Topics in Feedback Control for ITS Applications in Smart Cities

vorgelegt von
Dipl.-Ing.
Florian Häusler
geb. in Reinbek

von der Fakultät IV - Elektrotechnik und Informatik
der Technische Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
- Dr.-Ing -

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. Clemens Gühmann
Gutachter: Prof. Dr. Robert Shorten
Gutachter: Prof. Dr. Jörg Raisch
Gutachter: Prof. Dr. Joachim Taiber

Tag der wissenschaftlichen Aussprache: 13. Oktober 2016

Berlin 2016

Abstract

This thesis introduces new applications from the domain of Intelligent Transportation Systems to address some of the central challenges of urban road traffic. All applications exploit the potential of recent developments in the automotive field, in particular vehicular networks, new kinds of vehicles such as electric and full hybrid vehicles and results from networked control theory. By setting up a feedback loop around a large number of vehicles applications are able to address general road traffic problems in a very direct way - in a way which is not possible by per-vehicle measures. Domain-specific requirements (e.g. preserve drivers' privacy) are considered in the design of each application and its underlying control and optimization algorithm.

It is shown that multiple problems can be approached with a simple system architecture: On the one hand, there is a central infrastructure, which has access to general, system-wide information such as air pollution levels, traffic flow information, etc. On the other hand, there are vehicles - in particular full hybrids - which are equipped with advanced communication technology. Communication technology integrates both parts into a feedback control loop. It is shown that in many cases it is sufficient that the infrastructure broadcasts minimal information such that vehicles act in a collaborative way.

The introduced applications address the problem of air pollution in urban environments, optimize emission in a fleet, target the problem of supply variations in electric grids with hybrid vehicles, provide balancing solutions for electric vehicle charging stations and aim at utilizing given parking spaces in a more efficient way.

The performance is discussed with help of extensive simulations. (Limited) mathematical analyses for algorithmic approaches are included or cited.

It is considered that proposed applications might recommend behaviour which does not optimize individual utilities - while they do optimize other aspects such as the total utility of all users. Thus, the question of user acceptance is crucial. In this sense, selected user acceptance tests and surveys have been performed. Furthermore a technical platform is introduced which provides the possibility to test collaborative applications in the field (including user acceptance) without requiring a large number of vehicles.

Zusammenfassung

Die Dissertation beschreibt neue Anwendungen im Bereich Intelligent Transportation Systems, um zentralen Problemen des urbanen Straßenverkehrs zu begegnen. Die hier betrachteten Probleme können nur gelöst werden, wenn Fahrzeuge miteinander kooperieren und ihr Verhalten aufeinander abstimmen. Die entwickelten Lösungsansätze verbindet, dass sie die Möglichkeiten ausnutzen, die mit der Entwicklung neuer Fahrzeugtypen - insbesondere Elektrohybridfahrzeugen - und ihrer zunehmenden Vernetzung einhergehen. Dabei entsteht ein dem Internet verwandtes Systemumfeld auf das Ansätze aus dem Bereich Netzwerksteuerung mit wenigen Anpassungen übertragen werden können. Indem eine große Zahl von Fahrzeugen in einen übergeordneten Regelkreis integriert werden, können grundsätzliche Verkehrsprobleme in Smart Cities auf eine sehr direkte Art adressiert werden - viel direkter als es mit fahrzeugindividuellen, nicht-kooperativen Ansätzen möglich ist. Es wird gezeigt, dass dabei relevante, domänenspezifische Anforderungen (z.B. Sicherung der Privatsphäre) berücksichtigt werden können.

Dabei kann eine einfache Systemarchitektur genutzt werden: Auf der einen Seite steht eine zentrale Infrastruktur (z.B. Verkehrsmanagementzentrale), die auf entsprechende Informationen Zugriff hat (z.B. Luftbelastung, Verkehrsflussinformationen). Auf der anderen Seite stehen mit Kommunikationstechnologie ausgestatteten Fahrzeuge. Die entwickelten Anwendungen belegen, dass es in vielen Fällen ausreicht, einfache Informationen via Broadcasts durch die zentrale Infrastruktur zu verbreiten, so dass sich ein abgestimmtes, kollaboratives Verhalten der Fahrzeuge einstellt.

Auf diese Weise können mehrere Probleme im urbanen Verkehr adressiert werden. Dazu zählen die lokale Luftbelastung in städtischen Räumen, die Optimierung der Emission einer Fahrzeugflotte, Schwankungen im Stromangebot in Elektrizitätsnetzen und die ineffiziente Nutzung von Parkplätzen und Ladesäulen für Elektrofahrzeuge.

Alle Anwendungen werden mit Hilfe von umfangreichen Simulation validiert. Auf erweiternde mathematische Diskussionen und Herleitungen wird an entsprechender Stelle auf jeweilige Publikationen verwiesen.

Da die vorgestellten Lösungsansätze den systemweiten und nicht notwendigerweise den individuellen Nutzen von Teilnehmern optimiert, ist die Nutzerakzeptanz kritisch. Zu ausgewählten Anwendungen werden daher die Ergebnisse von Akzeptanzstudien vorgestellt. In diesem Zusammenhang wurde zudem eine technische Plattform entwickelt, die es ermöglicht kooperative Anwendungen im Feld zu testen, ohne auf eine große Anzahl von Fahrzeugen zurückgreifen zu müssen.

Acknowledgements

There are no proper words to convey my deep gratitude and respect for my thesis advisor and Doktorvater Prof. Robert Shorten, who has shared his ideas with adorable selflessness. He has been and still is a true inspiration. This does not only account for his wisdom, generous guidance, enthusiastic encouragement and unconventional approaches. This thesis represents only a small part of the impact he has had on my life and I will always be thankful for that.

My gratitude goes to Prof. Jörg Raisch for the support during the last year of the thesis. His empathy and targeted feedback could not been more constructive and helpful. I thank him and Prof. Joachim Taiber for readily accepting the task of being examiners.

I would also like to thank my colleagues at Fraunhofer FOKUS and Technical University of Berlin, in particular Ilja Radusch, who gave me the opportunity to do research on topics which were not directly covered by project work. I would also I like to thank my colleagues Sebastian Schwardt, Stefan Lobach, Kevin Rhinow, Bartosz Rytt, Kay Massow, Thomas Hecker, Robert Protzmann, Björn Schünemann and David Rieck.

Furthermore I like to thank for the collaboration, corrections, discussions, feedback and patient help in particular from Emanuele Chrisostomi and Arieh Schlote as well as Mahsa Faizrahnemoon, Wynita Griggs, Rodrigo Ordóñez-Hurtado, Jia Yuan Yu, Prof. Fabian Wirth, Yingqi Gu and Sebastian Zander.

I wish to express my heartfelt gratitude to my parents Ute and Nikolaus Häusler to whom I owe so much, as well as to my friends for their love and support over the years.

Finally, I would like to thank the love of my life Antje. She has encouraged me to work on the thesis and has asked me to take a break when it was needed. I thank her for the wonderful years and our wonderful daughters Leni and Maja. I love you. This thesis is for you.

Publications

The following papers were published or have been submitted for publication during the course of this thesis. The order of authors corresponds to the one on publications.

1. *Cooperative regulation and trading of emissions using plug-in hybrid vehicles*, in IEEE Transactions on Intelligent Transportation Systems, Volume 14, Issue 4, June 2013, pp. 1572-1585.

   Joint work from Arieh Schlote, Florian Häusler, Thomas Hecker, Astrid Bergmann, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

2. *A framework for optimal real-time emissions trading in large-scale vehicle fleets*, in Proceedings of the 9th ITS European Congress, 2013.
   The paper received the Best Scientific Paper Award.

   Joint work from Florian Häusler, Mahsa Faizrahnemoon, Arieh Schlote, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

3. *A framework for optimal real-time emissions trading in large-scale vehicle fleets*, in Intelligent Transport Systems (IET), Volume 9, Issue 3, April 2015, pp. 275-284.

   Joint work from Florian Häusler, Mahsa Faizrahnemoon, Arieh Schlote, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

4. *Stochastic park-and-charge balancing for fully electric and plug-in hybrid vehicles*, in IEEE Transactions on Intelligent Transportation Systems, Volume 15, Issue 2, November 2013, pp. 895-901.

   Joint work from Florian Häusler, Arieh Schlote, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

5. *Stochastically balanced parking and charging for fully electric and plug-in hybrid vehicles*, in Proceedings of the 1st International Conference on Connected Vehicles and Expo, 2012, pp. 341-342.

   Joint work from Florian Häusler, Emanuele Crisostomi, Arieh Schlote, Ilja Radusch and Robert Shorten.

6. *Cooperative regulation of emissions using plug-in hybrid vehicles*, in Proceedings of the 1st International Conference on Connected Vehicles and Expo, 2012, pp. 201-202.

   Joint work from Arieh Schlote, Florian Häusler, Thomas Hecker, Astrid Bergmann, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

7. *Smart Procurement Of Naturally Generated Energy (SPONGE) for PHEVs*, in Proceedings of the 4th International Conference on Connected Vehicles & Expo, October 2015.

   Joint work from Florian Häusler, Yingqi Gu, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

8. *Smart Procurement Of Naturally Generated Energy (SPONGE) for PHEVs*, accepted for publication in the International Journal of Control in 2016.

   Joint work from Yingqi Gu, Florian Häusler, Wynita Griggs, Emanuele Crisostomi, Ilja Radusch and Robert Shorten.

9. *Closed-loop flow regulation with balanced routing*, 3rd International Conference on Connected Vehicles and Expo, November 2014, pp. 1054-1055.

   Joint work from Florian Häusler, Rodrigo Ordóñez-Hurtado, Wynita Griggs, Ilja Radusch and Robert Shorten.

10. *A large-scale SUMO-based emulation platform*, in IEEE Transactions on Intelligent Transportation Systems, Issue 99, May 2015, pp. 1-10.

    Joint work from Wynita Griggs, Rodrigo. Ordóñez-Hurtado, Emanuele Crisostomi, Florian Häusler, Kay Massow and Robert Shorten.

11. *On the Design of Campus Parking Systems with QoS Guarantees*, IEEE in Transactions on Intelligent Transportation Systems, Issue 19, Number 5, 2016, pp. 1428-1437.

    Joint work from Wynita Griggs, Jia Yu, Fabian Wirth, Florian Häusler and Robert Shorten.

12. *A test architecture for V-2-X cooperative systems field operational tests*, in Proceedings of the 9th International Conference on Intelligent Transport Systems Telecommunications, October 2009, pp. 616-621.

    Joint work with Andrea Tomatis, Markus Miche, Florian Häusler, Massimiliano Lenardi, Thomas Michael Bohnert and Ilja Radusch.

Further, I supported and guided multiple Bachelor, Master and Diploma thesis. Some of the practical work done was motivated by the general topic of this thesis. Selected results integrated within this thesis are from the following works:

13. *Adaptive Urban Emission Control by Online Rerouting*, Master thesis at Technical University of Berlin in 2013 from Stefan Lobach. The work includes extensive simulations using a dispersion simulator. Results are included in Section 6.3.

14. *Entwicklung und Validierung von Anreizen zur Akzeptanzsteigerung von systemoptimierenden ITS Applikationen*, Master thesis at Technical University of Berlin in 2015 from Bartek Rytt. The work includes experiments in the context of routing. Some selected results are included in Section 9.4.1.

15. *Modellierung und wirtschaftliche und empirische Analyse eines kollaborativen Park Sharing Konzeptes*, Master thesis at Technical University of Berlin in 2013 from Kevin Rhinow. The work includes surveys in the context of user acceptance targeting parking concepts. Selected results are included in Section 9.4.2.

Contents

List of abbreviations and notations

ABS	Anti-lock Braking System
ACEA	European Automobile Manufacturers Association
ADMM	Alternating direction method of multipliers
AIMD	Additive Increase Multiple Decrease
CAM	Cooperative Awareness Message
CAN bus	Controller Area Network bus
CIS	Central ITS Station
CNG	Compressed Natural Gas
CO	Carbon monoxide
CO_2	Carbon dioxide
CPS	Cyber-Physical System
DENM	Decentralized Environmental Notification Message
EC	European Commission
ECU	Electronic Control Unit
ETSI	European Telecommunications Standards Institute
EU ETS	European Union Emission Trading System
EV	Electric Vehicle
FOT	Field Operation Test
G2V	Grid-to-Vehicle
GDP	Gross Domestic Product
GPS	Global Positioning System
HEV	Hybrid Electric Vehicle
HIL	Hardware-in-the-loop

HMI	Human Machine Interface
ICT	Information and Communication Technology
IEEE	Institute of Electrical and Electronics Engineers
ISO	International Organization for Standardization
ITS	Intelligent Transportation Systems
JAMA	Japan Automobile Manufacturers Association
KAMA	Korean Automobile Manufacturers Association
LDV	Light Duty Vehicles
LEZ	Low Emission Zone
LPG	Liquefied Petroleum
MIMD	Multiple Increase Multiple Decrease
NO_x	Oxigen of Nitrogen
O	Ozon
OBDII	On-board diagnostics interface
OSM	Open Street Map
P2P	Peer-to-Peer
PHEV	Plug-in Hybrid Electric Vehicle
PID	Proportional Integral Differential
PIS	Personal ITS Station
PM_{10}	Particular matter up to 10 micrometers in size
$PM_{2.5}$	Particular matter up to 2.5 micrometers in size
RED	Random Early Detection
RIS	Roadside ITS Station
RTT	Round Trip Time
SO_2	Primary sulfur dioxide
SO_3	Sulfur oxides

SUMO	Simulation of Urban Mobility (traffic simulation tool)
TMC	Traffic Message Channel
TPEG	Transport Protocol Experts Group
V2G	Vehicle-to-Grid
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-X
VANET	Vehicle Ad-Hoc Network
VIL	Vehicle-in-the-loop
VIS	Vehicle ITS Station
VOC	Volatile Organic Compounds
VSC	Vehicle Safety Consortium
WHO	World Health Organization

1. Introduction

Summary: *This opening chapter highlights some selected challenges in road traffic in the context of smart cities. It is shown that key issues are often caused by two main characteristics of road traffic: synchronization and the lack of context awareness.*
Three general developments are identified which make the proposed solutions very attractive: (1) communication technology applied to the automotive domain (2) vehicle technology (electric and hybrid vehicles) and (3) increasing public awareness in regard to eco-friendly mobility.
Proposed applications follow a different approach to those which are available today: they implement large-scale feedback strategies and integrate a huge number of traffic participants to establish a "collaborative" behaviour. These are able to address systematic problems, such as pollution, parking, congestion etc. This chapter provides relevant context for the detailed introduction and discussion of applications in subsequent chapters. It outlines the document structure and how selected developments and results of this work were developed in cooperation with other researchers.

1.1. Individual road traffic and smart cities

Smart City is a term which is still evolving and is used in many contexts. A common feature of many definitions is that ICT is playing a key role [36]. Furthermore, mobility is one of the central pillars of a smart city. Accordingly, the European Commission (EC) defines (smart) mobility "ICT supported, integrated transport and logistics systems" [111]. This understanding is equivalent with the one of Intelligent Transportation Systems (ITS) - a term which was found in the 1980s to underline the (expected) impact of ICT in traffic [177]: "ITS applies advanced communication, information and electronics technology to solve transportation problems such as traffic congestion, safety, transport efficiency and environmental conservation." [54] Despite the fact that there is a huge research community in ITS, road transport is still facing many challenges. The EC divides them in three basic groups: (1) congestion, (2) safety and (3) environmental pollution [34]. In fact, one might divide the latter in two distinct challenges: first securing human health from CO, particular matter etc. and second minimizing impact on global warming through CO_2. In addition, parking could also be identified as an additional problem as well as the integration of electric vehicles with the electric grid. One might think that there are straightforward solutions to the parking problem. As it will be outlined, parking gives rise to subtle challenges that can be addressed in a resource allocation context. But the expected increase of plug-in-hybrids and fully electric vehicles also cause new challenges. Selected ones (i.e. queueing at charging stations for electric vehicles) were not yet addressed at all.

In the remainder of this chapter, selected problems of road traffic are briefly described. The selection of problems are addressed by the developed applications outlined in greater detail in the subsequent chapters.

1.1.1. Congestion

Congestion is probably the most severe effect of road traffic [34], monitored in almost all urban environments. According to the Texas A&M Transportation Institute which publishes statistical data on urban mobility in the U.S., the average U.S. auto commuter in greater urban areas spent more than 1,100 Dollars per year because of congestion, excessed 24 Gallons ($\sim$ 90 litres) of gasoline and has a travel time that is increased by almost 30 % in 2011.

Basically, there are three ways to address the problem [175]: (1) extend the road infrastructure, (2) reduce the overall demand for road traffic and (3) improve efficiency of the road network usage. The most obvious way to tackle the congestion problem is to build more roads. It has been shown though, that extra road infrastructure is not as attractive as it seems at a first sight. The construction of new roads has general social and environmental effects and costs, which outperform potential benefits for the vehicle drivers [43]. On reason is, that in many cases - when traffic demand is elastic - the extension of road infrastructure does effect in more traffic demand - resulting again in congestion [121]. Furthermore, especially in urban environments, there is no space available to build any new roads. The reduction of overall demand for road traffic is assumed to be influenced by policy measures mainly, such as increasing the attractiveness of the public transport, ride sharing etc. Clearly, improving the efficiency of road network usage is the most attractive instrument. There are various approaches deployed and available. Simple examples are the deployment of traffic lights, (adaptive) speed limits, dynamic lane access, freeway incident management systems, freeway ramp metering and many more.

1.1.2. Parking

Parking is closely related to congestion as vehicles searching for parking spots represent a basic component of it. In San Francisco it is expected that about 30 % of vehicle traffic is caused by drivers who are searching for an open spot. According to the IBM Global Parking Survey from 2011, this is a representative share for many cities, both in developed and emerging economies. On a global level, 5 to 10 % of vehicle traffic worldwide is caused by parking search [100]. Moreover, driving around searching for open spots is associated with frustration, delay and pollution. Expected increases in auto commuting to urban areas and in car ownership will increase the problem. Besides congestion, land use is another problem. Los Angeles is a dramatic example: About 331 hectares are reserved for parking (including multi-story car parks etc.), while Los Angeles has a land area of 408 hectares [112]. Looking at electric vehicles, parking is even more critical as vehicles may run out of power while searching for available parking lots.

A state-of-the-art countermeasure for the pressing parking problem is to extent parking spaces. In addition to the necessary (large) investments, the approach faces the same problems and effects as the approach to build extra roads to relax congestion problems: the availability of (off-street) parking spaces promotes driving in general, even in cities with good public transport systems [97]. In this sense, recent legislative developments take the opposite approach: several European cities limit the number of parking spaces or make them more expensive which leads to a less attractive use of the car [97]. The report surveys multiple measures in this sense: Economic mechanisms (variable pricing, emission-based parking charges, extra taxes for parking space suppliers), regulatory mechanisms (parking supply caps, abolish regulations that forced house builders to construct parking spaces) and physical design mechanisms (e.g. bollards to ensure that people cannot park where they should not park).

One topic in this thesis concerns optimizing and making different use of parking facilities. It will be shown that there are smart solutions available that could be applied to the parking problem. They exploit the unbalanced demand in selected locations by supporting a collaborative use of car parks.

1.1.3. Environmental pollution

Pollution caused by road traffic is a well-discussed topic in the ITS domain, e.g. [50, 153, 80]. However, the problem of environmental pollution is still increasing: The European Environment Agency calculates that greenhouse gas emissions caused by traffic (including aviation and maritime transport) has increased by around 34 % between 1990 and 2008. Road traffic is a key issuer and accounts for about 17 % of the total EU's greenhouse gas emissions (CO_2) [4]. Clearly, the problem of greenhouse gas emission is a reasonable motivation to develop more efficient solutions.

Recent discussion include also other pollutants emitted by vehicles, such as sulphur dioxide, nitrogen dioxide and oxides of nitrogen and carbon monoxide, Benzene, particulate matter, (PM_{10} and $PM_{2.5}$), lead ozone and related NO and NO_2 [130]. In 2014, the WHO has released estimates on global death cause by air pollution (7 million), highlighting road traffic as a major emission source. In a survey from 2012, Yim and Barret estimate 5000 premature deaths caused by vehicle exhausts per year for the UK. In fact, that is much more than those, which were caused by road accidents (1850) [190].

Thus, the problem of air pollution could be divided into two rather distinct challenges: lowering the impact on (1) global warming (i.e. by CO_2) and (2) on human health (i.e. by CO).
The European Commission is well aware of the pollution problem and has set up policy countermeasures. Among other directives, there is the *European Union clean Air Directive* which aims at controlling PM_{10}. Furthermore, many cities have set up low emission zones, which determine inner city areas where not all vehicles have access to. Other policy instruments are the *License Plate Program*, total bans for vehicles or congestion pricing [186]. At the same time, modern vehicles emit much less than older vehicles due to increased pressure from legislative (e.g. emission-based taxes for vehicles in Germany since 2010) and general customer preferences for cleaner vehicles.

However, it is important to note that given approaches are open loop measures and the aggregate effect of vehicle emission is not tackled. The *aggregate effect* highlights the fact that the increasing number of vehicles outperforms the advances in fuel efficiency on a per-vehicle level.

1.1.4. Road safety

Safety continues to be the central objective of road transportation professionals. In 2004 almost 43,000 fatalities and 2.8 million injuries were recorded on U.S. highways. The deployment of new safety system such as *ABS* or air-bags helped to decline the number of fatalities over the past decade. Still, the main objective of automotive manufactures in the context of ITS is safety. The so-called *Vision Zero* shared by multiple nations and automotive manufactures clearly states that "Human life and health are paramount and take priority over mobility and other objectives of the road traffic system" [179].

While focussing on the related work in ITS, it should be noted that there are many (communication-based) applications developed in the community and that some have been already tested in field operational test (FOT) projects such as DRIVE C2X [158] or sim^{TD} [178]. Generally speaking, vehicle-to-x communication technology is used to enrich information to be fed to safety assistant systems.

1.1.5. Electric energy need

There are two complementary developments in the energy and the automotive sector which are often considered as part of the automotive research agenda: renewable energy sources and electric vehicles, e.g. [91, 109]. First, more and more electric energy is produced by renewable sources, such as solar or wind power. In Europe the share of renewable produced increased by 81.3 % between 2002 and 2012 [48]. Accordingly, an increasing part of the overall energy provisioning is subject to variable supply. This is a major challenge for the overall energy grid management as energy consumption does not necessarily follow the seasonal or daily temporal energy supply.

Therefore, a central pillar for future electricity supply management are storage capabilities, which could help to relax differences between electricity supply and demand. With the advent of fully electric vehicles and plug-in hybrid vehicles, batteries with large capacities are expected to be deployed on a large scale. The main idea is to integrate this capacity with the grid and use the batteries of vehicles to buffer energy where needed. There are various works which outline how the charging process could be managed efficiently in this regard. Available approaches are valid but are connected with high requirements to charging management.
In this thesis an alternative solution is proposed, which manages the electric power consumption while driving to balance the supply and demand of electric energy.

1.2. Two challenges: context-awareness and synchronization

Apart from traffic safety, all of the previously named challenges of road traffic in smart cities (congestion, parking, pollution, electric energy demand) could be - at least to some extent - explained by two more general issues in road traffic: synchronization and the lack of context awareness.

The approaches developed in the context of this thesis address these problems and exploit the most important development in the ITS community: the power of V2X communication. The advent of this technology enables completely new approaches to design road traffic applications for smart cities. Whereas system developers were limited by relevant system boundaries (e.g. of a singular vehicle, of a traffic light), system developers could now integrate various entities and develop applications that support true collaborative behaviour resulting in an overall benefit for all participants. This thesis considers two basic countermeasures for that: de-synchronization (or balancing) and cyber-physical systems.

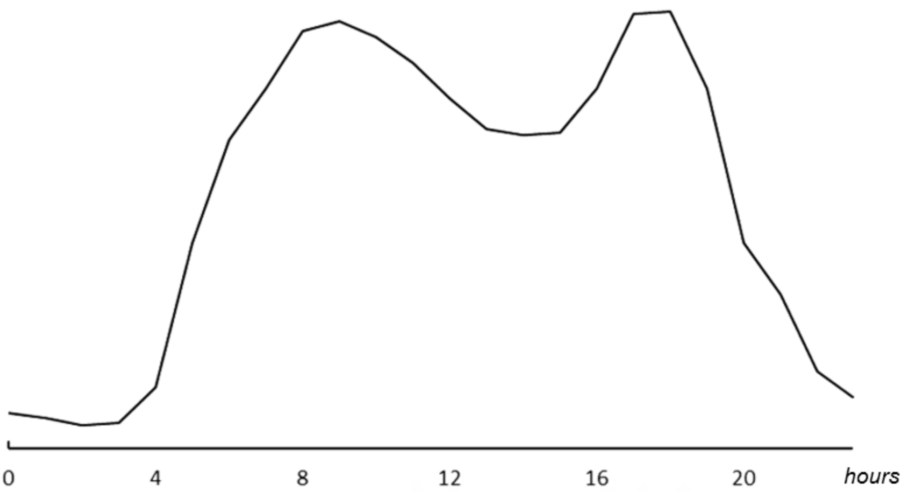

Figure 1.1.: Typical traffic density on a work day.

Synchronous consumption of resources is a challenge in many distributed systems. In road traffic, most commuters travel to work in the morning and back home in the afternoon. Figure 1.1 shows typical time series of traffic density on main roads in urban areas leading to traffic congestion during peak hours. An equivalent effect could be monitored in the parking domain. Parking space supply and demand has also a synchronous character which leads to inefficiencies: During work hours, parking spaces become short in business districts while there is an oversupply during weekends and nights. The opposite could be monitored in residential areas. The problem may also arise in the context of electric vehicle charging, where it would be a huge problem if available charging stations would not be used efficiently - especially in a roll out phase where the setup of the infrastructure is most critical. In the context of this thesis, balancing is a main concept to overcome problems caused by synchronized resource consumption. Furthermore, approaches are developed and discussed that aim at establishing a balance of resource supply and demand.

Today, context awareness in road traffic commonly addresses the design of infotainment and driver assistant systems to minimize driver distraction and information overload [134, 13] or to increase the efficacy of vehicular communication [3, 40]. However, context awareness has not yet been developed to a point that addresses core functionalities of smart city applications. To do this, one must introduce a notion of context awareness that goes beyond the limits of single vehicles. Specifically in this thesis, this is the integration with air pollution control and the electric grid. The term *Cyber-physical system* is often used to describe systems, that integrate computational power and physical entities at large scale, using open interfaces and networks which are suited to control particular features [184, 92]. This is what is envisioned here as well: This thesis presents approaches where a non-limited number of vehicles are connected to one integrated system. It is networked by open communication technology to control air pollution or features of the electric grid.

1.3. Recent enablers for novel smart city mobility applications

The approaches developed and presented within this thesis are possible to implement because of recent developments in ITS, the automotive domain as well as general social trends. This section gives a brief overview on the most relevant ones in the context of this work, namely vehicular communication, (plug-in) hybrid and fully electric vehicles and the increased willingness of governmental institutions to implement policies to support sustainable mobility. All serve as key enablers for proposed solutions.

1.3.1. Communication technology enablers

Communication technology has developed dramatically in recent years. That is also true for vehicular communication. More and more vehicles are equipped with communication technology already today and the networked or connected car is a broadly used term in the automotive domain. V2X communication technology - the communication between vehicles and infrastructure (V2I) and direct communication between vehicles (V2V) - are expected to be a major contributor to this concept [128]. V2X technology and exemplary applications have been proofed within large scale field operational tests (FOTs) - i.e. projects sim$^{\mathrm{TD}}$ or DRIVE C2X - to be beneficial in terms of road safety and traffic efficiency measures. The technology is expected to be deployed in very near future. General Motors CEO Mary Barra announced in 2014, that the Cadillac model CTS will be equipped with V2V communication technology in 2017 [63]. In parallel to such vehicular ad-hoc networks (VANETs), some vehicles and in-car equipment (such as navigation devices) are already equipped with communication technology and use common mobile (cellular) networks. Limited and defined information (e.g. defined in *TMC* or *TPEG*) is broadcasted from the infrastructure to vehicles via broadcasting services already today.

Generally speaking, communication has become an integral part of the mobility domain and is expected to play a dominant role. In the context of this thesis, it is the main enabler to implement new cyber-physical systems and collaborative applications, which integrate a large number of participants of an overall mobility system.

1.3.2. New kind of vehicles: EVs and PHEVs

Besides communication technology, there are also many developments in the field of alternative drive systems, vehicles powered by fuel cells, batteries, combinations of internal combustion engines, electric motors etc. Electric vehicles have a very long history. The first electric car was build in the 19th century - almost at the same time as the first automobile with an internal combustion engine. However, mainly because of range advantages, vehicles powered by internal combustion engines became much more successful and a de-facto standard.

In the context of increased environmental awareness, general political and economic objectives as well as relevant technical progress, electric vehicles experience a revival lately. Until 2020, the German government aims at having one million electric vehicles deployed [131]. More successful than electric vehicles (in terms of registration numbers) are hybrid vehicles. They combine the internal combustion engine with an electric motor powered by a rechargeable battery. Hybrid vehicles are divided into range extenders, plug-in hybrid vehicles, mild and full hybrid vehicles etc. In the context of this thesis, hybrid vehicles refer to plug-in full hybrid vehicles. Such vehicles could be recharged at charging stations and could run both in electric or internal combustion engine mode exclusively. The main advantage of hybrid vehicles is that they counteract the range problem of electric vehicles as they could be operated in the internal combustion engine mode in case the battery is empty.

The Toyota Prius is probably the best-known example for a full hybrid vehicle on the market. Its first version has entered the mass market as a (non-plug-in) HEV in 1997. Today almost all big vehicle manufacturers have hybrid vehicles in their fleet. It is assumed, that they will play a more and more important role in traffic during the next years.

In this thesis, it is shown that hybrid vehicles enable novel approaches for the smart city, when the mode choice is aligned among a greater number of vehicles. They provide a possibility to change their behaviour (here to change from electric to ICE mode or vice versa) without being invasive compared to speed adaptations, route changes etc. The two applications introduced in this chapter show how this capability could be exploited in a cyber-physical system.

1.3.3. Policies and regulations

Road traffic problems, such as outlined in Section 1.1 have attracted great attention among citizens and governments all over the world. Accordingly, governments or other public authorities have issued various directives and statutes to regulate them.

For example, the EC has issued the Clean Air Directive, which aims at controlling air pollution. Similar policies exist in the US with the Clean Air Act. According to

Donaghy et al., European citizens are more ready to accept command and control, and price measures (road pricing, fuel taxes etc.) to fight transport related challenges compared to the US [39]. Rather restrictive policies exist in Asia. In Shanghai for instance, the number of car registrations is limited per month. In Beijing, a licence plate recognition system has been set up: it prohibits to drive the car on selected weekdays based on the last digit of the license plate [188]. In Singapore vehicles are charged electronically for entry and use of roads within congested zones [39]. Low emission zones are implemented in various cities, not only in Europe but also Japan. In South America (Chile and Argentina), vehicle registrations are limited to vehicles fulfilling EURO emission norms.

Generally speaking, legislation and governments all over the world have started to implement policy measures to address urban road traffic problems. In regard to the feasibility to deploy proposed solution within this thesis, it could be stated, that the general support from authorities to fight relevant problems is available.

1.4. Contribution of this thesis

There are many problems related to road traffic for smart cities as introduced in Section 1.1. At the same time, recent developments in ITS (see Section 1.3) enable completely new approaches to address these challenges. Five of such applications are introduced in this thesis. All applications implement simple feedback loops around networked vehicles and share features of cyber-physical systems and balancing objectives.

The first application has the objective to balance charging stations' supply and demand for electric vehicles in a smart city such that queues are minimized. Queuing in front of charging stations is a problem which is not discussed in the community so far, even though this might be a "show-stopper" for a large-scale deployment of electric vehicles. Fast charging stations are helpful in general but are not solving the general problem. The proposed application addresses the queuing problem. The second application is closely related to that and targets parking by using available spots more efficiently. A related approach is targeting locally balanced vehicle routing. Most people consider routing on a global level (routes from an initial origin to the final destination). For system-centric routing, this approach often fails as it has undesired effects: imagine a dirty vehicle and zero-emission vehicle wish to be routed. Generally, one wishes to favour the zero-emission drivers and provide her or him optimal routes. Contrary to that, one wishes to minimize the emission of vehicles from a system-centric point of view. Thus, the shortest (best) route would be allocated to the “dirty vehicle” driver.

The fourth and fifth application exploit the power of full (plug-in) hybrid vehicles. By letting vehicles switch from one mode to the other in a smart way, it is shown that air pollution in cities could be controlled and minimized as well as that energy under- or oversupply in the electric could be regulated.

The approaches have in common that they respect the requirements and particular demands of vehicle drivers. That means that intrusiveness is kept at a very low level and that drivers maintain full control and freedom of choice. Technical constraints and privacy preserving features are considered as well. In effect, many applications have a distributed character. Inspired by results from network research, it is shown how distributed algorithms such as AIMD could be applied to the ITS context. It is shown how these approaches perform under the particular conditions and constraints in ITS. Comprehensive simulations and case studies validate the developments. Some (limited) analytical proofs are provided or referenced.

1.5. Structure of the document

The thesis is structured as follows: Chapter 1 describes the main motivation and provides the relevant context of this work. Chapter 2 goes into details on relevant constraints from the ITS domain that should be considered when smart city mobility applications are developed.

Part I (consisting of Chapters 3, 4 and 5) introduces three applications building up a cyber-physical system exploiting the power of full (plug-in) hybrid vehicles. More specifically, approaches are introduced to couple a large number of HEVs to their physical environment, thus creating a large scale cyber-physical system. It is outlined how networked hybrid vehicles could be used to control other aspects of a smart city such as air pollution (Chapter 3) and to balance the supply and demand of electric power (Chapter 4). The latter problem often occurs in electric grids where many renewable energy sources are integrated. Chapter 5 discusses an extension of the application to control air pollution with hybrid vehicles outlined in the previous Chapter 3. It implements a distributed optimization approach for hybrid vehicles which belong to a managed fleet. Here, vehicles switch from one mode to the other (internal combustion engine mode and electric mode) to maximize the utility of the fleet while respecting overall (fleets') constraints.

Part II (consisting of Chapters 6, 7 and 8) presents three applications implementing balancing approaches using distributed algorithms. They address challenges where the availability of resources and relevant demand diverge locally. The first application described in Chapter 6 exploits the instrument of routing to balance pollution and traffic flow. The second application addresses the problem of allocating electric vehicles to a limited number of charging stations (Chapter 7) in a way that all charging stations serve equal numbers of vehicles. Finally, in Chapter 8 it is shown how the utilisation of available vehicle parking spots could be increased by collaboration among parking spot providers (such as private landlords or companies) and parking spot users.

Part III (consisting of Chapters 9 and 10) is the final part of the thesis and discusses application evaluation aspects and provides concluding remarks for aforementioned applications. Chapter 9 introduces a novel simulation platform with hardware-in-the-loop features, which is used to assess networked, collaborative applications - both regarding technical features and user acceptance. With help of this platform, it is possible to analyse applications on a large scale integrating thousands of vehicles but only few (or even just one) human drivers. Furthermore, results from initial user acceptance studies are presented. The thesis closes with an outlook and conclusion in Chapter 10.

1.6. Joint work

It is important to note that several works outlined in this thesis are result of joint works with multiple collaborators. These are Robert Shorten, Arieh Schlote, Emanuele Chrisostomi, Mahsa Faizrahnemoon, Ilja Radusch, Jia Yuan Yu, Wynita Griggs, Rodrigo Ordóñez-Hurtado, Kay Massow, Stefan Lobach, Thomas Hecker, Astrid Bergmann, Kevin Rhinow and Bartek Rytt. In the opening introductions of each chapter, relevant joint works are highlighted.

2. ITS applications and requirements

Summary: *This chapter sets the scene by giving an overview of selected state-of-the-art applications in ITS research. It is shown that the applications proposed in this thesis have a different character than most available ITS applications as they implement real collaboration among vehicles to address broader, more general objectives. Furthermore, domain-specific design constraints for applications are outlined. All applications proposed in the subsequent chapters consider these requirements.*
Some aspects of the design constraints or requirements have been published in [75] - a joint work with Mahsa Faizrahnemoon, Emanuele Crisostomi, Arieh Schlote, Ilja Radusch and Robert Shorten - or were briefly named in a non-published collaborative work with Mahsa Faizrahnemoon as part of the research project TEAM.

2.1. Cooperative applications in ITS

Research in ITS has many facets. Surveys of the topic are mainly targeting communication protocols to enable V2V or V2I communication [30, 101, 104], including a broad discussion of security and privacy aspects [44, 16, 64] and ITS applications itself [88, 144, 71, 182]. There are many ways to group ITS applications. Many works allocate them to one of three groups: safety, convenience and efficiency [152], other distinguish more groups such as the ISO, which identifies eleven *service domains* (ISO 14813-1, also outlined in [182]).

Vehicular communication technology commonly refers to ad-hoc networks using industry standards given by IEEE and ETSI. Both define similar communication stacks (DSRC and ITS G5 respectively) and share the same physical and access layer (IEEE

802.11p). While such vehicular ad-hoc networks are not yet deployed on the mass market, cellular communication is used in the automotive domain already today.

The great majority of applications in ITS developed so far (mainly in the context of research projects) are active safety applications. Their objective is to avoid traffic accidents [88] (while passive safety approaches aim at minimizing the impact of accidents). The following list is just a small selection of various approaches and is inspired by the list of "high potential" safety-related applications defined by the Vehicle Safety Communications consortium (VSC) [28]:

- **Curve speed warning**: Contrary to the general speed limit, the recommended speed to pass curves may be much less and may depend on dynamic (e.g. weather) conditions. In regard to a networked vehicle setup, the required information could be provided by other vehicles or the infrastructure.

- **Forward collision warning**: CAM messages from the vehicle in front are used along with knowledge about the relevant ego-vehicle's status. The likelihood that they are on a course to collide is computed. Warnings are accordingly generated to the driver [41].

- **Emergency electronic brake light**: Similar to the forward collision warning, the application aims at warning the driver based on CAM messages received from vehicles in front. The idea here is to utilize the information to adapt a breaking event even before breaking lights of the vehicle in the line of sight are signalling [41].

- **Approaching Emergency Vehicle Warning**: While not being part of the list of "high potential" applications of the VSC, the application has been developed by multiple "networked vehicle"-projects such as sim^{TD} and DRIVE C2X. Emergency vehicles broadcast V2V awareness messages to other connected vehicles in the vicinity. The receiving vehicles warn their drivers.

- **Intersection collision warning**: Location and speed information among vehicles who approach a common intersection is exchanged via V2V or V2I to monitor other connected vehicles and decrease the risk of collisions. Drivers are warned and may receive driving instructions to pass the intersection in a safe way. The basic idea has been implemented in many projects, e.g. [32, 115, 172]. Related (sub-) applications are *stop sign violation warning*, *signal violation warning*, *left turn assistant* and *Stop sign movement assistant* [6]. These (sub-) applications belong to the applications named by the VSC.

- **Lane change warning**: Based on received CAM messages from nearby vehicles, the driver is alerted when there is a vehicle in the blind spot while the driver is about to change the lane. The application itself has been proposed and described in various contexts, among others in [5].

- **Hazardous location notification**: Connected vehicles or road side ITS stations signal (via DENM messages) to other vehicles about hazardous locations, such as obstacles on the road, construction works or any other hazardous conditions. Warning in the ego-vehicle are generated based on this information.

As stated before, ITS research is commonly focussing on demands of active safety and related requirements. Examples for such requirements are that V2V communication latencies should be below 100 milliseconds and lane-level position accuracy should be reached etc. A good overview is given in [71, 88].

The proposed applications in this thesis have a different character than the ones briefly outlined before. That is the reason why the majority of available requirements in given overviews are not helpful in the design phase of core algorithmic approaches within this thesis. Requirements, which are more relevant in this work's context are listed in Section 2.2.

The most interesting application group of the aforementioned three groups is the one that targets traffic management, control and efficiency. In most contexts, the group includes applications where cooperation among vehicles is not necessarily foreseen (e.g. *regulatory and contextual speed limit notification* [88] or *green light optimal speed advisory* [67]). The following list is limited to those applications where cooperation and aligned strategies of vehicles is exploited to reduce emission, improve traffic throughput etc. With this assumption in mind, the list becomes rather short. However, it must be stated, that there are recent developments, which aim at implementing collaborative strategies including feedback loops in ITS such as [142, 165, 166], as well as the work done within the *Green Transport and Communications Networks* project funded by the Science Foundation Ireland (SFI) and the European research project *TEAM*[1]:

- **Cooperative or collaborative adaptive cruise control**: Using V2V communication, vehicles exchange information about locations, speeds and headings with vehicles ahead and behind. Based on this information the application adjusts (or recommends) the speed. In addition, I2V communication is used to inform about speed limits such that these are respected. One of the objectives is to improve the overall traffic flow and to control overall traffic emission. The application is developed within the TEAM research project. Ordóñez-Hurtado et al. described an interesting version of the basic idea in [126].
- **Vehicle-highway automation system (platooning)**: The approach is an extension of the cooperative or collaborative adaptive cruise control. Vehicles exchange information with other vehicles in the vicinity to form a dense, high-speed platoon. The objective of the application is to lower air resistance of vehicles, safe energy and lower emission. The application has a long history

[1] See `www.collaborative-team.eu`

in ITS research and has been implemented by various research groups. [15] provides a dedicated survey on platooning projects in ITS.

- **Collaborative driving and merging**: The applications aims at improving traffic safety and increase traffic flow by helping two or more vehicles, which interact in a lane change or lane merging scenario. The application provides aligned driving recommendations to involved drivers [79].

- **Collaborative eco-friendly navigation**: The application includes a routing service and a route guidance engine. A centrally deployed routing engine calculates routes for all vehicles, which request a route. The routes are aligned in a way that optimizes the overall traffic throughput. The basic idea of the algorithm is introduced in [84]. It is compared to the routing approach in this thesis in Section 6.4. Implementations are done within the context of the TEAM project [79] and the Green Transport and Communications Networks project.

- **Collaborative smart intersection for intelligent priorities**: With help of V2I communication, the application aims at improving (public) transport service, by giving priority to selected vehicles (here public transport vehicles). The basic idea has been implemented already using other (less flexible) technologies such as optical sensors and induction loops. In this application, vehicles send their intended destination to an intersection. The traffic flow is regulated by the traffic light based on the number of vehicles that will follow in each direction [79].

- **Parking allocation**: Schlote et al. introduce a stochastic allocation scheme for connected vehicles to parking garages in [142]. The application provides recommendation for parking garages aiming at filling the different garages with vehicles in a balanced way.

2.2. Specific design constraints in ITS

As stated earlier, the main focus in ITS research has been put on the development of active safety applications (and the enabling communication technology). Thus, requirements collected so far (such as in [88, 71]) have an according focus. These kind of requirements are not repeated here.

Solutions for network resource allocation problems in the context of the internet are at a mature stage and do typically involve large scale decentralised control and optimization problems. This latter feature makes algorithms from this community attractive in a transportation context. Relevant ideas in this direction include the Kelly framework [90], RED [156, 27] and AIMD congestion control [149].

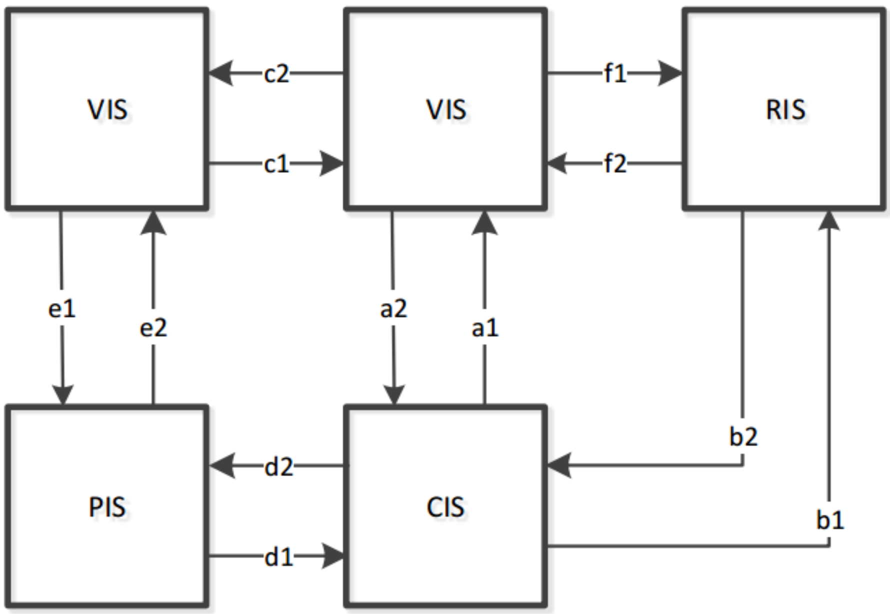

Figure 2.1.: Communication links between ITS stations.

However, when applying control and optimization algorithms in the context of ITS applications, domain-specific considerations should be considered - and respected. In this section, an overview on these particular constraints of ITS are presented. They have a great impact on the actual design and deployment of the central algorithmic mechanisms.

2.2.1. Processing capabilities

The basic architecture of ITS is shown in Figure 2.1. It includes the Vehicle ITS Station (VIS), the Roadside ITS Station (RIS), the Central ITS Station (CIS) and the Personal ITS Station (PIS). A general overview on the ITS (communication) architecture is available in the ETSI standardization documents[2] and many other publications. Regarding the processing capabilities, it is assumed, that the computing power of VIS, RIS and PIS is limited, while the CIS may have access to powerful machines and are capable to perform large and complex tasks. This needs to be kept in mind, when deciding where computing intensive jobs should be performed.

2.2.2. Communication aspects

As shown in Figure 2.1, various communication links in ITS exist. They are named from $a1$ to $e2$. In the following the most important ones are discussed.

[2]See `http://www.etsi.org/technologies-clusters/technologies/intelligent-transport`

There are various means of communication technologies to send information from the central authority (here CIS) to the VIS ($a1$). The most efficient way is probably via broadcast communication. In fact, infrastructure-to-vehicle (I2V) broadcasts are the most common and well-established communication link so far. The required technologies are already available in most vehicles, such as the traffic message channel (TMC), which is broadcasted using FM radio. A central feature of broadcasting is that the total communication effort is scalable as it does not depend on the number of vehicles in the fleet.

Obviously, another possibility to establish this link is by I2V unicast communication. That requires direct communication from the central infrastructure to individual users using cellular networks. Here, the communication effort depends on the number of vehicles directly. Sending information from the CIS to the VIS could also be done by traversing the RIS ($b1$ and $f2$). The communication link between CIS and RIS ($b1$) is commonly implemented using cellular communication or by wired broadband links. The link between RIS and VIS ($f2$) is established by an ITS G5 vehicular ad-hoc network.

The communication into the opposite direction (Vehicle-to-Infrastructure, V2I) is implemented by using cellular networks ($a2$) or via a roadside stations ($f1$ and $b2$). Obviously, the communication effort increases linearly with the number of vehicles. In some cases, the communication is implicitly available only. That is, when vehicles affect features of the environment, which could be monitored by the CIS. An exemplary example for that is when infrastructure sensors such as induction loops, air pollution sensors are used, which provide information about vehicles' states or effects.

The information exchange between VIS ($c1$, $c2$) and VIS or VIS and RIS ($f1$, $f2$) is again performed via ITS G5. It is assumed, that the communication might involve multiple via-nodes. That means that a message from one VIS to the other, or from one RIS to another one VIS (and the other way around) might involve multiple nodes in between, which store and forward the message (multi-hop communication). Moreover, it is assumed, that V2V communication is stateless.

In general, the exchange of information between entities could be divided into unidirectional and bidirectional communication: bidirectional V2V (or V2I) refers to a situation where vehicle i communicates with vehicle j if and only if j can communicate with i. It must be stated that in such vehicle ad-hoc networks connectivity is very dynamic as nodes are constantly moving, such that links are established and broken very often. Also, VISs and CISs have to share limited capacity of the local communication medium. Therefore, communication in VANETs is often unidirectional and the system developed should be able to handle this constraint.

2.2.3. Scalability

According to Neumann, a system is said to be scalable if it can handle additional users or entities without reasonable loss of performance or increase of administrative complexity [119]. Scalability is a crucial aspect in ITS design as one system shall be able to run with few and many participating vehicles. Furthermore, it must be assumed that the number of participating vehicles is not know. Such consideration have great impact on the general system architecture and on central control and optimization algorithms.

For example, while it might be possible to run complex calculations on powerful central machines for a few vehicles, it may be impossible if there are hundreds of thousands of vehicles. At the same time, the communication channel may turn out to be a bottleneck for large numbers of vehicles, if - for example - vehicles require a direct exchange of information with all other vehicles.

In the context of core application algorithms, two aspects are most crucial in regard to scalability. First, the number of vehicles is unknown. Two, the number of vehicles is varying. In this sense scalable algorithms are those, which work satisfactorily over large uncertainty ranges with fixed gains and parameters.

2.2.4. Non-homogeneous levels of vehicle participation

ITS is characterized by varying levels of vehicle participation as vehicles leave or enter the network at any time. Control and optimization approaches must be able to cope with that. The varying number of vehicles have a deep impact on the system behaviour. In consequence, common feedback control approaches are not applicable.

2.2.5. User acceptance

Not all valid control and optimization approaches are useful in ITS as they must be accepted by the participants.

In this context, only most critical aspects are named: First, users could not be forced to act in a defined way. I.e. when an optimization method defines a behaviour which is not accepted by the end user (like going a route which is much longer than the shortest one), the system has to cope with the case, that the user acts differently than desired and does not comply to the relevant instruction or recommendation. There a several studies targeting compliance and acceptance of advanced driver assistant systems (ADAS) especially in the context of routing, such as [21, 19]. Users may act non-conformant on purpose or by accident. In the context of the applications introduced in this work, applications should give the freedom to the driver to overrule the recommended actions at any time.

Second, users are expected to be very sensitive regarding privacy issues, here by sharing sensitive information with others - especially with a central authority.

2.2.6. Robustness and uncertainty

Robustness is the ability of the system to cope with malfunctioning parts of the system. Consequently, a basic requirement to all systems and algorithms is robustness. A failure of the system in the context of ITS could be a malfunction, a broken or intentionally misbehaving ITS station or a communication link.

There are various works that look into robustness features and there are multiple approaches in place: First, there are (coding) techniques, which help to detect and correct errors happened (to a received message). Second, the critical parts of a system (e.g. a central controller) might be replaced by another entity in case it breaks - the concept is called redundancy. Generally, a distributed character of the system architecture, where there is no central, critical system part like described in [70] increases the robustness of the system.

Also, the allocation (or optimization) system needs to cope with uncertainties, which are common for real-world applications in the context of ITS. Examples for works that handle uncertain parameters are [120, 102], where the road capacity is said to be unknown [120], where the demand for road network resources is assumed to be unknown and [106], where the utility functions of travellers is unknown and heterogeneous. A general overview of this topic in optimization is given in [14].

2.2.7. Delays in the feedback loop

Delays in a feedback loop cause major issues when designing a system controller. In ITS, there are communication delays and processing delays. Generally, communication delays in the context of road traffic applications could be neglected (in the context of the prosed applications) when modern communication technologies and direct communication links are assumed. At the same time, data from real world sensors (like loop detectors) could often not be accessed in a high frequency. Such real-world limitations might be explained by the fact that such systems were not designed for real-time control applications.

But there are also limitations defined on an application level. One example for that is routing: Assume a central routing system, which recommends routes where only links are passed with low pollution values. The delay problem here is, that vehicles might reach the relevant links on its route when conditions have changed.

2.2.8. Time discretisation and synchronization

All applications in ITS face a discrete time setting. This is an important statement, since most dynamic systems are described in continuous time using differential equations. In case solutions from the continuous time domain are transferred into a discrete time setting, Souza et al. outline in [155] the problem of classical discretisation methods: they often destroy the sparsity pattern of the original system.

When applying decentralized algorithms, time synchronization is a general issue [168]. Nevertheless, it could be assumed, that there are various solutions for the problem - especially for ITS related problems. Commonly, it is implemented by using the time provided by GPS, where time differences of only a small number of nanoseconds must be assumed [7]. Another way for time synchronization is done by communicating to a time server on the internet, where there is an accuracy in an order of milliseconds [114] implemented using the *Network Time Protocol* (NTP). For most of the applications in ITS (and all addressed in this thesis) these time differences cause no problem. Still there are also methods available for P2P-based synchronization. A comprehensive survey is provided in [167]. Given the numerous possibilities to synchronize the clock over all vehicles, it is assumed that the time synchronization is no constraint in the context of ITS.

2.3. Review of potential system architectures

Based on the ITS specific constraints outlined in the previous Section 2.2, four basic architectures are discussed: a centralized, a fully decentralized using V2V communication only, a partially decentralized using V2V and V2I communication and the partially decentralized design without V2V communication. Figure 2.2 illustrates these basic deployment groups.

In this section, all four configurations are discussed by referring to the aforementioned design constraints and requirements. Additionally, some selected optimization mechanisms are named which fit to the relevant setting.

2.3.1. Centralized architecture

The centralized setup is shown in Figure 2.2 (a). It is assumed, that a central authority communicates directly (unicast) to vehicles and that these give individual feedback. In combination with bidirectional communication with vehicles, this setup implies that the central infrastructure may do all the computation, may have perfect knowledge about vehicles' states, capabilities and could send individual controls to connected vehicles. While this comes with severe privacy concerns, the setup allows to implement a great number of control and optimization approaches (i.e. sub-gradient, Newton, primal-dual-interior point and other methods). Generally, all optimization

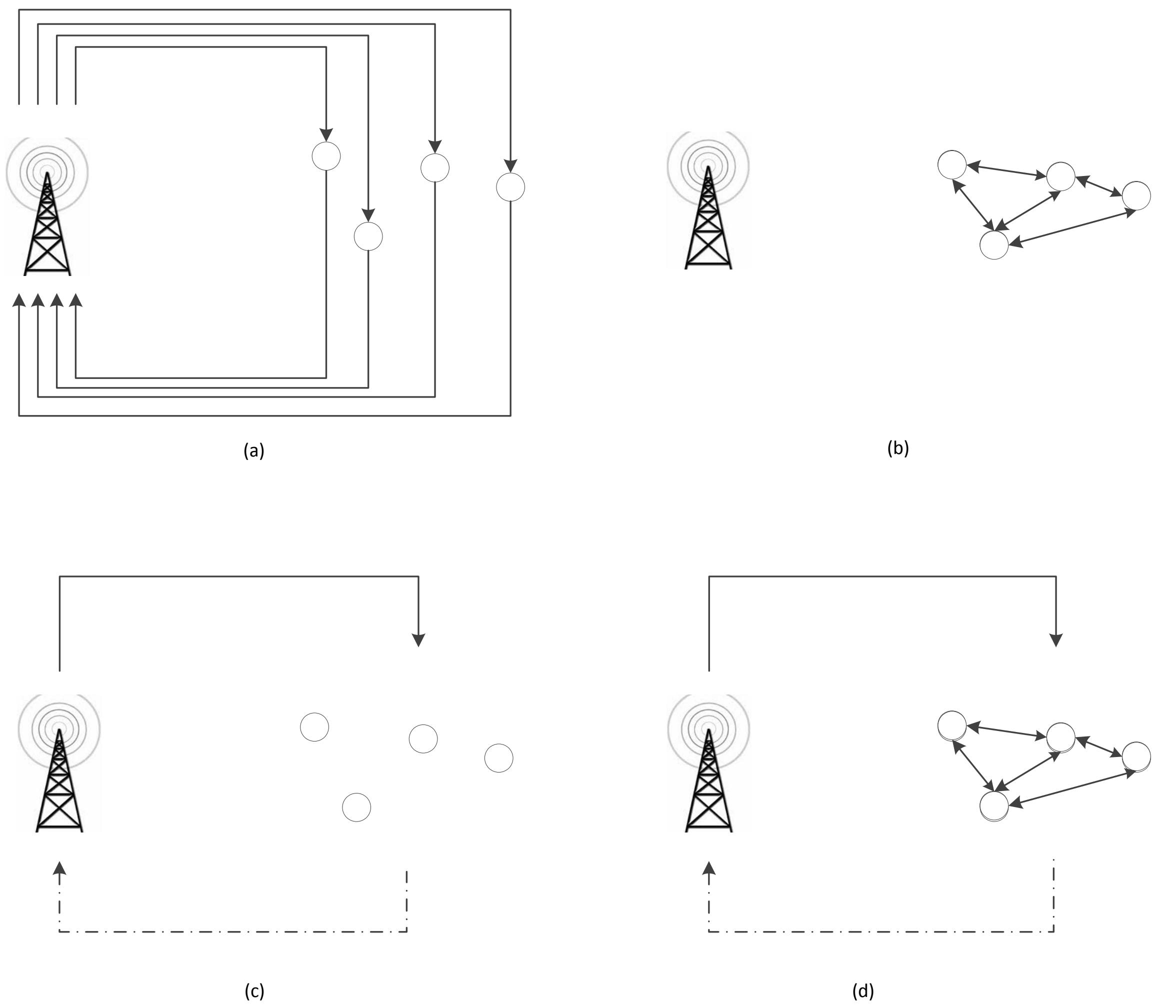

Figure 2.2.: Basic system architectures: (a) depicts the centralized setup, (b) the fully decentralized without any V2V, (c) the partially decentralized without V2V and (d) the partially decentralized with V2V. The transmission mast icon symbolizes an infrastructure entity, such as a traffic management centre or some local authority. The small circles represent vehicles equipped with a VIS. Arrows denote communication flows. Dashed lines stand for implicit, continuous lines for explicit communication flows.

methods could be implemented in such a scenario. Furthermore, the speed of convergence is not affected by the architectural constraints. At the same time, the communication effort is very high as a lot of information needs to be exchanged. It increases linearly with the number of vehicles. Another drawback is, that in case the central authority is not working properly (not considering a redundant setup) the approach fails completely.

