Safe, Efficient, and Fair

—

A Top-Down Approach to Inter-Vehicle Communication

Safe, Efficient, and Fair
—
A Top-Down Approach to Inter-Vehicle Communication

Inaugural-Dissertation

zur Erlangung des Doktorgrades der
Mathematisch-Naturwissenschaftlichen Fakultät
der Heinrich-Heine-Universität Düsseldorf

vorgelegt von
Daniel Baselt
aus Krefeld

Düsseldorf, Juni 2013

Aus dem Institut für Informatik
der Heinrich-Heine-Universität Düsseldorf

Gedruckt mit der Genehmigung der
Mathematisch-Naturwissenschaftlichen Fakultät der
Heinrich-Heine-Universität Düsseldorf

Referent: Prof. Dr. Martin Mauve
Heinrich-Heine-Universität Düsseldorf

Korreferent: Prof. Dr. Björn Scheuermann
Humboldt-Universität zu Berlin

Tag der mündlichen Prüfung: 19. Juli 2013

Abstract

Since the development of the first automobiles, individual traffic changed in many respects while the main objectives of transportation—safety and efficiency—have stayed the same. Perhaps the most obvious change is that traffic got denser. To support drivers coping with more and more cars that occupy the roads, car manufacturers developed aids of various kinds that assist driving according to those objectives. The assistants deployed today rely on information made available by a vehicle's on-board sensors and the decisions based on that data are made isolated from other cars. Cooperative decisions of vehicles are enabled by the active exchange of information using inter-vehicle communication. For example, drivers can coordinate with each other about route selections or can warn about hazards like the rear end of a traffic jam in time. The availability of low-cost wireless technology in recent years again inspired the network research community to develop novel protocols and algorithms for inter-vehicle communication. These were then used to build applications supporting the transportation objectives. The developed protocols are based on available or slightly modified standards; their suitability to the applications is limited by the technologies' constraints. Currently, car-to-car networks are designed as general-purpose networks, although the objectives for employing communication for exchanging car information are clear and very specific: optimize traffic safety and efficiency. In addition, the understanding of the application domain was not in focus in the depth that it deserves. With the knowledge available at the time work on this thesis was commenced, it was not possible to measure how well a protocol performs in view of the above objectives. There was in fact no information at all on how to relate available information in vehicles, how such information is exchanged, and the application objectives; the potential that inter-vehicle communication bears was unknown. But without a solid understanding of how these aspects interrelate, a protocol developer cannot be sure that what is pursued is the best—or even just a reasonably good—way of supporting road traffic by means of communication.

This thesis proposes to complement the existing research on inter-vehicle communication with a *top-down approach*. Instead of starting out with a given technology and then developing protocols bottom-up, this approach begins with a formal discussion of what does not change: the objectives of traffic. For these, an optimal behavior of cars and, with this, information demands are deduced that good protocols have to satisfy. Thereafter, according protocols are developed, followed by a suitable communication technology. Because this is an enormous task, a roadmap is presented that describes how to properly divide it. In this thesis, the approach is applied to two basic road topologies up to the point a protocol is obtained.

Followed by a comprehensive discussion of related work in the different areas of research that affect the top-down approach, the first steps of the road map are applied to a simplistic scenario. The influence of available information on the vehicles' behavior is considered. Different schemes for information exchange are analyzed including beaconing, a periodic information exchange. It turns out that a car's optimal behavior with beaconing is to drive with alternating periods of acceleration and deceleration which converge to a steady state.

The scenario is extended to multiple cars and analyzed for the optimal sending times of information for a cruise control application. This provides us with detailed knowledge about how beaconing should work in a reliable communication environment. The discussion continues with considering the effects of packet losses and it is found that not only the direct follower of a sending car leaves its steady state, but a chain reaction affects numerous upstream cars. To withstand this effect, and in particular to cope with consecutive losses, the steady state distances are enlarged by multiples with beaconing. The then proposed algorithm *Carrot* is able to detect losses implicitly without a need for acknowledgments and to react via a fast repetition of the missing beacon. Analytical and simulative evaluations show that Carrot is able to repeat beacons in fractions of a beaconing interval, so that the steady state distances can be chosen very close to the minimum.

Communication enables cars to cooperate and thereby allows for a third objective besides safety and efficiency that is not yet in focus of research for inter-vehicle communication: *fairness*. Towards this, this thesis contributes the definition of a fairness criterion for cooperation. This criterion is applied to the ordering of cars at a merging of two lanes. An analysis of the zipper merge, the only relevant merging scheme for today's cars if no lane is given the right of way, clearly shows an inherent unfairness. A coordination scheme is proposed that creates an optimal fair merge order. The scheme is adapted for distributed decisions with local knowledge and communication using a beacon-based approach. Evaluations which also accounts for unreliable wireless communication describe the influence of the ratio of participants on merging order fairness. The results show that the algorithm yields very good fairness even if only a small percentage of about 1 % of the cars follows the algorithm's guidance.

Zusammenfassung

Seit der Entwicklung der ersten Automobile hat sich der Individualverkehr in vielerlei Hinsicht gewandelt, allerdings sind die zentralen Ziele des Verkehrs – Sicherheit und Effizienz – gleich geblieben. Eine der sichtbarsten Veränderungen ist die Zunahme des Verkehrsaufkommens. Um auch bei den heute üblichen hohen Verkehrsdichten sicher und effizient fahren zu können, werden Fahrer moderner Fahrzeuge durch verschiedenste Assistenzsysteme unterstützt, welche auf Basis von Informationen arbeiten, die durch die im Fahrzeug vorhandenen Sensoren wie GPS, Radar und Lidar erzeugt werden. Die mit diesen Daten möglichen Entscheidungen werden isoliert von anderen Verkehrsteilnehmern getroffen. Erst durch den aktiven Austausch von Informationen zwischen Fahrzeugen können Entscheidungen auch gemeinschaftlich gefunden werden. So kann auf diese Weise die Wahl von Fahrtrouten koordiniert werden oder rechtzeitig vor Gefahren gewarnt werden, wie sie beispielsweise durch ein Stauende an einem uneinsehbaren Streckenteil entstehen. In den letzten Jahren hat die Forschung in diesem Bereich durch die Verfügbarkeit von kostengünstiger drahtloser Kommunikationstechnologie erneut Aufwind erhalten, es wurden neuartige Protokolle und Algorithmen für diese Technologien entwickelt, welche dann für die Konzeption von den Fahrer assistierenden Anwendungen verwendet wurden. Da jedoch die zugrundeliegenden Technologien bestehende oder lediglich leicht modifizierte Standards nutzen, ist die Anpassbarkeit der Protokolle begrenzt. Bisherige Entwürfe für Kommunikationssysteme der Fahrzeug-zu-Fahrzeug-Kommunikation setzen auf Protokolle, die für universelle Anwendbarkeit ausgelegt sind. Allerdings lassen sich die wesentlichen Anwendungen klar und präzise bestimmen: es geht um Verkehrssicherheit und die effiziente Nutzung verfügbarer Ressourcen. Hinzu kommt, dass ein definierbares Verständnis der Anwendungen bisher nicht im Fokus der Forschung in diesem Bereich war. Mit dem zum Beginn der Arbeit an dieser Dissertation verfügbaren Wissen und Werkzeugen war es nicht möglich zu bestimmen, wie gut sich ein Protokoll hinsichtlich von Anwendungszielen in einem absolutem Maßstab verhält. Es fehlten Informationen darüber, wie verfügbare Informationen, Methoden des Informationsaustausches und Anwendungsziele miteinander in Beziehung stehen. Dadurch war auch das tatsächliche Potenzial der Fahrzeug-zu-Fahrzeug-Kommunikation nicht bekannt, denn ohne eine

verlässliche Kenntnis dieser Beziehungen können Entwickler nicht sicher sein, dass ihre Kommunikationsprotokolle dabei helfen, die Ziele im Straßenverkehr bestmöglich zu erreichen.

Die vorliegende Arbeit beschreibt, wie die bestehenden Forschungsergebnisse in der Fahrzeug-zu-Fahrzeug-Kommunikation durch einen *Top-Down-Ansatz* ergänzt werden kann. Anstelle bestehende Standards vorauszusetzen und darauf Protokolle bottom-up aufzubauen, werden mit diesem Ansatz zunächst die bestimmenden Ziele im Straßenverkehr formal definiert, da diese sich im Vergleich zu Technologien längeren Bestand gezeigt haben. Für diese Ziele wird dann das optimale Verhalten von Fahrzeugen bestimmt und mit diesem Eigenschaften für gute Protokolle abgeleitet. Damit ist die Möglichkeit gegeben, entsprechende Protokolle und letztlich passende Kommunikationstechnologien zu entwerfen. Dies ist eine immense Aufgabe, weshalb in dieser Arbeit zunächst deren Aufteilung zu einem mehrschrittigen Vorgehen beschrieben wird. Danach wird der Top-Down-Ansatz auf zwei grundlegende Situationen im Straßenverkehr zur Entwicklung angepasster Protokolle angewendet.

Nach einer umfassenden Diskussion verwandter Arbeiten aus den verschiedenen Forschungsbereichen, die von einem Top-Down-Ansatz berührt werden, werden die ersten Schritte des Ansatzes mit einem einfachen Szenario diskutiert und der Einfluss verfügbarer Informationen auf das Verhalten von Fahrzeugen betrachtet. Unterschiedliche Methoden des Informationsaustausches werden analysiert, unter anderem ein idealisiertes Beaconing, ein Verfahren für den periodischen Austausch von Informationen. Es wird festgestellt, dass das optimale Verhalten eines Fahrzeugs ein abwechselndes Bremsen und Beschleunigen ist, welches zu einem stabilen Zustand strebt.

Darauffolgend wird das Szenario erweitert, um die optimalen Sendezeiten von mehreren Fahrzeugen für ein Geschwindigkeitsregelungssystem in einer Kolonne zu untersuchen. Diese Diskussion vertieft das Wissen darüber, wie Beaconing mit zuverlässiger Kommunikation funktionieren sollte. Darauf baut die anschließende Betrachtung von unzuverlässiger Kommunikation auf, bei der gezeigt wird, dass ein Paketverlust nicht nur einen direkten Hintermann in der Kolonne beeinflusst, sondern eine Kettenreaktion die stabilen Zustände zahlreicher folgender Fahrzeuge stört. Um diesem Effekt entgegenzuwirken, der besonders schwer bei mehreren Verlusten in Folge wiegt, wird bei Beaconing der stabile Zustand um ein Mehrfaches vergrössert. Mit dem dann vorgestellten Protokoll *Carrot* ist es möglich, Verluste implizit, also ohne zusätzliche Pakete, zu erkennen und auf diese mit einer schnellen Übertragungswiederholung zu reagieren. Analytische und simulative Auswertungen zeigen, dass Carrot Beacons in Bruchteilen eines Sendeintervalls wiederholt und damit ein stabiler Zustand nahe dem Minimum von verlustfreier Kommunikation gewährleistet werden kann.

Die Fahrzeug-zu-Fahrzeug-Kommunikation ermöglicht die Unterstützung eines weiteren Zieles im Straßenverkehr neben Sicherheit und Effizienz, welches bislang nicht im Fokus der Forschung lag: Ein weiterer Beitrag dieser Arbeit ist die Definition des Kriteriums *Fairness* für die Kooperation von Fahrzeugen. Dieses Kriterium wird auf die Reihenfolge von Fahrzeugen bei der Überfahrt eines Spurzusammenschlusses angewendet. Eine Untersuchung des Reißverschlussverfahrens, dem einzig Relevanten wenn keiner Spur die Vorfahrt gegeben ist, mit diesem Kriterium zeigt dabei klar eine grundlegende Unfairness. Ein Verfahren zur fairen Koordination wird vorgestellt, welches optimale Fairness erreicht. Dieses wird dann abgewandelt, um eine verteilte Entscheidungsfindung mit lokalem Wissen und Beacon-basierter Kommunikation zu ermöglichen. Auswertungen wurden mit unzuverlässiger Kommunikation durchgeführt, um den Einfluss der Ausstattungsrate der Fahrzeuge auf die Fairness der Reihenfolge von Überfahrten zu bewerten. Die Ergebnisse zeigen, dass das entwickelte Protokoll die Fairness deutlich erhöht, selbst wenn nur ein kleiner Anteil der Fahrzeuge von 1 % den Empfehlungen folgt.

Acknowledgments

Although this may sound a cliché, on a long journey—and exactly this is this now eventually finished thesis to me—you are likely to meet many different people. Luckily for me, by far the most of them that accompanied me at some time or another during my research were kind and helped me in manifold ways. I am deeply grateful to all them. However, there are a few people that deserve a special "thank you".

Certainly, the most did my Ph.D. supervisors have to endure, especially when I came up with new notes of pure wisdom (at least in my very own opinion) they had to work through. During our discussions, Martin stressed—several times—that doing a Ph.D. is not an apprenticeship, because a doctoral researcher already holds a university degree, but he taught me so much about scientific work that I must, honestly, view my years of doctoral research as such. Björn always gave feedback to my ideas that was so detailed and to the point even though the distance between our offices grew significantly during the past years; he read and commented my notes over and over and I can only admit that discussing with him is worth even the furthest way.

My colleagues helped making the years of work enjoyable and discussing with them was so productive it every time brought me another step closer to finishing this thesis. From encouragement at the start from Michael and Jedreck, over teaching and researching together with Markus and Norbert, up to sharing the office with me, as Yves mastered bravely. This all, indeed, would not have been enough without access to our great computer infrastructure which was cared for by Christian during my first years and was then passed on to Thomas. You two did a great job! This accounts for Sabine, too; without her organizing skills, I'm sure, none of my projects would have finished but stuck somewhere between the big gears of administration.

Another important part of my project's success is owed to my student helpers. Andre, Franz, and Raphael from the Volkswagen research projects, as well as Erzen, Nasir, and Tobias from the NRW car-to-car project, and—no less—Daniel, who investigated vehicle simulators close to the metal. Some of them wrote their theses under my supervision, and it was rewarding to see how they pushed topics forward that are closely related to this thesis. All the other students I supervised, without mentioning them all, also helped cutting the way for my research with so much energy that it was joyful for me teaching them.

However, I would not even have started this thesis if not my family—my parents, my wife—had promised to support me in this project. I'm so glad they did; and they really did support me wherever possible. Thank you.

Contents

List of Figures

List of Tables

Nomenclature

ACC	Adaptive/Autonomous Cruise Control
C2CC	Car-to-Car Communication
C2IC	Car-to-Infrastructure Communication
C2XC	Car-to-X Communication
CACC	Cooperative Adaptive Cruise Control
CALM	Communications, Air-interface, Long and Medium range
CAMP	Crash Avoidance Metrics Partnership
CCA	Cooperative Collision Avoidance
CPS	Cyber-Physical System
CSMA	Carrier Sense Multiple Access
DGPS	Differential Global Positioning System
DSRC	Dedicated Short-Range Communication
ETSI	European Telecommunications Standards Institute
FCFS	first-come first-served
GPS	Global Positioning System
IDM	Intelligent Driver Model
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
MAC	Medium Access Control
MANET	Mobile Ad-hoc Network
ns	network simulator
OFDM	Orthogonal Frequency Division Multiplexing
TCP	Transmission Control Protocol

TDMA	Time-Division Multiple Access
UMTS	Universal Mobile Telecommunications System
VANET	Vehicular Ad-hoc Network
VSC	Vehicle Safety Communications
WAVE	Wireless Access in Vehicular Environments
WSMP	WAVE Short Message Protocol

Chapter 1

Introduction

The main chapters of this thesis are partially based on co-reviewed international publications of which I was the main contributing author. The following introduction contains adapted and extended elements of the publication [BMS11]. Funding for parts of this thesis has been provided by the state of North Rhine-Westphalia and the European Union within the NRW-EU Ziel 2 program 2007-2013 (EFRE).

Since the development of the first automobiles, traffic has changed in many respects. To name only a few, a wide network of roads has been constructed from large highways down to small streets connecting every home. Cars have become faster, more comfortable, and, most obviously, more abundant. Along with traffic getting denser, manifold safety features have been developed during the past decades to assist drivers in critical situations and to prevent them from getting into such situations in the first place. In spite of these improvements, actually driving a car takes up most of the concentration of the person behind the steering wheel. This has not changed. And coupled with this, the objectives of transportation like safety and efficiency, along with their impediments represented by the physical constraints a car has to cope with, have stayed the same as well.

In spite of the technological evolution of the automobile, even the highest concentration level of drivers cannot guarantee safety; nor can it ensure driving efficiency with regard to traveling time or resource consumption at present. But there are aids of various kinds in modern car that address these problems: for example, such aids help drivers navigate streets, avoid traffic jams, keep a safe distance from cars in front of them, or change lanes. These applications, to name a few, rely on information that is made available by a car's sensors and gathered in knowledge bases. These aids make decisions independently of the information that other cars have; the decision process is autonomous.

The active exchange of information between cars enables drivers to decide on actions in a cooperative manner. In this way communication has the potential to dramatically improve road safety and traffic efficiency. Many road accidents could be avoided if drivers had better information about the status of other cars and the intentions of their drivers. For example, if drivers were warned about the back of a traffic jam in time or if a driver approaching an intersection had information about crossing traffic, these traffic situations could be much safer. Similarly, if drivers coordinated their selection of routes and their driving behavior by exchanging information, road traffic could be much more efficient.

This insight is nothing new: it has motivated researchers as well as car and communication equipment manufacturers to establish an entire research community over the past decades. Recently, the widespread and low-cost availability of wireless technology for local and long-range communication sparked a new interest in inter-vehicle communication, combined with the desire to apply this to the exchange of information between vehicles. Starting out with available or slightly modified communication technology, the inter-vehicle communication research community developed novel network protocols and distributed algorithms in order to support the exchange of information [ZC06, TWB10, CLL10]. These were then used to create applications such as traffic information systems [NDLI04, IW08] or intersection warning assistants [Con05]. It is quite likely that the applications developed in this way will help to avoid accidents and reduce the usage of resources to some extent.

However, despite these foreseeable benefits, research in this area today has not determined whether the developed protocols and applications really do make use of the technology's full potential—or whether they are merely scratching the surface. Regardless of which specific protocols and algorithms are proposed, questions remain: if other information was transmitted between vehicles, could even more accidents perhaps be prevented? Or might it be possible to reduce resource consumption even further if the data exchange took place in a different manner?

One main reason why these questions cannot be answered today is that the available understanding of the application domain itself is still very limited. In the areas of traffic theory and control theory, many related works already exist on how to model the behavior of traffic flows and how to control and optimize vehicle behavior, albeit from highly abstract points of view. For the latter, for example, communication is usually modeled as a source of uncertainty. It reduces the precision of information due to packet drops and other communication outages. Inter-vehicle communication, however, is more than pure information distribution: it enables interaction. Through the exchange of information, it is possible to influence the behavior of individual vehicles. At the time work on this thesis was commenced, there existed virtually no information at all on the interplay between available information in cars, how such information is exchanged, and traffic safety and efficiency. The ideal behavior of individual cars in order to optimize traffic safety and efficiency was unknown in a setting in which cars are able to communicate with their environment. Without a solid understanding of how these aspects interrelate, however, any approach to inter-vehicle communication can be a coarse heuristic at best: a protocol developer can do “something” and will likely obtain “some” improvement, but cannot be sure that what is pursued is the best—or even a reasonably good—way of assisting drivers by means of communication.

In some sense, current work on inter-vehicle communication is therefore similar to finding a good solution to an unspecified problem—the desired system behavior is not actually known, but developers are nevertheless busy specifying protocols and applications that aim to achieve it. The protocols found in this way are optimized for a given technology even though technology permanently changes and evolves. In the long run, the traffic objectives, and especially those of applications affecting traffic safety and efficiency, have to be of greater importance than the restrictions of a certain technology. We should be interested in what does not change; we have to understand what the objectives of safe and efficient transportation are as well as the properties which good protocols have to satisfy.

In the following chapters, how vehicles behave optimally to maintain safety, efficiency, and fairness is formally discussed. The term safety, to be more precise, is used to imply the absence of accidents at any time. Efficiency is aimed at road usage; each car has the objective to occupy the road for the shortest time possible while driving. The third aim is only possible with the exchange of information: cars cooperate to determine the fairest behavior. The formal discussion is accomplished by the development of a highly simplified model at the beginning of each of Chapters Chapters 3 to 5. Throughout each of these chapters, the respective model is altered stepwise to meet more realistic requirements. Although each of the models only describes an individual scenario, it will be shown how essential information for vehicles can be deduced from them and the properties of an optimal information dissemination strategy will be discussed. Communication protocols are proposed comprising these properties in Chapters 4 and 5. The knowledge about the optimal strategy allows us to evaluate the quality of the new protocols in absolute terms and relate them to already existing protocols. These metrics also serve future developers as a tool that shows which aspects of a protocol should be changed to obtain the highest benefit for the respective application. Thus this thesis presents how inter-vehicle communication is applied to traffic demands in a *top-down* manner: from application objectives to protocols.

With this short overview of the course in each main chapter in mind, the thesis is structured as follows. Chapter 2 contains a thorough overview of the work in related areas of research which are touched on by the top-down approach. Besides obvious areas such as literature on car-to-car communication (C2CC) protocol design and evaluation, a look is taken at existing car-following models from the field of traffic theory and at papers from control theory that model vehicle behavior for pursuing objectives. Finally, real-world experiments and the progress of research on fully automated driving are examined.

The breadth of related work areas illustrates that an exhaustive application of the top-down approach requires expertise on multiple research areas including traffic theory, wireless network protocols, and mathematics. The task of modeling traffic in all of its relevant details and of creating new communication protocols and perhaps even hardware platforms is too complex to be accomplished within a single thesis. Chapter 3 accordingly introduces the core idea of the top-down approach in an in-depth manner together with a road map that arranges the different steps of the approach. The first steps of the road map are applied to a simplified scenario. Starting with the creation of a model of cars driving in one lane, the two traffic objectives road usage and absence of accidents are discussed. Then, the influence of available information on car-following behavior is considered. The cars are given initial knowledge about the situation, and the two extremes of refreshing information are analyzed: complete, continuous knowledge in contrast to no new information at all. The findings of the latter are extended to describe periodic information updates, which is analogous to a protocol referred to as beaconing. Without networking layer influences such as delays or packet losses, it turns out that a following car's optimal behavior guided by information through beaconing is to drive with alternating periods of acceleration and deceleration towards a steady state with its predecessor.

With these first results from employing the top-down approach for two cars, the scenario is extended to multiple cars in Chapter 4. The optimal beacon sending times for a given bandwidth are determined with regard to a minimum sum of the cars' steady state distances. It is then turned closer to reality by omitting the objective of maximum efficiency in favor of enabling a steady following distance with constant speed. Again in view of a minimum sum of steady state distances, a car's sending times are considered and the optimum is determined. With sound knowledge of how beaconing should work in a reliable communication environment, the discussion continues with the introduction of delays and packet losses. The effect of a lost beacon on car distances is studied, and it is found that not only the direct follower of a sending car leaves its steady state, but a chain reaction is also triggered which affects numerous upstream cars. To withstand this loss effect, and in particular to cope with consecutive losses, the steady state distances are increased by multiples of the minimum steady state distances. With topology-related sending triggers, the then proposed algorithm *Carrot* is able to detect losses implicitly without the need for acknowledgments and to react via a fast repetition of the missing beacon. Analytical and simulative evaluations show that Carrot is able to repeat beacons in fractions of a beaconing interval so that the steady state distances can be chosen very close to the minimum.

In Chapter 5, the top-down concept is applied to a further central element of road networks: the merging of two lanes. First, the formal model for single lanes is extended to assist with mergings. An additional objective, the *fairness* of the merge order, is defined and subsequently applied to the zipper merge, the only relevant merging scheme for today's cars if no lane is given the right of way. The analysis clearly shows that the zipper merge is inherently unfair. But fairness is possible by employing a coordination scheme to control the merge order. This is demonstrated by the design of a merging algorithm that allows for an optimal fair merge order. The algorithm, however, relies on global knowledge of all cars, so it is adapted for local knowledge and communication using a beacon-based approach. Here, the algorithm decides on the merge order in a fully distributed manner. Beaconing is not reliable, as has been discussed in the previous chapter, and, in addition, it must be expected that the algorithm will not be built in or used in every car. To account for this, analytical as well as simulative evaluations describe the influence of the ratio of participants on the merging order fairness. The results show that the algorithm yields very good fairness even if only a small percentage of about 1 % of cars uses the algorithm and follows its guidance.

The thesis' concluding remarks are in Chapter 6. The appendix contains detailed proofs of the theorems and the lemmas that were too long for the running text.

The main contributions of this thesis are:

- The proposal of a novel approach to the design of inter-vehicle communication protocols.
- The application of the approach to cruise control in a platoon and the development of the Carrot protocol.
- The recommendation of the traffic objective fairness for a cooperative lane-merging scenario with design and evaluation of a protocol enabling fair merging using unreliable wireless communication.

Chapter 2

Related Work

A number of very distinct areas of research are involved in the attempt to model communication protocols for inter-vehicle communication with a top-down approach. This chapter covers works of the areas that come into play from the first step of understanding the application to the design and evaluation of specific protocols. It is started with an overview about existing papers that propose to do Car-to-Car Communication (C2CC) protocol design from the viewpoint of the applications that use this communication for exchanging information. All these publications differ from the thesis at hand by their objectives and methodologies—although many of the works motivate that an understanding of the application layer is important, the proposed approaches do not consider the application as extensive as done here. The concept of cyber-physical systems, which is explained thereafter, acknowledges the coupling between application and communication technology. Exemplary works of the hitherto mainly applied research direction, referred to as bottom-up in the following, are described to clarify in what sense it is oppositional to the top-down approach.

The applications modeled in the following chapters cover distinct basic parts of traffic scenarios. A short overview about possible application areas for C2CC related to this thesis is given in Section 2.2. The first application type described in detail is car following in a platoon as used for, e.g., cruise control systems. Not just in this section but throughout the whole thesis, it will come apparent that the development of a protocol with the top-down approach depends heavily on the modeling of the cars' behavior. The following top-down steps base on the car behavior and the choice of a proper model has therefore a major impact on them; that initial step of behavior modeling requires us to understand the motion of traffic in detail. In the field of traffic theory, various models have been proposed to describe the following of cars for different objectives. Existing car-following models are reviewed towards their suitability for a top-down approach in Section 2.2.1. Another related problem to the question of how to support drivers best is how to make cars drive on their own. Research groups work on partly as well as on fully automated cars that autonomously decide on their behavior. This particularly means that these cars usually do not decide by cooperation. Section 2.2.2 discusses examples of works on autonomous driving. The research in this field is supported by modeling vehicular behavior with control theory. A few works are outlined which describe the design of behavior controllers. A further kind of C2CC applications are those enabling cooperative behavior only possible through inter-vehicle communication; they get special attention in Section 2.2.3. In the literature are protocols especially for cooperation scenarios like lane mergings and intersections which are also outlined in this section.

In the area of inter-vehicle communication has been been a lot work on protocols, as explicated in Section 2.3. The layer model for networking protocols is briefly explained in Section 2.3.1. The network research community, together with car and equipment manufacturers, helped to standardize the physical and medium access layers and the coding of common messages. The standards with strong relation to this thesis are regarded in Section 2.3.2 where a focus is set on the European region. In the top-down approach, the medium access (MAC) is to be discussed after it has been understood how a communication protocol must look like to support an application. The protocols developed in this thesis were evaluated with existing medium access protocols; related work on designing MAC protocols is characterized in Section 2.3.3.

The application-layer protocol beaconing was early found to support many different applications, as it is shown in Section 2.4. This attracted many research groups to discuss adaption strategies that exploit characteristics of a Vehicular Ad-hoc Network (VANET), a type of C2CC within a local area, for enhancing network layer metrics. Among the adaptions are proposals to repeat beacons in various ways that are regarded in more detail in Section 2.4.2. A further method to enhance C2CC protocols is back-pressure; it has been employed in wireless networks to control the channel load by allowing packets to enter the network only after other packets left it, similar to water flowing through connected pipes. This is a related concept to the trigger-based scheme proposed in Chapter 4 and therefore a discussion of related work on back-pressure is covered in Section 2.4.3.

The evaluation of protocols in the C2CC research community is usually accomplished with simulations since experiments with real vehicles are very expensive. Simulators for computer networks as well as for vehicular movements are described in Section 2.5. The chapter closes by taking a glance at the rather few real-world experiments related to protocols and scenarios discussed in this thesis.

2.1 Protocol Design in Relation to Application Demands

In accordance to the top-down approach, several working groups of the C2CC research community recently stated that a thorough understanding of the application layer's goals is required. Only precise objectives enable a decision on proper metrics and working on optimal solutions. While basic work for an analytical approach on the scale of microscopic traffic behavior is made in this thesis, Gaugel et al. from Karlsruhe Institute of Technology discussed a simulative approach for the macroscopic level [GSEMHt11]. The future challenges for C2CC research were discussed in a more general range in a Dagstuhl seminar in 2010. There it has been stated that among the most pressing needs is a formal and fundamental understanding of the challenges and solutions. The seminar's results are summarized by Dressler et al. [DKO+11].

Designing C2CC protocols with respect to applications and their security demands is also proposed by Qian and Moayeri [QM08]. The authors argue that routing algorithms have to be aware about the applications running on a vehicle. They present a routing scheme that is intended to support multiple kinds of applications. However, this work does not explore the information requirements of vehicles in the same methodical and mathematical detail as it is done in this thesis.

The methodology of Tang and Yip resembles the top-down approach in looking at possible ways of cars and carrying out an analytical approach to model a car's behavior [TY10]. The authors consider the effects of delays that occur when sending wireless messages of collision avoidance applications, but in contrast to this thesis, only distinct ways of cars are regarded. This is due to largely differing intentions: the authors intend to understand effects of real-world delays while this thesis intends to understand how cars should behave for finding design guidelines for communication protocols to support car behavior through communication best.

2.1.1 Protocol Design in a Cyber-Physical System

The top-down approach considers a system in which physical, real-world components, the vehicles, interact with computation and digital communication. This is also referred to as a Cyber-Physical System (CPS). Lee describes the mutual interaction of physical processes and computing [Lee08]. He points out that, in contrast to general-purpose computing, the design of such systems requires to regard the connection of the physical and the computational domain to make use of a CPS's full potential. Fallah et al. characterize a VANET as a CPS and describe the tight coupling of physical vehicle dynamics, computing, and communication aspects [FHSK10]. They discuss this concept with regard to network performance and propose how a CPS should be designed and operated. The complexity of this task is explained by examples as the required work on the effect of physical vehicle dynamics. The interplay of the system components is discussed and a tightly coupled design is described. Simulative experiments are made and the results are compared to a system based on non-coupled design. While the paper relates the physical domain with communication, as done in this thesis, the authors focus on the influence of vehicular movements on network performance and do not close the loop back to the impact of information availability on vehicular behavior.

2.1.2 The Bottom-Up Approach to Protocol Design

The predominant methodology for C2CC protocol design in the literature is bottom-up: the design starts with a certain technology as given, for example communication hardware, link layer protocols, or network layer protocols. The existing publications then typically choose one of the following two directions to proceed: either they propose a network protocol and evaluate it with the help of simulations regarding issues of the network layer. Some works then tune the discussed protocol to fit vehicular movement characteristics better. Or, they exploit vehicular movements for creating a protocol that performs better on the network layer than chosen reference protocols. The downside of both bottom-up methodologies is that the communication protocols created by those ways are not focused on the demands of the applications that have to use it. Since evaluations usually concentrate on the network layer, the protocols show desirable network characteristics and fit—partly, at least—to a certain application; but it remains unclear how good the application performs with the protocols in terms of what is the *best* performance of this application. There are just no application layer metrics available to measure it. Even if a protocol is designed for a specific application, the behavior of that application is usually not specified in such detail as necessary to talk about fitness. In the remainder of this section, examples of related work belonging to the bottom-up methodology are discussed and the differences to a top-down approach are highlighted.

The starting point of a paper by ElBatt et al. [EGH$^+$06] is periodic broadcasting, i.e., sending a message to all cars capable to receive it, with DSRC, a standard for wireless networks that is explained in Section 2.3.2. The application that is enabled by the broadcasts, and that is used as motivation, is a collision warning system for driver awareness. The paper introduces several metrics with application-centric properties regarding a preferably steady reception interval, e.g., a latency metric that evaluates the time between reception of packets on the application layer. In contrast to latency measurements on the network layer, this captures the impact of packet losses on an application requiring periodic updates. Then, the authors propose to tweak the protocol's transmission range and its update rate to enhance network layer results.

Nzouonta et al. propose the protocol family RBVT (Road-Based using Vehicular Traffic information routing) which takes the road topology into account for routing purposes [NRWB09]. Simulations show that the new protocol performs better than Mobile Ad-hoc Network (MANET) protocols like AODV, OLSR, and GPSR[1] in terms of average delay and average delivery ratio. Although a special characteristic of VANETs is considered for the development of a better routing protocol, the authors do not regard how good the ideas fit to certain applications.

That routing protocols designed for MANETs perform fair in VANETs is also shown by Jaap, Bechler, and Wolf [JBW05]. The authors developed a freeway mobility model and evaluated the performance of MANET protocols in typical freeway traffic scenarios on the basis of network simulations. The metrics employed for the evaluation are network-centric and comprise, e.g., the routing overhead, the TCP[2] throughput, and the delivery ratio. They conclude that AODV performs best in most of the simulated traffic situations, followed by FSR and DSR, while TORA should not be applied for VANETs.

A further paper focused on examining the lower layers has been written by Yin et al. [YEY+04]. It evaluates whether DSRC is appropriate for broadcasting in VANETs. The authors look at the physical layer of DSRC in a simulation study in order to judge the link bit error rate performance. They develop a simulation testbed for a DSRC-based VANET where a broadcasting of packets is implemented as the authors expect a collision avoidance application would do.

Back to exploiting vehicle mobility for data dissemination, Wu et al., the authors of the protocol MDDV, explain the concept of cars having perfect knowledge and discuss in short, similar to a top-down approach, how the algorithm is influenced by this [WFGH04]. However, the perfect knowledge is not about the application but about the routing state. The evaluation concentrates on network properties like the message overhead and the delivery ratio.

The MUltihop Routing protocol for Urban vehicular ad hoc networks, or in short MURU, by Mo et al. exploits mobility information of vehicles in urban areas to avoid path breakages [MZMP06]. It is designed to be robust, while having a small overhead. The paper includes a theoretical analysis and describes extensive simulations to com-

[1] For an overview about position-based MANET routing protocols see the article of Mauve, Widmer, and Hartenstein [MWH01].

[2] TCP: Transmission Control Protocol [Pos81]

pare it to MANET protocols, but all metrics for evaluation are related to properties of the network and are not targeted at applications.

The broadcasting protocol ackPBSM (acknowledged Parameterless Broadcast in Static to highly Mobile ad hoc network protocols; proposed in a paper by Ros, Ruiz, and Stojmenovic) adapts the sending of messages to reduce redundancy in dense multi-hop settings [RRS09]. The protocol's suitability to vehicle topologies with movements along roads is evaluated with simulations and compared to other routing protocols. It is found that ackPBSM performs better than the competitors regarding the two routing metrics "number of retransmissions" and "percentage of receivers".

A simulation-based study is made by Cabrera, Ros, and Ruiz [CRR09]. The authors identified common issues in five VANET routing protocols which are considered as representatives of different routing approaches. The issues are related to wireless routing both in variable topologies and in specific topologies that are characteristic in car networks. The issues, however, refer to network-layer objectives only.

Little and Agarwal propose a multi-hop data dissemination scheme for VANETs which takes directional movement of the vehicles into account [LA05]. The authors discuss analytical performance bounds of the dissemination and do not limit the analysis to a certain technological constraint. But likewise, they do not measure the advantage a certain application might have by the enhanced dissemination.

These are only a few works in the area of protocol design for C2CC. Further papers that can be classified as bottom-up are discussed in the following sections of this chapter, where distinct aspects of the publications are highlighted.

2.2 Applications

This section discusses common types of C2CC applications and is organized into the areas platooning, autonomous driving, and cooperative driving. The boundaries of the first two application types are not sharp: a platooning application can be implemented to work in a fully autonomous manner. Cars driving in a platoon, however, is such a basic and intensively studied setting that it is discussed within an own subsection. Common to all application areas regarding the top-down approach is that the car behavior has to be modeled for understanding information requirements. This a highly complex task and a starting point has to be found before we are able to work out models of large road topologies. Therefore, applications need to be identified that are basic in that they are of low complexity regarding road topology and car interactions. Having chosen an application to model, the objectives that individual cars follow have to be determined.

Lübke made a survey on car-to-car and car-to-infrastructure communication (C2CC and C2IC, respectively; together referred to as C2XC) regarding standardization activities and differences between Europe and the US [Lüb08]. The paper defines the general application areas for C2XC to be safety, mobility, traffic efficiency, and entertainment; an emphasis is put on the importance of safety and traffic efficiency. In the survey's outlook, it is described that precise estimations about the potential of C2XC regarding the increase of traffic safety and efficiency were still lacking.

The Crash Avoidance Metrics Partnership (CAMP), a consortium of car manufacturers and the U.S. Department of Transportation, established the Vehicle Safety Communications (VSC) project to estimate the benefits of safety applications using C2CC. The partnership published a report on the task of identifying safety applications [Con05]. In it, more than 30 safety applications are described, of which eight are considered most relevant and thus get analyzed in detail. Also a rough estimation of communication requirements like information update frequency needs is given. From the described applications, most related to this thesis are the collision mitigation and avoidance systems for the forward driving area, for lane changing, and for turnings.

A goal common to all applications is successful market introduction. If a protocol for a C2CC application is to be brought as product on the markets, the initial market distribution is always the same: at first only very few cars use the protocol. This requires a protocol to be able to work with very few participating cars, too. Matheus et al. investigate the specific market properties of the C2CC technology and outline a strategy for market introduction [MMP+05]. The paper describes one substantial problem C2CC suffers from: the larger the distribution of the technology is, the larger is its value for every user. For that reason, early adopters of C2CC enabled vehicles cannot make full profit from the technology. The paper discusses required penetration rates for specific applications as well as possible solutions to get the required penetration. One demand of the authors is that every car should be equipped with a very basic set of communication hardware able to forward and route messages. That way, a high penetration rate can be achieved with minimum costs.

2.2.1 Platoon Driving

The consideration of how to drive "good" in a platoon is an important step in a top-down approach as this is a very basic element of traffic. Understanding and modeling the characteristics of platoon driving helps to apply the top-down approach to more complex applications. Cars driving one after another in a platoon can make use of an Adaptive Cruise Control (ACC) application to maintain for each vehicle a safe and comfortable distance to the car ahead [Nil95]. ACC is an advanced variant of a cruise control system which, in its basic form, manages the throttle of a car to keep a steady speed. An ACC system does not work with communication but with on-board sensors; a system that also integrates communication is termed Cooperative Adaptive Cruise Control (CACC). The research report by Shladover et al. describes that CACC enables a closer following of vehicles compared to a usual ACC system [SNC+09].

Vehicle mobility models for platooning are a central research field in the area of vehicular traffic sciences and many driver models have been proposed in the past that are targeted at simulating a human driver's behavior. It is not intended to give an extensive overview here, instead it is concentrated on a few representatives that are of most importance to this thesis. For a more detailed investigation, see, for example, Brackstone and McDonald who wrote a brief overview of the most important classes of car-following models [BM99]. A more recent description of car-following models employed in traffic simulators has been published by Olstam and Tapani [OT04].

Pipes proposed one of the first car-following models in the middle of the 20th century [Pip53]. It belongs to the class of non-linear models and employs the rule of thumb for safety distances: for every ten kilometer per hour driving speed, several meters are added to the safety distance. Non-linear models produce unrealistically high accelerations in certain traffic situations which renders them unsuitable for driver assistance systems. Treiber, Hennecke, and Helbing introduced the Intelligent Driver Model (IDM) which does not has this property [THH00, TH02]. Two directions for microscopic simulations were identified by Treiber, the complex approaches like the Wiedemann model [Wie74] and the simplified ones like the Optimal Velocity Model (OVM) which is described later on. IDM is intended to be both simple and yet realistic in terms of modeling characteristics. It is continuous and deterministic, which recommends it for ACC algorithms, and it models comfortable driving, including a continuous dynamic growth of accelerations. A safety distance is used that differs from the concept of being a physical boundary: if the distance between cars is shorter than the requested safety distance of IDM, no braking is performed by the rear vehicle as long as the front car accelerates.

Car-following models work with very different levels of granularity, differentiated into microscopic and macroscopic models. The former describes traffic by individual cars, while the latter considers traffic on lanes like water flowing through pipes. Obviously, the higher the granularity, the more precise is the model as well as the complexity and thereby the runtime of simulations. VANET research mostly uses models that run fast and that generate traffic "near" to real-world behavior. The simplifications cause models to have acceleration calculations with unrealistic results in special cases as noted earlier. An example for this is the OVM introduced by Bando et al. which is commonly used in traffic simulations [BHN+95]. It models the desired speed of a vehicle, a concept already used in the Pipes model, to depend on the spacing between the vehicles. To fix the acceleration issue, various research groups proposed enhancements like the Full Velocity Difference Model (FVDM; by Jiang, Wu, and Zhu [JWZ01]) and the Asymmetric Full Velocity Difference model (AFVD; by Gong, Liu, and Wang [GLW08]). The latter is a more general model with the OVM and FVDM being special cases of it.

While the models considered so far are continuous, there is also the type of cell-based models. The cell-based models abstract space into small areas, the cells, each of the size of a car. A commonly-known representative of the cell-based models is the Nagel-Schreckenberg following model [NS92]. It is space- and time-discrete and describes traffic flow with cellular automata. The discretization allows for comparably fast calculations and hence large-scale simulations. The project Autobahn.NRW by the land North Rhine-Westphalia in Germany uses this model to forecast traffic densities on German highways [CHPS04]. The traffic model defined in this thesis is highly abstracted, too, but continuous since the intention is to understand car information demands. These demands are not discrete and may be wrongly understood if interpreted as discrete by fitting them to cells of predefined sizes.

In the context of this thesis, interest is put on how cars *should* behave. Therefore, movement rules are created based on the physical constraints. The model obtained this way is provably optimal according to the defined objectives and is not aimed at mimicking realistic traffic behavior in the first place.

Car-following and communication

The choice of a car-following model influences the performance of communication schemes [SJ04]. This implies that choosing a certain car-following model can make a protocol perform better or worse compared to other protocols. If an application is used with a protocol it is not designed for, the protocol might not work as intended in every situation and the application can be expected to suffer from this, too. In a top-down approach, the protocol development starts at the application layer with defining a car's behavior; the protocol is, ideally, deduced from the behavior. Of course, the real-world behavior of drivers is different from the abstract behavior model defined in this thesis. Accordingly, the proposed communication schemes are also applied to settings with more realistic car-following models for evaluating the performance.

2.2.2 Autonomous Driving

Controlling car movement is an issue in the research area of automated driving, too. Applications like driver support systems for collision avoidance are considered as well as fully automated cars. Automation does not have to comprise cooperation by communication; indeed, one concept explored in detail in this area is autonomous driving, e.g., Adaptive Cruise Control systems are also referred to as Autonomous Cruise Control systems. There, a car decides on its behavior completely isolated based on information from on-board sensors only. This is the opposite to cooperativeness.

Pierowicz et al. describe crash scenarios and report on a system to avoid intersection collisions autonomously with technology like radar and DGPS (Differential Global Positioning System) which improves positioning accuracy over GPS [PPB+00]. A rear-end collision avoidance system is presented by Seiler, Song, and Hedrick [SSH98]. A key aspect of the paper is the discussion of two warning algorithms by Mazda and Honda. These are used to estimate the behavior of another car and thereby to decide when actions are required like informing the driver. Both algorithms base on kinematic equations and the Mazda algorithm uses a worst-case estimation. Ioannou and Chien propose an Autonomous Intelligent Cruise Control (AICC) to improve traffic flow and safety [IC93]. It is observed that automation of driving allows to eliminate human reaction times and, thereby, following distances may be reduced considerably. The modeling does not consider velocity differences of cars because tight vehicle following assured that velocities are nearly equal. This is in contrast to the model created in this thesis, in which similar speeds are only a special case.

The systems described so far rely on information created and made available by the own car. This already allowed researchers to design cars able to drive on their own with a computer making autonomous decisions. Although it appears obvious that a car controlled by a computer reacts faster and more precise than a human ever could, an autonomously deciding computer is still on its own. Information about events that are external of the own car is available only through on-board sensors. This limits the situation knowledge of the steering one—no matter if being a human or a computer—to that what these sensors detect. During the last two decades, several projects worked on the vision of driverless cars using inter-vehicle communication, from which two will be discussed briefly. A research report by Dabbous and Huitema is about the car-to-car part of PROMETHEUS (PROgraMme for a European Traffic of Highest Efficiency and Unprecedented Safety), a EUREKA project pointed at research on fully automated driving [DH88]. EUREKA is an European initiative that coordinates and funds research in different areas. Several different application scenarios for communication are discussed in this report like merging, overtaking, crossing, and platooning (called convoy driving). Within PROCOM, a sub-project of PROMETHEUS, the communication architecture is discussed by Heister et al. [HHR+91]. Different network layer protocols are considered and it is suggested to use the Internet Protocol to create a general-purpose network. Another groundbreaking project was PATH (Partners for Advanced Transportation TecHnology). Varaiya proposes a control system architecture for it [Var93], while Streisand and Walrand focus on the architecture of the communication system with regard to available communication technologies [SW92]. Radio communication was found to have too few bandwidth; technologies like 802.11 were not present at that time. Applications that profit from communication like cruise control in a platoon are discussed including necessary control primitives for initiating and monitoring application processes. Both works do not describe the development process for protocols as done in this thesis.

Modeling behavior with control theory

The scenarios looked at in the following chapters are rather simple compared to the highly complex dynamics of vehicles in traffic scenarios with advanced topologies. For understanding more sophisticated scenarios, tools from the area of control theory are expected to be of good use. A certain sub-area considered to be of interest is optimal control theory which can help to determine vehicle trajectories that optimize given objective functions. Kirk wrote an overview about optimal control theory [Kir04] and Fernández-Cara and Zuazua give an introduction to its mathematical beginnings [FCZ03].

Control theory enables to define application objectives and to model communication effecting the behavior of cars. Reinl and von Stryk discuss a methodology for modeling and optimizing a multi-vehicle system with connectivity constraints [RvS07]. Their paper considers cooperative control of vehicles under the objective functions of minimizing the energy consumption and the final time for monitoring an area. The communication network model is used to decide on connectivity between nodes; uninterrupted communication is necessary for the mission. The model is transformed to a linear programming problem for faster solving. However, deducing important knowledge for protocol design is not the aim of that paper.

In the PhD thesis of Fax, a cooperative control of vehicle formations is discussed under the influence of a communication network [Fax02]. He models information exchange via a connected graph representing the network topology. With that, the interplay between the network topology and the decentralized control of vehicles is explained by analyzing how the connection graph changes movement of controlled vehicles, which in turn changes the graph. The thesis thereby merges ideas from graph theory and control theory; it is not a matter of discussion how to design a communication protocol from this. The control theory is used to discuss string stability, which is about spacing error propagation and error amplification upstream a platoon in car-following scenarios. In the PATH project, the string stability of inter-vehicular spacing strategies has been analyzed for automated highway platoons, e.g., by Swaroop and Hedrick [SH99]. The cited paper focuses on constant spacing strategies and distinguishes them from variable spacing and hybrid spacing strategies. Constant spacing strategies do not regard a controlled vehicle's speed for specifying the following distance, while the variable strategies do. The achievable traffic capacity of constant spacing strategies is higher in comparison to variable spacing strategies. It is found that information about a reference vehicle is important for string stability, which in turn requires the ability to communicate between vehicles. The paper's higher aim is controller design for automatic vehicle following. In an earlier work within PATH by Darbha and Rajagopal, the effect of vehicle following control laws on macroscopic traffic flow is described regarding stability[DR98]. The authors also examined how the propagation of errors in spacing is influenced by the flow of information in a later work [DR05]. The papers discuss several following behaviors for developing controllers for automated vehicles; a top-down approach requires to model the behavior of vehicles, too, but without limiting the research to automated cars and with a focus on essential information to enable the intended behavior. The lower layers of inter-vehicle communication networks specified for PATH are considered, e.g., by Hedrick, Chen, and Mahal [HCM01] and by Hsu and Walrand [HW93]. Parameters like delays and sending frequencies, however, appear to be chosen rather arbitrarily instead of analytically.

Puri and Varaiya consider an Automated Vehicle/Highway System (AVHS) for platoons [PV05]. Three maneuvers are discussed: the merging of platoons, the splitting of platoons and the change of lanes of a single vehicle. The paper defines safety and then describes an approach for proving that a system is safe. Safety is defined by looking at the velocity difference of two vehicles and by determining whether the difference is larger or equal to a design parameter. Although this parameter does not fit to a top-down approach in that it is not reasoned analytically, the idea to prove safety of the system is closely related.

Motion safety for autonomous mobile vehicles is discussed by Macek et al. [MGFS08]. The authors present a hierarchical architecture that is composed of a mission manager, a route planner, a partial motion planner, and low level and hardware components executing the orders of the partial motion planning. The partial motion planning is concerned with collision avoidance while following the route segments given by the route planner by meeting three requirements: feasibility of dynamic constraints, goal convergence with the planned route, and safety. These parts resemble the goals of a top-down approach. Two safety levels are described, passive safety that guarantees the possibility to stop before a collision and passive *friendly* safety that ensures the availability of enough braking time also for other vehicles that cross a vehicle's way for a given maneuver. The hierarchical structure appears an appropriate tool for applying the top-down approach to more complex scenarios.

2.2.3 Cooperative Applications

C2CC enables a special kind of applications termed cooperative applications. With only having on-board sensors but no possibility to contact other cars, the knowledge of neighboring cars cannot be used and their driving intentions cannot be predicted but guessed at best. For that reason, cooperation of non-communicating cars is possible in an indirect manner only: there is no way to ensure that other cars will act as assumed. The exchange of information by communication allows a car to directly talk to another about what it knows and what it is going to do. Examples of cooperative applications are CACC systems, forward collision warnings, and merging assistants; the deliverable of the COMeSafety project's architecture describes many more for various aspects of traffic [BBE+09].

Knorr et al. explain how the severity of traffic jams can be reduced by using cars as sensors for measuring traffic properties [KBSM12]. The measurements are disseminated to upstream cars. If the measurements then indicate that a jam emerges, the upstream cars are able to adapt their following behavior. The adaption is to extend the gap to cars ahead; that way, the cars are less likely to intensify the jam which might lead to a traffic breakdown otherwise.

An example for a further kind of cooperative applications is shown by Müller, Uchanski, and Hedrick, where cooperative sensoring is employed to eliminate the necessity of roadside units [MUH03]. Their paper discusses how a set of communicating vehicles can estimate traffic conditions and road parameters by cooperative estimation without using a centralized infrastructure. The downside of the decentralized approach is that a rather high percentage of participating vehicles is required to be effective.

Protocols for lane mergings and intersections

The merging of two lanes is one of the scenarios modeled in detail in this thesis. Wang, Kulik, and Ramamohanarao discuss different merging strategies, give criteria for robust merging, and propose algorithms for communicating cars [WKR09]. Similar to the work done in this thesis, the modeling takes possible ways of cars into account for safety, the simulation is based on the IDM, and the evaluation is focused on robustness and merge throughput. Fairness, however, is not regarded. The stated model is different in that it does not combine objectives but treats them as competing. A precise relation of the goals is necessary, as the thesis at hand discusses in the following chapters, for obtaining an optimum vehicle behavior. Morla describes a slot-based road usage in which slots are maintained by cooperating vehicles and exemplifies this with a road merging scenario [Mor06]. The discussion is rather high-level and aspects like the merging order or the communication scheme that are proposed in the following chapters are not touched. It is suggested that the cooperation should fail if the inbound flows are high enough that jams will emerge. In contrast, a fair merging coordination scheme is particularly relevant in face of a traffic jam as argued later in this thesis.

