

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Theoretical Study of Backward Signal Coordination with Real-World Complications

Permalink

<https://escholarship.org/uc/item/8gt3t0kd>

ISBN

9798189270512

Author

Kong, Sili

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Theoretical Study of Backward Signal Coordination with Real-World Complications

by

Sili Kong

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Michael J. Cassidy, Chair

Carlos F. Daganzo

Steven N. Evans

Spring 2026

Theoretical Study of Backward Signal Coordination with Real-World Complications

Copyright 2026
by
Sili Kong

Abstract

Theoretical Study of Backward Signal Coordination with Real-World Complications

by

Sili Kong

Doctor of Philosophy in Civil and Environmental Engineering

University of California, Berkeley

Michael J. Cassidy, Chair

Backward signal coordination, or BSC for short, is a little-known but promising means of managing traffic in congested cities. When queues in one or more travel directions expand and spillover to street links upstream, there is little value in synchronizing traffic signals to what would be the forward motion of vehicles. Upon seeing green phases, drivers in these hyper congested conditions are constrained from moving forward. This is because links downstream are at least partially filled with standing queues. With BSC, in contrast, green times across neighboring intersections are synchronized to the backward-moving kinematic waves that propagate through queues. In this way, drivers receive green times only after downstream queues have begun moving forward. Far less green time is wasted as a result. As an added benefit, BSC reduces congestion's damaging footprint on a network by filling queued links to their brims, leaving fewer upstream links with standing queues.

The dissertation examines the effects of BSC in the presence of three complications. Each is commonly seen in real settings but has received little, if any, attention in the context of BSC. The three are as follows.

Complication 1: Differences in directional demands. The BSC strategy can be applied to a select travel direction along a hyper congested signalized street, meaning a street with spillover queues persisting across multiple consecutive links. Drivers traveling in the opposing direction, presumably the one with lower demand, must therefore confront traffic signals that are not synchronized to favor their (own) movements. It stands to reason that this sort of discriminating strategy can produce net benefits (e.g., reductions in networkwide travel times) when demand in the opposing, disfavored direction is small. But is the strategy still beneficial when demand in the disfavored direction is comparable to that in the favored (high-demand) direction?

Complication 2: Inhomogeneous network geometries. Although some prior studies of BSC considered arterials with nonuniform link lengths, it remains unclear how BSC performs

when a network includes short links with limited queue storage capacity.

Complication 3: Vehicle turning maneuvers. It is understood that protected phases that usher (left-) turning vehicles off synchronized streets can complicate the signal coordination problem. It is not known, however, to what extent protected left-turn phases and queue spillovers from left-turn pockets disrupt the BSC strategy, or whether BSC remains favorable when such spillovers occur.

The effects of these complications are separately studied on closed-loop beltways, the simplest facility type that retains the essential features of 2-dimensional street networks. This simple representation makes it easier to identify causal mechanisms between BSC and the complications of present interest. Analyses were performed by simulating traffic in microscopic detail, to maintain physical realism. Dimensional analysis was used as needed to reduce the number of essential inputs, so that analyses are thorough but do not succumb to the curse of dimensionality.

Findings are favorable, almost surprisingly so. As regards complication 1, we find that in the presence of directionally unbalanced demands, applying BSC to the heavy-flow direction improves a network’s overall performance, even on busy networks when demand in the unfavored travel direction approaches, but does not quite equal, that of the favored one.

As regards complication 2, we find that BSC can be especially effective in mitigating the problems that occur on hyper congested networks with short links. By completely filling the oversaturated links with standing queues, BSC postpones the onset of the gridlock process that begins when the tail of a queue spills over to its source, to become like a snake that has swallowed its tail. This phenomenon is known to occur on street networks, including grid networks.

As regards complication 3, we find that BSC remains beneficial in the presence of protected turn phases. We further find that BSC mitigates the disruptions caused when left-turn queues spill over from their turn pockets into adjacent through lanes. The spillovers reduce the usable space for storing queues and thereby cause the approach to behave as a de facto short link. BSC is then beneficial for the same reason it is favorable for complication 2.

We argue why the above findings should hold not just for congested beltways, but for street networks commonly seen in real cities. That said, we include discussion on how the findings can be applied in real-world settings, including major urban beltways, like those in Paris, Chicago, and elsewhere in the world.

I dedicate this dissertation to my family and all the people that have helped me along this journey.

Contents

Contents	ii
List of Figures	iv
List of Tables	v
1 Introduction	1
1.1 Background	1
1.2 Research Questions	4
1.3 Research Findings	5
1.4 Dissertation Roadmap	5
2 Literature Review	6
3 Experimental Setup	14
3.1 Study Sites: Idealized Beltways	14
3.2 Dimensional Analysis	16
3.3 Simulation Regimen	17
4 Imbalanced Directional Demands	18
4.1 Experimental Setup and Inputs	18
4.2 Outcomes and Implications	20
5 Inhomogeneous Link Lengths	22
5.1 Experimental Setup and Inputs	22
5.2 Simulation Results and Interpretation	23
6 Protected Left Turning Maneuvers	27
6.1 Experimental Setup	27
6.2 Outcomes and Implications	29
7 Findings and Future Research Directions	33
7.1 Summary of Findings	33

7.2 Future Research Directions	34
Bibliography	37

List of Figures

1.1	Signal Coordination using zero offsets (Adapted from Newell (1981))	2
1.2	Illustration of backward signal coordination	3
2.1	Example bandwidth for small, idealized system	7
2.2	Example coordination plan for a two-way arterial. Source: Newell (1989)	9
3.1	Examples of test sites used in the experiments	15
4.1	Simulation results	20
5.1	Geometric layouts of tested beltways: homogeneous benchmark, and two inhomogeneous beltways (moderately inhomogeneous and highly inhomogeneous). . .	23
5.2	Failure ratio ϕ across the three beltway configurations.	24
5.3	Time series of density for the shortest link (Link 1) and an upstream link (Link 3) under BSC and zero-offsets.	25
5.4	Average beltway density.	26
6.1	Link layout	28
6.2	BSC for two-way traffic, adapted from Figure 2.2	29
6.3	Trips completed vs. bi-directional demand.	30
6.4	Time series of density at turning pocket and through lane	31
6.5	Time series of density at link section upstream of turning pocket	32
6.6	Time series of density of short link from Figure 5.3	32
7.1	Real city beltways that can be used as testbeds	35

List of Tables

4.1	Supply-side inputs for simulation experiments	19
-----	---	----

Chapter 1

Introduction

Traffic signal coordination is widely used to reduce traffic delays and stops by adjusting signal offsets, i.e., the elapsed time between green initiations at neighboring intersections (FHWA, 2000). In undersaturated conditions where queues dissipate at every cycle and every intersection, synchronizing traffic signals to the forward motion of vehicles allows platoons to progress through the system at higher speeds with fewer stops. However, when residual queues in traffic grow long and spillback onto upstream links (known as hyper congested conditions), forward coordination becomes ineffective. This is because queue spillbacks block upstream traffic, which leads to wasted green time and degraded system performance.

Zero offsets, where signals serve opposing travel directions by simultaneously turning green across all intersections, help clear queues efficiently and is therefore a common approach to managing hyper congestion. However, recent studies suggest that Backward Signal Coordination (BSC), which synchronizes green phases to the backward-moving waves through residual queues, may provide greater benefits by addressing queue spillbacks. While promising, BSC has primarily been studied in idealized settings, free of commonly-occurring real-world complications.

This dissertation investigates the potential of BSC to improve traffic performance in hyper congested conditions, particularly when complications (i.e., directional demand imbalance, inhomogeneous link lengths, and protected left-turn phases) are present.

Before presenting the details of the research, this chapter provides an overview of the study's key components. Section 1.1 introduces the background and motivation for the dissertation. Section 1.2 outlines the research questions that this dissertation seeks to answer. Section 1.3 summarizes the main findings, and Section 1.4 describes the organization of the remaining chapters.

1.1 Background

Coordinating traffic signals along a street can improve vehicle flow, reduce the number of stops and thus reduce the severity of congestion (FHWA, 2000). Traffic engineers achieve

signal coordination, in large part, by adjusting the offsets (i.e., the time elapsed between the starting of the green times) between adjacent intersections (Roess et al., 2019).

Research has shown that the optimal choice of offsets depends on traffic conditions (Daganzo and Lehe, 2016). In undersaturated traffic, where vehicle queues at every intersection clear at the end of every green time, it makes sense to synchronize offsets to the forward motion of free-flow vehicles (Newell, 1981). But while it is often easy to achieve coordination in a single travel direction, it is often infeasible to do so in two opposing directions.¹ Thus, signal coordination is sometimes provided only to the direction with the heavier flows (Gordon et al., 1996).

