To the Graduate Council:
I am submitting herewith a dissertation written by Ruqing Huang entitled "Network Resilience Analysis and Traffic Control Strategies under Infrastructure Failures and Disaster Response." I have examined the final paper copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Civil Engineering.

Lee D. Han, Major Professor

We have read this dissertation
and recommend its acceptance:

Lee D. Han

Baoshan Huang

Shuai Li

Yang Zhou

Lingyao Li

Accepted for the Council:

Dixie L. Thompson

Vice Provost and Dean of the Graduate School

To the Graduate Council:
I am submitting herewith a dissertation written by Ruqing Huang entitled "Network Resilience Analysis and Traffic Control Strategies under Infrastructure Failures and Disaster Response." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Civil Engineering.

Lee D. Han, Major Professor

We have read this dissertation
and recommend its acceptance:

Lee D. Han

Baoshan Huang

Shuai Li

Yang Zhou

Lingyao Li

Accepted for the Council:

Dixie L. Thompson

Vice Provost and Dean of the Graduate School

(Original signatures are on file with official student records.)

# Network Resilience Analysis and Traffic Control Strategies under Infrastructure Failures and Disaster Response

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Ruqing Huang

May 2025

© by Ruqing Huang, 2025
All Rights Reserved.

*Dedicated to, grandpa.*

# Acknowledgements

Thanks for all the love I've received, from my family, friends and afar.

Special thanks to my advisor, Dr.Han—for the challenges, the guidance, and the spark.

***"Sache saisir de chaque moment sa nouveauté singulière, et n'y prépare pas tes joies."***

*— André Gide, Les Nourritures terrestres (1897)*

# Abstract

This dissertation investigates strategies to enhance the robustness and efficiency of transportation networks under severe disruptions, particularly infrastructure failures and disaster scenarios. It presents a comprehensive framework for monitoring and managing network resilience, and explores the potential of electric vehicles (EVs) to facilitate more effective traffic evacuation. Firstly, the dissertation develops a resilience assessment methodology to evaluate transportation network performance following infrastructure disruptions, exemplified by the closure of the Hernando de Soto Bridge on Interstate 40. Integrating the Resilience Triangle framework with advanced statistical methods, this study underscores the critical influence of large-scale queue distributions and the heavy-tailed characteristics of extreme congestion events. Secondly, the research introduces a percolation-based evacuation strategy designed to prevent rapid gridlock in scenarios of urgent evacuation. Through microscopic simulations, it demonstrates how strategically monitoring and regulating the inflow of vehicles can mitigate widespread spillback congestion and accelerate evacuation efficiency. Lastly, this dissertation evaluates how the unique characteristics of electric vehicles can enhance evacuation processes and overall network resilience. Empirical analyses highlight EVs' superior agility, including accelerated lane-changing capabilities and rapid acceleration compared to traditional gasoline-powered vehicles. Furthermore, a novel model predictive control (MPC) framework integrated with the MOBIL lane-changing model demonstrates how EV agility facilitates efficient lane-changing behaviors, enabling evacuees to utilize available traffic gaps effectively.

The study also explores EVs' implications for autonomous vehicle (AV) performance during evacuations, concluding from experimental data that EVs exhibit lower delays and greater precision in executing AV trajectories. Real-time analyses further illustrate that the operational advantages of EVs can reduce shockwave propagation and enhance traffic throughput during peak demand or emergency conditions. Collectively, the findings contribute novel insights into designing transportation networks capable of withstanding routine stresses and dynamically responding to significant disruptions. The outcomes of this research provide valuable guidance for transportation engineers, planners, and policymakers tasked with developing resilient, efficient, and future-oriented infrastructure systems.

# Table of Contents


**1 Introduction** **1**

1.1 Literature Review . . . . . 2

1.1.1 Resilience of Transportation Networks Under Infrastructure Failures . . . . . 2

1.1.2 Percolation Theory and Evacuation Traffic Control . . . . . 3

1.1.3 Potential of Electric Autonomous Vehicles for Improving Traffic Efficiency under Extreme Events . . . . . 4

1.2 Research Gap . . . . . 5

1.3 Dissertation Overview . . . . . 6

**2 Network Resilience Analysis in Response to Infrastructure Failures** **8**

2.1 Background . . . . . 8

2.2 Traditional Network Resilience Analysis — Triangle Method . . . . . 15

2.2.1 Corridor Travel Time . . . . . 17

2.2.2 Truck Miles Traveled . . . . . 19

2.3 New Metric of Resilience Through the Lens of Distribution . . . . . 21

2.3.1 Theoretical Background of Power-law Distribution . . . . . 21

2.3.2 Empirical Power-law Distribution of Congestion . . . . . 23

2.4 Theoretical Connections Between Power-law Distribution and Resillence 27

2.4.1 Brownian Motion with Drift Model . . . . . 28

2.4.2 Return (First Passage) Times to Zero . . . . . 29

2.5 Numerical Simulation Approach . . . . . 32
2.5.1 Hypothesis and Simulation Setup . . . . . 33
2.5.2 Replication of Power-law Congestion Distribution in Simulation 36
2.5.3 Simulating the Impacts of Truck Percentage and Ramp Lane . 38
2.5.4 Understanding the Impact of Truck Percentage on Network Resilience . . . . . 40

**3 Percolation-based Evacuation Control Strategies 43**
3.1 Experiment Setup . . . . . 44
3.1.1 Network configuration . . . . . 44
3.1.2 Evacuation trip definition . . . . . 45
3.1.3 Key simulation settings . . . . . 46
3.2 Equilibrium between loading and exiting . . . . . 47
3.2.1 The bounds of loading and exiting rates . . . . . 47
3.2.2 Equilibrium between loading and exiting . . . . . 48
3.3 Significance of Traffic Loading for Evacuation Efficiency . . . . . 50
3.3.1 Percolation of spillback jams . . . . . 52
3.4 Implications for Evacuation Traffic Monitoring and control . . . . . 53
3.4.1 Percolation of congestion with varying loading rates . . . . . 53
3.4.2 Percolation-based evacuation control strategies . . . . . 56
3.5 Implications for MFD-based Traffic Control . . . . . 59
3.5.1 Connections with MFD control . . . . . 60
3.5.2 Implementation for city networks . . . . . 61

**4 Potential of electric autonomous vehicles for improving traffic efficiency under extreme events 63**
4.1 Background . . . . . 64
4.1.1 Torque, power and accelerations . . . . . 64
4.1.2 EV Two-Regime Accelerating Model . . . . . 65

4.1.3 Limitations of ICE Engines and Advantages of EVs in Sustaining Power . . . 67
4.2 Trapezoidal Modeling of Accelerator Pedal Behaviors . . . 68
4.2.1 Pedal profile modeling . . . 68
4.2.2 Average Pedal Position During Acceleration . . . 71
4.3 Empirical characteristics of EV accel/decelerations . . . 73
4.3.1 Delay off standstill . . . 73
4.3.2 Acceleration power . . . 74
4.3.3 Regenerative braking . . . 76
4.3.4 Responsiveness of EVs . . . 77
4.4 Impact of EVs' Agility on Lane-Changing . . . 81
4.4.1 MPC in the spirit of MOBIL . . . 83
4.4.2 MPC Results and Impact Analysis . . . 90
4.4.3 Comparing EV and ICEV in lane change performance . . . 91
4.4.4 Impact of EV penetration on evacuating traffic throughput . . . 92
4.5 Outlook for EAV's Potential in Evacuation . . . 93

**5 Conclusion** **96**
5.1 Summary of Key Findings from Work 1: Network Resilience to Infrastructure Failure . . . 96
5.2 Summary of Key Findings from Work 2: Evacuation Traffic Monitoring and Control Strategy . . . 99
5.3 Summary of Key Findings from Work 3: EV Potential in Improving Evacuation Traffic . . . 100

**Bibliography** **102**

**Vita** **108**

# Chapter 1

# Introduction

Transportation networks serve as the backbone of modern economies and disaster response systems. The ability to ensure efficient and resilient traffic flow is crucial for economic stability, emergency evacuations, and technological advancements. However, these networks are inherently vulnerable to disruptions arising from infrastructure failures, natural disasters, and the adoption of new vehicle technologies. This dissertation investigates how transportation networks respond to such disruptions and explores strategies to enhance resilience and efficiency under extreme conditions.

The significance of this study lies in its interdisciplinary approach to transportation resilience, integrating perspectives from infrastructure resilience analysis, disaster evacuation traffic control, and the impact of electric autonomous vehicles on traffic flow. Existing literature has examined each of these three areas extensively: (1) resilience frameworks are used to measure and improve a network's capacity to withstand disruptions, (2) evacuation research focuses on managing demand and optimizing traffic flows during disasters, and (3) vehicle technology studies reveal how automation and electrification influence network performance. However, few studies synthesize these elements into a unified framework that addresses how new technologies and control strategies can strengthen transportation network resilience in the face of both routine and extreme stresses.