Mergings are considered as a basic road element for a top-down approach before intersections should be regarded because intersections can have more complex topologies and allow for more actions like car turnings. In spite of this, there are already several works that propose to use communication for coordination at intersections. Researchers have, e.g., thought about enhancing the throughput by supporting traffic light switching with communication. Here, vehicles surrounding the intersection tell the traffic lights how to behave (Gradinescu et al. [GGD+07]). Gokulan and Srinivasan adapt traffic light signals with distributed control using fuzzy logic [GS10]. The proposed approach copes with uncertainties as well as incidents and obstructions. Avin et al. describe traffic lights that are supported by virtual traffic lights which are visible only to communicating cars and help to increase intersection capacity by letting cars pass the bottleneck with maximum speed [ABHL12]. In addition to waiting time reduction, the ecological impact of traffic is a central concern for virtual traffic light research (see, e.g., Ferreira and d'Orey [Fd12]).

Those approaches use road-side units for coordination, while Ferreira et al. propose a distributed, self-organized traffic control [FFCa+10]. At this, cars are used as detectors for road usage instead of road-side hardware. The authors' idea is to let the closest car approaching the intersection embody a virtual traffic light. In contrast to the considerations in this thesis, again, fairness is not regarded.

Miller and Huang propose to build a peer-to-peer ad-hoc network for cooperation in an intersection collision avoidance system which makes autonomous decisions with input from on-board sensors [MH02]. These sensors, like radar and lidar, suffer from limitations as requiring opponent vehicles to be in line of sight. The paper motivates an application-oriented approach, although it does not exercise it in the formal depth as done in the thesis at hand. For example, neither possible ways of cars are regarded, nor fairness of the intersection passing order.

2.3 Standards and Protocols

The term *top-down* is inherently connected to the concept of dividing networking hard- and software into layers. For example, when specifying the behavior of cars, this work is done at the topmost layers, i.e., the application. The layer concept is explained in this section, followed by an overview of C2CC standards regarding the lower layers. Thereafter, protocol developments on the medium access layer are described.

2.3.1 Networking Layers

A common way to classify network protocols is to define a protocol stack that is divided into layers. A layer defines the tasks of a protocol, e.g., communication between neighboring nodes or modulation and demodulation of the data on the physical carrier signal. For a general-purpose network like the Internet it is common sense to regard the layers separately. That way, it is easily possible to let distinct protocols of one layer communicate using the same lower layers. This happens transparently: the upper layers do not have to be aware about mechanisms on a certain lower layer.

Network architectures for the Internet as well as current C2CC standards group the layers as defined by the Open Systems Interconnection (OSI) model [ITU94]. It defines seven layers, whereat the lower are closer to the physical hardware, while the upper are closer to the application. The lower four layers are responsible for the transport of data between the systems, thus in network research, the upper three layers (Application, Presentation, and Session) are often combined and treated as "application". The five remaining layers are, from the top, Application, Transport, Network, Data link, and Physical. For each layer, protocols were created, e.g., common protocols used for surfing the Internet are HTTP, TCP, and IP on the top three layers. These work on various data link and physical layers referred to as MAC (Medium Access Control) and PHY. In Internet-related research, the two layers are grouped to the Link Layer; this thesis, however, treats the two layers separately.

The proposal of a *top-down* approach literally means to specify one layer after the other from top to bottom. This is oppositional to the current *bottom-up* way most research approaches C2CC protocol design.

2.3.2 Standardization

There is a multitude of standards that specify wireless communication between vehicles. This section gives a short overview about a few of the standards with emphasis on those having importance in the following chapters. Although it might appear necessary to do an in-depth discussion since this thesis is about designing protocols, actual lower-layer properties are only of superficial interest. Neither the concept of the top-down approach, nor the scenarios that are analyzed with it regarding information demands require a sound understanding of the many standards all over the world.

In European research on C2XC, the technology of currently most importance is IEEE 802.11p [IEE12]. It originally was an amendment to the 2007 standard for wireless local area networks, IEEE 802.11-2007, for ad-hoc peer-to-peer communication with a slightly adapted frequency usage. IEEE 802.11p defines characteristics of the PHY and MAC layers. As the standard was intended to allow general-purpose networking like in a local area network, its interfaces enable various upper layer protocols, for example, IP (the Internet Protocol). In north America, IEEE 802.11 was the basis for the development of Dedicated Short-Range Communication (DSRC) [fTA10] that applies to PHY and MAC layers. On top of the layers defined by DSRC is the WAVE (Wireless Access in Vehicular Environments) protocol stack [Ins06a, Ins06b, Ins10, Ins11]. WAVE defines the WAVE Short Message Protocol (WSMP) that works on the network layer besides IP for the specific needs of C2XC. In Europe, WAVE is also standardized within CALM (Communications, Air-interface, Long and Medium range) M5 in ISO 21215 [fSI10]. While IEEE 802.11p is used for ad-hoc communication in a local area, cell-based communication like with UMTS (Universal Mobile Telecommunications System) [rGPPG99] enables wide-area access to the wired network via base stations at the road side.

2.3.3 Medium Access Control

A survey on medium access control in wireless sensor networks is given by Kredo II and Mohapatra [KIM07]. The authors group access control schemes in four groups: priority-based, traffic-based, clustering-based, and slotted TDMA. TDMA is short for Time Division Multiple Access: multiple nodes in a network access the shared medium by time division, i.e., at distinct times. All nodes in communication range have to use the same frequencies and the same coding of messages to communicate. In this thesis, considering TDMA in more detail appears obvious because necessary information for applications have to be identified and information needs arise at certain points in time.

There are several proposals to use TDMA-based schemes for medium access and for heartbeat applications in C2CC in the literature by, e.g., Saeed et al. [SAHK10], Omar et al. [OZL11], and Stanica et al. [SCB10]. Bilstrup et al. state that VANET applications need a finite channel access delay which is not possible with the probabilistic design of the CSMA (Carrier Sense Multiple Access) concept used in 802.11p [BUSB09]. Thus an adaption of a self-organizing TDMA algorithm, they term it STDMA, for VANETs is proposed which was originally designed for autonomous ship-to-ship reporting. Another approach is followed by Settawatcharawanit et al. with V-DESYNC, an adaption of DESYNC for VANETs [SCInR12]. DESYNC, proposed by Degesys et al., is a TDMA algorithm with self-organized desynchronization in which nodes try to achieve a steady state of equidistant time intervals between sending periodic messages [DRPN07]. In that state, the nodes sent their beacons always with the same interval. DESYNC neither has a mechanism to recognize missing packets at a target node implicitly, nor does it offer a fast recovery in case of a packet loss.

2.4 Beaconing

For the information demands of safety applications, the periodic single-hop broadcasting termed *beaconing* is typically used. Beaconing has also been proposed for building up topology information to support multi-hop or geographic broadcasting. The topology information is then used as an information base for other forwarding algorithms. It turned out, however, that this is not a useful application for beaconing because of intrinsic redundancy caused by building topology maps [FMR06, RFM08]. In this thesis, when it is talked about beaconing, single-hop periodic broadcasting for safety applications is meant. This section describes the "classic" beaconing and several ways proposed in the literature to enhance distinct properties like adapting sending times or repeating beacons.

The classic beaconing algorithm sends a packet with current car data in fixed intervals, usually within 1 Hzto 10 Hz; this range is defined by ETSI, the European Telecommunications Standards Institute [SA09]. Safety applications each have own update requirements, as pointed out by Robinson et al. [RCCL06], and the necessary information refreshing rate can even depend on an application's state. Besides periodic messages, it has been proposed to send instant alarm messages in case of emergency, e.g., by Benslimane [Ben04]. The evaluation whether there is a necessity for messages at certain situations is essential to this thesis.

The information received by beaconing enables applications to coordinate the movement of vehicles, e.g., for a lane merging system, which assists the driver in complex settings (as discussed by Wang, Kulik, and Ramamohanarao [WKR07]), or an adaptive cruise control, as discussed by Biswas, Tatchikou, and Dion with a chain car collision model [BTD06]. The latter paper analyzes the dissemination of messages in a geographic area by enabling nodes further away in the back of a sender to retransmit earlier. Cooperative Collision Avoidance (CCA) in platooning scenarios is chosen as application to evaluate the effect of naive and so-called intelligent broadcasting with different sending rates. The impact of C2CC for avoiding multi-car chain accidents is shown in detail; those accidents appear in the paper's model at certain sending rates. The forwarding algorithm uses implicit acknowledgments, a concept that is discussed later in this section. Furthermore, the effects of prioritized safety messages and packet delivery error rates are evaluated.

2.4.1 Adaptive Beaconing

A fixed beaconing rate increases the probability of packet collisions in dense vehicle settings because the channel load does not scale [vEKH10]. To control the channel load, adaptive frequency transmission schemes have been proposed with a focus on, e.g., a fair sharing of bandwidth [TMSH06], position error metrics [SLS+10, HFSK09], and estimations about neighboring vehicles with Kalman filters [RS07, RSKX07]. In this section, an overview about the proposals of adaptive beaconing schemes in the literature is made.

The aim of the Kalman filter integration is to achieve enough accuracy to enable cooperative collision warning systems. The cited papers show that update rates of about 2 Hz are sufficient to enable estimation within tolerable error bounds. The proposed estimation system uses on-board sensors like wheel speed sensors, but no near-field sensors, and it communicates the own state periodically via DSRC to other vehicles. The rather low update rate results from multiple estimators for car states that trigger a beacon only if estimation differences exceed a threshold. The measured system gets GPS (Global Positioning System) updates at 5 Hz and vehicle sensor updates at 20 Hz; these values represent reasonable maximum update rates for beacons. Position error metrics are considered in a proposal by Schmidt et al. about situation-adaptive beaconing [SLS+10]. The sending of messages is considered in relation to the movement of the own and the neighboring vehicles as well as macroscopic values of the situation. Different schemes with these criteria are evaluated for position errors and required channel load. A paper by Huang et al. describes how to advance vehicle position tracking with VANETs without creating network congestion [HFSK09]. This is achieved by using a rate control algorithm based on network conditions and a tracking error. The importance of considering the worst-case behavior of neighboring vehicles, i.e., the most critical situation for an observed use case, is discussed by Yousefi and Fathy [YF08]. A safety application should be able to handle that case in a predictable manner. The authors propose two metrics for the evaluation of the performance of safety applications, one is the beaconing rate and the other the effective range. The beaconing rate is the beacon delivery rate of a single vehicle to the vehicles in its surrounding. The effective range is the communication range in which a certain delivery probability threshold and an end-to-end delay is expected to be satisfied in the worst case. A further statement of the paper is that sending beacons with a lower frequency is better in terms of receive probabilities and delays than sending smaller messages with a higher frequency.

Exploiting other application characteristics of a vehicular network besides vehicle positions has also been studied for adapting sending rates. Among the analyzed traffic parameters are, e.g., the road topology [KEÖÖ04], the age of beacons [CGG+06], and the speed of vehicles [SFUH04]. The paper by Korkmaz et al. is concerned with the road topology, it compares single-hop and multi-hop beaconing [KEÖÖ04]. To accomplish this, it describes a multi-hop protocol that divides the road in a sender's transmission range into segments and chooses the vehicle in the farthest segment as forwarder. The road topology is regarded, too, by special forwarding rules at intersections. Regarding the age of beacons, Chiasserini et al. created a geographic multi-hop broadcasting algorithm that gives the forwarding of older safety messages a lower priority in favor of newer messages [CGG+06]. Saito et al. discuss a beaconing protocol that adapts the beaconing frequency with a vehicle's velocity [SFUH04]. The motivation for that adaption are network layer objectives like bandwidth consumption and reception probabilities.

A further idea to reduce the load on the channel uses piggybacking, as proposed by Mittag et al. [MTHH09]. The authors examine the effect of packing beacons into other packets in multi-hop scenarios. They show that the possible savings of multi-hop beaconing are difficult to exploit under non-perfect channel conditions and suboptimal relaying decisions. The header-to-payload ratio is found to influence the benefit of multi-hop beacons.

Besides the adaption of rates, the reduction of redundancy helps to conserve channel load. Yang et al. [YLZV04] discuss the elimination of redundant emergency warning messages for a cooperative collision warning protocol. The protocol objectives are chosen with regard to optimizing network layer issues: low-latency delivery of warning messages and effective congestion control policies. The redundancy is reduced by a rate decreasing algorithm for warning messages that incorporates a waiting time approach and that sets vehicles in different states to determine whether to retransmit a warning message.

Another degree of freedom to adapt beaconing is to control the transmission power. Torrent-Moreno, Santi, and Hartenstein state that beaconing requires a fair algorithm for bandwidth sharing which ensures to not overload the wireless channel and to give all vehicles an equal chance to disseminate their messages [TMSH06]. For the control of the channel load, the protocol D-FPAV is proposed that adjusts the nodes' transmission powers in a decentralized manner. The developing of D-FPAV starts with designing an algorithm with the assumption of complete knowledge about other senders, which resembles a top-down approach. The dissertation of Torrent-Moreno is about C2CC protocols capable of supporting the detection of dangerous driving situations and of informing other vehicles about them [Mor07]. In it, the differences between an information-centric forwarding and a packet-centric forwarding are discussed; the main difference of these strategies is the layer that is responsible for forwarding: either the network or the application. It is suggested to create a hybrid forwarding system. The conclusion of the thesis points out that no application-specific methodologies or metrics for evaluating C2CC protocols exist. This observation emphasizes the need of understanding the application area and further motivates the work presented in the thesis at hand. Another beaconing algorithm that adapts the transmit power is proposed by Yang, Guo, and Wu [YGW08]. The authors point out that, although the algorithm achieves similar performance as other protocols, it has the advantage of being fully compatible to 802.11 standards.

A classification of beacon parameters that can be influenced is given by van Eenennaam et al. [vEWKH09]. The authors create a channel busy time model for CACC and find that the solution space for beacons is three-dimensional, spanned by the number of nodes, the beaconing rate, and the beacon size. Van Eenennaam et al. describe a TDMA approach in a further paper which they term reactive beaconing and that exploits the placement of cars by an upstream approach [vEKH10]. In reactive beaconing, a car's sending right may be interpreted as a token passed upstream a platoon. When a car receives a beacon from its predecessor it waits a short duration and then sends its own beacon. The algorithm is targeted at enhancing network layer properties, as, e.g., effective frequency reuse.

Having a fixed interval B for sending updates, a straight-forward approach to prevent the same pair of cars sending repeatedly at the same time is randomizing the sending rate uniformly in $[0.5B, 1.5B]$, while maintaining the target interval on mean. This also avoids the synchronization of beacon sending times, a problem as originally described by Floyd and Jacobson and discussed with regard to wireless networks by Karp and Kung [FJ94, KK00].

2.4.2 Repeating Beacons

The repetition of packets is a further parameter in protocol design. Usually, single-hop beacons are not repeated: missing a beacon implies that no information is received within the whole sending interval. The need for a repetition, due to a beacon not received properly, is not easy to detect in a wireless network. The shared medium of all nodes in the network is the air and the medium quality is not constant over space. A central issue is that a collision, i.e., an interference of transmissions with similar signal-to-noise ratios, at one place in space cannot be detected at another. It has been argued in literature that a simple single-hop broadcasting is insufficient for VANETs because the messages are unacknowledged and prone to packet collisions, e.g., by Schnaufer et al. [SFTE06]. The repetition of broadcasts was studied by Xu et al. [XMKS04]; the authors propose layer two algorithms that repeat packets fast compared to the timespan in which a beacon is of interest. In order to do this, they extend the MAC layer with a state machine that controls the repetition mechanism. Another approach is proposed by Ali et al. in which the suitability of network coding in VANETs is analyzed [ABM11]. There, the redundancy of single-hop beacons is increased to enlarge the receive probability of a beacon—without plainly repeating it: the scheme allows to combine multiple beacons into one with only small overhead. A receiver of a such coded beacon is able to decode up to one missed beacon if the other encoded beacons were received successfully. For unicast communication, that means a sender targets a packet at one receiver, 802.11 uses an acknowledgement mechanism to indicate if a packet reached its destination [IEE12]. Beacons, however, are broadcasts, i.e., a packet is targeted at every possible receiver. Sending acknowledgments is not appropriate in that case because the sender does not know how many receivers there are. Forwarding of broadcasts over multiple hops by a simple repeating leads to a flood of packets which is referred to as the broadcast storm problem [NTCS99]. The flood causes so much traffic on the network, that throughput breaks down. There are various algorithms to avoid flooding or to reduce its impact, as in a paper by Alshaer and Horlait, in which the authors reduce the broadcast storm by adapting the probability of a node to forward a beacon in relation to the number of vehicles in range [AH05]. Williams and Camp give an overview about broadcasting by explaining, categorizing, and comparing different broadcasting techniques for MANETs [WC02].

The concept of implicit acknowledgments is used in a multihop broadcast algorithm by Biswas, Tatchikou, and Dion [BTD06]. If a car receives a broadcast that has to travel in a certain direction along a platoon of cars, it broadcasts the message until it hears another node in the target direction rebroadcasting the same message. Overhearing a rebroadcast of the message is treated as an implicit acknowledgment. Whatever repetition mechanism is used, a resending of the very same data means to send aged information: since the beacon was created, time has passed in which a car's state changed, e.g., by moving. However, resending a beacon with current information requires to include the data gathering process in the lower layer task of repeating a beacon: a cross-layer approach is necessary.

2.4.3 Back-Pressure

The concept of back-pressure in the context of multi-hop networks refers to a sending scheme for a flow of packets from a source to a drain. A node part of the flow's multi-hop routing path is allowed to send a packet to the next hop node only if that next hop signaled a successful reception of the packet previously sent to it. Zhai and Fang [ZF06] use back-pressure for alleviating shared medium effects in TCP like contention and congestion in multi-hop ad-hoc networks. One of the mechanisms of their proposal is a hop-by-hop back-pressure to keep nodes from sending additional packets in already congested networks areas by waiting until previously sent packets have been forwarded. Another application for back-pressure in mobile ad-hoc networks is congestion control for multicast communication as proposed by Scheuermann et al. [STL+07].

Despite the various approaches to optimize beaconing, none of the works discussed in this section analyzes the impact of the available information on traffic objectives as done in this thesis.

2.5 Simulations and Experiments

A way to show the suitability of a C2CC protocol is to implement and evaluate it within a realistic environment. Real-world experiments with vehicles, however, are costly for the developer as well as for researchers that want to verify the evaluations thereafter. A common solution in car-to-car research is thus to carry out simulations, often based upon existing and generally accepted simulators for network and road traffic behavior. The simulators bring basic functionalities that ease comparative tests, e.g., the network simulator ns offers an object-oriented network layer model and includes classes for 802.11-compatible simulations [NSN]. The traffic simulator SUMO, for example, enables to set-up road networks and vehicles that are controlled by the Krauss car-following model by default [BBEK11, Kra98].

Linking a traffic and a network simulator during runtime is necessary to enable both parts to influence each other and has become common practice in the last couple of years. Schumacher, Priemer, and Slottke used a simulator linking to evaluate a C2CC lane merging application and found that the application helps to reduce the cars' travel time [SPS09]. The Traffic and Network Simulation Environment (TraNS) was the first open-source framework for VANET simulations that linked ns in version 2 and SUMO. Its architecture is described by Piorkowski et al. [PRL+08]. Lochert et al. explain the linking of the tool VISSIM for the traffic simulation part with ns for the network part [LCS+05]. In addition, Matlab is appended to the described simulation framework for simulating VANET applications. This architecture has the drawback that the connection between Matlab and ns is costly so that not each car in the traffic simulation can get an own application simulator instance. In a paper by Eichler et al., a runtime linking of ns with the traffic simulator CARISMA is described [EOSK05]. The authors create a C2CC warning application as an ad-hoc agent for ns. For the evaluations in this thesis, an object-oriented framework is used that extends ns with models for roads and cars and allows different car-following algorithms [AW10].

Lerner et al. describe the design and implementation of a simulation environment for the coupling of macroscopic and microscopic traffic models [LHK+00]. The goal for this is to allow observing consequences of microscopic effects on the macroscopic level. The other way around may macroscopic settings be of value to decisions on the microscopic scale. The coupling of the models is achieved by attaching boundaries of micro- and macroscopic simulations to a wider road network. In this thesis, only low-complex situations are regarded, but coupling may help to handle larger road networks in future work.

Experiments with real cars are expensive and difficult to set-up such that situations occur which are interesting to the evaluation. Kerper, Kiess, and Mauve explain how to coordinate VANET experiments such that specific situations can be created [KKM09]. Volvo demonstrated a platooning application for cars that follow a leading truck autonomously on a highway. The application was developed within the SARTRE (Safe Road Trains for the Environment) project, as described by Robinson, Chan, and Coelingh [RCC10]. Experimental merging scenario studies with communication systems, part of the AUTOPIA project, showed that merging coordination is possible in a real setting with cars talking to each other about when and where to merge [MOP+11, MGVP11].

Chapter 3

The Top-Down Approach

As the discussion of related work on protocol design in the previous chapter pointed out, the design approach followed so far might be termed a bottom-up approach: the network is built first and the application demands are dealt with afterwards. This is very well illustrated by the standard 802.11p [IEE12] that was designed especially for C2CC. Its physical layer specification is a building block of virtually each current vehicular ad-hoc network architecture. However, 802.11p is an amendment to 802.11, a standard for the physical and the medium access layer of a wireless network commonly used today to connect, e.g., smartphones or mobile computers with the Internet. It was not designated to meet the very special requirements of safety applications of cars; on the contrary, it was designed as a general-purpose network. To understand what a general-purpose network is, a closer look at the Internet is taken as a well-known instance for this concept. At the time the architecture of the Internet was planned, it was hardly foreseeable what applications will emerge to make use of it [Cla88]. According to that, the Internet protocols were designed to support various application demands and were even redesigned to fit upcoming interests. The core objective for C2CC technology, though, is to support traffic safety and efficiency. We can intuitively describe the applications the network should be used for and, because of this, car networks should not be treated as general-purpose networks. Now, designing C2CC protocols and applications with the intent to build on 802.11p will enhance safety and efficiency to some extent but, after all, all applications created this way are limited by the specifications of that given technology. Therefore, instead of continuing to guess how to improve the current state of the art, the central proposal of this thesis is to start a top-down approach to inter-vehicle communication: in the following, it is advocated to consider the desired application behavior first before working towards a network that supports it in the best possible way.

This chapter describes the top-down approach in more detail. It is based on the publication [BMS11] and extends this paper with proofs of the theoretic statements.

A roadmap is explained that, in brief, starts with stating objectives, continues with modeling traffic behavior, and finishes with designing protocols for vehicles on a whole network of roads. The roadmap acknowledges the fact that we are dealing with a highly application-specific kind of network and therefore starts with the application's perspective. At the start, clear and concise objectives must be specified: minimize the number of accidents and minimize the resource usage. From these objectives, it is suggested to derive how each car should ideally behave. The next task, knowing the desired behavior of the cars, is to derive the information that needs to be present in each car in order to achieve this behavior. This should then allow a designer to infer algorithms, protocols, and communication technology choices satisfying these information needs.

Developing a top-down approach to inter-vehicle communication requires the solution of numerous hard and complex problems. It involves answering questions such as: given a specific situation, how should each vehicle behave, depending on the information it has on other vehicles? Or: given that an interdependency between the information present in a vehicle and its optimal behavior has been determined, how can this knowledge be leveraged to design protocols, algorithms, and communication technology to distribute the information between vehicles? All these issues cannot be addressed (or even touched) in a single thesis. Instead, it will require the long-term effort of many researchers with heterogeneous skills and backgrounds to successfully develop a comprehensive top-down approach to inter-vehicle communication.

In this chapter, it will only be looked at an almost trivial setting in order to provide a first glimpse at how such an approach might look like. The following chapters describe more elaborate, though still rather basic settings. The illustration of the roadmap is accomplished by means of a setting of two cars driving in a single lane. That setting is of minimal complexity in two dimensions: at least two cars are necessary to enable a discussion about the exchange of information. In addition, they drive in the same lane which limits interactions to the longitudinal direction, i.e., only a single dimension has to be considered. There is no need to model more complex actions of the cars like overtaking or turning. As soon as it is known how the cars optimally behave, it is looked at the impact available information in the cars has regarding their behavior. What has been learned from this is applied to evaluate a car's behavior with a protocol sending periodic broadcasts. This approach to C2CC protocol design is novel; no prior work formalized the behavior of cars with the aim to measure existing protocols and to create new, measurably suitable ones.

The remainder of this chapter is structured as follows: starting out with a brief general roadmap for top-down research on inter-vehicle communication, the key findings of the simple and very specific two-car example obtained with the approach will be summarized. The details of how the approach is applied to this example and what lessons can be learned by this are discussed next. Finally the chapter's main points are summarized. Proofs to the lemma and the theorems of this chapter may be found in the appendix of this thesis.

3.1 The Top-Down Roadmap

As discussed earlier, this thesis serves as a first step towards a comprehensive top-down approach to car-to-car communication protocol design. This section discusses a roadmap that appears suitable for the coordination of further research in this direction.

A top-down approach to inter-vehicle communication has to start with defining the objectives of this technology. The objectives of inter-vehicle communication are very specific: avoiding accidents and minimizing resource usage, in particular travel time, fuel and road capacity. Those goals may be conflicting; to keep the example in the remainder of this chapter simple, satisfaction levels for all but one objective are fixed and the remaining one is optimized. In the following chapters, multiple objectives are considered as accident absence, road usage, and fairness. But still, more elaborate approaches are conceivable that also integrate, e.g., fuel consumption or driving comfort. There exists related work on how objectives could be related, as in [TH01], where it is proposed to relate fuel consumption with traveling time and the number of transported passengers, in [Ste08], where a unified cost model for accidents and travel delays is analyzed, and in [AV11], where the approaching to a traffic light is optimized regarding trip time and fuel economy.

Given those objectives we need to understand how road traffic *should* look like under the assumption that each vehicle can exchange arbitrary information with each other vehicle arbitrarily fast. This will give us an understanding of optimal road traffic, i.e., a benchmark for all real systems, since we do not have to pay attention to constraints of certain communication systems.

In any real system, vehicles will only have limited information about each other. Therefore, after defining the objectives and modeling the vehicle behavior, the third step of a top-down approach should be dedicated to the investigation of the interplay between information availability and vehicle behavior. This step is all about determining how close to the optimal road traffic we can get, depending on the specific information that is available in each vehicle. Further, if the vehicle is controlled by a human driver, this step needs to account for the reaction time, the limits of human perception, and the driver's intentions and interests. Driver behavior analysis, also by means of incentives as described by game theory [Cam97], consequently is important to regard besides communication technology.

Ideally, at this point the best would be to derive optimal communication patterns and protocols from the results of the previous step. Yet, it is quite likely that this step will be very similar to finding an algorithm that solves a given problem, in that it cannot directly be derived, but additionally requires human creativity and intelligence to find a solution. Nonetheless the previous steps provided the protocol developer with design guidelines so that a solution can be stated accordingly, since, in contrast to the existing situation in inter-vehicle communication, a clear problem statement is present.

Performing a top-down approach in a complex and diverse setting like the road traffic of a whole city or country in one pass is likely to result in an unmanageable level of complexity. Instead, the top-down approach should first be used on specific sections of a road network. According to this, the settings discussed in this thesis are simple yet basic elements of road networks: a single road and a merging of two lanes. These two serve as basis for more complex settings. The more experience is gained with the top-down approach, the more complex can settings become. Scaling up to large road networks is expected to be possible by combining the findings for the individual elements that make up the whole road network. For example, the lane-merging can be altered to a description of an intersection of two lanes. Both topologies have similarities as they, e.g., comprise a point at which traffic flows meet. The objectives and at least parts of the optimal traffic behavior from the merging setting can therefore be transferred to an intersection. But in addition to crossing the lanes, a car also needs to be able to perform yet undefined actions like turnings. Subsequently, as soon as an understanding of a basic intersection is available, it can be extended to crossings of multiple lanes. A further step is to eventually connect elements to a road network. With that, we can begin to apply the findings to research on the macroscopic level at which vehicles are treated as flows in the network.

Now, with that roadmap in mind, it will be delved into the details of one specific example to see how the first steps in a small scale setting work.

3.2 Results in a Nutshell

To illustrate the top-down perspective on inter-vehicle communication that is advocated here, a scenario will be used that is simplified to the point of being almost trivial. This allows to focus on the key aspects of the idea. The scenario encompasses two cars, denoted by c_1 and c_2. They drive in the same direction in a single-lane road, where they cannot overtake each other. The first vehicle c_1 is driving in front of c_2 at a constant speed. The only information that c_2 has about c_1 is what is transmitted by c_1 via inter-vehicle communication. The question of how c_2 should behave is considered such that (a) regardless of how c_1 proceeds (e. g., even if c_1 decided to suddenly brake) c_2 has a sufficient safety distance to react without crashing into c_1, but at the same time (b) the distance between the two vehicles is minimal. The latter is motivated by the desire to minimize the usage of road space and thus of road capacity.

There is clearly a tradeoff here: if fine-grained, detailed, and frequent information about c_1 is provided to c_2, then c_2 will be able to follow c_1 more closely. The resource "road" can thus be used more efficiently. However, at the same time, more network resources must then be spent to deliver that feedback. If less network resources are used and c_2 is, consequently, provided with information less often, then a greater safety distance will be necessary.

If we understand this tradeoff, we are in a position to argue about how to spend network capacity best in order to support the application. In particular, we are then able to compare how well different schemes for information exchange make use of network capacity. Moreover—and maybe even more important—the same methodical approach can be used to determine an ideal baseline: assuming that c_2 at any point in time had perfect knowledge about c_1, how efficiently could the road then be used? If a given communication scheme comes close to the overall optimum obtained with perfect knowledge and at the same time minimizes communication requirements, then we know that we have designed a good inter-vehicle communication scheme not only in relative, but also in *absolute* terms.

In this spirit, first, the assumption is made that c_2 has perfect information at any time, and it is argued what this means for the required safety distance between the vehicles. Further, it is also assumed that the future behavior of c_1 is known to be constant, i.e., it continues to drive at a fixed speed. It will be seen that c_2 will first quickly approach c_1 and will then soon follow bumper-to-bumper at the same speed. While the latter is clearly an artifact of the unrealistic "perfect knowledge" assumption, the resulting behavior of c_2 nevertheless establishes the comparison baseline: it can then be argued how much is lost if using a given specific, more realistic communication pattern.

This general approach allows us to argue about proposed car-to-car communication protocols in a new way: if a given approach can be shown to come close to such an optimum performance limit, then it is known that it is a good solution to reduce accidents and resource usage. If no known protocol comes close to the derived bounds, then more work needs to be done—to find better protocols or to better understand the fundamental limitations by deriving tighter bounds on the achievable performance.

In the latter sense it will of course not suffice to consider only straightforward and idealized cases. In general, this will not result in reasonably tight performance bounds. It therefore is continued to narrow the problem down by considering another extreme case in which c_2 receives information about c_1 only once at time t^0. As a result, with more and more time passing, c_2 needs to be increasingly "careful", because its uncertainty about c_1's position and speed steadily increases after t^0. It will be argued that, in order to reliably avoid any accident, c_2 has to assume the worst possible case, namely that c_1 brakes with maximum deceleration right after t^0. The objective of accident absence forces c_2 to stop right behind that location where it *estimates* c_1 to have stopped.

Based on the foundations established through the discussed extreme cases, a scenario can then be modeled in which c_2 receives information on c_1 through some arbitrary transmission scheme. Essentially this will result in an optimal behavior of c_2 that is defined as a sequential application of the previously discussed single-information case caused by new information that may arrive at and change the behavior of c_2 before it comes to a stop. This general case allows to link the behavior of the vehicles to the way they communicate with each other. To be more precise, a metric is obtained that relates the time consumption of a way with a given communication algorithm to the provably best solution possible—the one with complete information. The use of this knowledge will be exemplified by assessing a frequently used way of transmitting information between vehicles: periodic beaconing. It is shown that the beaconing interval is a parameter that influences the possible minimum safe distance between two cars.

In summary, for the trivial setting discussed here, the key step is taken towards a top-down approach: it can be evaluated what a given transmission scheme can accomplish in relative as well as in absolute terms. It is exemplified how to turn from a blind search when looking for an information transmission scheme towards addressing a well-specified problem.

3.3 Model of a Two-Cars Scenario

The model that is created in this section is highly abstracted up to the point at which a concentration on the interaction of cars on a road is possible. The reduced complexity allows for explaining the procedure of deducing the essential information that cars require for an optimal behavior and that can thereafter be employed for evaluating communication protocols. Three main abstractions with regard to communication have been included in the model to develop a first understanding of the meaning of an optimal driving behavior: cars have complete knowledge about each other car, decision making is global, and decisions are perfectly applied without any delays or imprecisions. The motion of a car follows the physical model of Newton's laws of motion.

A setting is investigated consisting of a straight, single lane starting at position 0 and extending infinitely. In order to analyze the car behavior for optimality in a formal manner, a precise definition of the model's elements is necessary. Most of the definitions in this section will be employed throughout the thesis. The first thing to introduce is time.

Definition 1. Time. *Time is a value* $t \in \mathbb{R}$.

This trivial definition will be used consistently in all chapters of this thesis. A further, rather primitive description is used for a lane. A lane is considered to be one-dimensional, so it only defines the length of a lane.

Definition 2. Lane length. *The lane length is a value* $l \in \mathbb{R}^+$.

The progress of a car in a lane is a two times differentiable and monotonically increasing function in time. The former is motivated by the laws of kinetics, the latter ensures that cars not reverse direction. The progress of a car over time is termed a "way".

Definition 3. Way. *A way is a function* $w : \mathbb{R} \to \mathbb{R}_0^+$ *that is two times differentiable and monotonically increasing. At time* t*, the first-order derivate* $w'(t)$ *of a way is said to be the speed of the way at* t*. The second-order derivate* $w''(t)$ *is said to be the acceleration of the way at* t*.*

A car is described by four properties: its acceleration and deceleration bounds, the time it appears at the beginning of a lane, and its initial speed.

Definition 4. Car. $i \in \mathbb{N}$*. A car* c_i *is a four-tuple* $(A_i, D_i, t_i^0, w'(t_i^0))$*:*

- $A_i \geq 0$ *is the maximum acceleration,*
- $D_i \leq 0$ *is the minimum acceleration,*
- t_i^0 *is the point in time at which the car appears,*
- $w'(t_i^0) \geq 0$ *is the initial speed at* t_i^0*.*

A car is punctiform, it has no width or length. Let $i, j \in \mathbb{N}, i \neq j$*.* C *is a set of cars with a total order on the elements of the set, so that* $\forall\, c_i, c_j \in C : t_i^0 < t_j^0 \Rightarrow c_j <_c c_i$*.*

Two cars c_1 and c_2 enter the lane. They cannot overtake each other. The cars arrive at position 0 at times t_1^0 and t_2^0, respectively, with given initial velocities. Without loss of generality $t_1^0 < t_2^0$, i. e., c_1 is the earlier car, driving in front. Although cars are defined as the interacting objects here, the model can be applied to other types of vehicles as well if they can be characterized by the model's movement constraints. The maximum acceleration and deceleration capabilities of all cars are assumed to be identical and independent from their current speed, and are denoted by A and D, respectively. To keep things simple and the number of parameters reasonable, it will further be assumed that $D = -A$.

Definition 5. Acceleration bounds. *Given a set of cars C.*

$$\forall\, c_1, c_2 \in C : A_1 = A_2 = A, D_1 = D_2 = D, \textit{ and } A = -D.$$

In the setting considered here, it is assumed that the way for c_1 is always the same: it travels with constant speed after entering the lane. Its way can therefore be described as $w_1 = w_1'(t_1^0) \cdot (t - t_1^0)$, where $w_1'(t_1^0)$ is the initial—and accordingly constant—speed of c_1.

Interest is then put on assessing how c_2 should behave under varying conditions depending on the initial distance between the cars, their initial speeds, and with different assumptions about what c_2 knows about the car driving ahead of it. It will be started with a situation where c_2 has perfect knowledge about c_1 at any time, and the situation will be gradually developed towards a setting where c_2 does only learn about c_1 at discrete points in time—resembling a situation where beacons are received from c_1. This will, in the end, allow us to understand the impact of the beacon rate on the "safe" distance between the vehicles, and thus gives an idea of the interdependency between communication medium usage and road capacity usage.

In all these cases, the possible ways of car c_2 need to be considered. A two times differentiable and monotonically increasing function w_2 is a *valid way* for car c_2 if, speaking intuitively, it fits the parameters of that vehicle: the point in time when the lane is entered and the speed of the car at that point in time apply, and from this time on the way does not violate the maximum acceleration and deceleration capabilities given by $A \geq 0$ and $D = -A$. The following definition puts this more formally.

Definition 6. Validity of a way. *Let $i \in \mathbb{N}$. A way w is said to be valid for a car c_i iff*

$$w(t_i^0) = 0 \wedge \forall\, t \geq t_i^0 : D \leq w''(t) \leq A.$$

$\mathbb{W}_i$ is the set of all valid ways for a given car c_i.

3.3.1 The Objective

A valid way w_2 for car c_2 is said to be *accident absent* with a valid way w_1 for c_1 if the order of the vehicles in the lane does never change. That is, if $\forall\, t : w_1(t) \geq w_2(t)$. The case $w_1(t) = w_2(t)$ is not considered as an accident, since at this high level of abstraction cars do not have a length. This simplifies the formal reasoning.

Definition 7. Accident absence. *Given two cars c_1, c_2 with $c_2 <_c c_1$. Two valid ways w_1, w_2 for c_1, c_2 are said to be accident absent iff $\forall\, t \geq t_2^0 : w_1(t) \geq w_2(t)$.*

$i \in \mathbb{N}$. $\hat{\mathrm{W}}_i$ is the set of all accident-absent ways for a given car $c_i \in C$.

Only settings where at least one accident-absent way for c_2 exists are considered, i.e., where the initial condition—given by t_1^0, t_2^0, $w_1'(t_1^0)$, and $w_2'(t_2^0)$—does not inevitably lead to an accident. It is shown in Lemma 9 on Page 124 (in the proof to Theorem 1) that this holds if and only if c_2 does not change order with c_1 if c_2 starts braking with maximum deceleration immediately when entering the lane.

Accident absence is a constraint that has to hold at any time. As each setting is started in a safe state, this can be obtained easily by ordering c_2 to stop immediately and then to not move at all anymore. Although this is a valid way, a fundamental aim of traffic is not met by that behavior: transportation. For this reason, modeling behavior based solely on safety criteria is not enough; other goals have to make sure that movement takes place. This is ensured by modeling the road usage along the lane. The road usage is minimized by traveling the fastest, i.e., reaching every point in the lane in the shortest time, of course restricted by the constraint of accident-absent ways.

It is only possible to influence car c_2. It travels the fastest to every point in the lane if it always has the shortest distance possible to c_1. Our objective is thus to minimize the distance between both cars while guaranteeing that the ways are accident absent. Informally, given a certain algorithm to exchange information between cars, it is seeked to understand how closely c_2 can follow c_1 in the long run without risking an accident. In order to answer this question the possible ways for c_2 need to be identified which, given a certain level of knowledge about c_1, are guaranteed to remain accident absent. Among all these ways, those are interesting that are most efficient in terms of road space usage, i.e., traveling the fastest to every point in the lane. The road usage is measured by the consumed time to a point in the lane.

Definition 8. Time consumption of a way. *Let $i \in \mathbb{N}$. Given a car c_i, a way $w \in \mathbb{W}_i$, and $l \in \mathbb{R}_0^+$. Then $\tilde{t}(w,l)$ is the earliest point in time at which the way w reaches l:*

$$\tilde{t}(w,l) = \begin{cases} \arg\min_t (w(t) = l) & \text{if } \exists t \in \mathbb{R} > t_i^0 : w(t) = l \\ \infty & \text{else.} \end{cases}$$

The time consumption $k : (\mathbb{W}_i, \mathbb{R}_0^+) \to \mathbb{R}$ is $k(w,l) = \tilde{t}(w,l) - t_i^0$.

An accident-absent way is called *traveling-optimal*, or *optimal* in short, if it allows c_2 not to fall behind any other accident-absent way at any point in time. Now with the help of the time consumption of a way traveling optimality is defined.

Definition 9. Traveling optimality. *Let $i \in \mathbb{N}$. Given a car c_i and a way $w \in \hat{\mathbb{W}}_i$. The way w is said to be traveling optimal iff $\forall\, \hat{w} \in \hat{\mathbb{W}}_i, l \in \mathbb{R}_0^+ : k(w,l) \leq k(\hat{w},l)$.*

It is intuitively clear that, if c_2 has more precise and more up-to-date information on c_1, it will be able to follow c_1 more closely without risking an accident. By how much c_2 has to stay behind c_1, however, is not so trivial. Therefore, the existence and shape of optimal, accident-absent ways shall be explored under varying assumptions about c_2's knowledge in the following section. Further work based on the understanding gained from this model can take additional objectives into consideration, as it is done in later chapters.

3.4 Optimal Solution

First the case is considered where c_2 is omniscient: it always knows the precise location and speed of c_1. It also knows that the speed of c_1 will not change. The optimal behavior of c_2 is then to accelerate with A, the highest possible acceleration, until a point is reached where it has to decelerate with $D = -A$, the maximum deceleration, so that it will arrive at the same point as c_1, and will drive at the same speed as c_1. It will then follow c_1 with that speed. This way for c_2 is certainly accident absent if c_2 has correct information on the current and future behavior of c_1. Similarly, it is clear that the solution is optimal—no other accident-absent way would allow c_2 to be ahead of this solution at any point in time. The position of c_1 and c_2 over time in this setting is visualized in Figure 3.1. From these considerations the following conclusion may be drawn.

Theorem 1. *If c_2 continuously knows about c_1's present and future position and speed, then there is an optimal accident-absent way for c_2 where, after an initial transition*

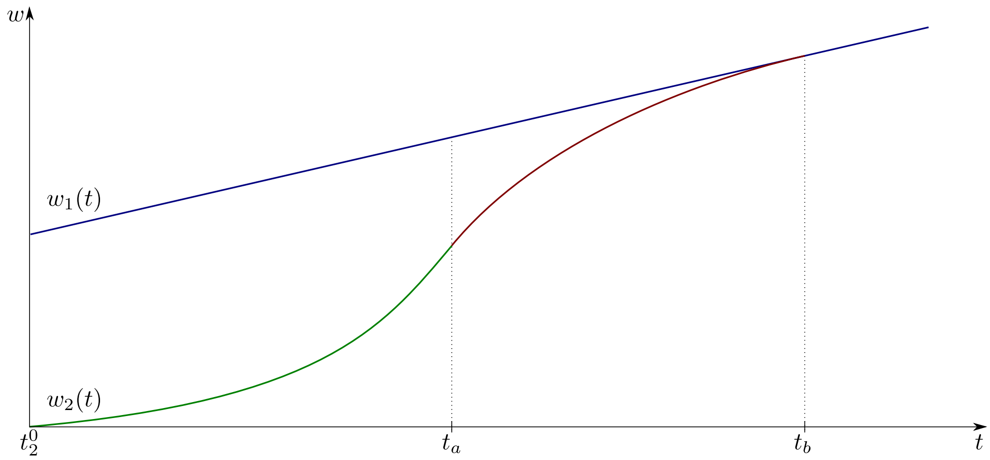

Figure 3.1: Way of the follower c_2 based on perfect knowledge about the first car c_1. t_a is the point in time when c_2 starts braking. At time t_b the steady state is reached.

period until $t_b = t_e + 2(\Delta w(t_e)/A)^{1/2}$*,* c_1 *and* c_2 *drive at the same speed and at the same point on the road.*

Here, $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/A$ *and* $\Delta w(t_e) = w_1(t_2^0) + A(t_e - t_2^0)^2/2$.

The proof to this theorem is shown in Appendix A.1. Of course, the result of the theorem is not very surprising, given the assumption of perfect knowledge about the present and future behavior of c_1. However, the theorem gives us a baseline for comparison on how much is lost if c_2 does not have that kind of perfect information. The "safety distance" at which c_2 needs to begin braking to avoid a collision is the distance c_2 needs to reduce its speed to that of c_1. This is different from the car-following models discussed in Section 2.2.1 that are intended to create a most realistic traffic behavior: the behavior described here focuses only on car braking abilities and does not regard driver reaction times.

3.4.1 One-Time Information

It is now looked at another extreme case: information about the current position and speed of c_1 is only available once when c_2 enters the scenario at t_2^0. c_2 will no longer know how c_1 behaves afterwards. From this point onwards, there will be increasing uncertainty about c_1's whereabouts on the side of c_2. As a consequence, c_2 will have to be increasingly "careful" to avoid potential accidents with c_1—it needs to take all the possible future ways of c_1 after time t_2^0 into consideration. It therefore needs to adjust to what is called in the following the worst-case behavior of c_1 within this space

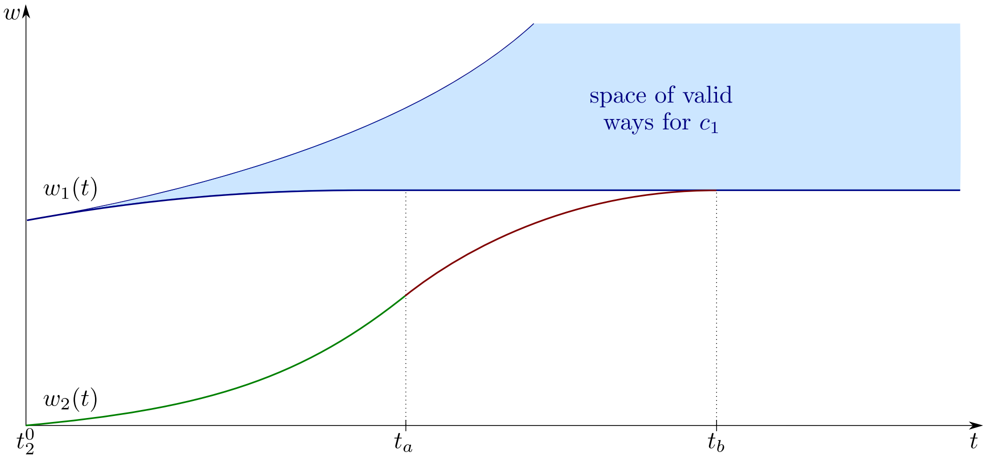

Figure 3.2: Way of the follower c_2 when information on the first car c_1 is available about one time only. t_a is the point in time when c_2 starts braking, t_b denotes when the steady state is reached.

of possible ways: braking with maximum deceleration immediately after c_2 entered the scenario and received the information on c_1.

As before, the optimal behavior of c_2 is to begin accelerating with acceleration A. Again, it switches to maximum deceleration at a certain point, which allows c_2 to come to a stop at the same location where the worst-case way for c_1, meant from the viewpoint of c_2, would have made c_1 stop. Thereafter, c_2 must retain a velocity of zero—clearly, under the given assumptions, c_2 does not know whether it is maybe standing at the same point as c_1, and therefore must not proceed any further. The initial strong acceleration phase before then switching to deceleration yields a way which is not only guaranteed to be accident absent, but is also optimal in the above defined sense: at any point in time it is further ahead in the lane than any other accident-absent way. The resulting behavior of c_2 is sketched in Figure 3.2. We, therefore, obtain the following theorem:

Theorem 2. *If c_2 learns about c_1's position and speed only once at time t_2^0, then c_2 must come to an halt at time*

$$t_b = t_2^0 - \frac{w_2'(t_2^0)}{A} + 2\sqrt{\left(1 + \frac{w_2'(t_2^0)}{A}\right)\frac{w_1(t_2^0)}{A} + 2t_e^2} \quad \text{with } t_e = \frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A}$$

in order to guarantee accident absence. The steady-state distance between both cars therefore increases unboundedly as time passes and c_1 proceeds on its way.

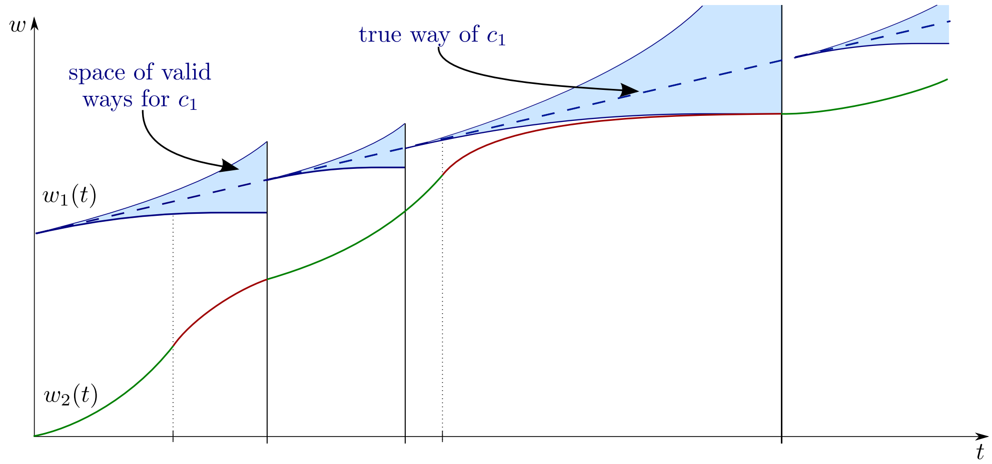

Figure 3.3: Way of follower c_2 when adapting to information updates about the first car c_1 that are received at arbitrary points in time (marked by the solid vertical lines).

The proof to this theorem is shown in Appendix A.2. Note that the transition period until time t_b differs from Theorem 1, because c_2 adapts to the estimated positions and speeds of a braking c_1. But again, this result is quite evident: c_2 cannot follow c_1 when it is "blind". Nevertheless, the specific behavior of c_2 as pointed out above provides us important hints on what happens if c_2 receives information not only once, but sporadically from time to time. So, let us consider this case next.

3.4.2 Arbitrary Transmission Schemes

Now, information about c_1 is received at c_2 through an arbitrary transmission scheme, i.e., at an arbitrary sequence of points in time. c_1 still drives with constant speed and its current state is revealed to c_2 whenever c_2 receives new information on c_1. In the time between two updates, in order to guarantee accident absence, c_2 has to assume the worst-case behavior of c_1 as described in the former subsection. After each update, c_2 has more current information on c_1 and can react accordingly. As a result, the behavior of c_2 will essentially be a sequence of maneuvers, each of them being equivalent to the one-time information case above. Figure 3.3 illustrates this.

The steady-state distance between the vehicles, if such a steady state exists at all, is determined by the respective transmission scheme. It is a metric to determine the quality of that scheme with respect to efficient road usage. In the following section it will be looked at one specific transmission scheme and this metric will be evaluated.

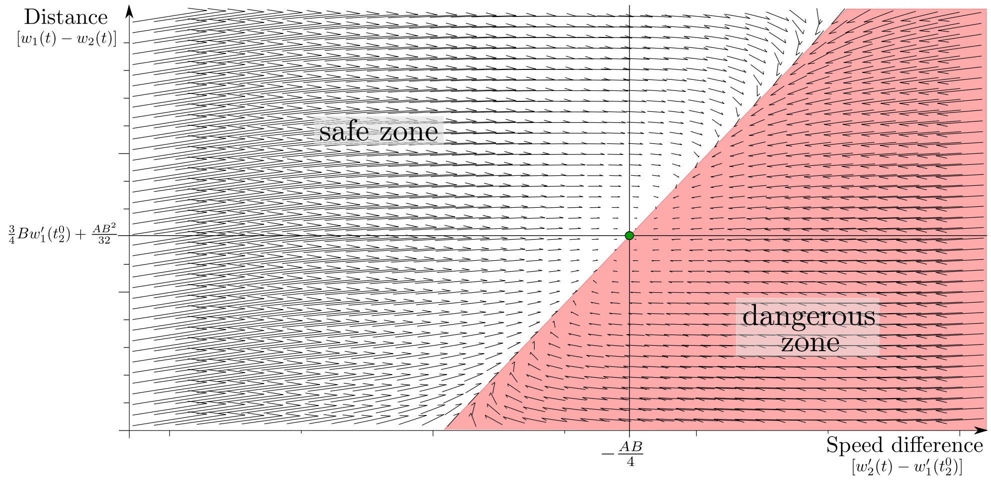

Figure 3.4: Plot of the distance between the cars in relation to the speed difference at the beginning of an interval with periodic beaconing. The steady state is at speed difference $-AB/4$ and distance $3Bw_1'(t_2^0)/4 + AB^2/32$ where B is the beacon interval length.