When residual queues persist at the end of each green time at some or all intersections, most of the gains produced by forward signal coordination disappear (Newell, 1981). This is because residual queues prevent vehicles from traveling at free-flow speed, and can even block vehicles from entering congested links downstream. This results in lost or unused green time.




Figure 1.1: Signal Coordination using zero offsets (Adapted from Newell (1981))

One of the methods explicitly designed for hyper congested conditions is zero offsets; see

¹ It is possible to achieve signal coordination in both directions when the bi-directional travel times are the same on all blocks.

Figure 1.1. Under this coordination scheme, green times for opposing travel directions are initiated simultaneously at all intersections. Compared to forward coordination, zero offsets can reduce the impact of lost green times by allowing residual queues at an intersection to clear before new platoons arrive from upstream. In the absence of traffic turning onto the links, queues are fully contained on the links, as shown in Figure 1.1.

Hence, zero offsets are known to have a greater positive impact on streets with short link lengths, where queue spillbacks and lost green times are more damaging (Daganzo and Lehe, 2016). However, queue spillover can still occur (Newell, 1981) under this scheme when traffic turns on to the links from cross streets.

A promising alternative scheme for hyper congested conditions lies in synchronizing the signal offsets to the backward-moving waves that propagate through the queues; see Figure 1.2. Under this Backward Signal Coordination (BSC) scheme, a green phase begins at an intersection when the tail of the queue on the link immediately downstream starts moving forward. Daganzo and Lehe (2016) presents a theoretical study of BSC, finding that the scheme offers better performance than zero offsets when queue spillbacks occur.




Figure 1.2: Illustration of backward signal coordination

An adaptive scheme, one that switches between forward and BSC with changing traffic

conditions, was shown to outperform forward and zero offsets when each was deployed on a one-way beltway (Daganzo et al., 2017), and on a grid network serving traffic in all directions (Sadek et al., 2022).²

These earlier works highlight the promise of BSC. Yet their findings come from idealized scenarios, and it is unknown how BSC will perform in more complex settings. Much like forward coordination, and for the same reasons, BSC can rarely be achieved in two opposing travel directions on the same street. This means that signal coordination in the opposite direction is constrained and bound to be suboptimal. It remains unclear how BSC performs when applied to the heavy-flow direction under imbalanced directional demands. This uncertainty is greatest when directional demands are nearly balanced.

In addition, past investigations on BSC (e.g., Daganzo et al., 2017) focused on rotationally symmetric networks. The effect of short links in heterogeneous networks with variable link lengths remains unclear.

Finally, another unknown regarding BSC is how it performs in the presence of protected left-turn phases with queues that spillover from the turn pockets to the beltway. Sadek et al. (2022) includes these maneuvers in its experiments, but the effects of left-turning maneuvers are not reported. Left-turning maneuvers are well known to disrupt forward coordination schemes (Little et al., 1981). It is not known to what extent protected left-turn phases and the spillovers they can generate disrupt BSC, or whether BSC remains favorable when such spillovers occur.

1.2 Research Questions

This dissertation addresses three complications that may affect the performance of BSC:

- Imbalanced directional demands. When demands are heavier in one travel direction, can applying BSC to that direction produce net benefits compared to the zero offsets plan? Can BSC still be effective when directional demands are nearly balanced?
- Inhomogeneous link lengths. How does BSC perform on beltways that include short links with limited queue storage capacities? Can BSC postpone the onset of gridlock in these cases?
- Turning maneuvers. Is BSC still favorable if protected left-turning maneuvers are present? Can BSC mitigate the failure mechanisms that arise when left-turn queues spillover into adjacent through lanes?

This research uses microscopic simulation to investigate these questions. The effects of the above complications are separately studied on closed-loop beltways, the simplest facility type that retains the essential features of two-dimensional street networks (Daganzo and Lehe,

² Sadek et al. (2022) shows how to deploy BSC on large networks to achieve coordination in two orthogonal travel directions.

2016). This simple representation makes it easier to identify causal mechanisms between BSC and the complications of present interest. Dimensional analysis was used to reduce the number of essential inputs, so that analyses are thorough but do not succumb to the curse of dimensionality.

1.3 Research Findings

In the presence of directionally-unbalanced demands, applying BSC to the heavy-flow direction improves overall network performance relative to zero offsets. This advantage persists unless directional demands are nearly equal.

BSC can be even more favorable on beltways with inhomogeneous physical link lengths because short links provide less space for storing queues. By compacting queues within this constrained space, BSC delays the adverse upstream propagation of queues, thereby enhancing beltway performance.

BSC also remains favorable in the presence of protected left-turn phases. Moreover, it mitigates the disruptions that arise when left-turn queues spill over from their pockets and block adjacent through lanes. This spillover reduces the usable space for storing queues, so the approach behaves as a de facto short link. By compacting queues, BSC enables this constrained space to hold more queued vehicles and shortens queue expansion. This can prevent or postpone gridlock.

1.4 Dissertation Roadmap

The literature is reviewed in Chapter 2. The experimental sites are described in Chapter 3. Chapters 4–6 evaluate BSC under three complications: imbalanced directional demands (Chapter 4), inhomogeneous link lengths (Chapter 5), and protected left-turning maneuvers with left turn spillover queues (Chapter 6). The findings are summarized, and future research directions are proposed in Chapter 7.

Chapter 2

Literature Review

Numerous publications address how one might coordinate network traffic signals. Many proposed strategies seek to carve out a “bandwidth” (i.e., green windows that span neighboring signals in a network) so that drivers might travel their routes uninterrupted by red phases, e.g., Webster (1958); Robertson (1969); Cohen (1983); Little et al. (1981); Tsay and Lin (1988); Gartner (1983); Gartner et al. (1991); Roess et al. (2011); Park et al. (2000).

An idealized example is shown via the time-space diagram in Figure 2.1. It entails three intersections, a single (upward) travel direction, and no vehicles turning on or off the arterial. Note the signal offsets in the figure, defined in each instance as the elapsed time between the initiation of green phases at an upstream intersection and the start of green at the intersection immediately downstream. By synchronizing the offsets to the forward motion of freely flowing vehicles, uninterrupted travel may be feasible for a simple scenario like the one in Figure 2.1. Uninterrupted travel becomes virtually infeasible for larger, more complex real-world networks, however. Reasons are presented as part of the literature review to follow.

We start by reviewing signal synchronization strategies for the simplest network forms, i.e., arterials. We review strategies for 1-directional arterials in Section 2.1, and 2-directional arterials in Section 2.2. We review strategies for 2-dimensional (e.g., grid) networks in Section 2.3. The discussion will reveal that many of the methods suffer from one or more of the following three pitfalls that reduce their effectiveness.

Pitfall 1: Myopic optimizations. The synchronization strategies almost invariably determine the green and red durations to be used at each signalized intersection on a network based upon the intersection’s local demands. Many of the strategies are adaptive, such that each intersection’s signal timing is periodically updated, as local conditions change. This kind of myopic (fragmented) optimization can steadily degrade network performance over the longer run. This is because drivers seek to improve their lot by changing routes as days go by (Smith, 1979). This gaming behavior tends to concentrate traffic on certain streets, adding to queues and delays on those links, while leaving other streets underused.

Pitfall 2: Unrealistic Objectives. Most strategies aim to enable all drivers to complete their trips without encountering a red phase. While desirable, the outcome is infeasible in




Figure 2.1: Example bandwidth for small, idealized system

all but the most idealized cases such as the case in Figure 2.1. By focusing on the infeasible, the strategies overlook compromise solutions that can be effective.

Pitfall 3: Poor adaptations. Though many synchronization strategies are said to be adaptive, most fail to adapt in suitable fashion when networks exhibit long queues that spill over to multiple street links. In these hyper congested conditions, most strategies continue to synchronize signals to the forward motion of vehicles, i.e., they attempt to usher traffic onto links already filled with queued vehicles. This constitutes little more than an exercise in futility.

We complete our literature review in Section 2.4. It describes efforts to address these pitfalls, in part or in whole, by formulating strategies suitable for hyper congested networks. We describe in Section 2.5 how findings from these latter efforts, and questions that remain unanswered in the literature, have informed the present research.

2.1 One-way Arterials

Even for the simple case in Figure 2.1, the green windows at intersections 1, 2, and 3 are not of equal duration. This is typical of what occurs when cross-street demands vary across intersections. The varying green window is a consequence of the myopic approach used to optimize signal timings that we call Pitfall 1. In this approach, green times at each intersection are allocated in proportion to the relative needs (vehicle demands and queue discharge flows) across all travel directions. Allocating greater (lesser) green duration to an intersection's cross streets in response to its local needs leaves less (more) green time to serve arterial traffic at the intersection.