Accordingly, this dissertation aims to bridge this gap by proposing and testing an integrated set of methods and strategies that collectively offer insights on preventing, mitigating, and recovering from disruptive events in transportation networks. Specifically, it explores how to i) **Quantify and model resilience** using both the Resilience Triangle and advanced percolation-based techniques; ii) **Manage evacuation traffic** under highly time-sensitive and congested scenarios, and iii) **Leverage electric and autonomous vehicle (EAV) technology** to enhance overall traffic stability and performance during routine and extreme events.

In doing so, this dissertation seeks to develop a more robust and adaptive transportation system, building on multiple methodological and theoretical foundations while highlighting critical interdependencies across these lines of inquiry.

## 1.1 Literature Review

### 1.1.1 Resilience of Transportation Networks Under Infrastructure Failures

Resilience in transportation networks refers to the system's ability to absorb, adapt to, and recover from disruptions while maintaining operational functionality (Ayyub, 2014; Sharma et al., 2018). Several interconnected concepts define resilience, including i) Reliability measures the probability of a transportation network maintaining expected performance under perturbations (Mattsson and Jenelius, 2015), ii) Vulnerability assesses the network's susceptibility to disruptions and their potential consequences iii) Robustness and Survivability characterize the system's capacity to sustain operations under stress without catastrophic failure (Faturechi and Miller-Hooks, 2015).

The Resilience Triangle Method (Bruneau et al., 2003) is commonly used to quantify how performance degrades and recovers over time, while percolation-based approaches model how disruptions fragment the network (Newman, 2018; Li et al.,

2021a). Yet, existing studies often fall short of capturing the complex, multi-scale mechanisms driving recovery and the real-time interventions that might mitigate cascading failures. Standardized resilience metrics, especially those leveraging real-time traffic data (Ganin et al., 2017; Gu et al., 2022), remain an open research question. Addressing these limitations requires synthesizing graph-theoretic perspectives, resilience triangles, and emerging data sources into a coherent framework for both localized and network-wide disruption scenarios.

### 1.1.2 Percolation Theory and Evacuation Traffic Control

Evacuations under disastrous conditions are time-sensitive, making the optimization of transportation network capacity a critical aspect of emergency management. In vehicle-based evacuation, network performance depends on traffic flow size; higher densities lead to non-linear increases in travel times, creating the risk of congestion or even complete deadlock.

Although there is a substantial body of work on evacuation demand modeling and traffic assignment (Intini et al., 2019; Cascetta, 2009; Murray-Tuite and Wolshon, 2013; Feng and Lin, 2022, 2021; Pel et al., 2010; Brown et al., 2009; Chiu et al., 2008; Barrett et al., 2000; Hsu and Peeta, 2015), most studies focus on long-term or large-scale planning horizons (e.g., hurricanes with multiple days of warning). By contrast, this dissertation examines rapid evacuation from a Central Business District (CBD) during an emergency—a scenario in which surge loading, if poorly managed, can trigger immediate gridlock (Alawadi et al., 2022).

One strategy to mitigate congestion is to manage the rate at which evacuees enter the network. Staged evacuations (Chen and Zhan, 2008; Sbayti and Mahmassani, 2006; Hsu and Peeta, 2014, 2013) distribute departures over time to reduce clearance times. However, many such models adopt aggregated or macroscopic representations of traffic flow, sacrificing detail on how congestion propagates in real-world scenarios. To address these limitations, this study employs a large-scale microscopic simulation

tool (SUMO) (Alvarez Lopez et al., 2018) to capture complex, link-by-link traffic behaviors during emergency evacuation.

Recent research has begun incorporating percolation theory into traffic flow analysis (Li et al., 2015; Laval, 2023; Laval and Menendez, 2023; Ambühl et al., 2023). Percolation theory offers an aggregated, network-level perspective on how congestion evolves, identifying critical thresholds where the network transitions from free-flow to gridlock. Unlike traditional link-by-link methods, percolation-based approaches reveal emergent behaviors—valuable for designing holistic evacuation strategies. For example, perimeter control rooted in Macroscopic Fundamental Diagram (MFD) theory has proven beneficial for network-level traffic management, yet its links to percolation theory remain underexplored. This dissertation thus investigates how percolation-based monitoring and control can be adapted and applied to urgent evacuation scenarios, hypothesizing that the complex system nature of evacuation traffic makes percolation-inspired strategies particularly relevant.

### 1.1.3 Potential of Electric Autonomous Vehicles for Improving Traffic Efficiency under Extreme Events

The rapid development and deployment of electric autonomous vehicles (EAVs) introduces a new dimension to network resilience and efficiency. Although Adaptive Cruise Control (ACC) systems have been widely studied (Li et al., 2021b; Makridis et al., 2020, 2021; Zhou et al., 2022b), their interaction with underlying vehicle performance—particularly the high torque and regenerative braking capabilities of electric vehicles (EVs)—has received comparatively little attention.

Empirical evidence about how EVs' physical performance (e.g., acceleration and deceleration rates, real-time torque measurements) affects traffic flow is sparse. One study by Lapardhaja et al. (2023) found hints that EV platoons under ACC may improve capacity, potentially due to instantaneous torque delivery and regenerative braking. However, that work did not include human-driven modes or detailed

CAN (Controller Area Network) data, leaving many performance-related questions unanswered.

This dissertation contributes new empirical data by deploying vehicle CAN-logger devices and open database (DBC) files (commaai, 2024) to measure torque, commanded acceleration, and other performance parameters in popular EV models. By integrating these high-resolution data with traffic simulations, it becomes possible to isolate the effects of EV hardware from ACC logic, shedding light on how EAVs might transform both regular and emergency traffic conditions. Understanding these dynamics is critical for designing next-generation traffic control strategies—especially under conditions where rapid maneuvers, coordinated braking, and reliable acceleration are paramount for preventing cascading failures.

## 1.2 Research Gap

Although resilience analysis, evacuation traffic control, and emerging vehicle technologies have each been the subject of extensive studies, there remains a critical need to integrate these domains into a cohesive framework. Resilience assessments often focus on either high-level metrics (e.g., the Resilience Triangle) or graph-based approaches (e.g., percolation), with limited attention paid to the underlying traffic dynamics that drive recovery processes. Meanwhile, evacuation research has predominantly emphasized travel demand modeling and staged evacuation planning, often using aggregated or macroscopic models that fail to capture the nuanced congestion patterns evident in large-scale emergencies. Emerging vehicle technologies, notably electric autonomous vehicles (EAVs), introduce distinct capabilities such as high torque delivery and regenerative braking that could significantly alter network performance; however, the lack of high-resolution empirical data leaves open questions about how these advantages might translate into real-world benefits—particularly under disruptive scenarios.

This dissertation addresses this gap by combining multi-scale resilience methodologies, percolation-based traffic control strategies, and in-depth empirical analyses of EAV performance. Through this interdisciplinary lens, the research aims to both clarify how networks behave under stress and propose actionable measures—ranging from infrastructure-focused interventions to advanced vehicle utilization—to improve overall system robustness. In doing so, it responds directly to the growing need for transportation systems that can withstand, adapt to, and quickly recover from disruptions, whether they stem from infrastructure failures, natural disasters, or technology-driven shifts in travel behavior.

## 1.3 Dissertation Overview

Collectively, the three core themes—transportation network resilience, evacuation traffic control, and EAV performance—form a cohesive story about how to prepare for and respond to disruptions.The overarching contribution of this dissertation is to show how these seemingly separate strands of research can be combined into an actionable framework. By examining how resilience metrics guide evacuation strategies, and by factoring in the capabilities of electric autonomous vehicles, the dissertation aims to provide both theoretical insights and practical tools for transportation planners, engineers, and policymakers. Ultimately, the findings from this research will help shape transportation networks that can quickly adapt to and recover from disruptions—from everyday incidents to large-scale disasters.

The dissertation is organized as follows:

i) **Chapter 2: Network Resilience Analysis in Response to Infrastructure Failure** This chapter develops a systematic framework for assessing and quantifying the resilience of transportation networks when faced with infrastructure disruptions. Drawing upon resilience metrics (such as the Resilience Triangle) and graph-based analyses, it identifies critical components,

explores percolation-based vulnerability, and demonstrates how multi-scale methods and real-time data can inform proactive interventions.

ii) **Chapter 3: Percolation-based Evacuation Control Strategy** Chapter 3 examines how different traffic loading rates affect large-scale evacuations. Through simulation and percolation-based analysis, this chapter reveals how excessive loading can trigger widespread spillback congestion, while maintaining an optimal inflow–outflow balance promotes efficient evacuation. Strategies rooted in percolation theory help traffic managers avoid catastrophic gridlock by monitoring and controlling loading levels. These insights set the stage for Chapter 4, where emerging vehicle technologies, particularly electric autonomous vehicles, are investigated as another avenue for enhancing network robustness and efficiency under extreme conditions.

iii) **Chapter 4: Potential of Electric Autonomous Vehicles for Improving Traffic Efficiency under Extreme Events** Finally, Chapter 4 explores how emerging vehicle technologies—specifically electric autonomous vehicles (EAVs)—could enhance overall network performance and resilience. By analyzing real-time vehicle data (via CAN-logger devices and open DBC files), it assesses how EV torque delivery, regenerative braking, and adaptive cruise control affect traffic flow. These findings offer insights into integrating EAVs into both everyday operations and disaster-response scenarios.

