143]. In [144], the focus is on computing when two vehicles sharing the same route should initiate a catch-up maneuver from a fuel economy perspective. The increased speed of one of the vehicles until the platoon is formed leads to an increased fuel consumption while the subsequent platooning phase lowers the fuel consumption. The study explores different parameters influencing this trade-off. In [191], strategies for merging vehicles en route into a platoon are analyzed. Vehicles are either allowed to increase speed, decrease speed or a combination of the two. The authors of [193] apply a consensus algorithm to join several vehicles into a platoon on a highway in a distributed way. The consensus approach is compared to the approach developed in [191]. The comparison is performed based on scenarios generated from real driving data on a public highway. The authors of

[129] consider fuel-optimal speed control of vehicles merging at intersections. The approach takes the inertia of vehicles into account and uses optimal-control theory and model-predictive control to solve the control problem.

The effect of traffic on platooning is explored in [253]. Travel times are modeled as stochastic quantities with a known distribution. In order to form platoons, vehicles can delay their start time. A penalty oriented on monetary costs for arriving early or late at the destination is derived. The optimal starting times in expectation considering both the arrival penalty and the fuel savings from platoon are given. Another approach to taking traffic on into account is presented in [143]. There, a macroscopic traffic model is considered. The vehicle that is ahead acts as a moving bottleneck, which can slow down the merging process. A similar approach is taken by [192] also considering the effect of changing the speed of a heavy-duty vehicle on traffic. It proposes a decision support system for the real-time operation of heavy-duty vehicle platooning. The focus of [192] is on optimizing the overall traffic flow rather than improving fuel consumption.

In [131] and [158], data mining techniques are used to identify vehicles with similar properties and similar routes. A ranking is created of matching vehicles and the top choices are displayed to the driver. The driver can then select a suitable platoon partner and form the platoon. Fuel consumption reduction is only implicitly considered in this approach.

Approaches using centralized optimization to coordinate platooning are considered by [138, 137, 205, 150, 167, 254]. In [138], the problem of fuel efficient platoon coordination is formulated as a mixed integer linear program. Combined routing and platooning is considered and it is assumed that vehicles wait at nodes of the network without additional cost. The problem is shown to be NP-hard and several heuristics are developed to obtain solutions. Information for departure times and destinations is considered to be known beforehand. In [137], further insight on the solution structure is modeled as constraints which allows to compute optimal solutions for more vehicles. The solution is tested both on an artificial network and a real highway network. In [205], the same optimization model as in [137] is being dealt with. While the focus of [137] lies on algorithmic performance, [205] concentrates on a case study on a grid road network with a traffic simulator. The effect of the maximum wait time is systematically investigated. Variable speed as part of the optimization is added in [150]. The method is analyzed in a case study. Furthermore, a heuristic clustering for large scale instances is considered, which groups vehicles with similar origin destination pairs and timing. The authors of [167] approach the problem posed in [138] using a genetic algorithm and test it on a simple graph reflecting an actual motorway network. The genetic algorithm is compared to a mixed integer linear programming solver and outperforms it in most cases. In [254], the formulation as a mixed integer linear program from [138] is extended with soft time-windows for the arrival. The formulation is evaluated in simulations and the sensitivity to disturbances in the time windows is estimated.

In [47], algorithms for grouping vehicles with the same route into platoons are developed and their complexity is analyzed. Different types of cost functions as

well as constraints on the platoon size and departure time windows are considered, leading to different complexity levels of the planning problem.

The authors of [136] present a distributed control framework. Vehicles approaching highway intersections adjust their speed as to form a platoon on the intersection when it is beneficial from a fuel economy perspective. The controller is simulated on a model of a motorway network and significant fuel savings are demonstrated for a moderate number of participating vehicles.

In [15], ride sharing and platooning is combined. Passengers board vehicles at dedicated stations. Before a vehicle leaves the station, it waits for some time in order to form a platoon with other vehicles leaving the station in order to save fuel. The trade-off between waiting time and energy consumption is explored under different policies.

In [100], platooning in the context of city logistics is explored. The study focuses on simulating platooning in a micro-simulation framework. In an urban scenario, the impact on delay of different policies to form platoons is tested.

The authors of [86] study departure time scheduling in a game theoretic setting. Drivers balance the increased probability of finding platoon partners when there are more platooning-enabled trucks on the road with increased travel time due to congestion, congestion tax, and deviation from the preferred departure time.

In [87], a cryptographic method is presented, which allows one fleet owner to query another fleet owner whether or not a vehicle of the the first fleet can platoon with any vehicle of the other fleet. This happens without the first fleet owner revealing routes and departure time to the second fleet owner.

There is significant potential for DATP in realistic settings if platooning-enabled vehicles are efficiently coordinated. Some aspects of such a platoon coordination have been investigated but there is still a research gap how an integrated global platoon coordination system should be designed. The following chapters are a contribution to closing that gap and making globally coordinated DATP systems a reality.

2.5 Summary

This chapter provides the background of the thesis. We motivate, that freight transport is the backbone of the industry and thus tightly coupled to economic prosperity. Freight transport systems are complex systems in which different modes of transport, ports, terminals and warehouse come together with the purpose of efficiently moving goods. In order to handle this complexity, various aggregation roles exist. Intricate planning problems arise on strategic, tactical, and operational level. We argue, that road transport is a particularly dynamic mode of transport with a large number of small companies, which can make advanced planning challenging.

Information and communication technology (ICT) has major impact on all modes of transport allowing for more reliable and efficient operation. We show, that there is a wide range of applications of ICT in road transport which can be

classified into technologies focusing on intelligent infrastructure, on-board automation, and connectivity. Platooning is a technology in which on-board automation and connectivity come together to decrease fuel consumption, to increase road capacity, and to partially automate the vehicle. It relies on a number of advanced sensor, communication, and vehicle control technologies, which are shared with other advanced driver assistance systems.

Platooning systems have first been developed with the focus of passenger vehicle automation. More recently, there has been a lot of interest in platooning for heavy-duty vehicles with the focus of reducing fuel consumption called driver assisted truck platooning. We list extensive evidence that heavy-duty vehicle platooning can indeed lead to significant reduction in fuel consumption due to the slipstream effect reducing aerodynamic drag forces.

We motivate that heavy-duty vehicle platooning being a cooperative technology needs some level of coordination. Platoon coordination is organized on three levels: local, regional, and global; where each level builds on top of the other. As discussed in this chapter, different aspects of this coordination have been investigated. However, there it is still an open question how to design a coordinated heavy-duty vehicle platooning system, which is addressed in this thesis.

Chapter 3

Coordinated Platooning System

IN this chapter, we introduce the architecture of the coordinated platooning system. As introduced in Section 1.3, the coordinated platooning system consists of platooning-enabled vehicles in a road network, fleet management systems, and a platoon coordination system. The platoon coordination system facilitates fuel-efficient platooning across multiple fleet management systems. In Section 3.1, we introduce the overall system architecture. As central part of the coordinated platooning system, the platoon coordinator interfaces with fleet management systems and data providers in order to solve the problem of guiding vehicles to their destinations in a fuel-efficient manner using platooning to reduce fuel consumption. In Section 3.2, the internal architecture of the platoon coordinator is introduced. In several steps, it computes platoon plans which are communicated through the fleet management systems to the on-board systems. The on-board system derives reference speeds for the vehicle controller and reference platoon configurations for the platoon manager from these plans.

3.1 System Architecture

Figure 3.1 shows a layered transport architecture. The road transport system is structured in four layers with the highest level of abstraction at the top and increasing detail towards the bottom. At the highest layer, the *service layer*, the existing planning systems for road transport reside in form of transport management systems (TMS) and fleet management systems (FMS). On this layer, goods flows are matched to vehicles and drivers. Typical time scales on this layer are in the order of hours to days. A wide range of complexity levels can be encountered on this layer ranging from manual planning to complex supply chain optimization.

On the layer below, the *strategic layer*, the platoon coordinator is situated. The platoon coordinator is the central system provided by a platoon service provider. Platoon service providers have been postulated in the literature [118] as organizations that provide crucial platooning services shared between road transport

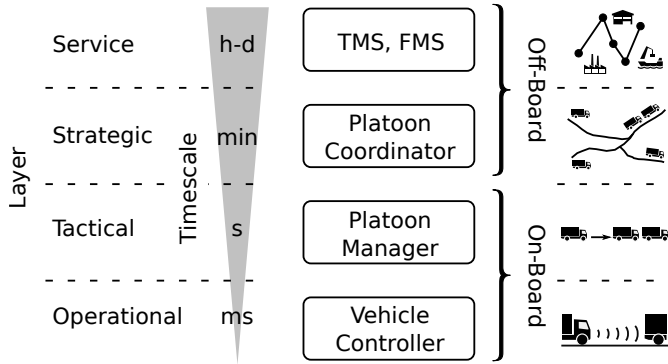


Figure 3.1: A layered control system architecture for coordinated platooning

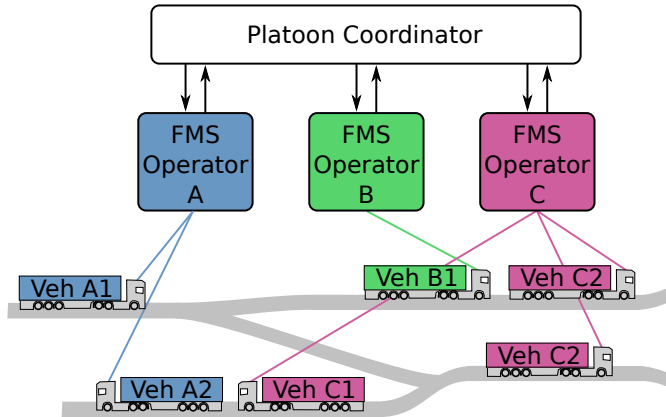


Figure 3.2: The coordinator receives assignment data from the fleet management systems of different transport operators and computes routes and speed profiles for the vehicles to form platoons en route.

providers, such as certification, insurance, and coordination. The coordination service, also known as match making, is what this thesis is concerned with. This service might be provided as a public service or through private enterprises.

At the *tactical layer*, the platoon manager resides. The platoon manager controls the formation, splitting, reorganization, and operation of the platoons according to platoon plans provided by the strategic layer. It does so by coordinating with the other vehicles using vehicle-to-vehicle (V2V) communication and a standardized protocol. The platoon manager sets reference speeds and reference inter-vehicle distances, which are then tracked by the vehicle controller on the operational layer.




Figure 3.3: Flow chart detailing the data flow between the components of the coordinated platooning system

The tactical information is also communicated to the driver in each vehicle. The driver takes a supervisory role and can overwrite the decisions of the platoon management.

At the *operational layer*, the vehicle controller tracks the speed and distance references from the tactical layer. It commands the engine actuators, the brake systems, and the gearbox. It provides an abstraction layer on top of the complex vehicle specific dynamics and presents a standardized interface to the tactical layer.

Remark 5. The service layer introduced here is sometimes divided into a strategic, tactical, and operational layer (see [68] and Section 2.1), which should not be confused with the layer structure introduced in this paper.

Remark 6. The layered architecture is inspired by the one presented in [17] consisting of a transport layer, a platoon layer, and vehicle layer. The service and strategic layer combined correspond to the transport layer, the tactical layer to the platoon layer, and the operational layer to the vehicle layer in [17].

The platoon coordinator interfaces with the FMSs of transport operators as illustrated in Figure 3.2. Each FMS controls the vehicles that belong to its transport operator. The platoon coordinator receives assignments from the FMSs as shown in Figure 3.3. An assignment $\mathcal{A} = (\mathcal{P}^S, \mathcal{P}^D, t^S, t^D, \mathcal{D})$ consists of a start position \mathcal{P}^S at which the vehicle starts its trip at start time t^S and a destination \mathcal{P}^D at which the vehicle is supposed to end the trip before deadline t^D . Additional data \mathcal{D} can be associated with the assignment such as type of the vehicle, constraints on the route etc. The index set of assignments that are considered by the platoon coordinator at any given point in time is denoted \mathcal{N}_c .

Remark 7. Complex transport missions with multiple stops are broken down into multiple assignments by the FMS. Start location and destination do not have to match the position at which the vehicle starts driving but can be, for instance, the points where the vehicle enters and leaves the highway network.

Remark 8. When t^S lies in the future, the respective FMS ensures that the vehicle position will be \mathcal{P}^S at time t^S , potentially using a preceding assignment for that purpose or directly sending the vehicle to that location without using the platoon coordinator. Thus, no interdependence of assignments for the same vehicle is taken into account by the platoon coordinator. If t^S is the current time or lies in the past, \mathcal{P}^S corresponds to a measurement of the position at that time.

For each assignment in \mathcal{N}_c , the platoon coordinator computes a platoon plan using information provided by the FMSs and data providers as illustrated in Figure 3.4. A platoon plan $\mathfrak{P} = (\mathbf{e}, \mathbf{t}, \mathbf{p})$ consists of a route \mathbf{e} , a time profile \mathbf{t} , and a platoon configuration profile \mathbf{p} . The route \mathbf{e} connects the start location \mathcal{P}^S of the corresponding assignment with the destination \mathcal{P}^D in the road network. The time profile \mathbf{t} is defined along the route and encodes when the vehicle should be at which location. The platoon configuration profile \mathbf{p} encodes in which platoon configurations the vehicle travels along the route.

Platoon plans are updated on a timescale of minutes in order to adapt to new assignments, new planning information, and deviation from the platoon plans. To this end, the platoon coordinator receives periodically updated vehicle positions (e, x) and platoon states π for the vehicles corresponding to the active assignments.

Remark 9. The platoon coordinator cannot interface directly with the vehicles' on-board systems as the layered architecture in Figure 3.1 might suggest. This is because the mapping between vehicles and transport assignments is only known by the respective FMS. Furthermore, the transport operator might want to check the generated platoon plans for its vehicles and potentially update or cancel the assignment when errors are detected. Note that each FMS only receives the platoon plans corresponding to the vehicles it manages.

Remark 10. The repeated solution of an optimization problem over a receding finite horizon has gained significant traction in the field of automatic control under the term model predictive control [155, 30].

The FMSs forward the platoon plans to the on-board systems. The objective of the on-board system is to track the plans locally on vehicle level. In this way, some disturbances can be compensated for at a faster time-scale without having to rely on the wireless communication link to the off-board system. An architecture overview of an on-board system is shown in Figure 3.5. The plan tracker translates the platoon plans received via vehicle-to-infrastructure (V2I) communication in combination with the locally measured vehicle position $e(t), x(t)$ into routing advice e^+ , reference speeds v^{ref} , and reference platoon configurations π^{ref} for the platoon manager. The platoon manager, in turn, sets, depending on the current platoon state, reference speeds or gaps for the vehicle controller and synchronizes with other vehicles using V2V communication. The vehicle controller accesses directly the engine and brake systems. Steering is either handled by the vehicle controller or by the driver. The steering functionality is not shown in Figure 3.5 for the sake of simplicity. For further details of the on-board system refer, for instance, to [245].




Figure 3.4: Each assignment consists of a start position and a destination in the network as well as a start time and arrival deadline. Platoon plans consist of a route, a speed sequence, and a time sequence giving a reference trajectory in time and space as shown in the image. When the position trajectories of two vehicles partially coincide, these vehicles can form a platoon and save fuel.

3.2 Platoon Coordinator Architecture

In this thesis, we focus in particular on the design of the platoon coordinator. The platoon coordinator receives new assignments, vehicle positions, platoon states,

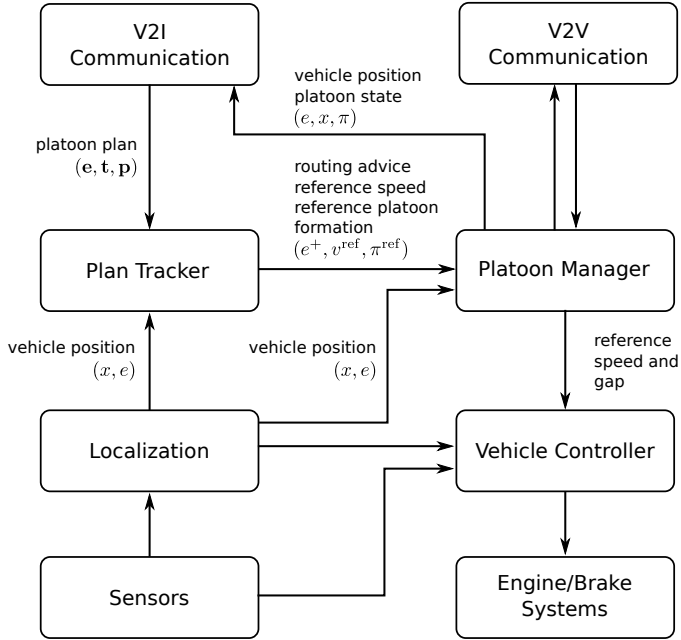


Figure 3.5: An overview of the on-board architecture of the vehicle

and data from other sources, such as weather and traffic information, in order to compute platoon plans and additional monitoring information for the FMSs. The platoon plans are tracked on-board the vehicles in order for every vehicle to reach its destination before the deadline and using platooning to reduce fuel consumption as described in Section 1.3. Figure 3.6 shows a flow chart detailing the computational steps and the data flow of the platoon coordinator. Dividing the computation of plans into these steps ensures tractability in both the design of planning algorithms and their execution.

- Route Computation/Map Matching and Assignment Update** The FMSs send two types of messages to the platoon coordinator. The first type of message contains new assignments. For a new assignment, a route e is either computed or it is matched to the road network representation of the platoon coordinator in case the FMS has already computed a route. The other type of message is a time-stamped update of the vehicle position and the platoon configuration. The time stamp is synchronized using, for instance, GPS as a clock. These messages are used to update the assignment replacing the start position \mathcal{P}^S with the reported position $(e(t^{ts}), x(t^{ts}))$ and the start time t^S with the time stamp of the position measurement t^{ts} . In most cases this does not require a recomputation of the route, but in some cases this

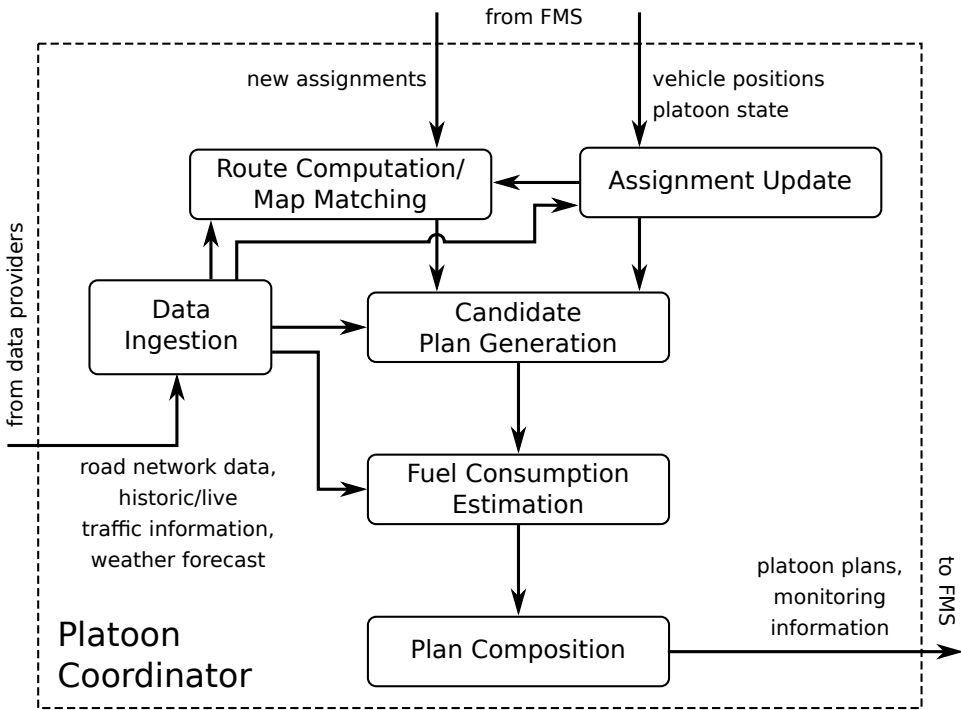


Figure 3.6: Flow chart illustrating components of the platoon coordinator and the data exchanged

might be necessary, for instance, due to a temporary road closure.

- Candidate Plan Generation** This stage computes a default platoon plan and adapted platoon plans for each assignment. The default plan involves no platooning. An adapted plan is such that one vehicle adapts to another which implements the default plan and the two vehicles drive in a platoon for some distance. For each assignment, a number of possible adapted plans are computed. Several data sources are used for the plans to be valid such as digital maps that indicate where platooning is possible according to infrastructure and legal restrictions, road topology, historic and live traffic data, weather information, and a dynamical model of the vehicle. More detail on methods for this stage are provided in Chapters 4, 5 and 7.
- Fuel Consumption Estimation** A fuel consumption model is used to estimate the fuel consumption of the default and the adapted plans. Adapted plans can lead to a lower fuel consumption as they include a phase during which vehicles drive in a platoon. Also, other cost factors can be considered such as driver wages, insurance costs, etc.

- **Plan Composition** At this stage the default and adapted plans, which are the result of the plan computation stage, are combined in a way that maximizes the combined fuel consumption reduction. In addition, the composed default and adapted plans can be jointly optimized. After this stage the system has computed a platoon plan for each assignment and this information is distributed to the respective FMSs along with monitoring information for the transport operators such as expected fuel savings. More detail on methods for this stage are provided in Chapter 6.
- **Data Ingestion** This component receives data relevant for the previous computation steps from several data providers and makes them available to the other components in the platoon coordinator. Relevant data sources are primarily static road network databases providing data such as legal speed limits, road slope, etc.; historic and live traffic measurements; and information on weather conditions.

Remark 11. There are a number of possible error scenarios such as an arrival time that is infeasible, or that no route for the vehicle can be found. These are omitted in order to keep the presentation concise.

Remark 12. Under the algorithms proposed in the following chapters, the route computation/map matching, assignment update, candidate plan generation, and fuel consumption estimation stages can be trivially parallelized on a per assignment basis. They can thus scale easily on modern computing infrastructures.

The most challenging stages from a design point of view are the candidate plan generation and the plan composition stages. The following chapters provide further detail and methods how to devise these two stages.

3.3 Summary

This chapter introduces the architecture of the coordinated platooning system. It is a hierarchical system with service, strategic, tactical, and operational layer. Each layer operates at a different time scale. Fleet management systems determine destinations and arrival deadlines for the vehicles they manage. The platoon coordinator computes platoon plans that let vehicles arrive at their destinations in time. The key feature of the platoon coordinator is that these plans instruct vehicles to platoon in a way that minimizes fuel consumption. Platoon plans are created and updated based on assignments and measured vehicle positions. They are synthesized in several computationally tractable steps: route computation/map matching, candidate generation, fuel consumption estimation, and plan composition. The platoon plans are forwarded by the fleet management systems to the vehicles' on-board systems. On board the vehicle, a plan tracker executes the plan sending reference signals to the on-board platoon manager and vehicle controller.

Chapter 4

Coordination Algorithms for Pairs of Vehicles

IN this chapter, we consider a pair of assignments that offers the possibility for platoon formation on the overlapping part of the corresponding routes. We derive a platoon plan for one vehicle, the coordination follower, for it to meet and drive in a platoon with another vehicle, the coordination leader, by computing an appropriate speed profile for the coordination follower. The coordination leader is assumed to track a default speed profile. We call such a platoon plan an adapted platoon plan. The structure of adapting one speed profile to the default speed profile of another vehicle is chosen in such a way that these profiles can be later composed resulting in platoon plans with more than two vehicles in a platoon. This way, a hierarchical structure is imposed which separates the complexity on vehicle and route level from the complexity arising from the large number of assignment combinations. Pairwise plans, such as the ones introduced in this chapter, are generated in the candidate plan generation stage introduced in Section 3.2 and combined in the plan composition stage.

In Section 4.1, we introduce the structure of adapted platoon plans consisting of a merge phase, a platooning phase, and a split phase. In Section 4.2, we consider two vehicles with the same route. One vehicle selects a speed that allows the two vehicles to meet and form a platoon. We derive how to select this speed in a fuel-optimal way. In Section 4.3, we extend this result to the case in which the two vehicles have different but overlapping routes. According to the structure introduced in Section 4.1, one vehicle adapts its speed profile in a way that allows it to meet the other vehicle on the common section of the routes and the two vehicles form a platoon. In Section 4.4, a method to computing adapted platoon plans is presented for a setting, in which the feasible speed range cannot be approximated as fixed. In that case, the fastest possible speed profile is computed from which the adapted platoon plan is derived.

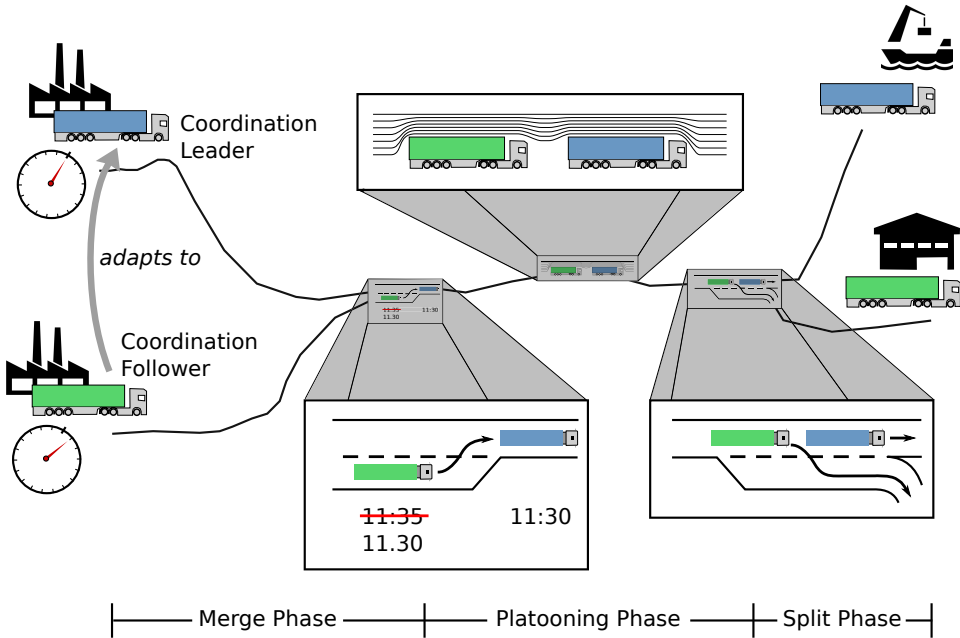


Figure 4.1: A coordination follower adapts its speed profile to a coordination leader in order to platoon on the common part of their routes.

4.1 The Structure of Adapted Platoon Plans

In this section, we introduce the structure of adapted platoon plans. We consider that one vehicle travels according to a default speed profile. We call this vehicle the coordination leader. The other vehicle, referred to as the coordination follower, adapts its speed profile in order to meet the coordination leader on the common part of the route, and platoon for some distance.

The meaning of the terms “coordination leader” and “coordination follower” becomes more apparent in Chapter 6 where adapted platoon plans are systematically combined with the goal of minimizing the combined fuel consumption by forming platoons.

Remark 13. The notion of a coordination leader/follower is different from the notion of a platoon leader/follower. A coordination leader/follower is a concept used in the composition of pairwise plans. A platoon leader, on the other hand, is the vehicle in the front of the platoon and the trailing vehicles in the platoon are platoon followers.

An adapted platoon plan consists of three phases illustrated in Figure 4.1. In the merge phase, the coordination follower keeps an average speed that allows it to meet the coordination leader on the common segment of their routes and merge into a platoon. Then the two vehicles platoon until they split up, followed by the third segment where the coordination follower drives according to a speed profile ensuring it arrives by its pre-specified deadline at its destination.

The computation of adapted speed profiles happens according to the steps illustrated in Figure 4.2. First, we identify if the routes overlap and if they do so we determine where exactly. If there is no overlap between the routes, no adapted plan can be computed. In case there are multiple disconnected overlaps between the routes, one of them is selected, for instance, taking first or the longest intersection, or computing an adapted plan for both intersections and discarding all but the most fuel efficient. Due to the hierarchical structure of road networks [13], the multiple intersections between two is not a great concern in practice. Next, the merge point along the routes is determined which has to lie both on the common section of the routes and must be feasible according to constraints on the vehicles' speeds. Due to the latter, there might be no feasible merge point on the common segment in which case there is no adapted plan. Similarly, the point for the coordination follower to split up from the coordination leader is determined. The latest such point is where the coordination leader's and the coordination follower's routes diverge. A fuel-optimal split point can lie before the end of the common route section to ensure on-time arrival or because the speed needed to arrive on time with a later split point would imply increased fuel consumption. Finally, whether or not the merge point lies before the split point decides if an adapted plan exists.

4.2 The Optimal Rendezvous Speed

In order to approach the problem of finding fuel-optimal merge and split points, consider two vehicles on the same route as depicted in Figure 4.3. The vehicles are initially separated by a distance Δd . Vehicle 0 drives at a default speed, which is denoted v_0 . Vehicle 1, which is behind vehicle 0, drives at a higher speed, denoted v_S^* . Since $v_S^* > v_0$, the distance between the vehicles decreases with time until the two vehicles meet and form a platoon. At this point, both vehicles continue driving in a platoon at default speed v_0 .

We want to select the rendezvous speed v_S^* in a fuel-optimal way, while v_0 is not altered. To this end, we introduce a linear affine fuel-model. The fuel consumption per distance traveled as platoon leader or alone is $f_0(v) = F^0 + F^1 v$, and the fuel consumption per distance traveled as platoon follower is $f_p(v) = F_p^0 + F_p^1 v$. We assume that the fuel consumption of a platoon follower is lower at default speed than if the vehicle was to travel alone, i.e., $F_p^0 + F_p^1 v_0 < F^0 + F^1 v_0$. It is reasonable to assume this since without this assumption there is no reason to form platoons for the sake of fuel consumption reduction. We assume that v_0 lies within the feasible

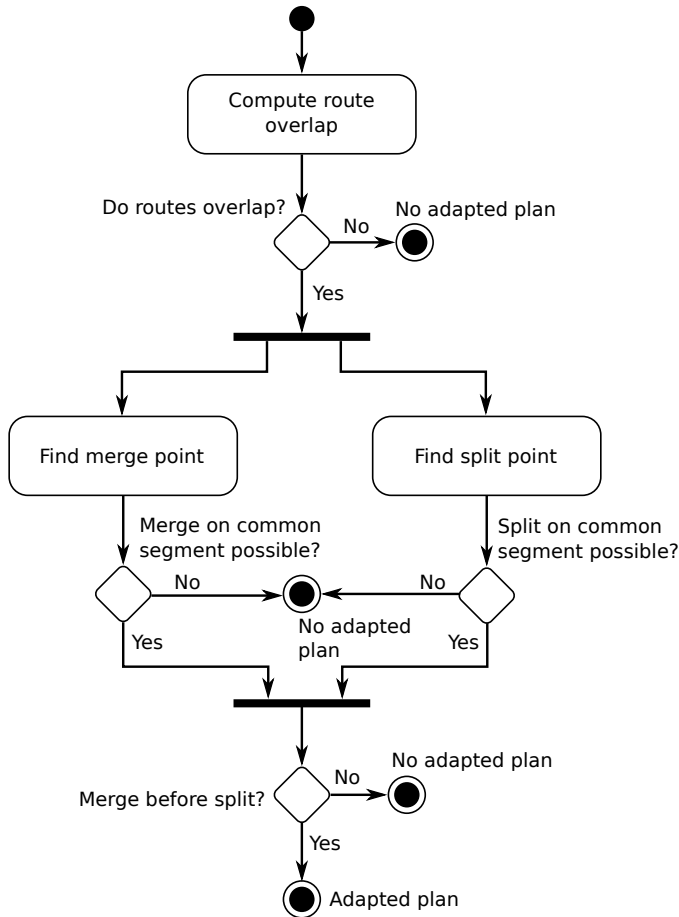


Figure 4.2: Activity diagram showing the steps in computing adapted platoon plans

speed-range, i.e., $0 < v_{\min} \leq v_0 \leq v_{\max}$. The optimal rendezvous speed v_S^* is also constrained to lie within the feasible speed range.

Remark 14. An affine fuel model is used here as it admits an analytical solution giving more insight than a data-driven one. Extending the result to other differentiable fuel models, such as higher order polynomials, is possible along the same lines of reasoning using numerical minimization.

This problem setting is related to the optimal catch-up schemes derived in [144]. In fact, the catch-up schemes from [144] have been combined with the methods of Chapter 6 in a simulation study presented in [35].




Figure 4.3: Two vehicles on the same route with inter-vehicle distance Δd . Vehicle 0 has speed v_0 and vehicle 1 has speed $v_S^* > v_0$. Since the speed of vehicle 1 is higher than the speed of vehicle 0, vehicle 1 will catch up with vehicle 0, and the two vehicles will form a platoon.

A similar scenario to the one described above is a setting where vehicle 1 is in front of vehicle 0. In that case, vehicle 1 selects a speed smaller than v_0 . This means that vehicle 0 will catch up to vehicle 1 instead, and the two vehicles can form a platoon. In the remainder of this section, we consider both the case in which vehicle 1 is behind vehicle 0 and the case that vehicle 1 is in front of vehicle 0.

We model this scenario on a road network with one road segment with road segment identifier e . The length of the road segment corresponding to e , i.e., $L(e)$, is assumed to be long enough to not impose any restrictions on where the two vehicles meet. The time, when the two vehicles start is denoted t^S , the time when they meet and start platooning is denoted t^M , and the time when they stop platooning is denoted t^{SP} .

The following proposition gives the optimal rendezvous speed v_S^* for vehicle 1.

Proposition 1. *Assume the following: The speed of vehicle 0 is constant v_0 with $v_0 \in \mathbb{R}$, $v_0 > 0$. The position of vehicle 0 at time t^S is $(e, x_0(t^S))$. The position of vehicle 1 at time t^S is $(e, x_1(t^S))$. Vehicle 1 platoons with vehicle 0 between time t^M and t^{SP} with $t^{SP} > t^M$. Vehicle 1 has constant speed v_S for time t^S to t^M and v_0 from time t^M to t^{SP} . The rendezvous speed v_S is constrained to the interval $[v_{\min}, v_{\max}]$.*

Then the rendezvous speed v_S^ that minimizes fuel consumption from time t^S to t^{SP} is given by*

$$v_S^* = \begin{cases} \max \left(v_0 \left(1 - \sqrt{1 - \frac{F_p^1}{F^1} + \frac{\Delta F^0}{F^1 v_0}} \right), v_{\min} \right) & \text{if } \Delta d < 0 \\ \min \left(v_0 \left(1 + \sqrt{1 - \frac{F_p^1}{F^1} + \frac{\Delta F^0}{F^1 v_0}} \right), v_{\max} \right) & \text{if } \Delta d > 0 \\ v_0 & \text{if } \Delta d = 0, \end{cases} \quad (4.1)$$

where $\Delta d = x_0(t^S) - x_1(t^S)$ and $\Delta F^0 = F^0 - F_p^0$.

Proof. Let $\Delta d^S = x_1(t^M) - x_1(t^S)$. Let $D_0 = x_1(t^{SP}) - x_1(t^S)$. Under the assump-

tion that $v_S \neq v_0$, we have the relation

$$\Delta d^S = \frac{v_S}{v_S - v_0} \Delta d. \quad (4.2)$$

At time t^M we have $x_0(t^M) = x_1(t^M)$. After the meeting point, both vehicles platoon at speed v_0 . Assume that vehicle 0 is the platoon leader once the two vehicles have merge into a platoon. Hence, the total fuel consumption of 1 up to some distance from the current position D_0 , which fulfills $D_0 > \Delta d^S$, becomes

$$f_0(v_S) \Delta d^S + f_p(v_0)(D_0 - \Delta d^S) = (f_0(v_S) - f_p(v_0)) \Delta d^S + f_p(v_0) D_0.$$

The fuel consumption of vehicle 0 is not affected by v_S . We see that the term $f_p(v_0) D_0$ is not a function of v_S , so the optimal rendezvous speed does not depend on the total distance traveled. In order to find the optimal v_S , we can therefore consider the remaining terms denoted as $f_r(v_S)$ and get with (4.2) and the definitions of f_0 , f_p

$$f_r(v_S) = (f_0(v_S) - f_p(v_0)) \Delta d^S = (F^1 v_S - F_p^1 v_0 + \Delta F^0) \frac{v_S}{v_S - v_0} \Delta d,$$

with $\Delta F^0 = F^0 - F_p^0$. We take the derivative of the above expression in order to find its extrema

$$\frac{\partial}{\partial v_S} f_r(v_S) = \frac{\Delta d}{(v_S - v_0)^2} (F^1 v_S^2 - 2F^1 v_0 v_S + F_p^1 v_0^2 - \Delta F^0 v_0).$$

In order to find the extrema \tilde{v}_S , we check where this expression is zero. We can assume that $\Delta d \neq 0$, otherwise $\Delta d^S = 0$, which means that the vehicles can directly start platooning. Therefore,

$$0 = (F^1 (\tilde{v}_S)^2 - 2F^1 v_0 \tilde{v}_S + F_p^1 v_0^2 - \Delta F^0 v_0) \quad (4.3)$$

$$\tilde{v}_S = v_0 \left(1 \pm \sqrt{1 - \frac{F_p^1}{F^1} + \frac{\Delta F^0}{F^1 v_0}} \right). \quad (4.4)$$

We have to differentiate between two cases. Either $\Delta d > 0$, which implies $v_S > v_0$, i.e., the coordination follower speeds up, or $\Delta d < 0$, which implies $v_S < v_0$, i.e., the coordination follower slows down. Otherwise Δd^S becomes negative. There are two solutions for \tilde{v}_S , one where $\tilde{v}_S > v_0$, and the other $\tilde{v}_S < v_0$. The appropriate one, depending on Δd , is \tilde{v}_S , the optimal unconstrained rendezvous speed.

