Stochastic models for road traffic control

Kovács, P.

Publication date
2016
Document Version
Final published version
License
Other

Stochastic Models for Road Traffic Control

Péter Kovács

Stochastic Models for Road Traffic Control

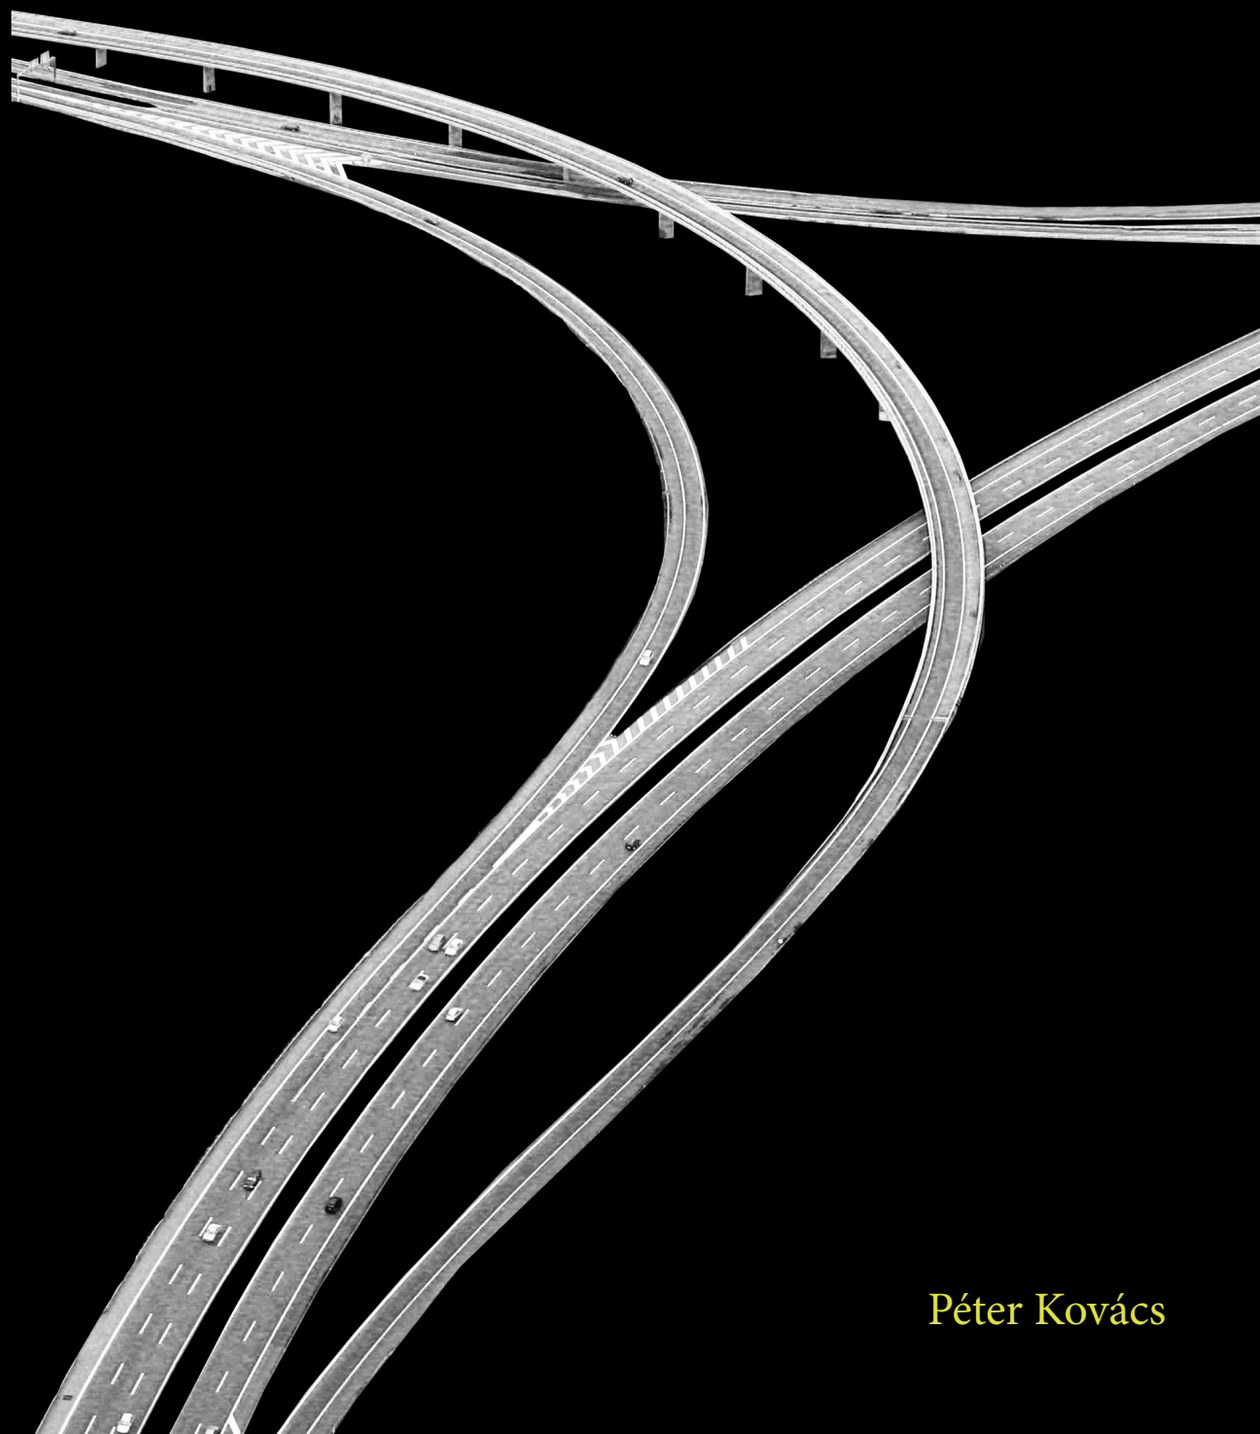

Péter Kovács

Stochastic Models for Road Traffic Control

ISBN 978-94-6295-555-4

Péter Kovács
Korteweg-de Vries Instituut voor Wiskunde
Faculteit der Natuurwetenschappen, Wiskunde en Informatica
Universiteit van Amsterdam

Printing & Lay Out: Uitgeverij BOXpress || proefschriftmaken.nl
Cover illustration: Archiko Kft.
Publisher: Uitgeverij BOXPress, Vianen

Stochastic Models for Road Traffic Control

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus prof. dr. ir. K. I. J. Maex
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op woensdag 14 december 2016, te 10:00 uur

door

Péter Kovács

geboren te Debrecen, Hongarije

Promotiecommissie

Promotors:	Prof. dr. R. Núñez-Queija	Universiteit van Amsterdam
Copromotor:	Dr. N. S. Walton	Universiteit van Amsterdam, University of Manchester
Overige leden:	Prof. dr. J. L. van den Berg	Universiteit Twente, TNO
	Dr. ir. M. A. A. Boon	Technische Universiteit Eindhoven
	Prof. dr. R. J. Boucherie	Universiteit Twente
	Prof. dr. D. T. Crommelin	Universiteit van Amsterdam
	Prof. dr. M. R. H. Mandjes	Universiteit van Amsterdam
	Dr. ir. E. M. M. Winands	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Acknowledgements

I wish to thank all the people who supported me during my time as a PhD candidate at the University of Amsterdam (UvA). First of all, I would like to thank my supervisor Sindo Núñez-Queija and my co-supervisor Neil Walton, who helped me in numerous ways during my research. They followed my work, sharing their wisdom, insights, questions, doubts and motivating words, whichever I needed at the time. I would also like to thank my family for being there for me to lift my spirits and reinforce my decisions, whenever I doubted myself. Furthermore, I would like to thank all my friends and colleagues who made my time at UvA joyful and filled it with memories that I will forever cherish.

Most of the work presented in this dissertation is the output of several fruitful collaborations while I was employed at UvA. More specifically, my thesis is based on a series of papers that I have written together with peers from within and outside of the university. Here I enlist all of these publications, detailing my role in the projects that lead to them.

- Chapter 2 is based on [1], which has been written with Wendy Ellens (UvA), Sindo Núñez-Queija and Hans van den Berg (University of Twente). Wendy also discussed this topic in her thesis [43]. I have shared an equal load with Wendy in performing the numerical experiments and in writing the paper with Sindo and Hans overseeing our work. The fluid results of Section 2.4 were only conjectured in the original publication. I have completed the formal proof, which is contained in Section 2.5, after the submission under Neil Walton's supervision.

- The work described in Chapters 3, 4 and 5 has been carried out in collaboration with a group from Swinburne University of Technology, Melbourne, Australia. The corresponding publications are [3] and [4]. In both cases I took the lead of the theoretical parts with a great deal of help from Neil, while the numerical experiments were lead by Tung Le under Hai Vu's supervision (both from Swinburne). The first paper, [4] was overseen by Lachlan Andrew (Swinburne)

and Serge Hoogendoorn (Delft University of Technology) as well, while Sindo consulted us on the second paper, [3].

- The first three sections of Chapter 7 correspond to [2], a paper with Neil and Gaurav Raina (Indian Institute of Technology Madras). Most of my effort has gone into Sections 2, 3 and 5 of this paper and the overall coordination, while Neil took the lead in Section 4. Section 7.4 of the same chapter describes further advancements in this research, which were mostly carried out by myself and Abhishek (UvA) with Gaurav and Sindo overseeing the progress.

I am grateful to all of my collaborators and hope the we can successfully engage in research projects in the future.

Péter Kovács

Amsterdam, 31 August 2016

Contents

Chapter 1

Introduction

Outline. Road traffic is so deeply imbedded in modern life that almost everyone has personal experience of it. Since freedom of movement feels natural to most of us, mostly we remember times when something in traffic went wrong. While this can involve personal accidents or other mishaps, we have all experienced and can recite occasions when traffic broke down. These breakdowns, while also being a nuisance to the individual, put a great cost on the economy and the environment alike. Therefore we expect controllers to provide us the opportunity of safe, smooth and uninterrupted travel by minimizing the possibility of traffic failure.

In recent years the engineering community has turned its attention more and more towards mathematical models to provide solutions for improved and more efficient road traffic management. The questions that arise present a great challenge to mathematicians, since the structures involved are *complex networks* with a lot of users making varying sorts of individual decisions. This dissertation investigates models that aim to describe segments of road traffic and control techniques that provide efficient means in easily implementable algorithms for traffic controllers. Throughout most of this work the focus is on two areas, namely *queueing theory* and *optimisation* in the fields of stochastics and operations research. The goal of this first chapter is to place the work presented in a societal and scientific (mathematical) context, whilst explaining its main ideas and contributions.

1.1 Historical overview

1.1.1 History of road traffic

The goal of this thesis is to introduce mathematical models for road traffic. In order to put our work into context we start this introductory chapter by a brief review of the evolution of road traffic. This section also serves to introduce the main technical terms regarding the road network that we shall use later on. We only discuss the most important inventions and developments, for a detailed overview we refer the reader to the book of Lay and Vance [82].

Early forms of roads and transportation

Many historians define the beginnings of civilisation as the time when groups of humans first erected settlements, which raised the possibility of an organised way of living, where different parts of the community would take part in different, specialised activities. Even temporary living areas required that the inhabitants followed the same paths over and over during their daily routine, therefore naturally forming the first *roads*. As these settlements evolved into the first towns and cities the communities grew so large that organised transportation of goods, mostly food, became a necessity. Archeological evidence suggests that the first constructed roads, such as stone paved roads in the Mesopotamian city of Ur and in Harappa in the Indus Valley or timber roads near Glastonbury, England, were already in use around 4000 BC, even before the wheel was invented. As the first towns and their networks of pathways extended, and in the same time bureaucratic practices developed, *urban planning* became the first method of regulating traffic. Designed cities were characteristic of the Minoan, Mesopotamian, Indus Saraswati and Egyptian civilisations of the 3rd millennium BC. The first recorded description of urban planning appears in the Epic of Gilgamesh, which was written around 2100 BC. The streets of many of these early cities were paved and laid out at right angles in a grid pattern, with a *hierarchy of streets* from major boulevards to residential alleys.

The next big step in the complexity of transportation networks was largely due to trade. The first long-range trade routes already appeared in the 3rd millennium BC, however they were water based. Land-based

transportation only became common practice with the domestication of the horse and the camel, widespread use of the wheel and when empires formed that could guarantee safety. The first directions where systematic trade began were along the Incense Route and the Silk Road, which both refer to an extensive network of maritime and caravan routes, the first connecting Egypt with Mesopotamia, the latter connecting the Mediterranean with East Asia and the Indian subcontinent. Another function of land based transportation was introduced in the Achaemenid Empire, when a postal system was formed, which was based on several relay stations, that helped to cement the rule of the central government by spreading orders and information faster. The Roman Empire developed these ideas further by constructing a large network of *paved roads* connecting its distant regions, which also allowed for faster movement of the military.

The inception of traffic control

From a technological standpoint no substantial development took place in transportation until the Industrial Revolution. The biggest change was perhaps that animal- or manpowered forms of land-based transportation were largely replaced by *mechanised vehicles*. The first working steam-powered vehicle was designed by Ferdinand Verbiest, a Flemish member of a Jesuit mission in China around 1672, however it is not known if Verbiest's model was ever built. Nicolas-Joseph Cugnot is widely credited with building the first full-scale, self-propelled mechanical vehicle in about 1769, when he created a steam-powered tricycle. In 1781 James Watt invented the first steam engine that could produce continued rotary motion, and in 1784 he patented the design for the first steam locomotive. Thus the first mechanised vehicles with widespread use were trains, which needed rails, not roads, however several means of traffic management were introduced in their context such as *time plans* or *safety regulations*.

Road networks still remained relevant, especially in urban areas as the construction cost of railways was relatively high and they would have been inefficient to use in cities. Furthermore there were technological and methodological advancements in road construction as well. In the late 18th, early 19th century Pierre-Marie-Jérôme Trésaguet and later Thomas Telford were the firsts since the Roman era to bring science into

paving, while the most impactful inventor was John Loudon McAdam. Macadam roads opened up the possibility of faster travelling, which together with the increasing population and the fast urbanization also introduced the need for traffic control. The first step on governmental level in this direction happened in the United Kingdom when in 1832, the Stage Carriage Act introduced the offense of endangering the safety of a passenger or person by "furious driving". The first *speed limits* were created there too, by a series of Locomotive Acts (1861, 1865 and 1878). The 1861 Act introduced a speed limit of 10 mph (16 km/h) on open roads in town, which was reduced to 2 mph (3 km/h) in towns and 4 mph (6 km/h) in rural areas by the 1865 "red flag act".

Another method for traffic control also emerged from the United Kingdom, namely the *traffic light*. The first, promoted by railway engineer John Peake Knight, was installed on 9 December 1868 outside the Houses of Parliament in London, although it used gas lights and had to be operated by a police constable. It introduced the red and green signals still in use today, however it did not operate for long because of an accident caused by gas leaking. Due to safety concerns the concept was abandoned until Lester Farnsworth Wire invented the first electric traffic light in 1912. His design, which also used red and green lights, was installed in Salt Lake City, Utah in the United States, which also became the home to the first *interconnected traffic light system* in 1917. Further developments were made by James Hoge, who introduced a buzzer, and William Potts, who added the yellow sign, to improve safety by providing warnings before signal changes. The first systems were controlled manually through switches from police stations and fire departments. Automatic control of interconnected traffic lights was introduced March 1922 in Houston, Texas. From the 1920s onwards the use of traffic lights became widespread all over the world.

From the first car to modern road traffic

More control on roads became a necessity due to the increasing number of fast moving vehicles. While steam-powered road vehicles were unable to sustain sufficient steam pressure for long periods, the emergence of the *internal combustion engine* provided the technology needed to build cars that could be of practical use. Although the first patent for the

internal combustion engine was granted to Nicéphore Niépce in 1807, it was not until the end of the 19th century and the inventions of Nikolaus Otto and Karl Benz that the creation was adapted to be used in road vehicles. Benz is acknowledged as the inventor of the *modern car* for his 1878 design, for which he was granted a patent a year later. The mass production of automobiles started in 1901, when Ransom Olds opened his Oldsmobile factory and was greatly expanded by Henry Ford beginning in 1913.

The widespread use of cars was further helped by other inventions and developments. In 1847 Robert William Thomson patented the first pneumatic tyre. This was further developed and put into mass production by John Boyd Dunlop in the late 1890s. Even more important was the improving quality of roads available. Developments in chemical engineering allowed for better materials to be used in the macadam process. In 1870 Edward de Smedt patented the use of asphalt concrete panels, while modern tarmac was patented by Edgar Purnell Hooley in 1901. This opened up the possibility for faster travelling speeds, which was especially useful in interurban transport. To further help the process *high-speed roads* were built, which were referred to as dual highways, although they bore little resemblance to the *highways* of today. The first of them opened in Italy between Milan and Varese in 1924, it now forms parts of the A8 and A9 motorways. The concept of highways, freeways or motorways (the term referring to this type of road changes from country to country and is defined by local statute or design standards most of the time) has developed over the years and varies geographically, however in the modern description the following defining properties apply. The road is divided with guard rails separating the directions. Highways have no at-grade intersections with other roads, railroads or multi-use trails, the crossing of freeways by other routes is typically achieved with grade separation either in the form of underpasses or overpasses. There is a *control of access*, where the term relates to a legal status which limits the types of vehicles that can use a highway, as well as a road design that limits the points at which they can access it. Access to freeways is typically provided only at grade-separated interchanges, though lower-standard right-in/right-out access can be used for direct connections to side roads. Speed limits are generally higher on highways and are occasionally nonexistent.

With all of the basic concepts and technologies that we connect with modern road transportation present, from the 1920s onwards the world has seen a rapid extension of the networks. Various types of design and safety regulations started to be enforced in most jurisdictions regarding both roads, vehicles and the behaviour of their users. At the same time innovators worked on reducing costs, travel times and increasing safety coming up with such a large scale of inventions that covering all of them is outside of the scope of this work. Therefore we only highlight those that we find the most important. The evolution of traffic signal systems, followed the route of automatisation. The first city to computerise its entire traffic light network was Toronto in Ontario, Canada, which accomplished this in 1963. The 1960s also saw the first signal system aimed at regulating the flow of freeway traffic, when *ramp metering* was deployed on the Eisenhower Expressway in Chicago in Illinois, United States. Another method for the same purpose is the use of *variable speed limits*. The first system of mechanically variable displays that allowed this was implemented in 1965 on the A8 motorway between Munich, Germany and Salzburg, Austria. In order to make full use of these technologies it is important for the traffic controller to have up to date information of the state of the roads. While the first algorithms governing the settings of the various signal systems were based on heuristics or statistical data, recent developments allow for an online measurement of the traffic flow. The most widespread of such innovations are *inductive loop detectors*. State of the art systems are capable of real time vehicle detection, which raises the possibility of adjusting variable traffic signals in a timely manner too.

Another recent development is the possibility of continouosly locating vehicles, while maintaining communication with them. The Global Positioning System (GPS), which was declared available for civilian use in 1996, has gained widespread recognition and use in the past decade. While its most manifest purpose is *route planning*, it also raises the possibility of traffic control through the individual by *user guidance*. It is widely believed that the future of transportation lies in autonomous vehicles, which first appeared in the 1980s and have been heavily researched since, albeit their use in public traffic is still in the experimental phase. These machines present the possibility of even more control of the individual actors of traffic and through that ever more

efficient methods for travelling.

The impact of road traffic

So far we have discussed the evolution of road traffic. Up until the 20th century, when modern extensive road networks were built, this process netted mostly positive results for society. The faster movement of people, goods and information clearly allowed for an increase in the quality of life, however as time progressed negative effects started to show as well. The most visible of these is the environmental impact that road traffic has. Air pollution from fossil (and some biofuel) powered vehicles can occur wherever vehicles are used and are of particular concern in congested city street conditions and other low speed circumstances. Emissions include particulate emissions from diesel engines, nitrogen oxides, volatile organic compounds, carbon monoxide and various other hazardous air pollutants. Motor vehicle traffic on roads generate noise, which have health effects on those living nearby. Urban runoff from roads and other impervious surfaces is a major source of water pollution, because rainwater and snowmelt picks up gasoline, motor oil, heavy metals, trash and other pollutants. Roads can act as barriers or filters to animal movement and lead to habitat fragmentation. The list goes on and on. As human population increases the need for transportation is bound to increase too. To avoid further damaging our living space the *efficient use of existing infrastructures* is of utmost importance and provides the goal of the traffic modelling and controlling community.

1.1.2 History of road traffic research

After discussing the ever-growing complexity of road traffic networks in Section 1.1.1, here we provide a brief overview of the history of its scientific investigation. This short section is based on a memoir by G. F. Newell [109], one of the first scientists to focus on the subject.

The *scientific investigation of road traffic* dates back to the middle of the 20th century. After the Second World War researchers in newly developing fields such as operations research, econometry or statistical physics were seeking areas where their freshly established theories and concepts (some of which emerged first due to military use during the

war) could be applied. With a boom in automobile production in the 1950s the field of *traffic flow and traffic control theory* presented itself as an obviously relevant choice. Often cited early publications in the subject include the works of M. Beckmann (economic models of various forms of transportation) [12], L.C. Edie (traffic delays at toll booths) [42], G. F. Newell (analytic model for freely-flowing highway traffic) [108], J. G. Wardrop (distinction between user optimal and system optimal route choice) [151], F. V. Webster (delays for a fixed-cycle traffic signal) [154] or the famous papers by M. J. Lighthill and G. B. Whitham in 1955 [85, 86] and P. L. Richards in 1956 [125], which described traffic waves and formulated the well-known Lighthill-Whitham-Richards model. These pioneering researchers often worked without being aware of each other's results. The first attempt to bring scientists interested in the topic resulted in an international symposium on "The Theory of Traffic Flow", which was organised by R. Herman in 1959 in Detroit, Michigan, United States. This was followed by a similar meeting organised by R. Smeed in 1963 in London, England and turned into a series of conferences that still continues in the present day under the title "International Symposium on Transportation and Traffic Theory" (ISTTT). On an organisational level the topic was first taken care of, when the Operation Research Society of America (ORSA now part of INFORMS) formed its Transportation Science section in 1962, which is responsible for the journals Transportation Science and Transportation Research. The research community have been active ever since, gaining more momentum in the past decade with the emergence of wireless technologies and the increase in computational capacity, adapting numerous results from other fields such as telecommunication or data networks. The work presented in this dissertation aims at furthering the knowledge gathered over the years.

1.2 Methodological background

The mathematical models that are developed for road traffic mainly investigate two aspects of the subject: description and control. While these are closely related, the methods applied to them cover a wide range of fields.

When focusing on describing the state of traffic, one has to consider many vehicles, whose drivers' decisions are mostly assumed to be unknown. These range from timings (i.e. when to join the system) through route planning (i.e. which roads to take) to setting the individual vehicles' speeds. The most common technique to take the variability of these decisions of different drivers into account is using *stochastic methods*. Models from *queueing theory* lend themselves to the subject of urban traffic. Motorway traffic on the other hand is usually described by models that emerged from statistical physics, such as *interacting particle systems*.

The decisions made by the traffic controller on the other hand are usually given by *operations research* techniques. They often cover finding an *optimal schedule*. In most cases optimality is different from the drivers' and the system's point of view. In most cases optimality is different from the drivers' and the system's point of view, see the famous result of Wardrop [151] for an example. There are concepts though, which help forming the control decisions, while also placing some requirements on them, such as *fairness* or *stability*.

In this section we give a brief overview of the methods and concepts mentioned above as they play an essential role in this dissertation.

1.2.1 Stochastic processes

A *stochastic process* is a collection of random variables representing the evolution of a (random) system over time. Time can be considered discrete when the changes in the state of the network are observed in a cyclic manner, or continuous when the focus is on the events themselves that drive the evolution.

Markov processes are an important type of stochastic processes (see e.g. [110] for an introduction). They consider systems that have countably many states and evolve between these states satisfying the *Markov property*, which means that the transition probabilities do not depend on the history of the system, only the current state. The Markov assumption greatly simplifies models that incorporate individual drivers' decisions, however it makes them mathematically tractable. As a consequence of the Markov property, the system is *memoryless*, which means that the times between transitions are independent from the past. This means that the time until the next event that changes the state of the system

is *distributed exponentially*.

A typical example of a Markov process is a *Poisson process* (see e.g. [128] for a rigorous definition). In this dissertation we will mostly use Poisson processes as *counting processes*, thus they will represent the total number of occurrences of events that happen during the time interval $[0, t]$, for example to count the number of vehicles that join a road network (either by arriving from outside the network or by starting a vehicle that was parked on one of the roads) over time. We do so because Poisson processes have two important features pertaining to them, namely that the number of occurrences over disjoint time intervals are independent, and that in the homogeneous case (or over time intervals when the rate of the process does not change) the timings of events on an interval, conditioned on a given number of occurrences, follow a *uniform distribution*. Both of these properties are realistic assumptions in most models. One of the defining properties of a Poisson process, namely that the number of occurrences over a time interval follows the *Poisson distribution* also simplifies calculation. The exact formula is given as

$$\mathbb{P}(N([s,t]) = n) = \frac{(\lambda(t-s))^n}{n!}\mathrm{e}^{-\lambda(t-s)},$$

where $n = 0, 1, \ldots$ and the process if homogeneous over $[s, t]$. The rate λ often lends itself to a simple interpretation of the model settings as it can correspond to actual measurements (e.g. average number of vehicles that pass a sensor over time).

A family of Markov processes are often referred to as *stochastic interacting particle systems*, for an extensive survey we point the reader to the book by Liggett [84]. These models stem from statistical physics, hence the name. Their state spaces are countable sets of sites which are possibly occupied by countably many particles. Transitions between states are a result of the movement of the particles, while their interaction is exhibited through the transition rates. Namely, the rates depend on the quantity of particles present at each site. The most prominent examples of such models are the asymmetric simple exclusion process (ASEP) and its totally asymmetric version (TASEP), which both possess the desirable property of having a product from stationary measure and thus being exactly solvable. Other important features of them are the

conservation of particles (i.e. particles cannot appear in vanish from the system other than entering or leaving through the boundaries) and the exclusion property (i.e. particles can only move into empty spaces, but not into a space occupied by another particle). These form the basis of establishing a one-to-one correspondence between particles and vehicles and using these models for road traffic related applications, predominantly to describe the traffic flow on road segments on a vehicle-to-vehicle scale.

Queueing models

The results of *queueing theory*, although they apply to purely deterministic models as well, are discussed in the context of stochastic processes. The general terminology refers to the basic elements of the models as jobs and servers, which we will keep in this short introduction. The traffic related applications lend themselves to the field naturally, as jobs can correspond to vehicles or sets of vehicles, whilst service can be associated with passing through a junction or a specified road segment.

A queueing model (see e.g. [126] for a detailed introduction) for a single queue (or a node in a queueing network) is usually described as $A/S/c$, where A corresponds to the time between job arrivals, S to the service time (or size) of a job and c to the number of servers. A common example is the $M/M/1$ queue. In this case M stands for Markovian or memoryless, which means that the interarrival and the service times are exponentially distributed, while their respective rates are usually denoted by λ and μ. Since the interarrival times are exponential and independent, jobs arrive according to a Poisson process. Trivially there is only 1 server in this case. Another widely used model is the $M/M/\infty$ queue. In this model each arriving job will be directly served, that is, the serving capacity is unlimited. Generally (unless stated otherwise) it is assumed that jobs are being served in the order of their arrival (FIFO, first-in-first-out) and that all arriving jobs are either being served directly (if there is an idle server available) or accommodated in a buffer.

For both examples the number of jobs in the system can be formed as a birth-death process. The corresponding birth rates and death rates are λ and μ respectively for the $M/M/1$ queue. For the $M/M/\infty$ queue, the birth rates are the same, whereas the death rates are $\mu_n =$

$n\mu$, where n denotes the number of jobs in the system. For birth-death processes the *stationary distribution* (also called steady-state or equilibrium distribution) of the number of jobs in the system can be calculated explicitly and gives the distribution over the system states in the long run. In general, for queueing models with exponential interarrival times and/or service times, detailed information is available about their stationary behaviour, this however ceases to be true in complex networks. Results that can be used in a broad range of models include *Little's law* [88], which states that the long-term average number of customers in a stable system, L is equal to the long-term average effective arrival rate, λ, multiplied by the average time a customer spends in the system (also called *sojourn time*), W, or expressed algebraically:

$$L = \lambda W.$$

Another prominent result is the so-called *PASTA property* [155], which stands for "**P**oisson **A**rrivals **S**ee **T**ime **A**verages", meaning that the probability of the state of the system as seen by an outside random observer is the same as the probability of the state seen by an arriving customer. Furthermore the method referred to as *mean value analysis (MVA)* [122, 123] is worth mentioning. It is a recursive technique for computing expected queue lengths, waiting time at queueing nodes and throughput in equilibrium for a closed system of queues.

For road traffic (and other) applications an important family of queueing models are *polling models*. They describe a system where a single server visits a set of queues in some order, see [137] for an introduction or [19] for more recent results. This is typically assumed to happen in a cyclic manner, a property that will be prominent in Chapter 5, since these models lend themselves easily to describe the behaviour of traffic lights.

Fluid scaling

In loose terms *fluid scaling* refers to speeding up time, while decreasing the effects of transitions by rescaling the jumps in the state space with both scalings happening linearly, in order to eliminate minor fluctuations and determine the direction of the long-term evolution of a stochastic

system. The term was first introduced by Kurtz [78] publishing a law of large numbers and central limit theorem for Markov chains. Formally, we consider a sequence of stochastic processes indexed by $n \in \mathbb{N}$, that are versions of the original process $X(t)$. To ensure that the initial values are equal we require $X^n(t) = nX(0)$. The fluid-scaled values are then given by

$$\overline{X}^n(t) = \frac{X(nt)}{n},$$

and if the sequence $\overline{X}^n(t)$ converges to some process $\overline{X}(t)$, which has continuous characteristics and a deterministic fluid input, then this limit is referred to as the *fluid limit* of $X(t)$.

Although in most cases the fluid limit of a process does not have an explicit form, determining the (ordinary or partial) differential equations driving its evolution often carries some benefits, which include, while not limited to, the following. In some cases the fluid-scaled versions of some model parameters can be easily associated with observable quantities of the system that we are aiming to model, which might help understanding the role of these parameters and validate numerical results. The differential equations describing the fluid limit can exhibit a specific behaviour (e.g. the existence of wave solutions in the fluid limits of traffic flow models) which explains the evolution on a macroscopic level. A highly important result of this type was presented by Dai [34], who proved that the stability of a family of queueing models is linked to the stability of the corresponding fluid-scaled models.

In general fluid scaling is a powerful tool to investigate the evolution and stability of stochastic processes, which also proves useful when analysing and comparing different policies by which decisions are made in complex networks.

Heavy traffic scaling

A key performance requirement for queueing systems is *stability*. Loosely speaking, stability means that the number of jobs in the system cannot grow unboundedly, or in other words, that the system is able to handle all service requests. In general stability may depend on the network structure, scheduling and routing decisions by the servers or the users, the arrival rates and service rates at the servers etc. When the load

approaches the stability limit, the system is said to be in the *heavy traffic* regime. Analysing the system in the heavy traffic regime can provide useful intuition as to how the system behaves when it is close to saturation, with typical results covering questions of optimal control, queue length approximations and state-space collapse. The first to study stationary queues in heavy traffic was Kingman [74]. The stability limit in the single server is approached as $\rho \uparrow 1$, with ρ denoting the load on the system. Kingman proved that the distribution of the stationary queue length, scaled with $1 - \rho$, converges to that of an exponential random variable.

The transient queueing processes of networks in heavy traffic also have received much attention in literature. In the limit, the scaled queue length processes essentially behave as reflected Brownian motion processes, see for example [13, 27, 59, 79] for different network models that obey such a law. Similar to the previously described fluid scaling regime, in this approach the entire evolution of the stochastic queueing process is investigated as the load reaches its bound.

In this thesis (in Chapter 5) we will only consider heavy traffic scaling limits of stationary distributions. For illustration, let us consider the standard $M/G/1$ queue in heavy traffic. The stability condition in this case is given by $\rho := \lambda\mathbb{E}[B] < 1$, with λ denoting the arrival rate and B a generic service time. The heavy traffic limit is reached as $\rho \to 1$. If we use X to denote the stationary queue length, the distribution of $(1-\rho)X$ has a proper limit [74]:

$$\lim_{\rho\to 1} \mathbb{E}\left[\mathrm{e}^{-s(1-\rho)X}\right] = \frac{1}{1 + \frac{1}{2}\lambda s \frac{\mathbb{E}[B^2]}{\mathbb{E}[B]}}.$$

In Chapter 5 we use this scaling approach for a polling system modeling traffic light controlled networks.

1.2.2 Operations research

Operations research (O.R.) is a discipline that deals with the application of advanced analytical methods to help make better decisions [68]. The decisions considered in this dissertation are *scheduling* decisions, i.e. they aim to achieve the *optimal* allocation of scarce resources to activities

over time [52]. In the context of road traffic the resource to be allocated is always the road capacity, whilst the method of allocation can mean granting access to a part of the network (e.g. green signals on traffic lights) or allowing for more extensive use of the network (e.g. increasing the speed limit, which increases the safety gap between vehicles and thus the space allocated to each vehicle). As the decisions need to be made in a sequential manner, we apply *policies*, which determine the allocation that is implemented. An important body of the scheduling literature is devoted to seeking policies that optimise the *performance* of a system. This can be expressed in terms of performance measures such as throughput or average sojourn times, etc. Another important notion is *fairness*, which relates to some level of "social justice" between the different actors in traffic as traffic controllers should aim to grant fair treatment to all of them. However, fairness is a subjective notion with a lot of research devoted to develop quantifying measures [11, 15, 45].

Policies based on convex optimisation

A common technique to allocate resources is to solve an optimisation problem on a utility function. The method has its roots in microeconomics, relating to the definition of fairness therein [57]. It relies on a social welfare function, which associates with each possible allocation the aggregate utility in the system. An allocation is called fair when it maximises the social welfare, i.e. the utility function. This can be formalised by introducing $x = (x_i)_{i\in\mathcal{I}}$, the allocation vector, where the index i runs over the set of actors $\mathcal{I}$. In this work $\mathcal{I}$ will be the set of queues of vehicles in the system in most cases. The allocation x is feasible if it belongs to a set S, which is often determined by the capacity of the system and possibly shaped by further constraints. If we denote the utility of each actor by U_i, the optimal allocation will be given by solving the optimisation problem

$$\text{maximise} \quad \sum_{i\in\mathcal{I}} U_i(x_i) \quad \text{over} \quad x \in S.$$

When the functions $U_i(\cdot)$ are strictly concave and the set S is convex and compact the optimiastion problem has a unique solution.

An important class of utility functions were introduced in [102], which

can be described as

$$U_i(x_i) = U_i^{(\alpha)}(x_i) = \begin{cases} w_i \log x_i & \text{if } \alpha = 1, \\ w_i \frac{x_i^{1-\alpha}}{1-\alpha} & \text{if } \alpha \in (0,\infty)\backslash\{1\}, \end{cases}$$

with $w_i > 0$, a weight assigned to each actor. The corresponding allocation is commonly referred to as the *weighted α-fair allocation*. The fact that these functions are increasing and strictly concave invokes fairness, as increasing the resources that are allocated to an actor, which would receive a low amount originally, yields a larger improvement in the aggregate utility. This class of allocations contains some popular allocation paradigms, for example if $\alpha \to \infty$, the so-called Max-Min Fair allocation (see for example [14] for a definition) arises, which is often conceived to be the most fair, as it maximises the minimal amount of resources to be allocated to any actor. The other extreme ($\alpha \to 0$) maximises the throughput on the other hand [73], although it can be highly unfair to certain actors by completely ignoring them. The case which will be studied in depth in this dissertation is the $\alpha = 1$ case, commonly referred to as the *Proportional Fair* allocation [70].

Another important set of policies are the so-called *Max-Weight* policies [73]. These emerge when

$$U_i(x_i) = w_i(x)x_i,$$

i.e. when the weights are determined by the state of the system, and the resources are allocated to the actor with the maximal weight, hence the name. An important example is the *Backpressure* policy, where the weights are determined according to the "pressure" an actor puts on the system.

Routing policies

A different type of scheduling problems arise when we consider the task of route planning for vehicles. This can be modelled as assigning jobs to servers in a queueing network. We discriminate between *static* (state-independent) and *dynamic* (state-dependent) routing policies. In addition, there are deterministic routing policies, which always make the

same decision in the same situation, and probabilistic routing policies, which contain some kind of randomness. For example, when there are two or more parallel queues serving jobs with exponential interarrival and service times, the optimal policy among static policies in terms of minimising the expected waiting time is a Bernoulli policy, which is a probabilistic policy assigning a fixed probability to each queue. Although they require the controller to be able to observe the system, dynamic policies outperform static policies. In the same example the optimal dynamic policy is *join-the-shortest queue* (JSQ) policy, which is deterministic. If we account for the fact that the service time may depend on the chosen server, we have to include weights in the routing decision, therefore arriving at the weighted JSQ policy. This dynamic policy is also optimal for general arrival processes [153], however the information required by it might not be available, for example when part of the traffic is *unobservable* for the controller. We dedicate Chapter 2 of the thesis to discuss potential routing policies in that case.

1.3 Overview of the thesis

In this dissertation we contribute to the development of mathematical modelling of road traffic. We investigate different aspects of the subject, namely the control of urban and interurban road traffic and of individual drivers. These themes are present in the chapters of the book in the following order.

In Chapter 2 we discuss a queueing model for a smart phone application giving routing guidance based on information received from its own pool of users. We describe several policies that can be implemented by the controller comparing their performance via simulation and fluid analysis.

Chapter 3 serves as a prelude to Chapters 4 and 5. Here we provide an overview on control policies for urban traffic light networks and introduce a basic model, which forms the basis of the theoretic discussion in the following chapters.

In Chapters 4 and 5 we investigate novel policies, which govern urban traffic light systems. The control methods in focus are the backpressure and the proportional fair scheme respectively. We describe the queueing processes performed by the vehicles, which develop under these policies,

proving the presence of some favorable attributes. Furthermore we evaluate the performance of these methods via simulation comparing the results both to each other and to those in the literature.

In Chapter 6 we describe models of traffic flow on highways, proposing a novel model based on a stochastically interacting particle system. We investigate the role of the transition rates of the process and evaluate the model's accuracy by comparing it to measured data from a British motorway.

In Chapter 7 we discuss novel policies for ramp metering and speed limit control on highways and the road segments leading up to them. The presented methods draw their ideas from the backpressure and the proportional fair schemes introduced for urban traffic. We highlight the strengths and weaknesses of these policies and evaluate their performances via simulation.

Chapter 2

Routing policies for individual drivers

Outline.

In this chapter we consider a route-guidance system that aims at minimising the average sojourn time of its users by giving advice based on information collected from them. We model this as a queueing system consisting of two queues and a controller that assigns newly arriving vehicles to the queues. Thus we can imagine the two servers as two alternate route choices, the service times as the time it would take a vehicle at traveling at the speed limit to go through its route and the waiting times as the delay caused by the traffic. We assume that in the system only part of the vehicles (those following the route guidance) can be observed and controlled, whereas the other vehicles cannot be observed explicitly and choose a queue independently. Thus, the controller has to perform its routing task based on partial information about the system state. In this chapter we investigate how to deal in an effective way with the partial observability and controllability of these systems, and how the performance depends on the penetration grade of the application (that is, the fraction of traffic that can be observed).

In short our model can be summarised as follows. It consists of two queues with independent exponential service times, serving two types of vehicles. Arrivals occur according to a Poisson process; a fraction α of the vehicles (type X) is observable and controllable. At all times the number of X type cars in each queue and their individual positions are known. Upon its arrival a router decides which queue the next X vehicle should join. The other vehicles (type Y, fraction $1-\alpha$) are non-observable and non-controllable. They randomly join a queue according to some static routing probability.

We show that under large loads the performance (in terms of the average sojourn time) of a simple policy that relies on little system information is close

to the weighted join-the-shortest-queue policy (w-JSQ), which is optimal in a fully controllable and observable system, by analysing deterministic fluid models that approximate the stochastic evolution in an appropriate asymptotic regime. This is due to the fact that in the asymptotic fluid limit the processes corresponding to the different policies converge to the same limiting process, a result that we prove in detail. We support this result by simulation, which also reveals that for heavily loaded systems a low penetration level suffices.

The remainder of this chapter is organised as follows. First, in Section 2.1 we briefly discuss relevant results in the literature. Second, in Section 2.2 we give a detailed description of the system model at hand, including the notation, and introduce and discuss the routing policies considered. Next, in Section 2.3, these policies are evaluated and compared by simulation. The analytic approach based on fluid approximations, especially useful for highly loaded systems, is described in Section 2.4, where we lay out the main theoretic result of our study discussing its consequences. The formalisation of the main result and its detailed proof is given in Section 2.5. Finally, in Section 2.6, we summarise the results of our study, provide conclusions and suggest directions for further research.

2.1 Literature review

Route guidance systems have been in the interest of the research community since the early 90's. With radar and mobile technology emerging and the U.S. Global Positioning System Policy of 1996 granting civilians access to GPS satellites, the positioning of vehicles and communicating with drivers in a real-time manner became available. Patents covering route guidance systems followed [67, 131] and the number of navigation devices in use has been steadily increasing since. Another spike in usage happened due to the appearance of smart phones that could serve this purpose [157, 120] and the emergence of applications that were built on this possibility [51, 9, 152]. The relevant literature covers the dynamic algorithms for route guidance [111, 20], addresses the question of a sufficient penetration level [69], consider the effects of drivers using route guidance systems on the network [18, 115] and proposes solutions for efficient use of network capacity [7]. These studies mostly apply simulations; to the best of our knowledge no rigorous model has been developed to consider the performance of the proposed methods.

The road traffic literature offers another point of view by considering the use of probe vehicles (which are the observable/controllable part of the traffic in our case) for traffic control extensively. Field experiments have been conducted, see e.g. [61]. The applications vary from estimations on queue lengths at traffic signals [32], to freeway travel-time prediction [28, 105] and detecting incidents [130]. The number of probe vehicles and the penetration level needed for these estimations to work well are also widely considered [25, 89, 135, 144]. Furthermore, there are systems in place that collect and process data like Vtrack [142] or CroTIS [127], and route-guidance applications based on such data, like Waze [152].

The aim of our work is to provide insight into the typical problems of route guidance whilst considering information received only from the users that subscribed to the application and thus serve as probe vehicles. We do this by intensively investigating and comparing the performance of various dynamic *routing* strategies for a simplified queueing system. Through extensive simulations we show that heuristic routing policies work quite well even for low penetration levels. Their performance (in terms of the average sojourn time) hardly improves further if the penetration level is increased beyond, roughly 25%. The performance of

these simple policies that rely on little system information appears to be close to the weighted join-the-shortest-queue (w-JSQ) policy, particularly so for heavily loaded systems, which is optimal in a fully observable and controllable system [64, 153]. This result is confirmed by the outcome of an approximative analytical study. In particular, assuming a highly loaded system, we show that the partial information available for the fluid models of the heuristic policies is enough to provide dynamics that give just as good performance as the fluid approximation of the optimal w-JSQ policy. We prove that in the limit the policies give identical trajectories.

To the best of our knowledge our approach to partial observability/controllability in queueing systems is novel, although similar models have been considered, see the works of Reiman [121] or Turner [145]. Kuri and Kumar [77] and Mitzenmacher [101], for example, discuss optimal routing for the case that delayed information on the queue lengths is available. In Guo et al. [55] the partial system information available concerns the service time distribution. Other papers, see e.g. [16], consider controllable (foreground) and uncontrollable (background) traffic, but there it is assumed that both traffic types are observable (i.e. the routing of foreground traffic is based on total state information).

2.2 Model description

2.2.1 The queueing system

We consider the model as a two-server queueing system, as depicted in Figure 2.1. The queues operate independently in a FIFO-manner with exponential service times of rates $\mu = (\mu_1, \mu_2)$. Vehicles arrive at the system according to a Poisson process with rate λ and are of one of the following two types.

Cars of the first type, called X vehicles, are controllable and observable. They make up for a fraction α of the total load. X vehicles thus arrive according to a Poisson process of rate $\alpha\lambda$. The parameter α is called the *penetration level*. X vehicles are *controllable*, because upon their arrival a router decides on which queue they should join. A *policy* is a set of rules according to which the routing decisions for X vehicles are taken. Saying that X vehicles are *observable*, we mean that at all times their

numbers in both queues (possibly including a vehicle that is in service) and their individual positions are known. The numbers of X vehicles are denoted by $X(t) = (X_1(t), X_2(t))$. The position of a vehicle in a queue is the number of vehicles that must be served before finishing the given vehicle, thus a vehicle currently in service is at position 1. We are especially interested in the positions of the last X vehicles in the queue, which will be denoted by $L^X(t) = (L_1^X(t), L_2^X(t))$. Should there be no X vehicles present in either queue, we set the corresponding component(s) of L^X to 0.

The second type of vehicles, Y cars, represent background traffic. They are non-controllable and non-observable. They arrive according to a Poisson process of rate $(1-\alpha)\lambda$ and join one of the two queues according to the static probabilities $p^Y = (p_1^Y, p_2^Y)$, where $p_2^Y = 1 - p_1^Y$. For the router the number of Y vehicles in each queue, which will be denoted by $Y(t) = (Y_1(t), Y_2(t))$, is unknown, thus leaving the total queue lengths, denoted by $Q(t) = (Q_1(t), Q_2(t))$, with $Q_i(t) = X_i(t) + Y_i(t)$, for $i = 1, 2$, unknown as well.

The loads on the queues depend on the routing policy and the penetration level. So as to conduct comparisons for α ranging from 0 to 1, we will measure the loads in each of the queues by the actual loads when $\alpha = 0$. For queue i this is given by $\rho_i^Y = p_i^Y \lambda / \mu_i$. When $\rho_1^Y = \rho_2^Y$ we use ρ^Y to denote either of them. For any choice of parameters for which the loads in both queues are below 1, it will thus be possible for the router to maintain the queue stable for all $\alpha \in [0, 1]$.

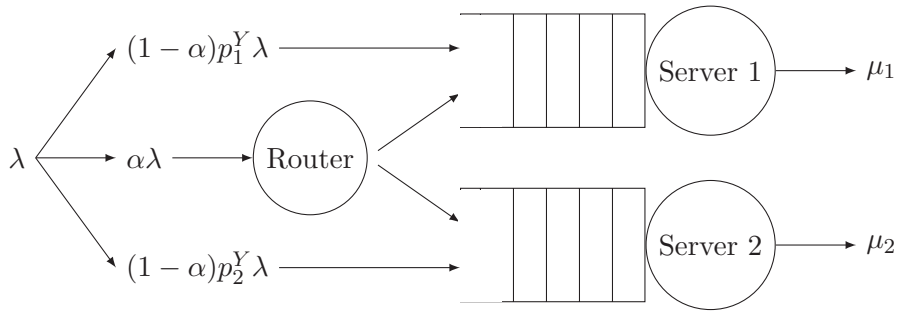

Figure 2.1. A schematic representation of the queueing system

2.2.2 Routing policies

In order to route incoming X vehicles to a queue, one can think of two types of policies: *dynamic policies* that use state information and *static policies* that do not. We consider the four dynamic routing policies introduced below.

- *Number of X vehicles* (denoted by #X): this policy sends an arriving X vehicle to the queue with the fewest X vehicles (possibly including an X vehicle currently in service).
- *Last X position* (abbreviated to LXP) compares the queues by the position of the last X vehicle and sends an arriving X vehicle to the queue with the lowest last X position.
- *Weighted last X position* (w-LXP) multiplies the last X position by the mean service time $1/\mu_i$ (giving an estimation of the sojourn time for queue i) and sends an arriving X vehicle to the lowest weighted last X position.
- *Estimated weighted last X position* (ew-LXP) does the same as w-LXP except that it does not assume μ to be known. The mean service time is estimated by dividing the sojourn times of the X vehicles lastly departed from queue i by their arrival position. The average is over the last w number of departed X users, with w a parameter of the policy.

For most of these policies no system information (e.g. the values of α, λ, p^Y and μ)) is assumed to be known by the router. Only in w-LXP μ is used. The first two policies are designed for symmetric queues (equal service rates), the last two for the asymmetric case.

In Section 2.3 we compare the performance of the above mentioned policies with respect to the expected sojourn time of the vehicles (the total time of a vehicle in the system). Let S_i denote the expected sojourn time in queue i and S_i^X the expected sojourn time in queue i over X vehicles only. We use the same symbols for the average sojourn times in the simulations. The averages over both queues are denoted by S and S^X for all, and for X vehicles only, respectively. In the policy comparison we also consider the following two reference policies:

- *Weighted join-the-shortest-queue* (w-JSQ): a dynamic policy that sends an arriving X vehicle to the queue with the fewest vehicles (X and Y types together). In case of unequal service times, the number of vehicles is weighted by the average service time, as for w-LXP. This reference policy assumes full state information (that is, also Y vehicles are observable) and knowledge of μ, but no other system information. Since w-JSQ uses more information than our policies, it will likely outperform our policies. Note that it may not be optimal for our partially controllable system, although it is optimal for fully controllable systems [64, 153].

- A *static policy* that applies probabilistic routing. It sends arriving X vehicles to queue i with a fixed probability p_i^X. The probabilities p_i^X are set such that the difference between S_1^X and S_2^X is minimised[1]. Note that the policy has the same goal as the policies discussed before: it aims at equalising the sojourn times in both queues. It can be verified that this policy is not the same as the static policy that minimises S. Although the static policy uses no state information, it may not always be outperformed by the dynamic policies, because the static policy uses system parameters $(\alpha, \lambda, \mu, p^Y)$ that are not (all) available to the dynamic policies.

Although all six policies introduced above assume partial controllability (i.e., only X cars can be routed), they differ in terms of the system and state information that is used. In Section 2.3 we compare their respective performance.

See Table 2.1 for a comparison. In some specific cases some of the policies are equivalent, that is, they make the same routing decisions. Obviously, LXP and w-LXP are the same if $\mu_1 = \mu_2$. Also, for $\alpha = 0$, all policies are equal, since there are no vehicles that can be controlled. For $\alpha = 1$, #X and LXP are the same, since the number of vehicles and the last position are equal if there are only X vehicles. In addition, in this case (for $\alpha = 1$), w-LXP and w-JSQ are the same, since there is full observability for all dynamic policies. For the static policy there is

[1] Simple calculations yield that this is achieved for $p_1^X = (\mu_1 - \mu_2 + \lambda - 2p_1^Y(1 - \alpha)\lambda)/(2\alpha\lambda)$, where S_1^X and S_2^X become equal, or $p_i^X \in \{0, 1\}$ if it is impossible to reach the desired equality.

Policy	#X	LXP	w-LXP	ew-LXP	w-JSQ	static
Purpose	*symmetric*		*asymmetric*		*reference*	
Partial						
controllability	yes	yes	yes	yes	yes	yes
observability	yes	yes	yes	yes	**no**	-
Uses state information						
number of X jobs	yes	yes	yes	yes	yes	**no**
position of X jobs	no	**yes**	**yes**	**yes**	no	no
state history	no	no	no	**yes**	no	no
Uses system information						
μ	no	no	**yes**	no	**yes**	**yes**
λ, α, p^Y	no	no	no	no	no	**yes**

Table 2.1. A comparison of the information used by the routing policies under consideration

no difference between S_i and S_i^X (for fixed i). Note that this is not the case for the dynamic policies (provided that $\alpha < 1$), because for these policies the routing of an X vehicle to a specific queue and the length of the queue are not independent.

A remark on the static policy

For the static policy (with routing probabilities p_i^X for the X vehicles), the system of Figure 2.1 consists of two independent $M/M/1$ queues with arrival rate $\lambda_i = \lambda(\alpha p_i^X + (1-\alpha)p_i^Y)$ for queue i. As a consequence, the expected sojourn time for queue i is given by $S_i = (\mu_i - \lambda_i)^{-1}$. Since the arrivals of X vehicles at queue i also follow a Poisson process, we have $S_i^X = S_i$, due to the PASTA property.

