Title

Impacts of Adaptive Cruise Control on Urban Congestion: A Queuing Perspective

Author

Lapardhaja, Servet

Publication Date

2023

Peer reviewed|Thesis/dissertation

Impacts of Adaptive Cruise Control on Urban Congestion: A Queuing Perspective

By

Servet Lapardhaja

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael Cassidy, Chair
Professor Alexander Skabardonis
Professor John Radke

Fall 2023

Impacts of Adaptive Cruise Control on Urban Congestion: A Queuing Perspective

Abstract

Impacts of Adaptive Cruise Control on Urban Congestion: A Queuing Perspective

by

Servet Lapardhaja

Doctor of Philosophy in Engineering – Civil and Environmental Engineering

University of California, Berkeley

Professor Michael Cassidy, Chair

The morning rush in two idealized settings are forecast into the future, assuming that most drivers will engage Adaptive Cruise Control (ACC) while driving their cars. To this end, careful measurements of ACC-equipped cars traveling on roads and a highway confirm an earlier finding reported by others: for a given traffic speed in ACC-rich congestion, the density tends to be smaller than in present-day congestion where all vehicles are manually operated. All else equal, the lower congested densities (i.e., larger car spacings), mean that queues will be less-densely-packed and will expand over longer physical distances in the future, as ACC-equipped vehicles become more prevalent. These uncompacted queues pose problems for urban areas, where queue storage is often already a problem during rush hours.

Simulations calibrated using the field-measured data confirm this concern and contradict optimistic predictions of how ACC may lead to a congestion-free future. In a setting with already moderately high congestion and long-street links inspired by Downtown Los Angeles, it is predicted that networkwide vehicle hours traveled (VHT) will increase by up to 12% compared to present-day levels, despite assumptions that are largely favorable to ACC. In a setting inspired by Midtown Manhattan, with short-street links and already high congestion levels, networkwide VHT is predicted to increase by as much as 87%. The higher bottleneck capacities often promised by advocates of ACC are shown to be irrelevant when spillover queues restrict flows from reaching those capacities. Interventions tested through simulations include adjusting onboard ACC controllers to produce lower jam spacings, to prevent a problematic future.

To my family, for their never-ending support.

Table of contents

List of figures

List of tables

Glossary

Abbreviation or Term	Description or Explanation
ACC	Adaptive Cruise Control
ACC-controlled	Operated using Adaptive Cruise Control
ACC-equipped	With onboard Adaptive Cruise Control
Capacity	The maximum flow that can be maintained on a road segment absent the influence of any queue that arrived from downstream
Car-following model	A mathematical representation that describes the behavior of individual vehicles in traffic by predicting how a vehicle follows another vehicle
Cycle	A sequence of high-low-high speeds in a field experiment
Density	The number of cars per unit length of a roadway or area, often measured in the number of cars per lane-kilometer (cars/lane-km)
EV	Electric Vehicle
Flow	The number of cars passing a point in a given time period, often measured in the number of cars per hour per lane (cars/h/lane)
Following-distance setting	Following-distance option, selected from a menu of short, medium, or long options, with certain vehicle models offering an extra-long option
Fundamental diagram	Bivariate plot of density and flow
GPS	Global Positioning System
Gridlock	A state of the system under which traffic in the entire network or a portion of the network comes to a complete standstill with zero (or minimal) flow
Headway	The time that elapses between the arrival of the leading vehicle and the following vehicle at a designated point, often reported in units of seconds
Human-controlled	Operated manually by human drivers
ICE	Internal Combustion Engine
IDM	Intelligent Driver Model
Jam	A state in traffic where vehicles are stopped
LiDAR	Light Detection and Ranging
OBD	Onboard Diagnostics
OD	Origin-Destination
Platoon	A group of vehicles traveling together in an experiment
RADAR	Radio Detection and Ranging

Spacing	The distance between the front bumper of the leading vehicle and the front bumper of the following vehicle, often reported in units of meters
Steady-state	A condition where each vehicle in the traffic stream maintains the same or nearly the same constant speed
USDOT	US Department of Transportation
VHT	Vehicle Hours Traveled

Acknowledgments

The culmination of my PhD journey symbolizes an extraordinary odyssey, one that has been enriched by the invaluable support, guidance, and inspiration of numerous exceptional individuals. As I pause to contemplate this monumental achievement, my heart brims with profound gratitude for those who have played a pivotal role in shaping my academic and personal growth.

First and foremost, I extend my deepest appreciation to my advisor, Professor Michael Cassidy. His unwavering commitment to my intellectual development, coupled with his wealth of expertise, has been the bedrock of my successful PhD journey. Under his guidance, I have honed my ability to formulate research problems, approach them systematically, and craft compelling narratives. Professor Cassidy's mentorship has bestowed upon me invaluable lessons that have transcended the academic realm.

In parallel, I extend my heartfelt gratitude to Dr. David Kan, an integral part of my PhD journey right from the outset. Dr. Kan's guidance, unflagging support, and the treasure trove of ideas he generously shared have broadened my horizons and enriched the trajectory of my research. Without his enduring collaboration over these years, my PhD would not have taken its current form.

My experience as a graduate student was profoundly enriched by the presence of exceptional colleagues, with one individual standing out distinctly—Dr. Jean Doig. His unwavering dedication and readiness to go the extra mile in aiding not only me but also our entire research group have left an indelible imprint on my academic voyage. His willingness to provide guidance and support in our shared quest for knowledge has rendered him an invaluable colleague.

In addition to these esteemed individuals, I must also pay tribute to the profound influence of Professor Alexander Skabardonis. Our initial encounter, at the inception of my PhD program, marked a pivotal moment in my academic journey.

Furthermore, I wish to extend my gratitude to Dr. Alexander Kurzhanskiy from California PATH. The opportunity to collaborate with Dr. Kurzhanskiy on various fascinating projects has been both enriching and enlightening.

Beyond the confines of academia, my parents, Esma Lapardhaja and Sali Lapardhaja, have stood as unwavering pillars of strength. Their selfless sacrifices and steadfast support have propelled me forward. Their unwavering faith in my potential has served as a perpetual wellspring of motivation, while their guidance has played a pivotal role in shaping my character. I fervently hope that this PhD stands as a testament to the profound impact of their love and wisdom.

This academic journey would have been an insurmountable challenge without the collective support of my family, friends, and colleagues who have steadfastly stood by my side. Each of you has played an indispensable role in my success, and I am eternally grateful.

Finally, I wish to express my heartfelt appreciation to the National Institute of Congestion Reduction for their generous support, without which this research would not have been possible. Your investment in my work has been the cornerstone of reaching this significant milestone.

1 Introduction

A vehicle equipped with Adaptive Cruise Control (ACC) features built-in sensors (e.g., RADAR or LiDAR) that continuously track the actions of the vehicle in front, coupled with a control mechanism that automatically adapts the speed of the vehicle in response to the lead vehicle's actions. When engaged, ACC technology automatically sustains car-following-distances appropriate for the vehicle's speed. These distances are based on the driver's preferred following-distance, selected from a menu of short, medium-, or long settings, with certain vehicle models offering an extra-long setting. ACC constitutes a fundamental component of present-day vehicle automation (Gunter et al., 2020; Makridis et al., 2021; Shang & Stern, 2021), and is an anticipated component for self-driving vehicles of the future (Tesla, 2023).

There have been objections raised against enabling ACC when traveling in congested, slow-moving traffic, primarily due to safety concerns (GMC, 2023; Honda, 2022; Sparks, 2022; Volvo, 2020). However, there are seemingly stronger advocates in favor of more unrestricted use of this technology. The US Department of Transportation (USDOT), for example, views ACC as a means to improve "throughput" and various other traffic performance metrics (USDOT, 2019). These assertions suggest part of ACC's value lies in its utilization in congested rush-hour traffic when good performance is most desirable. Drivers themselves appear receptive to this broader use of the technology, as indicated by a recent federal study (USDOT, 2022), where most participants expressed comfort with all following-distances generated by ACC, even at speeds as low as 40 km/h, and with the short setting as their preferred following-distance.

The prevailing trend appears to lean towards widespread adoption of ACC, whether traffic conditions are congested or not. Therefore, the present work assumes that drivers of ACC-equipped vehicles will likely utilize this technology regardless of traffic conditions, both in nearer term- and in more distant futures. Based on this assumption, it is illustrated why optimistic predictions regarding ACC's future impacts on traffic "throughput" or other performance metrics are not only overstated, but in the absence of intervention, why one should anticipate a significant increase in urban congestion due to ACC. The concern seems particularly pressing in older, denser cities with limited queue storage space and that are already experiencing considerable congestion during a rush.

In light of this concern, the dissertation utilizes careful measurements from presently conducted field experiments of ACC-controlled cars traveling on roads and a highway. The findings reveal that, for a given traffic speed, the density in ACC-rich congestion tends to be smaller than in present-day congestion, where all vehicles are fully human-controlled. All else equal, the smaller congested densities (i.e., larger vehicle spacings) mean that queues will be less compacted and will thus expand over greater physical distances in the future, as more and more ACC-controlled vehicles enter the scene. These uncompacted queues spell trouble for cities, where queue storage is often already a problem during a rush. Consequently, the improved "throughput" or higher bottleneck capacities, often promised by advocates of ACC are shown to be irrelevant when spillover queues restrict flows from reaching those capacities. Simulation-based case studies, inspired by two distinct urban settings—Downtown Los Angeles and Midtown Manhattan—and calibrated to the field-measured data, confirm this looming concern and run contrary to glowing predictions of how ACC may lead to a congestion-free future.

Before delving into the study, Section 1.1. presents the motivation driving this research. Section 1.2 describes the research questions that will be answered in this dissertation, and Section 1.3 provides the organization of chapters that follow.

1.1 Motivation

Careful measurements of ACC-controlled cars in this and previous research efforts indicate that, at a given traffic speed, ACC-rich congestion tends to exhibit lower densities compared to present-day congestion, where virtually all vehicles are human-controlled. The smaller congested densities (i.e., larger vehicle spacings) mean that queues will be less compacted and will thus expand over greater physical distances in the future, as more and more ACC-controlled cars enter the scene. These less compacted queues pose problems, particularly for older, densely populated cities, where queue storage is often already a problem during a rush. Higher bottleneck capacities commonly promised of ACC become irrelevant when queues that spillover from one link to the next constrain a bottleneck's flow from reaching those capacities. In turn, queue spillover triggers a vicious cycle of worsening congestion over time, resulting in diminishing trip-completion rates and soaring delays, a process called gridlock (Daganzo, 1996, 2007; Mahmassani et al., 2013). It becomes apparent that a determinant of future urban congestion will be ACC's impact on queue expansions, and unfortunately, the outlook in this regard appears unfavorable. With this concern in mind, the dissertation delves deeper into the implications of ACC-rich traffic on queue storage and urban congestion. This exploration is conducted through case studies loosely inspired by two distinct urban settings: Downtown Los Angeles and Midtown Manhattan.

1.2 Research Questions

Several pertinent questions arise concerning the expanding market presence of ACC. This dissertation endeavors to address the following research questions:

- What are the potential impacts of ACC's increasing market share on queue storage and urban congestion in the nearer term- and more distant futures?
- How does the utilization of ACC impact congestion in distinct urban settings, akin to Midtown Manhattan, an older city with short block lengths, and Downtown Los Angeles, a newer car centric city characterized by larger block lengths?
- What interventions could be implemented to manage future cases of ACC-rich traffic in urban areas?

1.3 Dissertation Organization

The remaining chapters of this dissertation address the research questions outlined in Section 1.2.

- Chapter 2 provides a literature review on both theoretical studies of ACC's impacts on traffic and field experiments, both of which offer insights for the present study. The chapter also discusses the potential implications in urban areas stemming from lower queued densities in ACC-rich traffic, laying the groundwork for subsequent chapters.
- Chapter 3 discusses the field experiments presently conducted to observe how ACC-equipped cars adjust their following-distances in response to their leaders, while traveling on real facilities. The experiments include ACC-equipped cars powered by internal combustion engines (ICE), and ACC-controlled, battery-powered electric cars (EVs).
- Chapter 4 describes the process of model calibration in the Aimsun simulation platform.
- Chapter 5 provides a detailed description of the approach used to simulate both present-day conditions and future scenarios in two urban settings inspired by Los Angeles and Manhattan. Networkwide performance predictions for nearer-term and more-distant futures are compared with the present-day. Interventions to address ACC-induced congestion are discussed.
- Chapter 6 concludes the dissertation by summarizing key findings and delves into potential challenges that could emerge in response to the proposed interventions. Moreover, this chapter outlines future research directions complementing the present study, aiming to offer further evidence of the findings of this dissertation and stimulate further discussions about interventions to manage ACC-induced congestion.