We can verify that this is indeed a minimum by considering the asymptotic behavior of $f_r(v_S)$ when v_S approaches $\pm\infty$ and when it approaches v_0 . Assume $\Delta d > 0$ so that $\tilde{v}_S > v_0$. We have

$$\begin{aligned} \lim_{v_S \rightarrow \infty} f_r(v_S) &= \infty, \\ \lim_{v_S \rightarrow v_0^+} f_r(v_S) &= \infty \end{aligned}$$

where we used that $f_0(v_0) > f_p(v_0)$ so that the term $f_0(v_0) - f_p(v_0)$ becomes positive, which is the prerequisite to save fuel by platooning. When we have $\Delta d < 0$, so that $\tilde{v}_S < v_0$, then

$$\begin{aligned}\lim_{v_S \rightarrow -\infty} f_r(v_S) &= \infty, \\ \lim_{v_S \rightarrow v_0^-} f_r(v_S) &= \infty.\end{aligned}$$

This shows that if $\tilde{v}_S > v_{\max}$, then $v_S^* = v_{\max}$, if $\tilde{v}_S < v_{\min}$, then $v_S^* = v_{\min}$, and $v_S^* = \tilde{v}_S$ otherwise.

In order to have real solutions for (4.3), we need

$$\begin{aligned}1 - \frac{F_p^1}{F^1} + \frac{\Delta F^0}{F^1 v_0} > 0 &\Leftrightarrow F_p^1 v_0 + F_p^0 < F^1 v_0 + F^0 \\ &\Leftrightarrow f_p(v_0) < f_0(v_0),\end{aligned}$$

which is the condition that the coordination follower saves fuel when platooning. The larger the difference $f_0(v_0) - f_p(v_0)$, the larger the absolute difference between v_0 and v_S^* , i.e., the longer the vehicles platoon.

Since the order of the platoon does not change the total fuel consumption of both vehicles, the result also holds in case vehicle 1 is the platoon leader and vehicle 0 the platoon follower. \square

Proposition 1 yields some insights into the computation of adapted platoon plans. Equation 4.1 shows that change in speed increases with the difference in fuel consumption between driving alone and as a coordination follower, as the term $\sqrt{1 - \frac{F_p^1}{F^1} + \frac{\Delta F^0}{F^1 v_0}}$ increases monotonically with ΔF^0 and decreases with F_p^1 . Furthermore, the rendezvous speed v_S^* is not dependent on how long the vehicles platoon which motivates computing the merge speed and the split speed independently. Another insight is, that should it not be possible to select a speed as high or low as the optimal merge speed without speed constraints, the speed closest to the optimal merge speed is optimal. In fact, using typical values for heavy vehicles on a motorway setting, oftentimes the rendezvous speed is determined by the speed constraints v_{\min} and v_{\max} .

4.3 Adapted Platoon Plans for a Fixed Speed Range

In this section, we detail how the individual computation steps shown in Figure 4.2 are implemented using the linear affine fuel model introduced in Section 4.2 and a fixed speed range $[v_{\min}, v_{\max}]$. The default speed profile of the coordination leader is a fixed speed v_0 . For the sake of simplicity, we assume in this section the vehicles to arrive at their destinations exactly on their respective deadlines. This can be extended to arrival before the deadline.

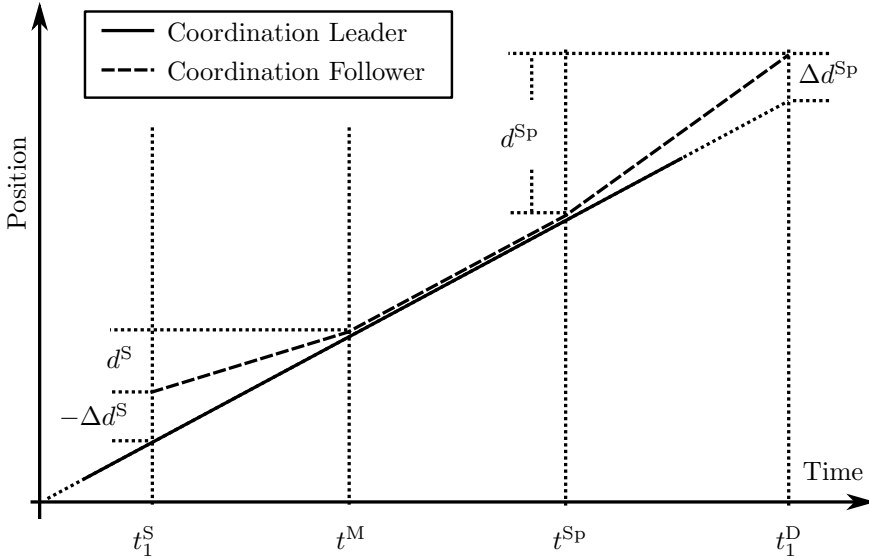


Figure 4.4: Speed profiles of the coordination leader and the coordination follower. The distance along the respective route with respect to a common reference point on the common part of the route is plotted over time. The coordination leader has a constant speed. In this example, the coordination follower drives slower at the beginning of its journey. Once it meets the coordination leader, the two vehicles platoon. At the end the coordination follower drives at an increased speed in order to make its deadline.

Remark 15. Having a fixed speed range from which the vehicle can choose is a simplification. Recall that these plans are locally tracked by the plan tracker on-board the vehicle. The speed profile in the plan is therefore to be interpreted as an average speed over some distance.

The route, which is computed in the route computation/map matching stage is a sequence of road segment identifiers $\mathbf{e} = (\mathbf{e}[1], \mathbf{e}[2], \dots, \mathbf{e}[N^A])$, where N^A is a notational convenience to refer to the number of elements in \mathbf{e} . In order to simplify notation, we define the distance $d_{\mathbf{e}}$ between two positions $(\mathbf{e}[i_1], x_1)$, $(\mathbf{e}[i_2], x_2)$ with respect to a route \mathbf{e} . Recall that a position $(e, x) \in \mathcal{E}_r \times \mathbb{R}$ in the road network is represented by a route segment identifier e and the distance from the start of the road segment x .

Definition 1 (Distance). Let i_1, i_2 be such that $N^A \geq i_2 \geq i_1$. Then,

$$d_{\mathbf{e}}((\mathbf{e}[i_1], x_1), (\mathbf{e}[i_2], x_2)) = \left| x_2 - x_1 + \sum_{i=i_1}^{i_2-1} L(\mathbf{e}[i]) \right|$$

Consider a coordination leader with index 0 and a coordination follower with index 1. Vehicles 0, 1 start at $\mathcal{P}_0^S = (e_0^S, x_0^S)$, $\mathcal{P}_1^S = (e_1^S, x_1^S)$ at time t_0^S , t_1^S and arrive at $\mathcal{P}_0^D = (e_0^D, x_0^D)$, $\mathcal{P}_1^D = (e_1^D, x_1^D)$ at time t_0^D , t_1^D , respectively. We denote the position at which the coordination leader and the coordination follower start platooning at time t^M as $\mathcal{P}^M = (e^M, x^M)$ and where they split at time t^{Sp} as $\mathcal{P}^{Sp} = (e^{Sp}, x^{Sp})$. These meeting points have to lie on the trajectory of the coordination leader with constant speed v_0 :

$$\begin{aligned} d_{e_0}(\mathcal{P}_0^S, \mathcal{P}^M) &= v_0(t^M - t_0^S), \\ d_{e_0}(\mathcal{P}_0^S, \mathcal{P}^{Sp}) &= v_0(t^{Sp} - t_0^S). \end{aligned}$$

When platooning with the coordination leader the planned trajectory of the coordination follower consists of three phases: from start to the meeting point with speed v_S , from meeting point to the split point platooning as platoon follower of 0 with speed v_0 , and from the split point to the destination with speed v_{Sp} . We define $d^S = d_{e_1}(\mathcal{P}_1^S, \mathcal{P}^M)$ and $d^{Sp} = d_{e_1}(\mathcal{P}^{Sp}, \mathcal{P}_1^D)$. We have the relations

$$\begin{aligned} d^S &= v_S(t^M - t_1^S), \\ d^{Sp} &= v_{Sp}(t_1^D - t^{Sp}). \end{aligned}$$

We define the virtual position difference at the start/end of the coordination follower's trajectory as

$$\begin{aligned} \Delta d^S &= d^S - (t^M - t_1^S)v_0, \\ \Delta d^{Sp} &= d^{Sp} - (t_1^D - t^{Sp})v_0, \end{aligned} \tag{4.5}$$

which are equivalent to Δd in Proposition 1. If $\Delta d^S > 0$ then $v_S > v_0$, if $\Delta d^S < 0$ then $v_S < v_0$, if $\Delta d^{Sp} > 0$ then $v_{Sp} > v_0$, and if $\Delta d^{Sp} < 0$ then $v_{Sp} < v_0$. Then, we can compute according to (4.4) the appropriate, fuel-optimal speed v_S^* for the first and the last phase. Proposition 1 considers that the two vehicles are initially separated. The same lines of reasoning apply in order to determine the optimal speed of the coordination follower during the last phase.

This derivation has not taken into account so far that the first possible point to merge is when the coordination leader's and the coordination follower's routes meet. If v_S^* leads to a distance from \mathcal{P}_1^S to the merge point that is too small, then the coordination leader selects a speed that lets the coordination leader and coordination follower merge at the position where the two routes meet, denoted here $(e^F, 0)$. This speed is

$$v_S = \frac{d_{e_1}(\mathcal{P}_1^S, (e^F, 0))}{t^M - t_1^S}.$$

The corresponding case might occur at split up, so that

$$v_{Sp} = \frac{d_{e_1}((e^L, L(e^L)), \mathcal{P}_1^D)}{t_1^D - t^{Sp}},$$

where $(e^L, L(e^L))$ is the position where the coordination leader's and the coordination follower's routes split up.

The first test if platooning is possible and beneficial is, whether the calculated merge point lies before the split point or not, i.e., whether

$$d^S + d^{Sp} < d_{e_1}(\mathcal{P}_1^S, \mathcal{P}_1^D).$$

The estimated fuel consumption for the coordination follower with the speed profile that is adapted for platooning with the coordination leader is

$$F = d^S f_0(v_S) + d^{Sp} f_0(v_{Sp}) + (d_{e_1}(\mathcal{P}_1^S, \mathcal{P}_1^D) - d^S - d^{Sp}) f_p(v_0). \quad (4.6)$$

If F is smaller than the fuel consumption that results from traveling alone at a constant speed, it is beneficial that the vehicles platoon.

The results of this section can be summarized as follows. The optimal speed profile of a coordination follower with index 1 to a coordination leader with index 0 consists of three phases with constant speed: v_S from t_1^S to t^M , then v_0 from t^M to t^{Sp} , and finally v_{Sp} from t^{Sp} to t_1^D , where coordination leader and follower platoon from time t^M to t^{Sp} .

Remark 16. In this model, the order of the platoon does change the combined fuel consumption in the platoon. It is convenient for the derivation to assume that the coordination follower is also platoon follower. Since we optimize the fuel consumption of all vehicles collectively, the results also apply when the coordination follower becomes the platoon leader.

4.4 Adapted Platoon Plans for Position-Dependent Speeds

In this section, we derive how adapted platoon plans as introduced in Section 4.1 can be computed considering that the maximum speed of a vehicle depends on the position in the road network. Heavy vehicles are less affected by local speed restrictions on highways compared to cars since in many countries the allowed maximum speed for heavy vehicles is lower than the posted speed limit in most parts of the highway network. However, it can still be important to consider these cases for accurate planning. Furthermore, the vehicle speed can be limited by dense traffic, which can be predicted and consequently modeled as a speed limit. Finally, heavy vehicles have small power-to-weight ratio and are thus limited in their speed when climbing steep uphill sections. By taking these effects into account, less deviations from the platoon plans occur.

We use the route segment identifiers of $\mathbf{e} = (\mathbf{e}[1], \mathbf{e}[2], \dots, \mathbf{e}[N^A])$ as a natural discretization for all data in the platoon plan as compared to, for instance, a discretization of time. This is convenient as most relevant properties are primarily a function of the road segment. We assume that a vehicle with assignment index n starts at position $\mathcal{P}_n^S = (\mathbf{e}_n[1], 0)$ and has destination $\mathcal{P}_n^D = (\mathbf{e}_n[N_n^A], L(\mathbf{e}_n[N_n^A]))$,

where $\mathbf{e}_n[1]$ is the first and $\mathbf{e}_n[N_n^A]$ the last element in the route computed for assignment n .

The time profile $\mathbf{t} = (\mathbf{t}[1], \mathbf{t}[2], \dots, \mathbf{t}[N^A + 1])$ according to which the vehicle is supposed to travel along the route is represented by a list of segment start times. For $i \in \{1, \dots, N^A\}$, $\mathbf{t}[i]$ is the reference time when the vehicle should start traversing the road segment identified by $\mathbf{e}[i]$. The first element $\mathbf{t}[1]$ is the start time of the vehicle and $\mathbf{t}[N^A + 1]$ is the arrival time of the vehicle.

The time profile can be converted into a speed profile \mathbf{v} approximating the speed on a road segment as constant. The length of a road segment is denoted $L : \mathcal{E}_r \rightarrow \mathbb{R}^+$. It is the distance a vehicle travels from the beginning to the end of the road segment. The speed profile is defined as

$$\mathbf{v}[i] = \frac{L(\mathbf{e}[i])}{\mathbf{t}[i+1] - \mathbf{t}[i]}, \quad i \in \{1, \dots, N^A\}.$$

The speed profile \mathbf{v} combined with the start time $\mathbf{t}[1]$ can be used to compute the corresponding time profile \mathbf{t} as

$$\mathbf{t}[i] = \mathbf{t}[1] + \sum_{j=1}^{i-1} \frac{L(\mathbf{e}[j])}{\mathbf{v}[j]}, \quad i \in \{2, \dots, N^A + 1\}. \quad (4.7)$$

Constraints such as maximum speeds are more conveniently expressed in the speed profile representation. For vehicles to platoon they have to be at the same location at the same time which is easier to express in terms of the time profile. Hence, we work with both representations and (implicitly) convert from one representation to the other whenever necessary in order to not violate the speed limit.

We derive time profiles based on a maximum speed profile $\bar{\mathbf{v}}$. The maximum speed profile corresponds to a vehicle that at any point drives as close to the maximum legal speed \mathbf{v}^{\max} as possible, i.e., that accelerates whenever possible and decelerates only as much as needed.

Definition 2. The maximum speed profile $\bar{\mathbf{v}}$ is defined for a speed limit \mathbf{v}^{\max} , maximum accelerations $\Delta\mathbf{v}^{\max}$, and minimum accelerations $\Delta\mathbf{v}^{\min}$ as

$$\bar{\mathbf{v}}[i] = \min(\mathbf{v}^{\text{fwd}}[i], \mathbf{v}^{\text{bwd}}[i]), \quad i \in \{1, \dots, N^A\}$$

where

$$\begin{aligned} \mathbf{v}^{\text{fwd}}[i+1] &= \min(\mathbf{v}^{\text{fwd}}[i] + \Delta\mathbf{v}^{\max}[i], \mathbf{v}^{\max}[i+1]), & \mathbf{v}^{\text{fwd}}[1] &= \mathbf{v}^{\max}[1] \\ \mathbf{v}^{\text{bwd}}[i-1] &= \min(\mathbf{v}^{\text{bwd}}[i] - \Delta\mathbf{v}^{\min}[i-1], \mathbf{v}^{\max}[i-1]), & \mathbf{v}^{\text{bwd}}[N^A] &= \mathbf{v}^{\max}[N^A]. \end{aligned}$$

The speed limit \mathbf{v}^{\max} covers legal restrictions that depend on the vehicle, road segment, and also the limiting effect of surrounding traffic. Legal restrictions are mostly straightforward to retrieve from databases. Predicting the effect of traffic is more involved and is based both on historic and real-time measurements in

combination with advanced prediction models. The maximum acceleration over road segment $e[i]$, $\Delta \mathbf{v}^{\max}[i]$ takes into account the limited power-to-weight ratio of heavy vehicles. This becomes particularly important on hilly roads where maximum engine power is not sufficient to keep the legal speed. It can be useful to model $\Delta \mathbf{v}^{\max}[i]$ as a function of $\mathbf{v}^{\text{fwd}}[i]$ using a dynamic vehicle model. The minimum acceleration $\Delta \mathbf{v}^{\min}[i] \leq 0$ over the i th segment is limited by safety and comfort considerations. The physical constraints of braking systems are typically only reached in emergency situations. Limits of the brake system in long downhill slopes can be taken into account by lowering the speed limit \mathbf{v}^{\max} on that part of the route. The forward speed profile \mathbf{v}^{fwd} models the effect that the vehicle is limited by the maximum acceleration $\Delta \mathbf{v}^{\max}$ when trying to gain a speed as close as possible to \mathbf{v}^{\max} . The backward speed profile \mathbf{v}^{bwd} models that the vehicle must decelerate in time in order to avoid exceeding the minimum acceleration $\Delta \mathbf{v}^{\min}$ or the speed limit \mathbf{v}^{\max} .

Speed profiles are computed by scaling the maximum speed profile $\bar{\mathbf{v}}$. The default speed profile is the speed profile used by the coordination leader or when the vehicle travels alone.

Definition 3. The default speed profile is defined as

$$\mathbf{v}^{\text{d}}[i] = \sigma \bar{\mathbf{v}}[i], \quad i \in \{1, \dots, N^{\text{A}}\},$$

where

$$\sigma = \max \left(\sigma_{\text{d}}, \frac{t^{\text{A}} - t^{\text{S}}}{t^{\text{D}} - t^{\text{S}}} \right),$$

with $t^{\text{A}} = \underline{t}[N^{\text{A}} + 1] \leq t^{\text{D}}$ being the arrival time according to the maximum speed profile and $0 < \sigma_{\text{d}} \leq 1$ being a design parameter. The time profile corresponding to \mathbf{v}^{d} according to (4.7) is denoted \mathbf{t}^{d} with $\mathbf{t}^{\text{d}}[1] = t^{\text{S}}$.

Remark 17. If $t^{\text{A}} > t^{\text{D}}$, the vehicle cannot meet its deadline, which means the assignment is ill-posed and the vehicle travels as fast as possible to its destination. The factor σ_{d} determines how much slower a vehicle travels compared to the maximum speed profile when the deadline permits it. The benefit of letting the vehicle travel slower is more room for the controller consisting of plan tracker, platoon manager, and vehicle controller to account for disturbances. It also allows other vehicles to catch up to the platoon. We consider σ_{d} being a choice of the transport operator.

Consider now a pair of vehicles with assignment indices n and m . We want to determine if it is possible to compute an adapted speed profile so that the two vehicles platoon during part of their journey. We denote the index of the first segment of the route that vehicle n has in common with vehicle m on overlapping part of the route as $\underline{N}_{n,m}^{\text{M}}$ and the last as $\overline{N}_{n,m}^{\text{Sp}}$, so $\mathbf{e}_n[\underline{N}_{n,m}^{\text{M}}] = \mathbf{e}_m[\underline{N}_{m,n}^{\text{M}}] = e^{\text{F}}$ and

$\mathbf{e}_n[\overline{N}_{n,m}^{\text{Sp}}] = \mathbf{e}_m[\overline{N}_{m,n}^{\text{Sp}}] = e^L$, where $\mathbf{e}_n, \mathbf{e}_m$ denote the routes for assignments n, m respectively. This means that

$$\mathbf{e}_n[\underline{N}_{n,m}^{\text{M}} + i] = \mathbf{e}_m[\underline{N}_{m,n}^{\text{M}} + i], \quad i \in \{0, \dots, \overline{N}_{n,m}^{\text{Sp}} - \underline{N}_{n,m}^{\text{M}}\}.$$

Note that $\overline{N}_{n,m}^{\text{Sp}} - \underline{N}_{n,m}^{\text{M}} = \overline{N}_{m,n}^{\text{Sp}} - \underline{N}_{m,n}^{\text{M}}$.

The adapted speed profile is computed in a way that the vehicle with the adapted speed profile meets the vehicle with the default speed profile on the common part of their routes in order to platoon until the two routes split as shown in Figure 4.5. For notational convenience, we introduce the offset in the route segment index on the common part of the route $\Delta N_{n,m} = \underline{N}_{m,n}^{\text{M}} - \underline{N}_{n,m}^{\text{M}}$.

Definition 4. The adapted speed profile of n adapted to the default speed profile \mathbf{v}_m^{d} of assignment m is denoted as $\mathbf{v}_{n,m}^{\text{a}}$ and the corresponding time profile denoted $\mathbf{t}_{n,m}^{\text{a}}$ are such that

$$\mathbf{v}_{n,m}^{\text{a}}[i] = \begin{cases} \sigma_{\text{M}} \overline{\mathbf{v}}_n[i] & \text{for } i \in \{1, \dots, N_{n,m}^{\text{M}} - 1\} \\ \mathbf{v}_m^{\text{d}}[i + \Delta N_{n,m}] & \text{for } i \in \{N_{n,m}^{\text{M}}, \dots, N_{n,m}^{\text{Sp}}\} \\ \sigma_{\text{Sp}} \overline{\mathbf{v}}_n[i] & \text{for } i \in \{N_{n,m}^{\text{Sp}} + 1, \dots, N_n^{\text{A}}\}, \end{cases}$$

with $\underline{\sigma} \leq \sigma_{\text{M}} \leq 1$ and $\sigma_{\text{d}} \leq \sigma_{\text{Sp}} \leq 1$, and with

$$\mathbf{e}_n[i] = \mathbf{e}_m[i + \Delta N_{n,m}], \quad \text{for } i \in \{N_{n,m}^{\text{M}}, \dots, N_{n,m}^{\text{Sp}}\},$$

and

$$\mathbf{t}_{n,m}^{\text{a}}[N_n^{\text{A}} + 1] \leq t_n^{\text{D}},$$

and

$$\mathbf{t}_{n,m}^{\text{a}}[N_{n,m}^{\text{M}}] = \mathbf{t}_m^{\text{d}}[N_{n,m}^{\text{M}} + \Delta N_{n,m}].$$

Figure 4.5 shows an example of an adapted plan. The coordination follower selects speeds higher than its default speed profile during the merge phase. During the platooning phase, coordination leader and follower drive according to the default speed profile of the coordination leader. Finally, during the split phase, the coordination follower drives according to its default speed profile since it is ahead of schedule from the higher speeds during the merge phase.

To check whether or not an adapted speed profile that fulfills the Definition 4 exists and computing one according to the procedure shown in Figure 4.2 if it exists is straightforward as such but entails many steps and a more detailed description is thus omitted here.

Remark 18. If the vehicle is parked at the beginning of its trip, the start time $\mathbf{t}_{n,m}^{\text{a}}[1] \geq t^{\text{S}}$ can be adjusted in order to minimize $|\sigma_{\text{d}} - \sigma_{\text{M}}|$, i.e., the deviation from the default speed profile during the first part of the adapted speed profile.




Figure 4.5: A coordination follower scales the default speed profile to platoon with a coordination leader on the common part of their routes.

Remark 19. The speed difference between consecutive elements of the adapted speed profile can lie outside the interval $[\Delta \mathbf{v}^{\min}, \Delta \mathbf{v}^{\max}]$ at parts of route where $\Delta \mathbf{v}^{\max} < 0$ and at the transition from road segment $\mathbf{e}_n[N_{n,m}^M - 1]$ to $\mathbf{e}_n[N_{n,m}^M]$ and from $\mathbf{e}_n[N_{n,m}^{\text{Sp}} - 1]$ to $\mathbf{e}_n[N_{n,m}^{\text{Sp}}]$. Negative maximum acceleration can occur, for instance, at the beginning of steep uphill sections. The resulting deviations are small and infrequent and can thus be compensated for by the on-board controller.

Remark 20. We neglect the small time gap between the vehicles in the platoon for the sake of planning as they are small compared to disturbances from traffic, road grade, etc. Once a platoon is formed, the platoon manager ensures cohesion of the platoon. The objective of the coordinator is to get vehicles close enough that platoons can be formed and small deviations from the time profiles are expected and they are dealt with by lower layers of control.

4.5 Summary

This chapter considers that a vehicle, a coordination follower, follows an adapted platoon plan so that it meets another vehicle, a coordination leader, during its journey and they platoon together. The speed profile of the coordination leader is fixed. An adapted plan consists of three phases. In the first phase the speed is set

to a value so the vehicles meet to form a platoon, which is the start of the second phase. During the second phase the vehicles platoon. At the end of the second phase, the vehicles split up and the coordination follower selects a speed that lets it arrive by its deadline at its destination.

Based on an affine fuel model, an analytical expression of the fuel-optimal speed for the first and the last phase is derived. Taking into account that platooning can only happen on the common part of the routes, we arrive at the fuel-optimal adapted plan for a fixed speed range. In order to handle space-dependent speed restrictions due to legal speed limits, traffic, and limited hill climbing capabilities of heavy-duty vehicles, a maximum speed profile is computed. By scaling the maximum speed profile, a default and an adapted speed profile is derived.

Chapter 5

Coordination Algorithms under Uncertain Travel Times

THIS chapter introduces coordination algorithms to form platoons en route that take the uncertainty of travel times into consideration. In deriving the pairwise planning algorithms introduced in Chapter 4, we assume that travel times of segments along the route can be deterministically planned. Despite progress in travel time prediction algorithms using the increasingly available travel time data and traffic flow models, much uncertainty remains predicting travel times up to several hours ahead as relevant for platoon coordination. While some of the effects of travel time uncertainty can be mitigated reactively by frequently updating plans using real-time feedback from the vehicles, it is advantageous to estimate the probability of two vehicles successfully merging into a platoon. The success probability for a merge to succeed is used when taking the decision whether or not two vehicles should platoon penalizing plans with a small estimated probability of success. A vehicle pair with small nominal benefit from platooning but a high chance of success can be more beneficial to pursue than one with large nominal benefit but small probability of success. For instance, a vehicle A could either platoon with vehicle B or C. Vehicles A and B have a longer part of their routes in common than A and C and the potential fuel consumption reduction for A and B is lower than of A and C. However, B passes through a densely populated area with frequent traffic jams resulting in high travel time uncertainty in that region. This results in a small chance of A and B successfully forming a platoon and therefore it is preferable A and C attempt to form a platoon on common part of their routes.

In Section 5.1, we model the problem of two vehicle merging at a highway intersection. Both vehicles pass a finite number of road segments until the merging point and the traversal time of a road segment is randomly distributed with known distribution. The distribution of travel times is affected by a reference speed, which is fixed for one of the vehicles. The other vehicle adapts the reference speed in order to maximize the probability of the two vehicles merging at the intersection.

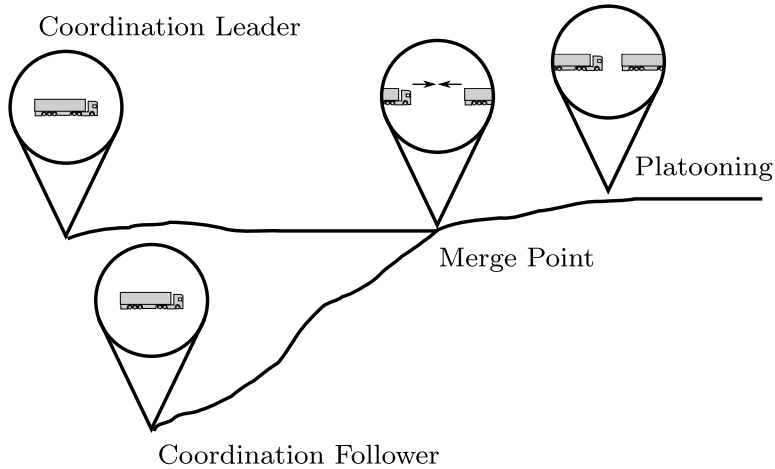


Figure 5.1: In the considered scenario, two vehicles are to merge into a platoon at the intersection of their routes by adapting their speed on the way leading to that intersection.

Section 5.2 shows how to control the speed in an optimal way using dynamic programming. It also derives a way to improve the computational efficiency of the dynamic programming algorithm when a small error in the computed value function can be accepted. It also considers the case when travel time distributions are correlated across adjacent road segments. In Section 5.3, we discuss how the influence of the reference speed can be modeled. An illustrative simulation example is presented in Section 5.4.

5.1 Modeling

Consider the scenario depicted in Figure 5.1 of two vehicles approaching an intersection at which they are supposed to merge into a platoon. The reference speed of the coordination leader is fixed. The speed of the other vehicle, the coordination follower, is controlled to maximize the probability of merging with the coordination leader. This setting of a coordination follower adapting to a coordination follower is similar to the setting considered in Chapter 4.

First, we model the movement of a single vehicle until the merge point. We consider that the route is partitioned into a finite number of segments. We consider discrete time and represent it as integers where the measurement unit is such that one increment corresponds to a sufficiently small discretization interval. The traversal time T^i of the i -th segment is a random variable. Let $t^i \in \mathbb{Z}$ be the time

the vehicle starts traversing the i -th route segment, in the following referred to as segment arrival time.

The arrival time at the next segment is the sum of the arrival time at the previous segment and the traversal time of the segment:

$$t^{i+1} = t^i + T^i. \quad (5.1)$$

The traversal time $T^i \in \mathbb{Z}$ is a random variable that is assumed only to be dependent on the reference speed at the i -th segment $v_{\text{ref}}^i \in \mathbb{V}$, which is considered to be a control input. The domain of \mathbb{V} is a finite set of reference speeds. It is assumed that $T_{\text{min}}^i \leq T^i \leq T_{\text{max}}^i$, see Figure 5.2, where, for instance, T_{min}^i can be derived from the maximum speed of a heavy-duty vehicle and T_{max}^i from the largest observed traversal time on that segment ever recorded.

Since T^i is assumed to depend only on the control input v_{ref}^i , eq. (5.1) describes a Markov decision process where t^i denotes the value of its state at the i -th stage and \mathbb{V} is the set of actions. Note that the stages in the decision process correspond to locations. Let $p_{T^i}(\tau|v_{\text{ref}}^i)$ denote the probability of $T^i = \tau$ conditioned on v_{ref}^i . The transition probability between state t^i to t^{i+1} is the probability that $T^i = t^{i+1} - t^i$, and thus the probability distribution of t^i can be recursively computed as

$$\begin{aligned} p_{t^{i+1}}(\mathbf{t}) &= \sum_{\tau=-\infty}^{\infty} p_{T^i}(\tau|v_{\text{ref}}^i) p_{t^i}(\mathbf{t} - \tau) \\ &= \sum_{\tau=T_{\text{min}}^i}^{T_{\text{max}}^i} p_{T^i}(\tau|v_{\text{ref}}^i) p_{t^i}(\mathbf{t} - \tau). \end{aligned} \quad (5.2)$$

Note that p_{T^i} can also be modeled conditioned on the segment arrival time t^i to reflect that travel time distributions are time dependent. Let t_{ℓ}^i denote the segment arrival time of the coordination leader at the i -th segment of its route. We consider that the reference speed of the coordination leader is given as v_{ℓ} and its start time t_{ℓ}^S is known meaning that $p_{t_{\ell}^i}(\mathbf{t}) = 1$ if $\mathbf{t} = t_{\ell}^S$ and $p_{t_{\ell}^i}(\mathbf{t}) = 0$ otherwise. Let N_{ℓ} be the index in the coordination leader's route at which the coordination leader and the coordination follower are supposed to meet. The probability distributions of $t_{\ell}^{N_{\ell}}$ are recursively computed from (5.2).

We assume that a coordination leader and a coordination follower can platoon if they arrive at the merge point with an absolute time difference of at most Δt , which is chosen small enough so that they can establish V2V communication and initiate a merge maneuver. Furthermore, being at this point close-by on the same route, both vehicles are subject to approximately the same traffic disturbances. In this case, the two vehicles can initiate a catch-up maneuver during which the vehicle in front selects a lower speed than the vehicle that catches up until the vehicles merge. This coordination phase is considered to be so small that its impact on the platooning benefit can be neglected.

The objective is to control the coordination follower so that the probability of platooning is maximized. Furthermore, we want to compute this probability

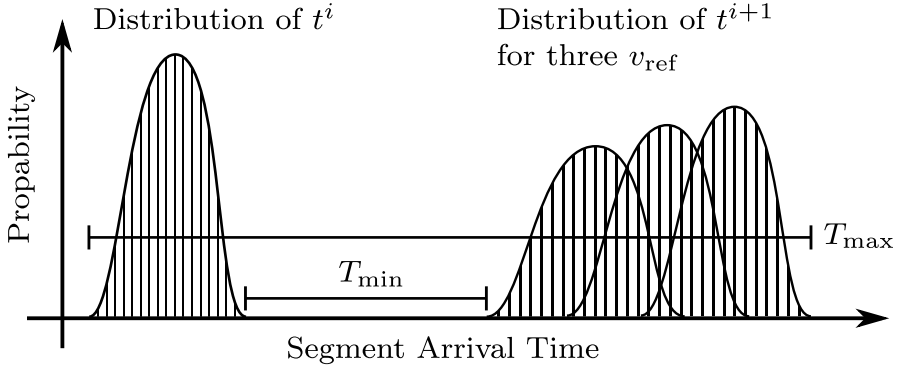


Figure 5.2: Illustration of how the distribution of t^i and t^{i+1} are related for different reference speeds v_{ref} .

explicitly as a basis for the decision whether the two vehicles should attempt to platoon or not. The probability of platooning conditioned on the arrival time of the coordination follower at the merge point $t_f^{N_f} = \mathbf{t}$ is

$$P_{\text{pl}}(\mathbf{t}) := \mathbb{P}(|t_\ell^{N_\ell} - t_f^{N_f}| < \Delta t \mid t_f^{N_f} = \mathbf{t}) = \sum_{\tau = \mathbf{t} - \Delta t}^{\mathbf{t} + \Delta t} p_{t_\ell^{N_\ell}}(\tau).$$

We want to compute policies $v_{\text{ref}}^i : \mathbb{Z} \rightarrow \mathbb{V}$, $i = 1, \dots, N_f - 1$ for selecting reference speeds for the coordination follower. These policies are to maximize the expected value of P_{pl} with respect to the probability distribution of the coordination follower's arrival time at the merge point and conditioned on that the follower's start time at the first segment is t_f^S .

$$\mathbb{E}_{t_f^{N_f}}(P_{\text{pl}} \mid t_f^1 = t_f^S) = \sum_{\mathbf{t} = -\infty}^{\infty} p_{t_f^{N_f}}(\mathbf{t} \mid t_f^1 = t_f^S) P_{\text{pl}}(\mathbf{t}). \quad (5.3)$$

The relation between $p_{t_f^1}$ and $p_{t_f^{N_f}}$ is given by (5.2). Equation (5.2) is also how the reference speed policies v_{ref}^i enter into (5.3).

5.2 Optimal Speed Control

In this section, we derive the optimal reference speeds using dynamic programming. At the same time, the probability of platooning $\mathbb{E}_{t_f^{N_f}}(P_{\text{pl}})$, as defined in (5.3), is computed. This information can be used to decide whether or not two vehicles should platoon. More specifically, expected fuel savings from platooning

would now be used instead of a predicted reduction in fuel consumption based on a deterministic model in the plan composition stage.

The formulated problem fits the framework of optimal stochastic programming [33], when defining the value function as

$$J^i(\mathbf{t}) = \mathbb{E}_{t_f^{N_f}}(P_{\text{pl}}|t_f^i = \mathbf{t}),$$

where $\mathbb{E}_{t_f^{N_f}}(P_{\text{pl}}|t_f^i = \mathbf{t})$ is the expected value of P_{pl} conditioned on that $t_f^i = \mathbf{t}$ and under the optimal policies v_{ref}^j , $j = i, \dots, N_f - 1$.

The value function at the final stage is accordingly

$$J^{N_f}(\mathbf{t}) = \mathbb{E}_{t_f^{N_f}}(P_{\text{pl}}|t_f^{N_f} = \mathbf{t}) = P_{\text{pl}}(\mathbf{t}), \quad (5.4)$$

and there is no stage cost.

The dynamic programming backwards recursion becomes

$$J^{i-1}(\mathbf{t}) = \max_{v_{\text{ref}}^i} \left(\sum_{\tau=-\infty}^{\infty} p_{T_f^{i-1}}(\tau|v_{\text{ref}}^i) J^i(\mathbf{t} + \tau) d\tau \right). \quad (5.5)$$

The probability of a successful merge when the coordination follower starts at time t_f^S is given by

$$J^1(t_f^S) = \mathbb{E}_{t_f^{N_f}}(P_{\text{pl}}|t_f^1 = t_f^S),$$

and the argument of the maximization in (5.5) yields the optimal policy for each stage.

Correlated Travel Time Distributions

This section describes how the previously derived method can be extended to the case where travel times are correlated between segments. Depending on the length of the segment and the number of exogenous explanatory variables T^i is conditioned on for prediction such as weather, time of the day, etc.; T^i might be dependent on segments that are geographically close.