In order for each X vehicle to be sent to the queue that minimises its expected sojourn time, p_i^X is chosen such that $S_1^X = S_2^X$ (or $p_i^X \in \{0, 1\}$ if it is impossible to reach the desired equality). This gives

$$p_1^X = (\mu_1 - \mu_2 + \lambda - 2p_1^Y(1-\alpha)\lambda)/(2\alpha\lambda). \tag{2.1}$$

2.3 Simulation results

In this section we investigate whether the dynamic policies with partial observability and controllability described in Section 2.2.2 perform better than the static policy and how close their performance (in terms of average sojourn times) is to that of w-JSQ. Furthermore, we want to assess the added value of using extra state and/or system information and how the performance depends on the penetration level α. To this end, we have performed many simulations in MATLAB [94] under various system settings with symmetric and asymmetric service times and background traffic, and with different loads and penetration levels α.

2.3.1 Policy comparison

Symmetric queues

We start by considering the case of equal service rates, $\mu_1 = \mu_2$, see Figure 2.3 for an example with $\mu_1 = \mu_2 = 1$. We compare the two simple dynamic policies based on the numbers of X vehicles (#X) and the positions of the last X vehicles (LXP) in the two queues, with the static policy and the full state information policy JSQ. Every data point in this plot, as well as in the following plots, is the result of a simulation experiment with 10 million vehicle arrival and departure events.

The performance of LXP is always slightly better than that of #X. This can be explained by the fact that LXP also takes into account the Y vehicles that have arrived before the last X vehicle, while #X only measures the number of X vehicles. The relative difference between the two policies is largest for moderate and somewhat high loads. Under very high loads ($\rho^Y \uparrow 1$) the dynamic policies show comparable performance. In comparison the performance of the static policy is clearly inferior. However, for low α it is better for the overall average sojourn time to apply the static policy. Note that the dynamic policies are always better for the X vehicles, and for all vehicles when the control level α is relatively large.

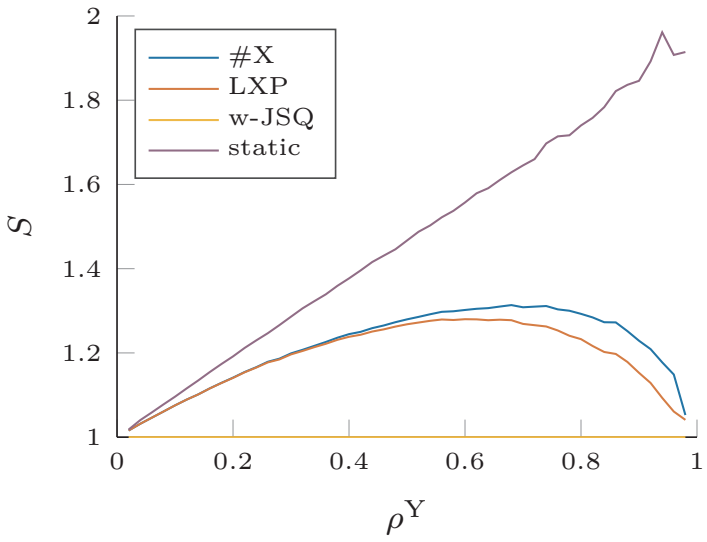

Figure 2.2. Relative average sojourn time over all X vehicles in both queues as a function of ρ^Y for several policies, $\alpha = 0.2$, $\mu = (1, 1)$ and $p^Y = (1/2, 1/2)$. The graph displays the relative sojourn times with respect to those of w-JSQ.

Asymmetric queues

We now allow for unequal service rates for the two queues describing the results of a similar set of experiments as in Section 2.3.1. In this case it is natural to use a weighted dynamic policy that takes the service rates into account. We can observe that (non-weighted) LXP also gives better results than #X for asymmetric queues (see Figure 2.4). Therefore we focus on LXP from now on, and consider its weighted versions w-LXP and ew-LXP. In the former the weights are given, in the latter they are estimated using data of the w lastly departed vehicles for each queue separately. We refer to w as the *window size* of ew-LXP and show results for values $w = 1$ and $w = 10$ in our graphs (denoted by ew-LXP(1) and ew-LXP(10) in the legends).

From the simulations in Figure 2.4 we see that, as expected, the weighted policies (w-LXP and ew-LXP) outperform the unweighted LXP. This turns out to be true across all values of α as can be observed in Figure 2.5. For practically relevant values of α (that is, relatively small

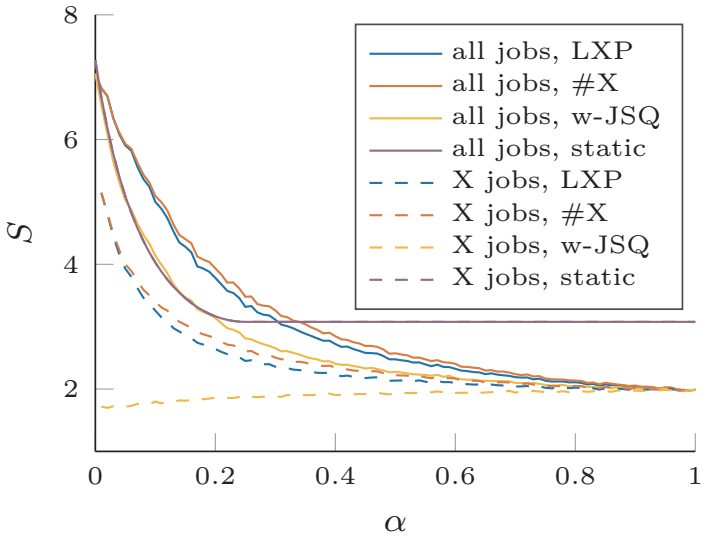

Figure 2.3. Average sojourn time over both queues as a function of α for several policies, $\rho_1^Y = 0.45$, $\rho_2^Y = 0.9$, $\mu = (1, 1)$ and $p^Y = (1/3, 2/3)$. We display both the average over all vehicles (solid), and over X vehicles only (dashed).

values) this even holds for a small window size of $w = 1$.

In practice the mean service times may vary over time. In this case the policy ew-LXP is particularly relevant. The recommended window size depends on the desired precision and on how fast μ is changing; the memory of the policy should refresh itself on a smaller time scale than the one on which μ changes. For a static μ, a window size $w = 10$ gives relatively accurate results as can be seen in Figure 2.4. In general, in order to get the same precision, w needs to be larger for lower loads (see Figure 2.4), because the queue lengths are smaller and therefore the average is taken over a smaller number of vehicles. In addition, we observe in Figure 2.5 that w needs to be larger for higher penetration levels (because the estimates of the average service time are using more overlapping data when there are fewer Y vehicles in the system).

The figures and observations discussed in this section refer to the case in which both queues are equally loaded by the Y vehicles (that is, $\rho_1^Y = \rho_2^Y$). Note that this does not imply that the queues are symmetric,

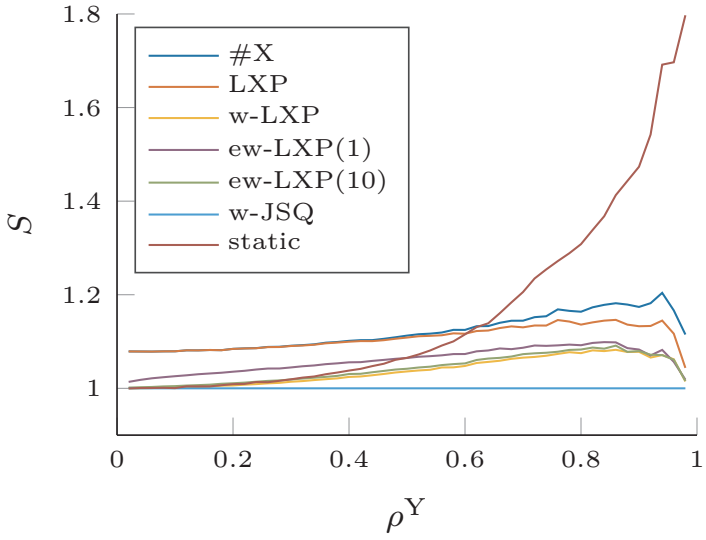

Figure 2.4. Relative average sojourn time over all vehicles and both queues as a function of ρ^Y for several policies, $\alpha = 0.2$, $\mu = (1, 2)$ and $p^Y = (1/3, 2/3)$. The sojourn times are relative to those of w-JSQ.

as the difference between the μ_i can be compensated by the p_i^Y. The observations turned out (in simulations not presented here) to be valid also for unequally loaded queues (e.g. if p^Y is symmetric, while μ is not). We noticed, however, that then the static policy outperforms our dynamic policies for low values of α (as we also observed in Section 2.3.1). So in that case the dynamic policies are only recommended if it is not possible to implement the static policy (because of a lack of system information). For the equally loaded scenario we have seen in Figure 2.4 that the weighted policies perform much better than the static policy and that the w-LXP policy is close to the optimal w-JSQ when ρ^Y approaches 1. This fact is explained in detail in Section 2.4.2.

2.3.2 Penetration level

One of the questions we raised was to determine the smallest penetration level α needed for effective control. To shed light on this question, we now study the performance of the various policies as a function of α. From

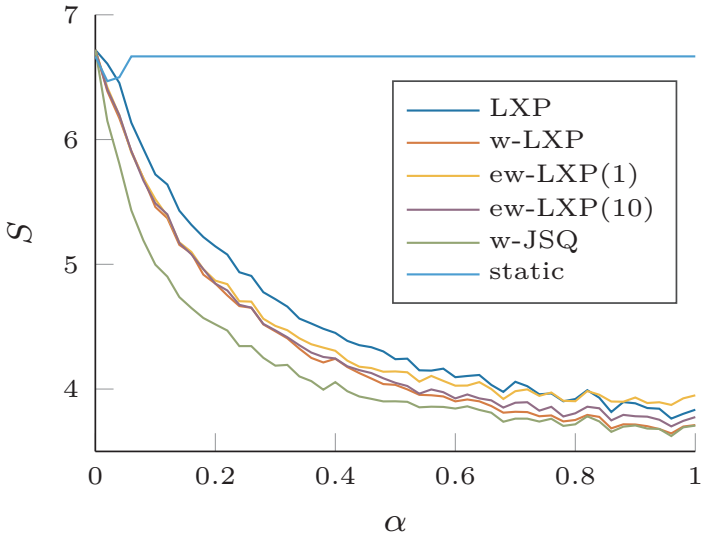

Figure 2.5. Average sojourn time over all vehicles and both queues as a function of α for several policies, $\rho_1^Y = \rho_2^Y = 0.9$, $\mu = (1, 2)$ and $p^Y = (1/3, 2/3)$

the figures presented in the preceding sections (Figures 2.3 and 2.5) we can conclude that the sensitivity of the performance to the penetration level is highest for small values of α (i.e. changes in α cause a larger change in the performance when α is small than when it is large), which can be understood as a "law of diminishing marginal returns". In most of the scenarios, the average sojourn time for w-LXP and ew-LXP is already rather close to its minimum if α is approximately 25%. The required penetration level depends on the system load; for lower loads a higher α is required and if the load caused by the Y vehicles is symmetric, a lower α suffices than if it is asymmetric (because more X vehicles are needed to compensate the 'wrong' choices of the Y vehicles).

First, let us inspect the average sojourn times S_1, S_2 for all vehicles (X and Y together). From Figure 2.6, which is representative for a large collection of graphs for a variety of scenarios, we see that for the dynamic policies the average sojourn times in both queues decrease when α increases. This means that increasing the level of control decreases the average time a vehicle spends in the system, independently of the

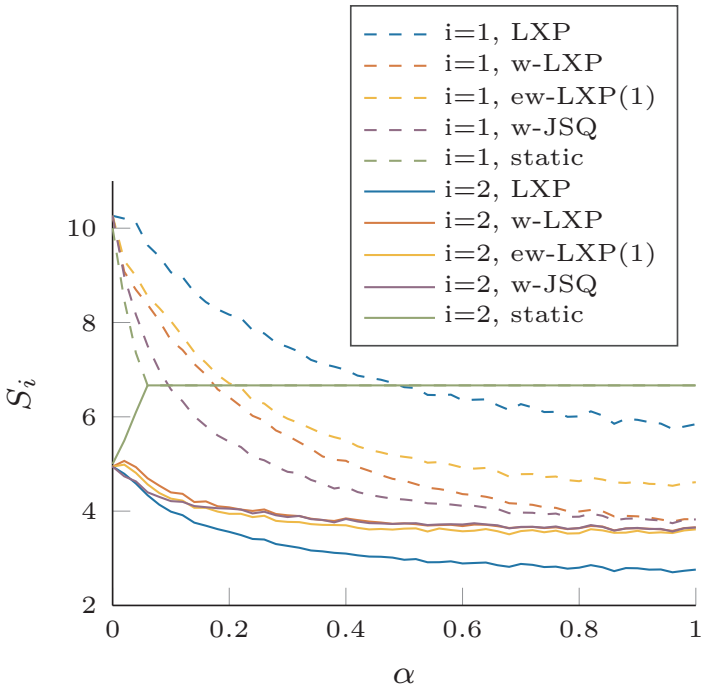

Figure 2.6. Average (over all vehicles) sojourn time as a function of α for both queues separately, $\rho_1^Y = \rho_2^Y = 0.9$, $\mu = (1, 2)$, and $p^Y = (1/3, 2/3)$

queue a vehicle is sent to. In contrast, for the static policy the sojourn time in one queue increases when more vehicles are being controlled (while the sojourn time in the other decreases as a function of α). As one can notice from Figure 2.5, this may result in non-monotonous average sojourn times over both queues. However, for the dynamic policies it is always beneficial to the overall average sojourn time to increase the amount of control.

If we consider X vehicles only, first we should note that the performance for X vehicles is always better than for an average (X or Y) vehicle. This is essential if vehicles can choose themselves whether to be in the X class (as in the road traffic example). For the static policy, the behaviour as a function of α is the same as before (when looking at both classes together). However, for w-JSQ the quantities S_1^X and S_2^X are now both *increasing*

in α. For this policy, increasing the percentage of controlled vehicles improves the overall performance, but deteriorates the performance for the controlled vehicles, because the advantage of being directed to the "right" queue becomes smaller if more vehicles are being directed to it. This reasoning does not hold for the other dynamic policies, because for those policies increasing α does not only mean more control (which is disadvantageous for the controlled vehicles), but also more information (which is advantageous). Consequently, the dynamic policies may show non-monotonous behaviour for one of the queues.

2.4 Fluid approximation

The policies that we have proposed and studied through simulation do not allow exact (analytical) performance analysis. In this section we approximate the behaviour of some of the policies under high loads using deterministic fluid models. For these deterministic fluid models we show that the partial information exploited by w-LXP is sufficient to give dynamics that yield just as good a performance as the fluid approximation of the "optimal" w-JSQ policy. As observed in Section 2.3.1 this result confirms the convergence of w-LXP to w-JSQ for loads approaching 1. A fluid approximation can often be shown to be the limiting process (*fluid limit*) of the original stochastic system under an appropriate scaling [126]. In this section we focus on the analysis of the fluid approximations themselves and also we provide a numerical validation study, while Section 2.5 contains a detailed technical proof of the result.

At a high level of abstraction, the fluid approximation may be seen to mimic the dynamics of the stochastic process at *large* system states; i.e., when the queueing processes move far away from the origin. Consequently, when the system load is low the applicability is limited to *transient* performance analysis, because large system states are rarely visited and therefore the fluid approximation does not reflect the typical behaviour of the process. When the system load is high, however, the stochastic process *does* typically move far away from the origin and the analysis of the fluid approximation provides valuable insight into the *stationary* behaviour of the stochastic process.

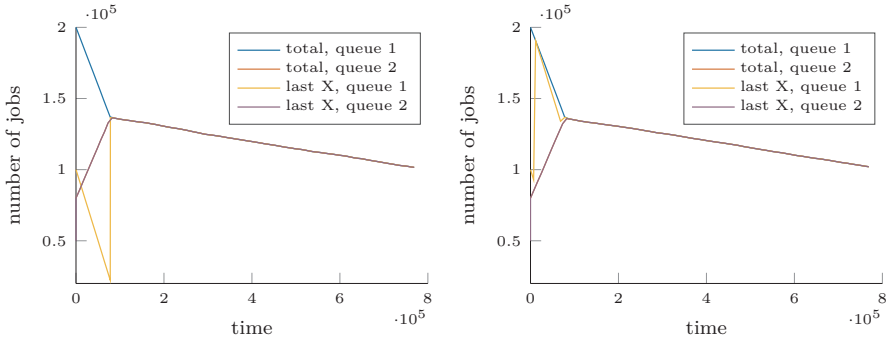

Figure 2.7. Simulated evolution starting from high initial queue sizes for JSQ (left) and LXP (right). In both figures the (red) curve for the number of (X plus Y) vehicles in queue 2 and the (purple) curve for the position of the last X vehicle in this queue coincide. After an initial period all four curves coincide.

2.4.1 High-load conditions and fluid limits

Since Little's Law applies to the queue lengths and the sojourn times irrespective of the policy, it is sufficient to concentrate on queue length dynamics to obtain the expectation of the sojourn time. Clearly, under all of the policies presented in Section 2.2.2 it is possible to describe the joint queue length processes as Markov processes if the state descriptor contains both queue lengths, *and* the positions of all vehicles, of both types, *plus* some information about the past (in the case of ew-LXP). Note that indeed, for some policies, there is a dependence of the dynamics on the positions of X vehicles, and the current value of a parameter, which is an estimate based on past system states. Therefore, the state space can be rather involved in general.

We focus our analysis on the w-JSQ and w-LXP policies. For these policies the random vector $M(t) = \left(Q_1(t), Q_2(t), L_1^X(t), L_2^X(t)\right)$ is Markovian. Let the Markov processes for the two policies be denoted by $M^{\text{w-JSQ}}$ and $M^{\text{w-LXP}}$ respectively.

The condition

$$\lambda/(\mu_1 + \mu_2) < 1, \tag{2.2}$$

is obviously necessary for these processes to be stable. In addition, we assume α to be sufficiently large, such that with appropriate routing, it

is possible to keep both queues from getting overloaded. Besides (2.2), we therefore require the load of Y vehicles on each queue to be below 1, i.e.,

$$p_i^Y(1-\alpha)\lambda < \mu_i, \tag{2.3}$$

for both $i = 1, 2$. The system is said to be in *heavy traffic* if the boundaries of condition (2.2) are approached, i.e. if $\lambda \uparrow \mu_1 + \mu_2$. If we impose that the system reaches heavy traffic, while satisfying (2.3) for both queues, we need

$$\alpha \geq 1 - \frac{1}{p_i^Y}\frac{\mu_i}{\mu_1 + \mu_2}, \tag{2.4}$$

for both $i = 1, 2$. In particular, if α is 0, it must be that $p_i^Y\lambda \uparrow \mu_i$ for both queues simultaneously, so that $\mu_2 p_1^Y = \mu_1 p_2^Y$.

Although we do not consider heavy-traffic limits, we assume that the system load is high, that is, λ is close to $\mu_1 + \mu_2$. Under such high-load conditions, the number of vehicles present in either queue will typically be very large and any *substantial* change in the queue lengths requires such a long time that the arrivals of both types virtually occur in a continuous fashion. This implies that in at least one of the two queues, there must be an X vehicle near the end of the queue. More precisely, for at least one of the two queues, the number of Y vehicles standing behind the last X vehicle in line is negligible compared to the total number of vehicles in the queue. This intuitive reasoning forms the rationale behind the fluid approximations proposed in Section 2.4.2. In Section 2.5 we prove that these (deterministic) fluid processes are the fluid limits of the original (stochastic) processes.

To facilitate the discussion, we end this section with the definition of the fluid limit, which formally describes the chosen scaling. If $\{M^{(c)}(t), t \geq 0\}_{c\in\mathbb{N}}$ is a sequence of Markov processes with $\|M^{(c)}(0)\|_1 = c$, then we define

$$\bar{M}^{(c)}(t) = \frac{M^{(c)}(ct)}{c}. \tag{2.5}$$

A *fluid limit* of this sequence is obtained by letting $c \to \infty$. Under suitable conditions, the limit often turns out to be a deterministic fluid process [126]. We show that this is also the case for the processes $M^{\text{w-JSQ}}$ and $M^{\text{w-LXP}}$.

2.4.2 Description and analysis of the fluid model

We now propose deterministic fluid processes $m^{\text{w-JSQ}}$ and $m^{\text{w-LXP}}$ to approximate the stochastic processes $M^{\text{w-JSQ}}$ and $M^{\text{w-LXP}}$ respectively, under large load conditions. Each of these processes consists of four components and, similar to our notation before, we generically write $m(t) = (q_1(t), q_2(t), l_1(t), l_2(t))$. As mentioned before, these processes will be formally justified in the next section.

For w-JSQ the evolution satisfies the following system of ordinary differential equations (ODEs). For conciseness we only report the derivatives for positive q_i, and l_i; at level zero the negative term in the derivative is to be removed, since there are no departures in this case.

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t} q_1^{\text{w-JSQ}} &= \begin{cases} \alpha\lambda + p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 < q_2^{\text{w-JSQ}}/\mu_2, \\ \alpha\lambda\mu_1/(\mu_1+\mu_2) + p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 = q_2^{\text{w-JSQ}}/\mu_2, \\ p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 > q_2^{\text{w-JSQ}}/\mu_2, \end{cases} \\
\frac{\mathrm{d}}{\mathrm{d}t} q_2^{\text{w-JSQ}} &= \begin{cases} p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 < q_2^{\text{w-JSQ}}/\mu_2, \\ \alpha\lambda\mu_2/(\mu_1+\mu_2) + p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 = q_2^{\text{w-JSQ}}/\mu_2, \\ \alpha\lambda + p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 > q_2^{\text{w-JSQ}}/\mu_2, \end{cases} \\
\frac{\mathrm{d}}{\mathrm{d}t} l_1^{\text{w-JSQ}} &= \begin{cases} \alpha\lambda + p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 < q_2^{\text{w-JSQ}}/\mu_2, \\ -\mu_1, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 > q_2^{\text{w-JSQ}}/\mu_2, \end{cases} \\
\frac{\mathrm{d}}{\mathrm{d}t} l_2^{\text{w-JSQ}} &= \begin{cases} -\mu_2, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 < q_2^{\text{w-JSQ}}/\mu_2, \\ \alpha\lambda + p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } q_1^{\text{w-JSQ}}/\mu_1 > q_2^{\text{w-JSQ}}/\mu_2. \end{cases}
\end{aligned}
\tag{2.6}
$$

These equations reflect that newly arriving X vehicles are routed to the weighted shortest queue. Note that the policy's rule in case the weighted queue lengths are equal is irrelevant, because any tie breaking rule forces the process to move along states with $q_1^{\text{w-JSQ}}/\mu_1 = q_2^{\text{w-JSQ}}/\mu_2$. It is worth emphasising that the last-X component of the weighted shortest queue also contains the effect of Y arrivals: both types of arrivals occur in a continuous fashion and are interleaved at the weighted shortest queue. Newly arriving Y vehicles thus push the last X position further back, than it would be in the absence of Y arrivals.

The above ODE lacks a description for the behaviour of the components $l_i^{\text{w-JSQ}}$ when the process hits the hyperplane $q_1^{\text{w-JSQ}}/\mu_1 = q_2^{\text{w-JSQ}}/\mu_2$. We

argued in Section 2.4.1 that for one of the two queues the last X position must be (almost) equal to the queue length. For w-JSQ this is the case for the weighted shortest queue. Indeed, new X vehicles are routed to the weighted shortest queue and the number of Y vehicles arriving in between two arriving X vehicles is negligible compared to the queue lengths. As a consequence, whenever the hyperplane $q_1^{\text{w-JSQ}}/\mu_1 = q_2^{\text{w-JSQ}}/\mu_2$ is hit, the last X position of the largest weighted queue also moves to the level of the corresponding queue length. Therefore we have to supplement the above ODEs with the following jumps:

$$l_i^{\text{w-JSQ}}(t+) = q_i^{\text{w-JSQ}}(t), \text{ for both } i, \text{ if } q_1^{\text{w-JSQ}}/\mu_1(t) = q_2^{\text{w-JSQ}}/\mu_2(t). \tag{2.7}$$

Here, by $t+$ we mean immediately after time t.

We now turn our attention to an approximating deterministic fluid system for w-LXP. Much of the discussion above holds in this case as well. The difference is due to the fact that the router now chooses the queue with the lowest weighted last X position, and only switches to the other queue when the weighted last X positions are equal. Thus the ODEs characterising the evolution of $m^{\text{w-LXP}}$ are — similarly to $m^{\text{w-JSQ}}$ — as follows:

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t} q_1^{\text{w-LXP}} &= \begin{cases} \alpha\lambda + p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } l_1^{\text{w-LXP}}/\mu_1 < l_2^{\text{w-LXP}}/\mu_2, \\ \alpha\lambda\mu_1/(\mu_1+\mu_2) + p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } l_1^{\text{w-LXP}}/\mu_1 = l_2^{\text{w-LXP}}/\mu_2, \\ p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } l_1^{\text{w-LXP}}/\mu_1 > l_2^{\text{w-LXP}}/\mu_2, \end{cases} \\
\frac{\mathrm{d}}{\mathrm{d}t} q_2^{\text{w-LXP}} &= \begin{cases} p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } l_1^{\text{w-LXP}}/\mu_1 < l_2^{\text{w-LXP}}/\mu_2, \\ \alpha\lambda\mu_2/(\mu_1+\mu_2) + p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } l_1^{\text{w-LXP}}/\mu_1 = l_2^{\text{w-LXP}}/\mu_2, \\ \alpha\lambda + p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } l_1^{\text{w-LXP}}/\mu_1 > l_2^{\text{w-LXP}}/\mu_2, \end{cases} \\
\frac{\mathrm{d}}{\mathrm{d}t} l_1^{\text{w-LXP}} &= \begin{cases} \alpha\lambda + p_1^Y(1-\alpha)\lambda - \mu_1, & \text{if } l_1^{\text{w-LXP}}/\mu_1 < l_2^{\text{w-LXP}}/\mu_2, \\ -\mu_1, & \text{if } l_1^{\text{w-LXP}}/\mu_1 > l_2^{\text{w-LXP}}/\mu_2, \end{cases} \\
\frac{\mathrm{d}}{\mathrm{d}t} l_2^{\text{w-LXP}} &= \begin{cases} -\mu_2, & \text{if } l_1^{\text{w-LXP}}/\mu_1 < l_2^{\text{w-LXP}}/\mu_2, \\ \alpha\lambda + p_2^Y(1-\alpha)\lambda - \mu_2, & \text{if } l_1^{\text{w-LXP}}/\mu_1 > l_2^{\text{w-LXP}}/\mu_2. \end{cases}
\end{aligned} \tag{2.8}$$

As before, if queue j has the lowest weighted last X position at time t, we must have $l_j^{\text{w-LXP}}(t) = q_j^{\text{w-LXP}}(t)$. The jumps on the switching plane

are now given by

$$l_i^{\text{w-LXP}}(t+) = q_i^{\text{w-LXP}}(t), \text{ for both } i, \text{ if } l_1^{\text{w-LXP}}/\mu_1(t) = l_2^{\text{w-LXP}}/\mu_2(t). \tag{2.9}$$

After a transient period the deterministic fluid approximations for both policies w-JSQ and w-LXP live on their switching planes, where respectively $q_1/\mu_1 = q_2/\mu_2$ and $l_1/\mu_1 = l_2/\mu_2$. On the switching planes, vehicles are sent to both queues, so that the last X positions are updated continuously and are equal to the queue sizes. It follows that (after a transient period) both fluid processes live on the same line ($q_1/\mu_1 = l_1/\mu_1 = q_2/\mu_2 = l_2/\mu_2$) and take the same decisions, both equalising the loads of the two queues.

For comparison, we also discuss the corresponding fluid approximation for the static policy. In this case the dynamics are very different; since there are no routing decisions the ODE system is much simpler:

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t} q_1^{\text{stat}} &= p_1^X \alpha\lambda + p_1^Y (1-\alpha)\lambda - \mu_1, \\
\frac{\mathrm{d}}{\mathrm{d}t} q_2^{\text{stat}} &= p_2^X \alpha\lambda + p_2^Y (1-\alpha)\lambda - \mu_2, \\
\frac{\mathrm{d}}{\mathrm{d}t} l_1^{\text{stat}} &= p_1^X \alpha\lambda + p_1^Y (1-\alpha)\lambda - \mu_1, \\
\frac{\mathrm{d}}{\mathrm{d}t} l_2^{\text{stat}} &= p_2^X \alpha\lambda + p_2^Y (1-\alpha)\lambda - \mu_2.
\end{aligned} \tag{2.10}$$

Here the static routing probabilities p_1^X and $p_2^X = 1 - p_1^X$ are given by (2.1). Since X vehicles are routed to both queues regardless of the state of the system, the last X positions clearly follow the queue lengths. Filling in the routing probabilities of (2.1), shows us that all derivatives are equal to $(\lambda - \mu_1 - \mu_2)/2$. As a consequence, under the static policy, the trajectories of the fluid process for different initial states are parallel to each other and basically all of the trajectories hit the axes away from the origin. As a consequence one of the queues empties before the other, thus leaving one of the servers idling, which indicates an inefficient use of capacity.

2.4.3 Numerical verification of the fluid approximations

We have conducted simulations to verify the suitability of the deterministic fluid approximations for large system loads. When the service rates are different, we observe the same behaviour as with equal μ's (after applying a correction for the different service rates). Thus we limit ourselves here to the symmetric case $\mu_1 = \mu_2$. We set the initial queue lengths and last X positions intentionally away from the switching curve, whilst choosing α such that it satisfies (2.4).

In our first set of experiments (see Figure 2.7) we illustrate the appropriateness of the linear deterministic approximations by plotting the trajectories over a very long time horizon and at very large system states. We only plot the simulation curves, as the deterministic approximations would be indistinguishable from them in the graphs. In these experiments we fixed $\alpha = 0.8$, $\mu = (1, 1)$, $p^Y = (1/2, 1/2)$ and $\rho^Y = \rho_1^Y = \rho_2^Y = 0.95$.

First we observe the plot of the trajectories of the queueing processes under JSQ. As long as the two queues have different sizes, all vehicles are routed to the shortest queue (queue 2 in this experiment); the last X position in queue 2 is therefore equal to the queue length in queue 2. For queue 1 we see that the last X position decreases faster than the queue length, since new Y vehicles are all placed *after* the last X vehicle. The position of the last X vehicle may hit zero and remain at zero until the two queues meet. At that point the last X position in queue 1 increases instantly to become equal to the queue length and from then on both queue lengths *and* the last X positions remain coupled forever.

Now we turn to the results under LXP. Again we start off at a very large state for all components, with queue 2 having the lowest queue length *and* the lowest last X-position (these are equal), i.e., the router is sending vehicles to queue 2. As soon as the last X position of queue 1 has dropped to the level of that of queue 2, it is increased to the size of queue 1. After a single X vehicle has been sent to queue 1, the next X vehicles are being sent to queue 2 again. This pattern repeats until the sizes of the two queues meet, from that point on all four components will have one single value. Note that, since the curve for the last X position in queue 1 can only bounce at a countable number of points, basically all X vehicles have been routed to queue 2, as was the case for JSQ. The point where the two queues meet is therefore exactly the same for both

policies.

Figure 2.7 considers extremely large system states, which will only rarely be visited. In Figure 2.8 we have plotted the trajectories for JSQ and LXP under a load as high as 0.995 and for system states that are more likely to be reached. Other parameters were set as $\alpha = 0.6$, $\mu = (1, 1)$, and $p^Y = (3/5, 2/5)$. We see that the trajectories meet quickly, after which the joint trajectories are well approximated by a linear trend.

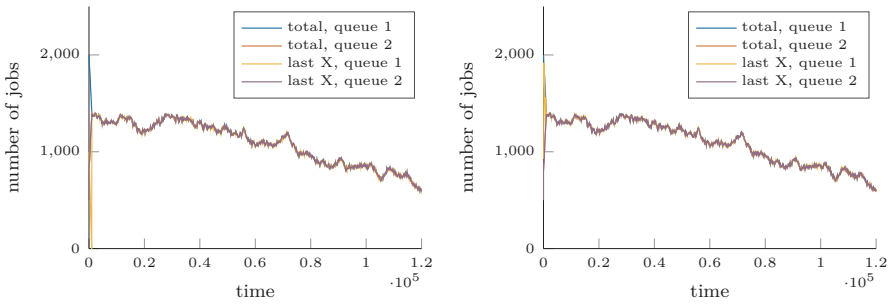

Figure 2.8. Simulated evolution under a high load for JSQ (left) and LXP (right). All curves coincide after a very brief initial period.

Our numerical experiments demonstrate that indeed the proposed fluid limit approximations closely follow the trajectories of the w-JSQ and w-LXP policies at large system states under high loads. The simulations also corroborate with our observations in Section 2.3, where we noticed that in heavy traffic w-LXP performs comparable to w-JSQ.

2.5 Formal proof of fluid convergence

In this section we provide a formal proof to the proposed fluid limits presented in Section 2.4.2. Firstly we formulate the result in the following theorem, where we denote the components of $\bar{M}^{(c)}(t)$ by $\bar{Q}_i^{(c)}(t)$ and $\bar{L}_i^{(c)}(t)$ respectively. Thus they can be obtained by taking the limit in (2.5) componentwise.

Theorem 2.1. *The sequence of stochastic processes* $\{\bar{Q}_i^{(c)}\}_{c\in\mathbb{N}}$ *is tight with respect to the topology of uniform convergence on compact time intervals in the space of continuous functions for* $i = 1, 2$*, while the*

sequence of stochastic processes $\{\bar{L}_i^{(c)}\}_{c\in\mathbb{N}}$ *is tight with respect to Skorokhod topology in the space of càdlàg functions for* $i = 1, 2$. *Moreover, any weakly convergent subsequence of* $\{\bar{M}^{(c)}\}_{c\in\mathbb{N}}$ *converges to a process satisfying fluid equations* (2.6) *and* (2.7) *under the w-JSQ policy, and* (2.8) *and* (2.9) *under the w-LXP policy respectively.*

Since the space of continuous functions (often refered to as space C) is a subspace of the space of càdlàg functions (which contains the functions that are everywhere right-continuous and have left limits, and is often referred to as space D), with the Skorokhod topology coinciding with the uniform topology over space C, the statement can be made for M in space D without separating its components. For details on space D we refer the reader to Chapter 3 of Billingsley's book [17], we closely follow its definitions and results for the rest of this section.

The proof will follow the subsequent steps. First, we take care of the processes $\bar{Q}_i^{(c)}$ by splitting them into a sum of terms for which we can show tightness by definition and then apply the triangle inequality to conclude that $\bar{Q}_i^{(c)}$ are tight. Secondly, we examine $\bar{L}_i^{(c)}$, for which we prove tightness by enlisting the possible events that could violate the tightness conditions and show that their probabilities vanish. Finally, we show that the weakly convergent subsequences, which exist due to tightness, indeed converge to the deterministic processes in (2.6), (2.7), (2.8) and (2.9), by examining their behaviour away from the corresponding switching planes first, then determining it on the switching planes.

In order to make our arguments formal, we introduce some variables describing the evolution of $M^{\text{w-JSQ}}$ and $M^{\text{w-LXP}}$. We take A_i^X and A_i^Y to be the arrival processes of X and Y jobs respectively to queue i, while we denote the departure process by D_i. We also introduce Δ_i, which gives the number of Y arrivals to queue i since the last X arrival to the same queue. We have

$$Q_i(t) = Q_i(0) + A_i^X(t) + A_i^Y(t) - D_i(t), \tag{2.11}$$

and

$$L_i(t) = (Q_i(t) - \Delta_i(t))^+ \tag{2.12}$$

giving the evolution of the components of $M(t)$ respectively, where $(x)^+$ equals x if $x \geq 0$, otherwise it is 0. We let n_i^X enumerate the X arrivals

to queue i and denote the random times when they happen by $T^X_{n,i}$. If we then denote by $N^X_i(t)$ the last X arrival to happen until time t, i.e.

$$N^X_i(t) = n, \quad \text{if} \quad T^X_{n,i} \leq t < T^X_{n+1,i},$$

then the last X arrival happened to queue i at $T^X_{N,i}(t)$. Thus, we can give the evolution of Δ_i as

$$\Delta_i(t) = A^Y_i(t) - A^Y_i(T^X_{N,i}(t)).$$

We also introduce $I_i(t)$ to be the indicator that X jobs are sent to queue i. With that notation we can write

$$A^X_i(t) = \mathcal{A}^X\left(\int_0^t I_i(s)\mathrm{d}s\right),$$

where $\mathcal{A}^X$ is a Poisson process with intensity $\alpha\lambda$.

2.5.1 Tightness

Tightness of the queue sizes

From Section 9 of [126] we see that in order to prove the tightness of a sequence $\{S^{(c)}\}_c$ in space C, we must prove for each $\epsilon > 0$ and $t > 0$ that

$$\lim_{\delta\to 0}\mathbb{P}\left(\sup_{\substack{u,v:u,v<t,\\ |u-v|<\delta}} \|S^{(c)}(u) - S^{(c)}(v)\|_1 \geq \epsilon\right) = 0. \tag{2.13}$$

We are going to show that (2.13) holds for $\bar{Q}^{(c)}_i(t)$ for each i. From (2.11) we have

$$\begin{aligned}
&\bar{Q}^{(c)}_i(t) = \bar{Q}^{(c)}_i(0) + \frac{1}{c}\left[A^X_i(ct) + A^Y_i(ct) - D_i(ct)\right] \\
=&\bar{Q}^{(c)}_i(0) + \frac{\alpha\lambda}{c}\int_0^{ct} I_i(s)\mathrm{d}s + \bar{X}^{(c)}_i(t) \\
&+(1-\alpha)p^Y_i\lambda t + \bar{Y}^{(c)}_i(t) - \mu_i t - \bar{Z}^{(c)}_i(t),
\end{aligned} \tag{2.14}$$

where we define

$$\bar{X}_i^{(c)}(t) = \frac{1}{c}\left(A_i^X(ct) - \alpha\lambda\int_0^{ct} I_i(s)\mathrm{d}s\right),$$

$$\bar{Y}_i^{(c)}(t) = \frac{1}{c}\left(A_i^Y(ct) - (1-\alpha)p_i^Y\lambda ct\right)$$

and

$$\bar{Z}_i^{(c)}(t) = \frac{1}{c}\left(D_i(ct) - \mu_i ct\right)$$

respectively.

By the triangle inequality it suffices to prove (2.13) holds for each term of the sum in (2.14) to prove that $\bar{Q}_i^{(c)}$ is tight. For $\bar{Q}_i^{(c)}(0) + (1-\alpha)p_i^Y\lambda t - \mu_i t$ (2.13) holds trivially. The processes $\bar{X}_i^{(c)}$, $\bar{Y}_i^{(c)}$ and $\bar{Z}_i^{(c)}$ are martingales, therefore we can apply Doob's L_2 inequality to them to obtain that

$$\begin{aligned}
&\mathbb{P}\left(\sup_{\substack{u,v:u,v<t,\\ |u-v|<\delta}} |\bar{X}_i^{(c)}(u) - \bar{X}_i^{(c)}(v)| \geq \epsilon\right)\\
\leq&\mathbb{P}\left(\sup_{u:u<t} |\bar{X}_i^{(c)}(u)| \geq \frac{\epsilon}{2}\right)\\
\leq&\frac{4}{\epsilon^2}\mathbb{E}\left[\left(\bar{X}_i^{(c)}(t)\right)^2\right] \leq \frac{4t\alpha\lambda}{c\epsilon^2} \to 0 \text{ as } c\to\infty,
\end{aligned} \tag{2.15}$$

and

$$\begin{aligned}
&\mathbb{P}\left(\sup_{\substack{u,v:u,v<t,\\ |u-v|<\delta}} |\bar{Y}_i^{(c)}(u) - \bar{Y}_i^{(c)}(v)| \geq \epsilon\right)\\
\leq&\frac{4}{\epsilon^2}\mathbb{E}\left[\left(\bar{Y}_i^{(c)}(t)\right)^2\right] = \frac{4t(1-\alpha)p_i^Y\lambda}{c\epsilon^2} \to 0 \text{ as } c\to\infty,
\end{aligned} \tag{2.16}$$

and finally

$$\begin{aligned}&\mathbb{P}\left(\sup_{\substack{u,v:u,v<t,\\|u-v|<\delta}}|\bar{Z}_i^{(c)}(u)-\bar{Z}_i^{(c)}(v)|\geq\epsilon\right)\\&\leq\frac{4}{\epsilon^2}\mathbb{E}\left[\left(\bar{Z}_i^{(c)}(t)\right)^2\right]=\frac{4t\mu_i}{c\epsilon^2}\to 0\ \text{ as } c\to\infty.\end{aligned}\tag{2.17}$$

For the integral term we can use the fact that $I_i(t)$ is an indicator, and thus bounded by 1 for all t. Therefore we have the Lipschitz-condition

$$\left|\frac{1}{c}\int_0^{cu}\alpha\lambda I_i(s)\mathrm{d}s-\frac{1}{c}\int_0^{cv}\alpha\lambda I_i(s)\mathrm{d}s\right|<\alpha\lambda|u-v|+2\alpha\lambda c^{-1}.$$

Thus with $\delta<\frac{\epsilon}{\alpha\lambda}$ (2.13) is satisfied. Therefore we have proven the tightness of the queue size processes.

Tightness of the last X positions

While $\bar{Q}_i^{(c)}(t)$ are tight in the space C, the same cannot be said about $\bar{L}_i^{(c)}(t)$. They are tight in the space D however. From [17] we know that in order to prove tightness in space D for a sequence $S^{(c)}$, it suffices to check the following conditions.

$$\begin{aligned}&(i)\quad\forall c,\eta>0\ \ \exists a\ \text{ such that}\\&\qquad\mathbb{P}\left(\sup_{0\leq u\leq t}|S^{(c)}(u)|>a\right)\leq\eta,\\&(ii)\quad\forall\eta,\epsilon>0\ \ \exists c_0,\delta>0\ \text{ such that if } c\geq c_0,\\&\qquad\mathbb{P}\left(w^*_{S^{(c)}}(\delta,t)\geq\epsilon\right)\leq\eta.\end{aligned}\tag{2.18}$$

Here, we define the modulus of continuity w^*:

$$w^*_S(\delta,t)=\inf_{\{t_j\}}\max_{0<j\leq r}w_S[t_{j-1},t_j),\tag{2.19}$$

where

$$w_S[t_{j-1},t_j)=\sup_{u,v\in[t_{j-1},t_j)}|S(u)-S(v)|,\tag{2.20}$$

and $0 = t_0 < t_1 < \ldots < t_r = t$, r is an integer, that can be chosen arbitrarily, with $t_j - t_{j-1} > \delta \; \forall j$, and $0 < \delta < t/2$, and the infimum in (2.19) runs over all possible finite sets of $\{t_j\}$.

When proving tightness, we first note that $\bar{L}_i^{(c)}(t) \leq \bar{Q}_i^{(c)}(t)$ at any time t. Furthermore the queue sizes are clearly dominated by a Poisson process with rate λ. A Poisson process with finite rate trivially satisfies condition (i) from (2.18), thus the same holds for $\bar{L}_i^{(c)}$.

Now we direct our attention towards condition (ii). First we note that the event $\{w^*_{\bar{L}_i^{(c)}}(\delta, t) \geq \epsilon\}$ happens if ϵ-sized differences occur in $\bar{L}_i^{(c)}$ more frequently than δ. Combined with the definition (2.5), we see that this is equivalent to $c\epsilon$-sized differences occuring in L_i more frequently than $c\delta$. As the last X position decreases between two X arrivals and increases upon an X arrival this can happen in two ways. Namely, either enough departures happen without an X arrival in time $c\delta$, such that the last X position decreases by more than $c\epsilon$, or the queue size changes so much, that within a $c\delta$-length interval the difference of the lowest X position and the highest queue size grows above $c\epsilon$, while a jump also happens. Let us point out, that the definition of w^* allows for choosing the interval boundaries, therefore the big jumps which happen after a long time without an X arrival to queue i (typically when the scheduler switches to i after some time spent sending X jobs to the other queue) can be chosen to happen on these boundaries, thus not changing w^*. Thus we need to examine the probability of the two possible events mentioned above.

Firstly, let us approximate the probability that the last X position in queue i decreases by a fixed number κ over a time interval of length τ. This event is conditioned on the fact that no X arrivals happen, which could depend on the last X position itself, however we will take 1 as an upper bound on its probability. Furthermore the last X position can only decrease if there are departures from the corresponding queue, which happen only if the queue is saturated. We will bound the probability of this condition by 1 as well. In this case the departures follow a Poisson

process of rate μ_i, thus we can write, that

$$\mathbb{P}\left(L_i(t) - L_i(t+\tau) \geq \kappa\right) \leq \sum_{k=\kappa}^{\infty} \exp(-\mu_i\tau)\frac{(\mu_i\tau)^k}{k!} < \sum_{k=\kappa}^{\infty} \exp(-\mu_i\tau)\frac{(\mu_i\tau)^\kappa}{\kappa!}\left(\frac{\mu_i\tau}{\kappa+1}\right)^{k-\kappa} = \exp(-\mu_i\tau)\cdot\frac{(\mu_i\tau)^\kappa}{\kappa!}\cdot\frac{\kappa+1}{\kappa+1-\mu_i\tau}. \tag{2.21}$$

We note that for the second inequality to hold, we need $\mu_i\tau < \kappa+1$. This condition will hold with a properly chosen δ as discussed below. In order to bound the modulus of continuity we seek to divide the time between two X arrivals to queue i into intervals of at least $c\delta$ length. Thus, the maximal interval length on which $c\epsilon$ departures should happen with a small probability is $2c\delta$. Therefore if we take $\kappa = \lceil c\epsilon \rceil$ and $\tau = 2c\delta$ in (2.21), we get

$$\begin{aligned}&\exp(-2\mu_i c\delta)\cdot\frac{(2\mu_i c\delta)^{\lceil c\epsilon\rceil}}{(\lceil c\epsilon\rceil)!}\cdot\frac{\lceil c\epsilon\rceil+1}{\lceil c\epsilon\rceil+1-2\mu_i c\delta} = \\ &\exp(-2\mu_i c\delta)\sqrt{\lceil c\epsilon\rceil}\left(\frac{2\mathrm{e}\mu_i\delta}{\epsilon}\right)^{\lceil c\epsilon\rceil}\mathcal{O}(1) \underset{c\to\infty}{\to} 0,\end{aligned} \tag{2.22}$$

if $2\mathrm{e}\mu_i\delta < \epsilon$, which gives the condition by which δ can be chosen in (2.18).

Secondly we discuss the case when X arrivals happen often enough, such that the time interval cannot be divided into parts, that only have X arrivals on their boundaries. This means, that X arrivals happen more often than $c\delta$. We can assume to put the interval boundaries on X arrivals, thus we are interested in the probability that the difference between the minimum and the maximum of L_i exceeds a given κ between t and $t+\tau$, provided that $L_i(t) = Q_i(t)$, if there are X arrivals happening. We point out that the all the events that change Q_i (i.e. arrivals and departures) happen at a slower rate than jumps in a Poisson process with rate $\lambda + \mu_i$. In order to have the values of L_i differ by at least κ on an interval of length τ, if at the beginning it was the same as Q_i, at least κ arrivals and departures must happen in the right order, thus we

can form the following bound.

$$\begin{aligned}
&\mathbb{P}\left(\sup_{u,v\in[t,t+\tau)}\{|L_i(t)-L_i(t+\tau)|\}\geq\kappa\right)\\
&\leq\sum_{k=\kappa}^{\infty}\exp(-(\lambda+\mu_i)\tau)\frac{((\lambda+\mu_i)\tau)^k}{k!}<\\
&\sum_{k=\kappa}^{\infty}\exp(-(\lambda+\mu_i)\tau)\frac{((\lambda+\mu_i)\tau)^\kappa}{\kappa!}\left(\frac{(\lambda+\mu_i)\tau}{\kappa+1}\right)^{k-\kappa}=\\
&\exp(-(\lambda+\mu_i)\tau)\cdot\frac{((\lambda+\mu_i)\tau)^\kappa}{\kappa!}\cdot\frac{\kappa+1}{\kappa+1-(\lambda+\mu_i)\tau}.
\end{aligned}\tag{2.23}$$

Now, similarly to (2.22) we can establish that taking δ such that $2\mathrm{e}(\mu_i+\lambda)\delta<\epsilon$ gives us the condition we are looking for, furthermore it also covers the first case. Therefore we can divide the $[0,t]$ interval into subintervals that are between δ and 2δ in length, during them either the first or the second case applies and all bigger jumps that occur once the scheduler returns to queue i after a longer vacation can be covered by the interval boundaries. This set of $\{t_j\}$, even though it might not give the infimum that is w^* by the definition (2.19) clearly satisfies condition (ii) in (2.18), thus the modulus of continuity does so as well, proving that $\bar{L}_i^{(c)}(t)$ is tight for all i.

2.5.2 Limits of weakly convergent subsequences

Limits of weakly convergent subsequences away from the switching plane

We have shown tightness for all components of $\bar{M}^{(c)}(t)$, thus it follows componentwise that the sequence of versions has a weakly convergent subsequence. We will denote the limit of this by $\tilde{m}=(\tilde{q}_1,\tilde{q}_2,\tilde{l}_1,\tilde{l}_2)$. What we are left to prove is that $\tilde{m}$ is the same deterministic process as the one described in (2.6) and (2.8) respectively. Let us formulate this proof with the assumption that the queues are saturated at all times.

First, let us point out that due to equations (2.15), (2.16) and (2.17), $\bar{X}_i^{(c)}$, $\bar{Y}_i^{(c)}$ and $\bar{Z}_i^{(c)}$ all converge to 0 in distribution. For the latter two this convergence does not depend on the visited states of M. For $\bar{X}_i^{(c)}$

this dependence is in the indicator $I_i(t)$, however since the indicator function is continuous away from the switching plane ($\tilde{q}_1/\mu_1 = \tilde{q}_2/\mu_2$ in the JSQ-case, $\tilde{l}_1/\mu_1 = \tilde{l}_2/\mu_2$ in the LXP-case), the following convergence holds on all trajectories that do not cross it,

$$\lim_{c\to\infty} \frac{1}{c}\int_0^{ct} I_i(s)\mathrm{d}s = \int_0^t \lim_{c\to\infty} I_i(M^{(c)}(cs))\mathrm{d}s = \int_0^t I_i(\tilde{m}(s))\mathrm{d}s, \quad (2.24)$$

where we have used the Skorohod Representation Theorem [17] in the first equation.

To describe the evolution of the limit of $\bar{L}_i^{(c)}$, we consider (2.12), which we can rewrite as

$$\bar{L}_i^{(c)}(t) = \left(\bar{Q}_i^{(c)}(t) - \bar{\Delta}_i^{(c)}(t)\right)^+, \quad (2.25)$$

where

$$\begin{aligned}
\bar{\Delta}_i^{(c)}(t) =& \frac{1}{c}\left[A_i^Y(ct) - A_i^Y(T_{N,i}^X(ct))\right] \\
=& \frac{1}{c}\Big[\left(A_i^Y(ct) - (1-\alpha)p_i^Y\lambda ct\right) \\
&- \left(A_i^Y(T_{N,i}^X(ct)) - (1-\alpha)p_i^Y\lambda T_{N,i}^X(ct)\right) \\
&+ \left((1-\alpha)p_i^Y\lambda(ct - T_{N,i}^X(ct))\right)\Big] \\
=& \bar{Y}_i^{(c)}(t) - \bar{Y}_i^{(c)}(T_{N,i}^X(t)) + (1-\alpha)p_i^Y\lambda\cdot\frac{ct - T_{N,i}^X(ct)}{c}.
\end{aligned} \quad (2.26)$$

We have already shown that $\bar{Y}_i^{(c)}(t)$ converges to zero in distribution. Since $T_{N,i}^X(t)$ is clearly a stopping time, and $\bar{Y}_i^{(c)}$ is a martingale, we can apply the Optimal Stopping Theorem to state that the second term converges to 0 as well. Thus the behavior of the limit is determined by the last term. For that let us examine the probability that the last X arrival to any of the two queues happened earlier than a given time. For any fixed $\tau > 0$ we have

$$\mathbb{P}\left(\max_i T_{N,i}^X(ct) < c(t-\tau)\right) \le \exp\{-\alpha\lambda c(t-\tau)\} \to 0 \quad (2.27)$$

as $c \to \infty$. On the other hand if it holds that there exists $\tau > 0$ such that $I_i(s) = 1$ if $s \in [c(t-\tau), ct)$ and $I_{i-}(s) = 0$, where by $i-$ we denote 2 if $i = 1$ and 1 if $i = 2$, then the index in which $T^X_{N,i}(ct)$ takes its maximum is i, thus the last term in (2.26) converges to 0 in distribution. For $i-$ however, $T^X_{N,i-}(ct)$ is smaller than the last time when I_i was 0 with probability 1. In the limit such τ exists, if the process is away from the switching plane due to the continuity of I_i, thus we can write that

$$\lim_{c\to\infty} \bar{\Delta}^{(c)}_i(t) = (1-\alpha)p^Y_i \lambda \int_\tau^t (1 - I_i(\tilde{m}(s)))\mathrm{d}s, \tag{2.28}$$

for some $0 \leq \tau < t$, where we have used (2.24).