2.3.2. Fully decentralized architecture with V2V communication

The fully decentralized or fully distributed architecture defines the complementary setup to the centralized one. The basic idea of the architecture is shown in Figure 2.2 (b). Communication is limited to links between neighbouring nodes. In the context of V2V communication by ITS G5, the neighbourhood is defined by the communication range, which could be assumed to be between 100 and 300 m [72].

V2V communication is not necessarily symmetric (see [160]). A central problem of such a setup is that vehicles have no knowledge about the global state [168] of the system, which should be controlled (e.g. a defined air pollution limit). It should be kept in mind that this is not necessarily a problem in all scenarios or applications, i.e. when resources are strictly limited like the number of parking spaces in the city. In such a scenario, the constraint could not be violated.

A suitable method for solving optimization problems in a fully decentralized way is consensus. Consensus problems have a long history in distributed computing [124] and have broad applications in networked systems, such as flocking, swarming, formation control, multi-agent coordination [56, 17, 82], distributed computing [160] and sensor networks [125].

An example of a consensus idea is given in the following. Assume one wishes to optimize the sum of all agents' (like vehicles) utilities $\sum_i u_i$, where u_i depends on the resource x_i allocated to the particular agent. With the constraint $x_i \leq C$, one may formulate a classic optimization problem:

$$\begin{aligned} &\text{maximize} \quad && \sum_i u_i(x_i) \\ &\text{subject to} \quad && x_i \leq C. \end{aligned} \tag{2.1}$$

Assume that the agent's utility functions are strictly convex. Roughly speaking, the approach follows the mathematical direction outlined below. The Lagrangian of the aforementioned optimization problem is given by:

$$L(x_1, x_2, ..., x_n, \lambda) = \sum_i u_i(x_i) - \lambda(\sum_i x_i - C) \tag{2.2}$$

where λ is the Lagrangian multiplier. The Lagrange dual function $g(\lambda)$ is defined as the minimum value of the Lagrangian:

$$g(\lambda) = \min_{u_i} L(x_1, x_2, ..., x_n, \lambda) = \min_{u_i} \sum_i f_i(u_i) - \lambda(\sum_i u_i - C). \tag{2.3}$$

In order to minimize the Lagrangian, the derivative is taken and it is

$$f_i'(u_i) - \lambda = 0 \qquad \forall\, i. \tag{2.4}$$

The solution is given by finding the allocation such that $\lambda^* = \frac{\partial f_i(u_i^*)}{\partial u_i} \quad \forall\, i$, subject to the linear constraint being satisfied. This problem could be formulated as a consensus problem, which leads a distributed solution, see [160] (assuming strict convexity of the utility functions). The consensus problem can be formulated as:

$$u_i(k+1) = u_i(k) + \epsilon \sum_{(i,j)N} (f_j'(u_j) - f_i'(u_i)). \tag{2.5}$$

In this approach each node communicates to its neighbours until they reach consensus. The limitation of this method is that the communication between the users should be symmetric to hold the constraint.

Besides, fully decentralized systems have various (general) features: The processing load could be partitioned among participating vehicles and communication is limited by the number of vehicles in communication range. Thus, it is highly scalable. However, processing capabilities at vehicles are limited and V2V communication is much more complex than V2I (or I2V, respectively). Furthermore, a fully decentralized setup is flexible to the number of participating vehicles and is robust to failures of individual nodes. Privacy is less critical as in a set-up with a central node, since the vehicles in communication range are changing over time. Convergence is generally rather slow.

2.3.3. Partially decentralized architecture without V2V communication

The basic idea here is to have a feedback control loop among a central entity and the distributed vehicles. In contrast to the centralized approach, it is assumed, that vehicles do the main work and are only supported by the central instance. Figure 2.2 (c) shows the basic set-up of this approach. The core advantage compared to the fully decentralized approach is that it is easy to share the overall system state with all participating vehicles with help of central entities' broadcasts.

Compared to the centralized approach, a positive aspect in regard to user acceptance is that approaches could be implemented, where vehicles do not need to share sensitive information - neither with other nodes in the neighbourhood, nor with the central authority. The basic architecture allows to implement a wide variety of control and optimization approaches, while not all methods could be applied (like in the centralized case). Among others, possible approaches are ADMM and the sub-gradient method.

2.3.4. Partially decentralized architecture with V2V communication

The partially decentralized setup with V2V communication is shown in Figure 2.2 (d). While having a rather complex communication architecture, the setup is very flexible: processing and communication could be shifted to where it is least costly.

The basic idea compared to the partially decentralized setup without V2V communication is that the link between the central authority and vehicles could be designed in a more loosely way (e.g. broadcasts of the central entity happen less often or with fewer details). At the same time, the approach aims at overcoming two central limitations of the fully distributed setup: First, it is simple to let the central instance share information about the overall system state. In effect the symmetric communication requirement for V2V communication could be relaxed. Second - with help of the central entity, which exchanges information with all vehicles - the speed of convergence could be increased (compared to the fully decentralized approach).

An example for such an implementation is described in the following. Let the basic problem be to allocate available resources, such that the aggregate utility of nodes is maximized. The problem can be formulated as a consensus problem leading to a decentralized solution. As said earlier, in the fully decentralized setup using the consensus approach, there is the problem to hold the constraint of the optimization problem as individual nodes could not know about the state of the overall system. The integration of the central authority overcomes this problem as it communicates with all nodes while it could be assumed, that it has information about the global state of the system (i.e. air pollution levels, traffic throughput etc.).

One can now formulate the problem as a consensus problem in which an integral control from the infrastructure is used to modify the approach (see also [95]):

$$x_i(t+1) = x_i(t) + \frac{1}{|S|} \sum_{j \in S \wedge j \neq i} (u_i'(x_i) - u_j'(x_j)) + \gamma(C - \sum_{j \in S} x_j(t)). \tag{2.6}$$

with S denoting the neighbourhood of the node i, including node i. The Equation (2.6) defines the update function for the resource allocation. The second term of the above formula guarantees that a consensus on the derivatives of the utility function is reached. This part of the algorithm could be implemented in a fully decentralised way. Vehicles would exchange the gradient of their utility functions u' with other vehicles - here the neighbours of vehicle i - and build the average. The third term is added to ensure that the the constraint $\sum x_i = C$ is satisfied. Here, the limiting value C - in case it is not a globally fixed value - needs to be made known to vehicles at least once in the process. In many cases C is fixed, such that there is no need to broadcast it periodically. In this case however, $\sum_i u_i$ is varying. Thus, it needs to be made shared with nodes on a regular basis.

An extension of this approach is introduced in [75], where it is shown that the speed of convergence could be improved by letting users communicate the maximum and minimum values of the derivatives of the utility functions to the central infrastructure at each time step. In return, these values are then broadcasted to all users who then augment their neighbourhood information to incorporate the new values.

Part I.

Cyber-physical mobility with PHEVs

3. Controlling air quality with PHEVs

Summary: *This chapter introduces the TwinLin application, which is an example for a cyber-physical ITS. It is shown how to integrate and how to exploit a large number of vehicles to control air pollution. TwinLin utilises the advent of vehicle communication technology and (full) hybrid vehicles to develop effective techniques to regulate local air emission generated by road vehicles. The basic idea is to place a feedback loop around all vehicles in a specific region, and use this loop to control aggregate emissions and - in effect - local air pollution.*
Important parts of the work have been published in [141], which is result of joint efforts with Robert Shorten, Arieh Schlote, Emanuele Chrisostomi, Astrid Bergmann, Thomas Hecker and Ilja Radusch. Furthermore, the approach has been presented and demonstrated at the Innovation Ireland / Technische Universität Berlin Science Colloquium in 2012.

3.1. Air pollution and the limits of state of the art countermeasures

The *TwinLin* application shows how hybrid, networked vehicles form a cyber-physical system (CPS). It addresses the problem of air pollution in urban areas. Before introducing the approach in Section 3.2, this section emphasizes the air pollution problem and relevant state-of-the-art countermeasures.

3.1.1. Air pollution in urban areas

Mobility is a central pillar of today's economics [53, 127]. According to *Eurostat*, 4.5 % of Europeans employees work in the sector. Furthermore, it accounts for 4.6 % of the Gross Domestic Product (GDP) in Europe. Moreover, it says that manufactures of transport equipment provide additional 1.7 % of the GDP and accounts for 1.5 % of the overall employment [47]. At the same time, there are negative impacts associated to mobility - especially in regard to road traffic - on global climate change, political dependencies, air quality, noise, land use and so forth associated with mobility. These could be summed up under the term *external costs* of traffic. In Germany for example, it is estimated, that the external costs of traffic are 40 billion Euros. The majority of those (about 90 %) are associated to road transport [110].

The European Environment Agency calculates that greenhouse gas emissions caused by traffic (including aviation and maritime transport) has increased by around 34 % between 1990 and 2008. During the same time, greenhouse gas emission caused by energy industries, has been reduced by about 9 %. In 2008, transport was responsible for about 25 % of the total EU's greenhouse gas emissions (CO_2) - more than two third of that are associated to road transport [4]. Clearly, the problem of greenhouse gas emission is a motivation to develop more efficient solutions and many efforts from the industry, the authorities are motivated by it.

However, the discussion about pollution caused by (road) traffic is not limited to greenhouse gas emission. Common vehicles with internal combustion engines also emit carbon monoxide (CO), volatile organic compounds (VOC), oxides of nitrogen (NO_x) and are causing photochemical smog and ground-level ozone (O_3). Moreover, their emissions contribute to the ambient air concentrations of sulfur oxides (SO_3), primary sulfur dioxide (SO_2), particulate matter (PM_{10}) and particulate matter ($PM_{2.5}$) [45, 154].

At the same time, there are two contrary trends. On the one hand, the level of air pollution in many urban environments is increasing (especially in mega cities in developing countries). On the other hand, the society in general is becoming more and more sensible regarding air pollution issues. It needs to be kept in mind that there are important differences between individual pollutants. While the emission of CO_2 should generally be decreased to lower its long-term and global impact on the greenhouse effect, other pollutants such as CO or particular matter have a more direct, local effect on the human health [26, 38]. This work introduces approaches to solve the latter problem by shifting the location of the emission from crowded areas to places, where the impact of pollution is less critical.

3.1.2. Regulations and activities from automotive manufacturers and authorities

There are many methods for urban air quality regulation. Selected ones are briefly introduced in the following.

Some countries are controlling access to areas with high population densities already. There are various approaches in place, such as low emission zones, pricing strategies, license plate programs and emission certificates.

A low emission zone (LEZ) defines an area - commonly city centres - where only those vehicles are allowed to enter which are classified as a low PM_{10} emitting vehicles, while high polluting vehicles are banned from entering such zones. Germany has been established 48 LEZs (called *Umweltzone*) [174]. On an EU-wide level, more than 150 LEZs have been implemented [185].

An alternative kind of driving restriction is called the *cell-based* approach. The basic idea here is to avoid letting vehicles to navigate through a third cell of the traffic network when they wish to travel from one cell to another. That means, that vehicles are allowed to drive within any cell, when it covers its departure or destination. On the contrary, they are forced to circumnavigate some third cell in between. As a result, those vehicles use ring roads. Athens, Amsterdam, Barcelona and Tokyo have set up such systems [185].

Pricing- or incentive-based policies have been established in London, Milan, Singapore and Stockholm. Here, the objectives to reduce pollution in the inner city is attempted by charging cars for entering the city [185].

A license plate program has been installed in various locations, especially in several South American cities like Mexico City (1989), Bogotá (1998) and Sao Paulo (1997), San Jose (2005), La Paz (2003) and Honduras (2008). On given weekdays it is prohibited to enter relevant parts of the city with vehicles which have certain license plates (e.g. if the last digit of the license plate is even or odd). The city of Beijing set up an equivalent restriction rule during the Olympic games in 2008 [185].

Further, the European Union plans to ban ICE driven vehicles from major cities by 2050 [86]. By 2030, it plans to reduce non-hybrid petrol or diesel-fueled vehicles by 50 % while carbon dioxide-free transport of goods in European metropolitan area is subsidized [86]. Thus, it could be assumed the number of hybrid vehicles will increase over the next years. The presented approach makes use of that. Moreover, it exploits the fact that communication technology is currently conquering the automotive environment and will soon be deployed to the majority of vehicles in the western world.

Another instrument to control the overall emission is the European Union Emission Trading System (EU ETS). The system is designed as a *cap and trade* approach. That means that the overall emission shall be decreased by issuing a the limited number of certificates that give the right to pollute to their owners. At the same time, the allocation of certificates should be optimized by the possibility to trade the certificates.

In parallel to such programs from the authorities, citizens demand cleaner vehicles from vehicle manufacturers. They are under constant pressure to improve their vehicles accordingly. One of the main instruments to accelerate the development of cleaner vehicles is most probably tax for gas, which reinforces consumers interest for cleaner vehicles. In Germany, the energy tax respects different types of fuels and they are also related to their environmental impact. For instance, in 2013 in Germany, tax on liquefied petroleum (LPG) was about 16.6 ct/kg, for compressed natural gas (CNG) 18.03 ct/kg, 65.45 ct/l for petrol, and 47.04 ct/l for diesel [49]. Moreover, the European Commission (EC) signed voluntary agreements with the European Automobile Manufacturers Association (ACEA), the Japan Automobile Manufacturers Association (JAMA) and the Korea Automobile Manufacturers Association (KAMA) to reduce CO_2 emission to 140 g/km for the relevant associations' fleets until 2008 and to 130 g/km until 2015 [35]. All these facts cause that car manufacturers invest a lot (about a third of their overall R&D budget) into activities to decrease fuel consumption or emission [180]. The example is typical for many approaches, which address singular vehicles. But in fact - at least in regard to many aspects - it addresses the pollution problem just indirectly. Air pollution is caused by multiple factors and a single vehicle is only one of it.

However, even though these measures are welcomed, they have serious weaknesses. Roughly speaking, they suffer from three main drawbacks [140].

- **Open loop measures**: All of the mentioned instruments are open loop approaches. Thus they are not adaptive to actual conditions. That means that the regulation does not consider if there is only one vehicle in a spatial area (e.g. in the middle of the night), or if there are millions of vehicles in the same area. Furthermore, they do not consider emission by other sources, such as households or the industry. But the combination of particular circumstances are often the reason for critical conditions. In Salt Lake City 2013, weather conditions caused that soot from car exhaust and industrial emissions promoted toxic smog, such that doctors were calling to declare a public health emergency. In 2015, Paris prohibited vehicles to access the inner city because of air pollution [1]. State-of-the-art instruments as outlined above are open loop measures, where the controller cannot compensate for changes in the system.

- **Per vehicle measures**: The degradation of air quality results from the aggregate effect of vehicles. Enforcing per-vehicle measures (unless all vehicles are banned) that target singular but not all vehicles in an integrated way, takes (almost) no account of this effect. In fact, the impact of the increasing number of cars (based on features of registered vehicles) outweighs the emission decrease of the individual vehicle [58]. That means that the overall emission caused by individual road traffic is increasing even though individual vehicles consume and emit less. Thus, per vehicle measures are not suited to control the aggregate level of pollution (which affects physical health or climate conditions).

- **Invasiveness**: All of the named instruments are highly invasive. They may even have unintended consequences for the city. Forcing vehicles away from particular zones may lead to congestion elsewhere and may cause negative socio-economic effects for the relevant area.

3.2. Objectives

The European Parliament set-up the directive 2008/50/EC on ambient air quality and cleaner air for Europe, which defines air quality standards and thresholds per pollutant to *protect human health and the environment as a whole* [130]. The directive defines thresholds for sulphur dioxide, nitrogen dioxide and oxides of nitrogen and carbon monoxide, benzene, particulate matter, (PM_{10} and $PM_{2.5}$) and lead ozone and related NO and NO_2 [130].

In the context of this work, the pollution control problem could be described in the following way: *Develop a control system such that the aggregate emission of HEVs is controlled in way such that given air pollution limits l_i are not exceeded.* This basic objective could also lead to an optimization problem, where the objective is to minimize the aggregate emission of vehicles. The basic motivation for the control approach (in contrast to the optimization objective) is the following assumption: it is beneficial to allow vehicles to emit pollutants as long as the given limit is not exceeded. Even though the aggregate emission in such a setting is assumed to be greater than in a minimization setting, the advantage of allowing more emission is that it generally provides more freedom, which could be exploited to increase the utility in regard to other aspects of the system (e.g. economic benefits, increased user acceptance, etc.).

Generally, air pollution has multiple features (sulphur dioxide, nitrogen dioxide and oxides of nitrogen and carbon monoxide, etc.). Accordingly, given limits target multiple pollutants. However, the objective here is limited to target a scalar value only.

More specifically, the objective is to control CO pollution to a reference value w^* with $w^* \leq l_{CO}$. The system is controlled to w^* (and not to l_{CO}) to guarantee that the actual pollution limit is not (or only during very few and short times) exceeded.

The basic idea of the embedded control loop is illustrated in Figure 3.1.The control algorithms are deployed to the central entity, here the *Central Control Center* and is fed with the difference $e(t) = w^* - y_b(t-1)$. The output of the control algorithm is the control variable $u(t)$, which is broadcasted to all HEVs in a specific region - the *Controlled Region*.The aggregate emission of the vehicles in this region is defined by $y_p(t) = \sum_{i=0}^{N} y_{q,i}(t)$. In case there is no additional polluting source such as households etc., it is $y(t) = y_p(t)$ with $d(t) = 0$; otherwise it is $y(t) = q(t) + d(t)$. The measured air pollution is called $y_m(t)$. The feedback element calculates the filtered emission value $y_b(t) = \lambda y_m(t) + (1-\lambda)y_m(t-1)$, which serves as the feedback signal.

The *TwinLin* application aims at providing a solution that overcomes the drawbacks of state-of-the-art countermeasures outlined in Section 3.1. Accordingly, it should:

- respect the aggregate effect of vehicle emission,
- grant that critical emission levels are not violated,
- be highly adaptive to actual air pollution levels and should be able to react fast to changing conditions,
- not be invasive in the sense that it should not force users to act in a specific (in particular non-attractive) way.

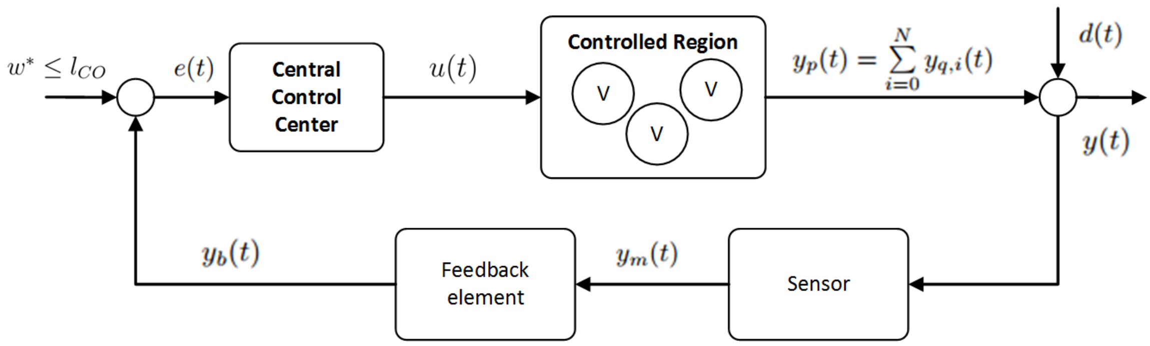

Figure 3.1.: Emission control loop.

Figures 3.2 and 3.3 depict the general objective to control the aggregate effect of vehicles. Assume, that vehicles are the only source of pollution and that there is a pollution level that should not be exceeded (depicted in both Figures 3.2 and 3.3 by the dotted line). The bold grey line indicates the emission by vehicles. Without any aggregate control, emission follows the number of vehicles (Figure 3.2). The high level objective of the aggregate control approach *TwinLin* is to keep the given emission threshold as shown in Figure 3.3.

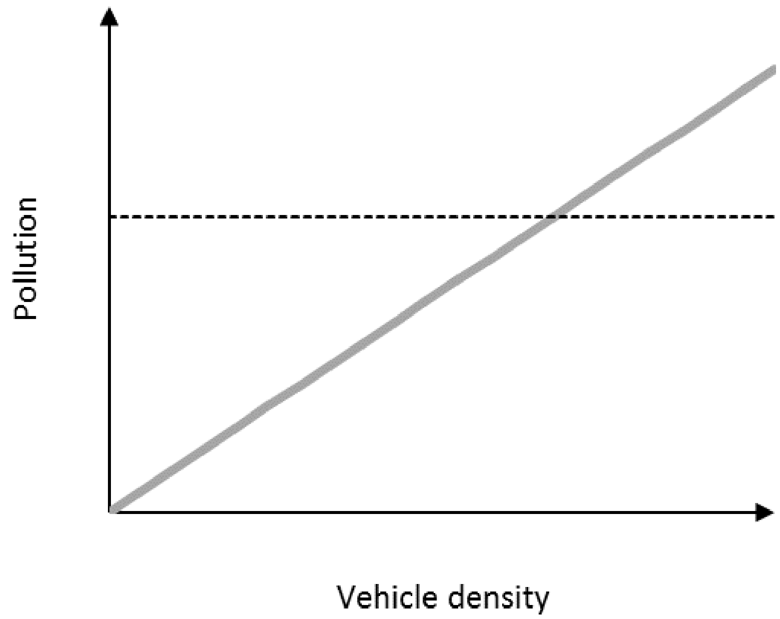

Figure 3.2.: Typically pollution levels follow vehicle density.

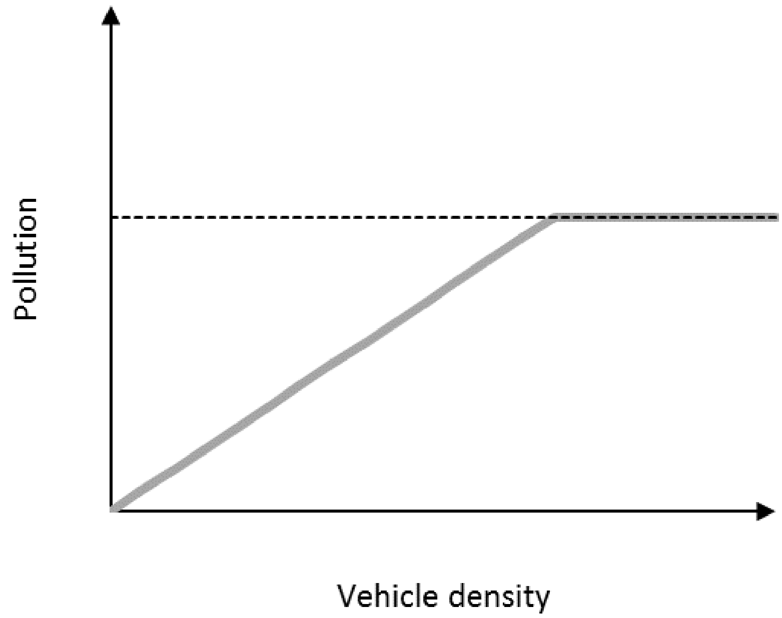

Figure 3.3.: By controlling the aggregate emission pollution is decoupled from vehicle density.

The key idea of the *TwinLin* application is to use the batteries of hybrid vehicles to uncouple the gasoline combustion from the emission geographically. In other words, the application allows that the location and time of vehicle emission is different from the time and location of the energy consumed for mobility. Compared with fully electric vehicles, which shift the emission (if any) rather statically to power plants (and to relevant production facilities[1]), hybrid vehicles are able to switch from ICE to EV mode and could thus decide where their emission has the smallest impact. That is highly valuable since vehicles are often driving through high density areas like inner cities, where additional emission is not wanted. By switching into EV mode, they shift the emission to regions where they could run in ICE mode (like industrial areas or sparsely populated regions).

In contrast to fleets of fully electric vehicles and vehicles with internal combustion engines only, a non-invasive control approach using full hybrids is possible, which allows to control the manner in which pollution is delivered into the environment [140]. By orchestrating the way in which a fleet of such vehicles switches modes, one is able to regulate pollution levels. The switching is orchestrated based on a function of aggregate pollution levels.

Based on estimated pollution level measures, a feedback loop is created that allows hybrid vehicles to change the powering mode according to emission characteristics. With such an orchestrated switching approach, aggregate emission (and pollution) could be controlled such that it remains below a safe, pre-described threshold.

3.3. Requirements and assumptions

Before discussing possible control approaches for the *TwinLin* application, this section provides a brief overview on basic assumption and requirement. They are based on the design constraints outlined in Section 2.2:

- **Aggregate emission**: Defined pollution levels should not exceed pre-described levels in certain geographical area.
- **Best effort behaviour**: The application should be able to cope with vehicles entering and leaving the controlled region.
- **External influences**: Vehicles' behaviour should respond to (uncontrolled) sources of emission.
- **Fairness**: *Dirty* cars (e.g. cars of Euro emission class 1) should not be able to pollute more than *clean* vehicles.

[1] In this thesis, the emission related to the full life cycle of vehicles - which goes beyond the emission related to the actual movement - is not regarded. In a general discussion about electric or hybrid vehicles, it must not be neglected that this does also cause pollution of its components [57].

- **Stability and convergence**: The application shall converge to a state where the desired pollution limit is met. It shall preserve this state over time.

- **Privacy**: The application should secure privacy of individual vehicles (drivers).

- **Scalability**: It is required that the system could handle a small number of vehicles as well as a large number of vehicles.

- **Freedom of choice**: The application should consider that fact, that participants should have the freedom to misbehave (in the sense that they do not comply to the recommended behaviour). In the context of the application, there might be some users who users, who wish to use one mode only and who would thus not follow a recommended behaviour. The requirement is that the application must not fail because of singular non-compliant participants.

The implementation of the concept comes with some assumptions. The most important ones are listed in the following.

- **Vehicle intention**: The objective of each vehicle is to use as little battery power as possible. This assumption seems counter-intuitive at the first sight. But there are some aspects that support this assumption. For example when access to certain parts of the city is restricted to zero emission vehicles, it might have the effect that HEVs wish to reserve relevant battery power to be able to access the zone in EV mode. More important is that this assumption comes with a worst case scenario in regard to the control problem: Vehicles with another interest than driving as often as possible in ICE mode (in the controlled area) could only lower the aggregate emission but not worsen it.

- **Shares of HEVs**: It is assumed that a considerable amount of full hybrid vehicles is deployed. In the simulations introduced later in this chapter, all vehicles are capable to switch from ICE mode to electric mode.

- **Communication capabilities**: All of the considered vehicles are able to receive messages sent by the traffic management centres (or the relevant infrastructure). Communication networks cover all relevant locations at any time. There is no explicit V2V communication and no bidirectional V2I communication needed.

- **Air quality information**: The central entity broadcasts control information based on air quality measures. It has access to perfect air quality data. Pollution levels in a geographic area can be measured or estimated without explicit communication from the vehicles.

- **Static pollution limit**s: Pollution limits are assumed to be fixed over time.

3.4. Traditional control approaches

The first control strategy uses a discrete implementation of the classical PID controller which is given by Equation (3.1):

$$p_{PID}(t) = K_p e(t) + K_i \int_0^t e(t) + K_d \frac{d}{dt} e(t) \tag{3.1}$$

with $e(t) = l^* - y_b(t)$, see Figure 3.1. The problem can viewed as a classical Lur'e problem [89, 150] when the system is (1) modelled by a sector-bounded plant with a non-linear, time-varying behaviour, (2) a noise term is added to the plant's output to account for stochastics and (3) an averaging filter is applied to the feedback signal. The stability of such systems can be guaranteed by selecting proper control gains [150].
However, applying the Lur'e framework is not satisfactory: it models a deterministic system with an additive noise term while the system to be controlled here is a true stochastic one. Therefore, an exact analysis (for a simple integral controller) has been developed in [140] and further analysed in [139]. Both references include the proof of stability assuming that the integral controller gain is $K < \frac{1}{y_p^{ICE}}$ with $y_p^{ICE} = \sum_i y_{q,i}^{ICE}(t)$ (where $y_{q,i}^{ICE}$ denotes the vehicle's emission which runs in ICE mode, see also Figure 3.1).

Simulations described in Section 3.8 will show further characteristics of the algorithm. Even though gains of the controller are fixed, the approach is capable to handle varying numbers of participating vehicles.

3.5. Multiple increase multiple decrease and random early detection

Alternative algorithms, which could be utilized to address a similar problem setup are the multiple increase multiple decrease (MIMD) and the Random Early Detection (RED) control algorithms. These approaches are widely used in the context of internet technology, e.g. [159, 52]. Generally, the algorithms are used to implement a better control and to avoid congestion. Network engineers wish to use the given resource (the network) in a most efficient way. That means the limited bandwidth shall be allocated to senders in a way, that the overall utility is maximized. In other words, network congestion shall be avoided while utilizing the available bandwidth most efficiently. The general setup is thus very similar to the pollution control problem.

3.5.1. Multiple increase multiple decrease (MIMD)

The first algorithm coming from the networking community is MIMD. Applied to the given context, the general goal is to keep the filtered emissions $y_b(t), t = 0, 1, ...$ within some interval $[y_{min}, y_{max}]$, see Figure 3.4.[2]

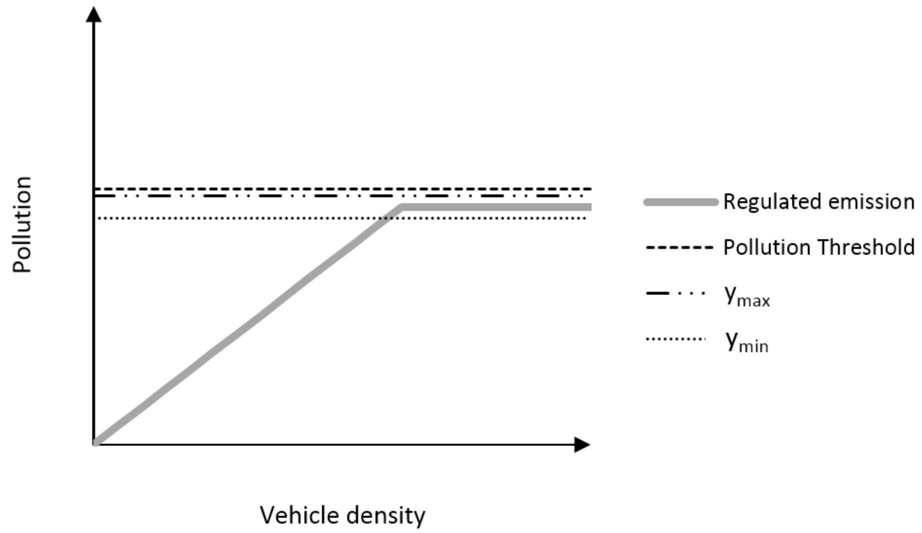

Figure 3.4.: Aggregate emission controlled with MIMD boundaries y_{max} and y_{min}.

With the assumption in mind that vehicles might choose the EV mode more often than assumed here, the resulting $y_b(t)$ is an upper bound for the emission. To this end two factors a_d and a_i satisfying $0 < a_d < 1 < a_i$ are chosen to apply the AIMD algorithm. In effect, the probability p is updated according to:

$$p(t+1) = \begin{cases} a_d p(t) & \text{if } y_b(t) > y_{max} \text{ with } 0 < a_d < 1 \\ a_i p(t) & \text{if } y_b(t) < y_{min} \text{ with } a_i > 1 \\ p(t) & \text{else.} \end{cases} \tag{3.2}$$

MIMD is designed for quick convergence and fast reaction to changes in the system and disturbances. However, the parameters a_d and a_i have to be adapted to the specific system. If they are too close to 1 then convergence will be very slow. If they are too far apart from 1, oscillations could be observed. A possible countermeasure is to choose a good filter for the emissions.

[2]The discerning reader may question the approach. Namely, why it is not aimed at minimizing the aggregate emission. In this approach, emission is allowed as long as the aggregate emission does not exceed critical limits. The motivation for that is that there might be parallel objectives which could be better achieved if vehicles are allowed to emit as much as possible.

3.5.2. Random early detection (RED)

An alternative approach to avoid congestion problems in the internet (while utilizing the available resources most efficiently) is the RED algorithm [156]. The basic idea of RED (which is intended to be deployed at network routers) is to adjust the probability to accept incoming packets based on the filling level of the router's buffer. This is done in a way that the acceptance rate decreases with an increasing buffer filling level. Applied to the ITS scenario, RED adjusts the probability of a car switching into EV mode according to a non-linear function of the average pollution levels y_b. An exemplary RED probability curve is depicted in Figure 3.5.

Similar to the MIMD algorithm, it is not only an upper bound y_{max} defined but also a lower bound y_{min} and the control objective in RED is to maintain pollution levels between these two thresholds. The argumentation for the approach to allow *as much emission as possible* could be transferred from the MIMD section (Section 3.5.1).

The motivation for the approach is the hypothesis that there exists - for a range of traffic densities - an equilibrium probability p^* corresponding to y^* with $y_{min} \leq y^* \leq y_{max}$ and that it is stable and that the system converges to it.

The system deployment architecture is again very simple: It is assumed that there is an infrastructure, which knows about the latest, measured pollution values $y_m(t)$. The probability that this is broadcasted to all the vehicles is adjusted according to an update function, such as defined by Equation (3.3) [140].

$$p_{EV}(t+1) = \begin{cases} 1 & \text{if } y_b(t) > y_{max} \\ 0 & \text{if } y_b(t) < y_{min} \\ p_{max}(t)\frac{y_b(t)-y_{min}}{y_{max}-y_{min}} & \text{else.} \end{cases} \tag{3.3}$$

The stability of RED has been widely studied and proven in the context of internet congestion control [157]. See for example [27], where RED is emulated in a distributed fashion. Unfortunately, no firm guidelines exist on configuring RED parameters (y_{max}, y_{min}, RED function f). That means that the probability curve is based on a range of vehicle densities.

Not all setups can be accommodated using a single curve; it would require adjustment of f based on feedback. There are approaches to do this, see for example [52, 29]. In this regard however, a concern is the performance of the algorithm in the presence of rapidly varying traffic densities.

A key character of a well-configured RED control is that it implements a fast converging and adaptive control system.

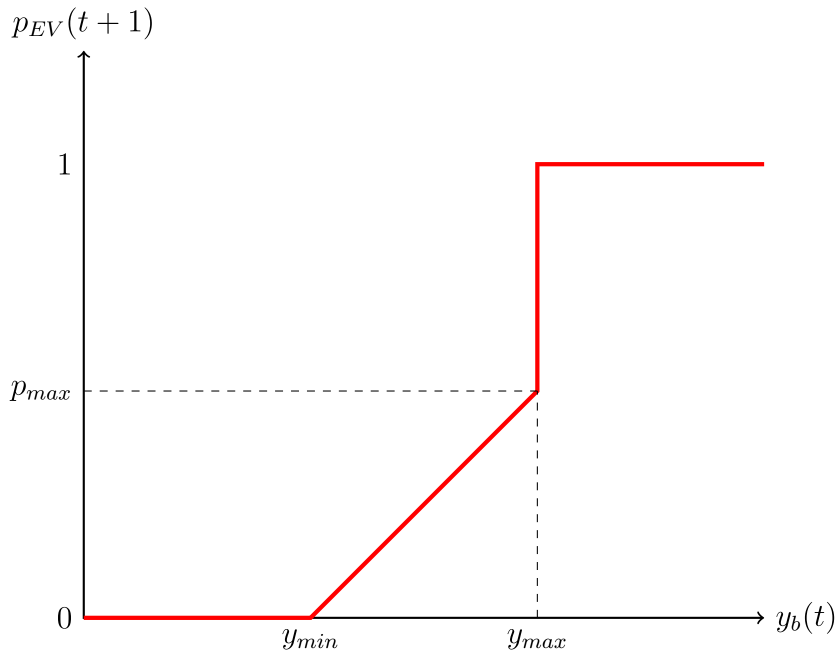

Figure 3.5.: Adjustment of $p(t)$ according to the (averaged) pollution measure $m_f(t)$ and with a RED-like feedback control loop.

Regarding the fairness aspect, the classic RED implementation, explicit fairness measures are not considered [55]. However, it is fair if senders have homogeneous parameters [103, 8]. Extensions of the classic RED approach exist which implement fairness also with heterogeneous clients, such as [103, 8].

The implementation here could be done in two ways, which lead to the same results: first, the infrastructure is running the update function f based on air pollution measurements and broadcasts the inferred probability p. The second way for an implementation is that the infrastructure broadcasts the air quality measures and vehicles run the update function f themselves.

3.6. Additive increase multiple decrease (AIMD)

The AIMD control approach is similar to the MIMD and RED algorithm introduced in Section 3.5.1. Again, the objective is to maintain the control value between predefined levels. The AIMD algorithm is defined by Equation (3.4).

$$p(t+1) = \begin{cases} \alpha p(t) & \text{if } y_b(t) > y_{max} \text{ with } 0 < \alpha < 1 \\ p(t) + p_{add} & \text{if } y_b(t) < y_{min} \text{ with } p_{add} > 0 \\ p(t) & \text{else.} \end{cases} \tag{3.4}$$

The basic idea for AIMD algorithm introduced in [31] consists of two parts: the additive increase part and the multiplicative decrease part. Applied to the *TwinLin* application, the probability to drive in ICE mode is constantly increased by an additive term as long as the pollution threshold is not exceeded. In the multiplicative decrease part, the probability is lowered by some multiplicative factor. Stability, convergence and fairness of AIMD have been widely discussed including the initial AIMD publication [31]. A central benefit of AIMD is that it is scalable in the sense that AIMD parameters do not depend on the number of participating vehicles.

The implementation of the algorithm for the *TwinLin* application is very different to traditional control approaches and to the implementation of the aforementioned algorithms MIMD and RED. The main reason for that is, that the way it is implemented here could be extended to implement network utility maximization based on the findings from [183]. It is briefly described in the following. The starting point for the implementation is, that it should aim at controlling the emission of vehicles over a particular time window of e.g. 10 seconds. Accordingly, a set value for the sum of emissions from all vehicles over that time window is defined. Each time the air pollution exceeds this limit, the central infrastructure broadcast a relevant event notification. Receivers of this broadcast (the hybrid vehicles) apply AIMD to an individual emission threshold. They update (lower) their (private) emission threshold by multiplying it with α. In case no broadcast is generated - as the overall emission threshold is not exceeded - vehicles increase their emission threshold with the additive term when the time window has ended. With a start of a new time window, the updated emission limit (or budget) is fully available again. Independent from the available budget, all hybrid vehicles switch to ICE mode at the beginning of the new time window. They last in ICE mode as long as they do not receive the binary information from the central infrastructure saying that the overall emission threshold is exceeded or as long as they have not exceeded their private emission budget.

3.7. Short comparison of proposed algorithms

As a short summary of the the previous discussion of PID, MIMD, AIMD and RED, the Table 3.1 is presented for a quick comparison.

	PID	MIMD	AIMD	RED
Fairness	+	+	++	+
Best effort	-	0	+	0
Stability & Convergence	++	+	++	++
Communication Effort	-	0	++	+
Scalability	-	-	+	0

Table 3.1.: Comparison of selected control algorithms in for aggregated pollution control. The "+" stands for a positive assessment, "0" for a unclear or non-applicable assessment and "-" for a negative assessment.

3.8. Simulations

This section describes the simulations performed to assess the pollution control approaches outlined in Sections 3.4, 3.5 and 3.6. The traffic simulation tool SUMO is used. If not stated otherwise, simulations aim at controlling CO emission. Note, that the approaches could easily be applied to any other pollutant or scalar set value. Emission is calculated using the HBEFA model or the emission modelling from [24], respectively.

3.8.1. Vehicle emission without control

The first simulation is performed without applying any control strategy. It is assumed that all vehicles run in ICE mode. It serves as a reference and underlines the basic motivation to uncouple the extend of emission from the number of vehicles in the network.

The simulation setup is based on an entirely artificial road network. Road traffic is generated in a grid-like network, where streets intersect orthogonally. It is illustrated in Figure 3.6. For the purpose of the simulations, cars enter at five different points

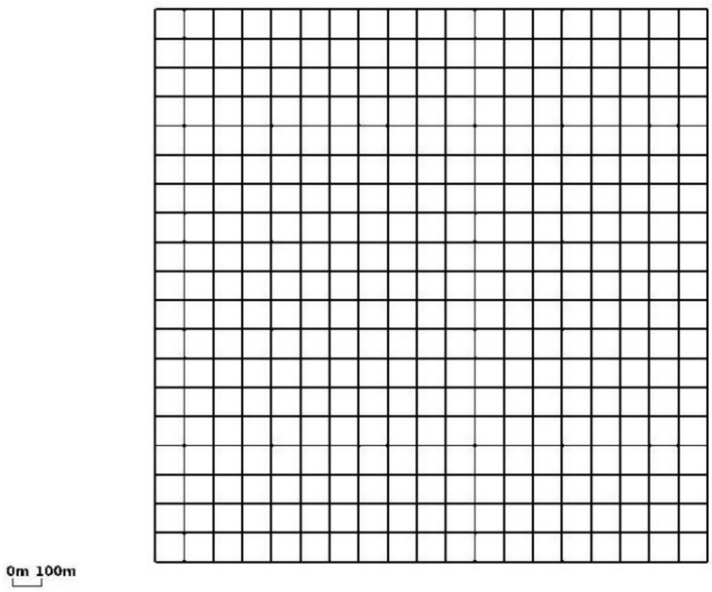

Figure 3.6.: Simple road network used for initial simulations.

in the network and drive along random routes. For 10,000 seconds the number of vehicles in the network is increased until 2000 cars have entered. The simulation continues for six hours. It is assumed that all cars are petrol electric hybrid cars with weight below 2.5 tonnes and with combustion engines larger than 2000 cc (this corresponds to an average passenger car). Further, a emission-class mix is defined for the vehicles which is based on the average distribution in Germany: Euro emission class

4 by 88 %, class 3 by 9 %, class 2 by 3 %, and class 1 by 0.4 %. All this information is included in the simulation scenario setup for the purpose of emission modelling. The evolution of the number of cars over time in the network is depicted in Figure 3.7.

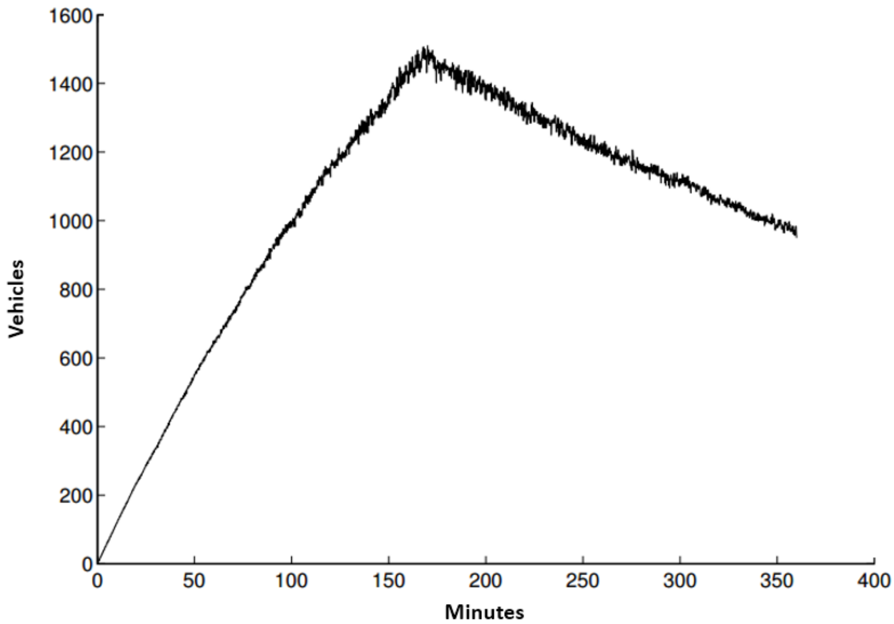

Figure 3.7.: Number of vehicles over simulation time for the simulation described in Section 3.8.1.

Figure 3.8 shows the evolution of emissions over time in this uncontrolled scenario. As initially stated, the amount of emission follows the number of vehicles in the network. This relationship would also hold, if there were some hybrid vehicles running in EV mode in a random fashion. With the help of the aforementioned control strategies, this relationship is broken. This will be shown with help of simulations introduced in the following sections.

3.8.2. Integral control of vehicle emission

The first control approach simulation uses an integral control strategy applied to the same simulation scenario as described in the previous section. A set value of 50 grams of CO emission per minute is defined. The evolution of the emissions and the corresponding evolution of the broadcasted probabilities is shown in Figure 3.9. Note that the picked pollution target values were selected arbitrarily but can easily be adapted to any other level.

The second simulation of the integral control approach is based on the same simplified grid as before. The main difference to the previous simulation setup, that congestion could be observed in the network. This simulation shows the impact of

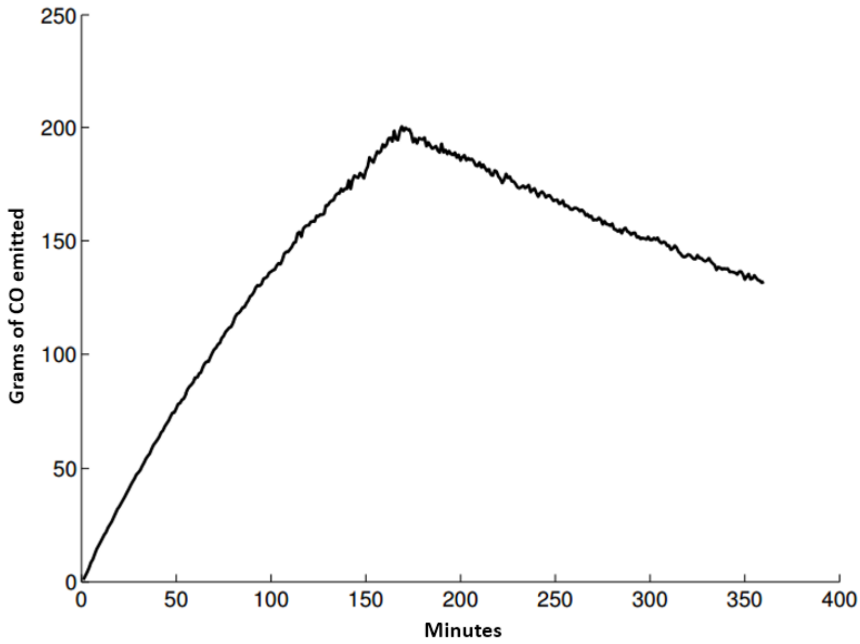

Figure 3.8.: Evolution of vehicle emission, who run in ICE only. It could easily be seen that it follows the number of vehicles in the simulation (see Figure 3.7).

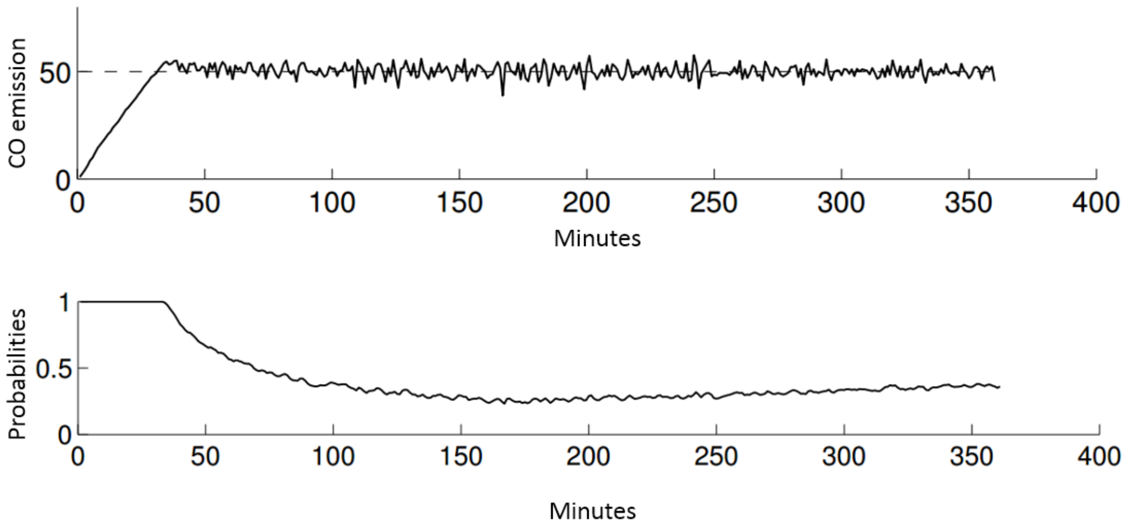

Figure 3.9.: Evolution of emission and broadcasted probabilities of the first integral control simulation described under the Simulation 2 heading.

traffic conditions on the performance (in fact it still works well) of the integral control approach. The number of vehicles over a simulation time of 2500 seconds is shown in Figure 3.10.

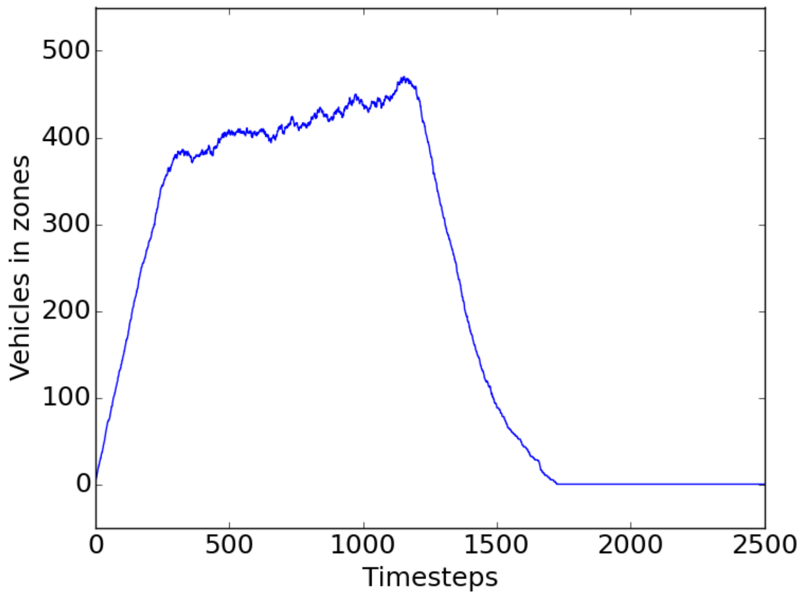

Figure 3.10.: Number of vehicles over time in the simulation setup where congestion is observed.

While the control value is maintained (see Figure 3.11), the critical reader may ask why the number of vehicles driving in ICE mode is increasing while the emission is maintained during the same time, see Figure 3.12. The reason for that is the observed congestion. The underlying emission model assumes that emission depends on vehicles' velocities. As shown in Figure 3.13 the congestion lowers the velocity of vehicles. Accordingly the mean emission of CO decreases as well (Figure 3.14).

The next simulation is based on a setup which is equivalent to the first simulation setup described within this section. The only difference is, that there is an extra source of CO emission added, which is not controlled. From minute 180 to 270, it contributes 30 grams of CO per minute. The performance of control approach with the external pollutant is depicted in 3.15. The light grey coloured series shows the aggregate emission (including the external source of emission), the black coloured series shows the emission of hybrid vehicles only.

3.8.3. Vehicle emission controlled by MIMD

Based on the same simulation setup as before (with and without external source of emission), the central controller uses a MIMD controller now. In contrast to the

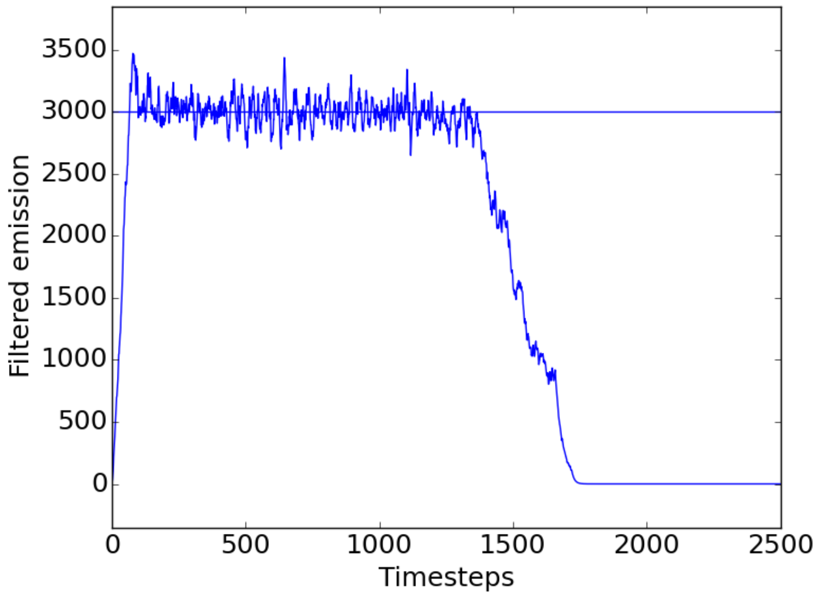

Figure 3.11.: The approach works well with varying traffic conditions - the control value is maintained.

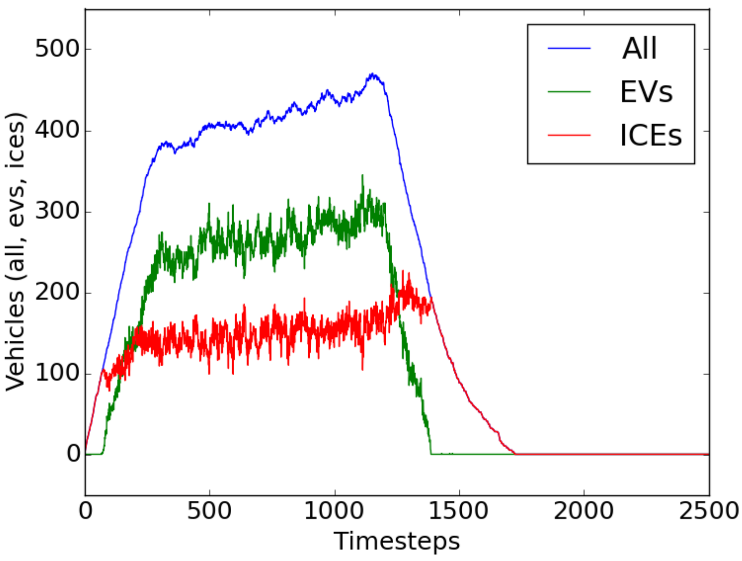

Figure 3.12.: Number of vehicles driving in EV and ICE mode over simulation time in the congested scenario.