The practical implication is important: the number of platooned vehicles that can travel uninterrupted through a system, including a simple system like in Figure 2.1, can be dictated by the critical intersection where the green window is smallest (intersection 3 in the figure). If arterial platoons cannot be fully accommodated at that intersection, residual queues will form and grow there. Hyper congestion occurs if the queue eventually fills the approach link and possibly others upstream.

Most existing synchronization strategies seek to soften the blow by maximizing bandwidths (Little, 1966; Little et al., 1981; Cohen, 1983; Tsay and Lin, 1988). A maximized bandwidth might still be too small to fully accommodate platoons, however. The concern is discussed in Newell (1981), but otherwise receives limited attention in the literature.

2.2 Two-way Arterials

We start the discussion using an idealized case in which the spacing (or more generally, the free-flow vehicle travel time) between intersections is uniform across the arterial; and no vehicles turn on or off the arterial. An example of this case is shown in Figure 2.2 (Figure 4.4 (a) in Newell, 1989). Note from the figure how platoons in both directions flow uninterrupted through the system. Further note how the bi-directional platoons in this idealized case overlap at every intersection. The figure shows that this overlap requires that “trip time... up and back between adjacent intersections” be an integer multiple of the cycle length used for all intersections in the system (Newell, 1989). The multiple used in Figure 2.2 is 1.

In a real-world system with typical link lengths and vehicle speeds, achieving this overlap at every intersection would almost invariably force the selection of an infeasibly short cycle length, i.e., one that cannot accommodate cross-street demands or pedestrian requirements; see again Newell (1989). Choosing a longer cycle length creates problems of its own. One can appreciate these by imagining Figure 2.2 to span a greater distance, and by envisioning additional (i.e., intervening) intersections to the ones shown in the figure. Consider how bi-directional platoons would not overlap nicely at the newly-added intersections. More green time would therefore be needed at those locations, if uninterrupted travel is to be provided in both directions. This can leave far too little cycle time available for cross-street vehicles and pedestrians.




Figure 2.2: Example coordination plan for a two-way arterial. Source: Newell (1989)

One can further use Figure 2.2 to appreciate the complication that arises when link lengths (and travel times on the links) vary. In theory, one might achieve overlapping bi-directional platoons at these intersections by using link-specific vehicle speeds for choosing offsets; see Little (1966) and Little et al. (1981). To this end, one might select higher (lower) speeds for longer (shorter) links. Problematically, this solution assumes that drivers somehow know how to change their speeds to align themselves with each link’s offset. The assumption seems heroic.

Traffic turning on and off the arterial further complicates the problem. The reader can refer again to Newell (1989) for discussion on this matter.

These complications are why we label efforts to synchronize signals bi-directionally as Pitfall 2. Given the insurmountable nature of these complications, signals are often synchronized in one travel direction only, leaving the opposing direction unsynchronized (Gordon et al., 1996). Favoring the direction with higher flow benefits more vehicles, which can reduce the time collectively spent in the system.

A variation on this idea is offered in Yang et al. (2015) for the case of a freeway interchange, i.e., the junction of an arterial and a freeway. In this special case, synchronization went to the higher-demand direction on the arterial. Those green phases were also synchronized with a left-turn phase that ushered traffic in the favored direction from the arterial to the freeway. Another variation is offered in Sadek et al. (2022) for the case of two-dimensional grid networks. These networks are discussed next.

2.3 Two-dimensional Networks

Further complications arise when synchronizing signals on grids and other 2-dimensional street networks. Although as noted in Sadek et al. (2022), one “can synchronize all N-S streets on a grid in one direction (say from N to S), choosing the offsets at all intersections of these streets fixes the offsets across the grid. Therefore, no additional adjustments can be made to guarantee that the intersecting set of streets is also synchronized.” The response to this complication has typically been to decompose a network and search for synchronization schemes for each sub-network. For example, Gartner and Stamatiadis (2004), Liu and Chang

(2011), and Hao and Dong-cheng (2020) propose methods that synchronize signals along select corridors within a network.

The works invariably report that their strategies bring (modest) improvements over various baseline solutions and for various case-study networks. But the methods still suffer from Pitfalls 1 and 2. The same is true of widely known computer packages, including SYNCHRO, HCS-Signals and PASSER II (Cubic Transportation Systems, 2024; McTrans Center, University of Florida, 2023; Texas Transportation Institute, 1991; Federal Highway Administration, 2009; Lieberman et al., 2000).

Work in Sadek et al. (2022) proposes a synchronization strategy for rectangular grids that resolves Pitfalls 1 and 2. As regards Pitfall 1, the strategy abandons efforts to synchronize signals for vehicles traveling in the lower-demand directions on a grid. The strategy instead synchronizes signals on all links that point toward clustered workplace destinations during morning periods, and on all links that point away from those clusters during the evenings.

The strategy achieves near-perfect synchronization for these higher-demand travel directions by resolving Pitfall 2 in a straightforward way. Rather than optimizing signal timings at each of the grid’s intersections in separate fashion, the strategy uses both a cycle length and a signal timing plan (i.e., green splits) that are common to all signals on the grid. On busy networks, one might choose the largest cycle length and green times practicable. By predefining the signal timings, only the offsets are optimized for the grid.

The optimization entails defining offsets relative to a reference intersection, defined as the intersection closest to the center of gravity of clustered workplace destinations. (A grid can have multiple said clusters and a distinct reference intersection for each cluster.) During the morning periods, the offset for an intersection j is simply the free-flow vehicle travel time on the shortest path between j and its reference intersection with modulo cycle length. Since all signals serving a single inbound direction share a common green, an inbound vehicle that travels through an intersection without turning at the start (end) of a green reaches the next intersection at the start (end) of its green, and so on down the line.

Hence, perfect forward progression is achieved on all streets that point in the directions toward a reference intersection. Since every inbound link on a grid lies along a shortest path to the reference intersection, inbound drivers in the morning see a red phase only in the wake of having made a turn maneuver (and at most only one such maneuver is required of any trip on a grid). The same near-perfect progression can be achieved in the evenings by setting the offsets to favor the opposing (outbound) travel directions.

2.4 Congested Networks

The strategies cited above are suitable for freely flowing or lightly congested network traffic conditions. The strategies perform poorly, however, on hyper congested links filled or nearly filled with spillover queues: Upon receiving green phases, drivers are blocked from moving forward by queues occupying the links downstream. Still, we see in the literature what we label Pitfall 3, i.e., attempts to usher traffic onto congested links; see widely used software packages such as TRANSYT-7F and Synchro (Wang et al., 2013; Chen et al., 2013).

A handful of studies have sought to address this pitfall. We review nascent, and unsuccessful attempts that used black-box solutions in Section 2.4.1. We describe alternative synchronization strategies that consider the physics of congested network traffic in Sections 2.4.2 and 2.4.3.

2.4.1 Emergent Methods

A growing line of machine-learning studies tackles signal coordination on congested networks, yet three stubborn gaps keep these models from leaving the lab.

First, their models hinge on data that are rarely available at network scale. As examples, Li et al. (2021) requires a cycle-by-cycle, lane-segment count of every vehicle, and Mon et al. (2024) builds both state variable and performance measurements on exact queue lengths at every approach to the intersections. Such noise-free, network-wide measurements would demand 100% detector coverage, zero calibration drift, and instantaneous transmission, which are conditions that are virtually absent outside test corridors. When even a fraction of these data are missing, noisy, or delayed, the algorithms can no longer form the states or compute the rewards they were trained on, and their performance can degrade sharply.

Second, even if all of the data were obtainable, machine-learning based strategies would have heavy computational burden. As examples, coordination schemes of Kuyer et al. (2008) and van der Pol and Oliehoek (2016) require every junction to exchange joint-action utilities with all neighbors at every decision epoch. The message-passing burden grows combinatorially with network size, making large-scale network deployment infeasible.

Third, the proposed strategies remain black-box controllers. As Schreiber et al. (2021) notes: “Deep reinforcement learning has shown potential for traffic signal control. However, the lack of explainability has limited its use in real-world conditions.” The internal logic of machine-learning strategies cannot be inspected, and the strategies’ responses under hyper congested conditions cannot be predicted or guaranteed. That opacity prevents analysts from diagnosing failures, verifying stability when spillbacks occur, or adapting the controller when demand patterns shift—and it provides no assurance that the learned policy will avoid the Pitfall-3 failure mode of releasing vehicles into downstream links that are already queue-filled under hyper-congestion.