Now we are ready to examine the limits of the components of $\bar{M}^{(c)}(t)$. From (2.14), (2.15), (2.16), (2.17) and (2.24), we can write that

$$\lim_{c\to\infty} \bar{Q}^{(c)}_i(t) = \tilde{q}_i(0) + p^Y_i(1-\alpha)\lambda t - \mu_i t + \alpha\lambda \int_0^t I_i(\tilde{m}(s))\mathrm{d}s,$$

from where we can deduce the differentiability properties for $\tilde{q}_i$. Namely for a sufficiently small h, such that $\tilde{q}_i(s) > 0$ for all $s \in [t-h, t)$, we have

$$\frac{\tilde{q}_i(t) - \tilde{q}_i(t-h)}{h} = p^Y_i(1-\alpha)\lambda - \mu_i + \frac{1}{h}\int_{t-h}^t \alpha\lambda I_i(\tilde{m}(s))\mathrm{d}s, \tag{2.29}$$

which is the same differentiability condition as the ones described in (2.6) and (2.8) for q_i away from the switching plane.

Similarly for $\tilde{l}$ we combine (2.14), (2.25) and (2.28) we obtain

$$\lim_{c\to\infty} \bar{L}^{(c)}_i(t) = \tilde{q}_i(\tau) - \mu_i t + p^Y_i(1-\alpha)\lambda \int_\tau^t I_i(\tilde{m}(s))\mathrm{d}s + \alpha\lambda \int_\tau^t I_i(\tilde{m}(s))\mathrm{d}s,$$

for some $0 \leq \tau < t$, when the process is away from the switching plane, unless $I_i(s) = 0\ \forall s \in [0, t]$, then

$$\lim_{c\to\infty} \bar{L}^{(c)}_i(t) = \tilde{l}_i(0) - \mu_i t.$$

Once again by taking a sufficiently small h, we can deduce the differenti-

ability properties by writing

$$\frac{\tilde{l}_i(t) - \tilde{l}_i(t-h)}{h} = -\mu_i + \frac{1}{h}\int_{t-h}^{t} \left(p_i^Y(1-\alpha)\lambda + \alpha\lambda\right) I_i(\tilde{m}(s))\mathrm{d}s,$$

which is the same differentiability condition as the ones described in (2.6) and (2.8) for l_i away from the switching plane.

Limits of weakly convergent subsequences on the switching plane

Now that we have shown that $\mathrm{d}\tilde{m}/\mathrm{d}t = \mathrm{d}m/\mathrm{d}t$ away from the switching plane under the respective policies we are left to prove the same holds on the switching plane as well. Firstly we focus on the JSQ policy. In this case $I_i(t)$ only depends on $Q(t)$, therefore the evolution of $\tilde{q}^{\text{w-JSQ}}(t)$ is independent of $\tilde{l}^{\text{w-JSQ}}(t)$. Furthermore due to (2.2), (2.3) and (2.4) it is clear that if $\tilde{q}^{\text{w-JSQ}}(t)$ is away from the switching plane, then it is moving towards it. Thus all trajectories end on the switching plane and cannot leave it, once they enter. However, regardless of the policy at any time

$$\frac{\mathrm{d}(\tilde{q}_1(t) + \tilde{q}_2(t))}{\mathrm{d}t} = \lambda - (\mu_1 + \mu_2) < 0, \tag{2.30}$$

when both $\tilde{q}_1, \tilde{q}_2 > 0$. Thus all trajectories have to point towards zero on the switching plane itself, without moving away from it. This can only be true if the evolution of $\tilde{q}^{\text{w-JSQ}}$ follows the corresponding part of (2.6). Let us assume now that the time spent between two X arrivals to queue i does not diminish in the limit, while the system is on the switching plane. This would mean that for a while all X arrivals are sent to $i-$ and during this time we would have

$$\begin{aligned}
\frac{\mathrm{d}\tilde{q}_i(t)}{\mathrm{d}t} &= p_i^Y(1-\alpha)\lambda - \mu_i, \\
\frac{\mathrm{d}\tilde{q}_{i-}(t)}{\mathrm{d}t} &= \alpha\lambda + p_{i-}^Y(1-\alpha)\lambda - \mu_{i-},
\end{aligned}$$

which contradicts that the trajectories do not move away from the switching plane. Thus there cannot be $\tau > t$, such that $I_i(s) = 0$ $\forall s \in [t, \tau)$, if the system is on the switching plane at time t, which

combined with (2.28) means that

$$\lim_{c\to\infty}\lim_{\tau\to 0}\bar{\Delta}_i^{(c)}(t+\tau)=0. \tag{2.31}$$

If we combine (2.31) with (2.25), we obtain

$$\lim_{\tau\to 0}\tilde{l}_i(t+\tau)=\tilde{q}_i(t),$$

for $i=1,2$, which shows that $\tilde{l}^{\text{w-JSQ}}$ evolves the same way as $l^{\text{w-JSQ}}$ does on the switching plane as described by (2.7). Therefore we can arrive at the conclusion that $\tilde{m}^{\text{w-JSQ}}\equiv m^{\text{w-JSQ}}$.

The proof for the weighted last X position policy follows a similar logic. Once again due to (2.2), (2.3) and (2.4), the evolution described by (2.29) pushes $\tilde{l}^{\text{w-LXP}}$ towards the switching plane. Let us assume that the system hits the switching plane at time t with $\tilde{l}_i^{\text{w-LXP}}/\mu_i > \tilde{l}_{i-}^{\text{w-LXP}}/\mu_{i-}$ before t, and the time spent after t between two X arrivals to queue i does not diminish in the limit. This, again, would mean that all X arrivals are sent to $i-$ and during a short time, making I_i zero. That however, is a contradiction, since this would cause $\mathrm{d}\tilde{l}_i/\mathrm{d}t=-\mu_i$, while pushing the system away from the switching plane, into the region where $\tilde{l}_i^{\text{w-LXP}}/\mu_i < \tilde{l}_{i-}^{\text{w-LXP}}/\mu_{i-}$, where $\mathrm{d}\tilde{l}_i/\mathrm{d}t=\alpha\lambda+p_i^Y(1+\alpha)\lambda-\mu_i$ as established earlier. Thus, there cannot be $t<\tau$, such that $I_i(s)=0\ \forall s\in[t,\tau)$, which combined with (2.28) provides the same behaviour as described in (2.31). Considering (2.25), this leads to the same behaviour for $\tilde{l}^{\text{w-LXP}}$ on the switching plane as described in (2.9). Furthermore, this means that if the system stays on the switching plane, then $\tilde{l}_i^{\text{w-LXP}}=\tilde{q}_i^{\text{w-LXP}}$ for both $i=1,2$, which causes $\tilde{q}_1^{\text{w-LXP}}/\mu_1=\tilde{q}_2^{\text{w-LXP}}/\mu_2$ on the switching plane. Due to (2.30) this can only be maintained if the evolution of $\tilde{l}^{\text{w-LXP}}$ follows the corresponding part of (2.8). Thus, once again, we can arrive at the conclusion that $\tilde{m}^{\text{w-LXP}}\equiv m^{\text{w-LXP}}$, which concludes the proof of Theorem 2.1.

We also have to note that our reasoning shows that in both cases the switching plane corresponds to $q_1/\mu_1=l_1/\mu_1=q_2/\mu_2=l_2/\mu_2$. This means that the switching planes coincide for the two policies, which implies that the trajectories for q are equivalent for the two policies aside form counatbly many discrete points, where only their derivatives differ.

2.6 Conclusion

Routing traffic based on partial state information is an important challenge in road traffic control by smartphone applications and on-line navigation systems, because these applications only have access to the position of their own users. In an attempt to tackle this challenge we have investigated routing policies for a multiserver queueing system based on partial state information. We took a novel approach in applying the routing control decisions only to a portion of the incoming vehicles, as we assumed that the router has control over just part of the traffic. In more detail, we considered a two-server queueing system in which part of the vehicles (the X vehicles) can be observed and controlled (routed to one of the queues), while the other vehicles (the Y vehicles) act as background traffic which is neither observable nor controllable. We analysed the performance of our policies (in terms of the average sojourn time) under the assumption of Poisson arrivals and independent exponential service times. We compared the results to join-the-shortest-queue (JSQ) and to a static policy that routes according to fixed routing probabilities using no state information, but full information about the system parameters.

An extensive simulation study revealed that the sojourn time as a function of the penetration level (i.e. the percentage of X vehicles) declines fast for small penetration levels, but slowly for large levels. As a consequence, with only a small percentage of controllable vehicles it is possible to obtain an average sojourn time which is close to the minimum value. We also observed in the simulations that a simple policy that sends arriving X vehicles to the queue with the fewest X vehicles performs quite well. If both queues are at least moderately loaded, it is much better than the static policy. A policy that does not base its decisions on the number of X vehicles in each of the queues, but on the position of the last X vehicle (that is, its distance to the server in number of vehicles) performs even slightly better. Obviously, if the service rates of the two queues are unequal, the last X positions of the queues need to be weighted by their average service times. If this average is unknown, or varying over time, an estimate of it based on the measured sojourn time of only a few recently departed X vehicles yields almost as good a performance.

For high loads, the performance of the weighted last X position policy

(w-LXP), which uses partial information, is close to the performance of weighted JSQ (w-JSQ), the optimal policy under full information. When the load of both queues converge to 1, the performance of w-LXP (without knowledge of Y vehicles in the system) is identical to that of w-JSQ (with full knowledge of Y vehicles in the system). We have investigated this remarkable result analytically by means of deterministic dynamical systems approximating the stochastic processes at large system states under high loads. We have proved the proposed dynamical systems to be the *fluid limits* of the original processes. We have shown that after a transitory phase the dynamical systems corresponding to w-LXP and w-JSQ have identical trajectories for the queue lengths, whatever the initial conditions are. This explains the similarity of the sojourn times under both policies. To justify using the analytic results from these dynamical systems to approximate the performance of high loaded queues we performed simulations which confirmed our findings.

Note that the routing policies proposed and investigated in the present work for the special case of two partially observable queues can easily be extended to similar systems with more than two queues, thus it can serve as a model of a road network with multiple possible paths.

Further research on the topic could take various directions. One possibility is to relax the modelling assumptions, by investigating the performance of the proposed routing policies under non-exponential service times at the queues. We would expect similarly good results in that case (particularly in heavy traffic), but the estimation of μ for the ew-LXP policy could be less accurate. Other relaxations or changes in the model could concern non-stationary inputs or state-dependent service rates. Other routing policies could also be considered, for instance policies that take into account more system information, such as the arrival rate of Y vehicles, or estimate the corresponding parameters. Note however, that under high loads not much gain in performance can be expected from such policies, as suggested by the fact that policy w-LXP, which is oblivious of Y vehicles in the system, performs just as good as w-JSQ with full state information.

Chapter 3

Optimal control in urban traffic

Outline. So far in the dissertation we have focused on giving individual advice to the drivers as a solution to reach the optimum in travel times and thus an efficient use of the existing infrastructure. The rest of the work considers reaching the same goal by different means, namely by optimal control of the devices that regulate the traffic flow (e.g. traffic lights, dynamic speed limit boards etc.). These are part of the infrastructure, therefore they effect all actors of traffic.

This chapter is a prelude to Chapters 4 and 5. In this part of the dissertation we discuss methods for optimal control of traffic light networks. The rest of this chapter is organised as follows. In Section 3.1 we provide a short introduction and discuss the relevant literature. In Section 3.2 we introduce a general model with notation and some modeling assumptions that are required in Chapters 4 and 5, where we discuss two optimality concepts and the ensuing algorithms for traffic light control based on the queueing model of Section 3.2.

3.1 Introduction

The structure of an urban area is very similar everywhere in the world. The space is divided into lots, which have direct or indirect connection to roads, streets or alleys that allow for the flow of traffic. These roads often intersect, whilst having relatively short sections between intersections. Due to the shorter distances, barring accidents, the traffic flow barely varies over a single section and is controlled only by speed limits most of the time. Therefore traffic jams occur due to network effects, i.e. due to the interaction of the traffic flows over single roads. This interaction takes place where the sections meet: at junctions. While the layout of intersections can vary between crossings, roundabouts or overpasses, the means of traffic control at them is restricted to giving priority, and thus free way to vehicles taking given directions. This can be done in a static way with signposts and regulations or in a dynamic way with traffic lights. At the inception of their use traffic lights were operated manually, later developing to automatized use with preprogrammed patterns. Modern technology allows for the measurement of traffic conditions and for the complex network of traffic lights to react to the results. As the junctions are the bottlenecks of an urban road network, controlling them efficiently is essential to avoid jams and to make way for free flowing traffic.

There has been much work in the past both on designing and optimising isolated or coordinated signals that reactively resolve congestion in the urban networks. Mainly, these can be categorised into two groups based on the information they operate with, see [58] for a comprehensive review. If a method aims at the optimisation of the cycle time, the offset between nearby intersections and the split of green times in different directions within a cycle based on expectations or statistics without changing these settings during the time when the traffic light operates, then it is referred to as *static control*, or sometimes "fixed time plan". The optimisation can be done in isolation or in a coordinated manner, for instance to create a so-called green wave where vehicles always arrive at intersections during the green time of the cyle. The most relevant static control methods can be found in [48, 47, 75, 154]. In contrast, if an algorithm uses online measurements from on-road detectors (e.g., inductive loops) to optimise signal timings on a cycle-to-cycle basis, then it is called a *vehicle-actuated* control method. For example commonly used implementations

of this type are SCOOT, [66]; UTOPIA [95]; and the hierarchical scheme RHODES, [100]. Methods that apply a combination of a fixed time plan and vehicle-actuated control also exist, one widely used example is SCATS, [90]. The algorithms we examine in Chapters 4 and 5 fall into the latter category. In particular, the length of traffic cycles will be determined either by a simulation study or a closed-loop optimisation which operates on a slower time-scale, while the green time allocation to service phases within a cycle will be done in a vehicle-actuated manner.

Other than categorising by the information used we can also discuss the proposed methods by the approach in optimisation they apply to determine the ideal signal plan. The approaches in the literature range from Mixed-Integer Linear Programming problems, see [40, 48, 47]; Linear Complementary Problem, see [35]; rolling horizon optimisation using dynamic programming, see [46, 60, 100], or its combination with online learning algorithms, see [24], store-and-forward models based on Mode Predictive Control (MPC) optimisation [6, 5, 83, 87, 139, 141, 140]. Many of these approaches formulate the problem in a way that is centralised and thus inherently not scalable, as a centralised approach can require exponential complexity in computations for a global optimal solution. As noted in [114], control strategies such as OPAC in [46] and RHODES in [100] are not real-time feasible for more than one intersection. In fact, they actually became a decentralised scheme via forced implementation at individual intersections that are heuristically coordinated over the network, see [49].

A scalable alternative to these centralised optimizations is to solve a set of loosely coupled optimisations, one for each intersection, with coupling provided by traffic conditions. Such decentralised schemes have the advantage that they are simple to implement. However, due to the fact that they mostly use detailed real-time local information for which the technology was in most cases not available until the emergence of wireless technology, they have been studied to a lesser extent in the traffic management literature. These schemes can be grouped based on the information used to form the criteria for optimality. The two natural approaches are to control the traffic lights based on either the expected number of vehicles to enter the intersection during the next cycle, or the difference in traffic load on the roads leading into the intersection and those leading out.

Notable among the first class is the work of [133], a so-called P0 policy and its variants [31, 134], followed by the work of [80] where the switching cost between phases is taken into account. The core idea of these methods is to allocate the green times such that each phase gets a time proportional to the number of cars expected to arrive at the intersection during its enactment. Beyond heuristic arguments, no formal proof of stability has been provided for this approach, in fact the lack of central control raises the possibility that intersections interact in unexpected ways and cause instability. To limit this, a stabilization mechanism was proposed in [81], however to our best knowledge the first formal proof was given in the paper [3] that Chapter 5 is built upon.

The road traffic literature on the second approach includes the works of [147], [156] and [158]. These are largely inspired by the so-called max-weight/backpressure algorithm which first appeared in [138]. Like the other approach, these methods do not require any *a priori* knowledge of the traffic demand, whilst maintaining the desirable property of rather simple stability results, even if those that hold the original for max-pressure do not naturally extend to road traffic scheduling. Problems may arise because the policy must be aware of the routes or turning ratios of cars and must communicate queue size information between junctions. Further, scheduling decisions are integral and so do not naturally lend to a fixed traffic cycle. Progress alleviating these specific issues can be found in [129] and [4], which serves as the basis for Chapter 4.

The decentralised optimisation techniques have strong routes in packet scheduling for communication networks. As mentioned before Chapter 4 discusses a *backpressure* algorithm, which aims at minimising the "pressure" put on congested parts of the network by comparing the difference of the number of vehicles queueing up in neighbouring intersections. The method originates from the work of [138] and appeared also in [97]. On the other hand Chapter 5 describes an algorithm that is based on the notion of *proportional fairness*. The concept was first introduced in [70]. It has been studied widely since and appeared in connection with different fields such as models of highway traffic, see [50], and [72]; bandwidth sharing, see [93]; or switch networks, see [148]. As the name suggests this concept invokes fairness, as it is the $\alpha = 1$ case of the family known as weighted α-fair allocations. Adapting these methods to urban road traffic poses a challenge since it has features (e.g. scheduling constraints

due to safety regulations) that do not necessarily allow for some of the conditions that make the analysis of the original methods simpler. Thus the algorithms proposed in Chapters 4 and 5 present some changes to these methods and extend their desirable features such as stability.

In the following section we describe the modeling framework in which the algorithms of Chapters 4 and 5 are set. There are slight differences between the assumptions that we make in the two methods, these are discussed as well.

3.2 Basic model and notation

We consider a network of traffic *junctions*, indexed by j. Each junction $j \in \mathcal{J}$ consists of a number of *in-roads*, indexed by $i \in \mathcal{I}_j$. Note that the $\mathcal{I}_j$ are mutually disjoint, and denote $\mathcal{I} = \cup_{j\in\mathcal{J}}\mathcal{I}_j$. A road with multiple lanes having different turning options (such as a left-turn only lane) is modeled as multiple in-roads, thus an in-road may model one or more lanes of traffic flow. Whether these traffic flows are conflicting or not is not considered in this setting. We use the inclusion $i \in j$ to indicate that in-road i is part of junction j, and we let $j(i)$ notate that the junction is used by in-road i. The in-roads of a junction receive green time in a *cyclic* manner, when phases of service are enacted. A *service phase* is a set of in-roads receiving green lights simultaneously. Service phases are indexed by σ, and we let Σ_j notate the set of phases that are used at junction j. The inclusion $i \in \sigma$ denotes the in-road being served during phase σ, and we associate a vector $\sigma = (\sigma_i)_{i\in j}$ with $\sigma_i = \mathbb{I}\{i \in \sigma\}$ for all $\sigma \in \Sigma_j$. The structure of the network is represented in the *links* in the system. A link is a pair of in-roads: after the service is completed, the vehicles from in-road $i \in j$ may join the queue at in-road $i' \in j'$ if the link ii' is included amongst the set of possible links $\mathcal{L} \subseteq \mathcal{I} \times \mathcal{I}$.

We introduce the following notation to address the cyclic nature of the system. Denote by $\{T_n^j\}_{j\in\mathcal{J},n\in\mathbb{N}}$ the sequence of *cycle lengths* for each junction and define the sequence $0 = t_0^j < t_1^j < t_2^j < \dots$ such that

$$t_{n+1}^j = t_n^j + T_{n+1}^j.$$

We will also use the notation

$$N_j(t) = \max\{n \in \mathbb{N} : t_n^j \leq t\}, \tag{3.1}$$

to count the number of cycles completed by time t at junction j. Thus the end of the last finished cycle is $t^j_{N_j(t)}$, and its length was $T^j_{N_j(t)}$. Under a fixed-time plan the length of the cycle at each junction, T^j will be fixed and we will denote the vector of cycle lengths by $T = (T^j)_{j\in\mathcal{J}}$. In Chapter 4 we further relax our model by assuming a common cycle length for every junction and use T as a scalar to denote it. This gives us the opportunity to model the system in discrete time, therefore in Chapter 4 we use a time-slotted approach and observe the system only at times $t = 0, 1, 2, \ldots$, which represent the given multiples of the common cycle length.

In order to adapt to some realistic safety requirements we impose further restrictions regarding the green time allocation. First of all, since pedestrians are not taken into account by other measures, to grant them regular service we require every phase to be enacted in every cycle. Furthermore, when switching between phases the lights are first set to yellow, and then for a short period they are all red to avoid accidents. During this period the vehicles do not receive service. We assume that this requires a fixed amount of time T_{switch} for every switch. The *effective service time* E_n^j – the period during which vehicles actually receive service – is thus given by

$$E_n^j = T_n^j - |\Sigma_j| \cdot T_{\text{switch}} \tag{3.2}$$

for $j \in \mathcal{J}$ and $n \in \mathbb{N}$. In Chapter 4 we account for this loss of service together with the effect that at the beginning of their green time vehicles need to speed up, thus the service that a queue receives during a phase is nonlinear. To do so we introduce the term L, which we call the loss of service. We model this effect in more detail in Chapter 5.

The traffic controller's job is to determine the cycle lengths and to allocate proportions of the effective service time to the separate phases. We will denote these proportions by $P = (P_\sigma^j)_{\sigma\in\Sigma_j, j\in\mathcal{J}}$. Furthermore, we introduce the *allocation vector* $y(P) = (y_i(P))_{i\in\mathcal{I}}$ which shows the

time proportions during which the in-roads receive service, thus

$$y_i(P) = \sum_{\sigma \in \Sigma_j} \sigma_i P_\sigma^j. \tag{3.3}$$

Consequently the set of allocation vectors lie in the following convex set

$$\mathcal{Y} = \Big\{ (y_i)_{i \in \mathcal{I}} \in (0,1)^{\mathcal{I}} : y_i = \sum_{\sigma \in \Sigma_j} \sigma_i P_\sigma^j, \sum_{\sigma \in \Sigma_j} P_\sigma^j = 1 \ \forall j \Big\}. \tag{3.4}$$

Observe that with this notation in-road i will get green light for $y_i E_n^{j(i)}$ amount of time in the given cycle.

The movement of vehicles throughout the system is modelled by their presence at the in-roads. Once receiving service they either leave the system or arrive at another in-road i before the next cycle of junction $j(i)$ begins, where they form a queue. The number of vehicles at each in-road is therefore accounted for by the queue length $Q_i(t)$ at time t. We denote the vector of queue lengths by $Q(t) = (Q_i(t) : i \in \mathcal{I})$. We assume that they can be measured and therefore are known at any time, although in Chapter 4 we allow for measurement errors and consider the controller's decisions to be based upon

$$\hat{Q}_i(t) = Q_i(t) + \delta_i(t), \tag{3.5}$$

where the error term $\delta_i(t)$ is bounded and independent of $Q_i(t)$ or the terms at other in-roads. We denote the vector of the error terms by $\delta(t) = (\delta_i(t) : i \in \mathcal{I})$.

The performance of our proposed policies is evaluated with respect to the evolution of the queues $Q(t)$. We use slightly different models in Chapters 4 and 5, however in both cases (without explaining in detail) we consider some arrival processes $A(t)$ and some departure processes $D(t)$, which give the basic notion

$$Q(t) = Q(0) + A(t) - D(t). \tag{3.6}$$

The departures are clearly bounded by the service received, which is modeled by the term $S(t)$, that is a function of the cycle lengths and the allocation vector. We assume that every vehicle which receives service in

a cycle leaves its current queue, therefore

$$D(t) = S(t) \wedge Q(t), \tag{3.7}$$

where $x \wedge y = \min\{x, y\}$. Thus the evolution of queueing systems and their performance is determined by the decisions that our policies make. In the following chapters we present these in detail, showing that they satisfy various notions of optimality.

Chapter 4

Backpressure based control for traffic light networks

Outline. This chapter introduces a decentralised signal control policy for urban traffic light systems. We consider the basic model described in Chapter 3, using the notations given in Section 3.2. The policy is based on the so-called *backpressure* algorithm, which is commonly used in packet scheduling for communication networks. We provide a formal proof that our proposed scheme, with fixed cycle time and cyclic phases, stabilises the network for any feasible traffic demand. We have also conducted simulations to compare our backpressure policy against other existing distributed control policies in various traffic and network scenarios. The ensuing numerical results suggest that the proposed policy can surpass other policies both in terms of network throughput and congestion.

The chapter is organised as follows. First, in Section 4.1 we discuss some details of the queueing dynamics that we did not cover in the basic model in Chapter 3. We then continue to describe the control scheme itself. The main results for stability are provided in Section 4.2, which can be interpreted as the proof that our policy is capable of stabilising the system for the largest possible set of arrival rates leading to sufficient throughput even in a congested network. We give a detailed proof of our main result in Section 4.2.3, which is aided by supplementary results that are discussed in Section 4.4. Section 4.3 presents the simulation results and contains a numerical comparison between our scheme and other existing policies. We demonstrate the benefits of the proposed cyclic phase backpressure signal control strategy. Finally, Section 4.5 concludes the chapter.

4.1 Cyclic phase backpressure traffic signal control

Throughout this chapter we use the notations and assumptions of the basic model described in Section 3.2. As already mentioned there, we use a time-slotted approach and observe the system only at times $t = 0, 1, 2, \ldots$, which represent the given multiples of a common cycle length T. We accumulate all the possible effects due to which no actual service is given (switching times, speeding up, etc.) into the term L, which we refer to as loss of service.

4.1.1 Queueing dynamics

As discussed before the queue length $Q_i(t)$ denotes the number of cars at in-road $i \in \mathcal{I}$ at the beginning of the tth traffic cycle. The decisions in the policy will be based on the measured queue lengths $\hat{Q}(t)$, which might differ from the actual value of $Q(t)$ by an error term $\delta_i(t)$, which is bounded and independent of $Q_i(t)$ or the terms at other in-roads.

We let $A(t) = (A_i(t) : i \in \mathcal{I}) \in \mathbb{Z}_+^{\mathcal{I}}$ denote the number of external arrivals at in-road i during the tth traffic cycle with the assumption that all the vehicles arrive at once and just before the end of the cycle. The expected number of arrivals or *arrival rate* into each in-road during the tth traffic cycle is defined as $\bar{a}_i(t) := \mathbb{E}[A_i(t)]$. Notice by allowing $\bar{a}_i(t)$ to vary as a function of time, we can model varying traffic demands which undoubtedly can change over the course of a day for instance. In cases where we choose arrival rates to be static and unchanging with time, then we will simply denote these arrival rates by a_i.

As we operate with a fixed, common cycle length, the only control decision to be made is to determine the proportion of the traffic cycle at junction j which is devoted to service phase σ, $P_\sigma^j(t)$. In order to use the allocated time efficiently, whilst maintaining service for all actors of traffic, we require for any policy and for all $j \in \mathcal{J}$, that

$$\sum_{\sigma \in \mathcal{S}_j} P_\sigma^j(t) = 1 - \frac{L}{T} \qquad \text{and} \qquad P_\sigma^j(t) > 0. \tag{4.1}$$

We assume that if the entire traffic cycle were devoted to provide service at in-road i, then vehicles would leave at rate μ_i. We recall, that the

indicator σ_i denotes if in-road i received service during phase σ. Thus, the product, $\mu_i \sigma_i P_\sigma^j(t)$ gives the expected number of cars to leave in-road i under service phase σ, provided the in-road is not emptied. Accordingly, as the term $S_i(t)$ describes the potential number of cars served from in-road i at junction j in traffic cycle t, its expectation must satisfy

$$\mathbb{E}[S_i(t)|Q(t), \delta(t)] = \sum_{\sigma \in \mathcal{S}_j} \mu_i \sigma_i P_\sigma^j(t), \tag{4.2}$$

where we note that the proportions $P_\sigma^j(t)$ are allocated according to the decision in the policy based on $\hat{Q}(t)$, thus the dependency on $\delta(t)$. The random variable $S_i(t)$ only gives the number of cars served if the junction does not empty. Thus, it may be possible for $S_i(t)$ to be greater than the queue size $Q_i(t)$. In this case, $Q_i(t)$ will be the number of cars served, thus the departures are given by the minimum of $Q_i(t)$ and $S_i(t)$ as given in (3.7).

When traffic is served it will move to neighbouring junctions. For $ii' \in \mathcal{L}$, we let $p_{ii'}(t)$ denote the proportion of cars served at in-road i that subsequently join in-road i' during the tth traffic cycle. We assume that cars within an in-road are homogeneous in the sense that each car at the junction has the same likelihood of joining each subsequent junction. We denote the expectation of $p_{ii'}(t)$ by $\bar{p}_{ii'}$. We further assume that this likelihood is constant over time and will not be altered by the queue lengths observed by cars within the network. Thus $[S_i(t) \wedge Q_i(t)]p_{ii'}(t)$ is the number of cars that leave inroad i and, next, join inroad i' provided the in-road does not empty. An important aspect that is novel in our approach is the estimation of the turning fractions, as previous studies have either assumed the turning fractions are either explicitly known or have been calculated prior to the implementation of the policy. In our scheme we emphasise that the turning fractions can be estimated using recent locally calculated information about traffic flows. For instance, if we form an estimate on the turning fractions based on the last k service cycles

$$\bar{q}_{ii'}(t) = \frac{1}{k} \sum_{\kappa=1}^{k} \hat{p}_{ii'}(t - \kappa), \tag{4.3}$$

where $\hat{p}_{ii'}(t)$ denotes the measurement result for $p_{ii'}(t)$. If $Q_i(t) = 0$ then

any estimate may be used to define $p_{ii'}(t)$. Given that turning fractions are stationary and independent of queue sizes, these estimates form an unbiased estimate of the underlying turning probabilities as long as the measurement error in $\hat{p}_{ii'}(t)$ has zero mean, since then

$$\bar{p}_{ii'} = \mathbb{E}[\bar{q}_{ii'}(t)|Q_i(t) > 0]. \tag{4.4}$$

Other rules incorporating historical data or more recent data could also be considered here. What is necessary is that $\bar{q}_{ii'}(t)$ provides an unbiased estimate of the underlying turning fractions of the vehicles for non-empty queues. It is even possible to use an inconsistent estimate of the turning fractions or for the proportions to change on a larger time scale, as long as the estimate is unbiased, independent of $Q(t)$ and its history.

Given a service policy $\{P(t)\}_{t=0}^{\infty}$, we can define the dynamics of our queueing model. In particular, we define for in-road i of junction j

$$Q_i(t+1) = Q_i(t) - S_i(t) \wedge Q_i(t) + A_i(t) + \sum_{i':i'i \in \mathcal{L}} [S_{i'}(t) \wedge Q_{i'}(t)] p_{i'i}(t). \tag{4.5}$$

Here we assume that cars first depart within a traffic cycle and then subsequently cars arrive from other in-roads.

4.1.2 The control policy

Now we are ready to give our proposed policy as follows.

1. At the beginning of each traffic cycle, form an estimate of the actual turning fractions $p_{ii'}(t)$ with the unbiased estimator $\bar{q}_{ii'}(t)$.

2. For each junction $j \in \mathcal{J}$, calculate the weight associated with each service phase at the junction as a function of the measured queue sizes $\hat{Q}(t)$ and the above defined estimated turning probabilities

$$w_\sigma(\hat{Q}(t)) = \sum_{i \in j} \mu_i \sigma_i \left(\hat{Q}_i(t) - \sum_{i':ii' \in \mathcal{L}} \bar{q}_{ii'}(t) \hat{Q}_{i'}(t) \right). \tag{4.6}$$

3. Given these weights, assign the following proportion of the common

cycle length to each phase σ in $\mathcal{S}_j$ within the next service cycle,

$$P_\sigma^j(t) = \frac{\exp\left\{\eta w_\sigma(\hat{Q}(t))\right\}}{\sum_{\pi\in\mathcal{S}_j}\exp\left\{\eta w_\pi(\hat{Q}(t))\right\}}, \tag{4.7}$$

for $j \in \mathcal{J}$ and where $\eta > 0$ is a parameter of the model.[1]

The weights defined in (4.6) are used in the backpressure policy as given by [138]. They can be viewed as a "pressure" a queue places on upstream queues, which is given by the weighted mean of the differences of the queue sizes. The larger the weight associated with a phase, the more important it is to serve the in-roads with green lights during that phase. Then those weights are used to calculate the portion of the traffic cycle for each phase according to (4.7). The distribution (4.7) gives each phase positive service, with more service given to the higher weight phases. As $\eta \to 0$, the service allocation tends to uniform, and as $\eta \to \infty$, the fraction of service given to the highest weight phase(s) tends to 1.

Notice that in contrast to max-weight policies which always serve the phase associated with the highest weight, the proposed policy ensures that each phase (and subsequently each in-road) receives non-zero service in each and every cycle. Thus this ensures a cyclic phase policy, while maintaining the property that higher weights result in higher proportions of allocated green time. Note the policy can be implemented in a decentralised way, after each junction has communicated queue sizes with its upstream in-roads, the phases can be calculated. This decentralisation has numerous advantages: it is computationally inexpensive, it does not require centralised aggregation of information and thus is easier to implement, and it increases the road networks robustness to failures.

The cyclic phase feature which we introduce to the backpressure policies is important from the users' point of view for various reasons. Firstly, the drivers usually expect an ordered phase sequence and anticipate traffic signal changes in advance. Secondly, the waiting time to receive service for any in-road is bounded in our policy while it could be arbitrary large for some in-roads in previous state-of-the-art distributed policies. It is particularly important when considering that pedestrian phases might

[1] Given this definition the weights do not need to be strictly positive.

also be initiated in parallel with the service phases for vehicles.

From an implementation point of view the policy is desirable since it does not require knowledge regarding the destination of each car within the road network, nor does it assume that the proportion of cars moving between links is known in advance. The policy estimates turning fractions and measures queue sizes in an on-line manner, and uses this adaptive estimates and the measurement results to inform the policy decision.

4.2 Stability of the cyclic phase backpressure control policy

4.2.1 Stability region and queueing stability

We define the stability region $\mathcal{A}$ of the network to be the set of arrival rate vectors $a = (a_i : i \in \mathcal{I}) \geq 0$, for which there exists a positive vector $\rho = (\rho_\sigma^j : \sigma \in \mathcal{S}_j, j \in \mathcal{J})$, namely the green time proportion devoted to the service phases in a cycle, and a positive vector $s = (s_i : i \in \mathcal{I})$, namely the departure rates, satisfying the constraints

$$a_i + \sum_{i':i'i\in\mathcal{L}} s_{i'}\bar{p}_{i'i} < s_i, \qquad \text{for each } i \in \mathcal{I}, \tag{4.8}$$

$$\sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \leq 1 - \frac{L}{T}, \qquad \text{for } j \in \mathcal{J}, \tag{4.9}$$

$$s_i \leq \sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \mu_i \sigma_i, \qquad \text{for } j \in \mathcal{J}, \tag{4.10}$$

where equation (4.8) represents the need for the accumulated arrival rates to be less than the potential departure rates, equation (4.9) guarantees the yellow and all-red periods at each cycle maintains sufficient time for switching and setup between phases, and equation (4.10) indicates the departure rates do not exceed the allocated service rates. We let $\bar{\mathcal{A}}$ denote the closure of the stability region, that is the set of rates $a = (a_i : i \in \mathcal{I}) \geq 0$ where the above inequalities in (4.8)–(4.9) may hold with equality. We also note that the random variables $A_i(T)$ and the assigned service time proportions $P_\sigma^j(t)$ are corresponding to a_i and ρ_σ^j and take their respective values from the sets a and ρ in the stable case.

Given the vector of queue sizes $(Q_i(t) : i \in \mathcal{I})$, we define the total queue size of the road network to be

$$Q^{\Sigma}(t) = \sum_{i \in \mathcal{I}} Q_i(t). \tag{4.11}$$

So $Q^{\Sigma}(t)$ gives the total number of cars within the road network. We say that a policy $P^j_{\sigma}(t)$ for serving cars at the junctions stabilises the network for a vector of arrival rates $(a_i : i \in \mathcal{I})$ if the long run average number of cars in the queueing network is finite, in particular,

$$\lim_{T \to \infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} Q^{\Sigma}(t)\right] < \infty. \tag{4.12}$$

This notion of stability originates from the theory of Markov chains, where (4.12) gives a necessary and sufficient condition for positive recurrence, for instance, see [98]. Our model does not assume that the underlying system is Markovian, thus recurrence cannot be defined. However by (4.12) we can have the same understanding of necessary and sufficient conditions for stability as in the previous literature, see [22, 132, 138]. So in the long run we expect there to be a finite number of cars within the road traffic network. If the road network was unstable then we would expect the number of cars within the system to grow over time. Thus we say that a policy is unstable for a vector of arrival rates $(a_i : i \in \mathcal{I})$ if

$$\lim_{T \to \infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} Q^{\Sigma}(t)\right] = \infty. \tag{4.13}$$

We note that if the queue size process was a Markov chain then definition of stable would be equivalent to the the definition of positive recurrence for that Markov chain. However, the process that we will define need not be a Markov chain hence we use the above definition.

4.2.2 Main theoretical results

First of all we state the following known result about the stability region $\mathcal{A}$ defined by (4.8), (4.9) and (4.10). We now refer to the demand induced by the arrival rates as the load on the network. In particular,

we show that any set of arrival rates outside the stability region must be unstable no matter what policy is used. Note that in practice, the traffic load is determined by an origin-destination (O-D) demand rather than a per-inroad arrival rate and turning fraction. If the O-D demand is stationary, then these quantities are also stationary, and the model correctly captures the load on each road. If the O-D demand is non-stationary, then we capture the first-order effects by allowing $A_i(t)$ to vary, but the assumption that $\bar{p}_{ii'}$ is constant is an additional modelling approximation.

Proposition 4.1. *Given that the arrivals at each time, $\{A(t)\}_{t=1}^{\infty}$, are independent identically distributed random variables with expectation a, it follows that if $a \notin \bar{\mathcal{A}}$ then any policy is unstable under these arrival rates, a.*

The previous proposition shows that the best a policy can do to stabilise the road traffic network is to be stable for all rates in $\mathcal{A}$. The following result shows that our policy is indeed stable for all arrival rates within the set $\mathcal{A}$.

Theorem 4.1. *Given that there exists an $\epsilon > 0$ such that for each traffic cycle t, $\bar{a}(t) + \epsilon\mathbf{1} \in \mathcal{A}$ then, for a constant $K > 0$, the long run average queue sizes of in-roads are bounded as*

$$\lim_{\tau\to\infty} \mathbb{E}\left[\frac{1}{\tau}\sum_{t=0}^{\tau-1} Q^{\Sigma}(t)\right] \leq \frac{K}{\epsilon} \tag{4.14}$$

and thus the policy is stable.

We provide the proof of these statements in Section 4.2.3. These results mean, that we can interpret our policy as being stable for the largest possible set of arrival rates. Thus the policy provides sufficient throughput in congested traffic as long as it is possible, reaching an efficient utilisation of the existing capacities. Note that Theorem 4.1 applies to time varying traffic levels. Although the stability region $\mathcal{A}$ is multidimensional, the intuition behind the traffic model can be understood by considering the scalar case. In that case, it corresponds to the expected number of arrivals in each time cycle, T, being bounded

above. If we interpret that bound as the traffic level during peak hour, the theorem applies to networks in which queues would remain stable even if peak hour extended indefinitely. We acknowledge that this is a stricter requirement than necessary, since the system can be stable in the long term even if queues build up during the peak hour, provided they empty sufficiently after the peak.

4.2.3 Proof of the main stability result

In this section we prove Theorem 4.1. In order to do that, we have to clarify some assumptions made about the stochastic elements of our model. Furthermore we need some technical lemmas, which are discussed in Section 4.4. We enumerate our assumptions as follows.

1. The number of vehicles that can be served from any in-road within a traffic cycle is bounded,

$$S_{\max} = \max_{t\in\mathbb{Z}_+, i\in\mathcal{I}} S_i(t) < \infty. \tag{4.15}$$

2. $(p_{ii'}(t) : ii' \in \mathcal{L})$ is stationary and independent of queue lengths $(Q_i(\tau) : i \in \mathcal{I})$ and the number of cars served at each queue $(S_i(\tau) : i \in \mathcal{I})$ for all $\tau \leq t$.

3. The matrix $I - \bar{p}$ is invertible.

4. Due to the previous point we have, for $ii' \in \mathcal{L}$

$$\mathbb{E}\left[S_i(t)p_{ii'}(t)|Q(t)\right] = \sum_{\sigma\in\mathcal{S}_j} \mu_i\sigma_i P_\sigma^j(t)\bar{p}_{ii'} \tag{4.16}$$

$$\mathbb{E}\left[S_i(t) \wedge Q_i(t)p_{ii'}(t)|Q(t)\right] = \mathbb{E}\left[S_i(t) \wedge Q_i(t)|Q(t)\right]\bar{p}_{ii'}. \tag{4.17}$$

5. The number of arrival $(A_i(t) : t \in \mathbb{Z}_+^I)$ is independent of the state of the queues in the road traffic network. Thus the average arrival rate into each junction can be defined as

$$\bar{a}_i(t) = \mathbb{E}[A_i(t)]. \tag{4.18}$$

6. The error term in the queue size measurement, $\delta(t)$ is bounded, i.e.

$$|\delta_i(t)| \leq \delta_{\max} \quad \forall i \in \mathcal{I}. \tag{4.19}$$

We now are ready to develop the proof of Theorem 4.1

Formal proof of Theorem 4.1

In order to prove Theorem 4.1, we need to prove some additional statements. First, we require a bound on the change in the (euclidean) distance of our queue sizes from zero. This is proven in the following proposition.

Proposition 4.2. *There exists a constant K^* such that*

$$\begin{aligned}
&\sum_{i\in\mathcal{I}} \mathbb{E}\left[\frac{1}{2}Q_i(t+1)^2 - \frac{1}{2}Q_i(t)^2 \middle| Q(t), \delta(t)\right] \\
\leq &\mathbb{E}\left[\sum_{i\in\mathcal{I}} Q_i(t)\bar{a}_i(t) - \sum_{j\in\mathcal{J}}\sum_{\sigma\in\mathcal{S}_j} P_\sigma^j w_\sigma(Q(t)) \middle| Q(t), \delta(t)\right] + K^*,
\end{aligned}$$

where $w_\sigma(Q(t))$ is defined by (4.6).

Proof. We can expand the left side through the following inequalities to reach the desired bound

$$\begin{aligned}
&\sum_{i\in\mathcal{I}} \mathbb{E}\left[\frac{1}{2}Q_i(t+1)^2 - \frac{1}{2}Q_i(t)^2 \middle| Q(t), \delta(t)\right] \leq \\
&\sum_{i\in\mathcal{I}} \mathbb{E}\left[Q_i(t)\left(A_i(t) - S_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right) \middle| Q(t), \delta(t)\right] + K_0 \leq \\
&\sum_{i\in\mathcal{I}} \mathbb{E}\left[Q_i(t)\left(\bar{a}_i(t) - \sum_{\sigma\in\mathcal{S}_{j(i)}} \mu_i\sigma_i P_\sigma^{j(i)}(t) + \sum_{i':i'i\in\mathcal{L}}\sum_{\sigma\in\mathcal{S}_{j(i')}} \mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\right) \middle| Q(t), \delta(t)\right] + \tilde{K}.
\end{aligned}$$

The first inequality can be reached by using Lemma 4.6 to expand the recursion (4.5) for some constant $K_0 > 0$. Then by using the inequalities

$$\begin{aligned}&\mathbb{E}\left[Q_i(t)S_{i'}(t)p_{i'i}(t)\middle|Q(t),\delta(t)\right]\\ \leq&\mathbb{E}\Big[Q_i(t)\sum_{\sigma\in\mathcal{S}_{j(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\Big|Q(t),\delta(t)\Big]+K_1\end{aligned}\tag{4.20}$$

and

$$\mathbb{E}\left[Q_i(t)A_i(t)\middle|Q(t)\right]=Q_i(t)\bar{a}_i(t)\tag{4.21}$$

with rearranging we can further expand by taking constant $\tilde{K}=K_0+|\mathcal{L}|K_1$ for instance. (4.20) is proven in Lemma 4.7, whereas (4.21) holds by definition (4.18). We can further rewrite the term by swapping the order of summation inside the expectation from a summation over queues on in-roads $\mathcal{I}$ and then as a summation over schedules $\mathcal{S}_j$ to the other way around, and using the bound on the error terms in (4.19).

$$\begin{aligned}&\sum_{i\in\mathcal{I}}\mathbb{E}\left[\frac{1}{2}Q_i(t+1)^2-\frac{1}{2}Q_i(t)^2\middle|Q(t),\delta(t)\right]\\ =&\mathbb{E}\left[\sum_{i\in\mathcal{I}}Q_i(t)\bar{a}_i(t)\middle|Q(t),\delta(t)\right]+\tilde{K}\\ -&\mathbb{E}\left[\sum_{i\in\mathcal{I}}\Big(\hat{Q}_i(t)-\delta_i(t)\Big)\left(\sum_{\sigma\in\mathcal{S}_{j(i)}}\mu_i\sigma_iP_\sigma^{j(i)}(t)-\sum_{i':i'i\in\mathcal{L}}\sum_{\sigma\in\mathcal{S}_{j(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\right)\middle|Q(t),\delta(t)\right]\\ \leq&\mathbb{E}\left[\sum_{i\in\mathcal{I}}Q_i(t)\bar{a}_i(t)-\sum_{j\in\mathcal{J}}\sum_{\sigma\in\mathcal{S}_j}P_\sigma^jw_\sigma(\hat{Q}(t))\middle|Q(t),\delta(t)\right]+K^*.\end{aligned}$$

By Lemma 4.8 and by the definition of the weights $w_\sigma(\hat{Q}(t))$ in (4.6), we have that

$$\begin{aligned}&\sum_{i\in\mathcal{I}}\hat{Q}_i(t)\left(\sum_{\sigma\in\mathcal{S}_{j(i)}}\mu_i\sigma_iP_\sigma^{j(i)}(t)-\sum_{i':i'i\in\mathcal{L}}\sum_{\sigma\in\mathcal{S}_{j(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\right)\\ =&\sum_{k\in\mathcal{J}}\sum_{\sigma\in\mathcal{S}_k}P_\sigma^k(t)\sum_{i\in k}\mu_i\sigma_i\left(\hat{Q}_i(t)-\sum_{i':ii'\in\mathcal{L}}\hat{Q}_{i'}(t)\bar{q}_{ii'}(t)\right)\\ =&\sum_{k\in\mathcal{J}}\sum_{\sigma\in\mathcal{S}_k}P_\sigma^kw_\sigma(\hat{Q}(t)).\end{aligned}$$

Furthermore we can bound the extra term arising from the measurement error as

$$\begin{aligned}
&\mathbb{E}\left[\sum_{i\in\mathcal{I}}\delta_i(t)\left(\sum_{\sigma\in\mathcal{S}_{j(i)}}\mu_i\sigma_i P_\sigma^{j(i)}(t)-\sum_{i':i'i\in\mathcal{L}}\sum_{\sigma\in\mathcal{S}_{j(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\right)\middle|Q(t),\delta(t)\right]\\
=&\sum_{i\in\mathcal{I}}\delta_i(t)\mathbb{E}\left[\left(\sum_{\sigma\in\mathcal{S}_{j(i)}}\mu_i\sigma_i P_\sigma^{j(i)}(t)-\sum_{i':i'i\in\mathcal{L}}\sum_{\sigma\in\mathcal{S}_{j(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\right)\middle|Q(t),\delta(t)\right]\\
\leq&\delta_{\max}\cdot|\mathcal{I}|\cdot S_{\max}(1+|\mathcal{L}|),
\end{aligned}$$

which by introducing $K^*=\tilde{K}+\delta_{\max}\cdot|\mathcal{I}|\cdot S_{\max}(1+|\mathcal{L}|)$ completes our proof. □

Now that we have proven the previous proposition, we are able to prove the main mathematical result of this paper, Theorem 4.1.

Proof of Theorem 4.1. By Proposition 4.2, we have that

$$\begin{aligned}
&\sum_{i\in\mathcal{I}}\mathbb{E}\left[\frac{1}{2}Q_i(t+1)^2-\frac{1}{2}Q_i(t)^2\middle|Q(t),\delta(t)\right]\\
\leq&\mathbb{E}\left[\sum_{i\in\mathcal{I}}Q_i(t)\bar{a}_i(t)-\sum_{j\in\mathcal{J}}\sum_{\sigma\in\mathcal{S}_j}P_\sigma^j w_\sigma(\hat{Q}(t))\middle|Q(t),\delta(t)\right]+K^*. \qquad (4.22)
\end{aligned}$$

Further, by Lemma 4.5, we know that

$$\sum_{\sigma\in\mathcal{S}_j}P_\sigma^j w_\sigma(\hat{Q}(t))\geq\max_{\sigma\in\mathcal{S}_j}\left\{w_\sigma(Q(t))\right\}-\delta_{\max}S_{\max}|\mathcal{I}|-\frac{1}{\eta}\log|\mathcal{S}_j|. \qquad (4.23)$$

Applying this bound to (4.22) and taking expectations, we have that

$$\begin{aligned}
&\sum_{i\in\mathcal{I}}\mathbb{E}\left[\frac{1}{2}Q_i(t+1)^2-\frac{1}{2}Q_i(t)^2\right]\\
&\leq\mathbb{E}\left[\sum_{i\in\mathcal{I}}Q_i(t)\bar{a}_i(t)-\sum_{j\in\mathcal{J}}\max_{\sigma\in\mathcal{S}_j}\{w_\sigma(Q(t))\}\right]+K
\end{aligned} \qquad (4.24)$$

where

$$K = K^* + \delta_{\max} S_{\max} |\mathcal{I}| + \sum_{j \in \mathcal{J}} \frac{1}{\eta} \log |\mathcal{S}_j|.$$

We focus on bounding the term (4.24) above. Recalling the definition of $Q^{\Sigma}(t)$, (4.11), observe that

$$\begin{aligned}
&\mathbb{E}\left[\sum_{j \in \mathcal{J}} \max_{\sigma \in \mathcal{S}_j} \{w_\sigma(Q(t))\} - \sum_{i \in \mathcal{I}} Q_i(t) \bar{a}_i(t)\right] \\
=&\mathbb{E}\left[Q^{\Sigma}(t) \times \left(\sum_{j \in \mathcal{J}} \max_{\sigma \in \mathcal{S}_j} \left\{w_\sigma\left(\frac{Q(t)}{Q^{\Sigma}(t)}\right)\right\} - \sum_{i \in \mathcal{I}} \frac{Q_i(t)}{Q^{\Sigma}(t)} \bar{a}_i(t)\right)\right] \\
\geq&\mathbb{E}\left[Q^{\Sigma}(t) \times \min_{u \geq 0: \sum_{i \in \mathcal{I}} u_i = 1} \left(\sum_{j \in \mathcal{J}} \max_{\sigma \in \mathcal{S}_j} \{w_\sigma(u)\} - \sum_{i \in \mathcal{I}} u_i \bar{a}_i(t)\right)\right] \\
\geq&\mathbb{E}\left[Q^{\Sigma}(t) \times \epsilon\right]. \qquad (4.25)
\end{aligned}$$

In the first inequality above, we substitute u_i for $\frac{Q_i(t)}{Q^{\Sigma}(t)}$ and minimize over u_i. In the second inequality, we apply our bound from Lemma 4.9,

$$\epsilon < \left(1 + \frac{L}{T}\right) \min_{u \geq 0: \sum_{i \in \mathcal{I}} u_i = 1} \left(\sum_{j \in \mathcal{J}} \max_{\sigma \in \mathcal{S}_j} \{w_\sigma(u)\} - \sum_{i \in \mathcal{I}} u_i \bar{a}_i(t)\right).$$

Applying (4.25) to (4.24), we gain a far simpler bound

$$\sum_{i \in \mathcal{I}} \mathbb{E}\left[\frac{1}{2} Q_i(t+1)^2 - \frac{1}{2} Q_i(t)^2\right] \leq -\epsilon \mathbb{E}\left[Q^{\Sigma}(t)\right] + K. \qquad (4.26)$$

Summing over $t = 0, ..., \tau - 1$, we have

$$\sum_{i \in \mathcal{I}} \mathbb{E}\left[\frac{1}{2} Q_i(\tau)^2 - \frac{1}{2} Q_i(0)^2\right] \leq -\epsilon \mathbb{E}\left[\sum_{t=0}^{\tau-1} Q^{\Sigma}(t)\right] + \tau K. \qquad (4.27)$$

Finally, rearranging and dividing by τ gains the required bound,

$$\lim_{\tau\to\infty}\mathbb{E}\left[\frac{1}{\tau}\sum_{t=0}^{\tau-1}Q^{\Sigma}(t)\right]\leq\frac{K}{\epsilon}+\lim_{\tau\to\infty}\frac{1}{\tau\epsilon}\mathbb{E}\left[\frac{1}{2}\sum_{i\in\mathcal{I}}Q_i(0)^2\right]=\frac{K}{\epsilon}<\infty.\ \square$$

4.3 Performance evaluation and design

4.3.1 Simulation settings

In this section we evaluate via simulation the performance of our proposed cyclic phase backpressure traffic signal control and compare its performance with a number of existing self-control (i.e. decentralised) schemes by [81], [133] and [156] as detailed below.