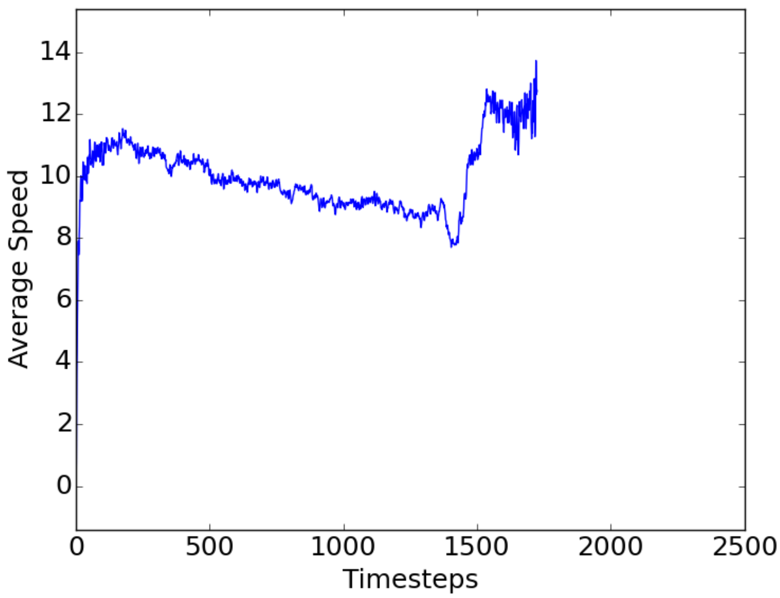

Figure 3.13.: Average velocity over simulation time. It is decreasing as the network becomes more and more congested as long as the number of vehicles increases (see Figure 3.12). When the number of vehicles decreases, the average velocity increases again.

integral control approach, the algorithm aims at to keep the value of the emissions between 40 and 60 grams every minute. Again, it is the central controller, which updates updates and broadcasts probabilities. Figure 3.16 shows the probabilities broadcasted and the aggregate emission without any external emission sources, Figure 3.17 with the external pollutant.

3.8.4. Vehicle emission controlled by RED

The next simulation uses an RED-like approach as introduced in Section 3.5.2. The central controller updates the broadcasted probability aiming at to keep the value of the emissions between 40 and 60 grams every minute - equivalent to the MIMD simulation before. The same network and flow of vehicles serves as the basis for the simulation again. Figure 3.18 shows the evolution of broadcasted probabilities and aggregate emission for the setup without the external pollutant, Figure 3.19 with the external pollutant (again with the same characteristics as before). The performance of RED is similar - even slightly better - than the MIMD approach.

3.8.5. Vehicle emission controlled by AIMD

As outlined in Section 3.6, the implementation of the AIMD approach differs from the implementation of RED, MIMD and the integral control. It is assessed with help of

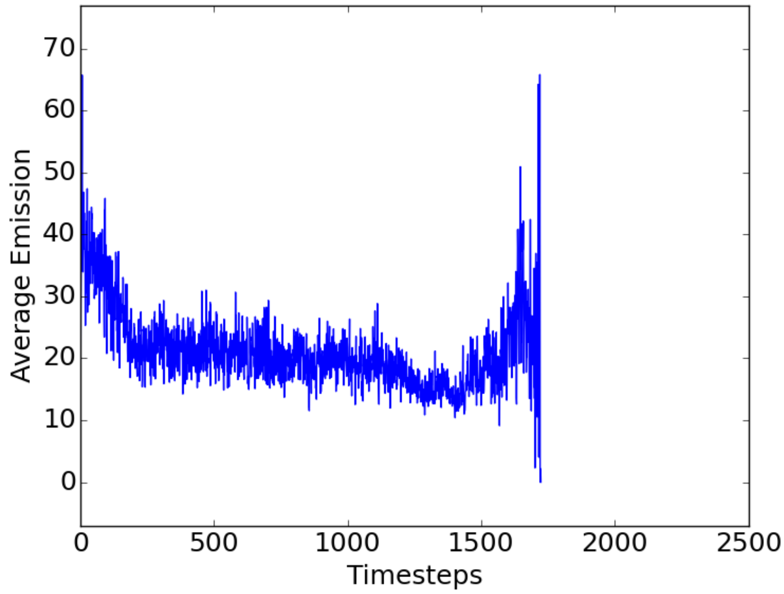

Figure 3.14.: Average emission of vehicles in ICE mode. The general evolution follows the average velocity (see Figure 3.13) as the emission model assumes that the emission depends on the vehicles' speeds. The great variance in emission at the end of the simulation is caused by the decreasing number of participating vehicles.

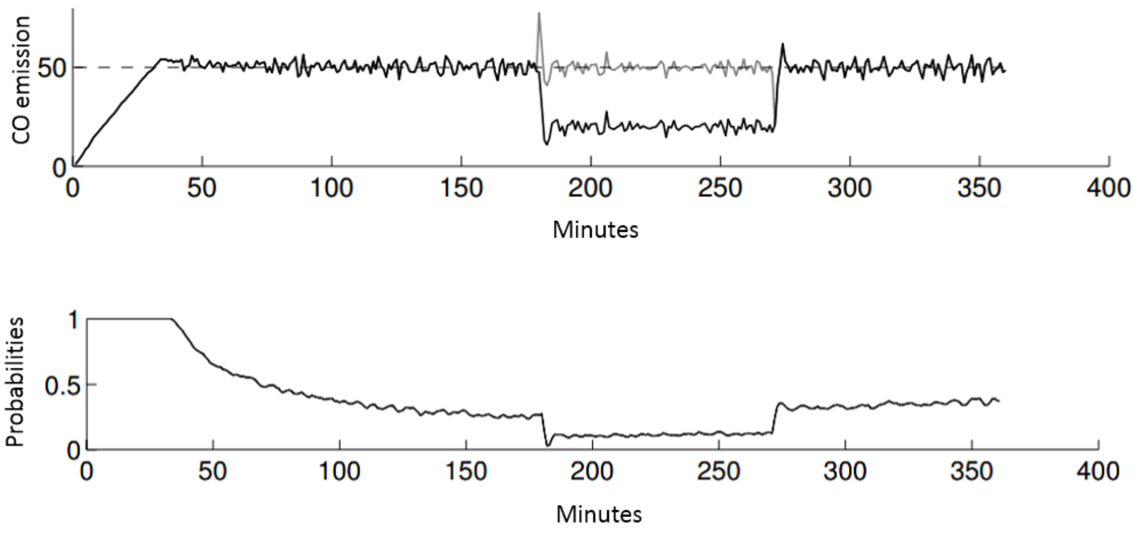

Figure 3.15.: Aggregate CO emission of the controlled hybrid vehicle fleet with external emission source using the integral controller described in Section 3.8.2.

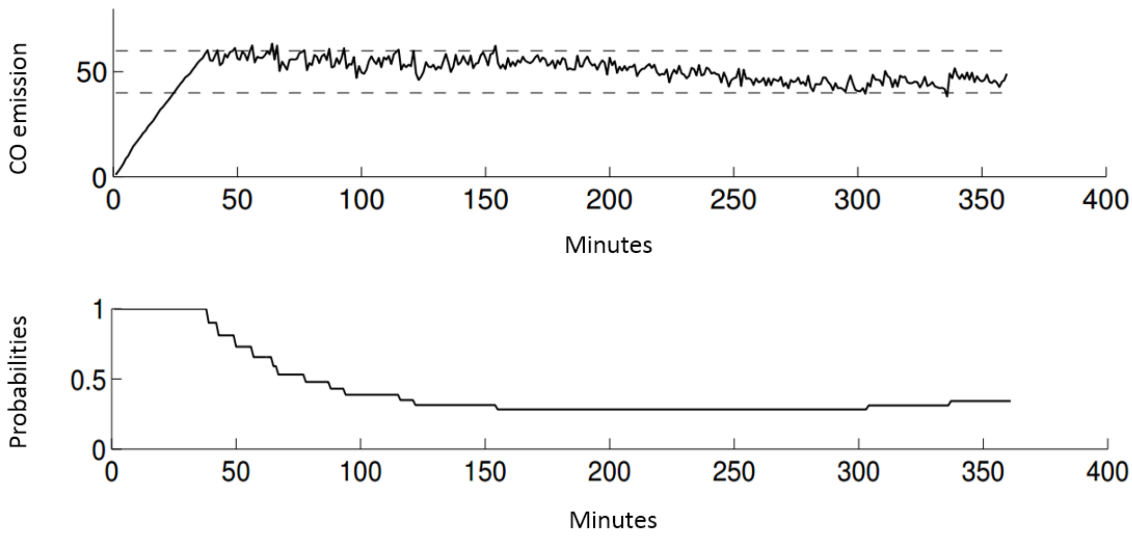

Figure 3.16.: Evolution of broadcasted probabilities and aggregate emission using MIMD (Simulation 3 described in Section 3.8.3).

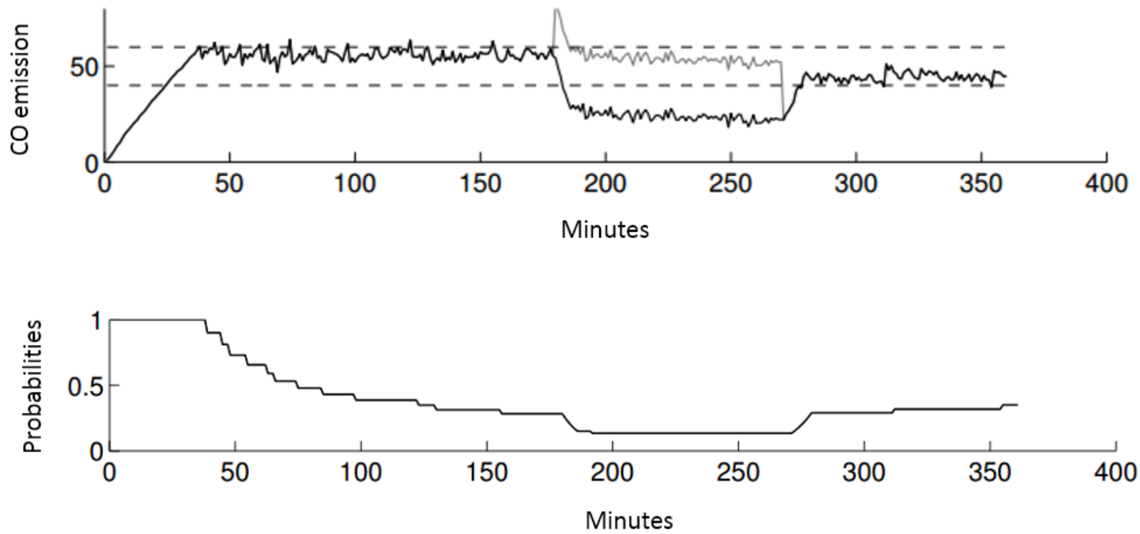

Figure 3.17.: Broadcasted probabilities to drive in ICE mode and aggregate emission including external source of emission using MIMD again (Simulation 3 described in Section 3.8.3).

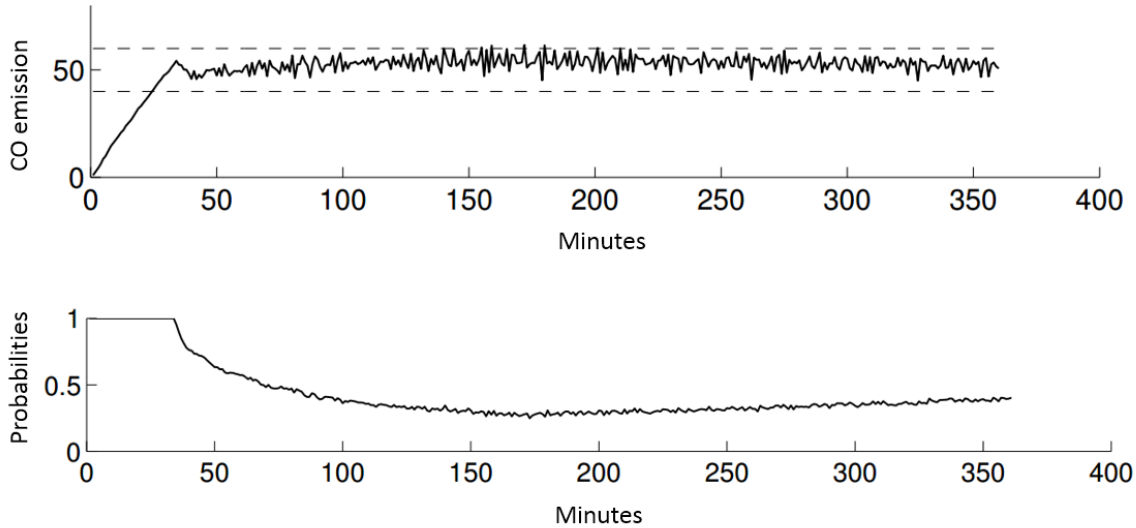

Figure 3.18.: Evolution of aggregate emission and probabilities of an individual vehicle using the RED-like approach.

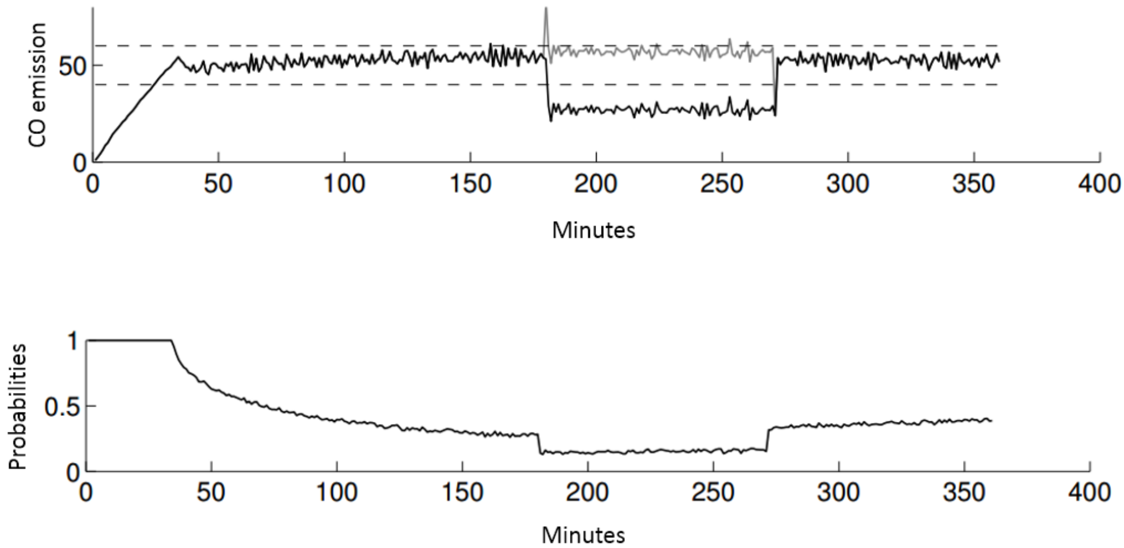

Figure 3.19.: Vehicle specific probabilities to drive in ICE mode and aggregate emission including external emission source.

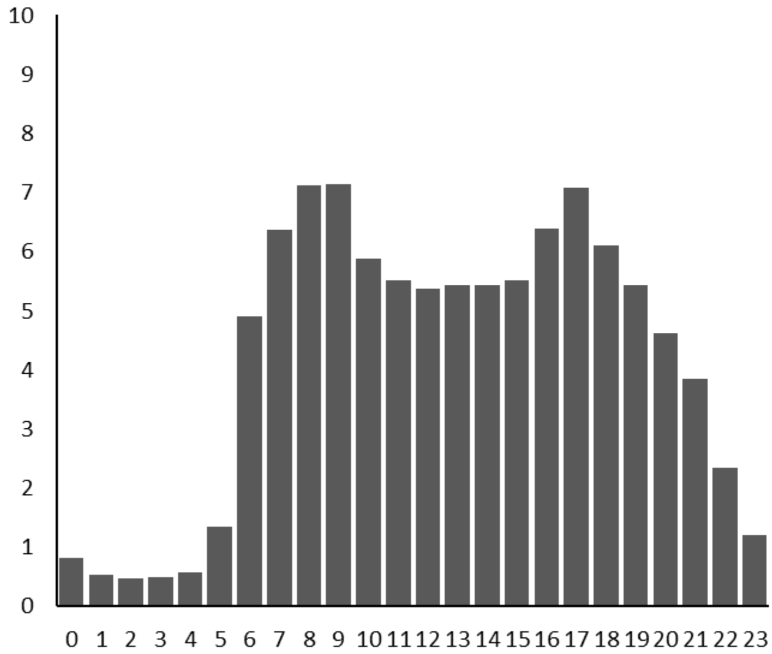

Figure 3.20.: Distribution of vehicles on a work day in the city centre of Berlin. It follows a characteristic density distribution of vehicles in urban areas.

simulation again. It is assumed that the distribution of vehicle density in the network follows a common characteristic in urban environments during working days. This feature is the basis for the following simulation scenario: the characteristic distribution of traffic load in Berlin has been taken from [59] and [60]. Figure 3.20 shows the distribution of passenger cars on the road over one working day. The distribution of European emission standards classes (EURO 4 by 88 %, EURO 3 by 9 %, EURO 2 by 3 %, and EURO 1 by 0.4 %) among passenger cars given for Berlin [65], is reflected in the distribution of vehicle in the simulation scenario as well. Equivalent vehicles (regarding weight and fuel type) per group are assumed; no heavy vehicles are considered. In addition, an even distribution of vehicles' velocities between zero and 60 km/h is assumed (which is similar to the Berlin velocity distribution [161]).

To work with a realistic number of vehicles, the distributions are applied to the mean number of vehicles per day in Berlin allocated to road network links, which is given by Berlin [60].

In this simulation, the objective is to limit the emission to a set value of 6000 mg for time window of 10 seconds. The additive term in AIMD is set to 1 mg CO, the multiplicative term to 0.9.

Figure 3.21 illustrates the sum of emission within a time window. The horizontal

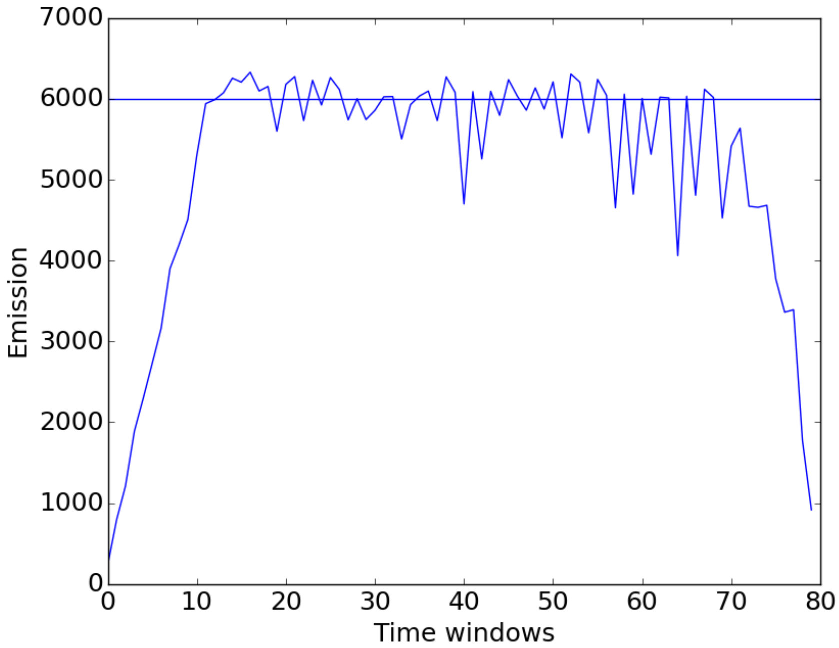

Figure 3.21.: CO emission controlled using AIMD.

line depicts the emission threshold. In most cases the threshold is hold. In case it is exceeded, the infrastructure broadcasts a congestion event to vehicles, see Figure 3.22.

It is assumed, that all vehicles choose to run in ICE mode as long as their emission budget is not exceeded. They do switch to EV mode when their budget is exceeded or, when they receive the broadcasts (congestion events) from the central infrastructure.

Figure 3.23 shows the evolution of private emission budgets over simulation time of three vehicles. These vehicles were randomly chosen. The structure of the series is typical for AIMD. Note that the outliers in the private emission budgets (e.g. shortly before timestep 400 for the emission budget illustrated in blue colour) is caused by the setup of the simulation: From time to time, vehicles exit the simulation and enter it again. While they are outside, the emission budget is set to zero. When they enter again, their emission budget is restored to the value before they exited the simulation.

Figure 3.24 shows the number of vehicles in ICE or EV mode. The aforementioned behaviour is clearly shown: At the beginning of a time window vehicles run in ICE mode while most vehicles have switched to EV mode while in the end.

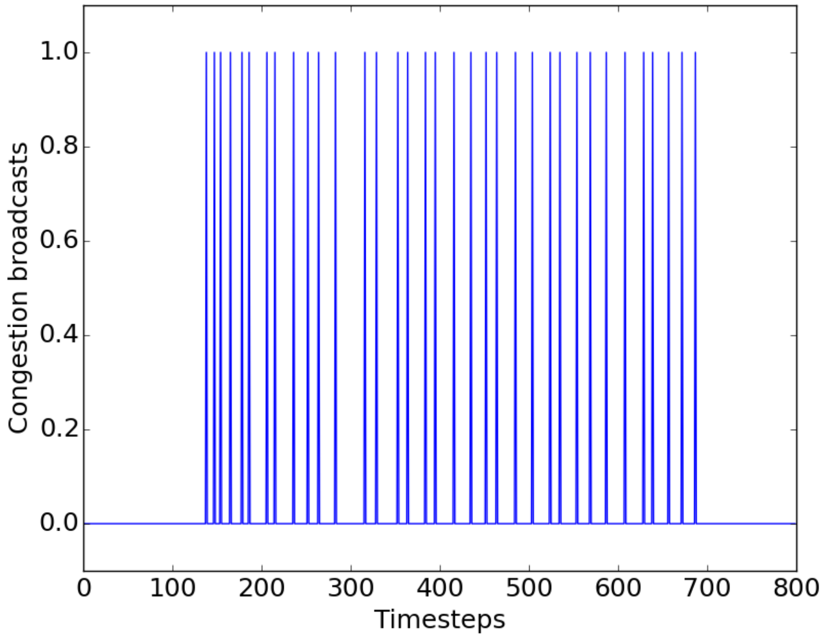

Figure 3.22.: Timing of congest events broadcasted by the centeral infrastucture.

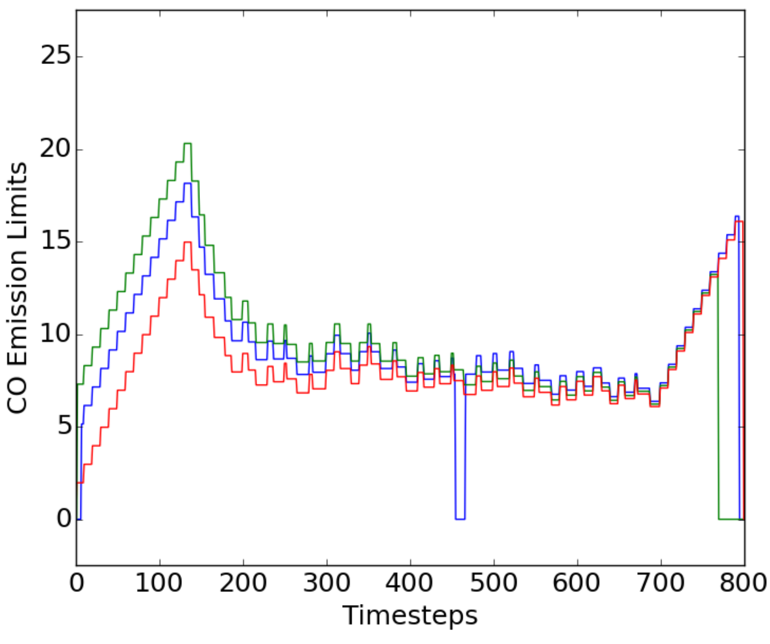

Figure 3.23.: Private emission limits of vehicles over simulation time in the context of the AIMD simulation.

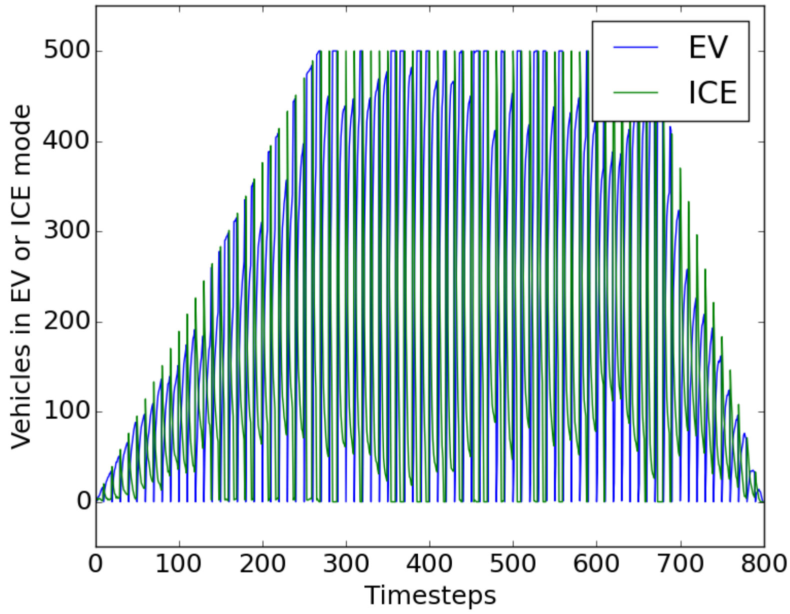

Figure 3.24.: Evolution of vehicle numbers in ICE or EV mode in the context of the AIMD simulation.

3.8.6. Controlling emission in small areas

Looking at all simulations done before, one may argue that the local air pollution may vary heavily in the controlled region. To fight this problem, the controlled area is divided into multiple, smaller zones, which could be controlled separately. In the following simulation, the controlled area around Rosenthaler Platz in Berlin is split up into nine (3x3) zones of equal size. The approach aims at controlling the air pollution in all sub-zones to one common value, such that the emission in the overall area is balanced. In Figure 3.25 the development of emissions in each zone is shown when per sub-zones are not individually controlled. As expected, the level of pollution differs among zones heavily. Figure 3.26 shows the development when each zone is controlled with an individual controller: The level of pollution is harmonized here among all the zones.

The approach could be adapted easily to control the air pollution to different values, e.g. one may wish to control air pollution in the vicinity of kindergartens, schools, hospitals to different levels than in less critical areas.

3.9. Utility functions and a concept of fairness

All algorithms and simulations have shown the effectiveness of pollution control strategies by achieving a desired level of pollution. This result was obtained by

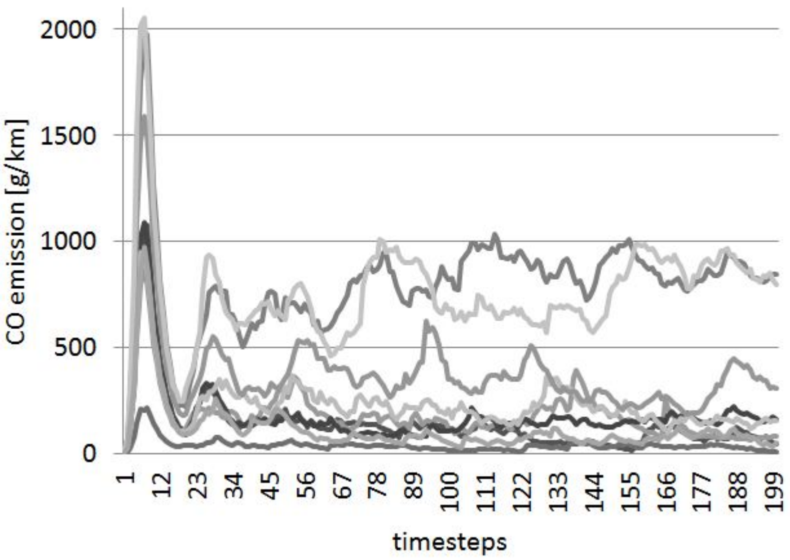

Figure 3.25.: Development of air pollution values in each sub-zone in the case where the complete area is controlled by one controllers only. The air pollution is not balanced over all areas of the controlled region.

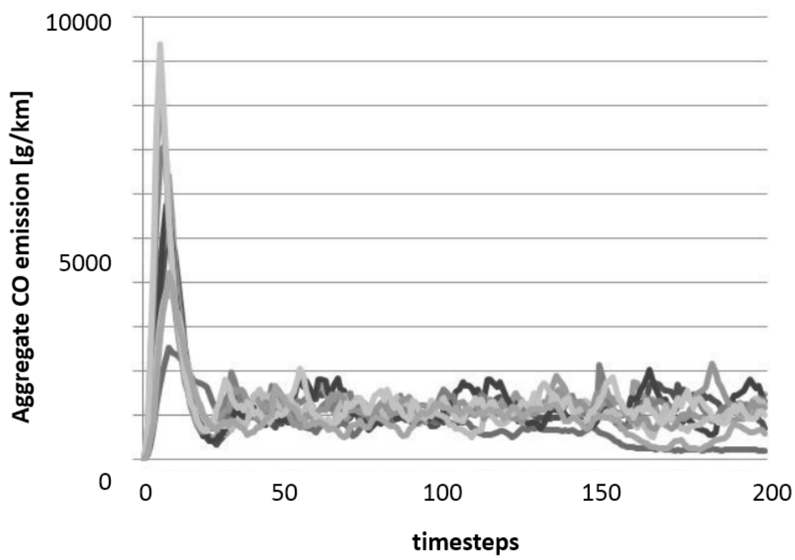

Figure 3.26.: Development of air pollution when each sub-zone is controlled by an individual controller. Air quality is harmonized in the complete region.

allowing only some hybrid vehicles to travel in ICE mode, according to a probability distribution that takes into account the distance between the current pollution level and the desired one.

Clearly, there are multiple ways in which cars can collaborate to achieve the desired pollution levels. The freedom can be exploited to address particular objectives. In the following, an approach is introduced to implement fairness aspects. In order to do so, a utility function for each vehicle is used to encapsulate the desired level of fairness between individual drivers. Accordingly, the utility of a vehicle is equivalent to its emission. Utility fairness thus means, that all vehicles shall be allowed to emit the same amount of pollution.

In the context of the implementation of PID, MIMD, AIMD and RED controllers, this amounts to tailoring the probability communicated to each vehicle according to the properties of that vehicle (e.g., the type of vehicle or the remaining charge in the battery) or with minimal communication requirements.

3.9.1. Emission class fairness

In the first example, the objective is to equalize the pollution per vehicle over all simulated emission classes. In other words, clean vehicles should be allowed to travel in ICE mode more often than dirty cars. Accordingly, each vehicle adjusts the probability of travelling in EV or ICE mode by considering its own nominal emission level (i.e., whether it is Euro 1, 2, or 3 or 4). The overall probability p_i of the i'th vehicle to travel in ICE mode is

$$p_i = p \cdot \widetilde{p}_{class} \ , \tag{3.5}$$

where p denotes the broadcasted probability (a higher p leads to more vehicles running in ICE mode) equal to all vehicles and where $\widetilde{p}_{class}$ depends on the emission class of the vehicle. According to the choice of probability ($\widetilde{p}_{class}$ is lower for dirty vehicles than for cleaner vehicles), dirty vehicles are less likely to be allowed to travel in ICE mode than cleaner vehicles. Figure 3.27 shows the adapted probabilities of travelling in ICE mode for each vehicle class. It could be seen that probabilities vary with the level of pollution. In general, dirty vehicles, i.e. those belonging to class Euro 1, are less likely to be allowed to travel in ICE mode than cleaner vehicles (which belong to Euro emission class 4).

3.9.2. Individual emission fairness

A refinement of the emission class fairness approach presented before is to take into account the actual pollution produced by each vehicle and to use this as a basis for the utility calculation for each vehicle: Generally, cars that spend more time driving in the network produce a more emission. This issue can be addressed in one of the two following ways:

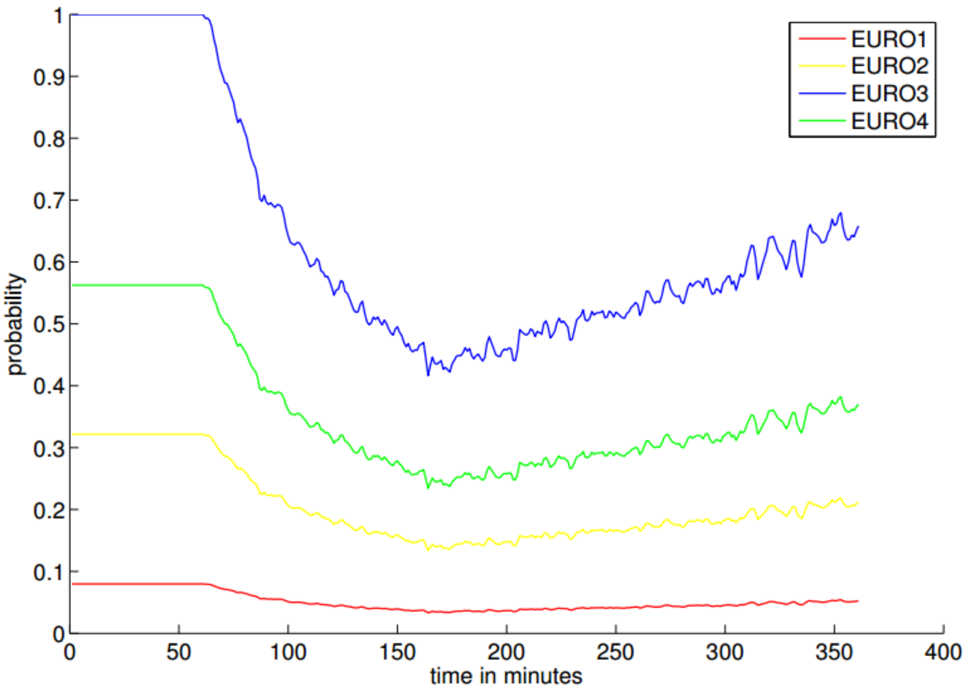

Figure 3.27.: Probabilities of travelling ICE mode depend on the emission class of the vehicle. The probability to travel in ICE mode is lower for dirty vehicles (belonging to Euro emission class 1) than for cleaner vehicles, which belong to Euro emission class 4.

1. Give car i a budget of $C_i \geq 0$ of some pollutant. The vehicle is allowed to expend this budget during a trip or during a certain time. Such budget could for example be bought by the driver in a pay-as-you-go fashion or there could be a subscription scheme to do this. The idea then is that the driver can use his ICE complying with the pollution control algorithm until he or she has used all his budget. This can be done in a strict fashion, where the compliance with the algorithm is strict until the budget is used. Another way of doing this is to adapt the broadcasted probability according to Equation (3.6).

$$p_i = p \cdot \widetilde{p}_{individualEmission} \tag{3.6}$$

   with

$$\widetilde{p}_{individualEmission} = \frac{\sum_{i=0}^{t} y_i(t)}{C_i}. \tag{3.7}$$

   If no budget remains, the driver has to use his electric drive or has to leave the zone or stop the vehicle in both cases.

2. The above approach relies on a infrastructure as it is needed to define budgets. In order to develop a distributed system without any infrastructure, a budget-free approach is required: Let each car calculate a moving average of its emissions $\bar{y}_i(t+1) = \lambda y_i(t) + (1-\lambda)\bar{y}(t)$. Then a factor $\widetilde{p}_{pollution}$ can be computed according to a common function or a function that depends on the characteristics of the individual car. The individual probability is then calculated according to

$$p_i = p \cdot \widetilde{p}_{pollution}. \tag{3.8}$$

The following simulations apply the first approach. The simulation set-up described in Section 3.8.1 is used again. Each vehicle receives a budget of 20 grams of CO to emit over all of its journey. The broadcasted probability is adjusted by each vehicle according to Equation (3.5). In Figure 3.28 it can be seen that the integral control algorithm still works well in this scenario. The corresponding evolution of the broadcasted probability is depicted in Figure 3.29. The evolution of individually calculated probabilities is shown in Figure 3.30. Figure 3.31 shows the evolution of individual probabilities according to Equation (3.6). In order to validate the result, Figure 3.31 shows the average individual probabilities per emission class. Es desired, cleaner vehicles (Euro emission class 4-vehicles) might drive much more often in ICE mode as dirty vehicles.

In order to show the fairness effect, the result is compared to the scenario where no fairness is implemented. Figure 3.32 shows the result of the reference scenario. It could easily been seen that the differences between vehicles belonging to different emission classes are very big. I.e. Vehicles of Euro emission class 1 emit about seven times more than vehicles belonging to emission class 3 or 4. When applying the fairness approach, the average emission is balanced among emission classes as shown in Figure 3.33.

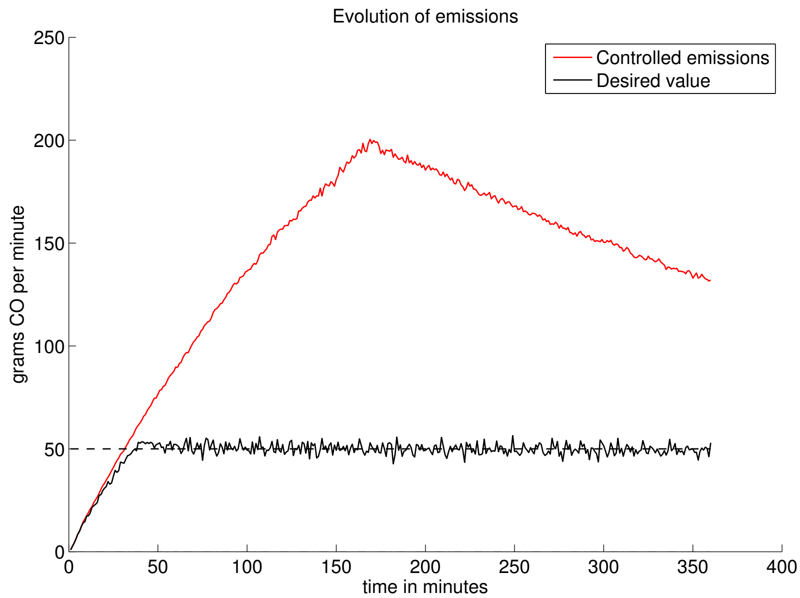

Figure 3.28.: Uncontrolled and controlled pollution in the individual pollution fairness scenario.

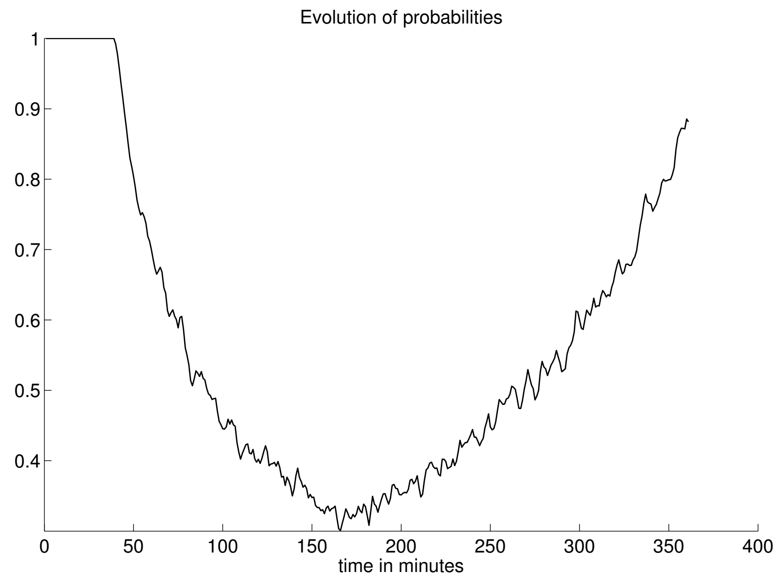

Figure 3.29.: Evolution of broadcasted probability, which is common for all vehicles.

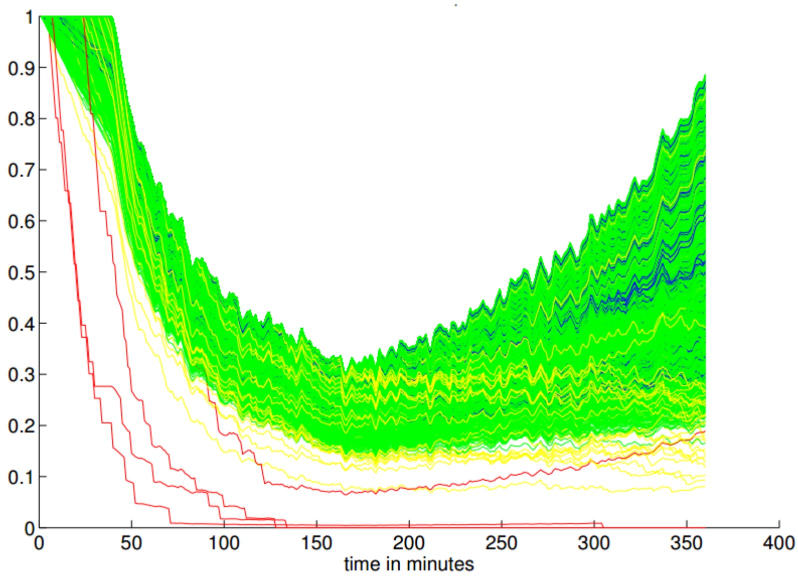

Figure 3.30.: Evolution of individual probabilities, it is calculated according to Equation (3.6).

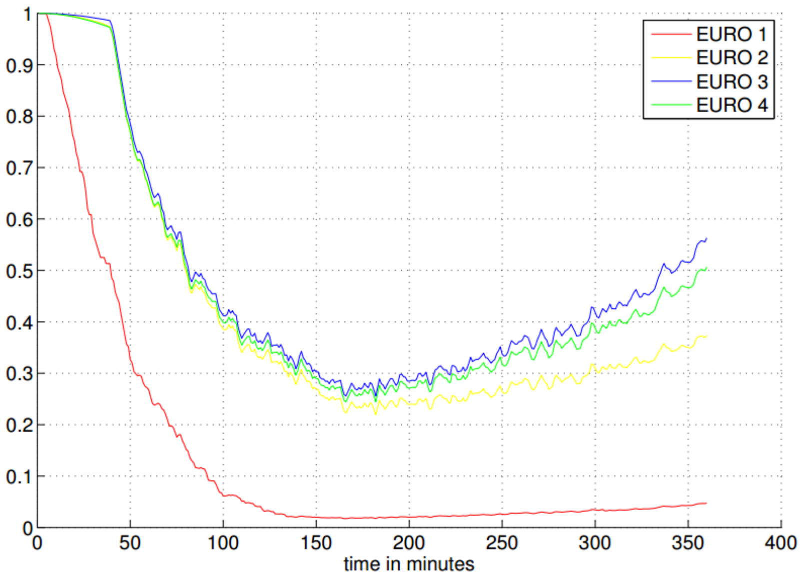

Figure 3.31.: Corresponding ensemble averages per Euro emission class: the average probability is different for each emission class.

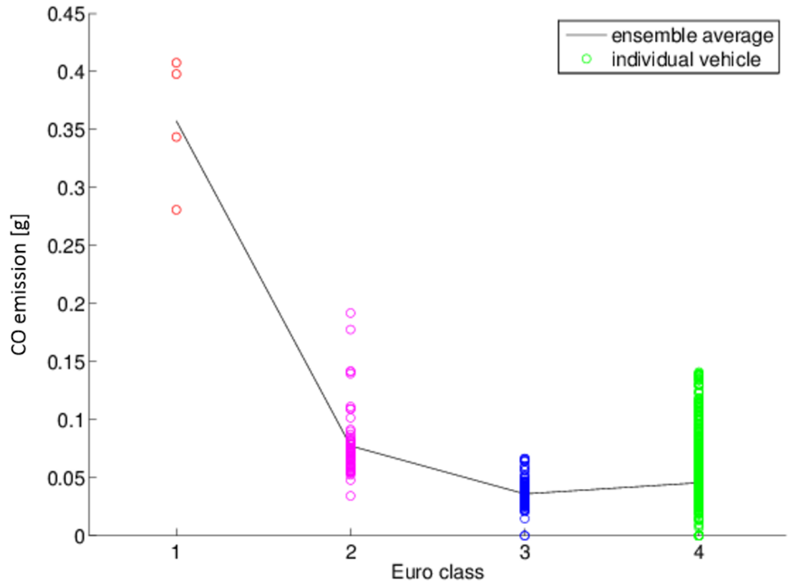

Figure 3.32.: Average emission per vehicle per second grouped into emission classes without any control.

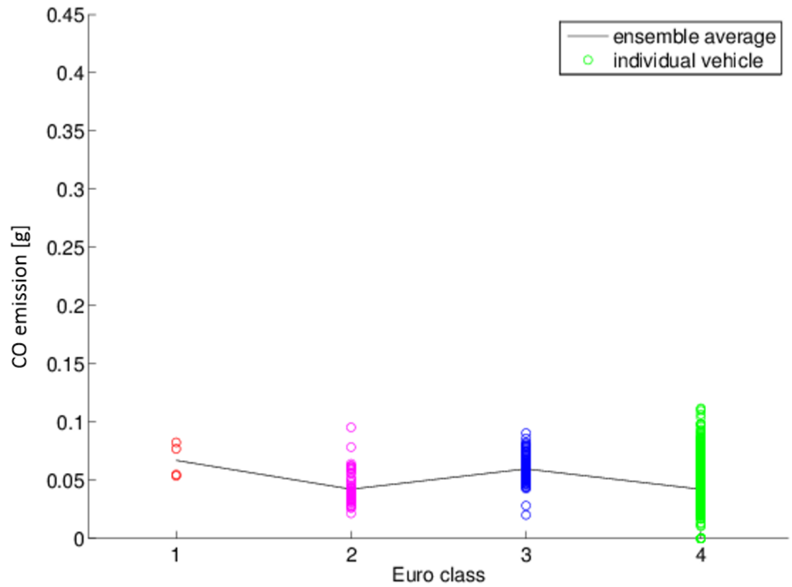

Figure 3.33.: Average emission per vehicle grouped into emission classes when the fairness extension is applied.

3.9.3. Case study: Integral control of vehicle emission implementing utility fairness

The basic idea of the approaches described in the previous section has been, that the vehicles of the hybrid fleet toss a coin according to the probabilities it has received from the infrastructure. In order to implement fairness it was proposed to include a factor which adapts the broadcasted probability considering the emission class of the vehicles or other vehicle-specific characteristics. This factor was vehicle-specific. However, finding this factor is very complicated and perfect fairness - in the sense that all vehicles are allowed to emit exactly the same - is almost impossible using this approach. In this extension, a perfectly fair implementation is introduced by referring to the concept of utility (again, the vehicle's utility is equivalent to its emission).

In contrast to the probability that defines whether vehicles should drive in EV or ICE mode, the utility (emission) is broadcasted here directly. As it is assumed, that the controlling infrastructure does not have any information about number of vehicles, an (integral) controller is employed to find the proper value for the value to be broadcasted y^*.

It is assumed, that vehicles know about their own (private) characteristics and could - based on the received emission value y^* - deduce their individual probability to drive in EV or ICE mode, respectively. The broadcasted value y^* corresponds to the expected value $E(y_i) = w_i \cdot y_i^{ICE}$ with the probability p_i to drive in ICE mode. The emission of a vehicle running in ICE mode is assumed to be y_i^{ICE}. It was said earlier, that y_i^{ICE} is known to the individual vehicle i. Thus, vehicles could easily calculate p_i.

The simulation of the approach is applied to the following scenario: it is based on a road network imported from OSM (around Rosenthaler Platz in Berlin, see Figure 3.34). The distribution of vehicle density is illustrated in Figure 3.35.

Figure 3.36 shows the evolution of average emissions of some randomly selected vehicles. Independent from the emission class or velocity of the selected vehicles, it could be seen that the average emission of vehicles converges. In this sense the approach leads to a perfectly fair result.

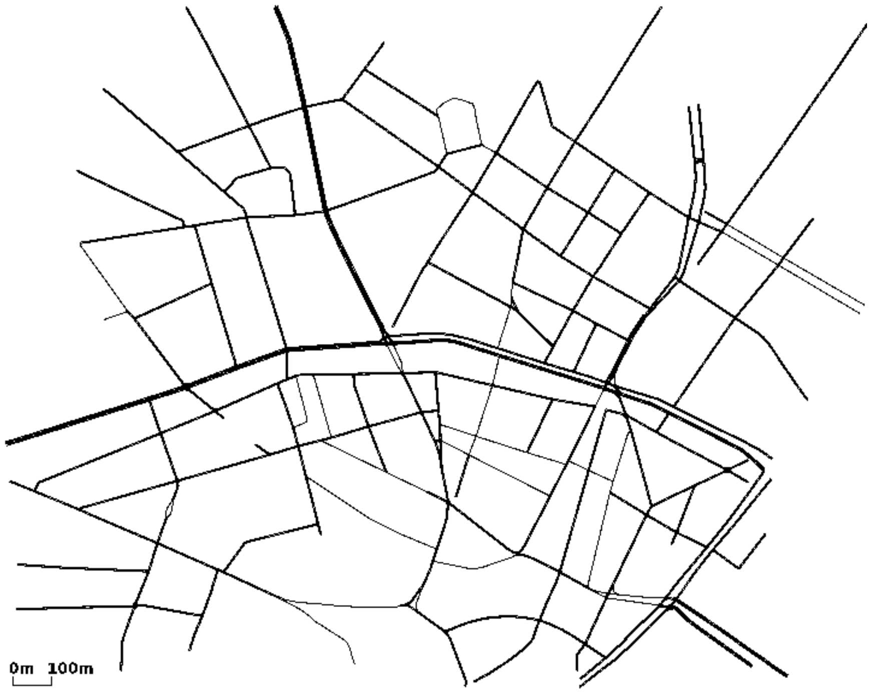

Figure 3.34.: Imported road network around Rosenthaler Platz in Berlin.

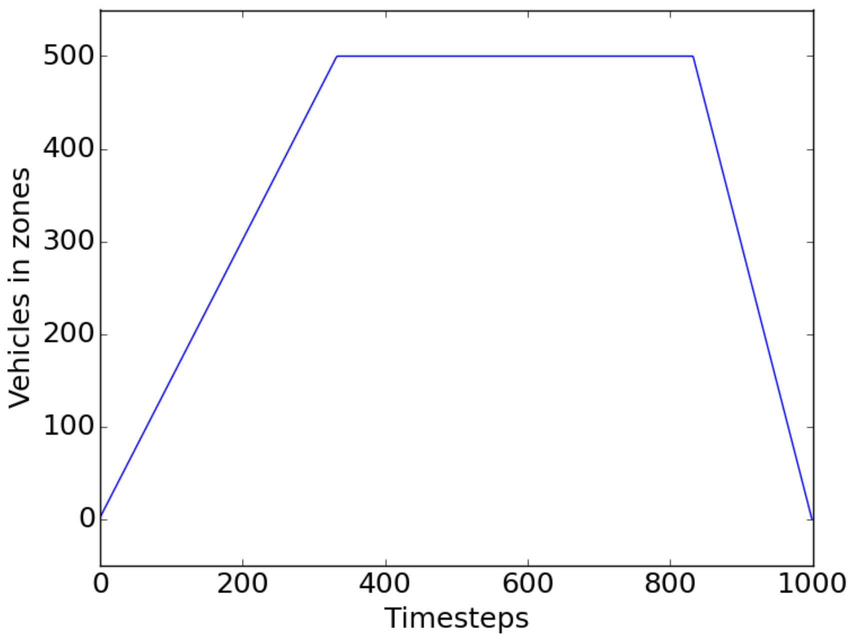

Figure 3.35.: Number of vehicles in the simulation over time.

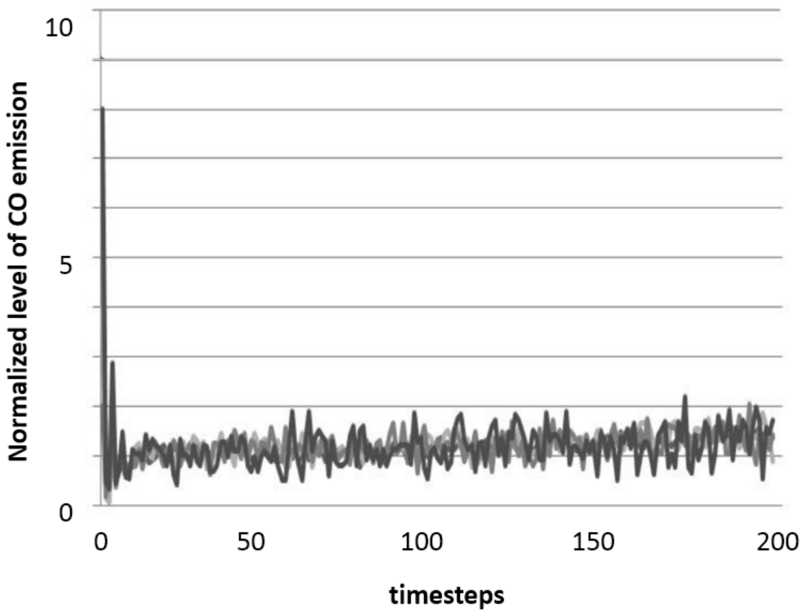

Figure 3.36.: The average emissions of vehicles converge - independent from their individual emission class or velocity profile.

4. Integrating the electric grid while driving PHEVs

Summary: *This chapter introduces the second cyber-physical ITS application. It is shown how to integrate a large number of vehicles with an electric grid, which suffers from unbalanced energy supply and demand.*
Central parts of the work are results of joint work with Emanuele Crisostomi, Yingqi Gu, Ilja Radusch and Robert Shorten. It has been published in [73, 77].

4.1. Renewable energy sources and electric vehicles in the smart grid

The smart grid is a hot topic in research and industry in recent years. The need for a smarter electric grid is mainly reasoned by the increasing fluctuation of energy supply caused by the ongoing integration of renewable energy sources - its share increased by 81.3 % between 2002 and 2012 [48] in Europe. In Germany about 25 % of energy has been generated from renewables in 2013. Until 2020, the share of wind energy shall exceed 20 % of the overall energy supply [146]. In Denmark, wind provides more than 30 % already today and it is foreseen to increase the share by more than 50 % of the overall demand by the year 2020 [113].

One of the most promising answers to the fluctuation of supply is the temporal storage of electric energy. Many researchers have proposed to use batteries of electric vehicles for that. Even though the number of electric vehicles is still low in general, there are countries like Norway where more than 10 % of sold cars are electric. Germany aims at having one million electric vehicles (including PHEVs and range extenders) deployed until 2020.

The main idea is to integrate the vehicles' batteries with the electric grid in order to damp electricity over- or undersupply. Roughly speaking, energy generated from renewable sources is stored when availability exceeds the energy demand and released again when needed. Amongst the available storage systems, the ability of EVs to act as an *aggregated* virtual battery for such purposes is of the most important arguments in favour of EV adoption. One of the main benefits of such a solution is that it would be automatically available without requirements of big investment in extra storage devices [105], [173] and [166]. All these proposed solutions wish to design the power supply in a smarter way.

Next to the energy supply aspects, one may also look at the demand side. Demand side management is a key pillar in the design of smart grids [108]. It could be divided in two basic principles: energy efficiency and conservation on the one hand and demand response programs or load shifting on the other [37]. In the context of this thesis the latter is of interest. It aims at shifting the demand from peak periods to off-peak times. McKinsey has found six aspects for implementing response programs: rates, incentives, access to information, utility controls, education or marketing respectively and customer insight and verification [37]. Load-adaptive pricing is the central idea and probably the most efficient instrument proposed. In parallel, it is (technology-wise) relatively easy to implement. However, pricing strategies might have critical socio-economic impacts. Moreover, it must be avoided that the new pricing schemes lead to new peaks in traffic demand. Finally, it requires utility controls (e.g. smart control of air conditioning) based on signalling from the grid operator comes mainly with technological challenges.

The advent of electric vehicles (and plug-in hybrid vehicles) does not only provide extra capacity that could be used for electricity buffering but it will also cause a huge demand with great variability. The most obvious way to adapt demand is to utilise energy when it is available. Clearly, in the case of EVs this is hardly possible as drivers would not accept electric vehicles which could only be charged during specific times. On the contrary, the grid should always serve the mobility needs of users as a primary objective.

To the best of my knowledge, all proposed approaches which integrate electric vehicles (and their batteries) with the smart grid so far target the charging process. Available approaches are valid; however in the case of PHEVs, it is also possible to control the vehicle while it is driving. At every instant of time vehicle owners have the choice whether to utilize the EV or ICE engine (or both in case of mild hybrids). This is the basic starting point for the *SPONGE* approach.

The *SPONGE* application exploits the flexibility of PHEVs to run in EV or ICE mode again. By doing this in a smart way, it couples the needs of the grid with those of the vehicle owner.

Different to the *TwinLin* application (presented in the previous Chapter 3), where a feedback loop has been set up around vehicles within a geographical area to control a local pollution, *SPONGE* connects vehicles in a temporal dimension: vehicles' behaviour is controlled over a longer time window (such as a day) with the effect that batteries are filled as required by the grid at a particular time.