2.4.2 Zero Offsets

As the name implies, this synchronization strategy initiates green times for opposing travel directions simultaneously at all intersections (Newell, 1981). The simultaneous synchronization for one travel direction (say northbound) is typically deployed for the opposing (southbound) direction as well. For cases when traffic turns on and off the street at modest rates, green times are typically identical across all signals (Newell, 1981). In these cases, through-moving vehicles travel along the street for a period equal to that of the green duration before stopping at a red. With light turning traffic, moreover, the strategy prevents queues from spilling-over to links upstream. This is because the number of vehicles stopped

at any approach link to a j th signalized intersection is approximately equal to the number that had been occupying that link at the onset of the red phase. That number can fit on the approach link during the red because the vehicle spacings are compressed in the jammed queue.

Queue spillovers can still occur with zero offsets if during the red phases vehicles turn onto a link in sufficient numbers. Consequently, Daganzo and Lehe (2016) found that on hyper congested links, zero offsets rarely outperform BSC. We discuss previous research on BSC next.

2.4.3 Backward Signal Coordination

The goal is to prevent vehicles from being blocked by queues on downstream links. To this end, the strategy synchronizes signals to the backward motion of the kinematic waves that propagate through (spillover) queues. As a result, voids (empty road spaces) appear at the upstream end of congested links in time to absorb flows arriving from upstream. The idea dates to work in Pignataro (1978) and Newell (1981), and later Rathi (1988), Quinn (1992), Abu-Lebdeh and Benekohal (1997), and Christofa et al. (2013). These works considered one-way traffic along a single arterial.

Work in Daganzo and Lehe (2016) was the first to explore how BSC might be implemented in a straightforward way using observable traffic features along a 1-way arterial. It also clarified that the effectiveness of BSC depends on both a link's physical length and the effective green time in each cycle. The work further showed how real-time measurements of traffic densities can be used to determine when to deploy the backward strategy; and how said deployment can increase link capacities. Follow-up work in Daganzo et al. (2017) demonstrated how signal offsets can be adaptively set to accommodate an arterial's evolving traffic states, i.e., from undersaturated to hyper congested and back again. More recent work in Sadek et al. (2022) extended the work to two-dimensional grids. It demonstrated how BSC can be periodically deployed inside select zones within cities, based upon the traffic densities averaged across all links within each zone.

2.5 Lessons Learned

Though many of the synchronization strategies in the literature suffer from the pitfalls we number 1 to 3, some works have addressed these pitfalls with good results. The present dissertation builds upon these past successes, particularly the recent successes in Sadek et al. (2022) and Daganzo et al. (2017).

Like in Sadek et al. (2022) and Gordon et al. (1996), the dissertation avoids Pitfall 1 by using a cycle length and phase timings that are common to all signals in our systems. And like in Sadek et al. (2022), the present work avoids Pitfall 2 by typically synchronizing signals in the heavy-demand direction along a street, while leaving drivers in the opposing direction to face unsynchronized signals.

Our primary focus is to explore how one might avoid Pitfall 3 via the use of BSC. To this end, we build upon successes in the literature while seeking to fill some of the present gaps.

Daganzo and Lehe (2016) showed that the potential for BSC to improve flow on a congested one-way arterial depends upon link lengths and effective green times. We extended the analysis to bi-directional beltways in Chapter 3.

Works cited in Section 2.4 further confirm the effectiveness of BSC in prioritizing movements in heavier-demand direction (Sadek et al., 2022). But systemwide improvements only occur when the opposing demands are sufficiently lower than in the favored direction. The benefits are shown in Sadek to diminish as directional demands approach parity. Yet, the simulated cases in Sadek entailed limited use of BSC: it was periodically deployed in central zones within cities that were subject to hyper congestion. Forward signal progression remained always in place elsewhere in the cities.

In short, other factors involved in the simulated case studies of Sadek et al. (2022) obscured the effects of opposing demands on BSC. Hence, the question remains: At what point are the differences in directional demands so small that using BSC to favor one travel direction over the other no longer reduces trip times systemwide? That question is explored in Chapter 4.

Christofa et al. (2013) studied BSC using a one-way arterial with link lengths that varied block by block. This sort of geometric inhomogeneity is common in real-world cases. But how these varied lengths impact the effectiveness of BSC is yet to be explored. We do so in Chapter 5.

Aside from Sadek et al. (2022), few studies have examined BSC under conditions that include left-turning maneuvers. Although Sadek et al. (2022) included those maneuvers, it did not determine whether protected left-turn phases and left-turn pocket spillovers diminish the benefits of BSC. This issue is examined in Chapter 6.

Finally, Sadek et al. (2022) shows that the beneficial features of deploying BSC on one-dimensional (arterial) networks can also be achieved on two-dimensional (grid) networks. This gives us license to perform our studies on the simplest representation of a two-dimensional network: A bi-directional closed-loop beltway, e.g., see Daganzo et al. (2017).

Chapter 3

Experimental Setup

The test sites used in the evaluations are presented in Section 3.1. Our means of limiting dimensionality in the analysis are described in Section 3.2. The simulation regimen is discussed in Section 3.3.

3.1 Study Sites: Idealized Beltways

The evaluations use three types of idealized, closed-loop beltways, one type for each factor of interest: (i) imbalanced directional demands, (ii) link lengths that vary block by block, and (iii) vehicular left-turn maneuvers with protected signal phases. All beltways serve traffic in two directions.

The rotationally symmetric beltway illustrated in Figure 3.1(a) is used for studying the Factor (i).¹ It is a beltway with eight bi-directional cross streets that intersect the beltway at 90-degree angles, and are equally spaced along the beltway circumference. To eliminate complications from vehicles' lane-changing maneuvers, both the beltway and its cross streets have a single lane in each direction.

The beltway in Figure 3.1 (b) is used for studying the Factor (ii).² This layout retains the features of the beltway used for Factor (i), except that this second beltway type includes a single short link.

The beltway in Figure 3.1(c) is used for studying Factor (iii). This beltway type has homogeneous link lengths and includes left-turn pockets at every intersection to serve both travel directions on the beltway.

As noted in Daganzo and Lehe (2016) and Daganzo et al. (2017), beltways reproduce the mechanism of urban gridlock because they form a closed system with finite link storage. Queue spillbacks can thus propagate from one intersection to the next and eventually wrap around the beltway. Once that happens, the beltway will soon become gridlocked:

¹The layout shown in Figure 3.1(a) is one representative example from a broader set of beltways used to examine Factor (i).

²This beltway layout is one example from a broader set of beltways used to examine Factor (ii).

queues prevent discharge at upstream intersections, circulating flow collapses, and vehicles can no longer advance—mirroring the systemwide paralysis that could occur in real street networks under extreme circumstances, such as a mass evacuation (Olmos et al., 2018). Our use of beltways also eliminates network features that have only modest effects on network performance but can obscure causal mechanisms.




Figure 3.1: Examples of test sites used in the experiments

All of a beltway’s intersections are controlled by pretimed traffic signals with common green times and a single cycle length. This arrangement is motivated by findings in Sadek et al. (2022), as reviewed in Chapter 2. We vary these signal settings parametrically. The specific variations are described in later chapters.

In all experiments for Factors (i) and (ii), the signals use an equal green split: 50% of the cycle for beltway traffic and 50% for cross streets. In the experiments for Factor (iii), a protected left-turn phase was added at every intersection to serve vehicles turning off the beltway. The specific timing adjustments for this set of experiments are described in Chapter 6.

Trip origins and destinations occur midblock on the beltway and at the extreme ends of each cross street, as shown by the circular markers in Figure 3.1(a). The fixed, time-independent OD tables used for each experiment produced a uniform pattern in which the beltway demand in a single (e.g., clockwise) travel direction was the same across all links. Trips originating on a cross street always turn, either rightward or leftward, onto the beltway. Trips always follow their shortest paths and never span more than four (out of eight) beltway intersections before exiting the system.

3.2 Dimensional Analysis

Even for the idealized sites described above, the effectiveness of BSC is dictated by several input variables:

- q_l and q_h — vehicular demands in the low- and high-demand directions, respectively (veh/h);
- B_l — the length of link l (m);
- V_w — the speed of backward-moving kinematic waves in queues (m/s), assumed independent of flow, as in Newell (1981, 1989);
- G — the common green duration to all beltway traffic (s).

We use dimensional analysis to reduce the number of input variables. Following the approach outlined by Daganzo (1997) (see ch. 3.2.2), variables are combined into unitless forms that capture their essential relationships while eliminating dependence on measurement units.

Hence, we denote the unitless variable

$$\delta = \frac{q_l}{q_h}, \quad \delta \leq 1 \quad (3.1)$$

as the directional demand ratio. Values of δ that are closer to zero indicate substantial directional imbalance, while values close to one denote directional balance.