By examining network resilience, evacuation strategies, and the role of innovative vehicle technologies, the dissertation provides a multi-dimensional view of how to design and manage transportation systems capable of withstanding and recovering from disruptive events.

# Chapter 2

# Network Resilience Analysis in Response to Infrastructure Failures

The first major study of this dissertation evaluates freight network resilience in response to infrastructure failures, using the 2021 Hernando de Soto Bridge closure as a case study.

## 2.1 Background

The Hernando de Soto Bridge, a major transportation link between Memphis, Tennessee, and West Memphis, Arkansas along Interstate 40, was shut down from May 11, 2021, to July 31, 2021, due to the discovery of a structural fracture. As a key corridor supporting substantial freight movement essential to both regional and national supply chains, its closure required rerouting traffic through alternative crossings. This diversion led to increased congestion and extended travel times on surrounding routes. The present study investigates the broader implications of such infrastructure failures on freight transportation networks and evaluates the system's resilience to major disruptions.

To conduct this analysis, a novel method was developed to assess network resilience by examining the statistical distribution of congestion levels throughout various phases of the bridge closure. Key aspects of this evaluation included measuring the frequency and duration of extreme congestion events, with particular emphasis on sustained traffic delays. This approach enabled a deeper understanding of both the severity and persistence of congestion, offering insights into the network's capacity to absorb and manage unexpected traffic surges. By quantifying these trends, the study outlines how quickly traffic patterns returned to normal and identifies critical congestion thresholds that might lead to substantial economic consequences from freight delays.

The research combines real-world data collection with structured traffic-pattern analysis, examining freight movement and travel conditions before, during, and after the bridge closure. Data from the Tennessee Department of Transportation's (TDOT) Radar Detection System (RDS) were utilized, which record traffic volume, speed, and occupancy at 30-second intervals across 150 monitoring points in Memphis. Extending the study beyond the closure period provided a robust comparison of pre-closure baseline conditions, immediate responses to the disruption, and the gradual return to normal traffic operations. By segmenting the analysis over time, the study identifies notable shifts in congestion patterns directly linked to the bridge outage and highlights any lasting effects on freight traffic.

To conduct a systematic and effective analysis of the bridge closure, a set of key assumptions guided the approach:

- **Traffic Rerouting:** It is assumed that most traffic previously relying on the Hernando de Soto Bridge was redirected to alternative crossings such as I-55, resulting in higher congestion and longer travel times on those routes.

- **Temporal Phases:** The analysis considered three primary time phases—pre-closure, closure, and post-closure—to track changes in congestion and freight

movement. This approach allowed a detailed examination of both disruption impacts and subsequent recovery.

- **Data Integrity:** The study assumed that the collected traffic data provided a comprehensive view of actual conditions, while acknowledging limitations in resolution and completeness.

The study period spans from March 1, 2021, to September 30, 2021, covering two months prior to the closure, the closure period itself (May 11, 2021, to July 31, 2021), and the subsequent recovery phase. This timeline ensures a thorough assessment of both immediate disruptions and longer-term recovery trends in Memphis-area traffic patterns.

**Figure 2.1** provides an area-wide perspective on traffic performance by depicting the average speed across approximately 150 detectors relative to a desired freeway speed (e.g., 70mph). The shaded red region highlights episodes when observed speeds dropped below this target, reflecting congestion or slowdowns that deviate from normal flow conditions.

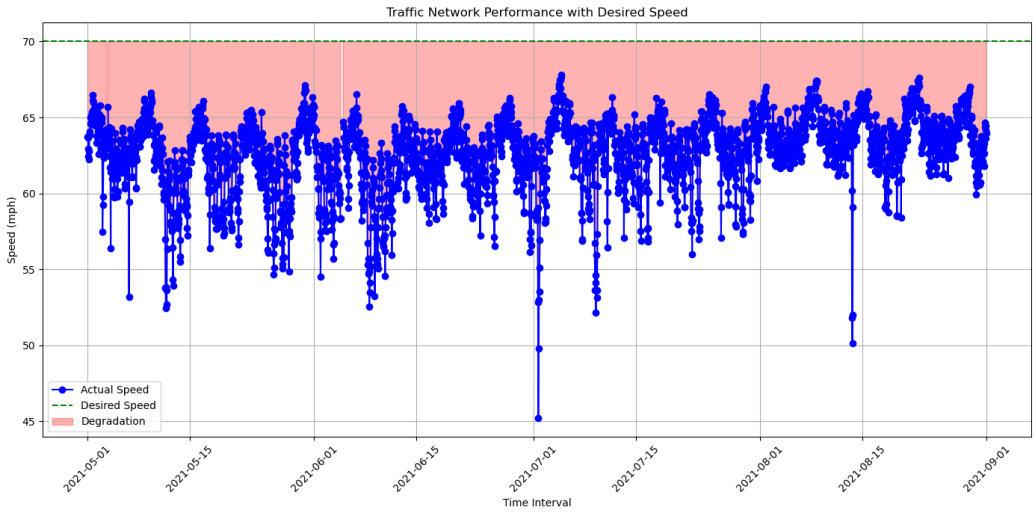



Figure 2.1: Traffic Network Performance with Desired Speed

**Figure 2.2** visualizes reveals how the Memphis network reacted unevenly to the I-40 bridge closure. Areas on the perimeter of the metropolitan region generally show little or no slowdown (green to blue circles), likely because they benefit from multiple alternate roads that disperse traffic. By contrast, segments around the Mississippi River–crossing corridor exhibit larger negative ratios (yellow to red), indicating notable speed reductions.The contrast explains why our analysis zeroes in on the river-crossing corridor:

- **Limited Redundancy** Unlike outlying roadways that channel traffic across a variety of parallel routes, the Mississippi River bridges (especially I-40 and I-55) represent critical links with fewer backup options. When one crossing is closed, the remaining corridor must handle a disproportionate load, making it more vulnerable to slowdowns.

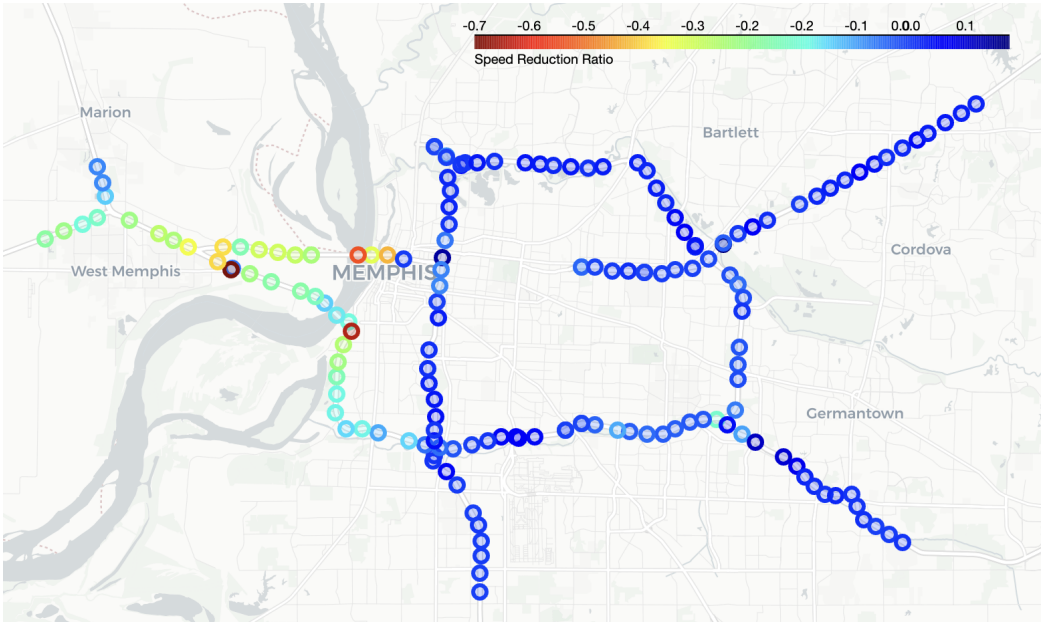


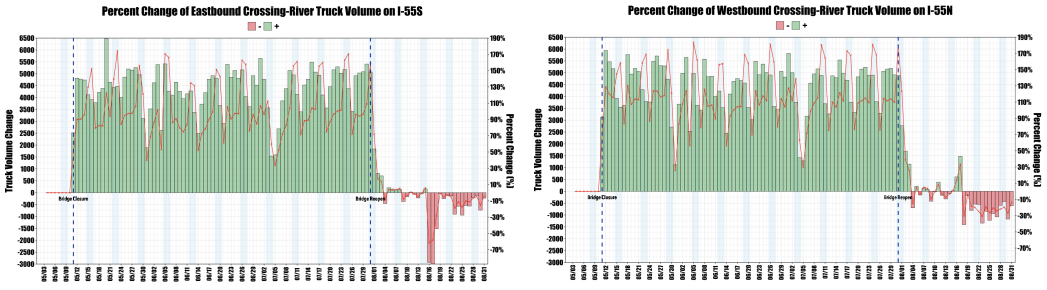

Figure 2.2: Traffic Network Performance with Detector Speed Reduction Ratio

- **Bottleneck Effects** With the I-40 bridge out of service, average speeds on the I-55 crossing dropped significantly, reflecting the corridor's limited capacity to absorb rerouted traffic. These sharp dips in the speed reduction ratio highlight where congestion intensified the most.

