Traffic Assignment for Urban Evacuation

Seyed Hassanali Emamzadeh

Thesis submitted to the University of Ottawa in partial fulfillment of the requirements for the degree of Master of Applied Science in Mechanical Engineering

Department of Mechanical Engineering
Faculty of Engineering
University of Ottawa



uOttawa

© Seyed Hassanali Emamzadeh, Ottawa, Canada, 2026

Abstract

Urban emergency evacuations impose sudden and extreme demands on road networks that differ fundamentally from normal operating conditions. When a large population must leave a defined area within a short time window, traffic demand surges abruptly, evacuation corridors are stressed toward their capacity, and the consequences of poor traffic management extend beyond inconvenience to a direct threat to public safety. Conventional traffic planning tools, which assume steady state demand, time invariant network conditions, and independently optimised intersections, are inadequate in this context.

This thesis develops and evaluates a two-stage procedure for traffic assignment and signal control under emergency evacuation conditions, in which a static assignment calibrates the signal plans and a microsimulation evaluates them. Macroscopic traffic flow theory, static and dynamic assignment methodologies, and signal timing design are systematically addressed, culminating in a microsimulation case study of the Sandy Hill neighbourhood in Ottawa, modelled in Aimsun Next. Evacuation demand is derived from 2021 Census data, with a documented extrapolation from the 2011 Census counts, through a chain of population, household, vehicle ownership, and compliance parameters, yielding 3,530 evacuating vehicles loaded over a two hour horizon by a Rayleigh departure time distribution. The network comprises 17 signalized intersections organised around a two branch evacuation corridor serving the two principal egress destinations, Highway 417 and the Macdonald–Cartier Bridge. Two fixed time signal control strategies are designed from the assigned peak interval flows and evaluated: a Webster baseline plan with pedestrian governed minimum greens, and a Coordinated Green Wave plan with a common corridor cycle and progressive offsets calibrated to the loaded corridor platoon speed of 25 km/h.

The comparison shows that the value of coordination depends on how heavily the corridor is loaded. While the network is lightly loaded the two plans are indistinguishable; through the peak of the Rayleigh loading profile the coordinated plan reduces average travel time per vehicle by 21.4 %, and as demand falls the two converge again. Over the whole run, across ten replications per plan with common random seeds, coordination reduces total system travel time by 9.5 %, stopped time by 15.8 % and mean queue by 33.5 %, every measure significant at the 0.1 % level.

The experiment is then repeated at three higher demand levels to test whether the benefit grows with load. It does, up to a point. At 1.25 times the base demand the reduction in total travel time is 34.9 %, at 1.5 times 63.9 %, and at twice the demand it falls back to 37.9 %. The reason is visible in the two plans separately: under coordination the travel time per vehicle is 96, 95 and 103 s across the first three levels, essentially unchanged while the demand rises by half, whereas under isolated timing it is 107, 145 and 285 s. At twice the demand both plans are beyond the corridor's capacity and the gap narrows because both are failing. At 1.5 times the difference is qualitative rather than quantitative: every coordinated run cleared the network within the horizon and no isolated run did.

Coordination also makes the outcome between six and twenty times more repeatable, and the margin widens with load. The saving arises not from added capacity but from the alignment of green times between junctions, and it is concentrated on the southbound route

to Highway 417, where isolated timing is liable to break down altogether. The results indicate that an evacuation signal plan should be judged by its performance at the crest of the demand surge, by how reliably it performs there, and by whether the network clears, rather than by its average delay over the evacuation period, which a steady-state assessment cannot provide.

Acknowledgements

First and foremost, I would like to express my deepest and most sincere gratitude to my supervisor, Dr. Mohammad Pirani. Without his unwavering support, continuous guidance, and immense patience, this work would simply not have been possible. I am profoundly grateful for his heartfelt encouragement, invaluable mentorship, and the scholarly insights he so generously shared with me throughout this journey. His dedication to teaching and research has been a constant source of inspiration.

I also owe a great debt of gratitude to my wonderful family. To my beloved parents and my dear sister, thank you for being my ultimate pillar of strength. Your unconditional love, endless sacrifices, and unwavering belief in my abilities have sustained me through every challenge. Words cannot fully express how grateful I am to have you by my side at every moment of my life; this achievement belongs to you as much as it does to me.

Contents

Abstract ii

Acknowledgements iv

List of Figures viii

List of Tables ix

1 Introduction 1
- 1.1 Background on Traffic Systems . . . 1
  - 1.1.1 Modelling Scales: Microscopic, Mesoscopic, and Macroscopic . . . 1
  - 1.1.2 Eulerian and Lagrangian Perspectives . . . 4
  - 1.1.3 Fundamental Diagram of Traffic Flow . . . 5
  - 1.1.4 Traffic Congestion . . . 8
- 1.2 Evacuation Traffic Management and Signal Control . . . 10
  - 1.2.1 Contraflow and Lane Reversal . . . 10
  - 1.2.2 Lane-Based and Conflict-Free Routing . . . 11
  - 1.2.3 Staged and Phased Evacuation . . . 11
  - 1.2.4 Intersection and Signal Control in Evacuations . . . 11
  - 1.2.5 Position of This Thesis . . . 12
- 1.3 Thesis Contributions and Organization . . . 12
  - 1.3.1 Research Objectives . . . 13
  - 1.3.2 Main Contributions . . . 13
  - 1.3.3 Thesis Organization . . . 14

2 Static Traffic Assignment 16
- 2.1 User Equilibrium (UE) . . . 16
  - 2.1.1 UE Example: Two Parallel Routes . . . 17
- 2.2 System Optimum (SO) . . . 18
- 2.3 Braess's Paradox in Traffic Networks . . . 20
- 2.4 Static Behavioural Principles in Evacuation . . . 22
  - 2.4.1 Behavioural and Practical Implications . . . 24

3 Dynamic Traffic Assignment 25
- 3.1 Instantaneous and Experienced Travel Times . . . 25
  - 3.1.1 The Physical Reality of Congestion: Spillback . . . 26

3.2 Limitations of Static Modelling . . . 27
3.2.1 The FIFO Limitation and Overtaking Behaviour . . . 28
3.2.2 The DTA Synthesis: Mesoscopic and Iterative Modelling . . . 30
3.3 Core Modelling Principles and Notation . . . 30
3.3.1 State Variables and Definitions . . . 30
3.3.2 Conservation of Flow and Admission Cap . . . 31
3.4 Dynamic User Equilibrium and Dynamic System Optimum . . . 33
3.4.1 Dynamic User Equilibrium (DUE) . . . 33
3.4.2 Dynamic System Optimum (DSO) . . . 33
3.5 Dynamic Assignment in a One-Shot Simulation . . . 34
3.6 Deterministic, Non Deterministic and Unobservable Elements . . . 36

4 Traffic Signal Control **38**
4.1 Terms Related to Signal Timing . . . 39
4.2 Traffic Flow Representation for Signal Analysis . . . 40
4.3 Yellow Interval and the Dilemma Zone . . . 41
4.3.1 Go Requirement . . . 42
4.3.2 The Dilemma Zone . . . 43
4.4 Cycle Length Determination: The Webster Method . . . 45
4.4.1 Interpretation of the Critical Flow Sum . . . 46
4.4.2 Saturation Flow . . . 46
4.5 Lost Time at Signalized Intersections . . . 46
4.5.1 Lost Time for a Phase . . . 47
4.5.2 Total Lost Time per Cycle . . . 47
4.5.3 Allocation of Effective Green Time . . . 48
4.5.4 Pretimed and Actuated Signal Control . . . 49
4.5.5 Signal Timing Parameters for Actuated Control . . . 49
4.6 Signal Coordination in Emergency Evacuation . . . 50
4.6.1 Coordination in Oversaturated Networks . . . 51
4.7 Cost-Based Evaluation of Signal Coordination Benefits . . . 52

5 Emergency Evacuation Simulation **54**
5.1 Origin–Destination Demand . . . 56
5.2 Static User Equilibrium Assignment . . . 63
5.3 Signal Timing Methodology . . . 64
5.3.1 Webster Baseline Plan . . . 65
5.3.2 Coordinated Green Wave Plan . . . 67
5.4 Results and Discussion . . . 69
5.4.1 Model inputs and reproducibility . . . 69
5.4.2 Definition of the performance measures . . . 71
5.4.3 Experimental protocol . . . 72
5.4.4 Network performance . . . 73
5.4.5 Performance over time . . . 74
5.4.6 Network Speed and Flow . . . 79
5.4.7 Travel time by destination . . . 80

5.4.8 Sensitivity to the Demand Level 81
5.4.9 Loss of Capacity During the Evacuation 85
5.4.10 Sensitivity to the Aggregation Interval 87
5.4.11 Where the delay is saved 88
5.4.12 Limitations of the comparison 89
5.5 Summary 90

6 Conclusions **92**
6.1 What the Thesis Establishes 92
6.2 Implications for Evacuation Traffic Engineering 94
6.3 Recommendations for Future Work 95
6.4 Conclusion 96

List of Figures

1.1 Microscopic and macroscopic views of traffic flow 3
1.2 Eulerian and Lagrangian viewpoints 5
1.3 Speed as the slope of the flow-density curve 6
1.4 The speed-density relationship 7
1.5 The speed-flow relationship 8
1.6 The three regimes of the fundamental diagram 9
1.7 Research framework 14

2.1 User equilibrium on two parallel routes 18
2.2 The Braess network 21
2.3 Network schematic 23

3.1 Instantaneous and experienced travel time compared 26
3.2 Spillback at a network bottleneck 27
3.3 Time discretization and flow evolution 31
3.4 One-shot dynamic assignment with incremental route updating 35

4.1 Signal timing for a two-phase cycle 39
4.2 The stopping requirement 42
4.3 The go requirement 43
4.4 The dilemma zone 44
4.5 Discharge rate during one signal phase 48

5.1 The Sandy Hill study network 55
5.2 The Rayleigh departure profile 62
5.3 Travel time by interval, base demand 76
5.4 Travel time by interval, 1.25 times demand 77
5.5 Travel time by interval, 1.5 times demand 78
5.6 Travel time by interval, twice the demand 79
5.7 Benefit of coordination against demand 82
5.8 Cumulative vehicles evacuated 84
5.9 Effect of a lane closure 86

List of Tables

3.1 Static and dynamic assignment compared 29

5.1 Signalised junctions of the Sandy Hill network 56
5.2 Vehicle ownership by household type 58
5.3 The demand derivation chain 60
5.4 Temporal distribution of evacuation demand 62
5.5 Volume to capacity ratios on the critical sections 64
5.6 Scenarios and experiments 71
5.7 Network performance over the whole run 74
5.8 Average travel time per vehicle by interval 75
5.9 Corridor speed and flow by interval 80
5.10 Travel time by destination 81
5.11 Performance at four demand levels 83
5.12 Proportion evacuated within the horizon 85
5.13 Effect of a lane closure on each signal plan 86
5.14 Effect of the aggregation interval 88
5.15 Where the delay is saved 89

Chapter 1

Introduction

Urban road networks that operate acceptably under everyday demand can fail abruptly when a large population must evacuate a defined area within a short time window. Analyzing and mitigating such failures requires a precise description of how traffic flow behaves as demand approaches and exceeds capacity. This chapter therefore proceeds in two parts. Section 1.1 introduces the mathematical foundations of traffic flow modelling, the modelling scales, observational perspectives, the fundamental diagram, and the congestion and shockwave phenomena that govern oversaturated networks, which are used throughout the thesis. Section 1.3 then presents the motivation, research objectives, and contributions of the thesis, and outlines the organization of the remaining chapters.

1.1 Background on Traffic Systems

To understand the complex behaviour of traffic, it is first necessary to introduce the mathematical frameworks used in traffic engineering. These frameworks describe movement dynamics ranging from local interactions between individual vehicles to aggregate flows at the network level. The choice of modelling approach depends on the level of detail required and the type of data available. This section examines traffic modelling from two complementary viewpoints: the scale of representation (microscopic, mesoscopic, and macroscopic) and the observational frame of reference (Eulerian and Lagrangian), and then introduces the fundamental diagram of traffic flow and the congestion phenomena that follow from it.

1.1.1 Modelling Scales: Microscopic, Mesoscopic, and Macroscopic

The fundamental distinction among traffic modelling scales lies in how the discrete nature of vehicles is treated. In *microscopic* modelling, the traffic system is viewed as a collection

of distinct particles (vehicles), each possessing a unique identity and behaviour. In this view, the behaviour of each driver is a direct function of the immediate surroundings: the instantaneous acceleration of a vehicle is a reaction to the spacing and relative speed with respect to the vehicle ahead. The mathematical formulation of these *car-following* models is typically based on a system of ordinary differential equations, one for each vehicle. Denoting the position of the i-th vehicle at time t by $x_i(t)$, with vehicle $i-1$ as its leader, a general car-following law takes the form

$$\ddot{x}_i(t) \;=\; f\big(s_i(t),\, \dot{x}_i(t),\, \Delta v_i(t)\big), \tag{1.1}$$

where $s_i(t) = x_{i-1}(t) - x_i(t) - \ell_{i-1}$ is the net spacing to the leading vehicle (with ℓ_{i-1} the leader's length), $\dot{x}_i(t)$ is the speed of vehicle i, and $\Delta v_i(t) = \dot{x}_{i-1}(t) - \dot{x}_i(t)$ is the relative speed between the leader and the follower. Classical instances of Eq. (1.1) include the Gipps model [7] and the Intelligent Driver Model [15].

This level of detail allows the simulation of phenomena such as lane changing, sudden braking, and driver heterogeneity (e.g., aggressive versus cautious drivers). However, the computational cost of microscopic models increases sharply with the number of vehicles, which often renders them impractical for the analysis of extensive networks.

In contrast, *macroscopic* modelling disregards the particle nature of vehicles and treats traffic analogously to a compressible fluid. Rather than tracking the positions of individual vehicles, the analysis concerns aggregate field variables defined over time and space. In this perspective, traffic is viewed as a continuum in which vehicle concentration varies continuously along the roadway. The primary advantages of the macroscopic approach are its computational efficiency and its ability to describe collective phenomena, such as the formation and propagation of shock waves, that unfold over long distances and extended periods. The two scales are contrasted schematically in Figure 1.1.

In the macroscopic setting, the state of traffic at location x and time t is described by the following field variables:

- **Density** $\rho(x,t)$: the number of vehicles per unit length of roadway (e.g., veh/km). Density expresses how crowded the traffic stream is at a given point in space and time.

- **Space-mean speed** $v_s(x,t)$: the average speed of the vehicles present on a road segment. The space mean is used in macroscopic traffic flow models because slower vehicles occupy the road for longer and therefore influence traffic conditions more strongly than faster vehicles. For a sample of n vehicles with individual speeds $v_1, \ldots, v_n$ observed on the segment during the measurement interval, the space-mean speed is the harmonic

mean

$$v_s \;=\; \frac{n}{\displaystyle\sum_{i=1}^{n} \frac{1}{v_i}}. \tag{1.2}$$

- **Flow** $q(x,t)$: the number of vehicles passing a fixed point per unit time (e.g., veh/h). Flow measures throughput and is related to density and space-mean speed through the fundamental identity given below.

These variables are linked by the fundamental relation

$$q(x,t) \;=\; \rho(x,t)\, v_s(x,t). \tag{1.3}$$

A closure relation based on the fundamental diagram can then be introduced as $q = F(\rho)$, with $F(0) = 0$. A detailed discussion of this closure and its modelling implications is given in Section 1.1.3.

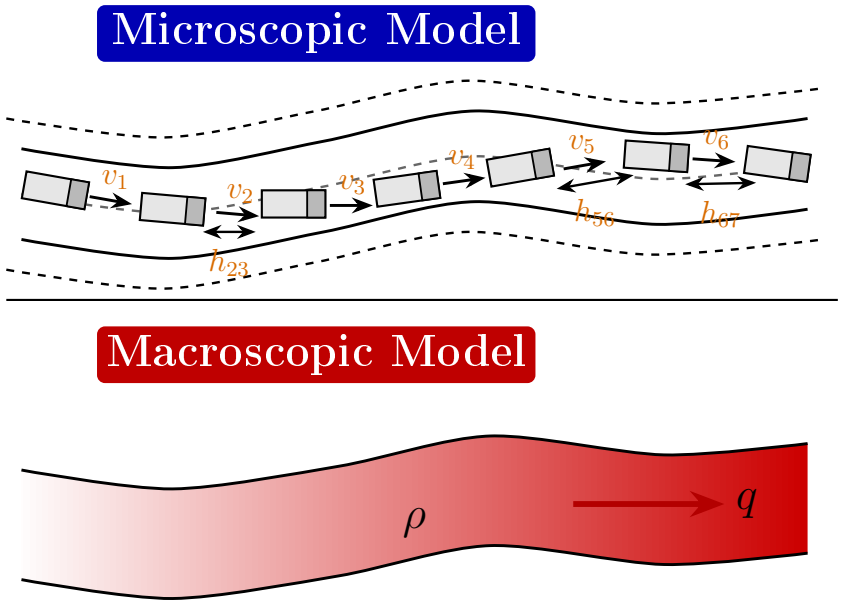


Figure 1.1: Microscopic and macroscopic views of the same traffic. Vehicles are tracked individually above, with speeds v_i and headways h_{ij}; below, the same traffic is a continuum of density ρ and flow q.

Mesoscopic Modelling: Bridging Individual Behaviour and Aggregate Flow

Mesoscopic traffic modelling occupies an intermediate position between the microscopic and macroscopic paradigms and is widely adopted in network level dynamic analyses [10]. In contrast to microscopic models, which explicitly simulate detailed vehicle-to-vehicle interactions,

and macroscopic models, which represent traffic as a fully continuous stream, mesoscopic formulations retain a discrete vehicle (or vehicle packet) representation while approximating interactions through aggregate behavioural and flow relationships. This hybrid structure enables the model to capture key congestion mechanisms, including queue formation, spill-back, and shock-wave propagation, at a substantially lower computational cost than fully microscopic simulation.

By discretizing time and space and tracking vehicle movements across network links, mesoscopic models provide the temporal sensitivity required for Dynamic Traffic Assignment (DTA), where the experienced travel time depends on the departure time and on time varying network states. Consequently, mesoscopic modelling offers a practical balance between behavioural fidelity and computational tractability: it is scalable to large networks and extended simulation horizons while preserving sufficient resolution to support route choice dynamics, departure time response, and equilibrium oriented dynamic analysis.

1.1.2 Eulerian and Lagrangian Perspectives

In addition to the scale of modelling, the manner in which traffic variables are observed and measured creates an important physical and mathematical distinction rooted in fluid mechanics. These two perspectives determine the types of sensors, data, and control strategies that can be employed; they are contrasted in Figure 1.2.

Eulerian Perspective

Named after Leonhard Euler, this perspective is anchored to *location*. The observer occupies a fixed point in space (such as the roadside) and records the traffic stream passing that point; in effect, this is a fixed observation window through which the number of passing vehicles is counted per unit time. Most traditional traffic control systems are built upon this perspective: inductive loop detectors embedded in the pavement and fixed surveillance cameras are prime examples of Eulerian instruments. The principal limitation of this view is that conditions are known only at the sensor locations; events occurring between sensors must be inferred through mathematical models.

Mathematically, the Eulerian description treats traffic variables as functions of space and time, and their local evolution is expressed through the partial time derivative at a fixed location,

$$\dot{\rho}_E(x,t) \triangleq \left.\frac{\partial \rho(x,t)}{\partial t}\right|_x \tag{1.4}$$

which represents the rate of change of traffic density at a fixed location in space.

Lagrangian Perspective

The Lagrangian perspective, attributed to Joseph-Louis Lagrange, is anchored to the *particle*. The observer is attached to a specific vehicle and moves with the flow: instead of asking how fast traffic moves at a given location, one asks how the speed of a particular vehicle evolves over time. With the emergence of technologies such as GPS and connected vehicles, this perspective has gained increasing importance. Floating Car Data (FCD), extracted from smartphones or in-vehicle navigation systems, is Lagrangian in nature. This approach permits direct measurement of the actual travel experience (such as delays accumulated along a route); its main challenge is that system-wide control requires a high penetration rate of connected vehicles and a paradigm shift from infrastructure-based to vehicle-based control.

In mathematical terms, the Lagrangian description follows a moving vehicle, and the change of a traffic variable along its trajectory is expressed through the material derivative

$$\frac{D\rho}{Dt} = \frac{\partial \rho(x,t)}{\partial t} + v_s(x,t)\,\frac{\partial \rho(x,t)}{\partial x}, \tag{1.5}$$

which combines the temporal and spatial variations experienced by an observer moving with the traffic stream.

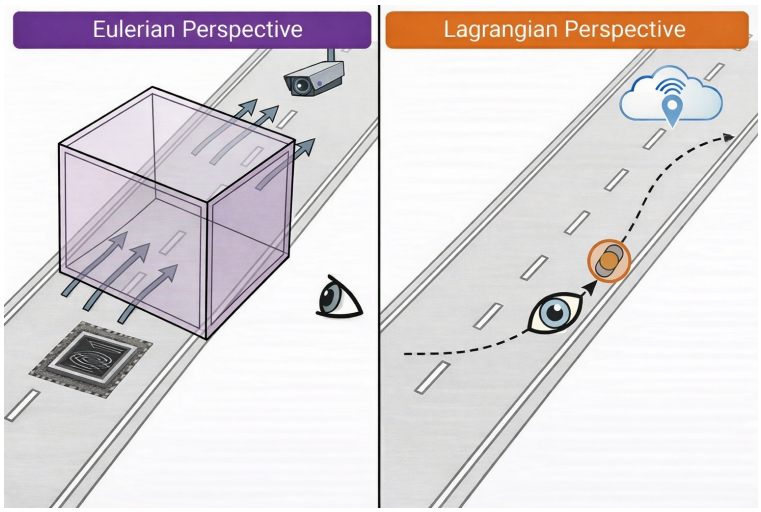


Figure 1.2: Comparison of Eulerian and Lagrangian viewpoints in traffic modelling. Figure prepared by the author.

1.1.3 Fundamental Diagram of Traffic Flow

The fundamental diagram of traffic flow describes the relationship among the three primary macroscopic traffic variables: density ρ, flow q, and space-mean speed v_s, which are linked

by the identity in Eq. (1.3). The fundamental diagram specifies how speed or flow varies as a function of density under different traffic conditions [8].

A central implication of the fundamental diagram is that traffic flow does not increase monotonically with density. Instead, flow initially increases as density rises, reaches a maximum at a critical density, and then decreases as congestion intensifies. This nonlinearity captures the transition from the free-flow to the congested traffic regime and forms the basis of macroscopic traffic flow theory.

Flow-Density Relationship

The flow-density relationship, often regarded as the fundamental diagram itself, is obtained by combining a speed-density function with the macroscopic identity in Eq. (1.3); flow can thereby be expressed as a function of density alone.

This relationship exhibits an inverted-U shape, as shown in Figure 1.3. Flow increases with density in the free-flow regime, attains its maximum value the capacity $q_{\max}$ at the critical density ρ_c, and then decreases in the congested regime as speeds decline. The decreasing branch of the curve reflects unstable traffic conditions and the onset of congestion.

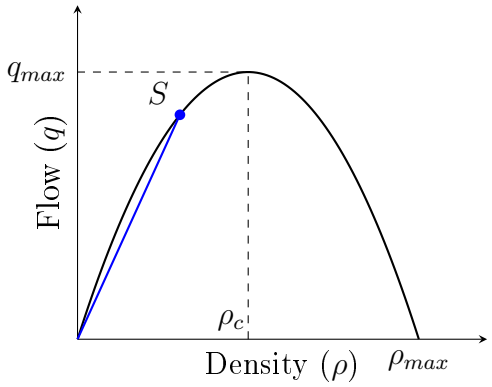


Figure 1.3: When traffic is at a specific state S on the q-ρ diagram, the corresponding average (space-mean) speed is given by the slope of the line from the origin to S; equivalently, $v_S = q_S/\rho_S$.

Speed-Density Relationship

The speed-density relationship describes how the space-mean speed decreases as density increases. At very low densities, vehicles travel near the free-flow speed, while at higher densities increased vehicle interactions reduce speeds. In general form,

$$v_s \;=\; V(\rho), \tag{1.6}$$

where $V(\rho)$ is a decreasing function of ρ. A classical linear model is the Greenshields relation [8],

$$v_s(\rho) \;=\; v_f \left(1 - \frac{\rho}{\rho_{\max}}\right), \tag{1.7}$$

where v_f is the free-flow speed and $\rho_{\max}$ is the jam density. This relationship is illustrated in Figure 1.4.

Substituting Eq. (1.7) into the macroscopic identity $q(\rho) = \rho\, v_s(\rho)$ yields

$$q(\rho) \;=\; \rho\, v_f \left(1 - \frac{\rho}{\rho_{\max}}\right) \;=\; v_f\, \rho - \frac{v_f}{\rho_{\max}}\, \rho^2, \tag{1.8}$$

which shows explicitly that, under the Greenshields model, flow is a quadratic function of density.

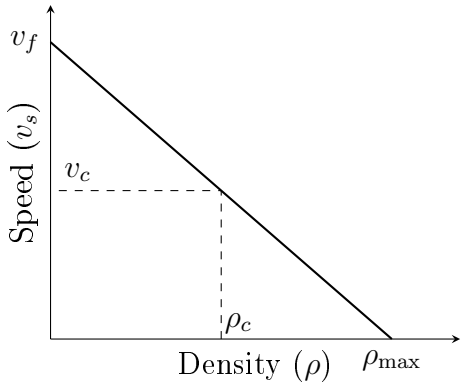


Figure 1.4: Greenshields linear speed–density relationship. The speed at the critical density ρ_c is denoted v_c and corresponds to operation at capacity.

Speed-Flow Relationship

The speed-flow relationship provides an alternative representation of the fundamental diagram by expressing speed as a function of flow, as shown in Figure 1.5. Although it is less commonly used for modelling purposes, it offers useful insight into operational conditions on a roadway: high speeds are associated with low flows under free-flow conditions, whereas lower speeds correspond to flows near capacity and to congested operation. Note that, except at capacity, each flow value corresponds to two distinct speeds one on the free-flow branch and one on the congested branch so the relationship is not a single valued function. This representation is particularly useful for interpreting field measurements obtained from traffic sensors.

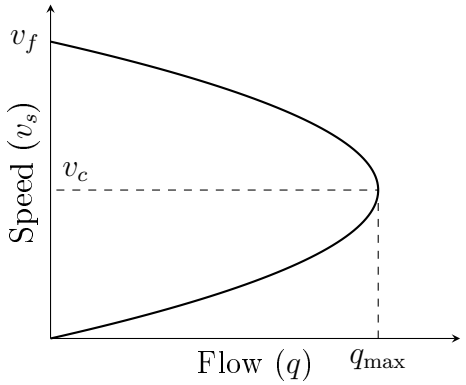


Figure 1.5: Speed-flow relationship derived from the Greenshields model. Each flow level below capacity $q_{\max}$ corresponds to two operating speeds: one in the free-flow regime and one in the congested regime.

Volume-to-Capacity Ratio

The volume-to-capacity ratio (V/C), also known as the degree of saturation, is a widely used indicator of congestion severity. It is defined as the ratio of traffic *demand* to capacity,

$$\text{V/C} = \frac{q_{\text{dem}}}{q_{\max}}, \tag{1.9}$$

where q_{dem} denotes the demand arrival rate over the analysis interval (the volume of traffic seeking to use the facility), and $q_{\max}$ denotes the capacity of the facility. The distinction between demand and measured throughput is essential here: because throughput is physically bounded by capacity, a ratio formed from *measured* flow can never exceed unity. Oversaturation is therefore expressed in terms of demand. Values of V/C well below one indicate uncongested operation; values approaching one indicate near-saturated conditions with diminishing reserve capacity; and $\text{V/C} > 1$ indicates an oversaturated state in which demand exceeds capacity and queues grow over time. The same definition applies at the level of an individual intersection approach, where the degree of saturation of a movement is the ratio of its arrival demand to the capacity provided by its signal timing; this approach-level form is used in the signal analysis later in this thesis.

1.1.4 Traffic Congestion

Traffic congestion is a macroscopic phenomenon that emerges when traffic demand approaches or exceeds the effective capacity of a roadway, that is, when the volume-to-capacity ratio of Eq. (1.9) approaches or exceeds unity. From the user's perspective, congestion is experienced primarily as a rapid increase in travel time in response to relatively small increases in traffic

volume. From a flow perspective, congestion is associated with high densities and unstable operating conditions.

This transition between operating regimes is consistent with the fundamental diagram shown in Figure 1.6, which relates traffic flow to density. At low densities, flow increases approximately linearly with density, corresponding to free-flow conditions. Beyond the critical density, however, flow reaches its maximum and subsequently decreases as density continues to increase. This regime, often referred to as supercritical or hyper-congested traffic, is characterised by reduced speeds, unstable flow, and the formation of congestion patterns such as queues and shock waves.

Together, these observations highlight that congestion is not merely a function of high demand, but rather a structural property of traffic systems operating near or beyond capacity, where small disturbances can trigger large degradations in performance.