2 Literature Review

This chapter provides background on the topics at hand. Section 2.1 focuses on the history and evolution of ACC. Section 2.2 examines theoretical studies in which car-following models were used to investigate the impacts of ACC on bottleneck capacity. The studies yielded contradictory findings, which underscore the need for more empirical study. Section 2.3. delves into more recent studies that eschewed theoretical models in favor of real-world field data. These reveal that ACC, when engaged with shorter following-distances, can increase queue discharge flows at isolated bottlenecks, but raise concerns about ACC's impacts on queue expansions in congested traffic. Section 2.4 examines the problems posed by less-densely-packed, expanded ACC-vehicle queues in rush hour traffic. It serves as a troubling precursor on what to expect from the dissertation's investigation into this matter.

2.1 Evolution of Adaptive Cruise Control

Since the inception of the automobile, various driver assistance features have been introduced to enhance safety and reduce driver fatigue. Many of these features are closely tied to the automation of driving tasks. One widely adopted automation feature is the cruise control system, which takes control of the vehicle's throttle and maintains the vehicle's speed at the driver's chosen value.

Recent developments in sensor technology have led to the introduction of ACC, representing an advanced iteration of traditional cruise control. While conventional cruise control simply maintains the driver-set vehicle speed, ACC goes a step further by also maintaining a suitable relative distance from the lead vehicle. When ACC detects that the lead vehicle is traveling at a slower speed than the driver's desired speed, the ACC system automatically decelerates the host vehicle to keep an appropriate car-following distance based on the driver's preferred following-distance setting. These preferences consist of options for short, medium-, or long following-distances, and in some vehicle models, an extra-long option is also available.

Nowadays, the ACC system is further extended to incorporate a "stop-and-go" feature, specifically designed for urban driving conditions. This extension provides the capability for the host vehicle to automatically come to a complete stop when it detects a stationary lead vehicle ahead.

2.2 Findings from Theoretical Studies

Initial studies on the impact of ACC in traffic have primarily centered around bottleneck capacity, i.e., queue discharge flow. These studies tend to be theoretical in nature, relying on car-following models originally designed for human-controlled driving (e.g., Gipps 1981; Treiber et al. 2000). The models were adapted, with parameters sometimes calibrated based on ACC-responses observed in field experiments, while in other cases, the models relied solely on speculative assumptions regarding parameter values.

When applied to individual, isolated bottlenecks, these models often predicted that higher levels of ACC adoption in future traffic will result in increased bottleneck discharge flows compared to what is currently observed (Goñi-Ros et al., 2019; Kesting et al., 2008; Papacharalampous et al., 2015; Talebpour & Mahmassani, 2016). For instance, Kesting et al. (2008) simulated a three-lane freeway section with an on-ramp bottleneck for a variety of ACC market shares. The study used the intelligent driver model (IDM) (Treiber et al., 2000), to model ACC-equipped cars using some untested parameters. Kesting et al. (2008) reported an increased bottleneck capacity when simulating a 25% ACC market share, resulting in the complete elimination of traffic congestion caused by the bottleneck. In the same vein, another study (Talebpour & Mahmassani, 2016) simulated a hypothetical one-lane freeway with an on-ramp located in the middle of the segment, using the primitive simulation model MIXIC (Van Arem & De Vos, 1997). Parameters values based solely on speculation (Van Arem et al., 2006) were used to simulate ACC-equipped cars. Simulations showed that as the ACC market share grew, bottleneck capacity increased.

Some less optimistic predictions of ACC impacts have also been reported in the literature. For instance, Vander Werf et al. (2002) studied the matter using Monte Carlo simulations with parameters for ACC controllers that were intended to represent best case scenarios for the capacity impacts of ACC. The simulated site was a single freeway lane consisting of a single-lane off-ramp followed immediately by a single-lane on-ramp. Based on the study's simulation results, the authors conclude that ACC, even under the most favorable conditions, does not have a significant impact on capacity. In yet another recent effort to investigate the impact of ACC on freeway bottleneck capacity, Shang and Stern (2021) calibrated the IDM car-following model using field data collected from ACC-equipped cars. For the analysis, a one-lane freeway with an on-ramp was simulated. The results of the study indicate that commercially available ACC cars lead to a reduction of bottleneck capacity by up to 35%.

The inconsistency in outcomes across the theoretical studies seems a result of using untested car-following models, often with parameters that lack empirical support. The predictions from these models have frequently served as the ground truth in previous research, enabling the exploration of the effects of ACC on various lane-changing maneuvers, string stability, and other traffic characteristics that were not considered in the small-scale field experiments performed to date (Bose & Ioannou, 2003; Ntousakis et al., 2015; Shang & Stern, 2021).

A notable concern arises due to the absence of testing these models against real data from traffic predominantly composed of ACC-equipped vehicles. Evaluations of numerous predictions furnished in a simulation study (James et al., 2019), which varies parametrically the mix of ACC- and human-controlled vehicles in traffic, lead us to question the realism of these models.

The study conducted in James et al. (2019) performed microscopic simulations in VISSIM, utilizing four car-following models: Autonomous Adaptive Cruise Control (AACC) (VanderWerf et al., 2001), the IDM (Treiber et al., 2000), the California PATH model (Milanés & Shladover, 2014), and the TU Delft model (Xiao et al., 2017). Pathologies can be observed in the resulting predictions. Initially, both the AACC and TU Delft model are shown logically, to predict that opting for short following-distances leads to higher discharge flows at bottlenecks. However, these discharge flows start to steadily decrease when the market share of ACC-equipped cars exceeds 50%, as if an excess of short following-distances somehow leads to adverse effects. Adding to the peculiarity, the calibrated but untested California PATH model nonsensically predicts that the selection of short following-distance settings (and thus small headways) produces a lower queue discharge flow than when choosing long following-distance settings. Of course, then, the California PATH model predicts that bottleneck capacities diminish as the share of ACC cars with short settings in the traffic stream steadily increases. Of further concern, the short setting results in a headway that is smaller than the average selected by human drivers and yet the model predicts that these ACC capacities are lower than in human-controlled traffic. Like the California PATH model, IDM predicts that steadily growing numbers of ACC-equipped vehicles with short following distances result in steadily lower queue discharge flows. To make matters worse, IDM predicts that ACC generates unrealistically large jam densities. These density predictions would be seen as positive news for cities. However, field-measured data present a different outcome, as discussed below and in other sections of this dissertation.

2.3 Findings from Field Experiments

Recent studies in Shi and Li (2021) and Li et al. (2022) utilized field data as the ground truth, rather than relying on predictions from car-following models. These latter two studies collected real-world ACC-response data from small fleets of 2 or 3 ACC-equipped cars operating in real-world settings.

The experiments in Shi and Li (2021) were conducted within a single lane on a four-lane segment of SR-56 in Florida. The study utilized two ACC-equipped Lincoln MKZs, model years 2016 and 2017. Geographic coordinates of each vehicle were recorded using GPS receivers, and data were collected for four following-distance settings, ranging from short to extra-long. During these experiments, the leader of the platoon initially accelerated to a pre-selected higher speed, subsequently decelerated to a lower speed that had been selected *a priori* from a menu reflecting various congestion levels, and then returned to the initial higher speed. The menu of the higher speeds ranged from 55 mph to 35 mph, while the menu of lower speeds ranged from 53 mph to 25 mph. Additionally, the study incorporated a second dataset from Gunter et al. (2020), which also utilized two ACC-equipped cars. This latter dataset included a range of high speeds from 75 mph to 65 mph and lower speeds from 55 mph to 35 mph, with data collected for two following-distance settings (short and long).

The study in Li et al. (2022) conducted field experiments employing a three-vehicle platoon on undisclosed public highways and rural roads. A GPS device was installed on each car to collect location and velocity data. The data were collected only for two following-distance settings: short and long. Similar to the above cited field experiments, the platoon leader initially traveled at a high speed that was pre-selected, then decelerated to a predefined lower speed, before eventually returning to the initial high speed. The menu of high speeds ranged from 70 mph to 65 mph, while the menu of lower speeds ranged from 45 mph to 35 mph. Li et al. (2022) also incorporated field data from Gunter et al.

(2020) and Makridis et al. (2021). The latter dataset encompassed high speeds ranging from 60 mph to 35 mph, and lower speeds ranging from 30 mph to 15 mph. The latter data were also collected for short and long following-distance settings.

Both Shi and Li (2021) and Li et al. (2022) used the data to construct bivariate plots of densities and flows, the inverses of field-measured car spacings and headways, respectively. The plots clearly indicate that the selections of short- or medium following-distances engendered queue discharge flows that were higher than those typically observed in present-day traffic conditions. Conversely, choosing longer following-distances resulted in smaller discharge flows.

These empirical findings, along with a prior finding that approximately 70% of ACC-equipped drivers select small- or medium following-distance, suggest that ACC is likely to have favorable impacts on isolated bottlenecks. Unfortunately, most cities typically contend with multiple bottlenecks that can interact in ways that hinder traffic flow when queues grow long (Carlson et al., 2010, 2014; Cassidy et al., 2002; Daganzo, 1998, 1998; Daganzo et al., 1999; Kim & Cassidy, 2012). The potential higher bottleneck capacities promised by ACC become irrelevant when queues spillover from one link to the next, constraining a bottleneck's flow from reaching those capacities. This suggests, that a significant factor influencing future urban congestion will be ACC's impact on queue expansions, and unfortunately, the outlook in this regard appears unfavorable, as indicated by clues found in the studies in Shi and Li (2021) and Li et al. (2022).

The density-flow data plotted in both those studies include measurements taken at slower speeds, commensurate with light to moderate congestion. To estimate jam densities, best-fit lines were extrapolated through those congested data. The resulting estimates varied, reaching as high as 90 cars/lane-km in Shi and Li (2021) and as low as 50 cars/lane-km in Li et al. (2022). This range encompasses jams characterized by noticeably uncompressed cars. For instance, the lower estimate suggests that each stopped car occupies a staggering 20 meters of longitudinal road space on average. These estimations have been influenced by errors, which could have arisen due to the heroic extrapolations made in the absence of severely congested data, and because some of the lower-speed data from Makridis et al. (2021) were collected under short-lived, non-steady-state conditions. Not surprisingly, a somewhat higher jam density of 101 cars/lane-km was separately estimated in Li et al. (2022) based on an unspecified number of jam spacings sampled from a Tesla Autopilot dataset.

Setting aside the errors, the above estimates underscore something of significance and concern. They suggest that, for any given speed in congestion, most of the densities generated by ACC appear to be lower than those observed when vehicles are solely under the control of human drivers. Notably, this finding seems to have been overlooked in earlier studies that primarily focused on discharge flows through isolated bottlenecks (see Section 2.2). In a prescient response to this finding on density, the authors of Li et al. (2022) caution readers about potential issues with queue storage in the future, when ACC-equipped cars become more prevalent. This caution is well-founded, for the reasons discussed in Section 2.4.

2.4 Queued Density and the Implications for ACC-Dominated Traffic

Issues are known to occur when vehicle queues expand significantly during a rush and spillover to multiple upstream links, e.g., Daganzo (1998). On freeways, for example, spillover queues can obstruct off-ramps and deprive them of exit flows, thereby adding to overall system delays (Cassidy et al., 2002; Kim & Cassidy, 2012; Newell, 1993). Even greater problems can emerge on city streets, where long queues often wind around short city blocks. These queues spread to other links, triggering a vicious cycle of worsening congestion over time, with ever diminishing trip-completion rates and skyrocketing delays, a process known as gridlock (Daganzo, 1996, 2007; Mahmassani et al., 2013).

As noted in Li et al. (2022), the lower densities (i.e., larger vehicle spacings) in congestion indicate that ACC-dominated queues in the future will be less tightly packed than at present. These expanded queues will give rise to more of the issues described above. With this concern in mind, this dissertation investigates the matter further.

3 Field Studies

A battery of field experiments were conducted as part of the present research, to measure how ACC-controlled cars adjusted their spacings in response to their leaders, while operating on real roads. These experiments were akin to those in Shi and Li (2021) and Li et al. (2022), but this time the field data: included measurements at lower speeds commensurate with more severe congestion and included numerous measurements of jammed (i.e., stopped) conditions; and all data were extracted from steady-state conditions. The resulting estimates of queued densities (the inverses of car spacings), while generally larger than those reported in the aforementioned studies, still turn out to be small enough to produce damaging spillover queues, as ACC market shares grow in the future; see Chapter 5.