- **Before–After Comparisons** By comparing average speeds before and after the closure, we can more precisely gauge how traffic shifted to I-55 and assess whether the network was able to compensate. Perimeter segments did not suffer as much because they could tap into alternate pathways, diluting the impact.

**Figure 2.3** presents the percentage change in truck volumes on the I-55 corridor (both eastbound and westbound) relative to a baseline week prior to the Hernando de Soto Bridge closure. This visualization highlights how freight activity shifted onto I-55 after the I-40 crossing became unavailable, and how quickly volumes reverted once I-40 reopened.

- **Immediate Impact and Sustained Traffic Increase** Beginning May 11, I-55 truck volumes surged sharply (green bars) in both directions. With I-40 closed, much of the region's freight traffic rerouted to I-55, underscoring its critical role as the main alternate crossing. These higher-than-baseline truck volumes persisted until I-40 fully reopened, suggesting that I-55 successfully bore an extended period of added freight demand.




(a) Eastbound (b) Westbound

Figure 2.3: Percentage Change of Cross-River Truck Volume

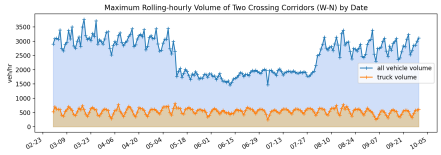
- **Return to Baseline Post-Reopening** Following the reopening of the I-40 bridge on August 2, truck traffic on I-55 quickly dropped—on some days dipping below the original baseline (red bars). This immediate traffic shift back to I-40 reflects the strong preference among freight operators for the more direct corridor once it became available, thereby relieving I-55 of its extra load.

- **Resilience and Infrastructure Redundancy** The ability of I-55 to absorb rerouted volumes during the closure underscores its function as a resilient fallback corridor. In a network sense, having at least one alternative crossing can mitigate a total shutdown of freight traffic. However, the sustained pressure on I-55 also illustrates the vulnerabilities inherent in a single crossing taking on all diverted freight. To maintain robust freight continuity in future disruptions, multiple redundancies are essential.

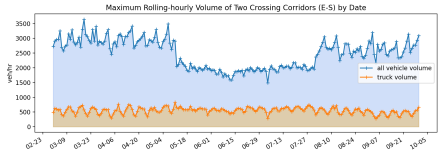
Next, we extend this corridor-level view by examining maximum rolling-hourly volumes and daily total volumes for both truck and overall traffic. These additional metrics shed light on capacity limits, load distribution across time, and the broader network's recovery trends once I-40 service resumed.

**Figure 2.4** compares maximum rolling-hourly traffic volumes (top row) and daily total volumes (with a 5-day rolling average, bottom row) for both east–south (E–S) and north–west (N–W) corridors. Several key patterns emerge regarding capacity, load management, and the network's post-closure recovery:

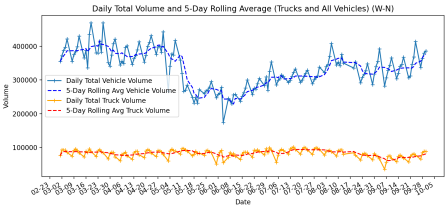
- **Immediate Post-Closure Drop** Immediately after the May 11 bridge closure, peak hourly volumes dropped significantly, especially in overall vehicle counts. This decline reflects reduced capacity for crossing the Mississippi River as traffic shifted to alternate routes. While truck volumes also fell, they remained relatively stable—highlighting the critical role of the I-55 corridor in freight movement.



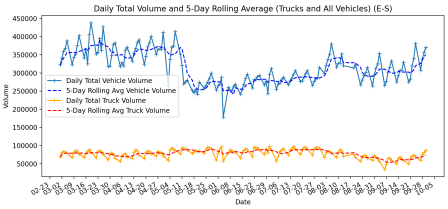

(a) Rolling Hourly Volume (W-N)




(b) Rolling Hourly Volume (E-S)




(c) Rolling Daily Volume (W-N)




(d) Rolling Daily Volume (E-S)

Figure 2.4: Rolling Hourly and Daily Cross-river Traffic Volume

- **Daily Volume Declines and Gradual Recovery** The daily volume plots show a system-wide dip in both truck and total vehicle counts after the closure. Once the Hernando de Soto Bridge reopened on August 2, these figures trended upward again. The 5-day rolling average helps confirm this steady normalization, demonstrating how drivers and freight carriers quickly resumed prior travel patterns once the corridor's capacity was restored.

- **Capacity and Load Management** By tracking the "maximum rolling-hourly volume," we gauge each corridor's highest load within a given hour. During the closure, these maximum volumes remained lower than normal, highlighting the challenge of accommodating peak-demand periods with one major crossing out of operation. The subsequent climb back toward pre-closure peaks, however, underscores the infrastructure's ability to handle typical demand levels once all lanes were available.

- **Resilience in Freight and Emergency Planning** Despite the overall disruption, truck volumes showed noteworthy consistency, reflecting the essential nature of freight movement through Memphis. This points to the resilience requirements of freight corridors, especially those crossing the Mississippi

River, which must continue supporting goods transport even under constrained conditions. Meanwhile, the smooth rebound of both hourly and daily volumes after reopening illustrates the network's adaptive capacity and confirms the value of redundant crossings. When critical links fail, the presence of alternative corridors can help maintain stability—vital for both routine operations and emergency scenarios.

Building on the preliminary examination of raw traffic patterns and corridor flows, the widely used Resilience Triangle Method is adopted for assessing the network's ability to absorb and recover from the I-40 bridge closure. This foundational framework quantifies both the immediate impact of a disruption and the subsequent recovery trajectory, thus offering a structured perspective on how effectively the system managed the loss of a major crossing. Subsequent sections then refine this perspective by introducing an expanded distribution-based approach to capture the nuances of congestion behavior not fully addressed by the Triangle Method alone.

## 2.2 Traditional Network Resilience Analysis — Triangle Method

The resilience triangle is a widely used framework for evaluating how quickly a transportation network can recover from disruptions. It primarily focuses on the temporal aspect of resilience, analyzing how long it takes for the system to return to normal after a disturbance. By examining performance degradation over time, the resilience triangle method provides fundamental insights into how infrastructure failures impact network stability and recovery.

Applying this method to the Hernando de Soto Bridge closure, we first assessed freight network resilience in the river-crossing corridor using a traditional time-series analysis. This approach allowed us to quantify the duration, severity, and overall effect of the disruption on traffic patterns. We examined key metrics such as

congestion duration, travel time recovery, and freight-flow restoration to determine how effectively the network withstood and adapted to the closure.

While the resilience triangle effectively captures the overall recovery trajectory, it does not provide deeper insights into the underlying distribution of disruptions across space and time. The traditional approach addresses the "when" and "how long" of recovery but lacks the granularity needed to assess localized variations and the frequency of extreme congestion events.

In urban highway networks, resilience is typically evaluated using multiple quantitative indicators rather than a single standardized measure. Researchers have introduced various metrics—delay, capacity, speed, travel time, and cost—to assess resilience in these networks (Serulle et al., 2011; Martinez-Pastor et al., 2016; Wu et al., 2016). The choice of metric depends on factors such as the nature of the disruption, mode of transportation, analytical framework, and study objectives, reflecting both the complexity of transportation systems and their diverse reactions to different perturbations.

The conventional method evaluates resilience by measuring the area beneath the performance curve over time. It provides insights into the severity and duration of service disruption, as well as the subsequent recovery process. Therefore resilience is often defined as:

$$R = \frac{\int_{t_0}^{t_1} P(t)\, dt}{(t_1 - t_0)P_0}, \tag{2.1}$$

where $P$ is the system performance metric, $t_0$ is the time the disruptive event occurs, and $t_1$ is when the system reaches full recovery.

**Figure 2.5** illustrates how system performance degrades following a disruptive event and then recovers over time. The y-axis represents the performance metric (or functionality), and the x-axis represents time. In this diagram:

- **Initial Performance (pre-disruption):** Represented by the top horizontal line.