We denote the variable

$$\lambda_l = \frac{B_l}{|V_w|G}, \quad (3.2)$$

as the unitless measure of link l 's queue storage capacity. The denominator in Eq. 3.2 is the distance that a backward wave can travel during an entire green duration. Larger values of λ_l indicate that link l has greater capacity to store queues.

Output from each evaluation is the total (bi-directional) throughput that can be accommodated on the beltway before inducing failure, defined by one of two possible events. The first possibility is gridlock, which begins when the queue wraps entirely around the beltway,

becoming like a snake that has swallowed its tail (Daganzo et al., 2017). The beltway’s queued flows then diminish over time until traffic grinds to a halt, i.e., to failure.

As a second possibility, a beltway fails if queues from side-street traffic fill their entire approach links. Each side-street link was sized to accommodate approximately 50 jammed vehicles. Under this failure condition, vehicles generated at side-street origins cannot enter the system.

We denote a beltway’s maximum bi-directional flow immediately preceding failure as Q_B under BSC, and Q_Z under zero offsets. The results are expressed as the unitless variable

$$\phi = \frac{Q_B}{Q_Z}, \quad (3.3)$$

which we call the failure ratio. Values of $\phi > 1$ indicate that BSC outperforms zero offsets, while $\phi < 1$ indicates the opposite.

Our experiments evaluate ϕ for various δ and λ_l combinations within ranges typical of real-world settings. To maintain physical realism, evaluations are performed via simulation, as described below.

3.3 Simulation Regimen

All experiments were implemented in Aimsun, a microscopic traffic simulation platform. For each (δ, λ_l) combination, 10 independent simulation runs were executed using different random seeds, and the resulting outcomes were averaged to reduce stochastic noise. Each simulation ran for 1 hour, a horizon long enough for either form of failure (gridlock or blocked side-street entry) to manifest itself.

Chapter 4

Imbalanced Directional Demands

This chapter explores the effect of Factor (i): directional demand imbalance on beltway performance. We do this by examining whether and when deploying BSC for the heavier-demand direction can raise a beltway’s failure threshold above that of the zero offsets strategy. The experiments thus seek to quantify the tradeoffs in the “one-sided” nature of BSC, which inherently favors one travel direction over the other.

Inputs used for the simulation experiments are presented in Section 4.1. Outcomes and implications are in Section 4.2.

4.1 Experimental Setup and Inputs

Bi-directional traffic was simulated on the rotationally symmetric beltway in Figure 3.1(a). Simulations were performed separately for BSC and for zero offsets. The parametric analysis for each coordination strategy considered a range of values for the directional demand ratio, δ , and each δ was analyzed under incrementally increasing bi-directional demands; see Section 4.1.1. We also parametrically varied the dimensionless queue storage capacities, λ ,¹ see Section 4.1.2.

4.1.1 Demand

Demands were specified to induce hyper congested traffic in the heavy-demand direction. We varied the directional demand ratio, δ , from 0.25 (highly imbalanced demands) to 1.00 (directionally balanced rates). For each value of δ , a baseline bi-directional demand of 960 veh/h was parametrically increased in increments of 10 veh/h. For each δ , the beltway’s maximum bi-directional flow immediately preceding failure was recorded, for BSC and for zero offsets.

¹Because all beltways used in this chapter are rotationally symmetric, the link-specific storage measure λ_l reduces to a common beltway value, denoted here by λ .

4.1.2 Supply

Supply-side inputs relating to the beltway’s geometry and traffic signal timing were selected to accommodate a peculiarity when maintaining BSC on a closed loop. To appreciate the issue, one might envision the tail of a backward kinematic wave as it emerges from an arbitrarily selected reference intersection on the beltway. Suppose that the wave emerges at the start of a cycle. The wave traverses the loop (against the flow of traffic) and eventually returns to its origin, the reference intersection.

Suppose that the wave’s return does not coincide with the start of another cycle but occurs at some other time in that cycle, say towards its middle. In this problematic case, the final link along the wave’s closed-loop path (i.e., the beltway link just downstream of the reference intersection) will be asynchronous with BSC. This resulting “weak link” would likely trigger failure (see again Section 3.2) in all experiments and would thus distort our findings.

We avoided this problem by satisfying what we call the “loop-closure constraint.” We did so by selecting suitable combinations of a beltway’s common link length, B , and the common green duration, G .² We parametrically varied B from 100 to 200 m, to align with real-world urban settings; see column 1 of Table 4.1. We chose reasonable (i.e., realistic) integer values of G for each B to satisfy the constraint, albeit approximately in some cases.³

These combinations of B and G were also selected to produce a reasonable range of λ , the beltway’s dimensionless queue storage capacity, as per Eq. (3.2) and for a wave speed, $V_w = -4.2 \text{ m/s}$ (-15 km/h). As shown in column 3 of the table, $0.43 < \lambda < 0.85$. The effect of these λ on beltway performance (failure threshold) is estimated in the following section.

Table 4.1: Supply-side inputs for simulation experiments

(1) B [m]	(2) G [s]	(3) λ
100	28	0.43
150	40	0.57
200	52	0.65
200	46	0.72
200	44	0.77
200	40	0.85

² Experiments use equal green splits (50/50) at all intersections with a common cycle length, see again Section 3.1.

³Some combinations of B and G did not satisfy the loop-closure constraint exactly, but discrepancies never exceeded 6 s. When a discrepancy occurred, it was evenly distributed to the offsets across all signals on the beltway.

4.2 Outcomes and Implications

Building on the demand–supply (δ - λ) inputs defined in Section 4.1, we compare the failure ratios, ϕ , between BSC and zero offsets. The outcomes from simulations are presented in Section 4.2.1 and their implications in Section 4.2.2.

4.2.1 Experimental outcomes

The surface shown in Figure 4.1 is that of failure ratio, ϕ , over the selected ranges of directional-demand ratio, δ , and dimensionless queue storage capacity, λ . Shading depicts the value of ϕ for each combination of δ and λ . Lighter shades (values above 1.0) indicate that BSC outperforms zero offsets. Darker shades (values below 1.0) indicate the opposite. The shaded surface thus delineates the δ - λ regions in which deploying BSC for the heavier-demand direction can raise a beltway’s failure threshold above that of the zero offsets.




Figure 4.1: Simulation results

Notice how maximum values of ϕ of around 1.14 occupy the lower left portion of Figure 4.1. These ratios favorable to BSC occur when directional demands are highly imbalanced ($\delta \lesssim 0.80$) and the beltway’s queue storage is quite limited ($\lambda \lesssim 0.52$). As the eye moves upward along the left-hand-side of the figure, toward higher λ , the ϕ progressively shrinks

toward 1. When moving rightward toward more balanced directional demands, the surface descends to $\phi = 1.0$ (at $\delta \approx 0.90$) and thereafter descends below 1.

4.2.2 Interpretation and Implications

The findings indicate that BSC outperforms zero offsets whenever directional demands are even modestly imbalanced, such that $\delta \lesssim 0.90$. In these cases, the benefits of providing BSC to the higher-demand direction outweigh any negative consequences from leaving the opposing (lower-demand) direction to confront unsynchronized signals. Figure 4.1 shows this to be true for the full range of dimensionless queue storage capacity, λ , tested.

Not surprisingly, the figure shows that the advantages of BSC are greatest when δ is well below 0.90, meaning when directional demands are highly imbalanced, and when λ is small. In these cases, hyper-congestion intensifies in the higher-demand direction, such that greater need exists for BSC. This pattern was previously identified in Daganzo and Lehe (2016).

When directional demands are more balanced, such that $\delta \gtrsim 0.90$, the benefits that BSC delivers to the higher-demand direction no longer compensate for the losses suffered in the opposing direction. In these cases, the failure ratio, ϕ , falls below 1, meaning that a zero offsets plan accommodates greater beltway flow and is thus the preferred option.

Consequently, adopting BSC for the heavier travel direction emerges as the optimal strategy for managing two-way traffic flow except when directional demands are nearly equal.

Chapter 5

Inhomogeneous Link Lengths

This chapter investigates Factor (ii): inhomogeneous link lengths by comparing the performance of BSC and zero offsets on beltways. The analyses in Chapter 4 considered beltways with uniform link lengths, where queue spillovers occur in a predictable fashion. Real street networks, however, often include short link lengths that provide less storage for queues. These short links are prone to queue spillover, which can degrade signal coordination and even trigger gridlock.

The simulation experiments in this chapter test whether short links diminish or amplify the benefits of BSC. Inputs used for the experiments are presented in Section 5.1. Outcomes and their mechanism are provided in Section 5.2.

5.1 Experimental Setup and Inputs

Figure 5.1 presents the three beltway configurations tested in the chapter. Unlike the homogeneous beltway studied in Chapter 4, two of the beltways studied in this chapter include a short link. The short link is noted as link 1 in the figure followed by seven longer links of equal length (which include the link labeled link 3 in the figure). The total beltway physical length is left unchanged from its homogeneous version.