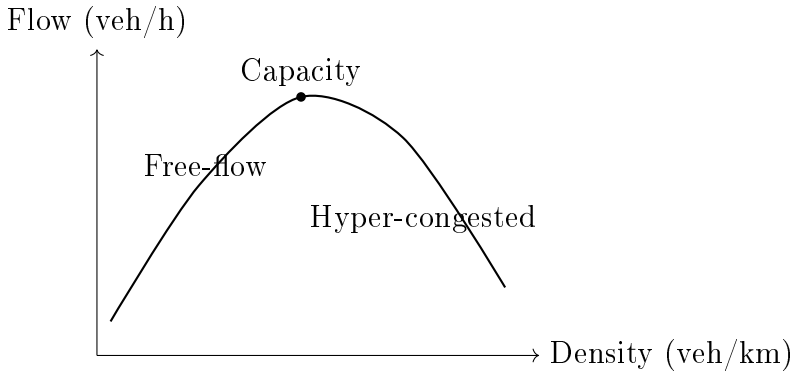


Figure 1.6: Fundamental diagram illustrating the free-flow and hyper congested traffic regimes, separated by the capacity point at the critical density.

Shock Waves in Traffic Flow

In traffic flow theory, a shock wave is the moving interface between two adjacent traffic states that differ in density and flow; it provides a compact description of how abrupt transitions, such as the onset or relaxation of congestion, propagate through space and time [9, 12]. On the flow–density plane, if the upstream and downstream states are represented by $P_1 = (\rho_1, q_1)$ and $P_2 = (\rho_2, q_2)$, respectively, the propagation speed of the interface is given by

$$w \;=\; \frac{q_2 - q_1}{\rho_2 - \rho_1}, \tag{1.10}$$

which is the Rankine–Hugoniot jump condition of the kinematic wave (Lighthill, Whitham, Richards) theory applied to traffic. Geometrically, w is the slope of the chord connecting P_1

and P_2 on the fundamental diagram. This expression links observable changes in flow and density to the motion of the congestion boundary, making shock-wave analysis a practical tool for interpreting and quantifying queue evolution in capacity constrained networks.

Types of Shock Waves

- **Backward propagating shock ($w < 0$):** the discontinuity moves upstream. This is the classic congestion formation front. It typically occurs when downstream supply drops (a capacity reduction) or when inflow demand exceeds the downstream discharge rate, creating a congested region that grows in the upstream direction.
- **Forward-propagating shock ($w > 0$):** the discontinuity moves downstream. This commonly represents a clearing (dissipation) front that arises when conditions improve downstream, for example, when a signal turns green, an incident clears, or demand decreases, so that the boundary between congested and free-flow traffic advances downstream.
- **Stationary shock ($w = 0$):** the discontinuity remains at a fixed location. This occurs when $q_1 = q_2$ while $\rho_1 \neq \rho_2$, that is, equal flows but different densities, a configuration often observed near a fixed bottleneck under steady conditions.

1.2 Evacuation Traffic Management and Signal Control

The traffic engineering response to a mass evacuation has been studied from two directions: the practice of evacuation operations, documented largely from hurricane evacuations in the United States, and the modelling and optimisation of evacuation networks. Broad overviews are given by Murray-Tuite and Wolshon [21], who review evacuation transportation modelling across research and practice, by Bayram [19], who surveys optimisation models for large-scale network evacuation planning, and by Pel et al. [20], who review how travel behaviour is represented in dynamic evacuation simulation models. This section summarises the four families of traffic management strategy most relevant to this thesis and then positions the thesis against them.

1.2.1 Contraflow and Lane Reversal

Contraflow reverses one or more lanes of the inbound carriageway so that they carry outbound traffic, increasing directional capacity away from the threat. Its use expanded considerably after the severe congestion of the 1999 Hurricane Floyd evacuation, and its planning

and operation are described by Wolshon [44, 23, 35]. Theodoulou and Wolshon [24] used microsimulation to evaluate alternative ways of loading traffic into the contraflow lanes, showing that the design of the entry points affects how effectively the added capacity is used. At the network level, Tuydes and Ziliaskopoulos [45] optimized the choice of lanes to reverse with a tabu-search heuristic, and Xie, Lin and Waller [46] formulated a bi-level model in which lane reversal and crossing elimination are chosen in the upper level while a cell-transmission-based dynamic traffic assignment represents the traffic response in the lower level.

1.2.2 Lane-Based and Conflict-Free Routing

Intersections are the principal source of delay in urban evacuation networks. Cova and Johnson [43] proposed lane-based evacuation routing, an integer extension of the minimum-cost flow problem that assigns routes so as to eliminate crossing conflicts at intersections and limit merging, accepting somewhat longer travel distances in exchange for uninterrupted flow. Xie and Turnquist [47] extended lane-based network optimisation with an integrated Lagrangian relaxation and tabu-search method, and Pillac et al. [18] developed a path-generation heuristic for evacuation planning that accounts for conflicts between evacuation paths. These approaches remove conflicts by *redesigning* the network rather than by retiming its signals.

1.2.3 Staged and Phased Evacuation

Rather than adding capacity, staged evacuation manages demand, releasing zones at different times so that the peak load on shared corridors is reduced. Liu, Lai and Chang [48] proposed a cell-based network optimisation model for staged evacuation planning, and Zhang, Spansel and Wolshon [26] examined the effect of phased evacuations on megaregion highway networks. Because staging acts on the timing of departures, it is closely related to the departure-time profile that this thesis models with the Rayleigh distribution (Section 5.1).

1.2.4 Intersection and Signal Control in Evacuations

In practice, evacuation traffic at critical intersections is frequently controlled manually by police at traffic control points, or by signal plans modified to favour the outbound movements; the state of practice is reviewed by Wolshon et al. [49]. The research literature on signal control is comparatively small. Chen, Chen and Miller-Hooks [50] observed that most evacuation studies had focused on freeways and that signal timing had received little systematic attention in urban evacuation planning; using a simulation model of Washington, D.C.,

they evaluated signal timing strategies for a no-notice evacuation and the trade-off between evacuation time and average delay. Parr and Kaisar [51] optimized the signal timing of a single critical intersection during an urban evacuation by dynamic programming, and Parr, Wolshon and Murray-Tuite [52] compared unconventional intersection control strategies to identify which is best suited to a given urban evacuation scenario. At the network level, Nassir et al. [53] developed a two-stage network-flow method that jointly determines evacuation routing and signal timing for real-sized urban networks. Simulation-based studies of urban evacuation have also tested traffic management policies more broadly, including the pipeline of Chen et al. [22], multimodal evacuation [25], and the evacuation of persons with mobility needs [1].

1.2.5 Position of This Thesis

Relative to this literature, the thesis makes a narrower and more applied contribution. It does not optimise routing and signal timing jointly, as the network-flow models do [53], nor does it redesign the network through contraflow or conflict-free routing [43, 46]. Instead, it asks whether a signal coordination strategy that is standard in everyday arterial operation, the fixed-time green wave, is beneficial when the demand is an evacuation surge, and it compares that strategy with the isolated Webster timing that would otherwise be in place, on a real 17-signal urban network under identical, census-derived demand. Its distinguishing feature is the temporal resolution of the comparison: the two plans are evaluated interval by interval through loading and clearance, which shows *when* during the evacuation coordination helps or hurts, a question that aggregate or static evaluations such as the average-delay trade-offs reported in [50] do not address. Contraflow and staged departures are not modelled here; their combination with signal coordination is identified as future work in Section 6.3.

1.3 Thesis Contributions and Organization

A central challenge in emergency evacuation planning is the gap between how traffic is *assigned* to a network and how that network is actually *controlled*. Traffic assignment models distribute vehicles across routes on the basis of travel time and capacity, while signal control governs how that capacity is allocated at each intersection; yet in the evacuation context these two components are rarely designed together. Static assignment ignores the time-varying buildup and clearance of queues, and signal timings calibrated for normal conditions are not tuned to serve the large directional platoons that move toward evacuation exits. This thesis fills that gap by developing and evaluating a procedure that links assignment informed

bottleneck diagnosis to corridor level signal coordination, and by quantifying, through microsimulation of a real Ottawa evacuation network, the operational benefit that coordination delivers, most notably in the network's clearance phase, a benefit that static assignment alone cannot reveal.

1.3.1 Research Objectives

The primary objective of this thesis is to develop and evaluate an integrated framework for traffic assignment and signal control under emergency evacuation conditions. Specifically, the thesis pursues the following objectives:

1. To analyse and compare static and dynamic traffic assignment methodologies, specifically the User Equilibrium and System Optimum formulations, in the context of emergency evacuation, identifying their respective assumptions, limitations, and behavioural implications.

2. To develop and apply signal timing methodologies, including the Webster method for cycle length optimisation and a coordinated green-wave strategy, to a real urban network under simulated evacuation demand.

3. To quantitatively evaluate the operational performance of the two signal control strategies through microsimulation, using metrics that include network travel time, corridor speed, throughput, and destination specific travel times.

1.3.2 Main Contributions

The main contributions of this thesis are as follows:

- **A documented, reproducible evaluation workflow for evacuation signal plans.** The thesis assembles, for the 17-signal Sandy Hill network in Ottawa, a traceable census-based derivation of evacuation demand, a static User Equilibrium assignment used to calibrate the signal plans (turning flows, bottleneck identification and progression speed), the design of a Webster baseline and a coordinated green wave plan, and their evaluation by microsimulation.

- **A controlled comparison of isolated and coordinated signal control under evacuation demand.** The two plans share demand, phase structure, lost time and pedestrian minima, and differ only in the coordinated design package stated in Section 5.3. Over ten replications per plan with common random seeds, the coordinated

plan reduced total system travel time by 9.5 %, mean queue by 33.5 % and stopped time by 15.8 %, each significant at the 5 % level, and reduced the spread of outcomes across replications by a factor of three to seven. The effect is not uniform in time: it is statistically indistinguishable from zero while the corridor is lightly loaded and reaches 26.7 % at the peak of the loading profile.

- **An argument for evaluating evacuation signal plans by their temporal performance.** The results indicate that the value of a signal plan in an evacuation depends on *when* it is measured, and that plans should therefore be judged by the timing of the congestion crest and by the total evacuation clearance time rather than by the average delay over the demand period, which a steady-state assessment cannot provide.

1.3.3 Thesis Organization

The overall structure of the thesis is illustrated in Figure 1.7, which connects evacuation demand to the traffic assignment and signal control components, with traffic flow outcomes feeding back as updated travel times and intersection delays.

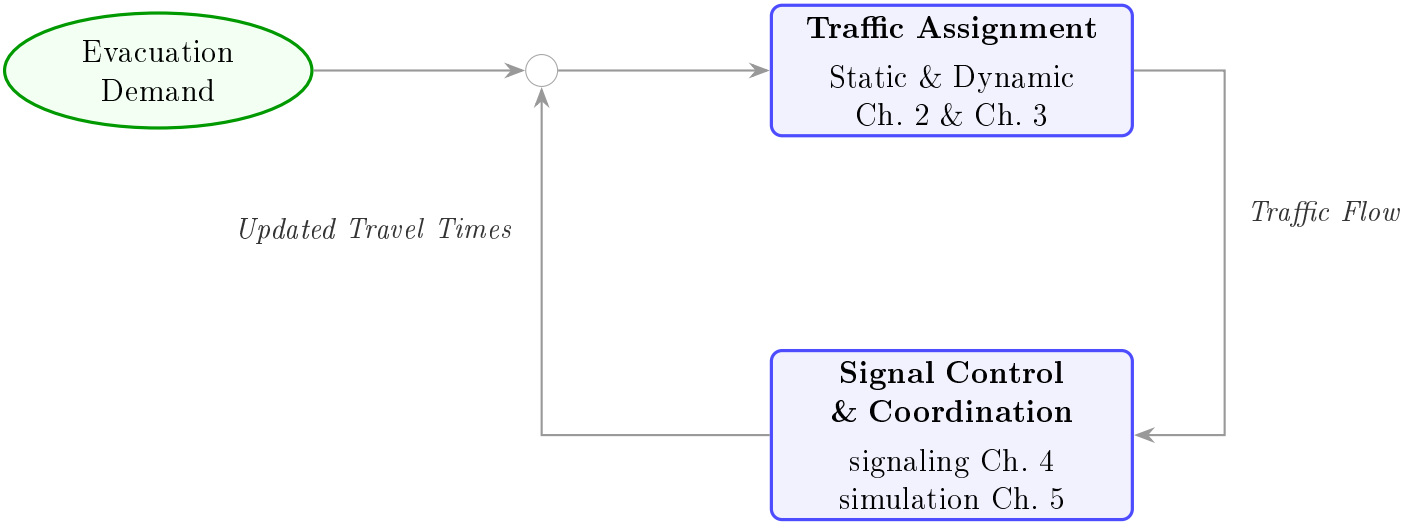


Figure 1.7: Research framework, connecting evacuation demand to traffic assignment.

The remainder of this thesis is organised as follows. Chapter 2 presents the theory of static traffic assignment, covering the User Equilibrium and System Optimum formulations, Braess's paradox, and the behavioural implications of each paradigm in the evacuation context. Chapter 3 extends this treatment to dynamic traffic assignment, introducing the distinction between instantaneous and experienced travel times, the limitations of static volume–delay functions, the formulation of the Dynamic User Equilibrium and Dynamic System

Optimum, and the one-shot stochastic route choice framework employed in the case study. Chapter 4 addresses traffic signal control, covering signal timing terminology, cycle length determination using the Webster method, lost time and effective green time allocation, and the design of coordinated green wave plans for emergency evacuation corridors. Chapter 5 presents the emergency evacuation simulation applied to the Sandy Hill network in Ottawa, describing the origin–destination demand matrix, the static User Equilibrium assignment, the signal timing methodology, and a detailed analysis of results from a microscopic simulation with fixed routes assigned by stochastic route choice, comparing the Webster and coordinated signal plans across network travel time, speed, flow, and destination specific performance metrics. Finally, Chapter 6 summarizes the contributions of the thesis and outlines directions for future work.

Chapter 2

Static Traffic Assignment

Static Traffic Assignment (STA) is a core method in transportation network modelling used to determine how a given travel demand is distributed across the links of a road network under steady state conditions. In STA, both the origin-destination (OD) demand and the network supply characteristics are treated as constant over the analysis period (e.g., a representative peak hour). As a result, STA provides a snapshot of average traffic conditions rather than explicitly modelling within period dynamics such as queue formation and dissipation.

The main purpose of STA is to compute consistent link flows and travel costs by recognising that travel times depend on congestion, and congestion depends on how many travellers choose each route. Drivers are assumed to select routes based on a performance measure (typically travel time or generalised cost), and the interaction between route choice and congestion is resolved through an equilibrium concept, most commonly the user equilibrium, which is discussed in detail in the following section. Practically, STA is widely used for planning and evaluation tasks such as network performance assessment, comparison of infrastructure alternatives, and estimation of link volumes needed for design and policy analysis.

2.1 User Equilibrium (UE)

User Equilibrium (UE), also known as Wardrop's first principle [16], is a descriptive concept in traffic modelling that explains how drivers distribute themselves across available routes when each traveller acts in their own self-interest. For a given origin–destination (OD) pair with multiple feasible routes, travellers, assumed to have full and accurate information about network conditions, compare route travel costs (most commonly travel time or a generalised cost) and choose the alternative with the least cost.

Because route travel time is a function of the traffic using that route, increased usage leads to greater congestion and, therefore, higher travel times. This congestion feedback

motivates travellers to shift away from slower routes toward faster ones, and the process continues until a stable state is reached in which no individual can reduce their own travel cost by unilaterally switching routes. At UE, all routes that are actually used between the OD pair have equal and minimal cost, and any unused route has a cost that is no lower than that of the used ones. The equivalent mathematical programming formulation of this condition was first given by Beckmann, McGuire, and Winsten [2], and its solution by convex programming is treated in detail by Sheffi [13].

2.1.1 UE Example: Two Parallel Routes

Figure 2.1 shows the simplest network in which UE can be understood clearly: a single OD pair with two parallel alternative routes. Let the total (fixed) travel demand between O and D be d (veh/h). Every traveller going from O to D must choose either Route 1 or Route 2. Denote the resulting route flows by q_1 and q_2, as labeled in the figure. Because the entire demand must be accommodated by these two routes, the first requirement is flow conservation:

$$q_1 + q_2 = d. \tag{2.1}$$

Congestion enters the problem through the fact that the travel time (or generalised travel cost) on each route depends on how many vehicles use that route. This is represented by increasing functions

$$t_1 = t_1(q_1), \qquad t_2 = t_2(q_2), \tag{2.2}$$

where $t_1(q_1)$ is the travel time on Route 1 and $t_2(q_2)$ is the travel time on Route 2. The essential modelling assumption is that both functions are nondecreasing: when a route carries more vehicles, it becomes slower due to congestion and interactions among vehicles.

The defining idea of UE is behavioural. Drivers are assumed to act selfishly (individually rationally): each driver prefers the route with the lower travel time. If, for a given split, Route 1 is faster than Route 2, then some drivers currently using Route 2 have an incentive to switch to Route 1. But as they switch, q_1 increases and $t_1(q_1)$ rises, while q_2 decreases and $t_2(q_2)$ falls. This switching process continues until no driver can gain by changing routes unilaterally. That stable point is the UE.

In this two route setting, the UE condition can be written compactly. If both routes carry positive flow at equilibrium, then their travel times must be equal; otherwise, drivers would abandon the slower route:

$$\text{if } q_1^* > 0 \text{ and } q_2^* > 0, \quad \text{then} \quad t_1(q_1^*) = t_2(q_2^*). \tag{2.3}$$

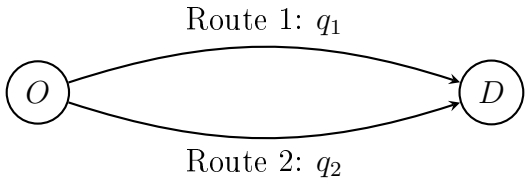


Figure 2.1: A simple two route parallel network connecting origin O and destination D, with route flows q_1 and q_2.

Finally, it is often convenient to eliminate one variable using the conservation constraint. From Eq. (2.1),

$$q_2 = d - q_1, \tag{2.4}$$

and the UE condition for the case in which both routes are used becomes

$$t_1(q_1^*) = t_2(d - q_1^*). \tag{2.5}$$

Conceptually, Eq. (2.5) states that the total demand d is split across the two routes such that the experienced travel times match. That matching is exactly what makes the split self-enforcing: once drivers see equal travel times on the used alternatives, nobody can reduce their own travel time by switching routes, so the assignment is at user equilibrium.

2.2 System Optimum (SO)

System Optimum (SO) traffic assignment, corresponding to Wardrop's second principle [16], is the system-level counterpart of User Equilibrium. While UE describes what happens when each traveller selfishly selects the route that minimises their own travel time, SO describes what a central planner (e.g., a traffic management center) would do if the goal were to minimise the *total travel time* experienced by all travellers in the network [29]. It represents the best achievable operating point from a network-efficiency perspective, but it does not necessarily arise naturally from individual route choice.

Consider a fixed OD pair with total travel demand d (veh/h). Let $i \in \mathcal{R}$ index the set of feasible routes connecting this OD pair. Denote by q_i the flow assigned to route i, and by $t_i(q_i)$ the corresponding route travel time (or generalised cost), assumed to be a nondecreasing function of route flow due to congestion effects. Feasibility requires that all OD demand is distributed across the available routes:

$$\sum_{i \in \mathcal{R}} q_i = d, \qquad q_i \geq 0 \quad \forall i \in \mathcal{R}. \tag{2.6}$$

The SO objective is to minimise total system travel time. The total travel time accumulated on route i equals the number of vehicles using that route times the travel time per vehicle, i.e., $q_i\, t_i(q_i)$. Summing over all routes yields the total travel time for this OD pair. The SO problem is therefore formulated as

$$\min_{\{q_i\}_{i\in\mathcal{R}}} \sum_{i\in\mathcal{R}} q_i\, t_i(q_i) \tag{2.7}$$

subject to the conservation constraint in Eq. (2.6). The objective in Eq. (2.7) is interpreted as minimizing the Total System Travel Time (TSTT), for this OD pair or, in a link-based formulation, for the entire network.

Lagrangian and first order optimality conditions. To solve Eq. (2.7) subject to Eq. (2.6), we use the method of Lagrange multipliers. The Lagrangian is

$$\mathcal{L}(\{q_i\},\lambda) = \sum_{i\in\mathcal{R}} q_i\, t_i(q_i) + \lambda\left(d - \sum_{i\in\mathcal{R}} q_i\right), \tag{2.8}$$

where λ is the multiplier associated with OD flow conservation. Differentiating Eq. (2.8) with respect to q_i and imposing stationarity gives

$$\frac{\partial\mathcal{L}}{\partial q_i} = t_i(q_i) + q_i\frac{dt_i(q_i)}{dq_i} - \lambda = 0. \tag{2.9}$$

This leads to the defining SO condition: the marginal (social) travel cost is equalised across all used routes,

$$m_i(q_i) \triangleq t_i(q_i) + q_i\frac{dt_i(q_i)}{dq_i} = \lambda, \qquad \forall i \in \mathcal{R} \text{ with } q_i > 0. \tag{2.10}$$

For unused routes ($q_i = 0$), the corresponding Karush Kuhn Tucker (KKT) conditions imply that their marginal cost cannot be lower than the common value λ (otherwise, it would be beneficial to assign flow to them).

The SO condition differs fundamentally from UE. Under UE, used routes equalise the private travel times $t_i(\cdot)$; under SO, used routes equalise the marginal costs $m_i(\cdot)$. The first component of m_i, namely $t_i(q_i)$, is the travel time experienced by an additional user choosing route i (the private cost). The second component, $q_i\,\frac{dt_i}{dq_i}$, represents the extra delay imposed on all other users already on route i when one more vehicle is added the classical congestion externality. Because each selfish user ignores this external delay, UE assignments tend to overuse some routes relative to the SO solution, thereby increasing total system travel time.

A crucial implication of SO is that the private route travel times need not be equal across used routes; it is the marginal costs that are equal. Consequently, an SO assignment may direct a portion of drivers onto a route that is slightly slower in terms of private travel time, because doing so reduces the overall congestion burden and improves total network performance. This is also why the SO solution serves as the efficiency benchmark for a network: its TSTT is a lower bound that no feasible assignment can improve upon, and reaching it in practice generally requires some form of management or incentive (e.g., congestion pricing, ramp metering, or system-oriented route guidance).

With UE and SO clearly distinguished, many counterintuitive phenomena in traffic networks become easier to explain. In particular, because UE reflects selfish route choice rather than system-wide optimisation, adding capacity (or a new link) can sometimes worsen network performance under UE by altering route choice incentives, even though the SO solution would not deteriorate in the same way. This mechanism underlies Braess's paradox, discussed in the following section.

2.3 Braess's Paradox in Traffic Networks

Braess's paradox [4] describes a counterintuitive situation in which adding a new road to a traffic network increases overall congestion instead of reducing it. This happens because drivers choose routes selfishly according to User Equilibrium, aiming to minimise only their own travel time without considering the effect of their decisions on others. When a new, faster link is introduced, many drivers switch to it, which can overload critical parts of the network and increase congestion on multiple links simultaneously. As a result, even though the network has more capacity, the total travel time experienced by all users becomes worse. Under System Optimum, the paradox would not occur, because traffic would be distributed to minimise total system delay rather than individual travel times. Braess's paradox therefore illustrates that individually rational behaviour does not always lead to socially optimal outcomes in transportation networks.

To demonstrate the paradox, consider the classical Braess network consisting of one origin (Start), one destination (End), and two intermediate nodes A and B (Figure 2.2). A total demand of $d = 4000$ drivers per hour wishes to travel from Start to End. Denote by x_{SA}, x_{AE}, x_{SB}, and x_{BE} the flows on the corresponding links. The link travel times (in minutes) are

$$t_{SA} = \frac{x_{SA}}{100}, \qquad t_{AE} = 45, \tag{2.11}$$

$$t_{SB} = 45, \qquad t_{BE} = \frac{x_{BE}}{100}, \tag{2.12}$$

where the coefficient 1/100 is the congestion sensitivity of the flow dependent links: each additional 100 veh/h adds one minute of travel time. (It is a cost-function slope, not a capacity; flow on these links is not bounded by 100.)

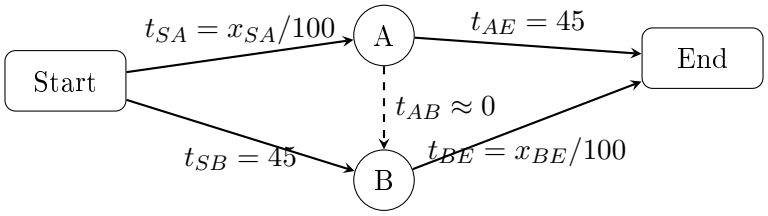


Figure 2.2: The Braess network. Links Start–A and B–End carry flow-dependent travel times; links A–End and Start–B are constant. The dashed link A–B is the addition considered in the modified network.

Initial network (without the A–B link). The available routes are Start-A-End and Start-B-End. Let a denote the flow on route Start-A-End and b the flow on route Start-B-End, so that $x_{SA} = a$ and $x_{BE} = b$. Demand conservation gives

$$a + b = d = 4000. \tag{2.13}$$

The route travel times are therefore

$$t_1 = \frac{a}{100} + 45, \qquad t_2 = \frac{b}{100} + 45. \tag{2.14}$$

At User Equilibrium, both used routes must have equal travel time:

$$\frac{a}{100} + 45 = \frac{b}{100} + 45 \quad \Longrightarrow \quad a = b = 2000. \tag{2.15}$$

The equilibrium travel time is therefore

$$t_1 = t_2 = \frac{2000}{100} + 45 = 65 \text{ minutes}, \tag{2.16}$$

so before the new road is added, every driver experiences a travel time of 65 minutes.

Modified network (with the A–B link). Now suppose a new road is built between A and B with negligible travel time, $t_{AB} \approx 0$. Drivers can now use the combined route Start–A–B–End. Under selfish route choice, suppose at equilibrium all 4000 drivers use this new

route. Then $x_{SA} = x_{BE} = 4000$, and the travel time on the new route is

$$t_{\text{new}} = \frac{4000}{100} + 0 + \frac{4000}{100} = 80 \text{ minutes.} \tag{2.17}$$

Checking the original routes under these link flows,

$$t_1 = \frac{4000}{100} + 45 = 85 \text{ minutes}, \qquad t_2 = 45 + \frac{4000}{100} = 85 \text{ minutes.} \tag{2.18}$$

Since both original alternatives are slower than 80 minutes, no driver has an incentive to deviate; the assignment in which all drivers use Start–A–B–End is therefore the new User Equilibrium, with $t_{\text{UE}} = 80$ minutes.

Paradoxical outcome. Before the new road, the equilibrium travel time was 65 minutes; after the new road, it is 80 minutes. Adding a zero cost link has thus increased the equilibrium travel time of every user by 15 minutes, and if the A–B link were closed or restricted, all drivers would benefit. This is Braess's paradox. It occurs because UE reflects selfish decision-making rather than system wide optimisation: under System Optimum, the new link would be left unused (or used only partially) so that total travel time is minimised. The paradox therefore highlights the fundamental behavioural difference between UE and SO, and demonstrates why unregulated infrastructure expansion can sometimes worsen congestion.

2.4 Static Behavioural Principles in Evacuation

This section is a qualitative illustration. No link performance functions, capacities or demand are specified for the network of Figure 2.3, and no equilibrium is computed on it; its purpose is to make concrete, in an evacuation setting, the distinction between the two behavioural paradigms formalised in Sections 2.1–2.2. The quantitative work of the thesis is carried out on the Sandy Hill network of Chapter 5.

To conceptualize the dynamics of static traffic assignment in a high stakes scenario, consider a simplified road network designed Figure 2.3 for an emergency evacuation. This model connects a high risk Origin (Evacuation Zone) to a Destination (Safe Area) via four parallel routes, each characterized by distinct physical and operational parameters:

- **Route 1 (Arterial):** A standard high capacity surface street with moderate free-flow speeds and a medium sensitivity to increased volume.
- **Route 2 (Freeway):** A high speed, multi lane facility engineered for high-volume throughput, representing the primary evacuation corridor.

- **Route 3 (Local Road):** A narrower, tertiary alternative with lower initial speeds and a high sensitivity to congestion.
- **Route 4 (Contraflow/Emergency):** A specialised operational strategy, such as reversed lanes or dedicated emergency only corridors, typically activated by authorities during critical events to expand available capacity.

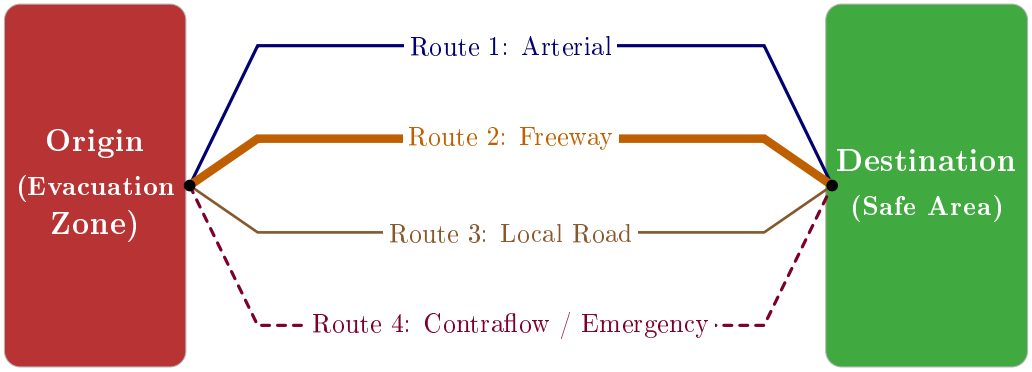


Figure 2.3: Simplified four-route evacuation network connecting origin to destination.