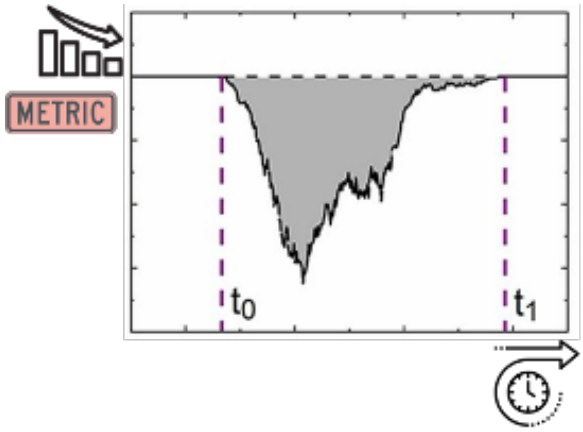



Figure 2.5: The Traditional Resilience Triangle

- **Disruption Impact:** A rapid drop in performance forms the left side of the triangle.
- **Recovery Phase:** A gradual return to near-original performance levels.
- **Shaded Area:** The total performance loss during disruption and recovery, illustrating the resilience "gap."

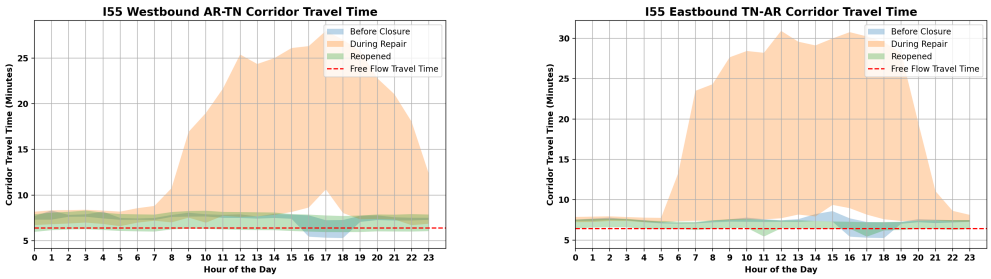
In traffic resilience studies, the shaded area indicates the cumulative performance loss. A smaller shaded area corresponds to a more resilient system, as it experiences minimal degradation and/or recovers quickly. A larger shaded area suggests greater vulnerability, with a significant performance loss persisting over a longer period.

### 2.2.1 Corridor Travel Time

The Resilience Triangle methodology captures the degree to which the I-55 corridor was impacted by the diverted traffic from the I-40 bridge closure. In this application, increased travel time during the disruption translates into a "flipped triangle," where the area under the curve represents the depth and duration of performance loss. A more expansive area indicates a more severe overall impact, encompassing both the magnitude and the persistence of the disruption.

The two panels in **Figure 2.6** illustrate the travel-time distributions along I-55 for both westbound and eastbound directions, captured by ribbons spanning the 15th to 85th percentiles. This format highlights the diurnal range of travel times across three distinct phases. Under normal conditions, the corridor remains near its free-flow travel time (red dashed line) with minimal variability. Once the Hernando de Soto Bridge closed, however, the orange ribbons expand both in height and width, indicating substantially elevated and more volatile travel times that persist through midday and into the early evening. After the I-40 bridge reopened, the blue ribbons contract and shift closer to baseline levels, reflecting the corridor's gradual return to stable operation.

Although the overall pattern of significant increase during closure and steady recovery afterward appears in both directions, subtle differences emerge. The westbound corridor occasionally exhibits an earlier peak, which may align with morning freight and commuter flows converging toward Memphis, whereas eastbound travel times can reach their highest levels later in the day. These directional nuances emphasize that even within the same corridor, congestion dynamics can vary by direction, influenced by distinct freight schedules, commuter patterns, and state-to-state traffic regulations.



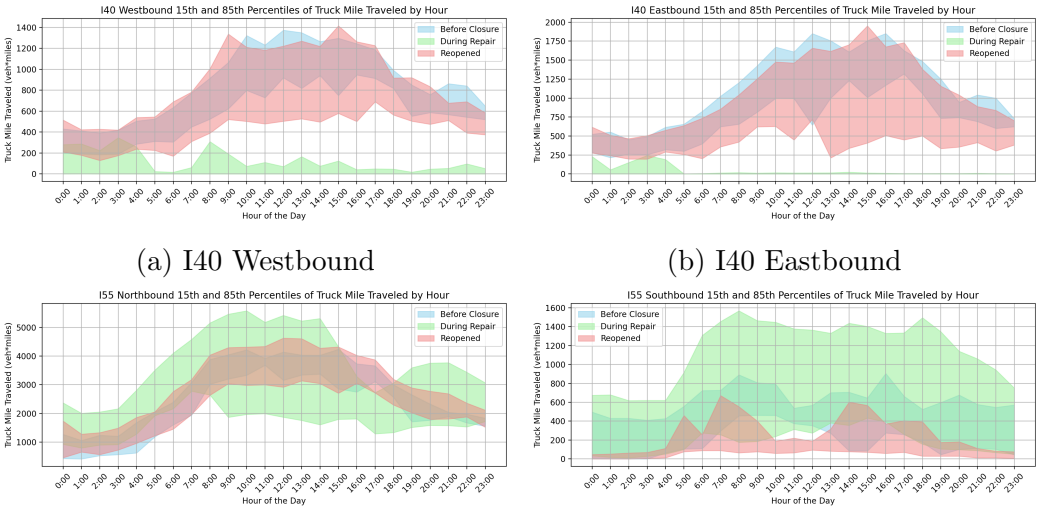

(a) Westbound (b) Eastbound

Figure 2.6: Bi-directional Corridor Travel Time

### 2.2.2 Truck Miles Traveled

**Figure 2.7** illustrates 15th-to-85th percentile ribbons of hourly truck-miles traveled for each direction of I-40 (top) and I-55 (bottom). The plots contrast three phases—Before Closure (blue), During Repair (green), and Reopened (red)—to demonstrate how truck flows adapted once the Hernando de Soto Bridge became unavailable, and subsequently when it returned to service.

- **I-40 Westbound and Eastbound (Closed Corridor)** Before closure, blue bands for both directions on I-40 remain stable, reflecting the corridor's typical free-flow conditions. During the closure, truck-miles dropped notably—especially from 6AM to 6PM—indicating that I-40 was unavailable and freight traffic had rerouted (the faint green bands represent construction traffic.). After reopening, red bands show a sharp recovery in hourly truck-miles, signaling that freight operations quickly returned to I-40. By late afternoon and evening, red-band values closely match or exceed pre-closure levels, underscoring both the restored corridor capacity and the freight sector's rapid response.




(a) I40 Westbound

(b) I40 Eastbound

(c) I55 Northbound (W-N)

(d) I55 Southbound

Figure 2.7: Truck Mile Traveled by Hour for Two Corridors

- **I-55 Northbound and Southbound (Functional Corridor)** The opposite trend appears on I-55. Before the closure (blue), truck traffic is relatively moderate, suggesting that many carriers favored I-40. However, once I-40 shut down, I-55 effectively absorbed that demand, causing the red bands to balloon from early morning into late afternoon. This increased truck mileage highlights the corridor's critical role in facilitating through-region freight movement. The red band (post-reopening) exhibits a partial reversion toward the lower, pre-closure (blue) levels, but it does not fully return to the original baseline. This slight elevation suggests that a portion of freight may have continued using I-55 even after I-40 came back online, possibly due to lingering congestion on I-40, newly established logistics patterns, or driver familiarity with the alternate route.

- **Resilience and Infrastructure Redundancy** Overall, these hourly truck-mile patterns demonstrate the adaptive resilience of Memphis's river-crossing network. I-55 carried the brunt of diverted freight during the closure, maintaining at least a usable corridor despite increased travel loads and variability. Meanwhile, the fast rebound on I-40 after reopening confirms that once primary infrastructure is restored, traffic typically gravitates back to its preferred corridor. The slight residual usage of I-55 post-reopening also suggests there may be longer-lasting shifts in routing behavior for some carriers, albeit on a smaller scale. From a planning and resilience perspective, these findings reaffirm the value of redundancy: I-55's capacity to handle drastically elevated truck miles helped avert a far more severe disruption, even though congestion and variability inevitably rose in the short term.

## 2.3 New Metric of Resilience Through the Lens of Distribution

### 2.3.1 Theoretical Background of Power-law Distribution

BBuilding on the theoretical framework by Pan et al. (2021), this advanced approach examines congestion metrics—such as the duration, length, and area of congestion events—that often follow a long-tail distribution. By analyzing these distributions, one can gain deeper insights into both frequent minor disruptions and rare but severe events, providing a more granular understanding of resilience in transportation systems. The effect of different tail exponents $\alpha$ on power-law distribution behavior is illustrated in **Figure 2.8**.