First, the self-control scheme in [81] aims at minimising the waiting times at each intersection anticipating future arrivals into those intersections instead of just efficiently clearing exiting queues as in a conventional μc priority rule, see [133]. However, since future traffic demand is not known, this scheme of [81], more or less *greedily* attempts to minimise the waiting time. When the setup time or the yellow traffic signal is ignored, this policy (referred to as greedy policy below) tends to allocate service to the phase that has the longest queue length. In contrast, our policy is non-greedy, as it ensures that an action at this time is not too suboptimal, regardless of what future traffic is like. Although it does not explicitly seek to minimise the waiting time, it is likely to result in lower waiting time and subsequently lower total travel time through a network than a greedy algorithm that does.

Second, the priority rule of the self-control scheme in [133] is approximately giving a green time split proportionally to the total number of vehicles on the in-roads and thus will be referred to as proportional scheme in the rest of this section (not to be confused with the proportional fair allocation scheme described in Chapter 5).

The third and final policy in [156] allocates green time to the phase that has the highest queue backlog differences between upstream queues and downstream queues, thus it will be referred to as non-cyclic backressure policy in this section. Note that although these benchmarks are in their genesis by the standards of currently implemented centralised schemes,

they are state-of-the-art among distributed schemes, and thus appropriate for the comparison. In the following list we shortly summarise the policies that we compare.

- Cyclic phase backpressure policy proposed in this work: We refer the reader to Section 4.1 for details.
- Non-cyclic backpressure policy [156]:
  1. At the beginning of each time slot, based on recent occurrences, form an estimate of the turning fractions according to

$$\bar{q}_{ii'}(t) = \frac{1}{k} \sum_{\kappa=1}^{k} p_{ii'}(t-\kappa) \tag{4.28}$$

     where k is a parameter of the model.
  2. For each junction $j \in \mathcal{J}$, calculate the weight associated with each service phase at the junction as

$$w_\sigma(Q(t)) = \sum_{i \in j} \mu_i \sigma_i \Big(Q_i(t) - \sum_{i' : ii' \in \mathcal{L}} \bar{q}_{ii'}(t) Q_{i'}(t) \Big). \tag{4.29}$$

  3. Given these weights, assign the whole service time of the next time slot to phase $\sigma^* \in \mathcal{S}_j$ where $w_{\sigma^*} > w_\sigma \ \forall \sigma \in \mathcal{S}_j$.
- Proportional policy [133]:
  1. At the beginning of each traffic cycle, calculate the weight associated with each service phase at each junction $j \in \mathcal{J}$ as

$$w_\sigma(Q(t)) = \sum_{i \in \sigma} Q_i(t). \tag{4.30}$$

  2. Given these weights, within the next service cycle assign an amount of time to each phase $\sigma \in \mathcal{S}_j$ that is proportional to

$$P_\sigma^j(t) = \frac{w_\sigma(Q(t))}{\sum_{\pi \in \mathcal{S}_j} w_\pi(Q(t))}.$$

- Greedy policy [81] with no switching time:

1. At the beginning of each time slot, calculate the weight associated with each service phase at each junction $j \in \mathcal{J}$ according to equation (4.30).
2. Given these weights, assign the whole service time of the next time slot to phase $\sigma^* \in \mathcal{S}_j$ where $w_{\sigma^*} > w_\sigma \ \forall \sigma \in \mathcal{S}_j$.

In this section we utilise an open source microscopic simulation package, SUMO (Simulation of Urban MObility), see [136], to study the above schemes in a small network of two intersections (Figure 4.1) and in a large network (Figure 4.2) that reassembles the Melbourne CBD(Australia) with about 70 intersections.

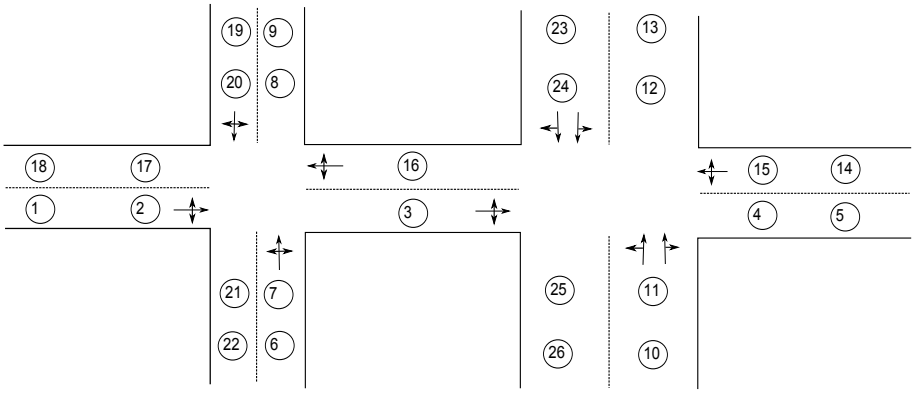

Figure 4.1. Small network topology.

The small network has 2 junctions consisting of several in-roads (numbered from 1 to 26 on the figure Figure 4.1). All the roads have bi-directional traffic with the North-South road going through the right intersection having double lanes. Direction of traffic movements on this network is indicated on each in-road leading to the junctions. The ingress queues, where vehicles enter the network, are assumed to be infinite and represented by a set of long links (i.e. links $\{1, 18, 9, 19, 13, 23, 5, 14, 10, 26, 6, 22\}$ on Figure 4.1). The cars immediately appear on the connecting in-roads inside the network, which are of finite capacity. Since the ingress queues are infinitely large, vehicles can enter the network even when there is a heavy congestion on the bottleneck link. All other links (i.e. links$\{2, 3, 4, 7, 8, 11, 12, 20, 21, 24, 25\}$ on

Figure 4.1) have the same length at 375 meters which can accommodate maximum 50 cars per lane.

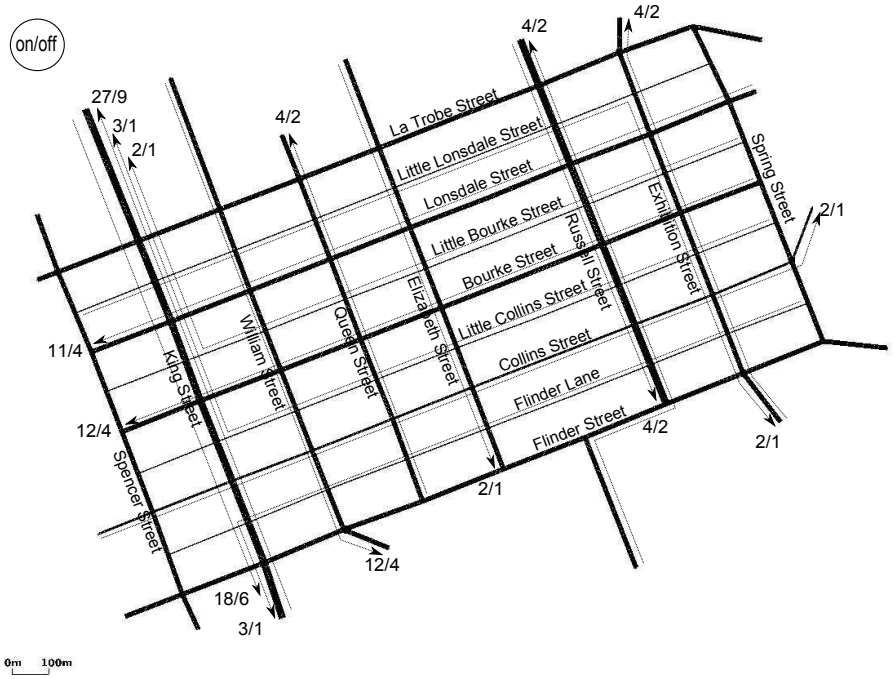

Figure 4.2. Large CBD network with demands.

The topology of the large CBD Melbourne network is shown in Figure 4.2. It consists of 73 intersections and 266 links. Most of the roads are bi-directional except for Little Lonsdale Street, Little Bourke Street, Little Collins Street and Flinder Lane which only have a single lane mono-directional traffic. King Street and Russell Street are the biggest roads in this scenario, each is modeled as 3 lanes each direction. Collins Street has one lane each direction. All other roads have two lanes each direction. The link lengths are varied between 106 meters for the vertical links and 214 meters and 447 meters for the horizontal links except for the ingress links at the edges.

Results are given in terms of the total number of vehicles in the network and the congestion level which is the average number of congested links in large network after long simulation runs using the different control

schemes. In all the studied scenarios, the exact queue lengths and turning fractions are observed directly from the simulation and used to make control decision in various policies. These variables are calculated using Matlab, based on the actual control algorithm and then are fed back into the SUMO simulation at every time step. We ignore switching times (i.e. transition between phases) in all control schemes in our study. This overhead can be incorporated into the simulation by extending the phase times. Nevertheless, the qualitative insights gained in this section would not change by that extension.

4.3.2 Performance study

Below we evaluate the performance of our cyclic phase backpressure scheme and compare it with other policies using fixed setting of routes in the studied networks using simulation. The cycle time of the cyclic phase backpressure policy and the proportional policy were set to 30 seconds, while the slot time of the non-cyclic backpressure policy and the greedy policy were set to 10 seconds in our simulation. Note that these timings were chosen rather arbitrary without any optimisation and applied for all the studied policies for comparison purposes. The optimisation of timings will be presented in the next sub-section (Sec. 4.3.3).

Small network

First, we study a small network scenario, for which the turning information and the arrival rates are indicated in Figure 4.3. In particular, the arrows indicate 10 routes with direction and demands (cars/minutes) in the peak and off-peak (i.e. on/off) time periods as shown in Figure 4.3. The η value introduced in (4.7) was set to 2.5. The main traffic flows are the ones with North-South direction of the second junction. The two backpressure policies give the majority of service time to the North-South phase of that junction which leads to heavy congestion on links 1, 2, 3, 14 and 15. On the other hand, the proportional policy and greedy policy put more balance between the service times depending on the queue lengths which creates more congestion in the North-South direction at the cost of having less congestion in East-West direction.

Results are shown in Figure 4.4 and Figure 4.5 where the total number

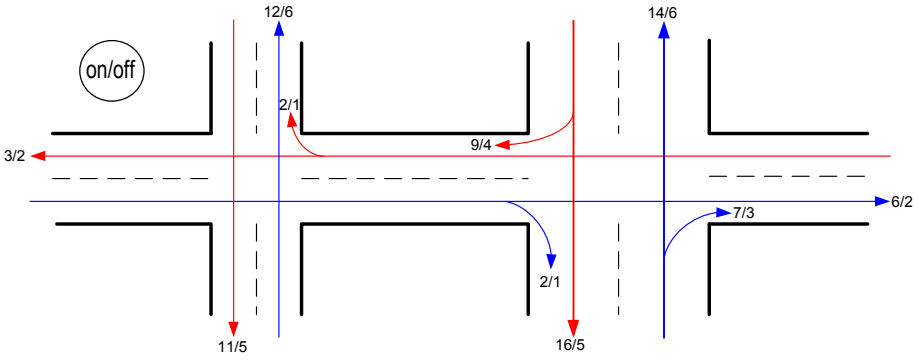

Figure 4.3. Small network with demands.

of vehicles in the network and in the congested link between the two junctions are plotted over time. Note that there are two in-roads between the two intersections but only link 3 is congested due to large traffic flows in the North-South direction at the second intersection.

Note that Figures 4.4 and 4.5 were based on the number of cars present at the times when control decisions were made which is 30 seconds for cyclic phase backpressure and proportional policies. In contrast, the average travel time depends on the waiting time on individual link which is an integral of queue size over continuous time. For this reason, intermediate queue size was also measured at 10 s intervals in the simulation, and the results differed by less than 2% in compare with the coarse sampling at once per cycle assuming linear interpolation. The resulting travel time values are reported in the next sub-section 4.3.3.

Observe that the cyclic phase backpressure control yields a lower number of total vehicles present in the network and thus results in higher number of vehicles reaching their destination (i.e. increased network throughput) during the whole simulation. This is due to the fact that in the two backpressure control schemes when the bottleneck link (link 3) is congested, less green time will be allocated to the East-West direction at the first junction. As a result more traffic can move through the North-South direction and the impact of a spill back from the second junction on the overall network throughput decreases. Similarly, the non-cyclic backpressure policy also outperforms the proportional policy and the greedy policy, since those two schemes allocate similar amount

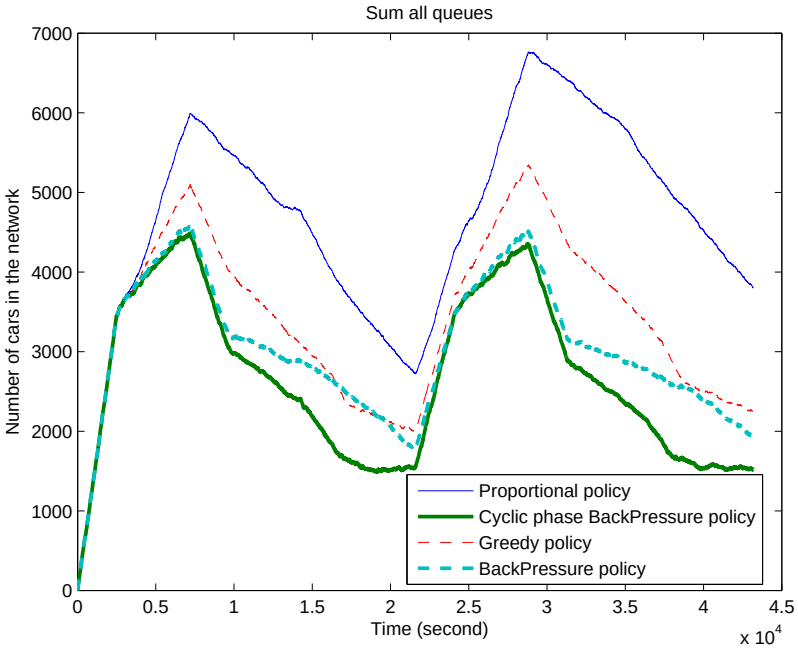

Figure 4.4. Total queue lengths for the small network using different policies. Cycle time for the cyclic phase backpressure policy and the proportional policy is 30 seconds. Slot time for the backpressure policy and the greedy policy is 10 seconds.

of green time to the East-West direction at the first junction despite the presence of a spilled back traffic and thus waste some of the green time.

We note that the performance of each policy largely depends on the cycle length in the cyclic phase backpressure policy and the proportional policy or the frequency of making decision in the non-cyclic backpressure policy and the greedy policy. For instance, Figure 4.6 represents the similar scenario where the cycle length of the cyclic phase backpressure policy and the proportional policy is increased to 60 seconds. While the queue evolution of the cyclic phase backpressure policy stays unchanged, there is significantly less congestion in the proportional policy. The optimality of cycle length and frequency of making decision will be studied in Section 4.3.3.

Large Melbourne CBD network

A similar study is performed with a large network with its turning information and arrival rates indicated in Figure 4.2. The parameter η is once again set to 2.5. In this setting, the King Street has the largest flows, thus, any flow that shares an intersection with King Street tends to be under-served especially the intersection between King Street and Lonsdale Street and the intersection between King Street and Bourke Street. Generally in the peak period, congestion in any link will cause spill-back which leads to further congestion in the neighbouring junctions. This can only be recognized by the two backpressure policies through comparing the in-road i and out-road i', and more service will be allocated

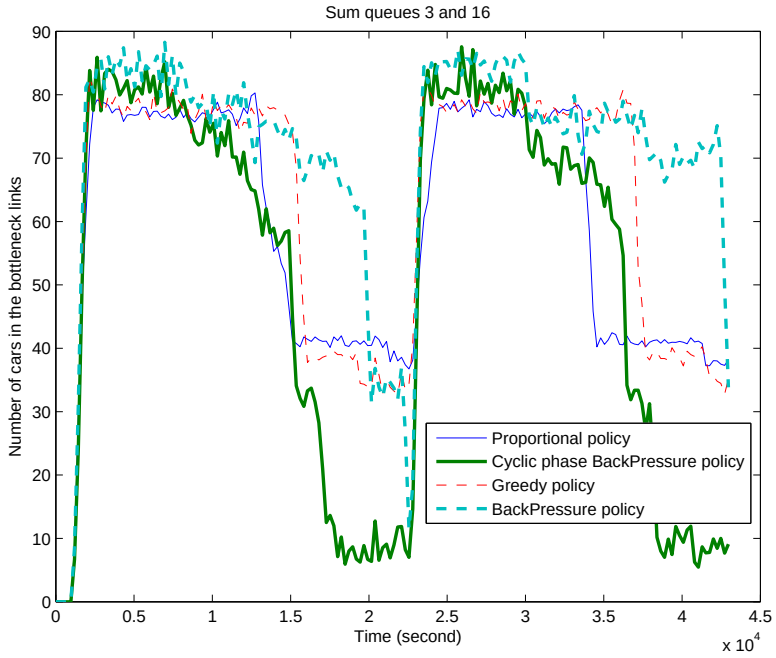

Figure 4.5. Congestion of bottleneck links for the small network using different policies. Cycle time for the cyclic phase backpressure policy and the proportional policy is 30 seconds. Slot time for the backpressure policy and the greedy policy is 10 seconds.

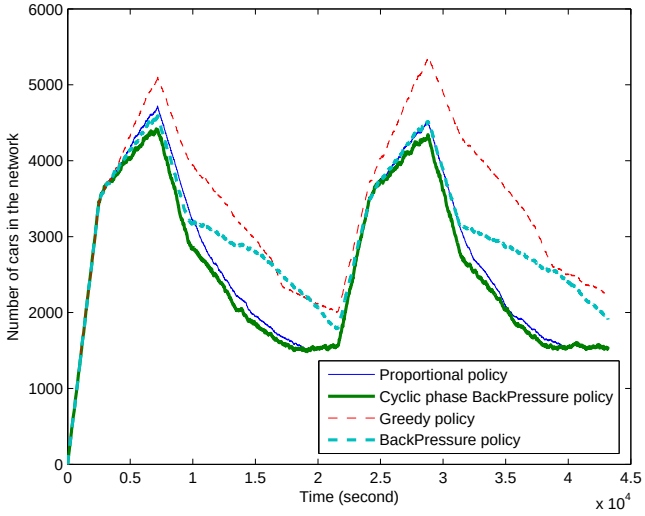

Figure 4.6. Total queue lengths for the small network using different policies. Cycle time for the cyclic phase backpressure policy and the proportional policy is 60 seconds. Slot time for the backpressure policy and the greedy policy is 10 seconds.

in this case to traffic flows on the less congested directions. In contrast, the proportional policy and the greedy policy only consider the queue lengths present at in-road i, and may waste some green time to the congested direction where traffic comes to a standstill due to the spill back.

Results for this scenario are shown in Figure 4.7 and Figure 4.8. As shown in Figure 4.7 the cyclic phase backpressure policy has the lowest total number of vehicles in the network, thus it provides the highest throughput, whereas the second highest throughput is provided by the non-cyclic backpressure policy. Similarly to the previous scenario, the two backpressure policies outperform the other two policies in case of heavy congestion because they take into account downstream queue lengths, and thus allocate resources (i.e. service phases) more efficiently.

Figure 4.8 plots the congested link over time. Herein a link is said to be congested at a certain time if its queue length is more than 85% of

the link capacity. It is clear that the two backpressure policies reduce the number of congested links significantly (i.e. less number of vehicles inside the network) resulting in higher network throughput.

4.3.3 Experimental parameter design

The cycle length in the cyclic phase backpressure policy and the proportional policy and the frequency of making decision in the non-cyclic backpressure policy and the greedy policy play a crucial role in the performance of the control schemes. A long cycle length or the low frequency of making decisions may be less efficient due to the fact that the queue might be depleted before the end of the service time. On the other hand,

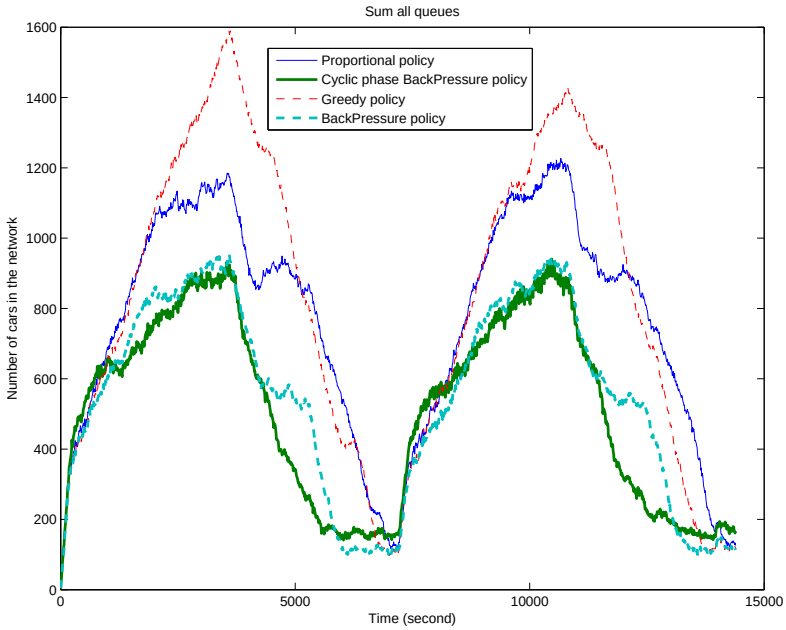

Figure 4.7. Total queue lengths for the large network using different policies. Cycle time for the cyclic phase backpressure policy and the proportional policy is 30 seconds. Slot time for the backpressure policy and the greedy policy is 10 seconds.

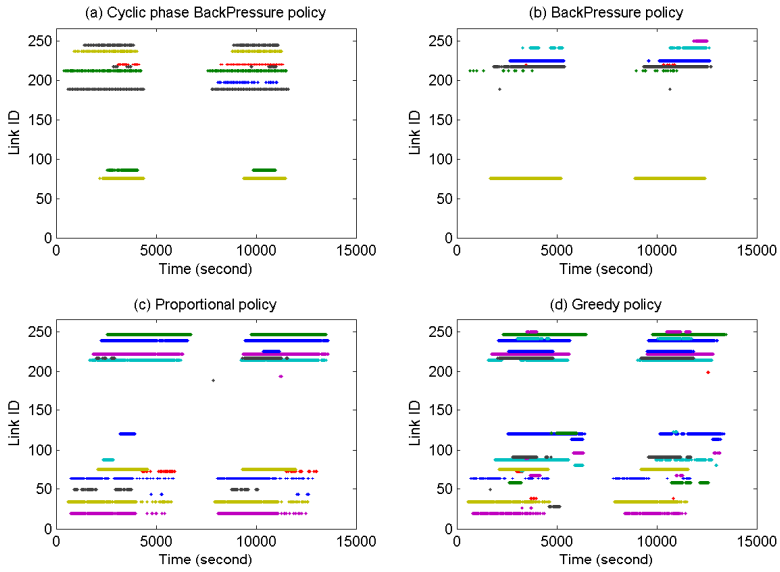

Figure 4.8. Congestion level for the large network using different policies. Cycle time for the cyclic phase backpressure policy and the proportional policy is 30 seconds. Slot time for the backpressure policy and the greedy policy is 10 seconds.

a short cycle length may reduce the overall capacity since the vehicles have to stop and accelerate more often. Note that the latter in fact represents a switching cost between phases even though the yellow traffic signal is not considered here. This subsection investigates the impact of the cycle time and decision making frequency on the throughput and congestion level of each scheme. We study both network topologies (the small network and the large CBD network) under the similar demand levels as in the previous subsection with different cycle times and decision frequencies. Particularly, for the cyclic phase backpressure policy and the proportional policy, the cycle length is set to $\{30, 60, 90, 120\}$ seconds, and for the non-cyclic backpressure policy and the greedy policy, a decision is made every $\{10, 30, 60, 90\}$ seconds, respectively.

Small network

For small network, the results are presented in Figure 4.9, Figure 4.10, and Figure 4.11. Figure 4.9 shows the average number of vehicles in the network plotted against different cycle times. Because a vehicle does not disappearfrom the network until it exits, a lower number of vehicles in the network equates to higher throughput of the same demand. In this scenario, all of the studied policies provide a similar throughput using their corresponding best setting. Furthermore, it can be observed that in a congested network higher cycle times tend to have better throughput because the traffic flows are less interrupted by the switching between phases. Nevertheless, the proposed cyclic phase backpressure is less sensitive to the changes of cycle time while producing a compatible throughput.

Figure 4.10 plots the average link densities versus link ID. It shows

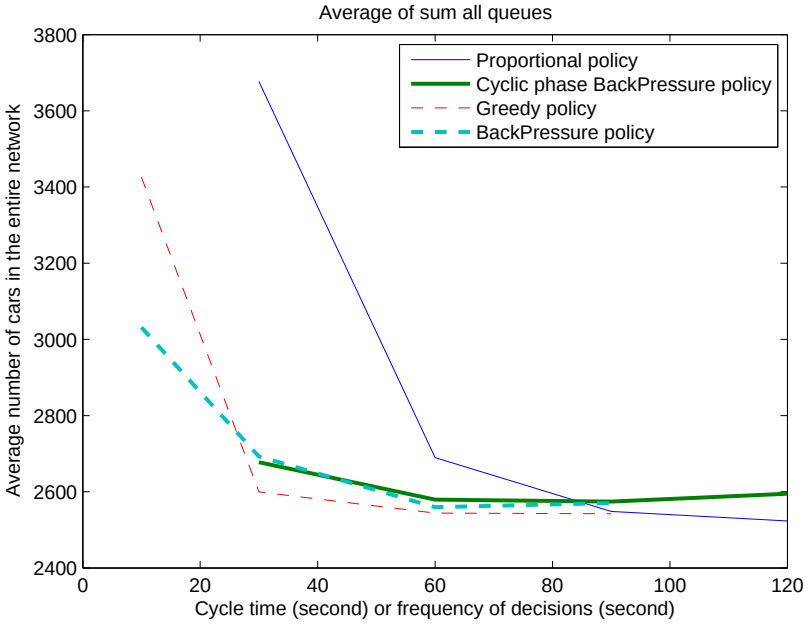

Figure 4.9. Average queue length vs cycle time/frequency of decision for the small network.

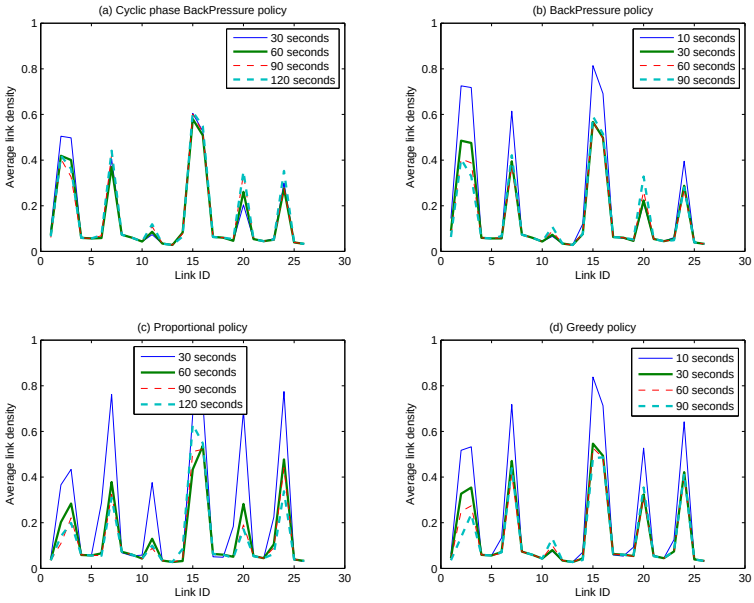

Figure 4.10. Average link density for the small network.

that congestions occur in the same set of links throughout all of the studied policies. In overall, the cyclic phase backpressure policy and the proportional policy have lower link density than the other two policies. Figure 4.11 shows the maximum link density over time pointing to when the congestions occur in the network. In all cases, congestions appear during the peak period and some portion of time during the off-peak period before the build-up traffic can be substantially drained. However, when the cycle time or the decision frequency is set too small (e.g. 10 seconds for the non-cyclic backpressure policy and the greedy policy and 30 seconds for the cyclic phase backpressure policy and the proportional policy), the congestions cannot be cleared at all except for the case, when the networks is operated under our cyclic phase backpressure policy.

Large Melbourne CBD network

The impact of the cycle time and frequency of decision making on network throughput and congestion level using different policies for a large network are investigated and discussed in this subsection. The results are shown in Figure 4.12, Figure 4.13 and Figure 4.14.

The average number of vehicles in the network for each setting is presented in Figure 4.12. Unlike the results in the small network, there clearly exists an optimal value for the cycle length or decision making frequency of each policy. In particular the optimal cycle length for proportional policy is 60 seconds, while the optimal cycle length/decision frequency for all other policies is 30 seconds.

Observe that the optimal cycle length in the large network scenario is shorter than that of the small network, which can be explained by the shorter link lengths and the bigger number of intersections on any

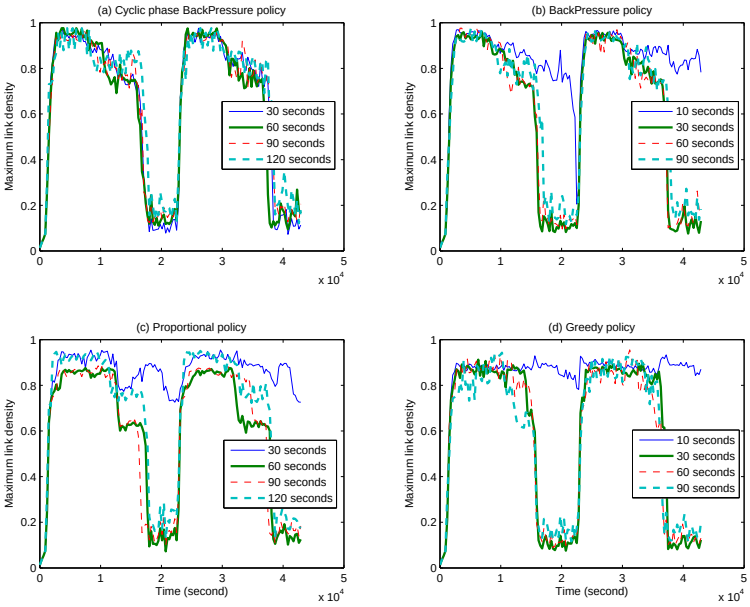

Figure 4.11. Maximum link density for the small network.

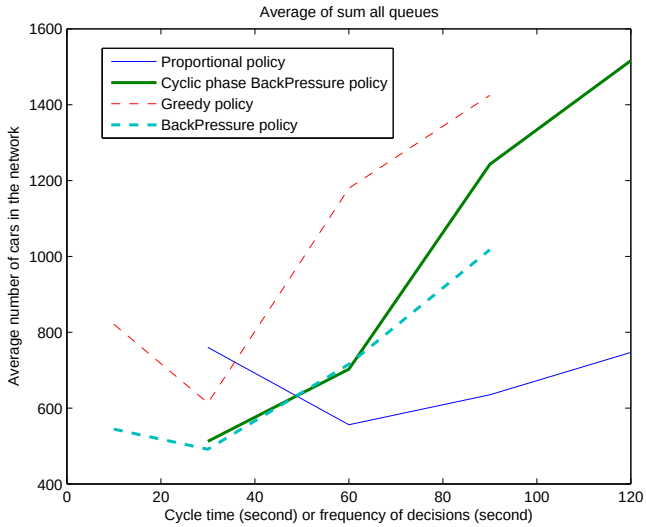

Figure 4.12. Average queue length vs cycle time/frequency of decision for the large network.

	Proportional	Cyclic phase BP	Greedy	BP
Avg. travel time	478.0	409.6	514.0	408.5

Table 4.1. Average travel time (in seconds) for the Melbourne CBD network using optimal setting for each policy (i.e. 60 seconds cycle length for the proportional policy, and 30 seconds cycle length for all other policies).

route. Both increase the interdependency between intersections and their performance as arriving traffic into any internal intersection is an output traffic from the others.

Furthermore, the average travel time of each policy in their best setting of cycle length and decision making frequency is shown in Table 4.1. There is a strong correlation between the average number of vehicles in the network and the average travel time through that network. In particular, the higher number of vehicles in the network results in the longer travel time and vice versa. It can be seen that the proposed cyclic phase backpressure policy has a competitive average travel time between

all the considered policies, and the results show that the backpressure-based policies yield a significantly better average travel time than that of the greedy policy. Note that this better average travel time has been achieved with the control decisions using the queue size measurements at discrete time intervals (once in every cycle) only as explained earlier.

Figure 4.13 plots the average link densities against the link ID. It shows that the congestion area is varied with different policy and with different parameter settings. Any cycle length/decision frequency setting other than the optimal setting obviously increases the congestion greatly.

Finally, Figure 4.14 shows the maximum link density over time. The cycle length plays a vital role to prevent congestion in this scenario. In the cyclic phase backpressure policy, the 30 second cycle length is undoubtedly outstanding. In other policies, the small cycle lengths seem to be better due to the short link lengths.

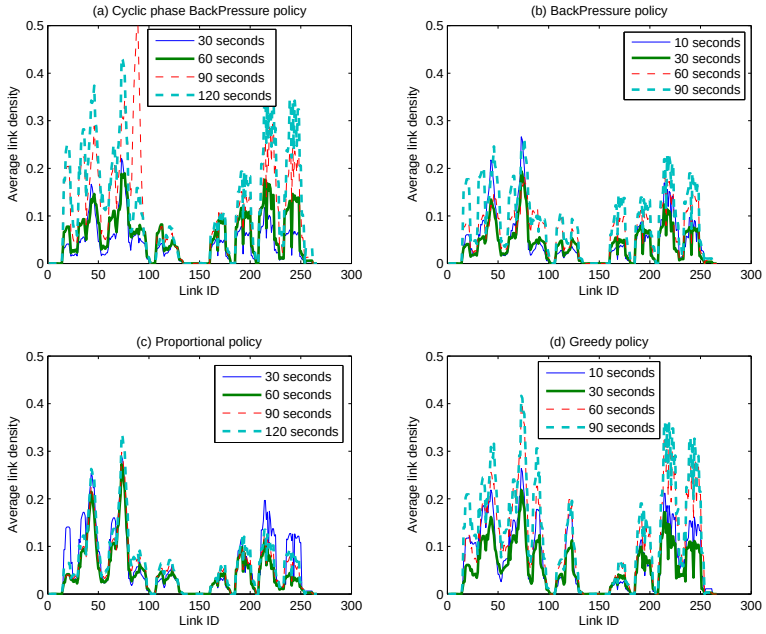

Figure 4.13. Average link density for the large network.

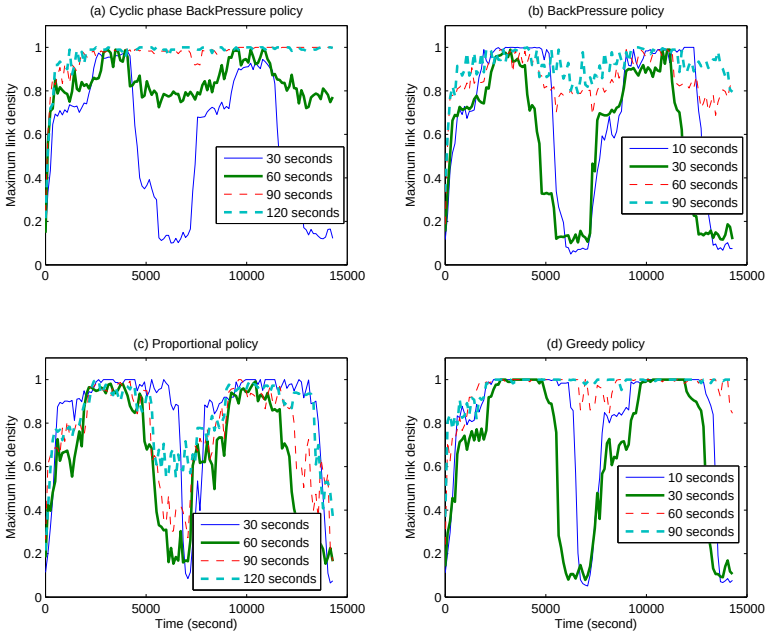

Figure 4.14. Maximum link density for the large network.

4.4 Supplementary results

In this section, we will prove a number of additional lemmas that are required for the main proofs. The following lemmas can be derived from our previous assumptions and the definitions in Section 4.1. The first lemma describes the difference of the weights caused by the error in measurement, whereas the second describes a general result on weights. The third lemma is a consquence of them, introducing a bound which is used multiple times in proving later statements.

Lemma 4.3.

$$w_\sigma(Q(t)) = w_\sigma(\hat{Q}(t)) - w_\sigma(\delta(t)), \tag{4.31}$$

where $w_\sigma(Q(t))$ *is defined according to* (4.6).

Proof. This statement is a straightforward consequence of the definition of the weights and the measurement error given in (3.5) and (4.6). Namely,

$$\begin{aligned} w_\sigma(Q(t)) &= \sum_{i\in j} \mu_i\sigma_i \left(Q_i(t) - \sum_{i':ii'\in\mathcal{L}} \bar{q}_{ii'}(t) Q_{i'}(t) \right) \\ &= \sum_{i\in j} \mu_i\sigma_i \left(\left(\hat{Q}_i(t) - \epsilon_i(t)\right) - \sum_{i':ii'\in\mathcal{L}} \bar{q}_{ii'}(t) \left(\hat{Q}_{i'}(t) - \delta_{i'}(t)\right) \right) \\ &= \sum_{i\in j} \mu_i\sigma_i \left(\hat{Q}_i(t) - \sum_{i':ii'\in\mathcal{L}} \bar{q}_{ii'}(t) \hat{Q}_{i'}(t) \right) \\ &- \sum_{i\in j} \mu_i\sigma_i \left(\delta_i(t) - \sum_{i':ii'\in\mathcal{L}} \bar{q}_{ii'}(t) \delta_{i'}(t) \right) \\ &= w_\sigma(\hat{Q}(t)) - w_\sigma(\delta(t)). \end{aligned} \tag{4.32}$$

□

Lemma 4.4. *Given weights $(w_y : y \in \mathcal{Y})$ with elements indexed by finite set $\mathcal{Y}$, we consider Y a random variable with the following probability of event y:*

$$P_y = \frac{e^{\eta w_y}}{\sum_{y'\in\mathcal{Y}} e^{\eta w_{y'}}} \tag{4.33}$$

then, the expected value of the weights w under this distribution obey the following inequality

$$\mathbb{E}w_Y \geq \max_{y\in\mathcal{Y}} w_y - \frac{1}{\eta}\log|\mathcal{Y}|. \tag{4.34}$$

Proof. In the following inequality, we note that the entropy of a distribution $H(P) = -\mathbb{E}\log P(Y)$ is maximized by a uniform distribution on

$\mathcal{Y}$, $H(U) = \log|\mathcal{Y}|$.

$$\begin{aligned}\mathbb{E}w_Y =& \frac{1}{\eta}\log\left(\sum_{y\in\mathcal{Y}} e^{\eta w_y}\right) + \frac{1}{\eta}\mathbb{E}\log P(Y) \geq \frac{1}{\eta}\log\left(\sum_{y\in\mathcal{Y}} e^{\eta w_y}\right) - \frac{1}{\eta}\log|\mathcal{Y}| \\ \geq& \frac{1}{\eta}\log\left(e^{\eta \max_{y\in\mathcal{Y}} w_y}\right) - \frac{1}{\eta}\log|\mathcal{Y}| = \max_{y\in\mathcal{Y}} w_y - \frac{1}{\eta}\log|\mathcal{Y}|,\end{aligned}$$

as required. □

Lemma 4.5.

$$\sum_{\sigma\in\mathcal{S}_j} P_\sigma^j w_\sigma\left(\hat{Q}(t)\right) \geq \max_{\sigma\in\mathcal{S}_j}\left\{w_\sigma\left(Q(t)\right)\right\} - \delta_{\max}S_{\max}|\mathcal{I}| - \frac{1}{\eta}\log|\mathcal{S}_j|. \quad (4.35)$$

Proof. We can prove this statement by applying Lemma 4.4 to $w_\sigma(\hat{Q}(t))$, expanding the terms according to Lemma 4.3 and using the bound on $\delta(t)$ in (4.19).

$$\begin{aligned}\sum_{\sigma\in\mathcal{S}_j} P_\sigma^j w_\sigma(\hat{Q}(t)) = \mathbb{E}w_\sigma(\hat{Q}(t)) &\geq \max_{\sigma\in\mathcal{S}_j}\left\{w_\sigma\left(\hat{Q}(t)\right)\right\} - \frac{1}{\eta}\log|\mathcal{S}_j| \\ &\geq \max_{\sigma\in\mathcal{S}_j}\left\{w_\sigma\left(Q(t)\right) + w_\sigma\left(\delta(t)\right)\right\} - \frac{1}{\eta}\log|\mathcal{S}_j| \\ &\geq \max_{\sigma\in\mathcal{S}_j}\left\{w_\sigma\left(Q(t)\right)\right\} + \min_{\sigma\in\mathcal{S}_j}\left\{w_\sigma\left(\delta(t)\right)\right\} - \frac{1}{\eta}\log|\mathcal{S}_j| \\ &\geq \max_{\sigma\in\mathcal{S}_j}\left\{w_\sigma\left(Q(t)\right)\right\} - \delta_{\max}S_{\max}|\mathcal{I}| - \frac{1}{\eta}\log|\mathcal{S}_j| \end{aligned} \quad (4.36)$$

□

Lemma 4.6 and Lemma 4.7 introduce bounds on the increments of the square of the queue sizes and their conditional expectations, while Lemma 4.8 indicates an allowed reordering of terms.

Lemma 4.6. *There exists a constant $K_0 \geq 0$ such that our queue size process,* (4.5)*, obeys the bound*

$$\frac{1}{2}Q_i(t+1)^2 - \frac{1}{2}Q_i(t)^2 \leq Q_i(t)\left(A_i(t) - S_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right) + K_0. \tag{4.37}$$

Proof. Firstly, the following bound holds for the queue size process, (4.5).

$$\begin{aligned} Q_i(t+1) &= Q_i(t) - S_i(t) \wedge Q_i(t) + A_i(t) + \sum_{i':i'i\in\mathcal{L}} [S_{i'}(t) \wedge Q_{i'}(t)]p_{i'i}(t) \\ &\leq Q_i(t) - S_i(t) \wedge Q_i(t) + A_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t) \\ &\leq \begin{cases} Q_i(t) - S_i(t) + A_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t), & \text{if } Q_i(t) \geq S_i(t), \\ A_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t), & \text{otherwise.} \end{cases} \end{aligned} \tag{4.38}$$

Let's consider the two cases above. Firstly, if $Q_i(t) < S_i(t)$ then, and according to the above bound, we have

$$\begin{aligned} \frac{1}{2}Q_i(t+1)^2 - \frac{1}{2}Q_i(t)^2 \leq& \frac{1}{2}Q_i(t+1)^2 \leq \frac{1}{2}\left(A_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right)^2 \\ \leq& \frac{1}{2}\left(a_{max} + S_{\max}(|\mathcal{I}|+1)\right)^2 \\ \leq& -Q_i(t)\left(S_i(t) - A_i(t) - \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right) && (4.39) \\ &+ S_{\max}\left(a_{max} + S_{\max}(|\mathcal{I}|+1)\right) && (4.40) \\ &+ \frac{1}{2}\left(a_{max} + S_{\max}(|\mathcal{I}|+1)\right)^2. \end{aligned}$$

In the final inequality, we use the fact that the term, (4.39), is bounded by the term (4.40).

Secondly, if $Q_i(t) \geq S_i(t)$ then, according to (4.38),

$$\begin{aligned}\frac{1}{2}Q_i(t+1)^2 - \frac{1}{2}Q_i(t)^2 \leq & -Q_i(t)\left(S_i(t) - A_i(t) - \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right) \\ & + \frac{1}{2}\left(S_i(t) - A_i(t) - \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right)^2 \\ \leq & -Q_i(t)\left(S_i(t) - A_i(t) - \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right) \\ & + \frac{1}{2}\left(a_{\max} + (|\mathcal{I}|+1)S_{\max}\right)^2.\end{aligned}$$

Thus defining

$$K_0 = \frac{1}{2}\left(a_{\max} + (|\mathcal{I}|+1)S_{\max}\right)^2 + S_{\max}\left(a_{max} + S_{\max}(|\mathcal{I}|+1)\right),$$

we see that in both cases, above, we have the required bound

$$\frac{1}{2}Q_i(t+1)^2 - \frac{1}{2}Q_i(t)^2 \leq Q_i(t)\left(A_i(t) - S_i(t) + \sum_{i':i'i\in\mathcal{L}} S_{i'}(t)p_{i'i}(t)\right) + K_0.$$

□

Lemma 4.7. *There exists a constant $K_1 > 0$ such that the following equality holds*

$$\begin{aligned}&\mathbb{E}\left[Q_i(t)S_{i'}(t)p_{i'i}(t)\middle|Q(t),\delta(t)\right] \\ \leq &\mathbb{E}\Big[Q_i(t)\sum_{\sigma\in\mathcal{S}_j(i')}\mu_{i'}\sigma_{i'}(t)P_\sigma^{j(i')}\bar{q}_{i'i}(t)\Big|Q(t),\delta(t)\Big] + K_1.\end{aligned} \tag{4.41}$$

Proof. First let us suppose that the queue has been empty over the last k time steps. Then, since a bounded number of cars arrive at the queue per traffic cycle, the queue size $Q_i(t)$ must be less than $(a_{\max} + \sigma_{max}|\mathcal{L}|)$. Clearly the above bound holds for any, $Q_i(t) \leq K_1 = S_{\max}(a_{\max} + \sigma_{max}|\mathcal{L}|)$. Now lets suppose $Q_i(t) \geq K_1$, we can take $Q_i(t)$

out the conditional expectation because it is known

$$\mathbb{E}[Q_i(t)S_{i'}(t)p_{i'i}(t)|Q(t),\delta(t)] = Q_i(t)\mathbb{E}[S_{i'}(t)p_{i'i}(t)|Q(t),\delta(t)].$$

The proportion of traffic $p(t)$ is independent of $S(t)$ and of $Q(t)$. So the expectation of $p_{i'i}(t)$ given $Q(t)$ (and $S(t)$) is its mean $\bar{p}_{ii'}$. So

$$\mathbb{E}[S_{i'}(t)p_{i'i}(t)|Q(t),\delta(t)]=\mathbb{E}[S_{i'}(t)\bar{p}_{i'i}|Q(t),\delta(t)]=\bar{p}_{i'i}\mathbb{E}[S_{i'}(t)|Q(t),\delta(t)].$$

Also (4.2) implies

$$\mathbb{E}[S_{i'}(t)|Q(t),\delta(t)] = \sum_{\sigma\in\mathcal{S}_{j(i')}} \mu_{i'}\sigma_{i'}P_\sigma^{j(i')}.$$

Also since $q_{i'i}(t)$ is an unbiased estimate of $p_{i'i}$ at time t, and independent of $Q(t)$ by assumption,

$$\bar{p}_{i'i} = \mathbb{E}[q_{i'i}(t)|Q(t)].$$

Substituting this all back in, we have

$$\begin{aligned}\mathbb{E}[Q_i(t)S_{i'}(t)p_{i'i}(t)|Q(t),\delta(t)]&= Q_i(t)\bar{p}_{i'i}\mathbb{E}[S_{i'}(t)|Q(t),\delta(t)]\\ &= Q_i(t)\mathbb{E}[q_{i'i}(t)|Q(t),\delta(t)]\sum_{\sigma\in\mathcal{S}_{j(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\\ &= \mathbb{E}\Big[Q_i(t)\sum_{\sigma\in\mathcal{S}_j(i')}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\Big|Q(t),\delta(t)\Big],\end{aligned}$$

thus the above inequality also holds in the case $Q_i(t) > K_1$ as required.□

Lemma 4.8. *The following equality holds for each measured queue size vector,*

$$\begin{aligned}&\sum_{i\in\mathcal{I}}\hat{Q}_i(t)\left(\sum_{\sigma\in\mathcal{S}_{j(i)}}\mu_i\sigma_iP_\sigma^{j(i)}(t) - \sum_{i':(i',i)\in\mathcal{L}}\sum_{\sigma\in\mathcal{S}_{j'(i')}}\mu_{i'}\sigma_{i'}P_\sigma^{j(i')}\bar{q}_{i'i}(t)\right)\\ &=\sum_{k\in\mathcal{J}}\sum_{\sigma\in\mathcal{S}_k}P_\sigma^k(t)\sum_{i\in k}\mu_i\sigma_i\left(\hat{Q}_i(t) - \sum_{i':ii'\in\mathcal{L}}\hat{Q}_{i'}(t)\bar{q}_{ii'}(t)\right).\end{aligned}$$

Proof. Although the following set of equalities is some what lengthy, the premise is fairly simple. We want to change to order of summation so that we first sum over junctions $\mathcal{J}$ instead of first summing over in-roads $\mathcal{I}$. These manipulations are as follows

$$\begin{aligned}
&\sum_{i\in\mathcal{I}} \hat{Q}_i(t) \left(\sum_{\sigma\in\mathcal{S}_{j(i)}} \mu_i \sigma_i P_\sigma^{j(i)}(t) - \sum_{i':(i',i)\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j'(i')}} \mu_{i'}\sigma_{i'} P_\sigma^{j(i')} \bar{q}_{i'i}(t) \right) \\
=&\sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} \hat{Q}_i(t)\mu_i\sigma_i P_\sigma^k(t) \mathbb{I}[i\in k] \\
-&\sum_{i\in\mathcal{I}} \sum_{i':i'i\in\mathcal{L}} \sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} \hat{Q}_i(t)\mu_{i'}\sigma_{i'} P_\sigma^k \bar{q}_{i'i}(t) \mathbb{I}[i'\in k] \\
=&\sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} P_\sigma^k(t) \sum_{i\in\mathcal{I}} \hat{Q}_i(t)\mu_i\sigma_i \mathbb{I}[i\in k] \\
-&\sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} P_\sigma^k \sum_{i\in\mathcal{I}} \sum_{i':i'i\in\mathcal{L}} \hat{Q}_i(t)\mu_{i'}\sigma_{i'}\bar{q}_{i'i}(t) \mathbb{I}[i'\in k] \\
=&\sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} P_\sigma^k(t) \left(\sum_{i\in k} \hat{Q}_i(t)\mu_i\sigma_i - \sum_{i'\in\mathcal{I}} \sum_{i:i'i\in\mathcal{L}} \hat{Q}_i(t)\mu_{i'}\sigma_{i'}\bar{q}_{i'i}(t)\mathbb{I}[i'\in k] \right) \\
=&\sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} P_\sigma^k(t) \sum_{i\in k} \mu_i\sigma_i \left(\hat{Q}_i(t) - \sum_{i':ii'\in\mathcal{L}} \hat{Q}_{i'}(t)\bar{q}_{ii'}(t) \right).
\end{aligned}$$

In the first equality above, we expand brackets. In the second equality, we reorder the summation so we first sum over junctions and then over schedules. In the third and fourth equality, we collect together terms for each in-road. □

Our last lemma gives a bound on an optimization problem to be used later.