We only consider correlation of travel times within the coordination leader's and the coordination follower's route. The travel times of the coordination leader's and the coordination follower's route are assumed to be independent. In the framework of dynamic programming, the above reasoning implies that we have to add upstream traversal times $\mathbf{T}^i = [T^{i-1}, T^{i-2}, \dots, T^{i-H}]$ with horizon length $H \geq 1$ to the state, which previously consisted only of the segment arrival time t , in order to retain the Markov property. The horizon length H depends on the probabilistic model of the travel times. Equation (5.1) is augmented to:

$$\begin{bmatrix} t^{i+1} \\ T^{(i-1)+1} \\ \vdots \\ T^{(i-H)+1} \end{bmatrix} = \begin{bmatrix} t^i + T^i \\ T^i \\ \vdots \\ T^{i-(H+1)} \end{bmatrix}.$$

A state $[\mathbf{t}, \mathbf{T}^{-1}, \mathbf{T}^{-2}, \dots, \mathbf{T}^{-H}]$ can be reached from $[\mathbf{t} - \mathbf{T}^{-1}, \mathbf{T}^{-2}, \mathbf{T}^{-3}, \dots, \mathbf{T}^{-(H-1)}, \tau]$ for any $\tau \in \mathbb{Z}$ with probability $p_{T^i}(\mathbf{T}^{-1} | \mathbf{T}^{-2}, \mathbf{T}^{-3}, \dots, \mathbf{T}^{-(H-1)}, \tau)$, so that (5.2) becomes

$$p_{(t^{i+1}, \mathbf{T}^{i+1})}(\mathbf{t}, \mathbf{T}) = \sum_{\tau=\mathbf{T}_{\min}^{i-H}}^{\mathbf{T}_{\max}^{i-H}} p_{T^i}(\mathbf{T}^{-1} | v_{\text{ref}}^i, \mathbf{T}^i = \mathbf{T}^{-}(\tau)) p_{(t^i, \mathbf{T}^i)}(\mathbf{t} - \mathbf{T}^{-1}, \mathbf{T}^{-}(\tau)),$$

where $p_{(t^i, \mathbf{T}^i)}$ denotes the joint probability distribution of t^i and \mathbf{T} , and where $\mathbf{T} = [\mathbf{T}^{-1}, \dots, \mathbf{T}^{-H}]$ and $\mathbf{T}^{-}(\tau) = [\mathbf{T}^{-2}, \dots, \mathbf{T}^{-(H-1)}, \tau]$.

Note the difference in notation between the random variable \mathbf{T}^i and a concrete value \mathbf{T} that \mathbf{T}^i can take. The terminal value is similar to (5.4)

$$J^{N_f}(\mathbf{t}, \mathbf{T}) = P_{\text{pl}}(\mathbf{t}),$$

where $\mathbf{T} \in \mathbb{Z}^H$ and where the distribution $p_{t_\ell^{N_\ell}}$ is computed by marginalizing the traversal time states.

The backwards recursion as in (5.5) changes to

$$J^{i-1}(\mathbf{t}, \mathbf{T}) = \max_{v_{\text{ref}}^i} \left(\sum_{\tau=-\infty}^{\infty} p_{T_f^{i-1}}(\tau | v_{\text{ref}}^i, \mathbf{T}) J^i(\mathbf{t} + \tau, \mathbf{T}^+(\tau)) d\tau \right),$$

where $\mathbf{T}^+(\tau) = [\tau, \mathbf{T}^{-1}, \dots, \mathbf{T}^{-(H-1)}]$.

While it is straightforward to keep a record of previous segment traversal times, measuring traversal times of segments before the start is not trivial. If we assume that the computation of the platooning probability and the policies happen shortly before the vehicles start driving, real-time information from other sources such as traffic sensors or other vehicles might be leveraged. Otherwise, travel time distributions that are not conditioned on segments before the start of the route have to be used.

It is well known that the complexity of dynamic programming increases exponentially with the size of the state, an effect known as the curse of dimensionality. However, if we assume that most correlation between segment traversal time is actually caused by traffic dynamics [221], we can reduce the state space size [178]. According to macroscopic traffic flow theory, traffic has mainly three states [126]: free flow, synchronized flow, and congested flow. Therefore, the elements of \mathbf{T}^i could potentially be discretized into these three regimes.

Efficient Computation of Optimal Control Policies

A challenge in using dynamic programming is finding ways of implementing the recursion and handling its complexity. There are two features making the problem considered in this paper computationally tractable. The first is the finite horizon of the problem and the second is the low dimensionality of the state space. Additionally, we can exploit the fact that unlike many other control systems, the objective of simultaneous arrival at the merge point does not have to be achieved at all cost. In case the merge fails, the problem can be resolved on the higher planning layer. Because of this property, it is reasonable to omit exploring state trajectories that lead to a successful merge but have low probability.

We show that J^i only has to be computed for an interval $[\underline{t}^i, \bar{t}^i]$ if a small error $\epsilon \geq 0$ on the computation of J can be accepted. Furthermore, the length of the interval, i.e., $\bar{t}^i - \underline{t}^i$ does not depend on the stage i . We define \underline{t}^i as

$$\underline{t}^{i+1} = \underline{t}^i + T_{\min}^i \Rightarrow \underline{t}^i = \underline{t}_f^1 + \sum_{j=1}^{i-1} T_{\min}^j,$$

with $\underline{t}^1 = \underline{t}_f^S$, with \underline{t}_f^S being the start time of the coordination follower. Similarly, we define \bar{t}^i as

$$\bar{t}^i = \bar{t}^{i+1} - T_{\min}^i \Rightarrow \bar{t}^i = \bar{t}^{N_f} - \sum_{j=i}^{N_f-1} T_{\min}^j,$$

where \bar{t}^{N_f} is selected large enough so that

$$J^{N_f}(t) \leq \epsilon \text{ for } t > \bar{t}^{N_f}, \quad (5.6)$$

for a given error tolerance $\epsilon \geq 0$.

Furthermore, we define an approximation of J^i denoted as \tilde{J}^i . It is initialized at $i = N_f$ by

$$\tilde{J}^{N_f}(\mathbf{t}) = \begin{cases} J^{N_f}(t) & \text{if } \mathbf{t} \in [\underline{t}^{N_f}, \bar{t}^{N_f}] \\ 0 & \text{if } \mathbf{t} \notin [\underline{t}^{N_f}, \bar{t}^{N_f}], \end{cases} \quad (5.7)$$

and analogously to (5.5) for $\mathbf{t} \in [\underline{t}^{i-1}, \bar{t}^{i-1}]$

$$\tilde{J}^{i-1}(\mathbf{t}) = \max_{v_{\text{ref}}^i} \left(\sum_{\tau=\underline{t}^i}^{\bar{t}^i} p_{T_f^{i-1}}(\tau - \mathbf{t} | v_{\text{ref}}^i) \tilde{J}^i(\tau) \right)$$

where the summation is rewritten in terms of the arrival time rather than the traversal time.

Note that segment arrival times $t^i \leq \underline{t}^i$ cannot be reached from $t_f^1 = t_f^S$, and therefore J^i and likewise \tilde{J}^i do not need to be computed for these times. Furthermore, $\tilde{J}^i(\mathbf{t}) = 0$ for $i = 1, \dots, N_f$ and $\mathbf{t} > \bar{t}^i$. The following result on the error between J^i and \tilde{J}^i holds

Proposition 2. *For given error tolerance $\epsilon \geq 0$ it holds that*

$$0 \leq J^i(\mathbf{t}) - \tilde{J}^i(\mathbf{t}) \leq \epsilon,$$

for all $i = 1, \dots, N_f$ and $\mathbf{t} \geq \underline{t}^i$.

Proof. The proof is done through induction over i . From the definition of \bar{t}^{N_f} in (5.6) and the definition of $\tilde{J}^{N_f}(\mathbf{t})$ in (5.7) it can be seen that the statement holds for $i = N_f$. Note that $J^i(\mathbf{t})$ is a probability and hence non-negative.

Assume that $0 \leq J^i(\mathbf{t}) - \tilde{J}^i(\mathbf{t}) \leq \epsilon$ holds. Then for $\mathbf{t} \geq \underline{t}^{i-1}$

$$\begin{aligned} & J^{i-1}(\mathbf{t}) - \tilde{J}^{i-1}(\mathbf{t}) \\ &= \max_{v_{\text{ref}}^i} \left(\sum_{\tau=-\infty}^{\infty} p_{T_f^i}(\tau | v_{\text{ref}}^i) J^i(\mathbf{t} + \tau) \right) - \max_{v_{\text{ref}}^i} \left(\sum_{\tau=\underline{t}^i}^{\bar{t}^i} p_{T_f^i}(\tau - \mathbf{t} | v_{\text{ref}}^i) \tilde{J}^i(\tau) \right) \\ &= \sum_{\tau=-\infty}^{\infty} p_{T_f^i}(\tau | v_1) J^i(\mathbf{t} + \tau) - \sum_{\tau=\underline{t}^i}^{\bar{t}^i} p_{T_f^i}(\tau - \mathbf{t} | v_2) \tilde{J}^i(\tau), \end{aligned}$$

where v_1 and v_2 are the optimal values of the two maximizations over v_{ref}^i . Since $p_{T_f^i}(\tau | v_{\text{ref}}^i) = 0$ for $\tau \leq T_{\text{min}}^i$, we can rewrite the above expression as

$$\begin{aligned} & \sum_{\tau=\underline{t}^i}^{\infty} p_{T_f^i}(\tau - \mathbf{t} | v_1) J^i(\tau) - \sum_{\tau=\underline{t}^i}^{\bar{t}^i} p_{T_f^i}(\tau - \mathbf{t} | v_2) \tilde{J}^i(\tau) \\ & \leq \sum_{\tau=\underline{t}^i}^{\infty} p_{T_f^i}(\tau - \mathbf{t} | v_1) J^i(\tau) - \sum_{\tau=\underline{t}^i}^{\bar{t}^i} p_{T_f^i}(\tau - \mathbf{t} | v_1) \tilde{J}^i(\tau) \\ & = \sum_{\tau=\underline{t}^i}^{\infty} p_{T_f^i}(\tau - \mathbf{t} | v_1) \left(J^i(\tau) - \tilde{J}^i(\tau) \right) \\ & \leq \sum_{\tau=T_{\text{min}}^{i-1}}^{\infty} p_{T_f^i}(\tau | v_1) \epsilon = \epsilon, \end{aligned}$$

where the first equality holds since v_2 maximizes the second sum.

Furthermore

$$\begin{aligned}
J^{i-1}(t) &= \sum_{\tau=\underline{t}^i}^{\infty} p_{T_{\bar{t}}^i}(\tau - \mathbf{t}|v_1)J^i(\tau) \\
&\geq \sum_{\tau=\underline{t}^i}^{\infty} p_{T_{\bar{t}}^i}(\tau - \mathbf{t}|v_2)J^i(\tau) \\
&= \sum_{\tau=\underline{t}^i}^{\infty} p_{T_{\bar{t}}^i}(\tau - \mathbf{t}|v_2) \left(\tilde{J}^i(\tau) + J^i(\tau) - \tilde{J}^i(\tau) \right) \\
&= \sum_{\tau=\underline{t}^i}^{\infty} p_{T_{\bar{t}}^i}(\tau - \mathbf{t}|v_2) \left(\tilde{J}^i(\tau) \right) + \sum_{\tau=\underline{t}^i}^{\infty} p_{T_{\bar{t}}^i}(\tau - \mathbf{t}|v_2) \left(J^i(\tau) - \tilde{J}^i(\tau) \right) \\
&\geq \sum_{\tau=\underline{t}^i}^{\bar{t}^i} p_{T_{\bar{t}}^i}(\tau - \mathbf{t}|v_2)\tilde{J}^i(\tau) = \tilde{J}^{i-1}(\mathbf{t}) \\
&\Leftrightarrow J^{i-1}(\mathbf{t}) - \tilde{J}^{i-1}(\mathbf{t}) \geq 0,
\end{aligned}$$

which concludes the proof. \square

This proposition states that we underestimate the probability of platooning by at most ϵ when using \tilde{J}^i instead of J^i . Choosing ϵ large means that $\bar{t}^i - \underline{t}^i$ is small which translates into small computational complexity but larger errors on the computation of J^i and vice versa, since \tilde{J}^i only needs to be computed in the interval $[\underline{t}^i, \bar{t}^i]$. For $\mathbf{t} > \bar{t}^i$, the chance of platooning with the coordination leader is smaller than ϵ . Once the vehicle reaches a segment later than \bar{t}^i , it would no longer try to platoon with this coordination leader and instead either try to join another coordination leader or drive alone. The computational complexity can potentially be significantly reduced by this approach depending on the distribution of $T_{\bar{t}}^{N\epsilon}$. In practice, we would expect that $T_{\bar{t}}^{N\epsilon}$ with high variance also leads to platooning probabilities so small that they can in any case be discarded regardless of the follower's arrival time at the merge point.

5.3 Modeling Controlled Travel-Time Distributions

A prerequisite for using the results derived in this paper is a good model of how traversal times are distributed, i.e., we need to be able to compute the probability distributions of the traversal times conditioned on the reference speed $p_{T^i}(\cdot|v_{\text{ref}}^i)$. In the context of commercial heavy vehicles, it is common that they report their position at regular intervals, mainly for fleet management purposes. A number of methods to model travel time distributions as functions of exogenously measurable factors such as time of the day or precipitation from such data has been reported in literature as discussed in the introduction. We introduce the reference speed as

an additional exogenous factor that is not present in the current road transport system.

Once coordinated platoon systems are widely deployed, obtaining travel time data with known reference speed will not be a technical challenge. However, in the initial phase of low penetration, it will be beneficial to use available data from uncontrolled vehicles to estimate the travel time distributions. In this case, the influence of control has to be modeled. We propose to pre-process the travel time data for different control values assuming that vehicles in regular traffic typically follow a speed profile that is determined by the maximum speed of the vehicle, legal restrictions, and external factors such as traffic and weather. Next, these data are used for modeling travel time distributions as if the data were obtained from measurements.

A practical way to control the vehicle speed is to set a reference speed on the adaptive cruise control system (ACC) that is present in many modern vehicles and that is a prerequisite for platooning. This system tracks a reference speed but ensures a safe gap from preceding vehicles. Assume we have a speed trajectory $v_{\text{free}}^i(x)$ on segment i for a vehicle driving a maximum possible speed as a function of the distance x traveled along the segment. So, the speed of the vehicle v^i will be the minimum of the set speed v_{ref}^i and speed profile $v_{\text{free}}^i(x)$ that the vehicle would drive without cruise control, i.e.,

$$v^i(x, v_{\text{ref}}^i) = \min(v_{\text{free}}^i(x), v_{\text{ref}}^i).$$

See Figure 5.3 for an illustration. Assume trajectories of $v_{\text{free}}^i(x)$ are have been measured. By computing the point-wise minimum, the controlled trajectories can be computed for a finite number of values of v_{ref}^i . This procedure generates samples of traversal times, which are used for training of the models that give the probability distributions of T^i denoted $p_{T^i}(\cdot | v_{\text{ref}}^i)$ using existing methods for modeling travel time distributions.

More elaborate models for pre-processing the trajectory data can also be considered. It is, for instance, possible to simulate a controller that tracks a specified traversal time but where the maximum speed of the vehicle is limited by v_{free} . In that case, unusually slow speeds at the beginning of the segment can be compensated by higher speeds in the end but not the other way round. With this kind of approach it is also possible to incorporate limits on the vehicles acceleration.

5.4 Simulations

Simulations are presented in this section in order to demonstrate the applicability of the derived results. The model published in [239] is adapted for modeling the traversal time distributions. In [239], speed distributions are modeled as the mixture of Gaussian distributions, i.e., $p_V(v) = W\mathcal{N}(v, \mu_1, \sigma_1) + (1 - W)\mathcal{N}(v, \mu_2, \sigma_2)$. We interpret one mode as corresponding to the free-flow and one to the congestion regime between which transition often happens suddenly ([126]).




Figure 5.3: The controlled speed profile $v^i(x, v_{\text{ref}}^i)$ is the pointwise minimum of the measured maximum speed profile $v_{\text{free}}^i(x)$ and the reference speed v_{ref}^i . The dashed curves show the controlled speed profile for different reference speeds v_{ref}^i , the red dashed line highlighting one of them for better readability.

We model the effect of control by setting the mean value of the free flow model μ_2 to v_{ref}^i , and we consider reference speeds in the range $v_{\text{ref}}^i \in [70 \text{ km/h}, 90 \text{ km/h}]$ with increments of 1 km/h. Furthermore, both Gaussian kernels in $p_V(v)$ are truncated individually to the range of $[10 \text{ km/h}, 100 \text{ km/h}]$. The measured speed value distributions in [239] contain some entries well above the heavy-duty vehicle speed limit of $60 \text{ mph} \approx 96.56 \text{ km/h}$. Limiting the maximum speed is justified considering that a centralized planning system would not recommend speeds above the legal speed limits. The speed is assumed to be constant over a segment, i.e., we have that the probability $P(T^i \leq \tau) = P(V \geq L^i/\tau)$, where L^i is the length of the i -th segment. The remaining parameters of the speed distribution are listed in Table 5.1. The first set of parameters corresponds to a reliable segment with high speeds and little variation in the speed. The other set of parameters corresponds to an unreliable segment with a high risk of small speeds due to congestion and a wide spread in possible speeds.

First, the following scenario is considered. The coordination leader's route consists of three segments with length 4 km, 4 km, 5 km respectively, and the leaders start at time $t_\ell^1 = 0$. The reference speed of the leader is 80 km/h. The coordination follower's route also consists of three segments with lengths 6, 4, 5 km, and the start time is computed so that the coordination leader and follower would meet if they kept a constant speed of 80 km/h. All segments are considered to be of

Table 5.1: Parameters of the speed distributions

Variable	Reliable	Unreliable
W	0.04	0.55
μ_1	64.45 km/h	38.64 km/h
σ_1	34.76 km/h	18.96 km/h
σ_2	8.22 km/h	9.96 km/h

the reliable type. The maximum tolerable error ϵ between J and \tilde{J} is set to 1%, and the maximum time-gap for platooning Δt is 0.01 h = 36 seconds. A time step corresponds 10^{-4} h = 0.36 s. The simulation was implemented using CPython 2.7 with Numpy and Scipy and this example takes less than 50 milliseconds to compute on a Core i3 processor using only one core.

Figs. 5.4 and 5.5 show the results from this scenario. Figure 5.4 shows the computed distributions of T_ℓ^i . We can see that the distribution of T_ℓ^i spreads out from segment to segment. Figure 5.5 shows the computed function \tilde{J}^i with and without optimal control as well as a visualization of the optimal control policies. The optimal control is able to significantly improve the probability of platooning by centering the arrival time distributions of the next segment at arrival times with high values for \tilde{J}^i . As \tilde{J}^i gets more spread out to the left, the transition from slow reference speeds to the highest reference speeds with increasing segment arrival times also becomes more spread out. The start time has been chosen here in a way that the two vehicles can easily meet in their reference speed range. This means that being unable to arrive sufficiently late is no issue and nothing could be gained from starting later. It is also possible to see how the optimal control is able to compensate if the coordination follower deviates from the trajectory that would be obtained by driving constantly at 80 km/h. The merge probability $\tilde{J}^0(t_f^0)$ equals 52.96% using the optimal control and 43.97% using the fixed reference speed. The interval $\bar{t}_i - \underline{t}_i$ is 11.51 times smaller than with $\epsilon = 0$, while the actual error on the platooning probability according to (5.3) from using \tilde{J}^i instead of J^i is 0.03% $\ll \epsilon = 1\%$.

Figure 5.6 shows \tilde{J}^i for a similar scenario as described above with the difference that the second segment in the coordination follower's route is unreliable. We can see that this causes a high risk of delay and thus much smaller values for \tilde{J}^i . Furthermore, the control policy selects higher reference speeds on the first segment compared to the previous scenario without an unreliable segment in order to compensate for a potential delay on the second segment. The merge probability $\tilde{J}^0(t_f^0)$ equals 32.62% using the optimal control and 24.62% using the fixed reference speed.

The two speed distributions are extreme cases of reliable and unreliable segments. In reality, there is a whole range of characteristics between these two extremes. Furthermore, it is unclear how a controlled vehicle would behave with re-




Figure 5.4: This plot shows the segment arrival time probability density functions of the coordination leader. At the beginning of the first segment, the start time is known and indicated by an arrow. The remaining densities are plotted on the vertical axis and are jointly scaled for presentation. The vertical dotted lines correspond to the maximum value of the first distribution and provide a reference for comparison of the distributions. The dashed line corresponds to the maximum speed of 100 km/h. At the end of the last segment, which is the merge point, also J^{N_f} is plotted in green. It is scaled so that a value of 1 corresponds to the level of the dotted line. The scale of $2\Delta t$ is indicated on the right side of the plot.

spect to traversal time distributions, but we can assume that the variability would probably be smaller. A vehicle that is controlled to follow a reference speed behaves more predictably as the control reduces the variability due to different driver characteristics. In addition, we have to take into account that the spot speed can vary more than the traversal time. Consider, for instance, a stop-and-go situation with regular shock-waves traveling upstream. In this case, vehicles will exhibit a large variability in the speed which will be averaged out over a longer distance. Nevertheless, the simulations demonstrate that the method can handle realistically sized instances of the problem and smaller variability would lead to even smaller




Figure 5.5: This plot shows \tilde{J}^i at the distances from the coordination follower's start point corresponding to $i = 1, \dots, N_f$ as a function of segment arrival time. The vertical dotted lines correspond to a level of 1 on the scale of \tilde{J}^i . The blue plots show \tilde{J}^i when the optimal control policy is implemented and the green curves when the reference speed is kept constant. The gray semi-transparent triangles visualize the control policy. The two rightmost corners of the triangle correspond to the 5- and 95 percentiles of the arrival times at a segment under the optimal control policy conditioned on that the start time at the previous segment corresponds to the leftmost corner of the triangle. The dashed lines correspond to \underline{t}^i and \bar{t}^i .

intervals of $\bar{t}^i - \underline{t}^i$ and thus faster computation times. Note also that these kind of computations lend themselves well to parallel processing, for instance, on graphics cards.

5.5 Summary

Travel times can be hard to predict accurately which can lead to failed platooning attempts decreasing the overall system performance. To address the issue, a




Figure 5.6: This plot shows, similar to Figure 5.5, \tilde{J}^i at the distances from the coordination follower's start point corresponding to $i = 1, \dots, N_f$ as function of segment arrival time, but for a scenario where the second segment is unreliable.

stochastic model is formulated in this chapter. In order to counteract the disturbances to the planned trajectory and to avoid overly pessimistic predictions, feedback control is introduced. In this model, control enters the system by affecting the traversal time distribution on segments leading towards the meeting point.

The control problem is solved by means of dynamic programming which also yields the meeting probability explicitly as an input to higher planning layers. Considering the maximum speed of a vehicle and allowing for a small error in computing the value function makes it possible to reduce the computational effort. The reduction is achieved by bounding the state-space in which solutions have to be computed. Simulations demonstrate the effectiveness of this approach.

Chapter 6

Coordination Algorithms for Many Vehicles

THIS chapter presents a systematic way of combining the pairwise platoon plans derived in the previous chapters assigning a plan to each vehicle. The problem of how to combine such plans into a fuel-efficient plan for all vehicles is expressed as a combinatorial optimization problem. Section 6.2 deals with the computation of exact solutions to this problem. The problem is formulated as an integer programming problem, which allows to use general purpose solvers to obtain a solution. Furthermore, we establish results on the solution structure of the problem, which can restrict the number of solutions to be explored and gives insight into the problem. Finally, the problem is proven to be NP-hard, which is commonly believed to imply that finding a global minimizer can be computationally expensive for general problem instances. This motivates developing the algorithm for computing heuristic solutions presented in Section 6.3. This algorithm can find sub-optimal solutions efficiently but it is not guaranteed to converge to a globally optimal combination of pairwise plans. Similar approaches are often used when dealing with NP-hard problems. Once the pairwise plans are combined into a platoon plan for all vehicles, it is possible to adjust the timing when platoons are formed and split, not altering where platoons are to be formed. Section 6.4 discusses how to do such adjustments in a way that minimizes fuel consumption.

6.1 Combining Pairwise Plans by Selecting Coordination Leaders

In Chapter 4 and 5, we introduce planning methods for two vehicles. The speed profile of one vehicle is fixed according to the default plan and the speed profile of the other vehicle is adjusted so the two vehicles platoon for some distance. The latter platoon plan is referred to as the adapted plan as introduced in Section 4.1. Due to the structure of these plans, it is possible to select one vehicle, the coordination

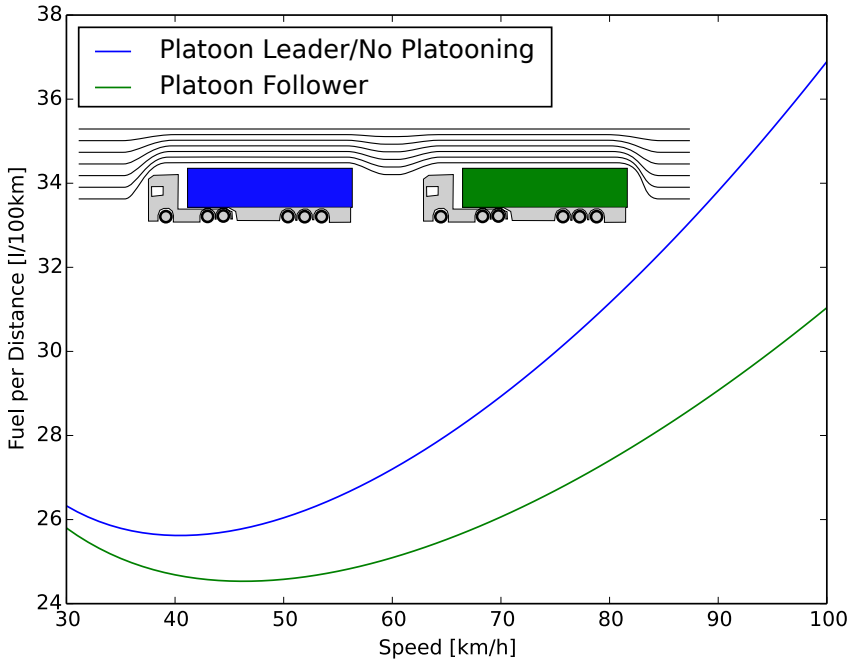


Figure 6.1: The fuel consumption per distance traveled as a function of speed with (green vehicle) and without (blue vehicle) the slipstream effect that lowers fuel consumption when platooning.

leader, to follow the speed profile of the default platoon plan and have more than one vehicle select the adapted plan, adapted to the coordination leader. Like this, platoons with more than two vehicles can be formed.

In order to guide the selection of adapted plans, the resulting fuel consumption needs to be estimated. Hereby, when considering an adapted plan of a vehicle, its fuel consumption is compared to the fuel consumption of its default plan and not platooning. We only need to model the difference in fuel consumption between default and adapted plan. The absolute fuel consumption is not needed for the presented method. The two factors that change the fuel consumption of the adapted plan compared to the default plan is the speed profile and that the two vehicles platoon for some distance.

Speed influences the fuel consumption in multiple ways as shown in Figure 6.1. When traveling at low speeds, smaller gears have to be selected causing increased friction losses per distance traveled and the correlation between speed and fuel consumption is negative. At high speeds, the quadratic increase of air-drag leads to a positive correlation between vehicle speed and fuel consumption.

The reduction of fuel consumption due to platooning is caused by a reduction in air-drag due to the small inter-vehicle distances. The dominant effect is a drag-reduction for the trailing vehicle. This effect is called slipstreaming. At very small distances, also an effect on the lead vehicle can be observed. This effect is smaller than the slipstreaming effect. These effects have been reported in a large number of studies using computational fluid dynamics, wind-tunnel experiments, and experiments with real vehicles. The results consistently show a reduction of fuel consumption for the trailing vehicles in the order of 10% due to slipstreaming [181].

In order to get a tractable solution to the planning problem, we estimate the fuel consumption reduction resulting from an adapted plan (e.g. Definition 4) of vehicle n , adapted to the default plan (e.g. Definition 3) of vehicle m , denoted $\Delta F(n, m)$. This assumes that only n and m platoon according to the adapted plan. Since the coordination leader's speed profile is not changed due to the adaptation, fuel consumption due to speed is only changed for the coordination follower. Furthermore, we assume that the effect of fuel saved during platooning due to the coordination follower is independent of the other vehicles in the platoon. Both physical models, data-driven models, or combinations of the two can be used to estimate ΔF .

Remark 21. It is also possible to consider other aspects than fuel in the cost function such as the trip duration or the risk of being delayed (see Chapter 5). It is also possible to add a fixed cost to join a platoon to a model a merge phase.

We now compute ΔF for all ordered pairs in \mathcal{N}_c . We are only interested in adapted plans that save fuel, i.e., for which ΔF is positive. We collect this information in a weighted graph that we call the coordination graph.

Definition 5 (Coordination Graph). The coordination graph is a weighted directed graph $\mathcal{G}_c = (\mathcal{N}_c, \mathcal{E}_c, \Delta F)$. Recall that \mathcal{N}_c represents the vehicles. $\mathcal{E}_c \subseteq \mathcal{N}_c \times \mathcal{N}_c$ is a set of edges, and $\Delta F : \mathcal{E}_c \rightarrow \mathbb{R}^+$ are edge weights, such that there is an edge $(n, m) \in \mathcal{E}_c$, if the adapted plan of n to m saves fuel compared to n 's default plan, i.e., $\mathcal{E}_c = \{(n, m) \in \mathcal{N}_c \times \mathcal{N}_c : \Delta F(n, m) > 0, n \neq m\}$.

Furthermore, we introduce the set of in-neighbors of a node $n \in \mathcal{N}_c$ as

$$\mathcal{N}_n^i = \{m \in \mathcal{N}_c : (m, n) \in \mathcal{E}_c\}$$

and the set of out-neighbors of n as

$$\mathcal{N}_n^o = \{m \in \mathcal{N}_c : (n, m) \in \mathcal{E}_c\}.$$

We define the maximum over an empty set to be zero, i.e., $\max_{n \in \emptyset}(\cdot) = 0$.

With these definitions, we are ready to formulate the problem of finding a fuel optimal set of coordination leaders \mathcal{N}_1 .

Problem 1. Given a coordination graph $\mathcal{G}_c = (\mathcal{N}_c, \mathcal{E}_c, \Delta F)$, find a subset $\mathcal{N}_1 \subset \mathcal{N}_c$ of nodes that maximizes

$$f_{ce}(\mathcal{N}_1) = \sum_{n \in \mathcal{N}_c \setminus \mathcal{N}_1} \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(n, m). \quad (6.1)$$

For a given set of coordination leaders \mathcal{N}_1 , a coordination follower \bar{n} implements the adapted plan to the best coordination leader, i.e., to $\arg \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(\bar{n}, m)$. If $(\bar{n}, \bar{m}) \in \mathcal{E}_c$ with $\bar{n} \in \mathcal{N}_c \setminus \mathcal{N}_1$ and $\bar{m} = \arg \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(\bar{n}, m)$, we say that \bar{n} is the coordination follower of \bar{m} and that \bar{m} is the coordination leader of \bar{n} . If there is no coordination leader for a coordination follower \bar{n} , then $\mathcal{N}_n^o \cap \mathcal{N}_1 = \emptyset$ and $\max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(\bar{n}, m) = \max_{m \in \emptyset} \Delta F(\bar{n}, m) = 0$. In this case the coordination follower does not platoon and implements its default speed profile. One coordination leader can have multiple coordination followers leading to platoons with more than two vehicles. The exact order of the platoon is left to the on-board systems to coordinate locally.

At this point, we have a combinatorial problem, whose solution allows us to group transport assignments in a fuel-efficient way. All continuous optimization is contained in the adapted plans. Since an adapted plan only involves computing the speed profile for one vehicle, deriving such adapted plans is a task that is feasible as demonstrated in Chapter 4.

The simplifying structure imposed by adapted plans comes with a price on the fuel-savings that can be achieved. In Section 6.4, we address this problem to some extent by jointly optimizing the speed profile of each cluster. Furthermore as introduced in Chapter 3, the platoon coordinator executes the optimization procedure repeatedly. Each time the coordination graph is updated based on the vehicle positions, platoon states, and new assignments. A coordination follower that joins a platoon during the first part of its journey can become coordination follower of another vehicle at a later point in time and drive in a platoon for the remaining part of its journey.

Remark 22. Realistically, there is a limit on the size of a platoon. This can be either handled by the platoon manager splitting up large platoons into several smaller platoons or by putting an additional constraint on the number of coordination followers of a coordination leader. Note however, that the platoon size can be smaller than the number of coordination followers of a coordination leader since different coordination followers can follow the same coordination leader at different parts of its route.

Remark 23. This approach is inspired by clustering algorithms. Clustering is a widely used tool for analysis of large data sets. Data are structured into a finite number of sets. Elements within the set are in some way related. K-means clustering is a popular technique in machine learning. An algorithm related to K-

means clustering is called K-medoids clustering [116, 117, 122], which inspired the formulation of Problem 1.

Remark 24. Other types of adapted plans than the ones presented in Chapter 4 can be handled in the framework, for instance, adapting the route of the coordination follower or including stops along the route.

6.2 Exact Solution to the Coordination Leader Selection Problem

In this section, we study the problem of selecting an optimal set of coordination leaders, i.e., solving Problem 1. Formulating the problem as an integer program let us employ general purpose solvers to compute optimal solutions. We also establish that Problem 1 is NP-hard.

Integer Programming Formulation

Problem 1 is a combinatorial optimization problem, which we can formulate as an integer programming problem. This makes it possible to use advanced general purpose solvers which can find both exact and heuristic solutions.

We do this by introducing a binary variable $x(e) \in \{0, 1\}$ for each edge $e \in \mathcal{E}_c$ in the coordination graph. This variable equals 1, if the corresponding edge contributes to the objective f_{ce} of Problem 1, i.e., if the corresponding pair-wise plan is selected. Otherwise $x(e) = 0$. A constraint for each node n in the coordination graph \mathcal{G}_c ensures that either exactly one outgoing edge is selected, which implies that n is a coordination leader, i.e., $n \in \mathcal{N}_c \setminus \mathcal{N}_1$. Otherwise an arbitrary number of incoming edges can be selected, which implies that n is a coordination follower, i.e., $n \in \mathcal{N}_1$. For notational convenience, we write $x(n, m)$ for $x((n, m))$ where $(n, m) \in \mathcal{E}_c$.

Proposition 3. *Given Problem 1 with coordination graph $\mathcal{G}_c = (\mathcal{N}_c, \mathcal{E}_c, \Delta F)$. The optimal solution $\{x(e) : e \in \mathcal{E}_c\}$ of the integer program*

$$\max_{\{x(e) : e \in \mathcal{E}_c\}} \sum_{e \in \mathcal{E}_c} \Delta F(e)x(e) \quad (6.2a)$$

s. t.

$$\sum_{m \in \mathcal{N}_n^o} x(n, m) + \frac{1}{|\mathcal{N}_n^i| + 1} \sum_{m \in \mathcal{N}_n^i} x(m, n) \leq 1 \quad \text{for } n \in \mathcal{N}_c \quad (6.2b)$$

$$x(e) \in \{0, 1\} \quad \text{for } e \in \mathcal{E}_c. \quad (6.2c)$$

is an optimal solution \mathcal{N}_1 to Problem 1, with

$$\mathcal{N}_1 = \{n \in \mathcal{N}_c : \exists(m, n) \in \mathcal{E}_c, x(m, n) = 1\}. \quad (6.3)$$

Proof. We denote the objective of the integer program as

$$f_{\text{IP}}(x) = \sum_{e \in \mathcal{E}_c} \Delta F(e)x(e),$$

and the mapping from solutions of the integer program to solutions of Problem 1 as

$$N_{\text{IP}}(x) = \{n \in \mathcal{N}_c : \exists(m, n) \in \mathcal{E}_c, x(m, n) = 1\}.$$

We define a mapping from solutions $\bar{\mathcal{N}}_1$ of Problem 1 to solutions of the integer program as

$$X_{\text{CL}}(\bar{\mathcal{N}}_1, n, m) = \begin{cases} 1 & \text{if } n \in \mathcal{N}_c \setminus \bar{\mathcal{N}}_1 \wedge m = \arg \max_{m \in \mathcal{N}_n^o \cap \bar{\mathcal{N}}_1} \Delta F(n, m) \\ 0 & \text{otherwise} \end{cases},$$

(6.4)

for some appropriate arg max operator in case the maximizer is not uniquely defined. We write $X_{\text{CL}}(\bar{\mathcal{N}}_1)$ to refer to the mapping $X_{\text{CL}}(\bar{\mathcal{N}}_1, n, m)$ for all $n, m \in \mathcal{N}_c \times \mathcal{N}_c$. We have from the definition of X_{CL} that

$$\begin{aligned} f_{\text{ce}}(\bar{\mathcal{N}}_1) &= \sum_{n \in \mathcal{N}_c \setminus \bar{\mathcal{N}}_1} \max_{m \in \mathcal{N}_n^o \cap \bar{\mathcal{N}}_1} \Delta F(n, m) \\ &= \sum_{(n, m) \in \mathcal{E}_c} \Delta F(n, m) X_{\text{CL}}(\bar{\mathcal{N}}_1, n, m) \\ &= f_{\text{IP}}(X_{\text{CL}}(\bar{\mathcal{N}}_1)). \end{aligned}$$

We show now that all solutions X_{CL} to the integer program are feasible. Consider (6.2b) for any $n \in \bar{\mathcal{N}}_1$, i.e., $n \notin \mathcal{N}_c \setminus \bar{\mathcal{N}}_1$. We denote $X_{\text{CL}}(\bar{\mathcal{N}}_1, n, m)$ as $\bar{x}(n, m)$ for readability. Then the sum

$$\sum_{m \in \mathcal{N}_n^o} \bar{x}(n, m) = 0,$$

and since $\bar{x}(n, m) \in \{0, 1\}$

$$\frac{1}{|\mathcal{N}_n^i| + 1} \sum_{m \in \mathcal{N}_n^i} \bar{x}(m, n) \leq \frac{|\mathcal{N}_n^i|}{|\mathcal{N}_n^i| + 1} \leq 1,$$

which implies that constraint (6.2b) is fulfilled.

Consider (6.2b) for any $n \notin \bar{\mathcal{N}}_1$, i.e., $n \in \mathcal{N}_c \setminus \bar{\mathcal{N}}_1$. Then

$$\frac{1}{|\mathcal{N}_n^i| + 1} \sum_{m \in \mathcal{N}_n^i} \bar{x}(m, n) = 0,$$

because $n = \arg \max_{\bar{n} \in \mathcal{N}_m^o \cap \bar{\mathcal{N}}_1} \Delta F(m, \bar{n})$ cannot hold for any m and

$$\sum_{m \in \mathcal{N}_n^o} \bar{x}(n, m) \leq 1,$$

as there is at most one $\arg \max_{m \in \mathcal{N}_c^o \cap \mathcal{N}_1} \Delta F(n, m)$. Consider an optimal solution x^* to the integer program. We show that $x^* = X_{\text{CL}}(N_{\text{IP}}(x^*))$. Consider n, m with $x^*(n, m) = 1$. This implies that $m \in N_{\text{IP}}(x^*)$ and $n \notin N_{\text{IP}}(x^*)$ due to constraint (6.2b). Furthermore, it holds that $m = \arg \max_{\bar{m} \in N_{\text{IP}}(x^*)} \Delta F(n, \bar{m})$ for an appropriate $\arg \max$ operator, otherwise x^* cannot be optimal.