For a variable $X$ representing congestion duration, length, or area, the long-tail (Pareto) distribution is typically expressed as: where:

- $x$ is the threshold value of interest for the congestion variable.
- $x_{\min}$ is the minimum threshold beyond which the power-law behavior applies.
- $\alpha$ is the tail exponent or shape parameter of the distribution.

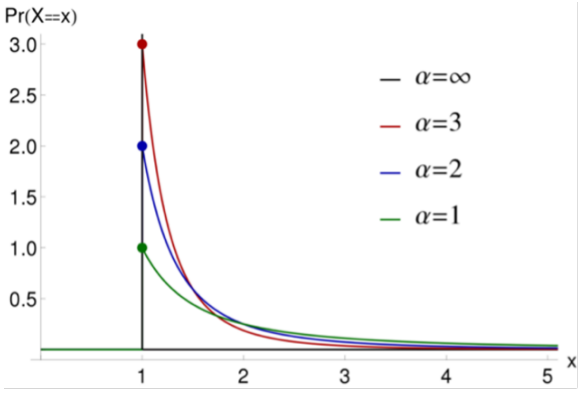



Figure 2.8: The Power-Law/Pareto Distribution

**Interpretations for Different Congestion Metrics:**

- **Congestion Duration** ($X = T$)**:** Most events are short-lived, but there is a non-negligible probability of very long-lasting congestion.
- **Congestion Length** ($X = L$)**:** Most congestion events affect small segments of the roadway, but some can stretch over large areas of the network.
- **Congestion Area** ($X = D$)**:** The total "area" of congestion in a time-space diagram (or the area under the resilience curve), which can remain small for most events but become large in rare, high-impact incidents.

The probability density function (PDF) of a continuous power-law distribution often takes the form:

$$f(x) = C\, x^{-\alpha}, \quad x \geq x_{\min}, \tag{2.2}$$

where $C$ is a normalizing constant, typically

$$C = (\alpha - 1)\, x_{\min}^{\alpha - 1},$$

ensuring that the total probability integrates to 1 over $([x_{\min}, \infty))$.

**Impact of the Tail Exponent $\alpha$**

The tail exponent $\alpha$ determines how likely extreme values are:

- **Smaller $\alpha$ (near 1):** Indicates high susceptibility to severe, long-lasting congestion events.
- **Moderate $\alpha$ (2–3):** Large events can still happen but are less frequent, indicative of a moderately well-managed network.
- **Larger $\alpha$ ($> 3$):** Suggests a more resilient system in which severe congestion events are rare and dissipate quickly.

In traffic systems, the exponent $\alpha$ thus serves as a key resilience indicator. Lower values correlate with a higher likelihood of long-duration congestion, whereas higher values indicate greater stability and quicker recovery.

### 2.3.2 Empirical Power-law Distribution of Congestion

The power-law distribution framework captures the occurrence of congestion events across a wide range of severities, from frequent minor disruptions to rare, high-impact events. This perspective is especially useful in freight-network resilience studies because it highlights both typical and extreme disturbances.

In analyzing resilience along the Memphis river-crossing corridors, we identified congestion "clusters" by focusing on time-space regions where speeds remained below a chosen threshold (e.g., 20 mph, 30 mph, or 40 mph in some cases).

**Figure 2.9** illustrates how these regions are grouped into clusters for delay analysis. By examining these clusters over time, we can measure metrics such as:

- **Cluster Size (Area)**: The integral of queue length over time.
- **Duration**: The total time a particular congestion cluster persists.
- **Spatial Extent**: How far the congestion spreads.

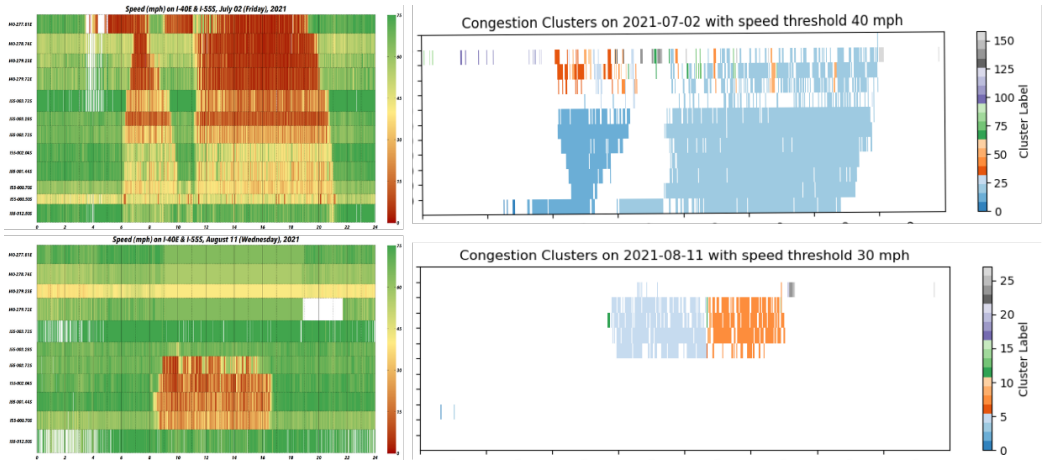



Figure 2.9: Traffic Congestion Clustering for Delay Analysis

This detailed view of congestion events enables a precise assessment of both the frequency and scale of disruptions, supporting a nuanced analysis of how effectively the system manages unexpected surges in traffic.

**Figure 2.10** shows how the closure influenced overall resilience. We analyze three primary phases:

i) **Before Closure (Mar. 1–May 10, slope $\approx -0.87$):** The moderately steep negative slope suggests congestion events typically remain contained. The network manages and disperses most disruptions efficiently under normal conditions.

ii) **During Closure (May 11–Jul. 31, slope $\approx -0.56$):** The shallower slope reflects a higher propensity for bigger or longer congestion clusters. With reduced capacity/routing options, congestion was more likely to persist and grow, reflecting diminished resilience during the closure.

iii) **After Reopening (Aug. 1–Sep. 30, slope $\approx -1.84$):** The very steep slope indicates a quick drop-off in the probability of large congestion events. Once capacity was fully restored, the network again dissipated congestion before it reached extreme levels, suggesting a robust recovery of resilience.

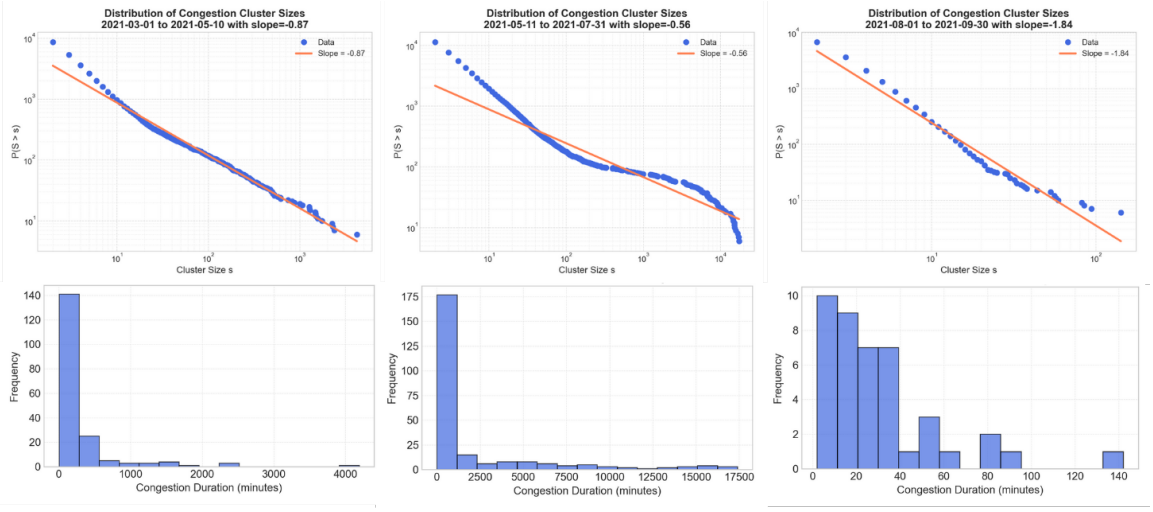



Figure 2.10: System Resilience in Three Phases

**Congestion Peaks During Bridge Closure:** Time-series plots of daily total delay and daily cluster size reveal a notable increase in congestion around May and June 2021, coinciding with the Hernando de Soto Bridge closure. This peak highlights the network's struggle to handle rerouted traffic, causing longer-lasting and more severe congestion clusters.

**Correlation Between Congestion Size and Total Delay:** Both measures peak and decline in tandem, reinforcing the notion that they capture the same underlying traffic condition: how extensively and persistently congestion affects the network.

**Gradual Recovery Post-Reopening:** After the bridge reopened on July 31, both total delay and cluster size dropped steadily. By September 2021, congestion indicators approached pre-closure levels, illustrating the system's capacity to adapt and return to stable operation.