Specifically, using energy forecasting models, it is possible for PHEVs to act as an energy sponge. They capture renewable energy as it becomes available and may release it again when it is required by the grid. To do this, weather forecast services can be used to make predictions of how much energy will be available from solar or wind power plants in the near future (e.g., next 24 hours). Based on these forecasts, cooperative strategies can be implemented to make space in a fleet of vehicles (in the batteries) for forthcoming energy. Essentially, vehicle owners allow the EMU (engine management unit) of the vehicle to be orchestrated (in a manner that is transparent to users) by a service that takes into account the electric grid's needs. By doing this, users fully utilise the clean energy as it becomes available and prevent clean energy from being wasted by ensuring that there is always enough capacity available for storage. Furthermore, they help balance energy supply and demand through active scheduling of energy sourcing for vehicles in a pro-active manner. It is also shown that this strategy has the potential to reduce the complexity burden of charging.

After outlining some basic limitations of given V2G approaches (Section 4.1) and a formal description of the problem statement (Section 4.2), three possible set-ups are outlined (Sections 4.3 - 4.5). A critical discussion of the approach is included later in this thesis in Section 10.3.

4.2. Problem statement and assumptions

For convenience the following set of assumptions are made. Each day k is separated into two time periods; a time period when vehicles are charging, and a time period when vehicles are not charging (and while they are possibly in transit). Furthermore:

1. It is assumed that a group of N PHEVs are participating in the application. That means that they might absorb available energy or may provide energy to the grid alternatively. As the overall discussion goes along the same lines. Nonetheless, the vehicle charging case is targeted if not stated otherwise.

2. Each participating vehicle (PHEV) is assumed to be able to run in EV-only or ICE-only mode (full hybrid vehicles). However, mild hybrid - which are able to

run both engines in parallel - could be utilized as well. In the remainder of this work, full hybrids are addressed if not stated otherwise.

3. A reliable day-ahead forecast of available renewable energy is available. This available energy to be provided by the grid is denoted by $w_{grid}(k)$. Although, the future horizon of optimisation can be longer than one day, weather forecasts might not be reliable enough to support optimal decisions over longer time periods, see [81, 191].

4. Vehicles can report their energy consumption over some period to a central entity.

In the following three different methods are introduced: *SPONGE* (Section 4.3), *Exact SPONGE* (Section 4.4) and *Optimized SPONGE* (Section 4.5).

4.3. Smart procurement of energy: SPONGE

In the *SPONGE* case, the objective is to ensure that

$$\sum_{i=1}^{N} D_i(k) \geq w_{grid}(k+1). \tag{4.1}$$

The electric energy dissipated by the i'th vehicle is denoted by $D_i(k)$. During the k'th day, the fleet acts like a sponge and shall make available at least enough space to absorb the available energy that is expected during the next charging period.

A central entity computes the desired electrical energy consumption and broadcasts some signal which is received by participating vehicles. Based on this signal, they are able to orchestrate the switching between EV and ICE mode, such that the regulation constraint is satisfied. The signal can be the probability to travel in EV mode rather than in ICE mode, or - in the case of mild hybrid vehicles - it can be the proportion of the traction torque that should be provided by the EV engine rather than from the ICE engine. The problem expressed by Equation (4.1) is called the basic *SPONGE* problem.

4.4. Smart procurement of energy: Exact SPONGE

The objective in the *Exact SPONGE* case is to make PHEVs travel in EV mode and deplete their batteries until the freed capacity match the expected energy that will be available from renewable sources exactly. This problem is denoted as *Exact SPONGE*. Its mathematical formulation is as follows:

$$\sum_{i=1}^{N} D_i(k) = w_{grid}(k+1). \tag{4.2}$$

The main advantage of the *Exact SPONGE* approach is that when the fleet of vehicles connects to the grid for recharging, there is no need for any charging control: every vehicle just charges until their batteries are full. The complete available energy will then be utilized, while all vehicles' batteries are fully charged.

4.5. Optimised access: Optimal SPONGE

The *Optimal SPONGE* case responses to the situation where certain vehicles shall be prioritised in terms of access to the incoming energy $w_{grid}(k+1)$. This prioritisation could be done according to an individual utility function $f_i(D_i(k))$. In accordance with this assumption, the above problem can be reformulated in an optimisation framework as:

$$\begin{aligned} &\text{maximize} && \sum_{i=1}^{N} f_i(D_i(k)) \\ &\text{subject to} && \sum_{i=1}^{N} D_i(k) = w_{grid}(k+1). \end{aligned} \tag{4.3}$$

The problem is most interesting when the f_i's represent a notion of utility. Further, the information about the utility functions is considered to be private and not to be revealed to a central entity or to other vehicles. The objective is to solve the problem in a privacy preserving manner. Note that the f_i's may be incorporated to represent various use cases such as the following:

1. OEM's may partner with utilities to provide a service where the price of energy is part of PHEV's owners car purchase plans. Those paying more upfront, may have prioritised access as it becomes available.

2. The f_i's could represent the price paid by an individual vehicle owner for energy access.

3. The f_i's could be used to penalise vehicles with a lower load factor (fewer passengers).

4. The f_i's could be used to penalise vehicles that drive close to schools, hospitals, etc.

In effect this presents a solution where - without any further control of the charging process - the complete available energy is utilized while the utility of charging vehicles is maximal.

Note that in some cases the optimisation problem might not has a feasible solution. For instance, when there are no vehicles on the road. Obviously, PHEVs cannot deplete their batteries to free capacity for the forthcoming energy. In such cases where the problem does not have a solution, a "best-effort" solution is of interest, where the closest feasible solution is achieved instead, see for instance [164].

4.6. Methods

There are many ways in which the problems specified in the previous section may be solved. One simple method that can be adopted is described in the following. To this end it is assumed that over the period when vehicles are not charging, denoted by $\theta(k)$, the aggregate electrical energy consumption of the fleet is required to be $w(k,t)$ with

$$\int_{\theta(k)} w(k,t)dt = w_{grid}(k+1), \tag{4.4}$$

where t denotes time (while k indexes days relevant for energy providing forecasts, t denotes the time relevant for the control loop). It is assumed that each vehicle is synchronised with a clock (possibly a GPS clock), and reports its consumption over the τ'th clock period as $D_i(\tau_i)$ to a centralised authority. This centralised authority aggregates the energy consumption from all PHEVs and broadcasts a signal to the vehicles depending on whether the aggregated consumption exceeds $w(k,t)$ or not. The probability whether the i'th vehicle travels in fully electric mode in the $\tau_i + 1$'th period depends on this broadcasted signal. The following two application use cases (fair energy consumption, utility maximization) can be used to orchestrate the fleet behaviour in an intelligent way.

4.6.1. Use case 1: Fair energy consumption (SPONGE and Exact SPONGE)

The fair energy consumption use case refers to the scenario when all vehicles participate in the *SPONGE* program in the same manner; namely, that they have the same probability to travel in EV mode (or ICE mode respectively). In the (fair) *SPONGE* case a simple proportional controller can be used.

Algorithm 1 SPONGE Algorithm

if $\sum_i^N D_i(k) < w_g(k+1)$ **then**
  $p_i^{EV}(k) = g_1(w_g(k+1) - \sum_{i=1}^N D_i(k))$
else
  $p_i^{EV}(k)$ is free to choose.

In this case, at every interval of time (e.g., every minute), a vehicle travels in EV mode with a probability $p_i^{EV}(k)$ which is an increasing function $g_1(\cdot)$ (see Algorithm 1) of the gap between the desired target of energy $w_{grid}(k+1)$ and the currently available space in the vehicles $\sum_{i=1}^{N} D_i(k)$. Note that if enough space has been vacated, vehicles are allowed to travel in any way they desire. Also note that, as already anticipated, even if all $p_i^{EV}(k)$'s are set to 1, the goal might not be accomplished if not enough vehicles are travelling in the time interval of interest.

The *SPONGE* problem usually takes place on a day-scale: for instance, vehicles are scheduled to spend a given quantity of energy during the day, and are then recharged at night time, when idle. However, in a practical scenario, it is more convenient to match the energy over a number of time windows during the day. This has a number of benefits: if the enough space has been vacated after a few hours, then cars travelling in the afternoon would be excluded from the programme. On the other hand, if the matching problem is split up in several time windows, then every single car, travelling at any time is equally involved in the programme. Another advantage is that the matching problem could be adjusted in every new time window taking into account new weather forecasts and previous matches. Accordingly, k refers to a shorter time window - e.g. five minutes - in Algorithm 1.

The request that vehicles have to travel in EV mode with some probability can be implemented in practice either by making a share of the vehicles travel in EV mode, or by adjusting the traction provided by the EV and ICE engine.

As for the *Exact SPONGE* case illustrated in Section 4.4, then a simple PI-controller can be adopted in the following manner to implement a fair solution:

Algorithm 2 *Exact SPONGE* Algorithm.

if $\sum_{i=1}^{N} D_i(k) < w_{grid}(k+1)$ **then**
  $p_i^{EV}(k) = g_2(w_{grid}(k+1) - \sum_{i=1}^{N} D_i(k))$
else
  $p_i^{EV}(k) = 0$

Recall that the main difference to the previous case is that the control objective is to *exactly* deplete the batteries of quantity $w_{grid}(k+1)$, while vehicles are not allowed to over-deplete their batteries. Although such a solution might penalise drivers (i.e., they are forced to travel in ICE mode to avoid over-depleting their batteries), it is very convenient for the grid, as it is possible to predict exactly how much energy will have to be delivered to the fleet of vehicles. But there may be good reasons for drivers to preserve a store of electric energy - some aspects are outlined in Section 10.3.

Note that the proposed approaches can be used to tackle many practical scenarios of interest. For instance, it could be assumed that a company provides a free battery-charging service to the PHEVs of its employees whenever there is enough power generated from some connected solar/wind plants. Although such a scenario gives rise to fair solutions, still personal constraints of single employees are not taken into account. In this perspective, the scenario can be made more complicated as described in the following subsection.

4.6.2. Use case 2: Utility maximisation (Optimized SPONGE)

Different to aforementioned scenarios, the optimisation outlined here required different (but still aligned) behaviour of users - i.e. different probabilities to drive in ICE or EV mode should be computed for different users while taking personal constraints into account. For instance, the decision to travel in EV mode represents a potential cost to the owner that can be represented as an increasing function of D_i, i.e., of the energy spent to travel in EV mode. For example, the cost of travelling in EV mode can be formulated as the likelihood to not have electrical power when it is needed, e.g., because some areas might be accessible only in EV mode for pollution reasons (for example, the so-called *Umweltzonen* in Germany[1]). Note that a similar discussion can be made in terms of discomfort of travelling in ICE mode.

This scenario allows the central infrastructure to explicitly take into account personal needs of PHEVs' owners. These are represented by individual utility functions $f_{u,i}(\bar{D}_i) = a_i \cdot \bar{D}_i^{\,2} + b_i$. There are many ways to solve the mathematical problem that arises. One possible way is to obtain the solution by formulating the problems as regulation problems with constraints, and to use these constraints to solve optimisation problems as they arise. One possible solution using and adapted version of the AIMD algorithm [183] is the approach described in 3.

Algorithm 3 *Optimal SPONGE* Algorithm.

if $\sum_{i=1}^{N} D_i(k) < w_{grid}(k+1)$ **then**
    $p_i^{EV}(k+1) = min\left(1, p_{add} + p_i^{EV}(k)\right)$
else
    $p_i^{EV}(k+1) = \beta \cdot p_i^{EV}(k)$
    with probability $\lambda_i(\bar{D}_i) = \frac{f_u'(\bar{D})}{\bar{D}} = a \cdot \bar{D}_i^2$,
    otherwise
    $p_i^{EV}(k+1) = min\left(1, p_{add} + p_i^{EV}(k)\right)$
    $k = k+1$

4.7. Simulations

Simulation results to show the efficacy of the proposed idea are presented briefly in the following. The simulations are performed using the traffic simulator SUMO [12] and the given TRACI interface again.

[1] http://gis.uba.de/website/umweltzonen/umweltzonen.php

4.7.1. SPONGE and Exact SPONGE

A map of a rural area near Hamburg, Germany, was extracted from Open Street Map to be used as the underlying street network (see Figure 4.1). The simulation runs for 1000 seconds (about 17 minutes) with about 600 PHEVs. A great share of vehicles start from random locations but share the same destination. This refers to the case when employees go to work using their PHEV vehicles. In such a setup the infrastructure (the employer) regulates the driving mode in order to meet the target of energy that will be available at the workplace to recharge the vehicles.

It is assumed that vehicles travel in ICE mode when they enter the simulation and that they are free to choose in which mode (ICE or EV) they wish to run when no indications from the infrastructure are received. Such indications make vehicles choose in a probabilistic manner to run in EV mode such that space in the batteries is freed. Once the target energy is matched, the infrastructure does not broadcast information any more and vehicles are allowed to travel in the mode they prefer. In this simulation it is assumed that vehicles choose - with equal probability - to run in EV or ICE mode. Figures 4.2- 4.5 show the simulation results for the first two algo-

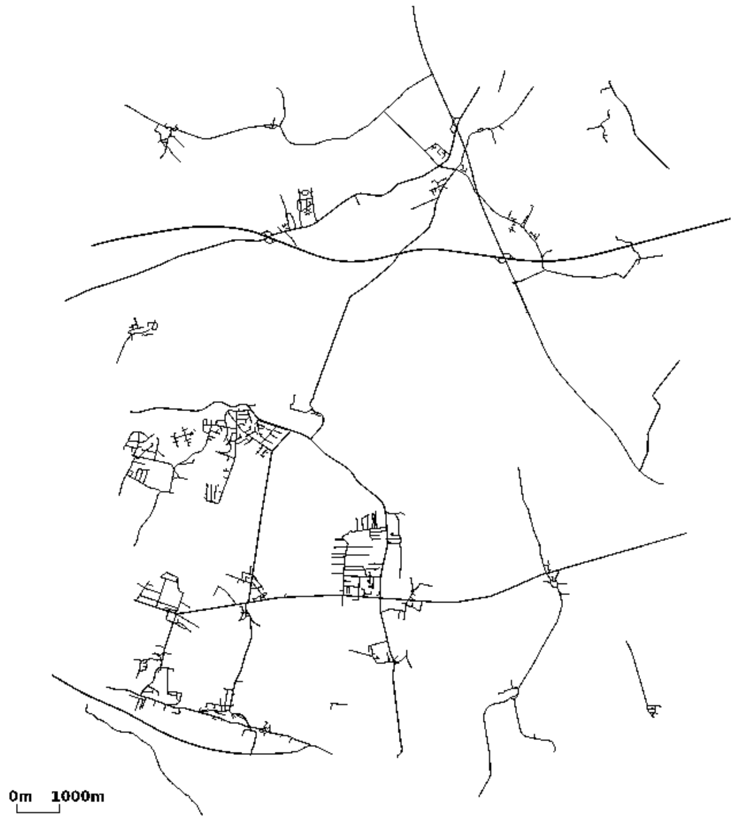

Figure 4.1.: Road network from an area near Hamburg, Germany imported from Open Street Map.

rithms (Algorithm 1 and Algorithm 2). There are four time windows of 250 seconds

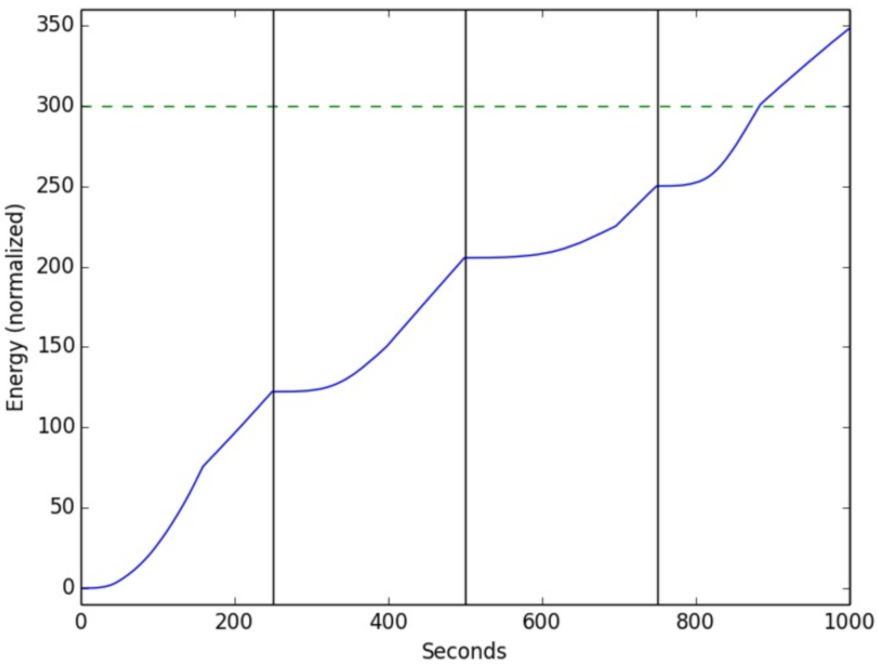

Figure 4.2.: The overall constraint is met after about 900 seconds.

each. It is assumed that vehicles should (at least) free a capacity of a scalar value, here 300 (some normalized value, visualized by the horizontal line in Figure 4.2).

The fleet of vehicles reach a situation to meet the infrastructure need before the end of the simulation (after about 900 seconds). When reaching this point, vehicles are free to choose their mode. That is the reason why the capacity increases and overachieves the infrastructure need, see Figure 4.2). Figure 4.3 shows the number of vehicles in the experiment running in ICE mode (blue, dotted line) and EV mode (green line).

Figure 4.4 shows how the capacity develops when the second approach is applied. Different to the previous simulation the approach guarantees that the grid's need is not overachieved but exactly matched.

Figure 4.5 shows again the according number of vehicles running in EV or ICE mode.

Note that the same pattern is repeated in every time window, until at the end of the simulation the relevant overall target is met. The advantage of dividing the 1,000 seconds into a number of smaller time windows of 250 seconds, is that both the vehicles that start their journey at the beginning and those at the end of the entire time frame, participate in the *SPONGE* programme.

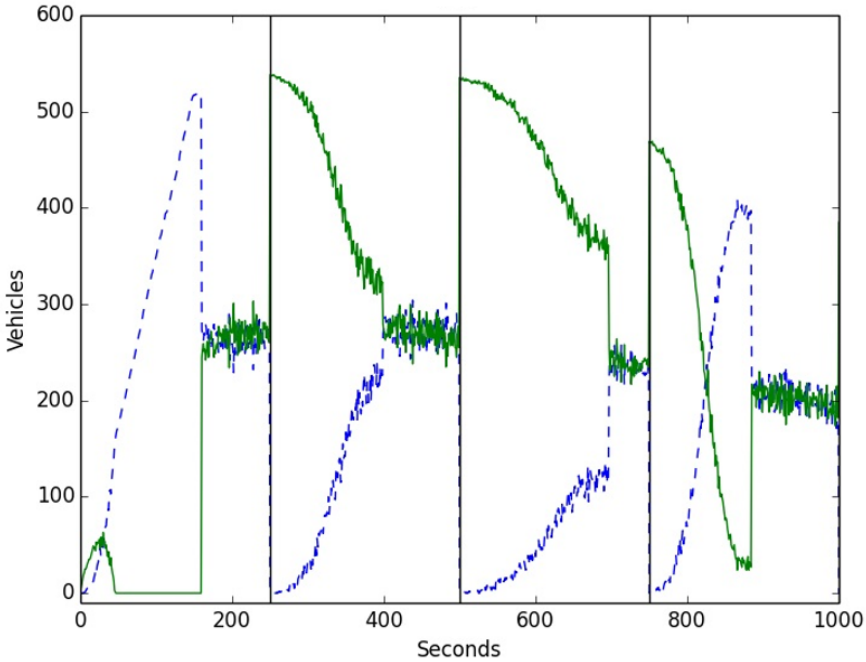

Figure 4.3.: Number of vehicles running in ICE mode (blue, dotted line) and EV mode (solid green line).

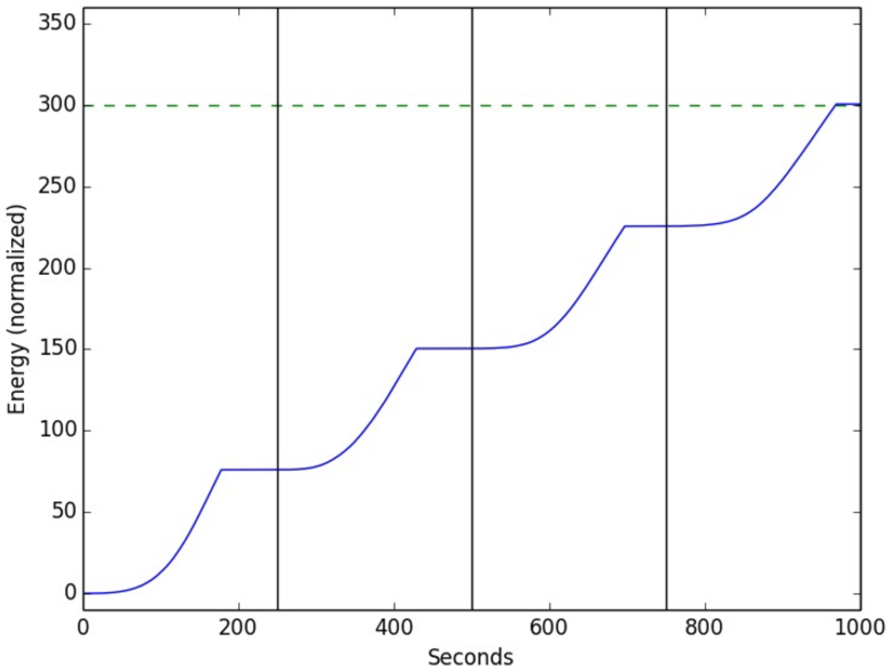

Figure 4.4.: The infrastructure need is exactly met.

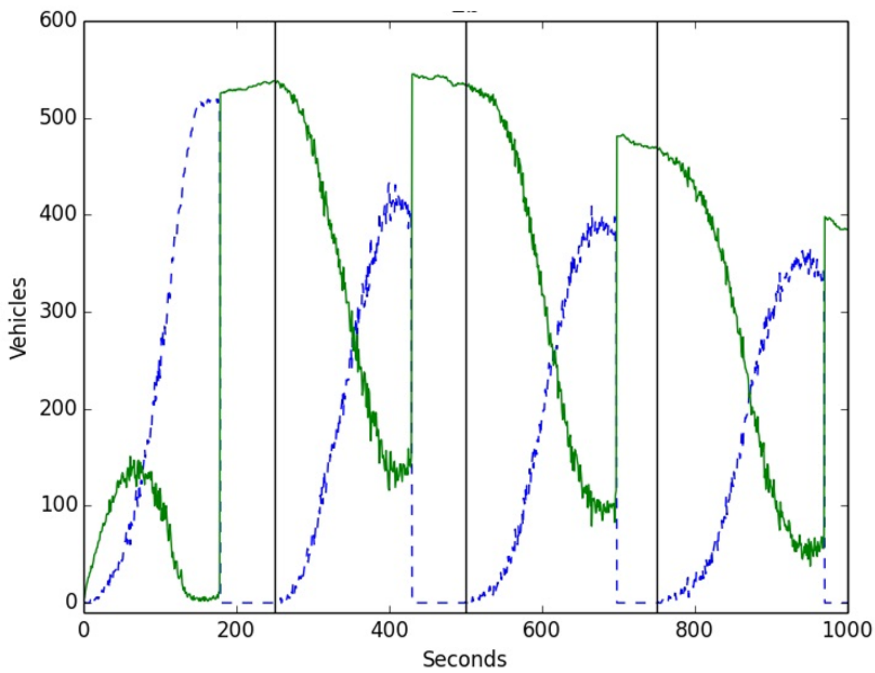

Figure 4.5.: Number of vehicles running in ICE mode (blue, dotted line) and EV mode when applying the *Exact SPONGE* approach.

4.7.2. Optimal SPONGE

The simulation results of the third approach (utility optimisation, see Section 4.5) are shown in Figures 4.6-4.9. The exact equality between vacated capacity in vehicles' batteries and expected forthcoming energy (Figure 4.6) is achieved by assigning different probabilities to travel in EV mode to different vehicles - according to individual utility functions. It is assumed that the convenience of vehicles in travelling in EV mode could be described through a convex quadratic function - here $f_i(\bar{x}_i) = a_i\bar{x}_i{}^2$ with $\bar{x}_i$ being the share of time running in EV mode up to the current simulation step. Parameters a_i are different for every vehicle. They are chosen in a random fashion from the interval [0,1]. The evolution of the utility functions of some randomly selected vehicles is shown in Figure 4.7. The optimal solution of the Problem (4.3) can be obtained by solving a consensus problem on the derivatives of the utility functions (more mathematical details together with a convergence proof can be found in [183]). However, Figure 4.8 shows that the utility functions converge. Finally, Figure 4.9 shows that the optimal solution is obtained by giving a different probability to travel in EV mode to each vehicle. The AIMD algorithm lets the probability of each vehicle to travel in EV mode linearly increase until the constraint is matched (and where a congestion signal is broadcasted). At that point, some probabilities decrease (back-off) in a multiplicative fashion to keep satisfying the constraint. Vehicles participate to the back-off step with a probability that is proportional to the derivative of their own utility function divided by the argument of the utility function (i.e., $\propto f_i'(\bar{x}_i)/\bar{x}_i$).

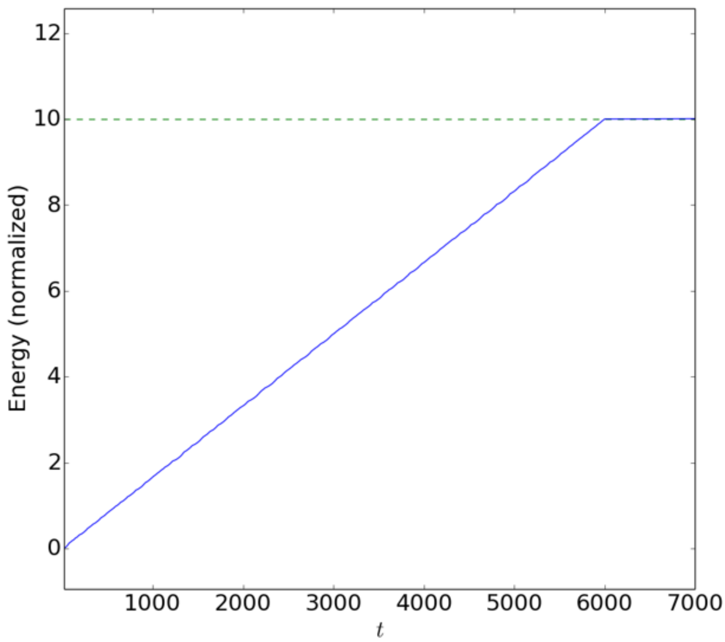

Figure 4.6.: Figure (a) shows that the overall constraint is again exactly satisfied at the end of the simulation. It is now achieved by maximising the sum of utility functions (values of the single utility functions are shown in Figure 4.7).

In this way, the optimal solution is obtained in a distributed way (i.e., without requiring V2V communication, or V2I communication). The only communication link goes from the infrastructure to the vehicles. The central entity broadcasts to all vehicles when the congestion event occurs. More details on the mathematical theory of the aforementioned algorithms, and more examples can be found in [183] and [33].

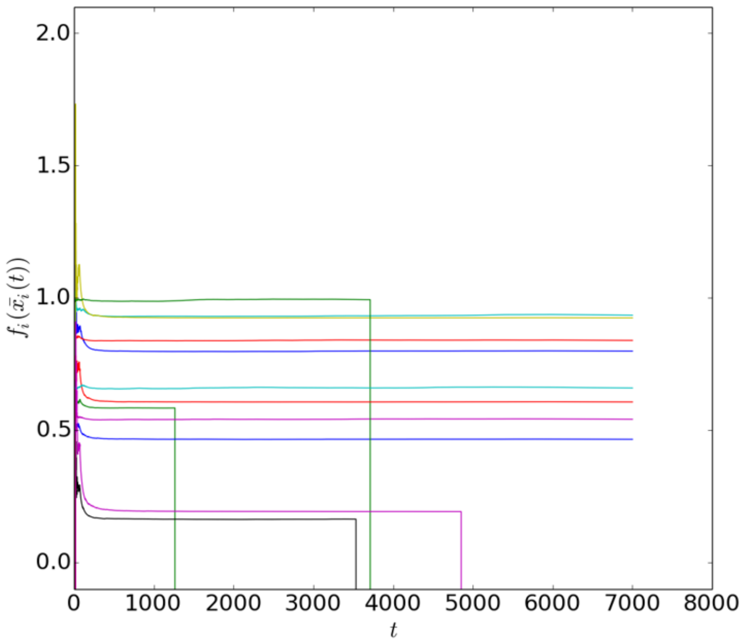

Figure 4.7.: Development of utility functions of some randomly selected vehicles.

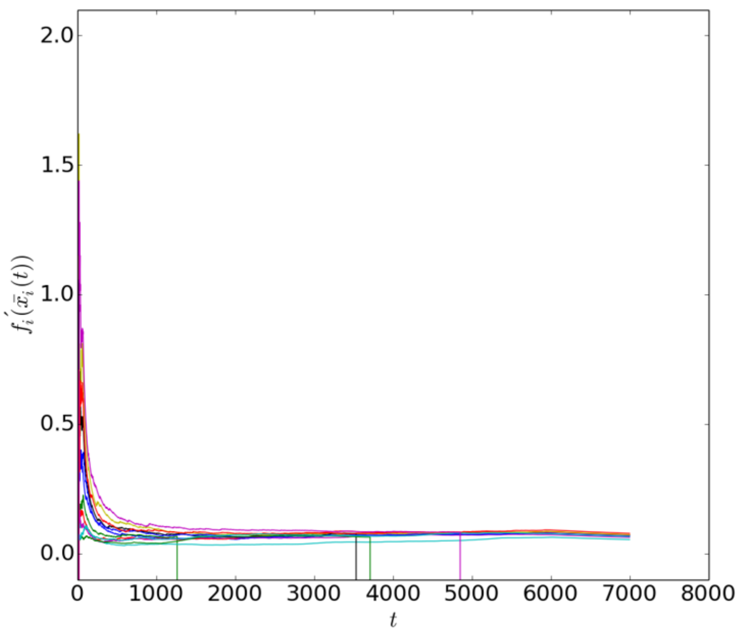

Figure 4.8.: The figure shows the utility derivatives of vehicles in the *Optimal SPONGE* case using AIMD. The optimal solution has been achieved as the derivatives of utilities converge.

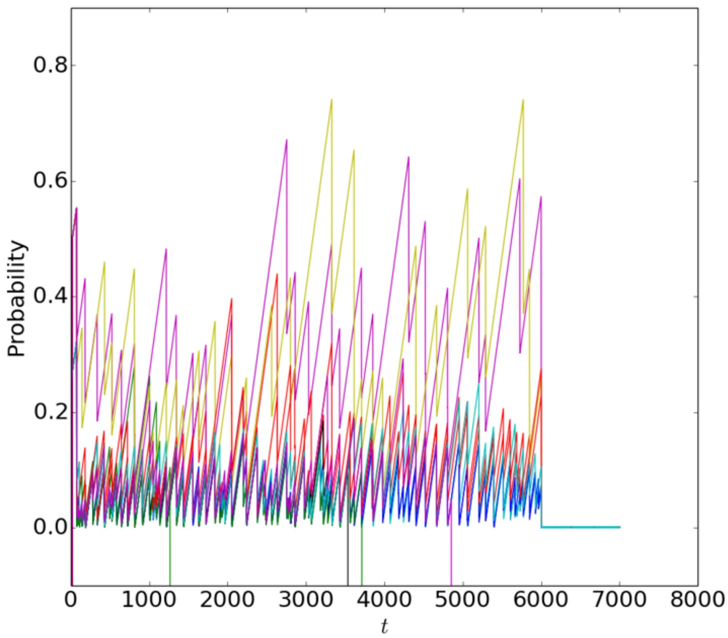

Figure 4.9.: The figure emphasises that the utility maximisation problem is solved by assigning different probabilities of travelling in EV mode to different vehicles, according to their utility functions. The probabilities of some randomly chosen vehicles are illustrated. Note that some vehicles are reaching their destination before the end of the simulation and exit the *SPONGE* programme in advance.

5. Optimized emission in a cap and trade system

Summary: *This chapter introduces an application, which is applied to a fleet of hybrid vehicles. It is assumed here that they aim at a common goal. That includes that vehicles accept to experience disadvantages in favour of advantages of other vehicles in the fleet. Such a set-up provides new extensions in regard to optimization.*
Furthermore, a method is introduced which exploits direct V2V communication, while aforementioned approaches in previous chapters were limited to communication links between a central infrastructure and vehicles (I2V communication). A second optimization method is introduced which is an extension of a an approach that was developed to implement fairness previously described in Section 3.9.1.
Central parts of the topic are result of a joint work with Mahsa Faizrahnemoon, Emanuele Crisostomi, Arieh Schlote, Ilja Radusch and Robert Shorten published in [76] and [75], which won the best scientific paper award at the European ITS Congress 2013. Section 5.3.2 refers to the work in the paper which was mainly driven by Mahsa Faizrahnemoon.

5.1. Objectives

The main objective of the TwinLin application introduced in Chapter 3 was to control air pollution in a specific geographic region by switching engines from ICE to EV (or vice versa). In contrast to that, the emission of a particular (closed) fleet should be controlled here - independent from their location. In parallel to the control objective, it is shown how emissions could be optimized at the same time.

A cap and trade emission system as it it introduced in the following could be interpreted as an extension of the *Kyoto Protocol* or the *Corporate Average Fuel Economy*, which is used in the US to provide incentives to vehicle manufactures to develop fuel-efficient models. In this sense it is assumed that fleet operators receive a budget (here a CO_2 emission budget). An exemplary allocation scheme may be based on the number of vehicles in the fleet.

The objective here is to present a method to regulate the aggregate effect (emission) of such a fleet of vehicles, such that the utility (benefit) to the vehicle owners (e.g. the fleet manager) is at all times maximized. While in this case the aggregate effect targets emission, the approach is presented in a general way and could be applied to many other objectives. In the application example, it is shown how a fleet of hybrid vehicles emits CO_2 respecting the given budget. At the same time, it is shown how to minimize the emission of a second pollutant - here CO. Accordingly, it is assumed that the fleet managers' utilities decrease with CO emission.

The basic assumption is that a fleet operator is faced with the choice of how to allocate the given budget between vehicles to maximize the utility. Clearly, there are many ways to achieve this, and this non-uniqueness creates an opportunity to regulate and optimize simultaneously [162, 163, 138].

Two basic algorithmic approaches for doing that are presented in this Chapter. Both are suited to be applied to the automotive domain as they respect specific constraints and requirements of ITS presented in 2.2.

5.2. Problem statement and assumptions

It is assumed that there is a fleet of N PHEVs. Each PHEV i has an individual utility function $f_{u,i}$. In the given context $f_{u,i}(c_i)$ defines the CO emission of a vehicle. There is the overall constraint relevant for the complete fleet, which says that the sum of (CO_2 emission) budgets c_i per vehicle should be equal to a given and static budget C.
This leads to the following optimization problem

$$min \sum_{i=1}^{N} f_{u,i}(c_i) \tag{5.1}$$

subject to

$$\sum_{i=1}^{N} c_i = C, \tag{5.2}$$

where the functions $f_{u,i}(c_i)$ map the scalar c_i to a utility scalar, and satisfy basic continuity assumptions. c_i is the pollution budget assigned to vehicle i and is - in case the pollution budget is fully utilized - equal to the emission of vehicle $y_{q,i}$. Note that in other contexts the optimization may be a maximization. The utility function $f_{u,i}(c_i)$ used here corresponds to the cost of the i'th vehicle using the allocation c_i.

Clearly, a rich repertoire of techniques exist to address with such problems and many of these have been successfully applied in other application domains [160]. With given reasonable assumptions on the $f_{u,i}(c_i)$ optimization approaches with a centralized set-up is always possible. Nevertheless, despite the availability of such techniques, solving problems that arise in an ITS context are not necessarily applicable due to a number of factors as outlined in Section 2.2.

Before proceeding with the description of how to find an optimal allocation of c_i, some important assumptions are listed in the following:

- The utility functions $f_{u,i}(c_i)$ are concave and increasing while the precise nature of the $f_{u,i}(c_i)$ may be unknown. However, individual vehicles have access to a measurement or a-posteriori calculation based on observable variables of $f_{u,i}(c_i)$ at each time step. That means that vehicles are have access to their own (private) utility information, i.e. by measuring their emission.
- All vehicles can communicate directly (or indirectly through a measurement) with a centralized infrastructure.
- The centralized infrastructure may broadcast information to the vehicles, but not communicate directly with each vehicle.

5.3. Algorithms

The previously introduced Equations (5.1) and (5.2) define a classical optimization problem. It could be solved with Lagrange multipliers. Let $H(c_1, \ldots, c_n, \lambda)$ be the Lagrangian associated with the problem

$$H(c_1, \ldots, c_n, \lambda) = \sum_{i=1}^{N} f_{u,i}(c_i) - \lambda(\sum_{i=1}^{N} c_i - C). \tag{5.3}$$

The solution to the maximization problem is given by finding the allocation so that

$$\lambda^* = \frac{\partial f_{u,i}(c_i^*)}{\partial c_i} \tag{5.4}$$
$$= g_i(c_i) \tag{5.5}$$

with $i = 1, \ldots, N$ subject to the linear constraint (Equation 5.2) being satisfied. The form of this equation suggests two approaches for solving the optimization problem.

First, it can be solved using feedback in combination with local on-vehicle computational power. Namely, by regulating the outputs of the vehicle functions $g_i(c_i)$ in such a way that they follow λ - with λ being broadcasted by a central infrastructure such that it becomes known to all vehicles. The task of the central infrastructure then becomes to find the appropriate multiplier λ^*. This can be solved by embedding the multiplier as part of a feedback loop as it is done in sub-gradient methods. This approach requires neither V2V nor I2V capabilities. It can deal with uncertainty in the utility functions, and requires almost no centralized computation.

A disadvantage of the approach is - like in any other centralized set-up - that it is not robust to failure of the centralized node. In addition to that, selecting proper control gains can be an issue.

The second approach is based on the fact that optimality is also reached when all vehicles achieve consensus regarding values of the functions $g_i(c_i) = \partial f_{u,i}(c_i)/\partial c_i$. This observation was first exploited by Stanojevic in [160] in the context of symmetric communication graphs. Later it was extended by Knorn for the asymmetric case and to the case where the outputs of all distributed nodes $g_i(c_i)$ are adjusted as part of a feedback loop [95]. The main benefits of these approaches are that no I2V capabilities are needed and that they can also deal with uncertainty in utility functions. In addition no centralized computation is required, and thus the technique is robust with respect to node failure.

In the following two subsections, the two approaches are briefly outlined.

5.3.1. Feedback control and local computation

The basic idea of the approach is to use an internal control to find the Lagrange multiplier λ^*. Accordingly, the central infrastructure uses the integral control update function (5.6)

$$\bar{\lambda}(t+1) = \bar{\lambda}(t) - \gamma(C - \sum_{i=1}^{N} c_i(t)), \tag{5.6}$$

where it is assumed that the central infrastructure either receives the information about the c_i from each vehicle (via V2I communication) and calculates or measures $\sum_{i=1}^{N} c_i$. It is assumed, that each vehicle receives the broadcasted $\overline{\lambda}(t)$ and updates its utility (implicitly) according to:

$$c_i(t+1) = c_i(t) - \epsilon_i(\bar{\lambda}(t) - \frac{\partial f_{u,i}(c_i(t))}{\partial c_i(t)}). \tag{5.7}$$

This corresponds to a sub-gradient method from optimization, see [25, 118]. The gains can be selected using methods from non-linear control theory. Figure 5.1 summarizes the control loop that is implemented by the application.

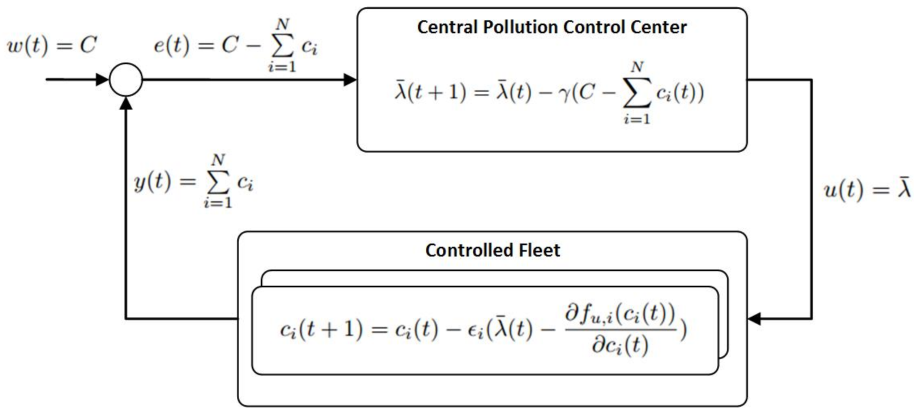

Figure 5.1.: Control loop in cap and trade system.

5.3.2. Implicit consensus with input

Stanojevic has shown in [160] that utility maximization problems can be formulated as a consensus problem leading to a completely decentralized solution. However, consensus based solutions require symmetric communication graphs to guarantee that it holds $\sum_{i=0}^{N} c_i(t) = C$ for all t. To overcome this problem, Knorn integrates a basic feedback error signal $C - \sum_{i=0}^{N} c_i(t)$, which is shared by the central infrastructure to yield a modified integral control of the following form [95]:

$$c_i(t+1) = c_i(t) + \epsilon_i(f'_{u,i}(t) - \frac{1}{n_i}\sum_{j} f'_{u,i}(t)) + \gamma_i(C - \sum_{q=1}^{N} c_q(t)) \;, \tag{5.8}$$

where j ranges over all neighbours of i and n_i is the number of neighbours j that can send information to vehicle i. $f'_{u,i}$ defines the derivative of $f_{u,i}$. In an ITS setup, it is a situation where each vehicle i exchanges the value of its utility derivative with its neighbours j (which are in communication range) and integrates the information to update the individual emission (budget).

The first term of Equation 5.8 achieves a consensus on the utility derivatives $f'_{u,i}(t)$. The second term ensures that the global constraint is satisfied. It exploits the broadcasted information from the central infrastructure which defines whether the global constraint is met. The stability and convergence properties of this algorithm, and the choice of ϵ_i and the γ_i, can again be determined using standard ideas from systems theory, see for example [96].

The contribution here is to present a modification of the algorithm to speed up convergence. For doing so at each time step, vehicles communicate the maximum and minimum values of $f'_{u,i}(c_i(t))$ from the set of derivative utilities known (from vehicle i and its neighbours) to the central infrastructure. The infrastructure then identifies the global maximum and minimum value. These two values are then broadcasted to all nodes (vehicles). They augment their neighbourhood set with these values and integrate in the consensus calculation (first part of Equation 5.8).

Clearly, it would also be possible to share the individual derivatives instead of the maximum and minimum values with the infrastructure. The hope here is to (1) have potential to decrease communication efforts and (2) preserve privacy.

This algorithm has the following two main features:

- As before other policies can be implemented by replacing the $f'_{u,i}(c_i(t))$'s in Equation 5.8 with an appropriate, alternative function. For example, if $f_{u,i}(c_i)$ would be used instead of partial derivatives, the policy implemented is a utility fairness policy.

- The burden placed on the infrastructure is slightly less than the first algorithm presented in the previous Section. The only task of the infrastructure is to broadcast a simple error signal (instead of updating Lagrangian multipliers).

5.4. Simulations

The following simulations illustrate the efficacy of both algorithms. In the first two simulations, a fleet of 20 hybrid vehicles have a total budget of CO_2 they are allowed to emit. The objective is to minimize the total quantity of emitted CO. Furthermore, as in [95] vehicles can regulate the quantity of emissions by deciding which fraction of the desired speed is supplied by the ICE or by the electric motor. This is consistent with a power split hybrid vehicle. Furthermore, emissions are calculated using a simple *average speed model*, with data taken from [23]. The scenario has an urban setting where vehicle speeds vary between 5 and 50 km/h. Then the amount of produced CO is related to the quantity of CO_2 according to:

- $CO = f(CO_2) = 2.3073 \cdot CO_2^2 + 12.25 \cdot CO_2 - 2.0531$
  for petrol Light Duty Vehicles (LDVs) with an ICE equivalent of emission class EURO1
- $CO = f(CO_2) = 0.8755 \cdot CO_2^2 + 2.1907 \cdot CO_2 - 0.24558$
  for petrol LDVs with an ICE equivalent of emission class EURO2
- $CO = f(CO_2) = 3.3209 \cdot CO_2^2 + 2.5657 \cdot CO_2 - 0.7812$
  for petrol LDVs with an ICE equivalent of emission class EURO3
- $CO = f(CO_2) = 0.53745 \cdot CO_2^2 + 1.3241 \cdot CO_2 - 0.1064$
  for petrol LDVs with an ICE equivalent of emission class EURO4

where $f(CO_2)$ corresponds to $f_u(c)$. The objective is thus

$$min \sum_{i=1}^{N} f_i(CO_2^{(i)}) \tag{5.9}$$

subject to

$$\sum_{i=1}^{N} CO_2^{(i)} = C, \tag{5.10}$$

where C is the overall CO_2 budget allocated to the fleet; $CO_2^{(i)} = c_i$ is the amount of CO_2 allocated to the i'th vehicle. Note that this problem is equivalent to the basic optimization problem described by Equations (5.1) and (5.2).

5.4.1. Utility minimization within a single class of vehicles

In the first simulation, a fleet of 20 vehicles belonging to EURO emission class 4, with slightly different parameters that are characteristic of the particular vehicle (e.g., brand, age). Figure 5.2 shows how the CO_2 budget is shared among the fleet of vehicles in order to minimize CO emissions. The parameters are $\gamma = 0.001$ and $\epsilon = 0.01$. Figure 5.3 shows that the fleet's overall budget is met.

5.4.2. Sensitivity of parameters

The following simulation shows the impact of different choices for γ and ϵ. The according simulation is performed with 2,000 vehicles. The objective is to show that if no information is communicated regarding the number of vehicles, the parameters which could be optimal with a smaller number of vehicles may give rise to undesired levels of emissions before the final goal is achieved. The convergence of the CO_2 emissions to the budget is depicted in Figure 5.4. It is easy to see, that the evolution of emission is completely different.

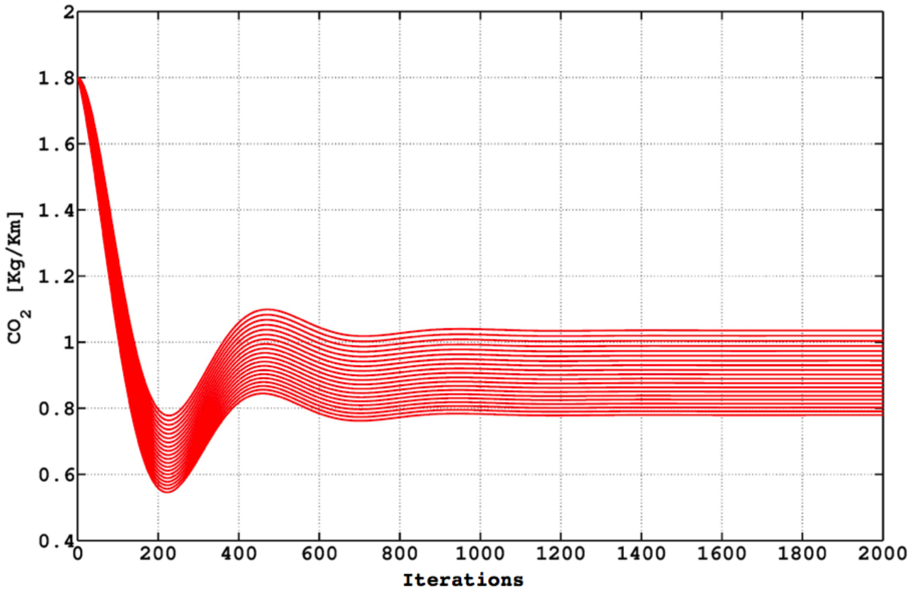

Figure 5.2.: The figure shows how the CO_2 budget is shared among the fleet of vehicles in order to minimize CO emissions. At the same time, the CO_2 budget is respected, as shown in Figure 5.3.

5.4.3. Utility minimization among different classes of vehicles

The third simulation embeds a fleet of 20 PHEVs again. It is assumed that the ICE in the vehicles belong to EURO 1, 2, 3 and 4 emission class. The objective is again to share a CO_2 budget among vehicles, while minimizing the overall emission of CO. As the emission of CO is very different from vehicles belonging to different classes, the approaches could not be applied without further adaptation. That is because the optimal solution would allocate unrealistic quantities of CO_2 to some vehicles, i.e. selected vehicles would not be able to travel at all while others would not be able to utilize their allocated budget.

For this reason, the proposed distributed algorithm is adapted such that the unknown quantities of interest (i.e., CO_2 budget) are forced to belong to the feasible sets for each vehicle. Note that the range of individually feasible CO_2 budget depends on the specific EURO class of the vehicle.

Figure 5.5 (a) illustrates how the same CO_2 budget is now shared among the different classes of vehicles, while Figure 5.5 (b) shows that the total level of CO_2 matches the desired one - and that the amount of CO is minimized at the same time.

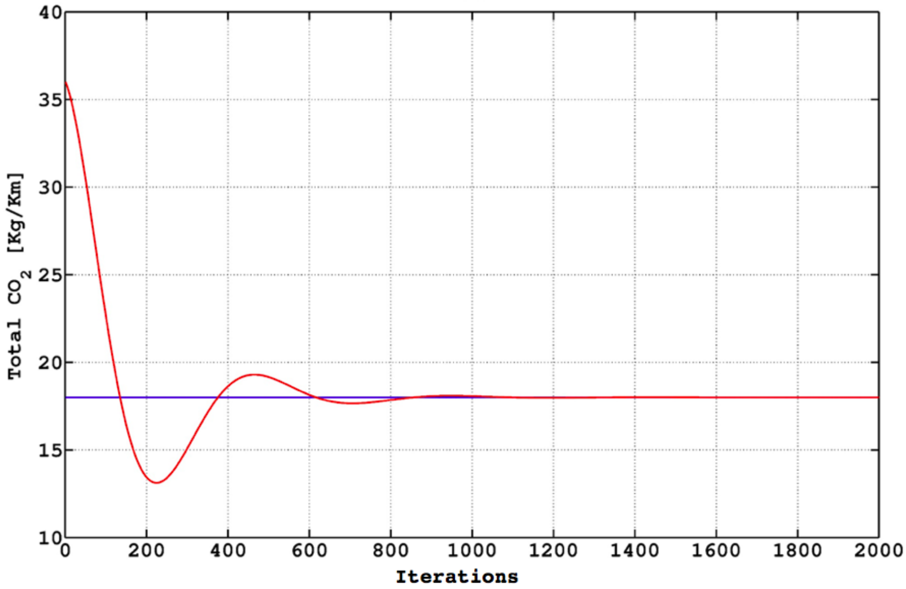

Figure 5.3.: While the CO_2 budget is shared among the fleet of vehicles in order to minimize CO emissions (Figure 5.2) this Figure shows how the CO_2 budget is respected.

5.4.4. Utility fairness

The objective of the previous simulations was to minimize the production of CO while keeping the total amount of CO_2 equal to some predefined threshold.

However, a key aspect regards the fairness of the previous solution. Note that while utility is maximized in these simulations, the solution suggested by the optimization may be very unfair. In some situations it may be of interest to make all vehicles produce an equal amount of pollution. This is easily accommodated in the framework.

This is shown in the next simulation, where vehicles produce the same quantity of CO, while still respecting the total CO_2 budget. While such a solution is not overall optimal for the environment (i.e., more CO than before is overall produced), it provides a solution that is fairer for vehicle owners. In the following, the two previously introduced algorithms are discussed. For simplicity, it is assumed that the utility function of node i is given by an (increasing and concave) function of the form $f_{u,i} = \alpha_i \log c_i$. The choice of this function is motivated by dispatch type problems where economic utility is an increasing function of allocated emissions. Recall, for such applications, it is $g_i(c_i(t)) = f_{u,i}(c_i)$. Figure 5.7 (a) depicts the performance of the algorithm *Feedback control and local computation* and Figure 5.7 (b) shows the unmodified algorithm *Implicit consensus with input* (b), and the performance of its modification Algorithm 2 (c).

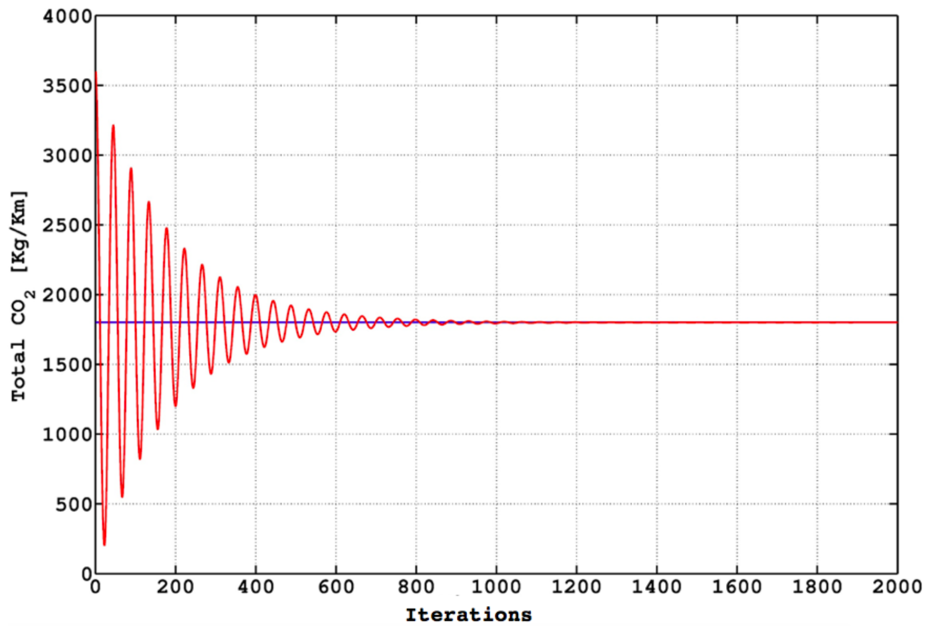

Figure 5.4.: Values of the optimal parameters depend on the number of vehicles, as wrong tuning of parameters may give rise to unsatisfactory responses.