3.4.3 Periodic Beaconing

A scheme for information transmission is considered that has been frequently proposed for inter-vehicle communication: periodic beaconing. In this special case, c_2, ideally, receives an information update periodically, i.e., at time $t_i, i \in \mathbb{N}_0$, with $\forall\, t_i : t_{i+1} - t_i =$ const. Again, it is assumed that c_1 is driving with constant speed. As in the preceding section, this is not known a priori by c_2. After the reception of a beacon, c_2 will once again behave as described for the one-time information case until the next beacon is received.

The first important question that needs to be addressed is: does a steady state exist? In fact, the following lemma holds.

Lemma 1. *If c_2 learns about c_1's position and speed periodically at $t_i, i \in \mathbb{N}_0$, with $\forall\, t_i : t_{i+1} - t_i = const$, then the sequence of distances between both cars and the speeds of c_2 at the points in time t_i converge to a steady state for $i \to \infty$.*

The proof to this lemma is shown in Appendix A.3. Figure 3.4 helps to get an intuitive understanding why it holds. Each point in the figure stands for one pair of speed difference (x axis, $w_2'(t_i) - w_1'(t_2^0)$) and position difference (y axis, $w_1(t_i) - w_2(t_i)$) at the beginning of a beaconing interval, i.e., when fresh information is received at c_2. For each pair of position and speed differences at one interval, the arrows in the figure point to the position and speed differences at the subsequent interval, given that c_2

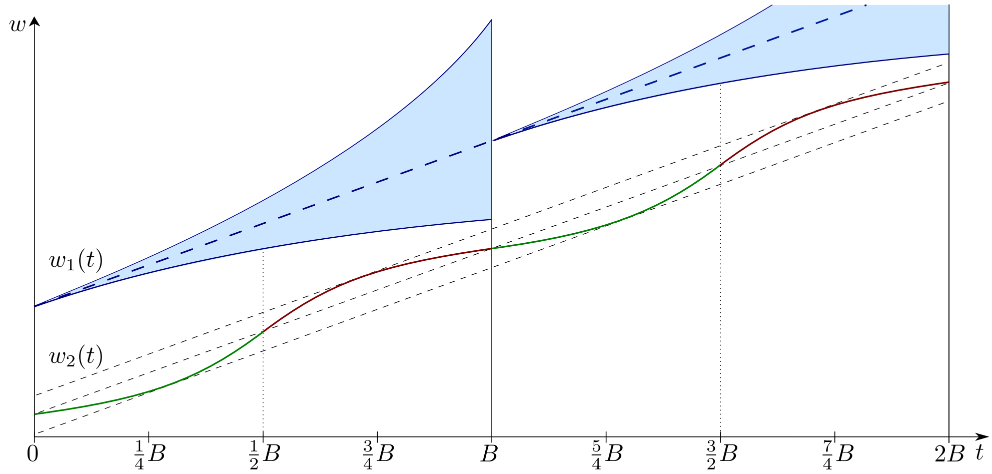

Figure 3.5: Way of the car c_2 in the steady state when adapting to periodic information updates about the first car c_1.

follows the strategy outlined in Subsection 3.4.2 about arbitrary transmission schemes. That is, the figure describes how the position and speed differences develop over time, from beacon interval to beacon interval. For points in the "dangerous zone" on the right hand side of the figure, it is not guaranteed that c_2 can prevent an accident. It can be seen that, if it is started anywhere outside this zone, then the (speed difference, distance) pairs at the beginnings of the beacon intervals will approach a steady state (marked by a circle), and they will never enter the dangerous zone.

The data in the figure originates from an actual example setting calculated with Matlab. The front car c_1 drives with a speed of $10\,\mathrm{m/s}$, and the acceleration capabilities are $A = -D = 1\,\mathrm{m/s}^2$. The time between updates is set to $B = t_{i+1} - t_i = 1\,\mathrm{s}$. The plotted data shows the central part of the calculated area; the background color and the target cross with the colored dot were drawn manually and are only meant for highlighting. Section A.3.1 explains how the arrows were calculated.

With the knowledge that the steady state exists, it can be turned towards describing how the steady state looks like. The positions of the cars in the steady state are depicted in Figure 3.5. It follows from straightforward calculations that the following theorem holds.

Theorem 3. *The steady state distance between c_1 and c_2 at the beginning of each beacon period is given by $3Bw_1'(t_2^0)/4 + AB^2/32$, where B is the beacon interval length (i. e., the time between information updates arriving at c_2).*

The proof to this theorem is shown in Appendix A.4.

3.4.4 Discussion and Further Steps

What does one gain from this simple example? After all, the setting has been simplified to the point of being almost trivial. Two important lessons can be learned from this exercise.

Lesson one: for the first time, since working on inter-vehicle communication was started, a clear reasoning about design decisions was possible. When talking about alternative beaconing schemes, it can now be thought about quantifiable advantages and disadvantages that are derived from clearly stated goals.

Lesson two: thinking in a top-down fashion is able to change the perspective on inter-vehicle communication. Now more fundamental questions are considered like: "when is transmitted information beneficial?" instead of questions such as: "how well does our new beacon scheme perform in a simulation/testbed setting?".

After understanding the influence of available information on the behavior of c_2, the next task will be to identify the information dissemination scheme that supports the knowledge requirements best. This implies to switch from the application level to the protocol layer. From the dissemination scheme, it then has to be went down further to the lower layers and, eventually, it has to be decided which kind of medium access and physical link technology is the right one for this application. This also means taking network and channel characteristics into account which were not yet regarded in this discussion. Even though these are all very important aspects of a top-down approach, they would quite obviously exceed the limitations of this small example. So, at this time, it is stopped after an outline of the first steps. The following chapters describe somewhat more complex settings building on the findings made here.

3.5 Conclusion

In this chapter the idea of top-down research in the area of inter-vehicle communication has been introduced. It has been argued that, in contrast to communication in general-purpose networks, the exchange of data between vehicles has clearly defined goals: preventing accidents and optimizing resource usage. Thus, research in this area should derive applications, protocols, and algorithms from those very specific goals, instead of building the network first and consider potential functionality that could be achieved with this network afterwards. A first glimpse is provided on how this top-down process could look like by examining an almost trivial example about the influence of available information on the behavior of a car. This is only a first step on a very challenging path that will be treaded along in the next chapters. But this chapter also showed that an extensive top-down approach comprising all levels from the application layer to the physical layer requires the effort of many research groups with very heterogeneous skills. However, meeting this challenge is extremely rewarding, too, since, in the end, finding the best possible solution might in fact save lives and protect valuable resources. This effort is likely to be very much worthwhile.

Chapter 4

Carrot—A Protocol For Platooning

The model of two cars driving in a single lane created in the previous chapter will be extended in the course of this chapter to support multiple cars. With this extension, the model is used to discuss traveling-optimal behavior in a platoon. It will be found in the following that, based on the model's objectives and knowledge exchanged via beaconing, cars after the leading one follow with a speed that oscillates with the beaconing frequency. This is similar to what has been discussed previously for two cars. The steady state distance between successive cars, however, is smaller than the distance between the first two cars. The minimum average road usage of a platoon, i.e., how much space it occupies, will be considered with regard to the cars' behavior. The optimal distribution of bandwidth among the cars for the minimum road usage is derived and proved. It will be shown that a fixed periodic beaconing with an equal distribution of bandwidth performs better than any other scheme. The findings of how to distribute bandwidth is suitable for an application that aids drivers keeping a steady distance between cars, e.g., it can be used within a Cooperative Adaptive Cruise Control (CACC) application.

As will be discussed, the objective of traveling optimality requires cars to accelerate and decelerate periodically. Towards a more realistic car following, the behavior is altered such that a constant following speed is possible. Traveling optimality is replaced for this with the goal road-usage efficiency. This lets the cars pursue a steady state with constant speed instead of one with oscillating speeds. The optimal behavior for road-usage efficiency will be discussed and, as in the traveling-optimal case, the minimum road usage in the steady state will be calculated.

While this discussion results in a behavior that is a bit more realistic, the information exchange is still far too simple with regard to a wireless network. It will therefore be turned from a scenario with lossless communication to one with unreliable communication; beacon packets may get lost. In addition, delays let the beacons age before reception. In this setting, the effect of a beacon loss on the behavior of cars driving in their steady states is analyzed. It will be found that a propagation of braking maneuvers upstream the platoon occurs. This severe consequence of a loss demonstrates that a simple periodic beaconing with fixed sending intervals does not perform well with lossy communication. Packet losses are frequent events in an ad-hoc wireless network based on IEEE 802.11p. For example, when two cars decide to send at about the same time, their packets collide: transmitting packets simultaneously causes the physical representations of the data packets to interfere which prohibits the intended receivers from decoding the packets. Mechanisms to prevent and handle collisions in C2CC like acknowledgments of successful receptions are not used at sending beacons, i.e., at broadcasting messages to multiple receivers. In addition, a fixed sending interval of periodic messages like beacons is prone to repeated collisions. Randomizing beacon sending times is a common means against this problem. But it will be shown that randomization causes unnecessarily large steady state distances, while collisions are not even avoided consequently.

Another way to avoid packet collisions will be proposed with the protocol *Carrot*. The core of this protocol is adaptive beaconing using time-multiplexing by passing a sending token upstream the platoon after the successful reception of a beacon. It exploits the application-specific property that all communicating parties are lined up in a row. Carrot's beacon sending rules base on a topology-related trigger mechanism to detect a missing beacon implicitly—without additional messages. The detection allows for a repetition of beacons very fast compared to usual sending intervals of beaconing schemes. As it is shown in the following, the fast repetition aids a car to slow down less than with beaconing. Carrot thereby mitigates the impact of a loss on upstream cars. In addition to a formal discussion, Carrot will be evaluated with simulations which show that Carrot allows for a smaller road usage at a target update interval than beaconing. Carrot works in a fully distributed manner with low communication and memory overhead and adapts dynamically to the number of cars in a platoon.

Simulations with the simulator ns demonstrate the suitability of Carrot to the application with wireless ad-hoc networking using IEEE 802.11p. They confirm that the intervals between successive beacons in lossy situations are shorter with Carrot than with fixed interval beaconing and with randomized interval beaconing, and through this, following at a shorter steady distance than with beaconing is possible.

This chapter is structured as follows: the formal model of a platoon scenario is described in Section 4.1. The optimal behavior for the model is derived in Section 4.2. The behavior for a steady state with constant speed is discussed in Section 4.3. Carrot is proposed in Section 4.4 and evaluated in Section 4.5. The chapter closes with a conclusion in Section 4.6.

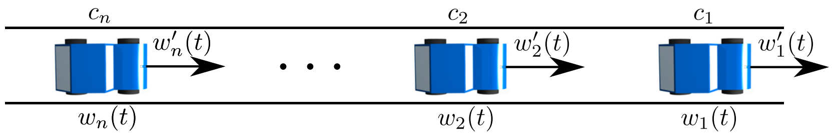

Figure 4.1: A simple sketch of n cars driving in a lane.

4.1 Model of a Platoon

A formal model of a platoon of cars is described in this section and the behavior objectives are defined. Most of the definitions made in Section 3.3 on Page 36, for the case of two cars, are applicable to this scenario, too, and will not be repeated here. The model's main characteristics are outlined in the following. The road topology consists of a single lane with $n > 1$ cars driving in it. The lane does not allow overtaking. The cars have no length and are identified by their order in the lane. The operator $<_c$ is used to describe that c_{i+1} (with $1 \leq i \leq n$) drives behind c_i: $c_{i+1} <_c c_i$. The car c_1 is the most ahead one and drives in the lane with a constant speed greater or equal zero. Figure 4.1 depicts the scenario with the leading car c_1 and $n-1$ cars following it. At t_i^0, car c_i appears at the origin of the lane; this is point 0. A car's position in the lane at a given time is defined by its way $w : \mathbb{R} \to \mathbb{R}_0^+$. A way is two times differentiable and monotonically increasing. Each car has the same maximum acceleration $A > 0$ and maximum deceleration $D = -A$.

Cars only get information about surrounding cars through communication. The communication is assumed to be lossless. Beacons contain information about a sender's state like its unique identifier, position in the lane, speed, and time of generation of this information. On receiving a beacon, a car stores the information contained in the beacon. If older information about the sender is already known at a car, this information is replaced by the new information. Information in a beacon is assumed to be precise; measuring errors are not regarded. Beacons have a fixed size and are sent in a single data frame on the data link layer. A car cannot send beacons faster than with the minimum interval between beacons. The minimum interval between sendings for all cars is B. The following definition describes how the sending time is related to the bandwidth of the communication channel.

Definition 10. Bandwidth. *The available bandwidth is* $1/B, B \in \mathbb{R}^+$*, i.e., it is normalized to beacons per second. Sending of information in a beacon consumes bandwidth. The bandwidth* $1/B_{0.1\,s}$ *to send one beacon in the interval of* $0.1\ s$ *is very small compared to the available bandwidth:* $1/B_{0.1\,s} \ll 1/B$*.*

The bandwidth has to be shared among the cars driving in a lane. For n cars, each car c_i allocates a fraction $1/B_i$ of the available bandwidth for its beacons. Time multiplexing allows a car to send beacons with the full bandwidth, although only at a portion of time. A beacon then occupies the medium for $B \in \mathbb{R}^+$ seconds. It is assumed that the bandwidth is assigned exclusively to the application in focus. Neither other, foreign applications, nor other senders require bandwidth.

4.1.1 Objectives

The two objectives *accident absence* and *traveling optimality* defined in the previous chapter will be used once more to discuss the optimal behavior of cars. The meaning of the objectives is now repeated briefly in an informal manner; refer to Section 3.3 for a more detailed description. Given two cars, accident absence of their ways means that their order never changes. It is required that accident-absent ways exist: each car can stop without an accident at least at its time of insertion in a lane. A traveling-optimal way for a car is as far or further ahead compared to any other accident-absent way for that car.

Later, in Section 4.3, traveling optimality is omitted in favor of a behavior goal that enables the cars to drive with constant speed in the steady state. This goal will be referred to as *road-usage efficiency* and is defined as follows.

Definition 11. Road-usage efficiency. *A car drives road-usage efficient if it drives on an accident-absent way that (a) reaches in minimum time the minimum distance to its predecessor that can be kept without acceleration or deceleration, and (b) maintains this minimum distance.*

The three objectives discussed so far are all related to application behavior. A communication objective that follows directly is: each car c_i receives a new beacon from its predecessor after a fix *target interval* of B_i seconds. This appears rather obvious, because beacons that arrive too late have a negative impact on the application: a car cannot keep its steady state in such a situation, as will be discussed later in this chapter. The metric used to evaluate the intervals between successive receptions of beacons sent by a car's predecessor is measuring the difference to the target interval.

4.2 Optimal Behavior for Traveling Optimality

The behavior analysis of a pair of cars with the goals accident absence and traveling optimality was discussed in Section 3.4.3. In this section, the scenario is extended first to a third car and then to a platoon of an arbitrary number of cars. The cars are each given a fixed, equal fraction of the available bandwidth for communication. Besides the behavior of the third car, the optimal beacon sending times in view of the cars' road usage are formally discussed. At the discussion of platoons with an arbitrary number of cars, the minimum road usage is considered in the situation when each car is in a steady state to its predecessor. It will be shown that non-periodic beacon schemes as well as schemes with non-equal sending intervals perform worse than periodic beaconing with equal sending times: they cause a less efficient road usage.

4.2.1 Sending Times of a Third Car

The optimal sending times of a third car are discussed towards minimizing the sum of steady state distances between the three cars. It will be assumed that the cars share the available bandwidth $1/B$ equally: c_1 and c_2 both use half of the bandwidth, that is, $1/B_1 + 1/B_2 = 1/B$.

c_1 drives with constant speed and sends beacons periodically. c_2's behavior in this situation was already discussed in the previous chapter. So only the influence of the sending times of c_2 regarding the behavior of the third car c_3 are yet unknown. The following theorem states the optimal sending times for minimizing the steady state distance between the second and the third car.

Theorem 4. *Given three cars, $c_1 <_c c_2 <_c c_3$, that drive in their respective steady states. c_1 sends updates at iB_1, $i \in \mathbb{N}_0$. The optimal beacon sending times for c_2 to minimize the steady state distance between c_2 and c_3 are $iB_1 + B_1/2$.*

The proof to this theorem is shown in Appendix A.5. The theorem states that if c_2 sends always exactly at the middle of c_1's sending interval, then the distance between c_2 and c_3 that travel in their steady states is minimized. At this time c_2 changes its acceleration from full acceleration to full deceleration. This result is informally explained in the following.

At the time c_2 changes to deceleration, it is fastest within c_1's sending interval and, therefore, the required safety distance between c_2 and c_3 is shortest. Moreover, c_2 brakes for the second half of the sending interval which equals c_3's worst-case expectation; its estimation error is zero regarding the speed and position of c_2. The estimation errors grow when c_2 accelerates again on receiving an update from c_1, until the next update is received at c_3. Right before this update is received at c_3, the estimation errors are smallest compared to any other sending time.

Besides the offset in time, the maneuvers of c_3 equal that of c_2; it also accelerates and decelerates for a half sending interval. This behavior is illustrated in Figure 4.2.

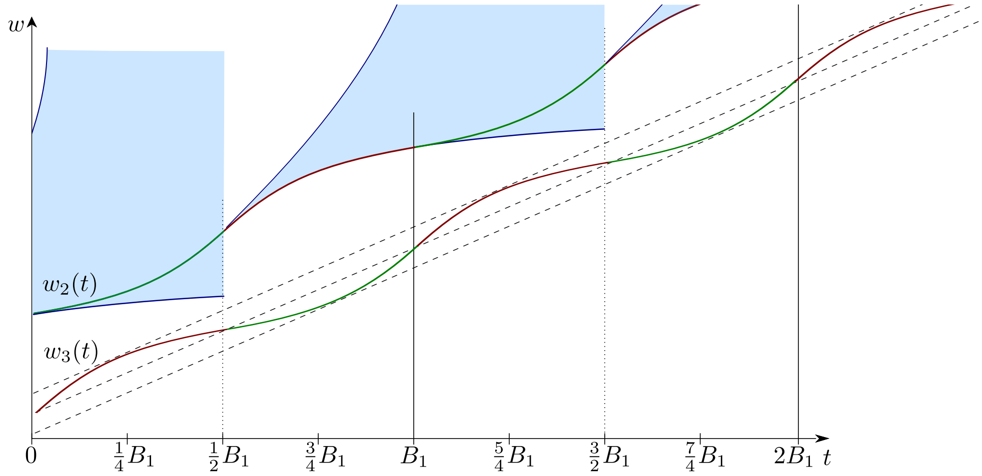

Figure 4.2: The positions of the second car c_2 and its follower c_3. Both cars are in their steady states that are enabled by receiving updates at the same, fixed rate from c_1 and c_2, respectively. The sending times have an offset of a half sending interval.

The distance $\Delta w_{2,3}(t)$ between the two cars in the steady state at c_2's sending times at $iB_1 + B_1/2$ can be expressed as follows[1]:

$$\Delta w_{2,3}\left(\frac{B_1}{2}\right) = \frac{1}{2}B_1 w'_{\text{min}} + \frac{1}{8}AB_1^2 = \frac{1}{2}B_1 w'_1\left(t_1^0\right).$$

The proof also shows that this distance, at the sending times of c_2, equals the average distance $\overline{\Delta w_{2,3}}$.

But the distance is not constant over time; it oscillates, as well as the cars' speeds. The oscillation of the cars' speeds cause an oscillation of the road usage of the platoon. c_1 drives with constant speed, and this speed is contained in each beacon sent to c_2. c_2 sends after constant intervals and so appears to c_3 as if driving with constant speed, too. The platoon's road usage depends on the speed oscillation of c_3. It oscillates around the road usage of a platoon with c_3 driving at constant speed (the same average speed as c_1 and c_2) at c_3's average distance to c_2. The road usage of the whole platoon is

[1] By using Equation A.3 on Page 151 and replacing $\Delta w_{\frac{B_1}{2}}(t)$ with $\Delta w_{2,3}(t)$

the sum of the average distance between c_1 and c_2 plus the distance between c_2 and c_3. The average average road usage of the platoon is

$$\begin{aligned}\overline{\Delta w_{1,3}} &= \overline{\Delta w_{1,2}} && +\overline{\Delta w_{2,3}} \\ &= \left(\frac{3}{4}B_1 w_1'\left(t_1^0\right) + \frac{AB_1^2}{32}\right) && +w_1'\left(t_1^0\right)\frac{B_1}{2} \\ &= \frac{5}{4}B_1 w_1'\left(t_1^0\right) + \frac{1}{32}AB_1^2.\end{aligned} \tag{4.1}$$

The platoon's road usage oscillates within

$$\left[\overline{\Delta w_{1,3}} - \frac{1}{32}AB_1^2, \overline{\Delta w_{1,3}} + \frac{1}{32}AB_1^2\right].$$

Now that the optimal sending times are determined and the observation was made that only the platoon's last car influences the oscillation of a platoon's road usage, the average road usage of a platoon of $n > 2$ cars is easily derived, so that the following theorem is obtained.

Theorem 5. *Given a platoon of $n > 2$ cars driving in their steady states. Let the sending times for all cars with uneven indexes be $(0 \mod B_1)$, and $(B_1/2 \mod B_1)$ for all cars with even indexes. The average road usage of the platoon is*

$$\overline{\Delta w_{1,n}} = \frac{1}{32}AB_1^2 + \left(\frac{3}{4} + \frac{(n-2)}{2}\right) w_1'(t_1^0)B_1.$$

Proof. The theorem is proved with complete induction. The distance with three cars has been shown above. Let the assumption hold for n cars. Their road usage is, with regard to Equation 4.1,

$$\overline{\Delta w_{1,n}} = \left(\frac{3}{4}B_1 w_1'(t_1^0) + \frac{AB_1^2}{32}\right) + (n-2)\left(w_1'(t_1^0)\frac{B_1}{2}\right).$$

$n \to n+1$: the nth car sends an update each time it begins to brake. The average distance between c_n and c_{n+1} is then equal to that of c_2 and c_3 (which was discussed earlier) and the average road usage of the platoon is:

$$\begin{aligned}\overline{\Delta w_{1,n}} + \overline{\Delta w_{n,n+1}} &= \left(\frac{3}{4}B_1 w_1'(t_1^0) + \frac{AB_1^2}{32}\right) + (n-2)\left(w_1'(t_1^0)\frac{B_1}{2}\right) + \left(w_1'(t_1^0)\frac{B_1}{2}\right) \\ &= \frac{1}{32}AB_1^2 + \left(\frac{3}{4} + \frac{(n-1)}{2}\right) w_1'(t_1^0)B_1.\end{aligned}$$

□

4.2.2 Other Update Schemes

Up to this point, a periodic sending of beacons was assumed in conjunction with an equal sharing of available bandwidth among the cars, i.e., all cars were able to send with the same update interval. It is shown in the following that these assumptions are optimal towards a minimum road usage of platoons.

Non-periodic update schemes

Assume the bandwidth is shared equally among the vehicles of a platoon of i vehicles (with $i \in \mathbb{N}$). With a non-periodic update scheme, no vehicle c_i, is able to send more often than with the minimum sending interval B_i. Sending non-periodically implies sending at least one beacon with a longer interval between two sendings than B_i. This causes the sender's follower to leave the steady state distance so that the distance between the vehicles grows. Even if the distance between other pairs of cars can be kept constant in this case, the sum of distances is larger. In addition, the non-periodicity avoids a steady state between successive cars and so it is obvious that no shorter distance between any two vehicles as with periodic communication is possible.

Non-equal bandwidth shares

So far, it was assumed that the bandwidth is shared equally among the cars. Now non-equal minimum sending intervals are discussed. Assume that the available bandwidth always equals the assigned bandwidth. Only the distribution of bandwidth to the vehicles can be changed. For three cars, the first two are able to share the available bandwidth $1/B$ like in the following. Let $1 > \varepsilon > 0$:

$$(1+\varepsilon)\frac{1}{B_1} + (1-\varepsilon)\frac{1}{B_1} = \frac{2}{B_1} = \frac{1}{B}.$$

Assume the three cars drive in their steady states. Let c_1 have a higher amount of bandwidth than c_2 so that it sends with the interval $B_1/(1+\varepsilon)$. c_2 is able to follow at a closer distance now, but has to send with less bandwidth, i.e., a higher interval $B_1/(1-\varepsilon) > 0$. So c_3 has a larger distance to c_2. The steady state distance of c_1 and c_2 is

$$\frac{3}{4}\left(\frac{B_1}{1+\varepsilon}\right) w_1'(t_1^0) + \frac{A}{32}\left(\frac{B_1}{1+\varepsilon}\right)^2.$$

Looking at the time $t \mod (B_1/(1+\varepsilon))$ between two updates of c_1, c_2 does not have to send at a fixed offset to c_1's sending time; the offset changes from one sending to the next by $B_1/(1-\varepsilon) \mod (B_1/(1+\varepsilon))$.

A steady state for c_3 is enabled if c_2 sends periodically at the same offset to c_1's sending times. Let $t_s \in [0, B_1/(1+\varepsilon))$ be the time of an update of c_2. Then a steady state exists for c_3 if either (1):

$$B_1/(1-\varepsilon) \mod (B_1/(1+\varepsilon)) = 0$$

or (2), $k \in \mathbb{N}$,

$$k \cdot B_1/(1-\varepsilon) \mod (B_1/(1+\varepsilon)) = B_1/(1-\varepsilon) \mod (B_1/(1+\varepsilon)),$$

i.e., after every k sendings, which are shifted by $B_1/(1-\varepsilon) \mod B_1/(1+\varepsilon)$ each, the update of c_2 contains the same speed again and a position that grew linearly with constant speed $w_1'(t_1^0)$ during that time.

In case (1), c_2 sends at integer multiples of c_1's sending interval. Then the optimal sending offset of c_2 is to send when c_2 begins braking. Already at the smallest multiple (i.e., 2), the distance is worse than sharing the bandwidth equally. For this case, the bandwidth share is $4/(3B_1) + 2/(3B_1) = 2/B_1$. It is now shown that the road usage is larger than in case of a bandwidth that is shared equally:

$$\begin{aligned}
&\frac{3}{4}\left(\frac{3B_1}{4}\right) w_1'(t_1^0) + \frac{A(3B_1/4)^2}{32} + \frac{1}{2}\left(\frac{3B_1}{2}\right) w_1'\left(t_1^0\right) = \frac{21}{16} B_1 w_1'\left(t_1^0\right) + \frac{9AB_1^2}{512} \\
>&\frac{3}{4} B_1 w_1'\left(t_1^0\right) + \frac{AB_1^2}{32} + \frac{1}{2} B_1 w_1'\left(t_1^0\right) = \frac{5}{4} B_1 w_1'\left(t_1^0\right) + \frac{AB_1^2}{32}.
\end{aligned}$$

This holds for $B_1 < 32w_1'(t_1^0)/(7A)$, which covers the usual application case: the steady state equations are based on the assumption $w_1'(t_1^0)/A \geq 1$, so it has to be $B_1 < 32/7$. If c_2 sends at larger multiples, the result is similar and B_1 can be even larger.

In case (2), c_3 is not able to hold a steady state distance over two subsequent sending intervals of c_2. The speeds sent by c_2 are on average the true average speed of c_2. The position sent on average is the position of c_1 minus the average distance between c_1 and c_2. Therefore, an average steady state distance between c_2 and c_3 based on c_2's updates cannot be better than the distance possible with c_2 being constantly on the average distance to c_1 and sending with $B_1/(1-\varepsilon)$. The sum of the average distances between the three cars in this case is larger than the sum with a bandwidth that is shared equally:

$$
\begin{aligned}
&\frac{3}{4}\left(\frac{B_1}{1+\varepsilon}\right) w_1'\left(t_1^0\right) + \frac{A}{32}\left(\frac{B_1}{1+\varepsilon}\right)^2 + \frac{3}{4}\left(\frac{B_1}{1-\varepsilon}\right) w_1'\left(t_1^0\right) + \frac{A}{32}\left(\frac{B_1}{1-\varepsilon}\right)^2 \\
=&\frac{6}{4} B_1 w_1'\left(t_1^0\right)\frac{1}{1-\varepsilon^2} + 2\frac{A}{32} B_1^2 \frac{1+\varepsilon^2}{1-\varepsilon^2} \\
>&\frac{5}{4} B_1 w_1'\left(t_1^0\right) + \frac{A}{32} B_1^2 \\
=&\frac{3}{4} B_1 w_1'\left(t_1^0\right) + \frac{AB_1^2}{32} + \frac{1}{2} B_1 w_1'\left(t_1^0\right).
\end{aligned}
$$

The inequality holds because $1/(1-\varepsilon^2) > 1$ and $(1+\varepsilon^2)/(1-\varepsilon^2) > 1$. So sending at non-equal intervals results in a larger average road usage of a three-cars platoon. For more than three cars, the additional bandwidth given to c_1 can be arbitrarily drawn from the other cars. This leads to a similar result, obviously, since each following car suffers the same from less bandwidth as shown with cars two and three. Another possibility is giving another car behind the first one more bandwidth while drawing that additional bandwidth from other, arbitrary cars. This is also not beneficial in sum, because the distance between the advantaged car and its follower shrinks less than the other distances grow. The average distance possible between the car with the additional bandwidth and its follower cannot be less than the optimal distance in the case that all cars had that additional bandwidth. At the same time, other distances have to grow: drawing bandwidth from an arbitrary car, the distance of that car to its successor is also at best a multiple of the optimal steady state distance with equally shared bandwidth. So the road usage of the platoon cannot be less than in the case of an equally shared bandwidth.

4.3 Optimal Behavior for Road-Usage Efficiency

Driving in a traveling-optimal manner causes an oscillation of the speeds of cars in a platoon. Of course, by choosing a way that is not traveling optimal, a car is also able to drive in a steady state with constant speed. For a cruise-control application, creating a non-oscillating behavior is clearly reasonable. Therefore, in this section, the optimal behavior regarding the objective *road-usage efficiency* is discussed. A car with this behavior reaches the steady state in an accident-absent manner in minimum time and then stays in that state.

This section is structured into the proposal of the optimal behavior regarding road-usage efficiency, followed by a discussion of this behavior. It will be found in Section 4.3.2 that B_i has to be equal at each car c_i to obtain a minimum sum of steady state distances; this is a similar result as in the previous chapter. Before turning to the algorithm description of Carrot in the next Section 4.4, the impact of a beacon loss on the steady state distance is shown and it is discussed how cars can maintain the steady state in case of losses.

A steady state is only possible with a periodic sending of beacons at a constant rate, so this is assumed in the following: a car c_i sends a beacon each B_i seconds. With this assumption, the behavior strategy of each car except the leader is now sketched before it is discussed in more detail.

A car is either on the minimum *safe* distance to its predecessor or further away; otherwise its way is not accident absent. The approaching phase to reach the minimum steady state distance consists of rapid changes of maximum acceleration and maximum deceleration: while not on a minimum safe distance, the car accelerates. Otherwise, it decelerates. As soon as the steady state can be reached, the car accelerates into it and then keeps a steady speed as long as it is not on the minimum safe distance.

The first car c_1 maintains the target speed. Each next car c_i enters the lane at a safe distance to its predecessor c_{i-1}. It accelerates until it is on the minimum safe distance. As discussed in the previous section about traveling-optimal behavior, this is an estimation by c_i, because it does not know how c_{i-1} accelerates between two subsequent beacons. Therefore, c_i has to stay on an accident-absent way to the worst-case way of c_{i-1}. It estimates c_{i-1} to begin braking at the time of sending the most recent beacon which c_i received. c_i will become faster than c_{i-1} during the approaching, because c_{i-1} enters the steady state after a couple of beacon intervals. Then, with each new beacon, c_i gets closer to c_{i-1}; we already know from previous discussions that its speed and distance converge to a steady state of oscillating speeds. c_i passes the minimum steady state distance during that, i.e., the minimum distance at which following with constant speed is possible. This distance can be reached via braking at the right time before c_i is at the minimum safe distance. So c_i is able to enter the steady state; this is in contrast to the steady state for traveling optimality which c_i can only converge to.

The steady state's minimum distance between the cars for a constant speed following is $B_1 w'_1$ with w'_1 being the constant speed of c_1. This will be shown in the following. The minimum steady state distance equals the distance that c_i drives within one interval with the same speed like c_{i-1}. After exactly one interval, c_i estimates that it has is to brake until a full stop after which it would be being bumper-to-bumper to c_{i-1}. But with the next beacon that arrives, c_i realizes that it is not on the minimum safe distance and therefore does not have to brake. The following equation describes the safe distance as estimated by c_2 after one interval. The index t means that a variable describes the true value, e.g., the true speed of a car, while a missing t denotes an estimated value.

$$\begin{aligned}
\Delta w(t_{i+1}) &= \frac{1}{2D}\left((w'_1(t_{i+1}))^2 - (w'_{2,t}(t_{i+1}))^2\right)\\
&= \frac{1}{2D}\left((w'_{1,t}(t_i) + DB_1)^2 - (w'_{2,t}(t_i))^2\right)\\
&= \frac{1}{2D}\left((w'_{1,t}(t_i))^2 - (w'_{2,t}(t_i))^2\right) + w'_{1,t}(t_i)B_1 + \frac{D}{2}B_1^2\\
&= 0 + w'_{1,t}(t_i)B_1 + \frac{D}{2}B_1^2.
\end{aligned}$$

The estimation error c_2 has about c_1 is $DB_1^2/2$:

$$\Delta w(t_{i+1}) = \Delta w_t(t_i) + \frac{D}{2}B_1^2.$$

Combining both equations, $\Delta w_t(t_i)$ is obtained:

$$\begin{aligned}
\Delta w_t(t_i) + \frac{D}{2}B_1^2 &= w'_{1,t}(t_i)B_1 + \frac{D}{2}B_1^2\\
\Leftrightarrow \Delta w_t(t_i) &= w'_{1,t}(t_i)B_1. \qquad (4.2)
\end{aligned}$$

So the necessary true distance is $\Delta w_t(t_i) = B_1 w'_{1,t}(t_i)$. Here, $w'_{1,t}(t_i)$ is the true constant speed of c_1, which is referred to as w'_1 above.

4.3.1 Discussion

It is now discussed why the car behavior is optimal regarding road-usage efficiency. It has to be shown that a car reaches the minimum steady state distance in an accident-absent manner in minimum time with that behavior. This is structured by looking at the distinct phases of acceleration, deceleration, and constant speed driving. A car accelerates with A while not in the steady state or on minimum safe distance. No other behavior lets the car be further ahead at a given time. If a car is at the minimum safe distance to its preceding car, it is at least as fast as with any other way. Being faster is not possible, otherwise it would be closer to the preceding car than the minimum safe distance which would not be accident absent. At the minimum safe distance, a car must brake for staying accident absent.

Because the stated behavior lets the car at least be as far ahead as with any other behavior, it approaches the steady state distance the fastest. The braking maneuver to enter the steady state starts earlier than braking is required for accident absence. If the car braked later, it would be at the steady state distance with too high speed. It could not maintain the distance and had to brake and accelerate again to be at steady state distance with the correct speed. This would take longer than braking perfectly into the steady state in the first place.

The braking maneuver to enter the steady state begins at time $t_{a,st}$. Braking is necessary because the car has to be faster to catch up to its predecessor; braking is stopped if the car is at the steady state distance and speed. For two cars c_1 and c_2 with $c_2 <_c c_1$, $t_{a,st}$ is calculated on receiving an update at t_i as follows:

$$t_{a,st} = t_e + \sqrt{\frac{\Delta w(t_e) - w_1'(t_i) B_1}{A}} \tag{4.3}$$

with

$$t_e = t_i + \frac{w_1'(t_i) - w_2'(t_i)}{A}$$

and

$$\Delta w(t_e) = w_1(t_i) - w_2(t_i) + \frac{A}{2}(t_e - t_i)^2.$$

The formulae are similar to those of Theorem 1 (about reaching the steady state with omniscience) because c_2 assumes c_1 to drive with constant speed. Let t_a be the time at which the cars would be at the *estimated* minimum safe distance if c_2 did not brake at $t_{a,st}$:

$$t_a = t_e - \frac{w_2'(t_e)}{A} + \sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}} \tag{4.4}$$

with

$$\Delta w(t_e) = w_1(t_i) - w_2(t_i) + A(t_e - t_i)^2.$$

t_e is calculated in the same manner as above. To stay accident absent, $t_{a,st}$ is to be compared to t_a and the minimum of the two has to be chosen to begin braking.

It was assumed so far that the distance between the vehicles is greater than the steady state distance at t_i, now the the other case is considered: let the distance be smaller than or equal to the steady state distance. Several cases to reach the steady state in minimum time have to be distinguished which are sketched in the following. If c_2 is faster than c_1 at t_i, it calculates if braking to the speed c_1 allowed it to enter the steady state. If so, c_2 simply brakes. If it was too near after braking to the speed of c_1, it brakes longer and accelerates such that it reaches the steady state. Otherwise, if was is too far away after braking to the right speed, it accelerates first and then brakes until it is again at the initial speed such that braking further on allows it to enter the steady state. If c_2 is slower than c_1 at t_i, it calculates if an acceleration is sufficient to enter the steady state directly. If so, c_2 accelerates. Otherwise, if it was too far away after acceleration, it accelerates longer and then brakes to enter the steady state. If is was too near, it accelerates for a shorter time, followed by another braking-accelerating phase to reach the steady state.

4.3.2 Minimum Steady State Road Usage

For multiple cars following each other with oscillating speeds in their steady states, it is stated in Theorem 4 on Page 53 that a car $c_i, i > 1$, should send with the same frequency as its predecessor c_{i-1} but with a sending offset of a half sending interval $(B_{i-1}/2)$ to minimize the average road usage. The bandwidth has to be shared equally among all cars to allow each car to send with that same frequency. Now, a bandwidth sharing rule for the case of constant speed following will be discussed.

Let $n \geq 2$ cars belong to a platoon and let $1/B$ be the available bandwidth for beaconing. The minimum distance for a car following another with steady speed is $B_1 w_1'$. The minimum sum of distances for all cars in the platoon is $(n-1)B_1 w_1'$ and $\forall i : B_i = B_1$. This will be shown with a complete induction.

Starting at $n = 2$ cars, the first car sends with the full bandwidth of $1/B$. Then the distance $B_1 w_1' = B w_1'$ is minimal.

Let the assumption hold for n cars. The platoon's road usage

$$\sum_{i=1}^{n-1} B_i w_1' = (n-1) B_1 w_1'$$

is minimal and $\forall i : B_i = B_1$.

For $n \to n+1$, the sum of distances is

$$\sum_{i=1}^{n} B_i w_1' = \sum_{i=1}^{n-1} B_i w_1' + B_n w_1' = \left(\sum_{i=1}^{n-1} B_i + B_n \right) w_1'.$$

The only parameter to adjust for finding the minimum road usage is the sending interval of the individual cars. Towards this, two cases are differentiated.

Case 1: the nth car is given more bandwidth than the equal share, while that extra share is drawn equally from all other cars. Let $\alpha \in [0, n-1]$:

$$(n-1)\frac{1 - \alpha/(n-1)}{B_1} + \frac{1+\alpha}{B_1} = \frac{n-1-\alpha}{B_1} + \frac{1+\alpha}{B_1} = \frac{n}{B_1} = \frac{1}{B}.$$

The road usage of the platoon is:

$$w_1' \left((n-1)\frac{B_1}{1-\alpha/(n-1)} + \frac{B_1}{1+\alpha} \right).$$

Deriving this for α shows that the local minimum in $[0, n-1]$ is at $\alpha = 0$, i.e., when each car is assigned the same bandwidth.

Case 2: the nth car is given less bandwidth than the equal share, while that extra share is added equally to the other cars. Let $\alpha \in [0, 1]$:

$$(n-1)\frac{1+\alpha/(n-1)}{B_1} + \frac{1-\alpha}{B_1} = \frac{n-1+\alpha}{B_1} + \frac{1-\alpha}{B_1} = \frac{n}{B_1} = \frac{1}{B}.$$

The road usage becomes:

$$w_1' \left((n-1)\frac{B_1}{1+\alpha/(n-1)} + \frac{B_1}{1-\alpha} \right).$$

The local minimum for $\alpha \in [0, n-1]$ is at $\alpha = 0$.

4.3.3 Effect of a Beacon Loss

We now know how optimal behavior looks like when sending beacons at a fixed interval B_1. So far, however, lossless communication was assumed. If a beacon is not received within the expected time, the steady state cannot be maintained. The effect of a single, lost beacon is explained in the following to show this.

Given a platoon of $n \geq 2$ cars with $c_n <_c \ldots <_c c_1$. Let all cars drive in their steady states and let t_0 be the time of sending an arbitrary beacon of the leading car c_1. c_2 receives that beacon and continues driving with steady speed. At time $t_0 + B_1$, c_2 estimates that it is on the minimum safe distance that would cause it to begin braking if no beacon was received. Although c_1 sends a new beacon, it is assumed that c_2 does not receive it this time. So c_2 estimates c_1 to be braking further on since t_0 and to be so close that braking is required for staying accident absent. On $t_0 + 2B_1$, the next beacon of c_1 is received successfully. c_2 braked during the whole interval which caused a distance growth of $(1/2)AB_1^2$. But now, due to new information that reveals that c_1 drove with constant time all the time, c_2 starts accelerating to approach the steady state again. This takes several beacon intervals in which each beacon is delivered successfully.

What happens to following cars? Obviously, the distance and speed changes of c_2 will influence following cars. The effect will be explained for the worst case for clarity reasons. The worst case regarding speed and distance difference to the steady state is when each car sends at the same time, i.e., the sending offset is zero. Of course, this worst case mainly is of theoretical interest because simultaneous sending times will lead to packet collisions on the wireless channel. The topmost line depicted in Figure 4.3 illustrates the sendings of beacons and the single loss happening at $t_0 + B_1$. In the third line of the figure, the reaction of c_2 is shown. Although all updates sent at $t_0 + 2B_1$ are delivered successfully, c_2 triggered a chain reaction by leaving the steady state: c_3 reads c_2's beacon telling that the true distance between them is less than the steady state distance. c_3, in fact, is on the minimum safe distance to c_2 and must begin braking. At each next beacon sending time, one additional car begins to brake for a duration of one sending interval. This process keeps traveling upstream the whole platoon and affects every car. It can also be observed that all cars react, one after another, at the same point of the lane: cars travel with speed $w_1'(t)$, beacons are sent with the interval B_1, and between two updates, each car moves by $w_1'(t)B_1$. However, without any additional lost beacon, the situation stabilizes again for each car after a few update intervals by entering the steady state again.

4.3.4 Withstanding Losses

The goal of road-usage efficiency demands to maintain a steady state distance. When packet losses occur due to unreliable communication, maintaining the minimum steady state distance is not possible. The distance does not allow for packet losses. To withstand losses without deceleration, a larger steady state distance is necessary. Clearly, at a given packet loss probability p, the number of consecutive losses is not limited, so that there is no guarantee that cars can always stay in their steady states. However, a reasonable robustness against losses is possible.

One needs to state a reasonable number of consecutive losses y at a random car that the platoon has to withstand without a single car leaving the steady state. It has been discussed in the previous section that a car drives the distance $w_1'B_1$ from one beacon to the next. So it is easy to see that the minimum steady state distance to endure y losses is $(1 + y)w_1'B_1$. The worst case of sending at the same (or a nearly equal) time can be avoided by inserting a sending time offset between successive cars. By this, the maximum speed deviation from the steady state in case of a loss is lessened and the effect on each additional upstream car gets smaller, too. A loss does not necessarily influence each following car of a platoon anymore.

The sending offsets between cars with beaconing with a fixed sending interval, however, are neither controllable nor predictable and thus the worst case has to be sup-

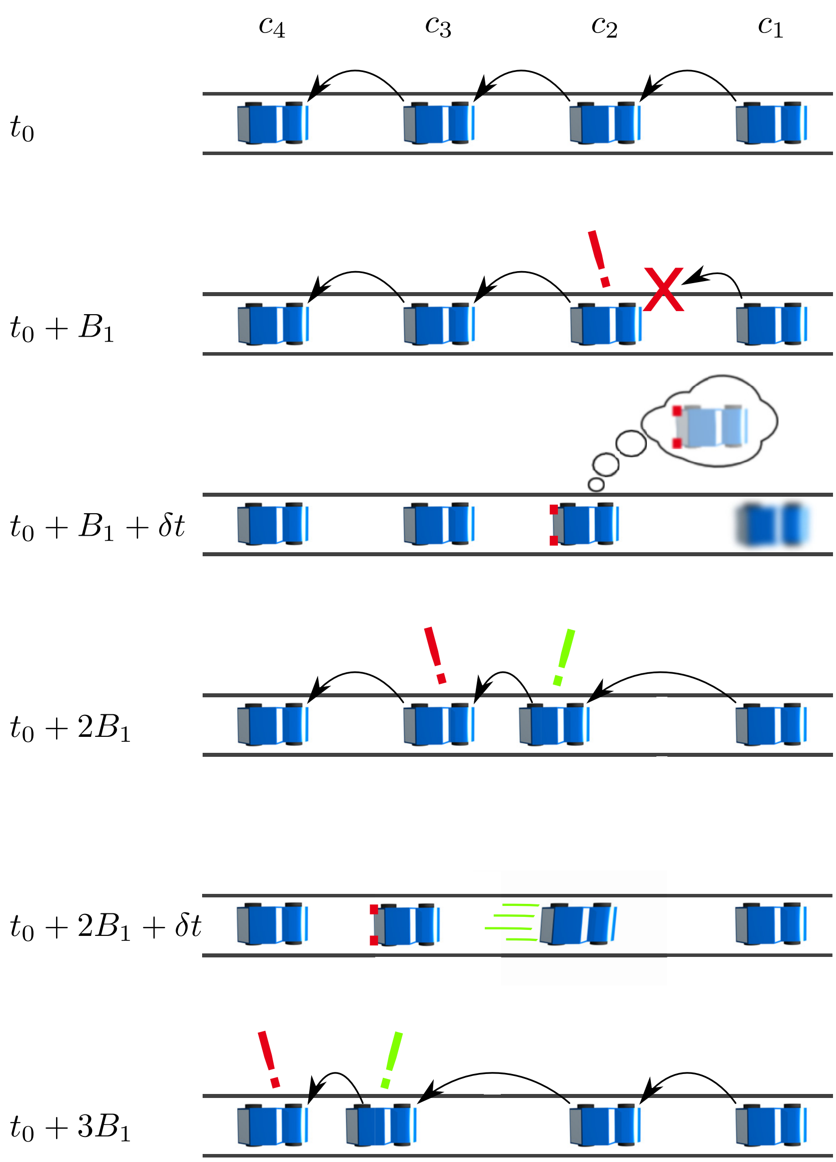

Figure 4.3: $\delta t \in [0, B_1)$. Propagation of braking events upstream a platoon caused by a single, lost packet.

ported. Repeated collisions are likely to occur with fixed interval beaconing and a randomization of the sending times has been proposed in the literature as discussed in Section 2.4.1. This solution leads to an even larger steady state distance: the randomization is modeled with $a \in (0, 1)$ and a uniform distribution for drawing the next sending time within $[1 - a, 1 + a]B_1$. This causes the minimum steady state distance to become extended to $(1 + a)w'_1 B_1$ without regarding losses. For supporting up to y losses, the distance increases to $(1 + y)(1 + a)w'_1 B_1$. In the following section, it is discussed how the algorithm *Carrot* allows for smaller steady state distances than this.

4.4 Carrot

Carrot is an adaptive beaconing algorithm for cars driving in a platoon using a Cooperative Adaptive Cruise Control (CACC) system. The core idea of Carrot is to let a car send a beacon directly after it received a beacon from its preceding car: the right to send is passed upstream one car by another. This avoids equal sending times and enables Carrot, in contrast to an unstructured sending scheme, to detect missing beacons as their transmission is mandatory for maintaining the steady state. A fast repetition algorithm is part of Carrot to mitigate distance growth when a missing beacon is detected. This section discusses Carrot up to the details of how it is initialized and how cars join and leave a platoon.

4.4.1 Algorithm

Beaconing as discussed in the literature has a fixed sending interval that is usually at 0.1 s or 1 s. These intervals are much larger than the minimum sending intervals possible with wireless communication technology, e.g., 802.11p. For the formal model, a minimum interval between beacons of B seconds has been introduced. Let the leading car of a platoon send with a target interval $B_t \gg B$. Each car in the platoon sends its beacon directly after it received a beacon from its predecessor.

The platoon leader has no car in front and so it never receives a beacon originating from downstream the lane. To find out whether a car is platoon leader, each car maintains a *leader sending time* t_l: each time t a car sends a beacon, its leader sending time is set to $t_l = t + B_t$. If a beacon from a downstream car is received, the receiver is not platoon leader and thus sets $t_l = \infty\,\mathrm{s}$.

An important part of Carrot is handling delays. For supporting delays, it is necessary to define an interval in which the beacon is created, sent, received, and, eventually, processed at the follower. Let this delay interval be b, $B_t \gg b > B$. The sender gathers information for its beacon each B_t seconds and the receiver adapts its behavior b seconds later. At this time, the receiver is able to start its own beacon sending procedure. So information gathering for a beacon, on receiving one from the preceding car with information from time t, happens at $t+b$. The minimum steady state distance with support for delays is $(B_t + b)w'_1$. The interval b does not contain time to apply behavior changes like a driver's reaction time which also cause an increase of the steady state distance. Modeling this in detail is not in scope of this thesis. In addition, the delays are not constant in a real environment as they depend on various parameters that change over time. The delay interval should be chosen as a robust "upper limit" that covers most common situations. Larger delays then cause a car to leave the steady state.

A beaconing approach termed reactive beaconing by the authors of [vEKH10] (that was described in Section 2.4.1) also uses an upstream sending scheme similar to Carrot. However, the realization differs due to another intention for employing structured sending times. While the design of Carrot focuses on application objectives, the reactive beaconing approach is aimed at enhanced network layer properties. By always assuming the worst-case behavior of the preceding car, Carrot enables safe following without out-of-order alarm messages that are discussed in related work (that is described in Section 2.4).

4.4.2 Fast Repetition

The sending scheme of Carrot, in contrast to beaconing, allows for an implicit detection of losses. The detection, as well as the fast repetition mechanism building on this, is described in the following.

The algorithm is extended with the *repetition time* t_r. If a car sends a beacon at t and the car knows about another car driving behind it, the repetition time is set to $t_r = t + 3b$ at which b is the delay interval as discussed above. The repetition time is set to $t_r = \infty$ if any beacon is received. The sending of a beacon is now triggered by a further condition besides receiving a beacon from downstream: a car sends a beacon if it either received a beacon from its preceding car or at repetition time. Before sending a beacon, Carrot always refreshes the information to be sent.

The choice of repeating after the multiple of three times the delay interval is based on the following arguments. Let c_{i+1} follow c_i. For setting the repetition time, it is exploited that c_i overhears c_{i+1}'s beacon and thereby knows that c_{i+1} received c_i's own beacon. If c_i sends at t, the beacon of c_{i+1} is expected to arrive the latest $2b$ seconds after t at c_i. Does the beacon not arrive until then, the sender detects a lost beacon and can react to this event. By directly re-sending at $t + 2b$, the beacon is repeated the fastest possible. Upstream nodes do not occupy the medium before the repetition is finished. This is beneficial if the loss is due to a medium being already congested. Prohibition of sendings at upstream nodes is an obvious means to prevent load on the medium. However, with wireless communication, it is possible that c_i does not get the beacon from c_{i+1}—but c_{i+1}'s subsequent follower c_{i+2} does: then c_i re-sends at the same time at which c_{i+2} sends. c_i and c_{i+2} have the same distance to c_{i+1} in the steady state. This causes a packet collision at c_{i+1}. c_{i+2} is now given a higher importance for sending to avoid c_i detecting a false negative at which a repetition is triggered though it is not necessary. Therefore, re-sending is altered to happen $3b$ seconds after a previous sending. As soon as c_i gets any beacon from behind, it stops repeating.