Lemma 4.9. *If* $a + \epsilon\mathbf{1} \in \mathcal{A}$ *then*

$$\epsilon < \left(1 - \frac{L}{T}\right) \min_{u\geq 0:\sum_{i\in\mathcal{I}} u_i = 1} \left(\sum_{j\in\mathcal{J}} \max_{\sigma\in\mathcal{S}_j} \{w_\sigma(u)\} - \sum_{i\in\mathcal{I}} u_i \bar{a}_i(t) \right).$$

Proof. By definition $a + \epsilon\mathbf{1} \in \mathcal{A}$ when there exists a positive vector $\rho = (\rho_\sigma^j : \sigma \in \mathcal{S}_j, j \in \mathcal{J})$ and a positive vector $s = (s_i : i \in \mathcal{I})$ satisfying

the constraints

$$a_i + \epsilon + \sum_{i':i'i\in\mathcal{L}} s_{i'}\bar{p}_{i'i} < s_i, \tag{4.42}$$

$$\sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \leq 1 - \frac{L}{T}, \qquad s_i \leq \sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \mu_i \sigma_i \tag{4.43}$$

for each $j \in \mathcal{J}$ and $i \in j$. We can express (4.42) more concisely in vector form as $a + \epsilon\mathbf{1} < s^\mathsf{T}(I - \bar{p})$. Notice the inverse of $(I - \bar{p})$ is the positive matrix $(I - \bar{p})^{-1} = I + \bar{p} + \bar{p}^2 + ...$ Thus we can equivalently express condition (4.42) as

$$(a + \epsilon\mathbf{1})^\mathsf{T}(I - \bar{p}) < s^\mathsf{T}.$$

We can now observe that if we replace s as above with

$$\tilde{s}_i = \sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \mu_i \sigma_i$$

then equations (4.42-4.43) must hold. In other words there exists a $\rho = (\rho_\sigma^j : \sigma \in \mathcal{S}_j, j \in \mathcal{J})$ such that

$$a_i + \epsilon + \sum_{i':i'i\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j(i')}} \rho_\sigma^{j(i')} \mu_{i'}\sigma_{i'}\bar{p}_{i'i} < \sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \mu_i \sigma_i, \quad \text{for } i \in j, \text{ with } j \in \mathcal{J}, \tag{4.44}$$

$$\sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \leq 1 - \frac{L}{T}, \qquad \text{for } j \in \mathcal{J}. \tag{4.45}$$

We now focus on the inequality (4.44). Rearranging it, the above holds when there exists ρ such that

$$\epsilon < \min_{j\in\mathcal{J}, i\in j} \left\{ \sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \mu_i\sigma_i - \sum_{i':i'i\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j(i')}} \rho_\sigma^{j(i')} \mu_{i'}\sigma_{i'}\bar{p}_{i'i} - a_i \right\}, \tag{4.46}$$

$$\sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \leq 1 - \frac{L}{T}, \qquad j \in \mathcal{J}. \tag{4.47}$$

The above statement can only hold if it also holds when we maximize

over ρ thus the following must hold

$$\epsilon < \max_{\rho:\forall j\in\mathcal{J},\sum_{\sigma\in\mathcal{S}_j}\rho_\sigma^j\leq 1-L/T} \min_{j\in\mathcal{J},i\in j} \left\{ \sum_{\sigma\in\mathcal{S}_j} \rho_\sigma^j \mu_i\sigma_i - \sum_{i':i'i\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j(i')}} \rho_\sigma^{j(i')} \mu_{i'}\sigma_{i'}\bar{p}_{i'i} - a_i \right\}. \tag{4.48}$$

Since the minimum of a finite set in $\mathbb{R}$ is equal to the minimum of the convex combinations of elements in that set,

$$\min_{i\in\mathcal{I}} \left\{ \sum_{\sigma\in\mathcal{S}_{j(i)}} \rho_\sigma^{j(i)} \mu_i\sigma_i - \sum_{i':i'i\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j(i')}} \rho_\sigma^{j(i')} \mu_{i'}\sigma_{i'}\bar{p}_{i'i} - a_i \right\} =$$

$$\min_{u\geq 0:\sum_{i\in\mathcal{I}}u_i=1} \left\{ \sum_{i\in\mathcal{I}} \sum_{\sigma\in\mathcal{S}_{j(i)}} \rho_\sigma^{j(i)} \mu_i\sigma_i u_i - \sum_{i\in\mathcal{I}} \sum_{i':i'i\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j(i')}} \rho_\sigma^{j(i')} \mu_{i'}\sigma_{i'}\bar{p}_{i'i}u_i - \sum_{i\in\in\mathcal{I}} a_i u_i \right\}.$$

Next by exactly the same argument used to prove Lemma 4.8, we have that

$$\begin{aligned} &\sum_{i\in\mathcal{I}} \sum_{\sigma\in\mathcal{S}_{j(i)}} \rho_\sigma^{j(i)} \mu_i\sigma_i u_i - \sum_{i\in\mathcal{I}} \sum_{i':i'i\in\mathcal{L}} \sum_{\sigma\in\mathcal{S}_{j(i')}} \rho_\sigma^{j(i')} \mu_{i'}\sigma_{i'}\bar{p}_{i'i}u_i - \sum_{i\in\mathcal{I}} a_i u_i \\ &= \sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} \rho_\sigma^k \sum_{i\in k} \mu_i\sigma_i \left(u_i - \sum_{i':ii'\in\mathcal{L}} u_{i'}\bar{p}_{ii'} \right) - \sum_{i\in\mathcal{I}} a_i u_i. \end{aligned}$$

Substituting this equality into (4.49), we have that (4.48) reads as

$$\epsilon < \max_{\rho:\forall j\in\mathcal{J},\sum_{\sigma\in\mathcal{S}_j}\rho_\sigma^j\leq 1-L/T} \min_{u\geq 0:\sum_{i\in\mathcal{I}}u_i=1} \left\{ \sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} \rho_\sigma^k \sum_{i\in k} \mu_i\sigma_i \left(u_i - \sum_{i':ii'\in\mathcal{L}} u_{i'}\bar{p}_{ii'} \right) - \sum_{i\in\mathcal{I}} a_i u_i \right\}.$$

Finally for any function $f(u,\rho)$, it holds that $\min_u \max_\rho f(u,\rho) \geq \max_\rho \min_u f(u,\rho)$. Thus if $\bar{a} + \epsilon\mathbf{1} \in \mathcal{A}$ then it must be true that

$$\epsilon < \min_{u\geq 0:\sum_{i\in\mathcal{I}} u_i = 1} \max_{\rho:\forall j\in\mathcal{J},\sum_{\sigma\in\mathcal{S}_j}\rho_\sigma^j \leq 1-L/T} \left\{ \sum_{k\in\mathcal{J}} \sum_{\sigma\in\mathcal{S}_k} \rho_\sigma^k \right.$$
$$\left. \sum_{i\in k} \mu_i \sigma_i \left(u_i - \sum_{i':ii'\in\mathcal{L}} u_{i'} \bar{p}_{ii'} \right) - \sum_{i\in\mathcal{I}} a_i u_i \right\}.$$

Finally, we note that the above maximization over ρ must be achieved at a value where, for each $j \in \mathcal{J}$, $\rho_\sigma^j = 1 - \frac{L}{T}$ for some σ, so

$$\epsilon < \left(1 - \frac{L}{T}\right) \min_{u\geq 0:\sum_{i\in\mathcal{I}} u_i = 1} \left\{ \sum_{j\in\mathcal{J}} \max_{\sigma\in\mathcal{S}_j} \left\{ \sum_{i\in j} \mu_i \sigma_i \right.\right.$$
$$\left.\left. \left(u_i - \sum_{i':ii'\in\mathcal{L}} u_{i'} \bar{p}_{ii'} \right) \right\} - \sum_{i\in\mathcal{I}} a_i u_i \right\}.$$

□

4.5 Conclusion

This chapter discussed a novel decentralised signal control policy based on the so-called backpressure algorithm, that does not require any *a priori* knowledge of the traffic demand and only needs information (i.e. queue size) that is local to the intersection. In contrast to other existing backpressure-based policies in which phases can form an erratic and unpredictable order resulting in potential unsafe operation, our scheme allocates a positive amount of time to each phase within the cycle, thus repeating them in a cyclic manner. Furthermore, no knowledge of the local turn ratios (turning fractions) is required in our control strategy, instead any unbiased estimator of the turning fractions can be utilised in the proposed scheme. We have formally proved the stability results of the proposed signal control policy even though the controllers are reacting based only on local information and demand in a distributed manner. The stability results indicate that our policy is stable for the largest possible set of arrival rates (or demand) that will provide sufficient

throughput even in a congested network.

Using simulation, we compared the performance of our cyclic phase backpressure policy against other well-known schemes in terms of network throughput and congestion level using both small and large network topology with fixed routings. The results showed that our cyclic phase backpressure policy tends to outperform other distributed policies both in terms of throughput and congestion. Although the performance of each policy varies widely depending on the parameter setting such as cycle length or decision frequency, under the optimal setting among the cases studied, the backpressure shemes with cyclic and non-cyclic operations had better throughput in comparison with the other policies.

Our study still left some open questions. Examples include the incentive to build up backlogs to increase the pressure, fairness for drivers and the possibility of better performance with varying cycle times. Chapter 5 introduces a control scheme that aims at answering some of these questions, whereas others are left for further research.

Chapter 5

Proportional fair control for traffic light networks

Outline. This chapter discusses another decentralised signal control policy for urban traffic light systems. Once again we consider the basic model described in Chapter 3, using the notations given in Section 3.2. The proposed scheme is similar in many ways to the one introduced in Chapter 4, there are some key attributes though that are different. The basic idea once again comes from schdeuling in communication networks, however it is based on the notion of *proportional fairness*. This model considers the possibility of changing cycle times, albeit on a different time scale then the decision to allocate proportions of said cycle time to the service phases. We propose an optimal scaling for the cycle times, whilst discussing its consequences on the stability region of the policy. Namely, we provide a formal proof that our proposed scheme stabilises the network for any feasible traffic demand. We have also conducted simulations to validate the chosen cycle lengths. Furthermore we numerically compare the performance of our policy to the a fixed cycle proportional fair scheme, that we have studied in Section 4.3, showing promising results.

The chapter is organised as follows. First, in Section 5.1 we introduce some properties of the model that further detail the seup given by the basic model in Chapter 3. Secondly, in Section 5.2 we describe the control scheme itself by providing the algorithm for green time allocation on a cycle-to-cycle basis and a discussion on the optimal cycle length. We also state our main results for stability. Section 5.3 presents a simulation study that aims at finding the optimal setting for the control parameters and also compares the performance of our model with previous results. We give detailed proofs to our theoretical results in Section 5.4, whilst Section 5.5 provides some supplementary results that aid the discussion and the proofs of Section 5.2 and Section 5.4. Finally, Section 5.6 concludes the chapter.

5.1 Queue size processes

The model that we discuss in this chapter is a more detailed version of the one presented in Section 3.2, thus the notation and assumptions introduced there apply here as well. One of the important features is that we let the vehicles follow a predefined, albeit unknown *route*. Let us denote the set of all possible routes by $\mathcal{R}$, and the routes themselves by $r \in \mathcal{R}$. A route is the sequence of in-roads that a vehicle visits as it navigates through the network. We will denote by $i \in r$ if route r goes through in-road i. If i is the first in-road on route r, we will denote that by $i = i_0^r$, and if i is the last in-road on route r, we will denote that by $i = i_l^r$. We will use the notations i_-^r and i_+^r for the in-roads that are, respectively, preceding and following in-road i on route r.

The cars in a traffic light network form a queueing network, where the vehicles waiting at every in-road form the queues. As discussed before we assume that the controlller of each junction has knowledge of the queue sizes present at each of its in-roads. We introduce the route-wise queue sizes, $X_{ir}(t)$, i.e. the number of vehicles queueing at in-road i which are following route r. Unlike $Q_i(t)$, the queue sizes $X_{ir}(t)$ are unknown to the controller. We denote the route-wise queue size vector by $X(t) = (X_{ir}(t))_{i\in\mathcal{I},r\in\mathcal{R}}$. By definition

$$Q_i(t) = \sum_{r:i\in r} X_{ir}(t). \tag{5.1}$$

We work with the assumption that at every in-road the cars following different routes are distributed homogeneously in the queue.

We define $\{A_{ir}(t)\}_{i\in\mathcal{I},r\in\mathcal{R}}$ as the route-wise arrival processes and similarly $\{D_{ir}(t)\}_{i\in\mathcal{I},r\in\mathcal{R}}$ as the route-wise departure processes. Thus the queue sizes develop as follows,

$$X_{ir}(t) = X_{ir}(0) + A_{ir}(t) - D_{ir}(t). \tag{5.2}$$

We note, that the route-wise arrivals described by $A_{ir}(t)$ can be external arrivals, if $i = i_0^r$, or internal arrivals if $i \neq i_0^r$. In the latter case they equal the departures of the previous in-road along route r. The external arrivals follow the process $A_r(t)$ for each route $r \in \mathcal{R}$. We assume, for simplicity, that these are independent Poisson processes and use the

vector notation $(a_r)_{r\in\mathcal{R}}$ for the arrival rates. Thus, by definition

$$A_{ir}(t) = \begin{cases} A_r(t) & \text{if } i = i_0^r, \\ D_{i^-_r}(t) & \text{if } i \neq i_0^r. \end{cases}$$

It follows that the long run arrival rate for in-road i is given by

$$a_i = \sum_{r:i\in r} a_r.$$

The departures $D_{ir}(t)$ are positive counting processes. They depend on both the properties of the junctions and the allocated green times. To model these effects we define the random variables S_{ir}^n as the number of cars served at in-road i during its junction's n^{th} cycle. Thus we have

$$D_{ir}(t_n^{j(i)}) = \sum_{m=1}^{n} S_{ir}^m \tag{5.3}$$

for each $r \in \mathcal{R}$, $i \in \mathcal{I}$, and $n \in \mathbb{N}$.

We assume that for every in-road there exists a maximum rate with which a steady stream of vehicles can be served, which is determined by the speed limits on the road segments. We will denote the vector of maximum rates by $\mu^{\max} = (\mu_i^{\max})_{i\in\mathcal{I}}$. When y_i is allocated to in-road i, this puts the following bound on the received service

$$\mu_i^{\max} y_i E_n^{j(i)} \geq \sum_{r:i\in r} S_{ir}^n,$$

for each $i \in \mathcal{I}$ and $n \in \mathbb{N}$. Further, we introduce the service function $s(\cdot) = (s_i(\cdot))_{i\in\mathcal{I}} : \mathbb{R}_+ \mapsto \mathbb{R}_+^{\mathcal{I}}$, where $s(t)$ represents the expected number of cars that can be served from each in-road whilst receiving a green light for t units of time. We assume $s(\cdot)$ to be continuous, almost everywhere twice differentiable, increasing and to have the following asymptotic behaviour

$$s_i(t) \sim \mu_i^{\max} t, \tag{5.4}$$

for each in-road i. This represents the idea that after an initial set-up phase, needed for the vehicles in the queue to speed up, the cars can move without interruption. With our assumption on the homogeneous

distribution of vehicles taking different routes among queues, we have, when $Q_i(t_n^{j(i)}) > 0$,

$$\mathbb{E}[S_{ir}^n | X(t_{n-1}^{j(i)}), Q^j(t_{n-1}^{j(i)})] = \frac{X_{ir}(t_{n-1}^{j(i)})}{Q_i(t_{n-1}^{j(i)})} s_i(y_i E_n^{j(i)}) \tag{5.5}$$

for $r \in \mathcal{R}, i \in \mathcal{I}$ and $n \in \mathbb{N}$. We define the average rate at which vehicles are served from each route-wise queue as

$$\mu_{ir}(t) = \frac{\mathbb{E}[S_{ir}^n | X(t_{n-1}^{j(i)}), Q^j(t_{n-1}^{j(i)})]}{T_n^{j(i)}}, \tag{5.6}$$

when $t_{n-1}^{j(i)} \leq t < t_n^{j(i)}$, thus $\mu_{ir}(t)$ is a piecewise constant function. We also let $\mu_i(t)$ be the rate at which cars leave the in-roads. From (5.5) and (5.6) we have

$$\mu_i(t) = \sum_{r: i \in r} \mu_{ir}(t) = \frac{s_i(y_i E_n^{j(i)})}{T_n^{j(i)}}, \tag{5.7}$$

for $t_{n-1}^{j(i)} \leq t < t_n^{j(i)}$. To ease our notation we will refer to $\mu_{ir}(t_{n-1}^{j(i)})$ as μ_{ir}^n. When we wish to make dependence of y and $T = T_n^{j(i)}$ explicit, we will write $\mu_i(y, T)$ for expression (5.7). We assume the actual number of cars leaving is close to its expectation in that there exists a constant κ such that

$$\mathbb{E}\left[\left(S_{ir}^n - T_n^{j(i)} \mu_{ir}^n\right)^2\right] \leq \kappa T_n^{j(i)} \tag{5.8}$$

for $r \in \mathcal{R}, i \in \mathcal{I}$ and $n \in \mathbb{N}$. This is a property which holds for numerous kinds of stochastic processes such as Poisson processes and renewal processes. Though the set of departures do not belong to these types due to dependence in the network structure, the basic properties of its evolution are similar on the level represented in (5.8).

5.2 Scheduling

The task of the traffic controller is to determine the cycle lengths and to allocate proportions of them to the service phases and switching. In this

section we first present an algorithm which schedules the green times in a fair way. Secondly we discuss the network's capacity given the lengths of cycles. Thirdly we investigate two polling models, which allocate service times in a similar manner to the proportional fair scheme, in order to determine which choice of cycle lengths would minimise the average waiting time in the system. Finally we summarise our suggested policy.

5.2.1 Proportionally fair allocation

We aim to allocate proportions of the cycle length to the service phases at each junction in a way that maximises throughput whereas maintaining service for all in-roads with vehicles on them. A well-regarded way to do so is using a proportional fair scheme. To fully utilise the resources and reach optimal throughput we have to maximise $s_i(t)$ for all $i \in \mathcal{I}$, however we have to consider the constraints

$$\sum_{\sigma \in \mathcal{S}_j} P_\sigma^j(t) = 1, \qquad \text{where} \qquad P_\sigma^j(t) \geq 0, \tag{5.9}$$

and

$$y_i(P) = \sum_{\sigma \in \Sigma_j} \sigma_i P_\sigma^j. \tag{5.10}$$

For the $(n+1)$th traffic cycle, the proportionally fair schedule is given by the solution to the following optimisation problem.

$$\begin{aligned} \text{maximise} \quad & \sum_{i \in j} Q_i(t_n^j) \log(s_i(E_{n+1}^j y_i)), \\ \text{subject to} \quad & \sum_{\sigma \in \Sigma_i} \sigma_i P_\sigma^j = y_i, \quad \forall i \in j, \\ & \sum_{\sigma \in \Sigma_j} P_\sigma^j = 1, \\ \text{over} \quad & y \geq 0, \;\; P_\sigma^j \geq 0, \; \forall \sigma \in \Sigma^j. \end{aligned} \tag{5.11}$$

This needs to be considered for each junction $j \in \mathcal{J}$. Solving the problem only requires local information, i.e. the knowledge of queue sizes present at the in-roads of said junction. An advantage of using such a scheme is

that the optimisation problem can be solved separately for each junction, thus requiring significantly less effort.

We denote the solution to the optimisation problem respectively by $P^* = (P^*_\sigma)_{\sigma\in\Sigma_j, j\in\mathcal{J}}$, $y^* = (y_i(P^*))_{i\in\mathcal{I}}$, $s^* = (s_i(y^*, T))_{i\in\mathcal{I}}$ and $\mu^* = (\mu_i(y^*, T))_{i\in\mathcal{I}}$ for further purposes. In general (5.11) does not have an explicit solution, although if s is given, it can be solved numerically. Since we need (5.11) to have a single optimal solution, $\log(s_i(E_j y_i(p)))$ has to be concave. This imposes the condition

$$(s_i'(t))^2 > s_i(t)s_i''(t) \quad \forall t > 0,\ \forall i \in \mathcal{I} \tag{5.12}$$

on the service function, which can be derived from the necessary condition of concavity,

$$(\log(s_i(t)))'' = \left(\frac{s_i'(t)}{s_i(t)}\right)' = -\left(\frac{s_i'(t)}{s_i(t)}\right)^2 + \frac{s_i''(t)}{s_i(t)} < 0, \quad \forall i \in \mathcal{I}$$

and the fact that s is positive and increasing. Any subexponential function is sufficient for (5.12) such as polynomial, piecewise linear or similar functions. In special cases (5.11) can even be solved explicitly, for example if $s_i(t) = vt\ \forall i \in \mathcal{I}$, i.e. the received service is a linear function of time with the same rate for all in-roads, and all in-roads are served during exactly one service phase, i.e. $|\Sigma_i| = 1\ \forall i \in \mathcal{I}$. Then the optimal schedule is given by

$$P^*_\sigma = \frac{\sum_{i\in\sigma} Q_i}{\sum_{i\in j} Q_i} \quad \forall \sigma \in \Sigma_j, \tag{5.13}$$

which demonstrates the proportional nature of the proposed algorithm. We note that the above scheme, (5.13), coincides with the P0 policy which was previously studied in the context of road traffic congestion, see [133].

5.2.2 Network capacity

The maximal throughput is constrained by the physical parameters of the network, which we represent by the vector $\mu^{\max}$. Thus the capacity

can be defined as

$$\mathcal{C} = \{(y_i \mu_i^{\max})_{i \in \mathcal{I}} : y \in \mathcal{Y}\},$$

which is convex due to $\mu^{\max}$ being constant and the convexity of $\mathcal{Y}$. The maximal service that could be allocated to the in-roads is given by the case when they receive uninterrupted green time, thus they form the extreme points of $\mathcal{C}$. These rates cannot be reached however, as we require switching in every cycle, which alone interrupts service and also invokes a slower setup phase in realistic service. To represent this we define the set of admissible service rates as

$$\mathcal{A}(T) = \{(\mu_i(y, T))_{i \in \mathcal{I}} : y \in \mathcal{Y}\},$$

which is clearly a subset of the capacity set. It depends on the cycle lengths since by its definition in (5.7), the vector of service rates $\mu = (\mu_i)_{i \in \mathcal{I}}$ is cycle-length dependent. This dependency is present not only due to switching times, but could also show in the service function as shown in the example in Section 5.5.2. The problem however can be solved by a sufficient choice of cycle lengths as stated in the following proposition. In the statement we use $\mathcal{M}^\circ$ to denote the interior of a set $\mathcal{M}$.

Proposition 5.1. *In the limit as $T \to \infty$, the set of admissible service rates $\mathcal{A}(T)$ reaches the capacity set $\mathcal{C}$, i.e. $\forall a \in \mathcal{C}^\circ$ there exists a vector of cycle lengths τ, such that if $T_j > \tau_j \ \ \forall j \in \mathcal{J}$, then*

$$a \in \mathcal{A}^\circ(T).$$

A proof of Proposition 5.1 can be found in Section 5.4.1.

We have not addressed the possibility of eliminating phases, and the extra switching they require, as another possible solution. However, from the traffic's point of view, this is a similar solution to increasing the cycle lengths since a longer cycle can be viewed as consecutive shorter cycles, which are all serving just one phase.

When choosing the cycle lengths we also have to consider the fact that longer cycles and thus longer service phases produce longer waiting periods for the vehicles on the in-roads receiving red light. There is clearly a trade-off between the network's stability and the average time vehicles spend in the system depending on the cycle lengths. In Section 5.2.3 we

discuss how long the cycles should be to minimise the average waiting times.

5.2.3 Optimal cycle length

In this section we investigate a polling model which represents a single junction. This provides an example where the optimal scaling of cycle lengths can be determined. The schedule in the model does not necessarily follow strictly the proportional fair allocation given by (5.11), instead it allocates service times based on the average queue lengths. However, due to similarities in the allocation method, the model gives us an opportunity to develop a rule for optimal cycle length setting.

Consider a single junction with N competing in-roads. In each schedule, only one in-road can be served. Let each in-road have Poisson arrivals with rates $\lambda_1, \lambda_2, \ldots, \lambda_N$ and exponentially distributed random green times $G_1, G_2, \ldots, G_N$. Let vehicles leave with rate $\mu_1, \mu_2, \ldots, \mu_N$ once receiving service, and define $\rho_i = \lambda_i/\mu_i$. In order to have similar behaviour to the proportionally fair policy the rates $\nu_1, \nu_2, \ldots, \nu_N$ of the green times are chosen to fulfill the following equation.

$$\frac{\frac{1}{\nu_i}}{\sum_{i=1}^{N} \frac{1}{\nu_i}} = \frac{\bar{Q}_i}{\sum_{i=1}^{N} \bar{Q}_i}, \tag{5.14}$$

where $\bar{Q}_1, \bar{Q}_2, \ldots, \bar{Q}_N$ are the respective average queue sizes. Let us denote the expected cycle length, i.e. the sum of the green times and switching times by τ. Thus

$$\tau = \mathbb{E}T = \mathbb{E}\left[\sum_{i=1}^{N} G_i + NT_{\text{switch}}\right] = \sum_{i=1}^{N} \frac{1}{\nu_i} + NT_{\text{switch}}. \tag{5.15}$$

To gain insight on the relation between the average queue lengths and the cycle time let us deduct a system of equations using Little's Law and PASTA (Poisson Arrivals See Time Averages), cf. [155]. If we denote the average waiting times by $\bar{W}_1, \bar{W}_2, \ldots, \bar{W}_N$, then (by Little's Law) we have the following equations for $i = 1, 2, \ldots, N$,

$$\bar{Q}_i = \lambda_i \bar{W}_i,$$

and, using PASTA in combination with the FCFS discipline in each queue,

$$\begin{aligned}\bar{W}_i =& \frac{1}{\mu_i}\left(1+\bar{Q}_i\right)\left(1+\nu_i\left(\sum_{j\neq i}\frac{1}{\nu_j}+NT_{\text{switch}}\right)\right)\\ &+\frac{T_{\text{switch}}\left(NT_{\text{switch}}+\sum_{j\neq i}\frac{1}{\nu_j}\right)}{\sum_{j=1}^{N}\frac{1}{\nu_j}+NT_{\text{switch}}}+\frac{\sum_{j\neq i}\frac{1}{\nu_j}\left(\frac{1}{\nu_j}+T_{\text{switch}}\right)}{\sum_{j=1}^{N}\frac{1}{\nu_j}+NT_{\text{switch}}}.\end{aligned}$$

From these we can derive the following expressions for the expected queue lengths given as a function of expected cycle length,

$$\bar{Q}_i(\tau)=\frac{1}{1-\rho_i\nu_i\tau}\times\left[\rho_i\nu_i\tau+\lambda_i\cdot\frac{{T_{\text{switch}}}^2+\sum_{j\neq i}\left(\frac{1}{\nu_j}+T_{\text{switch}}\right)^2}{\tau}\right], \tag{5.16}$$

for $i=1,2,\ldots,N$. From (5.16) it is easily seen that for stability we need the cycle length to fulfill

$$\tau > \frac{NT_{\text{switch}}}{1-\sum_{i=1}^{N}\rho_i}. \tag{5.17}$$

One would expect that T should be 'fairly close' to this limit in order to keep the expected queue lengths and subsequently the waiting times low. This would be the case if both the arrivals and the service times were deterministic as shown in Section 5.5.2. In this model however an optimal cycle length cannot be determined in the general case, only if for all $i=1,2,\ldots,N$ we have $\lambda_i=\lambda$ and $\mu_i=\mu$ and thus $\nu_i=\nu$ and $\rho_i=\rho$, and even then the calculations do not yield an explicit formula, see Section 5.5.2 for the $N=2$ case. On the other hand in this case we can derive the proper scaling for the cycle lengths as the system is in its heavy traffic limit as presented by the following proposition.

Proposition 5.2. *When $\rho\to 1/N$ the cycle length minimising the average waiting times has the following asymptotic behaviour*

$$T^2\sim\sum_{i=1}^{N}\bar{Q}_i. \tag{5.18}$$

Proposition 5.2 is proven in Section 5.4.2.

As we show in Section 5.3 setting the cycle lengths on a scale according to this "square root rule" provides pleasant results even when the actual service time allocation is determined by the proportional fair scheduling scheme. Thus we are ready to formulate our proposed traffic light control policy.

5.2.4 Policy

Given the described optimisation in (5.11) and the results in Proposition 5.1 and 5.2, the proposed algorithm for traffic light setting can be summarized as follows.

- Form an unbiased estimate of the expected queue sizes, $\widetilde{Q}$.[1]
- For a given time period[2] a priori set up a sequence of cycle lengths for each junction according to

$$T^j_{n+1} = c_j \sqrt{\sum_{i \in j} \widetilde{Q}_i(t^j_n)}, \tag{5.19}$$

  where c_j is a control parameter determined by the traffic controller.
- At the beginning of each cycle at all the junctions allocate green times to each service phase according to the solution of (5.11) based on the queue sizes present.

One of our main results is proving maximum stability of this scheduling algorithm. Namely we are going to prove that the capacity set $\mathcal{C}$ coincides with the stability region of the the ensuing stochastic network, i.e. that for any set of arrival rates $a = (a_i)_{i \in \mathcal{I}} \in \mathcal{C}^\circ$ there is a set of allocations y, under which the stochastic system is positive recurrent. Informally, this means that any demand that does not exceed the physical parameters of the network can be satisfied when the service is allocated by our policy. The stability proof is following the fluid limit approach of [34]. First we determine the fluid limits of the processes described in Section 5.1.

[1] For a suggested estimation method see Section 5.5.1

[2] This can be a day or a few hours for example.

To prove positive recurrence it suffices to show stability of these fluid limits. We are going to formalise these steps in Section 5.2.5 and give a full proof in Section 5.4.

5.2.5 Fluid stability

We can associate a fluid model with the network if we introduce the terms

$$q(t) = (q_i(t))_{i\in\mathcal{I}},$$

and

$$x(t) = (x_{ir}(t))_{i\in\mathcal{I},r\in\mathcal{R}}$$

as the fluid limits of $Q(t)$ and $X(t)$ respectively. If we introduce the notation

$$x_{i^r_- r} = \begin{cases} x_{i'r} & \text{if } i' = i^r_-, \\ x_{0r} & \text{if } i = i^r_0, \end{cases}$$

and similarly

$$x_{i^r_+ r} = \begin{cases} x_{i'r} & \text{if } i' = i^r_+, \\ x_{lr} & \text{if } i = i^r_l, \end{cases}$$

and for every route $r \in \mathcal{R}$, we can define auxiliary variables q_0, q_l and μ^*_0, μ^*_l to have

$$x_{0r} = \frac{q_0}{\mu^*_0} \cdot a_r, \tag{5.20}$$

and

$$x_{lr} = \frac{q_l}{\mu^*_l} \cdot a_r. \tag{5.21}$$

The fluid limit of the system is then governed by the ODE

$$\frac{\mathrm{d}}{\mathrm{d}t} x_{ir}(t) = \frac{x_{i^r_- r}(t)}{q_{i^r_-}(t)} \mu^*_{i^r_-}(q) - \frac{x_{ir}(t)}{q_i(t)} \mu^*_i(q), \quad \text{if } q_i(t), q_{i^r_-}(t) > 0, \tag{5.22}$$

with

$$q_i(t) = \sum_{r: i\in r} x_{ir}(t). \tag{5.23}$$

We formalise this result as follows. Let $\{X^{(c)}_{ir}\}_{c\in\mathbb{N}}$ be a sequence of versions of our route-wise queue size processes, where $\|X_{ir}(0)\|_1 = c$. We

define

$$\bar{X}_{ir}(t) = \frac{X_{ir}^{(c)}(ct)}{c}. \tag{5.24}$$

Our next proposition formally states that the only possible limit for $\bar{X}_{ir}$ is given by x_{ir}.

Proposition 5.3. *The sequence of stochastic processes $\{\bar{X}^{(c)}\}_{c\in\mathbb{N}}$ are tight with respect to the topology of uniform convergence on compact time intervals. Moreover, any weakly convergent subsequence of $\{\bar{X}^{(c)}\}_{c\in\mathbb{N}}$ converges to a Lipschitz continuous process almost everywhere satisfying fluid equations* (5.22), (5.23).

A proof of Proposition 5.3 can be found in Section 5.4.3.

Now we are ready to formalise the statement considering the stability of the system. The result is given by the following theorem.

Theorem 5.1. *If the set of arrival rates are such that*

$$a \in \mathcal{C}^{\circ}, \tag{5.25}$$

then the fluid limit in (5.22) *is stable, i.e. there exists a time $\tau > 0$ such that for every fluid model $\{x(t)\}_{t\in\mathbb{R}_+}$ satisfying* (5.22) *with $\|x(0)\|_1 = 1$,*

$$x_{ir}(t) = 0,$$

for all $t \geq \tau$ for each $i \in r$, and $r \in \mathcal{R}$.

As discussed in Proposition 5.1, cycle lengths can be chosen for any interior point of the capacity set such that the admissible set of rates will also contain that point. We work with the assumption that such a T is in place, which combined with (5.25) means that

$$a \in \mathcal{A}^{\circ}(T). \tag{5.26}$$

This can be ensured by the traffic controller by choosing the control parameters c_j in (5.19) correctly. Thus satisfying (5.26) also shows the correct traffic setting to follow.

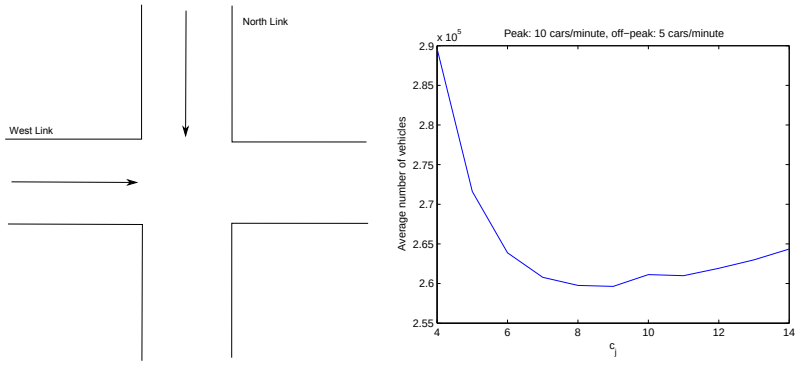

(a) Network topology. **(b)** Optimal input parameter c_j.

Figure 5.1. Single Intersection Network.

5.3 Simulation

In this section we investigate the optimal value of c_j, which was introduced in (5.19) and evaluate the performance of our proposed policy by comparing it with fixed cycle proportional policies.

5.3.1 Parameter validation

Based on our results in Proposition 5.2 and (5.19), a good estimation for the optimal value of c_j is given by

$$c_j = N \cdot \sqrt[2]{T_{\text{switch}}/\mu^{\max}}, \tag{5.27}$$

where N denotes the number of competing in-roads at junction j, T_{switch} denotes the fixed switching time and $\mu^{\max}$ denotes the maximal service rate of the in-roads.

The control parameter c_j plays a crucial role in determining the cycle lengths. A suboptimal value of c_j may shrink the network's capacity region as discussed in Section 5.2.2, which can result increasing congestion. In this section, we validate via simulation the optimal value of c_j as given in (5.27).

We have considered a single intersection network topology as shown in Figure 5.1a. Both the west and the east link were single lanes with

length of 2000 meters. The long link length ensured that the arriving vehicles were always able to enter the network, even in a congested period. The service rate when receiving green traffic light was capped at 20 cars per minute. We run the simulation with different values of c_j. For the demand, we considered the vehicles entering the network from the west and the north link, passing through the intersection before exiting the network. The arrivals followed a Poisson process with symmetric arrival rates. In each run we have simulated 10 hours of traffic including 10 peak periods and 10 off-peak periods alternately. The length for these periods was set as 30 minutes. The peak periods were considered to have 10 cars per minute arriving on average, whereas the arrival rate in off-peak period was 5 cars per minute.

The results are shown in Figure 5.1b. Note that the network parameters were $N = 2$, $T_{\text{switch}} = 6$, $\mu_{max} = 20$, hence according to (5.27) the estimated optimal value of c_j was $c_j = 8.5$. The figure plots the average number of vehicles present in the network throughout the whole run against different values of c_j. Since vehicles cannot disappear from the network through any other means, but exiting, a lower average number of vehicles can only occur due to shorter travel times. It can be observed that the optimal value for c_j is between 8 and 9 which is consistent with our estimate. For smaller values of c_j, the traffic flows are more frequently disrupted by switching, whereas at greater values of c_j, vehicles must wait longer for green traffic light. One can see from the results in 5.1b, that this indeed produces longer queues.

5.3.2 Performance study for varied cycle lengths

This section studies the performance of our proposed proportional fair policy by comparing it with the fixed cycle proportional policy in [4]. We have considered the Melbourne CBD network which is shown in Figure 5.2. It consists of 73 intersections and 266 links. Most of the roads are bi-directional except for Little Lonsdale Street, Little Bourke Street, Little Collins Street and Flinder Lane which only have a single lane mono-directional traffic. King Street and Russell Street are the biggest roads in this scenario, each is modeled as 3 lanes in each direction. Collins Street has one lane in each direction. All other roads have two lanes each direction. The link lengths are varied between 106 meters

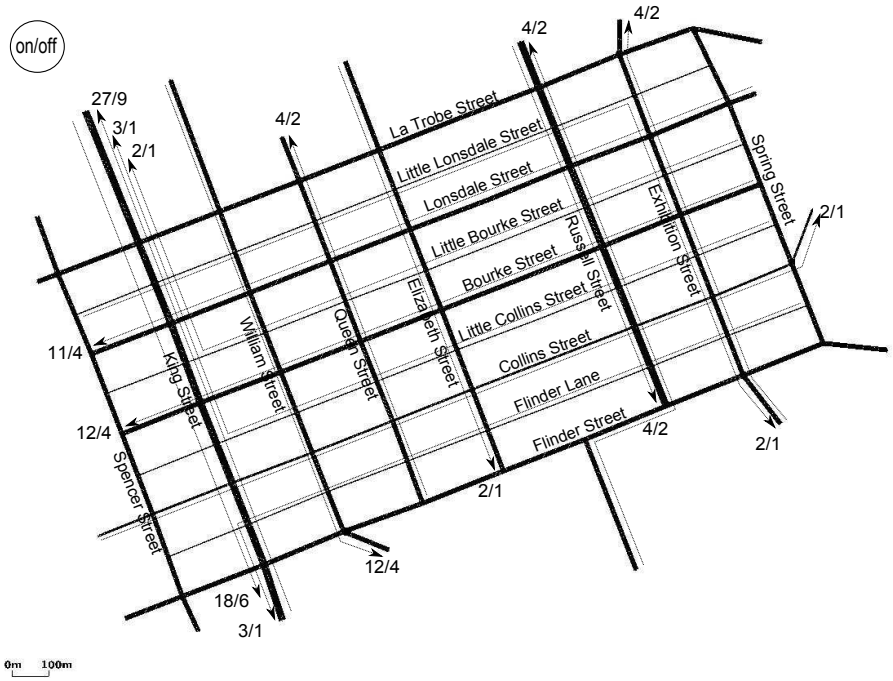

Figure 5.2. Large CBD network with demands.

for the vertical links and 214 meters and 447 meters for the horizontal links except for the ingress links at the edges. Each simulation run consisted of 4 hours including two peak periods and two off-peak periods alternately, each lasting 1 hour. The fix routes are indicated by the arrows in Figure 5.2 while the following numbers show the arrival rate in peak and off-peak periods.

We have run multiple simulations sequentially and the cycle lengths were pre-calculated before each run based on the ensemble average of queue lengths of previous runs and according to (5.19). Given the identical demands, the differences in queue evolution widely depended on the cycle lengths mentioned above and, in turn, the cycle lengths would adapt to the new queue evolution records. Naturally, we wanted to evaluate the proposed policy in convergence, that is when the updates of queue lengths no longer changed the cycle lengths. Particularly, we first ran the simulation 5 times with fixed 30 second cycle lengths to

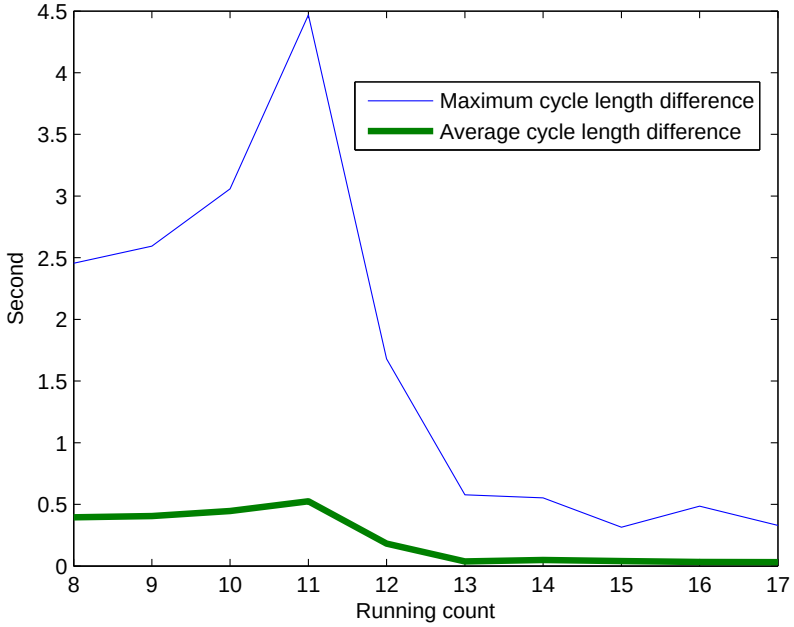

Figure 5.3. Convergence.

obtain the initial queue length records. Then from the 6th time on, the queue lengths were calculated according to (5.19) based on the ensemble average queue lengths of the 5 previous runs.

We estimated the convergence of our procedure for determining cycle lengths, see Figure 5.3. We evaluated the difference in cycle lengths at 4 fixed time points: 2500, 6000, 10000 and 12000 seconds respectively (see the vertical lines in Figure 5.4). At each time of these times we recorded the cycle lengths of each intersection. We then compared the cycle lengths of the current run with the average cycle lengths of the last 3 runs at each intersection and at each time point. The maximum difference and average difference are shown in Figure 5.3. Intuitively, the smaller values of those two quantities indicate better convergence. The results show strong convergence after the 13th run.

We compared the performance of our proposed policy with other two policies, namely fixed cycle proportional fair policies of cycle lengths with

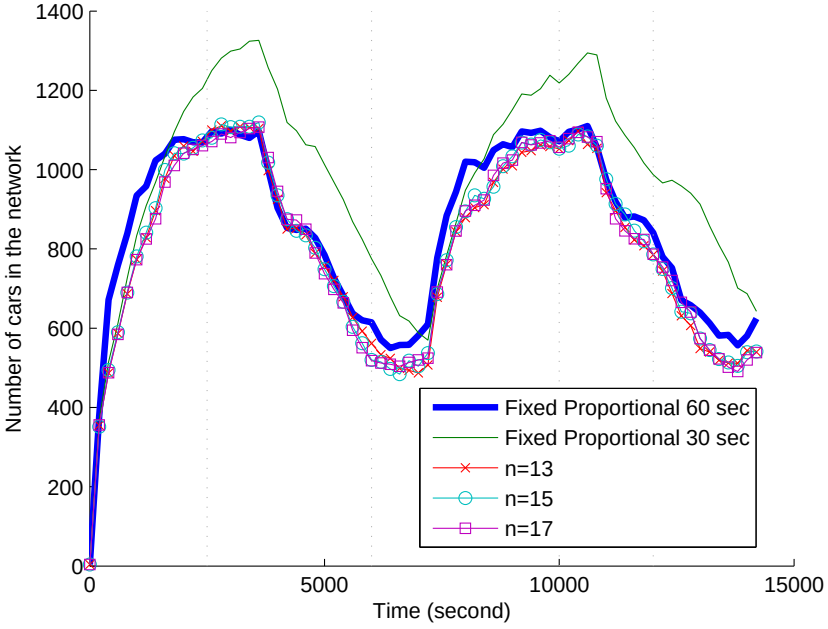

Figure 5.4. Total number vehicles in the network.

30 and 60 seconds. Note that 60 seconds was the optimal cycle length for the fixed cycle proportional policy as shown by simulation in [4]. The results are presented in Figure 5.4 in the form of number of vehicles in the network.

As shown in Figure 5.4, the proposed policy in convergence has significantly less congestion than the fixed cycle proportional fair policy with 30 second cycles and is slightly better than the fixed cycle proportional fair policy with 60 second cycles. This agrees with our theoretical results, since we aimed for minimising the queue lengths in the network.

Additionally, we have also repeated the whole simulation with a different a starting point, e.g. the fixed cycle proportional policy with 60 second cycles, and observed similar results in both convergence and performance, which indicates the robustness of our policy. Practically, the network is able to achieve its optimal cycle lengths by updating its cycle lengths based on historical queue length data regardless of the

network topology.

5.4 Proofs

In this section we formally prove our main theoretical results. Firstly, we prove Proposition 5.1, which claims that the set of admissible service rates reaches the capacity set in the limit as $T \to \infty$. Secondly, we prove Proposition 5.2 which justifies the square root rule given in (5.19). Then we prove convergence in the fluid limit as stated in Proposition 5.3. Finally, we end the section by providing a proof of stability according to Theorem 5.1.

5.4.1 Proof of Proposition 5.1

To prove Proposition 5.1, we have to look into the behaviour of s. By (5.4) it holds for any i, that

$$\theta_i(t) = \frac{\int_0^t (\mu_i^{\max} - s_i'(z))\mathrm{d}z}{t} \to 0 \quad \text{as } t \to \infty, \tag{5.28}$$

i.e. for any $\delta > 0$ there exists τ_i such that $\theta_i(t) < \delta$ if $t > \tau_i$. Now if we take the definition (5.7), we can derive the following bound

$$\begin{aligned}
\mu_i(T) =& \frac{s_i(y_i E^{j(i)})}{T^{j(i)}} \\
=& \frac{\int_0^{y_i T^{j(i)}} \mu_i^{\max}\mathrm{d}z - \int_0^{y_i T^{j(i)}} (\mu_i^{\max} - s_i'(z))\mathrm{d}z - \int_{y_i E^{j(i)}}^{y_i T^{j(i)}} s_i'(z)\mathrm{d}z}{T^{j(i)}} \\
\geq& y_i \mu_i^{\max} - \theta_i(T^{j(i)}) - y_i \mu_i^{\max} |\Sigma_j| \frac{T_{\text{switch}}}{T^{j(i)}}.
\end{aligned} \tag{5.29}$$

Now if we choose any $a \in \mathcal{C}^\circ$, then by definition of $\mathcal{C}$ there exists $\epsilon > 0$ such that

$$y_i \mu_i^{\max} - a_i > \epsilon \quad \forall i \in \mathcal{I}$$

for some choice of the allocation vector $y = (y_i)_{i \in \mathcal{I}}$. If we combine (5.28) with (5.29) and the fact that T_{switch} is constant, then we see that, for

this allocation, we can choose τ_i such that

$$\mu_i(\tau_i) > y_i \mu_i^{\max} - \epsilon > a_i.$$

By setting

$$\tau = \left(\tau_j := \max_{i \in j} \tau_i\right)_{j \in \mathcal{J}},$$

for any $T > \tau$ we have

$$\mu_i(T) - a_i > 0 \quad \forall i \in \mathcal{I},$$

which means that $a \in \mathcal{A}^\circ(T)$, and thus proves Proposition 5.1. □

5.4.2 Proof of Proposition 5.2

Since we are investigating the heavy traffic behaviour of the system, we have $\rho \to 1/N$. Using the symmetry, i.e., that $\tau/N = T_{\text{switch}} + 1/\nu$ by (5.15), we can rewrite (5.16) as

$$\bar{Q}_i(\tau) = \frac{\rho\nu\tau + \lambda \frac{{T_{\text{switch}}}^2 + (N-1)(\tau/N)^2}{\tau}}{1 - \rho\nu\tau} = \frac{\rho\nu\tau + \frac{\lambda T_{\text{switch}}^2}{\tau} + \frac{(N-1)\lambda\tau}{N^2}}{1 - N\rho - N\rho\nu T_{\text{switch}}}. \tag{5.30}$$

If we assume that $\tau \to \infty$, which should be the case by (5.48), then $\tau \gg T_{\text{switch}}$, and $1/\nu \sim \tau/N$ in the limit. Using this and (5.30), we have that

$$\bar{Q}_i \sim \frac{N\rho + \frac{\lambda T_{\text{switch}}^2}{\tau} + \frac{(N-1)\lambda}{N\nu}}{1 - N\rho - N\rho\nu T_{\text{switch}}} \sim \frac{(N-1)\lambda}{N} \cdot \frac{1}{(1 - N\rho)\nu - N\rho T_{\text{switch}}\nu^2},$$

which is an expression, for which a minimum in ν can be determined easily, by taking the derivative. This gives

$$\nu^* = \frac{1 - N\rho}{2N\rho T_{\text{switch}}},$$

which by (5.15) and some algebraic calculations gives

$$\tau^* = \frac{2N^2\rho T_{\text{switch}}}{1 - N\rho} + NT_{\text{switch}} = \frac{N(1 + N\rho)}{1 - N\rho} \cdot T_{\text{switch}}, \tag{5.31}$$

which clearly goes to infinity as $\rho \to 1/N$. Plugging this into (5.30) we get

$$\bar{Q_i}^* = \frac{2\lambda T_{\text{switch}}}{N(1+N\rho)} + \frac{1}{NT_{\text{switch}}} \cdot \tau^* + \frac{2\lambda}{(1+N\rho)T_{\text{switch}}} \cdot \frac{N-1}{N^3} \cdot \tau^{*2}. \quad (5.32)$$

Since $\tau^* \to \infty$ on the right side of (5.32) the dominant term is the last one. Thus by summing over the in-roads we get (5.18). Furthermore if we consider the coefficient of τ^{*2} and the fact that $\lambda = \rho\mu \to \mu/N$ and $\sum_{i=1}^{N} \bar{Q}_i = N\bar{Q}_i$, we justify the formula in (5.27). □

5.4.3 Proof of Proposition 5.3

From Chapter 9 of [126] we see that in order to prove the tightness of a sequence $\{Z^{(c)}\}_c$, we must prove for each $\epsilon > 0$ and $t > 0$ that

$$\lim_{\delta \to 0} \mathbb{P} \left(\sup_{\substack{u,v:u,v<t, \\ |u-v|<\delta}} \|Z^{(c)}(u) - Z^{(c)}(v)\|_1 \geq \epsilon \right) = 0. \quad (5.33)$$

The evolution of $\bar{X}_{ir}(t)$ is given by (5.2). Since we have only defined $D_{ir}(t)$ as a counting process, we have to extend this definition to continuous time to have a clean definition in (5.24). Thus, we set

$$D_{ir}(t) = \sum_{n=1}^{N_{j(i)}(t)} S_{ir}^n + \mu_{ir}(t) \cdot \left(t - t^{j(i)}_{N_{j(i)}(t)} \right) + \Delta_{ir}(t - t^{j(i)}_{N_{j(i)}(t)}), \quad (5.34)$$

where $\Delta_{ir}(t)$ represents the fluctuation in service during the cycles, which comes from the ordering of the phases and stochastic effects. It is clearly bounded as

$$|\Delta_{ir}(t)| \leq \mu_i^{\max} t. \quad (5.35)$$

We also point out that, as long as the cycle lengths are preset which is the case in our policy, by the definition in (3.1) $N_{j(i)}(t)$ is deterministic. Even more so $N_j(t)$ is bounded for all $j \in \mathcal{J}$ since

$$T_n^j \geq |\Sigma_j| \cdot T_{\text{switch}} \ \ \forall n, \ \forall j \in \mathcal{J},$$

and thus

$$N_j(t) \leq \left\lceil \frac{t}{|\Sigma_j| \cdot T_{\text{switch}}} \right\rceil. \tag{5.36}$$

Now we are ready to demonstrate that (5.33) holds for $\bar{X}_{ir}(t)$ in the case when $i = i_0^r$. The result then follows for all other cases since the arrivals there are internal, since they equal the previous in-roads departures, for which we show tightness here. By (5.2), (5.24) and (5.34) we have

$$\begin{aligned}
&\bar{X}_{ir}(t) = \bar{X}_{ir}(0) + \frac{1}{c}\left[A_r(ct) - D_{ir}(ct)\right] \\
=&\bar{X}_{ir}(0) + \frac{1}{c}\left[A_r(ct) - \sum_{n=1}^{N_{j(i)}(ct)} S_{ir}^n - \mu_{ir}(ct)\left(ct - t^{j(i)}_{N_{j(i)}(ct)}\right) - \Delta_{ir}\left(ct - t^{j(i)}_{N_{j(i)}(ct)}\right)\right] \\
=&\bar{X}_{ir}(0) + a_r t + \bar{L}_r(t) - \bar{M}_{ir}(t) - \frac{1}{c}\int_0^{ct} \mu_{ir}(s)\mathrm{d}s - \frac{1}{c}\Delta_{ir}\left(ct - t^{j(i)}_{N_{j(i)}(ct)}\right),
\end{aligned} \tag{5.37}$$

where we define the following terms,

$$\bar{L}_r(t) = \frac{1}{c}(A_r(ct) - a_r ct),$$

$$\bar{M}_{ir}(t) = \frac{1}{c}\sum_{n=1}^{N_{j(i)}(ct)} (S_{ir}^n - T_n^{j(i)}\mu_{ir}^n). \tag{5.38}$$

By the triangle inequality it suffices to prove (5.33) holds for each term of the sum in (5.37) to prove that $\bar{X}_{ir}$ is tight. For $\bar{X}_{ir}(0) + a_r t$ (5.33) trivially holds.