Figure 5.1: Geometric layouts of tested beltways: homogeneous benchmark, and two inhomogeneous beltways (moderately inhomogeneous and highly inhomogeneous).

Figure 5.1(a) presents a homogeneous beltway in which all links have a common dimensionless storage capacity of $\lambda = 0.78$. Figure 5.1(b) shows a moderately inhomogeneous beltway in which the short link labeled link 1 has a dimensionless storage capacity of $\lambda_1 = 0.48$ and links 2–8 have $\lambda_i = 0.81$. Figure 5.1(c) shows a highly inhomogeneous beltway in which Link 1 has a dimensionless storage capacity of $\lambda_1 = 0.39$, while links 2–8 have $\lambda_i = 0.85$.

Demand was loaded onto the beltways and cross streets in the same manner as in Chapter 4 and incrementally increased until failure (as defined in Section 3.2). The directional-demand ratio, δ , varied from 0.25 (highly imbalanced) to 1.0 (balanced). For each δ , the beltway's maximum bi-directional flow immediately preceding failure was recorded, for BSC and for zero offsets. Performance is again measured via the failure ratio, ϕ .

5.2 Simulation Results and Interpretation

Figure 5.2 presents the failure ratio, ϕ , across the three beltway configurations: the homogeneous site benchmark (checkerboard pattern in the figure), the moderately inhomogeneous beltway (grid pattern), and the highly inhomogeneous beltway (cross-hatched pattern). Recall that failure ratios greater than 1.0 indicate that BSC achieves a higher bi-directional flow prior to failure than zero-offsets.

The results show that BSC consistently outperforms zero-offsets on all three beltways until δ approaches unity and zero offsets become the better-performing strategy. Failure ratio, ϕ , is greater when directional demands are highly imbalanced. Note from Figure 5.2, the magnitude of the advantage depends on the severity of inhomogeneity. The capacity gain is most pronounced in the highly inhomogeneous case (cross-hatched), where the shortest link was only $\lambda = 0.39$. The moderately inhomogeneous case (grid, $\lambda = 0.48$) still shows that BSC typically outperforms zero offsets, though to lesser degrees. The homogeneous benchmark (checkerboard) is the beltway that benefits least from BSC.




Figure 5.2: Failure ratio ϕ across the three beltway configurations.

Figures 5.3 (a) and (b) confirm that BSC postponed failure relative to zero offsets. The time-series curves in both figures present links' densities. The curves were constructed from a single simulation run of traffic traveling in the counterclockwise (high-demand) direction of the highly inhomogeneous beltway. The solid curves show density in (short) link 1 and the dashed curves describe link 3 upstream. The curves in Figure 5.3(a) are the density produced by zero-offsets and the curves in Figure 5.3(b) by BSC.

We turn first to the zero-offset case in Figure 5.3(a). Note the sudden rise in Link 1's density at $time \approx 11$ min. The beltway queue filled (short) link 1 shortly thereafter and spilled over to upstream link 3 approximately 7 minutes later at $time \approx 18$ min. The density on both links thereafter increased, and typically above 100 veh/km.

Failure occurred some minutes later. Note from Figure 5.3 (a) how the density from link 1 jumped for a second time at $time \approx 32$ min. This is when the queue fully wrapped itself around the beltway and returned to its origin on link 1, i.e., the beltway queue now resembled a snake that swallowed its tail. The event brought flow on link 1 to a halt. Note how density plateaued. Wiggles then disappeared from the solid curve in the figure because the beltway vehicles came to a stop. The jammed state soon propagated to link 3, and gridlock occurred across the entire beltway (i.e., the system failed) shortly thereafter.

Figure 5.3 (b) reveals that something better occurred under BSC. The time-series curves in that figure show that the queue that formed on (short) link 1 still filled that link and spilled over to link 3 upstream. But with BSC, the queue never wrapped itself around the entire beltway to trigger gridlock (failure). This is evident in that the curves in Figure 5.3 (b) do not exhibit the second jump and wiggle-free plateaus seen in Figure 5.3 (a) under




Figure 5.3: Time series of density for the shortest link (Link 1) and an upstream link (Link 3) under BSC and zero-offsets.

zero offsets.

BSC prevented gridlock because it did a better job of compacting beltway queues. To see an example of this, average density was measured across all 8 links of highly inhomogeneous beltway. Measurement occurred over the interval from $t \approx 11$ min to $t \approx 30$ min, that is, from the time queues first arose to the onset of gridlock under zero offsets. Figure 5.4 shows that the average link density over this interval is higher under BSC than under zero offsets: 87.76 veh/km under BSC versus 71.06 veh/km under zero offsets, which is an increase of 23.51%.

This difference in density shows the preferred outcome occurred because BSC causes queues to compact and grow denser as described in Daganzo and Lehe (2016). The compaction enables links, whether short and long, to hold more queued vehicles. This, of course, shortens queue expansion. Additionally, by synchronizing green times with the movements of queues just downstream, BSC enables vehicles to reach their exit locations and leave the beltway at a higher rate.

In summary, Figure 5.2 shows that BSC outperforms zero offsets (except under balanced directional flows). Figure 5.3 and Figure 5.4 reveal the reason for BSC's superior performance.




Figure 5.4: Average beltway density.

Chapter 6

Protected Left Turning Maneuvers

The present chapter investigates conditions involving Factor (iii), vehicular left-turn maneuvers, since these can reduce intersection capacity and induce residual queueing (Lin et al., 1984), with detrimental effects on signal coordination. We examine cases in which left turns from the beltway to the side streets occur via exclusive turn pockets and under protected turn phases. We test whether BSC deployed in these cases can still raise a beltway's maximum bi-directional flow (prior to failure) above that of a zero-offset strategy. The setup used for the simulation experiments is described in Section 6.1. Outcomes are presented in Section 6.2.

6.1 Experimental Setup

The beltway introduced in Figure 3.1(c) was used for the present tests. It has a uniform link length, such that $B_l = 100 \text{ m}$ for all links l . Turn pockets of 50 m length were provided at every intersection along the beltway, and both for its clockwise and counterclockwise travel directions (see Figure 6.1).

Beltway demands were set equal in both directions (i.e., $\delta = 1$), with a baseline bi-directional demand of 960 veh/h that was parametrically increased in increments of 480 veh/h across experiments. In each experiment, 30% of vehicles were designated to turn left at specified intersections along the beltway. The other 70% ended their trips either by turning right at intersections to exit the beltway or upon reaching a mid-block location on the beltway. Each vehicle's destination was established upon its entry into the system.




Figure 6.1: Link layout

BSC was configured to favor the through movements in the two beltway directions equally. The only cycle length that can achieve this bi-directional coordination (given $B_l = 100$ m) is 48 s. The plan is illustrated for three beltway intersections in Figure 6.2.¹

In each experiment, the phase timings were common to every intersection on the beltway. The beltway through movements received 50% of the cycle length as green time (i.e., 24 s), the remaining green was separated between green time for side streets and left turning maneuvers from the beltway.

Two turn-phase designs were used in separate experiments to produce different propensities for left-turn pocket spillovers. In some experiments, a 7 s green interval was assigned to the turn phases to create a moderate propensity for spillover. A 5 s green interval was assigned in other experiments to create a high propensity for left-turn pocket spillover.²

These two distinct phase durations were evaluated for a range of demands. All other details of the experiments were performed as in Chapters 3 and 4.

¹As per discussion in Chapter 2, this bi-directional coordination requires that a backward wave's trip time "up and back between adjacent intersections be an integer multiple of the cycle length used for all intersections in the system" (Newell, 1989). The multiple used here is 1.

²Because startup lost time is approximately 3 s, increasing the turn-phase green from 5 s to 7 s substantially increased the number of left-turn vehicles that could be served within each cycle.





Figure 6.2: BSC for two-way traffic, adapted from Figure 2.2

6.2 Outcomes and Implications

The results in this section show that BSC can postpone system failures caused by left-turn pocket spillovers. This benefit arises because the spillovers block the adjacent through lane and functionally shorten the link's usable space for storing the queue. Thus, the link behaves as a de facto short link. As in Chapter 5, BSC is beneficial under this storage constraint because it compacts queues, thus enabling the constrained space to hold more queued vehicles. This retards queue expansion.

Figures 6.3(a) and 6.3(b) plot total trips completed within the 1-hour simulation period against total bi-directional demand. In both panels, the solid curve describes total trips completed under BSC and the dashed curve under zero offsets.



(a) High propensity for spillover.

(b) Moderate propensity for spillover.

Figure 6.3: Trips completed vs. bi-directional demand.