Each link represents an alternative route available to drivers travelling between these two points. The routes differ in physical and operational characteristics free-flow travel time, capacity, and sensitivity to congestion which together determine how attractive each route is under different traffic conditions.

As the density of evacuees increases, the physical space on each road becomes a finite resource. Vehicles begin to impede one another, illustrating the fundamental congestion effect: as the volume of traffic approaches a link's capacity, the travel time required to traverse it increases nonlinearly. The distribution of traffic across the four routes is then determined by the static behavioural paradigm governing the system: User Equilibrium or System Optimum.

Under the UE paradigm, evacuation behaviour is driven by individual rationality and self-interest. In this descriptive model, each driver acts as an independent decision-maker seeking to minimise their own personal travel time. In a crisis, this behaviour typically leads to the immediate saturation of the most attractive routes, such as the freeway or the arterial.

In contrast, the SO paradigm represents a normative, prescriptive perspective, reflecting the objectives of an emergency manager or a centralized traffic control authority. The goal of the SO formulation is not to satisfy the individual, but to minimise the Total System Travel Time, ensuring that the entire evacuation zone is cleared in the shortest possible time.

2.4.1 Behavioural and Practical Implications

The contrast between UE and SO provides insight into how individual behaviour shapes congestion and how network performance can be improved through coordination. Under UE, each traveller acts rationally from a narrow, individual perspective by choosing the route with the lowest travel time. However, every additional vehicle also increases the delay experienced by others, so the resulting equilibrium is stable but generally inefficient, a classic negative externality in which private and social costs diverge.

The SO solution, by contrast, internalises this externality by equalising marginal costs rather than private travel times, as shown in Eq. (2.10): flows are redistributed across routes so that the total delay experienced by all travellers is minimised. The gap between the total travel times under UE and under SO therefore measures the potential benefit of coordination, and represents an upper bound on the efficiency gains achievable through management or control strategies.

It should be said clearly how far this carries into the case study. The Sandy Hill model is loaded and evaluated under user equilibrium behaviour: the static assignment of Section 5.2 is a UE assignment, and the microsimulation routes vehicles by their own perceived travel times. No system optimal assignment is computed for Sandy Hill, and the thesis therefore does not report a price of anarchy for it. The SO formulation is used here as the reference point that explains *why* network management is needed at all, and the same logic carries into the dynamic case in Chapter 3, where the interval by interval DUE and DSO conditions of Section 3.4.2 extend the same contrast to a time varying network. That contrast motivates the instrument this thesis actually evaluates: signal control, which cannot dictate routes but can redistribute the capacity that determines them.

These behavioural mechanisms translate directly into practical measures. In an unmanaged network, UE is the natural outcome, as users respond myopically to congestion. To move the system closer to SO, planners can apply instruments such as congestion pricing, route guidance, or advanced traveller information systems that influence route choice. For example, Pigouvian tolls [11] that charge each driver an amount equal to the external delay they impose, $q_i \frac{dt_i}{dq_i}$, align private costs with the marginal social costs of Eq. (2.10), so that the user equilibrium under tolls coincides with the SO allocation. In the evacuation setting of the four-route network above, the same principle implies that efficient operation is achieved not by eliminating congestion, but by distributing it deliberately across the arterial, freeway, local, and contraflow alternatives. In this thesis, that redistributive role is played not by pricing but by traffic signal control and coordination, developed in the subsequent chapters.

Chapter 3

Dynamic Traffic Assignment

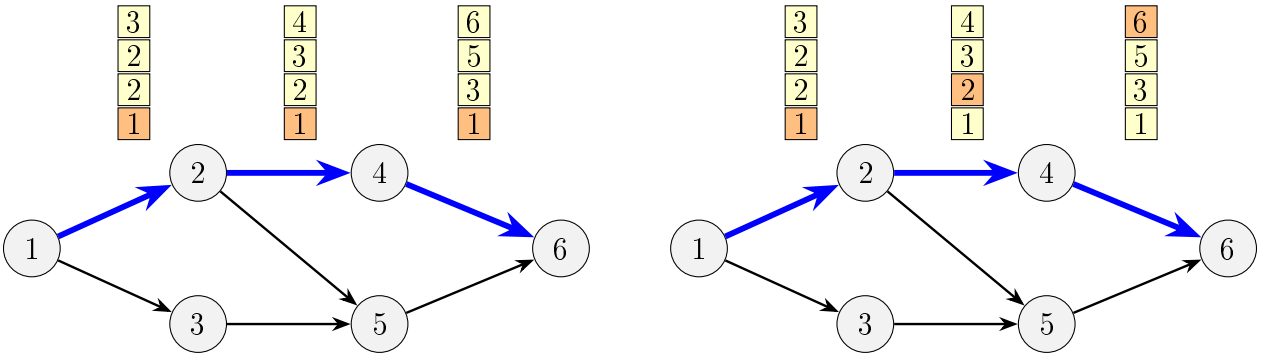
The static assignment framework of Chapter 2 represents network conditions as a single steady state, which is adequate for long-term planning but cannot describe how congestion builds, propagates, and dissipates within the analysis period. Dynamic Traffic Assignment (DTA) extends this framework by representing the interaction between travel choices, traffic flows, and travel times in a temporally coherent manner, so that route choice responds to network conditions as they evolve rather than to a fixed average [5, 10]. This temporal dimension is precisely what an evacuation analysis requires: the demand surge, the growth and spillback of queues, and the eventual recovery of the network are all inherently time varying phenomena. This chapter introduces the behavioural distinction between instantaneous and experienced travel times, examines the physical and mathematical limitations of static models that motivate the dynamic approach, develops a discrete time formulation with explicit queueing, and contrasts the Dynamic User Equilibrium and Dynamic System Optimum principles, closing with the one-shot simulation framework employed for rapid evacuation analysis.

3.1 Instantaneous and Experienced Travel Times

The core of DTA modelling lies in the distinction between *instantaneous* and *experienced* travel times. As illustrated in the left panel of Figure 3.1, Instantaneous Travel Time is essentially a "snapshot" of the network at the exact moment a traveller departs. It assumes that the congestion levels currently visible on a GPS app will remain frozen in time for the entire duration of the trip. While this is computationally simple, it is behaviourally "myopic" or short-sighted. It ignores the fact that by the time a driver reaches a downstream bridge or intersection, the traffic there will likely have changed significantly.

In contrast, the right panel of Figure 3.1 depicts the Experienced Travel Time. This

measure respects the "temporal causality" of movement. It calculates the time spent on the first link, updates the clock, and then looks at the traffic conditions on the next link at that future arrival time. In an evacuation, where thousands of vehicles are entering the network simultaneously, the experienced time is almost always higher than the instantaneous snapshot because the traveller is effectively "chasing" the buildup of congestion as they move deeper into the network.



(I) Route 1−2−4−6: $1+1+1=3$ (E) Route 1−2−4−6: $1+2+6=9$

Figure 3.1: Instantaneous and experienced travel time on route 1–2–4–6. The instantaneous computation (I) reads every link at departure; the experienced computation (E) reads each link on arrival.

3.1.1 The Physical Reality of Congestion: Spillback

The primary physical phenomenon that creates the gap between instantaneous and experienced travel times is *spillback*. In traditional static models, as illustrated in the upper panel of Figure 3.2, congestion is contained within a single link: even if a roadway segment is over capacity, the model simply increases the travel time of that specific link without affecting the links behind it. This representation is physically impossible, because in a real urban network vehicles are physical objects that occupy space.

As shown in the lower panel of Figure 3.2, DTA instead models the physics of *queue propagation*. When a bottleneck forms due to an incident or a capacity drop, a queue begins to grow; once that queue reaches the upstream end of a link, it spills back into the previous intersection, blocking traffic that may not even be heading toward the original bottleneck.

This spillback mechanism is the engine of dynamic congestion. Because it takes time for a spillback wave to travel backward through the network (the backward propagating shock waves introduced in Section 1.1.4), an instantaneous snapshot taken at the moment of departure may show a clear highway ahead, while the experienced reality for the driver is

a gridlocked network, reached precisely when the tail of a spillback wave that had not yet arrived at departure time finally meets the vehicle.

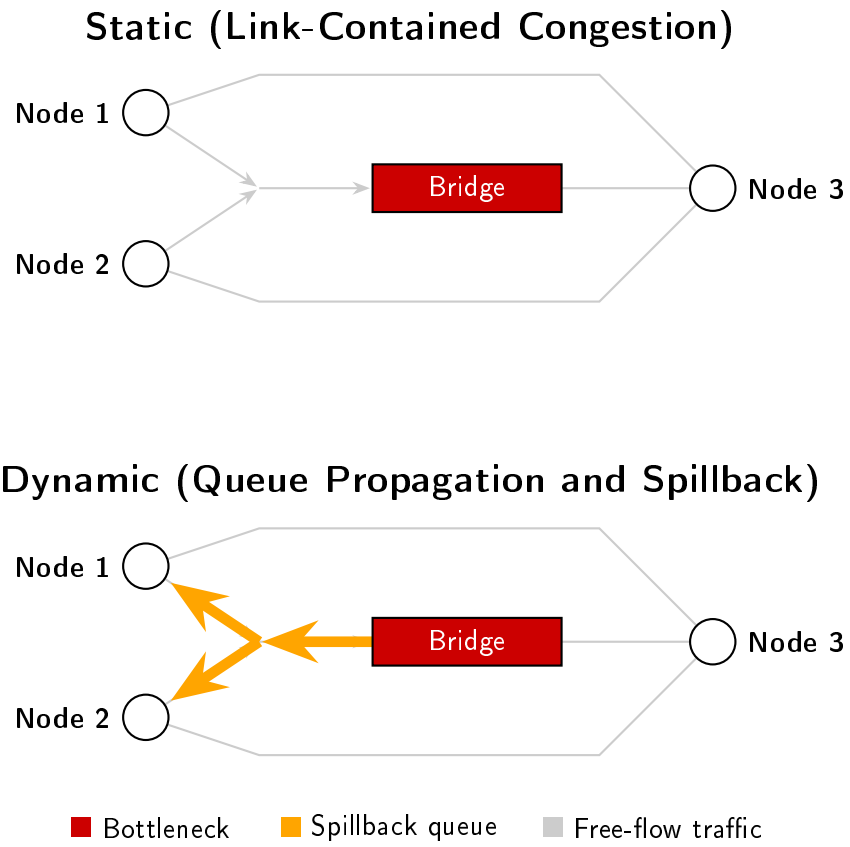


Figure 3.2: Spillback at a network bottleneck. Above, the static representation keeps congestion on the bottleneck link. Below, the queue propagates upstream and blocks approaches not bound for it.

3.2 Limitations of Static Modelling

While experienced travel time and spillback provide the behavioural basis for DTA, it is essential to examine the mathematical limitations of static models that necessitate this shift. In static assignment, link congestion is typically described by a *volume–delay function* (VDF), which treats travel time as a steady-state function of assigned volume.

Although computationally convenient, VDFs embody a critical physical contradiction, namely the unbounded V/C $>$ 1 phenomenon. In a static framework, inflow always equals outflow, so the model lacks a physical mechanism to restrict flow when a link reaches its capacity. Consequently, the assigned volume-to-capacity ratio can exceed unity without bound. As discussed in Section 1.1.3, a ratio above one is meaningful only when its numerator

is interpreted as demand: it indicates that the desired demand exceeds the feasible flow. The static model, however, propagates this excess volume downstream as though it had actually been served, which is physically impossible; vehicles that cannot be discharged must in reality accumulate. Because the assigned volume is not constrained by the physical supply, static models report travel times that do not account for this accumulation.

Furthermore, because VDF-based models operate on a single average condition over the analysis period, they are inherently unable to capture the rapid transitions typical of emergency evacuations, such as the demand peak, the buildup of queues, and the subsequent recovery. Within the DTA framework, the VDF is replaced by the fundamental diagram introduced in Section 1.1.3: rather than directly specifying a delay for a given volume, the fundamental diagram governs the propagation of congestion. This ensures that when a link reaches its physical capacity, the excess demand forms a queue that triggers the spillback mechanism described above, rather than allowing the assigned volume to grow without limit.

3.2.1 The FIFO Limitation and Overtaking Behaviour

Another fundamental modelling issue is the treatment of *first-in, first-out* (FIFO) behaviour. Static assignment models commonly impose this assumption at the link level: all vehicles using a given link are assigned the same travel time, regardless of their destination or turning movement at the downstream node. A direct implication is that overtaking is not represented.

In real traffic, however, overtaking is frequent and often critical to network performance. For instance, if a left-turn bay is severely oversaturated while the through lanes remain uncongested, vehicles queued for the turn may be overtaken by vehicles using the less congested through movement. In a static model, all of these vehicles are aggregated into the same link volume, resulting in an identical and often inflated travel time for everyone.

Models that move individual vehicles explicitly across separate lanes (lane-based DTA) can represent this non-ordered behaviour more realistically. By relaxing the strict link-level FIFO rule, DTA captures how different lanes experience different congestion states, ensuring that the computed experienced travel times reflect the tactical choices actually available to drivers in a congested network.

Table 3.1 summarizes the fundamental differences between static (VDF-based) and dynamic traffic assignment models across the physical and operational dimensions discussed above.

Table 3.1: Differences between static and dynamic traffic assignment.

Feature	**Static assignment**	**Dynamic assignment**
Physics	Volume–delay function (curve relating average travel time to volume)	Fundamental diagram (speed-density-flow relationship governing traffic dynamics)
Capacity	Soft limit; volume-to-capacity ratio can exceed one	Hard limit; explicit physical capacity constraint on flow
Flow equality	Inflow equals outflow at all times	Inflow may differ from outflow; queues and accumulation can form
Congestion	Represented as high travel time on a single link only	Represented as physical queues that can propagate upstream (spillback)
Lanes	Single "pipe" representation with FIFO enforced; no lane differentiation	Lane-specific behaviour; overtaking and lane-based effects can be represented
Best application	Long-term regional planning and strategic capacity studies	Operational analysis, signal and ramp control, intelligent transportation systems (ITS)

3.2.2 The DTA Synthesis: Mesoscopic and Iterative Modelling

The DTA methodology, particularly when implemented with mesoscopic simulation logic, offers a synthesis that addresses these limitations:

1. **Temporal coherence:** DTA explicitly models the accumulation of queues and the time variation of congestion using flow conservation, directly linking travel time to density.

2. **Equilibrium on experienced times:** DTA models apply iterative procedures, alternating network loading and assignment, that adjust traveller route choices until they converge to routes of minimal *experienced* travel time, accounting for queueing and system-wide congestion; this is the Dynamic User Equilibrium developed in Section 3.4.

3. **Scope and efficiency:** many DTA implementations adopt computationally efficient mesoscopic simulation logic, as introduced in Section 1.1.1. These models represent individual vehicles and their dynamic states through simplified car-following or traffic flow relationships, allowing traffic to be described from the corridor scale up to the regional scale, and over horizons of up to 24 hours, in contrast to the localized scope of detailed microscopic models.

3.3 Core Modelling Principles and Notation

We divide the finite analysis horizon into discrete, uniform time intervals, indexed by $k = 1, 2, \ldots, K$, each with a duration of Δ. This discretization converts the continuous time traffic evolution process into a tractable sequence of time steps, enabling explicit modelling of time-varying demand, capacity constraints, and congestion dynamics.

The network is represented at an aggregate level by a finite set of routes indexed by $i = 1, \ldots, R$; in the evacuation example developed in Chapter 2, $R = 4$, corresponding to the arterial, freeway, local, and contraflow alternatives. Since routes differ in travel-time and capacity characteristics, their performance evolves differently as traffic is assigned over time. Consequently, route assignment decisions in the interval k influence both the current travel conditions and the subsequent system state.

3.3.1 State Variables and Definitions

- k: time interval index (e.g., $k = 1$ is 0–10 min, $k = 2$ is 10–20 min).
- Δ: duration of the time interval (e.g., $\Delta = 10$ minutes).

- D_k: total demand (veh) released into the network in interval k.
- $C_{i,k}$: capacity of route i (veh/h) during interval k.
- $S_{i,k}$: maximum admission cap for route i in interval k (veh/interval).
- $Q_{i,k}$: entrance queue (veh) on route i at the start of interval k.
- $d_{i,k}$: flow (veh) requested for route i in interval k.
- $a_{i,k}$: flow (veh) admitted onto route i in interval k.

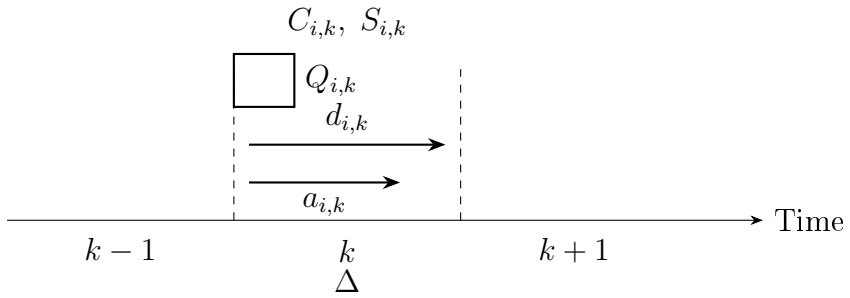


Figure 3.3: Time discretization for route i in interval k. The queue $Q_{i,k}$ carries in, the demand $d_{i,k}$ arrives within the interval, and the admitted flow $a_{i,k}$ enters subject to capacity.

Figure 3.3 illustrates the time-discretized structure of the model. Time is divided into intervals of equal length Δ, indexed by k. For each route i, $Q_{i,k}$ denotes the entrance queue at the beginning of interval k. During the interval, a demand $d_{i,k}$ is generated, of which $a_{i,k}$ vehicles are admitted to the route, subject to the route capacity $C_{i,k}$ and the admission cap $S_{i,k}$.

3.3.2 Conservation of Flow and Admission Cap

The primary constraint linking capacity and flow is the admission cap, obtained from the hourly capacity $C_{i,k}$ and the interval duration Δ (in minutes):

$$S_{i,k} = \frac{C_{i,k}}{60}\,\Delta. \tag{3.1}$$

For $\Delta = 10$ minutes, Eq. (3.1) simplifies to $S_{i,k} = C_{i,k}/6$. The admitted flow $a_{i,k}$ is then determined by the minimum of the available demand (queue carryover plus newly requested flow) and the admission cap:

$$a_{i,k} = \min\{\, Q_{i,k} + d_{i,k},\ S_{i,k} \,\}. \tag{3.2}$$

The Queue Recursion: Explicit Congestion Modelling

The evolution of the system state relies on the conservation equation, which determines the queue at the start of the next interval:

$$Q_{i,k+1} = \max\{0,\ Q_{i,k} + d_{i,k} - a_{i,k}\}. \tag{3.3}$$

This recursion is the fundamental principle linking one time step to the next and is the source of all dynamic effects: accumulation ($d_{i,k} > a_{i,k}$) is modelled explicitly, in contrast to static models.

The time experienced by a traveller is the sum of the time spent waiting in the entrance queue (queue delay) and the time spent traversing the road section (running time). The two components are developed in turn.

Running Time

We adopt a generalised BPR (Bureau of Public Roads) function [3], specifically its affine form (power parameter $\beta = 1$) for simplicity of the base model:

$$t_{i,k}(F_{i,k}) = t_{0,i}\left[1 + \alpha_i\left(\frac{F_{i,k}}{C_{i,k}}\right)\right], \tag{3.4}$$

where $t_{0,i}$ is the free-flow time of route i, α_i is a congestion sensitivity parameter, and $F_{i,k}$ is the flow rate (veh/h). Since the admitted flow $a_{i,k}$ is expressed in vehicles per interval, the flow rate is $F_{i,k} = a_{i,k}\,(60/\Delta)$; for $\Delta = 10$ min, $F_{i,k} = 6\,a_{i,k}$.

Queue Delay

The time spent waiting in the entrance queue is proportional to the queue size and inversely proportional to the service rate, which equals the capacity expressed per minute:

$$\text{queue delay} = \frac{Q_{i,k}}{C_{i,k}/60} = \frac{60\,Q_{i,k}}{C_{i,k}}. \tag{3.5}$$

Total Experienced Time

Combining the two components yields the experienced time (in minutes):

$$T_{i,k}^{\text{exp}}(a_{i,k}, Q_{i,k}) = t_{0,i}\left[1 + \alpha_i\left(\frac{6\,a_{i,k}}{C_{i,k}}\right)\right] + \frac{60\,Q_{i,k}}{C_{i,k}}. \tag{3.6}$$

Equation (3.6) is central to the formulation: the decision to admit $a_{i,k}$ affects the running time, while the state $Q_{i,k}$ determines the queue delay.

3.4 Dynamic User Equilibrium and Dynamic System Optimum

Within each interval k, the assignment of flows to routes is governed by one of two behavioural principles, which extend the static UE and SO concepts of Chapter 2 to the time-dependent setting.

3.4.1 Dynamic User Equilibrium (DUE)

Under DUE, travellers minimise their own experienced time; in equilibrium, no traveller can unilaterally switch routes and improve their travel time. The conditions are:

- **Equality condition:** on all used routes i ($a_{i,k} > 0$), the experienced time equals a common minimal time Λ_k:

$$a_{i,k} > 0 \quad \Longrightarrow \quad T_{i,k}^{\text{exp}} = \Lambda_k. \tag{3.7}$$

- **Inequality condition:** on unused routes j, the experienced time satisfies $T_{j,k}^{\text{exp}} \geq \Lambda_k$.

3.4.2 Dynamic System Optimum (DSO)

Under DSO, a central planner minimises the Total System Travel Time (TSTT)[27, 28] of all travellers currently using the network:

- **Objective:** minimise the total system cost in interval k,

$$Z_k = \sum_i a_{i,k}\, T_{i,k}^{\text{exp}}(a_{i,k}, Q_{i,k}). \tag{3.8}$$

- **Equality condition:** on all used routes i ($a_{i,k} > 0$), the *marginal travel cost* $M_{i,k}$ equals a common value μ_k:

$$a_{i,k} > 0 \quad \Longrightarrow \quad M_{i,k} = \mu_k. \tag{3.9}$$

The marginal cost is the derivative of the total cost on route i with respect to the admitted flow $a_{i,k}$:

$$M_{i,k} = \frac{\partial \left(a_{i,k}\, T_{i,k}^{\exp}\right)}{\partial a_{i,k}} = \frac{\partial}{\partial a_{i,k}} \left[a_{i,k} \left(t_{0,i} \left[1 + \alpha_i \frac{6\, a_{i,k}}{C_{i,k}} \right] + \frac{60\, Q_{i,k}}{C_{i,k}} \right) \right]. \tag{3.10}$$

Applying the product rule yields

$$M_{i,k} = t_{0,i} \left(1 + 2\alpha_i \frac{6\, a_{i,k}}{C_{i,k}} \right) + \frac{60\, Q_{i,k}}{C_{i,k}}. \tag{3.11}$$

Crucially, the coefficient on the flow term in Eq. (3.11) is $2\alpha_i$ double that of the DUE cost in Eq. (3.6) reflecting the negative externality imposed on other travellers, in direct analogy with the static marginal cost of Chapter 2. The queue term $60\, Q_{i,k}/C_{i,k}$ is carried unchanged because $Q_{i,k}$ is inherited from the previous interval and does not depend on $a_{i,k}$.

This interval-by-interval formulation should be read as a *myopic* approximation to a dynamic system optimum. Each interval k is optimised in isolation, so the effect of the admitted flow $a_{i,k}$ on the queues it leaves behind, $Q_{i,k+1}$ in Eq. (3.3), and hence on the travel times of all later intervals, is not priced. A full dynamic system optimum would minimise the total cost over the whole horizon, $\sum_k Z_k$, and would therefore charge each admitted vehicle for the delay it imposes on later intervals as well.

In large scale emergency evacuation scenarios, the computational burden of traditional iterative, equilibrium seeking methods often becomes prohibitive. When immediate decision support is required, the operational priority shifts from achieving exact mathematical convergence to obtaining rapid, feasible solutions. In such time constrained environments, where the speed of system resolution is as vital as its precision, DTA within a one-shot simulation framework provides a pragmatic alternative. This approach prioritizes computational speed and is well suited to the rapid scenario testing required in crisis management, as detailed in the following section.

3.5 Dynamic Assignment in a One-Shot Simulation

In dynamic traffic modelling, one-shot assignment is a non-iterative procedure in which vehicles are loaded and routed within a single simulation run. Its objective is to represent time-varying traffic propagation without repeated assignment reloading cycles for equilibrium search. Consequently, one-shot methods are computationally efficient and well suited to rapid scenario testing.

In its simplest form, one-shot assignment uses predefined routes that remain unchanged during the simulation. Vehicles are associated with routes at departure, and network loading

then produces the resulting link flows and travel times. Although this implementation is simple and fast, it does not capture adaptive route choice under evolving congestion; it is therefore more accurately interpreted as route-based network loading than as equilibrium oriented assignment.

A more flexible implementation is dynamic one-shot assignment, in which routes are updated incrementally using current traffic states. As illustrated in Figure 3.4, link conditions are periodically updated, instantaneous shortest paths are recomputed, and newly generated vehicles are assigned to the currently shortest routes. In some implementations, selected en-route vehicles may also re-evaluate their choices at decision nodes and switch routes when a shorter alternative becomes available.

Despite this added flexibility, one-shot assignments remain fundamentally myopic, because routing decisions are typically based on instantaneous information rather than on fully experienced travel times over the remaining trip. This may attract excessive demand to temporarily favourable corridors and amplify local congestion. For this reason, one-shot results should be interpreted with care in policy evaluation and long-horizon behavioural analysis.

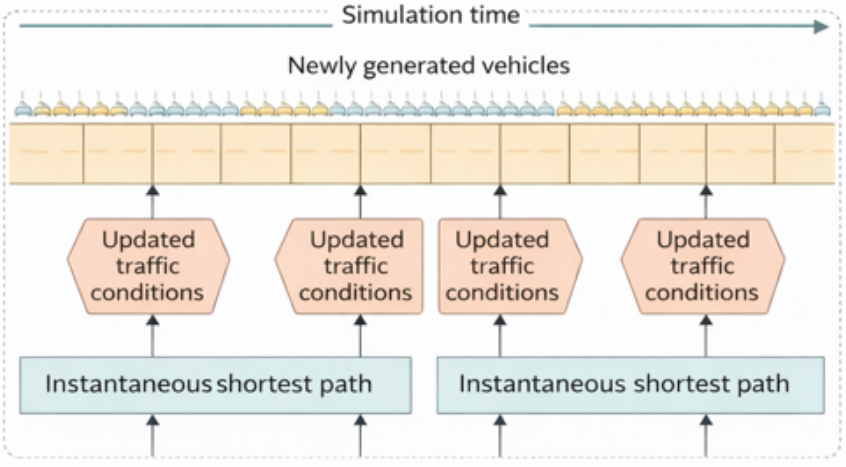


Figure 3.4: One-shot dynamic assignment with incremental route updating. Adapted by the author from Chiu et al. [5].

Overall, one-shot assignment is best described as network loading with incremental route updating. It is valuable for short-term operational analysis, sensitivity testing, fast comparison of management alternatives, and emergency evacuation planning; however, it does not, by itself, enforce equilibrium consistency between route-choice inputs and experienced travel-time outcomes.

3.6 Deterministic, Non Deterministic and Unobservable Elements

The framework in this thesis combines components that differ in how they respond to repetition. Separating them clarifies which results can be reproduced exactly, which carry a distribution, and which cannot be verified at all.

Deterministic components. A model is deterministic when the same inputs always produce the same output. The static assignment of Chapter 2 is deterministic in this sense: the demand matrix, the network and the cost functions fix a single equilibrium, and re-solving the problem returns the same link flows every time. The signal timing design of Chapter 4 is deterministic for the same reason. Webster's cycle length, the critical flow ratios and the green splits follow by arithmetic from the assigned flows, so the plan is fully specified once the flows are known. There is no randomness anywhere in that chain, and the offsets in the coordinated plan follow deterministically from the junction spacing and the progression speed.

This matters for what the design can be held to. If a reviewer recomputes the cycle at Rideau and King Edward from the flows reported in Table 5.1, they will obtain 55 s, not a value near it.

Non-deterministic components. The microscopic simulation is not deterministic. Vehicle generation, the assignment of desired speeds, lane-changing decisions and gap acceptance are all drawn from distributions, so the same scenario run twice with different random seeds produces different outcomes. The magnitude of that variation is not negligible here. Across the ten replications of the Webster plan alone, travel time per vehicle ranged from 100.8 to 115.7 s, and the number of vehicles actually generated varied from 3,476 to 3,619 against a nominal 3,530. That spread is of the same order as the effect being measured, which is why the comparison is made pairwise, with matched seeds, rather than between group means.

The non-determinism is not a defect to be eliminated. It is the property that makes the second principal finding of this thesis possible. Reporting a single run would give a point estimate and no indication of how dependable it is; reporting five gives both the mean and the spread, and the spread turned out to differ sharply between the two plans. At the demand peak the isolated plan spans ± 17.0 s per vehicle across replications against ± 3.3 s under coordination. An evacuation plan is executed once, without rehearsal, so that difference is arguably as important as the difference in the means.

Unobservable elements. A third category is neither deterministic nor random in a way the model can represent, because it cannot be measured at all. Several inputs to this study fall into it.