The triggering of beacons is related to the topology of the nodes in Carrot: a car sends a beacon if it received a beacon from its predecessor. Each beacon is broadcast and contains individual information, so there is no routing of beacons along the nodes. However, the trigger scheme causes a flow of beacons along a platoon in upstream direction. Controlling the sending of packets in a flow through a trigger mechanism is a common idea, as has been discussed in related work with, e.g., the back-pressure concept in Section 2.4.3. The mechanisms of those protocols were designed to improve network layer characteristics like a reduction of packet losses; in contrast to this, the trigger mechanism of Carrot is intended to reduce loss effects on the application layer.

Bandwidth usage

A car detects a loss if one of two distinct sendings fails: the sending of its own beacon (the correct case to repeat) and the following car's sending (the wrong case). A presumable loss is recovered by a repetition which is, in fact, a sending of a further beacon. This repetition of beacons requires bandwidth. At a given loss probability p, the probability that one of the two beacons gets lost is $1 - (1 - p)^2$. Assume that there is enough time for all necessary repetitions within one target beaconing interval B_t (which may be multiple as also repeated beacons can get lost). Carrot then sends $1/(1 - p)^2$ times the beacons compared to a fixed interval beaconing. Therefore, a comparison to fixed interval beaconing is fair only if both protocols use the same bandwidth. For this, the target sending interval of Carrot is enlarged to $B_t/(1 - p)^2$.

Minimum steady state distance

The minimum steady state distances of fixed interval beaconing and randomized interval beaconing, as discussed in Section 4.3.4, grow linearly with the number of consecutive losses y that have to be absorbed. The growth is discretized in steps of multiples of the target interval B_t. The distance growth of Carrot is linear, too, but the steps are of the size of the delay interval b:

$$\Delta w_{\text{Carrot}} = \left((1 + 3y)\, b + \frac{B_t}{(1-p)^2} \right) w_1'.$$

Delays have not been considered for the other beacon algorithms yet, so they are now analyzed with regard to delays, too. Fixed interval beaconing, as discussed earlier, enables a minimum steady state distance without sending delays at $(1+y)B_t w_1'$. Delays of size b cause the following distance to become $\Delta w_{\text{fixed}} = ((1+y)B_t + b)w_1'$. To see in which cases Carrot performs better than beaconing, the expressions of the steady state distance are put in relation. Let $p < 1$ and $y > 0$:

$$\begin{aligned} \Delta w_{\text{fixed}} &> \Delta w_{\text{Carrot}} \\ \Leftrightarrow ((1+y)B_t + b)w_1' &> \left((1+3y)b + \frac{B_t}{(1-p)^2} \right) w_1' \\ \Leftrightarrow 1 + y - \frac{3yb}{B_t} &> \frac{1}{(1-p)^2} \\ \Leftrightarrow p &< 1 - \sqrt{\frac{1}{1 + y - 3yb/B_t}}. \end{aligned}$$

The values of a usual application case are $B_t = 1\,\text{s}$ and $b = 1\,\text{ms}$. Consecutive losses happen with probability p^y, so y can be set to a rather small value depending on the desired level of robustness. For $y = 1$, the inequation shows that Carrot is preferable at a loss probability of up to 29 %. Figure 4.4 shows curves of the limiting loss probability (at which $\Delta w_{\text{fixed}} = \Delta w_{\text{Carrot}}$) over the beacon sending interval for distinct values of y; b remains fixed to 1 ms. When lower loss probabilities are expected, Carrot enables a shorter steady state following distance than beaconing with fixed sending intervals.

4.4.3 Joining

A car entering the lane does not know how many cars are driving in front of it. Simply starting to send at an arbitrary time is prone to bother other senders, so this should be avoided. As a solution to this, an initial phase is introduced. To cope with the limited number of cars that are able to send within a given target interval, a mechanism is proposed which dynamically extends the sending interval.

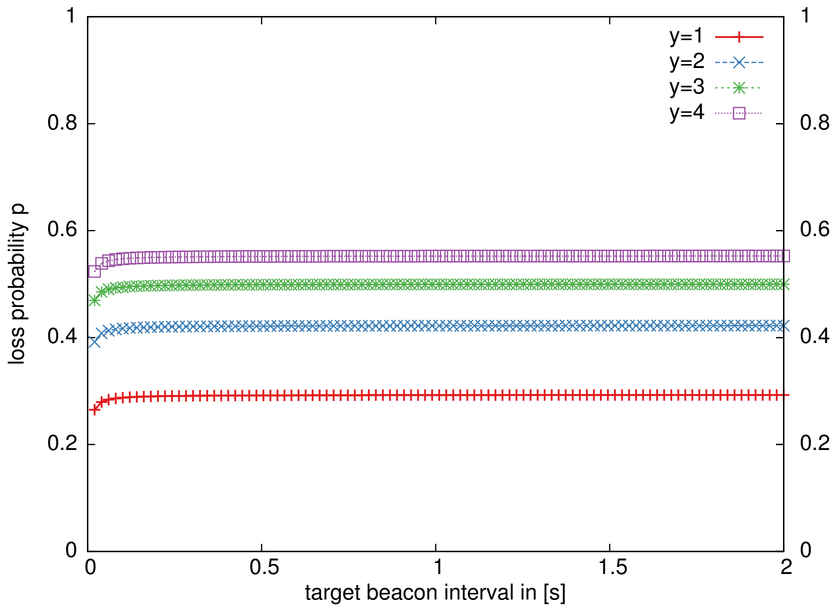

Figure 4.4: The loss probability at which Carrot is able to maintain the same steady state distance as fixed interval beaconing, plotted over the beacon interval and with multiple values of endurable consecutive losses y.

Initial phase

A car appearing in the lane at time t has an *initial sending time* $t_i = t + 2B_t$. As long as the initial sending time is later than the current time, beacon sending is suppressed. Both the leader sending time and the repetition time are set to $t_l = t_r = \infty$ initially. If a beacon from a downstream car is received and this beacon tells that the sending car is closer than any other known downstream car, the sender is stored as predecessor.

The initial phase ends at t_i or if the second beacon from the currently set predecessor is received; this ensures that the correct predecessor is found. In this case, t_i is set to a time earlier than the current time t, so that $t_i < t$. If no predecessor has been found until t_i, the leader sending time is set to the current time, $t_l = t$, to indicate that the car is platoon leader. Adding a new car within the platoon instead of at its end is possible by changing predecessor information such that the new car sets "its" predecessor at first. Next, the car in the platoon having the same predecessor considers the new car as its predecessor. Cooperation is required among the cars to create a gap for the new car between the correct two cars; this is not covered in more detail here. Merging cars from two lanes, however, is discussed in Chapter 5.

Interval extension

It is now explained how the sending times are adjusted with Carrot when there is at least one car too much to send in one target beacon interval. Each car sends beacons every B_t seconds. Between the sendings of two successive cars is a delay of b seconds. So this situation occurs with n cars if $(n-1)b > B_t$. The leader enlarges the beacon sending interval in this case by postponing its beacon: on receiving a beacon at time t, the leader checks if $t_l < t + 3b$. In that case, the leader sending time is adjusted to $t_l = t + 3b$. That way the leader sends after a silence of three sending times. The silence is long enough so that repetitions at the end of the platoon are supported as explained above.

4.4.4 Leaving

If a car leaves the platoon, two cars have to detect this: car c_i being directly in front of the leaving one and car c_j behind of it. c_i does not have to adapt its sending times to support the leaving operation: it stops repeating beacons as soon as it hears a beacon from an upstream car; this is independent of the sender's identifier. If the last car of the platoon leaves, however, c_i is the new last car. To stop repeating beacons, a timeout is required. c_i counts the number of its repetitions and stops repeating after sending a predefined number of sendings.

For c_j, the successor of the leaving car, two cases have to be distinguished. The leaving car may leave regularly by changing to another lane or the car's radio unit is defective. In the first case, c_j receives a beacon after the lane is left but before the leaving car is out of range. The lane change is encoded within the position information in this beacon. c_j is then able to change its predecessor from the leaving car to c_i, the next car in downstream direction. If there is such a c_i, it is known by beaconing. Otherwise, c_j is the new platoon leader. The case of a defective radio unit requires a special treatment. c_j does not know why it does not receive beacons anymore. Being limited to beacon-based knowledge, a dead radio unit cannot be differentiated from a dead car, so ignoring missing messages is not accident absent; c_j has to stop at the point it estimates the probably dead car to have stopped. An "about-to-leave" message, an obvious means in this situation, is not helpful since a leaving car might still become defective and stop before it leaves the lane.

In case knowledge from other sources is available, for example by including further sensors in the system and extending the behavior algorithm to use their measurements, waiting forever because of a dead radio unit can be prevented: the waiting car maintains a *dead time* t_d for the preceding car. At the the dead time, the radio unit of the car

leader sending time	t_l
repetition time	t_r
initial sending time	t_i
dead time	t_d

Table 4.1: Parameters of Carrot.

directly ahead is considered defective and the next downstream car is considered as new predecessor for communication. The car-following algorithm, of course, still has to treat the car with the defective radio as predecessor. The value of the dead time is set such that n repetitions of the preceding car are possible. n is recommended to be set to $n = B_t/b$. This means a car waits one target beacon interval before the preceding car's radio unit is treated as defective. Table 4.1 lists the parameters that are used by Carrot.

The leader's sending algorithm supports leaving without adaptions. If a car within the platoon leaves, the leader notices this indirectly: in case the cars finish sending before $t_l - 3b$, the leaving does not affect the next sending of the leader. In the other case, when there are more cars than fit into a target beaconing interval, the leader sends if it did not hear a beacon from an upstream car for $3b$ seconds.

4.4.5 Coping with the Sending Range

Carrot functions properly when the road usage of the platoon is larger than the communication range. It is assumed, though, that each two successive cars are in communication range. Let car c_i be the furthest car in upstream direction whose beacons are received by the leader. A beacon of c_i that arrives at the leader guarantees that the leader's next beacon is sent the earliest $3b$ seconds later. c_i repeats its beacon if necessary after $2b$ seconds, thereby delaying the leader's next sending again by $3b$ seconds. If the leader eventually sends and there exists another car behind c_i that sends at the same time, it is outside the leader's communication range. For both sending cars, it is estimated that their respective next follower is closer to them than to the other sender. In this setting, it is likely that the beacons are both received successfully at the respective followers.

In spite of this, c_1 sends before the last cars of the platoon received the beacon of the previous "round". Two sending tokens travel along the platoon simultaneously. Without unexpected delays, the updates do not have overlapping sending ranges since each distance between two cars is the same and the tokens travel with the same speed. In case that packet losses occur at cars still busy with the older sending round, the distance between the senders can shrink such that upstream cars sense communication of downstream cars and vice versa. In that case, the upstream cars stop sending to avoid packet collisions. This is advantageous compared to stopping the sending of the downstream cars, because the higher delay already causes the upstream cars to leave their steady states while the downstream cars are still in their steady states or at least closer them (because the downstream cars have more current information on their predecessors). The sending stop implies one missing update at each car being further upstream.

In the steady state, the distance between each two successive cars is equal. This can be exploited to adjust the transmission power of each sender in a way that reduces the number of cars suffering from interferences by a sending. For this, it is referred to related work like, e.g., [TMSH06] which discusses adjusting the transmission power in detail.

4.4.6 Carrot and Other Applications

An implementation of Carrot for use in a realistic vehicular environment has to cope with network effects caused by foreign nodes that use the same communication channel. In addition, other C2CC applications on the same node contend with Carrot for access to the communication medium. Carrot is evaluated on top of IEEE 802.11p in the following simulation and evaluation part because this is the standard technology for applications that use ad-hoc car-to-car communication. The simulation does neither contain foreign senders nor simulate other applications. This section discusses what is to be expected under such circumstances and suggests a possible countermeasure.

To enable a distributed access to the medium, IEEE 802.11p implements CSMA (Carrier Sense Multiple Access). This mechanism lets a car send only if the medium was sensed to be idle for a certain time. Should a collision during transmission occur (which is noted at the sender by a missing acknowledgment of the receiver), the packet is repeated. This repetition scheme is referred to as CSMA/CA (CSMA with Collision Avoidance). The repetition algorithm uses a randomized, exponential back-off scheme with unbounded delay [BUSB09]: transmission success is not guaranteed when sending after a back-off time due to other senders that can start sending at the same time. After a collision is detected, a resending of the packet is prepared and a new back-off value is chosen for this. If collisions happen repeatedly, the back-off time grows fast for each next repetition.

Although Carrot's beacons are sent by broadcasting, for which receivers do not send acknowledgments and so collisions cannot be detected, the carrier sensing time can still become long due to other traffic on the medium. The delay interval of Carrot introduced in Section 4.4.1 that is used to model carrier sensing delay (among other transmission and processing delays), however, is a constant that covers a "reasonable" carrier sensing duration. In other words, the carrier sensing may take longer than it is supported by the delay interval.

The delay of 802.11 CSMA under high traffic conditions is theoretically investigated in an article by Ziouva and Antonakopoulos [ZA02]. The delay is found to be around 100 ms in settings comparable to 802.11p in the cases of using acknowledgments or the RTS/CTS mechanism, a second method to avoid collisions. This delay is much longer than the delay interval used in the simulations section. A longer delay requires a larger steady state distance between cars. The other two beaconing algorithms investigated in this chapter suffer from a larger delay, too, but not as much as Carrot does: the delay interval is an important parameter for the maximum number of nodes supported with a beaconing interval. A further parameter, the additional distance to endure consecutive losses, also depends on the delay interval; the distance grows linearly with the delay interval.

A proprietary approach to get a grip on the carrier sensing delay issue would be to implement Carrot directly on MAC layer and to give its beacons priority over other 802.11p-based packets by means of a dedicated time slot or a shorter interframe spacing than for other packets. While such a new MAC especially for Carrot is obviously difficult to realize, a more general approach to change the MAC layer of 802.11p has been proposed by Bilstrup et al. [BUSB09]. The authors adapted the MAC scheme STDMA (self-organizing TDMA), which enables bounded delays, for use with vehicular networks. STDMA is already deployed for collision avoidance between ships. The drawback of STDMA is its need for synchronization: the number of time slots is dynamic and vehicles have to know about that number of slots as well as when a new frame starts. However, a solution like this is expected to enable a bounded delay for Carrot's delay interval.

4.5 Simulation and Evaluation

In this section, Carrot and the beaconing communication algorithms are evaluated with simulations for their suitability to a CACC application. A platoon is simulated in which each car is equipped with radio hardware for ad-hoc communication. Carrot is compared to the two beaconing algorithms with fixed sending interval and with uniformly randomized sending intervals. The simulator offers a lossless data pipe as well as 802.11p-like communication, including a delay model and a packet loss model.

Two setups with different car-following models are investigated: the first setup implements the road-usage efficient behavior for showing the effect of beacon loss as it was described in Section 4.3.4. The second setup employs the Intelligent Driver Model (IDM [THH00]). IDM is a lightweight car-following model for a single lane that calculates each car's acceleration individually. The model aims at simulating a realistic car behavior and does not choose extreme acceleration values in most situations, in contrast to the optimal behavior. The setup is used to evaluate the intervals between the reception of successive beacons. The simulation results show that the sending intervals with Carrot are nearly as constant as with the fixed interval beaconing. But in opposition to beaconing, Carrot avoids the hazard of repeated collisions. Randomized beaconing avoids most repeated collisions, but the reception intervals are not as constant as with Carrot.

4.5.1 Simulation Setup

For both the simulations of vehicle behavior and communication, the network simulator ns is employed in version 3.13. An object-oriented modeling of cars, following models, and communication units has been implemented based on the highway mobility framework [AW10] in version 2. This framework enhances ns with application elements like cars and highways. It has been adapted to suite the requirements of the scenario: the following and communication classes now offer more details and are easily exchangeable.

A following model is assigned to each car; it calculates a car's behavior until the next calculation step. IDM uses fixed simulation steps to update all cars' accelerations; the steps are set to 1 ms. With the road-usage efficient behavior, the next step is not predefined: a new calculation is made only at receiving a new beacon from a car's predecessor. Each car is attached a communication unit for sending and receiving with an initially defined communication algorithm. Sending and receiving is possible with the nanosecond granularity of ns.

Simulation parameters for the beacon loss evaluation

The simulated scenario consists of one lane and ten cars that are put in the lane at the simulation's start. The initial distance between two successive cars is 50 m and the initial speed of each car is 30 m/s. The maximum acceleration is set to $A = 6\,\mathrm{m/s}^2$ and the maximum deceleration is $D = -A$. These are rather high though still realistic values. The optimal behavior requires a car to use its maximum acceleration when it is not in its steady state. Each car starts at the same simulation time 0 s and the

Desired speed	5 m/s to 30 m/s
Acceleration	1.0 m/s^2
Deceleration	2.0 m/s^2
Minimum gap	2.0 m
Time headway	1.5 s
Delta exponent	4.0
Vehicle length	4 m
Simulation step	1 s

Table 4.2: Car-behavior parameters of IDM for the beacon interval evaluation.

simulation duration is 100 s. The behavior of the cars is based on the information they receive via beaconing.

Communication is realized with SimpleChannel, a concept of ns that pipes sent messages to the receiver without losses or delays. Updates are sent with fixed intervals of 1 s. Delays of 1 ms are regarded, so that a steady-state distance of 30.03 m is expected. One loss is forced after the simulation time of 30 s: car c_2 does not receive the next update of c_1, the platoon's leader. The loss time is chosen such that all cars have entered their steady states until then.

This simulation setup is evaluated concerning the distances between pairs of consecutive cars to depict the upstream effect at a beacon loss in the worst case.

Simulation parameters for the beacon interval evaluation

For evaluating the effective beacon intervals received at the cars, IDM is used. Before each simulation step, cars recalculate their constant acceleration during that step. The acceleration calculation of IDM uses knowledge about the distance to the car ahead and the speed difference to that car. The model assumes necessary information to be available; cars do not depend on beacons to get information about surrounding cars in this setup. The parameters of IDM are chosen to create realistic car following: the desired velocity equals the cars' maximum speed and is the same for all cars but varied for distinct runs. The acceleration is set to 1.0 m/s^2 and the deceleration to 2.0 m/s^2. The minimum gap between cars is 2.0 m, the time headway is 1.5 s, and the delta exponent is 4.0. Initially, one car is put in the lane and new cars enter the lane with a flow of 0.6 cars per second and a speed of 5 m/s. The simulation runs are finished after 1000 s of simulation time. Table 4.2 shows a list of the IDM parameters with their values.

WifiHelper	WIFI_PHY_STANDARD_80211a
Data mode	OfdmRate6MbpsBW10MHz
MAC type	AdhocWifiMac
Delay model	ConstantSpeedPropagationDelayModel
Loss model	NakagamiPropagationLossModel
TxPowerStart	33
TxPowerEnd	33
TxPowerLevels	1
RxGain	0
TxGain	0
EnergyDetectionThreshold	-71 dBm
CcaMode1Threshold	-91 dBm

Table 4.3: Communication parameters of the simulations with 802.11p. The string values describe classes and constants of ns.

As mentioned above, the communication algorithms are implemented using helper classes of ns. They are put on top of the channel and the network device implementations. These are the NqosWifiMacHelper class and the YansWifiChannelHelper class; the helpers are configured such that their setups are close to the 802.11p standard for sending in the control channel. OFDM (Orthogonal Frequency Division Multiplexing) is used at a rate of 6 Mb/s and a bandwidth of 10 MHz. The transmission power is set to 33 dBm without additional antennae gain. The sensitivity thresholds are based on Atheros 802.11a chipsets and the 802.11a standard, see [LRLK10] for more details on this. A constant speed propagation delay is added to the channel, as well as a loss model using a Nakagami-m propagation loss model chained after a Friis propagation loss model. This is a common combination. The Friis loss model is deterministic: all receivers within a certain threshold range of the sender get the packet. The Nakagami-m model is a probabilistic loss model that takes the output signal strength calculated by the Friis model and alters it with noise. That way, the deterministic results of the Friis model are added some salt: instead of a fixed transmission distance after which all beacons are lost, the Nakagami-m model enables a chance of receptions at nodes far away from the sender as well as a probability for losses at near nodes. Table 4.3 lists the communication parameters of the simulations.

The target beaconing intervals that are evaluated are 1 s and 0.1 s. Carrot is simulated with a minimum time between two beacons of 1 ms. This is a rather conservative value compared to the simulated delays which are around 150 μs. It has already been explained in Section 4.4.6 that delays can become much larger due to competing nodes and applications, so a parameter study is recommended for future work. Beaconing is simulated with a fixed sending interval and a randomized interval. With randomized beaconing, a car with an average sending interval B_t draws a number j uniformly distributed in $[-0.5, +0.5]$ at sending a beacon at time t and then schedules the next sending event to happen at $t + jB_t$.

The simulation runs are evaluated for the intervals between receptions of two successive beacons from a car's predecessor. The used metric describes the differences between simulated intervals and the target beacon interval are calculated. Optimally, there is no difference and this value is 0 s. In addition to the simulated differences, the number of sent beacons is counted.

4.5.2 Evaluation

The results of the two simulation setups are evaluated in separate sections.

Simulations with road-usage efficient behavior for loss evaluation

The distances between each pair of successive cars are shown in Figures 4.5 and 4.6. The former shows the communication with fixed interval beaconing and the single, forced loss of the leading car's beacon at its follower c_2 at simulation time 30.01 s (the first beacons are sent at simulation time 0.01 s and then beacons are sent in fix intervals of one second). The latter figure depicts a similar situation with communication via Carrot and a forced loss at 30.001 s. The steady-state distances between cars are 30.03 m each, including the additional distance to account for delays. The figures show an excerpt of the simulation starting shortly before the loss when each car has entered its steady state.

The distance changes after the loss are much lower with Carrot than with beaconing; they are hardly noticeable although the y-axis has been zoomed-in by a magnitude in Figure 4.6. The cars are in their steady states again after about half of the time. This shows that Carrot is much more robust against losses even without additional distances for withstanding losses that were proposed in Section 4.4.2. As expected for fixed interval beaconing, one car after another leaves the steady state; the simulations underline the description of the effect of a loss in Section 4.3.3. The figure also shows how the cars return to their steady states: two sending intervals after the loss, a new beacon is received at c_2 and the cars begin to approach the steady state again, one after the other. Each follower tries to reach a steady state with the predecessor's sent position and speed. Further studies should investigate the impact of transmitting

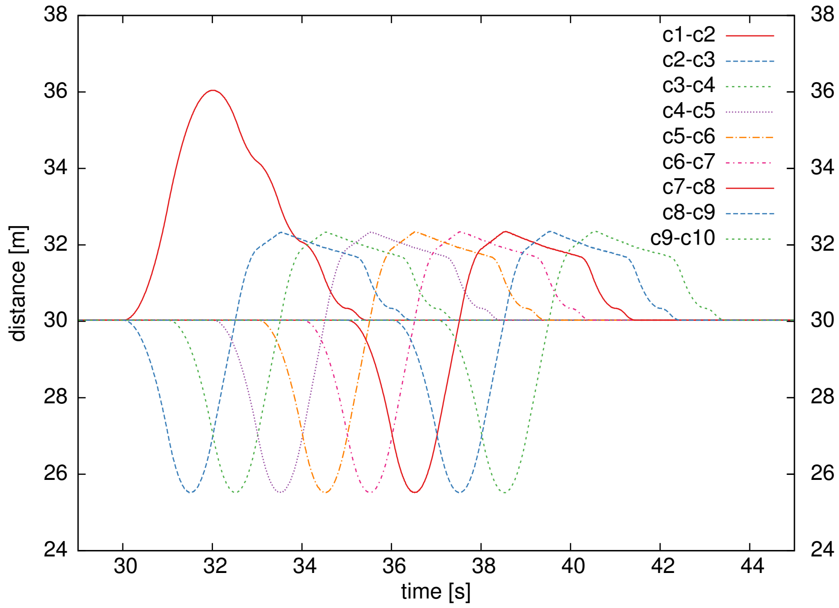

Figure 4.5: Distances between 10 cars with road-usage efficient behavior and communication using fixed interval beaconing. c_2 suffers a single beacon loss at 30.01 s.

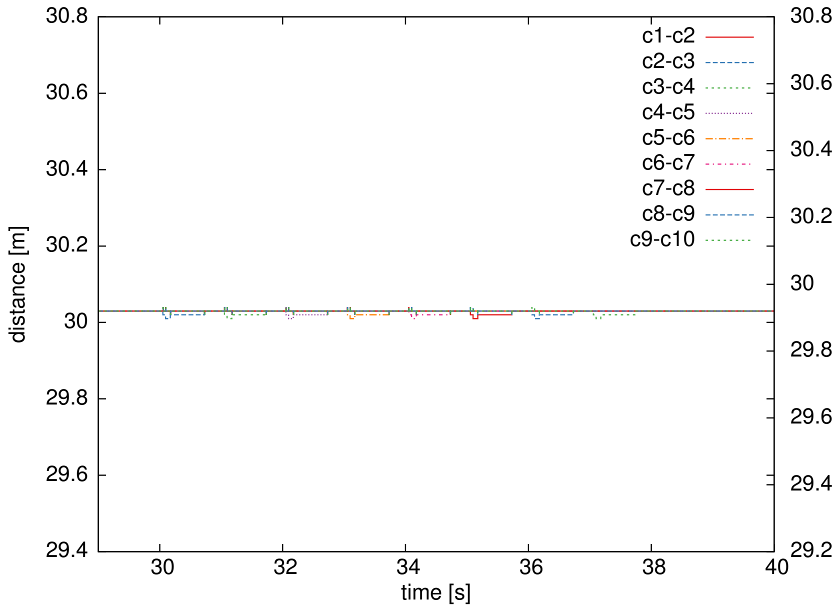

Figure 4.6: Distances between 10 cars with road-usage efficient behavior and communication using Carrot. c_2 suffers a single beacon loss at 30.001 s.

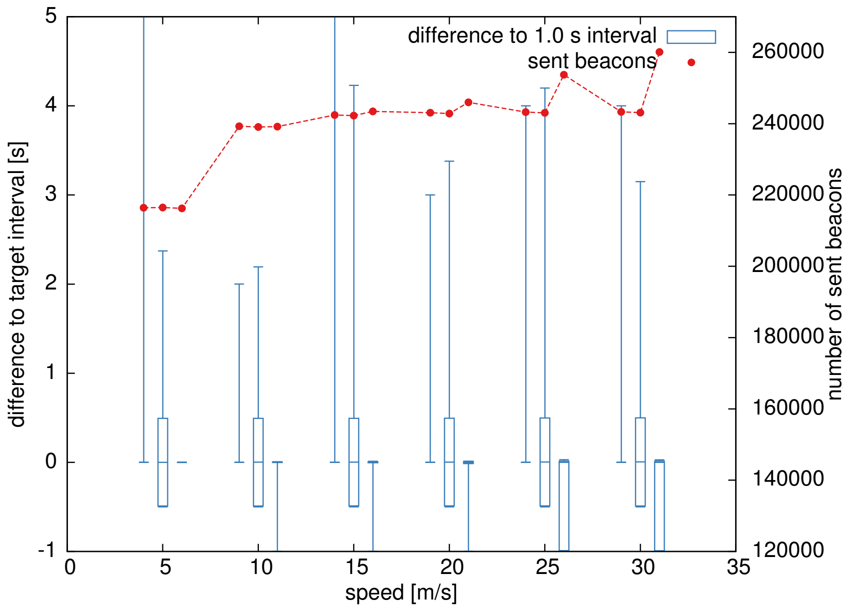

Figure 4.7: Difference to the target interval of 1 s of fixed beaconing (left), random beaconing (middle), and Carrot (right) at different speeds.

the platoon leader's speed, too, as it is suggested for string stability (as discussed in Section 2.2.2).

Simulations with IDM for the beacon interval evaluation

The results of Carrot and the two beaconing algorithms are compared at varying desired speeds of the IDM car-following model. The desired speed influences the car distances. Figure 4.7 shows the difference between the target interval and the actual intervals for fixed interval beaconing, randomized interval beaconing, and Carrot, each with a 1 s target interval. The differences are statistically evaluated and printed as candlesticks with 1 % and 99 % quantiles as box limits, minimum and maximum values as whiskerbars, and the measured median. For each simulation run with distinct desired speed and communication algorithm, the figure also shows the total number of sent beacons. These counts are represented by points for each run and are connected by a line to aid recognition.

The results of the runs with fixed interval beaconing show that the difference to the target interval is negligible at nearly all measurements. Only less than 1 % are not optimal. But lost beacons, though these are only a few, cause receiving intervals of multiples of the target interval (up to 7 s were measured, the maximum is not shown

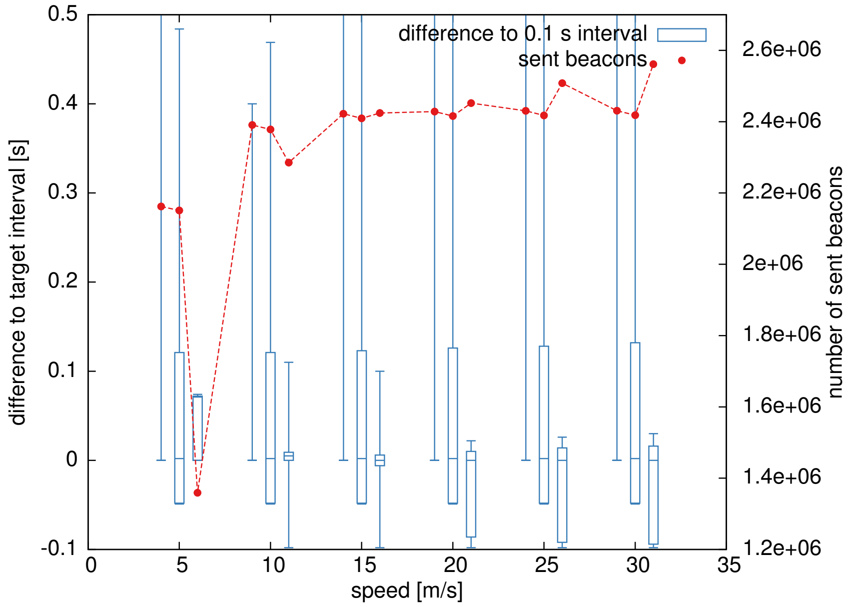

Figure 4.8: Difference to the target interval of 0.1 s of fixed beaconing (left), random beaconing (middle), and Carrot (right) at different speeds.

in the figure). As it is shown above at the other simulations' results, even a single loss is harmful since it can influence all cars upstream of the intended receiver. This is not tolerable.

The candles' boxes of the runs with randomized beaconing cover the whole interval possible with the random interval generation method and the medians are close to the target interval. These values are as expected. Like with fixed interval beaconing, only few beacons were lost, causing a difference to the target interval of several multiples. The total numbers of sent beacons are close to the numbers sent with fixed interval beaconing which is, of course, reasoned by the same target interval.

The medians and quantiles of interval differences with Carrot are close to optimal. A few beacon intervals are very short and the faster the cars drive, the more short intervals occur. This is due to the repetition mechanism that detects false negatives: it resends beacons if the sender thinks its previous beacon has not been received—which also happens if the sender does not get the receiver's beacon sent in response. It is noticeable that Carrot sends up to 7 % more beacons compared to the other algorithms. This effect has been discussed in Section 4.4.2. The overhead increases with the speeds of the cars: the higher the desired speed, the larger are the cars' distances. Larger distances cause more losses and, therefore, more repetitions are necessary.

Figure 4.8 shows the results of the simulations with a target interval of 0.1 s. The fixed interval beaconing and the randomized interval beaconing show a very similar behavior as with the target interval of 1 s. The shorter interval, however, leads to much more collisions as indicated by the boxes' longer upper parts of random interval beaconing. With fixed interval beaconing, repeated collisions happen in the majority of the simulations, creating interval differences of up to 495 s.

Again, Carrot performs much better, being nearly unaffected by the shortened target interval. The maximum interval differences are about the size of one target interval. An obvious deviation compared to the other results occurred at the desired speed of 5 m/s where two things are remarkable: the 99 % quantile is much higher than at the other runs with Carrot and the number of sent beacons is too small for the target interval (the expected number of beacons is given by fixed interval beaconing). This is caused by too many cars in reception range than Carrot can handle with the target interval: the cars are too close together at this speed for Carrot's minimum time between beacons of 1 ms. Consequently, the leader enlarges the sending interval and sends less often.

Concluding the results, the average sending intervals are nearly the same at all three algorithms. The maximum interval difference of the runs with Carrot is much less than those of the runs with fixed interval beaconing and randomized interval beaconing. Carrot sends a few more beacons than the fixed interval beaconing at higher speeds to maintain a steady sending interval.

4.6 Conclusion

In this chapter, a communication protocol for a CACC system was developed according to the top-down approach and analyzed formally as well as with the help of simulations. The application-level model of a single lane with a platoon of cars is an extension of the basic model created in Chapter 3. The intention to develop an algorithm for this application was to obtain design guidelines and understand the limits of beaconing. Through stating goals for the formal model and creating an optimal behavior, it was shown how to apply the top-down approach to a specific application in a setting with multiple cars. The resulting behavior enables cars to drive accident absent while minimizing the platoon's road usage. In a scenario with reliable communication, it was found for both a road-usage efficient behavior as well as a more realistic constant-speed following that beaconing with a fixed sending interval fits well to the application.

The effects of beacon losses on the road usage, however, are significant when using fixed interval beaconing. So it was discussed how the impact of losses can be decreased. Neither fixed nor randomized interval beaconing are able to handle losses. To allow for a steady state at least for a predefined number of consecutive losses, the following distances have to be increased by multiples. This in turn causes a decrease of the road usage. The adaptive beaconing algorithm Carrot was proposed that enables closer following distances. It structures the sending of beacons and, by this, enables an implicit detection of beacon losses. Carrot's fast repetition mechanism mitigates the effects of losses significantly; this was shown with formal reasoning and simulations. Carrot works in a distributed manner, needs only a couple of variables, and has a low computation overhead. The communication overhead through the repetition of beacons can be balanced with enlarging the sending interval. Even with such a longer interval, Carrot is preferable over fixed interval beaconing in usual application configurations.

The algorithm's name is inspired by the optimal car behavior in the formal model: like donkeys that stop moving unless a carrot is waved ahead of them, Carrot sends beacons to successive cars to make them speed up.

The next chapter extends the formal model further, now regarding multiple lanes. A distributed merging algorithm will be proposed that supports coordinating the merging order of cars from different lanes.

Chapter 5

Cooperative Lane Merging

The road topology of the scenarios discussed so far was as simple as possible: they involve only a single lane. In such a topology, it was shown for a cruise control application that there is no need for a car to interact with others than its direct predecessor and successor. Inter-vehicle communication, however, allows for exchanging information with cars further away than the direct neighbors, too. In this chapter, which is based on the publication [BKS+13] of which I was the main contributing author, it is made use of the ability to communicate with cars not in line-of-sight. The topology is extended for this to a merging of two lanes: this is still a basic and very common traffic scenario, although of higher complexity than just a single lane. As discussed in Chapter 2, lane merging assistants that employ communication have been proposed in related work and experimental studies showed that such systems are applicable to real cars. But in contrast to previous works will the merging assistant be developed top-down, starting with a formal modeling of the setting.

The formal model for a cruise control application created in the previous chapter will be enhanced in the following to let cars merge with the maximum flow possible. Defining behavior for a high flow, however, does not determine the order in which cars have to pass the merging: there are situations in which two cars are about to arrive at the bottleneck at (nearly) equal times and need to decide on who should drive first. A higher-order criterion will be introduced to decide on the merge order before the acceleration behavior for a maximum flow is calculated. This enables cars to calculate their time to merge first and adjust their order of arrival accordingly. The behavior to approach the merging is developed afterwards with regard to the merge order.

The criterion used for the merge order is *fairness*. Towards this, the question of how fairness of a merging process can be defined will be discussed. The notion of free-flow fairness will be introduced based on the waiting time a car is expected to suffer until it merges eventually. The waiting time is obtained by a reference time that describes the earliest time a car could be at the merging if it was able to drive freely in its lane.

This is referred to as the free-flow arrival time. This definition can already tell if a merging order is perfectly fair. To get more detailed information on how a certain merging order performs, a metric for unfairness is proposed. This will then be applied to the zipper merge order which drivers are required to follow by traffic rules when no lane's cars have the right of way. With zipper merge, each next-but-one merger originates from the same lane. If the flows of cars on the incoming lanes differ at least a bit and are high enough such that cars have to coordinate their merge order, the zipper merge creates grossly unfair results. Cars in the lane with the higher flow suffer a drawback by merging too late, while cars from the other lane merge earlier than at their fair positions. For congested situations, it will be shown that the unfairness grows unboundedly.

Despite the unfairness caused by traffic rules, optimal fairness is achievable; to show this, an algorithm will be proposed based on the differences of the cars' free-flow arrival times. The algorithm assumes each car to be omniscient and cooperative in view of following the optimal strategy. A more realistic setting will be considered thereafter, in which merely a subset of vehicles participates and wireless communication is unreliable. A distributed coordination algorithm will be proposed for this setting that enables cars to find a merge order with beacon-based communication. A formal analysis shows that it yields limited unfairness as soon as it is used by a fraction of cars. This is supported by simulative evaluations with varying participation ratios as well as varying incoming flows. The simulations demonstrate the advantage the algorithm brings when jams begin to emerge even if only 1 % of the cars participate.

This chapter is structured as follows. Section 5.1 describes the terms of the formal model and Section 5.2 specifies the fairness criterion. The fairness of common merge strategies is discussed in Section 5.3. In Section 5.4, it is shown how vehicles should behave ideally in order to reach a fair merging order. Afterwards, the cars' acceleration behavior to approach the merging is explained. It is then turned to a more realistic setting in Section 5.5, in which the most abstract assumptions are dropped such as omniscience of the involved cars. A merging algorithm for this setting is proposed that coordinates the cars by means of unreliable beaconing. The algorithm has been implemented using the simulator ns; the simulation results are described in Section 5.6. The findings on communication scheme design for merging scenarios are concluded in Section 5.7.

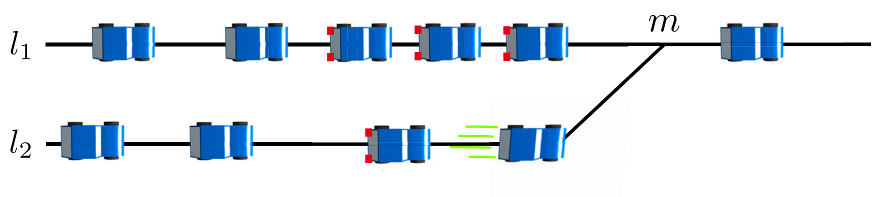

Figure 5.1: A sketch of a merging of the two lanes l_1 and l_2.

5.1 Model of a Merging Scenario

In order to analyze car behavior in a scenario in which two lanes merge into one, the formal model that was created in the previous chapter will be used as starting point for the following discussion. The model is extended with the intention to focus on the merge order while keeping it as simple as possible. The described topology consists of a main lane l_1 and an on-ramp lane l_2 which merges into the main lane. The lanes both start at a point 0. The on-ramp lane l_2 ends at a point m which is called the *merge point*, while l_1 continues infinitely. The longitudinal movement of cars is restricted by the lanes' maximum speed w'_{max} and a minimum speed of zero. The scenario is described by the following definition and depicted by Figure 5.1.

Definition 12. Two lanes with merging. $m \in \mathbb{R} > 0$. *The main lane* l_1 *is an interval* $[0, \infty)$. *The on-ramp* l_2 *is an interval* $[0, m)$. *Cars start at point* 0 *in one of the lanes.* w'_{max} *is the maximum speed allowed in both lanes.*

The on-ramp ends directly in the main lane at point m, so a car stopping at this point (or further ahead) is blocking cars in both lanes upstream of the merging. The merge point m equals the point to merge latest on a "real" merging: it is the point at which merging cars begin to leave the on-ramp and "touch" the main lane. The lanes, as well as cars driving in the lanes, are one-dimensional. A car c_i has a length L_i and drives in a distinct lane l_i.

Definition 13. Car with length. $i \in \mathbb{N}$*: A car* c_i *is a six-tuple* $(L_i, A_i, D_i, t_i^0, w'_{max}, l_i)$*:*

- $L_i > 0$ *is the length of the car .*
- $A_i \geq 0$ *is the maximum acceleration.*
- $D_i \leq 0$ *is the minimum acceleration.*
- t_i^0 *is the point in time at which the car appears.*
- $w'_{max} \geq 0$ *is the initial speed at* t_i^0*.*
- l_i *is the car's initial lane.*

The position of a car is denoted by the way $w_i(t \geq t_i^0)$. *This function describes the head of the car. The rear is at* $w_i(t \geq t_i^0) - L_i$. $i, j \in \mathbb{N}, i \neq j$. C *is a set of cars in a lane with a total order on the set, so that* $\forall c_i, c_j \in C : t_i^0 < t_j^0 \Rightarrow c_j <_c c_i$.

The simplification of each car having the same acceleration capability is adopted from the previously discussed models. All cars start with maximum speed and each car has the same length. A car in the on-ramp changes to the main lane with an instant transfer of its respective part touching the lane's end; this part appears in the main lane at m. A way for a car is valid, in addition to the rules of Definition 6 in Chapter 3, if its speed never exceeds the maximum speed of the car's current lane.

Definition 14. Validity of a way (with maximum speed). *Let* $i \in \mathbb{N}$. *A way* w *is said to be valid for car* c_i *in lane* l_i *with maximum speed* $w'_{\max}$ *if and only if*

$$w(t_i^0) = 0 \wedge \forall\, t \geq t_i^0 : D \leq w''(t) \leq A \wedge w'_i(t) \leq w'_{\max}.$$

$\mathbb{W}_i$ *is the set of all valid ways for a given car* c_i.

To keep track of the cars' initial order, the set of cars in a lane is considered to be a totally ordered set through the cars' times of appearance. Car c_i (for $i \in \mathbb{N}$) appears in one of the lanes at t_i^0. Only one car appears at the same time in the same lane. Cars are required to drive on accident-absent ways like in the previous chapters but the definition differs from Definition 7 in Chapter 3 because of the cars' lengths.

Definition 15. Accident absence (with car length). *Given two cars* c_1, c_2 *in the same lane with* $c_2 <_c c_1$. *Two valid ways* w_1 *and* w_2 *for* c_1 *and* c_2*, respectively, are said to be accident absent, if and only if* $\forall t \geq t_2^0 : w_1(t) \geq w_2(t) + L_1$.

Towards accident absence, a safety distance Δw_{safe} has to be maintained between two consecutive cars which is calculated as the difference of their minimum braking distances to a full stop. The existence of accident-absent ways with regard to this distance is described in Lemma 12 on Page 130.

The flow of cars per second is limited by several parameters of the model; the *limiting flow* Q at the merging is $Q = w'_{\max}/(L + \Delta w_{\text{safe}})$. Here, Q is defined with maximum speed because the model's safety distance is zero at equal speeds, so cars merge the closest when driving at the same speed. It follows directly that the flow is maximized at maximum speed. For merging behavior, this means that cars have to pass the merging at maximum speed and implies that each car needs the same time to pass the merging. The *maximum flow* $q_{\max}$ at the merging is an additional objective and defined as follows.

Definition 16. Maximum flow at the merging. *Given the flows q_1 and q_2 in the incoming part of lane l_1 and in lane l_2 as well as a limiting flow at the merging of $Q \geq 0$. The maximum flow at the merging is*

$$q_{\max} = \min \{q_1 + q_2, Q\} .$$

If the sum of the two in-flows is higher than the maximum flow at the merging, then cars have to wait for passing the merging and a jam emerges at least on one of the lanes.

5.2 Merging Order Fairness

In this section, the objective of fairness is introduced and it is discussed how fairness can be measured. Towards this, a key observation is that cars leave the merging in a certain order. Each car c_i has a position index $k_i \in \mathbb{N}$ in the sequence of cars leaving the merging. The fairness criterion refers to this sequence. The sequence is determined by the traffic scenario and the merging scheme. The unfairness metric that will be created describes the distance of a merging scheme's positions sequence to a perfectly fair order.

5.2.1 Free-Flow Fairness

To enable a fair merging order, it has to be discussed what is meant by "fair". Consider two cars that approach the merging and need to decide which one drives first. An intuitive approach would be to make this decision based on the cars' distances to the merging. This is not a good solution, though: depending on the length of the queue in each lane before the merging, two cars may be equally far away from the merging, while one of them has been waiting much longer than the other. It would be unfair to let a car with shorter waiting time pass first. The decision should therefore be based on the time that the cars spend waiting: a car with longer waiting time should be given preference.

But how to measure the waiting time of a car at a merging? Or, more specifically, when does a car start waiting? It will be argued that the waiting time of a car should be measured starting from the point in time at which a car would have arrived at the merging if it were not hindered by any other car. This is the earliest point in time at which a given car could possibly arrive at the merging; this point in time is termed the *free-flow arrival time*. The calculation of the free-flow arrival time is described in the following lemma.

Lemma 2. *The free-flow arrival time of car c_i, with an initial speed of $w'_{\max}$, at the merging at waypoint m is*

$$t_i = t_i^0 + \frac{m}{w'_{\max}}.$$

Proof. Car c_i enters one of the lanes at position 0 at time t_i^0 and starts at a speed of $w'_{\max}$. This is the maximum speed allowed in the lane, it hence keeps on driving at that speed. The time it takes to reach waypoint m from position 0 is $m/w'_{\max}$. □

A fair merge order based on free-flow arrival times lets a car pass earlier if it has an earlier free-flow arrival time. The following definition of fairness is based on this concept.

Definition 17. Free-flow fairness. *Given are two cars c_1, c_2 with free-flow arrival times of $\tilde{t}_1$ and $\tilde{t}_2$ and merging positions of k_1 and k_2, respectively. Without loss of generality let $\tilde{t}_1 < \tilde{t}_2$. The merge order is called fair if and only if $k_1 < k_2$.*

In other words, a merging order is considered fair if that car passes the merging first which could arrive there first in a free lane. Free-flow fairness can be seen as a form of first-come-first-served (FCFS) fairness which is also referred to as temporal fairness [Wie11]. Such a fairness definition fits well if a resource is used for the same amount of time by each consumer. This is considered a reasonable approximation in the case of a lane merging: the time it takes a vehicle to traverse the merging is always the same in the model. In reality, the duration will vary only slightly with the vehicle type. Fairness through FCFS is also employed at aviation [BBCF12].

With this concept of fairness, another reason to model the merging as a single point becomes obvious. A scenario with one point for merging has the same fair merge order like a scenario with multiple merge points. Multiple merge points, in the first place, enable cars starting in the same lane to overtake each other. They thereby merge in another order than in which they appeared, i.e., they do not merge in the order of ascending free-flow arrival times. This is less fair than an order without overtaking. The exclusion of orders requiring overtakings reduces the possible orders to the same as with a single merge point. Cars can adjust their relative positions in a scenario with just one merge then like in a scenario with multiple merge points, even though without merging yet: the merging is eventually performed at the merge point.

Fairness in network protocol design

While the concept of fairness is new to C2CC application design, it is common practice in network protocol design. Usual notions of fairness in that field are max-min fairness and proportional fairness. Both treat network traffic at the level of flows instead of distinct units (i.e., packets in this case). The proposed free-flow fairness for merging differs significantly from these fairness concepts in that it does not consider flows but individual cars.

Max-min fairness uses a water-filling approach in which the allocation of a resource, e.g., bandwidth, is equally increased stepwise at each consumer until one consumer is satisfied. The remaining resources then are shared among the other consumers by again increasing their allocations. This is repeated until either all consumers are satisfied or no resources are left. The mechanism maximizes the minimum share each consumer gets. Max-min fairness is applied to transmission power in VANETs to guarantee cars a minimum share of bandwidth for safety-related messages (in [TMSH05, TMSH06]).

If a max-min fair allocation was applied to a merging scenario, the flows on the incoming lanes would have to be considered. The allocation of the merging capacity would be made by water-filling it according to the flows. Should both flows be higher than half of the limiting flow of the merge, the flows were assigned an equal share of the capacity. A max-min fair approach does not regard waiting times of individual cars and hence does not result in an order which is fair in the sense of free-flow fairness.

Proportional fairness as described in [KMT98] employs a utility function to define fairness. A fair allocation maximizes the sum of utilities of all consumers. The utility function is used to model costs of a consumer for transmitting an amount of data per time unit, i.e., at a certain data rate. The bandwidth is shared proportionally to the consumers' utilities. TCP, the Transmission Control Protocol belonging to the Internet protocols (as discussed in Section 2.3.1), uses mechanisms to fair the usage of bandwidth in a manner similar to a proportional fairness criterion [KMT98].

When proportional fairness was implemented for a merging scenario, with incoming flows instead of data rates, the flows shared the bottleneck proportionally to the merging costs. To obtain a throughput share that is proportional to the flows, the costs needed to be related to delays of cars in the flows such that a higher flow was assigned a proportionally higher throughput than a lower flow. To enable a proportionally fair order, individual cars thus have to know about the flows on the lanes.

In contrast, free-flow fairness uses the free-flow arrival times to create a merge order. The arrival times are directly related to the flows: cars appear in the lanes according to the flows and the latest car that appears gets the latest free-flow arrival time compared to other cars already in the lanes. The merge order is proportional to the flows without the need to know the flows.

5.2.2 Measuring Unfairness

Based on the notion of free-flow fairness, the unfairness of a specific merge order in a given traffic scenario can be measured. To this end, the *position difference* is considered: the difference of a car's true merging position k_i to the position $\tilde{k}_i \in \mathbb{N}$ the car would have with a merging order given by free-flow fairness: $k_i - \tilde{k}_i$. For a group G of $n = |G|$ cars, the total *unfairness* within this group can be defined as

$$u := \sum_{i \in G} \left(k_i - \tilde{k}_i\right)^2,$$

the average unfairness $\overline{u}$ is given by

$$\overline{u} := u/n.$$

Squared differences are used in order to avoid problems with varying signs in the differences (cars merging too early versus cars merging too late), and to penalize large deviations from the fair order more strongly. The squaring allows minimizing the differences in the spirit of a least squares optimization.

It is assumed for a moment that the merging has been unused for a short time span because the queues at both incoming lanes were empty. It can be observed that in this case, no position interchanges between cars merging before this "gap" and cars merging afterwards can have occurred, according to the definition of free-flow fairness. Otherwise, one car would have arrived at the merging before the gap, but merged after it, which contradicts the assumption that both queues were empty. This constraint on permissible permutations simplify certain computations: in particular, if such gaps occur, the groups of cars can be considered between any pair of gaps independently. The unfairness of consecutive groups according to the above definition is summable.

5.3 Fairness of Merging Schemes

Two types of merging schemes are defined by traffic rules and for both, specific rules for merging exist. One type occurs if a lane ends in a major lane that is assigned the right of way. Cars in the major lane are allowed to traverse the merging without paying attention to cars in the ending lane. Cars in the minor lane are required to wait for a gap in which they fit safely. The second merging type is a merging of roads of equal order. The cars then have to do a zipper merge, a round robin scheme in which each next merger originates from another lane in circular order. This section discusses how the two merging schemes perform regarding free-flow fairness.

5.3.1 The Zipper Merge

It will be investigated how well zipper merge performs with respect to free-flow fairness as defined above. In particular, the long-term behavior of the average unfairness $\overline{u}$ is discussed. To this end, the traffic flows on the incoming lanes are considered. The average traffic flows on lanes l_1 and l_2 are denoted by q_1 and q_2, respectively.

For a first impression on how unfairness is discussed, very small traffic flows on both incoming lanes are assumed. Cars are so sparsely distributed that they do not influence each other while merging. Evidently, the cars will merge in the order given by their free-flow arrival times. The unfairness metric will therefore be zero.

Looking at fairness becomes more interesting when the flows are higher, so that cars start to interact and some cars must wait in order to merge. To assess such situations analytically, exponentially distributed interarrival times between cars are assumed, as also done in related work, e.g., [Mil61, RML02]. Two cases will be distinguished: the total incoming traffic flow in both lanes together may be higher or lower than the merging's limiting flow. In general, the prime interest is in cases where the incoming flows differ at least minimally (i. e., $q_1 \neq q_2$).