In the high-propensity-for-spillover regime shown in Figure 6.3(a), the two curves are nearly identical up to a bi-directional demand of 2880 veh/h. At higher demands, the curves diverge. Zero offsets first induce system failure at a bi-directional demand of 3600 veh/h, whereas BSC first induces system failure at 4080 veh/h. Over the demand range in which both strategies induce system failure, the failure ratio ranges from $\phi = 1.3$ (at a bi-directional demand of 4080 veh/h) to $\phi = 1.7$ (at a bi-directional demand of 4560 veh/h).

In the moderate-propensity-for-spillover regime shown in Figure 6.3(b), the two curves remain nearly identical up to a demand of 3840 veh/h. At higher demands, the curves again diverge. Zero offsets first induce system failure at a bi-directional demand of 4080 veh/h, whereas BSC first induces system failure at 4560 veh/h. Over the demand range in which both strategies induce system failure, the ratio ranges from $\phi = 1.24$ (at bi-directional demand of 4560 veh/h) to $\phi = 1.6$ (at bi-directional demand of 5280 veh/h). In both spillover regimes, BSC postpones system failure relative to zero offsets.

Figures 6.4–6.6 reveal how turning-pocket spillovers can create de facto short links. The figures show time series density measurements taken over three spatial segments: the turning pocket, the adjacent through-lane segment over the same length as the turn pocket, and the link segment immediately upstream of the pocket (see again Figure 6.1). Figure 6.4 shows densities averaged every minute from a representative simulation run with a high propensity for spillover. The solid curve shows turning-pocket density, and the dashed curve shows density in the adjacent through-lane segment.

Each time the solid curve reaches its maximum, i.e., the pocket's jam density of 200 veh/km, the pocket's storage is exhausted and spillover begins. The subsequent wiggles in the solid curve indicate alternating discharge and refill of the pocket. Pocket density drops when queued left-turn vehicles are discharged, but the pocket does not refill immediately because additional left-turn vehicles cannot enter the pocket while the through-lane queue

blocks the pocket entrance. The pocket refills only when the queue no longer blocks its entrance. This alternating discharge-and-refill process causes the pocket density to oscillate near its maximum.




Figure 6.4: Time series of density at turning pocket and through lane

The dashed curve in Figure 6.4 falls at the same instants that the solid curve reaches its peaks. An example of this is shown for time $T \approx 11$ min. This pattern occurs when left-turn spillovers periodically block access to the adjacent through-lane segment and leave that segment underutilized. For example, at $T \approx 11$ min, note how the adjacent through-lane segment is nearly empty of vehicles. Consequently, part of the link's usable storage becomes unavailable to through traffic. Thus, the link behaves as a de facto short link.

Figure 6.5 plots density in the link segment upstream of the pocket. That short upstream segment continues to store queued vehicles. The time-series curve in the figure exhibits a signature similar to the one seen in the density plot for a physically short link in Chapter 5 (this figure is reproduced here and labeled as figure 6.6). Note that in Figures 6.5 and 6.6, density initially rises with wiggles, indicating that density increases erratically in the early going, and then undergoes a sharp jump. The similar patterns in the two figures reveal

how the link constrained by queue spillovers in Figure 6.5 behaves similarly to the physically short link in Figure 6.6.




Figure 6.5: Time series of density at link section upstream of turning pocket




Figure 6.6: Time series of density of short link from Figure 5.3

Figures 6.4–6.6 therefore explain why BSC outperforms zero offsets. Left-turn pocket spillovers reduce usable storage by creating a de facto short link.

Chapter 7

Findings and Future Research Directions

The research findings are summarized in Section 7.1. Future research directions are discussed in Section 7.2.

7.1 Summary of Findings

This dissertation has investigated the performance of BSC in hyper congested signalized systems. Conventional forward progression quickly loses effectiveness once queues spill over and block upstream intersections. This is because vehicles at the upstream intersection cannot advance when their signals turn green. By contrast, BSC synchronizes green phases with the backward-moving waves in congestion, such that upstream intersections display green signals immediately after downstream queues clear the road by moving forward. This form of coordination reduces wasted green time and helps confine queues within the network's available storage space. The study used closed-loop beltways as test sites, and examined BSC under three conditions that are both common in practice and underexplored in the literature: directional demand imbalances, inhomogeneous block lengths, and protected left-turn maneuvers.

The experiments of Chapter 4 showed that BSC improves beltway performance whenever traffic demands are not directionally balanced. By prioritizing the higher demand direction, BSC increases the vehicle exit rate from the system prior to failure (i.e., gridlock or excessive formation of cross-street queues). The benefit of BSC is strongest when queue storage on the beltway is scarce and directional demands are highly unbalanced, because spillback in the high-demand direction is more damaging in those cases. Only when demand is nearly balanced in both directions does the benefit of BSC fade, as gains in one direction are offset by losses in the other.

Chapter 5 addressed the challenge of inhomogeneous link lengths, where short links with limited queue storage capacity act as weak points in the network. Under zero offsets, queues

that form on short links quickly propagate upstream, triggering gridlock. BSC proved more resilient: by compacting queues within the constrained space, it postponed the propagation of queues through the beltway.

Chapter 6 examined BSC in the presence of protected left-turn maneuvers. The results showed that BSC can postpone system failure when spillovers from left-turn pockets occur. These spillovers block the adjacent through lane and reduce usable queue storage space, causing the link to behave as a de facto short link. Under this storage constraint, BSC is beneficial for the same reason it is favorable on short links: it compacts queues, enables the constrained space to hold more queued vehicles, and shortens queue expansion.

Taken together, these findings show that BSC is a robust and adaptable strategy for hypercongested signalized networks. It improves system performance under directional demand imbalances, mitigates the vulnerabilities of short links and alleviates problems caused by spillbacks from protected left-turn pockets. The results point toward BSC's potential as a congestion management tool for real urban traffic.


7.2 Future Research Directions

The results of this dissertation highlight how BSC can improve traffic operations under conditions that are widespread in real cities. A natural next step is to extend the findings beyond idealized beltways to real-world urban networks. Two research thrusts appear most promising: applying the strategy to beltways in real cities and adapting it to urban street grids.


7.2.1 Applying Findings to Real City Beltways

Closed-loop beltways are fairly common in metropolitan areas around the world. In some cities, such as Tokyo and Beijing, older beltways have been converted into freeways, while in many European and American cities, beltways remain signalized. Examples include one-way loops in Milan and two-way loops in Dortmund and Moscow (see Figures 7.1(a), (b), and (c)). Even in grid-based cities without a formally designed ring, loop-like systems can be identified in the alignment of major arterials, as in Chicago (see Figure 7.1 (d)). These beltways often adapt to local topographic features, bending around rivers or built-up districts. The beltways serve a common function: to allow through-traffic to bypass the central business district, where narrow streets and heavy pedestrian activity constrain vehicle movements.


The simulation results presented in this dissertation suggest that BSC could be particularly effective on real-world beltways like those in Figure 7.1, where demand imbalances, turning maneuvers, and inhomogeneous link lengths all occur routinely. Each of those conditions was shown here to be manageable under BSC. A practical research direction is to calibrate BSC in detailed microsimulation models of existing real-world beltways—such as the Garden Ring in Moscow, or the Loop in Chicago—and to quantify performance gains




(a) Milan, Italy



(b) Dortmund, Germany



(c) Moscow, Russia



(d) Chicago, United States

Figure 7.1: Real city beltways that can be used as testbeds

under real-world traffic patterns. Such studies could bridge the gap between theory and practice and provide a direct pathway for pilot implementations.

7.2.2 Adapting the BSC for Real-world Urban Networks

Another promising direction is to adapt BSC for the grid networks that characterize many large cities. Sadek et al. (2022) propose a network-wide strategy that serves as a reference for future research and can be refined using the findings of this dissertation.

In recognition that signals cannot be synchronized in all travel directions in a grid, the study synchronized signal offsets to serve two orthogonal directions at each intersection: Links that serve inbound directions pointing toward destination-rich (e.g., employment) centers were synchronized in the morning rush and links serving outbound traffic were synchronized in the evening. The signal offsets were determined relative to the reference intersections, which resided near the center of the map for each destination-rich center in the network. By using an identical signal timing plan at all signalized intersections in the grid, it becomes possible to achieve orthogonal synchronization network-wide.

In the default coordination mode, offsets are synchronized to the forward motion of vehicles. When links inside the designated zone become hyper congested (i.e., average traffic density in the zone passes a preset threshold), and the mode changes from forward synchronization to BSC for all signals in the zone. Micro-simulations of numerous scenarios, including those involving distorted grids, show that the strategy is more effective in reducing network-wide vehicle-hours traveled than Synchro, which is commonly considered the gold standard in finding suitable signal coordination plans networkwide.