No evacuation of Sandy Hill has ever taken place, so the compliance rate η, the departure-time distribution and the destination split cannot be validated against observation. They are assumed, and each is stated as an assumption in Section 5.1 with the sensitivity of the result to it. The same applies to behaviour under stress: whether drivers in a real evacuation accept smaller gaps, follow more closely, or obey signals less reliably than the car-following model represents is not knowable from any available data.

The distinction between unobservable and merely unmeasured is worth keeping. Background traffic on King Edward Avenue is unmeasured in this study but it is observable in principle, and a count programme would supply it. Evacuation compliance is not: it can only be observed during an evacuation, and by then the plan has already been executed.

What this means for the conclusions. The three categories carry different kinds of claim. The signal design is deterministic, so it is reproducible and can be checked. The simulated performance is non-deterministic, so it is reported with confidence intervals and tested pairwise, and the claim is about a distribution rather than a number. The demand inputs are unobservable, so no absolute claim is made about how long Sandy Hill would take to clear.

What survives all three is the relative comparison. Both plans are evaluated on the same network, under the same demand realisation, with the same seeds. Whatever is unobservable in the demand is unobservable identically for both, and whatever is random in the simulation is paired away. The conclusion is therefore about the ordering of the two plans, which is defensible, rather than about the absolute evacuation time, which is not.

Chapter 4

Traffic Signal Control

In the previous chapters, the System Optimum concept described what a best possible traffic pattern looks like for the whole network, in a static sense (Chapter 2) or in a time-dependent sense (Chapter 3). In practice, municipalities cannot instruct individual drivers which route to take, but they can strongly influence how well the network performs by controlling intersections. Traffic signal control is therefore the main tool that planners and operators actually possess in the field. By choosing the order and timing of the green, amber, and red phases, signal control changes how much service each movement receives, affects delays and queue growth, and can push traffic operations toward a more efficient use of the available capacity. From this perspective, signal timing is one of the most direct means of moving traffic conditions closer to the System Optimum objective discussed earlier: within the research framework of Figure 1.7, signal control is the component through which traffic assignment and network performance are linked, as the timing plans shape the traffic flow that in turn updates the travel times and delays governing assignment.

Traffic signal control refers to the set of strategies and operational decisions used to manage traffic movements at intersections and along corridors through the timing of traffic signals. The primary objective of signal control is to allocate green time among competing traffic streams in a way that reduces delay, queue lengths, and unnecessary stops while improving overall safety and efficiency. This is achieved by adjusting key parameters such as cycle length and green splits, using either fixed-time plans or adaptive control strategies that respond to prevailing traffic conditions.

Under heavy or congested conditions, signal control plays a critical role in preventing queue spillback and delaying the onset of oversaturation. However, its effectiveness is ultimately limited by the physical discharge capacity and the available storage space at the intersection. When demand exceeds what the intersection can serve, even well-timed signals cannot remove congestion completely. For that reason, signal control is often deployed as the

first and most practical management measure, and it may need to be supported by geometric changes or other non-signal strategies in severely oversaturated networks.

4.1 Terms Related to Signal Timing

Several standard definitions are required for the analysis and design of signalized intersections; they are illustrated in Figure 4.1.

Cycle length the total time required to complete one full sequence of signal indications for all phases. It is denoted by C and expressed in seconds.

Phase a set of compatible traffic movements that receive the right-of-way simultaneously. Each phase consists of an effective green interval followed by a clearance interval.

Interval any distinct time segment within a signal cycle, including the green, yellow, or all-red intervals.

Change interval the yellow time together with the all red time, designed to provide safe clearance between conflicting movements.

All-red interval the period during which all approaches face a red indication, allowing vehicles already in the intersection to clear safely.

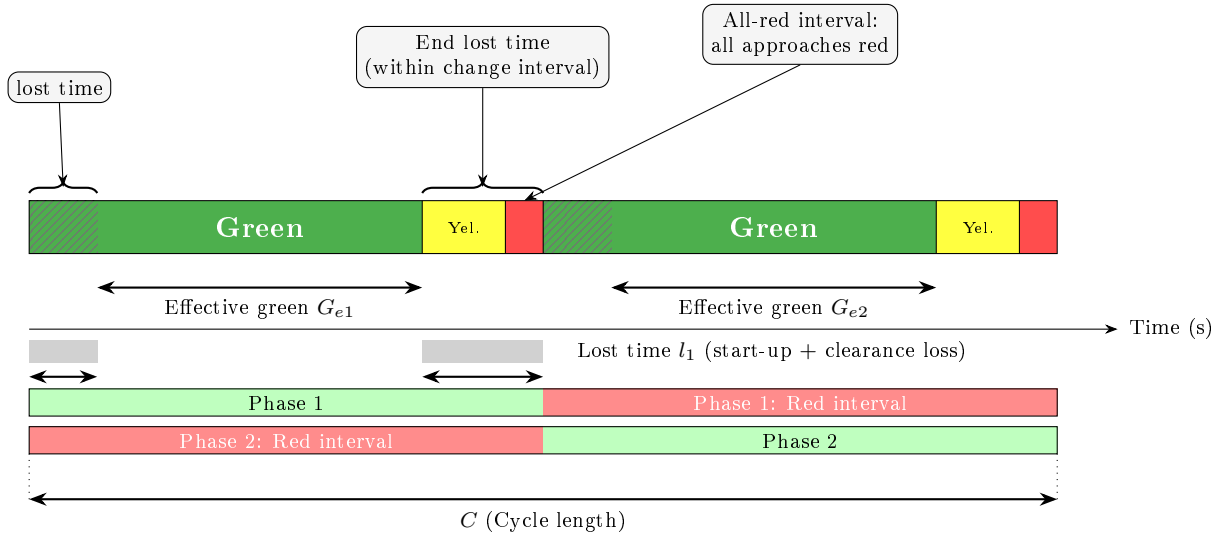


Figure 4.1: Signal timing for a two phase cycle, showing the effective green, change interval, and start up and clearance lost time.

4.2 Traffic Flow Representation for Signal Analysis

In signalized intersection analysis, traffic demand is often expressed in terms of through passenger cars. This convention is widely used because passenger cars travelling straight through an approach provide a simple and consistent reference for estimating capacity, saturation flow, and control delay [14]. If all vehicles behaved like passenger cars and all movements were through movements, the measured approach flow could be used directly in the core computations of signal design, such as determining the volume-to-capacity ratio, allocating green time among competing movements, and selecting an appropriate cycle length.

In practice, however, real traffic streams are heterogeneous. Approaches typically carry a mixture of passenger cars and heavier vehicles (e.g., buses and trucks) whose size, acceleration, and space requirements reduce discharge performance during the green interval. Moreover, a portion of the approach demand executes turning movements. Left and right turning vehicles can reduce the effective discharge rate because they may need to slow down, accept gaps, yield to pedestrians or conflicting flows, and follow tighter turning paths. As a result, the same measured flow rate (vehicles per hour) can impose very different operational demands depending on the vehicle composition and the share of turning traffic. It is therefore not sufficient to treat all vehicles and movements as identical when the goal is to obtain realistic estimates of capacity and delay.

To translate mixed traffic into a single comparable demand measure, signal analysis converts each traffic component into an equivalent passenger car flow. Two adjustment concepts are used. First, *passenger car units* (PCU) convert each non-passenger vehicle into the number of passenger cars that would produce approximately the same operational effect under similar conditions; a heavy truck, for example, may be represented by multiple passenger car units because it accelerates more slowly and occupies more effective space in the traffic stream. Second, *turning movement factors* account for the additional operational impact of turning vehicles compared with through vehicles on the same approach, representing how turning manoeuvres, yielding behaviour, and speed reductions alter the effective demand placed on the green time.

With these adjustments, an equivalent flow rate is defined as

$$q_{\text{eq}} = \sum_{j} q_j \, \text{PCU}_j \, f_j, \tag{4.1}$$

where q_j is the observed flow of traffic component j, PCU_j is the passenger car unit value assigned to the corresponding vehicle type, and f_j is the turning movement factor representing the operational effect of the corresponding movement (through, left turn, or right turn).

Conceptually, Eq. (4.1) states that each observed flow component is first converted into passenger car units, then adjusted to reflect movement specific effects, and finally aggregated to produce a single equivalent demand measure q_{eq} for the approach or movement group being analysed.

Example. Consider passenger cars P and trucks T, and three movements: through S, left L, and right R. Let the conversion factors be $\text{PCU}_P = 1.0$ and $\text{PCU}_T = 2.0$, and let the movement factors be $f_S = 1.0$, $f_L = 1.3$, and $f_R = 0.9$. Suppose the observed flows (veh/h) are

- through: $q_{P,S} = 600, \quad q_{T,S} = 80$;
- left: $q_{P,L} = 200, \quad q_{T,L} = 40$;
- right: $q_{P,R} = 150, \quad q_{T,R} = 20$.

The equivalent flow is then

$$q_{\text{eq}} = \underbrace{\big(600 \cdot 1.0 + 80 \cdot 2.0\big) \cdot 1.0}_{\text{through } (S)} + \underbrace{\big(200 \cdot 1.0 + 40 \cdot 2.0\big) \cdot 1.3}_{\text{left } (L)} + \underbrace{\big(150 \cdot 1.0 + 20 \cdot 2.0\big) \cdot 0.9}_{\text{right } (R)}. \tag{4.2}$$

Evaluating each term gives 760 PCU/h for the through movement, 364 PCU/h for the left turn, and 171 PCU/h for the right turn, so that $q_{\text{eq}} = 760 + 364 + 171 = 1295$ PCU/h.

The quantity q_{eq} is then used as the demand input for downstream signal calculations. Because it reflects both vehicle composition and turning behaviour, it provides a realistic basis for estimating how quickly vehicles can be served during green, how green time should be distributed among movements, and how cycle length choices influence the resulting delay and degree of saturation.

The following sections introduce the key timing specifications of traffic signals that govern how vehicles transition through an intersection, beginning with the yellow interval and clearance time and their influence on drivers' decision-making.

4.3 Yellow Interval and the Dilemma Zone

When the signal indication changes from green to yellow, approaching drivers must decide whether to stop or proceed through the intersection. The purpose of the yellow interval, and in many cases an additional all-red interval, is to provide a safe transition between conflicting phases. If the provided clearance time is insufficient, a region may exist in which drivers can

neither stop comfortably nor clear the intersection before the onset of red. This region is commonly referred to as the *dilemma zone* [6]. Its formal definition requires two preliminary concepts.

Stopping Requirement

Assume a vehicle approaches the stop line with speed u_0 at the onset of yellow, as shown in Figure 4.2. The driver experiences a perception–reaction time δ before initiating braking, and a comfortable deceleration rate a is assumed during braking.

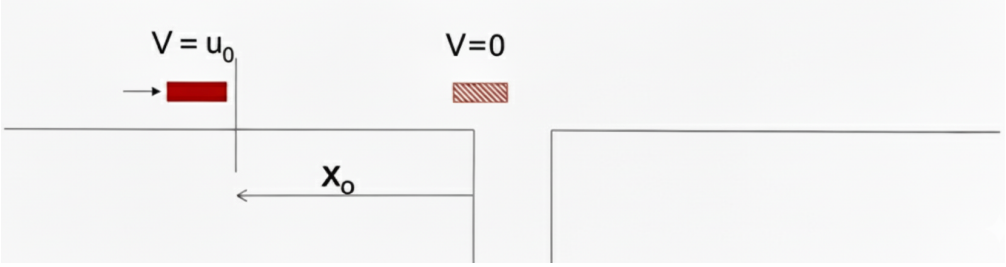


Figure 4.2: Stopping Requirement

Setting $v = u_0 - at = 0$ gives the braking time

$$t_{\text{brake}} = \frac{u_0}{a}, \tag{4.3}$$

and the corresponding braking distance

$$x_{\text{brake}} = \frac{u_0^2}{2a}. \tag{4.4}$$

The distance travelled during perception and reaction is

$$x_{\text{react}} = u_0 \delta. \tag{4.5}$$

Therefore, the minimum distance from the stop line required to stop safely and comfortably is

$$X_0 = u_0 \delta + \frac{u_0^2}{2a}. \tag{4.6}$$

4.3.1 Go Requirement

To proceed safely, the vehicle must be able to clear the conflict area before the red indication begins. Let W denote the width of the intersection (or conflict zone) and L_v the vehicle

length, as illustrated in Figure 4.3. If the minimum available clearance time is $\tau_{\min}$, a vehicle located at distance X_c from the stop line can clear the intersection if

$$u_0 \tau_{\min} \geq X_c + (W + L_v), \tag{4.7}$$

which yields the maximum clearing distance

$$X_c = u_0 \tau_{\min} - (W + L_v). \tag{4.8}$$

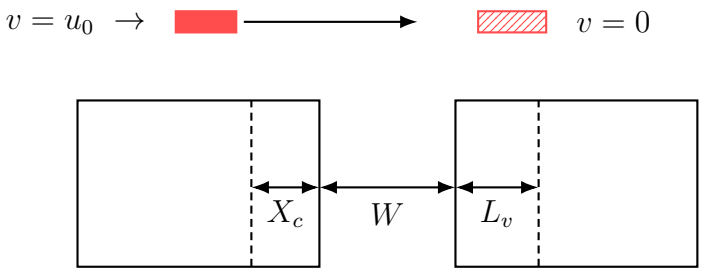


Figure 4.3: Go requirement: a vehicle at distance X_c from the stop line must traverse $X_c + W + L_v$ within the available clearance time $\tau_{\min}$ to clear the conflict zone of width W.

4.3.2 The Dilemma Zone

The dilemma zone is the region upstream of a signal in which a driver, at the onset of yellow, may be unable to make a safe and comfortable decision: stopping may require an uncomfortably high deceleration, while continuing may risk entering the intersection after the red begins. Using the definitions above, the dilemma zone exists when

$$X_c < X_0, \tag{4.9}$$

meaning that there is an interval of distances in which a driver is simultaneously (i) too close to stop comfortably, yet (ii) too far to clear the intersection before the end of yellow (and any clearance interval considered). The situation is illustrated in Figure 4.4. From an operational and safety perspective, this region is undesirable because it increases the probability of harsh braking, rear end conflicts, and red light violations.

A common design objective is therefore to eliminate the dilemma zone by selecting signal timing such that the must go and must stop boundaries coincide:

$$X_0 = X_c. \tag{4.10}$$

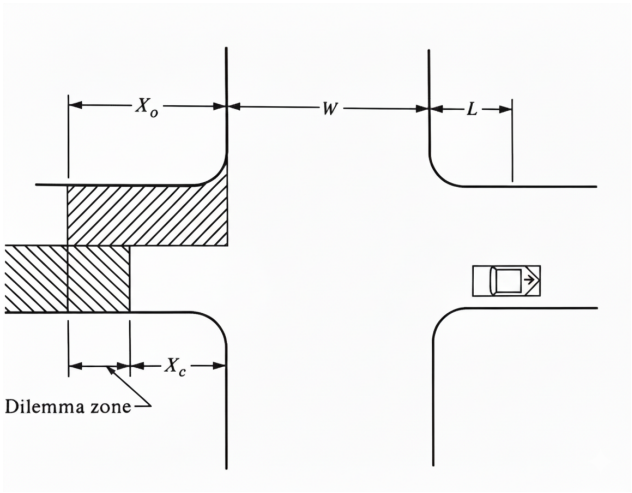


Figure 4.4: dilemma zone

Substituting Eqs. (4.6) and (4.8) into Eq. (4.10) yields the timing condition that removes the ambiguous region,

$$u_0\tau_{\min} - (W + L_v) = u_0\delta + \frac{u_0^2}{2a}, \tag{4.11}$$

and solving for the minimum clearance time required to eliminate the dilemma zone gives

$$\tau_{\min} = \delta + \frac{u_0}{2a} + \frac{W + L_v}{u_0}. \tag{4.12}$$

This expression has a clear physical interpretation: δ is the perception reaction component, $u_0/(2a)$ captures the kinematic contribution associated with the chosen comfortable deceleration level (a higher approach speed or a lower deceleration capability increases the required time), and $(W + L_v)/u_0$ is the time needed to traverse the intersection width plus one vehicle length at speed u_0.

For practical design, a is chosen to represent a comfortable deceleration rather than an emergency manoeuvre, so that the resulting timing does not force frequent harsh braking. A commonly adopted value is

$$a = 0.27\, g, \tag{4.13}$$

where g is the gravitational acceleration; expressing the deceleration as a fraction of g provides a convenient normalisation and reflects a moderate, comfortable braking level for typical passenger-car operation.

4.4 Cycle Length Determination: The Webster Method

One of the most widely used analytical approaches for determining the optimal cycle length of isolated, pretimed signalized intersections is the Webster method [17, 33]. The method was developed to minimise the average delay experienced by vehicles and is particularly suitable for intersections with relatively stable demand patterns.

The optimal cycle length, denoted by C_0, is given by

$$C_0 = \frac{1.5\,L + 5}{1 - \sum_{i=1}^{\phi} Y_i}, \tag{4.14}$$

where L is the total lost time per cycle (s), including startup lost times and all red intervals, and ϕ is the number of phases. Values reported in practice cluster around three to four seconds of lost time per phase, plus the all-red time. The spread reflects what is being counted: the start-up component depends on how quickly a queue reaches saturation flow, which varies with grade, vehicle mix and queue length, while the clearance component depends on how much of the yellow is actually used. The formal definition is developed in Section 4.5, and the value adopted for the case study is derived in Section 5.3.1 rather than taken from a published figure.

The constants 1.5 and 5 are empirical. Webster obtained them by fitting a curve to the cycle length that minimised total delay across a large set of simulated intersections, so Eq. (4.14) is an approximation rather than a closed-form optimum. A property of that fit is useful in design: the delay curve is flat near its minimum, and any cycle between roughly $0.75\,C_0$ and $1.5\,C_0$ gives delay within ten to twenty per cent of the optimal value. The exact value of C_0 is therefore less critical than its order of magnitude, which matters where operational minima or pedestrian crossing requirements raise the cycle above the computed optimum.

The parameter Y_i is the critical flow ratio for phase i, defined as the maximum ratio of approach flow to saturation flow among all traffic streams served during that phase:

$$Y_i = \max_{j} \frac{V_{ij}}{s_j}, \tag{4.15}$$

where V_{ij} is the flow on lane or movement j receiving the right-of-way during phase i, and s_j is the corresponding saturation flow. Only the critical movement of each phase enters the calculation, reflecting the fact that intersection performance is constrained by the most heavily loaded movement in each phase. Because Y_i is a ratio, the resulting design is insensitive to the absolute level of demand and depends only on how that demand is distributed across

the phases.

The denominator of Eq. (4.14) bounds the applicability of the method. As $\sum Y_i$ approaches unity the computed cycle diverges, reaching 170 s at $\sum Y_i = 0.90$ and over 1700 s at 0.99, and at $\sum Y_i = 1$ no solution exists. This is not a numerical artefact but a statement about the derivation: Webster assumes that each queue clears within the cycle it is served in, which is what makes a delay-minimising cycle well defined. Once demand exceeds capacity the queue does not clear, delay grows without bound regardless of the timing, and the design objective shifts from minimising delay to maximising throughput and managing which queues are permitted to grow. That regime is discussed in Section 4.6.1.

4.4.1 Interpretation of the Critical Flow Sum

The sum $\sum_i Y_i$ represents the overall degree of saturation of the intersection. As this value approaches unity, the intersection approaches oversaturated conditions, and the optimal cycle length in Eq. (4.14) increases rapidly. If $\sum_i Y_i \geq 1$, the Webster formula is no longer applicable, indicating that intersection demand exceeds the available capacity.

4.4.2 Saturation Flow

The saturation flow s_j represents the maximum discharge rate of movement j under ideal conditions. For a single through lane, a commonly used value is approximately 1800 passenger car units per hour per lane [14],

although this value varies with geometric and traffic conditions.

For practical implementation, the computed optimal cycle length is typically rounded to the nearest 5 seconds to facilitate signal operation and coordination. Overall, the Webster method provides a simple yet powerful analytical framework for cycle length determination and serves in this thesis as the baseline against which the coordinated control strategy is compared.

4.5 Lost Time at Signalized Intersections

At a signalized intersection, not all of the nominal green time contributes effectively to vehicle discharge. At the beginning of the green phase, startup delays occur due to driver reaction time and vehicle acceleration; similarly, near the end of the phase, part of the green and most of the yellow interval do not result in effective discharge. The total duration of these ineffective periods is referred to as *lost time*.

When the signal turns green, the discharge rate increases gradually until it reaches an approximately constant value known as the saturation flow, denoted by s (veh/s). The portion of the green phase during which the discharge rate remains close to s is called the effective green time, denoted by G_e (s). The number of vehicles discharged during the phase is given by the area under the discharge curve, approximated as

$$\text{discharged vehicles} \approx s\, G_e. \tag{4.16}$$

The actual green time for phase i, denoted $G_{a,i}$, is the green that drivers see on the signal head, from the moment it turns green until it switches to yellow. In real traffic, the first few seconds of green are not used at full efficiency, because drivers need a moment to react and vehicles need time to accelerate, so the queue does not discharge at the saturation rate immediately. The effective green time, $G_{e,i}$, is the part of the phase during which vehicles flow at, or close to, the saturation rate; this is the quantity used in capacity and delay calculations. In many cases, part of the yellow time still serves vehicles that can safely pass the stop line, and is therefore counted as part of the effective service time.

4.5.1 Lost Time for a Phase

The lost time for phase i is defined as the difference between the nominal time allocated to the phase and its effective green time:

$$l_i = G_{a,i} + \tau_i - G_{e,i}, \tag{4.17}$$

where $G_{a,i}$ is the actual green time excluding yellow, τ_i is the yellow time, and $G_{e,i}$ is the effective green time for phase i.

4.5.2 Total Lost Time per Cycle

For a signal cycle consisting of ϕ phases, the total lost time is

$$L = \sum_{i=1}^{\phi} l_i + R, \tag{4.18}$$

where R denotes the total all-red time during the cycle. Lost time is a critical parameter in signal timing design: it directly reduces effective capacity and plays a central role in cycle length optimisation through Eq. (4.14).

Figure 4.5 shows how vehicles are served during one signal phase. Immediately after the

light turns green, the discharge rate starts low as drivers react and vehicles accelerate, so flow builds up gradually. After a short startup period, the discharge rate becomes nearly constant at the saturation flow s. The period during which the flow stays close to s is the effective green G_e, which is the part of the phase that genuinely provides capacity. The early seconds of green and the transition at the end of the phase (yellow and all-red) are not fully productive and are treated as lost time. The total number of vehicles served corresponds to the area under the curve and is approximated by $s\,G_e$, as in Eq. (4.16).

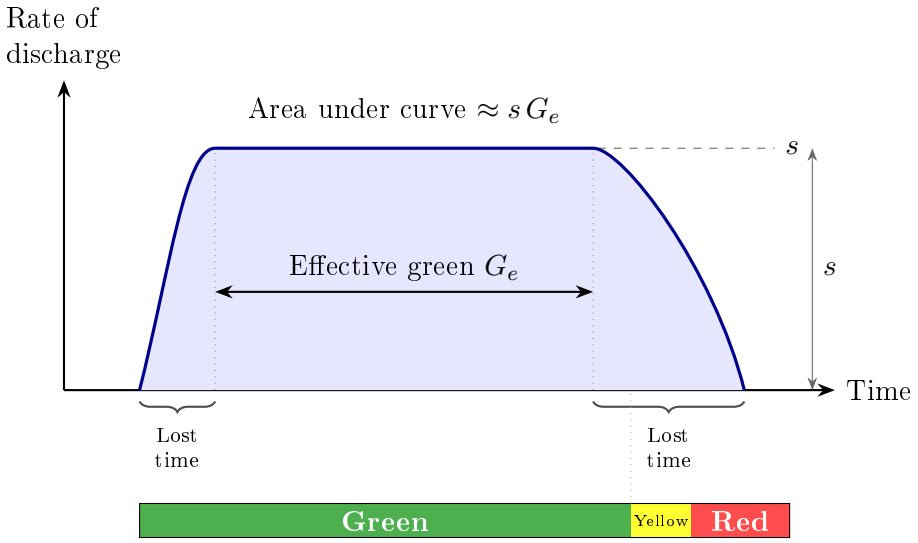


Figure 4.5: Discharge rate during one signal phase, with effective green G_e, saturation flow s, and lost time.

4.5.3 Allocation of Effective Green Time

Once the optimal cycle length has been determined using the Webster method, the next step in pretimed signal design is the allocation of effective green time among the signal phases, with the objective of distributing the available capacity so as to minimise overall vehicle delay.

Total Effective Green Time per Cycle

Only a portion of the signal cycle contributes to effective vehicle discharge. The total effective green time available per cycle is obtained by subtracting the total lost time from the cycle length:

$$G_{te} = C - L = C - \left(\sum_{i=1}^{\phi} l_i + R \right), \tag{4.19}$$

where C is the cycle length used in operation, typically obtained by rounding the optimal cycle length C_0 to the nearest 5 seconds.

Proportional Allocation of Effective Green Time

To minimise overall delay, the total effective green time is allocated among the phases in proportion to their critical flow ratios. The effective green time for phase i is

$$G_{e,i} = \frac{Y_i}{\sum_{m=1}^{\phi} Y_m} G_{te}. \tag{4.20}$$

This formulation ensures that phases serving higher critical demand receive a larger share of the available effective green time, while the effective green times sum to G_{te} over all phases.

4.5.4 Pretimed and Actuated Signal Control

Traffic signal control strategies are generally classified into pretimed and actuated systems. The selection between these approaches depends on traffic demand variability, operational objectives, and intersection characteristics.

Pretimed Signals

Pretimed signals operate with fixed cycle lengths, phase sequences, and interval timings, designed on the basis of typical or average historical traffic demand. Signal timing remains constant regardless of short-term fluctuations in traffic volumes. While pretimed control is simple and predictable, it may lead to inefficient operation when demand varies over time: green time is wasted during low-demand periods, and cycle failure can occur when traffic volumes exceed the capacity of a single cycle.

Actuated Signals

Actuated signals address demand fluctuations by means of traffic detectors installed on one or more approaches. Green time allocation is based on detected demand rather than on predetermined averages: green durations vary dynamically between specified minimum and maximum limits, allowing the signal to respond to real-time traffic conditions.

4.5.5 Signal Timing Parameters for Actuated Control

In actuated signal control systems, performance is strongly influenced by the selection of timing parameters and the placement of detectors [34]. These elements determine whether

queued vehicles and pedestrians can be served safely and efficiently.

Maximum Green Time

The maximum green time establishes an upper bound on the duration of a green phase, even when continuous demand is present. Once this limit is reached and a call exists on a conflicting phase, the green phase is terminated, a condition known as *max out*. In semi-actuated control, the maximum green time is often selected by judgement; in fully actuated control, it is commonly set so that the total green time is consistent with the optimal cycle length obtained from the Webster method under pretimed operation.

Minimum Green Time

The minimum green time is the initial portion of the green phase that must be fully served regardless of additional demand. It ensures that vehicles stored between the detector and the stop line are discharged and that minimum pedestrian crossing requirements are satisfied. It can be estimated as

$$G_{\min} = 4 + 2 \left\lfloor \frac{d_{\text{det}}}{6} \right\rfloor, \tag{4.21}$$

where $G_{\min}$ is the minimum green time (s), d_{det} is the distance between the detector and the stop line (m), 6 m is the assumed spacing between queued vehicles, 4 s is the assumed startup lost time, and 2 s is the assumed saturation headway. Equation (4.21) indicates that detectors placed farther upstream require longer minimum green times to discharge the stored queue.

Passage Time and Allowable Gap

Following the minimum green period, green extension depends on vehicle arrivals. The passage time is

$$P = \frac{d_{\text{det}}}{v_a}, \tag{4.22}$$

where P is the passage time (s), d_{det} is the detector setback distance (m), and v_a is the approach speed (m/s). If no vehicle is detected within one passage time after the previous actuation, the green phase terminates, a condition known as *gap out*.

4.6 Signal Coordination in Emergency Evacuation

In emergency evacuation scenarios, traffic demand surges abruptly, producing *oversaturated* traffic conditions. Under such circumstances, traditional signal timing methods, including

the Webster formula, may fail to provide an optimal evacuation strategy: while the Webster method is effective in undersaturated conditions, where it minimises delay, it does not account for the continuous queue formation and spillback effects typical of mass exits. To prevent gridlock and maximise the throughput of the evacuation network, *signal coordination* (green waves) becomes essential [30, 31].

The primary objective of signal coordination along an evacuation corridor is to synchronise consecutive intersections so that a large platoon of evacuating vehicles can traverse the network with minimal interruption. This is achieved by calculating the *offset*: the time difference between the start of the green phases at adjacent intersections.

For uniform progression along a one-way evacuation route, the ideal offset is primarily a function of the distance between intersections and the target evacuation speed:

$$t_{\text{offset}}^{0} = \frac{d_{\text{int}}}{v_p}, \tag{4.23}$$

where d_{int} is the inter-signal spacing (m) and v_p is the desired progression speed during evacuation (m/s).

In an evacuation, however, intersections are rarely empty when a new platoon arrives: a standing queue is often present at the downstream stop line due to local traffic or prior congestion. To ensure the arriving platoon does not meet a stationary wall of vehicles, the downstream signal must turn green earlier, so that the queue is discharged before the platoon arrives. The adjusted offset is

$$t_{\text{offset}} = \frac{d_{\text{int}}}{v_p} - \left(Q\,h + l_s\right), \tag{4.24}$$

where Q is the number of queued vehicles per lane, h is the discharge headway (typically about 2 s/veh), and l_s is the startup lost time (s).