The field experiments took place during off-peak hours when traffic was light, allowing the test cars to safely change their speeds, to emulate various levels of congestion. Several roads and a highway were utilized in these field experiments, as ACC responses are not influenced by the type of facility on which they operate (Xiao & Gao, 2010). Two separate sets of experiments were conducted, each using small platoons of ACC-equipped cars. One set involved ACC-equipped cars powered by internal combustion engines (ICE), while the other involved ACC-controlled battery-powered electric cars (EVs).

In each set of experiments, and as in previous experiments described in Section 2.3, the platoon leader initially traveled at a high speed, pre-selected from a menu of one or more free-flow speeds, each to emulate what might occur on a particular facility type. Subsequently, the lead car reduced its speed to a lower value, selected *a priori* from a menu reflecting various congestion levels. The lead car then returned to the initial high speed, and the entire sequence of prescribed speeds was regulated by the lead car's ACC controller onboard. We will refer to one sequence of high-low-high speeds as a "cycle".

Each combination of high and low speeds was tested for every available option for preferred following-distance, and each set of combined speeds and preferred following-distance was performed for a minimum of 12 cycles. To ensure that the data collected represented steady-state conditions, high- and low speeds were each maintained for 10s or more, and measurements recorded at short intervals (0.04s or 0.2s) were used in the analysis only when the speed of the leader and its follower was approximately the same. Averages were used to estimate the ACC response for each prolonged speed state.

In total, 1,264 data points were gathered across all experiments. It is worth noting that certain aspects of the experiments involving ICE-platoons differed from those involving EVs. These distinctions are described below, beginning with the ICE experiments.

3.1 ICE-Platoons

In this set of experiments, 3-car platoons were created using 2020 model-year Toyota Corollas powered by ICE and equipped with ACC (See Figure 1). These experiments were conducted at two locations in South Florida: an 8-km stretch in both directions of the US

441 Highway (between Atlantic and Boynton Beach Blvds.) in Delray Beach; and a 4-km stretch in both directions of Flying Cows Rd. (between US 98 Highway and Rustic Rd.) in Wellington.

Each experiment typically commenced on the shoulder of a facility, where each of the platoon's two follower-vehicles were positioned 1.4 meters apart from their respective leaders, as measured from the rear of the leader to the front of the follower. Once in a travel lane, the platoon leader followed a specific sequence: it accelerated to the speed of 88 km/h, decelerated to one of the predefined lower speeds from the menu of {72, 56, 40, 24, 0} (km/h), and then returned to the initial speed of 88 km/h, initiating a new cycle shortly thereafter. Decelerations to zero by the platoon leader were performed only when traveling on Flying Cows Rd., but not while on US 441. Speedometer readings for all three vehicles in the platoon were automatically logged at intervals of 0.2s, and the resulting estimates of distances traveled during each time step were recorded using Onboard Diagnostics (OBD) II data loggers (see Figure 2).

To limit systematic measurement errors in each vehicle's speedometer, which might have accumulated with distance traveled, a re-initialization procedure was implemented. The procedure involved re-establishing the 1.4-meter separation between platooned vehicles every two cycles and did so in the manner described above.

Figure 1 ACC-equipped ICE-powered cars

Figure 2 OBD II data logger used for collecting data from test vehicles

3.2 EV-Platoons

The second set of experiments involved 2-car platoons consisting of ACC-equipped EVs: a 2021 Toyota Camry as the lead vehicle and a 2022 Hyundai IONIQ 5 as the follower (see Figure 3). Notably, the Hyundai IONIQ 5 had an additional setting for preferred following-distance: extra-long. The experiments were carried out on remote, lightly trafficked sections of three public roads located in Northern California, specifically on approximately 10-km stretches of Pedrick, Robben, and Sikes Roads in the rural community of Dixon (see Figure 4).

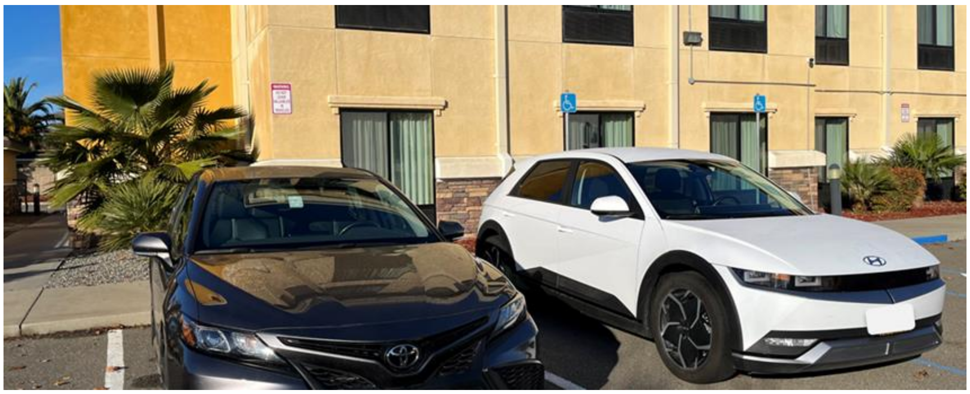

Figure 3 Test vehicles (left: 2021 Toyota Camry, right: 2022 Hyundai IONIQ 5)

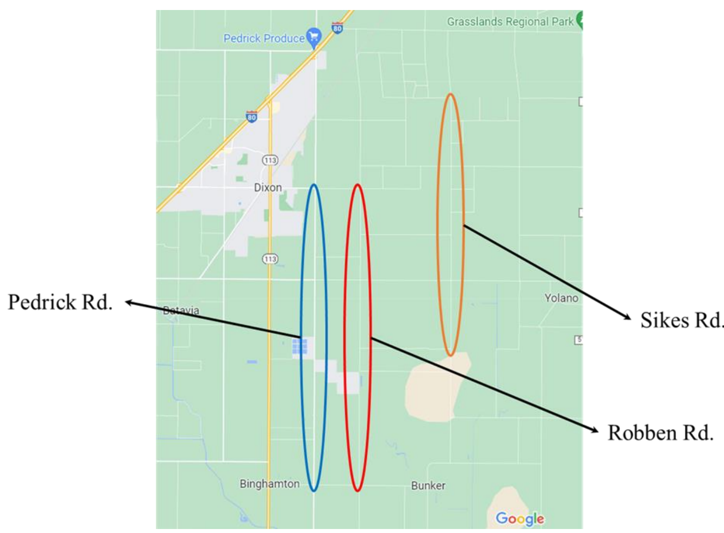

Figure 4 Test sites for EV ACC experiments

In this set of experiments, the lead vehicle initially decelerated from a pre-selected speed from a menu of {95, 88, 72, 56} (km/h) to one of the lower speeds from the menu of {72, 56, 20, 24, 0} (km/h), and then eventually returned to its initial high speed. During these experiments, lower-cost yet highly precise GPS devices manufactured by Racebox had by this time entered the marketplace and were now used to measure the geographic coordinates of each vehicle at 0.04s intervals (see Figure 5). These more accurate devices eliminated the need for periodic re-initialization of vehicle separations. Since the platoons consisted of only 2 cars instead of the previous 3, measurements were collected over additional cycles to ensure that the number of data points in each set of experiments (ICE and EV) remained comparable.

Figure 5 Racebox GPS mounted on test vehicle

3.3 Data Processing

Vehicle trajectories were constructed from the field-measured data and were used to extract joint estimates of density and flow, similar to the approach taken in Shi and Li (2021) and Li et al. (2022). To begin, a thorough examination was conducted to identify any potential accumulation of errors in the measurements of ICE-platoons. Figure 6 provides clarity in this regard. It presents median values of densities and flows derived from both the first cycles (illustrated by unshaded data points) and the second cycles (represented by shaded points). These data points are further categorized based on steady-state speeds (depicted by light lines radiating from the origin) and preferred following-distances. Best-fit polynomial curves are included in the figure to aid in visual inspection.

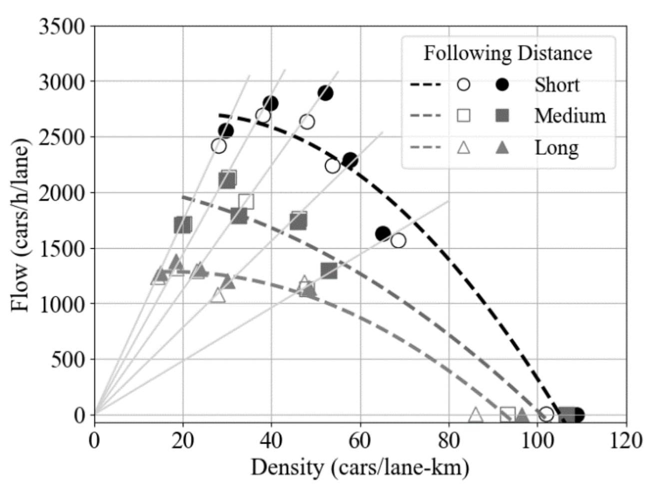

Figure 6 Median measurements for Cycle 1 (unshaded) and Cycle 2 (shaded) ICE Experiments

Figure 6 clearly illustrates that there are differences between the unshaded and shaded data points, representing the first and second cycles, both in their densities and flows. However, these differences are relatively small when compared to the displacements of the data points from their respective best-fit curves. In essence, the behavior of ACC responses demonstrates a degree of random variation. Consequently, the observed distinctions between the estimates from the first and second cycles, as depicted in Figure 6, seem within the expected range and do not suggest any systematic accumulation of measurement errors. Consequently, all the data collected during the field experiments, including those from ICE-platoons obtained in the second cycles, were used.

The entirety of the data is presented in Figures 7a and 7b for the ICE- and EV-platoons, respectively. The slightly enlarged and shaded data points represent median values for each group of data with the same speed and preferred following-distance setting, with best-fit curves shown for these medians.

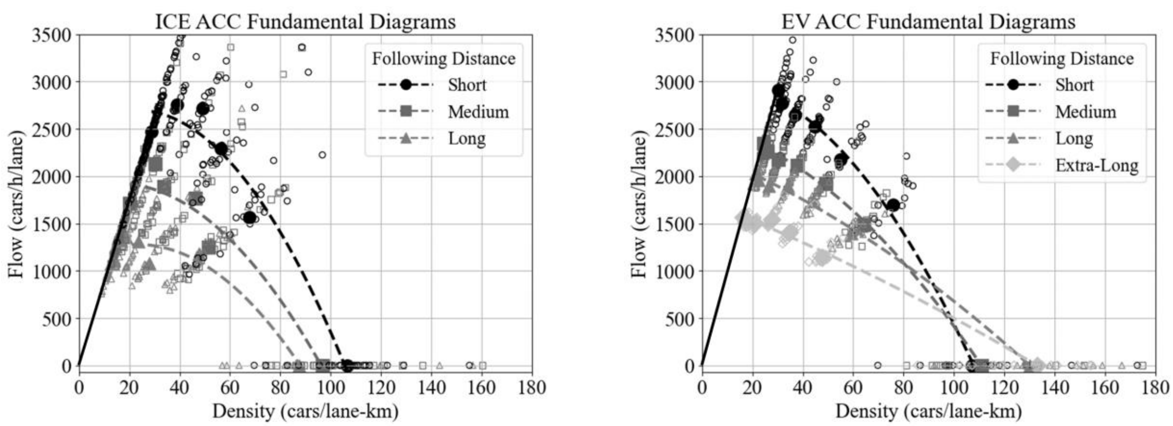

(a) (b)

Figure 7 Fundamental diagrams and field data: (a) ICE, (b) EV

Upon visual inspection of both figures, it is evident that selecting a short following-distance leads to queue discharge flows (represented by the peaks of the concave-shaped relations) greater than those typically observed in present-day traffic conditions, e.g., Seeherman and Skabardonis (2013); TRB (2016). This observation holds true even when the medium following-distance is chosen for the EV. These findings align with the results from Shi and Li (2021) and Li et al. (2022).

Turning our focus to jam densities, Figure 7a indicates that estimates for ICE-platoons range from 106 to 88 cars/lane-km, depending on the selected preferred following-distance. While most of these estimates are larger than those reported in Shi and Li (2021) and Li et al. (2022), they still fall below the values commonly reported for traffic composed solely of fully-human-controlled vehicles (Chiabaut et al., 2009; Hoogendoorn et al., 2015; Knoop & Daamen, 2017; Lárraga & Alvarez-Icaza, 2010; Li et al., 2022; Rakha et al., 2008). It is worth noting that the present estimates of jam density decrease as the settings for following-distance get larger. Consequently, settings that result in lower queue discharge flows also lead to lower congested densities, which in turn have a higher potential for queue spillovers.