For the integral we can use the fact that $\mu_{ir}(t)$ is bounded by $\mu_i^{\max}$ for all t, thus we have the Lipschitz-condition

$$\left|\frac{1}{c}\int_0^{cu} \mu_{ir}(s)\mathrm{d}s - \frac{1}{c}\int_0^{cv} \mu_{ir}(s)\mathrm{d}s\right| < \mu_i^{\max}|u - v| + 2\mu_i^{\max}c^{-1}.$$

Thus with $\delta < \epsilon/\mu_i^{\max}$ (5.33) is satisfied.

For Δ_{ir} we use the fact that it is bounded as described in (5.35), and

that we have pre-fixed cycle lengths. Thus

$$\left|\frac{1}{c}\Delta_{ir}(cu - t^{j(i)}_{N_{j(i)}(cu)}) - \frac{1}{c}\Delta_{ir}(cv - t^{j(i)}_{N_{j(i)}(cv)})\right| \leq \frac{2\mu_i^{\max}}{c} \sup_{n\leq N_{j(i)}} T_n^{j(i)} \to 0$$

as $c \to \infty$, since $\sup_{n\leq N_{j(i)}} T_n^j$ cannot be infinite by definition.

The process $\bar{M}_{ir}(t)$ can also be defined as

$$\bar{M}_{ir}(t) = \frac{1}{c}\sum_{n=1}^{\infty} \mathbb{I}_{\{n\leq N_{j(i)}(ct)\}}(S_{ir}^n - T_n^{j(i)}\mu_{ir}^n),$$

which definition is equivalent to (5.38) and helps our understanding of $\bar{M}_{ir}(t)$ as a martingale. We can look at S_{ir}^n as a series of random variables in $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$. Then $T_n^{j(i)}\mu_{ir}^n$ are their expected values respectively by (5.6). Since $N_{j(i)}(ct)$ is deterministic and bounded by (5.36), by using Doob's Optimal Stopping Theorem $\bar{M}_{ir}(t)$ is a martingale on $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$. For $m > n$,

$$\begin{aligned}
&\mathbb{E}\left[(S_{ir}^n - T_n^j\mu_{ir}^n)(S_{ir}^n - T_m^j\mu_{ir}^m)\right] \\
=&\mathbb{E}\left[\mathbb{E}[(S_{ir}^n - T_n^j\mu_{ir}^n)(S_{ir}^n - T_m^j\mu_{ir}^m)|\mathcal{F}_n]\right] \\
=&\mathbb{E}\left[(S_{ir}^n - T_n^j\mu_{ir}^n)\mathbb{E}[(S_{ir}^n - T_m^j\mu_{ir}^m)|\mathcal{F}_n]\right] = 0.
\end{aligned} \tag{5.39}$$

Thus we can apply Doob's L_2 inequality, (5.8) and (5.39) to get

$$\begin{aligned}
&\mathbb{P}\left(\sup_{\substack{u,v:u,v<t,\\ |u-v|<\delta}} |\bar{M}_{ir}(u) - \bar{M}_{ir}(v)| \geq \epsilon\right) \\
\leq&\mathbb{P}\left(\sup_{u:u<t} |\bar{M}_{ir}(u)| \geq \frac{\epsilon}{2}\right) \\
\leq&\frac{4}{\epsilon^2}\mathbb{E}\left[\left(\bar{M}_{ir}(t)\right)^2\right] = \frac{4}{c^2\epsilon^2}\sum_{n=0}^{N(ct)} \mathbb{E}\left[\left(S_{ir}^n - T_n^j\mu_{ir}^n\right)^2\right] \\
\leq&\frac{4}{c^2\epsilon^2}\sum_{n=0}^{N(ct)} \kappa T_n^j \leq \frac{4\kappa t}{c\epsilon^2} \to 0 \text{ as } c\to\infty.
\end{aligned} \tag{5.40}$$

Since the external arrivals follow a Poisson process, the process $\bar{L}_r(t)$

is also a martingale, so once again we can use Doob's L_2 inequality,

$$\begin{aligned}
&\mathbb{P}\left(\sup_{\substack{u,v:u,v<t,\\ |u-v|<\delta}} |\bar{L}_r(u)-\bar{L}_r(v)| \geq \epsilon\right)\\
\leq &\mathbb{P}\left(\sup_{u:u<t} |\bar{L}_r(u)| \geq \frac{\epsilon}{2}\right)\\
\leq &\frac{4}{\epsilon^2}\mathbb{E}\left[\left(\bar{L}_r(t)\right)^2\right] = \frac{4ta_r}{c\epsilon^2} \to 0 \text{ as } c\to\infty.
\end{aligned} \tag{5.41}$$

We have now established the tightness of each of the respective terms in (5.37). Thus our sequence of processes $\{X_{ir}^{(c)}\}_{c\in\mathbb{N}}$ are tight.

Therefore each subsequence of $\{X_{ir}^c(ct)/c\}_c$ has a weakly convergent subsequence: $X_{ir}^c(ct)/c \to \tilde{x}_{ir}(t)$, where $\tilde{x}_{ir}$ is some deterministic process. By the Skorohod Representation Theorem we may place these random variables on the same probability space and assume that convergence holds almost surely. Since $Q_i^c(ct)/c = \sum_{r:r\in i} X_{ir}^c(ct)/c$, we can define $\tilde{q}_i(t) = \sum_{r:r\in i} \tilde{x}_{ir}(t)$ to have

$$\frac{X_{ir}^c(ct)/c}{Q_i^c(ct)/c} \to \frac{\tilde{x}_{ir}(t)}{\tilde{q}_i(t)}. \tag{5.42}$$

Now we are left to prove that the limit is indeed given by the differentiability condition (5.22). Note that we have also proved that $\bar{L}_r$ and $\bar{M}_{ir}$ converge in distribution to zero for all $r\in\mathcal{R}$ and $i\in\mathcal{I}$ in (5.41) and (5.40) respectively. The same holds for Δ_{ir} as

$$\left|\frac{1}{c}\Delta_{ir}(ct - t^{j(i)}_{N_{j(i)}(ct)})\right| \leq \frac{\mu_i^{\max}}{c} \sup_{n\leq N_{j(i)}} T_n^{j(i)} \to 0. \tag{5.43}$$

By the piecewise constant property of μ_{ir}, we also have

$$\frac{1}{c}\int_0^{ct} \mu_{ir}(s)\mathrm{d}s = \int_0^t \mu_{ir}(cs)\mathrm{d}s.$$

We are going to describe the limit in the $i = i_0^r$ case. In the case where $i \neq i_0^r$ the same holds for both the arrivals and the departures as what is described below for the departures. By the definition (5.24)

the partitioning (5.37), and the convergences given by (5.40), (5.41) and (5.43) we have

$$\begin{aligned}\lim_{c\to\infty} \bar{X}_{ir}(t) &= \tilde{x}_{ir}(0) + a_r t - \lim_{c\to\infty} \int_0^t \mu_{ir}(cs)\mathrm{d}s \\ &= \tilde{x}_{ir}(0) + a_r t - \int_0^t \lim_{c\to\infty} \mu_{ir}(cs)\mathrm{d}s.\end{aligned} \tag{5.44}$$

In the latter equation we have used the Skorohod Representation Theorem. To arrive at the final conclusion we require that

$$\lim_{c\to\infty} \mu_{ir}(Q^c(cs)) \to \frac{\tilde{x}_{ir}(s)}{\tilde{q}_i(s)} \mu_i^*(\tilde{q}(s)). \tag{5.45}$$

This is demonstrated in Lemma 5.7, which is contained in Section 5.5.3.

We apply (5.45) to the integral in (5.44) to deduce the differentiability properties of $\tilde{x}$. In particular, for any t such that $\tilde{x}_{ir}(t) > 0$ and for h sufficiently small that $\tilde{x}_{ir}(s) > 0$ for all $t \le s \le t + h$, we have that

$$\begin{aligned}\frac{\tilde{x}_{ir}(t+h) - \tilde{x}_{ir}(t)}{h} &= a_r + \frac{1}{h}\int_t^{t+h} \lim_{c\to\infty} \mu_{ir}(cs)\mathrm{d}s \\ &= a_r + \frac{1}{h}\int_t^{t+h} \lim_{c\to\infty} \mu_{ir}(Q^c(cs))\mathrm{d}s \\ &= a_r + \frac{1}{h}\int_t^{t+h} \frac{\tilde{x}_{ir}(s)}{\tilde{q}_i(s)} \mu_i^*(\tilde{q}(s))ds,\end{aligned}$$

which gives the differentiability condition (5.22) as we take the limit $h \to 0$, thus $\tilde{x}_{ir} = x_{ir}$ and $\tilde{q}_i = q_i$ for all indices. □

5.4.4 Proof of stability

In this section we prove our main result, Theorem 5.1.

Consider the function

$$H(x) = \sum_{r\in\mathcal{R}} \sum_{i\in r} x_{ir} \log\left(\frac{x_{ir}\mu_i^*(q)}{q_i a_r}\right),$$

where we point out that $\mu_i^*(q)$ is the solution to the optimisation problem (5.11) for junction $j \ni i$, and thus its value depends on the other queue

sizes present at j. However, due to the properties of proportionally fair optimisation the partial derivatives do not show this dependence, as stated in the following lemma.

Lemma 5.4.

$$\frac{\partial H}{\partial x_{ir}} = \log\left(\frac{x_{ir}\mu_i^*(q)}{q_i a_r}\right).$$

Our next lemma puts a bound on $H(x)$.

Lemma 5.5. *The function $H(x)$ is positive, bounded when $\|x\|_1 = 1$ and minimised when $x = 0$.*

Furthermore we will use a technical lemma which goes as follows.

Lemma 5.6. *If u and v are two positive vectors with components indexed by $\mathcal{M}$, such that*

$$\sum_{m\in\mathcal{M}} u_m = \sum_{m\in\mathcal{M}} v_m,$$

then

$$\sum_{m\in\mathcal{M}} u_m \log\left(\frac{u_m}{v_m}\right) \geq \frac{1}{\sum_{m\in\mathcal{M}} u_m} \cdot \sum_{m\in\mathcal{M}} (u_m - v_m)^2.$$

Lemma 5.4, 5.5 and 5.6 are given proofs in Section 5.5.3. Now we can prove that $H(x)$ is a Lyapunov function for the fluid system (5.22). First, we show that the following equalities hold

$$\begin{aligned}\frac{\mathrm{d}H}{\mathrm{d}t} &= \sum_{r\in\mathcal{R}}\sum_{i\in r}\left(\frac{x_{i^r_- r}(t)}{q_{i^r_-}(t)}\mu^*_{i^r_-} - \frac{x_{ir}(t)}{q_i(t)}\mu_i^*\right)\log\left(\frac{x_{ir}\mu_i^*}{q_i a_r}\right)\\ &= -\sum_{r\in\mathcal{R}} a_r \sum_{i\in r}\frac{x_{ir}\mu_i^*}{q_i a_r}\log\left(\frac{\left[\frac{x_{ir}\mu_i^*}{q_i a_r}\right]}{\left[\frac{x_{i^r_+ r}\mu^*_{i^r_+}}{q_{i^r_+} a_r}\right]}\right).\end{aligned}$$

In the first equation we included $i = i_0^r$ too, which does not change the sum since by (5.20)

$$\log\left(\frac{x_{0r}\mu_0^*}{q_0 a_r}\right) = \log 1 = 0.$$

By incrementing the first terms of the summation such that the coefficients would become equal and multiplying and dividing by a_r we gained the final expression. Once again including $i = i_l^r$ does not change the sum, since by (5.21)

$$\log\left(\frac{x_{lr}\mu_l^*}{q_l a_r}\right) = \log 1 = 0.$$

Now we can bound the derivative of the proposed Lyapunov function as

$$\begin{aligned}\frac{\mathrm{d}H}{\mathrm{d}t} &\leq -\sum_{r\in\mathcal{R}}\left(a_r\left[\sum_{i\in r}\frac{x_{ir}\mu_i^*}{q_i a_r}\right]^{-1}\cdot\sum_{i\in r}\left[\frac{x_{ir}\mu_i^*}{q_i a_r}-\frac{x_{i_+^r r}\mu_{i_+^r}^*}{q_{i_+^r}a_r}\right]^2\right)\\ &\leq -\sum_{r\in\mathcal{R}}\left(\frac{a_r^2}{\tilde{\mu}}\sum_{i\in r}\left[\frac{x_{ir}\mu_i^*}{q_i a_r}-\frac{x_{i_+^r r}\mu_{i_+^r}^*}{q_{i_+^r}a_r}\right]^2\right)\leq -\epsilon,\end{aligned}$$

for some $\epsilon > 0$. In the first inequality we applied Lemma 5.6. In the second inequality we applied that

$$\sum_{i\in r}\frac{x_{ir}\mu_i^*}{q_i a_r}\leq\frac{\tilde{\mu}}{a_r},$$

where we introduced

$$\tilde{\mu} = |\mathcal{I}|\max_{i\in\mathcal{I}}\mu_i^{\max}.$$

To prove the third inequality we first look into the conditions under which the sum could equal zero. To have that for each in-road $i \in r$ on each route $r \in \mathcal{R}$ we would need

$$\frac{x_{ir}\mu_i^*}{q_i a_r}-\frac{x_{i_+^r r}\mu_{i_+^r}^*}{q_{i_+^r}a_r}=0.$$

This would mean that the terms are constant along each in-road i along a route, including the auxiliary i_0, for which the fraction was set as 1. Thus, this requires that

$$a_r = \frac{x_{ir}\mu_i^*}{q_i}\quad\forall i\in r,\ \forall r\in\mathcal{R}.$$

By summing over for each in-road, we get

$$\sum_{i\in\mathcal{I}} a_i = \sum_{r\in\mathcal{R}} \sum_{i:i\in r} a_i = \sum_{r\in\mathcal{R}} a_r = \sum_{i\in\mathcal{I}} \mu_i^*.$$

This contradicts (5.26), since μ_i^* clearly belongs to the boundary of $\mathcal{A}(T)$. Thus (5.4.4) cannot hold and $H(x)$ has a strictly negative drift. Since by Lemma 5.5 $H(x) = 0$ only if $x = 0$, for all $x(0)$ such that $H(x(0)) \leq h$ for some positive constant h, we have for all $t \geq h/\epsilon$ that

$$q(t) = 0.$$

By definition this means that the fluid system is stable, which concludes the proof of Theorem 5.1. □

5.5 Supplementary results

This section provides supplementary results (i.e. examples, calculations, lemmas), that we rely on for the rest of the chapter, however they are of lesser importance in comparison to the main results discussed so far.

5.5.1 On the estimation of queue lengths

In this section we shortly discuss the estimation of the cycle lengths $\widetilde{Q}$ introduced in Section 5.2.4. For our proposed policy to be stable the only condition needed on these estimators is that they have to be unbiased, i.e. $\mathbb{E}Q = \mathbb{E}\widetilde{Q}$. However, for practical purposes we suggest to use a "moving average", which is given as follows

$$\widetilde{Q}(U) = \sum_{u=1}^{Z} \alpha_u Q(U-u),$$

where $Q(U)$ denotes the measured queue size at sampling time U, Z denotes the number of samples considered in the estimation process and α_u denotes the weights put on each sample, with $\sum_{u=1}^{Z} \alpha_u = 1$. An example could be the following. To provide the estimation for a certain point of the day consider the samples from the previous days at the same time making a distinction between workdays and holidays

and set α to be linearly decreasing, which would mean that the most recent measurements would have the highest weights. We also point out that the cycle length settings described in Section 5.3 follow a similar method when using the results of previous runs to make the queue length estimations.

5.5.2 Examples

Example showing the dependencies of the set of admissible service rates

Take for all in-roads the following service function (with time units given in minutes),

$$s_i(t) = \begin{cases} 10t, & \text{if } t \leq 0.1 \\ 1 + 30(t - 0.1) & \text{if } t > 0.1, \end{cases}$$

which includes the mentioned setup phase. If we take the switching times $T_{\text{switch}} = 0.1$ min, then at a single junction with two competing in-roads, that both have 11 vehicles/min arrivals on average, the queues will build up if $T_j = 1.2$ min. The same junction can however be served, if $T_j = 2.4$ min. Even more interesting is the fact that in the $T_j = 1.2$ min case, if one in-road has 22 vehicles/min arrivals on average, while there is no traffic arriving to the other in-road, the junction can be served again. Thus the arrival vector $(11, 11)$ is outside of the set $\mathcal{A}(T)$, even though both $(22, 0)$ and $(0, 22)$ are in and $(11, 11)$ is their convex combination.

Example of a polling model with deterministic rates

Consider a junction with N competing in-roads, each having identical, deterministic arrival rates λ and identical constant service rates μ when receiving service. Scheduling green times G_i proportionally to the average queue lengths results in allocating $1/N$ of the effective time to each phase. Thus each in-road will be idle for

$$T - G_i = \frac{N - 1}{N} \cdot T + T_{\text{switch}}.$$

During the idle period a queue will build up, which will leave with all new arrivals during the green period. Thus at the end of the idle period,

the queue size present will be

$$Q^{\text{idle}} = \lambda \left(\frac{N-1}{N} \cdot T + T_{\text{switch}} \right),$$

which can empty out in

$$T^{\text{empty}} = \frac{\lambda}{\mu - \lambda} \cdot \left(\frac{N-1}{N} \cdot T + T_{\text{switch}} \right).$$

If the green time allocated is shorter than T^{empty}, then the queues never empty out, instead they grow infinitely large and consequently the waiting times become infinitely large too. If $G_i \geq T^{\text{empty}}$, then the queues will receive green light even after they emptied out which will result in all subsequent arriving vehicles leaving immediately with the queue size starting to grow from 0 again during the idle period. The average queue size in this case can be calculated by

$$\begin{aligned} \bar{Q} &= \frac{1}{T} \int_0^T Q(\tau) \mathrm{d}\tau = \frac{1}{T} \left[\frac{1}{2} \cdot \frac{\lambda\mu}{\mu - \lambda} \cdot \left(\frac{N-1}{N} \cdot T + T_{\text{switch}} \right)^2 \right] \\ &= \frac{T}{2} \cdot \frac{\lambda\mu}{\mu - \lambda} \left[\frac{T_{\text{switch}}}{T} + \frac{N-1}{N} \right]^2. \end{aligned}$$

Thus by Little's law the average waiting time is

$$\bar{W}(T) = \frac{\bar{Q}}{\lambda} = \frac{T}{2} \cdot \frac{\mu}{\mu - \lambda} \left[\frac{T_{\text{switch}}}{T} + \frac{N-1}{N} \right]^2, \tag{5.46}$$

if the system is stable. In order to have this quantity, we need $G_i \geq T^{\text{empty}}$, which happens if

$$T \geq \frac{\mu - \lambda}{\mu - N\lambda} \cdot 2NT_{\text{switch}}, \tag{5.47}$$

from where we can also see that $\mu > N\lambda$ is needed for stability. Since the function in (5.46) has its minimum at $T = \frac{N}{N-1} T_{\text{switch}}$ and is strictly increasing for bigger values of T, the vehicles have minimal average waiting time if there is an equality in (5.47). Thus in a deterministic system the shortest cycle length which ensures stability is the optimal

choice.

Determining the optimal cycle lengths for the symmetric case of the polling model in Section 5.2.3

Let us consider the polling model of Section 5.2.3. Take the $N = 2$ case in symmetry, i.e. assume that the two phases behave in the same way, $\lambda_1 = \lambda_2 = \lambda$, $\mu_1 = \mu_2 = \mu$, . This allows us to give a system of equations with (5.14), (5.16) and (5.15), from which ν can be eliminated. Assuming that τ takes the form of

$$\tau = \frac{2(T_{\text{switch}} + \epsilon)}{1 - 2\rho}, \tag{5.48}$$

thus assuring (5.17) with the perturbation ϵ being positive, we can express the expected queue lengths as a function of ϵ as

$$\begin{aligned}
\mathbb{E}Q(\epsilon) =& \frac{1}{2\mu(1-2\rho)^2\epsilon(T_{\text{switch}} + \epsilon)} \\
&\times \Big[4T_{\text{switch}}^2 + 4\epsilon^2 - 8\mu\rho^2 T_{\text{switch}}^3 + 6\mu\rho T_{\text{switch}}^3 + 8T_{\text{switch}}^3\mu\rho^3 \\
&-8\rho T_{\text{switch}}^2 - \mu T_{\text{switch}}^3 + \mu\epsilon^3 - 2\mu\rho\epsilon T_{\text{switch}}^2 \\
&-2\mu\rho\epsilon^2 T_{\text{switch}} + 4\mu\rho^2\epsilon T_{\text{switch}}^2 - 16\rho\epsilon T_{\text{switch}} \\
&+4\mu\epsilon^2 T_{\text{switch}} + 3\mu\epsilon T_{\text{switch}}^2 + 8\epsilon T_{\text{switch}} - 8\rho\epsilon^2 \Big].
\end{aligned}$$

Finding the minimum of this amongst positive values for ϵ leads to solving a fourth order polynomial, which is a doable task. We will save the reader from the lengthy symbolic expression. This tells us, that ϵ has a clear positive value, and is not just a small perturbation contrary to what could be expected from the result of the previous example.

5.5.3 Proofs for lemmas

In the proof of Proposition 5.2, we required the following lemma which we now prove.

Lemma 5.7. *In the case of $\tilde{x}_{ir} > 0$ the following holds for the fluid*

limit of $\mu_{ir}(Q^c(cs))$,

$$\lim_{c\to\infty} \mu_{ir}(Q^c(cs)) \to \frac{\tilde{x}_{ir}(s)}{\tilde{q}_i(s)}\mu_i^*(\tilde{q}(s)). \tag{5.49}$$

Proof. To prove our statement let us use the definition of μ_{ir}, which is given in (5.6). Thus

$$\mu_{ir}(Q^c(cs)) = \frac{X_{ir}^c(cs)}{Q^c(cs)} \frac{s_i(E_i y_i(Q^c(cs)))}{T_i}. \tag{5.50}$$

For the condition that $\tilde{x}_{ir} > 0$, we need $X_{ir}^c(ct) \to \infty$ as $c \to \infty$. Thus $Q_i^c(ct) \to \infty$ as $c \to \infty$, and the same holds for $\sqrt{Q_i^c(ct)}$. By the policy described in Section 5.2.4, we have $E_i(c) \to \infty$ as $c \to \infty$, since $E_i \propto \sqrt{Q_i}$. Also by the definition in (3.2), we have that

$$\frac{E_i(c)}{T_i(c)} \to 1, \tag{5.51}$$

as $c \to \infty$. Furthermore, by the assumption (5.4), we have that

$$\frac{s_i(ky)}{k} \to \mu_i^{\max}, \tag{5.52}$$

as $k \to \infty$. Finally, by Lemma A.3 in [71], we have that $y_i(\tilde{q})$ is continuous in $\tilde{q}$ for all indexes i with $\tilde{q}_i > 0$. Thus if $\tilde{x}_{ir} > 0$, we have

$$y_i(Q^c(cs)) = y_i\left(\frac{Q^c(cs)}{c}\right) \to y_i(\tilde{q}(s)), \tag{5.53}$$

from which we can deduce that

$$\mu_i(Q^c(cs)) = \frac{s_i^*(E_i y_i(Q^c(cs)))}{T_i} \to y_i^*(q(s))\mu_i^{\max} = \mu_i^*(q(s)). \tag{5.54}$$

If we combine the assumption in (5.5) with (5.54), we can conclude that (5.49) holds.s □

In this section we use the relative entropy, which is defined for two probability distributions u and v, which are both defined on the same finite set $\mathcal{M}$, as

$$D(u\|v) = \sum_{m\in\mathcal{M}} u_m \log\left(\frac{u_m}{v_m}\right). \tag{5.55}$$

We note that $D(u\|v)$ is strictly non-negative and the following bound also holds on it,

Lemma 5.8. *(Pinsker's Inequality)*

$$D(u\|v) \geq \sum_{m\in\mathcal{M}} |u_m - n_m|. \tag{5.56}$$

Proof of lemma 5.4

We prove this lemma by taking the derivatives from first principles. Let us use the notation $x^h_{ir} = x_{ir} + h$ and all other components of x remain the same. Naturally $q^h_i = q_i + h$, while $q^h_{i'} = q_{i'}$ for $i' \neq i$. Then for $h > 0$,

$$\begin{aligned}
&\frac{H(x^h) - H(x)}{h} \\
=&\frac{1}{h}\left[\sum_{r\in\mathcal{R}}\sum_{i\in r} x^h_{ir}\log\left(\frac{x^h_{ir}\mu^*_i(q^h)}{q^h_i a_r}\right) - \sum_{r\in\mathcal{R}}\sum_{i\in r} x_{ir}\log\left(\frac{x_{ir}\mu^*_i(q)}{q_i a_r}\right)\right] \\
\geq&\frac{1}{h}\left[\sum_{r\in\mathcal{R}}\sum_{i\in r} x^h_{ir}\log\left(\frac{x^h_{ir}\mu^*_i(q)}{q^h_i a_r}\right) - \sum_{r\in\mathcal{R}}\sum_{i\in r} x_{ir}\log\left(\frac{x_{ir}\mu^*_i(q)}{q_i a_r}\right)\right] \\
=&\frac{1}{h}\left[(x^h_{ir}\log x^h_{ir} - x_{ir}\log x_{ir}) + (q^h_i\log q^h_i - q_i\log q_i)\right] + \log\left(\frac{\mu^*_i(q)}{a_r}\right),
\end{aligned} \tag{5.57}$$

where the inequality derives from the fact that $\mu^*_i(q)$ is suboptimal for the proportional fair optimisation with the parameter choice q^h. On the other hand if we leave the first summation in the first equality the same,

but exchange $\mu_i^*(q)$ with $\mu_i^*(q^h)$, we can apply the same logic, thus

$$\begin{aligned}
&\frac{H(x^h) - H(x)}{h} \\
=&\frac{1}{h}\left[\sum_{r\in\mathcal{R}}\sum_{i\in r} x_{ir}^h \log\left(\frac{x_{ir}^h \mu_i^*(q^h)}{q_i^h a_r}\right) - \sum_{r\in\mathcal{R}}\sum_{i\in r} x_{ir}\log\left(\frac{x_{ir}\mu_i^*(q)}{q_i a_r}\right)\right] \\
\leq&\frac{1}{h}\left[\sum_{r\in\mathcal{R}}\sum_{i\in r} x_{ir}^h \log\left(\frac{x_{ir}^h \mu_i^*(q^h)}{q_i^h a_r}\right) - \sum_{r\in\mathcal{R}}\sum_{i\in r} x_{ir}\log\left(\frac{x_{ir}\mu_i^*(q^h)}{q_i a_r}\right)\right] \\
=&\frac{1}{h}\left[(x_{ir}^h \log x_{ir}^h - x_{ir}\log x_{ir}) + (q_i^h \log q_i^h - q_i \log q_i)\right] + \log\left(\frac{\mu_i^*(q)}{a_r}\right).
\end{aligned} \tag{5.58}$$

These two bounds can be derived the same way when $h < 0$. Since the $q \mapsto \mu^*(q)$ function is continuous by the properties of the service function, taking the limit as $h \to 0$ on these bounds imply that

$$\frac{\partial H(x)}{\partial x_{ir}} = \log x_{ir} - \log q_i + \log\left(\frac{\mu_i^*(q)}{a_r}\right), \tag{5.59}$$

as stated in Lemma 5.4. □

Proof of lemma 5.5

Observe that $H(x)$ can be expressed as linear combination of relative entropy terms as

$$\begin{aligned}
H(x) =& \sum_{i\in\mathcal{I}} q_i D\left(\left(\frac{x_{ir}}{q_i}\right)_{r\ni i} \| \left(\frac{a_r}{a_i}\right)_{r\ni i}\right) + \sum_{i\in\mathcal{I}} q_i \log\left(\frac{\mu_i^*(q)}{a_i}\right) \\
\geq& \sum_{i\in\mathcal{I}} q_i \log\left(\frac{\mu_i^*(q)}{a_i}\right) \geq 0.
\end{aligned} \tag{5.60}$$

The first inequality is a consequence of the positivity of relative entropy, whereas the second follows by the optimality of $\mu^*(q)$ and (5.26). From the entropy equality we see that $H(x)$ is continuous for $\|x\|_1 = 1$ and so bounded. Further the inequalities above hold with equality iff $x = 0$. □

Proof of lemma 5.6

In order to prove this lemma we define $\hat{u}_m$, and $\hat{v}_m$ respectively, as follows,

$$\hat{u}_m = \frac{u_m}{\sum_{m' \in \mathcal{M}} u_{m'}}. \tag{5.61}$$

Then by applying the definition of relative entropy on $\hat{u}_m$ and $\hat{v}_m$ and applying Pinsker's Inequality, we get

$$\begin{aligned} \sum_{m \in \mathcal{M}} u_m \log\left(\frac{u_m}{v_m}\right) &= D(\hat{u}\|\hat{v}) \sum_{m \in \mathcal{M}} u_m \\ &\geq \left(\sum_{m \in \mathcal{M}} |\hat{u}_m - \hat{v}_m|\right)^2 \sum_{m \in \mathcal{M}} u_m \\ &= \frac{1}{\sum_{m \in \mathcal{M}} u_m} \left(\sum_{m \in \mathcal{M}} |u_m - v_m|\right)^2 \\ &\geq \frac{1}{\sum_{m \in \mathcal{M}} u_m} \sum_{m \in \mathcal{M}} (u_m - v_m)^2, \end{aligned} \tag{5.62}$$

by rearranging after the second equality, applying (5.6) and bounding afterwards. □

5.6 Conclusion

This chapter discussed another decentralised signal control policy for traffic light systems. In comparison to the scheme presented in Chapter 4 this algorithm provided fair allocation of the capacities to drivers as it was based on the notion of proportional fairness. Furthermore we have incorporated more features, such as a detailing the effect of vehicles which need to speed up at the beginning of a green phase, that made the model more realistic and thus fit for implementation. We have investigated the possibility of variable cycle times and determined an optimal scaling, which can be used to set up a time plan for cycles, whilst maintaining optimal green time allocation on a cycle-to-cycle basis. By using the technique of fluid scaling we have prooved stability of the network under the proposed policy. This indicates that our method

keeps the queues in the network growing unboundedly for the largest possible set of arrival rates that will provide sufficient throughput even in a congested network. Using simulation we validated the proposed choice of cycle lengths. Furthermore we have compared the performance of our scheme to a well-known method of the literature, the P0-policy, which is also based on the same principles. The results are promising.

Our study still left some open questions. Examples include eliminating some modelling assumptions such as the Poisson property of external arrivals or the rather naive approach on the traveling times between junctions. Further research can also discuss ordering of the phases or the possibility of adjusting cycle lengths on a shorter time scale.

Chapter 6

Models for motorway traffic

Outline. The previous chapters discussed control policies for urban traffic networks. There the most important task is to control traffic at junctions efficiently. Interurban traffic presents different challenges as the traffic flow may show great variability between the locations of possible control. Another important feature, the control of access is closely connected to the main roads of the network, the *motorways*. As these are the roads where traffic flow can evolve for the longest parts without external influences, they are often the focus of the discussion. Our goal is to establish a connection between urban and interurban control mechanisms, a challenge that we deal with in Chapter 7. In order to do so however, we first need to discuss the evolution of traffic flow between the points of access, which is the purpose of this chapter. Furthermore, we present a novel model for highway traffic which provides an insight on the vehicle-to-vehicle level, whilst maintaining model parameters that make fitting to actual traffic scenarios a simple task.

The chapter is organised as follows. In Section 6.1 we present and discuss the relevant results of the road traffic literature. Section 6.2 introduces the so-called two lane model, providing some conditions that ensure its desirable properties and validating it by fitting to data acquired from the British highway system.

6.1 Literature overview

Traffic flows on the motorway

As motorways provide the arteries for interurban traffic in most places of the world they have been in the focus of the traffic modeling community from the very beginning, with the first description coming from Greenshields in 1935 [54]. They also allow for a different modeling approach of access points, which we will refer to as on-ramps and off-ramps for the rest of the work, there is no interruption for the traffic under normal circumstances. Therefore the analysis mostly focuses on three defining quantities, namely the density, the average speed and the flow of vehicles, for which we use the usual notation of the literature, denoting them by k, v and q respectively. By definition these parameters are related according to

$$q = k \cdot v.$$

However, the functional relation between the density and the average speed (and therefore the flow) is much more complicated, and is of interest. In the literature the way this relation is described is commonly reffered to as the fundmanental diagram. In the rest of Section 6.1 we describe its most used forms, discussing the models they emerge from.

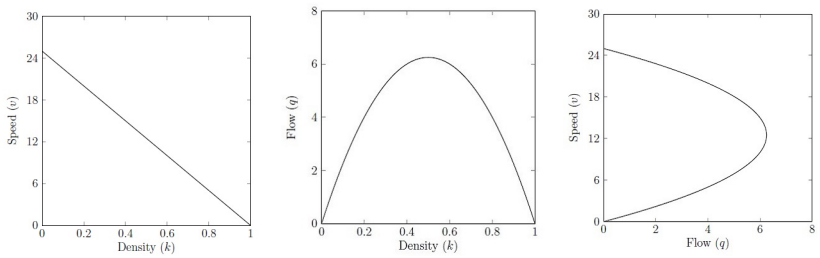

Figure 6.1. Fundamental diagram from the experimental data of Greenshields [54].

The seminal work of Greenshields [54] captured the empirical relation between flow, density and speed resulting in the first form of the fundamental diagram. From his measurements he concluded a simple linear relation between speed and density, resulting in parabolic relations for

the flow with the other two quantities. These are shown in Figure 6.1. The term fundamental diagram was first introduced by Haight [56], who used the relation found by Greenshields, to give the relation

$$q = v_f \left(k - \frac{k^2}{k_{\text{jam}}} \right), \tag{6.1}$$

where v_f is the desired maximum speed on a highway or freeflow speed and k_{jam} is the traffic jam concentration. The described relation is an example of a single-regime traffic model.

Model	Function
Greenshields	$v = v_f \left(1 - \frac{k}{k_{\text{jam}}}\right)$,
Greenberg	$v = v_m \ln\left(\frac{k_{\text{jam}}}{k}\right)$,
Underwood	$v = v_f \exp\left(-\frac{k}{k_c}\right)$,
Northwestern	$v = v_f \exp\left(-\frac{1}{2}\left(\frac{k}{k_c}\right)^2\right)$,
Drew	$v = v_f \left(1 - \left(\frac{k}{k_{\text{jam}}}\right)^{n+\frac{1}{2}}\right)$,
Pipes-Munjal	$v = v_f \left(1 - \left(\frac{k}{k_{\text{jam}}}\right)^{n}\right)$,
Kühne and Rödiger	$v = v_f \left(1 - \left(\frac{k}{k_{\text{jam}}}\right)^{a}\right)^{b}$,
Modified Greenshields	$v = v_0 + (v_f - v_0)\left(1 - \left(\frac{k}{k_{\text{jam}}}\right)^{n}\right)$,
Newell	$v = v_f \left(1 - \exp\left(\frac{\lambda}{v_f}\left(\frac{1}{k} - \frac{1}{k_{\text{jam}}}\right)\right)\right)$,
Del Castillo and Benitez	$v = v_f \left(1 - \exp\left(\frac{\lvert\omega_v\rvert}{v_f}\left(1 - \frac{k}{k_{\text{jam}}}\right)\right)\right)$,
Van Aerde	$k = \frac{1}{c_1 + \frac{c_2}{v_f - v} + c_3 v}$,
MacNicholas	$v = v_f \left(\frac{k_{\text{jam}}^n - k^n}{k_{\text{jam}}^n + c k^n}\right)$,
Wang et al.	$v = v_0 + (v_f - v_0)\left(1 + \exp\left(\frac{k - k_t}{\theta_1}\right)\right)^{-\theta_2}$.

Table 6.1. Speed-density relationships of single-regime traffic models

To give an overview of single-regime traffic models that have been developed since the Greenshields model, we adapted Table 6.1 from Wang et al.[150]. In the table we use to more constants other than v_f and

k_{jam}, namely v_m and k_m, which denote the speed and density at full capacity respectively, i.e., when the flow q obtains its theoretic maximum q_m. Wang et al. fit several single-regime models from Table 6.1 to empirical data. As they discuss several models introduced since, better capture realistic traffic flows. The Greenberg model [53], based on fluid-flow analogies, proposes a logarithmic speed-density relationship, which performs well under congested conditions as the speed goes to zero when approaching the jam density, but at low densities its results are not useful ($v \to \infty$ as $k \to 0$). To capture the behaviour at low densities Underwood proposed an exponential speed-density relationship, however that does not satisfy boundary conditions under congested conditions as it allows for positive speed above k_{jam} (to model a dense motorway the congestion density k_c is introduced). This limitation is also found in the Northwestern model [38], in which a bell-shaped speed-density relationship was proposed The simple formulation of Greenshields model was expanded upon in the Drew model [39] and the Pipes-Munjal model [117]. Both models introduce an extra parameter n, which allows for extra degrees of freedom in fitting the models to empirical data. This general form is also captured by the model of Kühne and Rödiger [76], and the so-called modified Greenshields model, which is an extension by Mahmassani et al. [92]. This latter model introduced a parameter v_0 denoting the minimal speed at jam density. Based on generating functions, Del Castillo and Benitez proposed a speed-density relationship based on v_f, k_{jam} and a jam wave speed ω_v in [36, 37]. They noted that their model can also be obtained from Newell's car-following model. Another form can be obtained from the Van Aerde car-following model [146], which contains three extra parameters c_1, c_2 and c_3, which emerge by solving the boundary conditions. It is shown by MacNicholas that a model similar to the Van Aerde model can be obtained with fewer parameters in his paper [91], where he makes a comparison between the two models. Wang et al. also proposed a novel model in [149]. Their logistic speed-density model based on five parameters, namely k_t, the density at which the transition between free-flowing and congested traffic occurs, θ_1 and θ_2, which are used to fit the traffic model and v_0 and v_f, which are defined as before. For a more concise overview and visual comparison of a selection of these models we direct the reader to Wang et al. [150].

A different approach is taken by Daganzo, who describes the relation in [33] as

$$q = \inf_{v}\{kv + R(v)\}, \tag{6.2}$$

where $R(v)$ is an upper bound to the average rate at which traffic can overtake any observer that travels with average speed v for a long time and is defined as

$$R(v) = \lim_{t \searrow 0} \inf_{P \in \mathcal{P}} \{\Delta(P) : v_P = v\}/t, \tag{6.3}$$

where $\mathcal{P}$ denotes the set of paths a driver can take and $\Delta(P)$ denotes the "cost" of one such path.

Hoogendoorn et al. generalized the fundamental diagram to incorporate the effects of unexpected events, such as accidents, giving the following formula in [63],

$$v(k) = \min\left\{v_f, \beta\left(\frac{1}{k} - \frac{1}{k_{\text{jam}}}\right)\right\}\left(1 - \frac{\sigma}{\sigma_0}\right), \tag{6.4}$$

where σ denotes the standard deviation of k and σ_0 denotes a scaling parameter for the spatial variation of the density. The model of Hoogendoorn et al. essentially describes two regimes, which differ in their speed-density relation. As such it is not the first multi-regime traffic model as several instances in the literature describe the use of multiple single-regime models at the same time (one for each regime) to obtain better fitting to data. A drawback of multi-regime models lies in finding the point at which regimes change. We list the most notable multi-regime models in Table 6.2, we note however, that any combination of single-regime models can serve as a multi-regime model, giving way to many more possibilities. The simplest, and most used multi-regime traffic model was introduced by Newell [106, 107]. It is called the Triangular model, due to the triangular shape of the flow-density plot. It introduces k_{bp}, which denotes the density at which traffic starts breaking down, hence calling it break-point density, setting it as the point, where the regime changes. Similar speed-density relationships are obtained with the modified versions of the Greenberg [38] and the Greenshields model [96]. All three models give a constant speed until the density reaches the

break-point, after which the speed decreases to zero or v_0. Edie proposed a model in [41] that combined the Underwood and the Greenberg model, using the first at low densities and the latter in congestion, which happens when the density reaches the congestion density k_c.

Model	Function	Regime
Triangular	$v = v_f,$	if $k \leq k_{bp},$
	$v = \frac{v_f k_c}{k_c k_{\text{jam}}} \left(1 - \frac{k}{k_{\text{jam}}}\right),$	if $k > k_{bp},$
Modified Greenshields	$v = v_f,$	if $k \leq k_{bp},$
	$v = v_0 + (v_f - v_0) \left(\frac{k_{\text{jam}} - k}{k_{\text{jam}} - k_{bp}}\right)^{\alpha},$	if $k > k_{bp},$
Modified Greenberg	$v = v_f,$	if $k \leq k_{bp},$
	$v = c \ln \left(\frac{k_{\text{jam}}}{k}\right),$	if $k > k_{bp},$
Edie	$v = v_f \exp\left(-\frac{k}{k_c}\right),$	if $k \leq k_c,$
	$v = c \ln \left(\frac{k_{\text{jam}}}{k}\right),$	if $k > k_c,$
Hoogendoorn	$v = v_f \left(1 - \frac{\sigma}{\sigma_0}\right),$	if $k \leq \frac{1}{\frac{v_f}{\beta} + \frac{1}{k_{\text{jam}}}},$
	$v = \beta \left(\frac{1}{k} - \frac{1}{k_{\text{jam}}}\right) \left(1 - \frac{\sigma}{\sigma_0}\right),$	if $k > \frac{1}{\frac{v_f}{\beta} + \frac{1}{k_{\text{jam}}}}.$

Table 6.2. Speed-density relationships of multi-regime traffic models

All the above models describe traffic at the lowest level of detail, thus they are often referred to as macroscopic models. Other examples of macroscopic models include the classical Lighthill-Whitham-Richards model [85, 86, 125] and the Payne model [116]. More detail is given in mesoscopic models, however there individual drivers are still not distinguished. The behaviour of drivers is instead characterised in terms of the probability density $f(x, v, t)$ of vehicles at position x with speed v at time t. Examples of mesoscopic models are headway distribution models [23] and gas-kinetic continuum models [118, 119]. The highest level of detail is obtained in microscopic models, where each individual vehicle is characterised by its position and behaviour over time. Well-known micropscopic models are the car-following model [21, 26], the cellular automata model [104] and the lane-changing model [8]. These mainly come from statistical physics and non-linear dynamics, often leading to

a system of differential equations. We highlight one such model, the so-called totally asymmetric simple exclusion process (TASEP), which often serves as a building block towards more complex models of the type [29]. It also experiences shockwaves as a possible solution to its hydrodynamic differential equations, which can serve as a model for the stop-go waves, that occur in congested traffic. Section 6.2 describes a novel model similar to TASEP, introducing more complexity to allow for modeling of multiple lanes of traffic flow.

Traffic control for motorways

We end our literature overview with a brief review of control schemes for motorway traffic. Since uninterrupted advancement is the preferred state of motorway traffic, most methods aim at preserving a free flow by avoiding congestion through limiting access to the road. This can be done, where external traffic can join the motorway, namely where lower-standard right-in/right-out access is used for direct connections to side roads. These road sections are often referred to as ramps, hence the name for the widespread solution for this type of control, ramp metering.

Ramp metering policies operate by controlling traffic streams from on-ramps to the motorway. A widely deployed policy is the ALINEA algorithm [112, 113]. The essential idea of this policy is to try and keep the occupancy downstream of the on-ramp to a pre-specified occupancy set point. Following an evaluation of a few strategies, the RWS strategy was adopted for deployment in the Netherlands [99]. In the RWS policy, as compared to ALINEA, only the calculation of the number of vehicles that may be allowed to enter the motorway is different. Otherwise, the basic rules for switching the system on and off are the same in both of these policies. An alternative idea is the Stratified Zone Metering strategy [143], which is based on the SCATS framework [90], and was first introduced in Minnesota. It has been developed over the years, see [44, 65]. One of the main features of the system are express lanes to which access can be gained by paying toll.

The performance of ramp metering strategies, which rely only on local measurements, is well analysed, see the results of Chu et al. for example [30]. Many of the improvements of the algorithms came from the insights gained through analysis and by applying ideas from different fields. For

insatance, a policy called Proportionally Fair Metering, inspired by rate control models developed for the Internet, has also been proposed [50]. We aim to further the discussion by investigating schemes that consider not only the traffic on the motorway, but also on the surrounding roads by the results presented in Chapter 7.

6.2 The two lane model

In this section we introduce a novel model that aims to give a better understanding of the fundamental diagram on the level of the vehicles. The construction is similar to that of the TASEP [84], extending it in order to be able to describe a motorway more realistically. Therefore we consider a continuous time Markov process, which runs on the following configuration space,

$$\Omega = \{\underline{\omega} = (\omega_i)_{i\in\mathcal{N}} : \omega_i \in I\} = I^{\mathcal{N}},$$

where $I = \{0, 1, 2\}$ is the set of the possible occupation numbers and $\mathcal{N} \subset \mathbb{Z}$. This way we consider two lanes of traffic, which gives the possibility of vehicles overtaking each other. The size of the system can change to reflect on the length of the motorway segment we aim to model. The evolution of the process is then given by transitions $(\omega_i, \omega_{i+1}) \to (\omega_{i+1} - 1, \omega_{i+1} + 1)$ with rate

$$\begin{aligned} r(\omega_i, \omega_{i+1}) =& r_{1,0}\mathbb{I}\{\omega_i = 1, \omega_{i+1} = 0\} + r_{1,1}\mathbb{I}\{\omega_i = 1, \omega_{i+1} = 1\} + \\ & r_{2,0}\mathbb{I}\{\omega_i = 2, \omega_{i+1} = 0\} + r_{2,1}\mathbb{I}\{\omega_i = 2, \omega_{i+1} = 1\}. \end{aligned} \tag{6.5}$$

We also allow vehicles to join the system on the left boundary with rate

$$\alpha(\omega_1) = \alpha_0\mathbb{I}\{\omega_1 = 0\} + \alpha_1\mathbb{I}\{\omega_1 = 1\}, \tag{6.6}$$

and on the right boundary with rate

$$\beta(\omega_N) = \beta_1\mathbb{I}\{\omega_N = 1\} + \beta_2\mathbb{I}\{\omega_N = 2\}. \tag{6.7}$$

These transitions happen independently conditionally on $\underline{\omega}(t)$. Thus at time t, $\underline{\omega}(t)$ gives the state of the system. If we let $\varphi : \Omega \to \mathbb{R}$ be a finite cylinder function (we later refer to possible choices of φ as a test function), i.e. a function that only depends on the occupation numbers of a finite number of sites, then the evolution of $\underline{\omega}(t)$ is a Markov process with the formal infinitesimal generator

$$(L\varphi)(\underline{\omega}) = (B\varphi)(\underline{\omega}) + \sum_{i=1}^{N-1} r(\omega_i, \omega_{i+1})(\varphi(\widetilde{\underline{\omega}}^i) - \varphi(\underline{\omega})), \tag{6.8}$$

where

$$\widetilde{\omega}_j^i = \begin{cases} \omega_j & \text{if } j \neq i, i+1 \\ \omega_i - 1 & \text{if } j = i \\ \omega_{i+1} + 1 & \text{if } j = i+1, \end{cases}$$

$(B\varphi)(\underline{\omega})$ describes the boundary terms with

$$(B\varphi)(\underline{\omega}) = \alpha(\omega_1)(\varphi(\widetilde{\underline{\omega}}^{(1)}) - \varphi(\underline{\omega})) + \beta(\omega_N)(\varphi(\widetilde{\underline{\omega}}^{(N)}) - \varphi(\underline{\omega})), \quad (6.9)$$

where α and β are as in (6.6) and (6.7) respectively,

$$\widetilde{\omega}_j^{(1)} = \begin{cases} \omega_j & \text{if } j \neq 1 \\ \omega_1 + 1 & \text{if } j = 1 \end{cases}$$

and

$$\widetilde{\omega}_j^{(N)} = \begin{cases} \omega_j & \text{if } j \neq N \\ \omega_N - 1 & \text{if } j = N, \end{cases}$$

and the expectations are taken via the measure of $\underline{\omega}$. We assume that the existence of dynamics can be established on a set of tempered configurations $\widetilde{\Omega}$ (given by the possible occupation numbers defined in $\widetilde{\omega}$), and we have the usual properties of the semigroup and the generator acting on nice functions on this set. We also assume that $\widetilde{\Omega}$ is of full measure w.r.t. the stationary measures to be defined in Section 6.2.1. In this case the construction can be based on functional analytic properties of the infinitesimal generator and the Hille-Yosida theorem, see [84]. We note that the process described here is similar to that of the so called 2-exclusion process, differing from it, because we do not impose the condition $r_{1,0} = r_{1,1} = r_{2,0} = r_{2,1}$.

We define our system this way, because we seek to have its stationary distributions in product form, since that represents the realistic idea, that parts of the highway far away from each other are occupied independently. In Section 6.2.1 we investigate the conditions on the transition rates that provide such a distribution. We do this in order to obtain a relation between the local density and the expected net current in the system. The rescaled versions of these functions then provide a relation between the density and the flow similar to those in the various forms of the fundamental diagram, whilst giving a microscopic reasoning on a vehicle-

to-vehicle level, which also considers the effects of different inflow rates. We describe these results in Section 6.2.2, where we also validate our model by comparing it to a set of measured data.