4.6.1 Coordination in Oversaturated Networks

When the network reaches oversaturation, the focus shifts from minimising delay to maximizing flow and managing queues. Coordination contributes in three ways:

1. **Preventing de facto red:** by clearing the downstream queue before the upstream platoon arrives, coordination prevents vehicles from being blocked by a full link even while their signal indication is green.

2. **Maintaining platoon integrity:** coordination keeps evacuating vehicles together, which reduces speed variance and lowers the probability of secondary collisions.

3. **Optimising throughput:** by aligning the effective green time across multiple intersections, the corridor operates as a single high-capacity pipeline rather than as a series of disconnected bottlenecks [32].

4.7 Cost-Based Evaluation of Signal Coordination Benefits

The benefits of traffic signal coordination are commonly assessed using a cost (or penalty) function that aggregates corridor-level inefficiencies into a single quantitative objective. The fundamental aim of coordination is to reduce the disbenefits associated with interrupted progression, such as frequent stops, excessive control delay, and unstable speed profiles. A widely used representation expresses total cost as a weighted combination of key performance indicators,

$$\text{Cost} = A\,N_{\text{stops}} + B\,D_{\text{total}} + \text{other terms}, \tag{4.25}$$

where N_{stops} is the total number of stops, D_{total} is the total delay, and A and B are user-specified coefficients reflecting the relative importance assigned to stops and delay. Depending on operational objectives, the additional terms may include travel time, fuel or energy consumption, emissions, speed variability, or reliability indicators.

From an operational perspective, well-designed offsets and coordinated timing plans reduce stops by aligning green initiations with expected platoon arrivals, thereby improving progression and lowering average control delay. By reducing stop-and-go driving, coordination also helps maintain a preferred operating speed along the corridor, improves perceived service quality, and generates secondary benefits such as reduced fuel consumption and emissions. These benefits are typically most pronounced under moderate demand levels (i.e., when intersections are not persistently oversaturated) and when platoon structures are stable enough for arrivals to remain within the intended green windows.

However, coordination effectiveness is not guaranteed, and several practical factors can diminish it. A primary limiting condition is inadequate capacity relative to demand: as demand approaches or exceeds capacity, queues form and spill back, and arrivals become governed by upstream discharge constraints rather than by the planned offset structure, making continuous progression difficult to maintain. Roadside frictions such as parking activity, bus stops, pedestrian crossings, frequent driveway access, and other lateral interactions further disrupt uniform flow, reduce effective speeds, and increase speed variability, weakening the assumption of a consistent progression speed. Coordination is also challenged in complex intersections and two-dimensional grid networks, where improving progression on one street

may conflict with the progression needs of cross streets, and where heterogeneous turning movements and phase patterns reduce the feasibility of stable platoon progression over long distances. Finally, speed variability is critical, because offsets rely on a reasonable concentration of vehicles travelling near the reference coordination speed; wide speed dispersion increases arrival dispersion, reduces the effective bandwidth, and raises the probability of stopping at downstream signals.

Building on the theoretical foundations of signal coordination and the challenges identified for oversaturated networks, the following chapter presents the simulation study. The operational efficiency of the coordinated signal strategy is evaluated against the Webster baseline for the emergency evacuation of the Sandy Hill network, with performance assessed through network travel time, corridor speed, throughput, and destination-specific travel times.

Chapter 5

Emergency Evacuation Simulation

The theoretical foundations developed in the preceding chapters establish the analytical basis for traffic flow modelling, static and dynamic assignment, and signal timing design. The present chapter translates those principles into a practical simulation framework applied to a real urban network in Ottawa, Canada. The central objective is to evaluate the operational effectiveness of two fixed time signal control strategies under emergency evacuation conditions: a *Webster* plan, which serves as the analytical baseline, and a *Coordinated Green Wave* plan that incorporates progressive intersection offsets and permissive right turn movements.

The Sandy Hill network encompasses 17 signalized intersections organised around a two branch evacuation corridor along King Edward Avenue, which bifurcates at Laurier Avenue: a northern branch directed toward the Macdonald–Cartier Bridge (Gatineau), and a southern branch directed toward Highway 417. Seven further signalized junctions on the crossing Rideau Street and Laurier Avenue alignments complete the network. As shown in Figure 5.1, the two branches carry the evacuating traffic toward the network's principal egress points. The network serves a census derived origin–destination (OD) demand of 3,530 vehicles over a two hour loading horizon, the derivation of which is presented in Section 5.1. Microsimulation was conducted in Aimsun Next [42], with performance metrics aggregated at 10 minute intervals across an 18 interval (180 minute) simulation window comprising the 120 minute demand loading period and a 60 minute clearance tail; intervals beyond 150 minutes, during which the network operates at free flow, are omitted from the result tables.

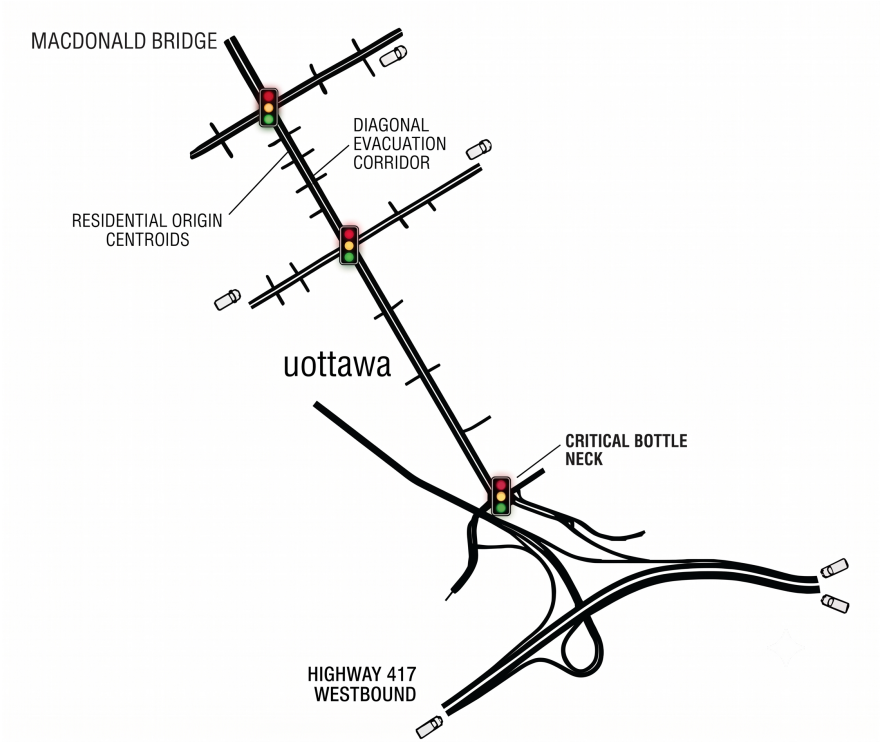


Figure 5.1: The Sandy Hill study network in Ottawa, showing the corridor toward Highway 417, the corridor toward the Macdonald–Cartier Bridge, and the signalised junctions of the evacuation network.

The seventeen signalised junctions of the network are listed in Table 5.1 with the identifiers used throughout this chapter and in the result tables. Ten of them lie on the King Edward evacuation corridor and carry the coordinated timing; the remaining seven are off-corridor and are identical in both signal plans.

Network representation. The network is held as a directed graph, with sections as arcs and the turning movements at each junction as the connections between them. This structure is built and maintained by Aimsun Next rather than specified by hand: importing the road geometry creates the sections, and the turning movements follow from the lane connectivity defined at each junction. Travel on an arc is governed by the microscopic car-following and lane-changing models rather than by a link cost function, and the permitted movements through a junction are those the turning definitions allow, restricted further by whichever signal group holds green.

Two consequences follow for how the results should be read. First, the connectivity used in the simulation is the connectivity of the imported geometry, so a movement that is physically possible on the street is available in the model unless a turn was explicitly removed. Second, because the same graph object serves both scenarios, the two signal plans operate on an

identical network: the arcs, the turning movements and the lane connectivity are shared, and only the signal group assigned to each phase differs between them.

Table 5.1: The seventeen signalised junctions, with the identifiers used in the text. Ten lie on the King Edward corridor and are retimed to the common 55 s cycle; the remaining seven are identical in both plans.

ID	Intersection	Webster cycle (s)
King Edward Avenue corridor, north to south		
J2710	Rideau Street and King Edward Avenue	55.0
J2844	Besserer Street and King Edward Avenue	40.0
J2678	Daly Avenue and King Edward Avenue	40.0
J2669	Stewart Street and King Edward Avenue	40.0
J2614	Wilbrod Street and King Edward Avenue	40.0
J2652	Laurier Avenue East and King Edward Avenue	43.4
J2639	Osgoode Street and King Edward Avenue	40.0
J2625	Somerset Street East and King Edward Avenue	40.0
J2914	Templeton Street and King Edward Avenue	40.0
J3307	Lees Avenue and King Edward Avenue	43.4
Off-corridor junctions		
J3095	Rideau Street and Friel Street	40.0
J3082	Rideau Street and Nelson Street	40.4
J2725	Rideau Street and Cumberland Street	40.4
J3122	Laurier Avenue East and Sweetland Avenue	40.0
J3423	Laurier Avenue East and Nelson Street	40.0
J2596	Laurier Avenue East and Copernicus Street	40.0
J3331	Laurier Avenue East and Cumberland Street	40.0

5.1 Origin–Destination Demand

A central methodological requirement of any credible evacuation simulation is that the total traffic demand be traceable to a documented, empirical population base rather than assigned by assumption. The present section establishes the total evacuating vehicle count denoted `vehicle` and constituting the key target variable through a four-step derivation grounded in 2021 Statistics Canada census data, the 2022 Ottawa Origin–Destination Survey, and published evacuation behaviour literature. The resulting count is then distributed temporally using the Rayleigh departure time model, whose theoretical properties are uniquely suited to mandatory near area evacuations.

Census-Based Population and Household Count

The 2011 Census of Population recorded a Sandy Hill residential population of 12,490 persons [36]. Applying the inner city Ottawa inter censal growth rate of approximately 5% over the 2011–2021 period the city wide rate was 8.9%, but transit served, high density inner city neighbourhoods consistently exhibit lower population growth than the municipal average [37] the 2021 Census-consistent residential population is estimated as

$$P = 12{,}490 \times 1.05 = 13{,}114.5 \approx 13{,}100 \text{ residents}, \tag{5.1}$$

consistent with the Environics Analytics neighbourhood profile derived from Statistics Canada 2021 dissemination area data [39].

Two points of this step should be made explicit. First, the 5% figure is an assumed rate, selected as roughly half of the 8.9% city wide 2011–2021 growth, on the reasoning that built out inner city neighbourhoods with a largely fixed housing stock grow more slowly than a municipality that is still expanding at its suburban edge. It is not measured for Sandy Hill. Second, the extrapolation is used because the neighbourhood boundary does not coincide with a census subdivision, so a direct 2021 count for the area is not published; a boundary consistent 2021 figure would have to be aggregated from the individual dissemination areas that make up the neighbourhood, which is identified in Section 6.3 as the preferred refinement. The effect of the assumption is bounded and close to linear: adopting the full city wide 8.9% instead would raise P to about 13,600 and, through the same chain, the evacuating fleet from 3,530 to roughly 3,660 vehicles, a difference of about 3.7%, which is smaller than the difference between the compliance bounds considered in Eq. (5.5).

The 2021 Census Profile for the Ottawa Census Subdivision reports an average household size for inner city census tracts of $\bar{s} = 2.1$ persons per private household [38]. Sandy Hill's demographic composition, including University of Ottawa students, embassy adjacent families, co-operative housing tenants, and subsidised housing residents, is structurally consistent with this inner city figure; the city wide Ottawa average of 2.4 persons per household, applicable to the full municipality, is inappropriate for this zone.

The total number of occupied private dwellings in Sandy Hill is therefore

$$H = \left\lfloor \frac{P}{\bar{s}} \right\rfloor = \left\lfloor \frac{13{,}100}{2.1} \right\rfloor = 6{,}238 \text{ households}. \tag{5.2}$$

Vehicle Ownership Rate per Household

The 2022 Origin-Destination Survey conducted by the TRANS Committee for the National Capital Region [40] reported an Ottawa wide average of 1.33 vehicles per household. The

downtown core and inner urban transects, however, exhibit materially lower ownership rates due to high transit accessibility, high student population fractions, and the prevalence of subsidised and co-operative housing. Consistent with the 2022 O-D Survey inner urban transect vehicle ownership distribution [40, 41], Sandy Hill is modelled using the following three category ownership decomposition:
The provenance of the three shares in Table 5.2 should be stated plainly. They are not published Sandy Hill values. They are a representative decomposition for an inner urban, transit accessible, student dominated neighbourhood, chosen so that the resulting average of 0.755 vehicles per household is materially below the Ottawa wide average of 1.33 reported by the 2022 O–D Survey [40], in the direction that the survey itself reports for downtown core households. The representative value of 1.7 vehicles assigned to the "two or more" category is likewise an assumption, chosen because a category that is open ended above two is dominated in practice by two vehicle households. The 35% car free fraction is supported by the documented prevalence of student rental and co-operative housing in the neighbourhood.

Table 5.2: Vehicle ownership distribution for Sandy Hill (inner urban transect, 2022 O-D Survey basis).

Ownership category	**Household share π_j**	**Vehicles v_j**	**Contribution $\pi_j v_j$**
Zero vehicles	0.35	0.0	0.0000
One vehicle	0.50	1.0	0.5000
Two or more	0.15	1.7	0.2550
Total	1.00		**0.7550**

Because the decomposition is assumed rather than measured, its influence should be judged by the aggregate it produces. Only the weighted mean $\overline{V}_{\text{hh}}$ enters the demand chain, and the fleet size scales linearly with it: a mean of 0.70 would give about 3,275 evacuating vehicles and a mean of 0.85 about 3,975, against the 3,530 adopted here. Replacing this decomposition with the published vehicle availability cross tabulation for the dissemination areas covering Sandy Hill would remove the assumption entirely and is identified as a refinement in Section 6.3. The weighted average number of vehicles per household is

$$\overline{V}_{\text{hh}} = \sum_{j} \pi_j \, v_j = 0.35(0) + 0.50(1.0) + 0.15(1.7) = 0.7550 \text{ vehicles/household}, \tag{5.3}$$

and the total vehicle pool available for evacuation is

$$V_{\text{pool}} = H \times \overline{V}_{\text{hh}} = 6{,}238 \times 0.7550 = 4{,}710 \text{ vehicles}. \tag{5.4}$$

Evacuation Compliance Rate

The compliance rate η represents the fraction of vehicle owning households whose members depart during the active two-hour loading horizon in response to a mandatory evacuation order. The evacuation behaviour literature consistently reports mandatory order compliance rates in the range 0.70–0.90 for urban North American events, with a modal value near 0.75 for events without extreme, immediate life safety threat [35, 21, 20]. The present study adopts a baseline compliance rate of

$$\eta = 0.75, \tag{5.5}$$

with sensitivity cases at $\eta = 0.60$ (lower-bound) and $\eta = 0.90$ (upper-bound) identified for future analysis as recommended in Section 6.3.

Two behavioural assumptions are embedded in this step and should be stated. The first is that a complying household evacuates with all of the vehicles it owns. This is the conservative choice for a network study, since it maximises the vehicle count and therefore the load on the corridor, and it is consistent with evidence that households in no notice evacuations tend to take multiple vehicles in order to keep the household together and to protect the vehicles themselves [21]. It is nevertheless an upper bound: if the two or more vehicle households, which contribute 0.255 of the 0.755 average, were to evacuate with one vehicle each, the fleet would fall to roughly 2,810 vehicles, about 20% below the value adopted.

The second assumption is that the 35% of households without a vehicle do not contribute to vehicular demand. In a student and transit dependent neighbourhood these residents are precisely the population least able to self evacuate, and in practice they would rely on transit, on ride sharing with neighbours, or on assisted evacuation [25, 1]. The present model therefore represents the vehicular evacuation only, and the corridor loading reported here should be read as excluding whatever additional demand a bus based or assisted evacuation of the car free population would place on the same signalised junctions. Extending the model to this multimodal case is identified in Section 6.3.

Total Evacuating Vehicle Count

Combining Equations (5.4) and (5.5), the total number of evacuating vehicles is

$$\texttt{vehicle} = \Big\lfloor V_{\text{pool}} \times \eta \Big\rceil_5 = \Big\lfloor 4{,}710 \times 0.75 \Big\rceil_5 = \Big\lfloor 3{,}532.3 \Big\rceil_5 = \boxed{3{,}530}, \tag{5.6}$$

where $\lfloor \cdot \rceil_5$ denotes rounding to the nearest multiple of five, consistent with the resolution of the Aimsun traffic demand profile entry. The complete derivation chain is collected in Table 5.3.

This figure is the base level of the analysis rather than its only level. Because the central claim of Chapter 5 concerns how the benefit of coordination varies with load, the whole experiment is repeated at 1.25, 1.5 and 2 times this demand, holding the departure profile, the network and the signal plans fixed. The derivation above therefore fixes the shape of the demand and the reference total; the scaling factors are applied to that total, and the four levels are compared in Section 5.4.8. The detailed analysis of intervals, sections and destinations is carried out at the base level; the higher levels are compared at the network level.

Table 5.3: Census-based demand derivation chain for `vehicle`.

Parameter	**Source / Basis**	**Value**
Residential population P	Stat. Canada 2021; Environics Analytics [39]	13,100 res.
Avg. household size $\bar{s}$	StatCan 2021, Ottawa inner-city CTs [38]	2.1 per/hh
Total households H	Eq. (5.2)	6,238 hh
Veh. per household $\overline{V}_{\text{hh}}$	Ottawa 2022 O-D Survey, inner-urban [40]	0.755 veh/hh
Vehicle pool V_{pool}	Eq. (5.4)	4,710 veh
Compliance rate η	Wolshon [35]; Pel et al. [20]	0.75
Total evacuating vehicles `vehicle`	**Eq.** (5.6)	**3,530 veh**

Rayleigh Departure Time Distribution

Model selection rationale. The temporal distribution of evacuation departures governs the evolution of queue length, degree of saturation, and signal performance at every junction. Two classes of models are commonly employed: (i) uniform loading, which assumes a constant departure rate over the horizon, and (ii) empirically calibrated departure distributions. Uniform loading is inappropriate for evacuation modelling because it suppresses the surge character of mass departures and systematically underestimates peak-interval queuing and saturation. The Rayleigh distribution is selected instead for three reasons. First, it is zero-valued at $t = 0$, ensuring no departures occur before issuance of the evacuation warning. Second, it rises steeply to a single unimodal peak and then decays smoothly as the compliant population is progressively depleted, behaviour that is qualitatively consistent with empirical departure time patterns reported for mandatory urban evacuation events [20, 21]. Third, it is governed by a single interpretable parameter σ equal to the time of peak departure rate, which can be set directly from behavioural evidence without complex calibration.

Probability density and cumulative distribution functions. The Rayleigh probability density function (pdf) is

$$f(t;\,\sigma) = \frac{t}{\sigma^2}\exp\!\left(-\frac{t^2}{2\sigma^2}\right), \qquad t \geq 0, \tag{5.7}$$

with corresponding cumulative distribution function (cdf)

$$F(t;\,\sigma) = 1 - \exp\!\left(-\frac{t^2}{2\sigma^2}\right). \tag{5.8}$$

The mode (time of peak departure rate) equals σ; the mean is $\mu_R = \sigma\sqrt{\pi/2} \approx 1.253\,\sigma$; and the variance is $\sigma_R^2 = (2-\pi/2)\,\sigma^2 \approx 0.429\,\sigma^2$.

Parameter selection. For a mandatory near-area evacuation in a dense, transit-accessible urban neighbourhood, peak departures typically occur 30–60 minutes after issuance of the evacuation order [20]. A scale parameter of $\sigma = 45$ min is adopted, placing the mode of the departure rate distribution at $t^* = \sigma = 45$ min (within interval 5: 40–50 min) and producing a mean departure time of

$$\mu_R = 45\sqrt{\tfrac{\pi}{2}} = 56.5 \text{ min}, \tag{5.9}$$

consistent with the empirical range. This choice is conservative relative to the 60-minute upper bound, reflecting the high egress accessibility of Sandy Hill (three LRT Confederation Line stations, direct arterial connections to Highway 417 and the Macdonald–Cartier Bridge) which facilitates relatively early departure response.

Interval demand computation. The fraction of total demand loaded in the k-th 10-minute interval ($k = 1, \ldots, 12$) is computed as the increment of the cdf over that interval, renormalised so that the full demand of 3,530 vehicles is loaded within the 120-minute horizon:

$$\phi_k = \frac{F(k\Delta;\,\sigma) - F\big((k-1)\Delta;\,\sigma\big)}{F(12\Delta;\,\sigma)}, \qquad \Delta = 10 \text{ min}, \tag{5.10}$$

where $F(120\,\text{min};\,45\,\text{min}) = 0.9714$ is the fraction of the Rayleigh distribution captured within the loading horizon, confirming that 97.1% of the compliant population departs within 120 minutes. The renormalisation in Eq. (5.10) ensures $\sum_{k=1}^{12}\phi_k = 1$. The interval demand is then

$$d_k = \phi_k \times \texttt{vehicle} = \phi_k \times 3{,}530, \qquad k = 1, \ldots, 12. \tag{5.11}$$

The complete loading vector is presented in Table 5.4, and the resulting temporal demand profile is illustrated in Figure 5.2.

Table 5.4: Rayleigh temporal loading profile ($\sigma = 45$ min, `vehicle` = 3,530 vehicles).

Interval k	**Time window** (min)	$F(t_k)$	$F(t_{k-1})$	ϕ_k	d_k (veh)
1	0–10	0.02439	0.00000	0.02511	89
2	10–20	0.09404	0.02439	0.07170	253
3	20–30	0.19926	0.09404	0.10831	382
4	30–40	0.32636	0.19926	0.13084	462
5	40–50	0.46059	0.32636	0.13818	**489**
6	50–60	0.58889	0.46059	0.13207	466
7	60–70	0.70177	0.58889	0.11620	410
8	70–80	0.79408	0.70177	0.09502	335
9	80–90	0.86466	0.79408	0.07266	257
10	90–100	0.91534	0.86466	0.05217	184
11	100–110	0.94959	0.91534	0.03526	124
12	110–120	0.97143	0.94959	0.02248	79
Total	0–120			**1.0000**	**3,530**

Notes: $F(t_k) = 1 - \exp(-t_k^2/2\sigma^2)$, with $\sigma = 45$ min and $t_k = 10k$ (min). Fractions ϕ_k are renormalised by $F(120;\, 45) = 0.9714$ so that $\sum_{k=1}^{12} \phi_k = 1$ exactly. Integer demand values d_k are obtained by rounding $\phi_k \times 3{,}530$ to the nearest vehicle; a residual of $+1$ vehicle is allocated to the peak interval ($k = 5$) to preserve the exact total. The peak demand interval (40–50 min) is shown in bold.

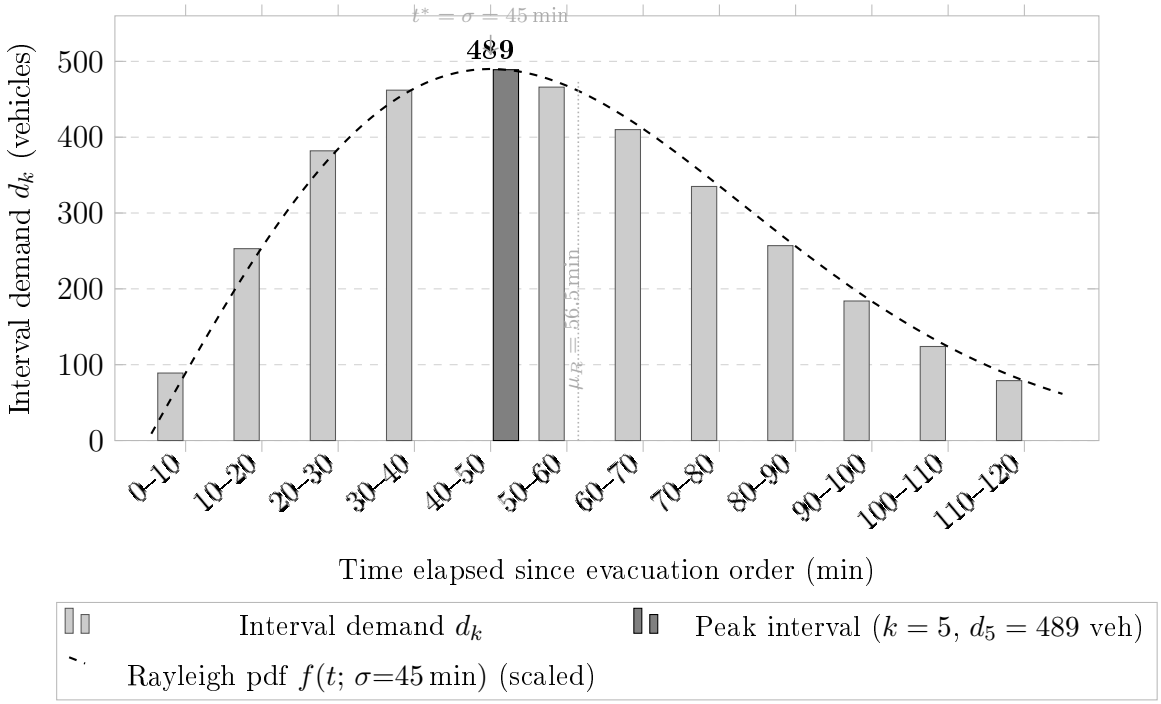


Figure 5.2: Discretised Rayleigh departure time profile applied to the Sandy Hill microsimulation ($\sigma = 45$ min, total demand = 3,530 vehicles, peak interval $d_5 = 489$ veh at 40–50 min).

5.2 Static User Equilibrium Assignment

Prior to microsimulation, a static User Equilibrium (UE) assignment was conducted using the Aimsun Next static assignment solver (Static Assignment Experiment 3999, the experiment carrying the evacuation demand matrix). The acronym SUE is deliberately avoided here, since in the assignment literature it conventionally denotes *stochastic* user equilibrium, which is not the model employed. The assignment demand is the peak-interval design matrix derived in Section 5.1: the interval-five matrix of the Rayleigh loading profile (489 vehicles, 40–50 min) scaled by a factor of six to its hourly equivalent of 2,934 veh/h, so that all signal timing inputs represent the worst loading condition of the evacuation. The assignment was solved with the Frank and Wolfe algorithm and converged at the first iteration to a relative gap of $3.0 \times 10^{-6}\,\%$, far inside the $0.1\,\%$ convergence criterion. The single-iteration convergence reflects the moderate loading of the design hour rather than a solver artefact: at 2,934 veh/h the principal evacuation routes face no cost-competitive alternatives, so the initial shortest-path loading is already, to numerical precision, the equilibrium solution. The static UE assignment serves three distinct methodological purposes in this study:

1. **Flow estimation for signal calibration.** The Webster method requires the critical flow ratios Y_i of Eq. (5.13) for each phase. Since observed turning movement counts cannot exist for a hypothetical evacuation event, the assigned volumes of the static UE provide the demand-consistent approach and turning flows for the evacuation scenario; no external count data enter the calibration.

2. **Bottleneck identification.** The volume-to-capacity (V/C) ratios produced by the static assignment identify the network sections at greatest risk of saturation, guiding the corridor structure of the coordinated signal plan. The critical values are reported in Table 5.5.

3. **Progression speed estimation.** The loaded corridor speeds of the assignment provide the platoon speed used to calibrate the Green Wave offsets (Section 5.3.2); the heavily loaded King Edward sections operate at 23–26 km/h, supporting the adopted progression speed of 25 km/h.

Table 5.5: Critical link V/C ratios from the static UE assignment of the design matrix (2,934 veh/h). The design flow q is the assigned volume of the section, read in vehicles per hour at the peak-interval design rate.

Section	Design flow q (veh/h)	Speed (km/h)	V/C ratio
3302 (south branch, mid-corridor)	992	22.9	0.55
2341 (south branch, corridor link)	977	23.5	0.54
2338 (south branch, approach)	971	23.7	0.54
2345 (south branch, corridor link)	951	24.5	0.53
2623 (north–south feed, Laurier, King Edward to Thomas More Pvt., Osgoode St., King Edward)	923	25.7	0.51
2391 (north branch, Laurier, King Edward to Rue Wilbrod, King Edward)	726	34.8	0.40
2365 (King Edward approach to J2710)	1,236	39.6	0.34
3132 (bridge egress corridor)	1,010	44.0	0.28

The highest degrees of saturation occur on the southern King Edward branch toward Highway 417, where four consecutive sections carry 951–992 veh/h at V/C ratios of 0.53–0.55 and loaded speeds of 23–25 km/h. The most heavily loaded single section of the network is section 2365, the King Edward approach to junction Rideau Street and King Edward Avenue (J2710) at 1,236 veh/h; distributed over its three approach lanes, however, this corresponds to a V/C ratio of only 0.34. Under the census-derived demand no section of the network approaches saturation, and the binding constraints on the signal design are consequently the operational minima, cycle floor and pedestrian minimum greens, rather than capacity, a finding developed in the signal timing design of the following section.