But then $X_{\text{CL}}(N_{\text{IP}}(x^*), n, m) = 1$. Consider n, m with $x^*(n, m) = 0$. This implies due to constraint (6.2b) either that $n \in N_{\text{IP}}(x^*)$ or that $m \neq \arg \max_{\bar{m} \in N_{\text{IP}}(x^*)} \Delta F(n, \bar{m})$.

But then $X_{\text{CL}}(N_{\text{IP}}(x^*), n, m) = 0$.

Assume now that there is an optimal solution x^* to the integer program for which $N_{\text{IP}}(x^*)$ is not an optimal solution to Problem 1. Let \mathcal{N}^* be an optimal solution to Problem 1. Then $X_{\text{CL}}(\mathcal{N}^*)$ is a feasible solution to the integer program with $f_{\text{IP}}(X_{\text{CL}}(\mathcal{N}^*)) = f_{\text{ce}}(\mathcal{N}^*)$. But then we have that $X_{\text{CL}}(N_{\text{IP}}(x^*)) = x^*$, and $f_{\text{IP}}(x^*) = f_{\text{ce}}(N_{\text{IP}}(x^*)) < f_{\text{ce}}(\mathcal{N}^*) = f_{\text{IP}}(X_{\text{CL}}(\mathcal{N}^*))$, which contradicts the assumption that x^* is an optimal solution to the integer program. \square

Remark 25. Constraints (6.2b) imply that

$$\sum_{m \in \mathcal{N}_c^o} x(n, m) \leq 1,$$

which constitutes a type 1 special ordered set constraint. In term of Problem 1 it corresponds to the structure that each vehicle adapts to at most one coordination leader. This structure can be used in general purpose integer programming solvers to speed up the optimization procedure.

Results on the Solution Structure

Next, we derive results on the structure of the optimal solution. These results can be used to limit the search space when searching for an optimal solution, for instance, when using a branch-and-bound algorithm. The first result on the structure of the optimal solution is an upper bound on the maximum number of coordination leaders, i.e., on the cardinality of the optimal \mathcal{N}_1 . It states that there is an optimal solution with at most $\lfloor |\mathcal{N}_c|/2 \rfloor$ coordination leaders.

Proposition 4. *There exists an optimal solution \mathcal{N}_1 to Problem 1 with $|\mathcal{N}_1| \leq \lfloor |\mathcal{N}_c|/2 \rfloor$.*

Proof. First of all, we note that Problem 1 is an unconstrained optimization problem and the optimization argument belongs to a finite set. Therefore, an optimal solution always exists.

The existence of an optimal solution \mathcal{N}_1 with $|\mathcal{N}_1| \leq \lfloor |\mathcal{N}_c|/2 \rfloor$ is proven by contradiction. Assume that every optimal solution \mathcal{N}_1 to Problem 1 fulfills $|\mathcal{N}_1| > \lfloor |\mathcal{N}_c|/2 \rfloor$. Then $|\mathcal{N}_c \setminus \mathcal{N}_1| < |\mathcal{N}_1|$. Hence, there is at least one $\bar{n} \in \mathcal{N}_1$ for which there

is no $n \in \mathcal{N}_c \setminus \mathcal{N}_1$ for which $\bar{n} = \arg \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(n, m)$. Thus, \bar{n} can be removed from \mathcal{N}_1 without decreasing $f_{ce}(\mathcal{N}_1)$, i.e., $f_{ce}(\mathcal{N}_1 \setminus \{\bar{n}\}) \geq f_{ce}(\mathcal{N}_1)$. This reasoning can be repeatedly applied until $|\mathcal{N}_1| \leq \lfloor |\mathcal{N}_c|/2 \rfloor$ with $f_{ce}(\mathcal{N}_1)$ no smaller than the optimal \mathcal{N}_1 . Thus, the smaller \mathcal{N}_1 is as well an optimal solution to Problem 1. This, however, contradicts the assumption. \square

This proposition helps when computing an optimal solution since coordination leaders sets with cardinality larger than $\lfloor |\mathcal{N}_c|/2 \rfloor$ do not have to be considered.

To add this result as a constraint to the integer programming formulation in Proposition 3, we have to introduce auxiliary variables $\{x^\ell(n) : n \in \mathcal{N}_c\}$, $x^\ell(n) \in \{0, 1\}$ that indicate if node n is a coordination leader, i.e., if $n \in \mathcal{N}_1$. This is ensured by adding the following constraints

$$x^\ell(n) \geq \frac{1}{|\mathcal{N}_n^i| + 1} \sum_{m \in \mathcal{N}_n^i} x(m, n) \quad (6.5a)$$

$$x^\ell(n) \leq 1 - \sum_{m \in \mathcal{N}_n^o} x(n, m) \quad (6.5b)$$

$$x^\ell(n) \leq \sum_{m \in \mathcal{N}_n^o} x(n, m) + \sum_{m \in \mathcal{N}_n^i} x(m, n). \quad (6.5c)$$

If node n is a coordination leader, then at least one element x in the sum $\sum_{m \in \mathcal{N}_n^i} x(m, n)$ equals 1 so that $x^\ell(n) > 0$ which implies $x^\ell(n) = 1$ since $x^\ell(n) \in \{0, 1\}$. Otherwise $\sum_{m \in \mathcal{N}_n^i} x(m, n) = 0$ and constraint (6.5a) is inactive. Conversely, when node n is a coordination follower, then $\sum_{m \in \mathcal{N}_n^o} x(n, m) = 1$ which implies $x^\ell(n) = 0$. Otherwise constraint (6.5b) is inactive. In the case $\sum_{m \in \mathcal{N}_n^o} x(n, m) + \sum_{m \in \mathcal{N}_n^i} x(m, n) = 0$, it makes no difference if $n \in \mathcal{N}_1$ so we exclude it from \mathcal{N}_1 with constraint (6.5c). Thus, the constraints (6.5a), (6.5b), and (6.5c) imply that $x^\ell(n) = 0$ if the node is a coordination follower and $x^\ell(n) = 1$ if it is a coordination leader.

Then Proposition 4 can be formulated as

$$\sum_{n \in \mathcal{N}_c} x^\ell(n) \leq \left\lfloor \frac{|\mathcal{N}_c|}{2} \right\rfloor.$$

The next result on the solution structure is that a node is either a coordination leader itself or at least one node in its two-hop out-neighbor set is a coordination leader. To this end, we define the set of two-hop out-neighbors of a node $\bar{n} \in \mathcal{N}_c$ as

$$\mathcal{N}_{\bar{n}}^{2o} = \mathcal{N}_{\bar{n}}^o \cup \bigcup_{n \in \mathcal{N}_{\bar{n}}^o} \mathcal{N}_n^o.$$




Figure 6.2: Example of a node's two-hop out-neighbor set. The gray circles that have a solid line represent the two-hop out-neighbor set of the node drawn as a circle with a dashed line. An optimal solution \mathcal{N}_1 to Problem 1 contains at least one of the gray-filled nodes.

Figure 6.2 shows an example of the set $\mathcal{N}_{\bar{n}}^{2o} \cup \{\bar{n}\}$.

Proposition 5. *Let \mathcal{N}_1 be an optimal solution to Problem 1. For each $\bar{n} \in \mathcal{N}_c$ with $\mathcal{N}_{\bar{n}}^{2o} \neq \emptyset$, we have that $\mathcal{N}_1 \cap (\mathcal{N}_{\bar{n}}^{2o} \cup \{\bar{n}\}) \neq \emptyset$.*

Proof. Assume $\mathcal{N}_{\bar{n}}^{2o} \neq \emptyset$. If $\bar{n} \in \mathcal{N}_1$, then clearly $\mathcal{N}_1 \cap (\mathcal{N}_{\bar{n}}^{2o} \cup \{\bar{n}\}) \neq \emptyset$. If $\bar{n} \notin \mathcal{N}_1$ and $\mathcal{N}_1 \cap \mathcal{N}_{\bar{n}}^{2o} = \emptyset$, then we can add any node in $\mathcal{N}_{\bar{n}}^{2o}$ to \mathcal{N}_1 and increase f_{ce} , which contradicts the assumption that \mathcal{N}_1 is an optimal solution to Problem 1. This is because for any $n \in \mathcal{N}_{\bar{n}}^{2o}$ it holds that $\max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(n, m) = \max_{m \in \emptyset} \Delta F(n, m) = 0$, but $\max_{m \in \mathcal{N}_{\bar{n}}^o} \Delta F(\bar{n}, m) > 0$. \square

Also this result can be formulated as a set of constraints to the integer programming formulation.

$$\sum_{m \in \mathcal{N}_{\bar{n}}^{2o} \cup \{\bar{n}\}} x^\ell(m) \geq 1 \quad \text{for } \bar{n} \in \{n \in \mathcal{N}_c : \mathcal{N}_n^{2o} \neq \emptyset\}.$$

Proposition 5 can be used to compute a lower bound on the number of coordination leaders in an optimal solution. Proposition 5 tells us that each union of a node and its two-hop out-neighbors contains at least one coordination leader, unless that node's two-hop out-neighbor set is empty. However, in most cases these sets overlap and one coordination leader is contained in the two-hop out-neighbor sets of several nodes. We can, nevertheless, select some of these sets so that the selected sets mutually do not intersect. A coordination leader cannot be contained in two of these sets.

Proposition 6. *Let \mathcal{N}_1 be an optimal solution to Problem 1 and let the set of sets $\underline{\mathcal{D}} \subset \{\{n\} \cup \mathcal{N}_n^{2o} : n \in \mathcal{N}_c, \mathcal{N}_n^{2o} \neq \emptyset\}$ be defined such that any two elements of $\underline{\mathcal{D}}$, $\mathcal{D}_1, \mathcal{D}_2 \in \underline{\mathcal{D}}$ have zero intersection, i.e., $\mathcal{D}_1 \cap \mathcal{D}_2 = \emptyset$. Then it holds that $|\mathcal{N}_1| \geq |\underline{\mathcal{D}}|$, and for every $d \in \underline{\mathcal{D}}$, it holds that $d \cap \mathcal{N}_1 \neq \emptyset$.*

Proof. From Proposition 5 it follows that $\mathcal{D}_1 \cap \mathcal{N}_1 \neq \emptyset$ and $\mathcal{D}_2 \cap \mathcal{N}_1 \neq \emptyset$. Since $\mathcal{D}_1 \cap \mathcal{D}_2 = \emptyset$ it holds also that $(\mathcal{D}_1 \cap \mathcal{N}_1) \cap (\mathcal{D}_2 \cap \mathcal{N}_1) = \emptyset$. This holds for any two elements $\mathcal{D}_1, \mathcal{D}_2$ in $\underline{\mathcal{D}}$. Thus, every element of $\underline{\mathcal{D}}$ contains at least one element of \mathcal{N}_1 .

Since this holds for any two elements in $\underline{\mathcal{D}}$, there is at least one unique element in \mathcal{N}_1 for every element in $\underline{\mathcal{D}}$, i.e., there are at least as many elements in \mathcal{N}_1 as in $\underline{\mathcal{D}}$. \square

The set $\underline{\mathcal{D}}$ is an independent subset of the set $\{\{n\} \cup \mathcal{N}_n^{2o} : n \in \mathcal{N}_c, \mathcal{N}_n^{2o} \neq \emptyset\}$. Maximal independent sets, i.e., sets $\underline{\mathcal{D}}$ where no element from $\{\{n\} \cup \mathcal{N}_n^{2o} : n \in \mathcal{N}_c, \mathcal{N}_n^{2o} \neq \emptyset\}$ can be added without violating that any two subsets have non-zero intersection, can be computed with a greedy algorithm. The problem of finding the maximum independent set—this is, the independent set with largest cardinality—is however NP-hard [182], so finding the largest value for the bound might not always be feasible.

Proposition 6 can be formulated as constraint to the integer program as

$$\sum_{n \in \mathcal{N}_c} x^\ell(n) \geq |\underline{\mathcal{D}}|.$$

The quality of this bound depends on the graph. Consider Figure 6.3. The optimal solution with the graph shown on the left side will have 4 or 5 coordination leaders, namely the middle layer of nodes. Adding the top node to the set of coordination leaders does not change the objective. All sets $(\mathcal{N}_n^{2o} \cup \{n\})$ include the node on the top of the graph, and therefore $|\underline{\mathcal{D}}| = 1$ for any choice of $\underline{\mathcal{D}}$. When the weights of the edges from the middle layer to the top node are changed in a way so that they are larger, than the edges from the bottom layer to the middle layer, then the top node becomes the only coordination leader and the bound is tight. On the other hand, the graph shown on the right-hand side of the figure will admit a tight bound regardless of the weights. For every pair of nodes that is connected by an edge, the top node becomes coordination leader. These pairs of nodes are the sets $(\mathcal{N}_n^{2o} \cup \{n\})$, which are all independent.

Worst Case Complexity

An important property to investigate with respect to algorithms for combinatorial optimization problems is the algorithm's worst case complexity. Like many combinatorial optimization problems, Problem 1 can be shown to be NP-hard. This means it is unlikely, even though not yet proven, that there can be an algorithm that solves every instance of the problem efficiently, meaning that the number of




Figure 6.3: Two different coordination graphs that illustrate how the usefulness of Proposition 6 depends on the coordination graph. An optimal solution of Problem 1 on both graphs has at least 4 coordination leaders. Proposition 6 shows that an optimal solution on the left graph has at least one coordination leader whereas an optimal solution on the right graph has at least 4 coordination leaders, which are drawn as gray-filled circles.

computation steps needed to compute the result cannot be upper bounded by a polynomial evaluated on the size of the input. The size of the input is measured in terms of number of edges and nodes in the coordination graph.

Proposition 7. *Problem 1 is NP-hard.*

Proof. We show the result by reduction of the optimization version of the set covering problem to Problem 1. The optimization version of the set problem covering is well known to be NP-hard. Reduction to a known hard problem is a common proof technique for this kind of result [66]. We do this by constructing a coordination graph \mathcal{G}_c for which there is a one-to-one correspondence between coordination leaders and selected sets for the cover. Then we show that the minimum number of leaders that corresponds to a set cover gives the maximum value for f_{ce} .

Consider the following set covering problem. We have a finite set \mathcal{U} . Furthermore, let \mathcal{S}_u be a family of subsets of \mathcal{U} with $\bigcup_{\mathcal{S} \in \mathcal{S}_u} \mathcal{S} = \mathcal{U}$. The problem is to find the smallest number of subsets in \mathcal{S}_u whose union is \mathcal{U} .

We construct the coordination graph as the one shown in Figure 6.4. We introduce a node for each element in \mathcal{U} . We denote the set of these nodes with \mathcal{N}_3 and let $\mu_3 : \mathcal{U} \rightarrow \mathcal{N}_3$ be a bijective mapping from the elements in \mathcal{U} to the nodes in \mathcal{N}_3 . We introduce a node for each element in \mathcal{S}_u . We denote the set of these nodes with \mathcal{N}_2 and let $\mu_2 : \mathcal{S}_u \rightarrow \mathcal{N}_2$ be a bijective mapping from elements in \mathcal{S}_u to nodes in \mathcal{N}_2 . Consider a node $n_2 \in \mathcal{N}_2$ that corresponds to the element $\mathcal{S} \in \mathcal{S}_u$. The in-neighbors of n_2 are $\mathcal{N}_{n_2}^i = \{\mu_3(\mathcal{S}) : \mathcal{S} \in \mu_2^{-1}(n_2)\}$. The weight of the corresponding edges is 1. We introduce an additional node N_1 . There is an edge from




Figure 6.4: Illustration of the graph used to prove that Problem 1 is NP-hard

each node in \mathcal{N}_2 to N_1 with weight 0.5. Clearly, this reduction is linear in the size of the input $\mathcal{U}, \mathcal{S}_u$.

Since N_1 has no out-neighbors, its membership in \mathcal{N}_1 can only increase $f_{ce}(\mathcal{N}_1)$. Since all nodes in \mathcal{N}_3 have no in-neighbors, adding a node in \mathcal{N}_3 to \mathcal{N}_1 can only decrease $f_{ce}(\mathcal{N}_1)$. Thus, the problem of finding the optimal \mathcal{N}_1 reduces to finding which nodes in \mathcal{N}_2 belong to \mathcal{N}_1 . In the optimal solution, each node in \mathcal{N}_3 has at least one out-neighbor in \mathcal{N}_1 . Otherwise we could add any out-neighbor of that node to \mathcal{N}_1 and increase $f_{ce}(\mathcal{N}_1)$ by at least 0.5. Therefore, $\{\mu_2^{-1}(n) : n \in \mathcal{N}_1 \cap \mathcal{N}_2\}$ is a set cover of \mathcal{U} . Otherwise there would be $u \in \mathcal{U}$ such that there is no $\mathcal{S} \in \{\mu_2^{-1}(n) : n \in \mathcal{N}_1 \cap \mathcal{N}_2\}$ with $u \in \mathcal{S}$. If such a u existed, $\mu_3(u)$ would be a node with no out-neighbor in $\mathcal{N}_1 \cap \mathcal{N}_2$. Furthermore, let $\bar{\mathcal{S}}_u \subseteq \mathcal{S}_u$ be a set cover of \mathcal{U} . Then $\{\mu_2(\mathcal{S}) : \mathcal{S} \in \bar{\mathcal{S}}_u\}$ has the property that $\bigcup_{n \in \{\mu_2(\mathcal{S}) : \mathcal{S} \in \bar{\mathcal{S}}_u\}} \mathcal{N}_n^1 = \mathcal{N}_3$, so any set cover has the property that all nodes in \mathcal{N}_3 have at least one out-neighbor in \mathcal{N}_1 . Each node in \mathcal{N}_2 contributes with 0.5 to the objective if it is not in \mathcal{N}_1 . Therefore, the optimal \mathcal{N}_1 contains a minimum number of nodes from \mathcal{N}_2 such that every node in \mathcal{N}_3 has at least one out-neighbor in $\mathcal{N}_1 \cap \mathcal{N}_2$. Since any $\mathcal{N}_1 \cap \mathcal{N}_2$ that fulfills this property maps to a set cover $\bar{\mathcal{S}}_u$ and vice versa, and since $|\mathcal{N}_1 \cap \mathcal{N}_2| = |\bar{\mathcal{S}}_u|$, we have that $\bar{\mathcal{S}}_u$ is the solution to the set covering problem. Thus, the NP-hard set-covering problem can be reduced to Problem 1, which shows that Problem 1 is NP-hard. \square

Algorithm 1 Iterative Algorithm to compute the set of coordination leaders \mathcal{N}_1 .

Input: \mathcal{G}_c

Output: \mathcal{N}_1

$\mathcal{N}_1 \leftarrow \emptyset$

while $\{n \in \mathcal{N}_c : \Delta u(n, \mathcal{N}_1) > 0\} \neq \emptyset$ **do**

 Select $\bar{n} \in \{n \in \mathcal{N}_c : \Delta u(n, \mathcal{N}_1) > 0\}$

if $\bar{n} \in \mathcal{N}_1$ **then**

$\mathcal{N}_1 \leftarrow \mathcal{N}_1 \setminus \{\bar{n}\}$

else

$\mathcal{N}_1 \leftarrow \mathcal{N}_1 \cup \{\bar{n}\}$

end if

end while

Exact solutions to NP-hard problems can be hard to compute, which is why heuristic and approximate solutions are often used. These algorithms compute sub-optimal solutions in a computationally efficient way. Heuristic algorithms for Problem 1 are developed in the next section.

6.3 Heuristic and Approximate Solutions to the Coordination Leader Selection Problem

In this section we present an algorithm that computes heuristic solutions to Problem 1. This is motivated by the result that Problem 1 is NP-hard, which implies that finding an optimal solution might take a very long time to find. Since Problem 1 is repeatedly solved as part of an automated feedback loop, it is important to find solutions quickly and to guarantee a result after a fixed amount of time.

Iterative Algorithm

Consider Algorithm 1. The input is a coordination graph \mathcal{G}_c and the output is a set of coordination leaders \mathcal{N}_1 . Initially \mathcal{N}_1 is an empty set. In each iteration, a node $\bar{n} \in \mathcal{N}_c$ is selected for which the objective function f_{ce} is increased if it is added to \mathcal{N}_1 or removed from \mathcal{N}_1 , and \mathcal{N}_1 is updated accordingly. The difference in f_{ce} when adding or removing a node in \mathcal{N}_c to or from the set of coordination leaders \mathcal{N}_1 is given by a function Δu . The algorithm iterates until no further increase of f_{ce} is possible.

The function Δu that measures how much is gained from switching whether \bar{n} belongs to \mathcal{N}_1 is defined as follows:

$$\Delta u(\bar{n}, \mathcal{N}_1) = \begin{cases} f_{ce}(\mathcal{N}_1 \setminus \{\bar{n}\}) - f_{ce}(\mathcal{N}_1) & \text{if } \bar{n} \in \mathcal{N}_1 \\ f_{ce}(\mathcal{N}_1 \cup \{\bar{n}\}) - f_{ce}(\mathcal{N}_1) & \text{otherwise .} \end{cases} \quad (6.6)$$

If $\bar{n} \notin \mathcal{N}_1$, we get

$$\begin{aligned} f_{\text{ce}}(\mathcal{N}_1 \cup \{\bar{n}\}) - f_{\text{ce}}(\mathcal{N}_1) = \\ \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_1} \left(\max_{m \in \mathcal{N}_n^o \cap (\mathcal{N}_1 \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(n, m) \right) \\ - \max_{m \in \mathcal{N}_{\bar{n}}^o \cap \mathcal{N}_1} \Delta F(\bar{n}, m). \end{aligned} \quad (6.7)$$

The sum over n covers nodes that can select \bar{n} as their new coordination leader. The last summand accounts for \bar{n} possibly not being a coordination follower any longer.

If $\bar{n} \in \mathcal{N}_1$, we get

$$\begin{aligned} f_{\text{ce}}(\mathcal{N}_1 \setminus \{\bar{n}\}) - f_{\text{ce}}(\mathcal{N}_1) = \\ \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_1} \left(\max_{m \in \mathcal{N}_n^o \cap (\mathcal{N}_1 \setminus \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(n, m) \right) \\ + \max_{m \in \mathcal{N}_{\bar{n}}^o \cap (\mathcal{N}_1 \setminus \{\bar{n}\})} \Delta F(\bar{n}, m). \end{aligned} \quad (6.8)$$

The sum over n covers nodes that can have \bar{n} as their coordination leader before the change. The last summand accounts for \bar{n} possibly becoming a coordination follower.

We consider two methods to select \bar{n} from the set $\{n \in \mathcal{N}_c : \Delta u(n, \mathcal{N}_1) > 0\}$. The first method is to select n in a greedy manner according to $\bar{n} = \arg \max_{n \in \mathcal{N}_c} \Delta u(n, \mathcal{N}_1)$. The second method is to choose \bar{n} randomly with equal probability from the set $\{n \in \mathcal{N}_c : \Delta u(n, \mathcal{N}_1) > 0\}$.

Algorithm 1 is guaranteed to converge in finite time. This is due to the number of possible subsets of \mathcal{N}_c being finite and thus the number possible assignments of \mathcal{N}_1 is finite. In every iteration $f_{\text{ce}}(\mathcal{N}_1)$ strictly increases, which means that \mathcal{N}_1 changes in every iteration and the same assignment for \mathcal{N}_1 never reoccurs. So in the worst case Algorithm 1 iterates over all subsets of \mathcal{N}_c before termination. It is also possible to interrupt the algorithm before termination and use the value of \mathcal{N}_1 at this point in the execution. It is easy to see that a coordination leader set \mathcal{N}_1 computed by Algorithm 1 fulfills the condition on the optimal solution stated in Proposition 5, i.e., that every union of a node and its two-hop out-neighbors contains at least one coordination leader.

Algorithm 1 can be efficient. Note for instance that the function Δu can be computed based on the sub-graph induced by the one- and two-hop neighbors of \bar{n} only. This means that the average complexity of computing Δu is a function of the average node degree but not of the number of nodes in the coordination graph. Furthermore, if a node is added to or removed from \mathcal{N}_1 , then only the Δu for the two-hop neighbors needs to be recomputed.

Simulations in Section 8.1 suggest that selecting \bar{n} in a greedy or a random manner makes little difference for the quality of the computed solution. However, greedy

Algorithm 2 One-Pass Algorithm**Input:** \mathcal{G}_c **Output:** $\underline{\mathcal{N}}_1$ $\underline{\mathcal{N}}_1 \leftarrow \emptyset, \overline{\mathcal{N}}_1 \leftarrow \mathcal{N}_c$ **for** $\bar{n} \in \mathcal{N}_c$ **do** $a \leftarrow f_{ce}(\underline{\mathcal{N}}_1 \cup \bar{n}) - f_{ce}(\underline{\mathcal{N}}_1), b \leftarrow f_{ce}(\overline{\mathcal{N}}_1 \setminus \bar{n}) - f_{ce}(\overline{\mathcal{N}}_1)$ **if** $a \geq b$ **then** $\underline{\mathcal{N}}_1 \leftarrow \underline{\mathcal{N}}_1 \cup \bar{n}$ **else** $\overline{\mathcal{N}}_1 \leftarrow \overline{\mathcal{N}}_1 \setminus \bar{n}$ **end if****end for**

node selection tends to lead to less iterations of the algorithm and is thus better suited for a serial implementation. Random node selection might be preferable for a parallel implementation due to the reduced need for synchronization.

One-Pass Algorithm

The approximation algorithm presented in [50] is applicable to Problem 1 and shown here as Algorithm 2. It has two theoretical advantages over Algorithm 1. Algorithm 2 executes $|\mathcal{N}_c|$ iterations while Algorithm 1 can potentially iterate over $2^{|\mathcal{N}_c|} - 1$ subsets with similar complexity per iteration. Furthermore, Algorithm 2 guarantees a (1/3)-approximation, while for Algorithm 1 no theoretical performance guarantees have been found.

Proposition 8. *Algorithm 2 is a (1/3)-approximation to Problem 1.*

Proof. We know from [50] that Algorithm 2 is a (1/3)-approximation to Problem 1, if f_{ce} is submodular. In particular, we need to show that for every $\mathcal{N}_A \subseteq \mathcal{N}_B \subseteq \mathcal{N}_c$ and $\bar{n} \in \mathcal{N}_c \setminus \mathcal{N}_B$ it holds that $f(\mathcal{N}_A \cup \{\bar{n}\}) - f(\mathcal{N}_A) \geq f(\mathcal{N}_B \cup \{\bar{n}\}) - f(\mathcal{N}_B)$.

We have that

$$f_{ce}(\mathcal{N}_1 \cup \{\bar{n}\}) - f_{ce}(\mathcal{N}_1) = \sum_{n \in \mathcal{N}_n^1 \setminus \mathcal{N}_1} \left(\max_{m \in \mathcal{N}_n^o \cap (\mathcal{N}_1 \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(n, m) \right) - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_1} \Delta F(\bar{n}, m).$$

It holds that

$$\begin{aligned} & \mathcal{N}_A \subseteq \mathcal{N}_B \\ & \Rightarrow \mathcal{N}_n^o \cap \mathcal{N}_A \subseteq \mathcal{N}_n^o \cap \mathcal{N}_B \\ & \Rightarrow \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_A} \Delta F(\bar{n}, m) \leq \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_B} \Delta F(\bar{n}, m) \\ & \Rightarrow - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_A} \Delta F(\bar{n}, m) \geq - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_B} \Delta F(\bar{n}, m). \end{aligned} \tag{6.9}$$

Furthermore

$$\begin{aligned}
& \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_A} \left(\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_A} \Delta F(n, m) \right) \\
& \quad - \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_B} \left(\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_B} \Delta F(n, m) \right) \\
& \geq \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_B} \left(\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_A} \Delta F(n, m) \right) \\
& \quad - \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_B} \left(\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_B} \Delta F(n, m) \right) \\
& = \sum_{n \in \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_B} \left(\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_A} \Delta F(n, m) \right. \\
& \quad \left. - \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) + \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_B} \Delta F(n, m) \right),
\end{aligned}$$

where the inequality holds since

$$\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_A} \Delta F(n, m) \geq 0,$$

and

$$\mathcal{N}_A \subseteq \mathcal{N}_B \Rightarrow \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_A \supseteq \mathcal{N}_{\bar{n}}^i \setminus \mathcal{N}_B.$$

We distinguish four cases:

1. $\bar{n} \neq \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m)$, $\bar{n} \neq \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m)$
2. $\bar{n} = \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m)$, $\bar{n} \neq \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m)$
3. $\bar{n} \neq \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m)$, $\bar{n} = \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m)$
4. $\bar{n} = \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m)$, $\bar{n} = \arg \max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m)$

We assume that there is a consistent way of selecting the $\arg \max$ in case there is no unique maximizer. In Case 1,

$$\begin{aligned}
\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) &= \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_A} \Delta F(n, m) \\
\max_{m \in \mathcal{N}_n^{\circ} \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) &= \max_{m \in \mathcal{N}_n^{\circ} \cap \mathcal{N}_B} \Delta F(n, m),
\end{aligned}$$

and thus

$$\begin{aligned} & \max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_A} \Delta F(n, m) \\ & - \max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) + \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_B} \Delta F(n, m) \\ & = 0. \end{aligned}$$

For case 2, we have that

$$\max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_B} \Delta F(n, m) = 0,$$

and clearly

$$\max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_A} \Delta F(n, m) \geq 0,$$

so that

$$\begin{aligned} & \sum_{n \in \mathcal{N}_{\bar{n}}^1 \setminus \mathcal{N}_B} \left(\max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_A} \Delta F(n, m) \right. \\ & \quad \left. - \max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) + \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_B} \Delta F(n, m) \right) \\ & \geq 0. \end{aligned} \tag{6.10}$$

Case 3 cannot occur since $\mathcal{N}_B \supseteq \mathcal{N}_A$ which implies that

$$\bar{n} \neq \arg \max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) \in \mathcal{N}_B,$$

but

$$\max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) > \Delta F(n, \bar{n}),$$

so

$$\bar{n} = \arg \max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m)$$

cannot hold. In case 4, we have

$$\max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) = \max_{m \in \mathcal{N}_n^\circ \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) = \Delta F(n, \bar{n}),$$

and since $\mathcal{N}_A \subseteq \mathcal{N}_B$

$$\max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_B} \Delta F(n, m) \geq \max_{m \in \mathcal{N}_n^\circ \cap \mathcal{N}_A} \Delta F(n, m),$$

and thus

$$\begin{aligned}
& \max_{m \in \mathcal{N}_n^o \cap (\mathcal{N}_A \cup \{\bar{n}\})} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_A} \Delta F(n, m) \\
& \quad - \max_{m \in \mathcal{N}_n^o \cap (\mathcal{N}_B \cup \{\bar{n}\})} \Delta F(n, m) + \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_B} \Delta F(n, m) \\
& \quad = \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_B} \Delta F(n, m) - \max_{m \in \mathcal{N}_n^o \cap \mathcal{N}_A} \Delta F(n, m) \geq 0
\end{aligned}$$

Combining (6.9) and (6.10) shows that f_{ce} is submodular. \square

Remark 26. The authors of [50] also present a randomized version of the algorithm that is a $(1/2)$ -approximation in expectation. In simulation results presented in [113], the theoretical advantages Algorithm 2 has over Algorithm 1 do not manifest.

Remark 27. Just as for Algorithm 1, the expressions $f_{ce}(\mathcal{N}_1 \cup \bar{n}) - f_{ce}(\mathcal{N}_1)$ and $f_{ce}(\mathcal{N}_1 \setminus \bar{n}) - f_{ce}(\mathcal{N}_1)$ can be computed based on the sub-graph induced by the two-hop out-neighbors of \bar{n} without computing f_{ce} explicitly (see (6.7) and (6.8)), which can significantly improve performance.

Remark 28. Algorithm 2 and 1 can be combined using the solution of Algorithm 2 as the initial assignment for \mathcal{N}_1 instead of the empty set. Since Algorithm 1 strictly increases the objective, the combined algorithm inherits the property of being a $(1/3)$ -approximation to problem to Problem 1. Algorithm 2 and 1 can also be used as a heuristic when solving the integer-program.

Having computed the set of coordination leaders, there is immediately a platoon plan for each vehicle. These plans can be further improved as discussed in the following section.

6.4 Joint Platoon-Plan Optimization

In this section, we derive how to jointly optimize the pairwise platoon plans that are selected by any algorithm solving Problem 1 such as Algorithm 1 or Algorithm 2. We do this by formulating a convex optimization problem with linear constraints for a group consisting of a coordination leader and its coordination followers. Hereby, the timing when platoons are assembled and split is adjusted while the locations where this happens stay the same. Vehicles that are not matched to any coordination leader and are not coordination leaders themselves just follow their default plans and are not considered in this section.

Remark 29. We present the derivation for a fixed, space-independent speed range as introduced in Section 4.3. This can be extended to space-dependent speed profiles. The main challenge lies in finding a convex analytical fuel model describing the aggregated behavior over larger segments.

Consider a coordination leader $n_1 \in \mathcal{N}_1$ and its followers

$$\mathcal{N}_{n_1}^{\text{fl}} = \{n \in \mathcal{N}_c \setminus \mathcal{N}_1 : n_1 = \arg \max_{m \in \mathcal{N}_1 \cap \mathcal{N}_c^n} \Delta F(n, m)\}.$$

This group of agents is denoted

$$\mathcal{N}_g = \{n_1\} \cup \mathcal{N}_{n_1}^{\text{fl}}.$$

We construct an ordered set of time instances $\mathbf{t} = (\mathbf{t}[1], \mathbf{t}[2], \dots)$. This set contains the start time and the arrival deadline of the coordination leader, and the merge times and the split times of its followers. We divide the distance traveled by the coordination leader from start to destination according to these time instances and get the distances $\mathbf{W}_{n_1}[i] = v_{\text{cd}} \times (\mathbf{t}[i+1] - \mathbf{t}[i])$ between these points, where v_{cd} is the speed of the leader according to its default plan. These are the distances between the points where coordination followers join or leave the platoon. Similarly, for a coordination follower $n \in \mathcal{N}_{n_1}^{\text{fl}}$, we have

$$\mathbf{W}_n = (v_S \times (t_n^{\text{M}} - t_n^{\text{S}}), \mathbf{W}_{n_1}[i_n^{\text{M}}], \dots, \mathbf{W}_{n_1}[i_n^{\text{SP}}], v_{\text{SP}} \times (t_n^{\text{A}} - t_n^{\text{SP}})).$$

The variables $t_n^{\text{S}}, t_n^{\text{M}}, t_n^{\text{SP}}, t_n^{\text{A}}$ denote the start time, merge time, split time, and arrival time of follower n according to its adapted plan. Recall that v_S is the speed of the coordination follower until the vehicle merges with the coordination leader and v_{SP} is the speed after the vehicle has split up from the coordination leader according to the adapted plan. The first element of \mathbf{W}_n is the distance along the route from start to the merge point. For the part of the route the coordination follower drives in a platoon with the coordination leader, the entries are the same as for the coordination leader. The indices $i_n^{\text{M}}, i_n^{\text{SP}}$ are defined accordingly. The last element of \mathbf{W}_n is the distance from the split point to the destination of the follower. We denote the number of elements in \mathbf{W}_n as N_n^{Y} .

Figure 6.5 illustrates the definition of \mathbf{W}_n . We introduce sequences $\mathbf{p}_n = (\mathbf{p}_n[1], \dots, \mathbf{p}_n[N_n^{\text{Y}}])$ that indicate on which segments of the journey the coordination follower is a platoon follower. If vehicle n is a platoon follower on the segment that corresponds to $\mathbf{W}_n[i]$ for some i , then $\mathbf{p}_n[i] = 1$. Otherwise, we have $\mathbf{p}_n[i] = 0$. For the coordination leader n_1 , we have $\mathbf{p}_{n_1} = (0, \dots, 0)$ and for a coordination follower $n \in \mathcal{N}_{n_1}^{\text{fl}}$, we have that $\mathbf{p}_n = (0, 1, 1, \dots, 1, 0)$.

The fuel consumption per distance traveled is modeled as function $f : \mathbb{R} \times \{0, 1\} \rightarrow \mathbb{R}$, where the first argument is the speed and the second argument whether or not the vehicle is a trailing vehicle in a platoon.

We express the speed and time sequence of vehicle $n \in \mathcal{N}_g$ as traversal times $\mathbf{T}_n = (\mathbf{T}_n[1], \dots, \mathbf{T}_n[N_n^{\text{Y}}])$ of the segments \mathbf{W}_n . The speed on each such segment remains constant and can be computed as

$$\mathbf{v}_n[i] = \frac{\mathbf{W}_n[i]}{\mathbf{T}_n[i]} \quad i \in \{1, \dots, N_n^{\text{Y}}\}.$$

The traversal times of the segments in all vehicles' routes are the optimization variables. Working with traversal times rather than the sequence of speeds \mathbf{v} allows us to state the optimization problem with linear constraints.




Figure 6.5: Illustration of how the sequences \mathbf{W}_n are defined. The red, dotted line represents the route of the coordination leader and the black, solid lines with arrows represent the routes of the coordination followers. The thin lines indicate the distances that the elements of \mathbf{W}_{n_1} correspond to.