In contrast, the pattern for tested EVs is reversed: Selecting shorter following-distance settings results in lower jam densities, as clearly illustrated in Figure 7b. When short and medium following-distances are chosen, the jam densities are comparable to those observed in ICE-vehicles using the same following-distance settings. However, setting the EV at long- and extra-long following-distances, produced jam densities of 130 and 133 cars/lane-km, respectively, which are nearly equivalent to what is typically observed in human-controlled traffic. This positive outcome appears to be attributed to the EV's regenerative braking system, i.e., when opportunities arose, the controller gradually rolled the car forward as it came to a halt, possibly to recharge the onboard battery.

4 Model Calibration

The Aimsun microscopic platform (Aimsun, 2017) was utilized to simulate network performance for settings inspired by Los Angeles and Manhattan. To emulate ACC-controlled cars, Aimsun parameters were calibrated to match the field data presented in Chapter 3. To emulate human-controlled cars, for which Aimsun is highly reputed (Ahmed et al., 2021; Panwai & Dia, 2005), parameters were carefully chosen from literature. Consequently, the choice of the Aimsun platform enabled simulations of both present-day conditions and of future times when ACC (and EVs) are expected to progressively dominate the market (Calvert et al., 2017; Litman, 2020; Tillema et al., 2017). The platform enabled this in a relatively simple fashion, and with good calibration results.

Section 4.1 outlines the calibration process employed for modeling ACC-controlled cars, accompanied by a discourse on the parameters selected for human-controlled cars. Section 4.2 provides the complete set of Aimsun input parameters used in the simulations, while Section 4.3 presents comparisons between Aimsun's predictions and the field data.

4.1 Model Calibration Process

The calibration process for Aimsun's car-following parameters to mimic ACC-controlled cars was meticulously executed for cases of homogeneous traffic, meaning that: model parameters were selected through separate simulation runs, always of traffic comprised entirely of one vehicle type (ICE-powered or EVs), and for which all cars were set to a single, same preferred following-distance (short, medium, long, or extra-long in the case of EVs). These simulation runs were conducted for every feasible combination of vehicle type and following-distance, resulting in the individual calibration of parameters for 7 cases of homogeneous ACC-controlled traffic.

For each case, parameters were fine-tuned through extensive trial-and-error iterations. The objective was to achieve a close match between the simulated outcomes and their respective fundamental diagram from the field data. Each simulation run’s outcomes were generated for a single travel lane spanning 1 km, with vehicles entering from the upstream end at a rate approaching the lane's capacity. To replicate slowdowns possibly caused by congestion from a downstream bottleneck, a speed limit was introduced at the lane’s midpoint. The speed limit varied parametrically at rates of {3, 6, 15, 25, 40, 50} (km/h) to represent different congestion levels. Congested densities and flows were jointly measured using a simulated vehicle detector placed just upstream of the posted speed limit.

Three parameters were calibrated via this process. The first parameter, referred to as "clearance" (measured in meters), is defined as the average jammed distance between the rear bumper of a lead vehicle and the front bumper of its follower. Additionally, two dimensionless parameters, namely "sensitivity" and "headway aggressiveness," were also calibrated. These parameters dictated the shape of the fundamental diagram's congested branch.

Human-controlled traffic was constrained to match a nearly piecewise-linear-shaped fundamental diagram. The fundamental diagram encompassed a free-flow speed of 50 km/h, an average queue discharge flow of 1,900- and 1,650 cars/h/lane, one to describe each of our two distinct urban settings (TRB, 2016), and a jam density of 138 cars/lane-km. The jam density was the smallest value found in the literature (Chiabaut et al., 2009). We refrained from utilizing larger jam density values found in other studies (Hoogendoorn et al., 2015; Knoop & Daamen, 2017; Lárraga & Alvarez-Icaza, 2010; Li et al., 2022; Rakha et al., 2008). This decision was made to ensure that our predictions regarding ACC's potential adverse effects on jam density would be the smallest within reason, thus conservatively addressing any potential queue storage issues anticipated in an ACC-dominated future.

4.2 Input Parameters Estimates

The values selected for Aimsun's input parameters are presented in Table 1 for each case of homogeneous-traffic. In the table's last row, Aimsun's default parameter values are provided as a reference. Values that are different from their respective defaults are denoted in bold italics in the table. The last column in Table 1 shows the maximum acceleration rates used for each car type. An acceleration rate of 4.5 m/s^2 was selected for EVs, based on the specifications in Hyundai (2022), to replicate the higher acceleration capabilities of EVs. For all other cases, the default value of 3 m/s^2 was retained.

Table 1 Input parameters

Vehicle Type (Following - Dist. Setting)	Clearance (m)	Sensitivity Factor	Headway Aggressiveness	Acceleration (m/s^2)
ICE ACC (Short)	***5.2***	***0.98***	***1***	3
ICE ACC (Medium)	***5.5***	***1.16***	***0***	3
ICE ACC (Long)	***6.4***	***1.62***	***-0.75***	3
EV ACC (Short)	***4.5***	***0.96***	***1***	***4.5***
EV ACC (Medium)	***4.2***	***1.10***	***0***	***4.5***
EV ACC (Long)	***3.0***	***1.20***	***-0.75***	***4.5***
EV ACC (Extra-Long)	***2.9***	***1.50***	***-1***	***4.5***
Human Driver (high discharge flow)	***2.3***	***1.10***	***-1***	3
Human Driver (low discharge flow)	***2.3***	***1.35***	***-1***	3
Aimsun Default	1	1	0	3

In the next section, the resulting agreements between the simulated and field-measured data for ACC-controlled traffic are presented. Additionally, comparisons are drawn against the fundamental diagrams that describe our human-controlled traffic.

4.3 Fundamental Diagrams and Comparisons

The fundamental diagrams produced via simulations are displayed in Figures 8a to 8d. In the case of ACC-controlled EVs, the fundamental diagrams are split across two figures (8b and 8c) for clarity. The shaded data points represent averages from 10 simulations. It is noteworthy, as shown in Figures 8a to 8c, that the simulated fundamental diagrams (solid curves) closely match with their empirically-estimated counterparts (dashed curves) presented earlier in Figures 7a and 7b. The largest differences are observed between the simulated and empirical fundamental diagrams when the following-distance for EVs is set at extra-large, as seen in Figure 8c. Nevertheless, the congested branch of this simulated diagram either matches or lies above the empirical branch, which will produce a conservative estimation of ACC's adverse effects in the forthcoming analyses.

Regarding the input to Aimsun, the jam density for the human-controlled fundamental diagram in Figure 8d exceeds those for the ACC cases. To make these distinctions more evident, we have replicated the lightly shaded area beneath the fundamental diagram, which features a queue discharge rate of 1,900 cars/h/lane from Figure 8d, in the three other figures.

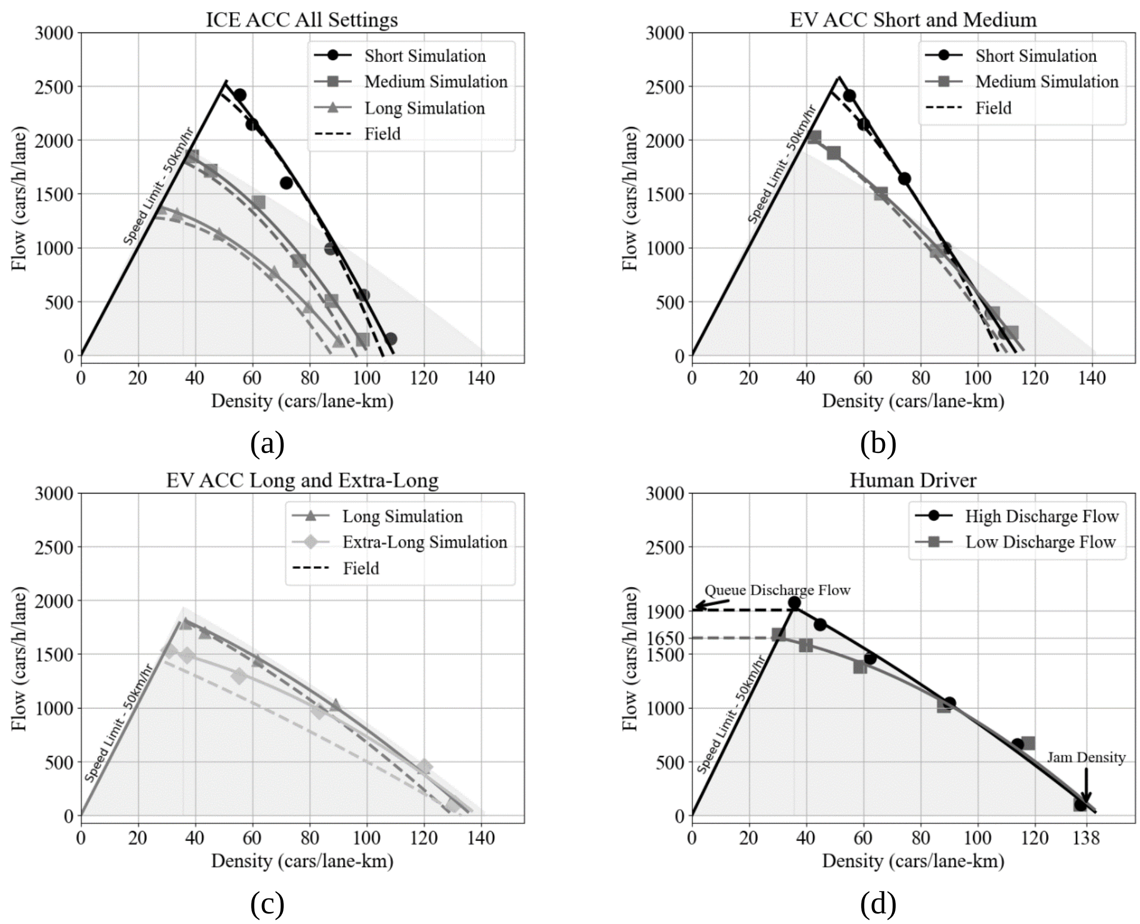

Figure 8 Simulated and empirical FDs. (a) ICE-ACC, (b) EV-ACC, short and medium following-distances, (c) EV-ACC, long and extra-long distances, (d) Human-controlled traffic

5 Illustrations of ACC Impacts

In this chapter, a preview of the impacts of ACC that one might expect in the nearer term- and more-distant futures is provided. Illustrations to this end draw inspiration from two distinct urban settings. One is reminiscent of Downtown Los Angeles, a car-oriented city characterized by longer block lengths and moderately high congestion. The other setting takes cues from Midtown Manhattan, featuring shorter blocks typical of older cities and more severe congestion. Street geometry in both the Los Angeles and Manhattan inspired settings was kept simple, again in the pursuit of conservative predictions regarding ACC-induced problems. Among these simplifications is the absence of turn pockets, which could have otherwise facilitated the spillover of less-compact ACC queues from turning traffic into adjacent lanes, hindering the flow of through-moving vehicles.

Projections for future demands are also conservative in nature. For instance, morning demands were set to approximately replicate the existing congestion levels in Los Angeles and Manhattan, with the assumption that these demands will not grow in the years to come. Forecasts regarding ACC adoption behavior are based on informed guesswork. However, as the subsequent sections reveal, it becomes evident why reasonable variations in these guesses are unlikely to skirt the problems anticipated in the nearer- term and more distant futures.

Inputs to our illustrations are presented in Sections 5.1 to 5.3. Section 5.4 delves into unsettling predictions for intervention-free futures, while findings that point toward promising interventions are offered in Section 5.5.

5.1 Simulated Networks

The street networks for both urban settings were idealized as homogenous square grids, consisting of 20 north-south (N-S) and 20 east-west (E-W) bidirectional arterials, each featuring two travel lanes per direction, as illustrated in the inset of Figure 9. Vehicles within these networks were allowed to make turning maneuvers from any approach. The spacing between the centerlines of all street links was 210 meters for the Los Angeles-inspired site and 170 meters for the Manhattan-inspired site. Additionally, a speed limit of 50 km/h was enforced on every link within the grid. Queue discharge flows were set at 1,900 cars/h/lane and 1,650 cars/h/lane for the Los Angeles and Manhattan-like grids, respectively.

For traffic control, pre-timed, two-phase traffic signals with unprotected left turns were assumed to control every intersection across the grid. All signals operated on a long 90s cycle, which is a good practice when managing congested traffic (see Koshi, 1989). To accommodate the symmetry within the network, equal green and amber splits of 45s were utilized for each phase. The signal phases were synchronized with zero offsets, a strategy known to be effective in congested settings (Daganzo & Lehe, 2016).