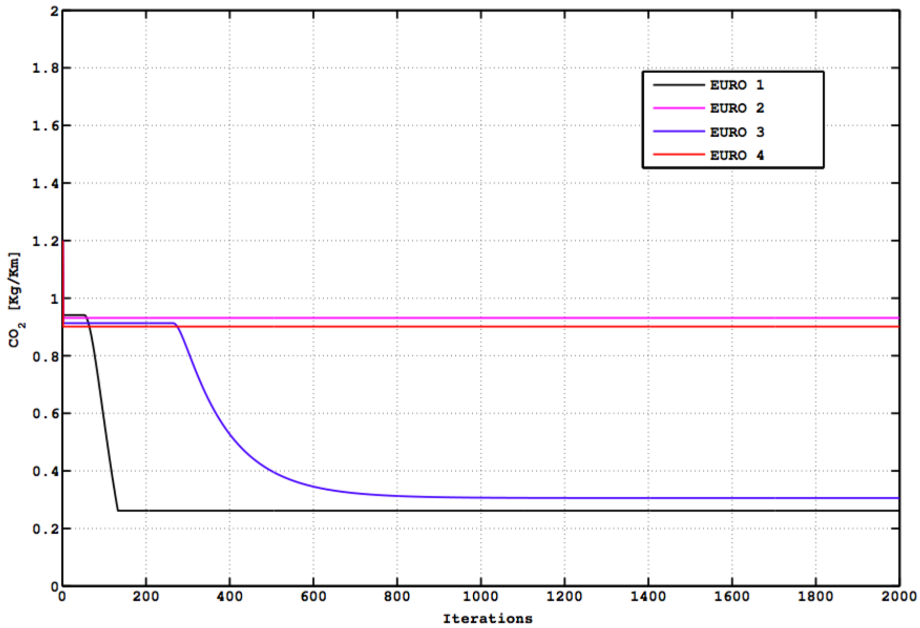

(a)

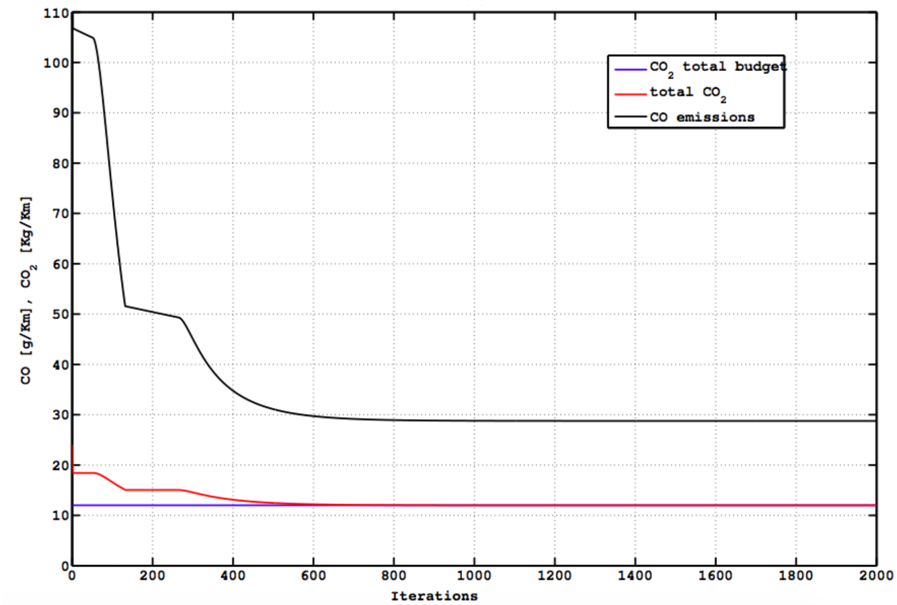

(b)

Figure 5.5.: Figure (a) shows how the CO_2 budget is shared among the fleet of vehicles of different EURO emission classes. At the same time, the CO emission is minimized, as shown in Figure (b).

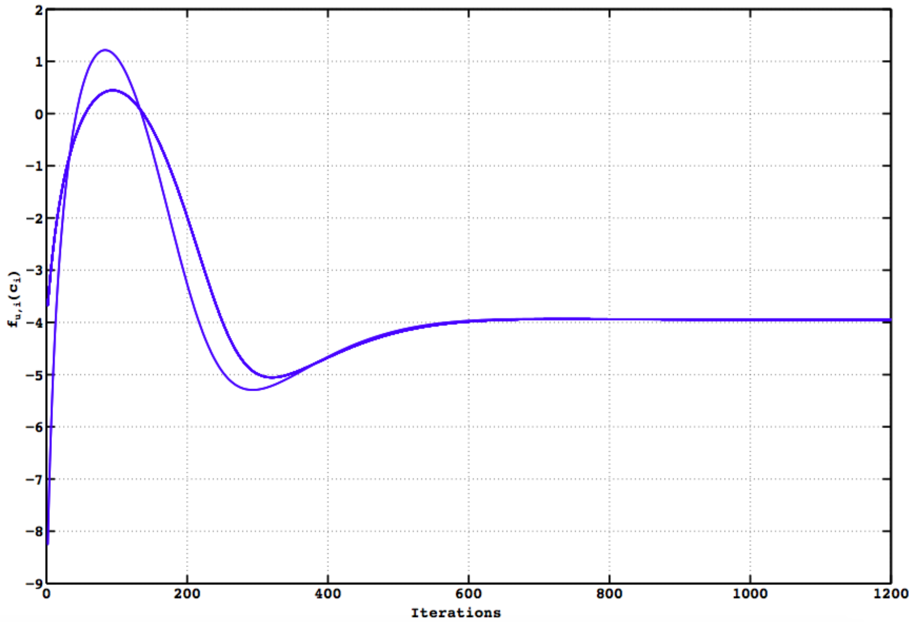

Figure 5.6.: The CO_2 budget is shared in such a way to equalize the production of CO per vehicle. The speed of convergence of the approach *Feedback control and local computation* is depicted. It is slow compared to the alternative approaches, shown in 5.7.

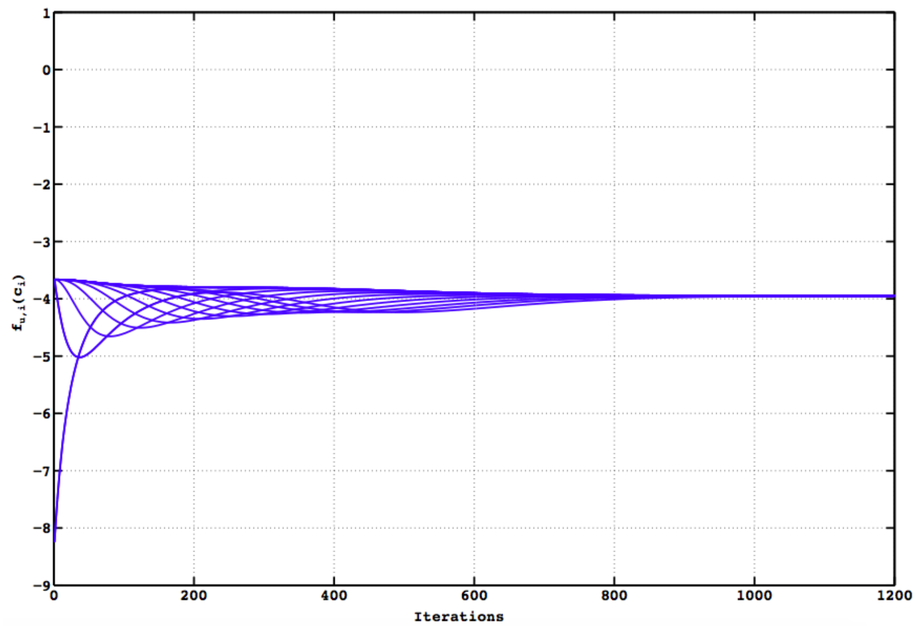

(a)

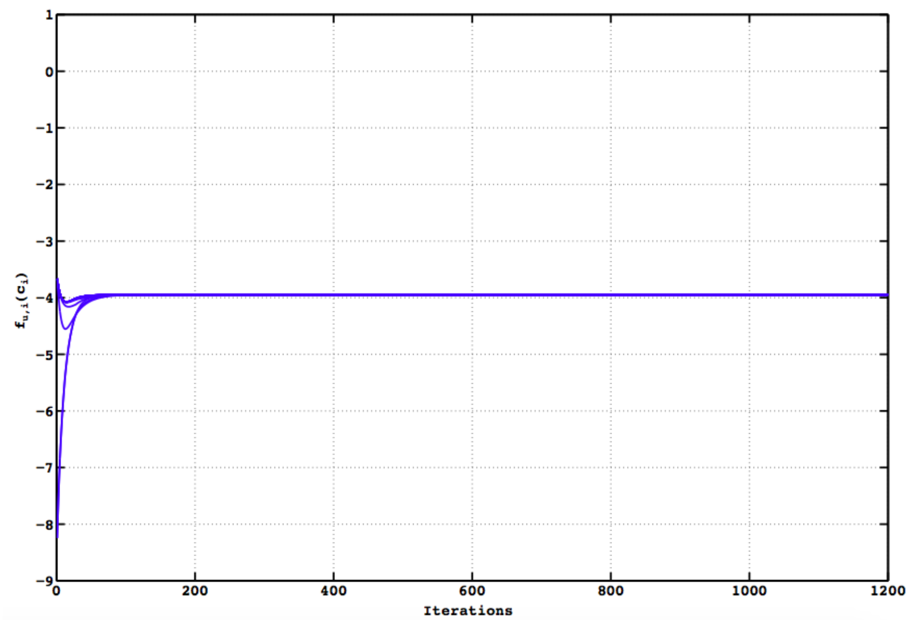

(b)

Figure 5.7.: Compared to the performance of the *Feedback control and local computation* algorithm shown in 5.6, the speed of convergence of the unmodified algorithm *Implicit consensus with input* shown in (a) and its modification (b) is much faster. In (a) vehicles communicate with one vehicle in the vicinity; in (b) they can communicate with one nearby vehicle and two additional randomly chosen vehicles.

Part II.

Balancing and de-synchronization

6. Emission-aware and balanced routing

Summary: *There are many approaches to address system-wide features by vehicle routing. Examples for such are traffic congestion, pollution, total travel time etc. To the best of my knowledge, state-of-the-art solutions do not consider realistic requirements, such as securing user privacy or their freedom of choice. The routing applications introduced in this chapter considers these aspects.*
The basic approach is a collaborative routing strategy under feedback to control air pollution. This approach is extended by a stochastic routing approach which balances features of the network.
The development of the routing applications is result of joint work with many authors. In particular, these are Stefan Lobach in regard to the pollution-aware routing application (Section 6.3) and Rodrigo Ordóñez-Hurtado, Wynita Griggs, Ilja Radusch and Robert Shorten in regard to the stochastic extension of the routing application (6.4, published in [78]).

6.1. Routing under feedback

The air pollution problem caused by road traffic has been outlined earlier (Section 1.1.3) and the TwinLin application (Chapter 3) was one solution to address the problem employing the power of PHEVs.

An alternative instrument does not require PHEVs but could be applied to any vehicle type: emission-aware routing. To implement that, the same architecture is exploited

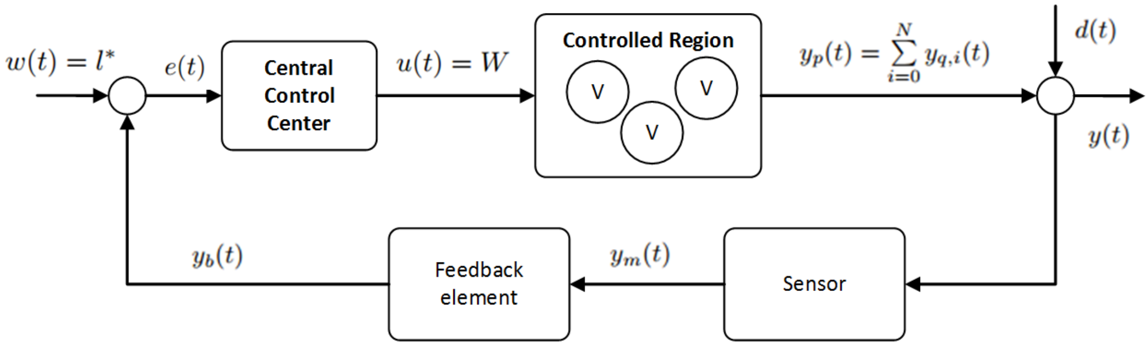

Figure 6.1.: Control loop for emission aware routing application.

as before: A feedback loop is constructed again around a central infrastructure, which is capable to broadcast information and participating, networked vehicles. Figure 6.1 is illustrating it. The control objective is that pollution should not exceed some predefined threshold l^*. The central infrastructure (control centre) broadcasts the costs $w_j \in W$ for each edge j in the network to all vehicles. This information is considered by a routing algorithm, which runs de-centrally at vehicles. $y_p = \sum_{i=1}^{N} y_{q,i}$ defines the emission of all vehicles. It is assumed again that the central authority is capable to measure the air pollution y_m (where $y_m = y_p$, if no other source of emission is considered, $d(t) = 0$) and may deduce information about vehicles' emissions. The feedback element applies a filter or some other preprocessing step to the measurement, such that y_b is fed back to the control algorithm.

6.2. Requirements and according edge cost function

Besides the general requirements targeting all proposed ITS applications (outlined in Section 2.2), routing applications have extra requirements that need to be considered. In contrast to previous applications, where switching the powering mode of a PHEVs is the main instrument, it could not be assumed that drivers are indifferent regarding different routes. That is because they vary in length and duration and thus are not equally attractive to users. Accordingly, the following considerations need to be kept in mind:

- Proposed routes should not be much longer than the shortest available route such that routes are acceptable to drivers.
- Routes should be allocated to drivers in a fair way in the sense that clean vehicles (i.e. vehicles with low emission) are not penalized. However, system-optimal and emission-aware routing do penalize clean vehicles on a general level. That is because - from a system-centric point of view - one wishes to have dirty vehicles as short as possible in the network. Thus, one would generally aim at allocating shortest routes to the dirtiest vehicles.

- Routes should be pollution-aware in the sense that critical pollution thresholds are not exceeded.

Generally, all routing algorithms depend on edge costs w, which commonly incorporate the edges' length, speed limits etc. These (ordinary) calculated edges' weights are denoted w_c. To influence routing, edge costs are adapted to the specific objective, namely that defined pollution limits are not exceeded. Pollution limits exist for various (here n) pollutants. Let l_i be the level for a pollutant i. It is the quotient of the measured pollution and the predefined pollutant threshold (which should not be exceeded). The adapted edges' weight w_a is calculated according to Equation (6.1).

$$W(e) = w_a = w_c(\alpha_k \phi)^{\frac{1}{n}\sum_{i=1}^{n} l_i}. \tag{6.1}$$

The factor ϕ scales the impact of pollution information. Implicitly, it is assumed, that all pollutants are equally interesting (or have equal impact on the overall air pollution) as a simple average is build. This could be easily adapted if a non-equal weighting of pollutants is more appropriate. The factor α is incorporated to implement some sort of fairness. The basic idea here is that fairness should be based on the individual contribution to air pollution. That means that a vehicle, which is cleaner than other vehicles receives shorter routes. In this sense, α is adapted based on the characteristics of vehicle k, here the emission class. Thus, the cleaner a vehicle is, the smaller α gets (the lower the emission class, the higher is α).

Equation (6.1) has two desired features:

- The impact on the edges' weights increases exponentially with the level of pollution. This is reasonable as low levels are hard to avoid.

- The impact of pollution is equal to ϕ, when the measured emission is equal to the predefined thresholds. That means that a shortest path algorithm using the pollution aware edge weights $W(e)$ will choose a route with no emission at all if it is ϕ times longer than the polluted route. This effect is in line with the general idea of pollution aware routing.

However, the latter example indicates a problem that might lead to route recommendations, which are not accepted by users as they imply detours. It is assumed here, that detours are accepted if they are less equal than ϕ times the shortest path. All found routes are rejected if they do not fulfil this condition. Clearly, this constraint implies, that the degree of local emission reduction is limited. However, it will be shown in the following simulations, that with routes allocated, which are at maximum $\phi = 1.5$ times longer than the relevant shortest path, local pollution values can be influenced considerably.

6.3. Simulations: Pollution-aware routing

The pollution-aware routing application has been implemented and simulated. Again, the traffic simulator SUMO is used. The simulation scenario is based on the city of Berlin. Figure 6.2 shows the underlying road network imported from OSM. The objective of the application here is to lower pollution around Alexanderplatz. Accordingly, the pollution threshold for this area is set much lower than for the other areas (cells). The relative level of pollution (actual pollution divided by pollution threshold) of the cells with lower limits should increase much faster and traffic should be routed away from these cells. Thus, by defining pollution limits of the individual grid cells, some system operator can actively influence the traffic flow and control emission locally. The area where the pollution limit should be lowered in the simulation is illustrated by the green rectangle in Figure 6.2. For realistic pollution information, a dispersion

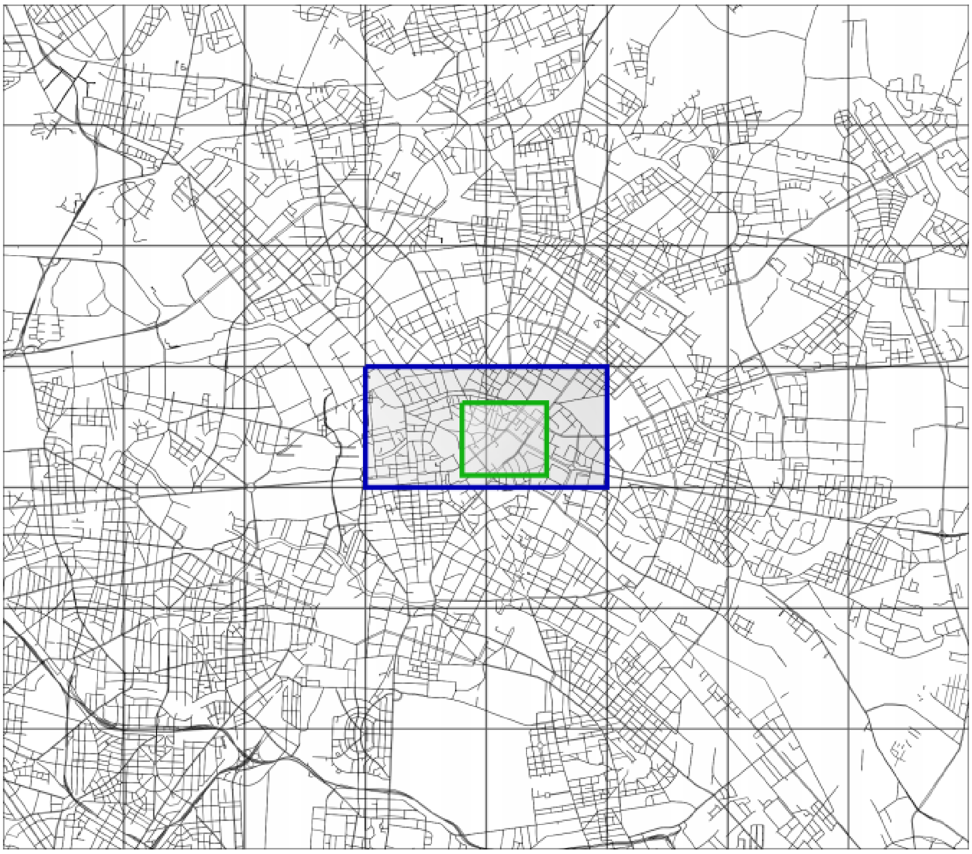

Figure 6.2.: Pollution around Berlin Alexanderplatz should be controlled [107]. The small green box is the area to be controlled. As the emission dispersion model is based on the cells - here 8 x 7 cells. The controlled area intersects two of these (blue framed).

model is integrated with SUMO. The emission dispersion model used was developed by Schmidt and Schaefer and is introduced in [143]. The model is fed with emissions by vehicles, which are provided by the traffic simulator SUMO. The model calculates the dispersion in a (fixed) horizontal grid. On a vertical level, the model differentiates between three layers: the 50 meters thick surface layer, the mixed layer of variable height and the t-layer, which is the upper part of the lower troposphere that goes up to 5000 meters. While all of these layers might be important for air movement and pollution transport, only the actual pollutant concentrations of the surface layer is of particular interest. This is because the main interest is the impact on human health in the given context.

The grid that the model is working with consists of square cells of 4 km^2 - Figure 6.2 shows the grid of cells. The model calculates the dispersion of given traffic emissions; it is also capable to calculate fixed single point emission sources and household emissions.

The emissions issued by the vehicles running in the SUMO traffic simulator is incorporated into the model periodically (each 30 simulated minutes). Thus, SUMO simulates traffic and determines the accumulated emissions for each cell for a time window of 30 minutes. The dispersion model assumes that the emission is equally distributed in the 30 minute time window. Based on this input, the simulator calculates how the emissions disperse during that time. This information (the pollution information per cell after 30 minutes) is then fed back into the emission aware routing application. This loop is repeated every 30 minutes.

The dispersion model is not only fed with the emissions from SUMO but also with information about building densities, water content, average altitude, vegetation, wind, and weather conditions. All of these information were provided for the city of Berlin. The given data refers to records from a day in late September 1997 with a daily average temperature of 22 degrees Celsius, a water temperature of 20 degrees Celsius, atmospheric pressure of 1012 millibars, a cloud cover of 30 percent and no wind. All of these features influence the dispersion [107].

The edges' weights are calculated according to Equation (6.1) with $\phi = 1$ and vehicle (Euro emission class) specific α, according to Table 6.1.

Emission class	1	2	3	4	5	6	EV
α	64	32	16	8	4	2	1

Table 6.1.: Values for α based on Euro emission class.

As shown in Figure 6.2, the green area around Alexanderplatz intersects with two grid cells from the dispersion simulator (highlighted with a blue border). Hence, if the pollution limits are altered for these two cells, the application should re-route traffic away from them, such that the emission is lowered and meets the pollution constraint.

In the first simulation, the pollutants CO, NO_X, PM_X, CO are regarded in an integrated way according to Equation (6.1).

In order to find suitable numbers of vehicles in the network, the following considerations were made: according to [83], about 8000 cars are passing Alexanderplatz on a business day per hour, about two third of them pass in north-south and others in east-west direction. In the given context, only vehicles with origins or destinations outside of the green rectangle are of interest. All others are not considered, as it is impossible to route these away from this area. Accordingly, the number of vehicles (8000) should be reduced. Another reason to reduce the number of simulated vehicles is that a high density traffic leads to unrealistic congestion in the simulation when working with plain imports from OSM. A traffic density of 2400 vehicles per hour turned out to lead to most realistic conditions in terms of average speed, congestion etc. With these numbers traffic flows as shown in Figure 6.3 (a) are generated for the simulation. Specifically, vehicles start from a random destination in one of the four rectangles and head to some random destination in the respective opposite rectangle. The simulation is done for a simulation time of five hours. For fairness analysis,

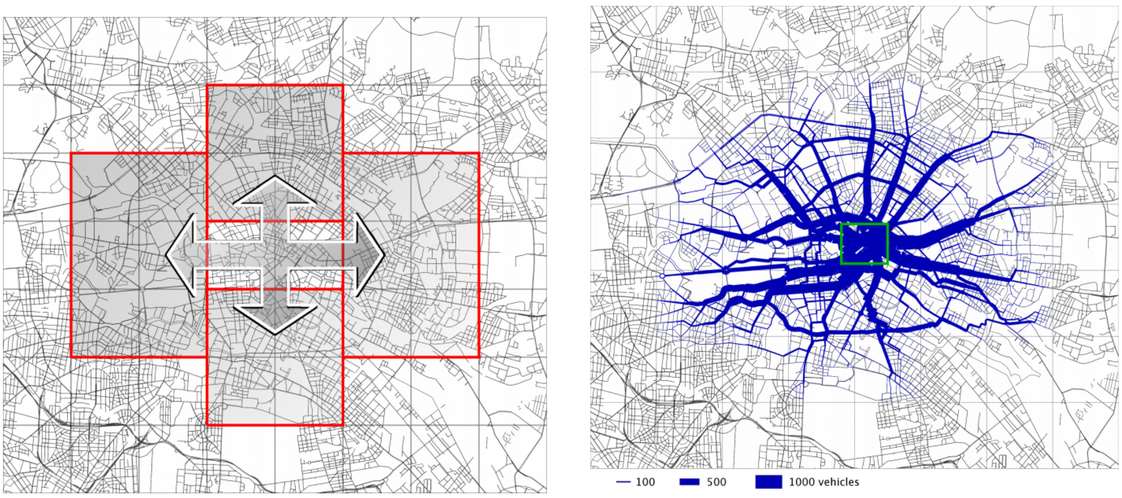

(a) (b)

Figure 6.3.: (a) shows the origins and destination of vehicles passing Berlin Alexanderplatz. (b) illustrates traffic densities if no adapted routing is applied [107].

vehicles of different emission classes are considered. According to the Kraftfahrzeug-Bundesamt [98], about 40.3 % of personal vehicles belong to emission class Euro 4, 19.6 % to Euro 5 and 0.1 % to Euro 6 with 70 % of the cars driving with gasoline, 29 % with Diesel and about 0.02% are electric vehicles. Vehicles belonging to Euro 1, Euro 2 and Euro 3 emission classes are not considered as they are - due to the

Berlin *Umweltzone* - not allowed to pass Alexanderplatz. For reference, Figure 6.3 (b) illustrates the road usage without any adapted routing. The boldness of the roads depicts the number of cars travelling via this road during simulation time. It is easy to see that the roads around Alexanderplatz are highly frequented. The road usages

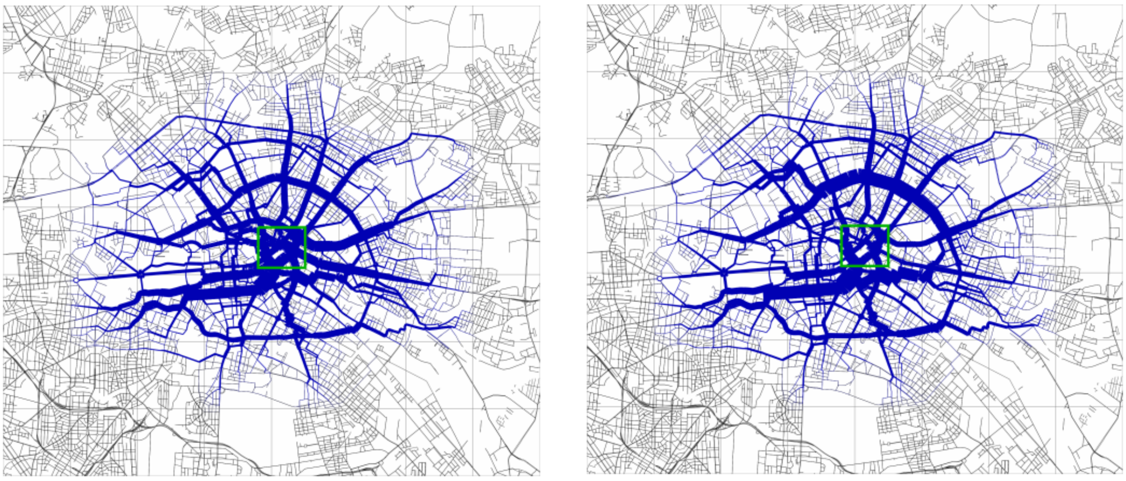

(a) (b)

Figure 6.4.: Figure (a) shows the road usages with 40 % of vehicles using the emission aware routing application. Figure (b) shows the road usages if all vehicles drive according the calculated routes from the application [107]. The results could be compared to the case where no adapted routing is applied, see Figure 6.3.

changes when the adapted routing is in place. Figure 6.4 shows the impact when the application is applied and used by 40 % of the vehicles (Figure 6.4 (a)) or when all vehicle use it, respectively (Figure 6.4 (b)).

Figure 6.5 shows the level of pollution (in regard to aggregate emission levels) if no emission-aware routing is applied. Even though the aggregate pollution thresholds are not exceeded (after five hours, the pollution intensity is about 65 % for the right cell and about 25 % for the left cell), the simulation results reveal weakness of the approach based on Equation (6.1). For reference Figure 6.6 is shown: Recall that the figure shows the results from a simulation run where no adaptive routing is applied. Both figures (a) and (b) show the pollution evolution for the aggregated, averaged pollution. Furthermore - and more interesting in regard to the weaknesses of the approach - it includes the development of pollution intensities per pollutant.

The limits of the approach could be seen in Figure 6.7, which shows the result from a simulation, where the adapted routing application is applied. Even though the control approach is applied the results show that while the controlled, aggregated pollution value holds the constraint (Figure 6.7 (a)), the level of NO_X exceeds its

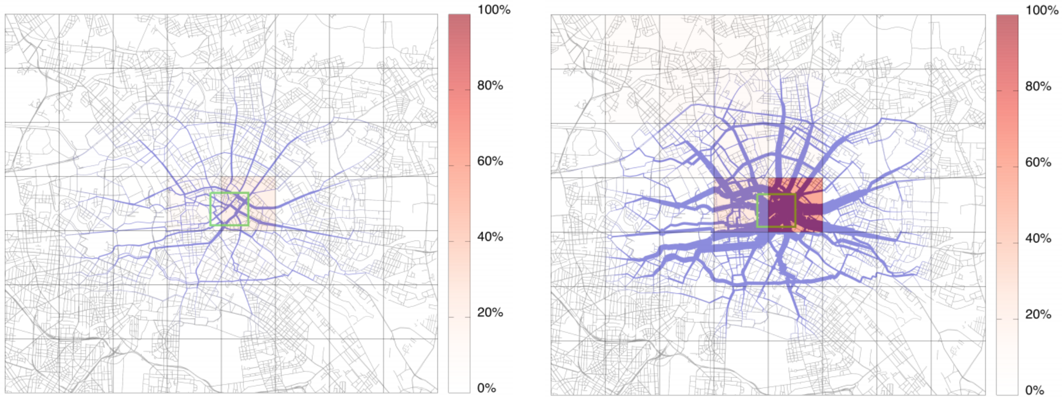

(a) (b)

Figure 6.5.: Level of pollution (quotient of measured pollution and pollution threshold) for each cell after one hour (a), and after five hours (b) [107]. The Figure visualizes the result from a simulation without applying the control application.

threshold (Figure 6.7 (b)). This could be directly explained by the initial assumption that traffic emissions are equally made of the pollutants CO, NO_X and PM_X. The assumption is also part of the key approach defined by Equation (6.1). Figures 6.6 (a) and (b) show that this is not the case. While CO and PM_X emissions is below 10 % of relevant thresholds at all times, NO_X emissions rose to 175 % of their limit. But since Equation (6.1) averages the pollution intensities the system only worked with "masked" feedback.

To countermeasure this problem, one may adapt the system, such that only the most critical pollutant is regarded and controlled. Thus, the simulation is repeated - now looking at NO_X pollution only. Figure 6.8 compares the effects on NO_X pollution of the adapted set-up (in regard to the critical threshold) with the results from the routing application when the aggregate emission is controlled. The figure shows, that the control of NO_X works well with the adapted approach (even if only 50% of vehicles were using the application). Contrary to that, the NO_X constraint is not held when the aggregate emission (100% ECS routing) is controlled.

As said earlier, emission classes of vehicles were recorded for fairness analysis. According to the adaptations of α (see Equation (6.1) and Table 6.1) it was desired that "clean" vehicles must not receive longer detours than "dirty" vehicles. Figure 6.9 shows that this objective is achieved.

Thus, in this particular case, feedback makes more social sense than optimization, which would result in an unfair allocation of routes. This dilemma is one of the starting points for the proposed system in Section 6.3.

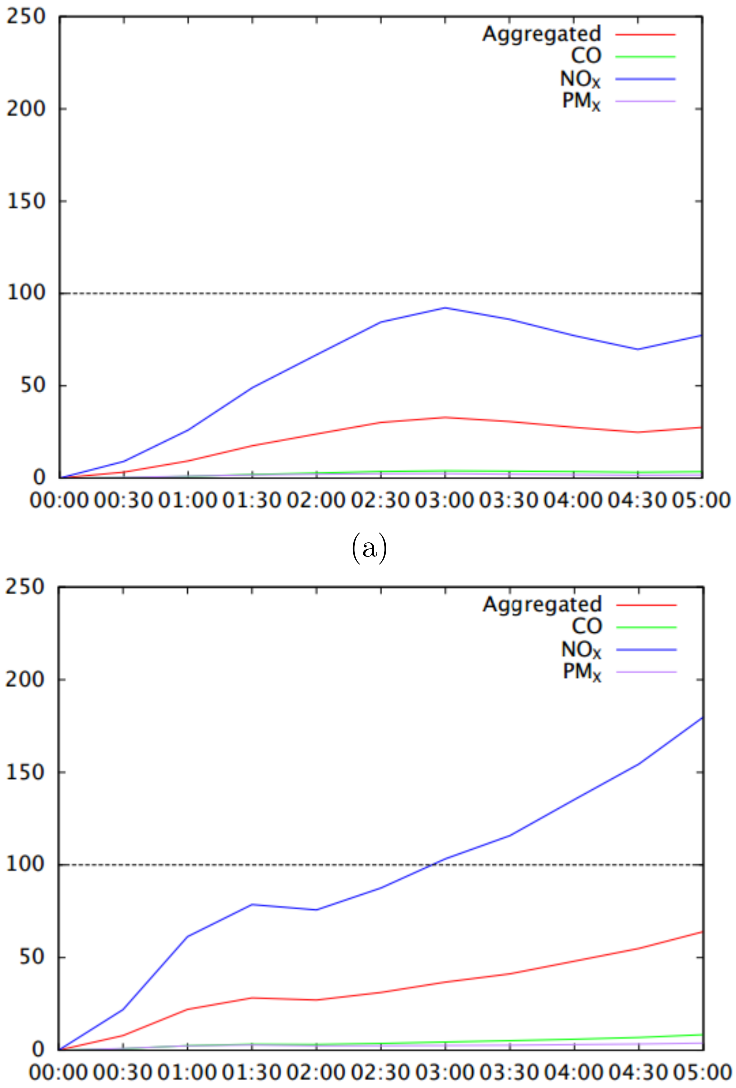

(a)

(b)

Figure 6.6.: The figure shows the level of pollution in regard to the given threshold for the left (a) and the right cell (b) of interest (see Figure 6.2) regarding the three recorded pollutants CO, NO_X and PM_X and the integrated, aggregated pollution level. The dashed line depicts the pollution threshold relevant for the aggregated pollution value as well as for the singular pollutants. The results show a simulation run, where no adapted routing was applied [107].

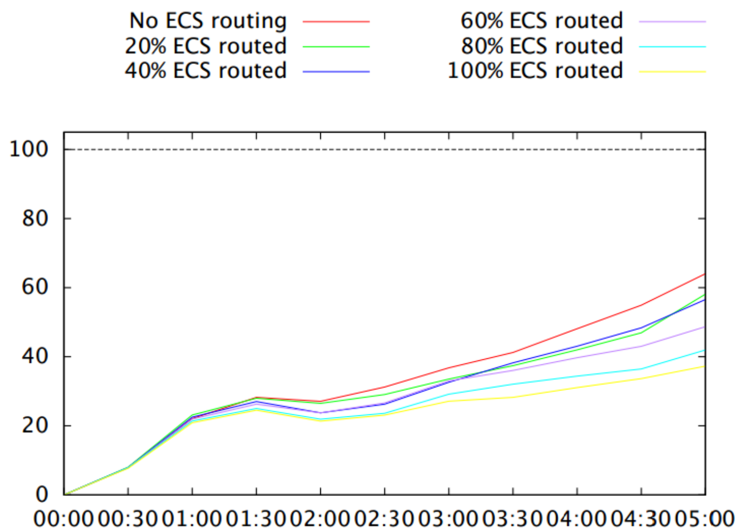

(a)

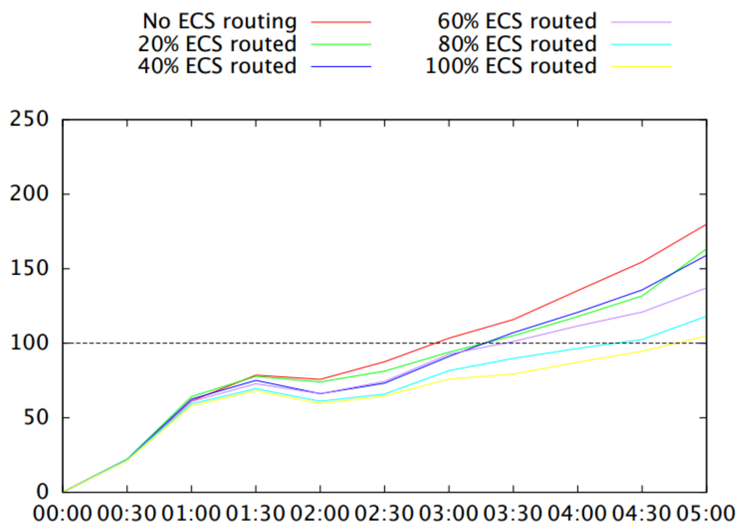

(b)

Figure 6.7.: Both Figures show the development of pollutant for the (more critical) right cell. Figure (a) shows the level of aggregated pollution (the integration of CO, NO_X and PM_X according to Equation (6.1)). The constraint for the aggregated pollution (the dashed line) is hold. Figure (b) shows the level of pollution for the pollutant NO_X. The NO_X pollution threshold (the dotted line) is not held [107].

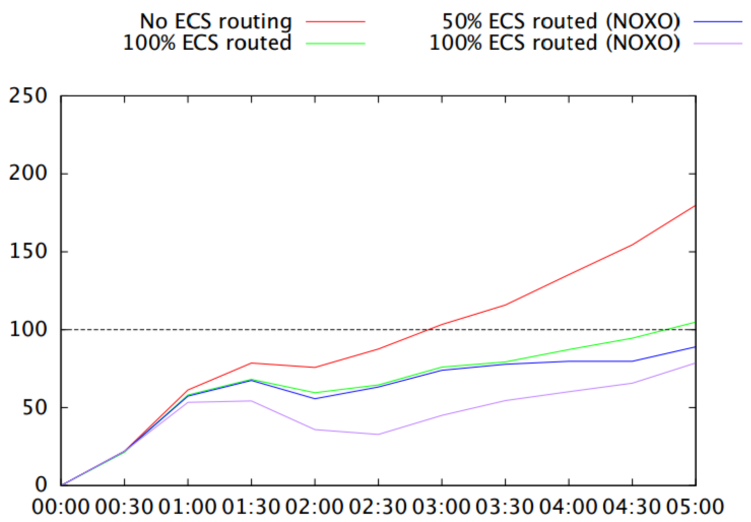

Figure 6.8.: The constraint for NO_X is held if more than half of the vehicles use the adapted routing application when the edge weights are adapted based on NO_X pollution measures only [107].

In addition to this dilemma, another possibility for advancement of the approach is the following: even though the emission-aware routing concept as presented above leads to an improved air quality in highly polluted areas, it could be assumed, that the system inherits flapping effects [123] like all other navigation systems that incorporate edge weights based on real-time information. The following Section 6.4 will introduce an approach to target the flapping problem (when congestions arise and disappear periodically) and avoids the fairness dilemma.

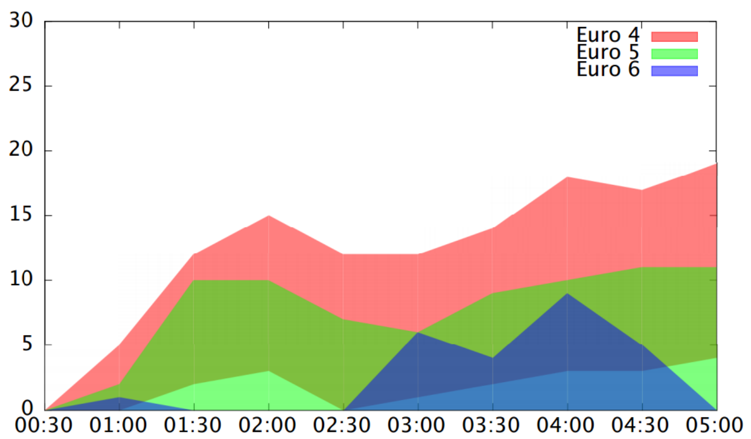

Figure 6.9.: "Clean" vehicles (vehicles with larger emission class numbers) receive shorter detours than "dirty" vehicles [107].

6.4. Stochastic balancing in routing

In the previous Sections 6.2 and 6.3 a pollution-aware routing system was introduced. It turned out that it comes with two drawbacks: (1) a fairness dilemma and (2) flapping problems. The fairness dilemma says that from a system-optimization point of view (in the sense of emission minimization), it would be beneficial if short (attractive) routes were allocated to "dirty" vehicles such that the overall emission is minimized. However, from a driver's point of view, it would be unfair if drivers of "clean" vehicles are penalized by receiving less attractive routes. By flapping, it is meant that the optimal (and allocated) route switches between alternatives back and forth (also known as ping-pong effect). Flapping effects occur, when all vehicles use common shortest path algorithms based on real-time edge costs.

In the approach presented in the following, a slightly different routing objective is chosen in regard to the scale. Here, a routing application is proposed that focusses on locally restricted routing. That means that the approach is not applicable to a city-wide or nation-wide routing. With such an adapted objective, the fairness dilemma could be neglected as the route's length are not dramatically different.

System-optimal routing approaches, which solve the flapping problem have been addressed in the literature before. I.e. Jahn et al. [84] proposed a central routing engine which calculates routes in an aligned way for all vehicles simultaneously. Such an approach has very different features compared to individual shortest path calculations used in ordinary navigation systems. For instance their calculation work on static data only. That means that the routing engine is fed with information regarding all vehicles, including information about individual origins and destinations. The algorithm does neither consider vehicles which enter the system over time, nor in general the changing conditions of the network. Moreover, it requires perfect control of all vehicles. These characteristics are the opposite to the requirements outlined in Chapter 2.2 and which guide the development here.

The basic idea of the stochastic balanced routing application shall be outlined with help of the following scenario: Assume there is a sudden local obstruction (e.g. an accident, a local pollution hotspot, etc.), which is on the route of vehicles. In the given context, a simplified view of the general routing problem is addressed. This is because it targets only those parts of the routes, which are in direct vicinity of the obstruction - it is called local routing. Further, it is assumed that affected vehicles proceed to follow the initially planned route when they have left the affected area.

The ad-hoc road capacity decrease initiates re-routing of affected vehicles to manage the reduced vehicular flow in the vicinity of the obstruction. This should be done in a way that such that flapping effects are avoided. In effect vehicles (with might share the same origin and destination) choose different routes and a balanced traffic load is established. Figure 6.10 illustrates the approach. It is different to common shortest path algorithms (such as Dijkstra-algorithm), which would identify a singular optimal route (here either s_A, s_B or s_C, see Figure 6.10 (b)) for all vehicles.

The problem is solved here by using a stochastic approach which leads to heterogeneous vehicle routes. Alternative routes are suggested to drivers in a manner that balances any kind of feature (e.g. pollution, congestion, etc.) along route options; here s_A, s_B or s_C. As routes are only generated for a locally restricted area, neither communication and transport delays nor the issue of fairness arise.

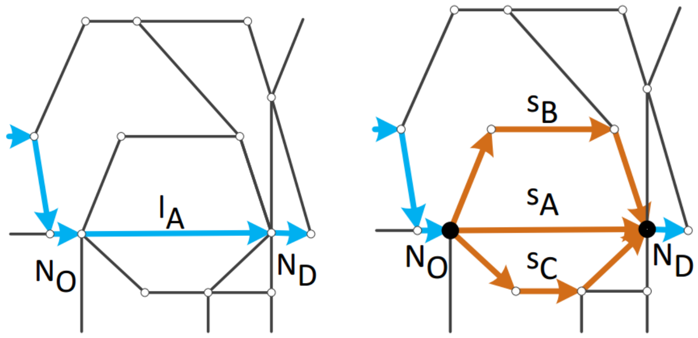

(a) (b)

Figure 6.10.: In Figure (a), the initially planned route is highlighted. It is assumed, that all vehicles wish to travel via the nodes N_O and N_D. At some point in time an obstruction occurs on link l_A, which connects N_O with N_D. In Figure (b), two extra routes are found, which connect N_O with N_D. The algorithm controls the load on route s_A and balances it on the alternative routes s_B and s_C.

The approach could be divided into two steps. First, (multiple) alternative routes are computed for each vehicle. Then, a route from the multiple alternatives is chosen, such that a balanced result could be monitored. Both steps are briefly outlined in the following, while further simulations are included in Chapter 9.3 to validate the approach.

In order to implement the stochastic balancing, it is necessary to have multiple routing options from which vehicles choose one in a smart and aligned fashion. Finding multiple route alternatives is known as the *k-shortest path* problem, where k stands for the number of alternatives. Approaches could be grouped into algorithms, which allow loops (edges could be passed twice or even more often) and those, where this is not allowed [46]. In the context of this work, loops do not make any sense. However, the latter group of k-shortest path finding algorithms are commonly impractical, as they are not efficient regarding computing costs [2]. This experience was also made by Lobach [107], who implemented an adapted version of the k-shortest path algorithm proposed by Yen [189] and applied it to the Berlin road network with the scenario described in section 6.3. However, since the routing problem is restricted to relatively small road networks in the given context, common k-shortest path algorithms could be applied. In addition to that, recent research has found more efficient route algorithms to find multiple (reasonable) routes from one origin to one destination, such as [62, 2]. Generally speaking, the problem of finding multiple routes for one origin destination pair is complex, but there are available solutions.

Assuming that multiple (reasonable) routes have been found, the next step for the vehicle is to choose one, such that traffic or pollution (or any other feature of interest) is balanced. It was already stated, that the basic idea is to do that in a stochastic way. Equation (6.2) defines the probabilities to choose route s_i.

$$P(X = s_i) = 1 - \frac{c_i}{\sum\limits_{i=1}^{n} c_i} \tag{6.2}$$

c_i defines the costs associated with route i. The costs may correspond to pollution, traffic density, congestion etc.

The simple simulation is performed using the hardware-in-the-loop simulation platform. The platform as well as the simulation of the balanced routing application is described in Chapter 9. Anticipating the results from there, it is stated that the approach balances traffic (or any other features) well.

6.5. Remarks

Finally some concluding remarks in regard to the proposed applications are briefly outlined.

6.5.1. Pollution-aware routing

It has been shown that emission-aware routing application implements a feedback loop around a central entity which has access to current network features (such as raffic flow, road link capacities, etc.) and networked vehicles. Pollution values are constantly measured and influence vehicle routes. Vehicle-specific emission features are considered. The route calculation is performed in a distributed way locally at each vehicle. The simulation results have shown that adapted routing helps to influence aggregate emission (and thus local pollution). However, the application implements a dilemma: "dirty" vehicles receive less attractive (longer) routes than "clean" ones. This is desirable from a fairness perspective but leads to an increased overall emission. Furthermore, the application leads to flapping effects (ping-pong effects). At the same time, the application meets most of the previously outlined requirements:

- **Privacy**: Vehicles share no information with other vehicles or a central entity. It is assumed that a central entity (such as a traffic management centre) measures pollution values and broadcasts this information to vehicles. This information is incorporated in the local (on-board) routing engine. Thus, it could be granted, that the application is not critical to any private information.

- **Scalability**: The number of vehicles does not affect the complexity or performance of the system as all calculations are done locally on-board.

- **Communication effort**: Communication is limited to the pollution information broadcasted from one central authority to vehicles. It could easily be controlled by an efficient selection of broadcasted data.

- **Stability and Convergence**: Due to the named flapping effects (ping-pong effects), pollution levels do converge on average over time only.

6.5.2. Stochastically balanced routing

Main drawbacks of the previous routing application are the existence of the fairness dilemma and the occurrence of flapping effect. Latter one leads to constrained convergence. The stochastically balanced routing application solve this problems. Note that - in contrast to the emission aware routing application - the approach is applicable to locally constrained routing only. At the same time, the application meets the general requirements:

- **Privacy**: It is assumed that vehicles share common information about the network status in regard to the balancing objective. That means that there might be a central infrastructure, which broadcasts real-time information about the level of pollution, traffic flow or any other objective of interest. The information is shared in a way that routing algorithms could apply them (i.e. on a per-edge level). In regard to privacy this information is not critical. The central infrastructure does not need any per-vehicle information, such as origins, destinations, emission classes etc. Such critical information is used in on-board systems only.

- **Scalability**: The approach is perfectly scalable as the communication effort does not depend on the number of vehicles. As the approach is locally restricted, the information to be shared about network states are limited. Processing at the central entity does also not depend on the number of vehicles. Most complex calculations are done in a distributed way locally at each vehicle. The complexity here is also independent from the number of participating vehicles.

- **Communication effort**: The communication effort is dependant on the size of the relevant network to be considered. Thus, it has an upper bound.

- **Best effort**: There is the requirement that vehicles should be able to enter and exit the system. This is not an issue.
- **Stability and convergence**: The approach converges to a stable equilibrium. This could also be seen in the simulation results shown in Chapter 9.

7. Balanced charging for EVs

Summary: *The basic idea of the balanced routing application introduced in the previous chapter is also the starting point for the balanced allocation of charging stations to electric vehicles. Motivated by the need to provide services to alleviate range anxiety of electric vehicles, the application targets the problem of balancing charging demand across a network of charging stations for electric vehicles. The objective is to reduce the potential for queueing at charging stations.*
The development of the balanced charging application is result of joint work with Emanuele Chrisostomi, Arieh Schlote, Ilja Radusch and Robert Shorten, which has been published in [74].

7.1. Queueing at charging station

Even though the history of electric vehicles is very long, they are still rarely seen in public road traffic. Due to several reasons such as increased environmental awareness, strategic political and economic considerations, noise pollution and technical progress, the concept of electric vehicles is experiencing a revival lately. Furthermore, electric vehicles are expected to play a major role in making the road transport system more environmentally friendly [132].

While benefits of zero emission vehicles are obvious, their mass deployment and adoption by users are uncertain due to a number of factors: (i) the limited availability of rare earth metals to construct the vehicles and batteries, (ii) safety concerns related to electromagnetic emissions, (iii) long charging times and (iv) the limited range. In this context the latter two issues - often discussed under the evocative title of range anxiety - are of particular interest as they are addressed by the application introduced in the following. Much of the current research and development is focussing on new battery types, optimal vehicle charging, in-vehicle energy management systems aiming at increasing vehicles' ranges. There are other approaches which propose a smart integration of electric vehicles with common internal combustion engines in a car sharing system [94]. Decreasing the charging time of electric vehicles' batteries is approached by fast charging stations, which can service average vehicles in about 30 minutes. Battery swapping is another approach prototypically implemented by Tesla recently [116]. Even though swapping is even faster than refuelling a common internal combustion engine, the concept is not expected to be deployed on a very large scale like (fast) charging stations.

However, in the presence of queuing, 30 minutes to recharge a vehicle at a fast charging station can rapidly become several hours - a scenario which is probably not accepted by any users.

The proposed solution avoids unbalanced demand for charging stations (and thus to avoid queuing) by exploiting V2X communication technology to implement a feedback loop again. The objective is to dynamically assign electric vehicles to charging stations in a manner that balances load in a geographic area. This algorithm can be considered as a hybrid of two related approaches: first, routing algorithms to balance the load on the charging network so that both the demands on the grid are balanced (as well as the expected waiting time at each charging point), and two, placing charging stations so as to balance the expected load on each charging point. The developed approach here is motivated by the limitations of both of these.

At the same time it respects the previously mentioned ITS specific requirements (scalability, privacy preserving, best effort behaviour, stability and convergence). It is shown that the proposed stochastic, decentralized solution is superior in many aspects while it reaches an almost equivalent performance as the optimal solution.

7.2. Related approaches: Routing for EVs and smart placing of charging stations

The problem of balancing the charging load across a set of charging stations can be tackled in a number of ways. Two basic instruments serve as a basis for the proposed solution. That is (1) routing to balance traffic load according to charging stations' availabilities or - the opposite case - (2) placing charging stations to service the expected demand from electric vehicles (or PHEVs) in a balanced fashion.

The first approach can be formulated as a classical routing problem where network-wide information is used to balance load [133]. Unfortunately, load balanced routing strategies place severe constraints on drivers by forcing them to take particular routes or charging stations. As in the case of pollution-aware routing, the willingness of drivers to follow recommendations is a significant issue in the design of such systems.

The second approach is a kind of dual of the routing approach: While vehicles were controlled previously and adapt to given charging station availabilities, the second approach adapts the placement of charging station based on given traffic densities. It can be formulated as a partitioning problem on a graph and can be solved by combining traffic density information with graph partitioning techniques such as the Voronoi partition. Such charging station positioning problems are already addressed in the operations research community [187].

While both approaches are interesting from a theoretical perspective, they suffer from a number of practical drawbacks: (1) traffic densities are temporal in nature. Consequently, an optimal placement of charging stations depends on changing traffic conditions, i.e. optimal places for charging stations during morning traffic might be completely different to the places optimal for evening traffic and (2) load balanced routing strategies place severe constraints on drivers by forcing them to take particular routes - or respectively - by forcing them to head to particular charging points.

Therefore, the goal is to investigate other balancing solutions: as before, the objective is to balance charging demand across a number of charging stations. The system should be adaptable to changing traffic densities. It should be implemented without centralised communication between infrastructure and vehicles. Furthermore, the system should not force drivers to take particular routes or to drive to particular charging stations. A further objective is to realise the system in a plug-and-play manner, which means that charging stations can enter and leave the system as they wish, without placing a reconfiguration requirement on a centralized infrastructure. Further, the solution shall come with minimal communication requirements that can be implemented using cellular networks. A key objective be scalable in the sense that the system can handle varying vehicle and charging stations numbers.

7.3. Stochastic balancing for electric vehicle charging

The starting point for the approach is the analogy with the mobile cellular network: electric vehicles correspond to mobile phones and charging points to base stations. Similar to the association of phones to base stations in communication range, electric vehicles are associated to charging points in the vicinity. In this sense, the city area is divided into cells as it is illustrated in Figure 7.1. The partitioning can be realised in many ways (e.g. using Voronoi diagrams). However, the main focus here is the car assignment problem within a cell which takes place over faster time-scales than the partitioning problem. The basic idea is to partition the relevant area into a number of cells where the quotient of the number of vehicles and charging stations is roughly constant. Note that the partitioning does not need to be fixed but may change with temporal or seasonal traffic patterns.

Within each cell, vehicles are associated to one charging station. Furthermore, within each cell, each (electric) vehicle is assigned to a charging station in a dynamic fashion with the objective to balance expected charging time within each zone.

As vehicles travel from one cell to another, a hand-over procedure is initiated and vehicles are assigned to a new charging station in the next cell. An alternative assignment strategy is to allocate vehicles with charging station only in case it requests a station. The latter strategy is discussed in more detail in the following. The advantage of both approaches is that there are no restrictions on vehicle mobility patterns; and the cells can be varied dynamically to reflect traffic density patterns.

Figure 7.1.: Partition of the Berlin city centre in cells.

A stochastic algorithm is proposed to achieve the following: within a cell, charging stations broadcast a "green" signal, indicating their availability to service (to charge) new electric vehicles. The broadcasting frequency is inversely proportional to the current queue length. The queue corresponds to the actual queue of vehicles waiting to be served at the relevant charging station. Alternatively, one may use the number of associated vehicles, the fraction of currently utilized charging stations (assuming here that charging stations may serve multiple cars at the same time), the expected time of occupancy and many others. Similarly, it is assumed that vehicles are capable of receiving and processing green signals with a frequency that is inversely proportional to the current filling level of their battery. If, at a given time, a vehicle receives and processes a green signal from a charging station, the vehicle is assigned to it. The objective of such a protocol is to associate (1) vehicles that are more in need of charge with (2) the charging points that have the shortest queue. The approach allows a fully decentralised set-up The actual implementation of the protocol is described in more detail in the following subsections. Figure 7.2 illustrates the basic idea.

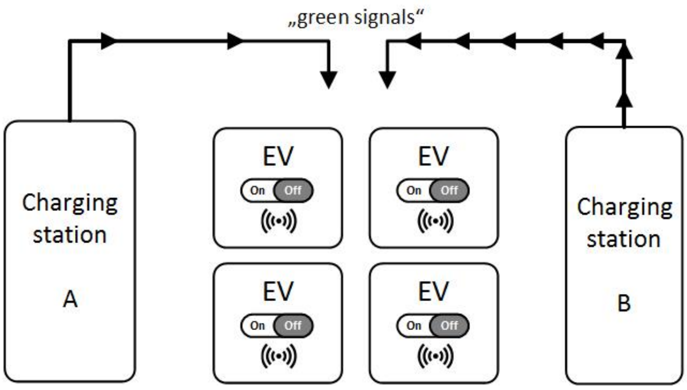

Figure 7.2.: General overview of the balanced charging application setup: Charging stations broadcast "green" signals indicating their availability to service EVs in a frequency which is inverse proportional to the queue at charging station A and B. The figure indicates that the charging station do not broadcast in the same frequency. EVs process received signals inverse proportionally to their battery filling level.

Such decentralised solutions are very attractive for a number of reasons: apart from general benefits (such as robustness in regard to node failures, less critical regarding privacy and security) the decentralised solution facilitate the possibility of implementing plug-and-play policies, where charging stations could be added and removed from the overall system without requiring a reconfiguration of other system's parts. This will be illustrated in a dedicated example.

One might question the implementation of the set-up in a probabilistic way rather than use a deterministic approach. The main benefit of the stochastic approach is that it could be implemented without requiring a communication network between charging station, which are not aware of queues at other stations. This leads to a truly decentralised solution.

7.3.1. Charging station subsystem

In the following a brief introduction of the system to be deployed at the charging station is introduced. At every time step the ith charging point communicates (broadcasts the *green signal*) its availability to accept a new vehicle with probability $p_{CS}^{(i)}$, which is a function of N_i – the number of vehicles currently queuing the charging station. Further, if no vehicle is waiting to be charged at the charging station, it shall be maximal, thus $p_{CS}^{(i)} = 1$. As the queue grows, $p_{CS}^{(i)}$ should decrease monotonically to 0. As charging stations with a shorter queue are more likely to accept a new car than ones with a longer queue, such a mechanism is able to balance the number of vehicles associated with each charging point.