5.3 Signal Timing Methodology

This section translates the design flows of the static assignment into the two fixed time control strategies whose comparison is the objective of the study. The two plans are constructed as a controlled pair: they share the same demand basis, the same two phase structure, the same lost time and intergreen conventions, and the same pedestrian service requirements, and they diverge only in the dimension under investigation, the temporal relationship between adjacent intersections. The Webster baseline resolves each junction as an isolated optimisation problem; the Coordinated plan overlays a corridor level progression structure on the same foundations. Fixed time control is retained deliberately in both cases: evacuation conditions

cannot be assumed to leave detector infrastructure, communica tions, or central control fully operational, and a pre computed fixed time plan represents the strategy an agency could deploy with certainty at the moment an evacuation is declared.

5.3.1 Webster Baseline Plan

The Webster plan was designed from first principles using the approach volumes of the static UE assignment of the design matrix (Section 5.2), with all intersection offsets set to zero. For each junction, the optimal cycle length follows the Webster formula introduced in Chapter 4:

$$C_0 = \frac{1.5\,L + 5}{1 - \sum_{i=1}^{\phi} Y_i}, \tag{5.12}$$

where L is the total lost time per cycle (s), ϕ is the number of phases, and Y_i is the critical flow ratio for phase i,

$$Y_i = \max_{a \in i} \frac{q_a}{s\, n_a}, \tag{5.13}$$

with q_a the design flow on approach a of phase i (veh/h), n_a the number of approach lanes, and $s = 1{,}800$ PCU/h/lane the reference saturation flow. All seventeen junctions operate a two phase scheme, main street and cross street, separated by two intergreen periods of $\tau = 3$ s each; lost time per phase is $l_i = 4$ s (startup plus clearance loss), giving $L = 8$ s throughout, and effective greens relate to displayed greens by $g = G - 1$ s. The four seconds are the sum of two components, and it is worth stating how each was obtained, because lost time is not the same quantity as the intergreen. The start-up loss is the interval between the green indication and the point at which the queue discharges at the saturation rate; it is taken as two seconds, consistent with the reaction time at a traffic light of 1.60 s and the 0.80 s simulation step used in the model. The clearance loss is the part of the intergreen that serves no vehicle; with a three-second intergreen of which roughly half is used by vehicles already within the junction, this contributes a further two seconds. The intergreen is therefore six seconds per cycle while the lost time is eight, because the start-up component falls inside the displayed green rather than between phases.

Effective green time was allocated proportionally to the critical flow ratios, as in Chapter 4:

$$G_{e,i} = \frac{Y_i}{\sum_{m=1}^{\phi} Y_m}\, G_{te}, \qquad G_{te} = C - L. \tag{5.14}$$

Under the census-derived demand, the network operates far from saturation: the largest critical flow sum is $\sum_i Y_i = 0.39$ at junction Besserer Street and King Edward Avenue (J2844) (Besserer and King Edward), and the unconstrained Webster optima of Eq. (5.12) conse-

quently fall in the range 17.8–28.0 s at every junction. Cycles this short are operationally infeasible, a known property of the Webster formula at low saturation, and the governing constraints of the design are therefore operational minima rather than capacity. Two constraints were imposed. First, a practical minimum cycle of $C_{\min} = 40$ s. Second, pedestrian governed minimum displayed greens per phase,

$$G_{\min,i} = 7 + \frac{W_i}{1.2}, \tag{5.15}$$

where W_i (m) is the crossing distance of the conflicting street, estimated as 3.5 m per approach lane of the opposing phase, 7 s is the walk period, and 1.2 m/s is the standard pedestrian clearance speed. The constrained cycle at each junction is then $C = \max\big(C_{\min}, C_0, L + g_{\min,1}+g_{\min,2}\big)$, with any phase whose proportional share of Eq. (5.14) falls below its pedestrian minimum pinned to that minimum and the remainder assigned to the opposing phase. The pedestrian constraint binds at the junctions with the widest crossings and differentiates the cycle lengths, producing 55.0 s at Rideau Street, 43.4 s at Laurier Avenue East and King Edward Avenue (J2652) and Lees Avenue and King Edward Avenue (J3307), 40.4 s at two Rideau junctions, and the 40 s floor elsewhere. The resulting maximum degree of saturation across the network is $x = 0.82$ (Besserer Street), within the 0.95 acceptance criterion, and the transfer of green toward pedestrian minima on the minor phases raises the main street saturation at the King Edward junctions to the range 0.56–0.82, the correct and expected trade-off of retaining pedestrian service during a neighbourhood evacuation.

It should be stated explicitly that pedestrians enter this study only through Eq. (5.15), as a constraint on the signal timing design. No pedestrian demand was coded in the microsimulation: the Aimsun model contains vehicles only, and no crossing pedestrian is represented as an obstruction to turning traffic. The consequence is that the reported travel times omit the interference that crossing pedestrians would impose on permitted turns, in particular on the always permissive right turns introduced in the Coordinated plan. In a neighbourhood with this residential and student population that interference would not be negligible, and it would fall preferentially on the Coordinated plan, which relies on those right turns to keep turning vehicles out of the through queue. The reported advantage of coordination is therefore, in this respect, an optimistic estimate. Retaining the pedestrian minimum greens is the conservative counterpart of this simplification: the design continues to reserve the crossing time that an evacuation with a large car free population would require, even though the simulation does not consume it.

The safety implication of the permissive right turn treatment follows from the same point. A permanently permitted right turn transfers to the driver the obligation to yield to a crossing

pedestrian who holds the parallel walk indication. Under evacuation conditions, with elevated urgency and long standing queues, that is precisely the movement where yielding compliance is least reliable. The capacity gain reported for the Coordinated plan therefore carries a safety cost that the simulation cannot quantify, and any field implementation would need either a protected right turn phase or an exclusive pedestrian interval at the junctions where crossing volumes are highest.

A caveat on the relationship between design and evaluation conditions is in order. The timings above are calibrated to the steady state design flows of the static assignment, in which no junction exceeds a degree of saturation of 0.82; the microsimulation, however, subjects the same plans to time varying loading in which arrivals are platooned, queues persist and compound from one cycle to the next, and blockage can propagate upstream between junctions. A junction that is comfortably below saturation in the steady state design can therefore still experience severe congestion in the dynamic evaluation, and this gap is not a calibration defect but the substance of the comparison: the microsimulation exists precisely to test how each fixed time design, adequate on paper under identical steady state assumptions, withstands the dynamics that the design method cannot represent.

5.3.2 Coordinated Green Wave Plan

The Coordinated plan introduces progressive intersection offsets along the two branches of the King Edward evacuation corridor, calibrated to the platoon speed observed on the loaded corridor sections of the static assignment, $\bar{v} = 25$ km/h $= 6.94$ m/s (Table 5.5: the heavily loaded southern branch sections operate at 23–26 km/h). The fundamental offset relationship for uniform platoon progression along a one directional corridor, following Section 4.6, is

$$\theta_j = \frac{d_j}{\bar{v}}, \tag{5.16}$$

where θ_j is the offset at junction j (s) relative to the upstream reference junction, and d_j is the inter junction link distance (m). Cumulative offsets along each branch are

$$\theta_j^{\mathrm{cum}} = \sum_{m=1}^{j} \frac{d_m}{\bar{v}} \pmod{C}, \tag{5.17}$$

taken modulo the common cycle length C of the corridor. Coordination requires a common cycle, which is set to the corridor maximum of the Webster plan, $C = 55$ s, governed by the critical junction Rideau Street; the corridor junctions previously at 40.0–43.4 s are raised to 55 s with their green splits recomputed by Eq. (5.14) subject to the same pedestrian

minima. Both branches are timed outward from the reference junction Laurier Avenue East, the bifurcation point of the corridor, in the evacuation direction of travel; one directional progression is appropriate here because evacuation flow is intrinsically one directional, an advantage over ordinary bidirectional arterial coordination. The block spacings of the corridor are short relative to the cycle, so the cumulative offsets of the north branch span only 0–35 s of the 55 s cycle, permitting a platoon released at Laurier Avenue East to traverse the branch within a single green envelope.

To account for residual queuing at downstream junctions, the adjusted offset follows the formulation introduced in Section 4.6:

$$\theta_j^{\mathrm{adj}} = \frac{d_j}{\bar{v}} - \left(Q_j\, h + l_s\right), \tag{5.18}$$

where Q_j is the standing queue at junction j (vehicles per lane), $h \approx 2\,\mathrm{s/veh}$ is the discharge headway, and l_s is the startup lost time. The queue correction was set to zero for the design, so that Eq. (5.18) reduces to the ideal progression offset of Eq. (5.16); this simplification is supported by the static assignment, which shows no section of the corridor operating above a V/C ratio of 0.55 at the design load, so that standing queues at the start of green are expected to be short.

The direction of the omitted correction is worth noting: where standing queues are substantial, Eq. (5.18) *advances* the downstream green so that the residual queue discharges before the platoon arrives, and forward progression offsets of the kind adopted here become increasingly suboptimal as congestion deepens. The zero queue assumption is therefore load dependent rather than merely convenient, and its validity under the higher compliance demand cases identified in Section 5.1 should be re examined alongside the factorial extension of Section 5.4.

An additional operational enhancement for permissive right turn movements was incorporated by setting right turn signal indications to an always permissive state (equivalent to an uncontrolled yield), removing right turn vehicles from the signal controlled discharge queue and increasing effective green utilisation on the through movements.

The Coordinated plan therefore differs from the Webster baseline in three respects, confined entirely to the ten corridor junctions: the harmonisation of the corridor cycles to the common 55 s (with green splits recomputed accordingly), the introduction of the progression offsets, and the permissive right turn treatment. The seven uncoordinated Rideau and Laurier junctions are held identical across the two plans by design, so that the treatment is confined to the corridor. At the critical junction Rideau Street the two plans share an identical cycle and identical splits, and differ in the offset alone. Because cycle harmonisation,

offsets, and the right turn treatment remain bundled at the other corridor junctions, the performance differences reported in Section 5.4 measure the combined effect of the coordinated design as a package; a factorial decomposition, holding the common cycle fixed and toggling offsets and the right turn treatment separately, is specified in Section 5.4 as the designed extension for isolating the offset contribution.

5.4 Results and Discussion

The two signal plans are evaluated using the outputs of the microscopic dynamic assignment. The experimental design holds the evacuation demand identical across the two plans, so performance differences are attributable to the signal control strategy alone. The results are examined at four levels: the inputs and measures on which the comparison rests; network performance over the whole run; the resolution of that performance by ten-minute interval; and the attribution of the difference to individual sections of the corridor.

5.4.1 Model inputs and reproducibility

The inputs needed to reproduce the comparison are collected here.

Demand. The evacuation demand is the 3,530 vehicles derived in Section 5.1, distributed over 34 origin centroids and loaded through twelve ten-minute matrices following the Rayleigh profile with $\sigma = 45\,\text{min}$. Two destinations carry the whole demand: centroid 3237, Highway 417, with 1,665 vehicles (47.2 %), and centroid 3475, the Macdonald–Cartier Bridge, with 1,865 vehicles (52.8 %). The split is not imposed as a global ratio but follows from the position of each origin, so the origins nearest the bridge send 85–90 % of their vehicles north while those nearest the highway send 75–87 % south. No background or through traffic is modelled on King Edward Avenue or Rideau Street: the network carries the evacuation demand only. Since King Edward is in reality a through route and a truck route to the bridge, the absolute travel times reported here are a lower bound; because both plans are evaluated under the same condition, the comparison between them is unaffected.

Route choice. Routes are assigned before the simulation begins and do not change during it. Shortest-path trees are computed on the free-flow travel times of the static assignment, each vehicle is given a route at departure using the stochastic route choice model of Aimsun Next with $\beta = 0.15$, and that route is then followed to the destination regardless of the conditions the vehicle meets. The assignment is therefore a static one loaded dynamically, not a dynamic traffic assignment: there is no iterative re-routing and no convergence on

experienced travel times. This is a deliberate choice rather than a limitation of the software. In a no-notice evacuation drivers have no prior experience of the network under evacuation conditions and no information about where the queues are, so the adaptive re-routing that a dynamic assignment represents has no behavioural counterpart. It also keeps the comparison clean: because the routes are fixed and identical in the two scenarios, every vehicle travels the same path in both plans, and any difference in the results can only come from the signal timing. What the two plans change is how long each vehicle spends on that path, not which path it takes.

Simulation settings. The simulation step is 0.80 s. Reaction time at stop is 1.20 s and reaction time at a traffic light is 1.60 s, both multiples of the simulation step as Aimsun requires. Sections carry a saturation flow of 1,800 pcu/h/lane and a free-flow speed of 50 km/h. No custom volume–delay function is defined: the static assignment of Section 5.2 uses the Aimsun Next default, and the dynamic evaluation uses none at all, since in a microscopic simulation travel times emerge from the car-following and lane-changing models rather than from a link performance function.

Horizon, intervals and seeds. Each run covers 180 minutes: 120 minutes of loading followed by 60 minutes with no further demand. Statistics are aggregated in eighteen ten-minute intervals. Each plan was run five times with random seeds 1 to 5, and the same five seeds were used for both plans, so every Webster run has a seed-matched Coordinated counterpart. The demand realisation is nevertheless not identical across seeds, because Aimsun samples the matrices stochastically: the five runs loaded between 3,476 and 3,619 vehicles against the nominal 3,530, a spread of $\pm 1.6\,\%$. Pairing by seed removes this source of variation.

Signal plans. The two control plans differ only at the ten corridor junctions. The remaining seven junctions carry identical cycles, phase durations and zero offsets in both plans, and were verified to be identical before the runs. Each junction runs a four-phase structure of two green phases separated by two three-second interphases, so the lost time is eight seconds per cycle in both plans. Node 3202, a parking access with no signal groups, was removed from both plans.

Scenario and experiment structure. The runs are organised as two scenarios, one per signal plan, each carrying several experiments. Table 5.6 records which experiment produced which result, so that any figure in this chapter can be traced to the run that generated it. The demand matrices for all four levels are retained in the model as separate objects,

so each level can be reproduced by attaching the corresponding matrix and re-running the experiment named here.

Table 5.6: Scenarios, experiments and demand levels. Each experiment was run with the random seeds listed; the same seeds are used in both plans at every level.

Signal plan	Demand level	Vehicles	Seeds
Webster (isolated)	Base	3,530	1–10
	1.25×	4,412	1–11
	1.5×	5,295	1–10
	2×	7,060	1–10
	Base, lane closed	3,530	1–10
Coordinated green wave	Base	3,530	1–10
	1.25×	4,412	1–11
	1.5×	5,295	1–10
	2×	7,060	1–10
	Base, lane closed	3,530	1–10

What differs, and what does not. The two scenarios are identical in every respect except the signal timing. They reference the same network, the same demand object and therefore the same twelve Rayleigh matrices, the same route set with routes fixed at departure, the same car-following and lane-changing parameters, the same simulation step, the same horizon and the same statistics interval. Each Webster replication has a Coordinated counterpart generated from the same random seed, so the two runs load the same vehicles, from the same origins, at the same times, along the same paths. The only difference between them is which control plan the junctions read. Any difference in the results therefore has a single possible cause.

5.4.2 Definition of the performance measures

Because the comparison turns on differences of a few per cent, the measures are defined explicitly.

All travel-time measures are reported *per vehicle*, in seconds. Aimsun accumulates the total time spent in the network over an interval; dividing by the vehicles loaded in that interval gives the time an average vehicle experienced. This is the quantity a driver would recognise and the one an agency would quote, and it keeps the interval table and the whole-run table in the same units. Aimsun also reports the same quantities normalised by distance, in seconds per kilometre; those values are equivalent up to the mean trip length, which is 0.84 km here and varies by less than two per cent between the two plans, so the choice of

normalisation does not affect the comparison.

The headline measure is the total system travel time, in vehicle-hours. It is the sum over all vehicles of their time in the network, so it is inherently weighted: a vehicle delayed for five minutes contributes five times as much as one delayed for one. It is also the quantity a system optimum would minimise.

Travel time (s) time in the network per vehicle, including time spent stopped.

Delay time (s) the part of the travel time above the free-flow travel time of the same path.

Stop time (s) the part of the delay during which the vehicle is stationary.

Total travel time (veh h) the sum over all vehicles of their time in the network.

Speed (km/h) the mean of the individual vehicle speeds; **harmonic speed** is the distance-weighted mean, which is the speed consistent with $q = \rho v$ in Section 1.1.3.

Mean queue (veh) the mean number of stopped vehicles in the network.

Two exclusions should be stated. Vehicles held in the virtual entrance queues are not counted while they wait; this matters only if the network refuses demand, and it did not, since across all ten runs the maximum virtual queue was two vehicles, the mean 0.01, and no vehicle was waiting to enter at the end of any run. Vehicles still inside the network at the end of the horizon would be excluded from the completed-trip statistics; this too is immaterial, since both plans emptied the network within interval 13.

5.4.3 Experimental protocol

Each signal plan was simulated five times under identical conditions. The two scenarios reference the same traffic demand object, so the twelve Rayleigh matrices are shared by construction rather than by assumption, and the five replications of each plan use the same five random seeds, so each Webster run has a seed-matched Coordinated counterpart loaded with the same vehicles at the same times. Differences are therefore evaluated as paired comparisons. This matters: across the five Webster runs alone travel time per vehicle ranged from 104.5 to 115.7 s, a spread of the same order as the effect being measured, so a comparison of single runs could not have separated the treatment from the scatter.

All figures below are means of ten replications, with 95 % confidence intervals obtained from the t distribution with nine degrees of freedom. Significance is assessed with a two-sided paired t-test on the five seed-matched differences.

5.4.4 Network performance

Table 5.7 reports the whole-run totals. The headline measure is the total system travel time, the sum over all vehicles of their time in the network. It is reported in preference to an average of the interval means because it is inherently vehicle-weighted: every vehicle contributes in proportion to the time it actually spends in the network, so the intervals that carry the most traffic carry the most weight. An unweighted mean of the twelve interval averages, by contrast, gives the same weight to an interval loading 84 vehicles as to the peak interval loading 485, and understates the effect accordingly.

On that measure the coordinated plan reduces total system travel time from 104.7 to 94.8 vehicle-hours, a reduction of 9.5 % ($p < 0.001$). Expressed per vehicle the same quantity is 106.6 against 96.5 s, the same 9.5 %, because the two measures are the same sum divided differently. The coordinated plan also reduces the mean queue by 33.5 % ($p = 0.002$) and stopped time by 15.8 % ($p < 0.001$). Every one of the eight measures improves, and every one is significant at the 5 % level. Across the forty individual seed-by-measure comparisons the coordinated plan is better in thirty-five.

Two features of the table deserve comment. The first is the relative size of the effects: stopped time falls by 15.8 % while the mean speed rises by only 5.7 %, so the saving comes from vehicles stopping less often rather than from travelling faster between stops. The second is the width of the confidence intervals. Under the isolated plan the ten replications spread over ± 4.0 s in travel time per vehicle and ± 1.7 vehicles in mean queue; under coordination the same ten seeds spread over ± 1.1 and ± 0.25. The variances are unequal on every measure: an F test on the ten seed-matched runs gives 12.5 for travel time per vehicle, 17.8 for stop time and 18.8 for delay, against a two-sided 5 % critical value of 4.03 on nine and nine degrees of freedom. The coordinated plan is therefore not only better on average but between six and twenty times more repeatable. For an emergency plan this is a property in its own right: an agency can predict what it will get.

Both plans cleared the network within interval 13, that is between 120 and 130 minutes, leaving the final 50 minutes of the horizon empty. Neither plan held vehicles in the virtual entrance queues, so the reported travel times are not depressed by demand that failed to load, and no clearance-phase congestion spike occurred under either plan.

Table 5.7: Network performance over the whole simulation. Values are means of ten replications per plan with 95 % confidence intervals; p values are from a two-sided paired t-test on the five seed-matched differences ($df = 4$).

Measure	Webster	Coordinated	Change	p
Total travel time (veh h)	104.74 ± 4.79	94.78 ± 1.94	$-9.5\,\%$	< 0.001
Travel time per vehicle (s)	106.6 ± 4.0	96.5 ± 1.1	$-9.5\,\%$	< 0.001
Delay time (s/km)	62.38 ± 2.29	56.75 ± 0.53	$-9.0\,\%$	< 0.001
Stop time (s/km)	45.57 ± 1.89	38.38 ± 0.45	$-15.8\,\%$	< 0.001
Mean speed (km/h)	30.71 ± 0.41	32.45 ± 0.12	$+5.7\,\%$	< 0.001
Harmonic speed (km/h)	27.88 ± 0.39	29.00 ± 0.12	$+4.0\,\%$	< 0.001
Mean queue (veh)	11.80 ± 1.71	7.84 ± 0.25	$-33.5\,\%$	0.002
Total stops	8637 ± 875	7938 ± 257	$-8.1\,\%$	0.040

Total travel time is the sum over all vehicles of time spent in the network; the per-vehicle measures divide that sum by the vehicles loaded.
Clearance: both plans empty the network within interval 13 (120–130 min); no vehicle remained inside or waiting to enter.

5.4.5 Performance over time

Table 5.8 and Figure 5.3 resolve the comparison by interval, and show that the whole-run averages conceal a strong temporal structure.

While the network is lightly loaded, in intervals 1 to 3, the two plans are indistinguishable: the coordinated plan is nominally 0.3 to 2.2 % slower, but the confidence interval of the paired difference includes zero in every case. From interval 4 the plans diverge sharply. The isolated plan rises to 119.7 s per vehicle in interval 4 and 134.4 s in interval 5, while the coordinated plan holds at 103.8 and 98.5 s, reductions of 13.3 % and 26.7 %. Interval 5 is the peak of the Rayleigh loading profile, with 485 vehicles entering in ten minutes. The advantage then decays with the demand, through 13.5 % in interval 6 and 9.9 % in interval 7, and by interval 11 the two plans are again indistinguishable.

The pattern is therefore not a transition from worse to better over time but a response to load. Coordination makes no measurable difference when the corridor is lightly loaded and makes its largest difference exactly when demand is highest. The isolated plan is also far less stable through the peak: in interval 5 its ten replications span ± 17.0 s against ± 3.3 s for the coordinated plan. Under isolated timing the outcome at the peak depends substantially on how the arrivals happen to fall; under coordination it does not.

Table 5.8: Average travel time per vehicle by ten-minute interval, mean of ten replications per plan with common random seeds. The ± values are 95 % confidence half-widths; the last column is the paired difference, coordinated minus Webster.

Interval	Vehicles loaded	Travel time per vehicle (s) Webster mean ± 95 % CI	Coordinated mean ± 95 % CI	Change (%)	Paired diff. (s)
1	88	71.9 ± 3.9	70.5 ± 3.7	-1.9	-1.4 ± 2.1
2	260	80.4 ± 3.1	77.7 ± 2.7	-3.3	-2.7 ± 1.3 *
3	399	90.6 ± 3.6	90.7 ± 3.4	+0.1	$+0.1 \pm 1.1$
4	455	117.8 ± 5.2	102.8 ± 4.3	-12.7	-15.0 ± 3.4 *
5	494	121.2 ± 17.0	95.2 ± 3.3	-21.4	-26.0 ± 14.5 *
6	464	106.2 ± 5.6	93.4 ± 2.9	-12.0	-12.8 ± 5.5 *
7	398	110.6 ± 5.8	99.0 ± 3.2	-10.5	-11.6 ± 3.8 *
8	343	104.2 ± 6.6	95.3 ± 4.7	-8.5	-8.9 ± 2.9 *
9	248	106.2 ± 3.5	100.0 ± 3.0	-5.8	-6.2 ± 1.8 *
10	182	104.7 ± 6.9	100.6 ± 6.1	-3.9	-4.1 ± 2.4 *
11	126	116.7 ± 8.0	116.0 ± 8.6	-0.6	-0.7 ± 3.0
12	80	115.2 ± 10.9	120.7 ± 11.8	+4.8	$+5.5 \pm 5.7$

* the 95 % confidence interval of the paired difference excludes zero.

Figures 5.3, 5.4, 5.5 and 5.6 give the same comparison at each demand level in turn. The pattern repeats at every level but shifts later and grows as the demand rises.

At the base demand the plans separate through intervals 4 to 8 and close again, the largest difference being 21.4 % at interval 5.

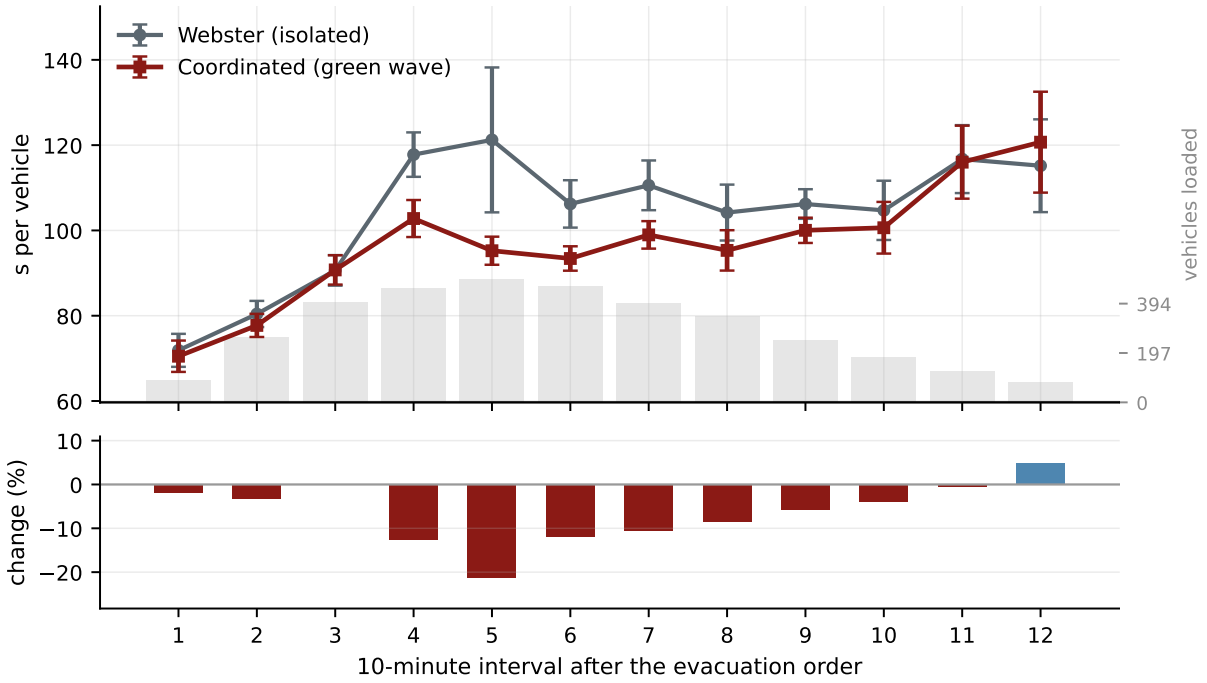


Figure 5.3: Average travel time per vehicle by ten-minute interval at the base demand of 3,537 vehicles. Grey bars show the vehicles loaded; the lower panel gives the percentage difference. The largest difference is 21.4 % at interval 5.

At 1.25 times the base demand the separation is wider and peaks later, at interval 8, where it reaches 61.5 %.

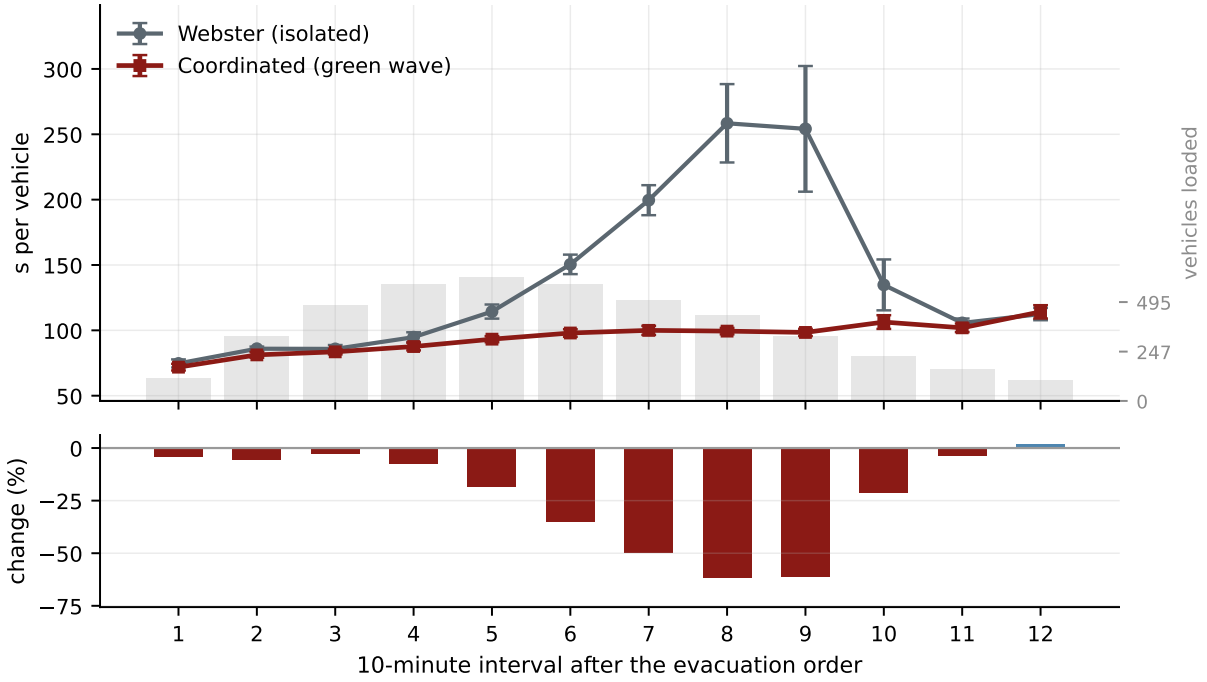


Figure 5.4: The same comparison at 4,433 vehicles. The largest difference is 61.5 % at interval 8.