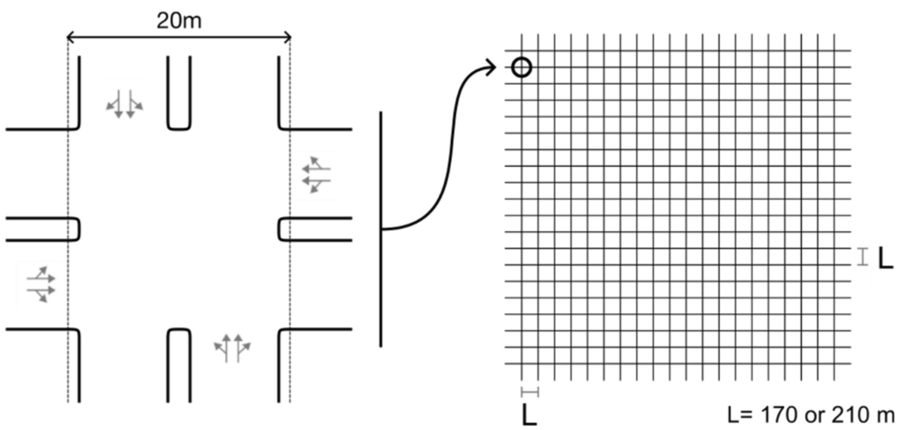

Figure 9 Representation of simulated networks and the street configuration

5.2 Travel Demands

The demands for car trips were set to approximately replicate the morning rush-hour congestion levels presently observed in Los Angeles and Manhattan. To achieve this, we quantified each city's congestion level as the maximum rise in average vehicle pace (the inverse of speed), normalized by the average free flow pace (see Figure 10 for example). We obtained these measurements by analyzing a large number of morning commute trips using Google Maps, similar to Doig et al. (2023). For reference, the estimated congestion levels in Los Angeles and Manhattan were 2.7 and 5.0, respectively.

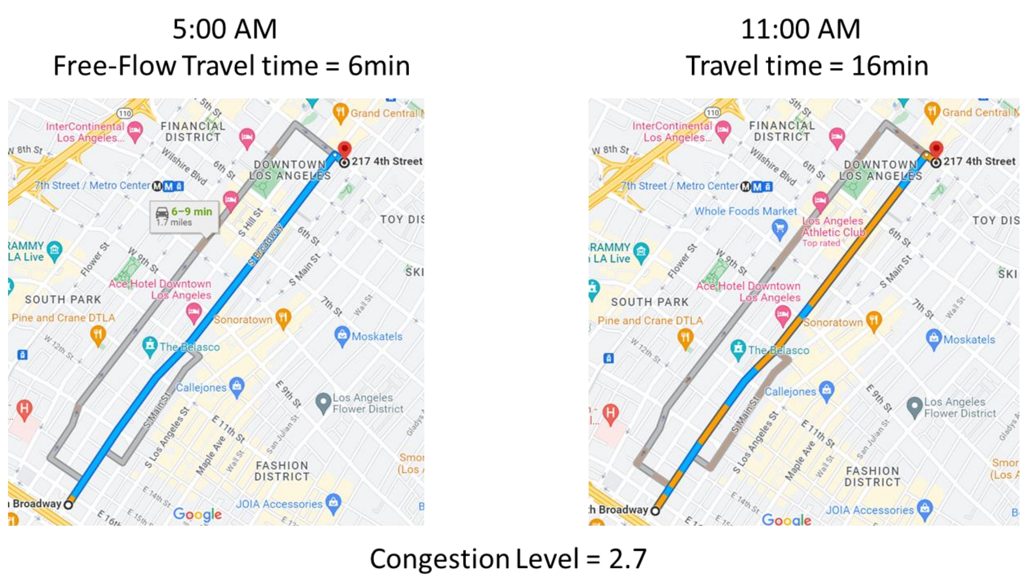

Figure 10 Estimation of congestion level in Los Angeles

The aggregate demand for each urban setting was assumed to follow a time-varying pattern illustrated by the respective piece-wise linear cumulative count curve in Figure 11. Note how the rates changed: starting at a lower rate, increasing at the 20-minute mark from the simulation's initiation, returning back to the lower rate at the 60-minute mark, and ultimately tapering off to zero by the 80-minute mark, ensuring that all cars could complete their trips by the end of the simulation. These higher and lower demand rates were adjusted accordingly for each urban setting to ensure that, when combined with a time-invariant Origin-Destination (OD) matrix reflecting trip fractions to represent the spatial distribution of demand, the simulated congestion levels for each city would align with their respective estimates derived from Google Maps.

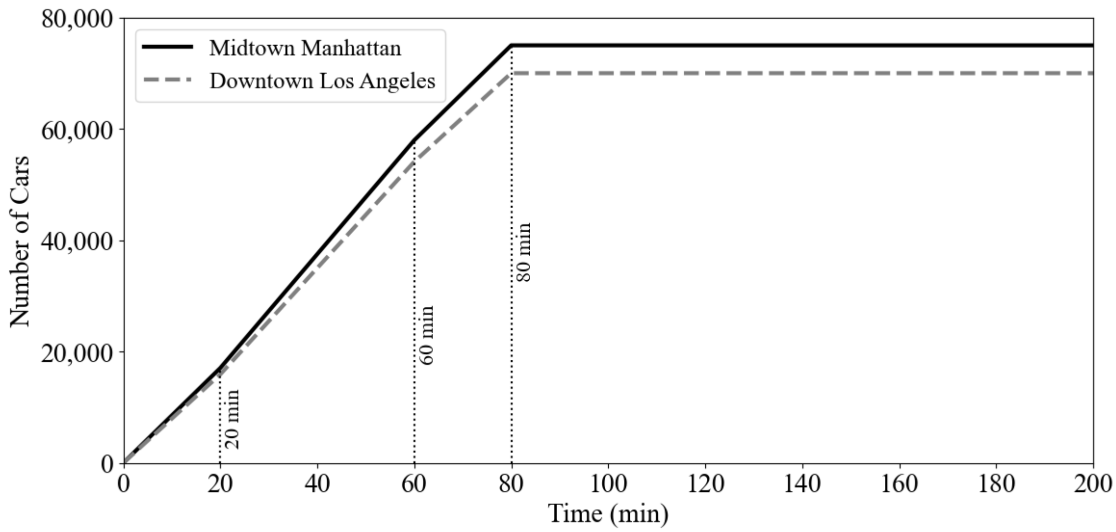

Figure 11 Time-varying aggregate demands

All trips in the simulations had origins and destinations located at intersections on the grids. Trip fractions were the same for all OD pairs in the matrix, meaning that OD pairs were uniformly and independently distributed over space, and trips between all OD pairs adhered to a Poisson process with an equal rate. Cars entered the network via one of the four departing links at their origin intersections when it was safe and possible to do so, with entry delays factored into all subsequent analyses. Upon reaching their destination intersections, cars were promptly removed from the network, without any additional time spent searching for parking.

5.3 ACC Adoption Forecasts

As vehicle automation continues to advance, the prevalence of ACC is expected to surge in the future. In fact, as of today, ACC is becoming increasingly ubiquitous, with a staggering 92% of newly sold cars either offering it as a standard feature or as an optional upgrade (Bartlett, 2021). Figure 12, sourced from Bansal and Kockelman (2017), provides a visual representation of a conservative forecast for the market share of ACC-equipped vehicles. According to the plot, it is forecasted that by 2045, the market share of ACC-equipped cars will surpass 80%. Consequently, in our simulations, we can reasonably assume that the adoption rates of ACC will increase over time (Calvert et al., 2017; Litman, 2020; Tillema et al., 2017). A growth is not only expected for ACC but for EVs as well, such that ICE-powered cars will likely decline in number (CARB, 2022; Choi, 2021; Lambert, 2021). The adoption of EVs is rising sharply as the global push for net-zero carbon emissions accelerates. The share of EV sales is anticipated to be well above 80% by 2040 in developed countries (Goldman Sachs, 2023), as seen in Figure 13. By taking those projections into account, our illustrations will follow the informed guessed-at forecasts displayed in Figure 14. We view these as optimistic, in the sense that other guesses, such as a higher prevalence of ICE-powered cars, for example, would produce bleaker prognostications. As seen in the figure, our projections encompass three futures: (i) a shorter-run future when half of all cars are equipped with ACC, and half of those are EVs; (ii) an exceptionally optimistic longer-run future when 90% of all cars are EVs with ACC, and (iii) a more moderately optimistic future when 80% of cars EVs with ACC. In all these futures, simulated outcomes are compared to present-day, where all cars are assumed to be operated by human drivers.

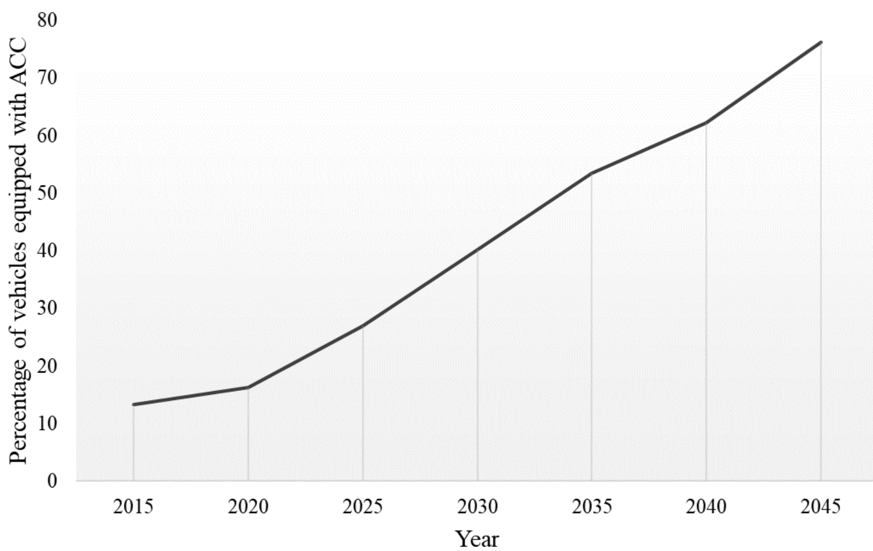

Figure 12 Forecast for the market share of ACC-equipped vehicles (Bansal & Kockelman, 2017)

Recall that drivers of ACC-equipped vehicles are assumed to always engage the technology. Additionally, these drivers are assumed to set a preferred following-distance that does not change during their trips. The distribution of preferred following-distances for ACC-driver populations is taken from the empirical study in Nowakowski et al. (2010). In this study, 50.4% of ACC-drivers selected a short following-distance, 18.5% opted for a medium distance, and 31.1% selected the long setting. For the purposes of our analyses, we assumed that drivers of ACC-equipped EVs select the long- and extra-long settings in equal proportions, accounting for a total of 31.1% of all ACC-drivers.

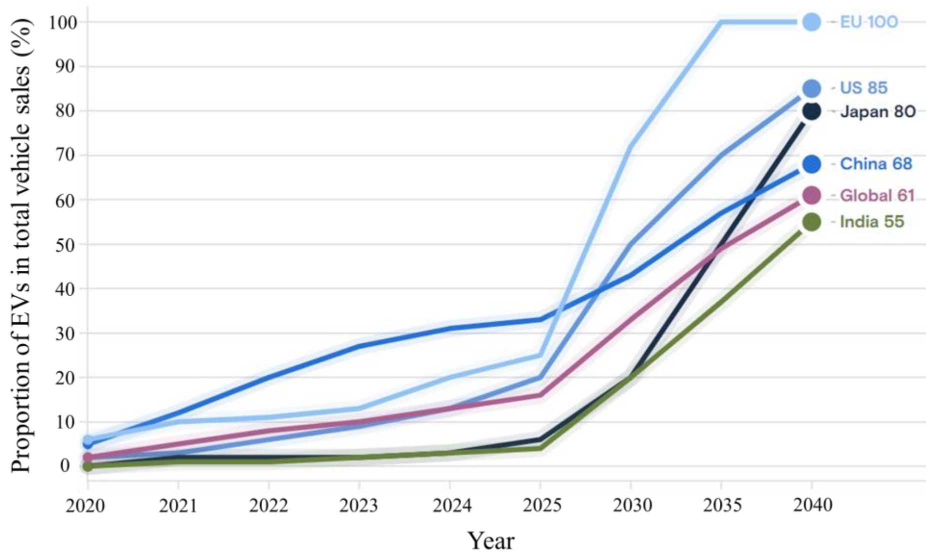

Figure 13 Projected EV sales as a proportion of total vehicle sales (Goldman Sachs, 2023)