With these definitions, we are ready to state the following problem:

Problem 2.

$$\min_{\{\mathbf{T}_n: n \in \mathcal{N}_g\}} \sum_{n \in \mathcal{N}_g} \sum_{i=1}^{N_n^v} f\left(\frac{\mathbf{W}_n[i]}{\mathbf{T}_n[i]}, \mathbf{p}_n[i]\right) \mathbf{W}_n[i] \quad (6.11a)$$

s.t.

for $n \in \mathcal{N}_g$:

$$\frac{\mathbf{W}_n[i]}{v_{\max}} \leq \mathbf{T}_n[i], \quad i \in \{1, \dots, N_n^v\} \quad (6.11b)$$

$$\frac{\mathbf{W}_n[i]}{v_{\min}} \geq \mathbf{T}_n[i], \quad i \in \{1, \dots, N_n^v\} \quad (6.11c)$$

$$t_n^S + \sum_{i=1}^{N_n^v} \mathbf{T}_n[i] \leq t_n^D \quad (6.11d)$$

and for $n \in \mathcal{N}_{n_1}^{\text{fl}}$:

$$t_n^S + \mathbf{T}_n[1] = t_{n_1}^S + \sum_{i=1}^{i_n^M - 1} \mathbf{T}_{n_1}[i] \quad (6.11e)$$

$$\mathbf{T}_n[1+i] = \mathbf{T}_{n_1}[i_n^M + i - 1], \quad i \in \{1, \dots, i_n^{\text{Sp}} - i_n^M + 1\}. \quad (6.11f)$$

The objective function (6.11a) equals the estimated combined fuel consumption of all active vehicles to complete their assignments. It is composed of the

fuel consumption of the coordination leader and the coordination followers. The coordination leader is considered to travel alone or to take the role as the platoon leader throughout its journey. The coordination followers travel alone on the first and the last segment of their journey. They become platoon followers in-between these segments.

There are two sets of constraints. The first set applies to all vehicles and ensures that the sequences \mathbf{T}_n correspond to valid platoon plans. In particular, the constraints (6.11b) and (6.11c) express that the trajectories stay within the allowed range of speed $[v_{\min}, v_{\max}]$. The constraints (6.11d) express that all vehicles arrive at their destinations before their individual deadlines t_n^D . The second set of constraints ensures that platooning happens as specified in the original pairwise plans. The constraints (6.11e) ensure that the coordination leader and each of its followers arrive at the same time at their respective merge point. The constraints (6.11f) ensure that the speed of the leader and the speed of the follower are the same when they are supposed to platoon.

When $f(T^{-1}, 0)$, $f(T^{-1}, 1)$ are convex in T for $T > 0$, then the objective (6.11a) is a sum of convex functions and hence convex. For instance, polynomials with arbitrary constant part and non-negative coefficients fulfill this requirement. Furthermore, all constraints are linear. Thus, Problem 2 is a convex optimization problem for which well developed numerical solvers are readily available [46, 22]. The optimization is initialized with the pairwise plans.

Remark 30. In this formulation, we assume that the coordination follower does not platoon with the coordination leader on the first and the last segment of its route. The formulation can be adapted to the case in which platoon follower and leader platoon on the first, the last, or both segments of the coordination followers route setting $\mathbf{p}_n[1]$ and $\mathbf{p}_n[N_n^y]$ to 1 accordingly as well as changing some indices appropriately.

6.5 Summary

This chapter describes ways of computing platoon plans for more than two vehicles. The key element in making the presented approach to coordinating heavy-duty vehicle platooning feasible for large numbers of vehicles is the systematic combination of default plans and adapted plans, building on the results from Chapter 4 and 5. Some vehicles, the coordination leaders, get their default plans assigned. The remaining vehicles use the most fuel efficient plan that is adapted to one of the coordination leaders or travel alone. The fuel efficiency of the adapted plans is based on becoming a platoon follower of a coordination leader during a part of the route. By selecting coordination leaders in a smart way, the fuel savings that result from the adapted plans are maximized.

We show, that the problem of optimally selecting coordination leaders can be formulated as an integer program allowing the use of general purpose solvers to obtain exact and heuristic solutions. The presented results on the structure of the

optimal solution help to reduce the search space and give insight into the problem. Optimally selecting coordination leaders is proven to be NP-hard, which means that any algorithm that computes exact solutions might have very long running times.

This motivates the design of algorithms that compute a good selection of coordination leaders efficiently but not necessarily the best one. One possible choice is an algorithm that starts from an empty set of coordination leaders and iteratively adds and removes leaders from that set, increasing the fuel savings in each iteration. Such an algorithm can compute a good result efficiently for large numbers of vehicles. Another option presented in this chapter is a one-pass $(1/3)$ -approximation algorithm, which iterates once over all assignments.

The combination of default and adapted plans can be further improved in order to reduce fuel consumption. By committing to which platoons to form and where they should be formed, it is possible to adjust the timing that leads to such platoons. Since adjusting the timing alters the speed profiles, it has an influence on the total fuel consumption. We derive how by using convex optimization, the timing can be adjusted in a fuel-optimal way.

Chapter 7

Computational Efficiency Improvements

IN this chapter, we introduce a scalable way of computing all pairs of assignments that have an overlapping route. Two vehicles can only form a platoon if they have at least one segment of their routes in common. Furthermore, they have to be able to arrive at the common segment at the same point in time. In the plan computation stage when computing the coordination graph \mathcal{G}_c as introduced in Chapter 6, all pairs of transport assignments that can potentially platoon have to be identified. The straightforward way of doing this is to compute for each pair of transport assignments individually the common road segment identifiers in their routes. If there are common segments, we can determine if the vehicles can form a platoon on any of those segments taking into account the start times, arrival deadlines, and speed constraints following a procedure for computing adapted platoon plans such as the ones introduced in Chapter 4. This procedure involves computing a number of set intersections, and this number scales quadratically with the transport assignment count. Such computation can become problematic for large vehicle fleets. Therefore, we introduce a computationally less expensive and scalable step to narrow down the set of candidate pairs.

The proposed approach is inspired by solutions to a related problem in computer graphics of finding the intersection of geometric objects [31, 77, 121, 65, 148]. When the number of objects is large, instead of checking all possible pairs for intersection, it is more efficient to identify a smaller set of candidate pairs. In our case, this set includes all pairs of vehicles that can potentially platoon, i.e., for which an adapted platoon plan exists. The pairs in the candidate set are then processed individually.

Section 7.1 associates each assignment with a sequence of time intervals and a sequence of two-dimensional positions. This information provides limits on the possible points in time a vehicle can be at a certain position as long as the vehicle travels according to a valid platoon plan. If two vehicles can be at the same position at the same time, they are candidates for platooning. Section 7.2 introduces




Figure 7.1: This plot illustrates the computation of the lower and upper bounds on the node arrival times \underline{t} , \bar{t} . The vehicle begins the trip at the start time t^S and at the start position \mathcal{P}^S . The lower bound \underline{t} can be achieved by the maximum speed profile from start to destination. The upper bound \bar{t} can be achieved by selecting the minimum speed up to the point where the vehicle arrives on deadline t^D when it follows the maximum speed profile until the destination.

the concept of feature extraction and culling. Features are computed based on sequences of positions and time intervals, and they are significantly less complex than these sequences. For some assignment pairs, it is feasible to efficiently rule out the possibility of platooning based on these features. An algorithm for such computation is called a classifier. Section 7.3 develops appropriate features and classifiers for the problem at hand. In Section 7.4, these classifiers are demonstrated in a simulation example.

7.1 Candidate Platoon Pairs

We start by defining a function that indicates whether platooning between two transport assignments is possible or not. This is the case if there is at least one common route segment in the routes of the transport assignments where the vehicles can platoon. To this end, we compute lower bounds \underline{t} and upper bounds \bar{t} on the segment arrival times \mathbf{t} . Overlapping time bounds on two consecutive segments indicate that the two transport assignments can potentially platoon. Recall, that the route e_i for vehicle i is represented as lists of route segment identifiers, i.e., edges in the road network graph $\mathcal{G}_r = (\mathcal{N}_r, \mathcal{E}_r)$, where $\mathcal{E}_r \subseteq \mathcal{N}_r \times \mathcal{N}_r$.

The lower bound $\underline{\mathbf{t}}$ on the segment arrival times is given by the time profile corresponding to the maximum speed profile as in defined in Section 4.4, Definition 2. The upper bound is defined as

$$\bar{\mathbf{t}}[i] = \min(\underline{\mathbf{t}}[i] + t^D - \underline{t}^A, \bar{\mathbf{t}}^*[i]) \quad \text{for } i \in \{1, \dots, N^A\},$$

where $\bar{\mathbf{t}}^*$ is the time sequence corresponding to the vehicle traveling at minimum speed $\underline{\sigma}\bar{\mathbf{v}}[i]$ for $i \in \{1, \dots, N^A\}$ regardless of the arrival deadline. The difference between arrival deadline t^D and earliest arrival time according to the maximum speed profile \underline{t}^A quantifies how much later compared to time obtained when following the maximum speed profile the vehicle can arrive at any route segment and not miss the deadline. Figure 7.1 illustrates the above definition of $\underline{\mathbf{t}}$ and $\bar{\mathbf{t}}$. Furthermore, each road segment is associated with the two-dimensional position of its beginning $\mathbf{P}_{\text{geo}}^S : \mathcal{E}_c \rightarrow \mathbb{R}^2$ and end $\mathbf{P}_{\text{geo}}^E : \mathcal{E}_c \rightarrow \mathbb{R}^2$. This can be, for instance, measured in longitude and latitude.

Remark 31. We follow here the framework from Section 4.4 but finding feasible time intervals or potentially an over-approximation for other assumptions on the vehicle dynamics is straightforward in most cases.

We introduce a function that indicates whether or not two transport assignments have the possibility to form a platoon.

Definition 6 (Coordination Function). The coordination function $C : \mathcal{N}_c \times \mathcal{N}_c \rightarrow \{0, 1\}$ has the following properties. Let $\underline{\mathbf{t}}_n, \underline{\mathbf{t}}_m$ be the lower bounds and $\bar{\mathbf{t}}_n, \bar{\mathbf{t}}_m$ the upper bounds on the node arrival times of transport assignments n and m . Then it holds that $C(n, m) = 1$, if there are indices a, b such that

$$\mathbf{e}_n[a] = \mathbf{e}_m[b]$$

and

$$\begin{aligned} [\underline{\mathbf{t}}_n[a], \bar{\mathbf{t}}_n[a]] \cap [\underline{\mathbf{t}}_m[b], \bar{\mathbf{t}}_m[b]] &\neq \emptyset \\ [\underline{\mathbf{t}}_n[a+1], \bar{\mathbf{t}}_n[a+1]] \cap [\underline{\mathbf{t}}_m[b+1], \bar{\mathbf{t}}_m[b+1]] &\neq \emptyset \end{aligned}$$

Otherwise $C(n, m) = 0$.

Comparing the routes and the time bounds in order to evaluate C , is straightforward but computationally expensive. We refer to this as the *exact algorithm*. In the remainder of this chapter, we derive a scalable method for computing the set of all possible platoon pairs $\mathcal{C} = \{(n, m) \in \mathcal{N}_c \times \mathcal{N}_c : C(n, m) = 1\}$. Instead of iterating over all elements in $\mathcal{N}_c \times \mathcal{N}_c$ and using the exact algorithm, we propose first efficiently computing an over-approximation $\hat{\mathcal{C}} \supset \mathcal{C}$ (see Figure 7.2) and then applying the exact algorithm.

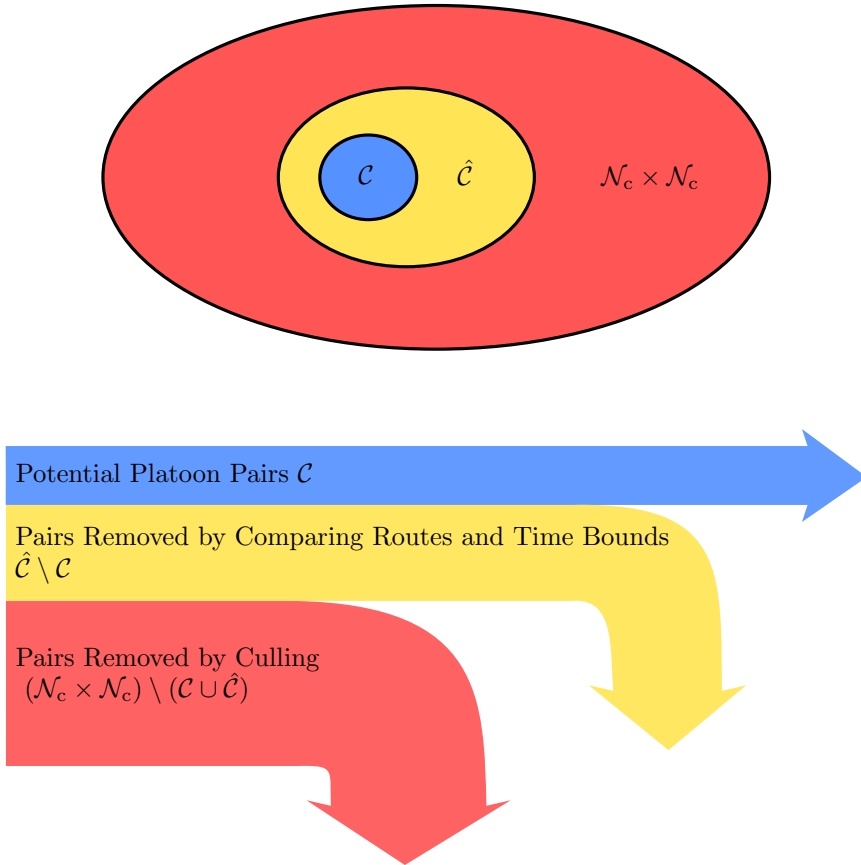


Figure 7.2: Instead of computing \mathcal{C} by directly iterating over all element in $\mathcal{N}_c \times \mathcal{N}_c$, we first compute an over-approximation of \mathcal{C} denoted $\hat{\mathcal{C}}$ in an efficient way.

Remark 32. The exact algorithm can be implemented computing the set intersection of the route segments identifiers and then comparing the time bounds on the intersecting road segments. Fast methods for computing set intersections that scale linearly in the size of the sets exist [73]. However, the number of assignment pairs scales quadratically in the number of assignments and can become so large that it is worthwhile studying a more efficient approach.




Figure 7.3: Each assignment’s route and time bounds are used to compute features, such as an interval.

7.2 Culling Candidate Platoon Pairs

The key idea of our approach is to extract features from the routes and time bounds $(\mathbf{e}, \mathbf{t}, \bar{\mathbf{t}})$ of the transport assignments, as illustrated in Figure 7.3, in order to compute $\hat{\mathcal{C}}$. These features can be more efficiently processed than $(\mathbf{e}, \mathbf{t}, \bar{\mathbf{t}})$. The features are designed in a way that no platooning opportunity in \mathcal{C} is excluded from $\hat{\mathcal{C}}$, so that \mathcal{C} can be computed from $\hat{\mathcal{C}}$ using the exact algorithm. However, there might be some additional elements in $\hat{\mathcal{C}}$ that do not actually correspond to platooning opportunities. We call these additional elements false-positives. The less false-positives there are in $\hat{\mathcal{C}}$, the faster the computation of \mathcal{C} from $\hat{\mathcal{C}}$ is. This approach is inspired by algorithms for detecting collisions between a large number of geometric objects [31, 77]. Figure 7.2 illustrates the relation between $\mathcal{N}_c \times \mathcal{N}_c$, $\hat{\mathcal{C}}$, and \mathcal{C} .

We consider two types of features. These are *interval features* and *binary features*. Interval features map each object to an interval. The corresponding classifier indicates an intersection between two objects if the intervals generated by the objects overlap. There are algorithms, such as [31, 77], that can compute this classifier for all object pairs more efficiently than checking each pair individually if the number of intersecting pairs is small. Binary features map each object to a boolean value. The corresponding classifier indicates an intersection between two objects if the feature holds true for both objects. In Section 7.3, we derive appropriate features for the problem of identifying candidate platoon pairs.

The classifiers are aggregated using boolean connectives. We formalize this in the remainder of the section. Let \mathcal{N} be a set of objects. We define a classifier as a function $c : \mathcal{N} \times \mathcal{N} \rightarrow \{0, 1\}$. If $c(n, m) = 0$, we call the combination of c and (n, m) a negative, and if $c(n, m) = 1$, we call it a positive. Let $g : \mathcal{N} \times \mathcal{N} \rightarrow \{0, 1\}$ be the ground truth, which can be computed by the exact algorithm. If, for a pair (n, m) , we have $g(n, m) = 0$ and $c(n, m) = 1$, we call it a false-positive, and if $g(n, m) = 1$ and $c(n, m) = 0$, we call it a false-negative. Our aim is to design classifiers that yield no false negatives for all elements of $\mathcal{N} \times \mathcal{N}$ and few false-positives that have to be processed by the exact algorithm in addition to the true-positives.

We can identify two types of basic classifiers that are combined in a specific way in order to achieve the above objective. A classifier c is *required* if

$$\neg c(n, m) \Rightarrow \neg g(n, m)$$

for all $n, m \in \mathcal{N} \times \mathcal{N}$. In some cases, we have to take into account a set of classifiers to conclude that g does not hold. A set of classifiers \mathcal{S}_c is *required* if

$$\neg \bigvee_{c \in \mathcal{S}_c} c(n, m) \Rightarrow \neg g(n, m)$$

for all $n, m \in \mathcal{N} \times \mathcal{N}$. It is straightforward to construct a required classifier from a required set of classifiers.

Proposition 9. *If a set \mathcal{S}_c of classifiers is required, then $\bigvee_{c \in \mathcal{S}_c} c$ is a required classifier.*

We can combine two required classifiers into one required classifier that performs no worse than any of the required classifiers it is combined of.

Proposition 10. *If c_1 and c_2 are required classifiers, then $c_{12}(n, m) := c_1(n, m) \wedge c_2(n, m)$ is a required classifier. Let $\bar{\mathcal{E}}_{12} = \{(n, m) \in \mathcal{N} \times \mathcal{N} : c_{12}(n, m) = 0\}$ be the set of negatives of c_{12} and let $\bar{\mathcal{E}}_1, \bar{\mathcal{E}}_2$ be the set of negatives for c_1 and c_2 respectively. Then $\bar{\mathcal{E}}_1 \subseteq \bar{\mathcal{E}}_{12}$ and $\bar{\mathcal{E}}_2 \subseteq \bar{\mathcal{E}}_{12}$.*

Proof. For c_{12} to be required, we need to show that $\neg c_{12}(n, m) \Rightarrow \neg g(n, m)$ for all $n, m \in \mathcal{N} \times \mathcal{N}$. We have

$$\begin{aligned} (\neg c_1 \Rightarrow \neg g) \wedge (\neg c_2 \Rightarrow \neg g) &= (c_1 \vee \neg c_1 \wedge \neg g) \wedge (c_2 \vee \neg c_2 \wedge \neg g) \\ &= c_1 \wedge c_2 \vee \neg g \wedge (\neg c_1 \wedge \neg c_2 \vee \neg c_1 \wedge c_2 \vee c_1 \wedge \neg c_2) \\ &= c_1 \wedge c_2 \vee \neg g \wedge (\neg c_1 \vee \neg c_2) \\ &= c_1 \wedge c_2 \vee \neg g \wedge \neg(c_1 \wedge c_2) \\ &= \neg(c_1 \wedge c_2) \Rightarrow \neg g \\ &= \neg c_{12} \Rightarrow \neg g. \end{aligned}$$

Let $(n, m) \in \bar{\mathcal{E}}_1$. Then from the definition of $\bar{\mathcal{E}}_1$ we have that $c_1(n, m) = 0$. We have that

$$c_{12}(n, m) = c_1(n, m) \wedge c_2(n, m) = 0 \wedge c_2(n, m) = 0.$$

It follows from the definition of $\bar{\mathcal{E}}_{12}$ that $(n, m) \in \bar{\mathcal{E}}_{12}$. Similarly, we see that any element of $\bar{\mathcal{E}}_2$ is an element of $\bar{\mathcal{E}}_{12}$. \square

In this manner, we can combine as many required classifiers as we want and have at our disposal. With each classifier we add, we potentially decrease the set of remaining candidates that need to be checked by the exact algorithm. There is a trade-off between doing more work to evaluate more classifiers and having less instances that have to be processed by the exact algorithm [148].

7.3 Features and Classifiers for Culling Platoon Pairs

In order to apply the results from Section 7.2, we need to specify appropriate features and classifiers based on these features for the problem of identifying candidate platoon pairs. Once we know how to compute appropriate features that yield required classifiers or required sets of classifiers, we can use the results from Section 7.2 to execute the culling phase. The remaining candidate pairs are passed on to the exact algorithm to compute \mathcal{C} . Hence, we derive a selection of features and corresponding classifiers in this section. In Section 7.4, we demonstrate these classifiers and combinations of them in a simulation example.

The first feature projects the possible trajectories in time and space on a line, which yields an interval. Formally, we define this feature as follows.

Definition 7. Let $\mathbf{p} \in \mathbb{R}^3$ be a three dimensional vector that defines the orientation of the line on which the trajectories are projected to. Then the associated interval feature is defined as

$$\mathcal{I} = [\min_{\mathbf{v} \in \mathcal{R}}(\mathbf{p}^T \mathbf{v}), \max_{\mathbf{v} \in \mathcal{R}}(\mathbf{p}^T \mathbf{v})] \quad (7.1)$$

with

$$\mathcal{R} = \left\{ \left[\begin{array}{c} \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}[1]) \\ \underline{\mathbf{t}}[1] \end{array} \right], \dots, \left[\begin{array}{c} \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}[N^A]) \\ \underline{\mathbf{t}}[N^A] \end{array} \right], \right. \\ \left. \left[\begin{array}{c} \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}[1]) \\ \underline{\mathbf{t}}[1] \end{array} \right], \dots, \left[\begin{array}{c} \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}[N^A]) \\ \underline{\mathbf{t}}[N^A] \end{array} \right] \right\}. \quad (7.2)$$

This feature is illustrated in Figure 7.4. The projection vector \mathbf{p} is a design choice. Proposition 10 allows us to combine arbitrarily many classifiers based on this kind of feature with different \mathbf{p} .

Next, we establish that if for a pair of transport assignments the intervals do not overlap, the coordination function is equal to zero. This allows us to define a required feature based on the overlap between these intervals.

Proposition 11. Let (n, m) refer to a pair of transport assignments. Let $\mathcal{I}_n, \mathcal{I}_m$ be the interval features according to (7.1) for the two transport assignments. Then $\mathcal{I}_n \cap \mathcal{I}_m = \emptyset \Rightarrow C(n, m) = 0$.

Proof. According to Definition 6, $C(n, m) = 1$ implies that there must be indices a, b such that $\mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a]) = \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_m[b])$ and $[\underline{\mathbf{t}}_n[a], \bar{\mathbf{t}}_n[a]] \cap [\underline{\mathbf{t}}_m[b], \bar{\mathbf{t}}_m[b]] \neq \emptyset$, where $\mathbf{e}_n, \underline{\mathbf{t}}_n, \bar{\mathbf{t}}_n$ and $\mathbf{e}_m, \underline{\mathbf{t}}_m, \bar{\mathbf{t}}_m$ are the routes and time bounds of transport assignments n, m respectively. We have

$$[\underline{\mathbf{t}}_n[a], \bar{\mathbf{t}}_n[a]] \cap [\underline{\mathbf{t}}_m[b], \bar{\mathbf{t}}_m[b]] \neq \emptyset \Leftrightarrow \underline{\mathbf{t}}_n[a] \leq \bar{\mathbf{t}}_m[b] \wedge \underline{\mathbf{t}}_m[b] \leq \bar{\mathbf{t}}_n[a].$$




Figure 7.4: Illustration of the projection feature. The illustration on the left shows how two routes (solid lines) are projected onto a line in the direction of the vector \mathbf{p} . The borders of the intervals are indicated with dashed lines. For illustration purposes the third dimension is omitted here. The resulting intervals are shown in the gray box on the bottom of the illustration. In this case the projection of the two routes does not overlap and we can conclude that these routes have no road segments in common. On the right illustration, a third route is added that overlaps with the first two and we can see that also the corresponding intervals overlap.

Let

$$\begin{aligned}\mathbf{p} &= [\mathbf{p}[1], \mathbf{p}[2], \mathbf{p}[3]]^T, \\ \mathbf{P}_{\text{geo}}^S &= \mathbf{P}_{\text{geo}}^S(\mathbf{e}_n[a]) = \mathbf{P}_{\text{geo}}^S(\mathbf{e}_m[b]), \\ P^0 &= [\mathbf{p}[1], \mathbf{p}[2]] \mathbf{P}_{\text{geo}}^S.\end{aligned}$$

We have

$$\begin{aligned}\underline{\mathbf{t}}_n[a] \leq \bar{\mathbf{t}}_m[b] \wedge \underline{\mathbf{t}}_m[b] \leq \bar{\mathbf{t}}_n[a] \\ \Rightarrow \min(\mathbf{p}[3] \underline{\mathbf{t}}_n[a], \mathbf{p}[3] \bar{\mathbf{t}}_n[a]) &\leq \max(\mathbf{p}[3] \underline{\mathbf{t}}_m[b], \mathbf{p}[3] \bar{\mathbf{t}}_m[b]) \\ \Rightarrow \min(\mathbf{p}[3] \underline{\mathbf{t}}_n[a] + P^0, \mathbf{p}[3] \bar{\mathbf{t}}_n[a] + P^0) &\leq \max(\mathbf{p}[3] \underline{\mathbf{t}}_m[b] + P^0, \mathbf{p}[3] \bar{\mathbf{t}}_m[b] + P^0) \\ \Rightarrow \min\left(\mathbf{p}^T \begin{bmatrix} \mathbf{P}_{\text{geo}}^S \\ \underline{\mathbf{t}}_n[a] \end{bmatrix}, \mathbf{p}^T \begin{bmatrix} \mathbf{P}_{\text{geo}}^S \\ \bar{\mathbf{t}}_n[a] \end{bmatrix}\right) &\leq \max\left(\mathbf{p}^T \begin{bmatrix} \mathbf{P}_{\text{geo}}^S \\ \underline{\mathbf{t}}_m[b] \end{bmatrix}, \mathbf{p}^T \begin{bmatrix} \mathbf{P}_{\text{geo}}^S \\ \bar{\mathbf{t}}_m[b] \end{bmatrix}\right) \\ \Rightarrow \min_{\mathbf{v} \in \mathcal{R}_n} (\mathbf{p}^T \mathbf{v}) \leq \max_{\mathbf{v} \in \mathcal{R}_m} (\mathbf{p}^T \mathbf{v}),\end{aligned}$$

with $\mathcal{R}_n, \mathcal{R}_m$ as in (7.2) for transport assignment n, m , respectively. Similarly, by swapping n and m , we can show that the conditions of the proposition imply that

$$\min_{\mathbf{v} \in \mathcal{R}_m} (\mathbf{p}^T \mathbf{v}) \leq \max_{\mathbf{v} \in \mathcal{R}_n} (\mathbf{p}^T \mathbf{v}).$$

The above two conditions combined imply that $\mathcal{I}_n \cap \mathcal{I}_m \neq \emptyset$. Thus

$$C = 1 \Rightarrow \mathcal{I}_n \cap \mathcal{I}_m \neq \emptyset,$$

or equivalently

$$\mathcal{I}_n \cap \mathcal{I}_m = \emptyset \Rightarrow C = 0.$$

□

Next, we introduce a binary feature that leads to a required classifier. This feature is based on the orientations of the individual road segments in a route. It is only useful if all segments in a route point approximately from start to goal location. Later on, we address the problem of outliers. Here, we derive a set of required classifiers each based on a binary feature from the orientation. The orientation $\Theta(e) \in [0, 2\pi]$ of a road segment with road segment identifier $e \in \mathcal{E}_r$ is the angle in polar coordinates of the vector $\mathbf{P}_{\text{geo}}^{\text{E}}(e) - \mathbf{P}_{\text{geo}}^{\text{S}}(e)$. We choose a partition of the interval $[0, 2\pi]$. Each element of the partition is related to one binary feature, which holds true if the orientation of at least one road segment in the route falls in the range of that element. When two routes overlap, there must be at least one road segment that has the same orientation. Figure 7.5 illustrates the classifier.

Proposition 12. *Let (n, m) refer to the pair of transport assignments. Let $\bar{\mathcal{P}}$ be a partition of $[0, 2\pi]$. If there is no element $I \in \bar{\mathcal{P}}$ and route segment identifiers in the routes of the transport assignments $\mathbf{e}_n[a]$, $\mathbf{e}_m[b]$ such that $\Theta(\mathbf{e}_n[a]) \in I$ and $\Theta(\mathbf{e}_m[b]) \in I$, then $C(n, m) = 0$.*

Proof. According to Definition 6, $C(n, m) = 1$ implies that there must be indices a, b such that $\mathbf{e}_n[a] = \mathbf{e}_m[b]$, where $\mathbf{e}_n, \mathbf{e}_m$ are the routes of transport assignment n, m respectively. For these it holds that $\Theta(\mathbf{e}_n[a]) = \Theta(\mathbf{e}_m[b])$. Since $\bar{\mathcal{P}}$ is a partition of $[0, 2\pi]$ and $\Theta(\mathbf{e}_n[a]) \in [0, 2\pi]$, there must be $I \in \bar{\mathcal{P}}$ with $\Theta(\mathbf{e}_n[a]) \in I$. Since $\Theta(\mathbf{e}_m[b]) = \Theta(\mathbf{e}_n[a])$, it follows that also $\Theta(\mathbf{e}_m[b]) \in I$. The proof follows from contradiction. □

In the following, we discuss how we can make the orientation-based classifier more efficient if we can disregard routes that overlap only over a short distance. Apart from the direct reduction in true positives, this approach will also reduce the false-positive rate of the classifiers, since some outlier route segments can be disregarded.

In order to cover the notion that there must be a minimum overlap in routes to be considered, we extend the definition of the coordination function (Definition 6).

Definition 8 (Minimum Distance Coordination Function).

A coordination function $C : \mathcal{N}_c \times \mathcal{N}_c \rightarrow \{0, 1\}$ according to Definition 6 requires minimum distance l_{\min} if the following properties hold: if for a pair (n, m) we have $C(n, m) = 1$, there must be a set of pairs of indices \mathcal{A} such that for all $(a, b) \in \mathcal{A}$ it holds that $\mathbf{e}_n[a] = \mathbf{e}_m[b]$, and $[\underline{\mathbf{t}}_n[a], \bar{\mathbf{t}}_n[a]] \cap [\underline{\mathbf{t}}_m[b], \bar{\mathbf{t}}_m[b]] \neq \emptyset$ and $[\underline{\mathbf{t}}_n[a+1], \bar{\mathbf{t}}_n[a+1]] \cap [\underline{\mathbf{t}}_m[b+1], \bar{\mathbf{t}}_m[b+1]] \neq \emptyset$.

Furthermore, we require

$$\sum_{(a,b) \in \mathcal{A}} \|\mathbf{P}_{\text{geo}}^{\text{E}}(\mathbf{e}_n[a]) - \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a])\|_2 \geq l_{\min}.$$




Figure 7.5: Illustration of the classifier based on the orientation. In this example the interval $[0, 2\pi]$ is partitioned into 20° intervals. The arrows on the left symbolize road segments of a route. The elements of the partition for which at least one road segment in the route has the same orientation are filled with gray.

We adapt the orientation-based classifier (Proposition 12) to exclude road segments of a total length less than l_{\min} . The approach is to calculate the fraction of route length that lies in each element of the partition. We can ignore the intersection with some elements of the partition as long as the lengths of the road segments whose orientation is contained in these elements sums up to a value less than $l_{\min}/2$. Figure 7.6 illustrates this approach.

Proposition 13. *Let (n, m) refer to a pair of transport assignments. Let $\bar{\mathcal{P}}$ be a partition of $[0, 2\pi]$. Let $\mathcal{I}_n \subseteq \bar{\mathcal{P}}$ and let $\bar{\mathcal{E}}_n \subseteq \mathcal{E}_n$, where*

$$\mathcal{E}_n = \{\mathbf{e}_n[i] : i \in \{1, \dots, N_n^A\}\},$$

such that for all $e \in \bar{\mathcal{E}}_n$, it holds that there exists $I \in \mathcal{I}_n$ with $\Theta(e) \in I$ and we have

$$\sum_{e \in \mathcal{E}_n \setminus \bar{\mathcal{E}}_n} \|\mathbf{P}_{\text{geo}}^E(e) - \mathbf{P}_{\text{geo}}^S(e)\|_2 < l_{\min}/2.$$

Similarly, by replacing n by m , we define \mathcal{I}_m for transport assignment m . If $\mathcal{I}_n \cap \mathcal{I}_m = \emptyset$, then $C(n, m) = 0$ with C according to Definition 8.

Proof. If $C(n, m) = 1$, then we have a set of pairs of indices \mathcal{A} such that for all $(a, b) \in \mathcal{A}$ it holds that $\mathbf{e}_n[a] = \mathbf{e}_m[b]$. Thus, it also holds that $\Theta(\mathbf{e}_n[a]) = \Theta(\mathbf{e}_m[b])$. Since $\bar{\mathcal{P}}$ is a partition of the image of $\Theta(\cdot)$, there is exactly one element $I \in \bar{\mathcal{P}}$ with

$\Theta(\mathbf{e}_n[a]) \in I$, and since $\Theta(\mathbf{e}_n[a]) = \Theta(\mathbf{e}_m[b])$, we have

$$\Theta(\mathbf{e}_n[a]) \in I \Leftrightarrow \Theta(\mathbf{e}_m[b]) \in I.$$

Furthermore, we have from Definition 8 that

$$\sum_{(a,b) \in \mathcal{A}} \|\mathbf{P}_{\text{geo}}^{\text{E}}(\mathbf{e}_n[a]) - \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a])\|_2 \geq l_{\min}.$$

Let $\bar{\mathcal{A}}_n$ be a set of the indices of the head nodes of road segment identifiers in $(\mathcal{E}_n \cap \mathcal{E}_m) \setminus \bar{\mathcal{E}}_n$ paired with the corresponding indices in route \mathbf{e}_m , with $\mathcal{E}_n, \mathcal{E}_m, \bar{\mathcal{E}}_n$ as defined in the proposition. These are the pairs of indices of the road segment identifiers in the common part of the route that are ignored in transport assignment n . Similarly, let $\bar{\mathcal{A}}_m$ be the index pairs that are excluded due to transport assignment m . We need to show now that \mathcal{A} is not empty without the pairs in $\bar{\mathcal{A}}_n$ and $\bar{\mathcal{A}}_m$, or in other words, that even if the features for either route ignore up to $l_{\min}/2$ of the common part of the route, there are still road segment identifiers left that let the set of classifiers indicate that the routes intersect. We have from the assumptions made in the proposition

$$\begin{aligned} \sum_{(a,b) \in \bar{\mathcal{A}}_n} \|\mathbf{P}_{\text{geo}}^{\text{E}}(\mathbf{e}_n[a]) - \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a])\|_2 &< l_{\min}/2, \\ \sum_{(a,b) \in \bar{\mathcal{A}}_m} \|\mathbf{P}_{\text{geo}}^{\text{E}}(\mathbf{e}_n[a]) - \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a])\|_2 &< l_{\min}/2, \end{aligned}$$

and from Definition 8 that

$$\sum_{(a,b) \in \mathcal{A}} \|\mathbf{P}_{\text{geo}}^{\text{E}}(\mathbf{e}_n[a]) - \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a])\|_2 \geq l_{\min}.$$

Thus,

$$\sum_{(a,b) \in \mathcal{A} \setminus (\bar{\mathcal{A}}_n \cup \bar{\mathcal{A}}_m)} \|\mathbf{P}_{\text{geo}}^{\text{E}}(\mathbf{e}_n[a]) - \mathbf{P}_{\text{geo}}^{\text{S}}(\mathbf{e}_n[a])\|_2 > 0,$$

and since this is a sum over positive elements, we deduce that $\mathcal{A} \setminus (\bar{\mathcal{A}}_n \cup \bar{\mathcal{A}}_m) \neq \emptyset$. But then there is $I \in \bar{\mathcal{P}}$ and $(a, b) \in \mathcal{A} \setminus (\bar{\mathcal{A}}_n \cup \bar{\mathcal{A}}_m)$ such that

$$\Theta(\mathbf{e}_n[a]) = \Theta(\mathbf{e}_m[b]) \in I,$$

and thus $\mathcal{I}_n \cap \mathcal{I}_m \neq \emptyset$. By contraposition it follows that $\mathcal{I}_n \cap \mathcal{I}_m = \emptyset \Rightarrow C(n, m) = 0$. \square

Propositions 11 and 13 define two families of classifiers that can be used to filter the set of candidate platoon pairs efficiently before applying the exact algorithm. These and potentially other suitable classifiers can be combined in various ways according to Propositions 9 and 10 in Section 7.2. Such a combined classifier creates a smaller candidate set but with each classifier added, the computational complexity increases. A good combination needs to be selected for applying this method and depends on the distribution of assignments that are coordinated.




Figure 7.6: This figure illustrates how the performance of the orientation classifier can be improved when overlaps of length less than l_{\min} can be excluded. The figure shows the histogram of two routes. The routes are sketched on the top of the figure. There are two elements in the partition that contain orientations from both routes corresponding only to a small fraction of the total route length. The classifier according to Proposition 12 indicates an intersection between these two route whereas the classifier according to Proposition 13 can exclude the few road segments with similar orientation.

7.4 Simulations

In this section, the method derived in this chapter is demonstrated in a realistic scenario. We show that the application of 6 classifiers can rule out 99% of the transport assignment pairs, leaving only 1% for the computationally expensive exact algorithm. The simulation setup is as follows. The start and goal locations are sampled randomly with probability proportional to an estimate of the population density in the year 2000 [204]. We limit the area to a large part of Europe, which is shown in Figure 7.7.