**Implications for Network Resilience:** The extended period of severe congestion during the closure underscores the importance of redundant routes and infrastructure capable of handling unexpected surges in traffic. The post-closure decline in congestion metrics indicates that restoring capacity and employing effective traffic management can significantly bolster a network's ability to recover from major disruptions.

In traffic-flow theory, *congestion duration* represents the time spent under congested conditions—periods when vehicle speeds fall below a specified threshold. Examining congestion duration in these corridors before the I-40 bridge closure shows how quickly the routes recovered from minor disruptions. This measure naturally aligns with the concept of *recovery time*: each congestion event begins with a buildup, reaches a peak, and then dissipates. By tracking congestion duration in different phases, one can assess how effectively the network rebounds after maximum impact. The shift from a moderate slope ($\approx -0.87$) pre-closure to a shallow slope ($\approx -0.56$) during closure, and finally to a very steep slope ($\approx -1.84$) post-reopening, provides

clear evidence of how severely the closure affected resilience and how effectively the network recovered.

A noteworthy observation from the post-closure period (August–September 2021) is that prolonged congestion clusters declined rapidly once the Hernando de Soto Bridge was reopened *and* a new ramp lane came into service (around mid-June). **Figure 2.11** illustrates that, shortly after this additional capacity became available, the peak sizes of congestion clusters dropped, particularly when compared to early June. The shallower tail during closure (slope $\approx -0.56$) began transitioning to a steeper slope (slope $\approx -1.84$) after full restoration. These changes strongly suggest that increasing roadway capacity (or improving network geometry) can alleviate large-scale, long-duration congestion.

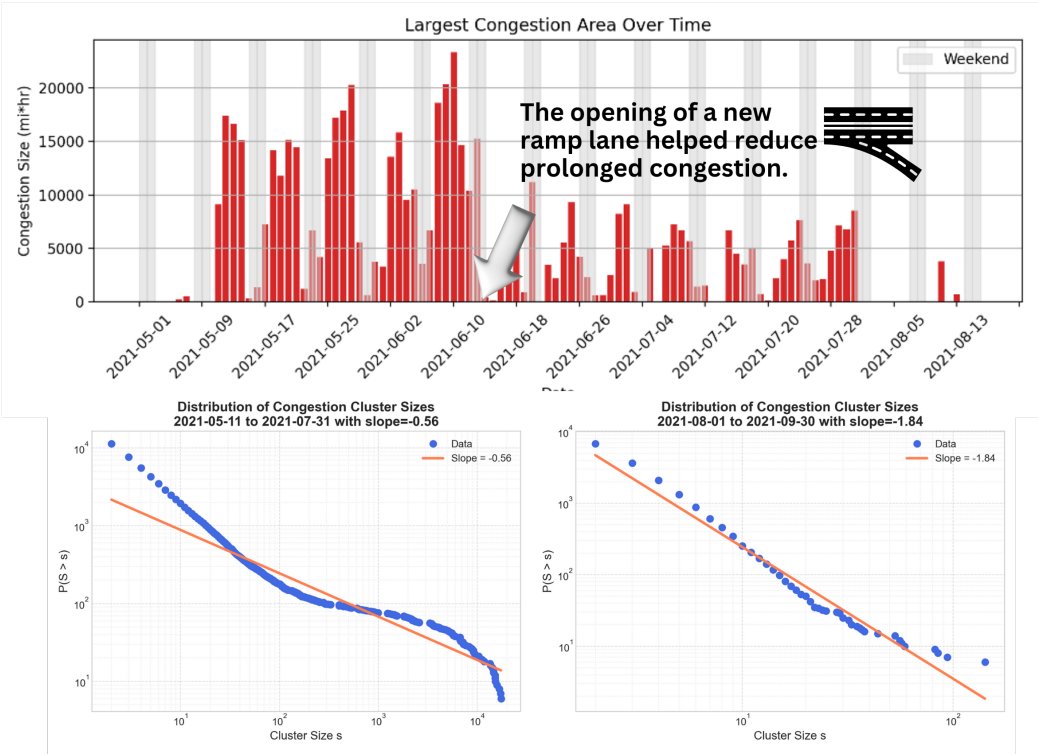



Figure 2.11: Traffic Congestion Clustering for Delay Analysis

While these empirical findings underline the effectiveness of adding a ramp lane to mitigate queuing and bottlenecks, many additional factors can affect resilience. **Truck percentages**, for instance, may shift congestion dynamics by introducing slower-moving, heavier vehicles that can exacerbate bottlenecks. **Network topology** changes, such as closing or adding specific links, can also redistribute traffic in ways that either improve or diminish resilience. Moreover, **traffic control measures** (e.g., lane restrictions for trucks) might further alter how quickly a network can recover from incidents.

## 2.4 Theoretical Connections Between Power-law Distribution and Resillence

In an earlier section, we noted that the time (duration), length (queue), and area (delay) often follow long-tailed distributions (e.g., power laws; see (Li et al., 2015)). By using congestion clusters, we can capture both the spatial and temporal dimensions of these disruptions—namely, contiguous regions in time–space diagrams where speeds remain below a set threshold.

- Time Dimension (Duration). One way to measure a cluster is by how long speeds remain under the threshold—yielding a "congestion duration" distribution.
- Space Dimension (Queue Length). Another measure is the maximum or average queue length within the cluster; however, queue length alone omits how that congestion evolves or persists over time.
- Spatio-Temporal Area (Delay). Finally, one can consider the total area of each congestion cluster in the time–space plane, effectively quantifying the "vehicle-hours" of slowdown for that event. This measure includes both how far the queue spreads and how long it persists.

From a resilience standpoint, identifying how likely it is for a queue to linger (and how severe that lingering can be) adds depth to standard measures like the Resilience Triangle. By quantifying both typical clearance periods and worst-case, heavy-tailed behavior, transportation planners can decide if they need simple, routine mitigations—or if the network is at risk of rare but catastrophic congestion that demands more robust interventions.

In this section, we analyze how the *return time distribution*—the distribution of times at which a traffic queue returns to zero length—depends on the nature of the arrival process. We begin by formulating a stochastic queue-length model in terms of Brownian motion with drift (Sec. 2.4.1), then investigate the corresponding first-passage-time distributions for both stationary (Sec. **??**) and non-stationary (Sec. **??**) arrival processes.

### 2.4.1 Brownian Motion with Drift Model

Consider a bottleneck with *downstream capacity* $q_D$ (vehicles per unit time). Let $q_t$ represent the *upstream flow* or arrival rate (vehicles per unit time). We denote by $L(t)$ the queue length at time $t$. The fundamental traffic-flow principle is that when $q_t > q_D$, the queue tends to grow; conversely, if $q_t < q_D$, the queue tends to shrink.

**Diffusion Approximation.** For large traffic volumes subject to random fluctuations, $L(t)$ can be approximated by a *Brownian motion with drift*:

$$dL(t) \;=\; \mu\, dt \;+\; \sigma\, dB(t), \tag{2.3}$$

where $B(t)$ is a standard Brownian motion, $\sigma$ is the volatility capturing stochastic fluctuations, and

$$\mu \;=\; q_t - q_D \tag{2.4}$$

is the *drift*, corresponding to the average net inflow minus outflow capacity.

### 2.4.2 Return (First Passage) Times to Zero

A key quantity in studying queueing behavior is the *first passage time* (or return time) to zero. Denote by $T$ the random time at which the queue first hits zero,

$$T \;=\; \inf\{\, t \geq 0 : L(t) \;=\; 0\}. \tag{2.5}$$

The distribution of $T$ reveals how quickly a queue can dissipate from a given initial condition $L(0) = L_0 > 0$.

**Negative Drift: Inverse Gaussian.** If $\mu < 0$ in (2.4), meaning $q_t < q_D$ on average, the queue has a net downward drift. In this case, it is well known (see, e.g., **?**) that $T$ follows an *inverse Gaussian* distribution whose density is given by

$$f_T(t) \;=\; \sqrt{\frac{\lambda_{\text{shp}}}{2\pi\, t^3}}\, \exp\!\Big(-\frac{\lambda_{\text{shp}}\,\big(t-\nu\big)^2}{2\,\nu^2\, t}\Big), \quad t > 0, \tag{2.6}$$

where

$$\nu \;=\; \frac{L_0}{|\mu|}, \qquad \lambda_{\text{shp}} \;=\; \frac{L_0^2}{\sigma^2}.$$

Hence, the bulk of the probability mass lies near $t = \nu$, with tail behavior reflecting possible fluctuations away from the mean drift.

**Critical Case: Zero Drift.** When $q_t = q_D$ so that $\mu = 0$, the queue is a pure diffusion with $\sigma > 0$. In this *critical* scenario, first passage times exhibit a heavy-tailed distribution, scaling like $t^{-3/2}$ for large $t$.