If $q_1 + q_2$ exceeds the limit Q, the unfairness grows to infinity if zipper merge is used. This can be seen as follows: with zipper merge, the same number of cars from each incoming lane will merge per time unit, as long as cars ready to merge are available on both lanes. Because the total traffic flow exceeds the limiting flow Q, an increasing backlog will build up on at least one lane (the one with the higher traffic flow). In the other lane, a backlog may or may not build up, depending on whether the traffic flow there exceeds $Q/2$; since traffic flows are assumed to be not equal, the backlog will at least be smaller, though, and the difference in length between the backlogs will become larger and larger over time. Cars arriving in the lane with lower traffic flow will therefore be able to pass more and more cars with earlier free-flow arrival times in the other lane. Consequently, position differences in the sense introduced above increase more and more, so that the average unfairness increases without limit.

The case where the total traffic flow on both incoming lanes does not exceed the limiting flow of the merging is slightly more difficult to understand. It will be shown that the long-term average unfairness gets higher and higher, the closer the total incoming flow $q = q_1 + q_2$ approaches the limiting flow Q. To obtain this result, queuing theory is used. The waiting in front of the merging is modeled with an M/D/1 queue. This queuing model describes exponential interarrival times and a constant merging time for each car at the single merge.

The so-called utilization rate of the merging is $\rho = q/Q$. Queuing theory gives us the expected number of cars waiting in front of the merging as

$$\rho^2/(2(1-\rho)),$$

see, e.g., [Coo81]. This is a total number for both lanes. At the limit of the utilization rate the expected number of waiting cars is infinite:

$$\lim_{\rho \to 1} \rho^2/(2(1-\rho)) = \infty.$$

For $\rho < 1$, the queues in front of the merging will, from time to time, be both empty. The expected portion of cars that see an empty queue is $1 - \rho$. These cars are group leaders. As argued above, the groups of cars between such gaps can be considered independently. The expected size for each of these groups is equal and finite. Since the first car of each such group will not be influenced by any other car, it merges at its free-flow arrival time. It therefore has a position difference of zero. The merging order of the other cars within the group will be a permutation of their fair merging order. For each other car of a group, the absolute value of the position difference will thus be smaller than the group size. Therefore, the average unfairness within each group is bounded above by the square of the group size. As a result, the long-term average unfairness is limited.

In general, the unfairness will be worse for higher incoming traffic flows and for higher differences between the two incoming flows. Moreover, the larger the difference between the traffic flows, the more quickly unfairness will build up.

To complement these analytical considerations, the zipper merging scheme has been simulated with different incoming flows in a simulation environment based on ns. The details of the simulation and the results are described in Section 5.6.

For exactly equal flows, zipper merge will be fair even over long time spans. However, equal flows are very unlikely to occur in practice—and even minimal deviations will lead to arbitrarily large unfairness over sufficiently long time spans. This raises the question of how we can leverage communication between cars to improve the merging fairness—and which fraction of cars needs to participate in order to yield good results. This question is discussed in the following sections of this chapter. But before this, fairness of a right-of-way scheme is discussed.

5.3.2 Right of Way

Cars in a lane with right of way are allowed to drive through the merging area without paying attention to the other lane's cars. Cars in the minor lane, on the other hand, have to wait for a gap.

In case of a now and then empty queue due to $\rho < 1$, the same arguments as for zipper merge apply, so that the long-term average unfairness is limited. The special case of exactly equal flows, however, does not result in a fair merge order. This becomes apparent by noticing that cars do not merge in an alternating order if a backlog on the minor lane emerges. A right-of-way scheme shows unfair behavior especially in case of the flow in the major lane being constantly higher than $Q/2$. Then no car from the minor lane is able to merge—starvation occurs and the average unfairness grows to infinity.

With $\rho \geq 1$, the right of way results in an unlimited long-term average unfairness for the same reasons as discussed with the zipper merge. So this scheme is not considered any further in the remainder of this chapter.

5.4 An Optimal Merging Scheme

In this section, it is discussed what the optimal behavior of cars regarding fairness has to be. An algorithm is specified that determines the optimal order with the assumption of omniscient cars that participate willingly. The omniscience applies to information regarding other cars including their future behavior. Moreover, the cars follow the given algorithm perfectly without any deviation or delay. The car behavior to get to the merging at the right time is developed with the two objectives accident absence and road-usage efficiency.

The discussion begins with a simple scenario in which only a single car drives in each lane. Then the complexity of the model is increased one step: multiple cars drive in one of the lanes and a single car has to merge between them at a fair position. Eventually, it is described how multiple cars that drive in both lanes have to merge.

One car in each lane

Let car c_1 drive initially in lane l_1 and car c_2 in lane l_2. The merge order of the cars regarding fairness is evident: that car has to traverse the merging first which is able to arrive there the earliest according to the cars' free-flow arrival times. Without loss of generality it is assumed that this car is c_1, so c_2 must merge second. The cars' order is obviously fair because it is in line with the definition of free-flow fairness.

Multiple cars in one of the lanes

The scenario is extended to a single car that has to merge between multiple cars driving in the other lane. Once again, the optimal behavior of cars regarding fairness is specified. Car c_1 drives initially in lane l_1 and $c_2, \ldots, c_n$ with $n \in \mathbb{N}$ in lane l_2. The correct two cars in l_2 have to be found by c_1 to merge between. The free-flow arrival times of all cars in l_2 form the list

$$T_2 = (t_2, \ldots, t_n).$$

The free-flow arrival time of c_1 and the list are put into a new list and that list's elements are ordered ascendingly. Then the ordered list describes by the order of free-flow arrival times between which two cars c_1 has to merge. The list's free-flow arrival times identify the cars uniquely except if two cars from different lanes have the same arrival time. In this case, a tie break needs to be applied: the car in l_1 is preferred. Cars in l_2 with earlier free-flow arrival times than c_1 merge without adaption. Then c_1 merges; cars with free-flow arrival times later than c_1 merge behind it in their order of appearance.

Multiple cars in both lanes

Finally, multiple cars driving in both lanes are considered. Let $m, n \in \mathbb{N}$: cars $c_1, \ldots, c_m$ drive initially in l_1, while $c_{m+1}, \ldots, c_{m+n}$ drive in l_2. In both lanes l_1 and l_2, the cars' free-flow arrival times are described by the lists T_1 and T_2, respectively. The lists are merged to a new list with an ascending order of the new list's elements. The order of free-flow arrival times in the list defines the merge order. Again, the tie break rule is applied if necessary.

5.4.1 Approaching the Merging

Now that a fair merge order can be determined, the cars' acceleration behavior is calculated such that they arrive at the merging with this order. In view of the two objectives accident absence and road-usage efficiency, the behavior for merging and for following are now separately discussed. The behavior for merging is specified first; it results a solution in which the merging is traversed the fastest and the shortest, so that upstream cars in both lanes are influenced minimally. Then a maximum road-usage efficient behavior is discussed which is similar to the behavior discussed for the platooning application but satisfies the constraints regarding merging time and merging speed. This means that optimizing the behavior at the merging is given a higher priority than the behavior in the distinct lanes. The formal model's optimal behavior for merging is specified by the objective of maximum merging efficiency:

Definition 18. Maximum merging efficiency. *Given a car c_i in lane l_i with free-flow arrival time t_i. Let car c_j be its predecessor with regard to free-flow fairness that merges at $\tilde{t}_j$. c_i can merge at $\tilde{t}_i$ at a minimum steady state distance to c_j with speed $w'_{\max}$. The behavior of c_i is said to be maximum merging efficient if and only if it merges at the maximum of t_i and $\tilde{t}_i$. The arrival at the merging is delayed with a traveling-optimal way.*

Merging efficiency demands to create a delay for merging such that the resulting way is traveling optimal, i.e., delaying is performed the latest possible. The behavior towards maximizing merging efficiency and for road-usage efficiency for constant-speed following introduced in Chapter 4 are now combined. At time t, a car c_i calculates its acceleration for merging efficiency to be $a_m(t)$ and for constant-speed following to be $a_f(t)$. The resulting acceleration is

$$a(t) = \min\{a_m(t), a_f(t)\}.$$

Choosing the minimum of both accelerations guarantees accident absence because if it is required to brake for following a car or for merging, $a(t)$ describes this. If no action for merging is necessary and driving with constant speed is possible, then $a(t)$ is zero.

The resulting behavior is sketched in the following. Cars start with a maximum speed at distances bounded by the minimum distance for constant-speed following. A car stays at this speed unless actions are required for merging efficiency upstream of the merging. These are always indicated by a deviation from the maximum speed through braking. The merging acceleration then overrides the steady state following acceleration.

In the remainder of this section, the acceleration $a_m(t)$ for a maximum merging-efficient behavior is specified. First, the cars' true arrival times, denoted by $\tilde{t}$, are determined. With this, it can be evaluated whether an adaption for merging is necessary; this is the case if the true arrival time is later than the free-flow arrival time.

The discussion of finding the true arrival times is started at a car c_i that is "first", i.e., it merges without a change in behavior. This requires that the merging is unused at its free-flow arrival time t_i as described in Section 5.2.2. c_i blocks the merging in the interval $[t_i, t_i + L/w'_{\max})$. Now, the following has to be repeated until all cars have been assigned a true arrival time: from the sorted list of free-flow arrival times of both lanes, choose the earliest free-flow arrival time belonging to car c_j that does not yet have a true arrival time (and prefer cars from l_1 in case of equality). Adapt its arrival

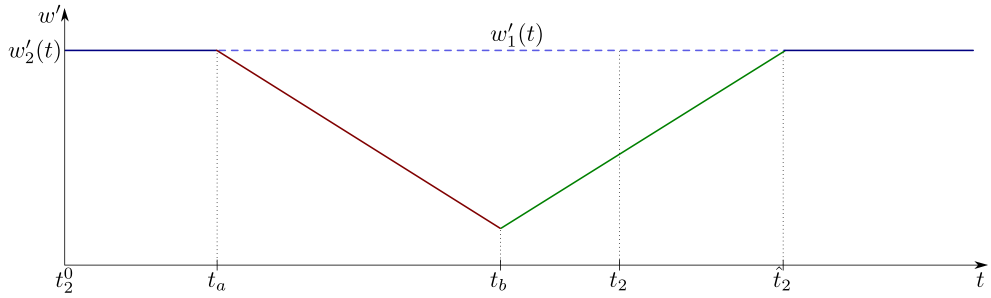

Figure 5.2: Speed of car c_2 that delays its arrival at the merging from its free-flow arrival time t_2 to the true arrival time $\tilde{t}_2$, just when car c_1 left the merging.

time to be at $\max\{t_j, t_{i,\text{free}}\}$ with $t_{i,\text{free}} = t_i + L/w'_{\max}$. At this time, c_i just left the merging. Since c_j has complete knowledge about the behavior of c_i, no additional safety distance required. If $t_j < t_{i,\text{free}}$, c_j has to delay its arrival such that it arrives exactly at $t_{i,\text{free}}$ with a speed of $w'_{\max}$. The cars' order, their distance, and c_j's speed are optimal regarding a merging efficient behavior. Next, replace c_i by c_j and c_j by the next car in the list of free-flow arrival times and repeat.

With the true arrival times assigned, all cars know when to be at the merging for merging efficiency. It is discussed next how a car has to behave to arrive just in time. A car that merges truly at its free-flow arrival time constantly keeps its acceleration $a_m(t)$ at zero: it enters its lane at maximum speed and stays at this speed all the time. If the true arrival time is later than the free-flow arrival time, the arrival is delayed with a deceleration phase followed by an acceleration phase to avoid a crash with the preceding car at the merging. Before delaying begins, a car stays at maximum speed the longest possible: the way has to be traveling optimal with regard to the validity of a way that was redefined in this chapter. Figure 5.2 illustrates this with the speeds $w'_1(t)$ and $w'_2(t)$ of the two cars c_1 and c_2. The latter has to delay merging to $\tilde{t}_2$ to merge directly behind c_1. The following lemma discusses how to obtain the times at which the acceleration for merging efficiency changes.

Lemma 3. *Given a car c_i with the initial speed of $w'_{\max}$ and a free-flow arrival time of t_i. Let the true arrival time at the merging point m be $\tilde{t}_i \geq t_i$. It exists an optimal way for c_i regarding merging efficiency in which it brakes from t_a to t_b, waits between t_b and t_c, and accelerates from t_c to $\tilde{t}_i$. It drives at constant speed from $\tilde{t}_i$ onwards.*

$$\textit{Here, } t_c = \tilde{t}_i - \min\left\{\left(\left(w'_{\max}(\tilde{t}_i - t_i^0) - m\right)/A\right)^{1/2}, w'_{\max}/A\right\},$$
$$t_b = t_c - \max\left\{0, \tilde{t}_i - \left(t_i^0 + m/w'_{\max} + w'_{\max}/A\right)\right\}, \textit{ and}$$
$$t_a = t_b + t_c - \tilde{t}_i.$$

The proof to this lemma is shown in Appendix A.6. It is possible with the formulae stated in the lemma that a car c_i decelerates to a full stop. Consecutive cars behind a stopping car may also have to stop, depending on the initial distances between the cars. The acceleration behavior for constant-speed following ensures stopping right behind a fully stopped car without additional gaps. If a car c_i stopped for merging efficiency and not for constant-speed following, it stopped such that it is able to accelerate constantly until the merging is reached exactly with maximum speed. But the next car waiting behind, let this be c_j, stopped at the rear bumper of c_i. That means c_j cannot accelerate constantly; it is one car length away from the perfect point. As soon as c_i accelerates, c_j also does because of acceleration behavior rules. If the merge order allows c_j to merge directly behind c_i, it travels the fastest way possible by accelerating to maximum speed and then keeping the speed until merging bumper-to-bumper with c_i. Otherwise, if merging directly behind c_i is not possible, c_j has to recalculate its behavior. In such a case, Lemma 3 is not sufficient because c_j may need to decelerate during acceleration already. A behavior that is more flexible in that it allows braking while not driving at maximum speed is discussed in Lemma 4.

Lemma 4. *Given a car c_i with a free-flow arrival time of t_i. Let the true arrival time at the merge point m be $\tilde{t}_i \geq t_i$, and let c_i have a safe distance to its predecessor at the current time $t \geq t_i^0$..*

Given the points in time $t_s \leq t \leq t_a \leq t_b \leq t_c \leq t_d \leq \tilde{t}_i$: there exists an optimal way for c_i regarding merging efficiency in which it accelerates in $[t_s \leq t, t_a)$, drives with constant speed in $[t_a, t_b)$, brakes in $[t_b, t_c)$, stops in $[t_c, t_d)$, and accelerates in $[t_d, \tilde{t}_i)$. It drives at constant speed from $\tilde{t}_i$ onwards.

The proof to this lemma is shown in Appendix A.7. In the proof, a behavior for c_i is created to show the existence of an optimal way. This behavior defines an acceleration for maximum merging efficiency; if braking for safety is required before traversing the merging, the behavior needs to be refreshed. The constructed way is road-usage efficient with regard to the new predecessor after merging.

5.5 Achieving Fairness with Communication

So, how should vehicles act in order to yield good free-flow fairness in a real system? As discussed above, optimal free-flow fairness will be achieved if each car merges at a position which matches the order of free-flow arrival times. If each car participated in the system, each driver were willing to obey the system's instructions regarding the merging order, and if each car had perfect knowledge about all other cars and their free-flow arrival times, then perfect fairness could easily be achieved.

In a real system, however, the above assumptions will not hold: not all cars will participate, not all drivers will always obey, and not all information will be known (due to unreliable wireless communication, limited communication range, etc.). This raises the question of how much we lose in terms of fairness if non-equipped or non-cooperative vehicles jump the queue or merge too late. Throughout the rest of this chapter, these assumptions will be dropped one by one, and the changes of the outcome are regarded. It will be shown that this, in fact, has only a minor impact on the achieved fairness.

At first the case will be considered in which not all cars participate in the system and/or not all drivers follow its instructions. These two are equivalent: a driver not willing to obey the guidance by the system may merge either too early or too late—just like a driver who does not receive instructions from the system. The case of limited system participation is now investigated analytically. It will be proved that even with only a small fraction of participating, cooperating cars fairness is achievable. Subsequently, in the next section, simulations will be used to also investigate the case in which unreliable beaconing is used and therefore perfect knowledge is no longer available.

Consider two arbitrary (participating or non-participating) cars A and B. Assume without loss of generality that A has an earlier free-flow arrival time than B, i.e., A should merge first. If A and B are in the same lane, they will merge in the correct order because A has the earlier free-flow arrival time and must therefore be ahead of B. If both A and B participate in the system, they will merge in the correct order regardless of their lanes. If, however, A and B are in different lanes and at least one of them does not participate, then it may, in principle, happen that they switch order, i.e., that B merges earlier than A.

It may be observed, though, that A and B may no longer switch order if there are cars which participate in the system and which "force" B to wait until after A has merged. More specifically, assume that there is a participating car X on A's lane behind A, and a participating car Y on B's lane before B, such that X has an earlier free-flow arrival time than Y. Then, B cannot merge before Y, Y will not merge before X, and X cannot merge before A; therefore, A and B are forced to merge in the correct order.

The more cars enter the system with free-flow arrival times in between A and B, the more likely it becomes that participating vehicles X and Y exist which fulfill the above conditions. Consequently, grossly unfair situations, where cars switch order that are "far apart" with respect to the global fair merge order, are very unlikely. Along the lines of these considerations, it is aimed to prove that high unfairness becomes very unlikely and that the expected unfairness is finite.

To this end, assume that a fraction r (with $0 < r < 1$) of cars participates in the system. Furthermore, let $\tilde{k}_A$ and $\tilde{k}_B$ be the merging positions of A and B, respectively, according to the globally fair merge order. Consequently, there are

$$d = \tilde{k}_B - \tilde{k}_A - 1$$

cars with a free-flow arrival time between those of A and B.

Now, the probability p that it is possible for A and B to merge in wrong order, depending on the number of cars d between A and B, is of interest. Those cars may be arbitrarily distributed to lanes l_1 and l_2. If the traffic flow is modeled like in Section 5.3 and traffic flows of q_1 and q_2 in the two lanes are assumed, an upper bound for p can be derived.

Recall that it is not possible for A and B to switch order if participating cars with the roles of X and Y above exist. Let $\delta = \lfloor d/2 \rfloor$, and assume without loss of generality that A's lane is l_1 with average inflow q_1. The probability that these cars exist is then bounded below by the probability that there is

1. at least one out of the first δ cars between A and B that is in A's lane l_1 and participates in the system (cf. X) and

2. at least one out of the last δ cars between A and B that is in B's lane l_2 and participates in the system (cf. Y).

The probabilities for these two conditions to hold are stochastically independent. They are given by

$$p_X = 1 - \left(1 - \frac{q_1}{q_1 + q_2} \cdot r\right)^{\delta}$$

and

$$p_Y = 1 - \left(1 - \frac{q_2}{q_1 + q_2} \cdot r\right)^{\delta},$$

respectively. If both conditions hold, then A and B cannot possibly merge in the wrong order. As a result, the probability p is bounded above by

$$p \le 1 - p_X \cdot p_Y.$$

With $q_{\min} := \min\{q_1, q_2\}$ it is easy to derive that

$$\begin{aligned}
p &\le 1 - \left(1 - \left(1 - \frac{q_{\min}}{q_1 + q_2} \cdot r\right)^{\delta}\right) \cdot \left(1 - \left(1 - \frac{q_{\min}}{q_1 + q_2} \cdot r\right)^{\delta}\right) \\
&\le 2 \cdot \left(1 - \frac{q_{\min}}{q_1 + q_2} \cdot r\right)^{\delta} = 2 \cdot \left(1 - \frac{q_{\min}}{q_1 + q_2} \cdot r\right)^{\lfloor d/2 \rfloor}.
\end{aligned}$$

Therefore, the probability that it is possible for A and B to switch order decreases exponentially with increasing "distance" d between them.

Now, what does this mean for the expected (un)fairness if a large number of vehicles merge? The total order in which cars actually merge is a permutation of the fair merge order. Assume that car c_i merges x positions too early (or too late) with respect to the totally fair merge order. Then, c_i's contribution to the total unfairness is x^2. In this situation, there must be a vehicle c_j which merges in the wrong order with respect to c_i and for which $|\tilde{k}_i - \tilde{k}_j| \ge x$; this follows from basic mathematical properties of permutations. The number of cars between c_i and c_j in the totally fair merge order is

$$d = |\tilde{k}_i - \tilde{k}_j| - 1 \ge x - 1.$$

Car c_i can be either x positions too early or x positions too late for its unfairness contribution to be x^2. In either of these cases, the probability that c_i switched places with the corresponding earlier or later vehicle is, by the arguments above, bounded above by

$$2 \cdot \left(1 - \frac{q_{\min}}{q_1 + q_2} \cdot r\right)^{\lfloor (x-1)/2 \rfloor}.$$

This is the probability to switch places for the case that c_j is x positions away. Before switching with c_j, c_i needs to switch with all vehicles in the other lane up to a distance of $x - 1$ positions. Is c_j further away, the probability to switch is lower because the number of vehicles c_i has to switch with is larger. So this equation bounds the probability for c_i to switch positions with any vehicle c_j that is x or more positions away regarding the fair order. To obtain an (admittedly coarse, yet sufficient) bound for the probability to be off by x positions *either* too early *or* too late, this expression is doubled:

$$p_x \leq 4 \cdot \left(1 - \frac{q_{\min}}{q_1 + q_2} \cdot r\right)^{\lfloor (x-1)/2 \rfloor}.$$

As a result, the expected contribution of c_i to the total unfairness is bounded above by

$$E[u_i] \leq \sum_{x=1}^{\infty} x^2 \cdot p_x.$$

As may be seen by d'Alambert's ratio test, this series converges, i.e., there is a finite upper bound. This upper bound for the unfairness contribution does neither depend on properties of the specific car nor on the total number of cars. Since the expected unfairness contribution of any car is bounded above by the same finite value, the expected average unfairness with the merging scheme is also finite—even for small fractions r of participating vehicles and for arbitrarily unbalanced flows q_1, q_2.

For zipper merge, as discussed in the previous section, the unfairness gets arbitrarily high even for slightly unbalanced inflows in the two lanes. When merging is supported by inter-vehicle communication in the way outlined here, in contrast, all vehicles—including non-participating ones—are kept from merging at grossly unfair positions. The resulting unfairness is finite even when only a small fraction r of cars actively participates in the system.

This is a very interesting trait of the mechanisms proposed here: they lay the foundation for one of the very few applications of inter-vehicle communication where substantial benefit results from even just a small technology penetration ratio. It will be shown that this trait does not only exist in the asymptotic limits as considered in the preceding section and here, but can likewise be shown in realistic simulation scenarios. To this end, a beacon-based algorithm is created in the following section.

5.5.1 A Beacon-based Algorithm

In the following, a distributed algorithm based on car-to-car communication is proposed. As stated above, the most simplifying assumptions of the formal model are removed: cars are not omniscient anymore, only a part of the cars is equipped with communication hardware, and equipped cars can decide whether to follow the algorithm's guidance or not.

Additional knowledge about the intended merging order of a car is distributed periodically by means of a single-hop broadcast containing at least a sender's free-flow arrival time, identifier, and lane. The assumed communication hardware is 802.11p, so the algorithm is confronted with delays, an unreliable packet transport as well as a limited communication range. Only a part of the cars is equipped with that hardware; a car is said to be a *participant* if it is equipped with radio hardware and willingly follows the guidance of the merging algorithm. To build up knowledge about other participants, a car stores received free-flow arrival times and links them with corresponding car identifiers and lanes. A participant merges only if all other participants merged that have an earlier free-flow arrival time. Until this is the case, the participant waits upstream of the merging and lets cars from the other lane pass. It is learned through beaconing if a participant traversed the merging. A timeout protects against starvation when no beacons from other cars are received anymore while waiting. The calculation of the free-flow arrival time is based on a predefined acceleration behavior starting from a distinct state of a car. This state is driving at maximum speed, before a car is first influenced by others cars' merging behaviors.

Fairness was not yet regarded in a situation in which the communication range is limited. The average unfairness is not affected fundamentally by this; it also is finite. To see this, the only interesting case to consider is the one in which the flows in the lanes are higher than the merging's limiting flow Q. The lanes merge at a common point, so cars in the last meters of both lanes are in communication range. If unfairness builds up, a car having an advantage by merging too early drives along disadvantaged cars in the part of the other lane in range. These cars have an earlier free-flow arrival time than the advantaged car but did not merge yet. It suffices to send one beacon from a disadvantaged participant to an advantaged participant to make it wait upstream the merging. Limited communication ranges are implemented in the simulations in the next section.

5.6 Simulation and Evaluation

The merging algorithm has been implemented using the network simulator ns. This section describes the evaluation of simulations about the influence of flow sizes and participation ratios towards unfairness. Before it is turned to the evaluation, the setup of the simulation environment is explained, including a description of the employed car behavior model.

5.6.1 Simulation Setup

The communication scheme and the car behavior have been implemented on top of the network simulator ns in version 3.13 and by using the Highway Mobility 2.0 framework [AW10]. This is the same environment that was described in Section 4.5.

Two different setups have been created. The first is targeted at investigating the impact of the traffic flow size. The flow sizes in the lanes are varied with a fixed ratio between the lanes' flows. The number of participating cars are set to constant, low rates. The second setup varies the number of participating cars while keeping the traffic flows constant. Both setups use the same car-following model that is described below in more detail. Cars, if equipped, communicate via an unreliable channel. The road network consists of two parallel lanes with the same length and a third lane that represents the road after the merging. The maximum speed in all lanes is 36 m/s which equals 129.6 km/h. The cars' interarrival times are exponentially distributed and the chance for a car to become a participant is randomized with a uniform distribution.

For communication, a beaconing scheme was implemented with randomized sending intervals: after sending a beacon, the next one is sent within $[1, 2]$ s. The 802.11a channel and network device implementation of ns are used. The parameters are chosen to match the IEEE 802.11p standard for sending in the control channel as described in Section 4.5.1: OFDM is used at a rate of 6 Mb/s and a bandwidth of 10 MHz, the transmission power is 33 dBm, the energy detection threshold is -71 dBm and the CCA mode 1 threshold is -91 dBm. A constant speed propagation delay and the Friis propagation loss model, linked with the Nakagami-m propagation loss model, are added to the channel. The Nakagami-m model is a probabilistic loss model that takes the output signal strength calculated by the deterministic Friis model and alters it with noise. That way, it becomes possible for nodes close to the sender to miss a packet, while nodes far away from the sender have a chance to receive it. Table 4.3 on Page 77 lists the communication model parameters.

Participants start sending beacons within a distance of 1 km to the merging and merge with regard to their free-flow arrival times if they know about each other. A participant learns via beaconing when a participant from the other lane traverses the merging. Packet losses may cause a participant to wait at the merging although all participants to wait for already merged. A waiting timeout avoids starvation in such situations. The timeout is calculated dynamically: if a car receives a message of a participant that has to merge first, it is estimated how long it will take until the sender traverses the merging in a congested setting. Cars that do not participate use the zipper merge scheme.

Car-following algorithm

To simulate car movement, the Intelligent Driver Model (IDM [THH00]) is used which were already described in the previous chapter's simulation Section 4.5. The model's

Desired speed	36 m/s
Acceleration	3.0 m/s^2
Deceleration	3.0 m/s^2
Minimum gap	2.0 m
Time headway	1.5 s
Delta exponent	4.0
Vehicle length	4 m
Simulation step	1 s

Table 5.1: Car behavior parameters of the IDM merging simulations.

parameters are chosen similarly to the simulation setup in Section 4.5.1. The desired velocity is the same in all simulations and equals the lanes' maximum speed of 36 m/s. The acceleration and deceleration values are set to 3.0 m/s^2. The minimum gap between cars is 3.0 m, the time headway is 1.5 s, the delta exponent is 4.0, and the length L of a car is 4 m. The simulation steps at which the cars recalculate their acceleration is set to 1 s. Table 5.1 lists the IDM parameters.

The merging behavior was also implemented using IDM. Since IDM is designed to follow a single car driving directly ahead, it is used twice for merging: for the car that is ahead in the own lane as well as for the car that has to be ahead after merging. When a car c_i has decided to merge behind a car c_j from the other lane, it calculates an acceleration value for c_j in addition to that for the car directly ahead of c_i. From both values, the minimum is chosen to ensure safe following. If c_j is far upstream of the merging, c_i treats it as if it had stopped in the merge. The acceleration value calculated by c_i in view of c_j then causes c_i to wait in front of the merging. After c_j merged, its true speed and location are used for calculation. If c_j is in the proximity of the merging, its speed value is gradually increased to enable a smooth transition between the two states.

Varying incoming flows

Multiple incoming flows were simulated for a duration of $11 \cdot 10^3$ s of simulation time. The lowest flow in lane l_1 was $0.025\,\mathrm{s}^{-1}$ and it was increased in steps of $0.025\,\mathrm{s}^{-1}$ up to $0.225\,\mathrm{s}^{-1}$. The flow in lane l_2 was always twice that of l_1's flow; the two flows in the incoming lanes had on average a ratio of $q_1/q_2 = 0.5$. Each parameter set was simulated 100 times with different random seeds for interarrival times. The lanes were long enough to avoid jam effects at the insertion point of cars (causing lower initial speeds or smaller flows when inserting cars).

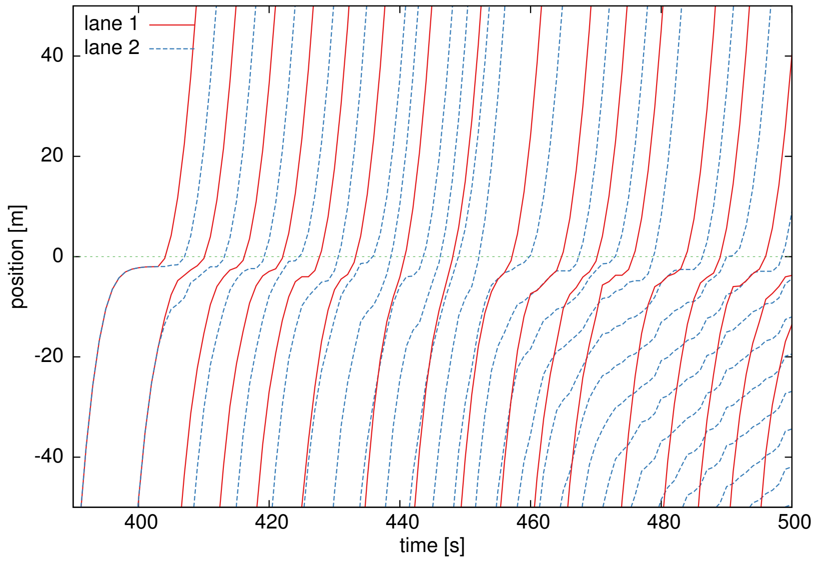

Figure 5.3: Plot showing the positions of cars on the y axis merging from different lanes in the interval [390 s, 500 s] of an exemplary simulation with $q_1 = 0.15\,\mathrm{s}^{-1}$, $q_2 = 0.3\,\mathrm{s}^{-1}$, zero participation, and the IDM car-following model. The horizontal line at $y = 0$ shows the position of the merging.

Varying participation ratios

The simulations were run with flows of $0.15\,\mathrm{s}^{-1}$ and $0.3\,\mathrm{s}^{-1}$ in lanes l_1 and l_2, respectively, and the average participation ratio was varied. The flow sizes were chosen to cause a jam upstream of the merging. Each setup was run for 100 times, each until 3000 cars merged. The lanes were chosen long enough to avoid a jam at the insertion point.

5.6.2 Evaluation

Before the results of the two setups are discussed, it is shown how the car-following model behaves at the merging. The positions and speeds of cars near the merging in the interval from 390 s until 500 s are depicted by the two Figures 5.3 and 5.4, respectively. The setup comprised the flows $q_1 = 0.15\,\mathrm{s}^{-1}$ and $q_2 = 0.3\,\mathrm{s}^{-1}$ without participants and IDM as car-following model. The position plot reveals that the cars do a zipper merge—as expected. If two cars arrive at about the same time at the merging, as it happens with the first two mergers and again at about 450 s, the cars slow down and determine who drives first. In such cases, the car in lane l_1 is preferred. The plot of the speeds pictures only the first 20 cars for clarity. As it was intended, the speeds

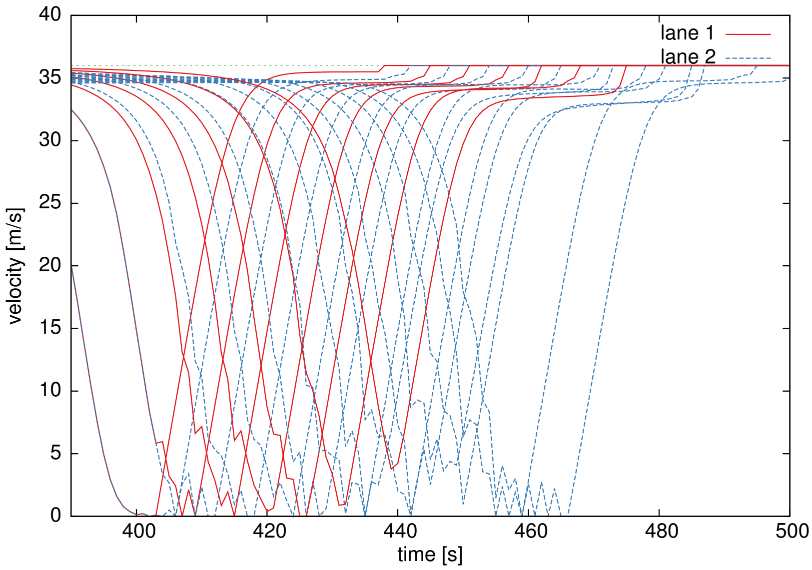

Figure 5.4: Plot showing the speeds in the interval $[390\,\mathrm{s}, 500\,\mathrm{s}]$ of the first 20 cars inserted in an exemplary simulation with $q_1 = 0.15\,\mathrm{s}^{-1}$, $q_2 = 0.3\,\mathrm{s}^{-1}$, zero participation, and the IDM car-following model. The horizontal line at $y = 36$ shows the lanes' maximum speed.

are limited by a minimum of zero and a maximum of 36 m/s. The cars stop shortly in front of the merging and then, if another vehicle is in front not having merged yet, approach this vehicle by driving a few meters before stopping again. This is repeated until it eventually is the car's turn to merge. A car's final acceleration to maximum speed indicates that this car reached the end of the merging area.

Varying incoming flows

Figure 5.5 shows the average unfairness after finishing the simulations, plotted over the sum of incoming flows on the x axis. The simulations were run without participants, which means all cars do the zipper merge, and with a participation ratio of 1 %. The boxes show the 10th percentile, the median, and the 90th percentile. The whiskerbars describe minimum and maximum values. It can be seen that the maximum flow at the merge is about 0.4 cars per second as the average unfairness grows by magnitudes when comparing the results with flows of $0.3\,\mathrm{s}^{-1}$ to $0.45\,\mathrm{s}^{-1}$. The results of runs with 0 % participants for flows larger than $0.45\,\mathrm{s}^{-1}$ are not plotted; the results are far beyond the bounds of the figure. The resulting unfairness with these parameters was merely limited by the fixed simulation duration and the consequently finite number of cars. The figure shows what was predicted by analytical means: with a zipper merge,

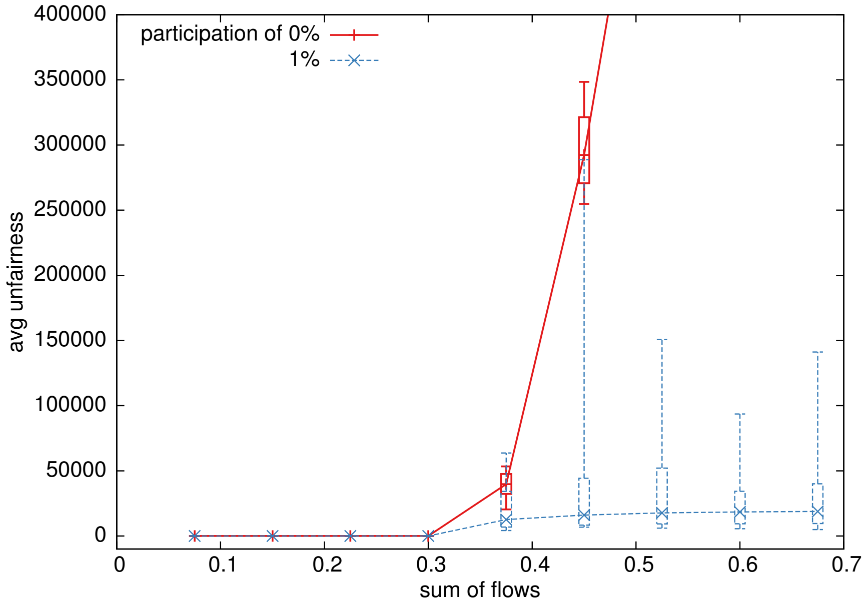

Figure 5.5: Average unfairness after long (but finite) runs with no participation and 1 % participation over different sums of flows in the incoming lanes.

the average unfairness grows without bound. In contrast, runs with 1 % participation converge to a maximum level of unfairness. The absolute position differences have also been calculated to examine how many positions a car is away from its fair place. In the simulations with 1 % participation and higher flows than the merging capacity, the average absolute position difference is about 100 cars.

Varying participation ratios

Figure 5.6 shows the average unfairness with a logarithmic scale on the y axis. For each run, the results of the first 3000 cars traversing the merge are shown in steps of 100 mergings. As in the set of simulations described previously, the boxes show the 10th percentile, the median, and the 90th percentile, while the whiskerbars describe minimum and maximum values. The unfairness of the runs with 100 % participants is constantly zero and is not visible here. Cars begin to build up a jam quickly after the first mergings. The average unfairness then grows monotonically with each merging car unless there are participants that meet at the merge. For all runs with participants, the average unfairness reaches a steady state. Higher participation ratios stabilize at lower states of average unfairness.

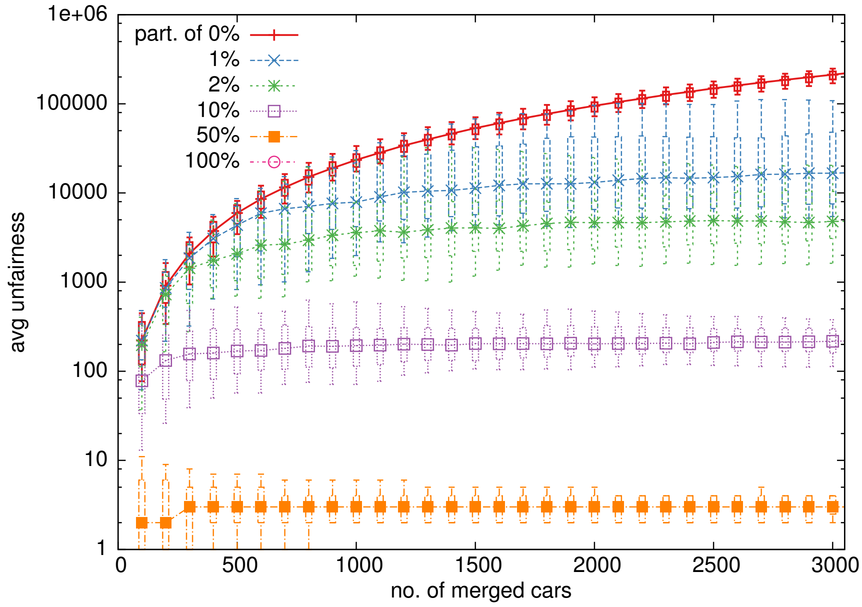

Figure 5.6: Average unfairness on y for a car merging at position x with different participation ratios, log scale.

Figure 5.7 shows how the position differences evolve from an initially fair state. Single runs with varying participation ratios were chosen. The plot depicts the position differences on the y axis over the cars' true merging positions on the x axis. I. e., a point at x coordinate 600 and y coordinate 200 means that the car that should have merged as the 600th car has merged 200 positions too early or too late. Only cars from lane l_1 are shown and the position differences are neither squared nor averaged. The plotted range is limited to the first 1500 mergers. Lane l_1 has a lower flow than lane l_2 and so a point in the figure describes a merging car's advantage. The figure shows that the differences return to a fairer state from time to time which is due to the meeting of two participants from different lanes. However, it can be noticed that the perfect position is not always reached. This is because participating vehicles can only ensure a perfectly fair merge order among themselves. But the the optimal *participant* to merge behind is not necessarily the optimal car when also taking non-participants into account. Non-participants are not able to delay merging for fairness and so position changes may still occur if non-participating vehicles are involved.

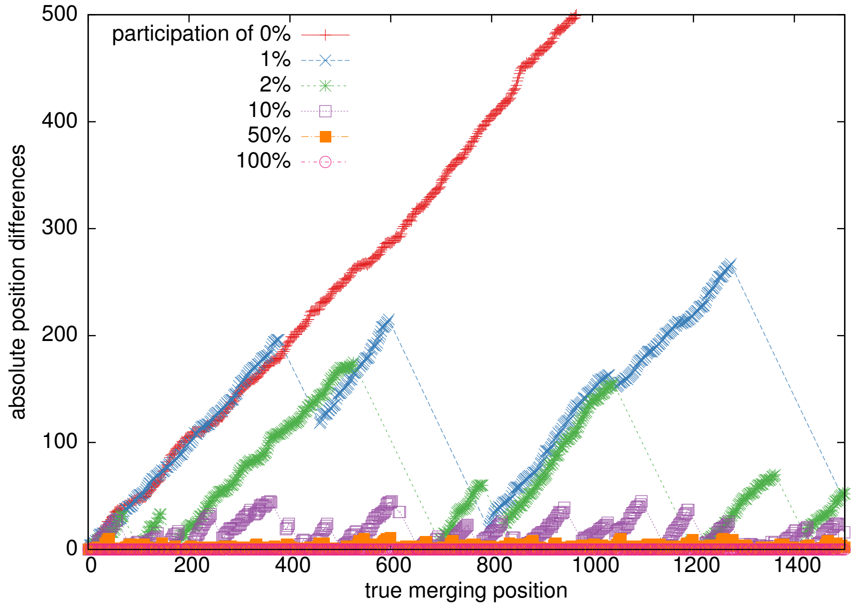

Figure 5.7: Position differences of cars from lane l_1. The differences are positive because the cars in lane l_1 merge too early regarding fairness. The differences decrease when two participants from different lanes meet.

5.7 Conclusion

In this chapter, an application to determine the merging order of cars from two distinct lanes has been discussed. In contrast to the previous chapters, the merging scenario is more complex in that it requires coordination of cars driving in lanes that are independent. The objective of fairness was applied to find a merging order. The notion of free-flow fairness was introduced together with a metric to measure unfairness. A formal modeling of the merging scenario was developed based on the modeling for the cruise-control application in the previous chapter. The optimal behavior to approach the merging was discussed in view of the goals accident absence and road usage efficiency. The merging rules of road traffic were analyzed towards their long-term performance using the unfairness metric and it was found that unfairness grows unboundedly when the flows are higher than the merging's capacity limit. An algorithm was proposed that enables an optimal fair merge order even in this situation. The algorithm was then adapted to work in a fully distributed manner in a setting in which cars do not have complete information about each other. To exchange required knowledge between cars, car-to-car communication technology has been employed with a simple beaconing protocol. To evaluate the algorithm, it has been implemented in a network simulator. Simulations were carried out with IDM and unreliable 802.11p-like communication to understand the influence of the height of incoming flows and of the participation ratio on the merging fairness. The results emphasize the formal analysis: the algorithm enhances fairness compared to the zipper merge in congested situations. This positive effect is noticeable even if only a very small portion of cars participates; a very desirable but rare feature for C2CC applications.

Chapter 6

Conclusions

The thesis at hand proposed to conduct communication protocol development for inter-vehicle communication *top-down*. This research direction is in direct contrast to the predominant way of protocol design found in the literature, which has consequentially been termed bottom-up. The top-down approach frees protocol design from the limitation of discussing a target application after a protocol's technological constraints were already set. In addition, an explicit modeling of specific applications provides a protocol developer with metrics that help to tweak the right parameters to enhance the application-layer performance of bottom-up protocols. After the statement of the top-down idea, supported by a research roadmap and a simple example, the approach was applied to two basic road settings and delivered, despite the abstract modeling, results in terms of how to design a protocol and how close we can get to optimal solutions with inter-vehicle communication. At the discussion of a cruise control application, it has been found that high importance should be given to delivering periodic broadcasts successfully. The algorithm Carrot was developed to accomplish this objective with the concept of an implicit packet loss detection and a fast repetition scheme. The second setting examined was the merging of two lanes. There, the notion of free-flow fairness has been introduced to enable vehicles from different lanes to merge in a fair order—something that was not thought of for C2CC before looking at mergings top-down. Analysis showed that the predominant strategy for merging, the zipper merge, does not perform well regarding fairness in congested situations, thus a beacon-based C2CC algorithm was created that yields fairness even if only few cars use it.

Before the details of the top-down approach were explained, a survey on related work about C2CC protocol design in Chapter 2 pointed out that related research follows the bottom-up way for protocol design and enhancement. Although many of the considered papers specify an application as use case for their protocols, and some even evaluate application-level metrics, the papers do neither define the application behavior in depth nor reason formally whether the chosen metrics are perfectly suited to measure the information requirements. The characteristics of applications are rather exploited to construct protocols working better than others in network layer terms. Papers from the field of control theory approach from the other direction by formally specifying vehicle behavior but do not pay the necessary attention to the details of information exchange by network protocols. A few works, however, discuss the outstanding importance of applications in a VANET and emphasize that there are connections between vehicle behavior, information exchange, and network protocols, although they do not start an extensive modeling of vehicle behavior as done in this thesis.

This appears to be reasoned by the complexity an all-embracing top-down approach involves as it was indicated by discussing very distinct areas of related work. The necessary steps to be taken for a top-down approach were proposed in a roadmap in Chapter 3, followed by an example of how it can be applied to derive knowledge about protocol design from modeling the behavior of cars. The different levels of knowledge considered in that example about car following revealed the existence of a steady state on the application level when sending periodic broadcasts.

The next step was made by observing the behavior of multiple cars driving in a lane in a row. The optimal beacon sending times for minimizing the road usage of such a platoon were determined in Chapter 4 for a cruise-control application. However, the optimality criterion regarding the road-usage efficiency objective required the cars to accelerate and decelerate ceaselessly. To get closer to a realistic behavior, the efficiency objective was reformulated to create a constant-speed following. With this, the optimal share of bandwidth for beaconing was shown. The impact of network layer effects was analyzed and it was found that beacons are to upstream cars what carrots are to donkeys: if a car does not get a regular glimpse at the car ahead, it just stops moving. To prevent cars from stopping, the algorithm Carrot was proposed which sends at the optimal times for maintaining the steady state and which mitigates the effects of packet losses by two mechanisms; cars using Carrot send beacons one after another in upstream direction using a topology-related trigger scheme in which a missing packet is signaled by a stop of the beaconing chain. On the detection of a loss, a fast repetition mechanism sends a beacon one more time, though with fresh information. That way, Carrot is able to maintain a much lower steady state distance between the cars than possible with a simple beaconing.

While cars in a platoon only need to interact with their direct neighbors ahead and behind, cars in two lanes that head towards a merging have to decide how to combine two independent streams of cars. It was explained why the decision on the merging order should be based on fairness. The free-flow fairness was introduced which reflects the idea that cars merge in a fair order if they merge according to their free-flow arrival times. These are virtual arrival times that are unbiased by other cars. It was shown that the zipper merging scheme performs bad regarding fairness in congested settings, while the proposed beacon-based algorithm enables perfect fairness under optimal conditions. The conditions were then worsened by assuming only fractions of the cars being able to use the algorithm; but still, the zipper merge was outperformed analytically as well as in simulations.

Summing up the results from this thesis, approaching protocol design top-down is absolutely feasible and, as this thesis already proved, can be expected to return valuable design guidelines. An exhaustive modeling of specific situations and applications, however, appear a laborious task. It demands expertise in various fields of research and engineering, as became apparent during the distinct top-down steps throughout the thesis, and clearly exceeds the capabilities of any specialized research group. The joint effort of working groups specialized in different research areas is required, like traffic theory, mathematics, networking standards, beaconing protocols, simulations, and experiments. After all, it is the way to go to complement bottom-up protocols with knowledge that is independent from technological constraints by concentrating on the objectives of transportation that did not change since the beginnings of individual traffic: safety, efficiency, and fairness.

Appendix

Appendix A

Proofs

This chapter contains the proofs to the theorems and the lemmata. Each of the following sections discusses a distinct proof, starting with a repetition of the statement. Some proofs are split into several propositions (that are proved directly after stating them). The proofs use the definitions of the respective chapters in which the theorems or lemmata are stated if not noted otherwise. Before the first proof is started, a consistent definition of the distance between two ways of cars is given, since this will be needed multiple times.

Definition 19. Distance. *The distance between two ways $w_1(t)$ and $w_2(t)$ at t is defined as: $\Delta w(t) = w_1(t) - w_2(t)$.*

A.1 Proof of Theorem 1

Theorem 1 states: *If c_2 continuously knows about c_1's present and future position and speed, then there is an optimal accident-absent way for c_2 where, after an initial transition period until $t_b = t_e + 2(\Delta w(t_e)/A)^{1/2}$, c_1 and c_2 drive at the same speed and at the same point on the road.*

Here, $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/A$ and $\Delta w(t_e) = w_1(t_2^0) + A(t_e - t_2^0)^2/2$.

There are two cars $c_1, c_2 \in C$. Car c_1 is driving with constant speed. Car c_2 is considered to always have precise knowledge about the position and speed of c_1. It also knows that the speed of c_1 will not change. In the following, the behavior of c_2 is reasoned and the value of t_b is derived. But first, some preliminary properties have to be discussed. It will be started with how a way with a constant second order derivative looks like.

Lemma 5. *Given $i \in \mathbb{N}$, $k \in \mathbb{R} \cap [D, A]$ and a vehicle c_i. $\forall t \geq t_i^0$: The second-order derivative of a valid way w of c_i be $w''(t) = k = const.$ Then $\forall t \geq t_i^0$:*

$$w(t) = w'(t_i^0) \cdot (t - t_i^0) + \frac{k}{2} \cdot \left(t - t_i^0\right)^2.$$

Proof. With $k \in \mathbb{R}$ and $w''(t) = k$ constant for c_i: at a point $t \in \mathbb{R} \geq t_i^0$, the derivative of the way of c_i (its speed) is

$$w'(t) = w'(t_i^0) + \int_{t_i^0}^{t} w''(\tau)d\tau = w'(t_i^0) + k \cdot (t - t_i^0)$$

And with this the way is obtained:

$$\begin{aligned} w(t) &= w(t_i^0) + \int_{t_i^0}^{t} w'(\tau)d\tau \\ &= w(t_i^0) + w'(t_i^0) \cdot (t - t_i^0) + \frac{k}{2} \cdot (t - t_i^0)^2 \end{aligned}$$

For every valid way, it is $w(t_i^0) = 0$ and $k \in [D, A]$, hence

$$w(t) = w'(t_i^0) \cdot (t - t_i^0) + \frac{k}{2} \cdot (t - t_i^0)^2$$

□

Note that this lemma can be used for describing a valid way of c_i that consists of concatenated parts of arbitrary intervals with constant acceleration. But for now, the Lemma is used for proving that no way of a car is faster than that with full acceleration at any point in time since the appearance of the car.

Lemma 6. *Given $i \in \mathbb{N}$, a car c_i, a valid way $\hat{w}$ with $\hat{w}''(t) = A$, and an arbitrary valid way w for c_i. $\forall t \geq t_i^0$:*

a) $w(t) \leq \hat{w}(t)$ and

b) $w(t) < \hat{w}(t) \Rightarrow \forall \hat{t} \geq t : w(\hat{t}) < \hat{w}(\hat{t})$.

Proof. Both statements will be proved by contradiction.