We calculate shortest routes with the Open Source Routing Machine [151]. If the route is longer than 400 kilometers, a 400 kilometers long subsection of the route is randomly selected. The maximum speed is constant and equals 80 km/h and we do not consider a minimum speed in the simulations. We set the start times t^S of half the assignments to 0 and sample the start times of the remaining assignments uniformly in an interval of 0 to 24 h. The first half is to account for assignments that are currently on the road while the other half is to account for assignments that are scheduled to depart later. The deadlines t^D are set in such a way that the




Figure 7.7: Population density map from which the start and goal locations are sampled. The brighter the pixel, the larger the population density in that area. Some areas outside Europe and areas without population are shown in blue. The horizontal axis in the image corresponds to the longitude and the vertical axis to the latitude. Therefore the shape looks different from the more familiar Mercator projection.

interval $\bar{\mathbf{t}}[a] - \underline{\mathbf{t}}[a] = 0.5h$ where a is any valid index. We consider the minimum length that two assignments have to overlap to be considered for platooning, l_{\min} , to be 20 km.

We implemented all features and corresponding classifiers that are described in Section 7.3, i.e., interval projection (Proposition 11) and minimum distance orientation partition (Proposition 13). Note that Proposition 12 is a special case of Proposition 11 with $l_{\min} = 0$. For interval projection we tested vectors of the form

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -\cos(\alpha) \\ \frac{-1}{\cos(50^\circ)} \sin(\alpha) \\ \frac{v_{\max} 180^\circ}{6371\pi} \end{bmatrix},$$

with $\alpha = 0, \pi/4, \dots, 7\pi/4$. The positions $\mathbf{P}_{\text{geo}}^{\text{S}}$, $\mathbf{P}_{\text{geo}}^{\text{E}}$ are expressed here as latitude and longitude and measured in degrees. The vectors parametrized by α are approximately orthogonal to a trajectory at maximum speed at the latitude of 50 degrees with heading angle α and should work well for trajectory pairs that have similar orientation, that cover the same area, and that are only separated by a small time margin. We refer to the corresponding classifiers in the following discussion as $c_{100}, c_{010}, c_{001}, c_{110}, c_{-110}, c_{\alpha 0}, \dots, c_{\alpha 7}$ respectively. For the orientation-based clas-

Table 7.1: Number of positives for different classifiers

None	499,500	c_{-110}	108,403	$c_{\alpha 4}$	134,019
c_{100}	104,380	$c_{\alpha 0}$	129,282	$c_{\alpha 5}$	107,287
c_{010}	101,542	$c_{\alpha 1}$	103,240	$c_{\alpha 6}$	105,883
c_{001}	208,896	$c_{\alpha 2}$	103,453	$c_{\alpha 7}$	109,934
c_{110}	98,343	$c_{\alpha 3}$	109,626	c_o	453,246

sifier, we use 100 equally sized cells to partition $[0, 2\pi]$. For each cell, the fraction of the route distance that falls in this cell is computed. Matches up to $l_{\min}/2$ starting in ascending order of route distance contained in the cells are excluded. We refer to this classifier as c_o .

This simulation focuses on demonstrating that the culling phase is able to filter out a significant amount of assignments before they are passed on to the exact algorithm. Therefore, we do not focus on optimizing the implementation for speed and refrain from reporting running times of the simulations as they might be misleading and we know from related work [148] that these computations can be performed fast enough for the problem at hand if the false-positive rate of the classifiers is small.

We test 1000 transport assignments. All classifiers are evaluated in parallel. Next, the sequence of classifiers that filters the most assignments at every stage is computed. The number of positives for each classifier is listed in Table 7.1. Figure 7.8 shows the number of remaining pairs at each stage, the ground truth, and the sequence of classifiers for this sample. The optimization of the classifier order would typically be done when the system is designed and is to some extent specific for the exact transport setting. In a running platoon coordination system the order in which classifiers are applied would remain fixed.

We can see in Figure 7.8 that two classifiers, c_{110} and $c_{\alpha 7}$, combined are able to reduce the number of pairs by one order of magnitude. The first classifier, c_{110} , only takes into account longitude and latitude of the routes. The second one, $c_{\alpha 7}$ is orthogonal to the first one, c_{110} , in the plane but also takes into account timing. The third classifier, $c_{\alpha 3}$, is also of the projection type, which is able to identify that a pair of assignments cannot platoon if they are geographically close but differ in timing, and it covers the opposite orientation compared to the previous classifier. The fourth classifier, c_{100} , covers a third direction in the plane. It is interesting to see that the fifth classifier, c_o , is the orientation-based classifier. Alone, it performs much worse than the other classifiers as can be seen in Table 7.1. Two transport assignments that take the same route in opposite directions and that “meet” on the way are impossible to identify as a negative with the projection based classifiers. The orientation-based classifier might be able to achieve that. The classifier that only takes into account start and arrival time, c_{001} , is selected last, since most of the




Figure 7.8: The number of remaining pairs when the classifiers are consecutively applied from left to right. The order the classifiers are chosen in a way that each stage removes as many pairs as possible. The classifier applied at each stage is indicated on the horizontal axis. The dashed line shows the ground truth from the exact algorithm.

cases it rules out are already covered by the classifiers $c_{\alpha 0}, \dots, c_{\alpha 7}$, and also because half the assignments start at the same time. We see that the benefit from adding more classifiers diminishes quickly as classifiers are added. All classifiers combined can reduce the number of pairs by two orders of magnitude and get within one order of magnitude from the ground truth. The false-positives are mostly very curvy routes that intersect geographically and are separated little in time in the area of the intersection. To be able to correctly identify such pairs as negatives is often not possible with the features presented in this chapter. Figure 7.9 shows an example of a false-positive. We get consistent results for different runs of the simulation.

7.5 Summary

Comparing the routes and the time bounds of a large number of assignments in order to find candidates for platooning is computationally expensive. This chapter presents a more efficient approach, which is to narrow down the set of candidates based on features. A feature is low dimensional data like a boolean truth value or




Figure 7.9: Example of a false-positive. The black lines show the routes of the two transport assignments. The start locations are marked with stars. The two routes do not overlap. However, the routes cannot be separated by a hyperplane, and since both routes are quite curvy the orientation-based classifier cannot conclude that these route do not overlap.

an interval that can be efficiently processed in the form of classifiers. The smaller set of candidates is then used as an input to computing fuel-efficient platoon plans for all vehicles that are coordinated.

We show how several classifiers can be combined to get even smaller sets of candidates, and, in some cases, classifiers have to be combined to be able to conclusively rule out that a pair of transport assignments is able to platoon. Two features and corresponding classifiers are derived. One is based on the projection of the route and time bounds onto a line. The other classifier is based on the intersection of route segment orientations with a partition of all possible orientations. The performance of orientation-based classifier is improved by assuming that the common part of the routes of two assignments must have a minimum length to be relevant for platooning. Simulations indicate that the method developed in this chapter can significantly narrow down the set of candidate platoon pairs.

Chapter 8

Evaluation

IN this chapter, we present three simulation studies and an experimental study to illustrate and verify the coordinated platooning system. Due to complexity of the system and its combinatorial nature, simulations and experiments are an important tool for evaluating the proposed coordinated platooning system. We use simulations to understand the large scale behavior of the system sacrificing some detail on vehicle level. Experiments are close to a real world use of the system. However, they are currently limited for practical reasons to a rather small number of vehicles.

In Section 8.1, we analyze the plan composition with Algorithm 1 on an artificially generated road network. Section 8.2 presents a simulation scenario in which assignments are generated according to a population density map and routes are computed in a real road network. In Section 8.3, the dynamic update of plans is simulated and sensitivity to disturbances is analyzed. In Section 8.4, we analyze results from an experiment in which three heavy-duty vehicles were coordinated to form a platoon on public highways.

8.1 Simulations on an Artificial Road Network

In this simulation study, we use an artificial road network, which can be useful for studying systematic parameter variations [79]. The road network was generated randomly by sampling 100 location points uniformly in a square of side length 800 km. All combinations of these locations were sorted by their Euclidean distance. Then, starting from the combination with the shortest distance, the combinations of locations were connected by two road segments (one in each direction) with length according to the Euclidean distance, if there had not already been a path between the two locations that was at most 1.5 times longer than the Euclidean distance. The road network is shown in Figure 8.1.

We evaluate the algorithm to select coordination leaders, Algorithm 1 with greedy and random node selection with Monte Carlo simulations. In the figures




Figure 8.1: Randomly generated road network used in the simulation. Start/destination nodes are marked by a blue circle.

and the following discussion we use “greedy”/“random” to refer to selecting the node with largest objective improvement Δu or a random node in each iteration with $\Delta u > 0$. Default and adapted plans are computed according to Section 4.3. We consider a nominal speed of $v_0 = 80$ km/h according to which arrival times are set, i.e., the constant speed of a coordination leader m is $v_m = 80$ km/h. For the linear affine fuel model fuel model, we consider that $F^0 = 1$, $F^1 = 1/(80 \text{ km/h})$, and $F_p^0 = 0.9F^0$, $F_p^1 = 0.9F^1$. We get that the fuel optimal rendezvous speeds according to (4.4) are 80 ± 35.777 km/h = (44.223 km/h, 115.777 km/h).

We conduct simulations for different numbers of assignments $|\mathcal{N}_c|$. For each number of assignments $|\mathcal{N}_c|$, 100 simulations are conducted. The starting times are sampled uniformly in the interval of $[0, 1]$ h. Start and destination nodes are randomly selected from a subset of 10 nodes marked with blue circles in Figure 8.1. The arrival times are calculated assuming a speed of 80. We set $v_{\min} = 70$ km/h, $v_{\max} = 90$ km/h. In each simulation, we run Algorithm 1 using either greedy or random node selection. Figure 8.2 shows a plot of the fuel savings compared to the case




Figure 8.2: Average relative fuel savings for different numbers of assignments $|\mathcal{N}_c|$

where all vehicles travel at a speed of 80 km/h and do not platoon. Additionally, we calculate fuel savings that would result from spontaneous platooning according to the following scheme. We assume that all vehicles travel with the nominal speed of 80 km/h. For each segment in the road network and for all vehicles traversing a particular road segment, we collect the times when the vehicles traverse the segment. Then, in the order of these points in time, we group vehicles into platoons in such a manner that the difference in the traversal time within the platoon is at most $0.01 \text{ h} = 36 \text{ s}$. We assume that these platoons can platoon over the whole road segment with no coordination phase. Figure 8.3 shows $|\Delta d^S|$ averaged over all coordination leader coordination follower pairs within one simulation and over the simulations for a specific $|\mathcal{N}_c|$. The quantity $|\Delta d^S|$ is a measure for how much the coordination follower adapts. Furthermore, we conduct simulations for different sizes of the band $[v_{\min}, v_{\max}]$ around the nominal speed of 80 km/h in which vehicles can select their rendezvous speed. $|\mathcal{N}_c|$ is fixed to 400. Figure 8.6 shows the relative fuel savings for different sizes of the band averaged over 100 simulations for each size.

In Figure 8.2, we can see that the relative fuel savings increase with the number of assignments. The relative fuel savings increase quickly with $|\mathcal{N}_c|$ for small values of $|\mathcal{N}_c|$ and then increase slowly for larger values. While the relative fuel savings with total gain keep increasing for large values of $|\mathcal{N}_c|$, they are almost constant for pairwise gain. Note that the relative fuel savings in the simulations are upper




Figure 8.3: Average $|\Delta d^S|$ for different numbers of assignments $|\mathcal{N}_c|$

bounded by 10%. To reach this upper bound, each vehicle needs to be a platoon follower during its entire journey. The difference between greedy and random node selection is small. For small $|\mathcal{N}_c|$, the ratio between the relative fuel savings due to the coordination algorithm and those due to spontaneous platooning is relatively large and gets smaller for increasing $|\mathcal{N}_c|$. Figure 8.3 shows that the average $|\Delta d^S|$ is relatively large for small $|\mathcal{N}_c|$ and that it drops quickly to smaller values as $|\mathcal{N}_c|$ increases. Figure 8.4 shows the average number of coordination leaders for different $|\mathcal{N}_c|$. We see that the number increases with $|\mathcal{N}_c|$ sub-linearly. The number of iterations until the algorithm terminates, shown in Figure 8.5, is almost proportional to the number of coordination leaders, the difference being that random node selection leads to a significantly higher number of iterations than greedy node selection. Figure 8.6 shows that fuel savings quickly increase for small $v_{\max} - v_{\min}$ and only increase moderately for larger bands.

We can conclude that the coordination of platooning is crucial for a small number of vehicles and can significantly improve the overall fuel savings for a large number of vehicles compared to spontaneous platooning. With coordination, already a small number of platooning-enabled vehicles can achieve significant reduction in fuel consumption. Random and greedy node selection give similar results but random node selection leads to more iterations. However, random node selection might be interesting for a distributed and parallel implementation since it does not require coordination amongst all nodes to determine which node is updated.




Figure 8.4: Average number of coordination leaders (CLs) for different numbers of assignments $|\mathcal{N}_c|$

Relatively small adjustments of the speed are sufficient to achieve most of the fuel savings possible. One should also keep in mind that fuel consumption per distance traveled is highly non-linear over the whole range of speeds a vehicle can attain and a first order fuel model is only accurate in a small range of speeds. Fuel consumption per distance traveled changes with speed in such a way that we expect large adjustments of the speed to be even less beneficial than predicted by the linear model.




Figure 8.5: Average number of iterations until the algorithm terminates for different numbers of assignments $|\mathcal{N}_c|$




Figure 8.6: Average relative fuel savings for different $v_{\max} - v_{\min}$ with $K = 400$

8.2 Simulations on a Real Road Network

In this section, we evaluate the computation of platoon plans on a real road network using Monte Carlo simulations. Additionally, we test the joint platoon plan optimization as discussed in Section 6.4. We show that the coordination of heavy-duty vehicle platooning can lead to significant reductions in fuel consumption compared to the current situation where vehicles do not platoon, as well as compared to spontaneous platooning where vehicles only form platoons if they happen to be in the vicinity of another.

We generate transport assignments randomly. The start and goal locations are sampled within mainland Sweden. The probability of an assignment starting or ending at a particular location is proportional to the population density [204], see Figure 8.7. The resolution is 0.1 degrees in longitude and latitude and the road network node that is closest to the sampled coordinate is chosen. We calculate the routes with the Open Source Routing Machine [151]. Assignments for which no route can be found are disregarded. If the route is longer than 400 kilometers, a 400 kilometers long subsection of the route is randomly selected. This is to take into account that merge points too far from the current position should not be considered for coordination since the uncertainty becomes too large due to traffic, new assignments, and rest periods of the driver. Start locations along the route are considered to take into account that the platoon coordinator frequently re-plans for assignments that are already en route and suspended for the driver to take a rest.

The fuel model is an affine approximation around 80 km/h of the analytical fuel model in [35]. We have for the fuel per distance traveled in kilograms diesel per meter

$$\begin{aligned} f_0(v) &= 8.4159 \cdot 10^{-6}v + 4.8021 \cdot 10^{-5} \\ f_p(v) &= 5.0495 \cdot 10^{-6}v + 8.5426 \cdot 10^{-5}. \end{aligned}$$

According to this model, the relative reduction in fuel consumption of a platoon follower is 15.9 percent at a speed of 80 km/h.

We consider a default speed of 80 km/h and we assume that the speed can be freely chosen between $v_{\min} = 70$ km/h and $v_{\max} = 90$ km/h throughout the entire journey. We sample the start time of the assignments uniformly in an interval of 2 hours and compute the arrival deadlines according to the default speed.

The pairwise plans are such that vehicles platoon as long a distance as possible. Once a coordination follower splits up from the coordination leader, it drives fast enough to arrive in time at its destination and at least at default speed. The split points are selected in a way that vehicles arrive in time. Thus, vehicles are guaranteed to meet their deadlines and the initial value for the joint platoon plan optimization is feasible. Figure 8.8 shows an example of the routes of a coordination leader and its coordination followers and where the coordination followers join and leave the platoon.




Figure 8.7: Population density map from which the start and goal locations are sampled. The brighter the pixel, the larger the population density in that area. Areas not belonging to mainland Sweden are shown in blue. The horizontal axis in the image corresponds to the longitude and the vertical axis to the latitude. Therefore the shape looks different from the Mercator projection commonly used to draw maps.

Analogous to Section 8.1, we compare our proposed platoon coordinator to fuel savings that arise from spontaneous platooning, i.e., that vehicles happen to get into each others vicinity and then spontaneously form platoons. To this end, we collect all the road segment arrival times according to the default plans for each road segment in the scenario. We sort these times and collect them in ascending order in groups of at most one minute difference in their road segment arrival time. We assume that each of these groups forms a platoon driving at default speed and that the default trajectory is not altered by the platooning. This is a generous estimate, since it neglects any kind of coordination effort, which would be present for time gaps up to one minute.

We implemented the platoon coordination algorithms in Python 2.7 and use CVXOPT [22] for convex optimization. The execution of Algorithm 1 takes less than a second for 2000 transport assignments. Even faster computation times could be achieved by optimizing the implementation.

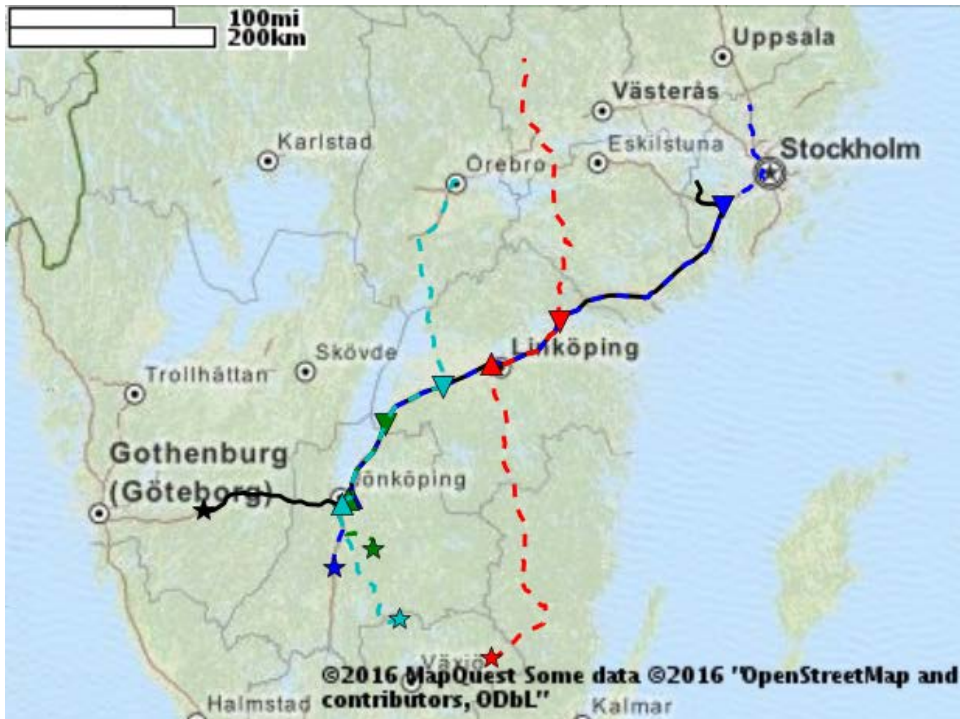


Figure 8.8: The routes of a platoon coordinator with four coordination followers. The route of the coordination leader is shown in black, the routes of the coordination followers are dashed. The beginning of a route is marked with a star. The merge point of a follower is indicated with an upwards-facing triangle and the split point with a downwards-facing triangle.

Each simulation consists of the following steps:

1. Random generation of transport assignments,
2. Computation of routes and default plans,
3. Computation of the coordination graph,
4. Computation of coordination leaders according to Section 6.3,
5. Joint platoon plan optimization according to Section 6.4.

We evaluate how different numbers of assignments affect the amount of platooning and the fuel savings relative to the default plans. For comparison, we compute




Figure 8.9: This plots visualizes the adjacency matrix of a coordination graphs with 100 assignments. Nonzero entries are indicated with a black or a red dot, each corresponding to an edge in the coordination graph. Edges whose corresponding plans are selected by the Algorithm 1 correspond to the red dots.

the fuel savings of spontaneous platooning. We run Algorithm 1 with greedy node selection. The results are averaged over 150 simulation runs.

Figure 8.9 visualizes an example coordination graph. In addition, it shows which assignments are selected in step 4). We can see that only a small fraction of assignment pairs can save fuel by forming a platoon. As the number of assignments grows, more opportunities are available for each assignment, which can translate into larger fuel savings [136].

Figure 8.10 shows the effect on the fuel savings when the numbers of transport assignments that are coordinated is varied. It is possible to make a number of observations based on these data. First of all, the fuel savings increase rapidly with the number of transport assignments when the absolute number of assignments is small. As more and more assignments are added, this trend stagnates and the relative fuel savings increase only slowly. Ideally, this should approach asymptotically the maximum fuel savings of 15.9% as the number of transport assignments goes to infinity, since then virtually every vehicle is a platoon follower for its entire journey. We can see a clear improvement in the fuel savings by the joint optimization of the platoon plans. Spontaneous platooning gives fuel savings that are less than half of what can be achieved by coordination. Also bear in mind that this is a generous

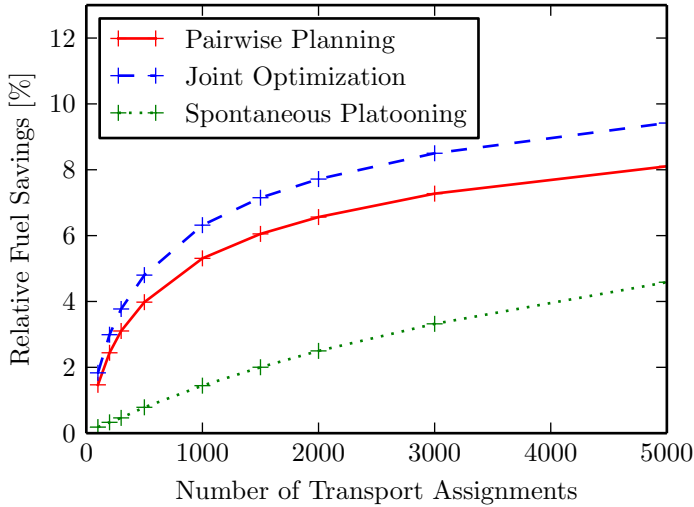


Figure 8.10: The relative fuel savings due to platooning compared to the default plans with varying numbers of assignments. The legend entries “Joint Optimization”/“Pairwise Planning” refer to the relative fuel savings with/without the joint optimization of the platoon plans. “Spontaneous Platooning” are the relative fuel savings based on the estimate of fuel savings due to spontaneous platooning.

estimate of fuel savings by spontaneous platooning implying that the real difference would probably be even larger.

We conclude that coordinated platooning can yield significant fuel savings and that coordination is crucial in leveraging these savings. For 2000 transport assignments starting over the course of two hours, we get 7.6% reduction in fuel consumption. A number of 2000 vehicles starting in that time interval on the area of mainland Sweden is a realistic number. The total distance traveled in the simulated scenario is in the same order of magnitude as the total distance traveled by domestic road freight transport in Sweden within two hours, assuming that traffic volume is equally spread over the year [5]. The density of the road freight traffic that is simulated is only a fraction of the total road freight traffic in countries with high population density.

Figure 8.11 shows how the distribution of platoon sizes changes with the number of transport assignments. Here, we evaluate both greedy and random node selection in Algorithm 1. We can see that the larger the number of transport assignments, the more distance is traveled in large platoons. For 2000 assignments, over half the distance traveled is in a platoon. Less than 5% of the distance is traveled in platoons with more than eight vehicles. This is promising since large

platoons might be difficult to control and thus the platoon coordinator would have to either prevent planning for larger platoons or during the sections of the route where the planned platoon size becomes too large, the platoon has to be split into smaller platoons. Since these large platoons only account for a small fraction of the distance traveled, this would not have too large an impact on the total fuel savings. The largest platoon formed has 28 vehicles. A noticeable effect occurs at a number of 200 transport assignments when more distance is traveled in relatively large platoons compared to the distribution with a number of 300 transport assignments. It seems that some kind of phase transition occurs at these points, where enough assignments are in the system to go from one coordination leader with many followers to having several coordination leaders that are better suited for their followers. This phenomenon is much more pronounced when greedy nodes selection is used in Algorithm 1 compared to random node selection.

The simulations show that computing plans for a large number of vehicles to form platoons is feasible with the methods outlined in this thesis. It motivates that real-time platoon coordination enables significant reductions in fuel consumption and might be the key to leveraging the full potential of heavy-duty vehicle platooning.




Figure 8.11: This figure shows the distribution of platoon sizes per distance traveled over the number of assignments in percent. The upper plot shows the results of greedy node selection whereas the lower plot shows those of random node selection in the clustering algorithm. To the right, the size of platoon is indicated for a platoon size up to eight. So, when the distance between the first and the second boundary from below is for instance at 20%, it means that 20% of the distance was traveled as member of a platoon of size 2.




Figure 8.12: The routes used are shown as black lines. The area considered for the road network and for assignment generation is indicated with a dashed blue line.

8.3 Dynamic Simulations and Sensitivity Analysis

This section presents simulation results with over 3000 active vehicles in the German road network. Here, the dynamic nature of the system is simulated, with new assignments becoming known to the system over time and plans being updated at regular intervals. The influence of several design parameters and disturbances on the proposed coordinated platooning system is tested. The simulations focus on behavior on strategic level.

Scenario Description

Different scenarios are created to test the coordinated platooning system. A scenario consists of assignments, design parameters, and disturbances. To generate assignments, pairs of start and goal locations are sampled randomly from a population density map in the geographical area of Germany shown in Figure 8.12. For each pair, a route is computed based on Openstreetmap data for the German road network as way to mimic the characteristics of real world traffic patterns. Only roads labeled “motorway”, “trunk”, “primary”, and “secondary”, as well as the

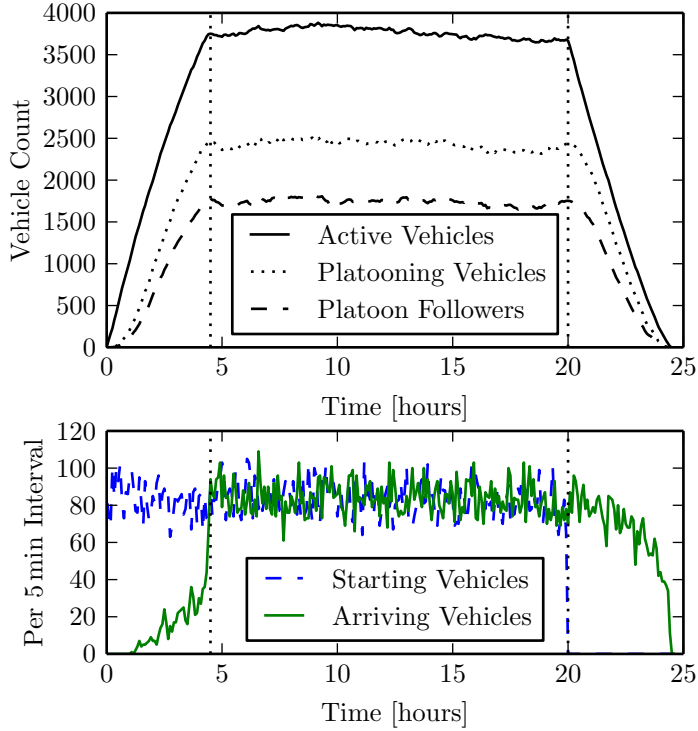


Figure 8.13: The upper plot shows the number of active vehicles, the number of vehicles in a platoon and the number of platoon followers. The lower plot shows the number of starting and arriving vehicles per 5 minute interval. Results are evaluated in the time window indicated by the dotted vertical lines.

corresponding segments are used, with a strong speed penalty on roads classified as “secondary”. Part of the route at the beginning and the end consisting of road segments labeled “primary” and “secondary” are removed. This is to take into account that platooning is likely to be limited to such roads. Assignments with less than 80 km traveled distance are rejected, and a random subroute 4.5×80 km is selected for assignments longer than 4.5×80 km. This models that a driver has to rest after 4.5 h of driving. The continuation of such a trip after a rest period would be registered by the fleet management system as a new assignment. The start of the subroute is then used as the start location \mathcal{P}^S and the end as the destination \mathcal{P}^D . Start times t^S are sampled uniformly over an interval of 20 h, in order to study the steady state behavior of the system. Deadlines t^D are computed assuming a fixed speed of 80 km/h.

We consider that the maximum speed profile is $\bar{v} \equiv 90$ km/h according to the formalism introduced in Section 4.4. Default plans are computed according to Definition 3 with $\sigma_d = 8/9$, i.e., the default speed is 80 km/h. Adapted speed profiles are computed according to Definition 3 and in way that the platooning distance is maximized, with $\underline{\sigma} = 7/9$, i.e., the minimum speed during the merge phase is 70 km/h. These are typical speeds of heavy vehicles on European overland routes. No adaptation of the start time is considered. Platoon plans are updated at a regular interval T_{upd} . The first time an assignment is considered in the planning process is at $t^S - T_{\text{hor}}$, where t^S is the assignments start time and T_{hor} is a design parameter we refer to as preview horizon. Should no plan be computed before the vehicle starts, the vehicle follows the default plan until the next update. The model for the fuel consumption estimation is a nonlinear analytical fuel model used in [35]. It is plotted in Figure 6.1. It gives a fuel consumption reduction of 12% at 80 km/h for each trailing vehicle in the platoon compared to traveling alone, assuming a 33% air-drag reduction for a fully loaded trailing vehicle. For the selection of coordination leaders, the Algorithm 1 with greedy leader selection is used and initialized with coordination leaders from the previous iteration.

There are two mechanisms to model disturbances. The first mechanism is to reduce the maximum speed to a value between 80–90 km/h. A platoon merge is considered to succeed if the difference in arrival time at the planned merge location is less than $0.5 \text{ km} / 80 \text{ km/h} = 22.5 \text{ s}$. This mechanism is intended to simulate smaller disturbances coming from traffic, incorrect information of the vehicles performance, etc. The other mechanism is that some vehicles end their trips before reaching the destination. This is to model that vehicles actually stop due to unexpected breaks by the driver, technical issues, but also that vehicle gets stuck in a severe traffic jam.

Simulation Results

Figure 8.13 shows the number of active vehicles over time for 20 000 vehicles starting over the course of 20 h. In this scenario $T_{\text{upd}} = 5$ min, $T_{\text{hor}} = 0$ min, and no disturbances are added. Since assignments have maximum duration $t^D - t^S$ of 4.5 h, the number of active vehicles converges after 4.5 h to approximately 3800 with some small fluctuations around that value. After 20 h, no more vehicles start and the number of vehicles declines to zero. The total simulated distance driven is 4 695 658 km, which is ca. 13% of the average daily distance traveled by road freight vehicles on inner German trips of more than 150 km distance [3]. Thus the simulated amount of coordinated heavy-duty vehicle traffic is in the right order of magnitude for the road network considered. The average distance of a route in the simulation is 234 km. We can see that the ratio of platooning vehicles to active vehicles is small just after time 0 h and large just after time 20 h. This can be explained by the structure of adapted platoon plans, which typically have a segment at the beginning and the end of the trip where vehicles drive alone. Therefore it takes some time before the first platoons form. The first vehicles that

Table 8.1: Aggregated simulation results from different scenarios

Scenario	T_{upd} [min]	T_{hor} [min]	N_{sim}	ΔF_{sim} [%]	P_{plf} [%]	P_{pla} [%]	P_{del} [%]
Spontaneous	5	0	20 000	1.41	11.82	22.11	0.00
Regular 1	5	0	20 000	5.09	46.87	65.23	0.00
Regular 2	30	0	20 000	4.76	43.93	61.39	0.00
Regular 3	60	0	20 000	4.35	40.25	56.66	0.00
Regular 4	300	0	20 000	1.73	15.99	23.68	0.00
Regular 5	5	10	20 000	5.21	47.88	66.61	0.00
Regular 6	5	30	20 000	5.26	48.27	67.51	0.00
Regular 7	5	60	20 000	5.27	48.20	67.52	0.00
Regular 8	60	60	20 000	5.25	48.12	67.04	0.00
Regular 9	300	300	20 000	5.09	46.36	64.71	0.00
Batch	-	-	20 000	5.05	45.85	63.88	0.00
Deadline	5	0	20 000	7.01	41.83	55.03	0.00
Simple F.M.	5	0	20 000	4.52	47.62	64.00	0.00
Deadl. Simple F.M.	5	0	20 000	5.84	49.20	64.21	0.00
Dropout 20%	5	0	20 000	4.83	44.43	62.67	0.00
Dropout 50% 1	5	0	20 000	4.40	40.40	58.19	0.00
Dropout 50% 2	60	60	20 000	4.44	40.76	59.09	0.00
Dropout 50% 3	300	300	20 000	3.95	36.23	53.45	0.00
Max. Speed 20%	5	0	20 000	4.96	45.21	63.39	5.72
Max. Speed 50% 1	5	0	20 000	4.74	42.70	60.61	14.41
Max. Sp. 50% Plan.	5	0	20 000	4.82	43.62	63.05	0.00
Max Speed 50% 2	60	60	20 000	4.77	42.75	61.62	14.49
Max Speed 50% 3	300	300	20 000	4.59	40.87	59.20	13.65
Spontaneous 1000	5	0	1 000	0.07	0.60	1.20	0.00
Regular 1000	5	0	1 000	1.03	10.49	19.93	0.00
Spontaneous 5000	5	0	5 000	0.40	3.35	6.59	0.00
Regular 5000	5	0	5 000	3.10	29.82	48.46	0.00
Spontaneous 10 000	5	0	10 000	0.79	6.61	12.80	0.00
Regular 10 000	5	0	10 000	4.02	37.79	57.11	0.00
Integer Program	5	0	20 000	5.18	48.19	65.92	0.00
One-Pass	5	0	20 000	4.39	41.14	60.05	0.00

finish after 20 h are close to the end of their trip and many do not platoon at this point in time. Between 4.5 h–20 h, the number of platooning vehicles and the number platoon followers is almost constant and follows the trend in the number of active vehicles.

Table 8.1 shows the results from the different simulations scenarios. The number of simulated assignments is denoted N_{sim} . Four statistics are computed for each scenario in the time window 4.5 h–20 h: the reduction of fuel consumption compared the each vehicle driving the simulated distance alone at a speed of 80 km/h, denoted ΔF_{sim} ; the percentage of distance driven as platoon follower P_{plf} ; the percentage of distance driven in a platoon P_{pla} ; and the percentage of assignments that get

delayed P_{del} . For all groups of scenarios with the same number of assignments, the start times and routes are the same. In the following, we discuss the scenarios in detail.

- **Spontaneous** This serves as a baseline scenario approximating fuel savings that can be obtained with local coordination, so-called spontaneous platooning. It is computed by rejecting all adapted plans that deviate from the default plan more than $0.5 \text{ km}/80 \text{ km/h} = 22.5 \text{ s}$, corresponding to an inter-vehicle distance of 500 meters at 80 km/h.
- **Regular 1–9, Batch** These scenarios differ in their values for the update interval T_{upd} and the preview horizon T_{hor} . In the “Batch” all plans were computed at once. We can see that more frequent plan updates and longer horizons lead to improved platooning benefits and higher platoon percentage. It is remarkable that “Regular 1” with no preview horizon outperforms the “Batch” scenario. The interpretation is that when plans are updated, a vehicle can be assigned to another platoon overcoming the limitation that an adapted plan can only adapt to one coordination leader. It also happens that a coordination leader and follower swap roles in an update so that both vehicles adapt.
- **Deadline, Simple F.M., Deadl. Simple F.M.** For “Simple F.M.”, the platoon planner uses a simple fuel model that does not take speed into account and assumes a fuel saving of 10% when platooning. It performs slightly worse than the regular fuel model highlighting the importance for accurate fuel consumption estimation. In the scenario “Deadline”, all deadlines are extended by half an hour. This gives much larger fuel savings than the “Regular 1” scenario but less distance platooned. In “Deadl. Simple F.M.”, the simple fuel model is used and more distance is traveled in platoons. This indicates that the platoon coordinator selects plans with a long merge phase at low speed and hence low fuel consumption in the “Deadline” scenario. Recall that default plans have a minimum speed of 80 km/h, while the speed during the merge phase is allowed to drop to 70 km/h. In “Deadl. Simple F.M.” the distance traveled as platoon follower is slightly higher, while the distance traveled in a platoon is slightly less than in “Regular 1”, which means that the increased flexibility is used to create larger platoons and save more fuel.
- **Dropout 20%, Dropout 50% 1–3** In “Dropout 20%”, 20% and in “Dropout 50% 1–3”, 50% of the vehicles end their trip at a random point on the route not known to the platoon coordinator, simulating unexpected events. Clearly, this negatively affects the fuel consumption reduction, but not dramatically, though. It also highlights the need for feedback. While “Regular 1” and “Regular 9” lead to approximately the same ΔF_{sim} , “Dropout 50% 1” performs better than “Dropout 50% 3” as the platoon coordinator can react to new situations.

- **Max. Speed 20%, Max. Speed 50% Plan., Max. Speed 50% 1–3** In these scenarios the maximum speed of 20% and 50% of the vehicles is reduced to a value between 80 km/h–90 km/h, that is sampled randomly for each vehicle. These disturbances lead to a mild reduction in fuel savings compared to the corresponding disturbance-free “Regular” scenarios. In the scenario “Max. Speed 50% Plan.” the true maximum speed of the vehicle is considered in the plan computation, which can partially compensate for the disturbance effects. A full compensation is not possible since the ability of vehicles to catch up to platoon partners is reduced. When the disturbance is not known to the platoon coordinator, some vehicles miss the deadline because a higher speed during the last part of the adapted plan is assumed than what is actually possible.
- **Regular {1000, 5000, 10 000}, Spontaneous {1000, 5000, 10 000}** These are scenarios with different numbers of assignments. Not surprisingly, the fuel savings from platooning increase with more assignments as there are more platoon opportunities. However, this value saturates as it is theoretically limited by the fuel consumption reduction of trailing vehicle in the platoon. We can also notice that the coordination significantly improves the fuel savings from platooning even for a small number of assignments.
- **Integer Program, One-Pass** These two scenarios are identical to the “Regular 1” scenario. The difference is in the algorithm used to solve the coordination leader selection (Problem 6.1) is different from Algorithm 1. For “Integer Program”, the exact solution is computed using a general purpose mixed integer programming solver. We see that the exact solution outperforms the greedy heuristic with a small margin in terms of fuel consumption reduction. In the scenario “One-Pass”, Algorithm 2 is used, which performs worse than the greedy heuristic, which is consistent with the results in [113].