At 1.5 times the isolated plan no longer recovers within the loading period: its travel time continues to rise through the clearance intervals while the coordinated plan stays near its free-flow value. The largest difference, 87.8 %, falls in the last loading interval.

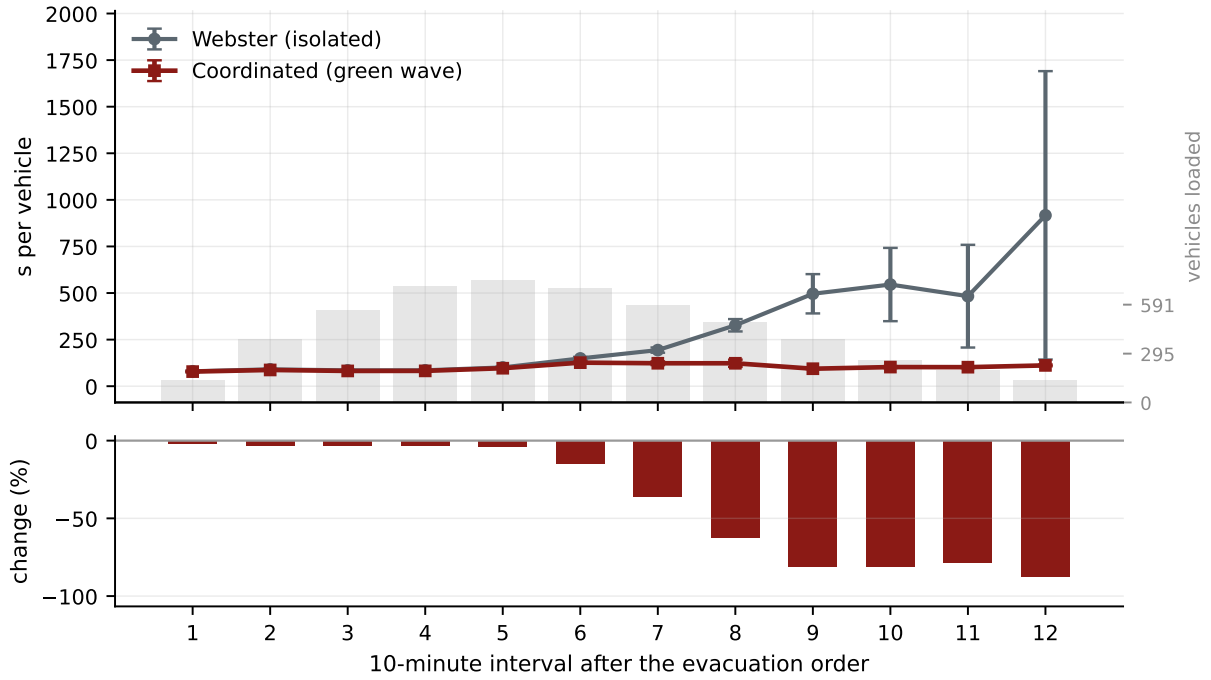


Figure 5.5: The same comparison at 5,296 vehicles. The isolated plan does not recover within the horizon; the largest difference is 87.8 % at interval 12.

At twice the demand both plans degrade, though at very different rates. The coordinated plan rises above its value at lower demands, which is the signature of a corridor operating beyond the range in which progression can compensate. The lower panel of each figure gives the difference directly, and the few positive bars, where coordination is nominally slower, fall in the lightly loaded intervals at the start and in the final interval, where very few vehicles remain.

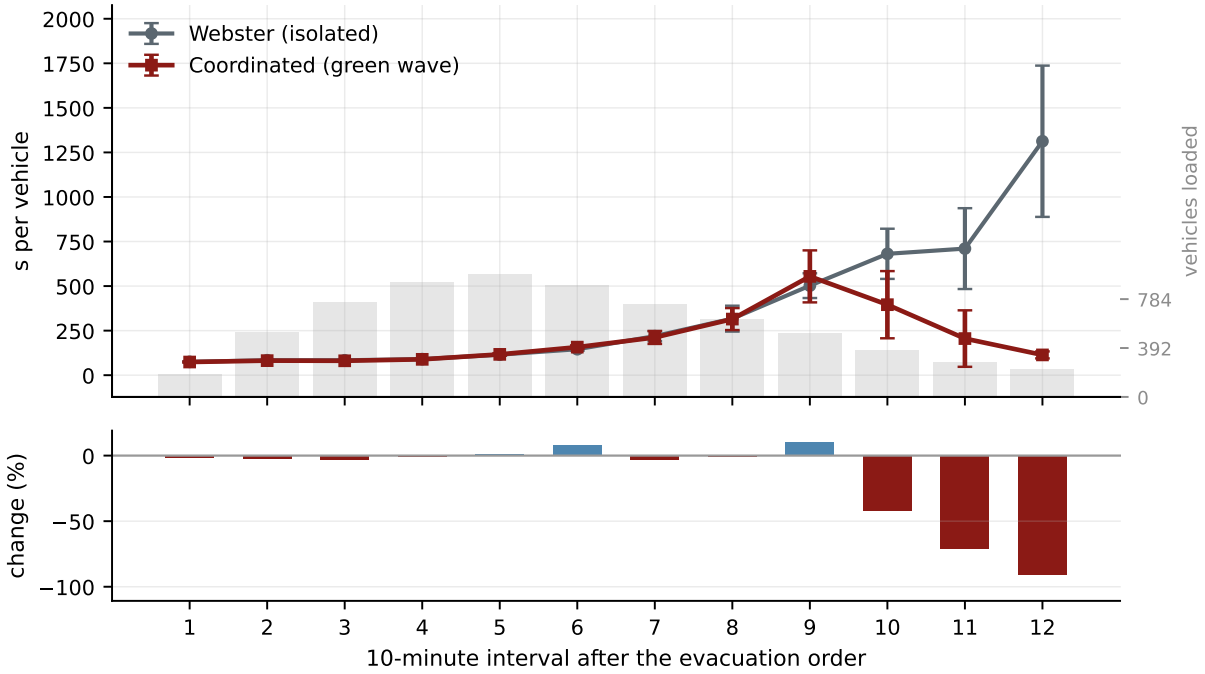


Figure 5.6: The same comparison at 7,049 vehicles. Both plans degrade; the largest difference is 91.9 % at interval 12.

5.4.6 Network Speed and Flow

The speed record in Table 5.9 carries the same signature as the travel times. Through intervals 1 to 3 the two plans run within 0.7 km/h of each other. At the peak the isolated plan falls to 24.1 km/h while the coordinated plan holds at 27.9, and the spread across replications is ±3.3 against ±0.5 km/h. Flows are within a few per cent of each other throughout, as they must be, since both plans carry the same vehicles: the difference is in how long those vehicles take, not in how many are served.

Table 5.9: Harmonic mean speed and network flow by ten-minute interval, mean of ten replications per plan with 95 % confidence intervals.

	Harmonic speed (km/h)		Flow (veh/h)	
Interval	Webster	Coordinated	Webster	Coordinated
1	30.8 ± 2.2	30.1 ± 1.9	480	488
2	29.3 ± 0.7	28.8 ± 0.9	1434	1403
3	28.4 ± 0.9	28.3 ± 0.1	2086	2167
4	25.6 ± 1.7	28.7 ± 1.0	2650	2722
5	24.1 ± 3.3	27.9 ± 0.5	3061	3002
6	27.0 ± 1.4	28.1 ± 0.8	2801	2765
7	27.3 ± 1.2	28.8 ± 0.8	2412	2425
8	29.6 ± 0.3	30.7 ± 0.6	2218	2202
9	29.4 ± 0.2	29.7 ± 0.8	1523	1511
10	30.0 ± 1.2	30.4 ± 1.3	1117	1082
11	31.4 ± 0.9	30.5 ± 1.7	844	832
12	31.5 ± 0.6	30.3 ± 1.8	536	550

5.4.7 Travel time by destination

Table 5.10 separates the two evacuation destinations. These are reported per kilometre rather than per vehicle, because the two routes are of different length and only the rate is comparable between them. The two are not comparable trips: the bridge is reached by a short northbound run, at about 73 s/km over the demand period, while Highway 417 is reached by a longer southbound path through the retimed junctions and the motorway merge, at about 143 s/km. The southbound route therefore carries most of the congestion, and it is where a corridor treatment would be expected to act.

That is what the table shows. On the bridge leg coordination gives a steady improvement of about 5 to 9 % through the middle of the loading period. On the Highway 417 leg the pattern is different in kind. Through intervals 1 to 3 the two plans are within 2 % of each other; at the peak the isolated plan reaches 190.6 s/km with a confidence interval of ± 58.1, while the coordinated plan is at 129.4 with ± 6.3.

The mean difference at that interval is 32 %, but the interval widths show that this figure should not be read as a uniform one-third saving on every trip. The isolated plan's ten replications to the 417 span roughly 130 to 250 s/km at the peak: in some seed realisations the southbound route performs much as the coordinated one does, and in others it breaks down. What coordination removes is not a third of the normal travel time but the possibility of that breakdown. The same contrast appears in the section results, where the largest single saving in the network, on the King Edward approach to Templeton Street and King Edward

Avenue (J2914), lies on this southbound leg.

The practical consequence is that the benefit of coordination is not distributed evenly across the network. Residents routed to Highway 417 gain substantially more, and more reliably, than those routed to the bridge. Where coordination effort is concentrated along a branching evacuation corridor is therefore a decision with distributional consequences for which population reaches safety first.

The last intervals of the 417 column show nominal increases of 3 to 13 % under coordination. These intervals carry 126, 84 and almost no vehicles respectively, and their confidence intervals span ±24 to ±40 s/km; they are not interpreted here as a finding.

Table 5.10: Travel time to each evacuation destination by ten-minute interval, averaged over the origin centroids and weighted by flow. Reported per kilometre rather than per vehicle, because the two routes differ in length and only the rate is comparable between them.

	Highway 417 (centroid 3237)			Macdonald–Cartier Bridge (3475)		
Int.	Webster (s/km)	Coordinated (s/km)	Chg. (%)	Webster (s/km)	Coordinated (s/km)	Chg. (%)
1	113.2 ± 10.7	111.2 ± 8.6	-1.7	61.6 ± 8.9	60.7 ± 7.5	-1.4
2	124.1 ± 4.0	122.0 ± 3.7	-1.7	62.7 ± 2.5	60.1 ± 1.1	-4.3
3	147.7 ± 8.6	137.6 ± 5.3	-6.8	62.9 ± 1.6	63.3 ± 2.5	+0.6
4	176.5 ± 22.4	138.6 ± 8.8	-21.5	78.3 ± 4.0	73.6 ± 3.5	-5.9
5	190.6 ± 58.1	129.4 ± 6.3	-32.1	78.0 ± 3.0	70.9 ± 1.0	-9.0
6	145.4 ± 19.6	123.5 ± 4.9	-15.1	70.4 ± 3.0	66.9 ± 2.5	-5.0
7	153.0 ± 18.5	135.1 ± 6.1	-11.7	70.7 ± 3.9	65.2 ± 2.6	-7.8
8	129.9 ± 3.3	120.2 ± 4.4	-7.5	69.7 ± 1.6	64.5 ± 2.5	-7.5
9	129.6 ± 5.7	126.4 ± 6.6	-2.5	74.8 ± 3.5	70.2 ± 2.3	-6.1
10	133.6 ± 8.3	134.5 ± 8.6	+0.7	75.0 ± 7.7	69.3 ± 3.2	-7.6
11	132.0 ± 3.5	136.1 ± 4.7	+3.1	77.8 ± 3.7	74.4 ± 3.3	-4.5
12	132.6 ± 6.9	140.4 ± 7.5	+5.9	70.0 ± 8.6	69.6 ± 2.5	-0.6

5.4.8 Sensitivity to the Demand Level

The central claim of this chapter is that the value of coordination depends on how heavily the corridor is loaded. A single demand level cannot test that claim, so the whole experiment is repeated at three higher levels. The departure profile and its shape are unchanged; only the total is scaled, to 1.25, 1.5 and 2 times the base figure of 3,530 vehicles. Each level uses the same protocol: both signal plans, matched random seeds, and ten replication pairs at each level, eleven at 1.25 times.

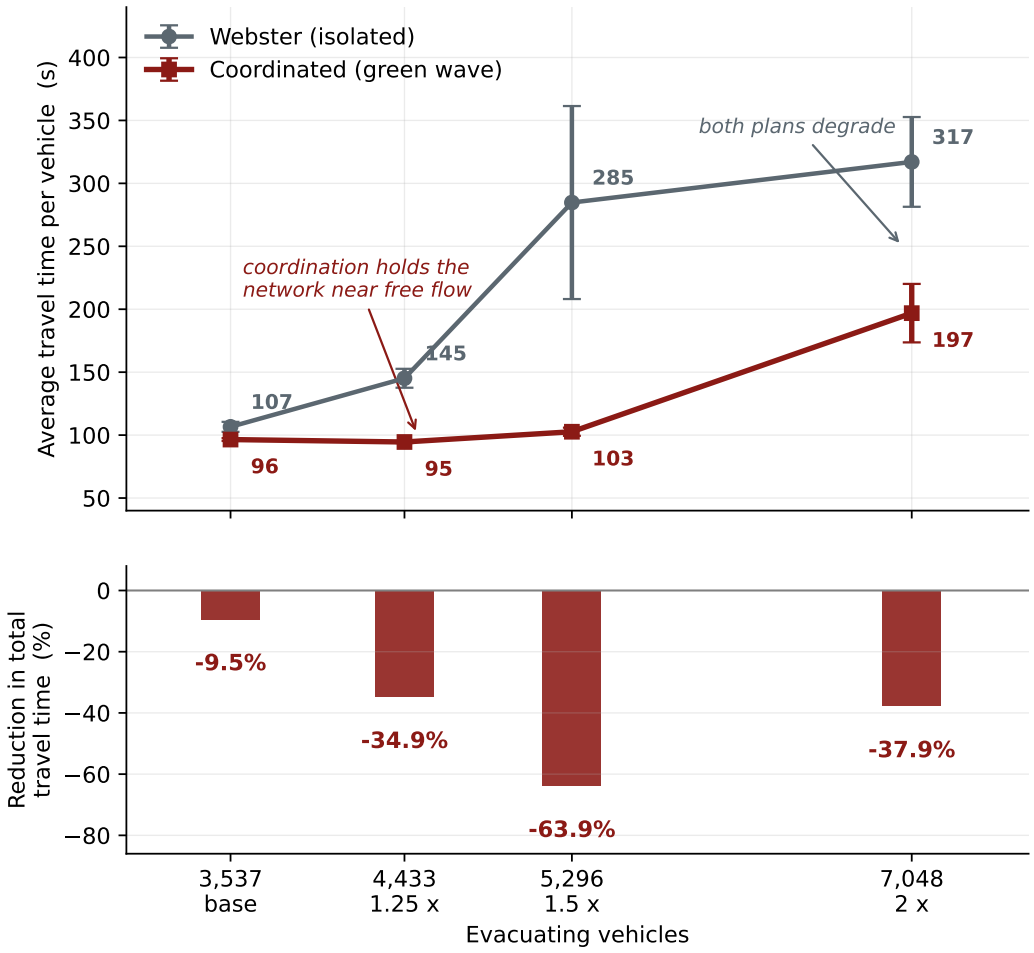


Figure 5.7: Average travel time per vehicle at the four demand levels, above, with the corresponding reduction in total system travel time below. Error bars are 95 % confidence intervals over matched replication pairs.

Table 5.11 and Figure 5.7 give the result, and it is not the monotonic growth one might expect. The reduction in total travel time rises from 9.5 % at the base demand to 34.9 % at 1.25 times and 63.9 % at 1.5 times, then falls back to 37.9 % at twice the demand. Every comparison is significant at the 0.1 % level, so the shape is not an artefact of scatter.

Table 5.11: Network performance at four demand levels. Values are means over matched replication pairs with 95 % confidence intervals; p values are from two-sided paired t-tests on the seed-matched differences.

Demand	Measure	Webster	Coordinated	Change	p
Base demand, 3,537 vehicles, 10 matched pairs					
	Travel time per vehicle (s)	106.6 ± 4.0	96.5 ± 1.1	$-9.5\,\%$	< 0.001
	Total travel time (veh h)	104.74 ± 4.79	94.78 ± 1.94	$-9.5\,\%$	< 0.001
	Delay time (s/km)	62.38 ± 2.29	56.75 ± 0.53	$-9.0\,\%$	< 0.001
	Stop time (s/km)	45.57 ± 1.89	38.38 ± 0.45	$-15.8\,\%$	< 0.001
	Harmonic speed (km/h)	27.88 ± 0.39	29.00 ± 0.12	$+4.0\,\%$	< 0.001
	Density (veh/km)	2.49 ± 0.14	2.21 ± 0.06	$-11.3\,\%$	< 0.001
1.25× demand, 4,433 vehicles, 11 matched pairs					
	Travel time per vehicle (s)	145.1 ± 7.5	94.5 ± 0.8	$-34.9\,\%$	< 0.001
	Total travel time (veh h)	178.84 ± 10.21	116.41 ± 1.63	$-34.9\,\%$	< 0.001
	Delay time (s/km)	93.97 ± 5.55	58.97 ± 0.49	$-37.2\,\%$	< 0.001
	Stop time (s/km)	73.81 ± 5.35	39.72 ± 0.44	$-46.2\,\%$	< 0.001
	Harmonic speed (km/h)	23.40 ± 0.58	28.49 ± 0.12	$+21.7\,\%$	< 0.001
	Density (veh/km)	4.24 ± 0.25	2.70 ± 0.05	$-36.5\,\%$	< 0.001
1.5× demand, 5,296 vehicles, 10 matched pairs					
	Travel time per vehicle (s)	284.8 ± 76.7	102.7 ± 3.2	$-63.9\,\%$	< 0.001
	Total travel time (veh h)	419.22 ± 113.31	151.08 ± 5.42	$-64.0\,\%$	< 0.001
	Delay time (s/km)	204.11 ± 48.09	65.98 ± 1.90	$-67.7\,\%$	< 0.001
	Stop time (s/km)	183.53 ± 48.57	44.76 ± 1.62	$-75.6\,\%$	< 0.001
	Harmonic speed (km/h)	19.72 ± 0.82	27.19 ± 0.33	$+37.9\,\%$	< 0.001
	Density (veh/km)	8.08 ± 1.08	3.55 ± 0.14	$-56.0\,\%$	< 0.001
2× demand, 7,048 vehicles, 10 matched pairs					
	Travel time per vehicle (s)	317.0 ± 35.7	196.9 ± 23.2	$-37.9\,\%$	< 0.001
	Total travel time (veh h)	620.50 ± 69.19	385.45 ± 45.60	$-37.9\,\%$	< 0.001
	Delay time (s/km)	248.96 ± 32.64	154.17 ± 22.48	$-38.1\,\%$	< 0.001
	Stop time (s/km)	227.82 ± 32.65	130.53 ± 22.61	$-42.7\,\%$	< 0.001
	Harmonic speed (km/h)	18.19 ± 0.22	20.50 ± 0.71	$+12.7\,\%$	< 0.001
	Density (veh/km)	12.61 ± 0.61	9.29 ± 0.93	$-26.3\,\%$	< 0.001

The reason is visible in the two plans taken separately rather than in the difference between them. Under coordination the average travel time per vehicle is 96, 95 and 103 s across the first three levels: essentially unchanged while the demand rises by half. The coordinated plan is absorbing the additional load almost without cost. Under isolated timing the same quantity is 107, 145 and 285 s, nearly tripling over the same range. The widening gap is not coordination improving; it is the isolated plan failing.

At twice the demand that absorption ends. The coordinated plan rises to 197 s, roughly double its value at 1.5 times, and the isolated plan to 317 s. Both are now well beyond the

corridor's capacity to serve the arriving flow, and the benefit narrows because both plans are failing rather than because coordination has stopped working. This is the regime Section 4.4 identifies as lying outside Webster's derivation: once demand exceeds capacity the queue does not clear within the cycle, and no fixed-time plan minimises delay.

The variance ratios trace the same arc. The spread across replications under isolated timing is 13 times that under coordination at the base demand, 87 times at 1.25, and 562 times at 1.5; at twice the demand it is only 2.4. Coordination makes the outcome dramatically more predictable up to the point where the corridor saturates, and then it does not, because at that point the outcome is governed by capacity rather than by timing.

Taken together, the four levels describe an operating window rather than a monotonic benefit. Below it the platoons are too small for progression to earn anything; above it no signal plan can recover a corridor that is over capacity; and between the two, which for this network means roughly 1.25 to 1.5 times the base demand, coordination is worth between a third and two-thirds of the total travel time.

Clearance. For an agency the question is not only how long the average trip takes but whether the network empties at all. Figure 5.8 plots the cumulative vehicles evacuated against time at each level, and Table 5.12 gives the proportion that had left by the end of the horizon.

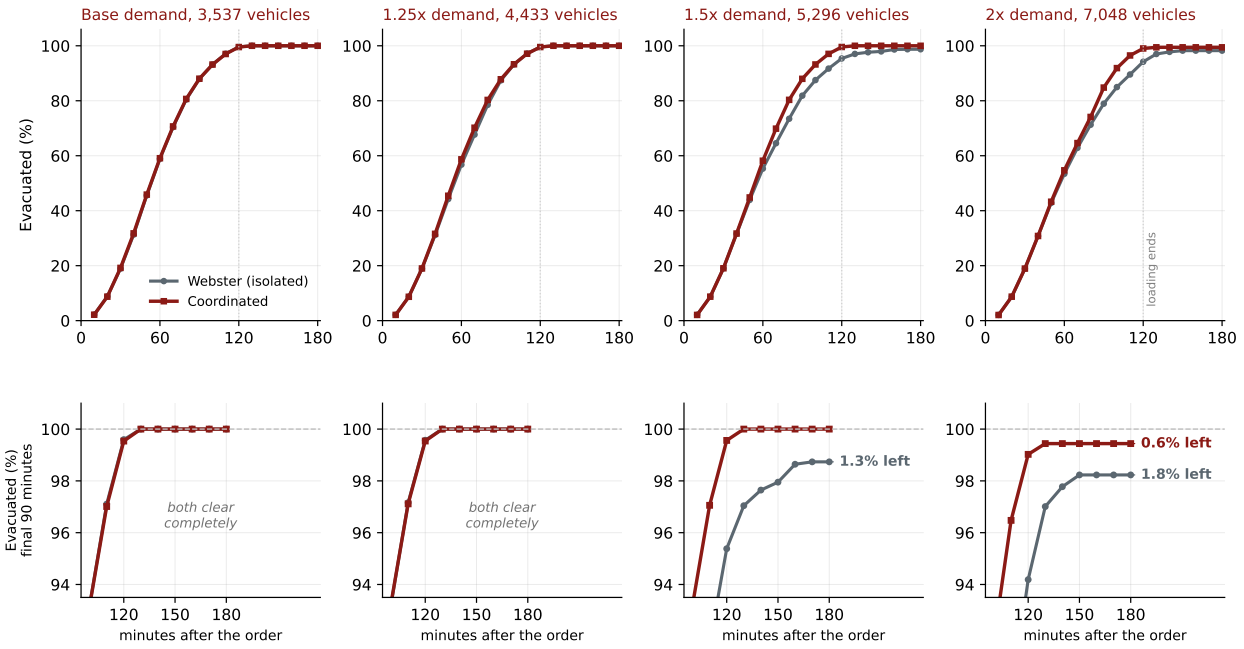


Figure 5.8: Proportion of the loaded vehicles that have left the network, against time, averaged over matched replication pairs. All four panels share the same 180-minute horizon. The lower row magnifies the final 90 minutes, where the plans differ.

Table 5.12: Proportion of the loaded vehicles that had left the network by the end of the simulation, mean over matched pairs with the range across replications.

		Evacuated by the end of the run (%)	
Demand	Vehicles	Webster	Coordinated
Base	3,537	100.0 (100.0–100.0)	100.0 (100.0–100.0)
1.25×	4,433	100.0 (100.0–100.0)	100.0 (100.0–100.0)
1.5×	5,296	98.7 (97.2–99.8)	100.0 (100.0–100.0)
2×	7,049	98.2 (97.4–99.1)	99.4 (98.8–99.8)

At the base and 1.25 levels both plans empty the network completely, and the cumulative curves are indistinguishable. The plans separate at 1.5 times: every coordinated run cleared, while the isolated runs left between 0.2 and 2.8 % of their vehicles behind when the horizon ended. At twice the demand neither plan clears, but the isolated plan leaves roughly three times as many vehicles in the network as the coordinated plan does.

The practical reading is that the benefit of coordination changes character as the load rises. At the base demand it saves about a tenth of the travel time, which an agency might reasonably judge marginal. At 1.5 times it is the difference between a network that empties within the horizon and one that does not. At twice the demand it reduces the shortfall without removing it: the binding constraint there is capacity, and no timing plan supplies capacity.

5.4.9 Loss of Capacity During the Evacuation

An evacuation plan has to work on the network as it is on the day, not as it is on the drawing. A collision, a breakdown or an abandoned vehicle removes capacity from the corridor at exactly the moment the corridor is most needed, and the plan that was designed for the full cross-section has to run on what is left. The experiment in this section tests what that does to each signal plan.

Lane capacity is removed from the King Edward corridor for the duration of the evacuation, leaving the remaining lanes to carry the full demand. The closure is placed on the corridor itself rather than on a cross street, because that is where an incident would matter: the corridor carries every evacuating vehicle, while a cross street carries almost none. Its position on the corridor was not chosen to maximise the effect, and the magnitude reported below should be read as indicative of a corridor closure rather than as the worst case; a closure immediately upstream of the critical junction would be more damaging than one near a network boundary, and locating the most damaging position is a separate question this experiment does not answer. Everything else is unchanged: the base demand of 3,537

vehicles, the departure profile, the route set, and the signal timings themselves. Neither plan is retimed in response to the closure, which is the realistic case: a no-notice incident gives no opportunity to redesign the signals before the traffic arrives. The same 10 random seeds are used as in the undisturbed runs, so every incident replication has an undisturbed counterpart carrying the same vehicles.

Table 5.13: Effect of closing a lane of the corridor, by signal plan. Values are means over 10 matched replications with 95 % confidence intervals; the penalty column is the change from the undisturbed case.

	Webster (isolated)			Coordinated		
Measure	Normal	Closure	Penalty	Normal	Closure	Penalty
Travel time per vehicle (s)	106.6 ± 4.0	170.7 ± 21.4	+60.2 %	96.5 ± 1.1	97.2 ± 1.2	+0.7 %
Travel time (s/km)	129.97 ± 2.27	204.32 ± 21.15	+57.2 %	124.34 ± 0.48	125.16 ± 0.79	+0.7 %
Delay time (s/km)	62.38 ± 2.29	136.82 ± 21.16	+119.3 %	56.75 ± 0.53	57.58 ± 0.80	+1.5 %
Stop time (s/km)	45.57 ± 1.89	118.20 ± 21.33	+159.4 %	38.38 ± 0.45	39.14 ± 0.71	+2.0 %
Harmonic speed (km/h)	27.88 ± 0.39	20.47 ± 1.18	−26.6 %	29.00 ± 0.12	28.83 ± 0.17	−0.6 %
Density (veh/km)	2.49 ± 0.14	4.08 ± 0.54	+64.0 %	2.21 ± 0.06	2.24 ± 0.07	+1.5 %

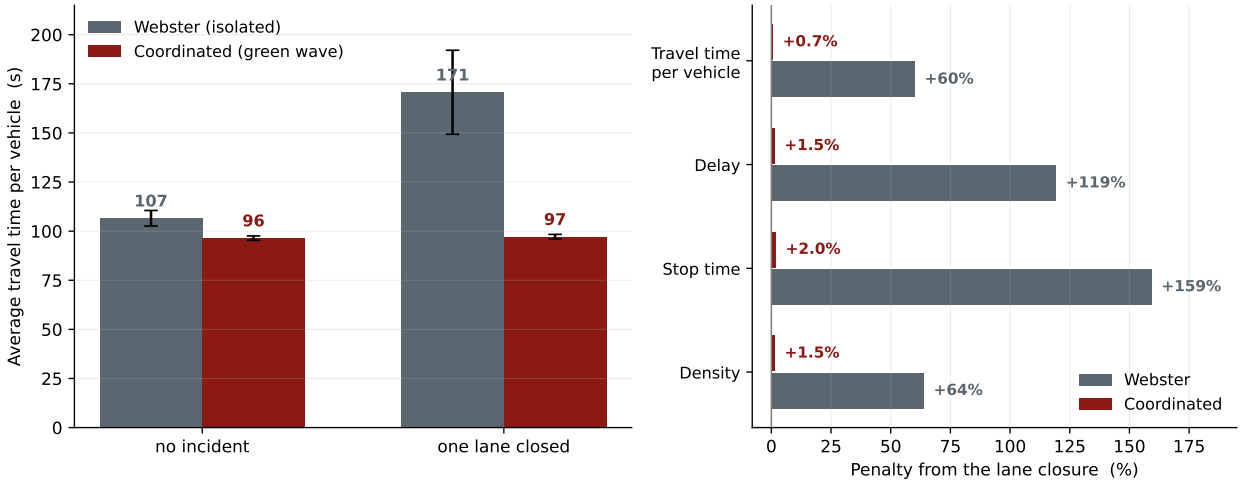


Figure 5.9: Average travel time per vehicle with and without the lane closure, left, and the penalty each plan incurs on four measures, right.