An exemplary function for $p_{CS}^{(i)}$ meeting this requirements is:

$$p_{CS}^{(i)} = 10^{-N_i} \tag{7.1}$$

The choice of Equation (7.1) is arbitrary; but it has some desired features. (1) A free charging station with a queue length $N_i = 0$, broadcasts the availability to accept a new vehicle at every instant. That means the probability that a vehicle is associated to this charging station is maximized. (2) The probability to associate a vehicle with a charging station where few vehicles are waiting is higher than the probability to associate it with a charging station with a longer queue. I.e. given two charging stations i and j with $N_i = N_j + 1$, then the charging station j broadcasts the green signal on average 10 times more frequently than i. The ratio 10:1 is arbitrary. However, it ensures that stations with shorter queues are much more likely than ones with fewer spaces to have new cars assigned.

The idea is that such a mechanism could be able to balance the number of vehicles associated with each charging point.

7.3.2. Electric vehicle subsystem

After having introduced the subsystem to be deployed at the charging station in the previous section, a brief overview on the corresponding vehicle subsystem is provided. Once the i'th electric vehicle is searching for a charging station it is assumed that this vehicle may process broadcasted green signals with a probability of $p_{EV}^{(i)}$ computed by Equation (7.2).

$$p_{EV}^{(i)} = 1 - \frac{e(i)}{M} \tag{7.2}$$

It is assumed, that green signals are discarded if the vehicle has no need to be charged. The current capacity of energy in the i'th vehicle's battery is denoted by $e(i)$. M is a maximum capacity of batteries. For the sake of simplicity, it is assumed that all vehicles have the same battery types. Equation (7.2) ensures that EVs with low charge are more likely to process green signals than vehicles with plenty of residual charge. That is done to avoid spending essential energy while waiting for an available charging point.

7.4. Analysis of stochastic balancing

A brief analysis of the algorithm is provided in the following. For simplification, three main assumptions are made:

- **Simultaneous offers (collisions) are neglected or are resolved**: If two EVs i and j sense a green signal from a charging station, then only one is assigned to it. Similarly, if an EV senses multiple green lights, then it is assigned to only one charging station (e.g. the closest one).
- **Higher order effects are neglected**: In the first part of the analysis it is assumed that the probability that an EV will get recharged from charging station i is

$$\frac{p_{CS}^{(i)}}{\sum\limits_{j=1}^{m} p_{CS}^{(j)}}, \tag{7.3}$$

  which is obtained by neglecting high order terms (i.e., considering only terms that are linear with respect to the probabilities).
- **Stationary**: In deriving waiting times for EVs and charging stations (see below in Section 7.4.2), a stationary assumption is made: in steady-state, the number of vehicles queuing at the charging stations is constant.

7.4.1. Convergence

As indicated before, the system is best described as a Markov chain and can be analysed in this framework. The analysis for a special case of three stations, which are capable to serve three EVs each, is given in the following.[1]

First, a step size for the chain (in seconds, here 1 second) is selected and the average arrival and departure rates are defined, here 0.003. This corresponds to a situation where every 333 seconds a vehicle arrives (and departs) at a station which can service three cars. Further, this corresponds to an average charging time of 1000 seconds or about 16 minutes - assuming all charging stations are fast charging stations. Then for $n_m = 3$ charging stations a state of the Markov chain is a vector of length n_m, which reports the number of EVs awaiting service at each station. The probabilities

[1] The method could also be applied to a larger number of vehicles but is difficult to write in a compact form.

of transition between the states can be computed from the arrival and departure rates and the probabilities of a vehicle ending up in each of the station, as approximated for example by Equation (7.3). The Markov chain is irreducible and aperiodic [122]. If the number of customers in each station is bounded, then the number of states is finite and one can use the transition matrix to compute the stationary distribution, which yields information about what fraction of time the chain spends in each state.

Figure 7.3 shows the fraction of times where the difference of the longest queue and the shortest queue in the system is equal to zero, one, two and three. This is computed from the aforementioned stationary distribution. It can easily be seen that states with a high difference in queue lengths (unbalanced situations) occur very rarely.

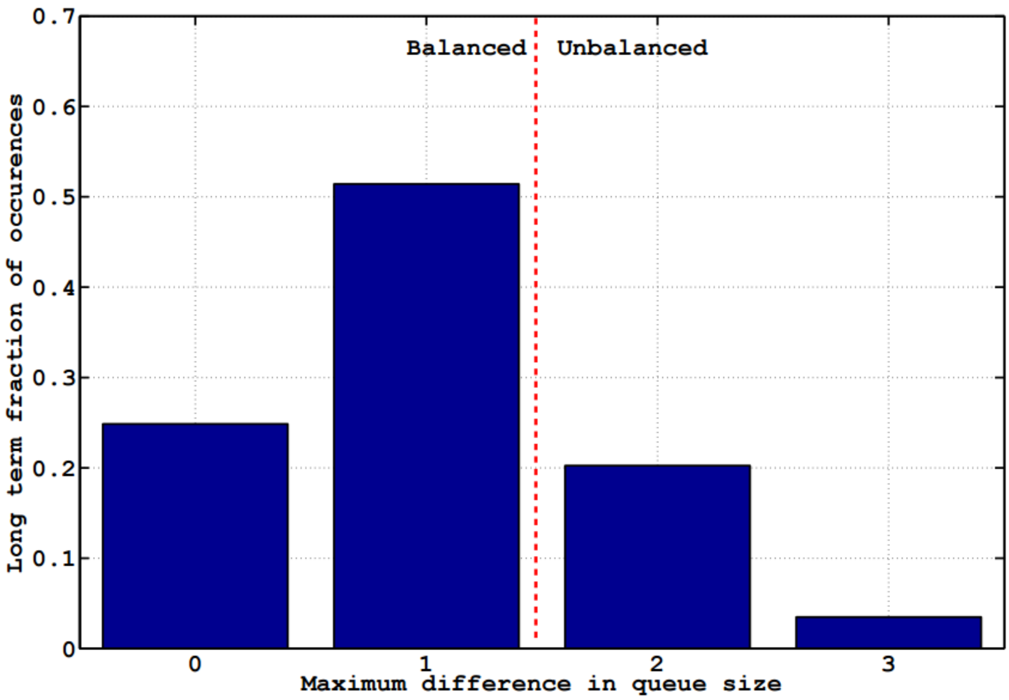

Figure 7.3.: Stationary distribution of queue length differences.

Besides this exact analysis for three charging stations, one may apply an approximate analysis and a plausibility argument. It is assumed that over each time step, each station i loses a fraction β of its vehicles N_i. In addition, each station obtains a fraction of arriving vehicles (the second term of Equation (7.4)).

$$N_i(t+1) = \beta_i N_i(t) + \frac{\alpha_i}{\sum_{j=1}^{n_m} \alpha_j} \left[\sum_{j=1}^{n_m} (1-\beta_j) N_j(t) \right] \tag{7.4}$$

The α's are given from the probability functions described earlier (Equations (7.1), (7.3)). This model has been used to analyse resource allocation network problems in the internet [148], and it is known that the states converge exponentially fast to the $[\alpha_1, \alpha_2, \ldots, \alpha_{n_m}]$ [93].

7.4.2. Waiting times

Another critical feature of the application is the waiting time before cars are assigned to a charging station. Thus, in the following analysis, the waiting time does not target the time until the process of charging is triggered but the time a vehicle needs to wait until it is associated to a charging station. That is the reason why information about the actual charging times does not need to be considered.

Waiting times of electric vehicles

A network of n_m identical charging stations is considered. At a certain time step one vehicle i requires battery recharge. Then the average waiting time $t_{wait}^{EV_i}$ before being assigned to a charging station is given by Equation (7.5):

$$t_{wait}^{EV_i} = \frac{1}{p_{EV}^{i} \cdot (1 - p_{red})}, \tag{7.5}$$

where

$$p_{red} = \prod_{j=1}^{n_m} \left(1 - p_{CS}^{(j)}\right), \tag{7.6}$$

which defines the probability that no charging station broadcasts a green signal. The waiting time of vehicle i $t_{wait}^{EV_i}$ is generally independent from the particular charging times or features of charging stations (i.e. whether they are fast charging stations or not).

Thus, the i'th EV gets assigned if it receives at least one green signal (probability $1 - p_{red}$) and if it is capable of sensing it (with probability $p_{EV}^{(i)}$). As the two events are independent, this occurs with probability $p_{EV}^{(i)}(1 - p_{red})$. The expected waiting time is given by Equation (7.5). Recall that the waiting times are defined by the frequency of the re-calculation of p_{EV} and p_{CS}. I.e. assume that both probabilities are recalculated every 5 seconds, then a calculated waiting time of 3 corresponds to 15 seconds. The quotient p_{EV}^{i} in Equation (7.5) is included to account for the fact that vehicles which request charging service should be regarded only.

Charging station waiting time

A network of n_m identical charging stations is considered again. It is assumed that EVs require charging, and that a new request is made as soon as the previous one is satisfied. Then, the average waiting time $t_{wait}^{CP_i}$ of the charging station i before receiving a new vehicle is given by Equation (7.7):

$$t_{wait}^{CS_i} = \frac{1}{\bar{p}_{EV} \cdot (1 - p_{red})} \sum_{j=1}^{n_m} \frac{p_{CS}^j}{p_{CS}^i} \tag{7.7}$$

where $\bar{p}_{EV}$ is the probability associated with an average request of energy according to Equation (7.2). The first term in Equation (7.7) generally accounts for describing the in-vehicle part - namely that a vehicle is associated to a (which-soever) charging station - the sum defines the inverse probability that it is associated to the charging station i. Equivalent to the calculation of the EV's waiting time, recall that the waiting time is defined by the frequency of recalculations of variables.

To conclude: the average time before a particular vehicle is assigned to a charging station is given by Equation (7.5). Here, it is further considered that the vehicle must be assigned to the particular charging station i rather than to another one (with probability given by Equation (7.3)).

7.5. Simulations

In the following, five simulations (including a prototype implementation using the integrated simulation environment VSimRTI [145]) are presented to illustrate the efficacy of the balancing approach.

Two autonomous applications are developed: one in-vehicle application and one deployed at the roadside station (IRS) - here attached to a charging station. They are integrated to form the balancing approach. The basic idea of the communication flow between the two is illustrated in Figure 7.4. All updates of calculations - applying

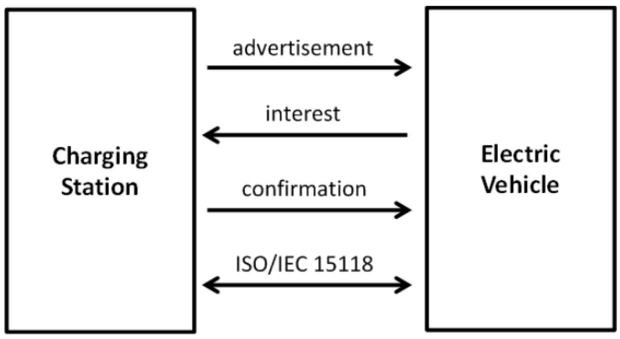

Figure 7.4.: Communication flow between vehicles and charging stations.

Equations (7.1) and Equation (7.2) - are performed once a second. Furthermore, broadcasts of green signals are issued with a frequency up to 1 Hz.

For the protocol implementation, a simple generic V2X message type is developed. As described earlier, the charging station computes the probability to broadcast its availability to accept a new vehicle (an advertisement, the green signal). The charging station broadcasts the green signal each time frame according to the probability defined by Equation (7.1). At the vehicle side, the application calculates the probability to listen to such advertisements. In return to the received green signal, the vehicle sends a notification of interest to exactly one of the advertising charging stations (i.e. to the closest one). The charging station, which receives the notification of interest, chooses one vehicle (here the vehicle who first sends a message) and sends a confirmation to that vehicle. All other incoming requests of vehicles are ignored, while the selected vehicle is added to the queue and a new broadcasting probability is calculated.

7.5.1. Balancing performance

The first simulation performs the stochastic balancing with ten charging stations.

The simulation setup is as follows: The simulation runs for 16 hours with ten charging station. In the beginning of the simulation, all charging stations are assumed to be empty. Vehicles enter the simulation and demand charging according to two Poisson distributions: To fill the queues at charging stations, the inter-arrival time is 90 seconds for the first 5 hours of the simulation. Once the queues are full, the arrival and departure rates are equalised by selecting an inter-arrival time of 100 seconds. Moreover, a fixed charge rate of 0.01kW/s per vehicle is assumed. The time required for charging is individually set for all vehicles according to a Poisson distribution with $\lambda = 0.001$ (thus the expected time of charging being equal to 1000 seconds or 16 minutes). This corresponds to assuming an average request of 10 kWh of energy. Thus, the number of vehicles entering and departing (the ones which have been charged) is balanced.

Figure 7.5 shows the distribution of vehicles at charging stations (the average number and the according variance).

In the second part of the simulation, the average number of charging vehicles is usually within three or four cars per charging station (i.e., 30-40 overall) and that the variance is always extremely low (about 0.36 on average).

7.5.2. Stochastic vs. deterministic balancing

The second simulation compares the proposed stochastic algorithm with a deterministic one. The deterministic implementation serves as a reference as its balancing performance is perfect and it is implemented in a way that the deterministic approach directly assigns entering vehicles to the charging station having the shortest

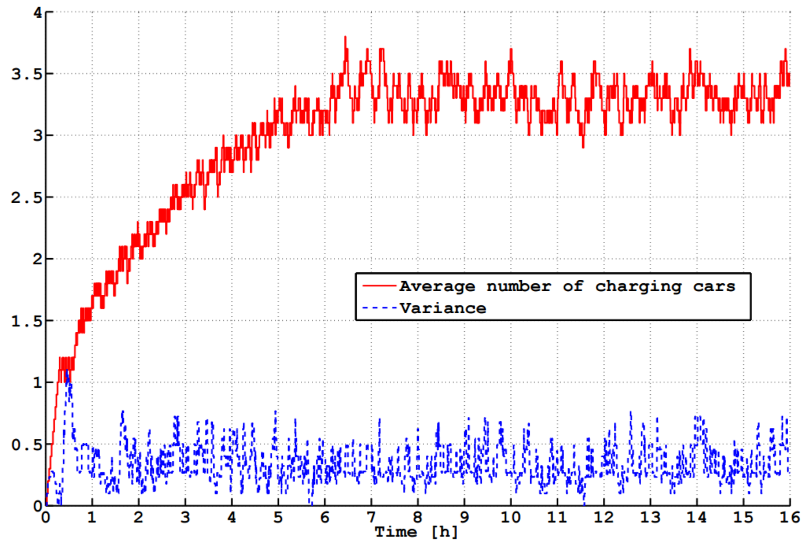

Figure 7.5.: At the beginning of the simulation, charging stations are assumed to be empty. After a few hours, the average population of cars oscillates between 30 and 40 cars overall. The figure shows the number of charging vehicles per charging station. The variance of queue length over the ten charging stations is 0.36.

queue.

Similar to the previous simulation, the charge request of energy ranges between 1 and 10 kWh, so that the maximum time for recharge (corresponding to a request of 10kWh of energy) is 1000 seconds (16 minutes). The simulation here runs for eight hours. It is also assumed that at the beginning of the simulation (8 AM) all charging stations are empty (note that Figure 7.6 is showing a temporal detail of the development and not the complete simulation time). Figure 7.6 illustrates the results of both approaches.

While the deterministic method is optimal regarding balancing performance, it requires a large amount of communication between all participants (cars, stations, cloud). However, as presented in Figure 7.6, the average number of vehicles waiting at charging points using the stochastic algorithm is very close to that obtained from a centralised solution. The two approaches also provide very similar performance in terms of the variance of the solution (variance equal to zero corresponds to exactly the same number of vehicles queueing at each charging point). Thus, the distributed approach comes with multiple benefits but limited disadvantages regarding the balancing performance.

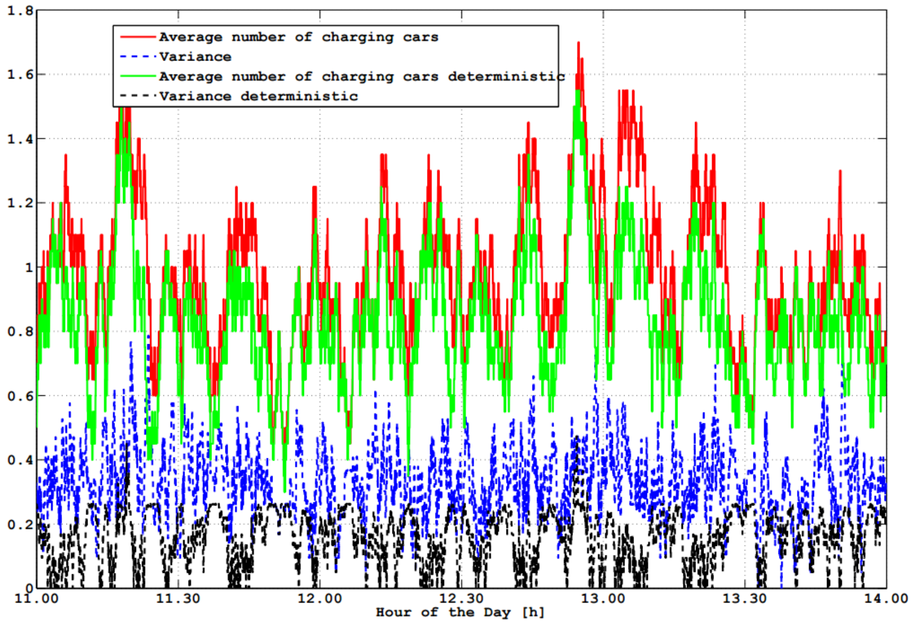

Figure 7.6.: Stochastic vs. deterministic balancing. The decentralised stochastic approach provides results very close to the deterministic centralised approach, both in terms of average number of queueing vehicles at each charging station and in its variance.

7.5.3. Waiting times of EVs

The next simulation (and its results) show that waiting times are within reasonable limits, provided that the number of charging points is well sized with respect to the number of EVs. The same setup as in the previous simulation is employed. Figure 7.7 shows the evolution of the waiting times of EVs - here from the service request to the assignment to a charging station. Recall that charging stations broadcasts green signals according to Equation (7.1), which is updated once a second in the simulation. Green signal are also issues less equal that once a second (the charging station i broadcasts one green signal per second if $p_{CS}^{(i)} = 1$).

Figure 7.8 is showing the according waiting time at charging stations. More precisely, it is the time between the vehicle's association to a charging station and the time when the vehicle is fully charged. The charging duration is up to 16 min. Keeping this in mind, the figure shows acceptable times on average. Variances could be lowered by integrating the amount of energy requested by associated vehicles into the communication protocol as well as in Equation 7.1.

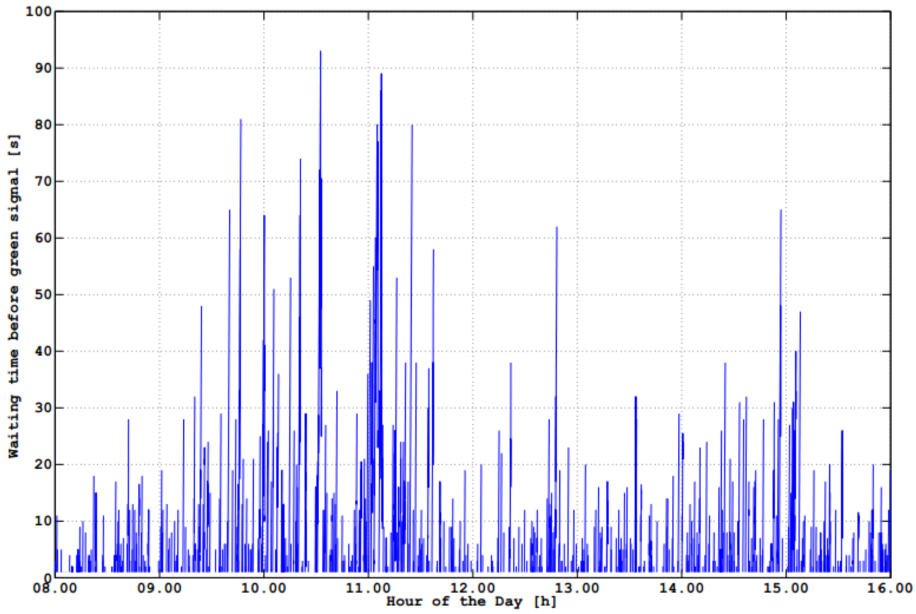

Figure 7.7.: EV waiting time: Illustration of the average time between a vehicle's charging request and the moment it gets associated with a charging station. A probabilistic algorithm facilitates vehicles that have a lower level of battery, but the waiting time never exceeds two minutes in the simulation example.

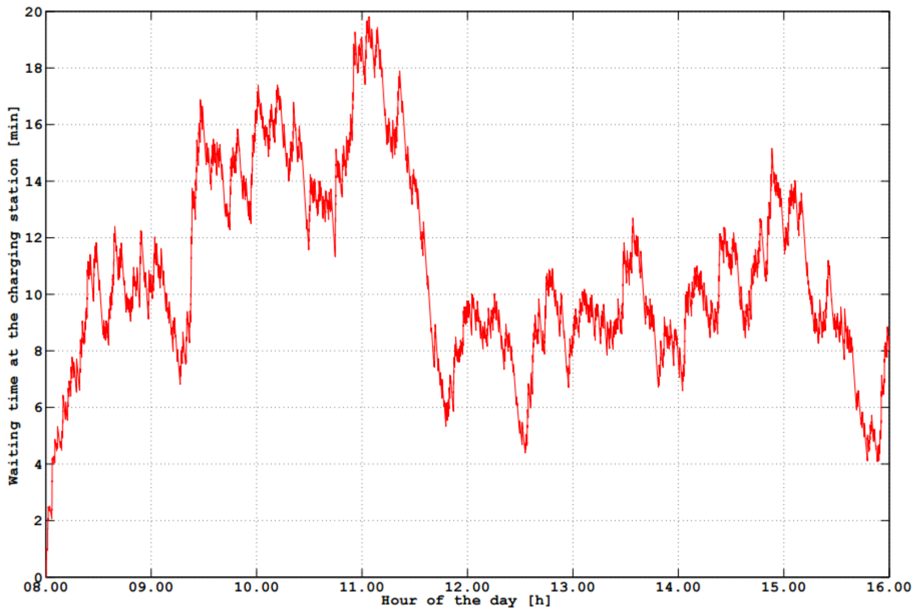

Figure 7.8.: Illustration of the average waiting time of EVs at charging stations including time for charging.

7.5.4. Plug-and-play behaviour

A key benefit of the approach is that it provides the opportunity of handling a varying number of charging stations in a manner that is completely transparent to vehicles and stations. Such a feature is beneficial in many scenarios, e.g. charging stations might be reserved for particular groups during office hours and be publicly available at other times.

In this simulation it is shown that the algorithm can be used to achieve such a functionality. It is assumed that there are only three charging stations in the network until 12 PM. From then on two extra stations are added and integrated which start offering their service. Figure 7.9 illustrates that the maximum difference between the longest and the shortest queue is usually very small (meaning that it is well-balanced), and is particularly large only when the new charging stations become available. That is because they are initially free. However, fairness is soon restored.

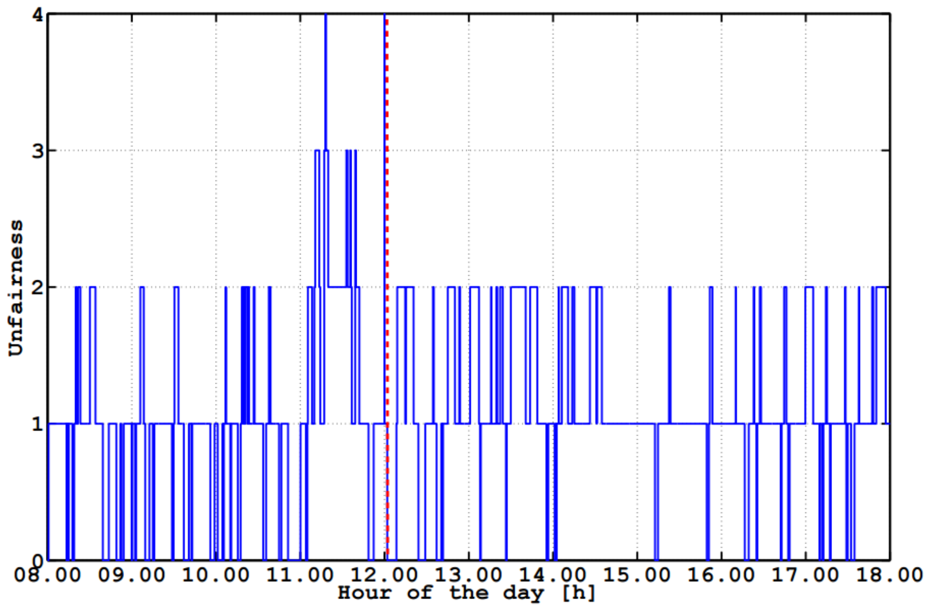

Figure 7.9.: Balancing performance in regard to the plug-and-play behaviour in terms of maximum difference between the busiest and the freest charging station, which could be interpreted as a fairness feature.

Figure 7.10 shows the development of average waiting times of EVs until they are being associated to a charging station. It is shown that in the moment the two extra charging stations are integrated (indicated by the vertical red, dashed line) the waiting times are decrease.

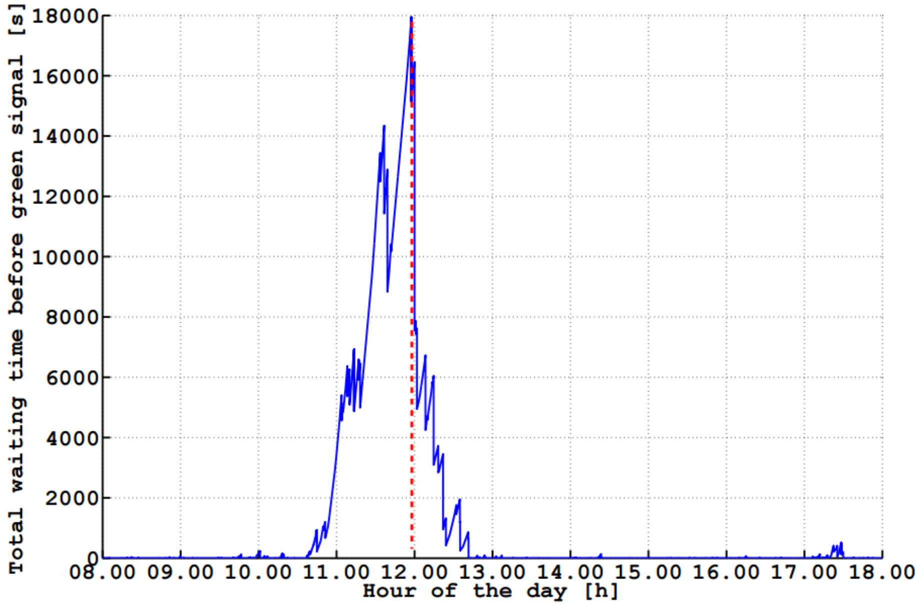

Figure 7.10.: The figure shows the aggregate time required before receiving a green signal, which is greatly reduced when the two new charging points are available - indicated by the red, dashed line.

7.5.5. Communication effort

As outlined in the protocol description, the complete vehicle assignment process has three steps. Figure 7.11 shows the number of sent messages per allocation. It depends on variable numbers of charging stations and electric vehicles.

Figure 7.11 shows that the number of exchanged messages increases linearly with deployed vehicles and charging stations. The approach has been simulated for each possible combination of numbers of electric vehicles and charging stations. The length of one time step is equivalent to the interval in which charging stations re-compute $p_{CS}^{(i)}$ and - if applicable - broadcast the "green signal". All other messages are generated in return to incoming messages, either of electric vehicles (notification of interest) or charging stations (confirmation message).

7.6. Remarks

The issue of queueing has been hitherto ignored in the study of electric vehicles, even though it can significantly impact the acceptability. Compared to deterministic methods to reduce queueing, the approach is fully distributed and has a plug-and-play character. At the same time the performance compares well. The algorithm can be deployed easily without any needed modification to assign parking spaces in an urban setting.

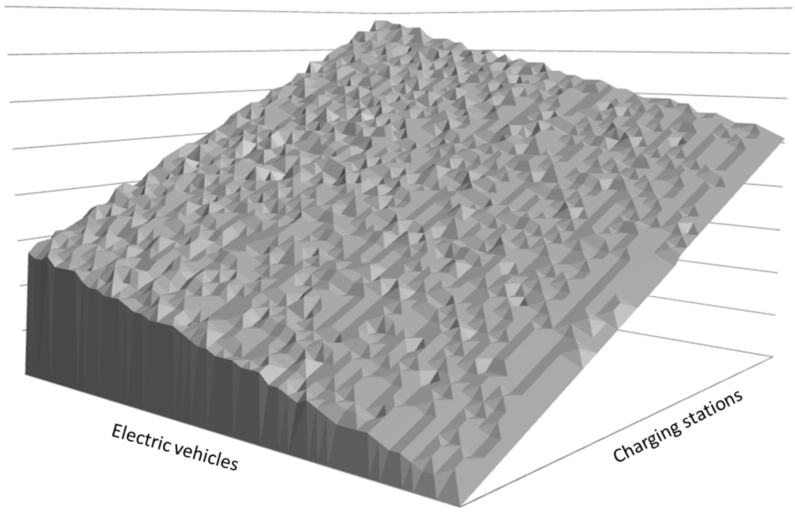

Figure 7.11.: Development of the mean number of exchanged messages based on the numbers of vehicles and charging stations.

Looking at general requirements, the application has the following features:

- **Privacy**: The only information that is shared is the interest to be charged. This seems uncritical and happens only on demand.
- **Scalability**: The approach is distributed. All calculations are performed independently at charging stations and do not depend on the number of vehicles in the network.
- **Communication effort**: The communication effort does not depend on the vehicles numbers. Furthermore, communication is done locally only.
- **Stability and Convergence**: As shown the queue length at charging stations is balanced.

8. Collaborative parking space sharing

Summary: *Parking is a key challenge in the context of urban road traffic. This chapter introduces two applications to address it: First, collaborative sharing of landlords' private parking spaces and second, trading parking spaces among individual landlords and other parties which manage parking spaces for many vehicles. Both applications show how to exploit given parking resources in a more efficient way.*
The development of the second application is result of joint work with Jia Yuan Yu, Wynita Griggs, Fabian Wirth and Robert Shorten. It has been published in [68].

8.1. The problem of parking

According to [151], about 30 % of today's inner-city road traffic is caused by vehicles searching for a parking spot. It was recently reported that over one year in a small Los Angeles business district, cars cruising for parking burned 47,000 gallons of gasoline and produced 730 tons of carbon dioxide [151]. Meanwhile, the consulting firm *McKinsey* recently claimed that the average car owner in Paris spends four years of his or her life searching for a parking space [42]. The parking problem associated with electric vehicles becomes even more acute. Due to the limited range of these vehicles, the marginal cost of expending energy to search for spaces may, in some cities, be prohibitively high. Thus, there is a real and compelling societal and economic need which puts extra pressure on the parking problem.

Increases in car ownership, inadequate public and private parking facilities, and synchronised demand (see Section 1.2), have led to serious (mainly temporal) mismatches in parking space supply and demand. Finding a parking space at certain times of the day is a non-trivial challenge. Furthermore, it is not only costly in terms of wasted time for the driver, but it also creates congestion and pollution affecting everyone. Thus, relaxing the parking problem improves both economic efficiency and environmental pollution.

Indeed, major companies and cities are responding to these challenges. For example, many city authorities, in order to safeguard resident on-street parking from people commuting to a city through the day, have initiated schemes such as resident permit parking, and a number of commercial initiatives have emerged in the parking area. *SFpark*[1], *Park2Gether*[2] and *ParkatmyHouse*[3] provide examples of city authorities and companies investing heavily in parking research and products within a smart city context.

Starting from the basic idea of these concepts (to share private parking spaces to realize a more efficient use of given spaces) two advanced applications are proposed. The first application described in Section 8.2 extend the idea of given platforms by letting landlords cooperate and rent out their spaces collaboratively to users. The second application incorporates private landlords with singular spaces, multi-spaces providers such as companies, users and employees. A feedback loop could be set up to establish a fair allocation of spaces to users.

8.2. Car park sharing of private landlords

The objective of given parking sharing platforms such as the ones named before is to use the number of given parking lots more efficiently. The basic assumption is that there are free parking spaces even if there are vehicles searching for available spots at the same time. That it is the case, when private parking spots could not be used by other vehicle drivers than the parking space owner or renter etc. Named platforms give parking spot owners the possibility to rent their spots to others for a given time window. They provide the functionality to connect (commonly individual) parking spot providers to users who require a space. Parking space providers could share information about the location, the renting time, price information etc. via these platforms. People who request a parking spot (potential *users*) could book a selected space for a defined time window. Administrative and payment service are commonly provided by the platform operator. But - until now - there is no link between the providers. This is the starting point for the first application.

[1]see `http://sfpark.org/about-the-project`
[2]see `https://www.park2gether.com`
[3]see `http://www.parkatmyhouse.com`

In order to increase the attractiveness of such parking space sharing systems, the following problem of available platforms is highlighted: vehicles, which have booked a parking space may find the space occupied even though it should be available according according the sharing agreement (given platforms ask users to call the platform provider in such a case, see FAQ from *Park2Gether*[4]). Focussing on the setting where landlords wish to use his or her parking space, two scenarios occur: (1) landlords arrive earlier than agreed and the space is (rightfully) occupied by the user and (2) the landlord arrives as agreed but the user has not freed the relevant space on time. In the survey outlined in detail in [135], it has been shown that potential users, who do generally support the idea to share their private parking spaces are afraid of exactly this scenario.

The basic idea to overcome this problem is to implement cooperation among individual private car park providers (landlords). It is proposed that these collaborate with other landlords in the vicinity (neighbours). Assume a landlord l_i owns a private parking space and has rented it to some user u_j on a particular day for some time window. When returning, the car space may still be occupied. In this case, the cooperation of landlords becomes beneficial as the owner may utilize available parking spaces from collaborating neighbours and park his or her car on their grounds. Assume landlord l_A arrives at home at time t and finds his or her space occupied by a user u_1. l_A will now search for an open space at collaborating neighbours. Assume the space of landlord l_B is free and l_A parks his vehicle at l_B's space.

The benefit of the system is reasoned by the scenario when the second landlord l_B returns home. Since landlord l_A occupies l_B's space, l_B will also search for a space at his or her neighbours - including the space of l_A. In some cases, u_1 has left l_A's space in the meantime, such that l_B could make use of it. In the conventional setup with a one-to-one relationship (Figure 8.1 (a)) between landlords and users, at least one of the two landlords l_A and l_B would not have had access to a parking space. With collaboration both landlords get a space. Generally, the same situation may occur for users when they arrive at a rented parking lot where the landlord (still) occupies his space. Figure 8.1 (b) illustrates the basic idea, which implements a many-to-many relationship between landlords and users.

The application could be deployed in many ways - it has no particular requirements to the technical sub-system. A simple solution would be provide a smartphone platform, where landlords and users exchange feedback on their timing constraints.

Problem definition and assumptions

The quality of service of car park sharing platforms should be increased by letting car park owners (landlords) cooperate. This should be discussed in the context of the

[4] `https://www.park2gether.com/faq#faq-4`

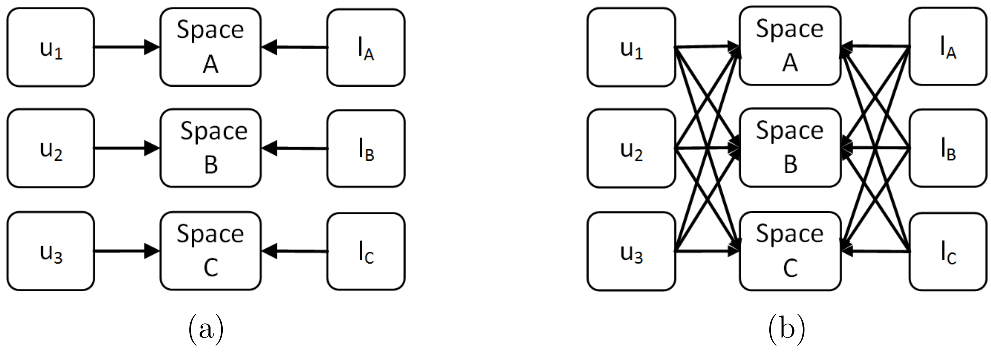

Figure 8.1.: Figure (a) illustrates the common setup for a parking space sharing system with a 1:1 relationship between landlords l_i and users u_j. Figure (b) shows the basic idea of the approach outlined with a many-to-many relationship between landlords and users.

problem which occurs when landlords arrive home and could not access their private parking space since a user has not yet vacated the space.

It is assumed that landlords offer their spaces to users in a time frame, which starts at t_s and ends at t_e. Thus, landlords plan to free their spaces during this time. Furthermore, it is assumed that the landlord provides this information some days in advance or defines a repetitive behaviour *like on every Tuesday the space is free from 9 AM to 5 PM*. However, from time to time, the landlord may not meet his own schedule, e.g. when being sick (according to [85] it is 17.6 days per year in Germany, which corresponds to about 10 % of working days). In addition, landlords may free their space later than scheduled or arrive back early for various reasons. Similarly, users of provided car spaces may arrive earlier than the space is provided or leave the space later than agreed.

8.2.1. Brief analysis

The benefit of the collaborative parking space sharing concept is briefly explained by comparing two scenarios: In the first scenario a user books a specific space from a specific landlord in a one-to-one relationship. In the second scenario, landlords collaborate and users book the right to use one of the parking spaces of collaborating landlords. Equal numbers of landlords and users are assumed. Then, the probability of the event that landlords need to find a parking spot on the street (as his or her parking space as well as those of neighbours are not accessible) shall be low. Thus, the probability corresponds to the probability that neither the landlord's space is available nor one of his neighbours'. It is assumed that the landlords' home coming times (arrival times) $T_{l,a}$ have a Gaussian distribution. The same applies to the departure time of users $T_{u,d}$, thus $T_{l,a} \sim N(\mu_a, \sigma_a^2)$ and $T_{u,d} \sim N(\mu_d, \sigma_d^2)$.

The one-to-one scenario is analysed first. The event of interest X (that landlords arrive home before the user has vacated the space) in this scenario occurs with probability $P(X)$ according to:

$$P(X) = P(T_a < T_d) \tag{8.1}$$

In order to analyse the complete set of possible events, cumulative distribution functions are regarded. Let F_a be the cumulative distribution of the arrival time of landlords and F_d the cumulative distribution of departure time of users. The departure time of users is equivalent with the time when the parking space is vacated. F_a is given by Equation (8.2), F_d by Equation (8.3).

$$F_a(T_a) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{T_a} e^{-\frac{1}{2}\left(\frac{t-\mu_a}{\sigma}\right)^2} dt \tag{8.2}$$

$$F_d(T_d) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{T_d} e^{-\frac{1}{2}\left(\frac{t-\mu_d}{\sigma}\right)^2} dt \tag{8.3}$$

It could be assumed that the expected time of the arriving users is commonly later than the expected value of the departing users. This assumption is meaningful for many reasons: the very initial idea is that a landlord rents his place to others during the time he or she is not requiring his private space (e.g. while he or she is at work); thus one may assume that the parking space is offered until the time that is (shortly) before the expected time of the landlord's return. This "buffer" time is called $t_{l,p}$. Furthermore, it is assumed that the user who rents the space expects to leave the parking spot before the time until when the space is rented. This additional buffer time is called $t_{u,p}$. It holds $\mu_a = \mu_d + t_{l,p} + t_{u,p}$. In practice the time buffer may comply to the commuting time, which means that the landlord as well as the user could start and end work at similar hours.
Now, suppose that $T_{l,a} \sim N(\mu_a, \sigma_a^2)$ and $T_{u,d} \sim N(\mu_d, \sigma_d^2)$. Then,

$$P(T_d > T_a) = P(T_d - T_a > 0) = 1 - P(T_d - T_a \leq 0) \tag{8.4}$$

By independence of T_a and T_d, $T_d - T_a$ is normally distributed as well [61], here

$$\frac{T_d - T_a - \mu}{\sigma} \sim N(0, 1). \tag{8.5}$$

with expected value

$$\mu := E(T_d - T_a) = \mu_d - \mu_a \tag{8.6}$$

and variance

$$\sigma^2 := Var(T_d - T_a) = \sigma_d^2 + \sigma_a^2 \tag{8.7}$$

Then, it is

$$P(T_d - T_a \leq 0) = P\left(\frac{T_d - T_a - \mu}{\sigma} \leq \frac{0-\mu}{\sigma})\right) = \Phi\left(\frac{-\mu}{\sigma}\right) \tag{8.8}$$

with Φ being the normal distribution $N(0,1)$. Thus it is

$$P(T_d > T_a) = 1 - P(T_d - T_a \leq 0) = 1 - \Phi\left(\frac{-\mu}{\sigma}\right). \tag{8.9}$$

A simple example is given in the following: a time window of 120 minutes is regarded, say between 4 PM and 6 PM - a plausible time window where users leave the space and landlords arrive at home. It is assumed that $T_d \sim N(\mu_d, \sigma_d^2) = N(50, 6)$ and $T_a \sim N(\mu_a, \sigma_a^2) = N(70, 8)$. This means that it is expected that the user of the landlord's parking space leaves the space at 4:50 PM on average (50 min after the time windows' beginning) and that the landlord arrives at home at 5:10 PM on average, both with according variances. Thus, the previously mentioned time buffer is 20 minutes long. Let $P(T_d > T_a)$ describe the probability that the landlord arrives before the user has left (such that the landlords could not use his or her parking space). According to Equation (8.9), it is:

$$P(T_d > T_a) = 1 - \Phi\left(\frac{-\mu}{\sigma}\right) \tag{8.10}$$

$$= 1 - \Phi\left(\frac{-(\mu_d - \mu_a)}{\sqrt{\sigma_d^2 + \sigma_a^2}}\right) \tag{8.11}$$

with

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{1}{2}t^2} dt. \tag{8.12}$$

In this example, $P(T_d > T_a)$ is 0.0228. When the time buffer decreases and/or variances increases, the probability increases as well. For instance, with $T_d \sim N(55, 6)$ and $T_a \sim N(65, 8)$, $P(T_d > T_a)$ becomes 0.1587. It is the same result if $T_d \sim N(50, 12)$ and $T_a \sim N(70, 16)$.

Thus, even in an ideal scenario, where a landlord and user share a parking spot, who have fitting time constraints and relatively low variances regarding their arrival and departure times, the probability that the landlord could not access his car park is not negligible.

The one-to-one scenario should now be compared to a scenario when landlords collaborate in the way it has been described earlier. For a comparability, the same assumptions are made but with two cooperating landlords and two users. There are the following situations where landlords could not access their parking space or the space from their neighbour (the other collaborating landlord) when coming home:

- Both users of both parking spots are still occupying the spaces. This probability is denoted $P_{u,u}$.
- One user is (still) occupying a parking space and the other landlord is occupying the other. This probability is denoted $P_{u,l}$.

$P_{u,u}$ is known from the previous analysis. It is $P_{u,u} = P(T_d > T_a)^2$. $P_{u,l}$ could also be inferred from the previous analysis. It is equivalent to the probability that one of the two users is (still) parking on one of the spaces (while the other has left already) multiplied by the probability that the other landlord comes home earlier.

Thus, the according probability $P_{u,u} + P_{u,l}$ is

$$P_{u,u} + P_{u,l} = P(T_d > T_a)^2 + P(T_d > T_a)(1 - P(T_d > T_a))(1 - \Phi(0)) \tag{8.13}$$

By applying the same example as before ($T_d \sim N(\mu_d, \sigma_d^2) = N(50, 6)$ and $T_a \sim N(\mu_a, \sigma_a^2) = N(70, 8)$, the probability is almost divided in half compared to the scenario without collaboration, it is 0.01165992. The benefit of the cooperation is increasing when the assumptions are relaxed, namely when the expected arrival and departure times of landlords and users have a greater difference and when according variances are greater. This will be shown in the following simulations.

8.2.2. Simulations: Collaborative parking networks vs. 1:1 allocations

In addition to the brief analysis outlined before, some simulations are performed. These are done to assess the performance of larger groups of collaborating landlords. By having multiple users and landlords, a uniform distribution of landlords' arrival times and users' departure times is assumed for a two hour time window T_w.

Furthermore, it is assumed, that landlords have an average commuting time of $\bar{t_c}$, which is uniformly distributed within the range $[max(0, \bar{t_c} - t_r), \bar{t_c} + t_r]$, where t_r defines the size of the range. This commuting time corresponds to the time buffer described in the previous subsection. Thus, the time window for the arrival time of landlords is shifted by $\bar{t_c} \pm \alpha \cdot t_r$, with $\alpha \in [0, 1]$ compared to the users' departure times.

Then, the probability $P(T_d > T_a)$ that the landlord arrives home while the user is still occupying the landlord's parking space is equal to:

$$P(T_d > T_a) = 1 - \frac{E(T_a)}{T_w}. \tag{8.14}$$

The following simulation is done over a time of 200 days (which corresponds to one year when looking at working days only) with 100 landlords and 100 users. The commuting time is individually set for each landlord - a uniform distribution over a range from 5 to 15 minutes serves as basis. In this setup, the analytically calculated probability for a 1:1 relationship between users and landlords is $0.58\bar{3}$ (using Equation (8.14)). The dashed line in Figure 8.2 determines this reference probability $P(T_d \leq T_a) = 1 - P(T_d > T_a)$. Furthermore, the Figure shows the average probability for landlords to find a spot either on the own parking space or on one of collaborating neighbours for varying numbers of collaborating neighbours.

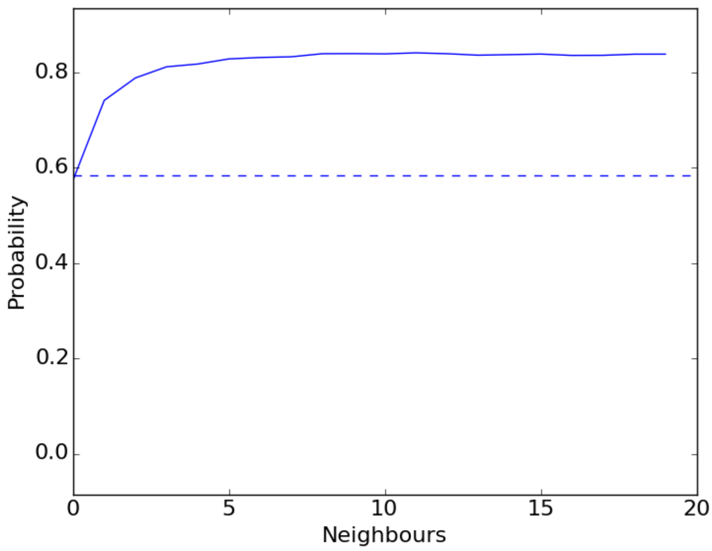

Figure 8.2.: Probability for landlords to get a parking space at own soil or from a neighbour. The dashed line determines the probability to get a parking space in a 1:1 relationship between landlords and users.

One may question, why the probability to find a parking converges to a value around 0.8 (and not 1). The basic idea of the application is that a landlord A could utilize parking spots from collaborating landlords (neighbours) - here landlord B, in case one users is (still) occupying A's parking spot. Furthermore, when B comes home (assuming B arrives later than A), his or her space is occupied by A. The core idea of the approach is that there is the chance that the user who occupied the space from A has departed in the meantime, such that B could make use of it.

However, it could be the case that there are more (collaborating) landlords arriving at home than users departing. In this scenario, there are landlords who have neither access to their own space nor to one of their neighbours' spaces. Figure 8.3 shows the result of a simple simulation with 100 landlords and 100 users. Both landlords and users depart according to a uniform distribution between minute 0 and 120. Landlords have a commuting time of 5 minutes. Figure 8.3 shows that there are times where there are more landlords arriving than users departing, resulting in landlords who do not get access to any parking spot.

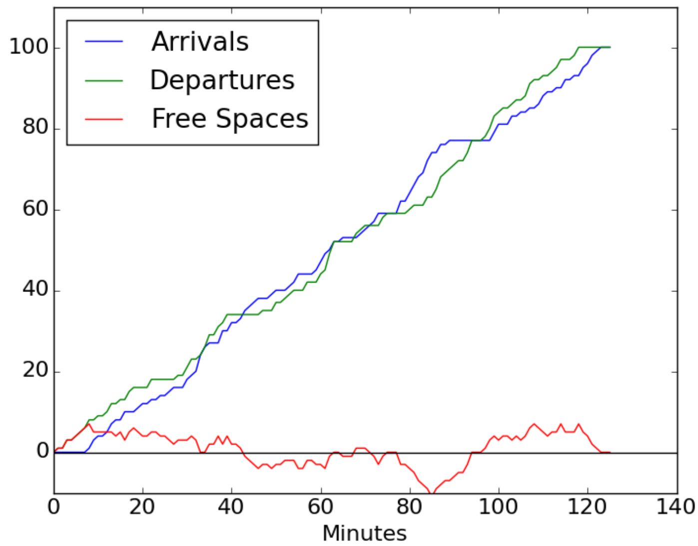

Figure 8.3.: Number of user departures and landlords arrivals constitute the maximum number of free parking spaces in the system.

Figure 8.4 shows the impact of the average commuting time. In this setup, the number of cooperating neighbours is assumed to be fixed, here six. From Figure 8.2 it could be seen that the benefit of having more collaborating neighbours than that is not beneficial anyway. Figure 8.4 shows the result of the simulation. As expected, the probability to get a space increases with the commuting time. The dashed line illustrates the equivalent probability P_{ref} for the reference scenario. It is calculated according to Equation (8.15).

$$P_{ref} = \begin{cases} \frac{1}{T_w+1} \sum_{t=0}^{T_w} \frac{t+t_c}{T_w} & \text{for } t + t_c \leq T_w \\ \frac{1}{T_w+1} \sum_{t=0}^{T_w} 1 & \text{for } t + t_c > T_w \end{cases} \tag{8.15}$$

8.2.3. The impact of reserved parking spaces

Another way to increase the probability to get a space (both at the landlord's own space or at one of his cooperating neighbours) could be achieved by reserving some spaces for landlords. That means that some spaces are not rented to users at all. Figure 8.5 shows the impact of the share of reserved spaces on the probability of landlords to get access to some parking space (to his or her own or to one of his or

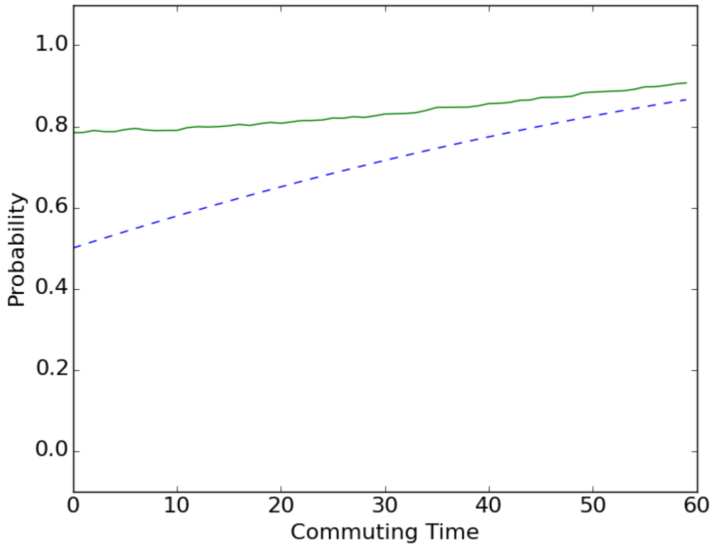

Figure 8.4.: Probability for landlords to get a parking space at own soil or from a neighbour having different commuting times.

her neighbours' private spaces). The simulation is done with an average commuting time (of landlords) of 10 minutes. Figure 8.5 shows that the impact of reserved parking spaces on the quality of service - in the sense of access to a parking space at the landlord's or his or her neighbours' spaces - does not increase much.

8.3. Trading parking spaces for better QoS

The second application incorporates operators of multiple parking spaces such as companies, shopping centres, university campuses (all summarized as mini-cities) but also (private) landlords. The scenario targets suburban environments where mini-cities and residential areas are co-located. In such a set-up, it could be assumed that the mini-city is short of parking spaces during particular time (i.e. business hours) while many nearby residents vacate their private spots at the same time. Thus, in principle there is an opportunity for the mini-city to use these vacated parking spaces of local landlords in the vicinity. In the context of the application introduced here, these parking resources are used in a trading mechanism implemented between a mini-city and private landlords in the vicinity, which helps to utilize resources in a more efficient way while preserving some defined quality of service.

Assume there is a mini-city with a total number of N parking spaces. It is surrounded by M private parking spaces from landlords, as illustrated in Figure 8.6. In the given application, it is assumed that landlords cooperate with the mini-city in a way that landlords lease their driveways to the mini-city during business hours while the land-

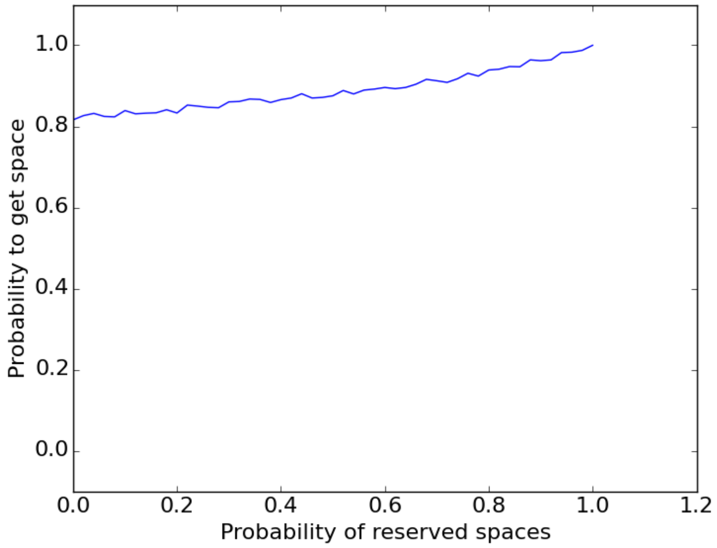

Figure 8.5.: Based on the share of reserved spaces (spaces not rented to users), the probability of landlords who get access to a space (on own grounds or at cooperating neighbours) increases.

lords expect to be away (e.g. at work). The time period in which the spaces from landlords are leased is $[0, W]$.

The application is discussed within the three subsections. In Section 8.3.1, the basic features of the system are regarded, name how to define proper numbers (N, M) of different parking spot groups such that the system implements some specific quality of service to mini-cities and landlords. In Section 8.3.2 shows how to establish a fair allocation of parking spaces to its users. Simulations are supporting the statements from these two subsections in Section 8.3.3.

8.3.1. Providing QoS to landlords in the the home-early event

The first question to address is how $Q \in \{0, 1, \ldots, N\}$ should be dimensioned in order to guaranty sufficient QoS to landlords for the home-early event (see below).

Assume there are M landlords who offer parking spaces. For each landlord a non-negative random variable $T^l_{i,a}$ is defined, which denotes the time at which he or she returns to his or her parking spot i. It is independent and identically distributed. Under normal circumstances, $T^l_{i,a}$ is greater than W ($[0, W]$ denotes the nominal rental window for every parking space). However, landlords may choose to come home earlier from time to time.

Further, it is assumed that each parking space has exactly one user per day; a non-

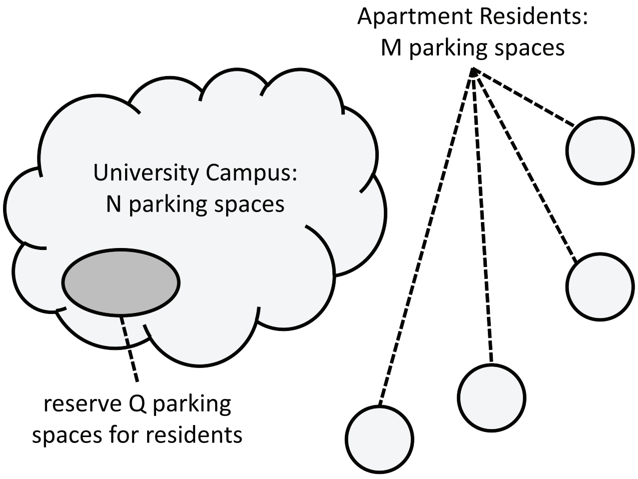

Figure 8.6.: The parking scenario: premium spaces are those on the university campus, whereas secondary spaces are those belonging to the apartment buildings.

negative random variable $T^u_{i,d}$ is defined, which denotes the departure time (from parking space i) of this user. Again, it is independent and identically distributed. In parallel to the assumption for landlords, under normal circumstances the user vacates the space on time, which implies $T^u_{i,d} < W$. Again, on rare occasions it is considered that this is not the case.