The simulations indicate that the proposed coordinated platooning system has the potential to make efficient use of platooning. It improves significantly compared to the case where platooning technology is used only with local coordination. By updating plans frequently, the system can efficiently coordinate platooning with a short preview horizon and under the presence of disturbances.

Remark 33. While the simulations do not focus on assessing the computational requirements of a platoon coordinator, they indicate that such a system is feasible to implement. Most time is spent in computing candidate adapted platoon plans, an operation, which can be highly parallelized and for which heuristics can be developed, especially, in a dynamic setting. Even the exact solution of Problem 1 with the general purpose mixed integer programming solver Gurobi 7.0 has without further tuning a maximum computation time of 12.88 s and a mean computation time of 6.97 s. These were evaluated in the simulation time interval 4.5 h–20 h and running the optimization on an Intel Pentium Core i3 processor with 8 GB of RAM.



Figure 8.14: Overview of the routes driven during the experiment. Stars mark the starting point of vehicles. The destination is the same for all three vehicles and located on the lower left of the map.

8.4 Experimental Evaluation

This section describes an experimental evaluation of the coordinated platooning system. In the scope of the COMPANION European research project [78], a demonstrator was developed, which can coordinate test vehicles on public roads. It implements the strategic, tactical, and operational layer as introduced in Section 3. We present results from an experiment with that demonstrator where three vehicles starting from different position are coordinated to form a platoon en route without stopping. Plan are computed on-the-fly by an off-board system and executed by the vehicle's on-board system. To our knowledge, it is one of the first demonstrations of en route heavy-duty vehicle platoon formation. The experiment shows the feasibility of the proposed coordination system for en route platoon formation under realistic conditions. It also highlights the need for sophisticated planning using high quality data sources and the importance of such a system to dynamically react to real-time information from vehicles and other data sources.

Experiment Scenario

The experiment was conducted on public motorways in Spain west of Barcelona in September 2016. Figure 8.14 shows a map with routes and the start positions of the assignments. It involved three tractor trailer combinations equipped with a custom on-board system, on-board human-machine interface [189], and vehicle-to-vehicle



Figure 8.15: View from the cockpit of the green vehicle while trailing the blue vehicle in the platoon. The on-board HMI can be seen in the dashboard. The computer on right side of the image is used to monitor and collect data from the on-board system.

and vehicle-to-infrastructure (mobile broadband) communication capabilities. Figure 8.15 shows a picture taken from the cabin of the vehicle whose route is plotted in green in Figure 8.14 during the platooning phase. All three vehicles started at different positions close to motorways and were given the same destination and the same arrival deadline. The initial start times were determined by the coordinator as part of the initial plan computation. The flexibility of freely adjusting the start time was necessary in order to create a situation where platooning would occur with only three vehicles. Since platooning capabilities are currently only present on a few test vehicles, experiments with larger numbers of vehicles are not feasible. The experiment took place in regular traffic around 10 a.m. in the morning on a weekday, i.e., under realistic conditions. The off-board system was running on multiple cloud computing instances.

Experimental Results

The initially computed plans entailed that first two vehicles, whose data are shown in blue and green in the figures, would merge at the motorway intersection depicted in Figure 8.16a and platoon until their destinations. On the way, the first two vehicles would pass the start location of the third vehicle, whose data are shown in purple in the figures. In the following, we refer to the vehicle as the blue, green, and purple vehicle according to the color used in the plots. The purple vehicle



(a) At the first merge point, both vehicles first have to pass a road toll booth before merging onto the same road in the lower left of the image.



(b) At the second merge point, the starting position of the red vehicle is located on the upper right of the image on a parking lot next to the highway. The lack of an acceleration lane made it necessary to start before the scheduled departure time.

Figure 8.16: Aerial photography of the two merge points with the GPS traces of the first two vehicles.




Figure 8.17: The distance driven along the vehicles' routes. The starting distance was adjusted in a way that two vehicles at the same location on the common section of the route have the same driven distance.

would join the platoon at that point shown in Figure 8.16b. Figure 8.17 shows the distance driven over time with a common reference on the overlapping part of the route. We can see that due to the alignment of the start times, no planned adaptation prior to the merge points had to be done. We also see that the relative distances of the vehicles in a platoon are very small compared to the total distance driven, which supports the assumption made in the planning that vehicles are at the same position when they platoon.

Figure 8.18 shows speed measured by the GPS receivers, the merge times according to the on-board system and the deviation from the latest computed plan by the off-board system. New plans were computed only when a vehicle exceeded a deviation of 30 s. We can see that the blue vehicle starts ca. 15 s too late and the second vehicle ca 17 s too early. This is partly because the starting point was on the motorway due to restrictions in the routing module and the off-board plan neglects the initial acceleration phase of the vehicle. However, also in a setting where the vehicle is operated by a human driver, such small deviations in the planned starting time are likely to occur. Both vehicles' on-board controllers manage to reduce the deviation considerably during their journey. The large spike in the de-




Figure 8.18: The measured speed and measured deviation from the plan of all three vehicles. The actual merge times according to the on-board system are indicated as vertical black lines.

viation of the blue vehicle at 1125s is a glitch in the deviation computation of the on-board system and triggered a re-computation of the plans. The purple vehicle starts 36s too early. This was done intentionally since the second merge point shown in Figure 8.16 has no acceleration lane and the back-end plan did not consider the initial acceleration phase needed. Furthermore, the driver has to wait for a gap in the stream of traffic to be able to enter the highway without acceleration lane. Normally, this would be resolved by the platoon coordinator updating the plans, however, due to the small number of test vehicles in the demonstration and the short distance driven after that point, this was handled by starting ahead of time and then having the on-board controller lower the vehicle speed compared to the plan. Finally, one can notice that strong occasional variations in speed in the second half of the experiment when all three vehicles are platooning. This is due to a minor time synchronization problem that sometimes triggered the platoon controller to increase the headway as a safety precaution. Since the purpose of the demonstration was the coordinated formation of platoons, this was not considered to be a problem for the experiment.

Figure 8.19 shows a detail of Figure 8.17. Additionally, the planned and the

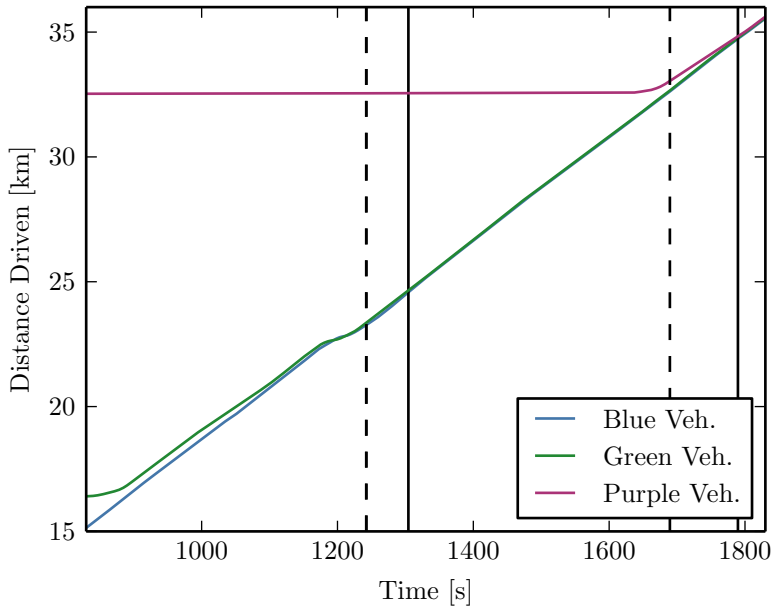


Figure 8.19: Zoom into Figure 8.17. Additionally the planned merge times are indicated with dashed black vertical lines and the actual merge locations according to the on-board system are indicated as vertical black lines.

actual merge time are indicated. We can see that in both cases the vehicles merge a little later than planned. The second merge happens later than expected since the vehicle has to start ahead of time, as previously discussed, and then has to wait for the other two vehicles by driving slower than what was planned. For the first merge, the blue vehicle gets delayed due to incorrect information on the speed restriction after the toll booth. It seems that this was caused by a wrong association of the current segment. This error would also explain the incorrect computation of the deviation. Figures 8.20 and 8.21 show more detail at this point. Already when arriving at the toll booth, it is slightly delayed. After passing the toll booth where the speed dips below 20 km/h, the back-end plan used the average speed driven by probe vehicles at this point while the on-board system recommended the maximum speed according to the incorrect legal speed limit of 30 km/h indicated in the map. Since the speed was controlled manually at this point passing the toll booth, the driver kept a low speed of 65 km/h. This speed was higher than the recommended 30 km/h in order to be safe in traffic but it was lower than the one in the platoon plan. Therefore, the vehicle got delayed and had to catch up to the blue vehicle. Furthermore, it is clearly challenging to accurately predict the speed profile at a

toll booth due to merging zones and possible waiting times highlighting the need for real-time feedback.

Another aspect that can be seen in Figures 8.20 and 8.21 is the update of plans in case of deviations larger than 30 s. The real deviation at this point would actually not be large enough to trigger a deviation, but due to the error in the computation of the deviation it exceeds the threshold for a short time. The time at which the deviation message and the vehicle position is sent is indicated in the plot by a vertical dashed line. The latest position update used in the new plan of the green vehicle is indicated with a green dashed vertical line. We can see that the planner assumes that the reported speed is kept until the beginning of the next segment, which becomes the first segment of the vehicle's route in the new plan. The computed plan assigns the green vehicle to be coordination leader and to follow a default speed profile to arrive at the destination in time. The blue vehicle gets the role of coordination follower and is supposed to select a higher speed to catch up with the green one. However, as indicated by the solid vertical lines, the blue vehicle receives the plan after this catch up phase due to delays in the computation and communication. Therefore, the updated plan does not have any effect and the merge is facilitated by the on-board controller.

Figures 8.20 and 8.21 show also how the on-board controller compensates for small deviations from the reference plan. The reference speed can be adjusted with up to ± 9 km/h, which proves sufficient to track the computed plans in the experiment.

Conclusions from the Experiment

We can conclude from this experiment that coordinated en route formation of vehicle platoons is feasible under realistic conditions. Additionally to the experiment discussed in this section, several test runs in Sweden and in Spain have been made, which suggest that the successful formation of platoons was not just a lucky coincidence. This complements the simulation results from Chapter 8, which features large number of vehicles at the expense of modeling less detail on the lower control layers, which is fully present in the experimental results.

The results also show that further research is needed on how much and what kind of data should be considered in the planning process. Doing a more detailed planning can be quite challenging, in particular, when it comes to considering the vehicle dynamics that depend, among others, on the vehicle's trailer, load, and engine characteristics and when influence of traffic should be taken into account. During a test run, a plan contained a segment with very low speed that was due to a traffic jam that was dissolved only a shortly before the vehicle passed that area without being detected in time by the traffic information system. The system recommended this low speed, which had to be overwritten by the driver to maintain traffic safety. Even static map data such as the legal speed limit can change with time and can thus be sometimes outdated. The alternative is to rely more on the

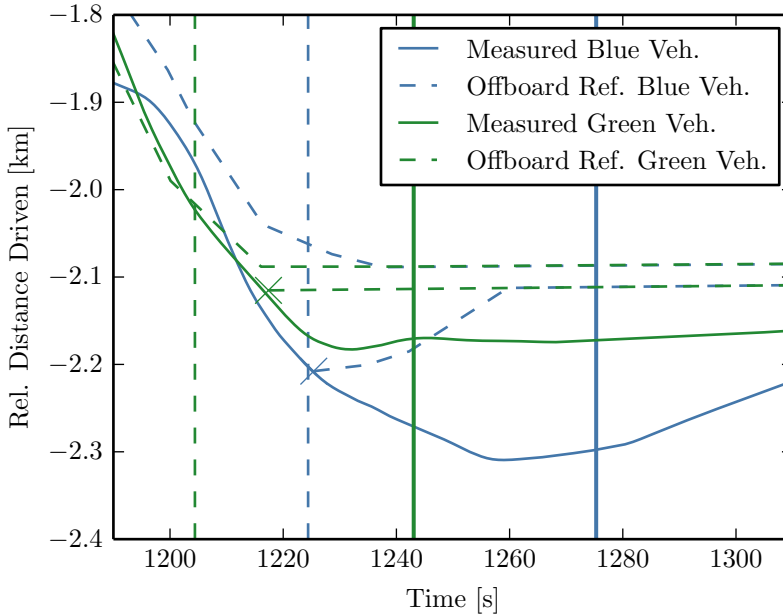


Figure 8.20: The first merge. This plot shows the distance driven relative to vehicle that would drive from time 0 with speed 74 km/h and is otherwise similar to Figure 8.17. The data are shown relative to a virtual vehicle for visual presentation reasons only. Additionally, the planned distance driven is shown with dashed lines for the original plan and the first update. The first data points of the update are marked with crosses. The vertical dashed lines indicate when the latest position was sent that was used in the replanning. The solid vertical lines indicate when the new plan was received on the respective vehicle. Note that the time scale of the left plot is different from the others.

ability to dynamically react to disturbances by means of real-time feedback from the vehicles.

It should also be mentioned that the nominal speed of the vehicles in the experiment was 16.7% lower than the maximum speed giving them the ability to both speed up and slow in down in order to adjust their timing. When a coordinated platooning system will be used commercially in the future, the default speed might be closer to the maximum which means that adjustments are mainly made by reducing the speed. This puts even more importance on the platooning systems ability to adapt and reconfigure plans across multiple vehicles since a vehicle that gets delayed can not simple compensate by increasing its speed.

At this point, I would like to thank the participants of the COMPANION




Figure 8.21: The first merge. The two upper plots show the measured speed, the on-board speed reference, and the speed reference according to the plan computed by the back end system. The plot on the bottom shows the measured deviation from the plan. In all three plots, the vertical dashed lines indicate when the latest position was sent that was used in the replanning. The solid vertical lines indicate when the new plan was received on the respective vehicle. Note that the time scale of the left plot is different from the others.

project [78] for their contributions in developing the demonstrator and performing the experiments. In particular, we like to mention Marcos Pillado and Stanislav Vovk for experiment planning and execution, Sergej Saibel, Thomas Friedrichs, Sönke Eilers, and David Spulak for implementation of the off-board system, Henrik Pettersson, Samuel Wickström and Shadan Sadeghian Borojeni for on-board system development, and Edzard Neumann and Thilo Schaper for work on route calculation and traffic estimation.

8.5 Summary

This section evaluates the platoon coordination methods developed in this thesis with three simulation studies and an experiment under realistic conditions. The first simulation study assesses the platoon plan computation algorithms on an artificial road network. It focuses on the influence of the number of coordinated vehicles and the flexibility to adapt the speed on the computed plans comparing two variants of the coordination leader selection algorithm. The second simulation study investigates the consequences of the computed plans on a real road network. Both studies show that the coordination algorithms are key to leveraging the fuel consumption benefits when there are only a limited number of platooning-enabled vehicles and that the algorithms can scale to traffic volumes of heavy-duty vehicles, as encountered in reality. The third simulation study evaluates the periodic plan updates the system performs and investigates the ability of the coordination system to handle disturbances. It shows that frequent replanning can compensate for the lack of preview information when vehicles start, improve the fuel consumption reduction, and it can compensate, to some extent, for unforeseen events such as a vehicle dropping out from a platoon.

The experimental evaluation shows data collected at a demonstration at which three heavy-duty vehicles were successfully coordinated to form a platoon en route on public highways under normal traffic conditions. The coordination was performed by a complete prototypical implementation of the coordination system introduced in this thesis and thus verifies the feasibility of the approach, but also highlights some aspects that need further development on the way to a commercial implementation of a coordinated platooning system.

Chapter 9

Conclusions and Future Work

THIS chapter concludes the thesis. Section 9.1 recapitulates and discusses the presented results. Section 9.2 provides some possible directions for future work on the topic of coordinated heavy-duty vehicle platooning.

9.1 Conclusions

This thesis develops and evaluates the system-wide coordination of heavy-duty vehicle platooning. Heavy-duty vehicle platooning is an emerging technology, which leverages recent developments in information and communication technology to address the major challenges freight transportation systems face. These challenges include the need to reduce emissions, congestion, and to improve safety. For heavy-duty vehicle platooning to have a significant impact, an effective and practical way to coordinate the formation of platoons is needed.

In order to integrate platooning in the complex road freight transport system, a hierarchical architecture is proposed. The platoon coordinator acts as a central system that receives assignments from fleet management systems and computes platoon plans maximizing the predicted fuel consumption reduction from platooning while ensuring timely arrival. These plans are forwarded to the vehicles' on-board systems where the platoon manager executes them. By updating the plans based on real-time data such as vehicle positions and platoon state, the system is resilient to disturbances and new assignments can be added on the fly. Thus, the proposed solution integrates well into the relatively decentralized and dynamic road freight transport system.

The synthesis of platoon plans is a complex problem. This problem is successfully approached by dividing the computation into tractable stages: route computation, candidate plan generation, fuel consumption estimation, and plan composition. This way, relevant shares of heavy-duty vehicle traffic encountered in reality can be handled by the coordination system. Furthermore, well developed methods for route computation in road networks can be used this way.

The candidate plan generation stage considers vehicle pairs, one vehicle adapting to the other in order to drive in a platoon for some distance. Platoon plans for vehicle pairs can be derived taking a great deal of detail into account, such as planning with fuel-optimal speed, location-dependent speed, and uncertain travel time. Methods for computing pairwise plans considering the most relevant factors are provided. Extending these methods to handle other constraints, factors, and degrees of freedom is possible along similar lines of reasoning. We discuss some relevant extensions in the next section.

In the plan composition stage, platoon plans for vehicle pairs are systematically combined based on the estimated fuel consumption for each individual pairwise plan. The task of composing the pairwise plans is formulated as a combinatorial optimization problem which is proven to be NP-hard and might be difficult to solve exactly. Therefore, efficient heuristic methods are proposed. Approaches for solving this combinatorial optimization problem are quite generic. Much detail is covered by the candidate plan generation stage and the fuel consumption estimation stage. Changing some of these details does not require any modifications to the plan composition stage. For instance, in this work, plans for vehicle pairs are combined in a way that minimizes the overall fuel consumption. Taking other measures such as driving time, emissions, and safety into account is straightforward as far as the plan composition stage is concerned. In the process of combining pairwise platoon plans, vehicles are clustered into coordination groups and the platoon plans for each cluster can be further improved, for instance, by jointly optimizing the speed profiles without changing where vehicles platoon.

One of the resource demanding steps in this approach is to identify all pairs of vehicles that can platoon based on their routes and relative timing. The computational effort of this step can be improved by extracting features from assignments and working with these low-dimensional features instead in order to identify candidate pairs. This approach is especially relevant when scaling platoon coordination over large areas where each vehicle can only potentially platoon with a small share of the vehicles in the network due to spatial separation.

The effectiveness of the method is demonstrated in three simulation studies and an experimental evaluation. Significant amounts of fuel can be saved by platooning. Coordinating platooning is shown to be crucial in fully exploiting the potential of heavy-duty vehicle platooning to reduce fuel consumption. The simulations demonstrate that the coordination approach can handle large amounts of vehicles as needed for platooning to have actual impact. Due to the dynamic and decentralized nature of road freight transport, long term planning can be difficult in practice. By updating plans frequently, such long term planning is not necessary with the proposed method. Repeated planning also helps to mitigate disturbances and to overcome some of the limitations imposed by combining pairwise plans.

The approach for platoon coordination was implemented in a demonstrator system in the scope of the COMPANION project. Three vehicles starting at different locations over 10 kilometers apart were successfully coordinated to form a three-vehicle platoon on the motorway. It was one of the first demonstrations of an

integrated platoon coordination system capable of forming platoons en route on public roads. The experimental results show that such a coordination system can be implemented based on readily available technology and infrastructure.

9.2 Directions for Future Work

There is a variety of ways to continue on the work presented in this thesis. One aspect that can be further investigated are the methods to compute pairwise plans. More advanced ways to compute the speed profile could be considered, for instance, taking vehicle dynamics and road topography explicitly into account as it is done already to improve fuel efficiency for individual vehicles. It is also possible to consider additional degrees of freedom for adaptation such as selecting the route from a small set of alternative routes and allowing the vehicle to stop along the way as part of the coordination. Also, the hard constraint on the arrival deadline could be replaced or complemented by a cost function penalizing deviations from the nominal arrival time. The presented approach to take travel time uncertainty into account for pairs of vehicles fits well into a reinforcement learning framework, where travel time distributions are learned over time along with the optimal speed controller.

On the level of coordination algorithms for many vehicles, further improvement is possible. Other heuristics can be applied to the coordination leader selection problem, for instance meta-heuristics such as simulated annealing, genetic algorithms, etc. Variations on the formulation are possible such as a limit on the number of coordination followers or allowing adapted plans where the coordination follower successively platoons with several coordination leaders. Also the interplay between the dynamic system evolution and the selection of coordination followers can be beneficial to investigate, for instance, ensuring that coordination followers do not prematurely switch coordination leader. There are opportunities to improve plans by local search, for example, by merging coordination followers already into sub-platoons before merging with the coordination leader.

One of the advantages of the approach taken in this thesis is that it is fairly generic. It relies only on a minimum of information about operational constraints of freight vehicles as part of the road freight transport system. However, integrating the fleet management systems and the platoon coordination system further might bring additional benefits at the cost of increased complexity. Integrating the planning of rest times would, for instance, allow vehicles in long-haulage transport to platoon over longer distances. Such an integration is even necessary when platooning is used as a way to extend the time a driver is allowed to operate the vehicle before taking a break. Testing different levels of integration will require to use realistic traffic patterns or historic fleet data in simulations.

An interesting question is also how platooning and platoon coordination should be embedded into the economic structure of road freight transportation. Which organizations should run platoon coordination systems and what happens if there are

several alternative coordination systems? What will be the business models behind platoon coordination? Due to the market structure of road freight transport, it is likely that vehicles in a platoon are not operated by the same company. Therefore, a scheme to balance the benefits from platooning might be needed. One approach is to introduce a market for platooning. How such a market should be designed is a challenging research question. In such a setting, there might be the need for an authentication scheme to ensure that the correct vehicles form a platoon and that platoon partners can be trusted to fulfill the technical requirements for platooning. Furthermore, there are privacy concerns as companies might be unwilling to share operational data. This problem might be possible to address to some extent on technical level.

As it is typically the case when implementing technical systems, a multitude of exceptional cases need to be considered and the presented platoon coordination system is no exception. In particular, ensuring correct integration of the system with different fleet management systems might be challenging. Also heavy vehicles overtaking each other is not permitted everywhere, which can be important to keep a certain reference speed. In a system like this, there can be significant communication delays or even communication failures, which need to be dealt with. There are also open questions related to the practical implementation of platooning in traffic. A commonly voiced concern is, for instance, that platoons can hinder other vehicles from entering the motorway. Therefore, platooning might be limited in some areas, which needs to be taken into account by the coordination system. Some of the practical aspects will first become apparent once platooning systems are more extensively tested in practice. As the adoption of platooning technology including the formation of platoon service providers is rapidly progressing at the moment, more insights on the challenges associated with implementing platooning technology will be available in near future. Efforts to standardize platooning systems and interfaces are needed in this process to enable inter-interoperability of the components from different parties in a coordinated platooning system.

Ultimately, platooning is the forerunner of automated driving technology. Higher levels of automation have the potential to fundamentally change the way road freight transport is operated. It is not unlikely that platoon coordination systems will grow into key enablers for the autonomous operation of freight transport vehicles. Such advanced coordination systems bear the potential to facilitate a more economic, reliable, and environmentally friendly transport system.

Bibliography

- [1] Verbundprojekt KONVOI: Entwicklung und Untersuchung des Einsatzes von elektronisch gekoppelten Lkw-KONVOIs, Abschlussbericht. Technical report, 2009.
- [2] White Paper on transport—Roadmap to a Single European Transport Area—Towards a competitive and resource efficient transport system. Technical report, European Commission, 2011.
- [3] Verkehr deutscher Lastkraftfahrzeuge [German goods vehicle traffic]. Technical report, German Federal Motor Transport Authority, 2014.
- [4] EU transport in figures – Statistical pocketbook 2015. Technical report, European Commission, 2015.
- [5] Lastbilstrafik 2014 Swedish national and international road goods transport 2014. Technical report, Trafikanalys, 2015.
- [6] European Truck Platooning, 2016. URL <https://www.eustruckplatooning.com/Press/default.aspx>.
- [7] Greenhouse gas emissions by source sector, June 2017. URL <http://ec.europa.eu/eurostat/data/database>.
- [8] Infographic: EU Roadmap for Truck Platooning, May 2017. URL <http://www.acea.be/publications/article/infographic-eu-roadmap-for-truck-platooning>.
- [9] Managing the Transition to Driverless Road Freight Transport. Technical report, OECD/ITF, 2017.
- [10] The Scania Report 2016. Technical report, Scania AB, 2017.
- [11] Annual detailed enterprise statistics for services, Jan. 2018. URL <http://ec.europa.eu/eurostat/data/database>.
- [12] Annual road freight transport by distance class with breakdown by type of goods (1 000 t, Mio Tkm, Mio Veh-km, 1 000 BTO), from 2008 onwards, Jan. 2018. URL <http://ec.europa.eu/eurostat/data/database>.

- [13] I. Abraham, A. Fiat, A. V. Goldberg, and R. F. Werneck. Highway Dimension, Shortest Paths, and Provably Efficient Algorithms. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 782–793. Society for Industrial and Applied Mathematics, 2010.
- [14] P. Adarsh, A. G. Cherian, and D. R. K. Arun, K. S. Numerical Investigation of Drag on a Trailing Aerodynamic Sedan Vehicle. *International Journal of Mechanical and Production Engineering*, 2(4):81–86, Apr. 2014.
- [15] A. Adler, D. Miculescu, and S. Karaman. Optimal Policies for Platooning and Ride Sharing in Autonomy-enabled Transportation. In *Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2016.
- [16] N. Agatz, A. Erera, M. Savelsbergh, and X. Wang. Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2):295–303, Dec. 2012.
- [17] A. Alam, B. Besselink, V. Turri, J. Mårtensson, and K. H. Johansson. Heavy-duty vehicle platooning towards sustainable freight transportation: A cooperative method to enhance safety and efficiency. *IEEE Control Systems Magazine*, 35(6):34–56, Nov. 2015.
- [18] A. Alam, A. Gattami, and K. H. Johansson. An experimental study on the fuel reduction potential of heavy duty vehicle platooning. In *13th International IEEE Conference on Intelligent Transportation Systems*, pages 306–311, Madeira, Portugal, 2010.
- [19] C. J. Almqvist and K. Heinig. European Accident Research and Safety Report 2013. Technical report, Volvo Trucks, 2013.
- [20] L. Alvarez and R. Horowitz. Safe Platooning in Automated Highway Systems. Research report, University of California, Berkeley, CA, USA, 1997.
- [21] M. Amoozadeh, H. Deng, C.-N. Chuah, H. M. Zhang, and D. Ghosal. Platoon management with cooperative adaptive cruise control enabled by VANET. *Vehicular Communications*, 2(2):110–123, Apr. 2015.
- [22] M. S. Andersen, J. Dahl, and L. Vandenberghe. CVXOPT: A Python package for convex optimization, 2013. URL <http://cvxopt.org>.
- [23] S. Arroyo and A. L. Kornhauser. Modeling Travel Time Distributions on a Road Network. In *11th World Conference on Transport Research*, 2007.
- [24] J. Axelsson. Safety in Vehicle Platooning: A Systematic Literature Review. *IEEE Transactions on Intelligent Transportation Systems*, 18(5):1033–1045, May 2017.

- [25] D. Banister and D. Stead. Impact of information and communications technology on transport. *Transport Reviews*, 24(5):611–632, Feb. 2004.
- [26] L. D. Baskar, B. De Schutter, J. Hellendoorn, and Z. Papp. Traffic control and intelligent vehicle highway systems: a survey. *Intelligent Transport Systems, IET*, 5(1):38–52, Mar. 2011.
- [27] L. D. Baskar, B. D. Schutter, and H. Hellendoorn. Optimal routing for automated highway systems. *Transportation Research Part C: Emerging Technologies*, 30:1–22, May 2013.
- [28] M. Baumgartner, J. Léonardi, and O. Krusch. Improving computerized routing and scheduling and vehicle telematics: A qualitative survey. *Transportation Research Part D: Transport and Environment*, 13(6):377–382, Aug. 2008.
- [29] I. Bäumlér and H. Kotzab. Intelligent Transport Systems for Road Freight Transport—An Overview. In M. Freitag, H. Kotzab, and J. Pannek, editors, *5th International Conference of Dynamics in Logistics*, pages 279–290. Springer International Publishing, 2017.
- [30] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, Mar. 1999.
- [31] J. Bentley and T. Ottmann. Algorithms for Reporting and Counting Geometric Intersections. *IEEE Transactions on Computers*, C-28(9):643–647, Sept. 1979.
- [32] C. Bergenhem, S. Shladover, E. Coelingh, C. Englund, and S. Tsugawa. Overview of platooning systems. In *19th ITS World Congress*, Oct. 2012.
- [33] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control: The discrete time case*. Academic Press New York, 1978.
- [34] D. Bertsimas and S. S. Patterson. The air traffic flow management problem with enroute capacities. *Operations Research*, 46(3):406–422, June 1998.
- [35] B. Besselink, V. Turri, S. van de Hoef, K.-Y. Liang, A. Alam, J. Mårtensson, and K. H. Johansson. Cyber-physical Control of Road Freight Transport. *Proceedings of the IEEE*, 104(5):1128–1141, Mar. 2016.
- [36] D. Bevly, C. Murray, A. Lim, D. R. Turochy, D. R. Sesek, S. Smith, L. Humphreys, G. Apperson, J. Woodruff, S. Gao, M. Gordon, N. Smith, S. Praharaaj, J. Batterson, R. Bishop, D. Murray, A. Korn, J. Switkes, S. Boyd, and B. Kahn. Heavy Truck Cooperative Adaptive Cruise Control: Evaluation, Testing, and Stakeholder Engagement for Near Term Deployment: Phase Two Final Report. Technical report, Auburn University, 2017.

- [37] A. K. Bhoopalam, N. Agatz, and R. Zuidwijk. Planning of truck platoons: A literature review and directions for future research. *Transportation Research Part B: Methodological*, 107:212–228, Jan. 2017.
- [38] C. Bierwirth and F. Meisel. A survey of berth allocation and quay crane scheduling problems in container terminals. *European Journal of Operational Research*, 202(3):615–627, May 2010.
- [39] K. Bimbraw. Autonomous cars: Past, present and future a review of the developments in the last century, the present scenario and the expected future of autonomous vehicle technology. In *12th International Conference on Informatics in Control, Automation and Robotics*, volume 01, pages 191–198, July 2015.
- [40] R. Bishop, D. Bevly, L. Humphreys, S. Boyd, and D. Murray. Evaluation and Testing of Driver-Assistive Truck Platooning. *Transportation Research Record: Journal of the Transportation Research Board*, 2615:11–18, 2017.
- [41] J. Blanch, T. Walter, and P. Enge. Satellite Navigation for Aviation in 2025. *Proceedings of the IEEE*, 100(Special Centennial Issue):1821–1830, May 2012.
- [42] C. Bonnet and H. Fritz. Fuel consumption reduction in a platoon: Experimental results with two electronically coupled trucks at close spacing. *SAE Technical Paper*, (2000-01-3056), Aug. 2000.
- [43] Y. Bontekoning and H. Priemus. Breakthrough innovations in intermodal freight transport. *Transportation Planning and Technology*, 27(5):335–345, Oct. 2004.
- [44] K. Boriboonsomsin, M. J. Barth, W. Zhu, and A. Vu. Eco-Routing Navigation System Based on Multisource Historical and Real-Time Traffic Information. *IEEE Transactions on Intelligent Transportation Systems*, 13(4):1694–1704, Dec. 2012.
- [45] M. T. Borzacchiello, I. Casas, B. Ciuffo, and P. Nijkamp. *Geospatial Technology and the Role of Location in Science*, chapter Geo-ICT in Transportation Science, pages 267–285. Springer Netherlands, Dordrecht, 2009.
- [46] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [47] N. Boysen, D. Briskorn, and S. Schwerdfeger. The identical-path truck platooning problem. *Transportation Research Part B: Methodological*, 109:26–39, Mar. 2018.
- [48] F. Browand and M. Hammache. The Limits of Drag Behavior for Two Bluff Bodies in Tandem. In *SAE 2004 World Congress & Exhibition*. SAE International, Mar. 2004.

- [49] F. Browand, J. McArthur, and C. Radovich. Fuel Saving Achieved in the Field Test of Two Tandem Trucks. Technical report, University of Southern California, 2004.
- [50] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. A Tight Linear Time (1/2)-Approximation for Unconstrained Submodular Maximization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, FOCS '12*, pages 649–658, Washington, DC, USA, Oct. 2012. IEEE Computer Society.
- [51] H.-C. Burmeister, W. Bruhn, Ø. J. Rødseth, and T. Porathe. Autonomous Unmanned Merchant Vessel and its Contribution towards the e-Navigation Implementation: The MUNIN Perspective. *International Journal of e-Navigation and Maritime Economy*, 1:1–13, Dec. 2014.
- [52] C. Caballini, S. Sacone, and M. Saeednia. Planning truck carriers operations in a cooperative environment. In *19th IFAC World Congress*, volume 47, pages 5121–5126, 2014.
- [53] L. Caltagirone, S. Torabi, and M. Wahde. Truck Platooning Based on Lead Vehicle Speed Profile Optimization and Artificial Physics. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pages 394–399, Sept. 2015.
- [54] P. Capros, A. D. Vita, N. Tasios, D. Papadopoulos, P. Siskos, E. Apostolaki, M. Zampara, L. Paroussos, K. Fragiadakis, N. Kouvaritakis, L. Höglund-Isaksson, W. Winiwarter, P. Purohit, H. Böttcher, S. Frank, P. Havlík, M. Gusti, and H. P. Witzke. EU Energy, Transport and GHG Emissions Trends to 2050. Technical report, European Commission, 2013.
- [55] P. Caramia, G. Lauro, M. Pagano, and P. Natale. Automatic train operation systems: A survey on algorithm and performance index. In *2017 AEIT International Annual Conference*, pages 1–6, Sept. 2017.
- [56] H. B. Celikoglu. Dynamic Classification of Traffic Flow Patterns Simulated by a Switching Multimode Discrete Cell Transmission Model. *IEEE Transactions on Intelligent Transportation Systems*, 15(6):2539–2550, Dec. 2014.
- [57] S. Cepolina and H. Ghiara. New trends in port strategies. Emerging role for ICT infrastructures. *Research in Transportation Business & Management*, 8:195–205, Oct. 2013.
- [58] R. Chandler, R. Herman, and E. Montroll. Traffic dynamics: Studies in car following. *Operations Research*, 6(2):165–184, Mar. 1958.
- [59] C. Chen, A. Skabardonis, and P. Varaiya. Travel-Time Reliability as a Measure of Service. *Transportation Research Record: Journal of the Transportation Research Board*, 1855:74–79, 2003.

- [60] M. Chen and S. Chien. Dynamic Freeway Travel-Time Prediction with Probe Vehicle Data: Link Based Versus Path Based. *Journal of the Transportation Research Board*, 1768:157–161, 2001.
- [61] S. I.-J. Chien and C. M. Kuchipudi. Dynamic Travel Time Prediction with Real-Time and Historic Data. *Journal of Transportation Engineering*, 129(6):608–616, Nov. 2003.
- [62] S. Clark and D. Watling. Modelling network travel time reliability under stochastic demand. *Transportation Research Part B: Methodological*, 39(2):119–140, Feb. 2005.
- [63] G. Clarke. Task A Collection and Analysis of Data on the Structure of the Road Haulage Sector in the European Union. Technical report, European Commission, 2014.
- [64] J.-P. B. Clarke, N. T. Ho, L. Ren, J. A. Brown, K. R. Elmer, K. Zou, C. Hunting, D. L. McGregor, B. N. Shivashankara, K.-O. Tong, A. W. Warren, and J. K. Wat. Continuous Descent Approach: Design and Flight Test for Louisville International Airport. *Journal of Aircraft*, 41(5):1054–1066, Sept. 2004.
- [65] J. D. Cohen, M. C. Lin, D. Manocha, and M. Ponamgi. I-COLLIDE: An interactive and exact collision detection system for large-scale environments. In *ACM Interactive 3D Graphics Conference*, pages 189–196. ACM, 1995.
- [66] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson. *Introduction to Algorithms*. MIT Press, 2nd edition, 2009.
- [67] T. G. Crainic, M. Gendreau, and J.-Y. Potvin. Intelligent freight-transportation systems: Assessment and the contribution of operations research. *Transportation Research Part C: Emerging Technologies*, 17(6):541–557, Dec. 2009.
- [68] T. G. Crainic and G. Laporte. Planning models for freight transportation. *European Journal of Operational Research*, 97(3):409–438, Mar. 1997.
- [69] A. Davila, E. del Pozo, E. Aramburu, and A. Freixas. Environmental Benefits of Vehicle Platooning. *SAE Technical Paper*, 2013-26-0142, Jan. 2013.
- [70] G. R. de Campos, P. Falcone, and J. Sjöberg. Autonomous cooperative driving: A velocity-based negotiation approach for intersection crossing. In *16th International IEEE Conference on Intelligent Transportation Systems*, pages 1456–1461, Oct. 2013.
- [71] E. Demir, T. Bektaş, and G. Laporte. A review of recent research on green road freight transportation. *European Journal of Operational Research*, 237(3):775–793, Sept. 2014.