(a) Assume $\exists t_x \geq t_i^0 : w(t_x) > \hat{w}(t_x)$. It follows

$$\int_{t_i^0}^{t_x} w'(\tau)d\tau > \int_{t_i^0}^{t_x} \hat{w}'(\tau)d\tau$$

Both ways are valid. They thus start at 0, so $t_x > t_i^0$.

$$\begin{aligned}
&\Rightarrow \exists t_y \in [t_i^0, t_x) : w'(t_y) > \hat{w}'(t_y) \\
&\Leftrightarrow w'(t_i^0) + \int_{t_i^0}^{t_y} w''(\tau)d\tau > \hat{w}'(t_i^0) + \int_{t_i^0}^{t_y} \hat{w}''(\tau)d\tau.
\end{aligned}$$

Because $w'(t_i^0) = \hat{w}'(t_i^0)$ and $\hat{w}''(\tau) = A$:

$$\Rightarrow \exists t_z \in [t_i^0, t_y) : w''(t_z) > A$$

But if $w''(t_z) > A$: w is not valid. Contradiction.

(b) Assume it holds $w(t) < \hat{w}(t)$ and $\exists \hat{t} \geq t : w(\hat{t}) \geq \hat{w}(\hat{t})$. Given $t \geq t_i^0 : w(t) < \hat{w}(t)$. It follows

$$\begin{aligned}
w(\hat{t}) \geq \hat{w}(\hat{t}) \Rightarrow \int_t^{\hat{t}} w'(\tau)d\tau &\geq \int_t^{\hat{t}} \hat{w}'(\tau)d\tau \\
\Rightarrow \exists t_x \in [t, \hat{t}) : w'(t_x) \geq \hat{w}'(t_x)& \\
\Leftrightarrow w'(t) + \int_t^{t_x} w''(\tau)d\tau \geq \hat{w}'(t) + \int_t^{t_x} \hat{w}''(\tau)d\tau.&
\end{aligned}$$

To check if the last inequation holds, it is started with considering the intervals. It is $\hat{w}''(\tau) = A$. If

$$\int_t^{t_x} w''(\tau)d\tau \geq \int_t^{t_x} \hat{w}''(\tau)d\tau,$$

then $\exists t_y \in [t, t_x) : w''(t_y) > A$. But $w''(t_y) > A$ means that w is not valid. w is assumed to be valid, so this is contradiction. So it must be $w'(t) \geq \hat{w}'(t)$ which implies

$$\int_{t_i^0}^{t} w'(\tau)d\tau \geq \int_{t_i^0}^{t} \hat{w}'(\tau)d\tau.$$

The speeds of both ways are equal at t_i^0, so this inequality can hold only if $\exists t_z \in [t_i^0, t) : w''(t_z) > A$. This is a contradiction to w being a valid way. Therefore, the assumption is wrong and it holds: $w(t) < \hat{w}(t) \Rightarrow \forall \hat{t} \geq t : w(\hat{t}) < \hat{w}(\hat{t})$. □

Similarly to that no way is faster than that with full acceleration, no way is slower at any point in time than that with maximum deceleration.

Lemma 7. *Given $i \in \mathbb{N}$, a car c_i, a valid way $\hat{w}$ with*

$$\forall t \geq t_i^0 : \hat{w}''(t) = \begin{cases} D \text{ if } \hat{w}'(t) > 0 \\ 0 \text{ else,} \end{cases}$$

and an arbitrary valid way w for c_i.

a) $w(t) \geq \hat{w}(t)$ and

b) $w(t) > \hat{w}(t) \Rightarrow \forall \hat{t} \geq t : w(\hat{t}) \geq \hat{w}(\hat{t})$.

Proof. Again, both statements will be proved by contradiction.

(a) Assume $\exists t_x \geq t_i^0 : w(t_x) < \hat{w}(t_x)$. It follows

$$\int_{t_i^0}^{t_x} w'(\tau)d\tau < \int_{t_i^0}^{t_x} \hat{w}'(\tau)d\tau$$

Both ways are valid. They thus start at 0, so $t_x > t_i^0$.

$$\Rightarrow \exists t_y \in [t_i^0, t_x) : w'(t_y) < \hat{w}'(t_y)$$
$$\Leftrightarrow w'(t_i^0) + \int_{t_i^0}^{t_y} w''(\tau)d\tau < \hat{w}'(t_i^0) + \int_{t_i^0}^{t_y} \hat{w}''(\tau)d\tau.$$

Because $w'(t_i^0) = \hat{w}'(t_i^0)$, only the integrals have to be considered. Let c_2 stop at time t_s: $\exists t_s \in [t_i^0, t_y]$, i.e., the part before and after stopping has to be discussed.

Part 1: interval $[t_i^0, t_s]$. $\forall \tau \in [t_i^0, t_s], \hat{w}''(\tau) = D$. It follows $\exists t_z \in [t_i^0, t_s) : w''(t_z) < D$. But $w''(t_z) < D$ means that w is not valid. This is a contradiction to the requirement that w is valid.

Part 2: assume c_i stops with the way $\hat{w}(t)$ at $t_s \in [t_i^0, t]$, so $\forall \tau \geq t_s : \hat{w}''(\tau) = 0 \wedge \hat{w}'(\tau) = 0$. Part 1 showed that $w(t) \geq \hat{w}(t)$ for $t \leq t_s$. Thus, $\exists t_z \in [t_s, t_y) : w(t_z) < \hat{w}(t_z) \wedge w'(t_z) < \hat{w}'(t_z) = 0$. But $w'(t_z) < 0$ means that w is not valid. Contradiction to the requirement that w is valid.

(b) Assume it holds $w(t) > \hat{w}(t)$ and $\exists \hat{t} \geq t : w(\hat{t}) < \hat{w}(\hat{t})$. Given $t \geq t_i^0 : w(t) > \hat{w}(t)$. It follows

$$w(\hat{t}) < \hat{w}(\hat{t}) \int_t^{\hat{t}} w'(\tau)d\tau \qquad < \int_t^{\hat{t}} \hat{w}'(\tau)d\tau$$
$$\Rightarrow \exists t_x \in [t, \hat{t}) : w'(t_x) < \hat{w}'(t_x).$$
$$\Leftrightarrow w'(t) + \int_t^{t_x} w''(\tau)d\tau < \hat{w}'(t) + \int_t^{t_x} \hat{w}''(\tau)d\tau.$$

To check if the last inequation holds, it is started with considering the intervals. Let c_2 stop at time t_s: $\exists t_s \in [t_i^0, t_y]$, i.e., the part before and after stopping has to be discussed.

Part 1: Interval $[t_i^0, t_s]$ with $\hat{w}''(\tau) = D$. If

$$\int_t^{t_x} w''(\tau) d\tau < \int_t^{t_x} \hat{w}''(\tau) d\tau,$$

then $\exists t_y \in [t, t_x) : w''(t_y) < D$. But with $w''(t_y) < D$, w is not a valid way. This is a contradiction to the requirement of w being valid.

Part 2: Assume c_i stops with the way $\hat{w}(t)$ at t_s, so $\forall \tau \geq t_s : \hat{w}''(\tau) = 0 \wedge \hat{w}'(\tau) = 0$. Part 1 showed that $w(t) \geq \hat{w}(t)$ for $t \leq t_s$. Thus to ensure $w(\hat{t}) < \hat{w}(\hat{t})$ holds: $\exists t_z \in [t_s, t) : w(t_z) = \hat{w}(t_z) \wedge w'(t_z) < \hat{w}'(t_z) = 0$. But if $w'(t_z) < 0$: w is not valid. Contradiction.

Now that it is known that the intervals are equal at best, consider $w'(t) < \hat{w}'(t)$. Here, also a t_s may exist at which c_i stops with $\hat{w}(t)$.

Part 1: Interval $[t_i^0, t_s]$. The initial speeds at t_i^0 are equal, so it holds $w'(t) < \hat{w}'(t) \Leftrightarrow \int_{t_i^0}^{t} w'(\tau) d\tau < \int_{t_i^0}^{t} \hat{w}'(\tau) d\tau$. This inequality can hold only if $\exists t_z \in [t_i^0, t) : w''(t_z) < D$. Contradiction to w being a valid way.

Part 2: $[t_s, t]$. Part 1 showed that $w(t) \geq \hat{w}(t)$ for $t \leq t_s$. Thus to ensure $w(\hat{t}) < \hat{w}(\hat{t})$ holds: $\exists t_z \in [t_s, t) : w(t_z) = \hat{w}(t_z) \wedge w'(t_z) < \hat{w}'(t_z) = 0$. With $w'(t_z) < 0$, w is not valid. This is a contradiction to the requirement that w is valid.

The assumption is thus wrong and it holds: $w(t) > \hat{w}(t) \Rightarrow \forall \hat{t} \geq t : w(\hat{t}) \geq \hat{w}(\hat{t})$. ☐

It also necessary to know what the minimum time span for c_2 is until it is equally fast as c_1, with the constraints that c_2 starts at t_2^0 and c_1 drives with a constant speed.

Lemma 8. *Given two cars c_1 and c_2 with $c_2 <_c c_1$, and two valid ways w_1 and w_2 for the respective cars. The speed of c_1 is constant. c_2 always knows the precise location of c_1 and that c_1's speed will not change. Then the minimum time when the cars' speeds are equal is*

$$t_e = t_2^0 + \left| \frac{w_1'(t_2^0) - w_2'(t_2^0)}{A} \right|.$$

Proof. To find t_e, two cases are considered regarding the relation of the cars' initial speeds.

Case $w_1'(t_2^0) \geq w_2'(t_2^0)$: c_2 has to accelerate with the maximum rate to minimize time t_e. The maximum acceleration rate is A. The initial speed of c_2 is $w_2'(t_2^0)$, the speed to reach is $w_1'(t_2^0)$. Thus to equalize the speeds in minimum time, it takes $(w_1'(t_2^0) - w_2'(t_2^0))/A$ seconds.

Case $w_1'(t_2^0) < w_2'(t_2^0)$: c_2 has to decelerate with the maximum rate to minimize time t_e. The maximum deceleration rate is $D = -A$. The initial speed of c_2 is $w_2'(t_2^0)$, the speed to reach is $w_1'(t_2^0)$. This gives $(w_1'(t_2^0) - w_2'(t_2^0))/D = -(w_1'(t_2^0) - w_2'(t_2^0))/A = |-(w_1'(t_2^0) - w_2'(t_2^0))/A| = |(w_1'(t_2^0) - w_2'(t_2^0))/A|$. □

It is demanded that ways are accident absent. But it has not yet been discussed if accident-absent ways for c_1 and c_2 exist.

Lemma 9. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two valid ways w_1 and w_2 for the respective cars. The speed of c_1 is constant. c_2 always knows the precise location of c_1 and that c_1's speed will not change. Accident-absent ways for c_1 and c_2 exist iff*

$$w_1(t_2^0) \geq \frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{2D}.$$

Proof. The proof is divided in two cases. Either c_1 is slower than or equally fast as c_2 at t_2^0 or c_1 is faster.

Case $w_1'(t_2^0) \leq w_2'(t_2^0)$: assume that accident-absent ways exist. Car c_2 appears on t_2^0 at $w_2(t_2^0) = 0$. At this time, c_1 is at $w_1(t_2^0) = w_1'(t_1^0)(t_2^0 - t_1^0)$. The speed of c_1 at t_2^0 is $w_1'(t_2^0) = w_1'(t_1^0)$ and the speed of c_2 is $w_2'(t_2^0)$. The minimum time t_e at which it holds $w_1'(t_e) - w_2'(t_e) = 0$ is described in Lemma 8. Accordingly, for $t \in [t_2^0, t_e)$: $w_2''(t) = D$. The distance of the cars changes from t_2^0 to t_e by

$$\int_{t_2^0}^{t_e} w_1(\tau) - w_2(\tau) d\tau$$

Changing this to an integral over the speed and with $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/D$:

$$\begin{aligned}
&\int_0^{t_e - t_2^0} A\tau d\tau \\
=&\int_0^{t_e - t_2^0} \frac{w_1'(t_2^0) - w_2'(t_2^0)}{t_e - t_2^0} \tau d\tau \\
=&\frac{w_1'(t_2^0) - w_2'(t_2^0)}{t_e - t_2^0} \frac{(t_e - t_2^0)^2}{2} \\
=&\frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{2D}
\end{aligned}$$

This equals the minimum initial distance of the cars, $w_1(t_2^0) - w_2(t_2^0) = w_1(t_2^0)$, so that accident-absent ways exist. It follows: accident-absent ways exist if $w_1(t_2^0) - w_2(t_2^0) \geq (w_1'(t_2^0) - w_2'(t_2^0))^2/(2D)$.

The other direction: given $w_1(t_2^0) = w_1(t_2^0) - w_2(t_2^0) \geq (w_1'(t_2^0) - w_2'(t_2^0))^2/(2D)$. The minimum time with $w_1'(t) - w_2'(t) = 0$ is $t = t_e$. This requires $t \in [t_2^0, t_e)$: $w_2''(t) = D$. Then, in that interval, the distance between the cars changes by $(w_1'(t_2^0) - w_2'(t_2^0))^2/(2D)$. The distance at t_e is, hence, $w_1(t_e) - w_2(t_e) \geq 0$. $\forall t \geq t_e : w_2''(t) = 0$, then c_2 remains at a location less equal to c_1, and c_2 drives with the same speed as c_1. The ways $w_1(t)$ and $w_2(t)$ are, consequently, accident absent. $w_1(t_2^0) - w_2(t_2^0) \geq (w_1'(t_2^0) - w_2'(t_2^0))^2/(2D) \Rightarrow$ accident-absent ways exist.

Case $w_1'(t_2^0) > w_2'(t_2^0)$: c_1 is faster than c_2 at t_2^0 and accident-absent ways exist $\Rightarrow w_1(t_2^0) \geq (w_1'(t_2^0) - w_2'(t_2^0))^2/(2D)$. The formula discussed by the previous case, $(w_1'(t_2^0) - w_2'(t_2^0))^2/(2D)$, now has a negative result. The initial distance between the cars, $w_1(t_2^0) - w_2(t_2^0) = w_1(t_2^0) - 0$, is positive. So it holds $w_1(t_2^0) \geq (w_1'(t_2^0) - w_2'(t_2^0))^2/(2D)$.

The other direction to proof is: $w_1(t_2^0) \geq (w_1'(t_2^0) - w_2'(t_2^0))^2/(2D) \Rightarrow$ accident-absent ways exist. Every initial distance greater or equal to zero allows accident-absent ways for the cars: a possible one is a way without acceleration. An initial distance less zero does not fit to the definition of a lane that starts at position 0; $w_1(t_2^0) < 0 = w_2(t_2^0)$ is not allowed. Accident-absent ways, for that reason, exist. $\square$

The lemmata discussed until now build the ground for the traveling-optimal and accident-absent behavior of c_2 that is described in the following lemma.

Lemma 10. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and let accident-absent ways w_1 and w_2 exist for the respective cars. The speed of c_1 is constant. c_2 always knows the precise location of c_1 and that c_1's speed will not change. $\forall t \geq t_2^0$: w_2 is traveling optimal iff*

$$w_2''(t) = \begin{cases} A \text{ if } t \in [t_2^0, t_a) \\ D \text{ if } t \in [t_a, t_b) \\ 0 \text{ if } t \geq t_b \end{cases}$$

with $t_a = t_e + (\Delta w(t_e)/A)^{1/2}$, $t_b = t_e + 2(\Delta w(t_e)/A)^{1/2}$, $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/A$, and $\Delta w(t_e) = w_1(t_2^0) + A/2 \cdot (t_e - t_2^0)^2$.

Proof. The proof is structured into three parts: first, it will be shown that the acceleration behavior follows from w_2 being traveling optimal. Then, the other way is shown, i.e., w_2 is traveling optimal, when the acceleration behavior is given. Finally, the shape of the variables t_a, t_b, t_e, and $\Delta w(t_e)$ is derived.

Part 1: given w_2 is traveling optimal. It will be shown that the acceleration behavior $w_2''(t)$ from the lemma follows.

Assume acceleration with A from t_2^0 onwards. To every point on the lane, w_2 is earlier or equally fast at this point compared to every other accident-absent way. A danger of accident is apparent if the distance between the vehicles is lower or equal than a *minimum braking distance* of an accident-absent way similar to what is stated in Lemma 9 as criterion for the existence of such ways. The minimum braking distance is calculated as in the proof of Lemma 9 with the adaption of t_2^0 to an arbitrary point in time:

$$w_1(t) - w_2(t) \geq \frac{(w_1'(t) - w_2'(t))^2}{2D}.$$

If no danger of accident is apparent, $t \in [t_2^0, t_a)$, $w_2''(t)$ has to be at maximum, which is A, as explained in Lemma 6. The acceleration continues until the point in time when that inequality becomes an equation: $w_1(t) - w_2(t) = (w_1'(t) - w_2'(t))^2/(2D)$. Be this time $t_a \geq t_2^0$. At t_a, the distance equals the minimum braking distance, and the full braking with D is required to maintain this minimum braking distance while $w_1'(t) < w_2'(t)$. No other way than the slowest possible as discussed in Lemma 7 is accident absent in this situation; this becomes apparent when regarding the construction of the minimum braking distance. The minimum braking distance as well as the distance are equal at t_a and grow equally during deceleration. So the speed of c_2 is reduced for $t \geq t_a$: $w_2'(t) = w_2'(t_a) + D(t - t_a)$ and $w_1'(t) = \text{const}$. The speed of c_2 approaches that of c_1. Let the speeds be equal at time t_b. Because of the construction of the minimum braking distance, the distance of c_1 and c_2 is zero at t_b. At this time, with w_2 being an accident-absent way, c_2 can only accelerate with $w_2''(t) \in [D, 0]$. w_2 is traveling-optimal, so acceleration at the maximum that is possible to result in an accident-absent way is chosen; it is, thus, 0. This is a steady state, $\forall t \geq t_b : w_2''(t) = 0$.

Ergo the behavior of c_2 proposed by the lemma follows if w_2 is traveling optimal. Figure A.1 illustrates the speeds within the different acceleration phases of c_2 for a setting with $w_1'(t_2^0) > w_2'(t_2^0)$.

Part 2: given the acceleration behavior for $w_2''(t)$ from the lemma, it will be shown that w_2 is traveling optimal.

Look at the acceleration phase for $t \in [t_2^0, t_a)$. In general, constant acceleration with A is traveling optimal, because no valid way with another acceleration function is further ahead at any time. But this strategy is not accident absent if the distance is shorter or equal than the minimum braking distance. The distance equals this minimum braking distance at time t_a, the only acceleration from then on resulting in an accident-absent way is with D. For that reason it is also the only behavior that results in a traveling-optimal way; remember that traveling-optimal ways are valid and accident absent according to Definition 9. At t_b, the speed of c_2 equals that of c_1 and

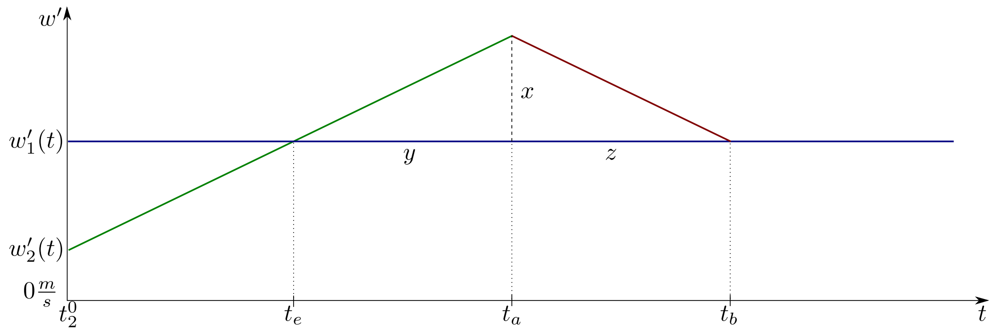

Figure A.1: Speeds of the cars if c_2 has perfect knowledge about c_1. At t_e, the speeds are equal for the first time since the start of c_2. t_a is the point in time at which c_2 starts braking. At time t_b, the steady state is reached.

their distance is zero. The change of c_2's acceleration from D to 0 results in c_2 staying at the same location as c_1 with the same speed. This clearly is accident absent and, moreover, a traveling-optimal way.

Part 3: having discussed both directions of the iff relationship, the values of t_a, t_b, t_e and $\Delta w(t_e)$ remain to be justified. Because the existence of the accident-absent ways w_1 and w_2 is required, it holds $w_1(t_2^0) - w_2(t_2^0) \geq (w'_1(t_2^0) - w'_2(t_2^0))^2/(2D)$; the initial distance is larger than or equal to the minimum braking distance. c_2 accelerates with A until this minimum distance, which depends on the current speeds of the vehicles, is reached at t_a. First, assume $w'_1(t_2^0) \geq w'_2(t_2^0)$ as in Figure A.1. On t_e, the speeds are equal, so the formula for the minimum braking distance results zero: $t_e = t_2^0 + (w'_1(t_2^0) - w'_2(t_2^0))/A$, see Lemma 8. For $t \in [t_2^0, t_e)$, the distance grows, because c_1 is faster than c_2. The distance at t_e has to be reduced to zero until t_b. This is performed by c_2 first accelerating further until t_a and decelerating with D until t_b. The slope of c_2's speed for $t \in [t_e, t_a)$ is $A = x/y$ and the slope for $t \in [t_a, t_b)$ is $D = -x/z$. With $A = -D$, according to Definition 5, it follows $z = y$. The distance changes from t_e to $t_e + y = t_a$ by the same amount as from t_a to $t_e + y + z = t_e + 2y = t_b$ which can be seen when integrating the speed of c_2 over those intervals. To obtain y, it is determined how long it takes to halve the vehicles' distance at t_e, $\Delta w(t_e)$, with $y = t_a - t_e$:

$$\begin{aligned}
\frac{\Delta w(t_e)}{2} &= \Delta w(t_e) + w'_1(t_e)y - w'_2(t_e)y - \frac{A}{2}y^2 \\
\Leftrightarrow \frac{\Delta w(t_e)}{2} &= \Delta w(t_e) + 0 - \frac{A}{2}y^2 \\
\Rightarrow y &= \sqrt{\frac{\Delta w(t_e)}{A}}.
\end{aligned}$$

y has to be positive and so the positive root is chosen. Since $y = t_a - t_e$, it is obtained $t_a = t_e + (\Delta w(t_e)/A)^{1/2}$. That is the expression for t_a that was searched for. Directly from knowing y, it can be stated $t_b = t_e + 2y = t_a + y = t_e + 2(\Delta w(t_e)/A)^{1/2}$.

$\Delta w(t_e)$ is easily derived with Lemma 5:

$$\begin{aligned}\Delta w(t_e) &= w_1(t_e) - w_2(t_e) \\ &= w_1(t_2^0) - w_2(t_2^0) + (w_1'(t_2^0) - w_2'(t_2^0))(t_e - t_2^0) + \left(0 - \frac{A}{2}\right)(t_e - t_2^0)^2\end{aligned}$$

With $w_2(t_2^0) = 0$ and $t_e - t_2^0 = (w_1'(t_2^0) - w_2'(t_2^0))/A$:

$$\begin{aligned}\Delta w(t_e) &= w_1(t_2^0) + (w_1'(t_2^0) - w_2'(t_2^0))\left(\frac{w_1'(t_2^0) - w_2'(t_2^0)}{A}\right) - \frac{A}{2}\left(\frac{w_1'(t_2^0) - w_2'(t_2^0)}{A}\right)^2 \\ &= w_1(t_2^0) + \frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{A} - \frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{2A} \\ &= w_1(t_2^0) + \frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{2A} \\ &= w_1(t_2^0) + \frac{A}{2}(t_e - t_2^0)^2.\end{aligned}$$

So far, the case $w_1'(t_2^0) \geq w_2'(t_2^0)$ has been discussed. In case of $w_1'(t_2^0) < w_2'(t_2^0)$, c_2 is regarded as if the speeds have been equal at $t_e < t_2^0$ and c_2 accelerated in $[t_e, t_2^0)$. This means $t_2^0 \in [t_e, t_a]$. It is still $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/A$; note that the absolute value bars in Lemma 8 introduced for braking with $D = -A$ are not applied here since c_2 does not brake. The formulae for t_a, t_b, and also $\Delta w(t_e)$ are valid as stated above. $\square$

With the proof of that lemma, all statements of Theorem 1 are proved. Behaving traveling optimal (and thereby accident absent, too) implies that there is no point in time where any other accident-absent way is further ahead. Being furthest ahead also means being closest to c_1. Thus this behavior of c_2 is also considered being optimal regarding the use of road space: the distance between the vehicles is the shortest possible any time.

A.2 Proof of Theorem 2

Theorem 2 states: *If c_2 learns about c_1's position and speed only once at time t_2^0, then c_2 must come to an halt at time*

$$t_b = t_2^0 - \frac{w_2'(t_2^0)}{A} + 2\sqrt{\left(1 + \frac{w_2'(t_2^0)}{A}\right)\frac{w_1(t_2^0)}{A} + 2t_e^2} \quad \textit{with } t_e = \frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A}$$

in order to guarantee accident absence. The steady-state distance between both cars therefore increases unboundedly as time passes and c_1 proceeds on its way.

One more time, it is looked at two cars c_1 and $c_2 \in C$. c_1 drives with constant speed, again, but c_2 only gets a single update about c_1's position and speed at t_2^0 instead of having perfect knowledge all the time. From the viewpoint of c_2, c_1 must choose an arbitrary valid way that fits to the information received at t_2^0. When c_2 chooses a way which is accident absent with the slowest valid way of c_1, that way of c_2 is accident absent with every valid way of c_1. This slowest way is discussed in Lemma 7; no valid way makes less progress than this one and has, for this reason, a shorter distance to a valid way of c_2. The set $\hat{\mathrm{W}}_2$ (of accident-absent ways of c_2) contains the ways that are accident-absent with all valid ways, and thus the slowest valid way, too, of c_1.

To reflect the changed estimated behavior of c_1, some of the lemmata stated for the proof of Theorem 1 have to be reformulated. The minimum time until the speeds are equal differs from the case of having perfect knowledge that is discussed in Lemma 8, because c_2 now assumes c_1 to be decelerating with D.

Lemma 11. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two valid ways w_1 and w_2 for the respective cars. Starting at t_2^0, the speed of c_1 decreases to zero with a deceleration of D and then remains zero. c_2 knows the precise location and speed of c_1 at t_2^0. Then the minimum time when the cars' speeds are equal is*

$$t_e = t_2^0 + \begin{cases} \frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A}, \textit{ if } w_1'(t_2^0) \geq w_2'(t_2^0) \\ \frac{w_2'(t_2^0)}{A}, \textit{ else.} \end{cases}$$

Proof. The time t_e until the cars' speeds are equal has to be minimized. For this, it will be differentiated between two cases of the initial relation of the speeds.

Case $w_1'(t_2^0) \geq w_2'(t_2^0)$. c_2 has to accelerate with the maximum rate to minimize time t_e. The maximum acceleration rate is A. c_1 has to accelerate with the minimum rate $D = -A$. The initial speed of c_2 is $w_2'(t_2^0)$, while that of c_1 is $w_1'(t_2^0)$. Because both vehicles accelerate with the same amount but opposite signs, the minimum time is exactly half of that in Lemma 8, which is $(w_1'(t_2^0) - w_2'(t_2^0))/A$, so it is $(w_1'(t_2^0) - w_2'(t_2^0))/(2A)$.

Case $w_1'(t_2^0) < w_2'(t_2^0)$. c_1 is slower than c_2 at t_2^0. c_1 decelerates with D until full stop at $t_s = -w_1'(t_2^0)/D$. To equalize the speeds, c_2 must also decelerate until the speeds are equal. This will be when c_2 came to a stop, since during the deceleration of both vehicles, the speed difference stays constant: both vehicles decelerate with D. For c_2 to stop, it takes at minimum $-w_2'(t_2^0)/D = w_2'(t_2^0)/A$. □

c_2 only knows the location and speed of c_1 at one time. After this time, c_2 has to stay accident absent with the slowest valid way of c_1. Here, Lemma 9 about the existence of accident-absent ways for c_2 is not applicable; therefore, it is adapted in the following lemma.

Lemma 12. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two valid ways w_1 and w_2 for the respective cars. The speed of c_1 is constant for $t \in [t_1^0, t_2^0)$. Starting at t_2^0, the speed of c_1 decreases to zero with a deceleration of D and then remains zero. c_2 knows the precise location and speed of c_1 at t_2^0. Accident-absent ways for c_1 and c_2 exist iff*

$$w_1(t_2^0) \geq \frac{1}{2D}\left((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2\right).$$

Proof. Once more, it will be differentiated between two cases regarding the initial relation of the cars' speeds.

Case $w_1'(t_2^0) < w_2'(t_2^0)$. Assume that accident-absent ways exist. Car c_2 appears on t_2^0 at $w_2(t_2^0) = 0$. At this time, c_1 is at $w_1(t_2^0) = w_1'(t_1^0) \cdot (t_2^0 - t_1^0)$. The speed of c_1 at t_2^0 is $w_1'(t_2^0)$ and the speed of c_2 is $w_2'(t_2^0)$.

The minimum time with $w_1'(t_e) - w_2'(t_e) = 0$ is described in Lemma 11 as $t_e = t_2^0 + w_2'(t_2^0)/A$. Accordingly, for $t \in [t_2^0, t_e) : w_2''(t) = w_1''(t) = D$. The location of c_2 at t_e is $-(w_2'(t_2^0))^2/(2D)$; c_1 is slower at t_2^0 and stops at $w_1(t_2^0) - (w_1'(t_2^0))^2/(2D)$. It must be

$$\begin{aligned}
& w_1(t_2^0) - \frac{(w_1'(t_2^0))^2}{2D} \geq -\frac{(w_2'(t_2^0))^2}{2D} \\
\Leftrightarrow \quad & w_1(t_2^0) \geq \frac{(w_1'(t_2^0))^2}{2D} - \frac{(w_2'(t_2^0))^2}{2D} \\
\Leftrightarrow \quad & w_1(t_2^0) \geq \frac{1}{2D}\left((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2\right).
\end{aligned}$$

The other direction: given $w_1(t_2^0) \geq (1/2D)((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2)$. The minimum time where $w_1'(t) - w_2'(t) = 0$ is $t = t_e$. This requires $t \in [t_2^0, t_e)$: $w_2''(t) = D$. The distances driven by the vehicles until a full stop are $-(w_1'(t_2^0))^2/(2D)$ and

$-(w_2'(t_2^0))^2/(2D)$, so the distance between the vehicles changes by $w_1(t_2^0) - w_2(t_2^0) + (-(w_1'(t_2^0))^2 + (w_2'(t_2^0))^2)/(2D)$. This expression must be greater than or equal to zero for an accident-absent way. With $w_2(t_2^0) = 0$, the initial formula results: $w_1(t_2^0) \geq (1/2D)((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2)$. $\forall t \geq t_e : w_2''(t) = 0$, i.e., c_2 remains stopped. For this reason, the way $w_2(t)$ is accident-absent with the way $w_1(t)$. $w_1(t_2^0) \geq (1/2D)((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2) \Rightarrow$ accident-absent ways exist.

Case $w_1'(t_2^0) \geq w_2'(t_2^0)$, c_1 is faster than c_2 at t_2^0. Accident-absent ways exist $\Rightarrow$ $w_1(t_2^0) \geq (1/2D)((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2)$. The formula discussed by the previous case, $(1/2D)((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2)$, is constructed as above and its result is zero or negative in this case. The initial distance between the vehicles, $w_1(t_2^0) - w_2(t_2^0) = w_1(t_2^0)$, is zero or positive. So it holds $w_1(t_2^0) \geq \frac{1}{2D}((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2)$.

The other direction to proof is: $w_1(t_2^0) \geq (1/2D)((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2) \Rightarrow$ accident-absent ways exist. c_1 has an equal or longer way to stop compared to c_2. The minimum distance, thus, is zero or negative. The initial distance must be zero or positive. That means, every allowed initial distance allows accident-absent ways for the cars: one possible strategy for both cars is a way with full deceleration until a full stop and then remain unaccelerated. An initial distance less than zero does not fit to the definition of a lane that starts at position 0, i.e., $w_1(t_2^0) < 0 = w_2(t_2^0)$ is not allowed. Consequently, accident-absent ways exist. □

With that lemma, all the tools are available to state a traveling-optimal way for c_2 in the case of receiving a single update about c_1. The behavior of c_2 described in the following lemma is traveling optimal and accident absent.

Lemma 13. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and let two valid accident-absent ways w_1 and w_2 exist for the respective vehicles. c_2 knows the precise location and speed of c_1 at t_2^0. $\forall t \geq t_2^0$: w_2 is traveling optimal iff*

$$w_2''(t) = \begin{cases} A \text{ if } t \in [t_2^0, t_a) \\ D \text{ if } t \in [t_a, t_b) \\ 0 \text{ if } t \geq t_b \end{cases}$$

with $t_a = t_e - \frac{w_2'(t_e)}{A} + \sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}}$, $t_b = t_s + 2\left(-\frac{w_2'(t_e)}{A} + \sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}}\right)$, $t_e = t_2^0 + \frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A}$, $t_s = t_2^0 + \frac{w_1'(t_2^0)}{A}$, *and* $\Delta w(t_e) = w_1(t_2^0) + A(t_e - t_2^0)^2$.

Proof. The proof is structured into three parts: first, it will be shown that the acceleration behavior follows from w_2 being traveling optimal. Then, the other way will be shown, i.e., w_2 is traveling optimal when the acceleration behavior is given. Finally, the shape of the variables t_a, t_b, t_s, t_e, and $\Delta w(t_e)$ will be derived.

Part 1: given w_2 is traveling optimal. It will be shown that the acceleration behavior $w_2''(t)$ from the lemma follows. c_2 *estimates* c_1 to choose an arbitrary valid way that fits to the initial location $w_1(t_2^0)$ and speed $w_1'(t_2^0)$. The set of accident-absent ways for c_2 consists of ways that are accident absent with each of the possible valid ways of c_1. If w_2 is accident absent with the slowest valid way of c_1 that is discussed in Lemma 7, it is accident absent with every other valid way of c_1, because no valid way of c_1 is behind that way at any time. Be $w_1(t)$ this way with $t_s = t_2^0 + w_1'(t_2^0)/A, \forall t \in [t_2^0, t_s) :$ $w_1''(t) = D$ and $\forall t \geq t_s : w_1''(t) = 0$. In the following, c_1 is regarded as if driving the way w_1.

The statement that w_2 is traveling optimal means: to every point on the lane, w_2 is earlier or at the same at this point compared to every other valid accident-absent way. A danger of accident is apparent when the estimated distance between the vehicles is lower or equal than the *minimum safe distance* of an accident-absent way similar to what is stated in Lemma 12 as criterion for the existence of such ways. The minimum safe distance is calculated as in the proof of Lemma 12 with the adaption of time t_2^0 to an arbitrary point in time:

$$w_1(t) - w_2(t) \geq \frac{1}{2D}((w_1'(t))^2 - (w_2'(t))^2).$$

If no danger of accident is apparent, $t \in [t_2^0, t_a)$, $w_2''(t)$ has to be at maximum, which is A, as explained by Lemma 6. The point in time is to be found when this inequality becomes an equation; $w_1(t) - w_2(t) = (1/2D)((w_1'(t))^2 - (w_2'(t))^2)$. Be this at time $t_a \geq t_2^0$. At t_a, the distance equals the minimum safe distance, and the full braking with D is required to maintain this minimum safe distance while $w_1'(t) < w_2'(t)$. No other way than the slowest possible one that is discussed in Lemma 7 is accident absent in this situation. This becomes evident when regarding the construction of the minimum safe distance. The minimum safe distance as well as the estimated distance are equal at t_a and grow equally during deceleration. The speed of c_2 is reduced for $t \geq t_a$: $w_2'(t) = w_2'(t_a) + D(t - t_a)$. c_1 brakes until full stop and remains unaccelerated, so the speed of c_2 equals that of c_1 for a point in time $t \geq t_a$, at which both cars have stopped. Be t_b this time.

Because of the construction of the minimum safe distance, the estimated distance of c_1 and c_2 is zero at t_b. For $t \geq t_b$, with w_2 being an accident-absent way, c_2 can only accelerate with $w_2''(t) \in [D, 0]$. But the speed of c_2 is zero, so $w_2''(t) < 0$ is no valid way. It must be $\forall t \geq t_b : w_2''(t) = 0$. This is the only accident-absent (and valid) way, and it

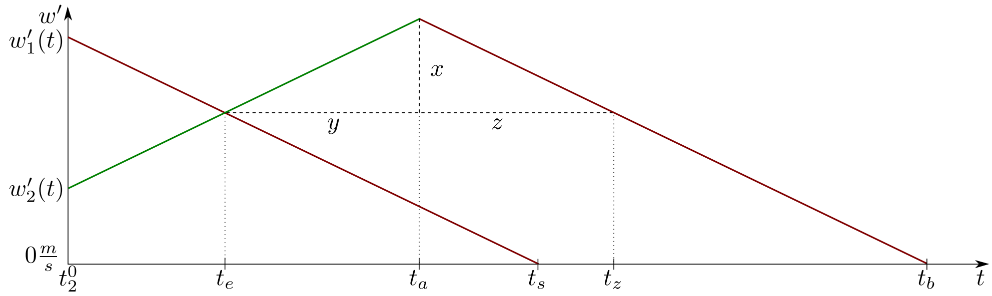

Figure A.2: Speeds of the cars if c_2 only has knowledge about c_1's location and speed at t_2^0. At t_e, the speeds are equal for the first time since the start of c_2. t_a is the point in time at which c_2 starts braking. At time t_b, the steady state is reached.

is, for that reason, also a traveling-optimal way. c_2 then is in a steady state in which it remains stopped. So the behavior of c_2 proposed by the lemma follows from w_2 being traveling optimal. Figure A.2 exemplary depicts the speeds in different acceleration phases of c_2 for the case of $w'_1(t_2^0) > w'_2(t_2^0)$.

Part 2: given the acceleration behavior of $w''_2(t)$ from the lemma, it will be shown that w_2 is traveling optimal. Look at the acceleration phase for $t \in [t_2^0, t_a)$. In general, constant acceleration with A is traveling optimal, because no valid way with another acceleration function is further ahead at any time. But this strategy is not accident absent if the distance is shorter than or equal to the minimum safe distance. If the estimated distance equals this minimum safe distance at t_a, the only acceleration resulting in an accident-absent way is with D. For that reason, it is also the only behavior that results in a traveling-optimal way. At t_b, the speed of c_2 equals that of c_1 and their distance is zero. The change of c_2's acceleration from D to 0 results in c_2 staying at the same location as c_1 with the same speed. This is, obviously, accident absent and, in addition, a traveling-optimal way.

Part 3: having discussed both directions of the iff relationship, the values of t_a, t_b, t_s, t_e and $\Delta w(t_e)$ still require to be justified. Because the existence of the accident-absent ways w_1 and w_2 is required, it holds $w_1(t_2^0) \geq (w'_1(t_2^0))^2/(2D) - (w'_2(t_2^0))^2/(2D)$, i.e., the initial distance is larger than or equal to the estimated minimum safe distance. c_2 accelerates with A until this estimated minimum distance, which depends on the current speeds of the cars, is reached at t_a. First, it is assumed that $w'_1(t_2^0) \geq w'_2(t_2^0)$ as in Figure A.2. On t_e, the speeds are equal, so the formula for the minimum safe distance results zero: $t_e = t_2^0 + (w'_1(t_2^0) - w'_2(t_2^0))/(2A)$, see Lemma 11. For $t \in [t_2^0, t_e)$, the distance grows, since c_1 is faster than c_2. The distance at t_e has to be reduced to zero until t_b for a traveling-optimal way. If c_2 braked to a full stop starting at t_e, both vehicles stopped at $t_s = t_2^0 + w_1'(t_2^0)/A$, but at different locations; their distance is then the same as at t_e. To behaving traveling optimally, c_2 reduces the distance at t_e, $\Delta w(t_e)$, to zero in minimum time. This is done by accelerating until t_a, followed by decelerating until $t_z \in [t_a, t_b]$. At t_z, c_2 again is at the same speed like at t_e. Also, the stopping distance required by c_2 at t_z equals that of c_1 at t_e.

c_2 is not at minimum safe distance at t_e but at t_a, hence acceleration is allowed. Then it decelerates to zero speed; c_2 brakes further from t_z until full stop and the distance is zero at t_b. The slope of c_2's speed for $t \in [t_e, t_a)$ is $A = x/y$ and the slope for $t \in [t_a, t_b)$ is $D = -x/z$. With $A = -D$, according to Definition 5, it follows $z = y$. c_2's distance to the stopping location of c_1 changes from t_e to $t_e + y = t_a$ by the same amount as from t_a to $t_e + y + z = t_e + 2y = t_z$. This can be seen by integrating the speed of c_2 over those intervals. To obtain y, it has to be determined how long it takes to halve the vehicles' distance at t_e. With $y = t_a - t_e$:

$$\begin{aligned}
& & \frac{\Delta w(t_e)}{2} &= \Delta w(t_e) - w_2'(t_e)y - \frac{A}{2}y^2 \\
&\Leftrightarrow & 0 &= y^2 + \frac{2w_2'(t_e)}{A}y - \frac{\Delta w(t_e)}{A} \\
&\Rightarrow & y &= -\frac{w_2'(t_e)}{A} + \sqrt{\left(\frac{w_2'(t_e)}{A}\right)^2 + \frac{\Delta w(t_e)}{A}}.
\end{aligned}$$

y must be positive, so the positive root is chosen.

$$y = t_a - t_e \quad \Rightarrow \quad t_a = t_e - \frac{w_2'(t_e)}{A} + \sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}}.$$

That is the expression for t_a used in the lemma. The difference of the stopping times is the same as the time interval $y + z$: $t_b - t_s = y + z = t_z - t_e$. With $t_z = t_e + 2y$, it is obtained

$$t_b = t_s + 2y = t_s + 2\left(-\frac{w_2'(t_e)}{A} + \sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}}\right).$$

$\Delta w(t_e)$ is easily derived with Lemma 5:

$$\begin{aligned}
\Delta w(t_e) &= \Delta w(t_2^0) + (w_1'(t_2^0) - w_2'(t_2^0))(t_e - t_2^0) + \frac{D}{2}(t_e - t_2^0)^2 - \frac{A}{2}(t_e - t_2^0)^2 \\
&= w_1(t_2^0) - w_2(t_2^0) + (w_1'(t_2^0) - w_2'(t_2^0))(t_e - t_2^0) + \frac{D - A}{2}(t_e - t_2^0)^2.
\end{aligned}$$

With $w_2(t_2^0) = 0$, $D = -A$, and $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/(2A)$:

$$\begin{aligned}
\Delta w(t_e) &= w_1(t_2^0) + (w_1'(t_2^0) - w_2'(t_2^0))\frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A} - A\left(\frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A}\right)^2 \\
&= w_1(t_2^0) + \frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{2A} - A\frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{4A^2} \\
&= w_1(t_2^0) + A\frac{(w_1'(t_2^0) - w_2'(t_2^0))^2}{4A^2} \\
&= w_1(t_2^0) + A(t_e - t_2^0)^2.
\end{aligned}$$

So far, the case $w_1'(t_2^0) \geq w_2'(t_2^0)$ has been discussed. In the other case of $w_1'(t_2^0) < w_2'(t_2^0)$, it is set $t_e < t_2^0$, i.e., $t_2^0 \in [t_e, t_a]$. The formulae for t_a and t_b are still valid for this case. Also, the formula of t_e works without adaption. $\Delta w(t_e)$ describes the distance for $t_e < t_2^0$. Because $w_1'(t_2^0) < w_2'(t_2^0)$ and c_2 accelerates in $[t_e, t_2^0]$, while c_1 decelerates, the distance at t_e is greater than at t_2^0. $\Delta w(t_e)$, therefore, also holds as discussed for the other case. □

Although now the behavior of c_2 has been proved, the proof is not yet finished, because in Theorem 2, t_b is described a bit different from the t_b in Lemma 13:

$$t_b = t_2^0 - \frac{w_2'(t_2^0)}{A} + 2\sqrt{\frac{w_1(t_2^0)}{A} + \frac{w_1'(t_2^0)w_2'(t_2^0)}{A^2} + 2t_e^2}.$$

This expression is obtained by stating t_b in the short form as in the proof of the previous lemma and transforming it:

$$\begin{aligned}
t_b &= t_s + 2y \\
&= t_2^0 + \frac{w_1'(t_2^0)}{A} + 2\left(-\frac{w_2'(t_e)}{A} + \sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}}\right) \\
&= t_2^0 + \frac{w_1'(t_2^0) - 2w_2'(t_2^0)}{A} - \frac{2A}{A}(t_e - t_2^0) + 2\sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}} \\
&= t_2^0 + \frac{w_1'(t_2^0) - 2w_2'(t_2^0)}{A} - 2\left(\frac{w_1'(t_2^0) - w_2'(t_2^0)}{2A}\right) + 2\sqrt{\frac{(w_2'(t_e))^2}{A^2} + \frac{\Delta w(t_e)}{A}} \\
&= t_2^0 - \frac{w_2'(t_2^0)}{A} + 2\sqrt{\frac{(w_2'(t_2^0) + A(t_e - t_2^0))^2}{A^2} + \frac{w_1(t_2^0) + A(t_e - t_2^0)^2}{A}} \\
&= t_2^0 - \frac{w_2'(t_2^0)}{A} + 2\sqrt{\frac{w_1(t_2^0)}{A} + 2(t_e - t_2^0)^2 + \frac{(w_2'(t_2^0))^2}{A^2} + \frac{(w_1'(t_2^0) - w_2'(t_2^0))w_2'(t_2^0)}{A^2}} \\
&= t_2^0 - \frac{w_2'(t_2^0)}{A} + 2\sqrt{\frac{w_1(t_2^0)}{A} + \frac{w_1'(t_2^0)w_2'(t_2^0)}{A^2} + 2(t_e - t_2^0)^2}.
\end{aligned}$$

A further thing that has to be adapted is that Theorem 2 defines $t_e = (w_1'(t_2^0) - w_2'(t_2^0))/(2A)$ instead of $t_e = t_2^0 + (w_1'(t_2^0) - w_2'(t_2^0))/(2A)$. To obtain the formula used in the theorem, t_e has to be changed to a relative value. The t_e of the theorem is redefined as $\hat{t}_e$ and it is stated: $t_e - t_2^0 = \hat{t}_e$. Now, all statements of Theorem 2 are proved.

A.3 Proof of Lemma 1

Lemma 1 states: *If c_2 learns about c_1's position and speed periodically at $t_i, i \in \mathbb{N}_0$, with $\forall\, t_i : t_{i+1} - t_i = const$, then the sequence of distances between both cars and the speeds of c_2 at the points in time t_i converge to a steady state for $i \to \infty$.*

At first, the existence of accident-absent ways is discussed in case of c_2 receiving multiple updates. With regard to the safe area in Figure 3.4 on Page 44, it will be shown that the line between the safe and the dangerous area will be reached by two accident-absent ways. After that, it will be seen that the ways then travel along this line, targeting a steady state that is described by the green point in the figure. It will then be proved that this state is never reached by a way if it does not start in it; the way converges to it.

In the following, it will often be referred to a *difference of minimum stopping distances*, so it will be defined in advance:

Definition 20. Difference of minimum stopping distances. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two valid ways w_1 and w_2 for the respective vehicles. The minimum distances to stop for the cars at time $t \geq t_2^0$ are $(w_1'(t))^2/(2D)$ and $(w_2'(t))^2/(2D)$. The difference of the minimum stopping distances at $t \geq t_2^0$ is*

$$\frac{1}{2D}\left((w_1'(t))^2 - (w_2'(t))^2\right).$$

Lemma 14. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two valid ways w_1 and w_2 for the respective vehicles. c_2 learns about c_1's position and speed at $t_i, i \in \mathbb{N}_0$. Without loss of generality let $t_0 = t_2^0$. Accident-absent ways for c_1 and c_2 exist iff $\forall t_i$:*

$$w_1(t_i) - w_2(t_i) \geq \frac{1}{2D}\left((w_1'(t_i))^2 - (w_2'(t_i))^2\right).$$

Proof. The lemma will be proved with complete induction over t_i. It is started at the first interval $[t_0, t_1)$, with $t_0 = t_2^0$. According to Lemma 12, accident-absent ways exist iff $w_1(t_2^0) \geq (1/2D)\left((w_1'(t_2^0))^2 - (w_2'(t_2^0))^2\right)$.

The induction hypothesis is that accident-absent ways exist iff $w_1(t_i) - w_2(t_i) \geq (1/2D)\left((w_1'(t_i))^2 - (w_2'(t_i))^2\right)$.

Now, it is shown that the lemma holds for t_{i+1}. For this, a way $w_{1,i}(t)$ is constructed for c_1 that has c_1's true state at t_i which was communicated to c_2. From then on, the state of c_1 is estimated by c_2 from the update at t_i within the interval $[t_i, t_{i+1})$: $w_{1,i}(t)$ equals the slowest valid way for c_1 as defined in Lemma 7 with the adaption to start braking at t_i instead of at t_2^0. $\forall t \geq t_i$: $w''_{1,i}(t) = D$ if $w'_{1,i}(t) > 0$, otherwise $w''_{1,i}(t) = 0$. No valid way is closer to c_2 at any time, i.e., if w_2 is accident-absent to this way, it is accident-absent with all possible valid ways of c_1.

The first direction: given accident-absent ways for c_1 and c_2. Let c_1 drive with any valid way. At time t_{i+1}, the distance between the ways must be accident absent with the old knowledge before the update of t_{i+1} is processed. The update reveals to c_2 that c_1 is at a larger or equal distance and faster or equally fast compared to the assumed worst case. The position and speed are replaced with the true values. It holds $w_1(t_{i+1}) \geq w_{1,i}(t_{i+1})$ and $w_1'(t_{i+1}) \geq w'_{1,i}(t_{i+1})$. Therefore,

$$\begin{aligned}
w_1(t_{i+1}) - w_2(t_{i+1}) &\geq w_{1,i}(t_{i+1}) - w_2(t_{i+1}) \\
&\geq \frac{1}{2D}\left((w'_{1,i}(t_{i+1}))^2 - (w_2'(t_{i+1}))^2\right) \\
&\geq \frac{1}{2D}\left((w_1'(t_{i+1}))^2 - (w_2'(t_{i+1}))^2\right).
\end{aligned}$$

The other direction: given $w_1(t_i) - w_2(t_i) \geq 1/(2D)((w_1'(t_i))^2 - (w_2'(t_i))^2)$. It will be shown that accident-absent ways exist. Assume c_2 brakes fully in the interval $[t_i, t_{i+1})$, then $w_{1,i}(t_{i+1}) - w_2(t_{i+1}) \geq 1/(2D)((w'_{1,i}(t_{i+1}))^2 - (w_2'(t_{i+1}))^2)$: the estimated distance at t_{i+1} is larger or equal compared to the difference of the minimum stopping distances calculated with the speeds at t_{i+1}. This is because the construction of the difference of the stopping distances: if both vehicles brake to a full stopp, their distance is reduced by this difference. With the update at t_{i+1}, $w_{1,i}(t) \leq w_1(t)$ and $w'_{1,i}(t) \leq w_1'(t)$. Therefore, $w_1(t_{i+1}) - w_2(t_{i+1}) \geq 1/(2D)((w_1'(t_{i+1}))^2 - (w_2'(t_{i+1}))^2)$, i.e., the distance is larger than the difference of the minimum stopping distances. Thus, at least one accident-absent way exists: c_2 brakes until a full stop starting at t_{i+1} and remains unaccelerated. □

It will be seen next that c_2 always reaches a minimum stopping distance to c_2 if c_1 travels with a constant speed.

Lemma 15. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two accident-absent ways w_1 and w_2 for the respective vehicles. Let $w_1'(t) = const$ and let w_2 be traveling optimal.*

c_2 learns about c_1's position and speed at $t_i, i \in \mathbb{N}_0$ with $\forall t_i : t_{i+1} - t_i = const.$ Without loss of generality let $t_0 = t_2^0$. $t \in [t_i, t_{i+1})$: let $w_{1,i}$ be the way of c_1 as estimated by c_2. $\exists t_i, \exists t_{a,i} \in [t_i, t_{i+1})$:

$$w_{1,i}(t_{a,i}) - w_2(t_{a,i}) = \frac{1}{2D}\left((w'_{1,i}(t_{a,i}))^2 - (w'_2(t_{a,i}))^2\right).$$

Proof. $t_{a,i}$ is considered to be the earliest time for that the equality of the lemma's formula holds, calculated with the knowledge c_2 gets from the ith update. The acceleration of the way $w_{1,i}(t)$ for c_1 is similar to the slowest valid way for c_1 as defined in Lemma 7. The way from that lemma is adapted because the interest is in time t_i and onwards. For this, as the state of c_1 at t_i is known, t_1^0 is simply replaced in the lemma by t_i. $\forall t \geq t_i$: $w''_{1,i}(t) = D$ if $w'_{1,i}(t) > 0$, otherwise $w''_{1,i}(t) = 0$. Three cases are differentiated based on the initial distance of the cars.