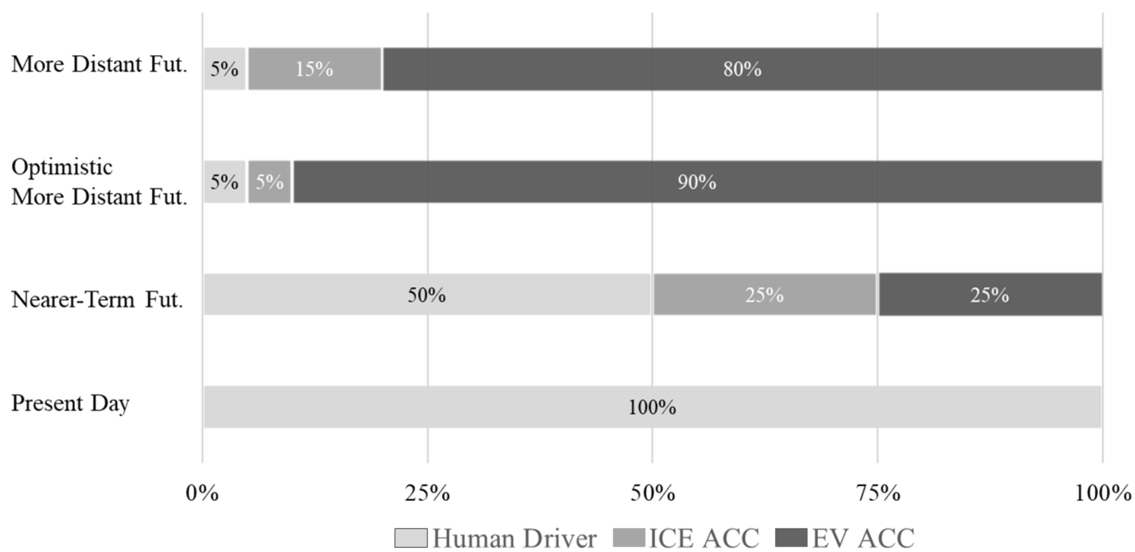

Figure 14 Adoption forecasts

5.4 Simulation Performance Predictions

The outcomes for the settings inspired by Los Angeles and Manhattan are depicted in Figures 15a and 15b, respectively. In these figures, the solid, bold curves represent the aggregate demands during the morning rush for each respective site. The thinner solid curves show cumulative exit counts for the present-day, while the dashed and dotted curves represent exit counts for the forecasted futures. Each exit curve is derived from an average of 10 simulations with distinct random seeds, and all outcomes presented in this study were similarly averaged in this fashion. The curves are plotted such that, the area between a demand curve and an exit curve corresponds to the networkwide vehicle-hours traveled (VHT) for the respective scenario, and the vertical displacements between these curves reflect the time-varying vehicle accumulations, which are also networkwide. Upon visual inspection of both figures, it becomes apparent that the predicted future traffic conditions are nearly always worse and most of the time much worse compared to the present-day for both settings.

From further examination of the figures, it is evident that the outlook is somewhat more optimistic for the Los Angeles-like setting. This can be attributed, at least in part, to its longer link lengths, which provide additional space for queue storage. Still, Figure 15a illustrates that the network's VHT and accumulation (indicative of congestion) are noticeably worsened in the nearer-term future, where ICE-powered cars are expected to still constitute a significant portion of the traffic. In this scenario, the VHT is projected to increase by 12% compared to the present-day level, with delays anticipated to rise by 25%.

Upon closer examination of the near-term exit curve in Figure 15a, the reasons why problems occur become evident. The declining slopes of the cumulative exit curve (i.e., network exit flows), start to manifest at approximately 60 minutes from the initiation of the simulations, coinciding with an accumulation of approximately 14,000 cars within the network. This diminishing exit flow is a conspicuous sign of the onset of gridlock, e.g., see Daganzo (1996, 2007); Mahmassani et al. (2013). It is worth noting how the gridlock situation gradually, albeit artificially, subsided as exit flows were restored when the demand abruptly dropped to zero at the 80-minute mark. It is important to mention that gridlock would have persisted, and its adverse effects exacerbated, had the demand diminished more gradually and realistically toward the end of the rush period.

The outlook for the more distant futures appears more promising for the Los Angeles inspired setting. In fact, the optimistic exit curve nearly aligns with the present-day curve to the extent that it is barely discernible in Figure 15a. The longer link lengths, and lower present-day congestion level of the Los Angeles-like network played a pivotal role. There were substantial periods when the network was not fully engulfed by residual queues. The higher queue discharge rates enabled by ACC, and the higher acceleration rates facilitated by EVs, increased exit flows during these periods. This positive effect largely offset the adverse impact stemming from ACC's less-compact queues. Consequently, the networkwide VHT predicted for the optimistic distant future closely matches that of the present-day.

Meanwhile, the exit curve for the less-optimistic distant future also closely approximates its present-day counterpart, albeit not as closely. This less optimistic future resulted in the formation of the least densely packed queues compared to all cases. Nevertheless, the damage done compared to the present-day is small, primarily attributable to the higher queue discharge flows and vehicle accelerations brought by the EVs present in that future.

While the predictions of less damaging distant futures for the Los Angeles inspired setting offer some solace, it is crucial to acknowledge that these predictions rely on several favorable assumptions. Thus, these predictions still fall significantly short of the benefits of ACC championed by its advocates.

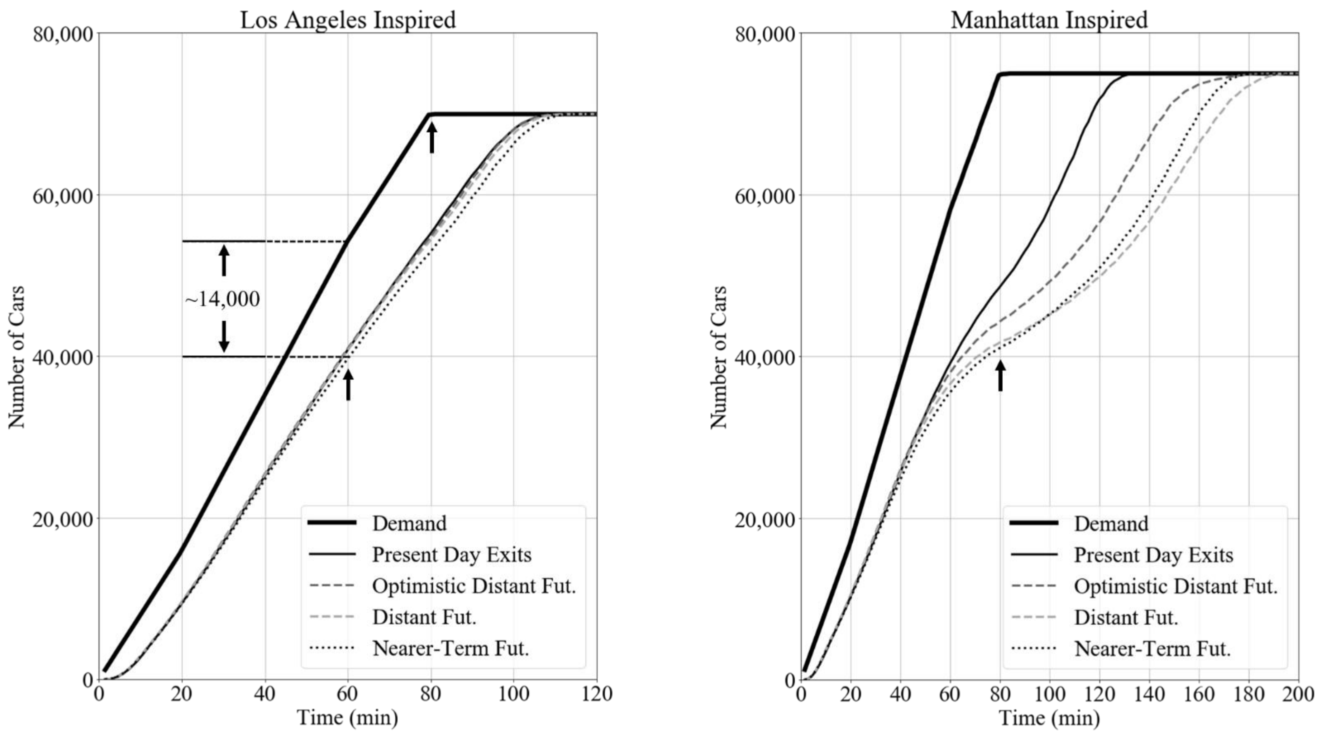

(a) (b)

Figure 15 Cumulative Curves for (a) Los Angeles-like and (b) Manhattan-like settings

The situation takes a bleaker turn for the Manhattan-inspired setting depicted in Figure 15b. This can be attributed to the shorter block lengths and limited queue storage space, which already contribute to high congestion levels in the present-day and exacerbate the problems with ACC predicted in the future. In Figure 15b, it is noticeable that up until around the 80-minute mark, the predicted damages for the various future scenarios ranked in a similar manner to those of the Los Angeles-like site, with the worst damage predicted for the nearer-term future. However, beyond this point, the Manhattan-like site became overwhelmed by queues to the extent that the benefits of increased discharge flows and accelerations no longer applied. Consequently, the less optimistic distant future, marked by the least compacted queues, emerged as the most concerning scenario for the Manhattan-like setting. The onset of gridlock, as evidenced more prominently in this case, led to a staggering 87% increase in networkwide VHT compared to present-day conditions. While the situation was somewhat improved in the optimistic distant future, VHT still rose by 46% relative to the present-day.

Table 2 presents a summary of the performance predictions for both the Los Angeles and Manhattan-inspired settings. This table offers a comparison of predicted networkwide VHT for present-day conditions and the three future scenarios, along with the corresponding percentage change in VHT compared to the present-day baseline.

Table 2 Summary of performance predictions for the Los Angeles and Manhattan-inspired settings

Scenario	Los Angeles Inspired		Manhattan Inspired	
	VHT	Change (%)	VHT	Change (%)
Present Day	15673	-	27990	-
Optimistic Distant Future	15623	~ 0%	40838	46%
Distant Future	16056	2%	40838	87%
Nearer-Term Future	17530	12%	50396	80%

5.5 Potential Interventions for a Different Future

One can question the ACC and EV adoption forecasts used in the work (see again Figure 14). Nevertheless, moderate deviations from these adoption forecasts are unlikely to significantly alter predicted outcomes, as evident from the results in Figures 15a and 15b. The nearer-term adoption forecasts differ considerably from those for the more distant future. Yet, when examining the shifts among exit curves in all three adoption scenarios, they appear relatively small compared to their deviations from the corresponding demand curves drawn in bold. This is a cause for concern, particularly for denser, more congested cities like Manhattan, regardless of any errors in the forecasts. Even more sprawling cities, such as Los Angeles, may also have valid reasons for concern, given that our predictions of future congestion are probably very loose lower bounds.

It becomes markedly evident that higher queue discharge flows, achievable through ACC, cannot avert impeding problems once street links are filled with queues. For instance, in Figure 16, we present the demand and present-day exit curves alongside two additional exit curves, all for the Manhattan-inspired setting. The dotted exit curve in this plot represents a somewhat dystopian future where all vehicles are EVs controlled by ACC, and all drivers opt for the short following-distance setting. In this case, the average queue discharge flow is the highest possible (significantly exceeding present-day rates), and the jam densities are relatively high (though lower than at present-day, as seen in Figures 8a and d). Yet, the dotted exit curve in Figure 16 demonstrates how conditions deteriorate over time, with a predicted 94% increase in networkwide VHT.

A more promising future lies in maximizing jam density to compact queued vehicles more tightly. To illustrate, the dashed curve in Figure 16 represents exit flow when all vehicles are EVs, ACC-controlled, and set for the extra-long following distance setting. While this setting reduces queue discharge flows compared to the short following-distance settings, it results in a jam density closer to present-day traffic (as seen in Figures 8c and d). Indeed, the dashed curve in Figure 16 lies closer with the present-day curve, with networkwide VHT predicted to increase by 14%. While this increase remains substantial, it represents a significantly better outcome compared to the 94% increase.