- [72] Q. Deng. A General Simulation Framework for Modeling and Analysis of Heavy-Duty Vehicle Platooning. *IEEE Transactions on Intelligent Transportation Systems*, 17(11):1–11, May 2016.
- [73] B. Ding and A. C. König. Fast Set Intersection in Memory. In *International Conference on Very Large Data Bases*, volume 4, pages 255–266. VLDB Endowment, Jan. 2011.
- [74] A. Driving and P. T. Force. White Paper: Automated Driving and Platooning Issues and Opportunities. Technical report, ATA Technology and Maintenance Council, 2015.
- [75] J. Eckhardt. European Truck Platooning Challenge 2016. Technical report, Challenge network, Apr. 2016.
- [76] R. Eddington. The Eddington Transport Study, Main report: Transport’s role in sustaining the UK’s productivity and competitiveness. Technical report, Great Britain Department for Transport, 2006.
- [77] H. Edelsbrunner and H. A. Maurer. On the Intersection of Orthogonal Objects. *Information Processing Letters*, 13(4/5):177–181, 1981.
- [78] S. Eilers, J. Mårtensson, H. Pettersson, M. Pillado, D. Gallegos, M. Tobar, K. H. Johansson, X. Ma, T. Friedrichs, S. S. Borojeni, and M. Adolfsen. COMPANION – Towards Co-operative Platoon Management of Heavy-Duty Vehicles. In *18th IEEE International Conference on Intelligent Transportation Systems*, pages 1267–1273, 2015.
- [79] R. Elbert and J.-K. Knigge. Analysis of Decentral Platoon Planning Possibilities in Road Freight Transportation Using an Agent-based Simulation Model. In *Simulation in Produktion und Logistik 2017*. kassel university press, 2017.
- [80] M. Ellis, J. I. Gargoloff, and R. Sengupta. Aerodynamic Drag and Engine Cooling Effects on Class 8 Trucks in Platooning Configurations. *SAE International Journal of Commercial Vehicles*, 8(2):732–739, Sept. 2015.
- [81] C. Englund, L. Chen, J. Ploeg, E. Semsar-Kazerooni, A. Voronov, H. H. Bengtsson, and J. Didoff. The Grand Cooperative Driving Challenge 2016: boosting the introduction of cooperative automated vehicles. *IEEE Wireless Communications*, 23(4):146–152, Aug. 2016.
- [82] European Commission. Transport. In *The European Union Explained*. Publications Office of the European Union, 2014.
- [83] European Parliament and Council. Regulation (EC) No 561/2006, 2006.
- [84] European Parliament and Council. Regulation (EC) No 661/2009, 2009.

- [85] Y. Fan and Y. Nie. Optimal Routing for Maximizing the Travel Time Reliability. *Networks and Spatial Economics*, 6(3):333–344, Sept. 2006.
- [86] F. Farokhi and K. H. Johansson. A game-theoretic framework for studying truck platooning incentives. In *16th International IEEE Conference on Intelligent Transportation Systems*, pages 1253–1260, Oct. 2013.
- [87] F. Farokhi, I. Shames, and K. H. Johansson. Private and Secure Coordination of Match-Making for Heavy-Duty Vehicle Platooning. In *20th IFAC World Congress*, volume 50, pages 7345–7350, 2017.
- [88] P. Fernandes and U. Nunes. Platooning With IVC-Enabled Autonomous Vehicles: Strategies to Mitigate Communication Delays, Improve Safety and Traffic Flow. *IEEE Transactions on Intelligent Transportation Systems*, 13(1):91–106, Mar. 2012.
- [89] M. Figliozzi, N. Wheeler, E. Albright, L. Walker, S. Sarkar, and D. Rice. Algorithms for Studying the Impact of Travel Time Reliability Along Multisegment Trucking Freight Corridors. *Transportation Research Record: Journal of the Transportation Research Board*, 2224:26–34, 2011.
- [90] H. Flämig. *Autonomous Driving*, chapter Autonomous Vehicles and Autonomous Driving in Freight Transport, pages 365–385. Springer, Berlin, Heidelberg, 2016.
- [91] T. Friedrichs, S. S. Borojeni, W. Heuten, A. Lüdtkke, and S. Boll. PlatoonPal: User-Centered Development and Evaluation of an Assistance System for Heavy-Duty Truck Platooning. In *8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, Automotive’UI 16, pages 269–276, New York, NY, USA, 2016. ACM.
- [92] T. Friedrichs and A. Lüdtkke. Modeling Situation Awareness: The Impact of Ecological Interface Design on Driver’s Response Times. In *Seventh International Conference on Advanced Cognitive Technologies and Applications*, pages 47–51, Mar. 2015.
- [93] T. Friedrichs, M.-C. Ostendorp, and A. Lüdtkke. Supporting Truck Platooning: Development and Evaluation of Two Novel Human-Machine Interface. In *8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, Oct. 2016.
- [94] S. Gao and I. Chabini. Optimal routing policy problems in stochastic time-dependent networks. *Transportation Research Part B: Methodological*, 40(2):93–122, Feb. 2006.
- [95] M. Gendreau, G. Laporte, and R. Séguin. Stochastic vehicle routing. *European Journal of Operational Research*, 88(1):3–12, Jan. 1996.

- [96] H. Gharavi, K. V. Prasad, and P. A. Ioannou. Special issue on Advanced Automobile Technologies. *Proceedings of the IEEE*, 95(2):325–327, Feb. 2007.
- [97] T. Gheysens and G. Van Raemdonck. Effect of the Frontal Edge Radius in a Platoon of Bluff Bodies. *SAE International Journal of Commercial Vehicles*, 9(2):371–380, Sept. 2016.
- [98] S. Grant-Muller and M. Usher. Intelligent Transport Systems: The propensity for environmental and economic benefits. *Technological Forecasting and Social Change*, 82(Supplement C):149–166, Feb. 2014.
- [99] H. Götz. Chapter 8 - Commercial vehicles. In W.-H. Hucho, editor, *Aerodynamics of Road Vehicles*, pages 295–354. Butterworth-Heinemann, 1987.
- [100] I. Haas and B. Friedrich. Developing a micro-simulation tool for autonomous connected vehicle platoons used in city logistics. In *20th EURO Working Group on Transportation Meeting*, volume 27, pages 1203–1210, Budapest, Hungary, Sept. 2017.
- [101] R. Hall and C. Chin. Vehicle sorting for platoon formation: Impacts on highway entry and throughput. *Transportation Research Part C: Emerging Technologies*, 13(5–6):405–420, Oct. 2005.
- [102] S. Hallé, B. Chaib-draa, and J. Laumonier. Car Platoons Simulated As A Multiagent System. In *4th Workshop on Agent-Based Simulation*, pages 57–63, Mar. 2003.
- [103] A. Hanelt, B. Hildebrandt, and D. Leonhardt. D8.4 Documentation on the socio-economic impacts the project results. Technical report, BEI St. Gallen, 2016.
- [104] K. Hansson. Data-Driven Analysis of the Fuel Saving Potential of Road Vehicle Platooning. Master’s thesis, KTH Royal Institute of Technology, 2013.
- [105] H. Hao and P. Barooah. Stability and robustness of large platoons of vehicles with double-integrator models and nearest neighbor interaction. *International Journal of Robust and Nonlinear Control*, 23(18):2097–2122, July 2013.
- [106] I. Harris, Y. Wang, and H. Wang. ICT in multimodal transport and technological trends: Unleashing potential for the future. *International Journal of Production Economics*, 159:88–103, Jan. 2015.
- [107] H. Hartenstein and K. P. Laberteaux. A tutorial survey on vehicular ad hoc networks. *IEEE Communications Magazine*, 46(6):164–171, June 2008.

- [108] J. Herrera, D. Work, R. Herring, X. Ban, Q. Jacobson, and A. Bayen. Evaluation of traffic data obtained via GPS-enabled mobile phones: The Mobile Century field experiment. *Transportation Research Part C: Emerging Technologies*, 18(4):568–583, 2010.
- [109] A. Hofleitner, R. Herring, and A. Bayen. Probability distributions of travel times on arterial networks: A traffic flow and horizontal queuing theory approach. In *91st Transportation Research Board Annual Meeting*, number 12–0798, Washington, DC, Jan. 2012.
- [110] A. Hooper and D. Murrey. An Analysis of the Operational Costs of Trucking: 2017 Update. Technical report, American Transportation Research Institute, Oct. 2017.
- [111] R. Horowitz and P. Varaiya. Control design of an automated highway system. *Proceedings of the IEEE*, 88(7):913–925, July 2000.
- [112] H. L. Humphreys, J. Batterson, D. Bevely, and R. Schubert. An Evaluation of the Fuel Economy Benefits of a Driver Assistive Truck Platooning Prototype Using Simulation. In *SAE 2016 World Congress and Exhibition*. SAE International, Apr. 2016.
- [113] E. Ihrén. Evaluating the Handling of New Assignments in a Global Truck Platooning System Using Large-Scale Simulations. Master’s thesis, KTH Royal Institute of Technology, 2017.
- [114] P. Ioannou and C. Chien. Autonomous intelligent cruise control. *IEEE Transactions on Vehicular Technology*, 42(4):657–672, Nov. 1993.
- [115] S. Iwan, K. Małeckı, and J. Korczak. Impact of Telematics on Efficiency of Urban Freight Transport. In J. Mikulski, editor, *Activities of Transport Telematics*, pages 50–57, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [116] A. K. Jain. Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, 31(8):651–666, June 2010.
- [117] A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.
- [118] R. Janssen, H. Zwijnenberg, I. Blankers, and J. de Kruijff. Truck Platooning Driving the Future of Transportation. Technical report, TNO Mobility and Logistics, 2015.
- [119] E. Jenelius and H. N. Koutsopoulos. Travel time estimation for urban road networks using low frequency probe vehicle data. *Transportation Research Part B: Methodological*, 53:64–81, July 2013.

- [120] J. Jespersen-Groth, D. Potthoff, J. Clausen, D. Huisman, L. Kroon, G. Maróti, and M. N. Nielsen. *Robust and Online Large-Scale Optimization: Models and Techniques for Transportation Systems*, chapter Disruption Management in Passenger Railway Transportation, pages 399–421. Lecture Notes in Computer Science. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [121] P. Jiménez, F. Thomas, and C. Torras. 3D Collision Detection: A Survey. *Computers and Graphics*, 25(2):269–285, Apr. 2000.
- [122] L. Kaufman and P. J. Rousseeuw. *Finding Groups in Data: An introduction to Cluster Analysis*. John Wiley & Sons, Inc., 2008.
- [123] P. Kavathekar and Y. Chen. Vehicle Platooning: A Brief Survey and Categorization. In *ASME Design Engineering Technical Conference*, volume 3, Jan. 2011.
- [124] S. Keese, W. Bernhart, N. Dressler, M. Baum, and W. Rentzsch. Automated Truck The next big disruptor in the automotive industry? Technical report, Roland Berger, 2016.
- [125] C. Kelleher. Report on the state of the EU Road Haulage Market. Technical report, AECOM Limited, 2014.
- [126] B. S. Kerner. *The Physics of Traffic*. Springer-Verlag Berlin Heidelberg, 2004.
- [127] U. Kiencke, L. Nielsen, R. Sutton, K. Schilling, M. Papageorgiou, and H. Asama. The impact of automatic control on recent developments in transportation and vehicle systems. *Annual Reviews in Control*, 30(1):81–89, 2006.
- [128] J. Kim and H. S. Mahmassani. A finite mixture model of vehicle-to-vehicle and day-to-day variability of traffic network travel times. *Transportation Research Part C: Emerging Technologies*, 46:83–97, 2014.
- [129] J. P. J. Koller, A. G. Colín, B. Besselink, and K. H. Johansson. Fuel-Efficient Control of Merging Maneuvers for Heavy-Duty Vehicle Platooning. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pages 1702–1707, Sept. 2015.
- [130] R. Kunze, R. Ramakers, K. Henning, and S. Jeschke. *Intelligent Robotics and Applications*, volume 5928 of *Lecture Notes in Computer Science*, chapter Organization and Operation of Electronically Coupled Truck Platoons on German Motorways, pages 135–146. Springer Berlin Heidelberg, 2009.
- [131] R. Kunze, R. Ramakers, K. Henning, and S. Jeschke. Efficient Organization of Truck Platoons by Means of Data Mining – Application of the Data Mining Technique for the Planning and Organization of Electronically Coupled Trucks. In *International Conference on Informatics in Control, Automation and Robotics*, pages 104–113, 2010.

- [132] S. Kwoczek, S. D. Martino, and W. Nejd. Stuck Around the Stadium? An Approach to Identify Road Segments Affected by Planned Special Events. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pages 1255–1260, Sept. 2015.
- [133] A. Ladino, A. Kibangou, H. Fourati, and C. C. de Wit. Travel time forecasting from clustered time series via optimal fusion strategy. In *2016 European Control Conference (ECC)*, pages 2234–2239, June 2016.
- [134] M. Lammert, A. Duran, J. Diez, K. Burton, and A. Nicholson. Effect of platooning on fuel consumption of class 8 vehicles over a range of speeds, following distances, and mass. *SAE International Journal of Commercial Vehicles*, 7(2):626–639, Sept. 2014.
- [135] D. Lang, T. Stanger, and L. del Re. Opportunities on Fuel Economy Utilizing V2V Based Drive Systems. In *SAE 2013 World Congress & Exhibition*. SAE International, Apr. 2013.
- [136] J. Larson, K.-Y. Liang, and K. H. Johansson. A Distributed Framework for Coordinated Heavy-Duty Vehicle Platooning. *IEEE Transactions on Intelligent Transportation Systems*, 16(1):419–429, Feb. 2015.
- [137] J. Larson, T. Munson, and V. Sokolov. Coordinated Platoon Routing in a Metropolitan Network. In *Seventh SIAM Workshop on Combinatorial Scientific Computing*, pages 73–82, 2016.
- [138] E. Larsson, G. Sennton, and J. Larson. The vehicle platooning problem: Computational complexity and heuristics. *Transportation Research Part C: Emerging Technologies*, 60:258–277, Nov. 2015.
- [139] R. Laxhamma and A. Gascón-Vallbona. D4.3 Vehicle models for fuel consumption. Technical report, Scania CV AB, 2015.
- [140] W. Levine and M. Athans. On the optimal error regulation of a string of moving vehicles. *IEEE Transactions on Automatic Control*, AC-11(3):355–361, July 1966.
- [141] L. Li, R. R. Negenborn, and B. D. Schutter. Intermodal freight transport planning – A receding horizon control approach. *Transportation Research Part C: Emerging Technologies*, 60:77–95, Nov. 2015.
- [142] S. E. Li, K. Li, and J. Wang. Economy-oriented vehicle adaptive cruise control with coordinating multiple objectives function. *Vehicle System Dynamics*, 51(1):1–17, Aug. 2013.
- [143] K. Y. Liang, Q. Deng, J. Mårtensson, X. Ma, and K. H. Johansson. The influence of traffic on heavy-duty vehicle platoon formation. In *2015 IEEE Intelligent Vehicles Symposium*, pages 150–155, June 2015.

- [144] K.-Y. Liang, J. Mårtensson, and K. H. Johansson. When is it Fuel Efficient for a Heavy Duty Vehicle to Catch Up With a Platoon? In *7th IFAC Symposium on Advances in Automotive Control*, volume 46, pages 738–743, Sept. 2013.
- [145] K.-Y. Liang, J. Mårtensson, and K. H. Johansson. Fuel-Saving Potentials of Platooning Evaluated through Sparse Heavy-Duty Vehicle Position Data. In *2014 IEEE Intelligent Vehicles Symposium*, pages 1061–1068, June 2014.
- [146] K.-Y. Liang, S. van de Hoef, H. Terelius, V. Turri, B. Besselink, J. Mårtensson, and K. H. Johansson. Networked Control Challenges in Collaborative Road Freight Transport. *European Journal of Control*, 30:2–14, July 2016.
- [147] J. Lioris, R. Pedarsani, F. Y. Tascikaraoglu, and P. Varaiya. Doubling throughput in urban roads by platooning. In *14th IFAC Symposium on Control in Transportation Systems*, volume 49, pages 49–54, 2016.
- [148] F. Liu, T. Harada, Y. Lee, and Y. J. Kim. Real-time Collision Culling of a Million Bodies on Graphics Processing Units. In *ACM SIGGRAPH Asia*, pages 154:1–154:8. ACM, 2010.
- [149] X.-Y. Lu and S. E. Shladover. Automated Truck Platoon Control. Technical report, University of California, Berkeley, 2011.
- [150] F. Luo, J. Larson, and T. Munson. Coordinated Platooning with Multiple Speeds. Technical report, Agronne National Laboratory, 2017.
- [151] D. Luxen and C. Vetter. Real-time routing with OpenStreetMap data. In *19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 513–516. ACM, 2011.
- [152] C. Macharis and Y. Bontekoning. Opportunities for OR in intermodal freight transport research: A review. *European Journal of Operational Research*, 153(2):400–416, Mar. 2004.
- [153] J. E. Manley. Unmanned surface vehicles, 15 years of development. In *OCEANS 2008*, pages 1–4, Sept. 2008.
- [154] G. Marchet, A. Perego, and S. Perotti. An exploratory study of ICT adoption in the Italian freight transportation industry. *International Journal of Physical Distribution & Logistics Management*, 39(9):785–812, 2009.
- [155] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, June 2000.

- [156] G. Mbydzenyuy. Arrival Times with Hours of Service Regulations for Truck Drivers-Tracks and Gaps from Current Research. In *18th IEEE International Conference on Intelligent Transportation Systems*, pages 2631–2636, Sept. 2015.
- [157] B. R. McAuliffe. Fuel-economy testing of a three-vehicle truck platooning system. Technical report, National Research Council Canada. Aerodynamics Laboratory, 2017.
- [158] P. Meisen, T. Seidl, and K. Henning. A Data-Mining Technique for the Planning and Organization of Truck Platoons. In *International Conference on Heavy Vehicles, Heavy Vehicle Transport Technology*, pages 389–402, 2008.
- [159] S. Melzer and B. Kuo. A closed-form solution for the optimal error regulation of a string of moving vehicles. *IEEE Transactions on Automatic Control*, 16(1):50–52, Feb. 1971.
- [160] M. Michaelian and F. Browand. Field Experiments Demonstrate Fuel Savings for Close-Following. Technical report, California Partners for Advanced Transportation Technology UC Berkeley, 2000.
- [161] F. Michaud, P. Lepage, P. Frenette, D. Letourneau, and N. Gaubert. Coordinated Maneuvering of Automated Vehicles in Platoons. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):437–447, Dec. 2006.
- [162] V. Milanés, S. E. Shladover, J. Spring, C. Nowakowski, H. Kawazoe, and M. Nakamura. Cooperative Adaptive Cruise Control in Real Traffic Situations. *IEEE Transactions on Intelligent Transportation Systems*, 15(1):296–305, Feb. 2014.
- [163] D. Mitra and A. Mazumdar. Pollution control by reduction of drag on cars and buses through platooning. *International Journal of Environment and Pollution*, 30(1):90–96, 2007.
- [164] S. Narayanaswami and S. Mohan. The roles of ICT in driverless, automated railway operations. *International Journal of Logistics Systems and Management*, 14(4):490–503, 2013.
- [165] G. Naus, R. Vugts, J. Ploeg, M. van de Molengraft, and M. Steinbuch. String-stable CACC design and experimental validation: a frequency-domain approach. *IEEE Transactions on Vehicular Technology*, 59(9):4268–4279, Sept. 2010.
- [166] D. Norrby. A CFD Study Of The Aerodynamic Effects Of Platooning Trucks. Master’s thesis, KTH, Mechanics, 2014.

- [167] A. Nourmohammadzadeh and S. Hartmann. The Fuel-Efficient Platooning of Heavy Duty Vehicles by Mathematical Programming and Genetic Algorithm. In C. Martín-Vide, T. Mizuki, and M. A. Vega-Rodríguez, editors, *5th International Conference on Theory and Practice of Natural Computing*, pages 46–57, Cham, Dec. 2016. Springer International Publishing.
- [168] C. Nowakowski, S. E. Shladover, X.-Y. Lu, D. Thompson, and A. Kailas. Cooperative Adaptive Cruise Control (CACC) for Truck Platooning: Operational Concept Alternatives. Technical report, UC Berkeley, Mar. 2015.
- [169] OECD/ITF. ITF Transport Outlook 2017. Technical report, OECD, 2017.
- [170] U. Ozguner, C. Stiller, and K. Redmill. Systems for Safety and Autonomous Behavior in Cars: The DARPA Grand Challenge Experience. *Proceedings of the IEEE*, 95(2):397–412, Feb. 2007.
- [171] R. M. Pagliarella. *On the Aerodynamic Performance of Automotive Vehicle Platoons Featuring Pre and Post-Critical Leading Forms*. PhD thesis, School of Aerospace, Mechanical and Manufacturing Engineering RMIT University, 2009.
- [172] M. Papageorgiou, C. Diakaki, V. Dinopoulou, A. Kotsialos, and Y. Wang. Review of road traffic control strategies. *Proceedings of the IEEE*, 91(12):2043–2067, Dec. 2003.
- [173] R. D. Pascoe and T. N. Eichorn. What is communication-based train control? *IEEE Vehicular Technology Magazine*, 4(4):16–21, Dec. 2009.
- [174] A. Perego, S. Perotti, and R. Mangiaracina. ICT for logistics and freight transportation: a literature review and research agenda. *International Journal of Physical Distribution & Logistics Management*, 41(5):457–483, 2011.
- [175] M. Pillado, L. García-Sol, J. Batlle, S. Sanchez, F. Freixas, S. Shadeghian, and T. Friedrichs. D7.1. Limited results of the on-board coordinated platooning system performance evaluation via physical testing. Technical report, Applus+ IDIADA, 2015.
- [176] J. Ploeg, N. van de Wouw, and H. Nijmeijer. Lp String Stability of Cascaded Systems: Application to Vehicle Platooning. *IEEE Transactions on Control Systems Technology*, 22(2):786–793, Mar. 2014.
- [177] G. K. Rajamani. CFD Analysis of Air Flow Interactions in Vehicle Platoons. Master’s thesis, RMIT University, 2006.
- [178] M. Ramezani and N. Geroliminis. On the estimation of arterial route travel time distribution with Markov chains. *Transportation Research Part B: Methodological*, 46(10):1576–1590, Dec. 2012.

- [179] H. Raza and P. Ioannou. Vehicle following control design for automated highway systems. *IEEE Control Systems Magazine*, 16(6):43–60, Dec. 1996.
- [180] J. Rice and E. van Zwet. A simple and effective method for predicting travel times on freeways. *IEEE Transactions on Intelligent Transportation Systems*, 5(3):200–207, Sept. 2004.
- [181] J. Roberts, R. Mihelic, M. Roeth, and D. Rondini. CONFIDENCE REPORT: Two-Truck Platooning. Technical report, North American Council for Freight Efficiency, 2016.
- [182] J. Robson. Algorithms for maximum independent sets. *Journal of Algorithms*, 7(3):425–440, Sept. 1986.
- [183] J.-P. Rodrigue. *Cities, Regions and Flows*, chapter Supply chain management, logistics changes and the concept of friction. Routledge London, 2012.
- [184] J.-P. Rodrigue, C. Comtois, and B. Slack. *The geography of transport systems*. Routledge, 2009.
- [185] J.-P. Rodrigue, J. Debie, A. Fremont, and E. Gouvernal. Functions and actors of inland ports: European and North American dynamics. *Journal of Transport Geography*, 18(4):519–529, July 2010.
- [186] M. Roeth. CR England Peloton Technology Platooning Test Nov 2013. Technical report, North American Council for Freight Efficiency, 2013.
- [187] K. Roy and C. Tomlin. Enroute airspace control and controller workload analysis using a novel slot-based sector model. In *American Control Conference*, June 2006.
- [188] S. Sadeghhosseini and R. Benekohal. Space headway and safety of platooning highway traffic. In R. Benekohal, editor, *1997 Conference on Traffic Congestion and Traffic Safety in the 21st Century*, pages 472–478. ASCE, June 1997.
- [189] S. Sadeghian Borojeni, T. Friedrichs, W. Heuten, A. Lüdtkke, and S. Boll. Design of a Human-Machine Interface for Truck Platooning. In *2016 CHI Conference on Human Factors in Computing Systems*, CHI EA '16, pages 2285–2291, New York, NY, USA, May 2016. ACM.
- [190] SAE On-Road Automated Vehicle Standards Committee. Surface Vehicle Recommended Practice. Technical report, SAE International, Sept. 2016.
- [191] M. Saeednia and M. Menendez. Analysis of Strategies for Truck Platooning. *Transportation Research Record: Journal of the Transportation Research Board*, 2547:41–48, 2016.

- [192] M. Saeednia and M. Menendez. A Decision Support System for real-time platooning of trucks. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems*, pages 1792–1797, Nov. 2016.
- [193] M. Saeednia and M. Menendez. A Consensus-Based Algorithm for Truck Platooning. *IEEE Transactions on Intelligent Transportation Systems*, 18(2):404–415, Feb. 2017.
- [194] S. Samaranayake, S. Blandin, and A. Bayen. A tractable class of algorithms for reliable routing in stochastic networks. *Transportation Research Part C: Emerging Technologies*, 20(1):199–217, Feb. 2012.
- [195] P. Sanders and D. Schultes. Engineering Fast Route Planning Algorithms. In C. Demetrescu, editor, *Experimental Algorithms: 6th International Workshop*, pages 23–36, Rome, Italy, June 2007. Springer Berlin Heidelberg.
- [196] N. B. Sarter and D. D. Woods. Pilot Interaction With Cockpit Automation: Operational Experiences With the Flight Management System. *The International Journal of Aviation Psychology*, 2(4):303–321, Nov. 1992.
- [197] B. Scholz-Reiter, K. Windt, and M. Freitag. Autonomous logistic processes: New demands and first approaches. In *37th CIRP international seminar on manufacturing systems*, pages 357–362, 2004.
- [198] M. Segata, B. Bloessl, S. Joerer, F. Dressler, and R. L. Cigno. Supporting platooning maneuvers through IVC: An initial protocol analysis for the JOIN maneuver. In *2014 11th Annual Conference on Wireless On-demand Network Systems and Services*, pages 130–137, Apr. 2014.
- [199] S. Shladover, D. Su, and X.-Y. Lu. Impacts of Cooperative Adaptive Cruise Control on Freeway Traffic Flow. *Transportation Research Record: Journal of the Transportation Research Board*, 2324:63–70, 2012.
- [200] D. C. Shoup. Cruising for parking. *Transport Policy*, 13(6):479–486, Nov. 2006.
- [201] R. Sims, R. Schaeffer, F. Creutzig, X. Cruz-Núñez, M. D’Agosto, D. Dimitriu, M. J. F. Meza, L. Fulton, S. Kobayashi, O. Lah, A. McKinnon, P. Newman, M. Ouyang, J. J. Schauer, D. Sperling, and G. Tiwari. *Climate change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate*, chapter Transport, pages 599–670. Cambridge University Press, 2014.
- [202] C. Siripanpornchana, S. Panichpapiboon, and P. Chaovalit. Travel-time prediction with deep learning. In *2016 IEEE Region 10 Conference*, pages 1859–1862, Nov. 2016.

- [203] J. Smith, R. Mihelic, B. Gifford, and M. Ellis. Aerodynamic Impact of Tractor-Trailer in Drafting Configuration. *SAE International Journal of Commercial Vehicles*, 7(2):619–625, Sept. 2014.
- [204] Socioeconomic Data and Application Center. Population Density Grid, v3, 2000, 2015.
- [205] V. Sokolov, J. Larson, T. Munson, J. Auld, and D. Karbowski. Platoon formation maximization through centralized routing and departure time coordination, Jan. 2017. URL <https://arxiv.org/abs/1701.01391>.
- [206] A. Spulber. Impact of Automated Vehicle Technologies on Driver Skills. *Connected Auto*, Fall 2016:46–57, Nov. 2016.
- [207] K. K. Srinivasan, A. Prakash, and R. Seshadri. Finding most reliable paths on networks with correlated and shifted log-normal travel times. *Transportation Research Part B: Methodological*, 66:110–128, Aug. 2014. Advances in Equilibrium Models for Analyzing Transportation Network Reliability.
- [208] T. Stanger and L. del Re. A model predictive Cooperative Adaptive Cruise Control approach. In *2013 American Control Conference*, pages 1374–1379, June 2013.
- [209] D. Steenken, S. Voß, and R. Stahlbock. Container terminal operation and operations research - a classification and literature review. *OR Spectrum*, 26(1):3–49.
- [210] G. Stefansson and K. Lumsden. Performance issues of Smart Transportation Management systems. *International Journal of Productivity and Performance Management*, 58(1):55–70, 2008.
- [211] H. Sternberg, T. Germann, and T. Klaas-Wissing. Who controls the fleet? Initial insights into road freight transport planning and control from an industrial network perspective. *International Journal of Logistics Research and Applications*, 16(6):493–505, Dec. 2013.
- [212] X. Sun, L. Munoz, and R. Horowitz. Highway traffic state estimation using improved mixture Kalman filters for effective ramp metering control. In *42nd IEEE International Conference on Decision and Control*, volume 6, pages 6333–6338, Dec. 2003.
- [213] D. Swaroop, J. Hedrick, C. Chien, and P. Ioannou. A comparison of spacing and headway control laws for automatically controlled vehicles. *Vehicle System Dynamics*, 23(1):597–625, July 1994.
- [214] D. Swaroop and J. K. Hedrick. String stability of interconnected systems. *IEEE Transactions on Automatic Control*, 41(3):349–357, Mar. 1996.

- [215] B. J. Tetreault. Use of the Automatic Identification System (AIS) for maritime domain awareness (MDA). In *OCEANS 2005*, volume 2, pages 1590–1594, Sept. 2005.
- [216] F. Torrey. Cost of Congestion to the Trucking Industry: 2017 Update. Technical report, American Transportation Research Institute, 2017.
- [217] S. Tsugawa. An Overview on an Automated Truck Platoon within the Energy ITS Project. *Advances in Automotive Control*, 7:41–46, 2013.
- [218] S. Tsugawa, S. Jeschke, and S. E. Shladover. A Review of Truck Platooning Projects for Energy Savings. *IEEE Transactions on Intelligent Vehicles*, 1(1):68–77, Mar. 2016.
- [219] S. Tsugawa, S. Kato, T. Matsui, H. Naganawa, and H. Fujii. An architecture for cooperative driving of automated vehicles. In *Intelligent Transportation Systems Conference*, pages 422–427, Dearborn, USA, 2000.
- [220] S. Tsugawa, S. Kato, K. Tokuda, T. Matsui, and H. Fujii. A cooperative driving system with automated vehicles and inter-vehicle communications in Demo 2000. In *IEEE Intelligent Transportation Systems Proceedings*, pages 918–923, Aug. 2001.
- [221] H. Tu, J. van Lint, and H. van Zuylen. Impact of Traffic Flow on Travel Time Variability of Freeway Corridors. *Transportation Research Record: Journal of the Transportation Research Board*, 1993:59–66, 2007.
- [222] V. Turri, B. Besselink, J. Mårtensson, and K. H. Johansson. Fuel-efficient heavy-duty vehicle platooning by look-ahead control. In *53rd IEEE Conference on Decision and Control*, pages 654–660, Dec. 2014.
- [223] B. van Arem, C. J. G. van Driel, and R. Visser. The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):429–436, Dec. 2006.
- [224] S. van de Hoef. *Fuel-Efficient Centralized Coordination of Truck Platooning*. KTH Royal Institute of Technology, June 2016. Licentiate Thesis.
- [225] S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Coordinating Truck Platooning by Clustering Pairwise Fuel-Optimal Plans. In *18th IEEE International Conference on Intelligent Transportation Systems*, pages 408–415, Sept. 2015.
- [226] S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Fuel-Optimal Coordination of Truck Platooning Based on Shortest Paths. In *American Control Conference*, pages 3740–3745, Chicago, IL, July 2015.

- [227] S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Computing Feasible Vehicle Platooning Opportunities for Transport Assignments. *14th IFAC Symposium on Control in Transportation Systems*, 49(3):43–48, May 2016.
- [228] S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Efficient Dynamic Programming Solution to a Platoon Coordination Merge Problem With Stochastic Travel Times. *20th IFAC World Congress*, 50(1):4228–4233, July 2017.
- [229] S. van de Hoef, K. H. Johansson, and D. V. Dimarogonas. Fuel-Efficient En Route Formation of Truck Platoons. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):102–112, Jan. 2018.
- [230] S. van de Hoef, J. Mårtensson, D. V. Dimarogonas, and K. H. Johansson. A Predictive Framework for Dynamic Heavy-Duty Vehicle Platoon Coordination. *ACM Transactions on Cyber-Physical Systems*, 2017. submitted.
- [231] J. van Lint, H. J. van Zuylen, and H. Tu. Travel time unreliability on freeways: Why measures based on variance tell only half the story. *Transportation Research Part A: Policy and Practice*, 42(1):258–277, Jan. 2008.
- [232] M. van Schijndel-de Nooij, B. Krosse, T. van den Broek, S. Maas, E. van Nunen, H. Zwijnenberg, A. Schieben, H. Mosebach, N. Ford, M. McDonald, D. Jeffery, J. Piao, and J. Sanchez. Definition of necessary vehicle and infrastructure systems for Automated Driving. Technical report, European Commission, 2011.
- [233] P. Varaiya. Smart cars on smart roads: problems of control. *IEEE Transactions on Automatic Control*, 38(2):195–207, Feb. 1993.
- [234] W. Vassallo. Freight market Structure and Requirements for Intermodal Shifts. Technical report, AMRIE, 2007.
- [235] P. Vegendla, T. Sofu, R. Saha, M. Madurai Kumar, and L.-K. Hwang. Investigation of Aerodynamic Influence on Truck Platooning. In *SAE 2015 Commercial Vehicle Engineering Congress*. SAE International, Sept. 2015.
- [236] A. Vinel, L. Lan, and N. Lyamin. Vehicle-to-vehicle communication in C-ACC/platooning scenarios. *IEEE Communications Magazine*, 53(8):192–197, Aug. 2015.
- [237] Y. Wang and M. Papageorgiou. Real-time freeway traffic state estimation based on extended Kalman filter: a general approach. *Transportation Research Part B: Methodological*, 39(2):141–167, Feb. 2005.
- [238] Y. Wang, V. S. Rodrigues, and L. Evans. The use of ICT in road freight transport for CO2 reduction – an exploratory study of UK’s grocery retail industry. *The International Journal of Logistics Management*, 26(1):2–29, 2015.

- [239] Z. Wang, A. Goodchild, and E. McCormack. Measuring Truck Travel Time Reliability Using Truck Probe GPS Data. *Journal of Intelligent Transportation Systems*, 20(2):103–112, Apr. 2016.
- [240] Z. Wang, G. Wu, P. Hao, K. Boriboonsomsin, and M. Barth. Developing a platoon-wide Eco-Cooperative Adaptive Cruise Control (CACC) system. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1256–1261, June 2017.
- [241] S. Watkins and G. Vino. The effect of vehicle spacing on the aerodynamics of a representative car shape. *Journal of Wind Engineering and Industrial Aerodynamics*, 96(6):1232–1239, June 2008. 5th International Colloquium on Bluff Body Aerodynamics and Applications.
- [242] A. Watts. Computational Characterization of Drag Reduction for Platooning Heavy Vehicles. Master’s thesis, Auburn University, 2015.
- [243] B. S. Westgate, D. B. Woodard, D. S. Matteson, and S. G. Henderson. Travel time estimation for ambulances using Bayesian data augmentation. *The Annals of Applied Statistics*, 7(2):1139–1161, June 2013.
- [244] M. Whaiduzzaman, M. Sookhak, A. Gani, and R. Buyya. A survey on vehicular cloud computing. *Journal of Network and Computer Applications*, 40(Supplement C):325–344, Apr. 2014.
- [245] S. Wickström. D3.1 Component Specifications for the Overall Architecture. Technical report, Scania CV AB, 2015.
- [246] H. Winner and M. Schopper. *Handbook of Driver Assistance Systems*, chapter Adaptive Cruise Control, pages 1093–1148. Springer International Publishing, Cham, 2016.
- [247] H. Wolf-Heinrich and S. R. Ahmed. *Aerodynamics of Road Vehicles*. Society of Automotive Engineers, 1998.
- [248] C.-H. Wu, J.-M. Ho, and D. T. Lee. Travel-time prediction with support vector regression. *IEEE Transactions on Intelligent Transportation Systems*, 5(4):276–281, Dec. 2004.
- [249] X. Yang, X. Li, B. Ning, and T. Tang. A Survey on Energy-Efficient Train Operation for Urban Rail Transit. *IEEE Transactions on Intelligent Transportation Systems*, 17(1):2–13, Jan. 2016.
- [250] R. Yoshimoto and T. Nemoto. The Impact of Information and Communication Technology on Road Freight Transportation. *IATSS Research*, 29(1):16–21, 2005.

- [251] M. Zabat, N. Stabile, S. Farascarioli, and F. Browand. The Aerodynamic Performance Of Platoons: A Final Report. Technical report, Institute of Transportation Studies, University of California Berkeley, Jan. 1995.
- [252] J. Zhang and P. A. Ioannou. Longitudinal control of heavy trucks in mixed traffic: environmental and fuel economy considerations. *IEEE Transactions on Intelligent Transportation Systems*, 7(1):92–104, Mar. 2006.
- [253] W. Zhang, E. Jenelius, and X. Ma. Freight transport platoon coordination and departure time scheduling under travel time uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, 98(Supplement C):1–23, Feb. 2017.
- [254] W. Zhang, M. Sundberg, and A. Karlström. Platoon coordination with time windows: an operational perspective. In *20th EURO Working Group on Transportation Meeting, EWGT 2017*, volume 27, pages 357–364, Sept. 2017.
- [255] Y. J. Zhang, A. A. Malikopoulos, and C. G. Cassandras. Optimal control and coordination of connected and automated vehicles at urban traffic intersections. In *2016 American Control Conference*, pages 6227–6232, July 2016.