Case $w_1(t_2^0) < (1/2D)((w'_1(t_2^0))^2 - (w'_2(t_2^0))^2)$ is not relevant: the ways have to be accident absent.

Case $w_1(t_2^0) = (1/2D)((w'_1(t_2^0))^2 - (w'_2(t_2^0))^2)$. Then it is $t_{a,0} = t_2^0$.

Case $w_1(t_2^0) > (1/2D)((w'_1(t_2^0))^2 - (w'_2(t_2^0))^2)$. Between the successive updates with the indices i and $i+1$, c_2 accelerates (because w_2 is traveling optimal), while c_1 is treated as if decelerating with the way $w_{1,i}(t)$. The thereby estimated difference of the minimum stopping distances grows during the interval, given by the following formula:

$$\frac{1}{2D}((w'_{1,i}(t_i) - A(t-t_i))^2 - (w'_2(t_i) + A(t-t_i))^2).$$

How does the estimated distance change? If c_2 is equally fast or faster compared to c_1 at t_i, then the estimated distance $w_{1,i}(t) - w_2(t)$ shrinks. Otherwise, the estimated distance grows at first until the vehicles are equally fast; from then on, the estimated distance shrinks. If c_1 stops decelerating at zero velocity, c_2 accelerates further. Obviously, it exists a time $t_{a,i}$, at which the estimated distance equals the estimated difference of the minimum stopping distances. But only if $t_{i+1} > t_{a,i}$, the estimated braking difference is reached within this interval $[t_i, t_{i+1})$.

What happens if an update is received at c_2? Let $i > 0$: on the arrival of an update at t_i, the information about c_1 is reset to the true values of driving with constant speed: the true c_1 is faster and further ahead than estimated.

In every following interval, the acceleration process is repeated. For this reason, the speed of c_2 is larger or equal to that of c_1 after a finite number of intervals of acceleration. From then on, the true distance shrinks at each successive interval. This causes the estimated minimum stopping distances difference at t_i to be larger than at t_{i-1}:

$$\frac{1}{2D}((w_1'(t_i))^2 - (w_2'(t_i))^2) > \frac{1}{2D}((w_1'(t_{i-1}))^2 - (w_2'(t_{i-1}))^2).$$

Because of this, the time from the update until the estimated difference of minimum stopping distances equals the estimated distance is shorter than in the previous interval: $t_{a,i-1} - t_{i-1} > t_{a,i} - t_i$. The time between t_i and $t_{a,i}$ is shorter at each successive interval until for an interval j: $t_{j+1} > t_{a,j}$, i.e., $t_{a,j} \in [t_j, t_{j+1})$. □

The arrows in the left, white zone (the safe one) of Figure 3.4 on Page 44 illustrate the movement: they point towards the right zone, i.e., from every interval to the next one, the speed difference $w_2' - w_1'$ grows as long as the border line is not reached by an arrow's top. The line between the left and right (red) zones describes that the vehicles are on a distance equal to the difference of their stopping distances. Being on the line is still an accident-absent way with the described behavior of c_2. An arrow starting on the line with a negative speed difference means that although c_2 is slower than the true c_1 at the beginning of that interval, the vehicles have been on the estimated difference of braking distances at the end of the previous interval.

From the previously discussed lemma, it is known that the figure's line between the safe and the dangerous zone will be reached by the cars' ways. Next, it is shown that once the cars are on this line at the end of one interval, they will be on it at the end of the next interval, i.e., the cars stay on that line.

Lemma 16. *Given two cars c_1 and c_2 with $c_2 <_c c_1$ and two accident-absent ways w_1 and w_2 for the respective vehicles. Let $w_1'(t) = const$ and let w_2 be traveling optimal. c_2 learns about c_1's position and speed at $t_i, i \in \mathbb{N}_0$ with $\forall t_i : t_{i+1} - t_i = const$. Without loss of generality let $t_0 = t_2^0$. $t \in [t_i, t_{i+1})$: let $w_{1,i}$ be the way of c_1 as estimated by c_2. If at the end of interval i:*

$$w_{1,i}(t_{i+1}) - w_2(t_{i+1}) = \frac{1}{2D}\left((w_{1,i}'(t_{i+1}))^2 - (w_2'(t_{i+1}))^2\right)$$

then at the end of interval $i+1$:

$$w_{1,i+1}(t_{i+2}) - w_2(t_{i+2}) = \frac{1}{2D}\left((w_{1,i+1}'(t_{i+2}))^2 - (w_2'(t_{i+2}))^2\right).$$

Proof. With regard to Figure 3.4, the lemma states that if an arrow points on the edge line, then the next arrow starting from the previous cars' states will also point on the line. When the difference of minimum stopping distances equals the estimated distance within an interval, it does so until the end of that interval: the arrow points on the line. It will be argued that this is again the case in the following interval.

Assume this is not true; the difference of minimum stopping distances equals the estimated distance at the end of interval $[t_i, t_{i+1})$ but not at the end of the next, $[t_{i+1}, t_{i+2})$. At the end of the interval $[t_i, t_{i+1})$, the position and the speed of c_1 are updated at c_2 to the true values. The difference of the *estimated* minimum stopping distances is thereby replaced by a smaller value: c_1 is faster than in the case of braking during the previous interval. The estimated distance is replaced by the larger true value: c_1 is truly further ahead.

c_2, driving traveling optimally in $[t_{i+1}, t_{i+2})$, accelerates until being on the difference of minimum stopping distances at $t_{a,i+1}$ which is calculated with the knowledge from the $(i+1)$th update of c_1. The assumption implies $t_{a,i+1} \geq t_{i+2}$. At the end of the interval, at t_{i+2}, the estimated speed of c_1 equals the estimated speed at the end of the previous interval (it braked for the same duration from the same initial constant speed), but c_2 is faster than at the previous interval's end (because it accelerated the whole interval). Let $B = t_i - t_{i+1}$: the speed of c_1 is estimated to linearly decrease from $w'_{1,i+1}(t_{i+1}) = w'_1(t_i)$ to $w'_{1,i+1}(t_{i+2}) = w'_1(t_i) - AB$. Hence the average speed of c_1 in that interval is $w'_1(t_i) - AB/2$. c_2 has an initial speed of $w'_1(t_i) - AB$ or higher because, at the previous interval's end, the estimated speed difference $w'_2(t_{i+1}) - w'_{1,i}(t_{i+1})$ is greater or equal to zero: the cars are on their minimum stopping distances difference at t_{i+1}. Furthermore, c_2 is assumed to accelerate the whole interval. Its average velocity is, thus, higher or equal $w'_1(t_i) - AB/2$. Because of this, the estimated distance between the cars from one interval's end to the next either stays the same or is shortened. The higher average speed of c_2 causes the difference of the estimated stopping distances to be larger at t_{i+2} than at the previous interval's end.

So, to conclude, at t_{i+1}, the estimated distance and the estimated stopping distance difference are equal. But the estimated distance is shorter or equal at t_{i+2} than at t_{i+1}, while the stopping difference is larger at t_{i+2} than at t_{i+1}: w_2 is not an accident-absent way. □

In Lemma 15, it is discussed that in each sequence of (speed difference, distance) pairs, there exists a pair for which the difference of the estimated minimum stopping distances equals the estimated distance at the end of the interval described by the pair. In Lemma 16, it is stated that once the equality of these two distances is given at the end of one interval, the equality of the two distances holds at the end of every following interval. Compared to Figure 3.4 which shows the interval beginnings, the pairs move along the edge line. The car c_2 estimates itself to be on a minimum distance for safe braking at the end of each interval. The lemma to proof in this section says that the cars strive to a certain point of the (speed difference, distance) pairs that are referred to as the steady state. Now it is possible to prove that initial assumption that is stated in Lemma 1 on Page 44.

Proof. The movement along the line shown in the figure is discussed, differentiated in three cases: 1) the cars are in the steady state and stay in this state; 2) the speed difference and distance are higher than in the steady state; 3) the speed difference and distance are lower than in the steady state. For this discussion, let $B = t_{i+1} - t_i$ and let $t_i + t_{a,i}$ be the earliest time ($\geq t_i$) where the cars are on the difference of minimum stopping distances, calculated with knowledge from the ith update of c_1. Note that $t_{a,i}$ is defined in this proof as a time relative to t_i.

Case 1: the cars are in the steady state and stay in this case. For now, it is simply assumed that such a state exists; a formal discussion about its characteristics can be found in the proof of Theorem 3. The term steady state means for this proof: if a sequence of (speed difference, distance) pairs contains the steady-state pair as element $i \in \mathbb{N}$, then every element $j \geq i$ equals that pair.

Case 2: the speed difference and distance are higher than in the steady state. It will be shown that the speed difference becomes less at the beginning of each next interval and converges to the steady state difference within an infinite number of intervals. The speed of c_2 evolves from one update to the next depending on whether the difference of minimum stopping distances is reached within an interval or not:

$$w_2'(t_{i+1}) = \begin{cases} w_2'(t_i) + BA & \text{if } t_{a,i} \geq B \\ w_2'(t_i) + t_{a,i}A + (B - t_{a,i})D & \text{else.} \end{cases}$$

The else part is simplified with $D = -A$:

$$w_2'(t_i) + t_{a,i}A + (B - t_{a,i})D = w_2'(t_i) + 2At_{a,i} - BA.$$

The condition $t_{a,i} \geq B$ can be tested by calculating if the estimated distance at the end of the interval is greater than or equal to the estimated difference of stopping distances, given that c_2 accelerates during the whole interval:

$$w_1(t_i) - w_2(t_i) + (w_1'(t_i) - w_2'(t_i))B + \frac{D-A}{2}B^2 \geq \frac{1}{2D}\left((w_1'(t_i) - AB)^2 - (w_2'(t_i) + AB)^2\right)$$

It is known that the edge line will be reached and once reached, it will not be left. Therefore, (speed difference, distance) pairs on the edge line are considered only. Then it holds $t_{a,i} < B$. Furthermore, it is $t_{a,i} < B/2$, because if otherwise: at t_i, c_2 is faster than in the steady state and the cars are on the difference of the estimated stopping distances at the end of the previous interval. c_2 approaches c_1 within an interval when accelerating for $B/2$ (or longer); the distance is shortened during the interval. So at the end of this interval, c_2 will be nearer but equally fast (or faster) compared to the previous interval, i.e., nearer to c_1 than the difference of the estimated minimum braking distances. Therefore, this way is not accident absent.

Hence, c_2 has to slow down in the interval, which requires $t_{a,i} < B/2$. The speed of c_2 is monotonically decreasing; $w'_2(t_{i+1}) < w'_2(t_i)$. It will now be argued that the speed of c_2 decreases until it equals the true speed of c_1 on an interval's average, which is possible after an infinite number of intervals; then $t_{a,i} = B/2$.

Be $(w'_2)_i$ the sequence of speeds of c_2 at the beginning of the intervals; the ith element of the sequence is $w'_2(i) = w'_2(t_i)$. Assume the speed at the beginning of an interval in the infinite is: $\lim_{i\to\infty}(w'_2)_i = w'_1(0) - (1/4)BA$. Given $\epsilon > 0$. With formula (A.3), it is $|(w'_2)_i - (w'_1(0) - (1/4)BA)| = |(w'_2(i) + 2At_{a,i} - BA) - w'_1(0) + (1/4)BA| = w'_2(i) - w'_1(0) + 2At_{a,i} - (3/4)BA < \epsilon$. This inequality holds iff $w'_2(i) < \epsilon + w'_1(0) - 2At_{a,i} + (3/4)BA$. Be $w'_2(i_0) < \epsilon + w'_1(0) - 2At_{a,i} + (3/4)BA$. Then $|(w'_2)_i - (w'_1(0) - (1/4)BA)| < \epsilon$ for each $i \geq i_0$. With $t_{a,i} < (B/2)$: $|(w'_2)_i - (w'_1(0) - (1/4)BA)| = w'_2(i) - w'_1(0) + 2At_{a,i} - (3/4)BA < w'_2(i) - w'_1(0) + (1/4)BA < \epsilon$. This last inequality holds iff $w'_2(i) < \epsilon + w'_1(0) - (1/4)BA$.

As the speed difference shortens, the distance shortens, too, because c_2 stays faster than c_1 on average within an interval. To what distance the sequence of distances at the beginning of an interval converges is discussed in the proof of Theorem 3.

Case 3: the speed difference and distance are less than in the steady state. It will be seen that the speed difference becomes larger at the beginning of each next interval and converges to the steady state difference after an infinite number of intervals. Proving this is done in a similar way to case 2. Again, it is $t_{a,i} < B$. But now, $t_{a,i} > B/2$. Assume otherwise: c_2 is slower than in the steady state at t_i and the vehicles are on the estimated difference of minimum stopping distances at the end of the previous interval. When c_2 accelerates for $B/2$ (or shorter), the distance grows from t_i to t_{i+1} and the cars will not be on the estimated difference of minimum stopping distances at the end of the ith interval, i.e., the way of c_2 is not traveling optimal.

Hence, $w'_2(t_{i+1}) > w'_2(t_i)$. The speed increases monotonically; so does the difference of minimum stopping distances. It will now be argued that the speed of c_2 increases until it equals the true speed of c_1 on an interval's average after an infinite number of intervals; then $t_{a,i} = B/2$.

Be $(w_2')_i$ the sequence of speeds of c_2 at the beginning of the intervals. Assume the speed at the beginning of an interval in the infinite is: $\lim_{i\to\infty}(w_2')_i = w_1'(0)-(1/4)BA$. Given $\epsilon > 0$. It is $|(w_1'(0) - (1/4)BA) - (w_2')_i| = |(w_1'(0) - (1/4)BA) - (w_2'(i) + 2At_{a,i} - BA)| = w_1'(0) - w_2'(i) - 2At_{a,i} + (3/4)BA < \epsilon$. This inequality holds iff $w_2'(i) > -\epsilon + w_1'(0) - 2At_{a,i} + (3/4)BA$. Be $w_2'(i_0) > -\epsilon + w_1'(0) - 2At_{a,i} + (3/4)BA$. Then $|(w_1'(0) - (1/4)BA) - (w_2')_i| < \epsilon$ for each $i \geq i_0$. With $t_{a,i} > B/2$: $|(w_1'(0) - (1/4)BA) - (w_2')_i| = w_1'(0) - w_2'(i) - 2At_{a,i} + (3/4)BA < w_1'(0) - w_2'(i) - (1/4)BA < \epsilon$. This last inequality holds iff $w_2'(i) > w_1'(0) - (1/4)BA - \epsilon$. □

A.3.1 Calculation Rules of Figure 3.4

The Figure 3.4 on page 44 shows a plot of points in the space of (speed difference, distance) pairs. The arrows describe how the pairs evolve from the beginning of one beaconing interval to the beginning of the next. Let these intervals be i and $i+1$. The calculation rules of the arrows are now explained. The speed is obtained by this formula:

$$w_2'(i+1) = \begin{cases} w_2'(i) + BA & \text{if } \Delta w(i+1) \geq \Delta w_{\text{safe}}(i+1) \\ w_2'(i) + 2At_{a,i} - BA & \text{else.} \end{cases}$$

The formula differentiates whether c_2 is on the difference of the minimum stopping distances at the end of the interval or not. The else part was simplified as shown in Equation A.3 of the previous proof. The condition of the cases can be calculated with knowledge about the ith interval:

$$\begin{gathered} \Delta w(i+1) \geq \Delta w_{\text{safe}}(i+1) \\ \Leftrightarrow w_1(i) - w_2(i) + w_1'B - w_2'(i)B - \frac{A}{2}B^2 \geq \frac{1}{2D}\left((w_1')^2 - (w_2'(i) + AB)^2\right). \end{gathered}$$

Here, w_1' is the constant speed of c_1. The distance between the cars is calculated like this:

$$\begin{aligned} w_1(i+1) - w_2(i+1) = & w_1(i) + w_1'B - w_2(i) - w_2'(i)B \\ & - \begin{cases} \frac{A}{2}B^2 & \text{if } \Delta w(i+1) \geq \Delta w_{\text{safe}}(i+1) \\ (2ABt_{a,i} - At_{a,i}^2 - \frac{A}{2}B^2) & \text{else.} \end{cases} \end{aligned}$$

This formula also distinguishes between the two cases if c_2 reached the estimated minimum braking distance to c_1 or not. The position $w_1(i)$ of c_1 at the beginning of the interval is its true position. The if-part describes the position development if c_2 does not have to brake. The else-part is simplified: $w_1(i+1)-w_2(i+1) = w_1(i)+w_1'B-w_2(i+1)$. It holds:

$$
\begin{aligned}
&w_2(i+1)\\
=&w_2(i) + w_2'(i)t_{a,i} + \frac{A}{2}t_{a,i}^2 + (w_2'(i) + At_{a,i})(B - t_{a,i}) + \frac{D}{2}(B - t_{a,i})^2\\
=&w_2(i) + w_2'(i)t_{a,i} + \frac{A}{2}t_{a,i}^2 + w_2'(i)B - w_2'(i)t_{a,i} + ABt_{a,i} - At_{a,i}^2 - \frac{A}{2}(B^2 - 2Bt_{a,i} + t_{a,i}^2)\\
=&w_2(i) + w_2'(i)B + 2ABt_{a,i} - At_{a,i}^2 - \frac{A}{2}B^2.
\end{aligned}
$$

This is also the true distance at the interval ends; the estimated distance at the end of an interval is shorter by $AB^2/2$ since c_1 is estimated to brake.

A.4 Proof of Theorem 3

Theorem 3 states: *The steady state distance between c_1 and c_2 at the beginning of each beacon period is given by $3Bw_1'(t_2^0)/4 + AB^2/32$, where B is the beacon interval length (i. e., the time between information updates arriving at c_2).*

Proof. The true speed of c_1 is constant. Let the minimum stopping time of c_1 be longer than the time between two beacons, $t_{i+1} - t_i = B$.

A steady state regarding the speed and location difference of the cars is easily constructed: c_2 is able to maintain the same speed and distance to c_1 at every beginning of an interval. Maintaining the same speed with the full acceleration/full deceleration behavior implies that c_2 accelerates the first half of the interval and decelerates the second half, because no other information is available to c_2 that could require intermediate acceleration or deceleration phases. Braking of c_2 at the interval's half is possible if the cars are on the difference of the minimum stopping distances. The braking then has to continue until the end of the interval, starting at $t_i + B/2$ until t_{i+1}.

To maintain a constant distance at the beginning of each interval, c_2 has to be at the same speed as the true way of c_1 on average within an interval; c_2 has to be slower than c_1 at the beginning, faster at the interval's half, and at the end as slow as at the beginning.

With $t_{a,i} = B/2$, c_2 must be as fast as the real c_1 at $t_i + B/4$ and $t_i + 3B/4$. At the latter time, the estimated distance equals the estimated difference of minimum stopping distances and, therefore, the true distance equals the estimated difference of minimum stopping distances plus the difference between the estimated position of c_1

and its true position. For $t \in [t_i, t_{i+1})$: let $w_{1,i}$ be the way of c_1 as estimated by c_2. Then the distance at $t_i + 3B/4$ is:

$$\begin{aligned}
&w_1\left(t_i + \frac{3B}{4}) - w_2(t_i + \frac{3B}{4}\right) \\
=&\frac{1}{2D}\left(\left(w'_{1,i}\left(t_i + \frac{3B}{4}\right)\right)^2 - \left(w'_2\left(t_i + \frac{3B}{4}\right)\right)^2\right) + \left(w_1\left(t_i + \frac{3B}{4}\right) - w_{1,i}\left(t_i + \frac{3B}{4}\right)\right) \\
=&\frac{1}{2D}\left(\left(w'_1(t_i) + \frac{3B}{4}D\right)^2 - \left(w'_1(t_i)\right)^2\right) - \frac{D}{2}\left(\frac{3B}{4}\right)^2 \\
=&w'_1(t_i)\frac{3B}{4}
\end{aligned}$$

With this, the distance at $t_i + B$ is easily calculated. c_1 and c_2 have the same speed at $t_i + 3B/4$, c_1 keeps the speed, and c_2 decelerates for $B/4$:

$$\begin{aligned}
w_1(t_i + B) - w_2(t_i + B) =&w'_1(t_i)\frac{3B}{4} - \frac{D}{2}\left(\frac{B}{4}\right)^2 \\
=&w'_1(t_i)\frac{3B}{4} + \frac{AB^2}{32}
\end{aligned}$$

Because of the steady state, this distance equals that of the beginning of the interval $[t_i, t_{i+1})$. This is verified by calculating the distance at t_i. For this, at first, the distance at $t_i + B/2$ is derived from the distance at $t_i + 3B/4$:

$$\begin{aligned}
&w_1\left(t_i + \frac{B}{2}\right) - w_2\left(t_i + \frac{B}{2}\right) \\
=&w'_1(i)\frac{3B}{4} - \left(w'_1\left(t_i + \frac{B}{2}\right)\frac{B}{4} - \left(w'_2\left(t_i + \frac{B}{2}\right)\frac{B}{4} + \frac{D}{2}\left(\frac{B}{4}\right)^2\right)\right) \\
=&w'_1(t_i)\frac{3B}{4} - \left(w'_t(t_i)\frac{B}{4} - \left(\left(w'_t(t_i) + \frac{AB}{4}\right)\frac{B}{4} + \frac{DB^2}{32}\right)\right) \\
=&w'_1(t_i)\frac{3B}{4} + \frac{AB}{4}\frac{B}{4} + \frac{DB^2}{32} \\
=&w'_1(t_i)\frac{3B}{4} + \frac{AB^2}{32}
\end{aligned}$$

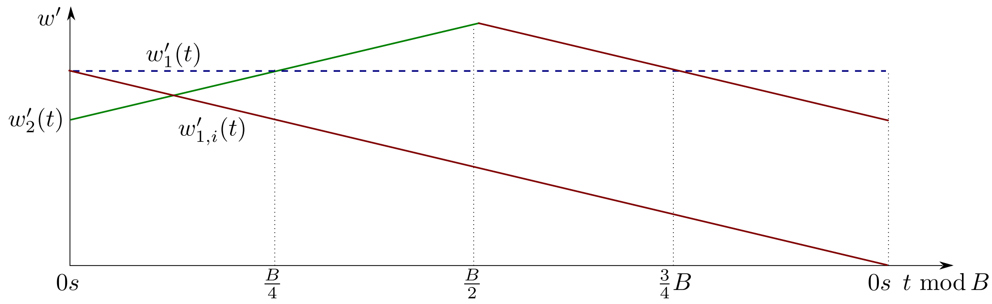

Figure A.3: Oscillating velocities of two cars in the steady state with periodic beaconing within one interval $[t_i, t_{i+1})$. $w'_1(t)$ is the true speed of c_1, while $w'_{1,i}(t)$ is the speed as estimated by c_2.

Note that the distance at $t_i + B/2$ equals that at $t_i + B$. Now the distance at t_i is obtained:

$$\begin{aligned}
&w_1(t_i) - w_2(t_i)\\
=&w_1\left(t_i + \frac{B}{2}\right) - w_2\left(t_i + \frac{B}{2}\right) - \left((w'_1(t_i) - w'_2(t_i))\frac{B}{2} - \frac{A}{2}\left(\frac{B}{2}\right)^2\right)\\
=&w'_1(t_i)\frac{3B}{4} + \frac{AB^2}{32} - \left(\left(w'_{1,t}(t_i) - \left(w'_{1,t}(t_i) - \frac{AB}{4}\right)\right)\frac{B}{2} - \frac{AB^2}{8}\right)\\
=&w'_1(t_i)\frac{3B}{4} + \frac{AB^2}{32}
\end{aligned}$$

This is the same distance as at the end of the interval, as has been expected. Figure A.3 shows the speeds of the vehicles within an interval of the steady state. With $w'_1(t_i) = w'_1(t_2^0)$, the distance stated in the theorem is obtained: $(3B/4)w'_1(t_2^0) + AB^2/32$. □

The distance between the cars can be described by their individual movement equations as a function of time. Let i be the time of the beginning of the ith interval. Then for $t \in [i, i + B)$ and with $D = -A$:

$$\begin{aligned}
&w_1(t) - w_2(t) = w_1(i) - w_2(i)\\
&+\begin{cases} \left(w'_1(i) - (w'_1(i) - \frac{1}{4}BA)\right)(t \mod B) - \frac{1}{2}A(t \mod B)^2 & \text{if } (t \mod B) < \frac{B}{2}\\ \left(w'_1(i) - (w'_1(i) + \frac{1}{4}BA)\right)\left((t \mod B) - \frac{B}{2}\right) + \frac{1}{2}A\left((t \mod B) - \frac{B}{2}\right)^2 & \text{else.}\end{cases}\\
&= w_1(i) - w_2(i)\begin{cases} +\frac{1}{4}BA\left(t \mod \frac{B}{2}\right) - \frac{1}{2}A\left(t \mod \frac{B}{2}\right)^2 & \text{if } (t \mod B) < \frac{B}{2}\\ -\frac{1}{4}BA\left(t \mod \frac{B}{2}\right) + \frac{1}{2}A\left(t \mod \frac{B}{2}\right)^2 & \text{else.}\end{cases}
\end{aligned}$$

Here, $w_1(i) - w_2(i) = (3/4)Bw_1'(i) + AB^2/32$. On deriving the distance function, the speed difference function (of the cars' true speeds) is obtained:

$$w_1'(t) - w_2'(t) = \begin{cases} -A(t \mod \frac{B}{2}) + \frac{1}{4}BA \text{ if } (t \mod B) < \frac{B}{2} \\ +A(t \mod \frac{B}{2}) - \frac{1}{4}BA \text{ if } (t \mod B) \geq \frac{B}{2} \end{cases}$$

On adding $w_1'(t_2^0)$ to this function, the speed function of c_2 in the steady state is obtained.

A.5 Proof of Theorem 4

Theorem 4 states: *Given three cars, $c_1 <_c c_2 <_c c_3$, that drive in their respective steady states. c_1 sends updates at iB_1, $i \in \mathbb{N}_0$. The optimal times for c_2 sending updates to minimize the steady state distance between c_2 and c_3 are $iB_1 + B_1/2$.*

Proof. Given that c_2 is in the steady state to c_1, the true behavior of c_2 is an oscillation of speed and distance to c_1 with frequency $1/B_1$. Theorem 3 states the distance in the steady state and Lemma 1 tells that if c_2 does not start in the steady state distance and speed, c_2 gets arbitrarily close to the steady state by converging to it.

c_3 is responsible for accident-absent driving with c_2 being ahead of it. c_3 estimates the acceleration behavior of its predecessor and has to rely on update-based knowledge through beacons. The estimation by c_3 about c_2 takes all possible ways into account and c_3 has to adapt to the possible ways of c_2 by choosing an accident-absent way. The behavior of c_2 to be assumed is the worst-case from the viewpoint of c_3. This is c_2 braking to a full stop directly after sending a beacon.

For c_3 to follow c_2 in a steady state, it requires periodic updates, too. c_1 and c_2 both send with the same bandwidth, i.e., their sending intervals are equal, but they have an offset. In contrast to the sending times of c_1, the offset of sending of c_2 to the time of sending of c_1 is important for the distance between c_2 and c_3. The minimum safety distance, at which c_3 has to brake, depends on the speeds of the cars. In contrast to the speed of c_1, c_2's speed is not constant between two beacon sendings with the interval B_1 although the fixed sending offset of c_2 regarding c_1 causes c_2 to always send the same speed within its updates.

From now on, the time is considered modulo the sending interval of c_1, $t \mod B_1$, such that c_1 sends at 0. $\forall t \geq t_1^0$: $w_1'(t)$ is the true, constant speed of c_1. The true speed of c_2 is an acceleration changing from A to D within each update interval at $B_1/2$.

Updates sent by c_2 contain the speed of c_2, which is in $[w'_1(t) - AB_1/4, w'_1(t) + AB_1/4]$ and the position, which is in $[w_1(t) - ((3/4)B_1 w'_1(t) + AB_1^2/16), w_1(t) - ((3/4)B_1 w'_1(t))]$; the values are oscillating with a frequency of $1/B_1$. Let $w'_{\min} = w'_1(t) - AB_1/4$ describe the minimum velocity of c_2. Then for $t \in [0, B)$:

$$w'_2(t) = \begin{cases} w'_{\min} + At & \text{if } t \in \left[0, \frac{B}{2}\right) \\ w'_{\min} + AB - At & \text{if } t \in \left[\frac{B}{2}, B\right). \end{cases} \tag{A.1}$$

The true average speed of c_2 is that of c_1: the location of c_2 grows linearly with $w'_1(t)B_1$.

Let t_s be the sending time of c_2. In order to proof the theorem it will be shown that the distance between c_2 and c_3 is minimum during the whole sending interval if c_2 sends at $B_1/2$. Towards this, the formula $\Delta w_{t_s}(t)$ describing the distance between the cars depending on c_2's sending time t_s will be derived at first. Several considerations are necessary until that formula is obtained; these are grouped in Part A of this proof. In Part B, the optimal t_s is determined.

Part A c_3 is required to have the same speed at t_s and at $t_s + B_1 \mod B_1$ to drive with the same true average speed as c_2 in the steady state. On the update at 0, c_3 starts accelerating until t_a and then decelerates until $t_s + B_1$. c_3 gets the true speed of c_2 at t_s and estimates the velocity of c_2 at t_a to be

$$w'_{2,\text{est}}(t_a) = w'_2(t_s) - \frac{1}{2}AB_1 = \begin{cases} w'_{\min} + At_s - \frac{1}{2}AB_1 & \text{if } t_s \in \left[0, \frac{B_1}{2}\right) \\ w'_{\min} + \frac{1}{2}AB_1 - At_s & \text{if } t_s \in \left[\frac{B_1}{2}, B_1\right). \end{cases}$$

The estimated distance between the cars at t_a depending on t_s is

$$\begin{aligned} \Delta w_{a,\text{est}}(t_s) &= \frac{1}{2A}(w'_3(t_s)^2 - w'_{2,\text{est}}(t_a)^2) \\ &= \begin{cases} -\frac{1}{2}At_s^2 + \frac{1}{2}AB_1 t_s - w'_{\min} t_s + B_1 w'_{\min} & \text{if } t \in \left[0, \frac{B_1}{2}\right) \\ -\frac{1}{2}At_s^2 + \frac{1}{2}AB_1 t_s + w'_{\min} t_s & \text{if } t \in \left[\frac{B_1}{2}, B_1\right). \end{cases} \end{aligned}$$

The distance error between the estimated and the true distance of c_3 for $t_s \in [0, B_1/2)$ is

$$\begin{aligned}
\Delta w_{3,err}(t_s) &= \int_{t_s}^{t_a} w_2'(t) - w_{2,\mathrm{est}}'(t)\mathrm{d}t \\
&= \int_{t_s}^{\frac{B_1}{2}} w_2'(t)\mathrm{d}t + \int_{\frac{B_1}{2}}^{t_s+\frac{B_1}{2}} w_2'(t)\mathrm{d}t - \int_{t_s}^{t_s+\frac{B_1}{2}} w_{2,\mathrm{est}}'(t)\mathrm{d}t \\
&= -At_s^2 + \frac{1}{4}AB_1^2
\end{aligned}$$

and similarly for $t_s \in [B_1/2, B_1)$:

$$\Delta w_{3,err}(t_s) = At_s^2 - AB_1t_s + \frac{1}{4}AB_1^2.$$

With the estimated distance and the distance estimation error, the true distance between c_2 and c_3 is

$$\begin{aligned}
\Delta w_a(t_s) &= \Delta w_{a,\mathrm{est}}(t_s) + \Delta w_{3,err}(t_s) \\
&= \begin{cases} -\frac{3}{2}At_s^2 + \frac{1}{2}AB_1t_s - w_{\min}'t_s + \frac{1}{4}AB_1^2 + B_1w_{\min}' & \text{if } t_s \in \left[0, \frac{B_1}{2}\right) \\ \frac{1}{2}At_s^2 - \frac{1}{2}AB_1t_s + \frac{1}{4}AB_1^2 + w_{\min}'t_s & \text{if } t_s \in \left[\frac{B_1}{2}, B_1\right). \end{cases}
\end{aligned}$$

The way of c_2 driving in the steady state is an integral of the velocity function in Equation A.1:

$$w_2(t) = \begin{cases} \frac{1}{2}At^2 + w_{\min}'t & \text{if } t \in \left[0, \frac{B_1}{2}\right) \\ -\frac{1}{2}At^2 + AB_1t + w_{\min}'t - \frac{1}{4}AB_1^2 & \text{if } t \in \left[\frac{B_1}{2}, B_1\right). \end{cases}$$

To indicate that the speed and the way of c_3 depend on the sending time t_s of c_2, an index is introduced for this. With this, the speed and the way are described as follows. If $t_s \in [0, B_1/2)$:

$$w_{3,t_s}'(t) = \begin{cases} -At + At_s + w_{\min}' & \text{if } t \in [0, t_s) \\ At - At_s + w_{\min}' & \text{if } t \in [t_s, t_s + \frac{B_1}{2}) \\ -At + At_s + AB_1 + w_{\min}' & \text{if } t \in [t_s + \frac{B_1}{2}, B_1), \end{cases}$$

$$w_{3,t_s}(t) = \begin{cases} -\frac{1}{2}At^2 + At_st + w_{\min}'t & \text{if } t \in [0, t_s) \\ \frac{1}{2}At^2 - At_st + w_{\min}'t + At_s^2 & \text{if } t \in \left[t_s, t_s + \frac{B_1}{2}\right) \\ -\frac{1}{2}At^2 + AB_1t + At_st + w_{\min}'t - AB_1t_s - \frac{1}{4}AB_1^2 & \text{if } t \in \left[t_s + \frac{B_1}{2}, B_1\right). \end{cases}$$

In the other case, if $t_s \in [B_1/2, B_1)$:

$$w'_{3,t_s}(t) = \begin{cases} At - At_s + AB_1 + w'_{\min} & \text{if } t \in [0, t_s - \frac{B_1}{2}) \\ -At + At_s + w'_{\min} & \text{if } t \in [t_s - \frac{B_1}{2}, t_s) \\ At - At_s + w'_{\min} & \text{if } t \in [t_s, B_1), \end{cases}$$

$$w_{3,t_s}(t) = \begin{cases} \frac{1}{2}At^2 - At_s t + AB_1 t + w'_{\min} t & \text{if } t \in \left[0, t_s - \frac{B_1}{2}\right) \\ -\frac{1}{2}At^2 + At_s t + w'_{\min} t - At_s^2 + AB_1 t_s - \frac{1}{4}AB_1^2 & \text{if } t \in \left[t_s - \frac{B_1}{2}, t_s\right) \\ \frac{1}{2}At^2 - At_s t + w'_{\min} t + AB_1 t_s - \frac{1}{4}AB_1^2 & \text{if } t \in \left[t_s, B_1\right). \end{cases}$$

Again for the case $t_s \in [0, B_1/2)$, it is $t_a = t_s + B_1/2 \in [B_1/2, B_1)$ and thus

$$w_2\left(t_s + \frac{B_1}{2}\right) = -\frac{1}{2}At_s^2 + \frac{1}{2}AB_1 t_s + w'_{\min} t_s + \frac{1}{8}AB_1^2 + \frac{1}{2}B_1 w'_{\min},$$
$$w_{3,t_s}\left(t_s + \frac{B_1}{2}\right) = \frac{1}{2}At_s^2 + w'_{\min} t_s + \frac{1}{8}AB_1^2 + \frac{1}{2}B_1 w'_{\min}.$$

For $t_s \in [\frac{B_1}{2}, B_1)$, it is $t_a - B_1 = t_s - B_1/2 \in [0, B_1/2)$ and it follows

$$w_2\left(t_s - \frac{B_1}{2}\right) = \frac{1}{2}At_s^2 - \frac{1}{2}AB_1 t_s + w'_{\min} t_s + \frac{1}{8}AB_1^2 - \frac{1}{2}B_1 w'_{\min},$$
$$w_{3,t_s}\left(t_s - \frac{B_1}{2}\right) = \frac{1}{2}At_s^2 + AB_1 t_s + w'_{\min} t_s - \frac{3}{8}AB_1^2 - \frac{1}{2}B_1 w'_{\min}.$$

The distance between the cars at time 0 depending on t_s is

$$\Delta w_{t_s}(0) = \begin{cases} \Delta w_a(t_s) - w_2(t_s + \frac{B_1}{2}) + w_{3,t_s}(t_s + \frac{B_1}{2}) & \text{if } t_s \in \left[0, \frac{B_1}{2}\right) \\ \Delta w_a(t_s) - w_2(t_s - \frac{B_1}{2}) + w_{3,t_s}(t_s - \frac{B_1}{2}) & \text{if } t_s \in \left[\frac{B_1}{2}, B_1\right) \end{cases}$$
$$= \begin{cases} -\frac{1}{2}At_s^2 - w'_{\min} t_s + \frac{1}{4}AB_1^2 + B_1 w'_{\min} & \text{if } t_s \in \left[0, \frac{B_1}{2}\right) \\ -\frac{1}{2}At_s^2 + AB_1 t_s + w'_{\min} t_s - \frac{1}{4}AB_1^2 & \text{if } t_s \in \left[\frac{B_1}{2}, B_1\right). \end{cases}$$

Now with this, $\forall t \in [0, B_1)$ the distance between c_2 and c_3 is given by

$$\Delta w_{t_s}(t) = \Delta w_{t_s}(0) + w_2(t) - w_{3,t_s}(t). \tag{A.2}$$

The function is subdivided regarding t_s as in the cases above; first for $t_s \in [0, B_1/2)$:

$$\Delta w_{t_s}(t) = \begin{cases} -\frac{1}{2}At_s^2 - w'_{\min}t_s + \frac{1}{4}AB_1^2 + B_1w'_{\min} + At^2 - At_st & \text{if } t \in [0, t_s) \\ -\frac{3}{2}At_s^2 - w'_{\min}t_s + \frac{1}{4}AB_1^2 + B_1w'_{\min} + At_st & \text{if } t \in \left[t_s, \frac{B_1}{2}\right) \\ -\frac{3}{2}At_s^2 - w'_{\min}t_s + B_1w'_{\min} - At^2 + AB_1t + At_st & \text{if } t \in \left[\frac{B_1}{2}, t_s + \frac{B_1}{2}\right) \\ -\frac{1}{2}At_s^2 - w'_{\min}t_s + \frac{1}{4}AB_1^2 + B_1w'_{\min} - At_st + AB_1t_s & \text{if } t \in \left[t_s + \frac{B_1}{2}, B_1\right). \end{cases}$$

And for $t_s \in [B_1/2, B_1)$:

$$\Delta w_{t_s}(t) = \begin{cases} -\frac{1}{2}At_s^2 + AB_1t_s + w'_{\min}t_s - \frac{1}{4}AB_1^2 + At_st - AB_1t & \text{if } t \in \left[0, t_s - \frac{B_1}{2}\right) \\ \frac{1}{2}At_s^2 + w'_{\min}t_s + At^2 - At_st & \text{if } t \in \left[t_s - \frac{B_1}{2}, \frac{B_1}{2}\right) \\ \frac{1}{2}At_s^2 + w'_{\min}t_s - \frac{1}{4}AB_1^2 + AB_1t - At_st & \text{if } t \in \left[\frac{B_1}{2}, t_s\right) \\ -\frac{1}{2}At_s^2 + w'_{\min}t_s - \frac{1}{4}AB_1^2 - At^2 + AB_1t + At_st & \text{if } t \in [t_s, B_1). \end{cases}$$

Part B Now having an expression for $\Delta w_{t_s}(t)$, the optimal sending times for minimizing the steady-state distance are discussed. It will be shown that the optimal time for c_2 sending updates to minimize the steady-state distance between c_2 and c_3 is $t_s = B_1/2 + iB_1$ by

$$\Delta w_{\frac{B_1}{2}}(t) \leq \Delta w_{t_s}(t) \quad \forall t_s \in [0, B_1), t \in [0, B_1).$$

For $t_s = B_1/2$, the equation becomes

$$\Delta w_{\frac{B_1}{2}}(t) = \begin{cases} \frac{1}{8}AB_1^2 + \frac{1}{2}B_1w'_{\min} + At^2 - \frac{1}{2}AB_1t & \text{if } t \in \left[0, \frac{B_1}{2}\right) \\ -\frac{3}{8}AB_1^2 + \frac{1}{2}B_1w'_{\min} - At^2 + \frac{3}{2}AB_1t & \text{if } t \in \left[\frac{B_1}{2}, B_1\right). \end{cases} \tag{A.3}$$

With that the theorem is proven in two cases.

Case 1: let $t_s \in [0, B_1/2)$.

For $t_s = 0$, the minimum distance is larger than the maximum distance of $t_s = B_1/2$:

$$\Delta w_{0,\min}(t) > \Delta w_{\frac{B_1}{2},\max}(t).$$

The distance for $t_s = 0$ is constant, so the minimum distance equals the distance at $t = 0$. The maximum distance for $t = B_1/2$ is at $t = (3/4)B_1$, when the vehicles are equally fast, because c_2 decelerates, while c_3 is in its acceleration phase. It is obtained:

$$B_1w'_{\min} + \frac{1}{4}AB_1^2 > \frac{1}{2}w'_{\min}B_1.$$

$t_s \in (0, B_1/2)$ is expressed as $B_1/2 - \varepsilon$ with $0 < \varepsilon < B_1/2$, so it is $\Delta w_{B_1/2-\varepsilon}(t)$

$$= \begin{cases} \frac{1}{8}AB_1^2 + \left(\frac{1}{2}B_1 - \frac{1}{2}\varepsilon + t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} + At^2 - \frac{1}{2}AB_1t & \text{if } t \in \left[0, \frac{B_1}{2} - \varepsilon\right) \\ -\frac{1}{8}AB_1^2 + \left(\frac{3}{2}B_1 - \frac{3}{2}\varepsilon - t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} + \frac{1}{2}AB_1t & \text{if } t \in \left[\frac{B_1}{2} - \varepsilon, \frac{B_1}{2}\right) \\ -\frac{3}{8}AB_1^2 + \left(\frac{3}{2}B_1 - \frac{3}{2}\varepsilon - t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} - At^2 + \frac{3}{2}AB_1t & \text{if } t \in \left[\frac{B_1}{2}, B_1 - \varepsilon\right) \\ \frac{5}{8}AB_1^2 - \left(\frac{1}{2}B_1 - \frac{1}{2}\varepsilon + t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} - \frac{1}{2}AB_1t & \text{if } t \in \left[B_1 - \varepsilon, B_1\right). \end{cases}$$

All four intervals are considered in the following to show that the distance is larger than for $t_s = B_1/2$. First, $t \in [0, B_1/2 - \varepsilon)$ is evaluated:

$$\Delta w_{\frac{B_1}{2}-\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) = \frac{1}{2}AB_1\varepsilon - \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon + At\varepsilon > 0,$$

because $a, B_1, \varepsilon > 0$ and $w'_{\min}, t \geq 0$. Moreover, $B_1 > B_1/2 > \varepsilon$, so $(1/2)aB_1\varepsilon - (1/2)a\varepsilon^2 > 0$. The next interval is $t \in [B_1/2 - \varepsilon, B_1/2)$:

$$\Delta w_{\frac{B_1}{2}-\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) = -\frac{1}{4}AB_1^2 + \frac{3}{2}AB_1\varepsilon - \frac{3}{2}A\varepsilon^2 + w'_{\min}\varepsilon + AB_1t - At\varepsilon - At^2$$

This is derived:

$$\left(\Delta w_{\frac{B_1}{2}-\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t)\right)' = AB_1 - A\varepsilon - 2At.$$

To find the maximum, this is set equal to zero, $(\Delta w_{B_1/2-\varepsilon}(t) - \Delta w_{B_1/2}(t))' = 0$, which results in $t = B_1/2 - \varepsilon/2$. From

$$\Delta w_{\frac{B_1}{2}-\varepsilon}\left(\frac{B_1}{2} - \varepsilon\right) - \Delta w_{\frac{B_1}{2}}\left(\frac{B_1}{2} - \varepsilon\right) \geq \frac{1}{2}AB_1^2 + AB_1\varepsilon - \frac{3}{2}A\varepsilon^2 > \frac{1}{2}AB_1(B_1 - \varepsilon) > 0$$

and

$$\Delta w_{\frac{B_1}{2}-\varepsilon}\left(\frac{B_1}{2}) - \Delta w_{\frac{B_1}{2}}(\frac{B_1}{2}\right) > \varepsilon(aB_1 - \frac{3}{4}aB_1 + w'_{\min}) > 0,$$

it follows $\Delta w_{B_1/2-\varepsilon}(t) - \Delta w_{B_1/2}(t) > 0$ for $t \in [B_1/2 - \varepsilon, B_1/2)$. The next interval is $t \in [B_1/2, B_1 - \varepsilon)$:

$$\begin{aligned} &\Delta w_{\frac{B_1}{2}-\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) = \frac{3}{2}AB_1\varepsilon - \frac{3}{2}A\varepsilon^2 + w'_{\min}\varepsilon - At\varepsilon \\ \geq &A\varepsilon\left(\frac{3}{2}B_1 - \frac{3}{2}\varepsilon - t\right) > A\varepsilon\left(\frac{3}{2}B_1 - \frac{3}{2}\varepsilon - (B_1 - \varepsilon)\right) = A\varepsilon\left(\frac{1}{2}B_1 - \frac{1}{2}\varepsilon\right) > 0 \end{aligned}$$

since $t < B_1 - \varepsilon$. And eventually the fourth interval is $t \in [B_1 - \varepsilon, B_1)$:

$$\begin{aligned}\Delta w_{\frac{B_1}{2}-\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) &= AB_1^2 - \frac{1}{2}AB_1\varepsilon - \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon - 2AB_1t + At\varepsilon + At^2 \\ \left(\Delta w_{\frac{B_1}{2}-\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t)\right)' &= -2AB_1 + A\varepsilon + 2At.\end{aligned}$$

Setting $(\Delta w_{B_1/2-\varepsilon}(t) - \Delta w_{B_1/2}(t))' = 0$ yields $t = B_1 - \varepsilon/2$ which in this case is the lowest point of the function. With

$$\begin{aligned}&\Delta w_{\frac{B_1}{2}-\varepsilon}\left(B_1 - \frac{\varepsilon}{2}\right) - \Delta w_{\frac{B_1}{2}}\left(B_1 - \frac{\varepsilon}{2}\right) \\ =&\frac{1}{2}AB_1\varepsilon - \frac{3}{4}A\varepsilon^2 + w'_{\min}\varepsilon \geq \frac{1}{2}A\varepsilon\left(B_1 - \frac{3}{2}\varepsilon\right) \overset{\varepsilon<\frac{B_1}{2}}{>} \frac{1}{2}A\varepsilon\left(B_1 - \frac{3}{4}B_1\right) > 0\end{aligned}$$

it holds $\Delta w_{t_s}(t) > \Delta w_{B_1/2}(t)$ for all $t_s \in [0, B_1/2)$ and $t \in [0, B_1)$.

Case 2: let $t_s \in (B_1/2, B_1)$. Here, $t_s = B_1/2 + \varepsilon$, $0 < \varepsilon < B_1/2$, so $\Delta w_{B_1/2+\varepsilon}(t)$

$$= \begin{cases} \frac{1}{8}AB_1^2 + \left(\frac{1}{2}B_1 - \frac{1}{2}\varepsilon + t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} - \frac{1}{2}AB_1t & \text{if } t \in [0, \varepsilon) \\ \frac{1}{8}AB_1^2 + \left(\frac{1}{2}B_1 + \frac{1}{2} - t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} + at^2 - \frac{1}{2}aB_1t & \text{if } t \in \left[\varepsilon, \frac{B_1}{2}\right) \\ -\frac{1}{8}AB_1^2 + \left(\frac{1}{2}B_1 + \frac{1}{2}\varepsilon - t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} + \frac{1}{2}AB_1t & \text{if } t \in \left[\frac{B_1}{2}, \frac{B_1}{2} + \varepsilon\right) \\ -\frac{3}{8}AB_1^2 - \left(\frac{1}{2}B_1 - \frac{1}{2}\varepsilon + t\right)A\varepsilon + \left(\frac{1}{2}B_1 + \varepsilon\right)w'_{\min} - At^2 + \frac{3}{2}AB_1t & \text{if } t \in \left[\frac{B_1}{2} + \varepsilon, B_1\right). \end{cases}$$

Again, all four intervals are considered. At first, let $t \in [0, \varepsilon)$. Then

$$\Delta w_{\frac{B_1}{2}+\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) = \frac{1}{2}AB_1\varepsilon - \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon + At\varepsilon - At^2 > 0$$

since $\varepsilon < B_1/2$ and $t < \varepsilon$, because $t \in [0, \varepsilon)$. For $t \in [\varepsilon, B_1/2)$:

$$\Delta w_{\frac{B_1}{2}+\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) = \frac{1}{2}AB_1\varepsilon - At\varepsilon + \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon > 0$$

because $t < B_1/2$. For $t \in [B_1/2, B_1/2 + \varepsilon)$:

$$\begin{aligned}&\Delta w_{\frac{B_1}{2}+\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) = \frac{1}{2}AB_1^2 + \frac{1}{2}AB_1\varepsilon + \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon + At^2 - AB_1t - At\varepsilon \\ \Rightarrow &\left(\Delta w_{\frac{B_1}{2}+\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t)\right)' = A(2t - B_1 - \varepsilon).\end{aligned}$$

The minimum of this function is at $t = B_1/2 + \varepsilon/2$ and has the value

$$\Delta w_{\frac{B_1}{2}+\varepsilon}\left(\frac{B_1}{2} + \frac{\varepsilon}{2}\right) - \Delta w_{\frac{B_1}{2}}\left(\frac{B_1}{2} + \frac{\varepsilon}{2}\right) = w'_{\min}\varepsilon \geq 0.$$

For $t \in [B_1/2+\varepsilon, B_1)$, it is

$$\begin{aligned}\Delta w_{\frac{B_1}{2}+\varepsilon}(t) - \Delta w_{\frac{B_1}{2}}(t) &= -\frac{1}{2}AB_1\varepsilon - \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon + At\varepsilon \\ \geq -\frac{1}{2}AB_1\varepsilon - \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon + \frac{1}{2}AB_1\varepsilon + A\varepsilon^2 &= \frac{1}{2}A\varepsilon^2 + w'_{\min}\varepsilon > 0.\end{aligned}$$

With that, it holds $\Delta w_{t_s}(t) > \Delta w_{B_1/2}(t)$ for all $t_s \in (B_1/2, B_1)$ and $t \in [0, B_1)$. Combining the findings of the two cases, it is shown that

$$\Delta w_{\frac{B_1}{2}}(t) \leq \Delta w_{t_s}(t) \quad \forall t_s \in [0, B_1), t \in [0, B_1).$$

□

A.6 Proof of Lemma 3

Lemma 3 states: *Given a car c_i with the initial speed of $w'_{\max}$ and a free-flow arrival time of t_i. Let the true arrival time at the merge point m be $\tilde{t}_i \geq t_i$. It exists an optimal way for c_i regarding merging efficiency in which it brakes from t_a to t_b, waits between t_b and t_c, and accelerates from t_c to $\tilde{t}_i$. It drives at constant speed from $\tilde{t}_i$ onwards.*

$$\begin{aligned}\textit{Here, } t_c &= \tilde{t}_i - \min\left\{\left(\left(w'_{\max}(\tilde{t}_i - t_i^0) - m\right)/A\right)^{1/2}, w'_{\max}/A\right\}, \\ t_b &= t_c - \max\left\{0, \tilde{t}_i - \left(t_i^0 + m/w'_{\max} + w'_{\max}/A\right)\right\}, \textit{ and} \\ t_a &= t_b + t_c - \tilde{t}_i.\end{aligned}$$

Proof. Car c_i needs to be at the merge point m at time $\tilde{t}_i$ instead of at the free-flow arrival time t_i. Assume $\tilde{t}_i > t_i$, i.e., a delay is required. To build up a delay, c_i brakes and then accelerates until it is again at the speed $w'_{\max}$. This is performed the latest possible for traveling optimality: c_i merges instantly after delaying is finished with the maximum speed. At the merging, c_i has a car ahead and is directly on minimum steady-state distance to it, which is the smallest distance for a constant speed following. The way driven by c_i until m including braking and accelerating is

$$\begin{aligned}m = {}& w'_{\max}(t_a - t_i^0) + w'_{\max}(t_b - t_a) \\ &+ \frac{D}{2}(t_b - t_a)^2 + \left(w'_{\max} + (t_b - t_a)D\right)\left(\tilde{t}_i - t_c\right) + \frac{A}{2}\left(\tilde{t}_i - t_c\right)^2.\end{aligned}$$

A waiting time between decelerating and accelerating is required if c_i brakes to a full stop. It is assumed at first that no full stop is required, thus $t_b = t_c$. With $D = -A$, the acceleration phase and deceleration phase are equally long: it is $\tilde{t}_i - t_c = \tilde{t}_i - t_b = t_b - t_a$ and $\tilde{t}_i = 2t_b - t_a$. It is obtained

$$\begin{aligned} m &= w'_{\max}(t_b - t_i^0) + (w'_{\max} + (t_b - t_a)D)(t_b - t_a) \\ &= w'_{\max}(\tilde{t}_i - t_i^0) + (\tilde{t}_i - t_b)^2 D \\ \Leftrightarrow \tilde{t}_i - t_b &= \pm\sqrt{\frac{m - w'_{\max}(\tilde{t}_i - t_i^0)}{D}}. \end{aligned}$$

Because $t_b \leq \tilde{t}_i$, the positive root is chosen. This equation is solved for t_b:

$$\Rightarrow t_b = \tilde{t}_i - \sqrt{\frac{w'_{\max}(\tilde{t}_i - t_i^0) - m}{A}}.$$

The time intervals to accelerate and brake are the same due to $D = -A$, thus

$$t_a = t_b - (\tilde{t}_i - t_c) = t_b + t_c - \tilde{t}_i.$$

Now, it is assumed that a full stop is required. The expression for t_b determined so far becomes t_c. t_c is the time at which acceleration is started. c_i can fully stop at most, hence t_c minimally is $t_{c,\min} = \tilde{t}_i - w'_{\max}/A$. With that, t_c can be stated as

$$t_c = \tilde{t}_i - \min\left\{\sqrt{\frac{w'_{\max}(\tilde{t}_i - t_i^0) - m}{A}}, \frac{w'_{\max}}{A}\right\}.$$

t_b is the time c_i has finished decelerating. If it has stopped, it must wait for $t_c - t_b$ before accelerating. t_c is already known and also, by $\tilde{t}_i - t_c$, the equally long time interval for deceleration, $t_b - t_a$. The waiting time is thus obtained by subtracting the potential arrival time of c_i by driving without waiting (by setting $t_b = t_c$) from the true arrival time $\tilde{t}_i$. Because the waiting time cannot be negative, it can be expressed as

$$t_c - t_b = \max\left\{0, \tilde{t}_i - \left(t_i^0 + \frac{m - 2(w'_{\max})^2/(2A)}{w'_{\max}} + 2\frac{w'_{\max}}{A}\right)\right\},$$

where $2w'_{\max}/A$ is the time required to stop fully and accelerate again, while the term $(m - 2(w'_{\max})^2/(2A))/w'_{\max}$ describes the time to drive the distance from $w_i(t_i^0) = 0$ to the waypoint to begin decelerating until full stop; this distance is $m - 2(w'_{\max})^2/(2A)$, i.e., the distance needed for a speed change from 0 to $w'_{\max}$ is subtracted twice from m. The equation is simplified and solved for t_b:

$$t_b = t_c - \max\left\{0, \tilde{t}_i - \left(t_i^0 + \frac{m}{w'_{\max}} + \frac{w'_{\max}}{A}\right)\right\},$$

t_a is easily obtained by subtracting the deceleration interval from the time at which c_i ends the deceleration. Because the deceleration interval equals the acceleration interval, it is $t_a = t_b - (\tilde{t}_i - t_c) = t_b + t_c - \tilde{t}_i$.