The equivalent problem as discussed in the previous parking space sharing application (Section 8.2) is regarded: A landlord gets home but his parking spot is not (yet) vacated be the user. Recall that there are two events which lead to this problem: The *home-early* event E_i is due to the landlord needing the space during the day (while the user of the space is still there). The *overstay* event O_i is due to the fact that the daytime user overstayed (and the landlord arrives home before the user has vacated the space). The two events are defined as follows:

- Home-early event: $E_i \triangleq \{T^l_{i,a} \in [0, W]\} \cap \{T^l_{i,a} < T^u_{j,d}\}$
- Overstay event: $O_i \triangleq \{W < T^l_{i,a}\} \cap \{T^l_{i,a} < T^u_{j,d}\}$

It is assumed that $T^l_{i,a}$ and $T^u_{j,d}$ are mutually independent.

The first step is to quantify the probability of O_i: $p(O_i) \leq p(T^u_{j,d} > W)$. It is the probability of users staying longer than agreed.

Let $F_{T^u_d}$ denote the probability distribution of $T^u_{i,d}$ (the time when the user vacates the space i) and let $F_{T^l_a}$ denote the probability distribution of $T^l_{i,a}$ (the time when the landlord requests the space i). Under the independence assumption, it is:

$$p(O_i) = \int_{t=W}^{\infty} (F_{T^l_{i,a}}(t) - F_{T^l_{i,a}}(W)) \, dF_{T^u_{i,d}}(t). \tag{8.16}$$

Thus, the probability $p(O_i)$ of the overstay event at a specific parking space i is equivalent to the probability of arrived landlords' after the agreed time window while the relevant parking space user is still on the space (given by $F_{T^u_{i,d}}(t)$). Equation (8.16) is deduced from the $p(O_i) = p(W < T^l_{i,a} < T^u_{j,d})$. The proof is part of the publication [68].

In the following, the formula to define a proper dimensioning of the reserve parking spaces Q from N for the home early event E_i (where $T^l_{i,a} \leq W$) is developed where some specific QoS met. The QoS is defined by the probability $p(M,Q)$ of the event that more than Q spaces are needed to accommodate all landlords who come earlier than agreed (the event E_i):

$$p(M,Q) = p\left(\sum_{i=1}^{M} X_{E_i} > Q\right), \tag{8.17}$$

where X_{E_i} denotes a Bernoulli random variable taking value 1 when event E_i occurs, and value 0 otherwise. We first characterise the probability of the event E_1 in terms of the probability distributions of $T^l_{i,a}$ and $T^u_{j,d}$.

The next step is to quantify the probability of the event E_i. Let $F_{T^l_a}$ denote the probability distribution of $T^l_{i,a}$ and $F_{T^u_d}$ of $T^u_{i,d}$. Under the assumption that $T^l_{i,a}$ and $T^u_{i,d}$ are independent and identically distributed, it is

$$p\left(E_1\right) = \int_{t=0}^{W} (F_{T^u_d}(W) - F_{T^u_d}(t)) dF_{T^l_a}(t) + F_{T^l_a}(W)(1 - F_{T^u_d}(W)). \tag{8.18}$$

Equation (8.18) is deduced from $p(\{T^l_{i,a} < T^u_{i,d}\} \cap \{T^l_{i,a} \in [0,W]\})$. The proof is given in the publication [68].

The probability that Q from N reserve spaces at the mini-city is not enough to serve the landlords who come home early (event E_i) is calculated by:

$$p(M,Q) = \sum_{k=Q}^{M} \binom{M}{k} \phi^k (1-\phi)^{M-k} \tag{8.19}$$

with $\phi \triangleq p\left(E_1\right)$. Equation (8.19) is straight forward. It sums up the probabilities for all events where Q is not sufficient. The binomial coefficient is included to account for permutation of the two possible incidents: (1) the k landlords who do not have access to their parking spots in the home-early case (the probability is ϕ^k) and (2) the $M-k$ landlords who have (the probability is $(1-\phi)^{M-k}$)). The sum starts with Q as Q landlords will have access to the reserve parking spots - a scenario which is not targeted here. Under the assumption that $T^l_{i,a}$ and $T^u_{i,d}$ are independent and identically distributed, $p(M,Q)$ is a random variable characterised by ϕ. The probability can be mitigated by increasing the parameter Q. Given a QoS target $p(M,Q)$, Equation (8.19) can be used to determine the corresponding value of Q needed to achieve it.

8.3.2. Fair and optimal allocation of premium spaces to mini-city users

One may assume that users (i.e. employees of the mini-city) who require parking at the mini-city have monthly subscriptions, which allow to park at the mini-city. However, such a subscription does not necessarily guarantee a free parking space since commonly more subscriptions are sold than there are actual spaces. That is, if users arrive "too late" and the mini-cities parking spaces is already full.

Such "first come, first served" systems can be inefficient. Consider the example of subscribers who have certain duties which lead to later arrival times (e.g. parents who need to take their children to school). In the following part, an approach is presented that establishes a fair access to parking spaces of the mini-city for all users over a long-term period. In this context, fairness means that all users should have the same chance to get access to one of the $N - Q$ premium parking space at the mini-city (independent from the arrival times).

Assume that the mini-city has been able to obtain enough parking spaces from landlords such that $(N - Q) + M$ is greater than or equal to the number of mini-city members who might require a parking space. The approach is similar to the AIMD optimisation approach exploited in the *SPONGE* application described in Section 4.5. The basic idea was first introduced in [183]. For simplicity of exposition, it is assumed that $Q = 0$ here, but the results generalise in a straightforward fashion.

Fair access to parking spaces could be implemented in many ways. Here, it is assumed that each user i has a cost function $f_i : [0, 1] \to \mathbb{R}$, which depends on $z \in [0, 1]$ where z denotes the premium space allocation frequency of the relevant user. In other words, z defines the probability that a user gets a premium parking space at the mini-city. In the case where users' cost function is equal, a fair allocation would result in the scenario where all subscribers would have the same chance to get access to a premium space.

Generally, the cost function could be used to represent various applications to respect different needs of users, i.e. it might represent the amount a user is willing to pay or depend on the vehicle's emission class. Given individual cost functions, the aim is to design a system that achieves overall minimal cost for users. Thus, the optimal allocation of resources is:

$$\underset{z_1,\dots,z_n \in \mathbb{R}}{\text{minimize}} \sum_{i=1}^{n} f_i(z_i) \tag{8.20}$$

$$\text{subject to} \sum_{i=1}^{n} z_i = N, \tag{8.21}$$

$$z_i \geq 0, \quad i = 1, \dots, n.$$

The objective in the next step is to find the optimal allocation per user (z_i^*). The proposed algorithm for solving the parking allocation problem can be summarised as follows: It is assumed that - for each day - each user is allocated to one of the premium spaces with a certain probability. In order to realize a distributed architecture, it is again implemented a way that each vehicle tosses a coin to decide whether it shall have access to a premium spot or not. This is done on a daily basis. The according probability for the probabilistic assignment is calculated on every subsequent day individually for each user i:

$$p_i(t+1) \triangleq p(X_i(t+1) = 1) = \Gamma(t)\, \frac{\overline{X}_i(t)}{f_i'(\overline{X}_i(t))}. \tag{8.22}$$

$X_i(t)$ denotes the event whether the algorithm allocates a premium parking space to user i ($X := 1$) or not ($X := 0$). Furthermore $\overline{X}_i(t)$ denotes the average allocation to a premium parking spot for the i'th user up to the t'th day. $\Gamma(t)$ is a common value for all vehicles and it is broadcasted by the mini-city. It is based on past utilisations. Here, $\Gamma(t)$ is chosen such that $p_i(t) \in (0,1)$ for all $i = 1, \dots, n$ and all $t \in \mathbb{N}$. It is determined in a time-varying manner as it also influences the demand for premium spaces. The following example illustrates the performance of the algorithm.

Assume that all cost functions f_i are continuously differentiable and strictly convex. Thus, the optimal point $z^* \in \mathbb{R}^n$ is unique. Furthermore, it is assumed that the optimal point z^* has only positive entries. This assumption guarantees that the algorithm is well defined for every user.

The objective is to control access to premium spaces in such a way that the average utilisation for each user approaches the optimal value z_i^*; for large t it should be

$$\overline{X}_i(t) \approx z_i^* \tag{8.23}$$

with $\overline{X}_i(t) = \frac{1}{t+1}\sum_{j=0}^{t} X_i(j)$ subject to the (loose) capacity constraint $\sum_{i=1}^{n} X_i \approx N$. That is, all premium spaces are occupied on average. Again, privacy of users should be preserved. That means that $\overline{X}_i$ and f_i should not be revealed to any other users during the course of the optimisation. The algorithm presented here extends the ideas of [183], as it should be ensured that the utilisation variables $X_i(t)$ sum to N, or at least to a value close to it. Moreover, the resource to be allocated in this setting is atomic as opposed to arbitrarily divisible. These differences require changes to the algorithm presented in [183].

At each time t, each user i determines a probability $p_i(t)$ and sets

$$X_i(t+1) = \begin{cases} 1 & \text{with probability } p_i(t+1), \\ 0 & \text{with probability } 1 - p_i(t+1), \end{cases} \tag{8.24}$$

where all users make this probabilistic choice independently of other users or previous decisions. The evolution of the probabilities is governed by Equation (8.22). Note that each user i can determine its own probability with the exclusive knowledge of its own past utilisation $\overline{X}_i(t)$ and his or her (private) cost function f_i. If at a certain time t each $p_i(t)$ is fixed then the expected utilisation of the premium spaces is

$$E\left(\sum_{i=1}^{n} X_i(t+1)\right) = \sum_{i=1}^{n} p_i(t+1) = \Gamma(t) \sum_{i=1}^{n} \frac{\overline{X}_i(t)}{f_i'(\overline{X}_i(t))}. \tag{8.25}$$

To ensure optimal utilisation of the premium spaces and avoid overbooking, the expected utilisation should be (slightly) below the given number of premium parking spaces. Denoting this number by $N_E \leq N$, $\Gamma(t)$ is adjusted so that the expectation in Equation (8.25) tracks N_E. As the expectation is unknown, one may use the observed utilisation as an estimator for this. Taking a simple error regulation approach,

$$\Gamma(t+1) = \Gamma(t) + \alpha \left(N_E - \sum_{i=1}^{n} X_i(t) \right). \tag{8.26}$$

The overall system is now prescribed by the dynamics of X_i given by Equation (8.24), the dynamics of p_i given by Equation (8.22) and the dynamics of Γ as described in Equation (8.26).

8.3.3. Simulations of fair allocation

In the following, simulations show the fair allocation of parking spaces based on the approach outlined in the previous Subsection 8.3.2. The simulation has the following features: there are 900 users who compete for 450 premium parking spaces. Assume that each evening, users are assigned to a parking space as described above (with the scalar $\Gamma(t)$ determined using a PI controller). For simplicity, users have one of three arbitrary strictly convex cost functions: $f_1(z) = z^4/4$, $f_2(z) = z^6/6$, and $f_3(z) = z^8/8$.

Figure 8.7 shows the average premium space allocation $\bar{X}_c(t)$ achieved for each class c of vehicles. It could easily been seen that the values converge to different values per class ($\bar{X}_1(t) \approx 0.36$, $\bar{X}_2(t) \approx 0.52$, $\bar{X}_3(t) \approx 0.62$).

Furthermore, Figure 8.8 shows that the average utilisation of premium spaces converges to the target utilisation of 450.

Figure 8.9 shows the evolution of the controller gain, which is based on past utilisations. The figure shows that the gain converges. It is determined with the PI controller.

Finally, Figure 8.10 shows that the cost function derivatives do actually achieve consensus asymptotically.

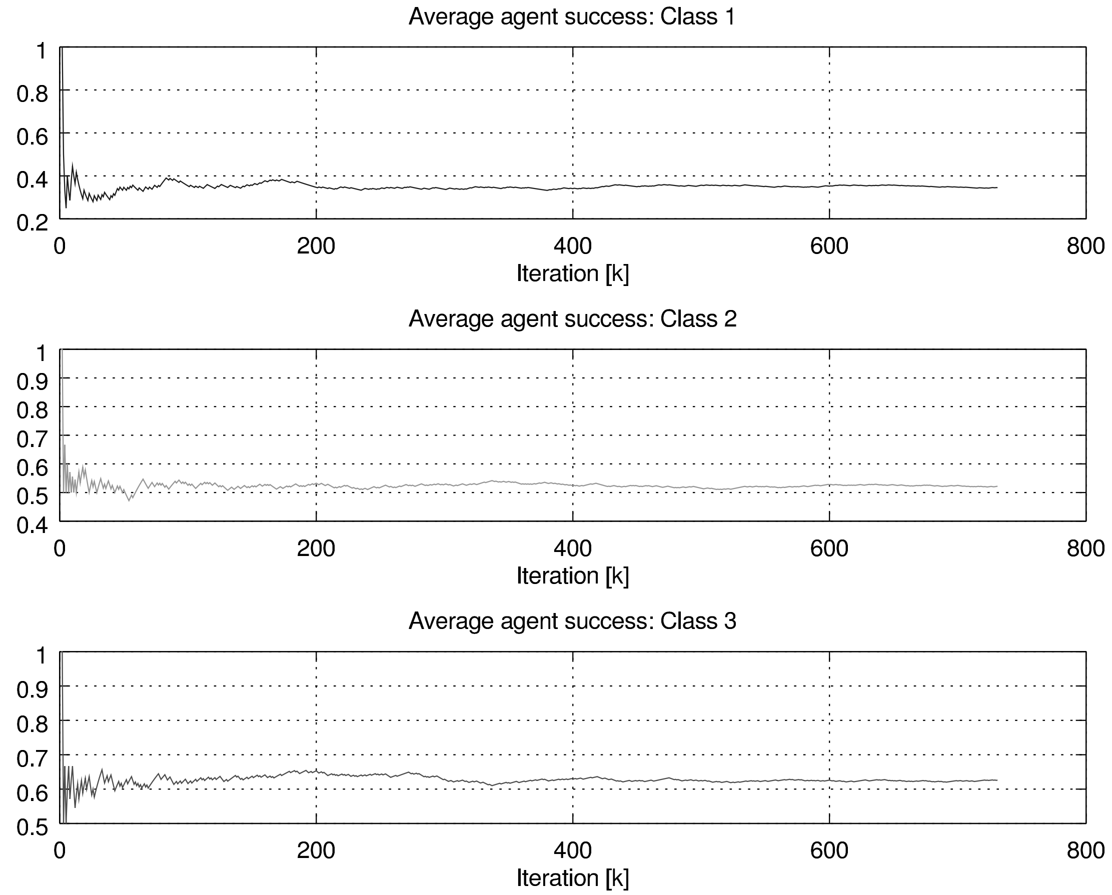

Figure 8.7.: Instantaneous allocation $\bar{X}_i(t)$ for three different vehicle classes. It could be seen that the probabilities to be assigned to a premium space converge to different values per vehicle class. It could be shown that the associated costs - calculated with help of individual cost functions f_1, f_2 and f_3 - converge to a common value.

Further discussion on the heuristic which suggests that the algorithm works in a general setting is outlined in [68].

8.4. Remarks

Two implementations of a collaborative parking sharing theme have been introduced. The first one introduced an extension to common car sharing themes by letting landlords collaborate. The approach is introduced in Section 8.2. The second application targets a the special scenario of a mini-city (e.g. a company with many employees) which trades parking spaces with landlords in the vicinity. Features of the two approaches are briefly discussed in the following.

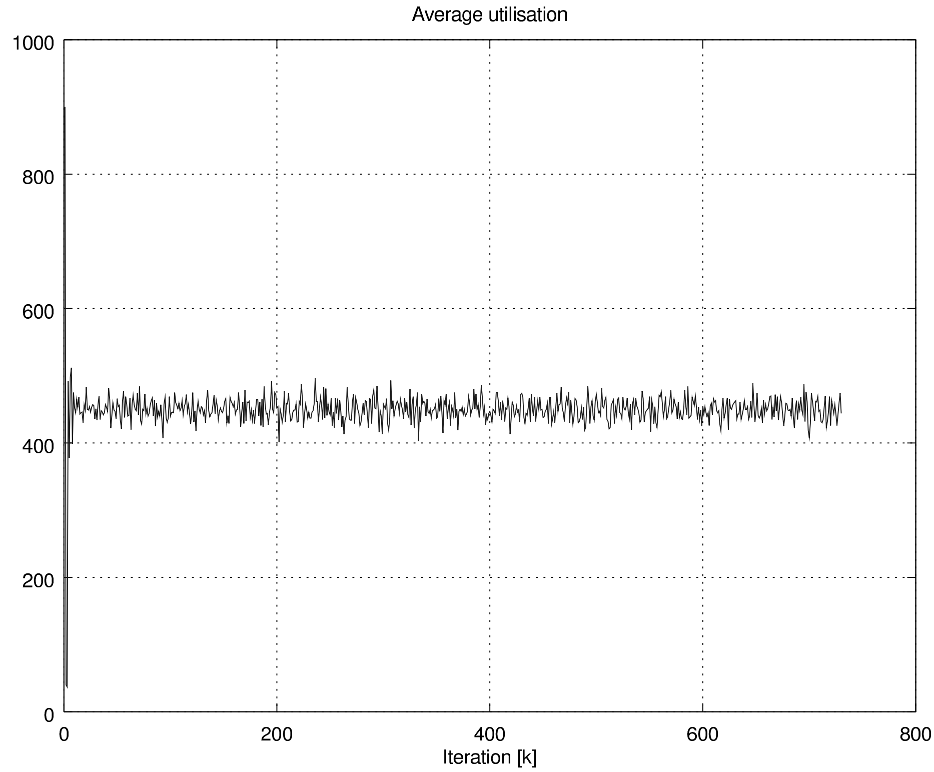

Figure 8.8.: The figure shows the average utilisation of premium spaces over time. The horizontal axe defines the simulation time, the vertical axe the average number of utilizes premium spaces. It is concentrated around the target value of 450, which corresponds to the number of given premium parking spaces in the simulation.

8.4.1. Car park sharing with private landlords

The first application is simple extension of given parking sharing platforms. Landlords with private parking space cooperate with other landlords and rent out their spaces collaboratively. In contrast to other applications no feedback loop is implemented. Features of the application are:

- **Privacy**: Users do not communicate with other users or an authority. It is not tracked who or when a private parking lot is used.
- **Scalability**: Scalability is limited as the number of parking spots in a area (which is in the vicinity of a landlord's home) is restricted. I.e. landlords would not accept alternative parking spaces which are miles away.
- **Communication effort**: The approach does not make use of any communication capabilities.
- **Stability and Convergence**: No feedback loop is implemented. Thus, stability and convergence is not relevant in this application.

Topics in the context of user acceptance are further discussed in Section 9.4.2.

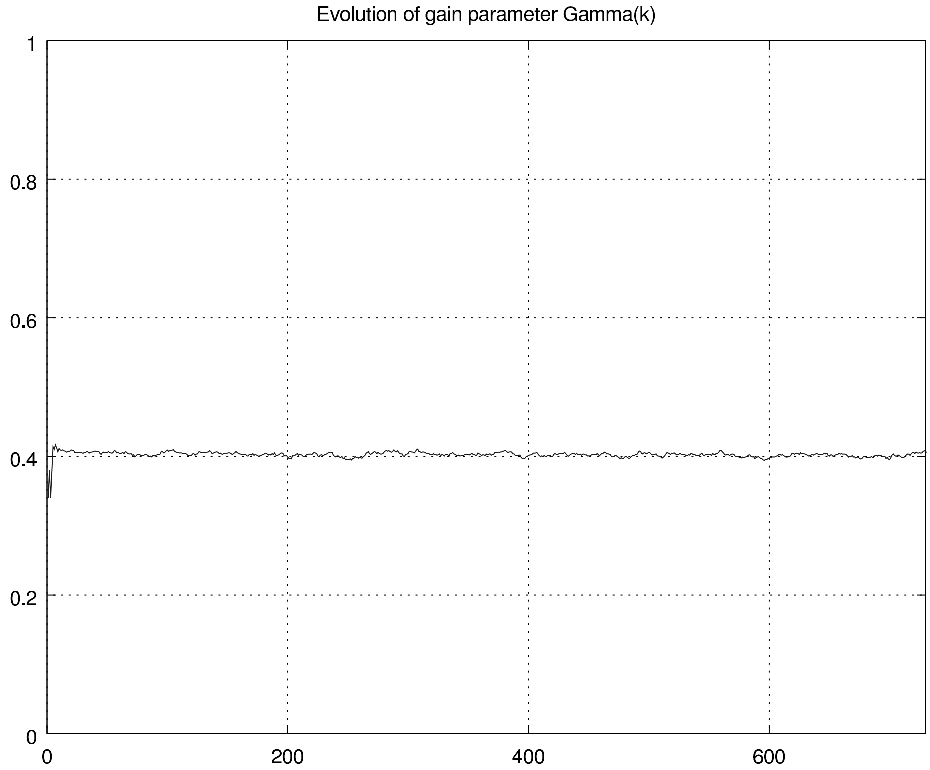

Figure 8.9.: The figure shows the evolution of $\Gamma(t)$ over simulation time. The horizontal axe defines the simulation time, the vertical axes the gain's value. It shows that the gain, which is determined with the PI controller, converges.

8.4.2. Trading mini-city's parking spaces for better QoS

The second parking application connects operators of multi-space car parks (mini-cities), its users and landlords in the vicinity. In order to implement fairness it establishes a feedback loop that leads to fair and optimal allocation of users to premium parking spaces of mini-cities. The proposed allocation is a general scheme, which can be applied to many problems of the form of Equation (8.20) (i.e. assigning spaces in overhead bins on passenger planes, seats in trains, etc.). At the same time, the application meets most of the previously outlined requirements:

- **Privacy**: Users do not communicate with other users, nor do they reveal any state information or cost information to a central authority (here the mini-city).
- **Scalability**: The number of vehicles does not affect the complexity or performance of the system as all calculations are done in a distributed way.
- **Communication effort**: Communication is limited to broadcasts of a network-wide constant, which is updated once a day.
- **Stability and Convergence**: It has been shown by simulations that the average success to get a premium space converges for each users with equal cost

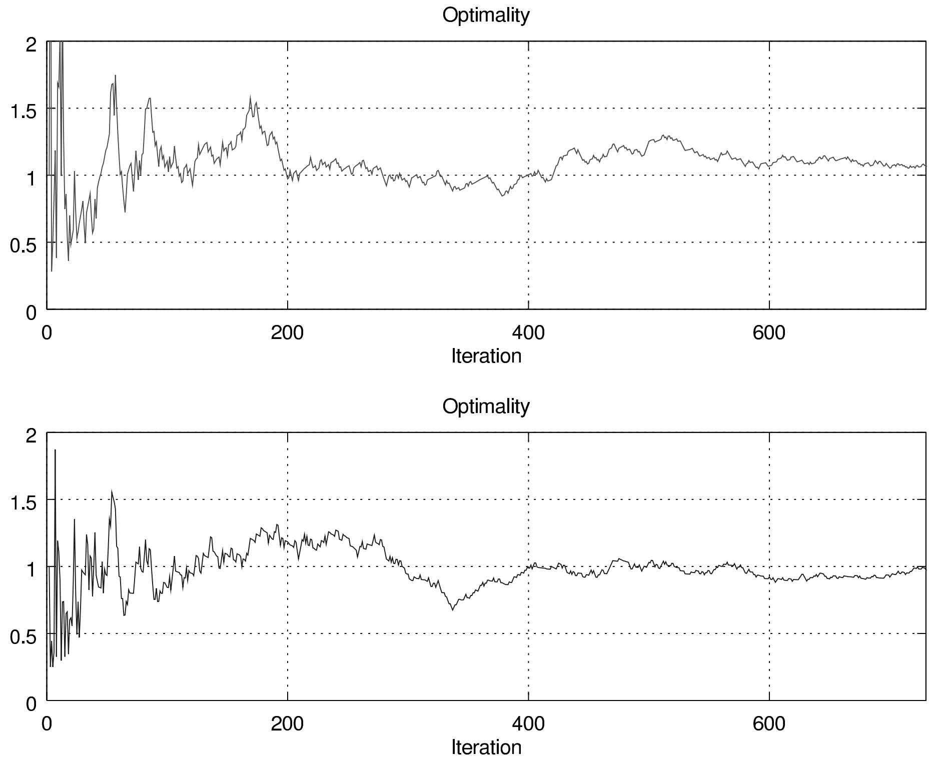

Figure 8.10.: Convergence to optimality. The cost function derivatives achieve consensus asymptotically. The upper Figure shows the quotient from the first vehicle class cost function derivatives and the second class. The lower plot shows the quotient from the second vehicle class cost function derivatives and the third vehicle class cost function derivatives. As desired, both quotients are around 1, which means that the relevant cost function derivatives converged to the same value.

functions.

However, it needs to be kept in mind, that the analysis of the QoS to landlords addresses the home-early event only.

Part III.

Testing and Conclusion

9. Testing and acceptance of applications

Summary: *This chapter discusses user acceptance of applications introduced in the previous chapters and introduces a testing framework dedicated for collaborative ITS applications with a large number of participating vehicles. The framework support both aspects: technical and user acceptance testing.*
Important parts of the chapter are results of joint efforts with other authors. Selected work on testing environments for ITS outlined in Section 9.1 was jointly conducted with Andrea Tomatis, Markus Miche, Massimiliano Lenardi, Thomas Bohnert and Ilja Radusch, see [171]. The work on the platform which integrates a traffic simulator with real-world vehicles (Section 9.2) was conducted together with Robert Shorten, Wynita Griggs, Rodrigo Ordóñez-Hurtado, Kay Massow and Emanuele Chrisostomi and has been published in [69] and [78]. The prototype vehicle itself was setup together with Robert Shorten, Arieh Schlote, Thomas Hecker and Astrid Bergmann. Moreover, selected parts of the user acceptance discussion outlined in Section 9.4 are inspired by joint work with Bartek Rytt [136, 137] and Kevin Rhinow [135].

9.1. Testing and acceptance of collaborative ITS applications

In the previous chapters, ITS applications were introduced which depend on the integration and collaboration of many vehicles and drivers. Simulations were conducted to validate the basic performance of the algorithms. However, such simulations cannot accommodate for all of the complexities, uncertainties, technical issues, driver attitudes and responses that might arise in the real world [170]. But there are other evaluation environments for ITS applications. All of them have with particular capabilities, strengths and weaknesses.

Testing environments range from full simulations (see Section 9.1.2), where all aspects are simulated, to field operation tests (see Section 9.1.1), where the developed system is assessed in the environment where it will be deployed to in the final step. This section provides a brief overview on selected evaluation environment. In addition to the aforementioned this includes a hardware-in-the-loop (Section 9.1.3) and a driving simulation setup (Section 9.1.4) as well as a brief overview on the capabilities of user surveys and interviews (Section 9.1.5).

9.1.1. Field operational tests

Testing ITS applications in field operational tests (FOT) employing a large number of vehicles leads to most meaningful verification results as the system under test are embedded into a realistic environment. Accordingly, FOTs should be the final validation environment of ITS developments [66]. At the same time it (1) comes with the most challenging (technical and organisational) requirements and (2) has also some functional limitations.

Technical challenges are related to the complex test setup, which requires a dedicated test framework with test operators who monitor tests, a test data logging and management system, dedicated on-board systems and so forth. Key outcome of the funded FOT research projects PRE-DRIVE C2X, DRIVE C2X and sim^{TD} has been such a testing infrastructure [171].

The functional limitations mainly relate to the limited knowledge about and control of real-world features. In effect, test results may always be affected by features of the system, which are not (or not precisely) known. That makes it also complicated (if not impossible) to reproduce tests (and results). Furthermore, full control about the environment is not possible. An example are traffic conditions on public roads, where the influence is limited.

The most critical aspect of FOTs is related to the required (monetary and organization) efforts. FOTs require enormous (monetary) efforts, especially when hundreds of vehicles are needed to demonstrate general impacts on urban traffic flows, air pollution etc. Furthermore, FOTs require access and control of traffic infrastructure (such as traffic lights) but also permissions from relevant authorities. In most cases this is impractical and prohibitively expensive - even large, funded FOT projects such as DRIVE C2X and sim^{TD} (where there were more than one hundred vehicles involved) had to upscale results.

However, all concepts outlined in this thesis assume a huge number of participating vehicles. Therefore, FOTs are practically impossible to conduct.

9.1.2. Full simulation

In a fully simulated environment the complete environment of the system under test (here the ITS application) is simulated. That may include vehicle dynamics, communication links and partners, sensors, actuators, the traffic network and infrastructure etc. including all dependencies. Even though the setup of a simulation could become very complex, the main advantage is the (theoretical) full control of all system's parts and the reproducibility of tests.

In the validation chain, the full simulation has two main objectives: First, it is suited to pre-evaluate certain aspects of the implementation (e.g. correct interface implementations) in order to minimize the extent of costly tests in a real-world environment. Second, the simulation is suited to run tests, which could not even be evaluated during the FOT, e.g. traffic efficiency effects with thousands of vehicles.

However, a key limitation of the full simulation is the lack of a realistic driver behaviour and its impacts.

9.1.3. Hardware-in-the-loop

A hardware-in-the-loop (HIL) setup connects an (embedded) system using existing interfaces with a HIL-simulator which mimics the real environment of the system. Thus, a HIL system is a combination of physical and virtual prototype. The authors in [51] define the HIL approach as "a setup that emulates a system by immersing faithful physical replicas of some of its subsystems within a closed loop virtual simulation of the remaining subsystems". The definition highlights the closed loop character of a HIL: the (embedded) system is fed with data which emulates the real-world environment. It responses and interacts with the emulated environment. The main benefits are outlined in [51]: cost effectiveness, rapid prototyping, fidelity and verisimilitude[1], simulation speed, repeatability, non-destructive nature, comprehensiveness, safety and concurrent systems engineering.

HIL simulators are widely used in the automotive industry - especially in the context of testing and validating ECUs [9] and other components. The latest HIL-platforms also aim at driver assistant systems. One example for these is the VEHIL (vehicle hardware-in-the-loop [66]) platform, which aims at testing Advanced Driver Assistant Systems (ADAS), such as adaptive cruise control, stop-and-go, forward collision warning, pre-crash systems, blind spot systems, etc. The HIL system *X-in-the-loop* is

[1] "Verisimilitude is a philosophical concept that distinguishes between the relative and apparent (or seemingly so) truth and falsity of assertions and hypotheses." [181]

a development from Karlsruhe University, which could embed a complete vehicle as a system under test into a larger traffic scenario (e.g. crossings). The vehicle is "driven" on a roller test bench. The HIL simulator *VIL* (Vehicle-in-the-loop) developed within an dissertation for AUDI from Bock [18] follows a virtual reality approach to assess driver behaviour in the context of ADAS. The test could be performed within public space on open roads. The basic idea of the *VIL* approach is similar to the one developed within the context of this thesis and introduced later.

9.1.4. Driving simulation

There are many real-time driving simulators on the market - in particular for research and training. In the context of testing and evaluating ITS applications, they are mainly used to assess driver behaviour and the acceptance of Human Machine Interface (HMI) design. All driving simulators aim at providing an environment for test drivers which is as realistic as possible. Implementations range from simple graphics on a computer monitor to high performance driving simulators with more than six degrees of freedom, 360 degree visualizations etc. An example for an advanced driving simulator is the one from University of Stuttgart [10].

The main benefit of driving simulators is their versatility. They can be easily configured to assess a variety of questions in the context of human factors research. However, it is impossible to mimic the real world in all its complexity. Thus, there will always be the issue of validity, i.e. to what extent behaviour in a simulator corresponds to that in the real life. A handicap of such setups is that many users become sick in those simulators [117]. Another constraint for using advanced simulators are their costs, i.e. the aforementioned driving simulator from the University of Stuttgart was acquired for more than 7 million Euro [10].

9.1.5. Survey and interviews

The question of user acceptance is most often tested with help of interviews and questionnaires. Many research projects make use of this method like PRE-DRIVE C2X [11], TEAM[2] and MyWay[3][129]. A main benefit is that it could be performed without the need for a working prototype. Furthermore it is cost efficient. It is highly efficient when assessing driver opinions and acceptance.

However, results from surveys and interviews where users have not experienced the proposed solution are not very precise.

[2]See `www.collaborative-team.eu`
[3]See `myway-project.eu`

9.2. Low-cost HIL simulation platform

Considering the aforementioned evaluation methods in the given context (the proposed applications described within this thesis), a HIL platform and additional user studies are best suited. Given HIL approaches in the automotive domain mainly address the assessing of non-collaborative ADAS applications, where the aggregate effect of many vehicles is not targeted. Given HIL solutions outlined in Section 9.1.3 are either not suited to assess the introduced applications or come with relatively high burdens to integrate. Therefore, a light-weight and low-cost HIL platform is proposed that focusses on the particular needs to assess collaborative ITS applications, while it could be implemented and applied fast comes with limited costs. Nevertheless it should be suited to let users experience large-scale collaborative applications such as the ones outlined within this thesis.

9.2.1. HIL platform requirements

All proposed applications in this thesis exploit the collaboration of hundreds or thousands of networked vehicles. With help of the HIL platform, users should be able to experience collaborative applications (including the large scale effects and interaction with many other vehicles) in a real-world environment. Similar to the approach from Bock [18] (see Section 9.1.3), there should be at least one vehicle which could be driven on public roads and which is equipped with the application.

The vehicle (including the driver) should be embedded in a real-time feedback loop with a HIL simulator, which mimics traffic conditions and the collaborative behaviour of thousands of vehicles. Since all proposed applications make use of vehicular communication, the vehicle should be equipped with according technologies.

9.2.2. HIL platform description

According to the outlined requirements of a HIL platform - and the need for a low-cost and light-weight solution - the platform described in the following has been set up. A general overview of the platform architecture is shown in Figure 9.1.

The HIL simulator mimics hundreds of vehicles; the aforementioned micro-simulator SUMO is used. The HIL simulation is fed with a road network imported from OSM. The imported region corresponds to the location of the real-world vehicle. Vehicles in the simulation are controlled using the given TraCI (Traffic Control Interface) interface [176]. The interface is used to integrate the application, which is deployed to every simulated vehicle and (if applicable), to the central authority ("the traffic management centre"). Applications are affecting the behaviour of the vehicles, both the real, experimental vehicles as well as the simulated vehicles in the simulation.

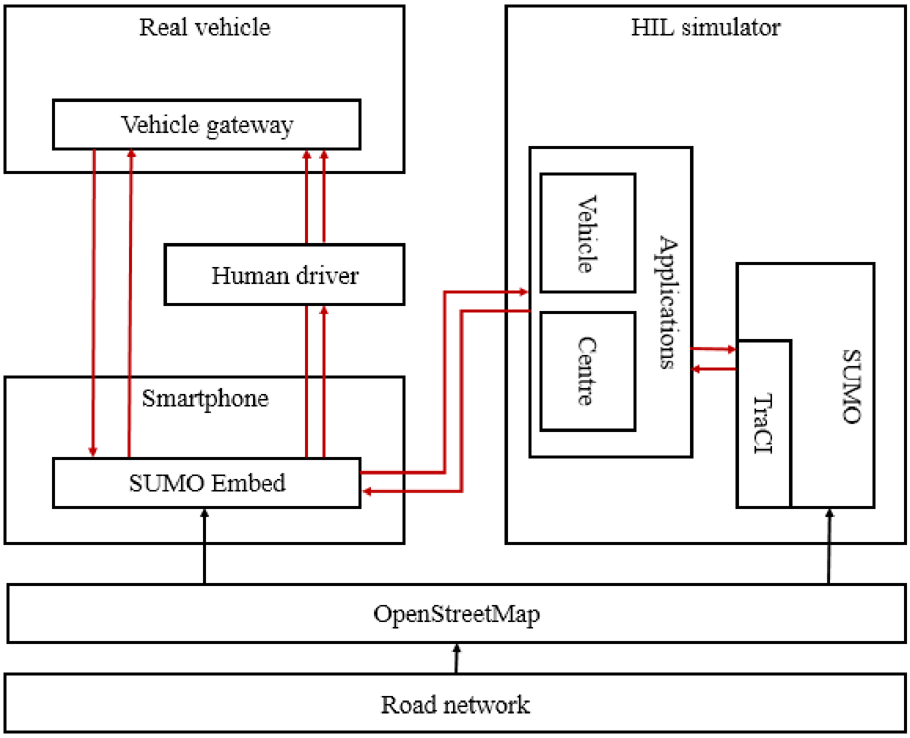

Figure 9.1.: Main components of the platform. The red arrows highlight the dynamic connections (i.e. the feedback loop) between the components.

The real vehicle including the vehicle gateway is described in detail in Section 9.2.3. The vehicle has been equipped with an Android smartphone (Samsung Galaxy S III mini). It relays, over a cellular network (3G), periodic information from the vehicle's on-board computer (e.g. the speed of the vehicle) to the HIL simulator with SUMO. Furthermore, it receives messages from the HIL simulator and shows them on the smartphone's display to provide instructions to the driver (such as alternative route recommendations). The JAVA application running on the device is called *SumoEmbed* and is described further below.

The implemented platform integrates one real-world vehicle in the testing scenario. However, the number of integrated or embedded vehicles is not limited. The key feature of the platform is the provisioning of the feedback loop which is setup among the real vehicle and the simulated vehicles (and - if applicable - a simulated, central authority). To establish that, information to synchronize the real-world vehicle's state with the simulation (and thus all simulated vehicles) is exchanged each second. The red arrows in Figure 9.1 highlight the links of the feedback loop.

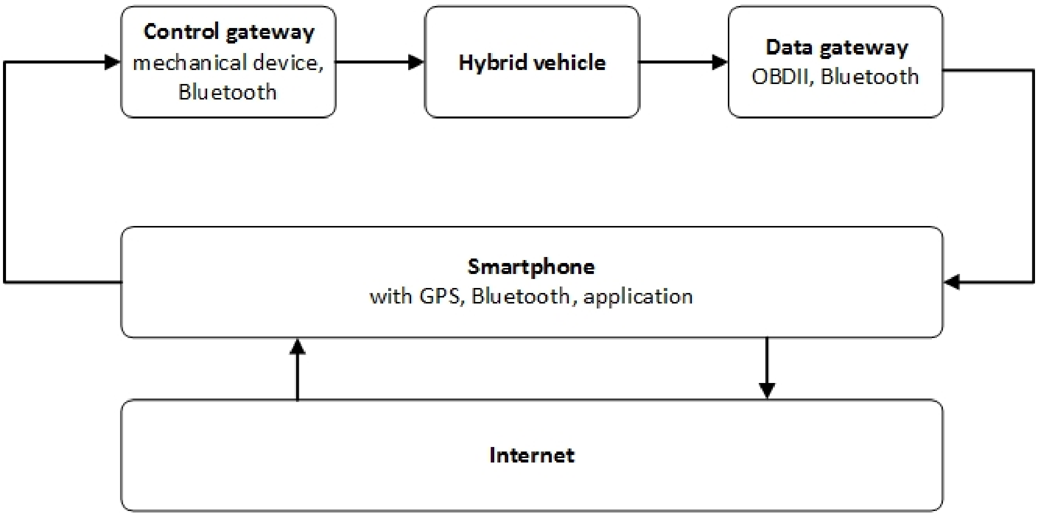

Figure 9.2.: Logical setup of the proof-of-concept vehicle.

9.2.3. The networked and hybrid car

The *Networked* or *Connected Car* is a popular term to describe future developments of vehicle technology. All applications discussed in this thesis make use of the core character of the networked car: communication technology. In order to emulate a fully integrated networked vehicle, a proof-of-concept vehicle has been set up. It is equipped with a Bluetooth gateways which could be interfaced by any modern smartphone. Thus, on the one hand there is access to (selected) dynamic vehicle data and controls (via the *control gateway*), on the other hand there are capabilities to connect to the internet via the cellular network. Some applications - i.e. TwinLin (Section 3) and Sponge (Section 4) - make use of (full) hybrid vehicles. Accordingly, the vehicle testing platform is a hybrid car - a 2008 Toyota Prius, depicted in Figure 9.3.

This setup allows the vehicle to interact with its environment in a smart manner. A simple overview of the setup is illustrated in Figure 9.2. A film demonstrating the operation of the vehicle can be found at `http://www.hamilton.ie/aschlote/twinLIN.mov`. In the following a brief overview on the basic components of this networked, hybrid vehicle is given.

Experimental car platform: 2008 Toyota Prius

The standard edition Prius is a hybrid car equipped with an internal combustion engine (ICE) and an electrical engine. Both are connected to the power train via differential gears and both can drive the vehicle on their own. The vehicle's electronic systems decide automatically which engine to use in order to get the most energy efficient driving with respect to the current driving situation. Furthermore, the Prius is equipped with an "EV-mode button". Pressing this button forces the car to switch to pure electric driving mode, if the battery level allows to. The EV mode is turned off if (1) the button is pushed again, (2) the vehicle's speed rises above 45 km/h, or (3) the battery level is to low.

Figure 9.3.: The experimental vehicle platform, a 2008 Toyota Prius.

Vehicle control gateway: Switching the driving mode

To be able to overwrite the Prius' default behaviour and to control the powering mode (EV- or ICE mode), a small mechanical device was developed and mounted to the "EV-mode" button. The device presses the button and thus changes the driving mode - just like the human driver could do. Accordingly, the vehicle still overrides the control (as long as the vehicle runs slower than 45 km/h and the battery has enough electric power). The device serves as a vehicle control gateway - it is equipped with a Bluetooth interface and can be controlled by the smartphone via that.

Vehicle data gateway: Accessing vehicle data

In order to get information about the current vehicle powering mode, the smartphone connects to the data gateway to the vehicle CAN bus. The vehicle's OBDII-interface is exploited here. A standard, commercially available hardware module has been used to interface the standardized OBDII plug and forwards data to connected Bluetooth devices - here the smartphone.

As a result, the smartphone has access to selected in-vehicle bus data such as driving mode (EV mode or ICE mode), battery level, emission levels, speed, etc.

Vehicle ITS Station: Android smartphone and dedicated application

The aforementioned Android smartphone serves as the Vehicle ITS Station. It is equipped with GPS and connects to the vehicle control and data gateway via Bluetooth, and to the internet via cellular networks. Furthermore it serves as an application unit, which means that it provides a software platform, where JAVA applications could be deployed to.

9.3. Proof of concept application deployment: Balanced routing

The HIL platform and several of its applications are described in [69]. Here, the deployment of the balanced routing application is described as a proof of concept (both for the platform and the application, which has been described in Section 6.4). Recall that the basic idea of the application is to route vehicles in a stochastic way such that features of the network (e.g. local pollution, traffic flow, etc.) are balanced.

9.3.1. Simulation setup

In order to illustrate the balancing approach, the following simulation scenario is setup. At the beginning of the simulation run there is no obstruction and vehicles go freely from the start node to the end node according to some route planning algorithm (like shortest path). Here, this route is the inner lap (route s_1 in Figure 9.4 (a)). As long as no traffic incident is monitored (up to timestep 120), no control approach is applied and all vehicles take this route s_1.

The traffic incident occurs 120 time steps after the start and causes a decrease in the capacity on this route (the red link visualized the relevant bottleneck in Figure 9.4 (d)). In consequence, it is defined that not all upcoming vehicles should stick to the route s_1 passing the obstruction but choose from one of the two other possible routes s_2 (Figure 9.4 (b)) and s_3 (Figure 9.4 (c)), with probabilities p_2 and p_3. It is defined that the traffic flow at the bottleneck should be controlled, here there should be three vehicles associated to this route s_1.

The probabilities p_1, p_2 and p_3 define the route choices of the vehicles. They are set according to a feedback control loop for the vehicular flow F_1 on s_1 and a balancing algorithm for vehicular flows F_2 and F_3 on s_2 and s_3 respectively. The vehicular flow F_i is defined as the number of vehicles travelling on the route i.
For this particular setup the controller for p_1 is defined by

$$p_1(k) = \begin{cases} 0, & \text{if } e(t) < 0, \\ 0.5, & \text{otherwise,} \end{cases} \tag{9.1}$$

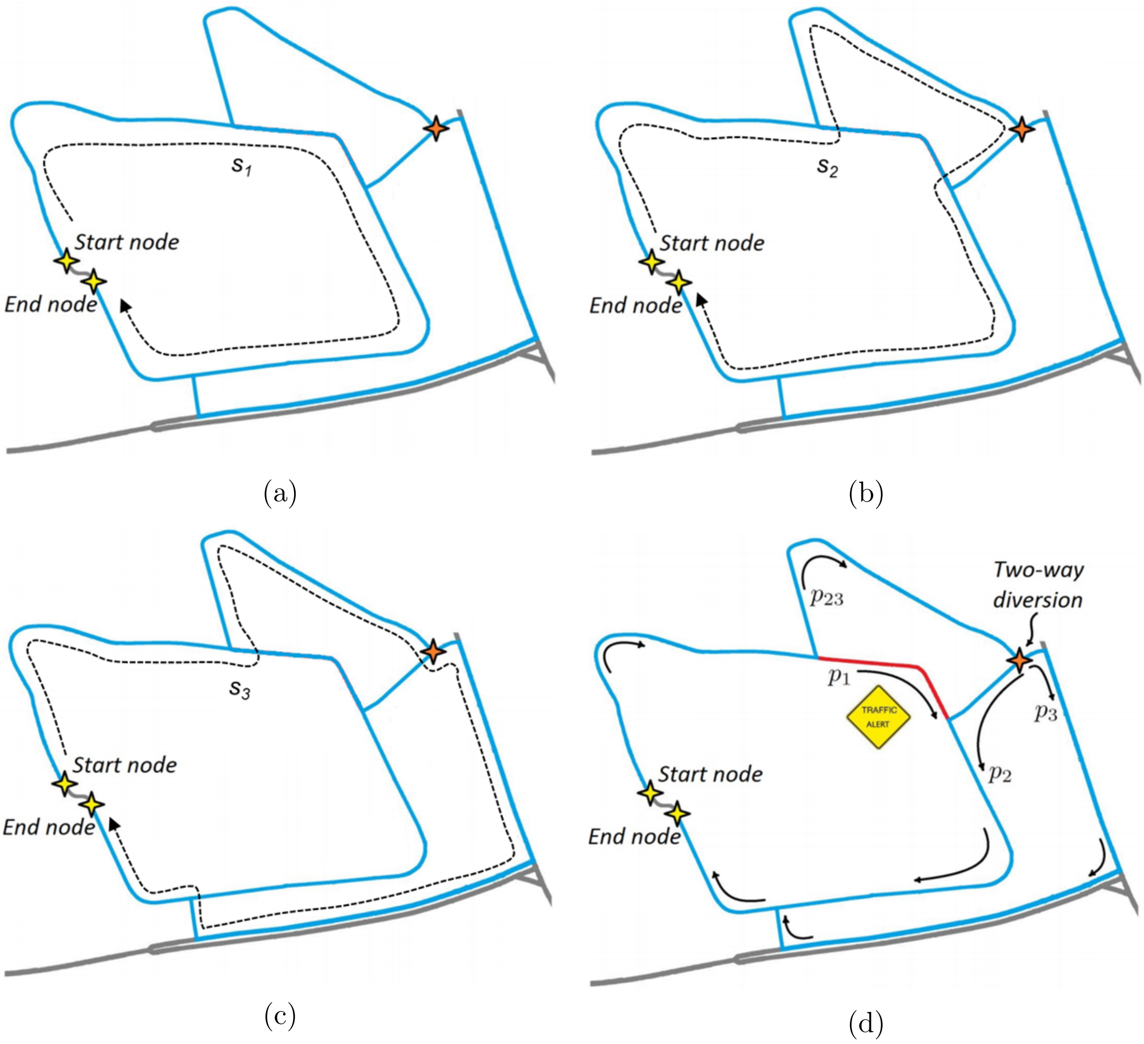

Figure 9.4.: Figure (a) shows the inner route s_1, Figure (b) shows route s_2, Figure (c) illustrates the longest route s_3. Figure (d) shows the setup with turn probabilities p_1, p_{23}, p_2 and p_3 as well as the location of the local incident, which affects the capacity of the red link.

where $e(t)$ is the difference between the set point for the traffic flow on s_1 and the actual traffic flow. The objective of this algorithm is to control the flow of vehicles on route s_1.

The second part of the setup targets the balancing control approach and is defined by

$$p_{23}\left(k\right) = p_2(t) + p_3(t) = 1 - p_1\left(t\right), \tag{9.2}$$

where $p_{23} = p_2 + p_3$ is chosen in a way such that p_2, p_3 guarantee that F_2 and F_3 are balanced. This corresponds to the initial objective. This balancing is done by letting vehicles - which have not chosen s_1 - choose between the route alternatives s_2 and s_3 stochastically. The probabilities are chosen in such a way that the traffic flows F_1, F_2 (here the number of vehicles on both routes s_i for $i \in 2, 3$) are balanced (here equal). Accordingly, the probability $P(X = s_i)$ to choose route s_i with $i \in 2, 3$ is calculated according to:

$$p_i = \alpha\left(1 - \frac{F_i}{F_2 + F_3}\right) \text{ with } \alpha = \begin{cases} 1 & \text{if } p_1 = 0 \\ 0.5 & \text{otherwise.} \end{cases} \tag{9.3}$$

9.3.2. Simulation results

The experiment was conducted with the networked vehicles driving through the North Campus of the National University of Ireland Maynooth (NUIM). The SUMO simulation was run for 278 seconds (timesteps) with an initial vehicular flow of 30 cars per minute released at the start node. The results from the test are shown in Figure 9.5 and Figure 9.6. The experimental vehicle driving on the campus is integrated in the number of vehicles.

74 timesteps after the beginning of the simulation, first vehicles reach the crossing where they could choose whether they follow route s_1 or not. As there is no capacity decrease yet, all choose to go through s_1. This behaviour lasts until timestep 120, when the obstruction occurs.

From timestep 121 on the control and the subsequent balancing approach are applied. Note that vehicles are associated with s_1 while they pass the obstructed link (coloured red in Figure 9.4 (d)) and that they are associated with s_2 after having passed this section. That is the reason for having vehicles associated with s_2 before timestep 120 (and the initial imbalance of F_2 and F_3).

In Figure 9.5, it can be seen that F_1 or $\bar{F}_1$ (the average of F_1 with a moving window of the last 50 timesteps - it is input to the control approach) is properly controlled and converges to the set value $r_1 = 3$.

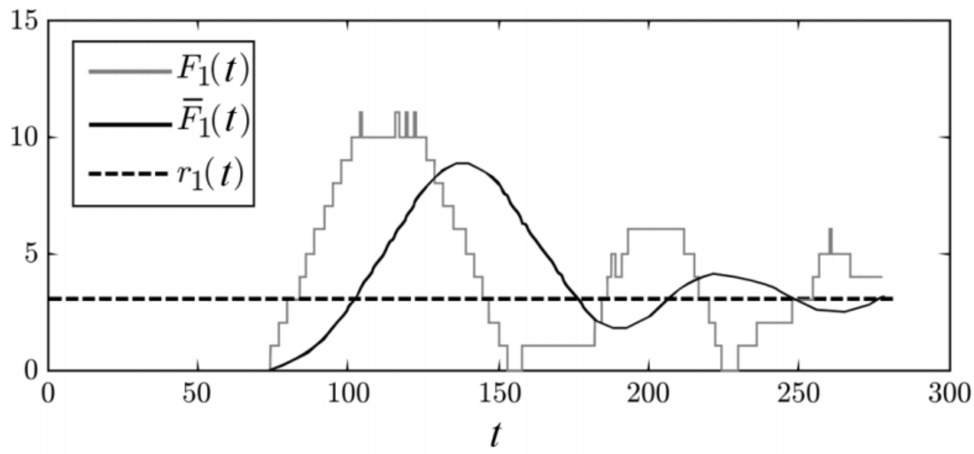

Figure 9.5.: Evolution of the traffic flow F (here number of vehicles associated to a route s_1). The traffic flow $F_1(t)$ corresponds to the number of vehicles associated to route s_i at time instant t. $\bar{F}_1$ is the average of F_1 with a moving window of the last 50 timesteps. It is input to the control algorithm. The set point - the desired value of vehicles going through route 1 is given by r_1.

In Figure 9.6 shows the result for the subsequent balancing approach for routes s_2 and s_3. The initial imbalance (see above) of F_2 and F_3 is equalized over time. F_2 and F_3 are well-balanced roughly from timestep 220.

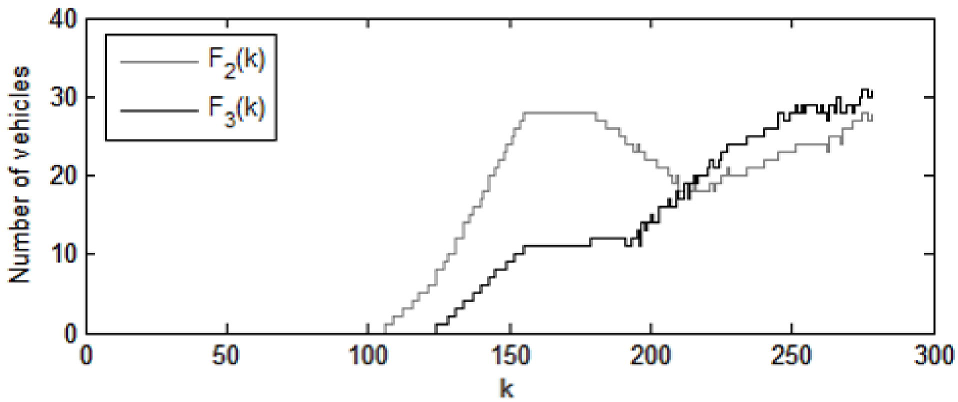

Figure 9.6.: Evolution of the vehicular traffic, where $F_i(k)$ is the vehicular flow through route s_1 and s_2 at instant k, $\bar{F}_1$ is the average of F_1 with a moving window of the last 50 timesteps, and r_1 is the set point for F_1. The vertical axis defines the number of vehicles on the relevant route s_1 or s_2.

The results show that the general control approach defined in Equation 9.3 performs well. The advantage of testing the re-routing system with the platform is that it allows real drivers to gain a feel of the technology being tested. The experimental vehicle was equipped with a smartphone with a simple application which displayed the routes the driver shall follow. Initial feedback from (friendly) users were positive. However, non-participating traffic on the routes where a challenge such that test drivers were performed during off-hours.

9.4. User acceptance of collaborative ITS applications

A key requirement for the proposed ITS applications is that they do not impose choices on the driver. In all applications, the user receives rather recommendations than instruction he or she has to follow. Accordingly, all applications are robust to users who do not follow advices. However, the majority of users must comply such that applications lead to the desired impacts. Therefore, it is essential that users accept the applications even though they may come with (temporal) disadvantages for themselves compared to their normal behaviour. This hypothesis is one of the key starting points of this thesis. It has been supported by one of the key results of Rytt [136]: In a brief survey in the context of the work of this thesis it was found that the majority of people would accept detours if it serves the social welfare (more details on the survey are found in Section 9.4.1).

This is especially true for the balanced routing application (Section 6.4) where users might be requested to take a longer route than a traditional shortest path routing algorithm would propose. The same applies to the collaborative car park sharing concept (Section 8.2) where the acceptance of letting others park at own grounds as well as the possibility that the private parking space might be occupied is required by the relevant owner.

The key instrument in TwinLin (Section 3) as well as the SPONGE application (Section 4) is the engine switching (EV to ICE or vice versa). Here is implicitly assumed, that users are indifferent to the engine type choice. User acceptance in this context is thus not as critical as in the context of the two aforementioned applications. Thus, the following discussion on user acceptance focusses on the routing and the parking application.

9.4.1. Acceptance of balanced routing

In the aforementioned preliminary survey, the main goal has been to get an idea whether drivers accept assistant systems' instructions, which do not optimize the individual users' egocentric objectives but a community goal - as it is the case in the balanced routing application. Bonsall and Joint analysed the driver compliance in the context of common navigation devices in [22] in general. Compliance is defined as the ratio of followed assistant system instructions and all given instructions. According to the study, navigation instructions were rejected at 27 % of junctions and only 35 % of trips were conducted as initially recommended by the navigation device.