6.2.1 Stationary distribution

As discussed, we look for a stationary distribution which is of the form

$$\mathcal{P}(\underline{\omega}) = \bigotimes_{i \in \mathcal{N}} \mu(\omega_i), \tag{6.10}$$

where $\mu(\omega_i)$ is of the form

$$\mu(\omega_i) = \begin{cases} 1/(1+\theta+\vartheta) & \text{if } \omega_i = 0 \\ \theta/(1+\theta+\vartheta) & \text{if } \omega_i = 1 \\ \vartheta/(1+\theta+\vartheta) & \text{if } \omega_i = 2. \end{cases} \tag{6.11}$$

Due to normalisation we can assume this form without loss of generality.

To obtain the conditions on the transition rates $r_{1,0}$, $r_{1,1}$, $r_{2,0}$, $r_{2,1}$, boundary rates α_0, α_1, β_1, β_2 and parameters θ, ϑ which provide us with such a stationary distribution, we need the following steps. Firstly, we equate the expectation of the generator of the system to zero,

$$0 = \mathbb{E}[L(\varphi(\underline{\omega}))] = \mathbb{E}\left[B\varphi(\underline{\omega}) + \sum_{i=1}^{N-1} r(\omega_i, \omega_{i+1})(\varphi(\widetilde{\underline{\omega}}^i) - \varphi(\underline{\omega})) \right], \tag{6.12}$$

as that is the condition of stationarity. Here we will apply a change of measure to be able to replace $\varphi(\underline{\omega})$ with $\varphi(\widetilde{\underline{\omega}}^i)$, allowing us to have the right side of (6.12) in a form, where under the expectation we have a sum multiplied by $\varphi(\widetilde{\underline{\omega}}^i)$. Stationarity cannot depend on the chosen test function, therefore it is obtained if the sum equals zero. As we want to allow the system to be of any finite size, in order to have the desired result we need the sum to be telescopic, i.e. to have all but the boundary terms cancel in the bulk. Then the nonzero terms left over can cancel with the terms coming from B, establishing the relations between the transition rates and the boundary rates.

Following the steps described above we rewrite part of the generator

as follows,

$$\begin{aligned}
&\mathbb{E}\left[\sum_{i=1}^{N-1} r(\omega_i,\omega_{i+1})(\varphi(\widetilde{\underline{\omega}}^i)-\varphi(\underline{\omega}))\right] = \\
&\quad \sum_{i=1}^{N-1}\Big\{\mathbb{E}[r(\widetilde{\omega}_i+1,\widetilde{\omega}_{i+1}-1)\varphi(\widetilde{\underline{\omega}}^i)]-\mathbb{E}[r(\omega_i,\omega_{i+1})\varphi(\underline{\omega})]\Big\} = \\
&\quad \sum_{i=1}^{N-1}\left\{\mathbb{E}\left[r(\widetilde{\omega}_i+1,\widetilde{\omega}_{i+1}-1)\varphi(\widetilde{\underline{\omega}})\frac{\mu(\widetilde{\omega}_i)\mu(\widetilde{\omega}_{i+1})}{\mu(\omega_i)\mu(\omega_{i+1})}\right]-\mathbb{E}[r(\widetilde{\omega}_i,\widetilde{\omega}_{i+1})\varphi(\widetilde{\underline{\omega}})]\right\} = \\
&\mathbb{E}\left[\sum_{i=1}^{N-1}\left(r(\widetilde{\omega}_i+1,\widetilde{\omega}_{i+1}-1)\frac{\mu(\omega_i)\mu(\omega_{i+1})}{\mu(\widetilde{\omega}_i)\mu(\widetilde{\omega}_{i+1})}-r(\widetilde{\omega}_i,\widetilde{\omega}_{i+1})\right)\cdot\varphi(\widetilde{\underline{\omega}})\right].
\end{aligned} \tag{6.13}$$

In the first expectation we needed to change variables to incorporate $\widetilde{\underline{\omega}}^i$ which we completed by multiplying with a Radon-Nykodym term. In the meantime we changed notation in the second term (from $\underline{\omega}$ to $\widetilde{\underline{\omega}}$) without actually changing variables. Since we want the generator to be zero no matter what the test function is, we look for solutions where the sum in (6.13) is telescopic and the remaining terms cancel out with the boundary terms. We can extend the summed term in (6.13) according to (6.5) and (6.11):

$$\begin{aligned}
&r(\widetilde{\omega}_i+1,\widetilde{\omega}_{i+1}-1)\frac{\mu(\widetilde{\omega}_i+1)\mu(\widetilde{\omega}_{i+1}-1)}{\mu(\widetilde{\omega}_i)\mu(\widetilde{\omega}_{i+1})}-r(\widetilde{\omega}_i,\widetilde{\omega}_{i+1}) = \\
&r_{1,0}\mathbb{I}\{\widetilde{\omega}_i=0,\widetilde{\omega}_{i+1}=1\}+r_{1,1}\mathbb{I}\{\widetilde{\omega}_i=0,\widetilde{\omega}_{i+1}=2\}\frac{\theta^2}{\vartheta}+ \\
&r_{2,0}\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=1\}\frac{\vartheta}{\theta^2}+r_{2,1}\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=2\}- \\
&r_{1,0}\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=0\}-r_{1,1}\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=1\}- \\
&r_{2,0}\mathbb{I}\{\widetilde{\omega}_i=2,\widetilde{\omega}_{i+1}=0\}-r_{2,1}\mathbb{I}\{\omega_i=2,\omega_{i+1}=1\} = \\
&r_{1,0}(\mathbb{I}\{\widetilde{\omega}_i=0,\widetilde{\omega}_{i+1}=1\}-\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=0\})+ \\
&r_{1,1}(\mathbb{I}\{\widetilde{\omega}_i=0,\widetilde{\omega}_{i+1}=2\}\frac{\vartheta}{\theta^2}-\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=1\})+ \\
&r_{2,0}(\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=1\}\frac{\theta^2}{\vartheta}-\mathbb{I}\{\widetilde{\omega}_i=2,\widetilde{\omega}_{i+1}=0\})+ \\
&r_{2,1}(\mathbb{I}\{\widetilde{\omega}_i=1,\widetilde{\omega}_{i+1}=2\}-\mathbb{I}\{\omega_i=2,\omega_{i+1}=1\}).
\end{aligned} \tag{6.14}$$

Since both $\widetilde{\omega}_i$ and $\widetilde{\omega}_{i+1}$ can only have values 0, 1 and 2, we may exchange the indicators in (6.14) with polynomials:

$$
\begin{aligned}
&r_{1,0}(\mathbb{I}\{\widetilde{\omega}_i = 0, \widetilde{\omega}_{i+1} = 1\} - \mathbb{I}\{\widetilde{\omega}_i = 1, \widetilde{\omega}_{i+1} = 0\}) + \\
&r_{1,1}(\mathbb{I}\{\widetilde{\omega}_i = 0, \widetilde{\omega}_{i+1} = 2\}\frac{\theta^2}{\vartheta} - \mathbb{I}\{\widetilde{\omega}_i = 1, \widetilde{\omega}_{i+1} = 1\}) + \\
&r_{2,0}(\mathbb{I}\{\widetilde{\omega}_i = 1, \widetilde{\omega}_{i+1} = 1\}\frac{\vartheta}{\theta^2} - \mathbb{I}\{\widetilde{\omega}_i = 2, \widetilde{\omega}_{i+1} = 0\}) + \\
&r_{2,1}(\mathbb{I}\{\widetilde{\omega}_i = 1, \widetilde{\omega}_{i+1} = 2\} - \mathbb{I}\{\omega_i = 2, \omega_{i+1} = 1\}) = \\
&r_{1,0}\left[-\frac{1}{2}(\widetilde{\omega}_i - 1)(\widetilde{\omega}_i - 2)\widetilde{\omega}_{i+1}(\widetilde{\omega}_{i+1} - 2) + \frac{1}{2}\widetilde{\omega}_i(\widetilde{\omega}_i - 2)(\widetilde{\omega}_{i+1} - 1)(\widetilde{\omega}_{i+1} - 2)\right] + \\
&r_{1,1}\left[\frac{\theta^2}{4\vartheta} \cdot (\widetilde{\omega}_i - 1)(\widetilde{\omega}_i - 2)\widetilde{\omega}_{i+1}(\widetilde{\omega}_{i+1} - 1) - \widetilde{\omega}_i(\widetilde{\omega}_i - 2)\widetilde{\omega}_{i+1}(\widetilde{\omega}_{i+1} - 2)\right] + \\
&r_{1,0}\left[\frac{\vartheta}{\theta^2} \cdot \widetilde{\omega}_i(\widetilde{\omega}_i - 2)\widetilde{\omega}_{i+1}(\widetilde{\omega}_{i+1} - 2) - \frac{1}{4}\widetilde{\omega}_i(\widetilde{\omega}_i - 1)(\widetilde{\omega}_{i+1} - 1)(\widetilde{\omega}_{i+1} - 2)\right] + \\
&r_{1,0}\left[-\frac{1}{2}\widetilde{\omega}_i(\widetilde{\omega}_i - 2)\widetilde{\omega}_{i+1}(\widetilde{\omega}_{i+1} - 1) + \frac{1}{2}\widetilde{\omega}_i(\widetilde{\omega}_i - 1)\widetilde{\omega}_{i+1}(\widetilde{\omega}_{i+1} - 2)\right]
\end{aligned}
\tag{6.15}
$$

We can reorder these terms to get the coefficients for $\widetilde{\omega}_i^2\widetilde{\omega}_{i+1}^2$, $\widetilde{\omega}_i^2\widetilde{\omega}_{i+1}$, $\widetilde{\omega}_i\widetilde{\omega}_{i+1}^2$, $\widetilde{\omega}_i\widetilde{\omega}_{i+1}$, $\widetilde{\omega}_i^2$, $\widetilde{\omega}_{i+1}^2$, $\widetilde{\omega}_i$ and $\widetilde{\omega}_{i+1}$, which we will denote by $A_1, , \ldots, A_8$ respectively. In order to have a telescopic sum, we need

$$A_1 = A_2 = A_3 = A_4 = 0, \tag{6.16}$$

$$A_5 = -A_6 \tag{6.17}$$

and

$$A_7 = -A_8. \tag{6.18}$$

By reordering (6.15) we get:

$$A_1 = \frac{\theta^2}{4\vartheta}\, r_{1,1} - r_{1,1} + \frac{\vartheta}{\theta^2}\, r_{2,0} - \frac{1}{4}\, r_{2,0}, \tag{6.19}$$

$$A_2 = -\frac{1}{2}\, r_{1,0} - \frac{\theta^2}{4\vartheta}\, r_{1,1} + 2r_{1,1} - \frac{2\vartheta}{\theta^2}\, r_{2,0} + \frac{3}{4}\, r_{2,0} - \frac{1}{2}\, r_{2,1}, \tag{6.20}$$

$$A_3 = \frac{1}{2}\, r_{1,0} - \frac{3\theta^2}{4\vartheta}\, r_{1,1} + 2r_{1,1} - \frac{2\vartheta}{\theta^2}\, r_{2,0} + \frac{1}{4}\, r_{2,0} + \frac{1}{2}\, r_{2,1}, \tag{6.21}$$

$$A_4 = \frac{3\theta^2}{4\vartheta}\, r_{1,1} - 4r_{1,1} + \frac{4\vartheta}{\theta^2}\, r_{2,0} - \frac{3}{4}\, r_{2,0}, \tag{6.22}$$

$$A_5 = r_{1,0} - \frac{1}{2}\, r_{2,0}, \tag{6.23}$$

$$A_6 = -r_{1,0} + \frac{\theta^2}{2\vartheta}\, r_{1,1}, \tag{6.24}$$

$$A_7 = -2r_{1,0} + \frac{1}{2}\, r_{2,0}, \tag{6.25}$$

$$A_8 = 2r_{1,0} - \frac{\theta^2}{2\vartheta}\, r_{1,1}. \tag{6.26}$$

Now notice that from (6.19) and (6.22)

$$A_4 - 3A_1 = \frac{\vartheta}{\theta^2}\, r_{2,0} - r_{1,1}.$$

Thus for (6.16) to hold, we need

$$r_{2,0} = \frac{\theta^2}{\vartheta}\, r_{1,1}. \tag{6.27}$$

If we substitute (6.27) into (6.20) and (6.21) we find that for (6.16) to hold, we need

$$r_{1,0} + r_{2,1} = r_{2,0}. \tag{6.28}$$

By substituting (6.27) into (6.23), (6.24), (6.25) and (6.26) we find that (6.17) and (6.18) automatically hold. Thus (6.27) and (6.28) are the necessary and sufficient conditions on the transition rates in order to

have a stationary product distribution. Substituting into (6.13) we get

$$\begin{aligned}
&\mathbb{E}\Bigg[\sum_{i=1}^{N-1}\Big(r(\widetilde{\omega}_i+1,\widetilde{\omega}_{i+1}-1)\frac{\mu(\omega_i)\mu(\omega_{i+1})}{\mu(\widetilde{\omega}_i)\mu(\widetilde{\omega}_{i+1})}-r(\widetilde{\omega}_i,\widetilde{\omega}_{i+1})\Big)\cdot\varphi(\underline{\widetilde{\omega}})\Bigg]=\\
&\mathbb{E}\Bigg[\sum_{i=1}^{N-1}\Big((\widetilde{\omega}_i^2-\widetilde{\omega}_{i+1}^2)(r_{1,0}-\frac{1}{2}r_{2,0})+(\widetilde{\omega}_i-\widetilde{\omega}_{i+1})(-2r_{1,0}+\frac{1}{2}r_{2,0})\Big)\varphi(\underline{\widetilde{\omega}})\Bigg]=\\
&\mathbb{E}\Bigg[\Big((\widetilde{\omega}_1^2-\widetilde{\omega}_N^2)(r_{1,0}-\frac{1}{2}r_{2,0})+(\widetilde{\omega}_1-\widetilde{\omega}_N)(-2r_{1,0}+\frac{1}{2}r_{2,0})\Big)\varphi(\underline{\widetilde{\omega}})\Bigg].
\end{aligned}\tag{6.29}$$

Furthermore, we also need sufficient boundary conditions such that the boundary terms and the remaining term from the telescopic sum cancel out. To obtain these we take the form of the boundary term in (6.9) and repeat the steps we used by the transition rates in the bulk for the left and the right term separately:

$$\begin{aligned}
&\mathbb{E}\Big[\alpha(\omega_1)\big(\varphi(\underline{\widetilde{\omega}}^{(1)})-\varphi(\underline{\omega})\big)\Big]=\\
&\mathbb{E}\Big[\big(\alpha(\widetilde{\omega}_1-1)\cdot\frac{\mu(\widetilde{\omega}_1-1)}{\mu(\widetilde{\omega}_1)}-\alpha(\widetilde{\omega}_1)\big)\varphi(\underline{\widetilde{\omega}})\Big]=\\
&\mathbb{E}\Big[\big(\alpha_0\mathbb{I}\{\widetilde{\omega}_1=1\}\frac{1}{\theta}-\alpha_0\mathbb{I}\{\widetilde{\omega}_1=0\}+\\
&\qquad\alpha_1\mathbb{I}\{\widetilde{\omega}_1=2\}\frac{\theta}{\vartheta}-\alpha_1\mathbb{I}\{\widetilde{\omega}_1=1\}\big)\varphi(\underline{\widetilde{\omega}})\Big]=\\
&\mathbb{E}\Big[\big(-\alpha_0\frac{1}{\theta}\widetilde{\omega}_1(\widetilde{\omega}_1-2)-\alpha_0\frac{1}{2}(\widetilde{\omega}_1-1)(\widetilde{\omega}_1-2)+\\
&\qquad\alpha_1\frac{\theta}{2\vartheta}\widetilde{\omega}_1(\widetilde{\omega}_1-1)+\alpha_1\widetilde{\omega}_1(\widetilde{\omega}_1-2)\big)\varphi(\underline{\widetilde{\omega}})\Big]=\\
&\mathbb{E}\Big[(\widetilde{\omega}_1^2(-\frac{1}{\theta}\alpha_0-\frac{1}{2}\alpha_0+\frac{\theta}{2\vartheta}\alpha_1+\alpha_1)+\\
&\qquad\widetilde{\omega}_1(\frac{2}{\theta}\alpha_0+\frac{3}{2}\alpha_0-\frac{\theta}{2\vartheta}\alpha_1-2\alpha_1)-\alpha_0)\varphi(\underline{\widetilde{\omega}})\Big]
\end{aligned}\tag{6.30}$$

and

$$
\begin{aligned}
&\mathbb{E}\Big[\beta(\omega_N)\big(\varphi(\widetilde{\underline{\omega}}^{(N)}) - \varphi(\underline{\omega})\big)\Big] = \\
&\mathbb{E}\Big[\big(\beta(\widetilde{\omega}_N + 1)\cdot \frac{\mu(\widetilde{\omega}_N + 1)}{\mu(\widetilde{\omega}_N)} - \beta(\widetilde{\omega}_N)\big)\varphi(\widetilde{\underline{\omega}})\Big] = \\
&\mathbb{E}\Big[\big(\beta_1 \mathbb{I}\{\widetilde{\omega}_N = 0\}\theta - \beta_1 \mathbb{I}\{\widetilde{\omega}_N = 1\} + \\
&\qquad \beta_2 \mathbb{I}\{\widetilde{\omega}_N = 1\}\frac{\vartheta}{\theta} - \beta_2 \mathbb{I}\{\widetilde{\omega}_N = 2\}\big)\varphi(\widetilde{\underline{\omega}})\Big] = \\
&\mathbb{E}\Big[\big(\beta_1 \frac{\theta}{2}(\widetilde{\omega}_N - 1)(\widetilde{\omega}_N - 2) + \alpha_0 \widetilde{\omega}_N(\widetilde{\omega}_N - 2) - \\
&\qquad \beta_2 \frac{\vartheta}{\theta}\widetilde{\omega}_N(\widetilde{\omega}_N - 2) - \beta_2 \frac{1}{2}\widetilde{\omega}_N(\widetilde{\omega}_N - 1)\big)\varphi(\widetilde{\underline{\omega}})\Big] = \\
&\mathbb{E}\Big[\big(\widetilde{\omega}_N^2(\frac{\theta}{2}\beta_1 + \beta_1 - \frac{\vartheta}{\theta}\beta_2 - \frac{1}{2}\beta_2) + \\
&\qquad \widetilde{\omega}_N(-\frac{3\theta}{2}\beta_1 - 2\beta_1 + \frac{2\vartheta}{\theta}\beta_2 + \frac{1}{2}\beta_2) + \theta\beta_1\big)\varphi(\widetilde{\underline{\omega}})\Big].
\end{aligned}
\tag{6.31}
$$

If we substitute (6.29), (6.30) and (6.31) into (6.8) we get

$$
\begin{aligned}
\mathbb{E}[L(\varphi(\widetilde{\underline{\omega}}))] = \mathbb{E}\Big[\big(&- \alpha_0 + \theta\beta_1 + \\
&\widetilde{\omega}_1^2(-\frac{1}{\theta}\alpha_0 - \frac{1}{2}\alpha_0 + \frac{\theta}{2\vartheta}\alpha_1 + \alpha_1 + r_{1,0} - \frac{1}{2}r_{2,0}) + \\
&\widetilde{\omega}_1(\frac{2}{\theta}\alpha_0 + \frac{3}{2}\alpha_0 - \frac{\theta}{2\vartheta}\alpha_1 - 2\alpha_1 - 2r_{1,0} + \frac{1}{2}r_{2,0}) + \\
&\widetilde{\omega}_N^2(\frac{\theta}{2}\beta_1 + \beta_1 - \frac{\vartheta}{\theta}\beta_2 - \frac{1}{2}\beta_2 - r_{1,0} + \frac{1}{2}r_{2,0}) + \\
&\widetilde{\omega}_N(-\frac{3\theta}{2}\beta_1 - 2\beta_1 + \frac{2\vartheta}{\theta}\beta_2 + \frac{1}{2}\beta_2 + 2r_{1,0} - \frac{1}{2}r_{2,0})\big)\varphi(\widetilde{\underline{\omega}})\Big].
\end{aligned}
\tag{6.32}
$$

Thus the conditions on the boundary rates are given by setting all the coefficients in (6.32) to zero. It can be seen easily that these are equivalent with

$$
\alpha_0 = \theta\beta_1, \tag{6.33}
$$

$$\alpha_1 = \frac{\vartheta}{\theta}\,\beta_2, \tag{6.34}$$

$$\alpha_0\left(\frac{1}{\theta}+1\right) - \alpha_1 = r_{1,0}, \tag{6.35}$$

and

$$\alpha_0 + \frac{\theta}{\vartheta}\,\alpha_1 = r_{2,0}. \tag{6.36}$$

Altogether (6.27), (6.28), (6.33)-(6.36) give all the necessary and sufficient conditions to have a stationary distribution as in (6.10). This also shows that the parameters θ, ϑ, $r_{1,0}$, $r_{2,1}$ and α_0 define exactly one stationary solution of the system.

6.2.2 Fluid dynamics

In this section we investigate the functions that describe the movement of vehicles along the motorway. Firstly, we discuss the microscopic level, then we apply results from the literature to determine the macroscopic, rescaled counterparts of these functions. These bare a direct relation to the quantities considered in the fundamental diagram, thus they also provide means to numerically verify our model.

Firstly, we define the local density as

$$\rho := \mathbb{E}\omega_i = \frac{\theta + 2\vartheta}{1+\theta+\vartheta}, \tag{6.37}$$

where we introduce the canonical sum

$$Z := 1 + \theta + \vartheta. \tag{6.38}$$

We consider Z as a function of ρ. It has to be noted that $Z(0) = 1$ and $Z(\rho) \to \infty$ as $\rho \to 2$. With (6.37) and (6.38) we have two equations on the parameters θ and ϑ, thus they can be expressed as functions of ρ and Z as

$$\theta = (2-\rho)\cdot Z(\rho) - 2, \tag{6.39}$$

$$\vartheta = (\rho-1)\cdot Z(\rho) + 1. \tag{6.40}$$

Now, we can express the expected net current as a function of the local density. By definition it is given as

$$\mathcal{J} := \mathbb{E}r(\omega_i, \omega_{i+1}) = \frac{1}{Z(\rho)^2}(\theta\, r_{1,0} + \theta^2\, r_{1,1} + \vartheta\, r_{2,0} + \vartheta\theta\, r_{2,1}),$$

which can be extended by (6.27), (6.28), (6.39) and (6.40) to

$$\begin{aligned}
\mathcal{J} &= \frac{\vartheta}{Z(\rho)^2}\left(\frac{\theta}{\vartheta}\, r_{1,0} + \frac{\theta^2}{\vartheta}\, r_{1,1} + r_{2,0} + \theta\, r_{2,1}\right) \\
&= \frac{\vartheta}{Z(\rho)^2}\left(\frac{\theta}{\vartheta}\, r_{1,0} + 2r_{2,0} + \theta\, r_{2,1}\right) \\
&= \frac{\vartheta}{Z(\rho)^2}\left(\frac{\theta}{\vartheta}\, r_{1,0} + 2(r_{1,0} + r_{2,1}) + \theta\, r_{2,1}\right) \\
&= \frac{\vartheta}{Z(\rho)^2}\left(\frac{\theta + 2\vartheta}{\vartheta}\, r_{1,0} + (\theta + 2)\, r_{2,1}\right) \\
&= \frac{\rho}{Z(\rho)}\, r_{1,0} + \frac{\vartheta(\theta + 2)}{Z(\rho)^2}\, r_{2,1} \\
&= \frac{\rho}{Z(\rho)}\, r_{1,0} + \frac{\vartheta(2 - \rho)Z(\rho)}{Z(\rho)^2}\, r_{2,1}.
\end{aligned}$$

At this point, we note that choosing the proper function $Z(\rho)$ and parameters $r_{1,0}$, $r_{2,1}$, α_0 and α_1 gives exactly one solution for ρ by (6.35) and (6.36) and thus determine a stationary point on the $\mathcal{J}(\rho)$ curve. Namely, from (6.35) and (6.36) we get

$$\mathcal{J}(\rho) = \frac{\rho}{Z(\rho)}\, r_{1,0} + \frac{(2 - \rho)Z(\rho)}{Z(\rho)^2}\, \frac{\alpha_0\alpha_1 r_{2,1}}{(r_{1,0} + \alpha_1 - \alpha_0)(r_{1,0} + r_{2,1} - \alpha_0)}. \tag{6.41}$$

One can obtain the macroscopic counterparts of the microscopic functions by applying fluid scaling (also referred to as Eulerian scaling or hydrodynamic limit in the literature). In vague terms this could be done by considering versions $\mathcal{P}^M$ of the stationary measure in (6.10), where both the spatial terms (given by the numbering of the sites i) and the time-related terms (appearing indirectly in the generator) are scaled by M^{-1}. This naturally requires the extension of the system size N, which is possible. Following results of the literature we can say that our model

belongs to a family of models satisfying a conservation law of the form

$$\partial_t k(t,x) + \partial_x q(k(t,x)) = 0, \quad t \geq 0, \quad x \in \mathbb{R},$$

where t and x are the rescaled time and space variables, and k and q denote the density and flow as usual, functioning as the macroscopic counterparts of ρ and $\mathcal{J}$ respectively. For more details on the hydrodynamics of this family of models we direct the reader to the works of Rezakhanlou [124] and Bahadoran [10]. We continue our discussion by comparing data measured on a British highway with the fundamental diagram given by our model with a specific choice of the $Z(\rho)$ function.

6.2.3 Statistical analysis

In this section we present the results of an initial numerical analysis, which aims to verify our model and provide a better understanding for its parameters. Namely, we use a least mean square method to fit the parameters to a set of data obtained by Highways England, an agency of the Department for Transport of the Government of the United Kingdom [62]. The measurements were taken at Junction 9 on the British M25 highway in June 2013.

As mentioned in Section 6.2.2 we need to choose the function $Z(\rho)$ to fix the relation between the local density and the net current. We do this by setting

$$\vartheta = \kappa\theta^2, \tag{6.42}$$

where we introduced the parameter κ, representing the idea that θ describes the probability of a vehicle being present at a short piece of road, whilst κ gives the correlation between multiple vehicles. This fixes $Z(\rho)$ as $Z(\rho)(\theta) = 1 + \theta + \kappa\theta^2$. Furthermore (6.42) allows us to rewrite the conditions in (6.27), (6.34) and (6.36) as

$$\kappa r_{2,0} = r_{1,1},$$

$$\alpha_0 = \kappa\theta\beta_2,$$

and

$$r_{2,0} = \alpha_0 + \frac{1}{\kappa\theta}\alpha_1$$

respectively. This way α_0 (or α_1 as they can be expressed as functions of each other due to (6.35) and (6.36)) gives the density for any fixed triplet of $\{\kappa, r_{1,0}, r_{2,1}\}$, introducing a direct correlation between the arrival rate at the ingress and the density on the motorway segment.

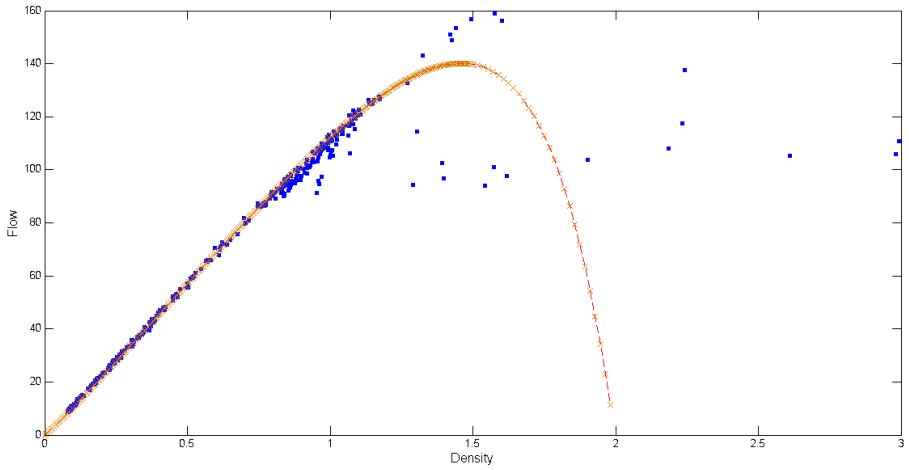

Figure 6.2. Fitted curve for density-flow relation.

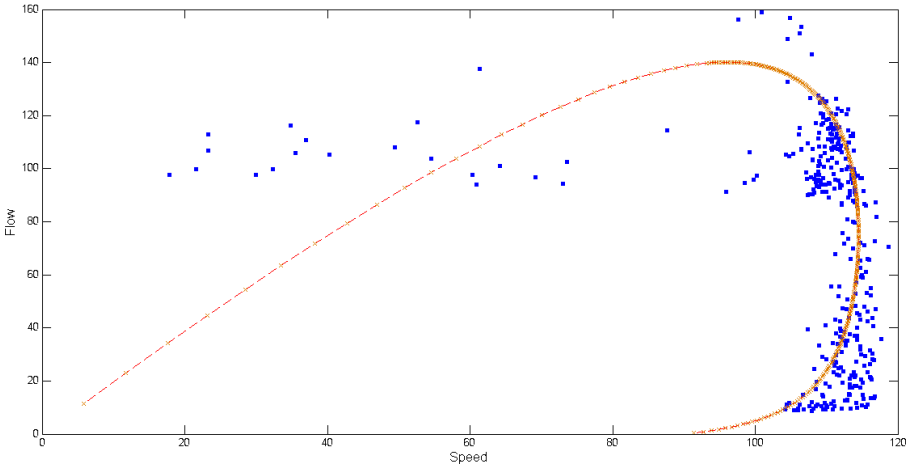

Figure 6.3. Fitted curve for speed-flow relation.

The set of data that we used contained information on the length of

the motorway segment, the number of vehicles passig by every hour and their average speed. In order to compare our model's predictions with the measurement results we had to determine local densities. This was done by considering an average headway by which we divided the motorway segment into slots. The local density was then given by dividing the hourly flow with the average speed and the number of slots. The values taken for the average headway were between 120 and 150 meters. The parameter fitting was done by the Levenberg-Marquardt algortihm [103] in MATLAB. Typical results for the density-flow and the speed-flow relations are shown in Figure 6.2 and Figure 6.3 respectively. In the parameter sets, that provided the best fit, κ was between 2.5 and 3, whilst the typical values for $r_{1,0}$ and $r_{2,1}$ were around 110 and 70 respectively. This suggests a connection between $r_{1,0}$ and the average speed of the free flow regime and between $r_{2,1}$ and the average speed before the system collapses due to congestion. However, the measurements mostly provided data from freeflowing traffic, thus a more extensive study would be required.

6.2.4 Conclusion

This chapter presented a novel single-regime traffic flow model, that provided a probabilistic description of its dynamics on the vehicle-to-vehicle level. By establishing some natural requirements of the system (e.g. the dynamics do not depend on the length of the motorway segment) and its stationary measure we have determined several conditions connecting the parameters of the model. This allowed us to obtain a simple density-flow function on the microscopic level, given by (6.41). By fluid scaling this can be extended to the macroscopic level, showing that our model belongs to a well-studied family of models with Burgers-type dynamics. Furthermore we have also compared the form of the fundamental diagram given by our model to a set of data collected on a British highway. The initial results show a nice fit and suggest natural explanations for the system parameters, however further research should be conducted to obtain conclusive results.

Besides an extensive statistical study there are other directions in which this research should be continued. The model can be naturally extended to include more than two lanes of traffic. Another future goal is

providing an analysis of the system's transient evolution, which may serve as a basis for predicting congestion. This would ultimately provide means for efficient traffic control for motorways, a subject that is discussed in a different modeling environment in Chapter 7.

Chapter 7

Traffic control for ramps leading to motorways

Outline. This chapter investigates the points where urban and interurban traffic flows meet, namely the ramps where motorways can be accessed from urban roads. The task of the controller is multifaceted here, as both the specifics of motorway traffic and the queueing dynamics of the surrounding urban roads need to be considered. Thus we propose joint control policies, which apply the ideas presented in Chapters 4 and 5 in this different environment. For most of the chapter we take a rather simple approach in modeling the evolution of traffic flows on the motorway, but we also discuss a scheme, which takes the complex behaviour represented by the fundamental diagram into account.

The chapter is organised as follows. In Section 7.1 we introduce the basic model that describes the evolution of the system that we investigate. Section 7.2 then presents two control methods reprising the notions of pressure, on which the scheme presented in Chapter 4 was built, and fairness, which served as the core idea for Chapter 5. We discuss the ensuing optimisation problems and show methods which provide solutions to them allowing for simple implementation. Section 7.3 contains a a simulation study, which compares the two policies in terms of performance, highlighting the scenarios in which their utilisation is advantegeous. The first part of the chapter takes a rather naive approach in modeling the evolution of traffic flows on the motorway. We propose a model and a control scheme that considers a more complex behaviour in Section 7.4, discussing the challenges presented therein. Finally, we conclude the chapter in Section 7.5.

7.1 Basic model

We consider a motorway connecting urban areas. One of the goals in traffic management is to preserve the free flow of vehicles on the motorway and thus, where possible, avoid traffic jams. In this case this is vital, as a traffic collapse on the motorway causes a breakdown in the surrounding areas as well. The connection to the urban network is modelled in our system by ramps leading traffic on and off the motorway, that join with secondary urban roads. For simplicity we consider one-direction of a motorway, as depicted in Figure 7.1. The system consists of a set of segments, or sections, denoted by $\mathcal{L} = \{1, 2, \ldots, L\}$. We assume that there is an on-ramp at the beginning of each segment. The traffic on a segment depends on both the traffic joining from the previous segment, and from the traffic released from the on-ramp at the ingress of the segment. With each on-ramp $l \in \mathcal{L}$, we associate a queue, $M_l(t)$, which represents the number of vehicles waiting at the ramp. For the on-ramps we also define a rate, Λ_l, at which the vehicles leave the ramp and join the motorway. This represents the traffic volume multiplied by the speed of the traffic entering the motorway. We assume that Λ_l can be controlled by the traffic controller. For the off-ramps, we define the proportion of vehicles leaving the motorway from an off-ramp, p_l for $1 \leq l \leq L$, which might change with time, but is not a control decision. From this we can obtain the flow on each section of the motorway. Letting Γ_l denote the traffic flow on section l of the motorway, we have the following relation,

$$\Gamma_l = \Lambda_l + (1 - p_l)\Gamma_{l-1}, \tag{7.1}$$

for $l = 1, ..., L$ and we let $\Lambda_0 = \Gamma_0$ for the first section on the road segment $l = 1$. We assume each motorway segment has a maximum capacity C_l. If Γ_l, the flow on segment l, exceeds this rate then the motorway will go into a congested state, and this will lead to significant delays to drivers. Thus, we model this as a hard constraint that must at all times be satsified by the traffic controller. We ackonwledge that this is a rather naive approach due to the complex nature of traffic flow represented by the fundamental diagram as we discussed in Chapter 6. We address this question in Section 7.4.

The connection of the road segments to the urban network is established

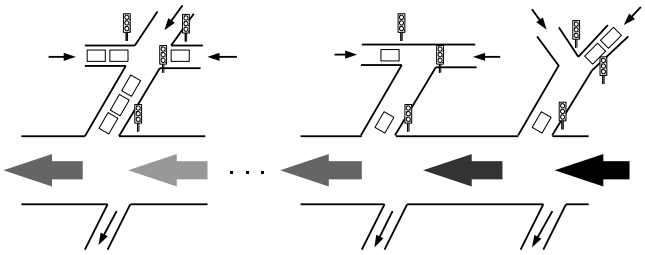

Figure 7.1. Motorway with on- and off-ramps and junctions with secondary roads at the end of each on-ramp.

as follows. Each road segment l has an on-ramp, and at the end of the ramp there is a junction. At each junction there are a number of adjacent secondary roads. We assume that the ramps are sufficiently far from each other, such that the traffic entering these junctions are independent from each other, i.e. the urban surroundings are not further connected. The structure of the urban network follows the earlier parts of the thesis. We define a set of in-roads $\mathcal{I}$, which represent one or more lanes of traffic at the junctions at the start of each on-ramp. An in-road, i, being present at a junction, l, will be denoted by the inclusion $i \in l$. We define the queue sizes $Q_i(t)$ for each $i \in \mathcal{I}$. As we do not intend to model farther parts of the urban network, we assume that the in-roads have infinite capacity, which in reality transfers to the idea that the vehicles in queue wait on subsequent secondary roads if the in-road itself would be full. We introduce the service rates $\sigma_i(t)$ for every $i \in l$, representing the amount of service recieved by a junction i, during time period t. During each time period, at every junction, the traffic controller allocates a proportion of green time to the different roads leading into the on-ramps. If T_l is the maximum permissable green time allocated to the cars at the junction, then the rates σ_i must satisfy the contraints

$$\sum_{i \in l} \sigma_i \leq T_l, \tag{7.2}$$

for all $l \in \mathcal{L}$.

This describes the basic primitives that we require to measure and control traffic within this framework. We now discuss the schemes for

joint control of motorways and urban road traffic.

7.2 Joint control schemes

We consider two decentralised control proceedures for the control of congestion on a motorway, its on-ramps and the surrounding roads. In particular, we define a time-slotted scheme governed by a proportional fair and a backpressure-based controlling mechanism. For a discussion on the background of these notions we refer the reader to Chapter 3.

7.2.1 Proportional fair control

Proportional fair control represents the idea that everyone involved in the traffic should recieve the same service, i.e. the available capacities should be distributed in a way that is proportional to the flows on each link. It is a natural extension of Nash's bargaining solution and can easily described by an optimisation problem. Here, we describe how the concept can be applied for joint control at urban intersection connections of on-ramps and within the motorway segments. To do so, we consider the following proportional fair optimisation problem,

$$\begin{aligned}
&\text{maximise} && \sum_{i\in\mathcal{I}} Q_i \log \sigma_i + \sum_{l\in\mathcal{L}} M_l \log \Lambda_l && && (7.3)\\
&\text{subject to} && \sum_{i\in l} \sigma_i \le T_l, && && (7.4)\\
& && \Gamma_l \le C_l, && l\in\mathcal{L}, && (7.5)\\
& && \Gamma_l = (1-p_l)\Gamma_{l-1} + \Lambda_l, && l\in\mathcal{L}, && (7.6)\\
&\text{over} && \Lambda_l \ge 0, \Gamma_l, \sigma_i \ge 0, && && (7.7)
\end{aligned}$$

with the restriction that if $Q_i = 0$, then $\sigma_i = 0$, and similarly if $M_l = 0$, then $\Lambda_l = 0$ is needed.

Although the control problem may look difficult to solve, provided that certain queue length and flow measurements are feasible, then a distributed implementation is possible. In particular, at each junction

prior to an on-ramp, one can solve the following optimisation:

$$\text{maximise} \qquad \sum_{i\in l} Q_i \log \sigma_i \tag{7.8}$$

$$\text{subject to} \qquad \sum_{i\in l} \sigma_i \leq T_l, \tag{7.9}$$

where a quick calculation yields a solution,

$$\sigma_i = \frac{Q_i}{\sum_l Q_i}. \tag{7.10}$$

In other words, one should serve each in-road at a junction for a time that is proportional to the number of cars that are present at the in-road. Then, along each junction, we optimise the function

$$\text{maximise} \qquad \sum_{l\in\mathcal{L}} M_l \log \Lambda_l \tag{7.11}$$

$$\text{subject to} \qquad \Lambda_l + (1-p_l)\Gamma_l \leq C_l, \qquad l \in \mathcal{L}, \tag{7.12}$$

$$\text{over} \qquad \Lambda_l \geq 0. \tag{7.13}$$

We note, as observed in [50], that if a junction l is a bottleneck, i.e.

$$\sum_{l'\in\mathcal{L}} \Lambda_{l'} A_{l'l} = C_l, \tag{7.14}$$

where $A_{l'l} = \prod_{k=l}^{l'-1}(1-p_k)$ denotes the proportion of traffic from on-ramp l' that passes through motorway section l, then the solution to the above optimisation problem is such that

$$\Lambda_{l'} A_{l'l} = C_l \frac{m_{l'}}{\sum_{k\leq l} m_k}. \tag{7.15}$$

In other words, the proportion of the flow out of on-ramp l, that passes through the bottleneck C_l, is proportional to the amount of traffic on the on-ramp. Notice that the smaller the proportion of traffic that the downstream flow places on the bottleneck junction, the more the flow that can be released from the downstream flow. This is different from the distribution achieved by a waterfilling procedure applied to the

constraints (7.5-7.6).

Although closed form solutions to the optimisation problem (7.11-7.13) are not always possible, it is straightforward to optimise its dual, which corresponds to the following optimisation:

$$\text{minimise} \qquad \sum_{l\in\mathcal{L}} C_l q_l - \sum_{l\in\mathcal{L}} M_l \log\Big(\sum_{l'\in\mathcal{L}} A_{ll'} q_{l'}\Big) \tag{7.16}$$

$$\text{over} \qquad q_l \geq 0, \quad j \in \mathcal{L}. \tag{7.17}$$

Since the above optimisation is concave and over a product space, the optimisation can be resolved by the following distributed gradient descent iteration on dual variables $q = (q_l : l \in \mathcal{L})$,

$$q_l' = \Big[q_l - \gamma\Big(C_l - \sum_{l'\in\mathcal{L}} \Lambda_{l'}(q) A_{l'l}\Big)\Big]_+, \tag{7.18}$$

where γ is a control gain, and $\Lambda_l(q)$ is the rate associated with q as expressed by the relation

$$\Lambda_l(q) = \frac{Q_l}{\sum_{l'\in\mathcal{L}} A_{ll'} q_{l'}}. \tag{7.19}$$

We could replace the parameterised $\Lambda(q)$ above with the observed rate of traffic at the motorway section, and thus arrive at a distributed implementation of the proportional fair control.

7.2.2 Backpressure control

An alternative adaptive control policy is the backpressure policy, based on the idea first described by Tassiulas and Ephremides [138]. The core concept of backpressure schemes lies in distributing the load (or "pressure") on different parts of the network in an efficient way. Namely, the policy favors the flows which put less pressure on the queues down the road and allocates less service to those which would send more vehicles to a congested part of the network. The adaptation to our examined control problem can be given through solving the following maximisation

problem,

$$\begin{aligned}
&\text{maximise} && \sum_{l\in\mathcal{L}}\left(\sum_{i\in l}(Q_i - M_l)\sigma_i + M_l\Lambda_l\right) && && (7.20)\\
&\text{subject to} && \sum_{i\in l}\sigma_i \le T_l, && && (7.21)\\
& && \Gamma_l \le C_l, && l\in\mathcal{L}, && (7.22)\\
& && \Gamma_l = (1-p_l)\Gamma_{l-1} + \Lambda_l, && l\in\mathcal{L}, && (7.23)\\
&\text{over} && \Lambda_l \ge 0, \Gamma_l, \sigma_i \ge 0, && && (7.24)
\end{aligned}$$

with the restriction that if $Q_i = 0$, then $\sigma_i = 0$, and similarly if $M_l = 0$, then $\Lambda_l = 0$ is needed. This optimisation can be separated as follows. For each junction of on-ramp l, serve in-road $i \in l$ if

$$Q_i = \max_{i'} Q_{i'} > M_l, \tag{7.25}$$

otherwise, do nothing. For the motorway, we have the following linear program,

$$\begin{aligned}
&\text{maximise} && \sum_{l\in\mathcal{L}} M_l\Lambda_l && && (7.26)\\
&\text{subject to} && \Gamma_l \le C_l, && l\in\mathcal{L}, && (7.27)\\
& && \Gamma_l = (1-p_l)\Gamma_{l-1} + \Lambda_l, && l\in\mathcal{L}, && (7.28)\\
&\text{over} && \Lambda_l \ge 0, \Gamma_l \ge 0. && && (7.29)
\end{aligned}$$

Standard procedures such as the Simplex algorithm will solve this optimisation in a straighforward manner. We note that this is the same method that is referred to as the greedy algorithm in [50]. We extend the idea of the backpressure scheme to the control of the on-ramps in Section 7.4.

7.3 Performance evaluation

This section presents a simulation study, which compares the performance of the proportional fair and the backpressure control policies. The

numerical experiments were implemented in the MATLAB environment.

We investigate two network settings: one with ramps that are 2 km apart from each other, and one where they are 6 km apart. In both cases, the network consists of 10 segments with on- and off-ramps at the beginning of each segment. At each on-ramp there is a set of adjacent in-roads, the maximal number of the in-roads at a junction is 4. During the simulations we keep track of the number of vehicles on each motorway segment, on-ramp and in-road. These numbers are updated according to the time that has passed, the flow on the motorway, and the rates σ, Λ that are assigned. To simplify our setting, we use cycles with a fixed length of 1 min, a fixed speed on the motorway of 120 km/h (thus assuming that the system remains in the free flow regime), and linear relations between the density on the motorway segments and the flow, and between the vehicles leaving the in-roads and the green time assigned to them. Furthermore, we did not take into account the length of the on-ramps, the time it takes for vehicles on them to speed up, or the possible spill-back effects in case they become jammed. We assume that arrivals from the preceding part of the motorway, and also at the queues on the adjacent in-roads, follow a Poisson process, where we leave the rate of the process as a simulation parameter. We also assume that every vehicle on the last segment of the motorway can leave the network in an uninterrupted manner.

We investigate the above network settings under the following set of parameters. For both the shorter and the longer segment lengths we conduct simulations for 200 cycles, i.e. for 200 minute-long time periods. We take 4 different sets of arrival rates. We call the arrival rates that represent the boundary of the network capacity as peak rates. With this conevntion the 4 examined setes of arrival rates are enlisted as follows,

- Peak rates on both the motorway and the in-roads
- Peak rate on the motorway and 60% of the peak rate on the in-roads
- Peak rates on the in-roads and 60% of the peak rate on the motorway
- 60% of the peak rates everywhere.

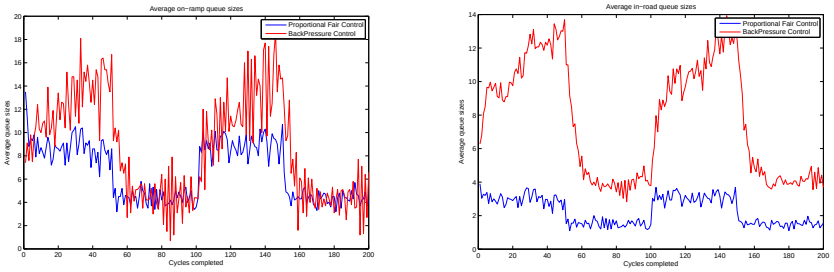

Figure 7.2. Queue sizes on the on-ramps and on the adjacent in-roads with peak inflow rate on the motorway, and 60% of the peak inflow rate at the in-roads with 6 km long motorway segments.

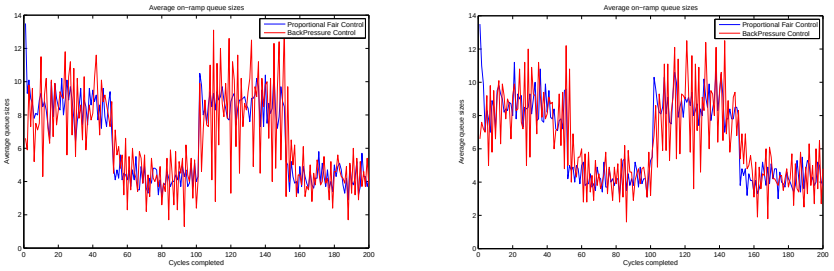

Figure 7.3. Queue sizes on the on-ramps with low inflow rates at various road segment lengths.

In each of these cases, we simulate 50 cycles with the given arrival rates, 50 cycles with half of the given arrival rates, again 50 cycles with the given rates, and again 50 cycles with half the rates. In our simulations the proportionally fair and the backpressure control policies received the same initial settings and the same number of arriving vehicles throughout the given time period.

The results show that the proportional fair controller outperforms backpressure with regards to queue sizes on both the on-ramps and on the adjacent in-roads in almost every scenario; see Figure 7.2. The only exception can be seen with low inflow rates, where the fluctuation of the queue sizes measured under backpressure control is higher, thus resulting in unfavourable behaviour; see Figure 7.3. On the other hand, the

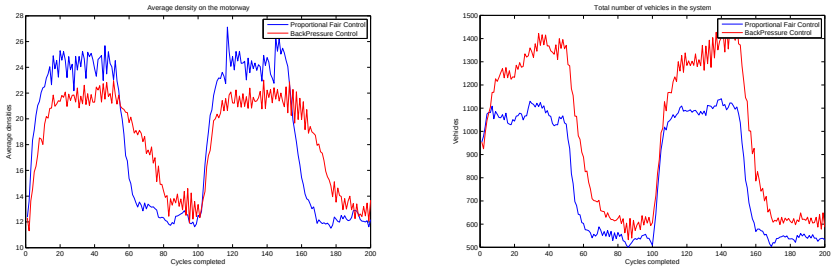

Figure 7.4. Motorway density and total vehicle number in the system at peak inflow rates.

backpressure policy keeps the average density throughout the motorway lower, a phenomenon which is even more emphatic at higher rates. This is a desirable property since it helps avoiding congested states, specially during peak hours. In terms of the total number of vehicles in the network, the proportional fair control performs better once again, see Figure 7.4, which translates into shorter journey times for individual motorists. The observed results do not seem to be influenced significantly by the lengths of the motorway segments, although intuitively longer segments could reduce the effect of traffic control.

In summary, our investigation suggests that the proportional fair controller is more efficient in most cases. However, if the traffic is close to congestion levels the backpressure policy is more effective in avoiding a traffic jam on the motorway. The results are not completely unexpected as some of them represent the basic concept of these schemes. However, there are phenomena, such as spill-over effects or accidents, which were not considered in the simulation, that could change these results. Most importantly the complexity of the evolution of the traffic flow on the motorway has to be considered, a problem we aim to address in the following section.

7.4 The box method

In this section we discuss a control policy for the motorway part that utilises the knowledge of the complex behaviour of the traffic flows

represented by the fundamental diagram[1], whilst extending the idea of the backpressure algorithm. The main idea is based on the observation, that because of the nonlinear dependence between the flow rate and the average speed of the vehicles or the flow rate and the density of the cars, direct control of the flow is not possible by controlling only the speed or the number of the joining vehicles. Thus a complex scheme which controls both is needed.

Modelling assumptions

We assume the structure of the motorway is the same as described in Section 7.1, with the exception that here the number of vehicles entering the motorway from the on-ramps can be controlled (by traffic lights for example). Furthermore, we allow for setting speed limits dynamically by signal boards at the beginning of each segment, after traffic joined in from the on-ramp. These speed limits thus effect both vehicles joining from the previous motorway segment and from the on-ramps.

As one can deduct from the fundamental diagram, controlling speed alone is not sufficient as lowering the speed limit may result in increasing density and therefore more congestion. On the other hand, controlling the rate at which external traffic joins the motorway through the on-ramps does not directly translate to controlling the density on the segments of the motorway, since the speed at which the cars pass by determines when they pass from one segment to the other. This, combined with the fact, that the critical density at which traffic starts to collapse is hard to determine, makes density control a hard task. However, in order to achieve efficient use of the motorway, one should aim to control the flow on the segments, since maximising that would mean throughput optimality. Experiments (upon which the fundamental diagram is built) show that there is a maximal flow that can pass a road segment, let us denote it by $q^{\max}$. Our proposed algorithm aims at maintaining the flow at this level. For simplicity we assume that the motorway is homogeneous in the sense, that the maximal flow is constant over the segments.

[1]For a discussion we refer the reader to Section 6.1.

The control policy

To help the explanation of the control method, we introduce artificial entities into the system that we will call boxes, hence the name. They represent gaps that can be filled by vehicles, we will not allow more than one vehicle per box. We allow these boxes to change their size, denoted by x_l for segment l, according to the speed limits, but we fix their flow to equal $q^{\max}$ at all times. This means that if the speed limit is set to v_l for segment l, then

$$x_l = \frac{v_l}{q^{\max}}. \tag{7.30}$$

As we know from the fundamental diagram, $q^{\max}$ can be reached at a given network speed, which we will denote by v^{fund}. We use the notation $v_l^{\max}$ for the highest possible speed limit over a segment. We assume both v^{fund} and $v_l^{\max}$ to be positive multiples of a granularity v_g and allow our speed limits to take values between them that are also positive multiples of v_g. Our goal will be setting speed limits and contolling the traffic lights at the on-ramps such that the boxes have sizes and are "filled" with vehicles in a manner that corresponds to optimal throughput. Below we discuss a control policy that produces such results.