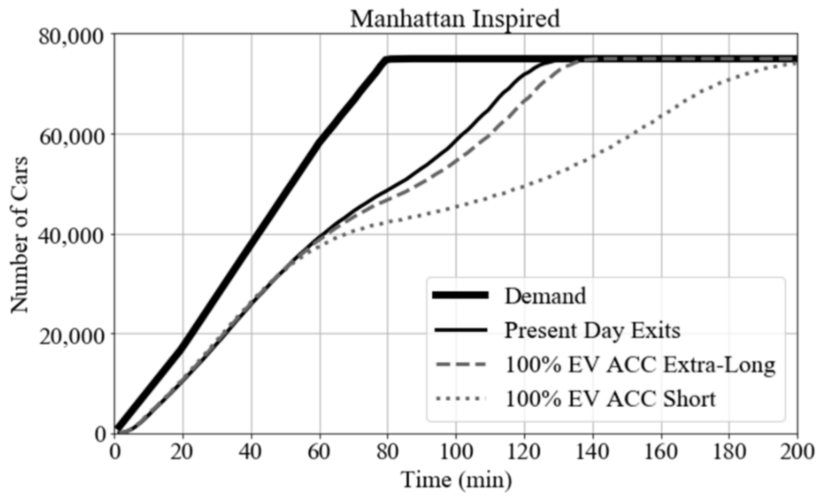

Figure 16 Cumulative Curves Diagrams for Manhattan-like settings

In light of the more promising outlook resulting from maximizing jam density to pack queued vehicles more tightly, a viable intervention strategy becomes apparent. This strategy entails adjusting the settings of ACC controllers to opt for smaller spacings in congestion, particularly when jammed. Simulations were carried out after recalibrating ACC-equipped vehicles to adopt jam spacings akin to present-day values, regardless of the preferred following-distance settings chosen by drivers. Throughout this recalibration process, discharge flows remained consistent with those depicted in Figures 8a to 8c. The results, as portrayed in Figure 17, showcase significant improvements across both nearer-term and more distant future scenarios, for the Manhattan-inspired setting, and especially for our optimistic distant future scenario. Networkwide VHT in this scenario is anticipated to decrease by 39% compared to present-day levels. Table 3 furnishes a summary of performance predictions for the Manhattan-inspired setting, highlighting the potential VHT reductions due to this intervention.

Table 3 Performance predictions with ACC spacing intervention for the Manhattan-inspired setting

Scenario	Manhattan Inspired	
	VHT	Change (%)
Present Day	27990	-
Optimistic Distant Future	17170	-39%
Distant Future	18054	-36%
Nearer-Term Future	23180	-17%

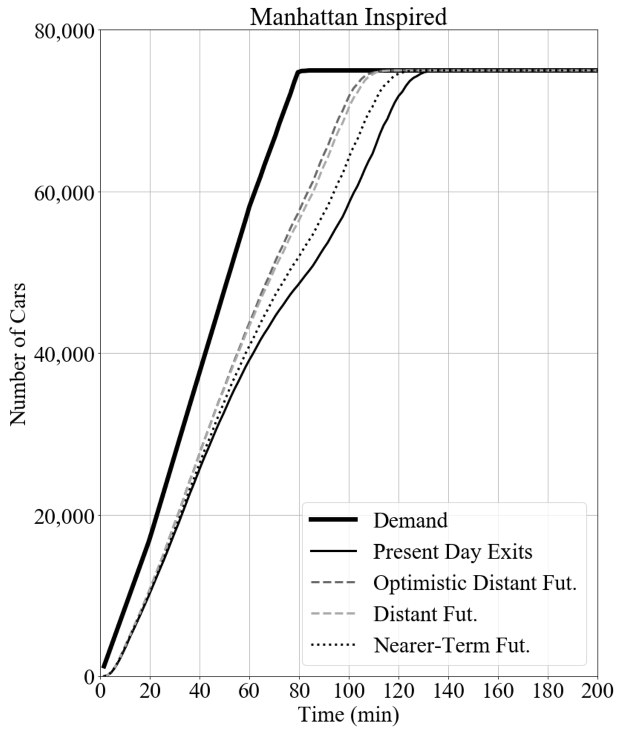

Figure 17 Cumulative Curves with ACC spacing intervention for Manhattan-inspired setting

6 Conclusion

In conclusion, this dissertation raises a red flag by revealing the adverse impacts of ACC-rich traffic on queue storage and urban congestion. Through careful field measurements of ACC-controlled cars traveling on roads and a highway, this study confirms that queues formed by ACC-controlled vehicles exhibit lower densities compared to those comprised entirely of human-operated vehicles. These uncompacted queues pose problems for urban areas, where queue storage is often already a problem during rush hours. Simulation-based case studies loosely inspired by two distinct urban settings—Downtown Los Angeles and Midtown Manhattan—confirm this concern and contradict optimistic predictions from advocates of how ACC may lead to a congestion-free future. The higher bottleneck capacities commonly promised of ACC become irrelevant when queues that spillover from one link to the next constrain a bottleneck's flow from reaching those capacities.

In this concluding chapter of the dissertation, we recapitulate the key findings, suggest some interventions and examine potential challenges that may arise from those interventions. Furthermore, we outline future research avenues that complement the current study. These research avenues involve conducting additional simulation studies tailored to specific urban settings, considering their unique characteristics. Additionally, we propose naturalistic driving experiments with ACC-equipped cars that go beyond the experiments conducted in the present study and previous related efforts. The objective of these future research avenues is to offer further evidence of the findings presented in this dissertation and stimulate further discussions about potential interventions to manage ACC-induced congestion.

6.1 Summary of Findings

Field measurements of steady-state conditions, with numerous observations of jammed states, have confirmed that queues formed by ACC-controlled vehicles exhibit lower densities compared to those comprised entirely of human-operated vehicles. Specifically, the reduction in jam density within a platoon of ACC-equipped ICE cars can vary from 23% to 38%, while for ACC-equipped EVs, the reduction in jam density ranges from 4% to 23%, depending on the preferred following-distance settings selected by drivers. Simulations utilizing calibrated car-following models in two idealized settings illustrate that these lower-density queues indeed pose impending problems for urban areas in the future, as ACC's market share continues to grow. Unsurprisingly, queue storage issues were forecast to have the most pronounced adverse effects in our Manhattan-inspired setting, characterized by shorter block lengths and already-high levels of congestion. What may be both surprising and disconcerting is the magnitude of the damage predicted (up to 87% increase in networkwide VHT), especially since some of our assumptions were quite conservative in this respect.

Remedying this disconcerting future could involve implementing restrictions on the use of ACC when driving in congested traffic, especially on city streets. Such a prohibition would be akin to the earlier calls in (GMC, 2023; Honda, 2022; Sparks, 2022; Volvo, 2020), yet could be justified not only for safety reasons, but also on societal grounds too, i.e., to curb congestion and its externalities. Nevertheless, imposing restrictions on a technology feature for which drivers have paid good money might encounter resistance and prove challenging to enforce. Furthermore, it would signify a regression in the march toward greater vehicle automation and halt any notion of automation as a panacea for congestion. Thus, we anticipate a degree of reluctance when considering such interventions.

A more feasible intervention could involve adjusting ACC controllers to opt for smaller spacings in congested conditions, particularly when vehicles are jammed, e.g., at red lights. Simulations were conducted after recalibrating ACC-equipped cars to adopt jam spacings similar to present-day values, regardless of the preferred following-distance settings selected by drivers. As a result of this adjustment to ACC controllers, networkwide VHT reductions were predicted for all future scenarios compared to present-day. The most substantial drop in networkwide VHT was predicted in the most optimistic longer-run future for the Manhattan-inspired setting, with reductions up to 39% in VHT. While this outcome may be less impressive than what has been promised by advocates (USDOT, 2019), it offers a hopeful prospect. Additionally, our previously conservative assumptions (e.g., concerning link geometry and future demand levels), may be causing us to underestimate the benefits of these fine-tuned ACC controllers.

Nonetheless, mandating shorter vehicle spacings during congested traffic could raise concerns and bring the issue of liability to the forefront, with automobile manufacturers expressing apprehensions about their customers' safety when faced with shorter following-distances. In short, there may be no straightforward solution to the currently identified looming concerns.

For these and other reasons, we recognize that the red flag raised in the dissertation will not be the final word on this subject. In the initial stages, it is foreseeable that proponents of technology and various stakeholders will likely advocate for more extensive studies, potentially encompassing locations such as Los Angeles, Manhattan, and other settings. This could indeed be a positive development, provided that any subsequent analyses adhere to the care and objectivity as in this dissertation. Nonetheless, it is important to acknowledge that conducting further investigations will necessitate time, which may not be on our side. This is particularly pertinent considering the steady growth in the market share of ACC (Calvert et al., 2017; Litman, 2020; Tillema et al., 2017), and recalling vehicles in the future to update their ACC controllers retroactively could present substantial challenges.

Therefore, any further assessments ought to be expedited. This, in turn, may require USDOT to shift its role from being an enthusiastic advocate to a more cautious assessor and potentially a regulator. The likelihood of these transformations remains uncertain, particularly in light of a recent USDOT publication indicating that the responsibility for balancing ACC's impact on safety and road network performance is now placed on manufacturers; see USDOT (2022), p. 20. This transfer of authority raises concerns for us and appears to be a worrisome relinquishment of oversight. Perhaps leadership may instead come from state and local governments or potentially from other parts of the world. Nevertheless, wherever it arises, effective leadership is urgently required.

6.2 Future Work

The work presented in this dissertation raises a red flag by revealing the adverse impacts of ACC-rich traffic on queue storage and urban congestion. The exploration of these impacts was conducted through careful field experiments and through simulation-based case studies loosely inspired by two distinct urban settings: Downtown Los Angeles and Midtown Manhattan. These urban settings were simplified as homogeneous square grids, while incorporating certain assumptions (e.g., for link geometry and future demand levels) favorable to ACC.

The dissertation thus makes a contribution in its own right and paves the way for additional research in this area. Decision-makers in cities may seek a more detailed understanding of how the future might unfold in their unique urban environments. As a next step, to this end, decision-makers might stimulate their cities with their unique characteristics. This task can be reasonably straightforward given decision-makers' presumed access to input data regarding street geometry, projected future demands, etc.

However, it is important to acknowledge that microscopic traffic simulations, while valuable, can only approximately emulate the dynamics of individual vehicles, via use of car-following and lane changing models. Thus, the most definitive findings will come from real-world field experiments that go beyond the experiments conducted in the present study and in previous related efforts. To this end, researchers might conduct naturalistic driving experiments involving commercial ACC-equipped vehicles. In these experiments, the participant drivers would be instructed to activate ACC while travelling through various facilities (e.g. city streets, highways, freeways), at different times of the day and under varying traffic (e.g. free-flow, moderate congestion, severe congestion) and environmental conditions (e.g. rainy, foggy, sunny). All vehicles would be equipped with technology for collecting vehicle spacings and transmitting that data to the analysts as vehicles perform their day-to-day trips. The objective of these experiments is to address any skepticism that may exist among advocates of ACC regarding the small-scale controlled field experiments conducted by this study, and offer further evidence of the findings presented in Chapters 3 and 5.

If advocates of ACC remain unconvinced by the results of these naturalistic driving experiments, an alternative and more dramatic experiment may be considered. It would involve gathering a substantial number of ACC-equipped vehicles to collectively drive designated city streets during a rush hour. Due to the expanded ACC-vehicle queues, it would be essential to acknowledge the associated risks, as this experiment could lead to extreme congestion and even gridlock. Therefore, meticulous planning and an alternative course of action would need to be in place to terminate the experiment if conditions approach an unrecoverable gridlock. The results of this experiment would undoubtedly be eye-opening for advocates of ACC and would hopefully stimulate further discussions about interventions.

In the quest for those interventions, one emerges as seemingly obvious: opting for smaller spacings when vehicles are jammed. However, implementing this intervention is probably easier said than done. Drivers are likely to feel uncomfortable when they are robotically driven at smaller spacings, and mandating such spacings may introduce safety concerns. The question of liability might thus take center stage for automobile manufacturers, who may also worry about the comfort perceived by their customers when confronted with shorter following-distances.

One potential solution that might accommodate driver comfort and safety while achieving smaller jam spacings is to fine-tune the controllers so that ACC-equipped vehicles gradually roll forward as they come to a halt, similar to how ACC-equipped EVs come to a complete stop at long- and extra-long following-distance settings; see again Section 3.3. However, it is uncertain whether drivers would be receptive to these adjustments. Therefore, conducting additional naturalistic driving experiments, like those mentioned above but with ACC-equipped cars with fine-tuned controllers, and offering participants the freedom to engage or disengage ACC based on their comfort, is necessary to assess driver receptiveness to this intervention.

If adjusting the controllers in this way faces opposition, cities should then consider, as a last resort, more drastic measures, such as segregating ACC-equipped vehicles by restricting them to specific lanes on streets or confining their use to designated streets when ACC is engaged. However, cities may hesitate to implement such measures for several reasons. One significant concern is the potential opposition from stakeholders, such as the USDOT and vehicle manufacturers, who often fund research in the field of vehicle automation and could view these measures a step backward in the march toward greater vehicle automation. Furthermore, ACC-equipped vehicle owners may also strongly object to restrictions on their vehicle use. Additionally, substantial enforcement challenges are associated with these restrictions. Ensuring that drivers comply with the restrictions would demand additional resources and rigorous enforcement measures, which might pose practical difficulties for city authorities. Considering all these obstacles, cities might be cautious about implementing such measures to manage ACC-induced traffic problems and may need to seek alternative approaches that fit their specific needs. Nonetheless, urgent and effective decision-making in this matter seems needed.