More interesting are the reasons Bonsall and Joint have identified for compliant or non-compliant behaviour. These are briefly outlined in the following. It needs to be kept in mind that other surveys - such as Kantowitz et. al [87] - come to other compliance measures, while they basically agree on the basic compliance reasoning:

1. The extent to which a navigation instruction is corroborated by, or in conflict with local evidence about the alternative. That means for instance that the compliance sinks at a turning option when a traffic jam is detected at the alternative, which is recommended by the system.

2. History of instructions: In case previous instructions were non-satisfactory, drivers' compliance decreases.

3. Mismatch between the system's route choice criteria and those of the driver.

4. Familiarity of the user with the environment: It has been shown that the compliance rate sinks with an increase of the user's familiarity about the area.

5. Driver's age: Younger people tend to reject advices more often than older people.

6. Driving experience: Drivers with more experience reject advices more often than drivers with little driving experience.

7. Modest drivers: Drivers with more modest route criteria tend to accept route recommendations more often.

In the context of the proposed routing application, the mismatch between the system's and the user's route choice criteria is particularly interesting. That is because the application obviously does not aim at optimizing the user's benefit but focusses on the control of network features. All other of the reasons listed above target features of ordinary navigation systems, which are not related to the specific question of the proposed routing application.

Bonsall took a deeper look into the mismatch question in [20]. Ordinary navigation systems calculate routes and aim at minimizing travel time or generalized costs - as a weighted combination of time and distance. In fact, between 75 % to 90 % of all route choice decisions are made on this basis. However, other criteria such as road type, congestion, stop sign and traffic signal avoidance are mentioned. For more details the reader is referred to [20] and the citations within.

The particular question whether user would accept detours in favour for a community benefit (which could be applied to the emission aware routing and the balanced routing) has been targeted by the aforementioned survey [136]. Relevant outcomes are outlined in the following.

The aforementioned survey was conducted during two weeks in July and August 2012. More than one hundred people (121) participated and gave feedback in a web-based questionnaire. Respecting the relatively small number of participating users, main features of the set is well-suited: the number of male and female users is balanced, 55 % of users are between 25 and 40 years, 87 % stated that they have a valid drivers licence and 71 % that they have a navigation system.

The most important result of the survey is related to the question *"How much time would you accept for a detour that benefits the overall traffic conditions?"*.
It shows that 74 % of users would generally accept detours if it benefits the community (74 % accept a detour in general if it benefits the overall traffic situation, 75 % if it benefits the environment). 67 % of users would accept a detour that causes up to 3 extra minutes assuming a normal travel time of 30 minutes. The feedback of users supports the basic statement that people would accept an individual disadvantage (detour) if it benefits the community welfare.

Based on the promising outcome of the survey, a second survey was conducted with real world experiments in January and February 2015 with ten participants [137]. In the first week six participants joined the experiments, in the second week another four. Most participants were employees who commuted by car to work during the experiment. Alternative routes were manually generated and provided to the individual user the night before the trip was conducted. The route alternative respected the result form the first survey: no route alternative was causing more than five minutes of extra time compared to the ordinary calculated route to the same destination. Each user received different route recommendation each day. In the evening users were asked to provide feedback whether they followed the recommended route or not.

Main objective of the second study has been to analyse users' motives and the influence of incentives that are applied to increase user compliance.

The experiments closed with a focus group discussion with participants. Focus group discussions are a standard method for validating proposed functionalities in HMI research and software development [99]. They are often employed in early system development.

Routes were recommended with different objectives. These are:

- **Safety of children:** Each user receives routes which pass less schools and kindergartens than the original, ordinary route.

- **Traffic safety:** Routes are recommended that do not only pass less schools and kindergartens but also hospitals, construction sites, complex crossings and locations with many pedestrians etc.

- **Environment**: Users receive alternative routes which avoid passing environmental hotspots such as recreation areas.

The acceptance of the routes relate directly to the algorithmic objective of the route finding algorithm. In the context of incentives, they are called functional incentives.

Apart from the three aforementioned routing objectives, at one day participants were asked to take a particular route, where no system-centric objectives were named. When looking at the results in terms of followed routes, there is no impact of the routing objectives on user compliance. However, a key requirement for functional incentives could not be granted during the experiment: credibility. It turned out in the focus group discussion, that users questioned the impact of their route choices on the relevant system feature - mainly because users knew that only few user participated in the experiment. Furthermore generated routes were not always correct (e.g. alternative routes passed schools that were not considered during route calculation but were known by participants). However, results from the web-based questionnaire (that was performed in parallel to the experiments) have shown that the safety of children is most important functional incentives to users in from regarded list.

The survey targeted also other incentives in regard to the balanced routing application such as gamification, the use of nudges (a broad overview on the use of nudges is given in [169]), and non-monetary incentives.

On a general level, the real-world experiment confirmed the first survey: Users are willing to follow routes which optimize some system-centric feature. The real-world experiments have further shown that the use of incentives could help to increase the compliance of users. More information on the survey can be found in [137]. A broad overview on the use of incentives to fight traffic congestion is provided in [192] and the references within.

9.4.2. Acceptance of collaborative parking

The attractiveness of the collaborative parking applications with private landlords (see Section 8.2) was further analysed in a survey published in [135], which was conducted in the context of the work this thesis. Among other aspects the survey was intended to help specify user requirements to the application such as:

- **General attractiveness**: How attractive is the general concept for landlords?

- **Space occupancy**: To what extent will landlords accept that their own space is occupied when they need it?

- **Free space reachability**: Assuming the space on own grounds is not accessible. How far may alternative spaces be away for landlords?

92 people participated in the survey. The general idea to share the own parking space was evaluated ambivalently. Only a third of the participants were interested in joining such a park sharing system.[4] About 75 % of people who were not interested said that the main drawback is the possibility of not getting access to their own parking space. A second reason was that people were afraid that allocated parking spaces would not meet their requirements (i.e. that they are to small etc.).

Furthermore, participants were asked if they would require to get access to their space even though they have rented it to other users during the time. About half of the participants said that they would want the own space to be free in 95 % of cases. That confirms that there are many people who are not willing to rent out their parking spaces at all.

Assuming the own space is not accessible people were asked how far the next available spot should be. Figure 9.7 illustrates the accepted distances to the next available parking spot assuming the own space is occupied by a user. The great majority of people would not accept alternative parking spaces which are more than 200 m away from their own space. That implies that the system would only be accepted if many landlords in the vicinity would join such a system.

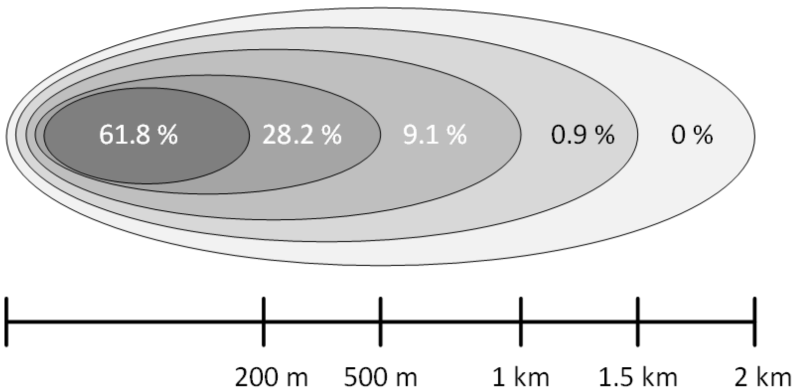

Figure 9.7.: Feedback regarding accepted distances to the next available parking spot if the landlord's own space is occupied. Percentages define the share of users who accept the illustrated distance.

[4] 31 % of people were interested to join such a concept. However, it must be stated, that 27 % of people were not experiencing a parking problem, such that the concept could not be applied. Independent from their own circumstances, over 60 % said that the concept is attractive on a general level.

10. Outlook and conclusion

Summary: *In this chapter the main contributions of the thesis and potential extensions of previously introduced applications are briefly outlined.*

10.1. Collaborative vehicular behaviour through feedback loops

Several collaborative applications have been introduced in this thesis. All exploit the power of new developments in the automotive domain, in particular vehicular communication and new types of vehicles (in particular PHEVs). Networked vehicles are utilized to set-up a feedback loop around a large number of vehicles and some infrastructure (e.g. a traffic management centre). On the basis of such feedback loops, applications were introduced which establish collaborative behaviour among vehicles and address key challenges of smart cities, which could not be solved if vehicles would not align their actions.

It has been shown that true collaborative behaviour of vehicles is generally possible with the help of a simple system architecture: networked vehicles and a central infrastructure, which has access to real-time information about network features and broadcasting capabilities. Based on broadcasted information provided by some central entity, vehicles deduce an aligned collaborative behaviour.

A key feature of the proposed applications is that they meet key specific requirements of the ITS domain:

- **Non-Invasiveness**: All applications require an adaptation of the behaviour of drivers, e.g. adapted routes, adapted selection of parking spaces and charging stations, etc. At the same time, drivers should not be forced to any behaviour. All introduced applications respect and accept these requirements. That means that all approaches work even though not all users follow the recommended behaviour.

- **Privacy**: Securing the privacy of users is a central requirement for acceptance of applications. Thus, sensible information about user preferences and intentions should not be shared with any party. All application are designed in a way such that this kind of information is processed locally only.

- **Varying scale and communication topology**: Vehicular networks are characterized by changing communication topology as vehicles move and enter and exit the system over time. Furthermore, the number of vehicles in the network change dramatically (e.g. rush-hour vs. a midnight scenario). All approaches are based on algorithms which could cope with that. That means that there is no need to change their parametrization based on the number of collaborating vehicles.

Before proceeding with discussions on a per application level, it is important to note that the proposed solutions support a basic assumption: The advent of new technologies in the automotive domain (in particular communication technology and HEVs) enable a new way of approaches: Large-scale collaboration among vehicles. The possibility to control a large number of vehicles is a very powerful instrument and makes it possible to fight problems in a very direct way, where conventional approaches (per-vehicle measures) are not suited at all. I am convinced that collaborative behaviour using closed-loop feedback strategies among vehicles are suited to change the way road traffic is designed and to solve (or at least relax) multiple of most pressing urban problems.

In the following, selected results, impacts and potential extensions of these applications are discussed. Furthermore, an environment is briefly discussed which supports the assessment of the proposed applications.

10.2. Emission control with cooperative HEVs

The TwinLin application (Chapter 3) controls traffic related emission in a geographic area. It is a novel approach to exploit the potential of HEVs in a smart city. A proof-of-concept implementation with (cellular) V2I communication was introduced using an experimental networked vehicle described in Section 9.2.3. A number of simulations showcased the efficacy of the approach. The basic idea has been extended to guarantee a notion of fairness, such that all vehicles are allowed to emit the same amount of pollution.

The core idea could be applied to a number of derivative scenarios. I.e., it is possible to relax the assumption that all cars are within a small geographic area and instead assume that a company operates a fleet of vehicles (e.g. a taxi or a delivery company). Equivalent to a defined local pollution limit that serves as a constraint, one can give the company an emissions budget that it can allocate to its vehicles. This approach has been discussed in Chapter 5. It has been extended with the help of a generalized trading framework in a way than the fleet operator can solve an optimisation problem to allocate emission permits to cars based on some utility functions.

In a next step, it could be investigated how similar approaches can be used to control emissions for conventional vehicles. For example, when considering an area within a city in which pollution is to be controlled and this area is only accessible through traffic light controlled junctions, it is possible to choose a policy for admission control such that the fraction of green time for traffic lights on streets into the area is proportional to a broadcasted speed limit (assuming that the emission of vehicles depends on the speed).

10.3. Grid control with cooperative HEVs

The SPONGE application utilizes the possibilities of HEVs again. In this context it exploits the fact that vehicles can absorb (naturally) generated electrical energy from the grid by switching from EV mode to ICE mode in a smart way. From a theoretical perspective, such a problem can be easily formulated and solved using e.g. AIMD. It helps to align the behaviour of vehicles in a distributed way. From a practical point of view, the technology influence the engine mode is already available.

The SPONGE solution has the potential to change the "charging paradigm". Until today, research on electric vehicle charging has focussed on how to share the available energy among the connected fleet of vehicles (in real-time) in a manner that is compliant with the desires of the EV owners, the constraints of the grid, and the available power. Note that in this case, there might arise problems in the power grid to accept the unexpected load, with the ultimate possibility of causing thermal overload of network components, low voltages at sensitive locations of the network, and increased phase unbalance [81]. Even ignoring this, the required optimisations often place severe constraints on the EV owners in the form of inconvenient charging profiles. In the proposed approach, the control is done while driving.

The discerning reader may ask why the individual vehicle owner should not simply deplete the electric energy completely before switching to ICE mode. There are many reasons for doing this. First, in some engines, electrical power and ICE are combined (mild hybrid electric vehicles) to reduce overall consumption, or for other objectives of interest (e.g., extend the lifetime of the battery) [147]. Second, access to certain parts of the city may be restricted to zero emission vehicles. Thus, maintaining a particular level of electrical energy for this purpose is meaningful. Finally, depleting the battery beyond the energy levels available during the next charging period, may lead to a situation where the battery is not filled during the $k+1$'th charging period. Thereby, the ICE may need to be engaged prematurely in driving, thus leading to unnecessary emissions and increased fuel consumption.

In a step, one should integrate a reliable weather forecast in the overall system, in order to make better estimations on the available energy from renewables and thus take improved decisions about when to switch from one mode to another mode. Furthermore, the application should be tested with help of an experimental car (the experimental car utilized as part of the HIL platform, see Section 9.2.3) is a HEV without plug-in charging capabilities.

10.4. (Local) feedback-based routing

The first routing application (described in Section 6.2) introduced a feedback-based routing approach, where vehicles were routed along an environmentally friendly route instead of a conventional (fastest or shortest) route. However, the implemented routing engine considered user acceptance aspects by guaranteeing that routes do not become much longer than the conventional ones. With help of simulations, it has been shown that the emission of vehicles was shifted to other locations, such that environmental hotspots were avoided and relieved.

Generally, hotspots could be freely defined based on the application's objectives and contexts; examples are kindergartens, hospitals, recreation areas, safety hotspots etc. Unfortunately, flapping effects could not be avoided in the approach.

In order to overcome this drawback, a second approach was developed (described in Section 6.4). It introduced the idea of stochastic routing combined with feedback control. By simulations and on-street experiments using the HIL platform, it was shown that the application is suited to control and balance traffic in a locally limited area (at all points in time and without flapping effects).

In a next step, the valid size of the area should be discussed. It has turned out in an initial evaluation, that the simulation with hundreds of vehicles (and many more routing requests) requires great amount of processing capabilities.

10.5. Charging EVs and parking in the smart city

The issue of queueing of electric vehicles at charging stations has been ignored in the research of electric vehicles so far, even though it can have serious impacts on the acceptance of the concept general of electric vehicles. The introduced approach presents a novel idea to address this challenge. A key feature of the (stochastic) approach is that there is no need for a central instance which monitors the status of charging stations or electric vehicles.

Compared to deterministic methods, the distributed and stochastic approach is scalable and has a plug-and-play character, which means that charging stations could enter and leave the system without any adjustments of the system. The balancing performance compares well with deterministic alternatives.

Further work may investigate different utility functions for charging stations to balance needs other than power load, and consider the issue of communication delays.

The general starting point for the proposed parking solutions has been to utilize private or restricted parking spaces during times, when they are not occupied. Different to available solutions, the proposed approach guarantees a high quality of service. In this context that means that the original parking space owner (here landlords or enterprises) have access to their space when they request it - even though they might request the space during times when they rented out the space or when the car space renter overstays.

The next step should be to validate the attractiveness to users. While a simple survey has already been performed on selected aspects (see Section 9.4.2), further proofs should be given.

10.6. Validation of smart city applications

There are many ways and environments that support a validation process of ITS applications (see Chapter 9). In many cases, validation addresses technical questions. However, in the context of the proposed applications, user acceptance might be the more critical aspect.

A simple and light-weight hardware-in-the-loop (HIL) simulation platform for emulating large-scale intelligent transportation systems was presented, which embeds a real vehicle into the microscopic traffic simulator SUMO. A goal of the platform is to provide drivers a way to experience of what it would feel like to participate in large-scale, feedback-based, connected vehicle applications. It allows ITS developers to better examine real driver reactions and acceptance of the proposed (and related) applications. The routing application described in Section 6 was tested using the platform and served as one (of multiple, see [69]) proof-of-concept application.

The HIL platform could be enhanced in multiple ways. The driver interface in the networked car could be improved, i.e. by the integration of map-matching functionalities, by integration with other simulators and by the addition of further embedded real vehicles, including a plug-in hybrid electric vehicle.

Finally, a user survey and experiments have been performed, which have shown inconclusive results regarding the acceptance of collaborative ITS applications, i.e. the routing and parking application. Even though all proposed applications respect key domain-specific requirements, user acceptance is a critical feature which should be further investigated. Finally, additional analysis should be performed in regard to business models, legal questions and deployment scenarios for each application.

Bibliography

[1] Paris imposes car restrictions to fight pollution. `http://www.bbc.com/news/world-europe-32016730`, 2015.

[2] Ittai Abraham, Daniel Delling, Andrew V. Goldberg, and Renato F. Werneck. Alternative routes in road networks. *Journal of Experimental Algorithmics (JEA)*, 18:1–3, 2013.

[3] Christian Adler, Stephan Eichler, Timo Kosch, Christoph Schroth, and Markus Strassberger. Self-organized and context-adaptive information diffusion in vehicular ad hoc networks. In *Wireless Communication Systems, 2006. ISWCS'06. 3rd International Symposium on*, pages 307–311. IEEE, 2006.

[4] European Environment Agency. Roadmap to a Single European Transport Area. `http://ec.europa.eu/transport/strategies/facts-and-figures/putting-sustainability-at-the-heart-of-transport/index_en.htm`, August 2013.

[5] F. Ahmed-Zaid, F. Bai, S. Bai, C. Basnayake, B. Bellur, S. Brovold, G. Brown, L. Caminiti, D. Cunningham, H. Elzein, et al. Vehicle safety communications–applications (VSC-A) final report. Technical report, 2011.

[6] Saif Al-Sultan, Moath M. Al-Doori, Ali H. Al-Bayatti, and Hussien Zedan. A comprehensive survey on vehicular ad hoc network. *Journal of network and computer applications*, 37:380–392, 2014.

[7] D.W. Allan, N. Ashby, C.C. Hodge, and Hewlett-Packard Company. *The Science of Timekeeping.* Application note. Hewlett-Packard, 1997.

[8] Farooq M. Anjum and Leandros Tassiulas. Fair bandwidth sharing among adaptive and non-adaptive flows in the Internet. In *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1412–1420. IEEE, 1999.

[9] Ing Kiriakos Athanasas, C. Bonnet, H. Fritz, and C. Scheidler. VALSE–VALidierung Sicherheitskompetenter Elektronischer Assistenz-und Fahrfunktionen. *VDI-Berichte*, pages 261–280, 2003.

[10] Gerd Baumann, Anne Piegsa, and Christoph Liedecke. Der neue Fahrsimulator der Universität Stuttgart. In *ASIM Workshop*, 2010.

[11] M. Bechler et al. PRE-DRIVE Implementation and Evaluation of C2X Communication Technology, Deliverable D1.4. *PRE-DRIVE Std*, 2009.

[12] Michael Behrisch, Laura Bieker, Jakob Erdmann, and Daniel Krajzewicz. SUMO–Simulation of Urban MObility. In *The Third International Conference on Advances in System Simulation (SIMUL 2011), Barcelona, Spain*, 2011.

[13] Francesco Bellotti, Alessandro De Gloria, R. Montanari, N. Dosio, and Domenico Morreale. Comunicar: designing a multimedia, context-aware human-machine interface for cars. *Cognition, Technology & Work*, 7(1):36–45, 2005.

[14] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization.* Princeton University Press, 2009.

[15] Carl Bergenhem, Steven Shladover, Erik Coelingh, Christoffer Englund, and Sadayuki Tsugawa. Overview of platooning systems. In *Proceedings of the 19th ITS World Congress, Oct 22-26, Vienna, Austria (2012)*, 2012.

[16] Jitendra Bhatia and Bhumit Shah. Review On Various Security Threats & Solutions And Network Coding Based Security Approach For Vanet. *International Journal of Advances in Engineering & Technology*, 6(1), 2013.

[17] Vincent Blondel, Julien M. Hendrickx, Alex Olshevsky, and J. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *IEEE Conference on Decision and Control*, volume 44, page 2996. IEEE; 1998, 2005.

[18] Thomas Bock, Markus Maurer, and G. Farber. Vehicle in the Loop (VIL)–A new simulator set-up for testing Advanced Driving Assistance Systems. In *Driving Simulation Conference, North America 2007 (DSC-NA 2007)*, 2007.

[19] Peter Bonsall. The influence of route guidance advice on route choice in urban networks. *Transportation*, 19(1):1–23, 1992.

[20] Peter Bonsall. The influence of route guidance advice on route choice in urban networks. *Transportation*, 19:1–23, 1992.

[21] Peter W. Bonsall and Matthew Joint. Driver compliance with route guidance advice: The evidence and its implications. In *Vehicle Navigation and Information Systems Conference, 1991*, volume 2, pages 47–59. IEEE, 1991.

[22] P.W. Bonsall and M. Joint. Driver compliance with route guidance advice: The evidence and its implications. In *Vehicle Navigation and Information Systems Conference, 1991*, volume 2, pages 47 – 59, October 1991.

[23] P.G. Boulter, Tim J. Barlow, and Ian S. McCrae. Emission factors 2009: Report 3-exhaust emission factors for road vehicles in the United Kingdom. *TRL Published Project Report*, 2009.

[24] P.G. Boulter, T.J. Barlow, I.S. McCrae, and TRL Limited. *Emissions Factors 2009: Report 3 - Exhaust Emission Factors for Road Vehicles in the United Kingdom*. Published project report ; PPR356. TRL, 2009.

[25] Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[26] Robert D. Brook, Barry Franklin, Wayne Cascio, Yuling Hong, George Howard, Michael Lipsett, Russell Luepker, Murray Mittleman, Jonathan Samet, Sidney C. Smith, et al. Air pollution and cardiovascular disease a statement for healthcare professionals from the expert panel on population and prevention science of the american heart association. *Circulation*, 109(21):2655–2671, 2004.

[27] Lukasz Budzisz, Rade Stanojevic, Arieh Schlote, Fred Baker, and Robert Shorten. On the Fair Coexistence of Loss- and Delay-Based TCP. *IEEE/ACM Trans. Netw.*, 19(6):1811–1824, 2011.

[28] CAMP Vehicle Safety Communications Consortium and others. Vehicle safety communications project: task 3 final report: identify intelligent vehicle safety applications enabled by DSRC. *National Highway Traffic Safety Administration, US Department of Transportation, Washington DC*, 2005.

[29] Kartikeya Chandrayana, Biplab Sikdar, and Shivkumar Kalyanaraman. Scalable configuration of red queue parameters. In *High Performance Switching and Routing, 2001 IEEE Workshop on*, pages 185–189. IEEE, 2001.

[30] Jasmine Chennikara-Varghese, Wai Chen, Onur Altintas, and Shengwei Cai. Survey of routing protocols for inter-vehicle communications. In *Mobile and Ubiquitous Systems: Networking & Services, 2006 Third Annual International Conference on*, pages 1–5. IEEE, 2006.

[31] Dah-Ming Chiu and Raj Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN systems*, (1):1–14, 1989.

[32] Hong Cho and Byeongwoo Kim. Cooperative Intersection Collision-Warning System Based on Vehicle-to-Vehicle Communication. 2014.

[33] Emanuele Chrisostomi, M. Liu, M. Raugi, and Robert N. Shorten. Plug-and-play Distributed Algorithms for Optimised Power Generation in a Microgrid. *IEEE Transactions on Smart Grid*, 5(4):2145–2154, 2014.

[34] European Commission. *Towards Fair and Efficient Pricing in Transport: Policy Options for Internalizing the External Costs of Transport in the European Union: Green Paper : Document Drawn Up on the Basis of COM(95) 691 Final.* Number Nr. 2 in Bulletin of the European Union: Supplement. Office for Official Publications of the European Communities, 1996.

[35] European Commission. Regulation (EC) No. 443/2009 of the European Parliament and of the Council of 23 April 2009 setting emission performance standards for new passenger cars as part of the Community's integrated approach to reduce CO2 emissions from light-duty vehicles. `http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:140:0001:01:EN:HTML`, August 2013.

[36] R.P. Dameri and C. Rosenthal-Sabroux. *Smart City: How to Create Public and Economic Value with High Technology in Urban Space*. Progress in IS. Springer, 2014.

[37] Brandon Davito, Humayun Tai, and Robert Uhlaner. The smart grid and the promise of demand-side management. *McKinsey & Company*, 2010.

[38] Douglas W. Dockery, C. Arden Pope, Xiping Xu, John D. Spengler, James H. Ware, Martha E. Fay, Benjamin G. Ferris Jr., and Frank E. Speizer. An association between air pollution and mortality in six US cities. *New England journal of medicine*, 329(24):1753–1759, 1993.

[39] Kieran Donaghy, Georg Rudinger, and Stefan Poppelreuter. Societal trends, mobility behaviour and sustainable transport in Europe and North America. *Transport Reviews*, 24(6):679–690, 2004.

[40] Valentin Dumitrescu and Jinhua Guo. Context assisted routing protocols for inter-vehicle wireless communication. In *Intelligent Vehicles Symposium, 2005. Proceedings. IEEE*, pages 594–600. IEEE, 2005.

[41] Tamer El Batt, Siddhartha K. Goel, Gavin Holland, Hariharan Krishnan, and Jayendra Parikh. Cooperative collision warning using dedicated short range wireless communications. In *Proceedings of the 3rd international workshop on Vehicular ad hoc networks*, pages 1–9. ACM, 2006.

[42] Wim Elfrink and R. Kirkland. The smart-city solution. *McKinsey Quarterly*, 2012.

[43] R.H.M. Emmerink. *Information and Pricing in Road Transportation.* Advances in spatial science. Springer-Verlag, 1998.

[44] Richard Gilles Engoulou, Martine Bellaïche, Samuel Pierre, and Alejandro Quintero. Vanet security surveys. *Computer Communications*, 44:1–13, 2014.

[45] U.S. Environmental Protection Agency (EPA). Automobile Emissions: An Overview. `http://www.epa.gov/otaq/consumer/05-autos.pdf`, August 2013.

[46] David Eppstein. Finding the k shortest paths. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 154–165. IEEE, 1994.

[47] Eurostat. Roadmap to a single european transport area. `http://ec.europa.eu/transport/strategies/facts-and-figures/transport-matters/index_en.htm`, August 2013.

[48] Eurostat. Renewable energy statistics-electricity generated from renewable sources, 2014.

[49] Excise. Excise — Wikipedia, the free encyclopedia, 2010. [Online; accessed 01-May-2015].

[50] Asif Faiz. *Automotive air pollution: Issues and options for developing countries*, volume 492. World Bank Publications, 1990.

[51] H. Fathy, Z. Filipi, Jonathan Hagena, and J. Stein. Review of hardware-in-the-loop simulation and its prospects in the automotive area. *Ann Arbor*, 1001:48109–2125, 2006.

[52] Wu-chang Feng, Dilip D. Kandlur, Debanjan Saha, and Kang G. Shin. A self-configuring RED gateway. In *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1320–1328. IEEE, 1999.

[53] John G. Fernald. Roads to Prosperity? Assessing the Link between Public Capital and Productivity. *American Economic Review*, 89(3):619–638, September 1999.

[54] Lino Figueiredo, Isabel Jesus, J.A. Tenreiro Machado, J. Ferreira, and J.L. Martins De Carvalho. Towards the development of intelligent transportation systems. In *Intelligent Transportation Systems*, volume 88, pages 1206–1211, 2001.

[55] Sally Floyd and Van Jacobson. Random early detection gateways for congestion avoidance. *Networking, IEEE/ACM Transactions on*, (4):397–413, 1993.

[56] Randy A. Freeman, Peng Yang, and Kevin M. Lynch. Distributed estimation and control of swarm formation statistics. In *American control conference*, volume 7. Citeseer, 2006.

[57] Stephan Fuchs, Korbinian Lohr, and Markus Lienkamp. The relevance of the production process for the life cycle of CO2 emissions of electric vehicles. *ATZproduktion worldwide*, 5(3):34–38, 2012.

[58] Jan Fuglestvedt, Terje Berntsen, Gunnar Myhre, Kristin Rypdal, and Ragnhild B. Skeie. Climate forcing from the transport sectors. *Proceedings of the National Academy of Sciences*, 105(2):454–458, January 2008.

[59] Senatsverwaltung für Stadtentwicklung und Umweltschutz Berlin (Ed.). Umweltaltlas Berlin, Karte 07.01 Verkehrsmengen (1995). `http://www.stadtentwicklung.berlin.de/umwelt/umweltatlas/d701_03.htm`.

[60] Senatsverwaltung für Stadtentwicklung und Umweltschutz Berlin (Ed.). Umweltaltlas Berlin, Karte 07.01 Verkehrsmengen (2011).

[61] S. Gailmard. *Statistical Modelling and Inference for Social Science*. Analytical Methods for Social Research. Cambridge University Press, 2014.

[62] Robert Geisberger, Moritz Kobitzsch, and Peter Sanders. Route planning with flexible objective functions. In *ALENEX*, volume 10, pages 124–137. SIAM, 2010.

[63] Tom Geller. Car talk. *Communications of the ACM*, 58(3):16–18, 2015.

[64] Matthias Gerlach, Horst Rechner, and Tim Leinmüller. Security framework for vehicular applications. In *Third International Workshop on Vehicle-to-Vehicle Communications*, volume 2007, 2007.

[65] Gesellschaft für Emissionsmanagement und Beratung GmbH (GEMB). Informationsportal zur Deutschen Umweltplakette. `http://www.umwelt-plakette.de`.

[66] Olaf Gietelink, Jeroen Ploeg, Bart De Schutter, and Michel Verhaegen. Development of advanced driver assistance systems with vehicle hardware-in-the-loop simulations. *Vehicle System Dynamics*, 44(7):569–590, 2006.

[67] Martin Goralczyk, Jens Pontow, Florian Häusler, and Ilja Radusch. The automatic green light project-vehicular traffic optimization via velocity advice. In *Adjunct Proceedings of the First International Conference on The Internet of Things*, 2008.

[68] Wynita Griggs, J. Y. Yu, Fabian Wirth, Florian Häusler, and Robert N. Shorten. On the Design of Campus Parking Systems With QoS Guarantees.

IEEE Transactions on Intelligent Transportation Systems, 17(5):1428–1437, May 2016.

[69] Wynita M. Griggs, Rodrigo H. Ordóñez-Hurtado, Emanuele Crisostomi, Florian Häusler, Kay Massow, and Robert N. Shorten. A Large-Scale SUMO-Based Emulation Platform. *Intelligent Transportation Systems, IEEE Transactions on*, 16(6):3050–3059, 2015.

[70] Nadim Haque, Nicholas R. Jennings, and Luc Moreau. Scalability and robustness of a market-based network resource allocation system, 2005.

[71] Hannes Hartenstein and Kenneth Laberteaux. *VANET vehicular applications and inter-networking technologies*, volume 1. John Wiley & Sons, 2009.

[72] Hannes Hartenstein and Kenneth P. Laberteaux. A tutorial survey on vehicular ad hoc networks. *Communications Magazine, IEEE*, 46(6):164–171, 2008.

[73] Florian Häusler, Emanuele Crisostomi, Ilja Radusch, and Robert Shorten. Smart Procurement Of Naturally Generated Energy (SPONGE) for PHEV's. *arXiv preprint arXiv:1501.02931*, 2015.

[74] Florian Häusler, Emanuele Crisostomi, Arieh Schlote, Ilja Radusch, and Robert N. Shorten. Stochastic park-and-charge balancing for fully electric and plug-in hybrid vehicles. *Intelligent Transportation Systems, IEEE Transactions on*, 15(2):895–901, April 2014.

[75] Florian Häusler, Mahsa Faizrahnemoon, Emanuele Crisostomi, Arieh Schlote, Ilja Radusch, and Robert Shorten. A framework for optimal real-time emissions trading in large-scale vehicle fleets. In *9th ITS European Congress, Dublin, Ireland*, 2013.

[76] Florian Häusler, Mahsa Faizrahnemoon, Emanuele Crisostomi, Arieh Schlote, Ilja Radusch, and Robert Shorten. A framework for real-time emissions trading in large-scale vehicle fleets. *Intelligent Transport Systems, IET*, 9(3):275–284, 2015.

[77] Florian Häusler, Yingqi Gu, Wynita Griggs, Emanuele Crisostomi, Ilja Radusch, and Robert Shorten. Cooperative control and smart procurement of naturally generated energy (SPONGE) for PHEVs. In *2015 International Conference on Connected Vehicles & Expo (ICCVE 2015)*, Shenzhen, P.R. China, October 2015.

[78] Florian Häusler, Rodrigo H. Ordóñez-Hurtado, Wynita M. Griggs, Ilja Radusch, and Robert N. Shorten. Closed-loop flow regulation with balanced routing. In *Connected Vehicles and Expo (ICCVE), 2014 International Conference on*, pages 1054–1055, Nov 2014.

[79] Florian Häusler (main editor). Deliverable D1.0 TEAM users, stakeholders and use cases (Part A). 2006.

[80] Mayer Hillman and Ian Christie. Pollution, transport and energy. *Policy Studies*, 11(1):66–78, 1990.

[81] B. Hodge and Michael Milligan. Wind power forecasting error distributions over multiple timescales. In *Power and Energy Society General Meeting, 2011 IEEE*, pages 1–8. IEEE, 2011.

[82] Ali Jadbabaie, Jie Lin, and A. Stephen Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *Automatic Control, IEEE Transactions on*, 48(6):988–1001, 2003.

[83] H. Jahn and J. Krey. Mobilität der Stadt - Berliner Verkehr in Zahlen. `http://opus.kobv.de/zlb/volltexte/2012/12983/pdf/Mobilitaet_Berlin_2010.pdf`, 2011.

[84] Olaf Jahn, Rolf H. Möhring, Andreas S. Schulz, and Nicolás E. Stier-Moses. System-optimal routing of traffic flows with user constraints in networks with congestion. *Operations Research*, 53(4):600–616, 2005.

[85] Jan Drebes. Zahl der Krankheitstage stark gestiegen. `http://www.rp-online.de/wirtschaft/krankheit-zahl-der-krankheitstage-stark-gestiegen-aid-1.4728571`, 2014.

[86] S. Kallas, European Commission. Directorate General for Mobility, and Transport. *White Paper on Transport: Roadmap to a Single European Transport Area–towards a Competitive and Resource-efficient Transport System*. Publications Office of the European Union, 2011.

[87] Barry H. Kantowitz, Richard H. Hanowski, and Susan C. Kantowitz. Ergonomics and safety of intelligent driver interfaces (Chapter: Driver reliability requirements for traffic advisory information. pages 1–22. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1997.

[88] Georgios Karagiannis, Onur Altintas, Eylem Ekici, Geert Heijenk, Boangoat Jarupan, Kenneth Lin, and Timothy Weil. Vehicular networking: A survey and tutorial on requirements, architectures, challenges, standards and solutions. *Communications Surveys & Tutorials, IEEE*, 13(4):584–616, 2011.

[89] Christopher M. Kellett, Robert N. Shorten, and Douglas J. Leith. Sizing internet router buffers, active queue management, and the Lur'e problem. In *Decision and Control, 2006 45th IEEE Conference on*, pages 650–654. IEEE, 2006.

[90] F.R. Kelly. Mathematical modeling of the internet. In *Proceedings of ICIAM 99*, Edinburgh, UK, 1999.

[91] Willett Kempton and Amardeep Dhanju. Electric vehicles with V2G. *Windtech international*, 2(2):18, 2006.

[92] Kyoung-Dae Kim and Panganamala R. Kumar. Cyber–physical systems: A perspective at the centennial. *Proceedings of the IEEE*, 100(Special Centennial Issue):1287–1308, 2012.

[93] C. King, Robert N. Shorten, Fabian R. Wirth, and Mehmet Akar. Growth conditions for the global stability of high-speed communication networks with a single congested link. *Automatic Control, IEEE Transactions on*, 53(7):1770–1774, 2008.

[94] Christopher King, Wynita Griggs, Fabian Wirth, and Robert Shorten. Using a car sharing model to alleviate electric vehicle range anxiety. In *Proceedings of the Sixteenth Yale Workshop on Adaptive and Learning Systems*, pages 130–135, 2013.

[95] Florian Knorn. *Topics in cooperative control*. PhD thesis, National University of Ireland, Maynooth, 2011.

[96] Florian Knorn, Martin Corless, and R. Shorten. Results in cooperative control and implicit consensus. *International Journal of Control*, 84(3):476–495, 2011.

[97] Michael Kodransky and Gabrielle Hermann. Europe's parking U-Turn: from accommodation to regulation. *Institute for Transportation and Development Policy*, 2011.

[98] Kraftfahr-Bundesamt. Personenkraftwagen am 1. Januar 2013 nach ausgewählten Merkmalen (Teil 2). `http://www.kba.de/DE/Statistik/Fahrzeuge/Bestand/Ueberblick/2013/2013_b_bestandsbarometer_teil2_absolut.html?nn=793588`, 2014.

[99] Richard A. Krueger and Mary Anne Casey. *Focus groups: A practical guide for applied research*. Sage, 2009.

[100] Amélie le Fauconnier, Eric Gantelet, et al. The time looking for a parking space: strategies, associated nuisances and stakes of parking management in france. 2006.

[101] Fan Li and Yu Wang. Routing in vehicular ad hoc networks: A survey. *Vehicular Technology Magazine, IEEE*, 2(2):12–22, 2007.

[102] Hao Li, Michiel C.J. Bliemer, and Piet H.L. Bovy. Optimal toll design from reliability perspective. In *Proceedings of the 6th Triennnial Conference on Transportation Analysis, Phuket, Thailand*. Citeseer, 2007.

[103] Dong Lin and Robert Morris. Dynamics of random early detection. In *ACM SIGCOMM Computer Communication Review*, volume 27, pages 127–137. ACM, 1997.

[104] Yun-Wei Lin, Yuh-Shyan Chen, and Sing-Ling Lee. Routing Protocols in Vehicular Ad Hoc Networks: A Survey and Future Perspectives. *J. Inf. Sci. Eng.*, 26(3):913–932, 2010.

[105] Chunhua Liu, K.T. Chau, Diyun Wu, and Shuang Gao. Opportunities and challenges of vehicle-to-home, vehicle-to-vehicle, and vehicle-to-grid technologies. *Proceedings of the IEEE*, 101(11):2409–2427, 2013.

[106] Hong K. Lo, X.W. Luo, and Barbara W.Y. Siu. Degradable transport network: travel time budget of travelers with heterogeneous risk aversion. *Transportation Research Part B: Methodological*, 40(9):792–806, 2006.

[107] Stefan Lobach. Adaptive urban emission control by online rerouting. Master thesis, Technical University of Berlin, Germany, December 2013.

[108] T. Logenthiran, D. Srinivasan, and Tan Zong Shun. Demand side management in smart grid using heuristic optimization. *Smart Grid, IEEE Transactions on*, 3(3):1244–1252, Sept 2012.

[109] Henrik Lund and Willett Kempton. Integration of renewable energy into the transport and electricity sectors through V2G. *Energy policy*, 36(9):3578–3587, 2008.

[110] Christiane Malina and Frauke Fischer. The impact of low emission zones on PM10 levels in urban areas in Germany, 2012.

[111] Catriona Manville, Gavin Cochrane, Jonathan Cave, Jeremy Millard, Jeremy Kevin Pederson, Rasmus Kåre Thaarup, Andrea Liebe, Matthias Wissner, Roel Massink, and Bas Kotterink. Mapping smart cities in the EU, 2014.

[112] Michael Manville and Donald Shoup. Parking, people, and cities. *Journal of Urban Planning and Development*, 131(4):233–245, 2005.

[113] M. Marinelli, F. Sossan, G.T. Costanzo, and H.W. Bindner. Testing of a Predictive Control Strategy for Balancing Renewable Sources in a Microgrid. *Sustainable Energy, IEEE Transactions on*, 5(4):1426–1433, Oct 2014.

[114] Miklós Maróti, Branislav Kusy, Gyula Simon, and Ákos Lédeczi. The flooding time synchronization protocol. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 39–49. ACM, 2004.

[115] James A. Misener, Raja Sengupta, and Hariharan Krishnan. Cooperative collision warning: Enabling crash avoidance with wireless technology. In *12th World Congress on ITS*, volume 3. Citeseer, 2005.

[116] Tesla Motors. Tesla battery swapping station. `http://www.teslamotors.com/batteryswap`.

[117] Ronald R. Mourant and Thara R. Thattacherry. Simulator sickness in a virtual environments driving simulator. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 44, pages 534–537. SAGE Publications, 2000.

[118] Angelia Nedić and Asuman Ozdaglar. Subgradient methods for saddle-point problems. *Journal of optimization theory and applications*, 142(1):205–228, 2009.

[119] B. Clifford Neuman. Scale in distributed systems. *ISI/USC*, 1994.

[120] ManWo Ng and S. Travis Waller. Dynamic Route Choice Model in Face of Uncertain Capacities. In *Transportation Research Board 88th Annual Meeting*, number 09-1718, 2009.

[121] Robert B. Noland and William A. Cowart. Analysis of Metropolitan Highway Capacity and the growth in vehicle miles of travel. *Transportation*, 27(4):363–390, 2000.

[122] James R. Norris. *Markov chains*. Number 2008. Cambridge University Press, 1998.

[123] Yasuhiro Ohara, Manav Bhatia, Nakamura Osamu, and Jun Murai. Route flapping effects on ospf. In *Applications and the Internet Workshops, 2003. Proceedings. 2003 Symposium on*, pages 232–237. IEEE, 2003.

[124] Reza Olfati-Saber and Richard M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *Automatic Control, IEEE Transactions on*, 49(9):1520–1533, 2004.

[125] Reza Olfati-Saber and Jeff S. Shamma. Consensus filters for sensor networks and distributed sensor fusion. In *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*, pages 6698–6703. IEEE, 2005.

[126] Rodrigo H. Ordóñez-Hurtado, Wynita M. Griggs, Kay Massow, and Robert N. Shorten. Intelligent speed advising based on cooperative traffic scenario determination. In *Optimization and Optimal Control in Automotive Systems*, pages 77–92. Springer, 2014.

[127] Kaan Ozbay, Dilruba Ozmen-Ertekin, and Joseph Berechman. Contribution of transportation investments to county output. *Transport Policy*, 14(4):317 – 329, 2007.

[128] Alexander Paier, Refi-Tugrul Güner, and Wolfgang Brückler. V2X cooperative systems – on the way to next generation ITS. *Prace Naukowe Politechniki Warszawskiej. Transport*, pages 401–410, 2013.

[129] Kate Pangbourne, Judith Masthoff, Paula Forbes, Silvia Gabrielli, Simon Wells, Cristina Pou, Athanasios Ballis, and Florian Häusler. Influencing travel behavior through segmentation and personalized messages. The work has not yet been submitted., 2015.

[130] European Parliament and of the Council. Directive 2008/50/EC of the European Parliament and of the Council of 21 May 2008 on ambient air quality and cleaner air for Europe. `http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2008:152:0001:01:EN:HTML`, August 2013.

[131] Bernd Propfe, Danny Kreyenberg, Joerg Wind, and Stephan Schmid. Market penetration analysis of electric vehicles in the German passenger car market towards 2030. *International Journal of Hydrogen Energy*, 38(13):5201–5208, 2013.

[132] Kejun Qian, Chengke Zhou, Malcolm Allan, and Yue Yuan. Modeling of load demand due to EV battery charging in distribution systems. *Power Systems, IEEE Transactions on*, 26(2):802–810, 2011.

[133] Hua Qin and Wensheng Zhang. Charging scheduling with minimal waiting in a network of electric vehicles and charging stations. In *Proceedings of the Eighth ACM international workshop on Vehicular inter-networking*, pages 51–60. ACM, 2011.

[134] Andry Rakotonirainy. Design of context-aware systems for vehicle using complex systems paradigms. In *Workshop on safety and context in conjunction with CONTEXT'05*, pages 464–475, 2005.

[135] Kevin Rhinow. Modellierung und wirtschaftliche und empirische Analyse eines kollaborativen Park Sharing Konzeptes. Master thesis, TU Berlin, December 2013.

[136] Bartek Rytt. Compliance bei der Nutzung von kooperativen Assistenzsystemen. Bachelor thesis, TU Berlin, August 2012.

[137] Bartek Rytt. Entwicklung und Validierung von Anreizen zur Akzeptanzsteigerung von systemoptimierenden ITS Applikationen. Master thesis, TU Berlin, February 2015.

[138] A. Schlote, Florian Häusler, Thomas Hecker, Astrid Bergmann, Emanuele Crisostomi, Ilja Radusch, and Robert N. Shorten. Cooperative Regulation of Emissions Using Plug-in Hybrid Vehicles. In *Proceedings of the 2012 International Conference on Connected Vehicles and Expo*, ICCVE '12, pages 201–202, Washington, DC, USA, 2012. IEEE Computer Society.

[139] Arieh Schlote. *New Perspectives on Modelling and Control for Next Generation Intelligent Transport Systems*. PhD thesis, National University of Ireland, Maynooth, 2014.

[140] Arieh Schlote, Florian Häusler, Thomas Hecker, Astrid Bergmann, Emanuele Crisostomi, Ilja Radusch, and Robert Shorten. Cooperative regulation and trading of emissions using plug-in hybrid vehicles. *Intelligent Transportation Systems, IEEE Transactions on*, 14(4):1572–1585, 2013.

[141] Arieh Schlote, Florian Häusler, Thomas Hecker, Astrid Bergmann, Emanuele Crisostomi, Ilja Radusch, and Robert N. Shorten. Cooperative Regulation and Trading of Emissions Using Plug-in Hybrid Vehicles. In *IEEE Transactions on Intelligent Transportation Systems*, pages 1–14. IEEE Intelligent Transportation Systems Society, 2013.

[142] Arieh Schlote, Candice King, Emanuele Crisostomi, and Robert Shorten. Delay-tolerant stochastic algorithms for parking space assignment. *Intelligent Transportation Systems, IEEE Transactions on*, 15(5):1922–1935, 2014.

[143] Matthias Schmidt and Ralf-Peter Schäfer. An integrated simulation system for traffic induced air pollution. *Environmental Modelling & Software*, 13(3):295–303, 1998.

[144] Elmar Schoch, Frank Kargl, Matthias Weber, and Tim Leinmuller. Communication patterns in VANETs. *Communications Magazine, IEEE*, 46(11):119–125, 2008.

[145] Björn Schünemann. V2X simulation runtime infrastructure VSimRTI: An assessment tool to design smart traffic management systems. *Computer Networks*, 55(14):3189–3198, 2011.

[146] Frank Sensfuss, Mario Ragwitz, and Martin Wietschel. Fluktuationen der Windenergie und deren Vorhersagbarkeit bei einem verstärkten Ausbau des Off-Shore Anteils in Deutschland bis 2020. *Proceedings IEWT, Vienna*, 2003.

[147] M.L. Shaltout, D. Chen, A.A. Malikopoulos, and S. Pannala. Multi-disciplinary decision making and optimization for hybrid electric propulsion systems. In *IEEE International Electric Vehicle Conference (IEVC), Florence, Italy*, 2014.

[148] Robert Shorten, Chris King, Fabian Wirth, and Douglas Leith. Modelling TCP congestion control dynamics in drop-tail environments. *Automatica*, 43(3):441–449, 2007.

[149] Robert Shorten, Fabian Wirth, and Douglas Leith. A positive systems model of TCP-like congestion control: asymptotic results. *IEEE/ACM Trans. Netw.*, 14(3):616–629, June 2006.

[150] Robert Shorten, Fabian Wirth, Oliver Mason, Kai Wulff, and Christopher King. Stability criteria for switched and hybrid systems. *SIAM review*, 49(4):545–592, 2007.

[151] Donald C. Shoup. Cruising for parking. *Transport Policy*, 13(6):479–486, 2006.

[152] Rodrigo Silva, Satoru Noguchi, Thierry Ernst, Arnaud de La Fortelle, and Walter Godoy Junior. Standards for Cooperative Intelligent Transportation Systems: a Proof of Concept. In *The Tenth Advanced International Conference on Telecommunications (AICT)*, 2014.

[153] Kenneth A. Small and Camilla Kazimi. On the costs of air pollution from motor vehicles. *Journal of Transport Economics and policy*, pages 7–32, 1995.

[154] Kenneth A. Small and Camilla Kazimi. On the Costs of Air Pollution from Motor Vehicles. University of california transportation center, working papers, University of California Transportation Center, 1995.

[155] Matheus Souza, Jose C. Geromel, Patrizio Colaneri, and Robert Shorten. On the Discretisation of Sparse Linear Systems. In *Proceedings of the European Control Conference 2014*. IEEE, 2014.

[156] R. Srikant. *The Mathematics of Internet Congestion Control*. Systems & Control. Birkhäuser, 2004.

[157] Rayadurgam Srikant. *The mathematics of Internet congestion control*. Systems & control. Birkhäuser, Boston, 2004.

[158] R. Stahlmann, A. Festag, A. Tomatis, I. Radusch, and F. Fischer. Starting European field tests for Car-2-X communication: the DRIVE C2X framework. In *Proceedings 18th ITS World Congress and Exhibition*, 2011.

[159] Rade Stanojevic. Router-Based Algorithms for Improving Internet Quality of Service. PhD Thesis, Hamilton Institute, NUI Maynooth, Ireland., 2007.

[160] Rade Stanojevic and Robert Shorten. Load balancing vs. distributed rate limiting: a unifying framework for cloud control. In *Communications, 2009. ICC'09. IEEE International Conference on*, pages 1–6. IEEE, 2009.

[161] H. Steven. Comparison of different European databases with respect to road category and time periods (on peak, off peak, weekend). `http://www.unece.org/fileadmin/DAM/trans/doc/2010/wp29grpe/WLTP-DHC-03-04e.ppt`, 2010.

[162] S. Stüdli, E. Crisostomi, R. Middleton, and R. Shorten. A flexible distributed framework for realising electric and plug-in hybrid vehicle charging policies. *International Journal of Control*, 85(8):1130–1145, 2012.

[163] S. Stüdli, Emanuele Crisostomi, Richard Middleton, and Robert Shorten. AIMD-like algorithms for charging electric and plug-in hybrid vehicles. In *Electric Vehicle Conference (IEVC), 2012 IEEE International*, pages 1–8. IEEE, 2012.

[164] Sonja Stüdli, Emanuele Crisostomi, R. Middleton, and Robert Shorten. Optimal real-time distributed V2G and G2V management of electric vehicles. *International Journal of Control*, 87(6):1153–1162, 2014.

[165] Sonja Stüdli, Emanuele Crisostomi, Richard Middleton, Julio Braslavsky, and Robert Shorten. Distributed Load Management Using Additive Increase Multiplicative Decrease Based Techniques. In *Plug In Electric Vehicles in Smart Grids*, pages 173–202. Springer, 2015.

[166] Sonja Stüdli, Wynita Griggs, Emanuele Crisostomi, and Robert Shorten. On optimality criteria for reverse charging of electric vehicles. *Intelligent Transportation Systems, IEEE Transactions on*, 15(1):451–456, 2014.

[167] Bharath Sundararaman, Ugo Buy, and Ajay D. Kshemkalyani. Clock synchronization for wireless sensor networks: a survey. *Ad Hoc Networks*, 3(3):281–323, 2005.

[168] Gerard Tel. *Introduction to distributed algorithms*. Cambridge University Press, 2000.

[169] Richard H. Thaler and Cass R. Sunstein. Nudge: improving decisions about health, wealth, and happiness. 2008. *Newhaven: Yale*, 2009.

[170] Tessa Tielert, Moritz Killat, Hannes Hartenstein, Raphael Luz, Stefan Hausberger, and Thomas Benz. The impact of traffic-light-to-vehicle communication on fuel consumption and emissions. In *Internet of Things (IOT), 2010*, pages 1–8. IEEE, 2010.

[171] Andrea Tomatis, Markus Miche, Florian Häusler, Massimiliano Lenardi, Thomas Michael Bohnert, and Ilja Radusch. A test architecture for V-2-X cooperative systems field operational tests. In *Intelligent Transport Systems Telecommunications,(ITST), 2009 9th International Conference on*, pages 616–621. IEEE, 2009.

[172] Lai Tu and Chung-Ming Huang. Forwards: A map-free intersection collision-warning system for all road patterns. *Vehicular Technology, IEEE Transactions on*, 59(7):3233–3248, 2010.

[173] Mosaddek Hossain Kamal Tushar, Chadi Assi, Martin Maier, and Mohammad Faisal Uddin. Smart microgrids: Optimal joint scheduling for electric vehicles and home appliances. *Smart Grid, IEEE Transactions on*, 5(1):239–250, 2014.

[174] Umweltbundesamt. Umweltzonen in Deutschland. http://gis.uba.de/website/umweltzonen/umweltzonen.php, August 2013.

[175] Matteo Vasirani and Sascha Ossowski. A Market-Inspired Approach for Intersection Management in Urban Road Traffic Networks. *J. Artif. Intell. Res. (JAIR)*, 43:621–659, 2012.

[176] Axel Wegener, Michał Piórkowski, Maxim Raya, Horst Hellbrück, Stefan Fischer, and Jean-Pierre Hubaux. TraCI: an interface for coupling road traffic and network simulators. In *Proceedings of the 11th communications and networking simulation symposium*, pages 155–163. ACM, 2008.

[177] Richard J. Weiland and Lara Baughman Purser. Intelligent transportation systems. *Transportation in the new millennium*, 2000.

[178] Christian Weiß. V2X communication in Europe – From research projects towards standardization and field testing of vehicle communication technology. *Computer Networks*, 55(14):3103–3119, 2011.

[179] John Whitelegg and Gary Haq. Vision Zero: Adopting a target of zero for road traffic fatalities and serious injuries. *Stockholm Environment Institute*, 2006.

[180] Tobias Wiesenthal, Guillaume Leduc, J. Koehler, Wolfgang Schade, and Burkhard Schade. Research of the EU automotive industry into low-carbon vehicles and the role of public intervention. Working Papers 058727, Institute for Prospective and Technological Studies, Joint Research Centre, 2010.

[181] Wikipedia. Verisimilitude — Wikipedia, the free encyclopedia, 2015. [Online; accessed 24-September-2015].

[182] Bob Williams. *Intelligent transport systems standards.* Artech House, 2008.

[183] Fabian Wirth, Sonja Stüdli, Jia Yuan Yu, Martin Corless, and Robert Shorten. Nonhomogeneous Place-Dependent Markov Chains, Unsynchronised AIMD, and Network Utility Maximization. *arXiv preprint arXiv:1404.5064*, 2014.

[184] Wayne Wolf. Cyber-physical Systems. *Computer*, 42(3):88–89, March 2009.

[185] Hendrik Wolff and Lisa Perry. Policy monitor. *Review of Environmental Economics and Policy*, 4(2):293–308, 2010.

[186] Hendrik Wolff and Lisa Perry. Policy monitor trends in clean air legislation in europe: Particulate matter and low emission zones. *Review of Environmental Economics and Policy*, 4(2):293–308, 2010.

[187] Owen Worley, Diego Klabjan, and Timothy M. Sweda. Simultaneous vehicle routing and charging station siting for commercial electric vehicles. In *Electric Vehicle Conference (IEVC), 2012 IEEE International*, pages 1–3. IEEE, 2012.

[188] Jun Yang, Ying Liu, Ping Qin, and A.A. Liu. A review of Beijing's vehicle lottery: short-term effects on vehicle growth, congestion, and fuel consumption. *Energy policy*, 2014.

[189] Jin Y. Yen. Finding the k shortest loopless paths in a network. *Management Science*, 17(11):712–716, 1971.

[190] Steve H.L. Yim and Steven R.H. Barrett. Public health impacts of combustion emissions in the United Kingdom. *Environmental science & technology*, 46(8):4291–4296, 2012.

[191] Jie Zhang, Bri-Mathias Hodge, Anthony Florita, Siyuan Lu, Hendrik F. Hamann, and Venkat Banunarayanan. Metrics for evaluating the accuracy of solar power forecasting. In *3rd International workshop on integration of solar power into power systems, London, England*, 2013.

[192] Chenguang Zhu, Jia Shuo Yue, Chinmoy V. Mandayam, Deepak Merugu, Hossein Karkeh Abadi, and Balaji Prabhakar. Reducing road congestion through incentives: a case study. In *Transportation Research Board 94th Annual Meeting*, number 15-0403, 2015.