We consider control decisions to be made in a cyclic manner with a common cycle time T at all intersections. We note that, since the flow of boxes is fixed at $q^{\max}$, in each cycle there are exactly $q^{\max}T$ boxes arriving at a segment independently from the average speed on thethe previous segment. We assume that the number of cars that are going to join a road segment from the previous part of the motorway during the next cycle can be measured. We define the pressure that a queue puts on the system as a function of the queue size and the fraction of leaving vehicles as

$$\pi_l = f(p, M). \tag{7.31}$$

The function f can be chosen depending on the goals of the controller, we will work with choice $f(p, M) = (1 - p_l)(M_l - M_{l+1})$, which is based on the backpressure scheme. We assume that the number of vehicles $N_l(t)$ moving to the adjoint segment from an off-ramp during the t-th cycle can be determined by setting the traffic light green for a given proportion of the cycle time, with a maximum of $N^{\max}$. A tighter bound will be given by the number of empty boxes, which we denote by $n_l(t)$.

This is not only affected by the number of vehicles that continue their journey on the motorway, which we denote by $n_l^{\text{join}}(t)$, but also by any change in the speed limits. Due to equation (7.30), if v_l is increased, x_l is increased linearly as well, however the flow of the boxes is fixed at $q^{\max}$, which can only be accomplished if their number is decreased accordingly. By the same logic a decreasing speed limit allows for an increase in the total number of boxes. Thus, as the number of boxes arriving from the previous segment is constant, increasing the speed limit can only be allowed if there are enough empty boxes to delete from the system. On the other hand, we can generate extra boxes, and thus allow more vehicles on the motorway, by decreasing the speed limit. With these considerations the policy can be given as follows.

- We keep track of the fraction of boxes left empty during the previous w cycles. The window size w is a fixed parameter of the model. This quantity is denoted by $\alpha_l(t)$ and calculated as

$$\alpha_l(t) = \sum_{t-w}^{t-1} \frac{n_l(s) - N_l(s)}{N^{\max}}. \tag{7.32}$$

- At time t we update the speed limit according to the following rule.
  - If $\alpha_l(t) < \alpha_1(v_l(t-1))$ and $M_l(t) > M_l^{\text{low}}(t)$, then we decrease the speed limit, $v_l(t) = v_l(t-1) - v_g$ and generate $n_l^{\text{diff}}(v_l(t), v_l(t-1))$ boxes.
  - If $\alpha_l(t) > \alpha_2(v_l(t-1))$ and $M_l(t) < M_l^{\text{up}}(t)$, then we increase the speed limit, $v_l(t) = v_l(t-1)+v_g$ and delete $n_l^{\text{diff}}(v_l(t), v_l(t-1))$ boxes.
  - Otherwise we let $v_l(t) = v_l(t-1)$, and $n_l^{\text{diff}}(t) = 0$.

- We determine $n_l(t)$ as

$$n_l(t) = q^{\max}T - n_l^{\text{join}}(t) + n_l^{\text{diff}}(t). \tag{7.33}$$

- We determine the fraction

$$\eta_l(t) = \min\Big\{ \exp\Big(\gamma\pi_l(t)\Big), 1\Big\}. \tag{7.34}$$

- We let
$$N_l(t) = \min\left\{\lfloor n\eta_l(t)\rfloor, M_l(t), N^{\max}\right\} \tag{7.35}$$
vehicles on the motorway from on-ramp l.

The parameters α_1 and α_2 are possibly speed-dependent marks set by the controller, which indicate, when the fraction of the boxes filled during previous cycles, suggests a different speed limit. Similarly $M_l^{\text{up}}(t)$ and $M_l^{\text{low}}(t)$ are upper and lower thresholds in the queue sizes respectively. If the queue is longer than the former, then increasing the speed limit and thus removing boxes, i.e opportunities to join the motorway from an already long queue can cause congestion on the adjoint secondary roads. On the other hand, generating extra boxes for a short queue by slowing down the vehicles on the motorway is unnecessary and might lead to longer travel times, especially for vehicles joining from other, more congested on-ramps. Another parameter, which can be determined by the controller is γ, which reflects the level of fairness in the system. We discuss its role in detail below.

Initial results

We propose an extensive simulation study in order to validate our model and evaluate the performance of the proposed control scheme under different parameter settings. As this study is currently ongoing, we refer the reader to future works on the subject, here we only present initial results to facilitate the discussion on fairness in the system. Our simulations only cover a limited set of cases, therefore clear conclusions cannot be drawn from the results.

The initial analysis was focused on the role of the parameter γ. As it determines how the traffic load is distributed among the on-ramp queues, it allows the controller to set the level of fairness in the system. The extreme case of $\gamma = 0$ represents a greedy policy, where all vehicles are let on the motorway up to its capacity, regardless of how this affects upstream queues. Figure 7.5 presents the simulation results when $\gamma = 0.005$, thus the system is essentially following the greedy policy. While we see the lowest queue sizes in this case, the flow highly varies over different segments and is clearly suboptimal. Figure 7.6 shows the more realistic $\gamma = 0.5$ setting. In this case the queue sizes are a little higher, although

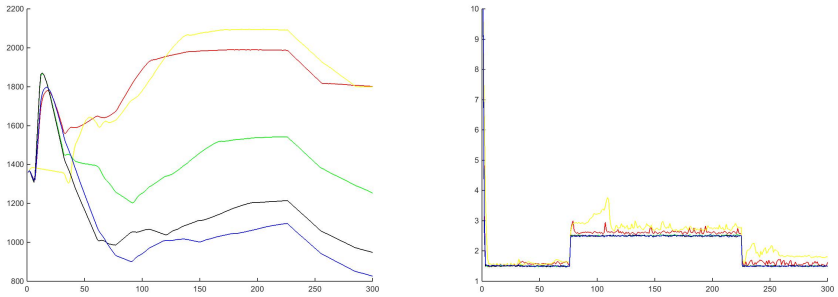

Figure 7.5. Flows (left) and on-ramp queue sizes (right) for different segments on a motorway with $\gamma = 0.005$.

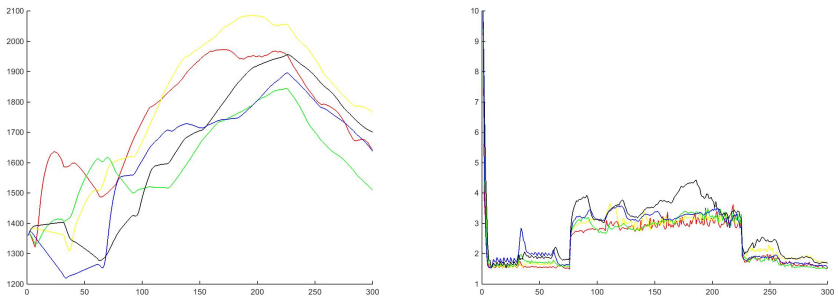

Figure 7.6. Flows (left) and on-ramp queue sizes (right) for different segments on a motorway with $\gamma = 0.5$.

still managable. The flow on the other hand is similar over the segments, remaining close to its maximum. The other extreme case happens if $\gamma \to \infty$, when only queues with more vehicles than their upstream neighbours receive service. We emulated this, by setting $\gamma = 50000$. The results ae shown in Figure 7.7. In this case the flow was clearly suboptimal, while some of the queues became congested over time, making this setting the worst performance-wise.

As our simulations show, changing one of the parameters can greatly alter the performance of the policy. This can be explained easily by the highly complex nature of the system. In order to understand the importance of each parameter and analyse the performance of the system

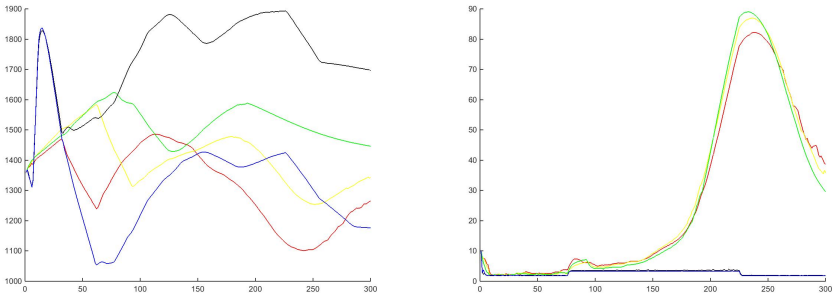

Figure 7.7. Flows (left) and on-ramp queue sizes (right) for different segments on a motorway with $\gamma = 5000$.

an extensive simulation study is to be carried out.

7.5 Conclusion

In this section we conducted an investigation of two control policies in the context of joint control of motorways and the surrounding roads; namely, a proportional fair and the backpressure based control policy. Both schemes are distributed which is a desirable property for implementation. We showed how the ensuing optimisations can be separated leading to simpler problems, that can then be solved by a gradient descent algorithm or a linear program.

To compare the performance of the schemes, we conducted numerical experiments on simple network settings. Our experiments highlighted trade-offs that may exist between a proportional fair controller and backpressure. For example, backpressure keeps the motorway less congested when operating under high loads, whereas the proportional fair policy leads to lower total vehicle numbers resulting in lower sojourn times.

Since we focussed on simple scenarios, there are many avenues left for further investigation. One of our ongoing research projects focuses on incorporating the complex traffic behaviour described by the fundamental diagram. Namely, the box method presents a combined control scheme for ramp metering and dynamic speed limits, that aims at maximising the traffic flow of the motorway. In the policy we have included ideas

from the backpressure scheme to distribute the load among the on-ramp queues. As the system shows a great degree of complexity an extensive simulation study is in plan for the future.

Further investigations may also consider taking account of local fluctuations in the density and the flow, and not just account for averages, since they can also be the cause of collapsing traffic [63]. It could also be useful to better understand the impact of the correlations between the traffic signals at the end of the on- and off-ramps, or relax the assumption that junctions at the ramps belonging to different segments are far enough to be independent.

There is clearly a need for a better theoretical understanding, and comprehensive performance evaluation, of the underlying models. The use of distributed schemes, such as the ones described here, provide an opportunity to analyase ways of controlling complex road networks, which may eventually lead to more efficient traffic control methods in the future.

Summary

Stochastic models for road traffic control

In this dissertation we made use of the theories of *stochastic processes* and *operations research* to develop models and methods to be applied for the analysis and control of road traffic networks. Three subjects were considered: individual routing, urban traffic light networks and motorways. Each chapter in the thesis covers a specific topic relating to one of these subjects. Chapter 2 discusses individual routing. For urban traffic light networks a general modeling framework is introduced in Chapter 3, which serves as a basis for the control methods described in Chapters 4 and 5. The last part of the dissertation is about inter-urban traffic focusing on motorways. Chapter 6 introduces a model that describes the evolution of traffic flow on a motorway, while Chapter 7 discusses the connection between urban road networks and motorways, presenting different control strategies for their traffic flows. A detailed summary of the results is given below.

Individual routes

Chapter 2 introduces a mathematical model for a route-guidance system that aims at minimising the average sojourn time of its users by giving advice based on information collected from them. Our model is a queueing system consisting of two queues and a controller that assigns newly arriving vehicles to the queues. Furthermore, it is assumed that only part of the vehicles (type X) use the route-guidance application, thus only they can be observed and controlled, whereas the other vehicles (type Y) cannot be observed explicitly and randomly choose a queue themselves. Arrivals occur according to a Poisson process, where a fixed fraction α (the penetration level) of the vehicles belongs to type X. At all times the number of X-type vehicles in each queue and their individual positions are known, however no other information is available to the controller about the traffic. Therefore, the controller has to perform its routing task based on partial information about the system state.

Our research was driven by a number of practically relevant ques-

tions: 1) what penetration level α is needed, 2) which policy should be implemented at the router, and 3) what is the added value of having more system information? To gain answers we have conducted extensive simulations and an analytical study focusing on the system under large loads. We have found that under such conditions the performance (in terms of the average sojourn time) of a simple policy that relies on little system information is close to the *weighted join-the-shortest-queue* policy, which is *optimal* in a fully controllable and observable system. The simple policy estimates queue lengths with the position of the last X vehicle in each queue. This result was obtained by analysing deterministic fluid models that approximate the stochastic evolution: in the *fluid limit* the processes corresponding to the different policies were shown to converge to the same limiting process. Our results were supported by simulation, which also revealed that for heavily loaded systems a low penetration level suffices to achieve comparable performance as in a fully controlled setting.

Urban road traffic

The second major subject of the dissertation is the optimal control of traffic light networks, which is covered in Chapters 3-5. These are modeled as queueing systems too, the main features of these models are described in Chapter 3, which we summarise next. Vehicles are assumed to be observable at all times, thus the lengths of the queues in front of the traffic lights are known. The vehicles leave the queues once the light turns green, which happens in a cyclic manner. Their route choices are made unbeknownst to the controller, whose task is to allocate green times and thus service to the queues for the different directions at the traffic lights. The allocation, which has to obey some restrictions for example for safety reasons, is made cycle-to-cycle, assigning fractions of the total cycle length to the phases enacted by the traffic light. There are two decentralised control policies introduced in this dissertation: the backpressure scheme in Chapter 4 and the proportional fair scheme in Chapter 5. Each gives the allocation as a solution of an optimisation problem.

The focus of the research was on the performance of these policies. In both cases we have shown that the methods ensure maximum stability,

which means that they provide sufficient throughput for the largest possible set of arrival rates (or demand). These proofs were made by relying on a *Lyapunov* function argument for the backpressure policy and by an analysis of the fluid limit of the queueing processes for the proportional fair policy. We have also discussed further possibilities, which might alter the performance of these methods. For example for the backpressure policy we have included the possibility of measurement errors, and for the proportional fair policy we allowed for changing cycle lengths and determined an optimal scaling for them. We have also conducted simulations to compare the performance of these algorithms to each other and to other well-known schemes available in the road traffic literature. The results were promising as the two policies performed comparably to each other (the backpressure was slightly better, but the proportional fair scheme has the advantage of an easier implementation), while outperforming other algorithms, for example the P0 policy.

Interurban road traffic

The last part of the dissertation discusses the traffic flow of roads connecting urban areas. Chapter 6 introduces a model, which describes the evolution of the traffic flow on a two lane motorway on a car-to-car basis. Here, vehicles are modeled as interacting particles which move in the same direction. Their movement is given by a stochastic process, in which their interactions are captured by the transition rates. Chapter 7 discusses the joint control of motorways and the connecting urban areas. This is exercised through ramps, which are the road segments connecting the different types of roads in the network. The ensuing model considers the motorway traffic as a continuous flow and the vehicles on the surrounding roads as parts of a queueing system. The joint control schemes for this network are built upon the proportional fair and the backpressure algorithms introduced earlier. The latter applies a greedy policy for the motorway, however in Section 7.4 an alternative method is introduced as well, which applies the main idea of the backpressure policy to the motorway part as well.

Our research took different directions in the different chapters. In Chapter 6 we have investigated natural features of the model and the ensuing conditions on the transition rates. Furthermore we have determ-

ined the *hydrodynamic* relations in the system and compared them via statistical methods to a set of data measured at a British motorway. On the other hand, in Chapter 7 we focused on the performance of the control policies listed above. We highlighted their advantages and disadvantages both by analysis and simulation. One of the main differences could be observed in terms of *fairness*, which is highly discussed topic in the literature.

The research presented in this dissertation highlighted many ways to effectively control road traffic networks using mathematical techniques and – particularly in the last two chapters – offers many possibilities for further study of road networks through mathematical models.

Samenvatting

Stochastische modellen voor wegverkeer regulering

In dit proefschrift hebben we gebruik gemaakt van *stochastische processen* en *mathematische besliskunde* om modellen en methodes te ontwikkelen voor de beschrijving en regulering van wegverkeersnetwerken. We hebben drie centrale thema's behandeld: individuele routering van voertuigen, verkeersstromen in stedelijke netwerken met verkeerslichten en verkeersstromen op snelwegen.

In Hoofdstuk 2 behandelen we individuele routering, een onderwerp dat in de literatuur veel aandacht krijgt door de toenemende mogelijkheden om bij de routering informatie te gebruiken over de locatie van voertuigen. Voor de modellering van stedelijke netwerken met verkeerslichten beschrijven we een algemeen raamwerk in Hoofdstuk 3, dat de basis is voor de reguleringsmethodieken die in Hoofdstukken 4 en 5 beschreven worden. Het laatste deel van het proefschrift gaat over interstedelijk verkeer op snelwegen. Hoofdstuk 6 introduceert een model dat de dynamiek van snelwegverkeer beschrijft. Tenslotte behandelt Hoofdstuk 7 regelstrategieën voor de verbinding van stedelijke netwerken en snelwegen. We beschrijven deze thema's nu in iets meer detail.

Routeren van voertuigen

Hoofdstuk 2 beschrijft een "navigatie systeem" gericht op het minimaliseren van de gemiddelde reistijd van zijn gebruikers, gebruik makend van locatie informatie die van de voertuigen zelf wordt ontvangen. Het wiskundige model is een wachtrijsysteem met twee wachtrijen en een routeerder die aankomende voertuigen naar een van de twee wachtrijen stuurt. We nemen aan dat slechts een deel van de voertuigen (we duiden deze voertuigen met type X aan) de apparatuur hebben om enerzijds locatie informatie te verschaffen en anderzijds gebruik te kunnen maken van de routeringsadviezen. De overige voertuigen (type Y) kunnen niet individueel worden waargenomen en kiezen zelf onafhankelijk een van de twee wachtrijen (volgens vaste kansen). Aankomsten van nieuwe voertuigen geschieden volgens een Poisson proces en daarvan behoort een

vaste fractie α (de penetratiegraad) tot type X. De routeerder weet alleen van de type X voertuigen hoeveel er in elk van de twee wachtrijen staan en welke posities ze bezetten. De type Y voertuigen zijn niet zichtbaar en dus moet de routeerder werken op basis van onvolledige informatie.

Ons onderzoek werd gedreven door een aantal praktisch relevante vragen: 1) hoe hoog moet de penetratiegraad α zijn opdat de routeerder in staat is om de reistijden (van type X voertuigen) effectief te verbeteren, 2) wat is de optimale routeringsstrategie en 3) welke voordelen zijn er nog te behalen met aanvullende informatie over het verkeer. Om deze vragen te beantwoorden, hebben we uitgebreide simulaties uitgevoerd en een analytische studie onder de aanname dat het systeem zwaar belast is. Onze bevindingen waren dat als de drukte in het system hoog is, dan behaalt een eenvoudige strategie met beperkte informatie vrijwel even goede resultaten (in termen van lage reistijden) als de zogenoemde *weighted join-the-shortest-queue* strategie die *optimaal* is wanneer de routeerder ook over volledige informatie beschikt over de type Y voertuigen. De eenvoudige strategie schat de totale wachtrijlengtes op basis van de de positie van het laatste type X voertuig in elk van de rijen. Dit optimaliteitsresultaat voor de eenvoudige strategie werd behaald met behulp van een deterministisch vloeistof model waarvan we hebben aangetoond dat het de stochastische modellen van beide strategieën goed benadert onder een zogenaamde *fluid* limiet. De fluid schaling is een veelgebruikte techniek om complexe stochastische processen te bestuderen. In onze analyse laten we zien dat de fluid limieten van beide strategieën - de eenvoudige op basis van partiële informatie en de optimale die gebruik maakt van volledige informatie - dezelfde zijn. Deze resultaten worden ook ondersteund door onze simulatie-experimenten waaruit blijkt dat in heel drukke systemen, een lage penetratiegraad voldoende is om goede resultaten te boeken.

Stedelijk wegverkeer

Het tweede thema in het proefschrift is de optimale regulering van verkeerslichten in stedelijk verkeer en wordt behandeld in Hoofdstukken 3-5. We gebruiken weer wachtrijsystemen om deze netwerken te modelleren. In Hoofdstuk 3 worden de belangrijkste kenmerken beschreven. We nemen aan dat alle voertuigen waarneembaar zijn, dus we kennen de

wachtrijlengtes bij alle verkeerslichten. Voertuigen verlaten de wachtrijen wanneer ze groen licht krijgen; dit gebeurt op een cyclische manier. Hun route keuzes worden zelfstandig gemaakt en zijn niet van tevoren bekend bij de beheerder van de verkeerslichten. Deze bepaalt wanneer welk verkeerslicht groen wordt en dus welke verkeersstromen kunnen rijden. De toewijzing van groentijden moet voldoen aan verschillende voorwaarden, die onder andere ingegeven worden door veiligheidsoverwegingen. Per cyclus wordt bepaalt welk deel elk verkeerslicht toegewezen krijgt als groentijd. Hiertoe wordt de cyclus opgedeeld in fases, zodanig dat conflicterende stromen niet binnen eenzelfde fase groentijd krijgen. In dit proefschrift bestuderen we twee decentrale reguleringsmechanismen: het "backpressure" mechanisme in Hoofdstuk 4 en de "proportional fair" regel in Hoofdstuk 5. In beide wordt de toewijzing van groentijden gevonden met behulp van de oplossing van een wiskundig optimaliseringsprobleem.

In ons onderzoek hebben we ons gericht op de effectiviteit van deze strategieën in termen van wachtrijlengtes en reistijden. Voor beide strategieën hebben we bewezen dat ze maximale stabiliteit garanderen, dat wil zeggen dat als er voor een bepaalde belasting van het netwerk een strategie bestaat die het netwerk stabiel houdt, dan doen de backpressure en proportional fair strategieën dat ook. In onze wiskundige bewijzen hebben we twee verschillende technieken gebruikt. Voor het backpressure model hebben we gebruik gemaakt van een argument met behulp van een geschikt gekozen *Lyapunov* functie en voor het proportional fair model gebruikten we (de reeds eerder genoemde) fluid limiet schaling. We hebben ook uitbreidingen beschouwd op onze basismodellen om de robuustheid van onze resultaten te onderzoeken. In het backpressure model hebben we de resultaten uitgebreid voor het geval waarin er meetfouten kunnen worden gemaakt bij het bepalen van de wachtrijlengten. In het proportional fair model hebben we onderzocht wat het effect zou zijn van een dynamisch aan te passen cycluslengte en hebben we een optimale schaling van de cyclus ten opzichte van de (gemiddelde) rijlengte bepaald. Tenslotte hebben we voor deze algoritmen met behulp van simulaties laten zien dat hun effectiviteit vergelijkbaar, of beter is dan dat van alternatieve bestaande algoritmen (bijvoorbeeld de P0 strategie). In het algemeen doet backpressure het iets beter dan proportional fair, maar deze laatste heeft het voordeel dat deze gedistribueerd kan worden geïmplementeerd (dus zonder dat een globaal overzicht van alle wachtrijen in

het netwerk nodig is om de groentijden van een knooppunt te bepalen).

Verkeersstromen op snelwegen

In het laatste deel van het proefschrift bestuderen we de verbinding van stedelijke gebieden. Hoofdstuk 6 presenteert een model dat de dynamiek beschrijft van verkeersstromen op een tweebaanssnelweg . Voertuigen worden hierbij gerepresenteerd als interacterende deeltjes die in dezelfde richting bewegen. Hun verplaatsing wordt gereguleerd door een stochastisch proces, waarvan de transitie intensiteiten zo gekozen zijn om de interacties goed te beschrijven. Hoofdstuk 7 gaat in op het gelijktijdig reguleren van de snelweg en de stedelijke gebieden die worden verbonden. Deze regulering vindt plaats bij de op- en afritten die de verbinding vormen tussen de verschillende typen wegen in het netwerk. Het resulterende model beschrijft het snelwegverkeer als een continue stroom en de voertuigen in de omliggende gebieden als onderdeel van een wachtrijsysteem. De gezamenlijke regulering van dit netwerk maakt gebruik van de eerder geïntroduceerde proportional fair en backpressure algoritmes. In Sectie 7.4 wordt ook een backpressure strategie gebruikt voor het snelweggedeelte (in het basismodel gebruiken we daarvoor een naieve "greedy" strategie).

We hebben verschillende onderzoeksrichtingen beschouwd in de diverse hoofdstukken. In Hoofdstuk 6 hebben we natuurlijke karakteristieken van het model onderzocht en de bijbehorende voorwaarden op de transitie intensiteiten. Verder hebben we de *hydrodynamische* vergelijkingen bepaald om het systeem te beschrijven en deze met statistische methoden vergeleken met een data set van metingen op een Britse snelweg. In Hoofdstuk 7 richtten we ons juist op de effectiviteit (performance) van de voornoemde reguleringsstrategieën. We benoemen de voor- en nadelen, zowel via wiskundige analyse als simulatie. Een van de belangrijkste verschillen zien we in de *eerlijkheid* (fairness) van het systeem, wat momenteel veel aandacht krijgt in de literatuur.

Het onderzoek dat in dit proefschrift wordt beschreven, heeft ons veel geleerd over de mogelijkheden om verkeersnetwerken effectief te reguleren door middel van wiskundige technieken en - in het bijzonder in de laatste twee hoofdstukken - biedt legio mogelijkheden voor verder onderzoek van wegverkeernetwerken door middel van mathematische modellen.

Publications of the author

Publications related to this dissertation

[1] Wendy Ellens, Péter Kovács, Rudesindo Núñez-Queija and Hans van den Berg. 'Routing policies for a partially observable two-server queueing system'. In: *Proceedings of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering). 2016, pp. 111–118.

[2] Péter Kovács, Gaurav Raina and Neil Walton. 'Ramp and signal control: Where motorways and urban roads meet'. In: *2015 7th International Conference on Communication Systems and Networks (COMSNETS)*. IEEE. 2015, pp. 1–6.

[3] Péter Kovács, Tung Le, Rudesindo Núñez-Queija, Hai Vu and Neil Walton. 'Proportional green time scheduling for traffic lights'. Submitted for publication. 2016.

[4] Tung Le, Péter Kovács, Neil Walton, Hai Vu, Lachlan Andrew and Serge Hoogendoorn. 'Decentralized signal control for urban road networks'. In: *Transportation Research Part C: Emerging Technologies* 58 (2015), pp. 431–450.

References

[5] K. Aboudolas, M. Papageorgiou and E. Kosmatopoulos. 'Store-and-forward based methods for the signal control problem in large-scale congested urban road networks'. In: *Transportation Research Part C: Emerging Technologies* 17.2 (2009), pp. 163–174.

[6] K. Aboudolas, M. Papageorgiou, A. Kouvelas and E. Kosmatopoulos. 'A rolling-horizon quadratic-programming approach to the signal control problem in large-scale congested urban road networks'. In: *Transportation Research Part C: Emerging Technologies* 18.5 (2010), pp. 680–694.

[7] J. Adler, G. Satapathy, V. Manikonda, B. Bowles and V. Blue. 'A multi-agent approach to cooperative traffic management and route guidance'. In: *Transportation Research Part B: Methodological* 39.4 (2005), pp. 297–318.

[8] K. Ahmed, M. Ben-Akiva, H. Koutsopoulos and R. Mishalani. 'Models of freeway lane changing and gap acceptance behavior'. In: *Transportation and traffic theory* 13 (1996), pp. 501–515.

[9] *Apple Maps*. `www.apple.com/ios/maps`.

[10] C. Bahadoran, H. Guiol, K. Ravishankar and E. Saada. 'A constructive approach to Euler hydrodynamics for attractive processes. Application to k-step exclusion'. In: *Stochastic processes and their applications* 99.1 (2002), pp. 1–30.

[11] S. Baruah, N. Cohen, G. Plaxton and D. Varvel. 'Proportionate progress: A notion of fairness in resource allocation'. In: *Algorithmica* 15.6 (1996), pp. 600–625.

[12] M. Beckmann, C. McGuire and C. Winsten. *Studies in the Economics of Transportation*. Tech. rep. 1956.

[13] S. Bell, R. Williams et al. 'Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy'. In: *The Annals of Applied Probability* 11.3 (2001), pp. 608–649.

[14] D. Bertsekas, R. Gallager and P. Humblet. *Data networks*. Vol. 2. Prentice-Hall International New Jersey, 1992.

[15] D. Bertsimas, V. Farias and N. Trichakis. 'The price of fairness'. In: *Operations research* 59.1 (2011), pp. 17–31.

[16] S. Bhulai, G. Hoekstra, J. Bosman and R. van der Mei. 'Dynamic traffic splitting to parallel wireless networks with partial information: A Bayesian approach'. In: *Performance Evaluation* 69 (2012), pp. 41–52.

[17] P. Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.

[18] P. Bonsall. 'The influence of route guidance advice on route choice in urban networks'. In: *Transportation* 19.1 (1992), pp. 1–23.

[19] M. Boon. *Polling models: from theory to traffic intersections*. Eindhoven University of Technology, 2011.

[20] J. Bottom. *Consistent anticipatory route guidance*. Massachusetts Institute of Technology, 2000.

[21] M. Brackstone and M. McDonald. 'Car-following: a historical review'. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 2.4 (1999), pp. 181–196.

[22] M. Bramson. *Stability of queueing networks*. Springer, 2008.

[23] D. Buckley. 'A semi-poisson model of traffic flow'. In: *Transportation Science* 2.2 (1968), pp. 107–133.

[24] C. Cai, C. Wong and B. Heydecker. 'Adaptive traffic signal control using approximate dynamic programming'. In: *Transportation Research Part C: Emerging Technologies* 17.5 (2009), pp. 456–474.

[25] M. Cetin, G. List and Y. Zhou. 'Factors affecting minimum number of probes required for reliable estimation of travel time'. In: *Transportation Research Record: Journal of the Transportation Research Board* 1917 (2005), pp. 37–44.

[26] R. Chandler, R. Herman and E. Montroll. 'Traffic dynamics: studies in car following'. In: *Operations research* 6.2 (1958), pp. 165–184.

[27] H. Chen and D. Yao. *Fundamentals of queueing networks: Performance, asymptotics, and optimization*. Vol. 46. Springer Science & Business Media, 2013.

[28] M. Chen and S. Chien. 'Dynamic freeway travel-time prediction with probe vehicle data: Link-based versus path-based'. In: *Transportation Research Record: Journal of the Transportation Research Board* 1768 (2001), pp. 157–161.

[29] D. Chowdhury, L. Santen and A. Schadschneider. 'Statistical physics of vehicular traffic and some related systems'. In: *Physics Reports* 329.4 (2000), pp. 199–329.

[30] L. Chu, H. Liu, W. Recker and M. Zhang. 'Performance evaluation of adaptive ramp-metering algorithms using microscopic traffic simulation

model'. In: *Journal of Transportation Engineering* 130.3 (2004), pp. 330–338.

[31] R. Clegg, A. Clune and M. Smith. 'Traffic signal settings for diverse policy goals'. In: *Proc. of PTRC* P445 (2000), pp. 93–104.

[32] G. Comert. 'Simple analytical models for estimating the queue lengths from probe vehicles at traffic signals'. In: *Transportation Research Part B: Methodological* 55 (2013), pp. 59–74.

[33] C. Daganzo and N. Geroliminis. 'An analytical approximation for the macroscopic fundamental diagram of urban traffic'. In: *Transportation Research Part B* 42 (2008), pp. 771–781.

[34] J. Dai. 'On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models'. In: *The Annals of Applied Probability* (1995), pp. 49–77.

[35] B. De Schutter and B. De Moor. 'Optimal traffic light control for a single intersection'. In: *European Journal of Control* 4.3 (1998), pp. 260–276.

[36] J. Del Castillo and F. Benitez. 'On the functional form of the speed-density relationship I: general theory'. In: *Transportation Research Part B: Methodological* 29.5 (1995), pp. 373–389.

[37] J. Del Castillo and F. Benitez. 'On the functional form of the speed-density relationship II: empirical investigation'. In: *Transportation Research Part B: Methodological* 29.5 (1995), pp. 391–406.

[38] J. Drake, J. Schofer and A. May. 'A statistical analysis of speed-density hypotheses. in vehicular traffic science'. In: *Highway Research Record* 154 (1967).

[39] D. Drew. *Traffic flow theory and control*. Tech. rep. 1968.

[40] Y. Dujardin, F. Boillot, D. Vanderpooten and P. Vinant. 'Multiobjective and multimodal adaptive traffic light control on single junctions'. In: *IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE. 2011, pp. 1361–1368.

[41] L. Edie. 'Car-following and steady-state theory for noncongested traffic'. In: *Operations Research* 9.1 (1961), pp. 66–76.

[42] L. Edie. 'Traffic delays at toll booths'. In: *Journal of the operations research society of America* 2.2 (1954), pp. 107–138.

[43] W. Ellens. *Stochastic Methods for Measurement-based Network Control*. University of Amsterdam, 2015.

[44] B. Feng, J. Hourdos and P. Michalopoulos. 'Improving Minnesota's stratified ramp control strategy'. In: *Transportation Research Record: Journal of the Transportation Research Board* 1959 (2006), pp. 77–83.

[45] E. Friedman and S. Henderson. 'Fairness and efficiency in web server protocols'. In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 31. 1. ACM. 2003, pp. 229–237.

[46] N. Gartner. 'OPAC: A demand-responsive strategy for traffic signal control'. In: *Transportation Research Record* 906 (1983), pp. 75–81.

[47] N. Gartner, J. Little and H. Gabbay. 'Optimization of Traffic Signal Settings by Mixed-Integer Linear Programming: Part I: The Network Coordination Problem'. In: *Transportation Science* 9.4 (1975), pp. 321–343.

[48] N. Gartner, J. Little and H. Gabbay. 'Optimization of Traffic Signal Settings by Mixed-Integer Linear Programming Part II: The Network Synchronization Problem'. In: *Transportation Science* 9.4 (1975), pp. 343–363.

[49] N. Gartner, F. Pooran and C. Andrews. 'Implementation of the OPAC adaptive control strategy in a traffic signal network'. In: *IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE. 2001, pp. 195–200.

[50] R. Gibbens and F. Kelly. 'An investigation of proportionally fair ramp metering'. In: *Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on*. IEEE. 2011, pp. 490–495.

[51] *Google Maps*. `www.google.com/maps`.

[52] R. Graham, E. Lawler, J. Lenstra and R. Kan. 'Optimization and approximation in deterministic sequencing and scheduling: a survey'. In: *Annals of discrete mathematics* 5 (1979), pp. 287–326.

[53] H. Greenberg. 'An analysis of traffic flow'. In: *Operations research* 7.1 (1959), pp. 79–85.

[54] B. Greenshields, W. Channing, H. Miller et al. 'A study of traffic capacity'. In: *Highway research board proceedings*. Vol. 1935. National Research Council (USA), Highway Research Board. 1935.

[55] P. Guo, W. Sun and Y. Wang. 'Equilibrium and optimal strategies to join a queue with partial information on service times'. In: *European Journal of Operational Research* 214 (2011), pp. 284–297.

[56] F. Haight. *Mathematical theories of traffic flow*. Tech. rep. 1963.

[57] R. Hall and M. Lieberman. *Microeconomics: Principles and applications*. Cengage Learning, 2012.

[58] A. Hamilton, B. Waterson, T. Cherrett, A. Robinson and I. Snell. 'The evolution of urban traffic control: changing policy and technology'. In: *Transportation Planning and Technology* 36.1 (2013), pp. 24–43.

[59] M. Harrison. 'Heavy traffic analysis of a system with parallel servers: asymptotic optimality of discrete-review policies'. In: *Annals of applied probability* (1998), pp. 822–848.

[60] J. Henry, J. Farges and J. Tuffal. 'The PRODYN real time traffic algorithm'. In: *IFAC/IFIP/IFORS Conference on Control in Transportation Systems*. Sept. 1983, pp. 307–312.

[61] J. Herrera, D. Work, R. Herring, X. Ban, Q. Jacobson and A. Bayen. 'Evaluation of traffic data obtained via GPS-enabled mobile phones: The mobile century field experiment'. In: *Transportation Research Part C: Emerging Technologies* 18 (2010), pp. 568–583.

[62] *Highways England*. `http://data.gov.uk/`.

[63] S. Hoogendoorn, V. Knoop, H. van Lint and H. Vu. 'Applications of the generalized macroscopic fundamental diagram'. In: *Traffic and Granular Flow'13*. Springer, 2015, pp. 577–583.

[64] A. Hordijk and G. Koole. 'On the optimality of the generalized shortest queue policy'. In: *Probability in the Engineering and Informational Sciences* 4 (1990), pp. 477–487.

[65] J. Hourdos, M. Janson, D. Levinson and G. Parikh. *MnPASS Modeling and Pricing Algorithm Enhancement*. Center for Transportation Studies, University of Minnesota, 2015.

[66] P. Hunt, D. Robertson, R. Bretherton and R. Winton. *SCOOT-a traffic responsive method of coordinating signals*. Tech. rep. 1981.

[67] T. Ikeda, M. Yoshii, Y. Doi and K. Mitoh. *Adaptive in-vehicle route guidance system*. US Patent 5,031,104. 1991.

[68] *INFORMS*. `www.informs.org`.

[69] D. Kaufman, R. Smith and K. Wunderlich. 'An iterative routing/assignment method for anticipatory real-time route guidance'. In: *Vehicle Navigation and Information Systems Conference, 1991*. Vol. 2. IEEE. 1991, pp. 693–700.

[70] F. Kelly. 'Charging and rate control for elastic traffic'. In: *European transactions on Telecommunications* 8.1 (1997), pp. 33–37.

[71] F. Kelly and R. Williams. 'Fluid model for a network operating under a fair bandwidth-sharing policy'. In: *Annals of Applied Probability* (2004), pp. 1055–1083.

[72] F. Kelly and R. Williams. 'Heavy traffic on a controlled motorway'. In: *Probability and Mathematical Genetics: Papers in Honour of Sir John Kingman* (2010).

[73] F. Kelly and E. Yudovina. *Stochastic networks*. Vol. 2. Cambridge University Press, 2014.

[74] J. Kingman. 'On queues in heavy traffic'. In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1962), pp. 383–392.

[75] W. Kraft. *Traffic Engineering Handbook*. Institute of Transportation Engineers, 6th Edition, 2009.

[76] R. Kühne and M. Rödiger. 'Macroscopic simulation model for freeway traffic with jams and stop-start waves'. In: *Proceedings of the 23rd conference on Winter simulation*. IEEE Computer Society. 1991, pp. 762–770.

[77] J. Kuri and A. Kumar. 'Optimal control of arrivals to queues with delayed queue length information'. In: *IEEE Transactions on Automatic Control* 40 (1995), pp. 1444–1450.

[78] T. Kurtz. 'Limit theorems for sequences of jump Markov processes approximating ordinary differential processes'. In: *Journal of Applied Probability* 8.2 (1971), pp. 344–356.

[79] H. Kushner. *Heavy traffic analysis of controlled queueing and communication networks*. Vol. 47. Springer Science & Business Media, 2013.

[80] S. Lämmer and D. Helbing. 'Self-control of traffic lights and vehicle flows in urban road networks'. In: *Journal of Statistical Mechanics: Theory and Experiment* 2008.04 (2008), P04019.

[81] S. Lämmer and D. Helbing. 'Self-stabilizing decentralized signal control of realistic, saturated network traffic'. In: Citeseer. 2010.

[82] M. Lay and J. Vance. *Ways of the World: A History of the World's Roads and of the Vehicles that Used Them*. Rutgers University Press, 1992.

[83] T. Le, H. Vu, Y. Nazarathy, Q. Vo and S. Hoogendoorn. 'Linear-quadratic model predictive control for urban traffic networks'. In: *Transportation Research Part C: Emerging Technologies* 36 (2013), pp. 498–512.

[84] T. Liggett. *Interacting particle systems*. Vol. 276. Springer Science & Business Media, 2012.

[85] M. Lighthill and G. Whitham. 'On kinematic waves. I. Flood movement in long rivers'. In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. Vol. 229. 1178. The Royal Society. 1955, pp. 281–316.

[86] M. Lighthill and G. Whitham. 'On kinematic waves. II. A theory of traffic flow on long crowded roads'. In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. Vol. 229. 1178. The Royal Society. 1955, pp. 317–345.

[87] S. Lin, B. De Schutter, Y. Xi and H. Hellendoorn. 'Fast model predictive control for urban road networks via MILP'. in: *IEEE Transactions on Intelligent Transportation Systems* 12.3 (2011), pp. 846–856.

[88] J. Little. 'A proof for the queuing formula: L= λ W'. in: *Operations research* 9.3 (1961), pp. 383–387.

[89] R. Long Cheu, C. Xie and D. Lee. 'Probe vehicle population and sample size for arterial speed estimation'. In: *Computer-Aided Civil and Infrastructure Engineering* 17 (2002), pp. 53–60.

[90] P. Lowrie. 'The Sydney coordinated adaptive traffic system-principles, methodology, algorithms'. In: *International Conference on Road Traffic Signalling* 207 (1982).

[91] M. MacNicholas. 'A simple and pragmatic representation of traffic flow'. In: *Symposium on The Fundamental Diagram*. Vol. 75. 2011, pp. 161–177.

[92] H. Mahmassani, J. Dong, J. Kim, R. Chen and B. Park. 'Incorporating weather impacts in traffic estimation and prediction systems'. In: *US Department of Transport, Washington* (2009).

[93] L. Massoulié. 'Structural properties of proportional fairness: stability and insensitivity'. In: *The Annals of Applied Probability* (2007), pp. 809–839.

[94] *MATLAB*. `www.mathworks.com/products/matlab`.

[95] V. Mauro and C. Taranto. 'Utopia'. In: *Control, computers, communications in transportation* (1990), pp. 245–252.

[96] A. May. *Traffic flow fundamentals*. 1990.

[97] N. McKeown, A. Mekkittikul, V. Anantharam and J. Walrand. 'Achieving 100% throughput in an input-queued switch'. In: *IEEE Transactions on Communications* 47.8 (1999), pp. 1260–1267.

[98] S. Meyn and R. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, 2009.

[99] F. Middelham and H. Taale. 'Ramp metering in the Netherlands: an overview'. In: *Rijkswaterstaat, AVV Transport Research Centre* (2006).

[100] P. Mirchandani and L. Head. 'RHODES: A real-time traffic signal control system: architecture, algorithms, and analysis'. In: *Transportation Research Part C: Emerging Technologies* 9.6 (2001), pp. 415–432.

[101] M. Mitzenmacher. 'How useful is old information?' In: *IEEE Transactions on Parallel and Distributed Systems* 11 (2000), pp. 6–20.

[102] J. Mo and J. Walrand. 'Fair end-to-end window-based congestion control'. In: *IEEE/ACM Transactions on Networking (ToN)* 8.5 (2000), pp. 556–567.

[103] J. Moré. 'The Levenberg-Marquardt algorithm: implementation and theory'. In: *Numerical analysis*. Springer, 1978, pp. 105–116.

[104] K. Nagel. 'Particle hopping models and traffic flow theory'. In: *Physical review E* 53.5 (1996), p. 4655.

[105] C. Nanthawichit, T. Nakatsuji and H. Suzuki. 'Application of probe-vehicle data for real-time traffic-state estimation and short-term travel-time prediction on a freeway'. In: *Transportation Research Record: Journal of the Transportation Research Board* 1855 (2003), pp. 49–59.

[106] G. Newell. 'A simplified theory of kinematic waves in highway traffic, part I: General theory'. In: *Transportation Research Part B: Methodological* 27.4 (1993), pp. 281–287.

[107] G. Newell. 'A simplified theory of kinematic waves in highway traffic, Part II: Queueing at freeway bottlenecks'. In: *Transportation Research Part B: Methodological* 27.4 (1993), pp. 289–303.

[108] G. Newell. 'Mathematical models for freely-flowing highway traffic'. In: *Journal of the Operations Research Society of America* 3.2 (1955), pp. 176–186.

[109] G. Newell. 'Memoirs on highway traffic flow theory in the 1950s'. In: *Operations Research* 50.1 (2002), pp. 173–178.

[110] J. Norris. *Markov chains*. Vol. 2. Cambridge university press, 1998.

[111] M. Papageorgiou. 'Dynamic modeling, assignment, and route guidance in traffic networks'. In: *Transportation Research Part B: Methodological* 24.6 (1990), pp. 471–495.

[112] M. Papageorgiou, H. Hadj-Salem and J. Blosseville. 'ALINEA: A local feedback control law for on-ramp metering'. In: *Transportation Research Record* 1320 (1991).

[113] M. Papageorgiou, H. Hadj-Salem and F. Middelham. 'ALINEA local ramp metering: Summary of field results'. In: *Transportation Research Record: Journal of the Transportation Research Board* 1603.1 (1997), pp. 90–98.

[114] M. Papageorgiou, C. Diakaki, V. Dinopoulou, A. Kotsialos and Y. Wang. 'Review of road traffic control strategies'. In: *Proceedings of the IEEE* 91.12 (2003), pp. 2043–2067.

[115] Y. Pavlis and M. Papageorgiou. 'Simple decentralized feedback strategies for route guidance in traffic networks'. In: *Transportation science* 33.3 (1999), pp. 264–278.

[116] H. Payne. 'Models of freeway traffic and control.' In: *Mathematical models of public systems* (1971).

[117] L. Pipes. 'Car following models and the fundamental diagram of road traffic'. In: *Transportation Research* 1.1 (1967), pp. 21–29.

[118] I. Prigogine and F. Andrews. 'A Boltzmann-like approach for traffic flow'. In: *Operations Research* 8.6 (1960), pp. 789–797.

[119] I. Prigogine and R. Herman. *Kinetic theory of vehicular traffic*. Tech. rep. 1971.

[120] T. Rayburn. *Route planning system for mobile telecommunications*. US Patent 6,937,869. 2005.

[121] M. Reiman. 'Some diffusion approximations with state space collapse'. In: *Modelling and performance evaluation methodology*. Springer, 1984, pp. 207–240.

[122] M. Reiser. 'Mean-value analysis and convolution method for queue-dependent servers in closed queueing networks'. In: *Performance Evaluation* 1.1 (1981), pp. 7–18.

[123] M. Reiser and S. Lavenberg. 'Mean-value analysis of closed multichain queuing networks'. In: *Journal of the ACM (JACM)* 27.2 (1980), pp. 313–322.

[124] F. Rezakhanlou. 'Hydrodynamic limit for attractive particle systems on 417-1417-1417-1'. In: *Communications in mathematical physics* 140.3 (1991), pp. 417–448.

[125] P. Richards. 'Shock waves on the highway'. In: *Operations research* 4.1 (1956), pp. 42–51.

[126] P. Robert. *Stochastic networks and queues*. Springer, 2003.

[127] T. Roopa, A. Iyer and S. Rangaswamy. 'CroTIS: Crowdsourcing based traffic information system'. In: *IEEE International Congress on Big Data (BigData Congress)*. 2013, pp. 271–277.

[128] S. Ross et al. *Stochastic processes*. Vol. 2. John Wiley & Sons New York, 1996.

[129] K. Savla, E. Lovisari and G. Como. 'On maximally stabilizing traffic signal control with unknown turn ratios'. In: *IFAC Proceedings Volumes* 47.3 (2014), pp. 1849–1854.

[130] V. Sethi, N. Bhandari, F. Koppelman and J. Schofer. 'Arterial incident detection using fixed detector and probe vehicle data'. In: *Transportation Research Part C: Emerging Technologies* 3 (1995), pp. 99–112.

[131] L. Seymour, M. Barnea and A. Kirson. *Driver preference responsive vehicle route guidance system*. US Patent 6,212,470. 2001.

[132] R. Sipahi, S. Lämmer, D. Helbing and S. Niculescu. 'On stability problems of supply networks constrained with transport delay'. In:

Journal of Dynamic Systems, Measurement, and Control 131.2 (2009), p. 021005.

[133] M. Smith. 'A local traffic control policy which automatically maximises the overall travel capacity of an urban road network'. In: *Traffic Engineering & Control* 21.6 (1980), pp. 298–302.

[134] M. Smith. 'Dynamics of route choice and signal control in capacitated networks'. In: *Journal of Choice Modelling* 4.3 (2011), pp. 30–51.

[135] K. Srinivasan and P. Jovanis. 'Determination of number of probe vehicles required for reliable travel time measurement in urban network'. In: *Transportation Research Record: Journal of the Transportation Research Board* 1537 (1996), pp. 15–22.

[136] *SUMO*. `http://sumo-sim.org/`.

[137] H. Takagi. *Analysis of polling systems*. MIT press, 1986.

[138] L. Tassiulas and A. Ephremides. 'Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks'. In: *IEEE Transactions on Automatic Control* 37.12 (1992), pp. 1936–1948.

[139] T. Tettamanti and I. Varga. 'Distributed traffic control system based on model predictive control'. In: *Periodica Polytechnica ser. Civil. Eng.* 54.1 (2010), pp. 3–9.

[140] T. Tettamanti, I. Varga and T. Péni. 'MPC in urban traffic management'. In: *Model Predictive Control. SciYo* (2010).

[141] T. Tettamanti, I. Varga, B. Kulcsár and J. Bokor. 'Model predictive control in urban traffic network management'. In: *16th Mediterranean Conference on Control and Automation*. 2008, pp. 1538–1543.

[142] A. Thiagarajan, L. Ravindranath, K. LaCurts, S. Madden, H. Balakrishnan, S. Toledo and J. Eriksson. 'VTrack: Accurate, energy-aware road traffic delay estimation using mobile phones'. In: *ACM Conference on Embedded Networked Sensor Systems (SenSys)*. 2009, pp. 85–98.

[143] N. Thompson and S. Greene. 'Ramp Metering for the 21st Century: Minnesota's Experience'. In: *ITS America. Meeting (7th: 1997: Washington, DC.). Merging the transportation and communications revolutions: conference proceedings*. 1997.

[144] S. Turner and D. Holdener. 'Probe vehicle sample sizes for real-time information: the Houston experience'. In: *IEEE Vehicle Navigation and Information Systems Conference (VNIS)*. 1995, pp. 3–10.

[145] S. Turner et al. 'A join the shorter queue model in heavy traffic'. In: *Journal of Applied Probability* 37.1 (2000), pp. 212–223.

[146] M. Van Aerde. 'Single regime speed-flow-density relationship for congested and uncongested highways'. In: *74th Annual Meeting of the Transportation Research Board, Washington, DC*. vol. 6. 1995.

[147] P. Varaiya. 'The Max-Pressure Controller for Arbitrary Networks of Signalized Intersections'. In: *Advances in Dynamic Network Modeling in Complex Transportation Systems*. Springer, 2013, pp. 27–66.

[148] N. Walton. 'Concave switching in single and multihop networks'. In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 42. 1. ACM. 2014, pp. 139–151.

[149] H. Wang, J. Li, Q. Chen and D. Ni. 'Logistic modeling of the equilibrium speed–density relationship'. In: *Transportation research part A: policy and practice* 45.6 (2011), pp. 554–566.

[150] H. Wang, J. Li, Q. Chen and D. Ni. 'Speed-density relationship: From deterministic to stochastic'. In: *The 88th Transportation Research Board (TRB) Annual Meeting. Washington, DC*. 2009.

[151] J. Wardrop and J. Whitehead. 'Correspondence. Some theoretical aspects of road traffic research.' In: *Proceedings of the Institution of Civil Engineers* 1.5 (1952), pp. 767–768.

[152] *Waze*. `www.waze.com`.

[153] R. Weber. 'On the optimal assignment of customers to parallel servers'. In: *Journal of Applied Probability* 15 (1978), pp. 406–413.

[154] F. Webster. *Traffic signal settings*. Tech. rep. 1958.

[155] R. Wolff. 'Poisson arrivals see time averages'. In: *Operations Research* 30.2 (1982), pp. 223–231.

[156] T. Wongpiromsarn, T. Uthaicharoenpong, Y. Wang, E. Frazzoli and D. Wang. 'Distributed traffic signal control for maximum network throughput'. In: *IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE. 2012, pp. 588–595.

[157] H. Yoneyama, A. Yashiki and T. Kodama. *Route guidance system for provding a mobile station with optimum route data in response to a guidance request together with base station data indicative of an identification of a base station*. US Patent 5,187,810. 1993.

[158] R. Zhang, Z. Li, C. Feng and S. Jiang. 'Traffic routing guidance algorithm based on backpressure with a trade-off between user satisfaction and traffic load'. In: *Vehicular Technology Conference (VTC Fall), 2012 IEEE*. IEEE. 2012, pp. 1–5.