If $\tilde{t}_i = t_i$, then no delay is required and it is $t_c = t_b = t_a = t_i$. □

A.7 Proof of Lemma 4

Lemma 4 states: *Given a car c_i with a free-flow arrival time of t_i. Let the true arrival time at the merge point m be $\tilde{t}_i \geq t_i$, and let c_i have a safe distance to its predecessor at the current time $t \geq t_i^0$.*

Given the points in time $t_s \leq t \leq t_a \leq t_b \leq t_c \leq t_d \leq \tilde{t}_i$: there exists an optimal way for c_i regarding merging efficiency in which it accelerates in $[t_s \leq t, t_a)$, drives with constant speed in $[t_a, t_b)$, brakes in $[t_b, t_c)$, stops in $[t_c, t_d)$, and accelerates in $[t_d, \tilde{t}_i)$. It drives at constant speed from $\tilde{t}_i$ onwards.

Proof. The existence of an optimal way is shown by constructing it. Even if c_i is not at maximum speed at t because it braked for safety, it just began delaying earlier than optimal regarding merging efficiency. Since it braked because of being on a safety distance, its way is road-usage efficient. It is still able to merge at $\tilde{t}_i$ because the car directly ahead, causing c_i to brake, braked either for merging efficiency or for safety reasons that were caused by a merging efficient way of another car. Braking without need does not exist in the model.

The behavior for the optimal way that will be built consists of three cases depending on when braking for delaying the arrival at the merging has to start and it looks as follows:

(1) if $\hat{t}_i \geq \tilde{t}_i$ then

$t_a = t_{\max}, \quad t_b = t_c = t_d = \hat{t}_i.$

(2) else if $\hat{w}_i(\tilde{t}_i) \leq m \vee \breve{m} \geq (w'_{\max})^2/A$ then

$t_a = t_{\max}, \quad t_d = \tilde{t}_i - \min\left\{((w'_{\max}(\tilde{t}_i - \check{t}_{\max}) - \breve{m})/A)^{1/2}, w'_{\max}/A\right\},$

$t_c = t_d - \max\left\{0, \tilde{t}_i - (t_{\max} + \breve{m}/w'_{\max} + w'_{\max}/A)\right\}, \quad$ and $t_b = t_c + t_d - \tilde{t}_i;$

(3) else

$t_b = t_a = t_s + t_y/4 + \breve{m}/(At_y);$

(3a) if $w'_i(t_a + t_y/2) < 0$ then $t_b = t_a = t_s + (\breve{m}/A)^{1/2};$

$t_c = t_a + \min\{t_a - t_s, t_y/2\}, \,$ and $t_d = \max\left\{t_c, \tilde{t}_i - w'_{\max}/A\right\};$

$$
\begin{aligned}
\text{Here, } \hat{t}_i &= t_s + w'_{\max}/A + \breve{m}/w'_{\max},\\
t_s &= t - w'_i(t)/D,\\
t_{\max} &= t_s + w'_{\max}/A,\\
\breve{m} &= m - w_i(t_s) - (w'_{\max})^2/(2A),\\
t_y &= \tilde{t}_i - t_s - w'_{\max}/A, \text{ and}\\
\hat{w}_i(\tilde{t}_i) &= w_i(t_{\max}) + w'_{\max}(\tilde{t}_i - t_{\max})/2 - A(\tilde{t}_i - t_{\max})^2/16.
\end{aligned}
$$

Now it is discussed how this behavior is created. It is assumed that the incoming lanes are long enough so that accident-absent ways for newly inserted cars exist. c_i is at speed $w'_i(t)$. To simplify the discussion, c_i is assumed to have had stopped at $t_s \leq t$ and then to have accelerated until t. This time is calculated as $t_s = t - w'_i(t)/A$. The position of c_i at time t_s is $w_i(t_s) = w_i(t) - (A/2)(t - t_s)^2$. The earliest time $\hat{t}_i$ the merging can be reached from the state at t_s is obtained by accelerating until $w'_{\max}$ and then driving the remaining space to the merging at that speed. $\hat{t}_i$ does not have to equal the free-flow arrival time t_i because until $t > t_i^0$, c_i may have braked already for safety reasons. $\hat{t}_i$ is calculated like this: the time for c_i to reach the maximum speed $w'_{\max}$ is $t_{\max} = t_s + w'_{\max}/A$. When c_i has that speed it is at point

$w_i(t_{\max}) = w_i(t_s) + (A/2)(t_{\max} - t_s)^2$. The remaining lane length to the merging is $\breve{m} = m - w_i(t_{\max})$ and the remaining time to reach m from $w_i(t_{\max})$ is $\breve{m}/w'_{\max}$. Thus,

$$\hat{t}_i = t_s + \frac{w'_{\max}}{A} + \frac{\breve{m}}{w'_{\max}}.$$

How does a way with a delaying phase have to look like regarding the objectives? If a delay is necessary before merging, i.e., $\tilde{t}_i > \hat{t}_i$, that delay is created by decelerating and accelerating and thereby driving slower than required for safety. The way is built such that c_i merges exactly at $\tilde{t}_i$ at the maximum speed and a minimum steady state distance to its predecessor regarding a fair merge order. Then this way creates an optimal behavior regarding maximum merging efficiency. As such a delaying phase can be started at multiple times, the question of when to insert the delaying phase has to be discussed. This is solved by choosing a traveling-optimal way; this demands to brake the latest possible to stay furthest ahead anytime. According to this is the strategy to insert the delaying phase the latest possible at which the merging can still be reached with $w'_{\max}$. Three cases when the delaying has to start are differentiated.

Case (1): no additional delay is necessary, i.e., $\hat{t}_i \geq \tilde{t}_i$. Starting at time t_s, it takes $t_a = t_{\max}$ seconds to accelerate to $w'_{\max}$. It takes another $\breve{m}/w'_{\max}$ seconds to reach the merging at time $\hat{t}_i$.

Case (2): The delay fits completely in the interval between $t_{\max}$ and $\tilde{t}_i$. This is the case if the minimum distance driven in $[t_{\max}, \tilde{t}_i)$, starting and finishing with speed $w'_{\max}$, is less than or equal to $m - \breve{m}$. Let $t_b = t_{\max} + (\tilde{t}_i - t_{\max})/2$. The minimum distance is obtained by braking from $t_{\max}$ until t_b and accelerate from t_b until $\tilde{t}_i = t_{\max} + 2(t_b - t_{\max}) = 2t_b - t_{\max}$. Let $t_x = \tilde{t}_i - t_b = t_b - t_{\max}$ and let the way $\hat{w}_i(t)$ describe that behavior:

$$\hat{w}_i(\tilde{t}_i) = w_i(t_a) + w'_{\max}\frac{t_x}{2} + \frac{D}{2}\left(\frac{t_x}{2}\right)^2 + \left(w'_{\max} + D\frac{t_x}{2}\right)\frac{t_x}{2} + \frac{A}{2}\left(\frac{t_x}{2}\right)^2.$$

With $D = -A : \hat{w}_i(\tilde{t}_i) = w_i(\max) + w'_{\max}t_x - At_x^2/4$. This equation does not catch the special case of c_i braking longer than to a full stop. Adding a test for this, the condition for (2) becomes: $\hat{w}_i(\tilde{t}_i) \leq m \vee \breve{m} \geq (w'_{\max})^2/A$. The calculation of the points in time t_b until t_d is now very similar to Lemma 3's points t_a to t_c. Only two minor changes are necessary: replace t_i^0 with $t_{\max}$ and m with $\breve{m}$.

Case (3): c_i has to begin delaying already before $t_{\max}$, i.e., during the acceleration phase, because a larger delay than possible with case (2) is necessary. At first it is assumed that no full stop during the delaying phase is required. The phase is created

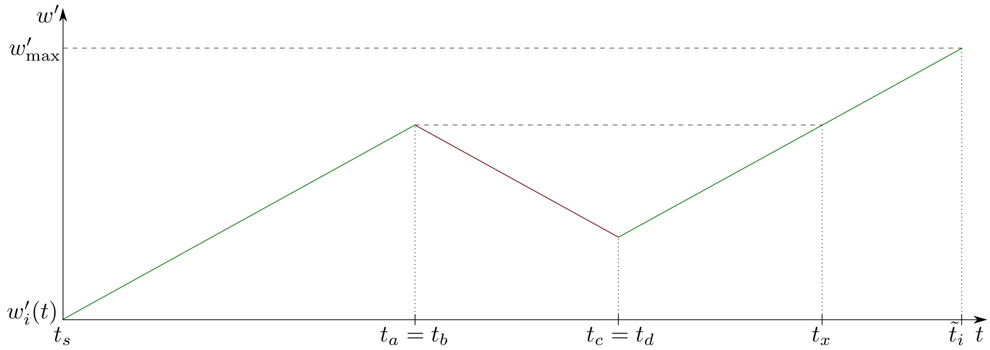

Figure A.4: Speed $w'_i(t)$ of car c_i that has to begin delaying its merging time before the acceleration to maximum speed is finished. c_i starts fully stopped and does not have to stop again in this case.

by inserting a deceleration and an acceleration phase. Figure A.4 shows the speed $w'_i(t)$ of car c_i in such a situation.

The acceleration phase is split into $[t_s, t_a)$ and $[t_x, \tilde{t}_i)$. During these intervals, the distance of $(w'_{\max})^2/(2A)$ is driven. The remaining distance $\breve{m}$ is given by driving during the delay interval $[t_a, t_x)$. The delay interval is defined by $t_x - t_a = \tilde{t}_i - (t_s + w'_{\max}/A)$. Now the only variable left to be determined is t_a such that $\breve{m}$ is driven during $[t_x, t_a)$. The distance driven in that interval equals the average speed $\overline{w}$ of c_i in the interval multiplied by the interval's length $t_x - t_a$:

$$\begin{aligned} \breve{m} &= \overline{w}(t_x - t_a) \\ \Leftrightarrow \breve{m} &= \frac{w'_i(t_c) + w'_i(t_a)}{2}(t_x - t_a). \end{aligned}$$

Let $t_y = t_x - t_a = \tilde{t}_i - (t_s + w'_{\max}/A)$ and with $w'_i(t_c) = w'_i(t_a) + D(t_x - t_a)/2$, the equation becomes:

$$\breve{m} = \left(w'_i(t_a) + \frac{D}{4} t_y \right) t_y.$$

With $w'_i(t_a) = w'_i(t_s) + A(t_a - t_s)$ and $w'_i(t_s) = 0$, it follows:

$$\breve{m} = \left(A\,(t_a - t_s) + \frac{D}{4} t_y \right) t_y.$$

Replacing $D = -A$, it is obtained

$$t_a = t_s + \frac{1}{4} t_y + \frac{\breve{m}}{A t_y}.$$

There is no phase of driving with maximum speed before deceleration begins, thus $t_b = t_a$. It is switched from deceleration to acceleration exactly in the middle of the delaying phase: $t_c = t_a + (t_x - t_a)/2 = t_a + t_y/2$. Eventually, $t_d = t_c$.

Now, it is assumed that a full stop during deceleration is required. In that case, c_i intends to drive the farthest possible from t_s until t_c, waits until t_d, and then accelerates until $\tilde{t}_i$. Here, $w_i(t_c) = w_i(t_s) + \breve{m}$, $t_b = t_a, t_c - t_a = t_a - t_s$, and $t_d = \tilde{t}_i - w'_{\max}/A$. t_d is later than t_c in the case of a stop and it is at a minimum. Combining both cases of stopping at t_c and of not stopping at t_c, the result is

$$t_d = \max\left\{t_c, \tilde{t}_i - \frac{w'_{\max}}{A}\right\}.$$

As described above, $t_c - t_a = t_a - t_s$, thus $t_c = 2t_a - t_s$. t_c is at a minimum if c_i stops. In an equation that comprises the full-stop case and the non-stop case t_c can be expressed as

$$t_c = \min\left\{2t_a - t_s, t_a + \frac{t_y}{2}\right\} = t_a + \min\left\{t_a - t_s, \frac{t_y}{2}\right\}.$$

To get t_a, the equation $w_i(t_c) = w_i(t_s) + \breve{m}$ is reformulated:

$$w_i(t_c) = w_i(t_s) + m - w_i(t_s) - \frac{(w'_{\max})^2}{2A} = m - \frac{(w'_{\max})^2}{2A}$$

and then it is set equal to the kinematic equation of the position c_i has at t_c:

$$m - \frac{(w'_{\max})^2}{2A} = w_i(t_s) + \frac{A}{2}(t_a - t_s)^2 + A(t_a - t_s)(t_c - t_a) + \frac{D}{2}\left(t_c - t_a\right)^2.$$

With $t_c = 2t_a - t_s$ and $D = -A$, t_a can be isolated like in the following:

$$\begin{aligned}
m - \frac{(w'_{\max})^2}{2A} &= w_i(t_s) + \frac{A}{2}(t_a - t_s)^2 + A(t_a - t_s)(t_a - t_s) + \frac{D}{2}(t_a - t_s)^2 \\
\Leftrightarrow m - \frac{(w'_{\max})^2}{2A} &= w_i(t_s) + A(t_a - t_s)^2 \\
\Leftrightarrow (t_a)_{1,2} &= t_s \pm \sqrt{\frac{m - w_i(t_s) - (w'_{\max})^2/(2A)}{A}} = t_s \pm \sqrt{\frac{\breve{m}}{A}}.
\end{aligned}$$

Because $t_s \leq t_a$, the positive root is chosen. To decide whether a full stop is required or not, t_a is simply calculated as if no stop was necessary. Then it is tested in (3a) if the speed is negative at t_c: $w'_i(t_a + t_y/2) < 0$. If so, a full stop is necessary. □

Bibliography

Own Publications

[ABM11] Nasir Ali, Daniel Baselt, and Hermann De Meer. Analysis of car-to-car beaconing with network coding. *Electronic Communications of the EASST*, 37:1–13, March 2011.

[BKS+13] Daniel Baselt, Florian Knorr, Björn Scheuermann, Michael Schreckenberg, and Martin Mauve. Merging lanes—fairness through communication. *Submitted for publication*, 2013.

[BMS11] Daniel Baselt, Martin Mauve, and Björn Scheuermann. A top-down approach to inter-vehicle communication. In *VNC '11: Proceedings of the 3rd Annual IEEE Vehicular Networking Conference*, pages 123–130, November 2011.

[KBMS11] Markus Koegel, Daniel Baselt, Martin Mauve, and Björn Scheuermann. A comparison of vehicular trajectory encoding techniques. In *MedHoc-Net '11: Proceedings of the 10th Annual Mediterranean Ad Hoc Networking Workshop*, pages 87–94, June 2011.

[KBSM12] Florian Knorr, Daniel Baselt, Michael Schreckenberg, and Martin Mauve. Reducing traffic jams via VANETs. *IEEE Transactions on Vehicular Technology*, 61(8):3490–3498, October 2012.

[SBM09] Michael Stini, Daniel Baselt, and Martin Mauve. Feasibility of Incremental Contract Signing. Technical Report TR-2009-001, Computer Science Department, Heinrich Heine University, Düsseldorf, Germany, September 2009.

Other References

[ABHL12] Chen Avin, Michael Borokhovich, Yoram Haddad, and Zvi Lotker. Optimal virtual traffic light placement. In *FOMC '12: Proceedings of the 8th ACM International Workshop on Foundations of Mobile Computing*, pages 6:1–6:10, July 2012.

[AH05] Hamada Alshaer and Eric Horlait. An optimized adaptive broadcast scheme for inter-vehicle communication. In *VTC '05-Spring: Proceedings of the 61st IEEE Vehicular Technology Conference*, volume 5, pages 2840–2844, May 2005.

[AV11] Behrang Asadi and Ardalan Vahidi. Predictive cruise control: Utilizing upcoming traffic signal information for improving fuel economy and reducing trip time. *IEEE Transactions on Control Systems Technology*, 19(3):707–714, May 2011.

[AW10] Hadi Arbabi and Michele C. Weigle. Highway mobility and vehicular ad-hoc networks in ns-3. In *WSC '10: Proceedings of the 42nd Conference on Winter Simulation*, pages 2991–3003. Society for Computer Simulation International, December 2010.

[BBCF12] Cynthia Barnhart, Dimitris Bertsimas, Constantine Caramanis, and Douglas S. Fearing. Equitable and efficient coordination in traffic flow management. *Transportation Science*, 46(2):262–280, May 2012.

[BBE+09] Richard Bossom, Roberto Brignolo, Thierry Ernst, Knut Evensen, Alexander Frötscher, Wolfgang Höfs, Juhani Jääskeläinen, Zeljko Jeftic, Paul Kompfner, Timo Kosch, Ilse Kulp, Antonio Kung, Abdel-Kader Mokaddem, Andreas Schalk, Elisabeth Uhlemann, and Christian Wewetzer. European ITS communication architecture, overall framework, proof of concept implementation. Technical Report D31 Version 2.0, COMeSafety, March 2009.

[BBEK11] Michael Behrisch, Laura Bieker, Jakob Erdmann, and Daniel Krajzewicz. SUMO – simulation of urban mobility: An overview. In *SIMUL '11: Proceedings of the 3rd International Conference on Advances in System Simulation*, pages 63–68, October 2011.

[Ben04] Abderrahim Benslimane. Optimized dissemination of alarm messages in vehicular ad-hoc networks (vanet). In *High Speed Networks and Multimedia Communications*, volume 3079 of *Lecture Notes in Computer Science*, pages 655–666. Springer Berlin Heidelberg, 2004.

[BHN+95] Masako Bando, Katsuya Hasebe, Akihiro Nakayama, Akihiro Shibata, and Yuki Sugiyama. Dynamical model of traffic congestion and numerical simulation. *Physical Review E*, 51(2):1035–1042, February 1995.

[BM99] Mark Brackstone and Mike McDonald. Car-following: a historical review. *Transportation Research Part F: Psychology and Behaviour*, 2:181–196, December 1999.

[BTD06] Subir Biswas, Raymond Tatchikou, and Francois Dion. Vehicle-to-vehicle wireless communication protocols for enhancing highway traffic safety. *Communications Magazine, IEEE*, 44(1):74–82, January 2006.

[BUSB09] Katrin Bilstrup, Elisabeth Uhlemann, Erik G. Ström, and Urban Bilstrup. On the ability of the 802.11p MAC method and STDMA to support real-time vehicle-to-vehicle communication. *EURASIP Journal on Wireless Communications and Networking*, 2009(1), January 2009.

[Cam97] Colin F. Camerer. Progress in behavioral game theory. *The Journal of Economic Perspectives*, 11(4, Fall):167–188, November 1997.

[CGG+06] Carla-Fabiana Chiasserini, Rossano Gaeta, Michele Garetto, Marco Gribaudo, and Matteo Sereno. Efficient broadcasting of safety messages in multihop vehicular networks. In *IPDPS '06: Proceedings of the 20th IEEE International Parallel and Distributed Processing Symposium*, April 2006.

[CHPS04] Roland Chrobok, Sigurdur F. Hafstein, Andreas Pottmeier, and Michael Schreckenberg. OLSIM: Future traffic information. In *ASM '04: Proceedings of the IASTED International Conference on Applied Simulation and Modelling*, June 2004.

[Cla88] David D. Clark. The design philosophy of the DARPA internet protocols. In *SIGCOMM '88: Proceedings of the 1988 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 106–114, August 1988.

[CLL10] Yuh-Shyan Chen, Yun-Wei Lin, and Sing-Ling Lee. A mobicast routing protocol in vehicular ad-hoc networks. *Mobile Networks and Applications*, 15:20–35, February 2010.

[Con05] The CAMP Vehicle Safety Communications Consortium. Vehicle safety communications project task 3 final report - identify intelligent vehicle safety applications enabled by DSRC. Technical Report HS-809 859, Crash Avoidance Metrics Partnership and National Highway Traffic Safety Administration (CAMP), March 2005.

[Coo81] Robert B. Cooper. *Introduction to queueing theory.* London: Arnold, 2nd edition, 1981.

[CRR09] Victor Cabrera, Francisco J. Ros, and Pedro M. Ruiz. Simulation-based study of common issues in VANET routing protocols. In *VTC '09-Spring: Proceedings of the 69th IEEE Vehicular Technology Conference*, April 2009.

[DH88] Walid Dabbous and Christian Huitema. *PROMETHEUS: Vehicle to Vehicle Communications.* EUREKA PROMETHEUS project, August 1988.

[DKO+11] Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz, and Lars Wischof. Research challenges in intervehicular communication: Lessons of the

2010 Dagstuhl seminar. *IEEE Communications Magazine*, 49(5):158–164, May 2011.

[DR98] Swaroop Darbha and Kumbakonam R. Rajagopal. *Intelligent Cruise Control Systems and Traffic Flow Stability.* California Partners for Advanced Transit and Highways (PATH) research report. California PATH Program, Institute of Transportation Studies, University of California at Berkeley, December 1998.

[DR05] Swaroop Darbha and Kumbakonam R. Rajagopal. Information flow and its relation to the stability of the motion of vehicles in a rigid formation. In *ACC '05: Proceedings of the 2005 IEEE American Control Conference*, pages 1853–1858, June 2005.

[DRPN07] Julius Degesys, Ian Rose, Ankit Patel, and Radhika Nagpal. DESYNC: Self-organizing desynchronization and tdma on wireless sensor networks. In *IPSN '07: Proceedings of the 6th International Conference on Information Processing in Sensor Networks*, pages 11–20, April 2007.

[EGH+06] Tamer A. ElBatt, Siddhartha K. Goel, Gavin Holland, Hariharan Krishnan, and Jayendra S. Parikh. Cooperative collision warning using dedicated short range wireless communications. In *VANET '06: Proceedings of the 3rd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 1–9, September 2006.

[EOSK05] Stephan Eichler, Benedikt Ostermaier, Christoph Schroth, and Timo Kosch. Simulation of car-to-car messaging: Analyzing the impact on road traffic. In *MASCOTS '05: Proceedings of the IEEE Computer Society's 13th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems*, pages 507–510, September 2005.

[Fax02] J. Alexander Fax. *Optimal and Cooperative Control of Vehicle Formations.* PhD thesis, California Institute of Technology, Pasadena, California, 2002.

[FCZ03] Enrique Fernández-Cara and Enrique Zuazua. Control theory: History, mathematical achievements and perspectives. *Anuario de la Sociedad Española de Matemática Aplicada*, 26:79–140, 2003.

[Fd12] Michel Ferreira and Pedro M. d'Orey. On the impact of virtual traffic lights on carbon emissions mitigation. *IEEE Transactions on Intelligent Transportation Systems*, 13(1):284–295, March 2012.

[FFCa+10] Michel Ferreira, Ricardo Fernandes, Hugo Conceição, Wantanee Viriyasitavat, and Ozan K. Tonguz. Self-organized traffic control. In *VANET '10: Proceedings of the 7th ACM International Workshop on VehiculAr Inter-NETworking*, pages 85–90, September 2010.

[FHSK10] Yaser P. Fallah, ChingLing Huang, Raja Sengupta, and Hariharan Krishnan. Design of cooperative vehicle safety systems based on tight coupling of communication, computing and physical vehicle dynamics. In *ICCPS '10: Proceedings of 1st ional Conference on Cyber-Physical Systems*, pages 159–167, April 2010.

[FJ94] Sally Floyd and Van Jacobson. The synchronization of periodic routing messages. *IEEE/ACM Transactions on Networking*, 2(2):122–136, April 1994.

[FMR06] Roberta Fracchia, Michela Meo, and Dario Rossi. VANETs: To beacon or not to beacon? In *AutoNet '06: Proceedings of the 1st IEEE Workshop on Automotive Networking and Application*, December 2006.

[fSI10] International Organization for Standardization (ISO). Intelligent transport systems – communications access for land mobiles (CALM) – M5, 2010.

[fTA10] American Society for Testing and Materials (ASTM). Standard specification for telecommunications and information exchange between roadside and vehicle systems – 5 GHz band dedicated short range communications (DSRC) medium access control (MAC) and physical layer (PHY) specifications, 2010.

[GGD+07] Victor Gradinescu, Cristian Gorgorin, Raluca Diaconescu, Valentin Cristea, and Liviu Iftode. Adaptive Traffic Lights Using Car-to-Car Communication. In *VTC '07-Spring: Proceedings of the 65th IEEE Vehicular Technology Conference*, pages 21–25, April 2007.

[GLW08] Huaxin Gong, Hongchao Liu, and Bing-Hong Wang. An asymmetric full velocity difference car-following model. *Physica A Statistical Mechanics and its Applications*, 387:2595–2602, April 2008.

[GS10] Balaji P. Gokulan and Dipti Srinivasan. Distributed geometric fuzzy multiagent urban traffic signal control. *IEEE Transactions on Intelligent Transportation Systems*, 11(3):714–727, September 2010.

[GSEMHt11] Tristan Gaugel, Felix Schmidt-Eisenlohr, Jens Mittag, and Hannes Hartenstein. A change in perspective: information-centric modeling of inter-vehicle communication. In *VANET '11: Proceedings of the 8th ACM International Workshop on VehiculAr Inter-NETworking*, pages 61–66, September 2011.

[HCM01] J. Karl Hedrick, Yu-Han Chen, and Sonia Mahal. Optimized vehicle control/communication interaction in an automated highway system. Research Report UCB-ITS-PRR-2001-29, Institute of Transportation Studies, UC Berkeley, January 2001.

[HFSK09] Ching-Ling Huang, Yaser P. Fallah, Raja Sengupta, and Hariharan Krishnan. Information dissemination control for cooperative active safety applications in vehicular ad-hoc networks. In *GLOBECOM '09: Proceedings of the IEEE Global Telecommunications Conference*, pages 4085–4090, November 2009.

[HHR+91] Bernhard Heister, Th. Hellmich, Carl-Herbert Rokitansky, Bernhard Walke, and Wei Zhu. Implementation of inter-vehicle and vehicle-beacon communication protocols for short-range mobile radio networks. In *Proceedings of the 5th PROMETHEUS Workshop*, page 10, October 1991.

[HW93] Ivy Pei-Shan Hsu and Jean Walrand. Communication requirements and network design for IVHS. PATH Working Paper UCB-ITS-PWP-93-18, Institute of Transportation Studies, UC Berkeley, November 1993.

[IC93] Petros Ioannou and Cheng-Chih Chien. Autonomous intelligent cruise control. In *IEEE Transactions on Vehicular Technology*, volume 42, pages 657–672, November 1993.

[IEE12] 802.11-2012 – standard for information technology – telecommunications and information exchange between systems – local and metropolitan area networks–specific requirements part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, 2012.

[Ins06a] Institute of Electrical and Electronics Engineers (IEEE). *1609.1-2006 – Trial-Use Standard for Wireless Access in Vehicular Environments (WAVE) – Resource Manager*, October 2006.

[Ins06b] Institute of Electrical and Electronics Engineers (IEEE). *1609.2-2006 – Trial-Use Standard for Wireless Access in Vehicular Environments – Security Services for Applications and Management Messages*, July 2006.

[Ins10] Institute of Electrical and Electronics Engineers (IEEE). *1609.3-2010 (Revision of IEEE Std 1609.3-2007) – Standard for Wireless Access in Vehicular Environments (WAVE) – Networking Services*, December 2010.

[Ins11] Institute of Electrical and Electronics Engineers (IEEE). *1609.4-2010 (Revision of IEEE Std 1609.4-2006) – Standard for Wireless Access in Vehicular Environments (WAVE) – Multi-channel Operation*, February 2011.

[ITU94] ITU: International Telecommunication Union. *Recommendation X.200 : Information technology – Open Systems Interconnection – Basic Reference Model: The basic model*, July 1994.

[IW08] Khaled Ibrahim and Michele C. Weigle. CASCADE: Cluster-based accurate syntactic compression of aggregated data in vanets. In *GLOBE-*

COM '08: Proceedings of the IEEE Global Telecommunications Conference – Workshops, November 2008.

[JBW05] Sven Jaap, Marc Bechler, and Lars Wolf. Evaluation of routing protocols for vehicular ad hoc networks in typical road traffic scenarios. In *EUNICE '05: Proceedings of the 11th Open European Summer School*, Colmenarejo, Spain, July 2005.

[JWZ01] Rui Jiang, Qingsong Wu, and Zuojin Zhu. Full velocity difference model for a car-following theory. *Physical Review E*, 64(1), July 2001.

[KEÖÖ04] Gökhan Korkmaz, Eylem Ekici, Füsun Özgüner, and Ümit Özgüner. Urban multi-hop broadcast protocol for inter-vehicle communication systems. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 76–85, October 2004.

[KIM07] Kurtis Kredo II and Prasant Mohapatra. Medium access control in wireless sensor networks. *Computer Networks: The International Journal of Computer and Telecommunications*, 51(4):961–994, March 2007.

[Kir04] Donald E. Kirk. *Optimal Control Theory: An Introduction*. Dover Publications, April 2004.

[KK00] Brad Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *MobiCom '00: Proceedings of the 6th Annual ACM International Conference on Mobile Computing and Networking*, pages 243–254, August 2000.

[KKM09] Markus Kerper, Wolfgang Kiess, and Martin Mauve. Coordinated VANET experiments: a methodology and first results. In *VANET '09: Proceedings of the 6th ACM International Workshop on VehiculAr InterNETworking*, pages 109–116, September 2009.

[KMT98] Frank P. Kelly, Aman K. Maulloo, and David K.H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, March 1998.

[Kra98] Stefan Krauß. *Microscopic Modeling of Traffic Flow: Investigation of Collision Free Vehicle Dynamics*. PhD thesis, Deutsches Zentrum für Luft- und Raumfahrt, September 1998.

[LA05] Thomas D. C. Little and Ashish Agarwal. An information propagation scheme for VANETs, September 2005.

[LCS+05] Christian Lochert, Murat Caliskan, Björn Scheuermann, Andreas Barthels, Alfonso Cervantes, and Martin Mauve. Multiple simulator interlinking environment for inter vehicle communication. In *VANET '05: Proceedings of the 2nd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 87–88, September 2005.

[Lee08] Edward A. Lee. Cyber physical systems: Design challenges. In *ISORC '08: Proceedings of the 11th International Symposium on Object-Oriented Real-Time Distributed Computing*, May 2008.

[LHK+00] Georg Lerner, Almut Hochstaedter, Ronald Kates, Cesim Demir, Jürgen Meier, and Andreas Poschinger. The interplay of multiple scales in traffic flow: Coupling of microscopic, mesoscopic and macroscopic simulation. In *ITSWC '00: Proceedings of the 7th World Congress and Exhibition on Intelligent Transportation Systems and Services (ITS)*, November 2000.

[LRLK10] Jeongkeun Lee, Jiho Ryu, Sung-Ju Lee, and Ted Taekyoung Kwon. Improved modeling of IEEE 802.11a PHY through fine-grained measurements. *Computer Networks*, 54(4):641–657, March 2010.

[Lüb08] Andreas Lübke. Der aktuelle Stand der Car-to-X Kommunikation (*in German language*). In *VDE-Kongress 2008*, November 2008.

[MGFS08] Kristijan Macek, Dizan Alejandro Vasquez Govea, Thierry Fraichard, and Roland Siegwart. Safe vehicle navigation in dynamic urban scenarios. In *ITSC '08: Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems*, October 2008.

[MGVP11] Vicente Milanés, Jorge Godoy, Jorge Villagrá, and Joshué Pérez. Automated on-ramp merging system for congested traffic situations. *IEEE Transactions on Intelligent Transportation Systems*, 12(2):500–508, June 2011.

[MH02] R. Miller and Qingfeng Huang. An adaptive peer-to-peer collision warning system. In *VTC '02-Spring: Proceedings of the 55th IEEE Vehicular Technology Conference*, pages 317–321, May 2002.

[Mil61] Alan J. Miller. A queueing model for road traffic flow. *Journal of the Royal Statistical Society. Series B (Methodological)*, 23(1):64–90, 1961.

[MMP+05] Kirsten Matheus, Rolf Morich, Ingrid Paulus, Cornelius Menig, Andreas Lübke, Bernd Rech, and Will Specks. Car-to-car communication – market introduction and success factors. In *ITS '05: Proceedings of the 5th European Congress and Exhibition on Intelligent Transport Systems and Services*, June 2005.

[MOP+11] Vicente Milanés, Enrique Onieva, Joshué Pérez, Jorge Villagrá, Jorge Godoy, Javier Alonso, Carlos González, Teresa de Pedro, and Ricardo García. Autopia program advances: How to automate the traffic? In *EUROCAST '11: Proceedings of the 13th International Conference on Computer Aided Systems Theory*, volume 2, pages 374–381, February 2011.

[Mor06] Ricardo Morla. Sentient future competition: Vision of congestion-free road traffic and cooperating objects. In *EWSN '06: Proceedings of the 3rd European Conference on Wireless Sensor Networks*, February 2006.

[Mor07] Marc Torrent Moreno. *Inter-Vehicle Communications: Achieving Safety in a Distributed Wireless Environment: Challenges, Systems and Protocols*. PhD thesis, University of Karlsruhe, Faculty of Computer Science, Karlsruhe, 2007.

[MTHH09] Jens Mittag, Florian Thomas, Jérôme Härri, and Hannes Hartenstein. A comparison of single- and multi-hop beaconing in VANETs. In *VANET '09: Proceedings of the 6th ACM International Workshop on VehiculAr Inter-NETworking*, pages 69–78, September 2009.

[MUH03] Steffen Müller, Michael Uchanski, and J. Karl Hedrick. Cooperative estimation using vehicle communications. In *Vehicle System Dynamics: International Journal of Vehicle Mechanics and Mobility*, volume 39, pages 121–133, February 2003.

[MWH01] Martin Mauve, Jörg Widmer, and Hannes Hartenstein. A survey on position-based routing in mobile ad-hoc networks. *IEEE Network Magazine*, 15(6):30–39, November 2001.

[MZMP06] Zhaomin Mo, Hao Zhu, Kia Makki, and Niki Pissinou. MURU: A multi-hop routing protocol for urban vehicular ad hoc networks. In *MobiQuitous '06: Proceedings of the 3rd Annual International Conference on Mobile and Ubiquitous Systems: Networking & Services – Workshops*, pages 1–8, July 2006.

[NDLI04] Tamer Nadeem, Sasan Dashtinezhad, Chunyuan Liao, and Liviu Iftode. TrafficView: traffic data dissemination using car-to-car communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 8:6–19, July 2004.

[Nil95] Lena Nilsson. Safety effects of adaptive cruise control in critical traffic situations. November 1995.

[NRWB09] Josiane Nzouonta, Neeraj Rajgure, Guiling Wang, and Cristian Borcea. VANET routing on city roads using real-time vehicular traffic information. *IEEE Transactions on Vehicular Technology*, 58(7), 2009.

[NS92] Kai Nagel and Michael Schreckenberg. A cellular automaton model for freeway traffic. *Journal de Physique I France 2*, 2(12):2221–2229, 1992.

[NSN] NSNAM. ns-3: network simulator 3. http://www.nsnam.org.

[NTCS99] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu. The broadcast storm problem in a mobile ad hoc network. In *MobiCom '99: Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 151–162, July 1999.

[OT04] Johan J. Olstam and Andreas Tapani. *Comparison of Car-following Models*. VTI meddelande. Swedish National Road and Transport Research Institute (VTI), 2004.

[OZL11] Hassan A. Omar, Weihua Zhuang, and Li Li. VeMAC: A TDMA-based MAC protocol for reliable broadcast in VANETs. In *INFOCOM WKSHPS '11: Proceedings of the IEEE Conference on Computer Communications Workshops*, April 2011.

[Pip53] Louis A. Pipes. An operational analysis of traffic dynamics. *Journal of Applied Physics*, 24(3):274–281, March 1953.

[Pos81] Jon Postel. Transmission control protocol. RFC 793 (INTERNET STANDARD), September 1981. Updated by RFCs 1122, 3168, 6093, 6528.

[PPB+00] J. Pierowicz, B. Pirson, A. Bittner, M. Lloyd, and E. Jocoy. Intersection collision avoidance using its countermeasures. Technical Report HS-809 171, U.S. Department of Transportation, National Highway Traffic Safety Administration, 2000.

[PRL+08] Michal Piorkowski, Maxim Raya, Ada Lugo, Panos Papadimitratos, Matthias Grossglauser, and Jean-Pierre Hubaux. TraNS: Realistic joint traffic and network simulator for VANETs. *ACM SIGMOBILE Mobile Computing and Communications Review*, 12(1):31–33, 2008.

[PV05] Anuj Puri and Pravin Varaiya. Driving safely in smart cars. In *ACC '05: Proceedings of the 2005 IEEE American Control Conference*, pages 3597–3599, June 2005.

[QM08] Yi Qian and Nader Moayeri. Design of secure and application-oriented VANETs. In *VTC '08-Spring: Proceedings of the 67th IEEE Vehicular Technology Conference*, pages 2794–2799, May 2008.

[RCC10] Tom Robinson, Eric Chan, and Erik Coelingh. Operating platoons on public motorways: An introduction to the SARTRE platooning programme. In *ITSWC '10: Proceedings of the 17th World Congress on Intelligent Transportation Systems and Services (ITS)*, pages 109–116, October 2010.

[RCCL06] Craig L. Robinson, Lorenzo Caminiti, Derek Caveney, and Kenneth P. Laberteaux. Efficient coordination and transmission of data for cooperative vehicular safety applications. In *VANET '06: Proceedings of the 3rd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 10–19, September 2006.

[RFM08] Dario Rossi, Roberta Fracchia, and Michela Meo. VANETs: Why use beaconing at all? In *ICC '08: Proceedings of the IEEE International Conference on Communications*, pages 2745–2751, May 2008.

[rGPPG99] 3rd Generation Partnership Project (3GPP). Ts 21.101: Technical specifications and technical reports for a UTRAN-based 3GPP system, 1999.

[RML02] M. Rudack, M. Meincke, and M. Lott. On the dynamics of ad hoc networks for inter vehicle communications (IVC). In *ICWN '02: Proceedings of the International Conference on Wireless Networks*, June 2002.

[RRS09] F.J. Ros, P.M. Ruiz, and I. Stojmenovic. Reliable and efficient broadcasting in vehicular ad hoc networks. In *VTC '09-Spring: Proceedings of the 69th IEEE Vehicular Technology Conference*, April 2009.

[RS07] Shahram Rezaei and Raja Sengupta. Kalman filter-based integration of dgps and vehicle sensors for localization. In *IEEE Transactions on Control Systems Technology*, volume 15, pages 1080–1088, November 2007.

[RSKX07] Shahram Rezaei, Raja Sengupta, Hariharan Krishnan, and Guan Xu. Reducing the communication required by DSRC-based vehicle safety systems. In *ITSC '07: Proceedings of the 10th International IEEE Conference on Intelligent Transportation Systems*, pages 361–366, September 2007.

[RvS07] Christian Reinl and Oskar von Stryk. Optimal control of multi-vehicle-systems under communication constraints using mixed-integer linear programming. In *RoboComm '07: Proceedings of the 1st International Conference on Robot Communication and Coordination*, pages 1–8, October 2007.

[SA09] Gérard Ségarra and Martin Arndt. TR 102 638 v1.1.1; Intelligent Transport Systems (ITS); vehicular communications; basic set of applications; definitions, 2009.

[SAHK10] Rashid A. Saeed, Mahamat A. Abakar, Aisha A. Hassan, and Othman O. Khalifa. Design and evaluation of lightweight IEEE 802.11p-based TDMA MAC method for road side-to-vehicle communications. In *ICCCE '10: Proceedings of the 2nd International Conference on Computer and Communication Engineering*, pages 1–5, May 2010.

[SCB10] Razvan Stanica, Emmanuel Chaput, and André-Luc Beylot. Comparison of CSMA and TDMA for a heartbeat VANET application. In *ICC '10: Proceedings of the 45th IEEE International Conference on Communications*, pages 1–5, May 2010.

[SCInR12] Tossaphol Settawatcharawanit, Supasate Choochaisri, Chalermek Intanagonwiwat, and Kultida Rojviboonchai. V-DESYNC: Desynchronization for beacon broadcasting on vehicular networks. In *VTC '12-Spring: Proceedings of the 75rd IEEE Vehicular Technology Conference*, pages 1–5, May 2012.

[SFTE06] Sascha Schnaufer, Holger Füßler, Matthias Transier, and Wolfgang Effelsberg. Vehicular ad-hoc networks: Single-hop broadcast is not enough. In *WIT '06: Proceedings of the 3rd International Workshop on Intelligent Transportation*, pages 49–54, March 2006.

[SFUH04] Masashi Saito, Mayuko Funai, Takaaki Umedu, and Teruo Higashino. Inter-vehicle ad-hoc communication protocol for acquiring local traffic information. In *ITSWC '04: Proceedings of the 11th World Congress on Intelligent Transportation Systems and Services (ITS)*, October 2004.

[SH99] Darbha Swaroop and J. Karl Hedrick. Constant spacing strategies for platooning in automated highway systems. *Journal of Dynamic Systems, Measurement, and Control*, 121(3):462–470, September 1999.

[SJ04] Amit Kumar Saha and David B. Johnson. Modeling mobility for vehicular ad-hoc networks. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 91–92, October 2004.

[SLS+10] Robert K. Schmidt, Tim Leinmüller, Elmar Schoch, Frank Kargl, and Günter Schäfer. Exploration of adaptive beaconing for efficient intervehicle safety communication. *IEEE Network Magazine*, 24(1):14–19, January 2010.

[SNC+09] Steven E. Shladover, Christopher Nowakowski, Delphine Cody, Fanping Bu, Jessica O'Connell, John Spring, Susan Dickey, and David Nelson. Effects of cooperative adaptive cruise control on traffic flow: Testing drivers' choices of following distances. PATH Research Report UCB-ITS-PRR-2009-23, Institute of Transportation Studies, UC Berkeley, March 2009.

[SPS09] Henrik Schumacher, Christian Priemer, and Eric N. Slottke. A simulation study of traffic efficiency improvement based on car-to-x communication. In *VANET '09: Proceedings of the 6th ACM International Workshop on VehiculAr Inter-NETworking*, pages 13–22, September 2009.

[SSH98] Peter Seiler, Bongsob Song, and J. Karl Hedrick. Development of a collision avoidance system. In *Object Detection, Collision Warning and Avoidance Systems*, pages 97–103. SAE International, February 1998.

[Ste08] Seiji S.C. Steimetz. Defensive driving and the external costs of accidents and travel delays. *Transportation Research Part B*, 42(9):703–724, November 2008.

[STL+07] Björn Scheuermann, Matthias Transier, Christian Lochert, Martin Mauve, and Wolfgang Effelsberg. Backpressure multicast congestion control in mobile ad-hoc networks. In *CoNEXT '07: Proceedings of the 3rd International Conference on Emerging Networking Experiments and Technologies*, pages 265–275, December 2007.

[SW92] Susan L. Streisand and Jean Walrand. A communication architecture for IVHS. Institute of Transportation Studies UCB-ITS-PRR-92-10, Institute of Transportation Studies, UC Berkeley, December 1992.

[TH01] Helmut Tschöke and Hanns-Erhard Heinze. Einige unkonventionelle Betrachtungen zum Kraftstoffverbrauch von PKW (*in German language*). In *Magdeburger Wissenschaftsjournal*, volume 1-2, pages 11–18, January 2001.

[TH02] Martin Treiber and Dirk Helbing. Realistische Mikrosimulation von Straßenverkehr mit einem einfachen Modell (*in German language*). In *ASIM '02: Proceedings of the 16th Symposium Simulationstechnik of the Arbeitsgemeinschaft Simulation*, September 2002.

[THH00] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805–1824, August 2000.

[TMSH05] Marc Torrent-Moreno, Paolo Santi, and Hannes Hartenstein. Fair sharing of bandwidth in VANETs. In *VANET '05: Proceedings of the 2nd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 49–58, September 2005.

[TMSH06] Marc Torrent-Moreno, Paolo Santi, and Hannes Hartenstein. Distributed fair transmit power adjustment for vehicular ad hoc networks. In *SECON '06: Proceedings of the 3rd Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks*, volume 2, pages 479–488, September 2006.

[TWB10] Ozan K. Tonguz, Nawaporn Wisitpongphan, and Fan Bai. DV-CAST: a distributed vehicular broadcast protocol for vehicular ad hoc networks. *IEEE Wireless Communications*, 17:47–56, April 2010.

[TY10] Antony Tang and Alice Yip. Collision avoidance timing analysis of DSRC-based vehicles. *Accident Analysis & Prevention*, 42(1):182–195, January 2010.

[Var93] Pravin Varaiya. Smart cars on smart roads: Problems of control. *IEEE Transactions on Automatic Control*, 38:195–207, February 1993.

[vEKH10] Martijn van Eenennaam, Georgios Karagiannis, and Geert J. Heijenk. Towards scalable beaconing in vanets. In *ERCIM '04: Proceedings of the 4th Workshop on eMobility*, pages 103–108, May 2010.

[vEWKH09] Martijn van Eenennaam, Wouter Klein Wolterink, Georgios Karagiannis, and Geert J. Heijenk. Exploring the solution space of beaconing in VANETs. In *VNC '09: Proceedings of the 1st Annual IEEE Vehicular Networking Conference*, pages 1–8, October 2009.

[WC02] Brad Williams and Tracy Camp. Comparison of broadcasting techniques for mobile ad hoc networks. In *Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking & computing*, pages 194–205, June 2002.

[WFGH04] Hao Wu, Richard Fujimoto, Randall Guensler, and Michael Hunter. MDDV: a mobility-centric data dissemination algorithm for vehicular networks. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 47–56, October 2004.

[Wie74] Rainer Wiedemann. *Simulation des Straßenverkehrsflusses (*in German language*)*. PhD thesis, Institut für Verkehrswesen (IfV; Institute for Transport Studies), University of Karlsruhe, 1974.

[Wie11] Adam Wierman. Fairness and scheduling in single server queues. *Surveys in Operations Research and Management Science*, 16(1):39–48, January 2011.

[WKR07] Ziyuan Wang, Lars Kulik, and Kotagiri Ramamohanarao. Proactive traffic merging strategies for sensor-enabled cars. In *VANET '07: Proceedings of the 4th ACM International Workshop on Vehicular Ad Hoc Networks*, pages 39–48, September 2007.

[WKR09] Ziyuan Wang, Lars Kulik, and Kotagiri Ramamohanarao. Robust traffic merging strategies for sensor-enabled cars using time geography. In *GIS '09: Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 362–371, November 2009.

[XMKS04] Qing Xu, Tony Mak, Jeff Ko, and Raja Sengupta. Vehicle-to-vehicle safety messaging in DSRC. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 19–28, October 2004.

[YEY+04] Jijun Yin, Tamer ElBatt, Gavin Yeung, Bo Ryu, Stephen Habermas, Hariharan Krishnan, and Timothy Talty. Performance evaluation of safety applications over DSRC vehicular ad hoc networks. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 1–9, October 2004.

[YF08] Saleh Yousefi and Mahmood Fathy. Metrics for performance evaluation of safety applications in vehicular ad hoc networks. *Transport*, 23(4):291–298, 2008.

[YGW08] Lin Yang, Jinhua Guo, and Ying Wu. Channel adaptive one hop broadcasting for VANETs. In *ITSC '08: Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems*, volume 11, October 2008.

[YLZV04] Xue Yang, Jie Liu, Feng Zhao, and Nitin H. Vaidya. A vehicle-to-vehicle communication protocol for cooperative collision warning. In *MobiQuitous '04: Proceedings of the 1st Annual International Conference on Mobile and Ubiquitous Systems: Networking & Services*, pages 114–123, August 2004.

[ZA02] Eustathia Ziouva and Theodore Antonakopoulos. Csma/ca performance under high traffic conditions: throughput and delay analysis. *Computer Communications*, 25(3):313–321, 2002.

[ZC06] Jing Zhao and Guohong Cao. VADD: Vehicle-assisted data delivery in vehicular ad hoc networks. In *INFOCOM '06: Proceedings of the 25th Annual Joint Conference of the IEEE Computer and Communications Societies*, April 2006.

[ZF06] Hongqiang Zhai and Yuguang Fang. Distributed flow control and medium access in multihop ad hoc networks. *IEEE Transactions on Mobile Computing*, 5(11):1503–1514, November 2006.

Index