The present findings motivate future refinements to these ideas pertaining to BSC. Multiple research directions might be worth pursuing. For example, zones where BSC is deployed might be sized and configured in real time, in response to real-time changes in network traffic conditions. Perhaps BSC could even be implemented on a link-by-link basis as needed. Methods in Daganzo (2015) for switching between forward and BSC might be used to this end. Future research to adapt the ideas to irregular (non-grid) networks might also bear fruit, as might studies on networks where signals are sparse. Reinforcement learning and other tools from computer science might have roles to play in this regard.

Finally, evaluating the dissertation's ideas in real-world street grids could be valuable. A first step might entail micro-simulation of real cities or portions thereof. Perhaps it would be beneficial to engage BSC on major streets when they become hyper congested but not on minor ones. Success in these efforts might lead to real-world deployments.

Bibliography

- Abu-Lebdeh, G. and Benekohal, R. F. Development of traffic control and queue management procedures for oversaturated arterials. *Transportation Research Record*, 1603(1):119–127, 1997.
- Chen, S., Xu, H., and Liu, H. Timing oversaturated signals: What can we learn from classic and state-of-the-art signal control models. *Journal of Transportation Systems Engineering and Information Technology*, 13(1):97–110, 2013. doi: 10.1016/S1570-6672(13)60093-8.
- Christofa, E., Argote, J., and Skabardonis, A. Arterial queue spillback detection and signal control based on connected vehicle technology. *Transportation research record*, 2366(1): 61–70, 2013.
- Cohen, S. L. Maxband program for arterial signal timing plans. *Transportation Research Record*, (906):82–90, 1983.
- Cubic Transportation Systems. *Synchro/SimTraffic User’s Guide*. Cubic Transportation Systems, Ottawa, Canada, 2024. Version 12; user manual for Synchro Studio traffic signal timing software.
- Daganzo, C. F. *Fundamentals of Transportation and Traffic Operations*. Emerald, Oxford, UK, 1997.
- Daganzo, C. F. and Lehe, L. J. Traffic flow on signalized streets. *Transportation Research Part B: Methodological*, 90:56–69, 2016.
- Daganzo, C. F., Lehe, L. J., and Argote-Cabanero, J. Adaptive offsets for signalized streets. *Transportation research procedia*, 23:612–623, 2017.
- Federal Highway Administration. Signal timing under saturated conditions. Technical Report FHWA-HOP-09-008, U.S. Department of Transportation, Washington, DC, 2009. URL <https://ops.fhwa.dot.gov/publications/fhwahop09008/>.
- FHWA. *Traffic Flow Theory: Volume 1 - Traffic Flow Fundamentals*. Federal Highway Administration, 2000.

- Gartner, N. H. Opac: A demand-responsive strategy for traffic signal control. *Transportation Research Record*, (906):75–81, 1983.
- Gartner, N. H. and Stamatiadis, C. Progression optimization featuring arterial-and route-based priority signal networks. *Journal of Intelligent Transportation Systems*, 8(2):77–86, 2004.
- Gartner, N. H., Assmann, S. F., Lasaga, F., and Hou, D. L. A multi-band approach to arterial traffic signal optimization. *Transportation Research Part B: Methodological*, 25(1):55–74, February 1991.
- Gordon, R. L., Reiss, R. A., Haenel, H., Case, E., French, R. L., Mohaddes, A., Wolcott, R., et al. Traffic control systems handbook. Technical report, United States. Federal Highway Administration. Office of Technology Applications, 1996.
- Hao, W. and Dong-cheng, Y. Network green-wave band model permitting relaxation of constraints. *China Journal of Highway and Transport*, 33(3):184, 2020.
- Kuyer, L., Whiteson, S., Bakker, B., and Vlassis, N. Multiagent reinforcement learning for urban traffic control using coordination graphs. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '08)*, pages 849–856, Estoril, Portugal, 2008. International Foundation for Autonomous Agents and Multiagent Systems.
- Li, Z., Yu, H., Zhang, G., Dong, S., and Xu, C. Network-wide traffic signal control optimization using a multi-agent deep reinforcement learning. *Transportation Research Part C: Emerging Technologies*, 125:103059, 2021.
- Lieberman, E. B., Chang, J., and Prassas, E. S. Formulation of real-time control policy for oversaturated arterials. *Transportation Research Record*, 1727(1):77–88, 2000. doi: 10.3141/1727-10.
- Lin, H.-J., Machemehl, R. B., Lee, C. E., and Herman, R. Guidelines for use of left-turn lanes and signal phases. Technical Report Report 258-1, Research Project 3-18-80-258, Center for Transportation Research, The University of Texas at Austin, Austin, TX, 1984. In cooperation with Texas State Department of Highways & Public Transportation and FHWA.
- Little, J. D. The synchronization of traffic signals by mixed-integer linear programming. *Operations Research*, 14(4):568–594, 1966.
- Little, J. D., Kelson, M. D., and Gartner, N. H. Maxband: A versatile program for setting signals on arteries and triangular networks. *Transportation Research Record*, 795:40–46, 1981.

- Liu, Y. and Chang, G.-L. An arterial signal optimization model for intersections experiencing queue spillback and lane blockage. *Transportation research part C: emerging technologies*, 19(1):130–144, 2011.
- McTrans Center, University of Florida. *Highway Capacity Software (HCS) User Manual*. McTrans Center, Gainesville, FL, USA, 2023. User guide for HCS software analyzing intersections, roadways and networks.
- Mon, E. E., Ochiai, H., and Aswakul, C. Application of traffic light control in oversaturated urban network using multi-agent deep reinforcement learning. In *IEEE Access*, 2024. IEEE Access Article.
- Newell, G. Blocking effects for synchronized signals. In *Proc., 8th International Symp. on Transportation and Traffic Theory*, University of Toronto Press, Toronto, Canada, 1981.
- Newell, G. F. Theory of highway traffic signals. Technical Report UCB-ITS-RR-89-7, Institute of Transportation Studies, University of California, Berkeley, Berkeley, CA, April 1989. URL <https://escholarship.org/uc/item/7zn2b9bc>.
- Olmos, L. E. et al. Macroscopic dynamics and the collapse of urban traffic. *Proceedings of the National Academy of Sciences*, 115(50):12654–12661, 2018. doi: 10.1073/pnas.1800474115.
- Park, B., Messer, C. J., and Urbanik, T. Enhanced genetic algorithm for signal-timing optimization of oversaturated intersections. *Transportation Research Record*, 1727(1):32–41, 2000.
- Robertson, D. I. Transyt: A traffic network study tool. Laboratory Report LR 253, Road Research Laboratory, Crowthorne, Berkshire, 1969.
- Roess, R. P., Prassas, E. S., and McShane, W. R. *Traffic Engineering*. Pearson, Upper Saddle River, NJ, 4th edition, 2011.
- Roess, R. P., Prassas, E. S., and McShane, W. R. *Traffic Engineering*. Pearson Education, 5 edition, 2019.
- Sadek, B., Doig Godier, J., Cassidy, M. J., and Daganzo, C. F. Traffic signal plans to decongest street grids. *Transportation Research Part B: Methodological*, 162:195–208, 2022.
- Schreiber, L. V., de O. Ramos, G., and Bazzan, A. L. C. Towards explainable deep reinforcement learning for traffic signal control. *Proceedings of the International Conference on Machine Learning (ICML) Workshop on RL for Real-World Systems*, 2021.
- Smith, M. J. The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological*, 13(4):295–304, 1979.

- Texas Transportation Institute. *PASSER II-90 User's Guide*. Texas AM Transportation Institute, College Station, TX, USA, 1991. Program user's guide for arterial traffic signal progression software.
- Tsay, H. S. and Lin, L. T. New algorithm for solving the maximum progression bandwidth. *Transportation Research Record*, (1194):15–30, 1988.
- van der Pol, E. and Oliehoek, F. A. Coordinated deep reinforcement learners for traffic light control. In *Proceedings of the NIPS 2016 Workshop on Learning, Inference and Control of Multi-Agent Systems (LFM 2016)*, pages 1–8, Barcelona, Spain, 2016.
- Wang, P., Jones, L. S., Yang, Q., and Gurupackiam, S. Cell transmission model based traffic signal timing in oversaturated conditions. *Journal of Central South University*, 20(4): 1129–1136, 2013. doi: 10.1007/s11771-013-1594-8.
- Webster, F. V. Traffic signal settings. Road Research Technical Paper 39, Road Research Laboratory, London, 1958.
- Yang, X., Cheng, Y., and Chang, G.-L. A multi-path progression model for synchronization of arterial traffic signals. *Transportation research part C: emerging technologies*, 53:93–111, 2015.