Bibliography

Ahmed, H. U., Huang, Y., & Lu, P. (2021). A Review of Car-Following Models and Modeling Tools for Human and Autonomous-Ready Driving Behaviors in Micro-Simulation. *Smart Cities*, *4*(1), Article 1. https://doi.org/10.3390/smartcities4010019

Aimsun. (2017). *Aimsun Next*. https://www.aimsun.com/

Bansal, P., & Kockelman, K. M. (2017). Forecasting Americans' long-term adoption of connected and autonomous vehicle technologies. *Transportation Research Part A: Policy and Practice*, *95*, 49–63.

Bartlett, J. (2021, November 4). *How Much Automation Does Your Car Really Have?* Consumer Reports. https://www.consumerreports.org/cars/automotive-technology/how-much-automation-does-your-car-really-have-level-2-a3543419955/

Bose, A., & Ioannou, P. A. (2003). Analysis of traffic flow with mixed manual and semiautomated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, *4*(4), 173–188. https://doi.org/10.1109/TITS.2003.821340

Calvert, S., Schakel, W., & Van Lint, J. (2017). Will automated vehicles negatively impact traffic flow? *Journal of Advanced Transportation*, *2017*.

CARB. (2022, August 25). *California moves to accelerate to 100% new zero-emission vehicle sales by 2035* [California Air Resources Board]. https://ww2.arb.ca.gov/news/california-moves-accelerate-100-new-zero-emission-vehicle-sales-2035

Carlson, R. C., Papamichail, I., & Papageorgiou, M. (2014). Integrated feedback ramp metering and mainstream traffic flow control on motorways using variable speed limits. *Transportation Research Part C: Emerging Technologies*, *46*, 209–221. https://doi.org/10.1016/j.trc.2014.05.017

Carlson, R. C., Papamichail, I., Papageorgiou, M., & Messmer, A. (2010). Optimal mainstream traffic flow control of large-scale motorway networks. *Transportation Research Part C: Emerging Technologies*, *18*(2), 193–212. https://doi.org/10.1016/j.trc.2009.05.014

Cassidy, M. J., Anani, S. B., & Haigwood, J. M. (2002). Study of freeway traffic near an off-ramp. *Transportation Research Part A: Policy and Practice*, *36*(6), 563–572.

Chiabaut, N., Buisson, C., & Leclercq, L. (2009). Fundamental diagram estimation through passing rate measurements in congestion. *IEEE Transactions on Intelligent Transportation Systems*, *10*(2), 355–359.

Choi, J. (2021, January 5). *Massachusetts to require 100 percent of car sales to be electric by 2035*. The Hill. https://thehill.com/policy/energy-environment/532684-massachusetts-to-require-100-percent-of-car-sales-to-be/

Daganzo, C. F. (1996). The nature of freeway gridlock and how to prevent it. *Transportation and Traffic Theory*, 629–646.

Daganzo, C. F. (1998). Queue spillovers in transportation networks with a route choice. *Transportation Science*, *32*(1), 3–11.

Daganzo, C. F. (2007). Urban gridlock: Macroscopic modeling and mitigation approaches. *Transportation Research Part B: Methodological*, *41*(1), 49–62.

Daganzo, C. F., Cassidy, M. J., & Bertini, R. L. (1999). Possible explanations of phase transitions in highway traffic. *Transportation Research Part A: Policy and Practice*, *33*(5), 365–379. https://doi.org/10.1016/S0965-8564(98)00034-2

Gipps, P. G. (1981). A behavioural car-following model for computer simulation. *Transportation Research Part B: Methodological*, *15*(2), 105–111.

GMC. (2023). *About Adaptive Cruise Control*. https://www.gmc.com/support/vehicle/driving-safety/driver-assistance/adaptive-cruise-control

Goldman Sachs. (2023, September 14). *Electric Vehicles are Forecast to Be Half of Global Car Sales by 2035*. https://www.goldmansachs.com/intelligence/pages/electric-vehicles-are-forecast-to-be-half-of-global-car-sales-by-2035.html

Goñi-Ros, B., Schakel, W. J., Papacharalampous, A. E., Wang, M., Knoop, V. L., Sakata, I., van Arem, B., & Hoogendoorn, S. P. (2019). Using advanced adaptive cruise control systems to reduce congestion at sags: An evaluation based on microscopic traffic simulation. *Transportation Research Part C: Emerging Technologies*, *102*, 411–426.

Gunter, G., Gloudemans, D., Stern, R. E., McQuade, S., Bhadani, R., Bunting, M., Delle Monache, M. L., Lysecky, R., Seibold, B., & Sprinkle, J. (2020). Are commercially implemented adaptive cruise control systems string stable? *IEEE Transactions on Intelligent Transportation Systems*, *22*(11), 6992–7003.

Honda. (2022). *Adaptive Cruise Control (ACC) with Low Speed Follow*. https://techinfo.honda.com/rjanisis/pubs/OM/AH/AT202222IOM/enu/details/131229047-67924.html

Hoogendoorn, S. P., Knoop, V. L., van Lint, H., & Vu, H. L. (2015). Applications of the Generalized Macroscopic Fundamental Diagram. In M. Chraibi, M. Boltes, A. Schadschneider, & A. Seyfried (Eds.), *Traffic and Granular Flow '13* (pp. 577–583). Springer International Publishing. https://doi.org/10.1007/978-3-319-10629-8_65

Hyundai. (2022). *IONIQ 5 Technology*. HYUNDAI MOTORS. https://www.hyundai.com/worldwide/en/eco/ioniq5/technology

James, R. M., Melson, C., Hu, J., & Bared, J. (2019). Characterizing the impact of production adaptive cruise control on traffic flow: An investigation. *Transportmetrica B: Transport Dynamics*, *7*(1), 992–1012.

Kesting, A., Treiber, M., Schönhof, M., & Helbing, D. (2008). Adaptive cruise control design for active congestion avoidance. *Transportation Research Part C: Emerging Technologies*, *16*(6), 668–683.

Kim, K., & Cassidy, M. J. (2012). A capacity-increasing mechanism in freeway traffic. *Transportation Research Part B: Methodological*, *46*(9), 1260–1272.

Knoop, V. L., & Daamen, W. (2017). Automatic fitting procedure for the fundamental diagram. *Transportmetrica B: Transport Dynamics*, *5*(2), 129–144. https://doi.org/10.1080/21680566.2016.1256239

Koshi, M. (1989). Cycle time optimization in traffic signal coordination. *Transportation Research Part A: General*, *23*(1), 29–34. https://doi.org/10.1016/0191-2607(89)90137-4

Lambert, F. (2021, November 10). *Countries and automakers agree to go all-electric by 2040 in weak new goal set at COP26*. Electrek. https://electrek.co/2021/11/10/countries-automakers-agree-go-all-electric-by-2040-weak-new-goal-cop26/

Lárraga, M. E., & Alvarez-Icaza, L. (2010). Cellular automaton model for traffic flow based on safe driving policies and human reactions. *Physica A: Statistical Mechanics and Its Applications*, *389*(23), 5425–5438. https://doi.org/10.1016/j.physa.2010.08.020

Li, T., Chen, D., Zhou, H., Xie, Y., & Laval, J. (2022). Fundamental diagrams of commercial adaptive cruise control: Worldwide experimental evidence. *Transportation Research Part C: Emerging Technologies*, *134*, 103458.

Litman, T. (2020). *Autonomous vehicle implementation predictions: Implications for transport planning*.

Mahmassani, H. S., Saberi, M., & Zockaie, A. (2013). Urban network gridlock: Theory, characteristics, and dynamics. *Procedia-Social and Behavioral Sciences*, *80*, 79–98.

Makridis, M., Mattas, K., Anesiadou, A., & Ciuffo, B. (2021). OpenACC. An open database of car-following experiments to study the properties of commercial ACC systems. *Transportation Research Part C: Emerging Technologies*, *125*, 103047.

Milanés, V., & Shladover, S. E. (2014). Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data. *Transportation Research Part C: Emerging Technologies*, *48*, 285–300.

Nowakowski, C., O'Connell, J., Shladover, S. E., & Cody, D. (2010). *Cooperative adaptive cruise control: Driver acceptance of following gap settings less than one second*. *54*(24), 2033–2037.

Ntousakis, I. A., Nikolos, I. K., & Papageorgiou, M. (2015). On Microscopic Modelling of Adaptive Cruise Control Systems. *Transportation Research Procedia*, *6*, 111–127. https://doi.org/10.1016/j.trpro.2015.03.010

Panwai, S., & Dia, H. (2005). Comparative evaluation of microscopic car-following behavior. *IEEE Transactions on Intelligent Transportation Systems*, *6*(3), 314–325. https://doi.org/10.1109/TITS.2005.853705

Papacharalampous, A. E., Wang, M., Knoop, V. L., Ros, B. G., Takahashi, T., Sakata, I., van Arem, B., & Hoogendoorn, S. P. (2015). *Mitigating congestion at sags with adaptive cruise control systems*. 2451–2457.

Rakha, H., Farzaneh, M., Arafeh, M., & Sterzin, E. (2008). Inclement Weather Impacts on Freeway Traffic Stream Behavior. *Transportation Research Record*, *2071*(1), 8–18. https://doi.org/10.3141/2071-02

Shang, M., & Stern, R. E. (2021). Impacts of commercially available adaptive cruise control vehicles on highway stability and throughput. *Transportation Research Part C: Emerging Technologies*, *122*, 102897.

Shi, X., & Li, X. (2021). Constructing a fundamental diagram for traffic flow with automated vehicles: Methodology and demonstration. *Transportation Research Part B: Methodological*, *150*, 279–292.

Sparks, E. (2022, May 24). *Is Adaptive Cruise Control a Safe Feature to Use?* Eckell Sparsks Attoneys at Law. https://www.eckellsparks.com/2022/05/24/adaptive-cruise-control-safe-feature-use/

Talebpour, A., & Mahmassani, H. S. (2016). Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research Part C: Emerging Technologies*, *71*, 143–163.

Tesla. (2023). *Autopilot and Full Self-Driving Capability*. https://www.tesla.com/support/autopilot

Tillema, T., Gelauff, G., van der Waard, J., Berveling, J., & Moorman, S. (2017). *Paths to a self-driving future: Five transition steps identified*.

Treiber, M., Hennecke, A., & Helbing, D. (2000). Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, *62*(2), 1805–1824. https://doi.org/10.1103/PhysRevE.62.1805

USDOT. (2019). *Advanced Driver Assistance Systems (ADAS): Cooperative Adaptive Cruise Control (CACC)*. https://www.itskrs.its.dot.gov/sites/default/files/executivebriefings/2019/EB03-%20ACC%20DRIVER%20ASSISTANCE_%2005_16_19_FINAL.pdf

USDOT. (2022). *Preferred Following Distance as a Function of Speed—Function-Specific Automation (Level 1) Applications*. https://highways.dot.gov/research/publications/safety/FHWA-HRT-22-107

Van Arem, B., & De Vos, A. (1997). *The microscopic traffic simulation model MIXIC 1.3*.

Van Arem, B., van Driel, C. J. G., & Visser, R. (2006). The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics. *IEEE Transactions on Intelligent Transportation Systems*, *7*(4), 429–436. https://doi.org/10.1109/TITS.2006.884615

Vander Werf, J., Shladover, S. E., Miller, M. A., & Kourjanskaia, N. (2002). Effects of adaptive cruise control systems on highway traffic flow capacity. *Transportation Research Record*, *1800*(1), 78–84.

VanderWerf, J., Shladover, S., Kourjanskaia, N., Miller, M., & Krishnan, H. (2001). Modeling Effects of Driver Control Assistance Systems on Traffic. *Transportation Research Record*, *1748*(1), 167–174. https://doi.org/10.3141/1748-21

Volvo. (2020, March 19). *Limitations for adaptive cruise control*. https://www.volvocars.com/en-om/support/car/xc40/19w17/article/53495f6e8386b6b9c0a8015150beb4c1

Xiao, L., & Gao, F. (2010). A comprehensive review of the development of adaptive cruise control systems. *Vehicle System Dynamics*, *48*(10), 1167–1192. https://doi.org/10.1080/00423110903365910

Xiao, L., Wang, M., & van Arem, B. (2017). Realistic Car-Following Models for Microscopic Simulation of Adaptive and Cooperative Adaptive Cruise Control Vehicles. *Transportation Research Record*, *2623*(1), 1–9. https://doi.org/10.3141/2623-01