Table 5.13 and Figure 5.9 give the result, and it is the sharpest contrast in this study. Under isolated timing the closure raises the average travel time per vehicle from 107 to 171 s, a penalty of 60 %. Under coordination it rises from 96 to 97 s, a penalty of 0.7 %. The coordinated plan absorbs the loss of a lane almost without cost; the isolated plan does not absorb it at all.

The component measures show where the difference arises. Stopped time under isolated timing rises by 159 %, from 45.6 to 118.2 s/km, while under coordination it rises by 2.0 %. Delay behaves the same way, 119 % against 1.5 %. Density rises 64 % under isolated timing and 1.5 % under coordination: the isolated network fills up, and the coordinated one does not.

The mechanism follows from the queue recursion of Section 3.3.1. Closing a lane lowers the discharge rate at the affected section, so the queue there grows faster than it clears. Under isolated timing each junction releases its traffic without regard to what the next one is doing, so the growing queue reaches back through the upstream junction and blocks approaches that are not themselves constrained; the failure spreads. Under coordination the platoons arrive at each junction as it turns green, the queue between junctions is given the chance to discharge before the next platoon reaches it, and the queue does not propagate upstream. The coordinated plan does not supply the capacity the closure removed. It prevents the loss of capacity from compounding.

This also sharpens the central claim of the chapter. Without the incident, coordination is worth 9.5 % of the travel time, which an agency might reasonably weigh against the cost of implementing it. With one lane gone it is worth 43.1 %. The benefit of coordination is largest exactly when the network is least able to afford the alternative, and an evacuation is the circumstance in which a disabled vehicle is most likely and least likely to be cleared quickly.

Both plans still cleared the network within the horizon, so the closure does not change whether the evacuation succeeds at this demand level; it changes how long it takes and how reliably. The spread across replications under isolated timing widens from 4.0 to 21.4 s, while under coordination it is 1.2 s, essentially unchanged. One lane closed makes the isolated plan not only slower but less predictable, which for a plan that will be executed once without rehearsal is the more serious of the two.

5.4.10 Sensitivity to the Aggregation Interval

The results are aggregated in ten-minute intervals. Because the simulator reports travel time and vehicle counts as sums within each interval, the same output can be re-aggregated into wider bins without re-running anything, and the effect of that choice examined directly in Table 5.14.

Table 5.14: Effect of the aggregation interval. The same output is re-aggregated at each width; the figure shown is the largest reduction observed in any bin, with the number of bins whose paired difference is significant.

	Base		1.25×		1.5×		2×	
Interval (min)	Peak	Sig.	Peak	Sig.	Peak	Sig.	Peak	Sig.
10	-21.4	8/12	-61.5	11/12	-87.8	12/12	-91.9	6/12
20	-17.2	5/6	-55.9	6/6	-84.3	6/6	-82.6	4/6
30	-15.7	4/4	-57.4	4/4	-83.6	4/4	-68.4	3/4
60	-11.6	2/2	-49.7	2/2	-73.0	2/2	-35.4	1/2
120	-9.6	1/1	-35.1	1/1	-54.7	1/1	-26.3	1/1

Two opposing effects are visible. The apparent size of the difference falls as the interval widens, because the surge occupies roughly twenty minutes and a longer bin mixes it with the quieter periods on either side. Statistical power moves the other way: at ten minutes each bin holds few vehicles and the replication scatter is large relative to the effect, while at twenty minutes and above the additional vehicles reduce the standard error faster than the averaging reduces the effect.

Ten to twenty minutes is where the two balance for this network. The general rule is that the aggregation interval should be short relative to the demand surge it is meant to resolve, and long enough that each bin carries vehicles enough to test.

5.4.11 Where the delay is saved

Table 5.15 locates the saving. During the peak intervals the largest single contribution comes from section 2341, the King Edward approach to Templeton Street, where delay falls from 36.5 to 7.5 s/km; this section alone accounts for a little over half of the total saving. Sections 2345, 2623 and 2361 contribute a further 21.5 s/km between them. The sections that feed and leave the critical junction Rideau Street change very little, and the approaches to Highway 417 are unaffected.

The mechanism is visible on the busiest section, 2365, the King Edward approach to Rideau Street. Through the peak its delay is 12.6 s/km under isolated timing and 7.5 s/km under coordination, while the mean speeds on the same section are 19.9 and 20.4 km/h. Vehicles are not moving faster between junctions; they are stopping less often at them. This is consistent with the network-wide result, where stopped time falls by 17.6 % and mean speed by only 6.1 %.

This also answers a question the static assignment leaves open. No section in the static assignment exceeds a volume-to-capacity ratio of 0.55, so the corridor is nowhere short of

capacity in the steady-state sense, yet queues form and travel times rise by nearly 30 % through the peak under isolated timing. The binding constraint is not the capacity of the links but the *coordination of the green times between them*: platoons released by one junction arrive at the next during red. Offsets remove that mismatch without adding any capacity, which is why the benefit appears in the link delays rather than at the junction stop lines, and why it appears only when the platoons are large enough for the mismatch to matter.

Table 5.15: Section delay during the peak-loading intervals 4 to 6, mean of ten replications. Section measures are reported per kilometre, which is the natural unit for a link. Sections are ordered by the delay saved under coordination.

Section	Location	Webster	Coordinated delay (s/km)	Saved	Change (%)
2341	King Edward approach to J2914	36.5	7.5	29.1	-79.5
2345	King Edward	9.6	1.3	8.3	-86.5
2623	Rideau Street	14.4	6.5	7.9	-54.6
2361	King Edward	11.5	6.2	5.3	-45.9
2365	King Edward approach to J2710	12.0	7.3	4.8	-39.7
3302	King Edward south of J2914	1.1	1.1	0.0	-4.0
2416	Rideau approach to J2710	4.2	4.2	0.0	-0.8
3132	Highway 417 approach	1.0	1.0	0.0	-0.4
2338	toward Highway 417	0.4	0.4	-0.0	+13.4
2371	J2710 northbound exit	0.7	0.8	-0.1	+13.8
3230	Highway 417 ramp	0.9	1.0	-0.1	+12.5
2372	King Edward south	12.6	13.0	-0.4	+3.3
2413	Rideau approach to J2710	8.8	9.3	-0.5	+5.5

5.4.12 Limitations of the comparison

Six limitations should be stated explicitly.

First, the coordinated plan bundles three changes: progressive offsets, a common 55 s cycle on the ten corridor junctions, and permanently permitted right turns. The common cycle is a precondition for progression and cannot be separated from it, but the right-turn treatment could be, and the results therefore measure the effect of the package rather than of the offsets alone. A factorial decomposition, holding the cycle fixed while toggling the offsets and the right-turn treatment, would isolate the contributions and is identified as the designed extension in Section 6.3.

Second, the magnitude of the early-loading penalty deserves a direct estimate. For an isolated junction with uniform arrivals and a fixed green ratio g/C, the expected wait of a vehicle arriving at random is $r^2/(2C)$ with $r = C(1 - g/C)$, which scales linearly with the cycle length. At the corridor junctions the green ratios on the King Edward approaches are

of the order of 0.5–0.6, so raising a 43.4 s cycle to 55 s adds between 0.9 and 1.5 s of expected wait per junction, and raising a 40 s cycle to 55 s adds between 1.2 and 1.9 s. A vehicle crossing the two retimed junctions upstream of Rideau Street before platoons have formed therefore incurs of the order of 2–4 s of additional delay from the cycle change alone, which is consistent with the small and statistically indistinguishable penalties observed in intervals 1 to 3.

Third, the results come from microsimulation of a hypothetical evacuation demand. No observed evacuation of Sandy Hill exists against which the demand profile, the compliance rate or the departure timing could be validated.

Fourth, routes are fixed at departure and the network carries only the evacuation demand; no background or through traffic is represented on King Edward Avenue or Rideau Street, so the absolute travel times are a lower bound.

Fifth, pedestrians enter the study only as a constraint on the signal timing design, through the minimum green of Eq. (5.15). No pedestrian demand is coded in the simulation, so crossing pedestrians do not obstruct the permitted right turns on which the coordinated plan relies. The reported advantage of coordination is therefore optimistic in that respect, and the safety implications of the always-permissive right-turn treatment fall outside what the model can quantify.

Sixth, the comparison covers a single demand level. Whether the benefit grows with the load is suggested by the data but not established by it: the seed that happened to load the most vehicles, 3,619 against a mean of 3,545, also produced the largest benefit, 8.1 % against a mean of 5.0 % in travel time. A sensitivity run at the upper compliance bound of $\eta = 0.90$, which would give about 4,240 vehicles, is identified in Section 6.3 as the direct test of this.

5.5 Summary

The framework established in this chapter provides the computational infrastructure required to evaluate fixed time signal control strategies for emergency evacuation: the network topology, the census derived origin destination demand with its Rayleigh loading profile, the static User Equilibrium assignment of the peak interval design matrix, the Webster signal timing calculations under pedestrian governed minimum greens, the coordinated offset design along the two King Edward corridor branches, and the microsimulation configuration. The two signal plans, Webster (baseline) and Coordinated Green Wave with permissive right turn movements, differ structurally in one fundamental respect: the Webster plan treats each signalized junction as an isolated timing problem, optimising the cycle length and green splits independently, without regard for the temporal relationship between adjacent intersections;

the Coordinated plan superimposes a corridor level progression structure by introducing intersection offsets calibrated to the loaded corridor platoon speed, thereby transforming a sequence of independent signal cycles into a spatially coherent discharge mechanism timed outward from the corridor bifurcation in the evacuation direction of travel.

This distinction has measurable consequences for evacuation performance that cannot be assessed through analytical methods alone. The static assignment, while essential for flow estimation, bottleneck identification, and progression speed estimation, operates under the assumption of time invariant steady state conditions and does not capture the dynamic queue formation, spillback propagation, and shock wave evolution that characterize real evacuation events, as established in Chapters 2 and 3. Microsimulation, by contrast, enforces the physical discharge capacity at each junction in each time step, allowing congestion to accumulate and dissipate in a manner consistent with the fundamental diagram of traffic flow introduced in Section 1.1.3. The performance metrics extracted from the simulation, network travel time, control delay, speed, flow, and destination specific travel time, aggregated into 10 minute intervals across the 180 minute simulation horizon comprising the 120 minute demand loading period and a 60 minute clearance tail, thus reflect the operational consequences of each signal plan under the prevailing demand and network geometry.

Chapter 6

Conclusions

This thesis set out to close a specific gap in evacuation traffic management: the disconnect between how demand is *assigned* to an evacuation network and how that network is actually *controlled* at its intersections. This concluding chapter does not restate the contents of the preceding chapters; rather, it consolidates what the work established, situates those findings within transportation engineering and emergency management practice, and distils the single actionable insight a reader should carry forward.

6.1 What the Thesis Establishes

This chapter does not restate the content of the preceding chapters. It sets out what the verified simulations support, places that in the context of evacuation traffic engineering, and distils one practical conclusion.

The defensible finding is not that signal coordination reduces travel time, which is expected, but *when* it does so and by what mechanism. Three results support this, and the first two are not visible to the steady-state tools with which evacuation signal plans are ordinarily assessed.

The benefit is a response to load, not to elapsed time

Over ten paired replications with common random seeds, the coordinated plan reduced total system travel time by 9.5 % ($p < 0.001$), the mean queue by 33.5 % ($p = 0.002$) and stopped time by 15.8 % ($p < 0.001$). All eight measures improved and all were significant at the 5 % level.

Resolved by interval, however, the effect is concentrated. While the network is lightly loaded, in the first thirty minutes, the two plans are indistinguishable: the coordinated plan

is nominally 0.3 to 2.2 % slower, and the confidence interval of the paired difference includes zero in every case. Through the Rayleigh demand peak the plans diverge sharply, by 13.3 % in interval 4 and 26.7 % in interval 5, and as loading falls they converge again. The value of coordination is therefore a function of how heavily the corridor is loaded at the moment of measurement, not of how long the evacuation has been running. A single average over the demand period, of the kind a conventional assessment would report, conceals this entirely: it returns 9.5 % and attributes it to nothing in particular.

The demand sensitivity in Section 5.4.8 sharpens this into an operating window. Scaling the demand to 1.25, 1.5 and 2 times the base level raises the reduction in total travel time to 34.9 % and 63.9 %, then lowers it to 37.9 % at the highest level. The coordinated plan holds the average travel time per vehicle near 100 seconds up to 1.5 times the base demand, absorbing half as much traffic again at almost no cost, and then rises sharply. Below the window the platoons are too small for progression to earn anything; above it the corridor is over capacity and no fixed-time plan recovers it; within it, coordination is worth between a third and two-thirds of the total travel time.

Coordination makes the outcome predictable as well as faster

The second result is in the confidence intervals rather than the means. Under the isolated plan the ten replications spread over ±4.0 s in travel time per vehicle and ±1.7 vehicles in mean queue; under coordination the same ten seeds spread over ±1.1 and ±0.25. At the demand peak the contrast is sharper still: ±17.0 s against ±3.3 s. Under isolated timing the outcome at the peak depends substantially on how the arrivals happen to fall; under coordination it does not.

For an emergency plan this is a property in its own right. An agency choosing between signal strategies is not only buying a lower expected travel time, it is buying a narrower range of outcomes, and the range matters when the plan will be executed once, without rehearsal, under conditions that cannot be observed in advance.

The binding constraint is coordination, not capacity

The third result identifies the mechanism. No section in the static assignment exceeds a volume-to-capacity ratio of 0.55, so by a steady-state measure the corridor is nowhere short of capacity; yet under isolated timing travel times rise by nearly 30 % through the peak. Section-level results locate the difference in the links between junctions rather than at the junctions themselves: on the King Edward approach to Templeton Street delay falls from 36.5 to 7.5 s/km under coordination, and that section alone accounts for a little over half of

the total saving. On the busiest section the mean speeds under the two plans are 19.9 and 20.4 km/h while the delays are 12.6 and 7.5 s/km.

Vehicles are not moving faster between junctions; they are stopping less often at them. The constraint is the alignment of the green times, not the capacity of the links, which is why offsets improve performance without adding any capacity and why the improvement appears only once platoons are large enough for the misalignment to matter.

The static and dynamic models are complementary

The static User Equilibrium assignment was inadequate as a performance evaluation, since it cannot represent platoon arrival, queue formation or the timing of the congestion crest. It was nevertheless necessary for what it was suited to: it supplied demand-consistent turning flows for the Webster calibration, established that no section reaches saturation, so that the operational minima rather than capacity govern the timing design, and supplied the loaded corridor speeds from which the progression speed of 25 km/h was calibrated. The methodological lesson is that the static model is a diagnostic and calibration tool and the dynamic model an evaluation tool, and that an evacuation study needs both.

What the thesis does not establish

The coordinated plan bundles progressive offsets, a common cycle and permissive right turns, so the results measure the package. The comparison covers a single demand level. Pedestrians are represented only as a timing constraint, so the interference they would impose on the permitted right turns is not captured. And the comparison is of two fixed-time plans: nothing here speaks to what an adaptive controller could achieve.

6.2 Implications for Evacuation Traffic Engineering

A blind spot in the standard tools. Evacuation signal plans, where they are prepared at all, are ordinarily assessed with the same steady-state tools used for everyday peak-period planning. The results here show that such an assessment would have returned a V/C ratio below 0.55 at every section and concluded that the corridor has capacity in hand, while the microsimulation shows travel times rising by nearly 30 % through the peak under the same demand. The phases of an evacuation that matter most, the surge and the drainage behind it, are precisely the phases a steady-state assessment cannot see.

What coordination is for in an evacuation. In ordinary operation signal coordination is a convenience and efficiency measure: fewer stops, less fuel, smoother progression. In an evacuation it is a queue-management mechanism whose purpose is to keep the corridor discharging through the demand surge. That is a different objective, and it is measured differently: by the height of the congestion crest and by the spread of outcomes across repetitions, not by the mean delay over the whole period. On the mean-delay criterion the coordinated plan here returns 9.5 % at the base demand, which an agency might reasonably judge not worth the effort; on the crest and repeatability criteria it returns 21.4 % and a six to twentyfold reduction in variability, and at half again the demand it is the difference between a network that clears and one that does not. Evacuation signal design deserves to be treated as its own problem rather than as a special case of peak-hour timing.

A spatial asymmetry with distributional consequences. The saving is not spread evenly across the network. More than half of it occurs on one approach, and the sections leading to Highway 417 are essentially unaffected. Where coordination effort is concentrated along a branching corridor is therefore itself a policy decision, with consequences for which population reaches safety first.

6.3 Recommendations for Future Work

The findings suggest several directions, ordered from the most immediately tractable to the longer-term.

Decomposition of the coordinated package. The coordinated plan changes three things at once. A factorial design, holding the common 55 s cycle fixed while toggling the progression offsets and the permissive right-turn treatment, would separate their contributions. Since the replication protocol is now in place, this requires only two further scenario pairs.

Loss of capacity at higher demand. The lane-closure experiment of Section 5.4.9 was run at the base demand, where both plans still cleared the network. Repeating it at 1.25 and 1.5 times the demand, where the isolated plan is already close to failing, would establish whether coordination continues to absorb the loss of a lane once the corridor is itself near capacity, or whether the two effects compound.

Sensitivity to progression speed. The offsets are calibrated to a single platoon speed of 25 km/h. Dispersion around that speed spreads the platoon and erodes the progression

bandwidth, and the present results do not establish how quickly. Re-running the coordinated plan with perturbed desired-speed distributions would bound the robustness of the design.

Departure-time behaviour and multimodal demand. The Rayleigh profile is a shape assumption, not an observed one, and the 35 % of households without a vehicle are excluded from the demand entirely. Replacing the profile with an agent-based departure model, and adding the transit or assisted evacuation of the car-free population, would address the two largest behavioural simplifications in the demand chain.

Pedestrians. Pedestrians constrain the timing design but are not simulated, so the interference they would impose on the permitted right turns, and the safety cost of the always-permissive treatment, are outside what the present model can quantify. Coding pedestrian demand would close this gap.

Adaptive and real-time control. The coordinated plan is fixed-time and designed offline. Since the benefit is concentrated at the demand peak, an adaptive controller that recognised the onset of the surge and applied progression only while it lasts would in principle capture the gain without the early-loading penalty. Whether that is achievable under the communications conditions of a real evacuation is itself a research question.

6.4 Conclusion

The outcome of an urban evacuation depends on the co-design of traffic assignment and signal timing rather than on optimising either in isolation, and the tools conventionally used to evaluate evacuation plans are blind to the benefit that matters most. On a real 17-signal network under census-derived demand, coordinating the signals along the escape corridor reduced total system travel time by 9.5 %, reduced the mean queue by a third, and made the outcome six to twenty times more repeatable. The benefit was negligible while the corridor was lightly loaded and largest at the demand peak, and it arose not from added capacity but from the alignment of green times between junctions. An evacuation signal plan should therefore be judged by how it performs at the crest of the surge, and by how reliably it performs there, not by its average delay.

Bibliography

[1] M. J. Alam, M. A. Habib, and D. Husk, "Evacuation planning for persons with mobility needs: A combined optimization and traffic microsimulation modelling approach," *International Journal of Disaster Risk Reduction*, vol. 80, p. 103164, 2022.

[2] M. Beckmann, C. B. McGuire, and C. B. Winsten, *Studies in the Economics of Transportation*. New Haven, CT: Yale University Press, 1956.

[3] Bureau of Public Roads, *Traffic Assignment Manual for Application with a Large, High Speed Computer*. Washington, D.C.: U.S. Department of Commerce, Bureau of Public Roads, Office of Planning, Urban Planning Division, 1964.

[4] D. Braess, "Über ein Paradoxon aus der Verkehrsplanung," *Unternehmensforschung*, vol. 12, pp. 258–268, 1968. English translation: D. Braess, A. Nagurney, and T. Wakolbinger, "On a paradox of traffic planning," *Transportation Science*, vol. 39, no. 4, pp. 446–450, 2005.

[5] Y.-C. Chiu, J. Bottom, M. Mahut, A. Paz, R. Balakrishna, T. Waller, and J. Hicks, "Dynamic traffic assignment: A primer," Transportation Research Board, Washington, D.C., Transportation Research Circular E-C153, 2011.

[6] D. Gazis, R. Herman, and A. Maradudin, "The problem of the amber signal light in traffic flow," *Operations Research*, vol. 8, no. 1, pp. 112–132, 1960.

[7] P. G. Gipps, "A behavioural car-following model for computer simulation," *Transportation Research Part B: Methodological*, vol. 15, no. 2, pp. 105–111, 1981.

[8] B. D. Greenshields, "A study of traffic capacity," in *Proceedings of the Highway Research Board*, vol. 14, Washington, D.C., 1935, pp. 448–477.

[9] M. J. Lighthill and G. B. Whitham, "On kinematic waves. II. A theory of traffic flow on long crowded roads," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 229, no. 1178, pp. 317–345, 1955.

[10] S. Peeta and A. K. Ziliaskopoulos, "Foundations of dynamic traffic assignment: The past, the present and the future," *Networks and Spatial Economics*, vol. 1, no. 3–4, pp. 233–265, 2001.

[11] A. C. Pigou, *The Economics of Welfare*. London: Macmillan, 1920.

[12] P. I. Richards, "Shock waves on the highway," *Operations Research*, vol. 4, no. 1, pp. 42–51, 1956.

[13] Y. Sheffi, *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1985.

[14] Transportation Research Board, *Highway Capacity Manual, Sixth Edition: A Guide for Multimodal Mobility Analysis*. Washington, D.C.: The National Academies Press, 2016.

[15] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, 2000.

[16] J. G. Wardrop, "Some theoretical aspects of road traffic research," *Proceedings of the Institution of Civil Engineers, Part II*, vol. 1, no. 3, pp. 325–362, 1952.

[17] F. V. Webster, "Traffic signal settings," Road Research Laboratory, London, Road Research Technical Paper No. 39, Her Majesty's Stationery Office, 1958.

[18] V. Pillac, P. Van Hentenryck, and C. Even, "A conflict-based path-generation heuristic for evacuation planning," *Transportation Research Part B: Methodological*, vol. 83, pp. 136–150, 2016.

[19] V. Bayram, "Optimization models for large scale network evacuation planning and management: A literature review," *Surveys in Operations Research and Management Science*, vol. 21, no. 2, pp. 63–84, 2016.

[20] A. J. Pel, M. C. J. Bliemer, and S. P. Hoogendoorn, "A review on travel behaviour modelling in dynamic traffic simulation models for evacuations," *Transportation*, vol. 39, no. 1, pp. 97–123, 2012.

[21] P. Murray-Tuite and B. Wolshon, "Evacuation transportation modeling: An overview of research, development, and practice," *Transportation Research Part C: Emerging Technologies*, vol. 27, pp. 25–45, 2013.

[22] Y. Chen, S. Y. Shafi, and Y.-F. Chen, "Simulation pipeline for traffic evacuation in urban areas and emergency traffic management policy improvements through case studies," *Transportation Research Interdisciplinary Perspectives*, vol. 7, p. 100210, 2020.

[23] B. Wolshon, "Planning for the evacuation of New Orleans," *ITE Journal*, vol. 72, no. 2, pp. 44–49, 2002.

[24] G. Theodoulou and B. Wolshon, "Alternative methods to increase the effectiveness of freeway contraflow evacuation," *Transportation Research Record*, vol. 1865, no. 1, pp. 48–56, 2004.

[25] H. Naghawi and B. Wolshon, "Performance of traffic networks during multimodal evacuations: Simulation-based assessment," *Natural Hazards Review*, vol. 13, no. 3, pp. 196–204, 2012.

[26] Z. Zhang, K. Spansel, and B. Wolshon, "Effect of phased evacuations in megaregion highway networks," *Transportation Research Record*, vol. 2459, no. 1, pp. 101–109, 2014.

[27] S. Peeta and H. S. Mahmassani, "System optimal and user equilibrium time-dependent traffic assignment in congested networks," *Annals of Operations Research*, vol. 60, no. 1, pp. 81–113, 1995.

[28] A. K. Ziliaskopoulos, "A linear programming model for the single destination system optimum dynamic traffic assignment problem," *Transportation Science*, vol. 34, no. 1, pp. 37–49, 2000.

[29] V. Morandi, "Bridging the user equilibrium and the system optimum in static traffic assignment: A review," *4OR*, vol. 22, pp. 89–119, 2024.

[30] J. D. C. Little, M. D. Kelson, and N. H. Gartner, "MAXBAND: A program for setting signals on arteries and triangular networks," *Transportation Research Record*, vol. 795, pp. 40–46, 1981.

[31] N. H. Gartner, S. F. Assmann, F. Lasaga, and D. L. Hou, "A multi-band approach to arterial traffic signal optimization," *Transportation Research Part B: Methodological*, vol. 25, no. 1, pp. 55–74, 1991.

[32] C. Zhang, Y. Xie, and N. H. Gartner, "AM-Band: An asymmetrical multi-band model for arterial traffic signal coordination," *Transportation Research Part C: Emerging Technologies*, vol. 58, pp. 515–531, 2015.

[33] R. P. Roess, E. S. Prassas, and W. R. McShane, *Traffic Engineering*, 5th ed. Hoboken, NJ: Pearson, 2019.

[34] P. Koonce et al., *Traffic Signal Timing Manual*, Report No. FHWA-HOP-08-024. Washington, D.C.: Federal Highway Administration, 2008.

[35] B. Wolshon, "Evacuation planning and engineering for Hurricane Katrina," *The Bridge*, vol. 36, no. 1, pp. 27–34, 2006.

[36] Statistics Canada, *Census Profile, 2011 Census of Population*. Ottawa: Statistics Canada Catalogue no. 98-316-XWE, 2012.

[37] Statistics Canada, "Population and dwelling counts: Canada, provinces and territories, census divisions and census subdivisions," 2021 Census of Population, Statistics Canada Catalogue no. 98-10-0002-01, 2022.

[38] Statistics Canada, *Census Profile, 2021 Census of Population: Ottawa, Census Subdivision*. Ottawa: Statistics Canada Catalogue no. 98-316-X2021001, 2022.

[39] Ottawa Neighbourhood Study, "Sandy Hill neighbourhood profile," University of Ottawa, demographic estimates derived from the Statistics Canada 2021 Census via Environics Analytics, 2023. [Online]. Available: https://www.neighbourhoodstudy.ca

[40] TRANS Committee, *2022 Origin–Destination Survey: National Capital Region*. Ottawa: TRANS Committee (City of Ottawa, Ville de Gatineau, National Capital Commission, OC Transpo, STO), 2023.

[41] G. Gower, "Transportation trends in Ottawa: vehicle ownership," online commentary on the 2022 TRANS Origin–Destination Survey, 2024.

[42] Aimsun, *Aimsun Next Users Manual*, version 24.0.2. Barcelona, Spain: Aimsun SLU, 2024.

[43] T. J. Cova and J. P. Johnson, "A network flow model for lane-based evacuation routing," *Transportation Research Part A: Policy and Practice*, vol. 37, no. 7, pp. 579–604, 2003.

[44] B. Wolshon, "'One-way-out': Contraflow freeway operation for hurricane evacuation," *Natural Hazards Review*, vol. 2, no. 3, pp. 105–112, 2001.

[45] H. Tuydes and A. Ziliaskopoulos, "Tabu-based heuristic approach for optimization of network evacuation contraflow," *Transportation Research Record*, no. 1964, pp. 157–168, 2006.

[46] C. Xie, D.-Y. Lin, and S. T. Waller, "A dynamic evacuation network optimization problem with lane reversal and crossing elimination strategies," *Transportation Research Part E: Logistics and Transportation Review*, vol. 46, no. 3, pp. 295–316, 2010.

[47] C. Xie and M. A. Turnquist, "Lane-based evacuation network optimization: An integrated Lagrangian relaxation and tabu search approach," *Transportation Research Part C: Emerging Technologies*, vol. 19, pp. 40–63, 2011.

[48] Y. Liu, X. Lai, and G.-L. Chang, "Cell-based network optimization model for staged evacuation planning under emergencies," *Transportation Research Record*, no. 1964, pp. 127–135, 2006.

[49] B. Wolshon, E. Urbina Hamilton, M. Levitan, and C. Wilmot, "Review of policies and practices for hurricane evacuation. II: Traffic operations, management, and control," *Natural Hazards Review*, vol. 6, no. 3, pp. 143–161, 2005.

[50] M. Chen, L. Chen, and E. Miller-Hooks, "Traffic signal timing for urban evacuation," *Journal of Urban Planning and Development*, vol. 133, no. 1, pp. 30–42, 2007.

[51] S. A. Parr and E. Kaisar, "Critical intersection signal optimization during urban evacuation utilizing dynamic programming," *Journal of Transportation Safety & Security*, vol. 3, no. 1, pp. 59–76, 2011.

[52] S. A. Parr, B. Wolshon, and P. Murray-Tuite, "Unconventional intersection control strategies for urban evacuation," *Transportation Research Record*, no. 2599, pp. 52–62, 2016.

[53] N. Nassir, M. Hickman, H. Zheng, and Y.-C. Chiu, "Network flow solution method for optimal evacuation traffic routing and signal control with nonuniform threat," *Transportation Research Record*, no. 2459, pp. 54–62, 2014.