**Positive Drift: Exponential Decay in Survival.** When $\mu > 0$ (i.e., $q_t > q_D$), the queue drifts upward on average, making the event of returning to zero less likely. Although the queue can still occasionally return to zero due to randomness, the *survival probability* (i.e., probability that $L(t)$ remains at zero up to some time)

exhibits an exponential decay with rate

$$\alpha \;=\; \frac{\mu^2}{2\,\sigma^2}. \tag{2.7}$$

Equivalently, one may view $\alpha$ as the exponential rate in the large-$t$ behavior for the distribution of $T$.

When the upstream flow exceeds the downstream capacity, the queue length process becomes *transient*. In this regime, the drift drives the system away from zero, so the process *typically grows*, and a return to zero becomes a rare event.

**Exponential Decay of Return Times.** For those rare sample paths where the queue does return to zero, the first passage (or return) time distribution exhibits an *exponential tail*. In particular, if $T$ denotes the first passage time to zero, it can be shown that, for large $t$, the survival probability decays as

$$\mathbb{P}(T > t) \sim \exp\Big(-\frac{\mu^2}{2\,\sigma^2}\,t\Big), \tag{2.8}$$

where

$$\alpha = \frac{\mu^2}{2\,\sigma^2}$$

serves as the *exponential decay rate*. This exponential tail indicates that the probability of extremely long return times decreases rapidly, which is typical of *light-tailed* distributions. However, it is important to note that the overall probability of a return occurring (starting from zero) is less than one in this regime due to the positive drift.

**Comparison with the Critical Case.** In the *critical case* where $q_t = q_D$ (and hence $\mu = 0$), the queue behaves as a pure diffusion process:

$$dL(t) = \sigma\, dB(t).$$

Here, the first passage time distribution is well known to have a *heavy tail*:

$$f_T(t) \sim t^{-3/2} \quad \text{for large } t. \tag{2.9}$$

This power-law decay implies that, although the process is recurrent (i.e., it returns to zero with probability one), the mean return time diverges. The algebraic decay in (2.9) is much slower than the exponential decay in (2.8), meaning that the likelihood of very long return times is significantly higher in the critical case than in the positive drift regime.

**Implications for Traffic Flow.** In practical terms, when the demand exceeds capacity:

- The system is predisposed to queue buildup, with the process being *transient*—returns to an empty state are rare.
- For the few instances when the queue does clear, the distribution of the return time is governed by an exponential decay with rate $\alpha = \mu^2/(2\,\sigma^2)$.
- In contrast, at the critical balance ($q_t = q_D$), although the queue eventually dissipates almost surely, the heavy (power-law) tail implies that the queue clearance time can be extremely long with non-negligible probability.

Thus, while the exponential tail in the positive drift regime ensures that the probability of exceedingly long return times is small, the very low overall probability of return (due to the persistent buildup) means that traffic systems operating in this regime are at high risk of sustained congestion. It is noteworthy that the nature of the tail not only reflects the typical behavior of the system but also has implications for the design of traffic control or resillence management strategies. For instance, in regimes with $\mu > 0$, the focus might be on interventions that reduce the effective drift (e.g., by increasing downstream capacity or modulating upstream flow) to shift

the system closer to criticality, where, despite the heavy-tailed clearance times, the certainty of eventual queue dissipation is guaranteed.

## 2.5 Numerical Simulation Approach

To systematically investigate the role of such factors, we will employ a **microscopic traffic simulation** using SUMO (Simulation of Urban MObility) (Alvarez Lopez et al., 2018). This open-source platform provides a flexible environment for modeling individual vehicle movements under various traffic-control regimes and infrastructural configurations. By adjusting parameters such as:

- *Truck percentages* (e.g., 5%, 15%, 25% of total flow),
- *Network layouts* (road closures, additional lanes, alternative corridors),
- *Control strategies* (e.g., dedicated truck lanes, signal timing changes),

we can examine how each scenario influences the shape of congestion distributions. Specifically, SUMO allows us to track when and where congested clusters form, how long they persist, and how large they become. We will then compare the resulting cluster-size distributions, delay patterns, and resilience metrics across scenarios to determine which strategies most effectively reduce the frequency and severity of extreme congestion events.

This section presents the numerical simulation results designed to evaluate freight network resilience under varying conditions of vehicle composition, network topology, and traffic control strategies. The simulation framework is built in Simulation of Urban MObility (SUMO). A major advantage of microscopic traffic simulation is the ability to capture complete information on the evolving queue length, duration, and total area of congestion clusters. In real-world data collection, these dimensions can be incomplete or only partially observable. The use of numerical simulation allows for precise tracking of queue length, formation, and dissipation dynamics, enabling

the identification and quantification of congestion clusters across three dimensions: spatial extent (length), temporal duration (time), and cumulative severity (area). The analysis focuses on a critical segment near the I-55 bridge, where congestion behavior is sensitive to both infrastructure constraints and traffic flow demand itself.

### 2.5.1 Hypothesis and Simulation Setup

Heavy vehicles often have slower accelerations, longer lane-change times, and larger headway requirements. The central hypothesis is that a higher truck share causes congestion clusters to form earlier and dissolve later, weakening network resilience. Specifically:

- **Lane-Change Delays** Trucks require larger gaps to merge, effectively creating rolling blockages when traffic is dense.
- **Accelerations & Start-Off Times** Heavy vehicles accelerate more slowly than passenger cars, causing lengthened discharge times once a queue forms.
- **Moving Bottlenecks** Two or more trucks traveling parallel at lower speeds can block faster cars behind them, slowing the overall stream and potentially expanding the spatial extent of congestion.

#### Key Simulation Setup

The simulation framework adopts a 24-hour weekday traffic profile, featuring distinct rush-hour peaks (e.g., 6 AM–9 AM, 4 PM–7 PM) that exceed nominal capacity. Demand is generated in 5-minute intervals using Poisson random variables, allowing for realistic fluctuations around each target flow level. Throughput is tracked at multiple loop detectors along the corridor to capture both demand flow (vehicles entering) and downstream flow (vehicles exiting).

**Network Layout and Geometry** The simulated network focuses on the I-55 corridor near the Mississippi River crossing, incorporating essential on/off ramps,

merge points, and lane configurations that mirror real-world conditions, as shown in **Figure 2.12**. In the base scenario, each roadway segment maintains two through-lanes per direction, with a single-lane on-ramp for vehicles entering I-55 upstream of the bridge. This ramp merges into the rightmost lane, creating the potential for localized bottlenecks—particularly when a high proportion of trucks is present. Several loop detectors are placed at strategic points (e.g., upstream demand entry, merging area, and downstream exit) to track both flow rates and queue formation over time.

**Loop Detectors** To accurately monitor the traffic dynamics at key points in the network, virtual loop detectors were embedded into the simulation environment. These detectors, positioned strategically along ramp lanes and mainline segments, enable the collection of detailed temporal flow data, including vehicle counts and composition at high temporal resolution, as shown in **Figure 2.13**.

In particular, detectors were placed:

- **Upstream of the I-55 ramp merge**, to record incoming demand flow (i.e., the total number of vehicles—passenger cars and trucks—attempting to enter the bottleneck).

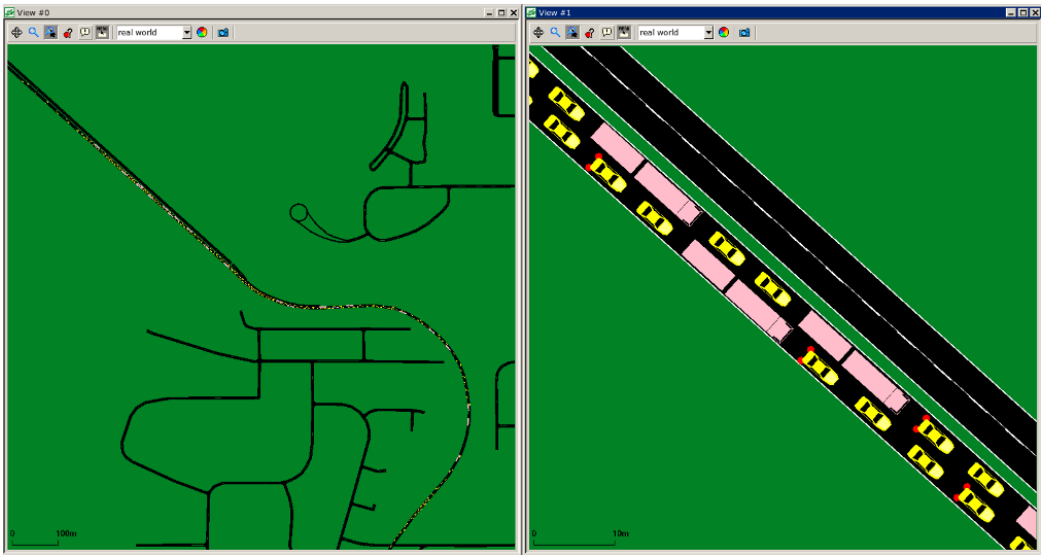
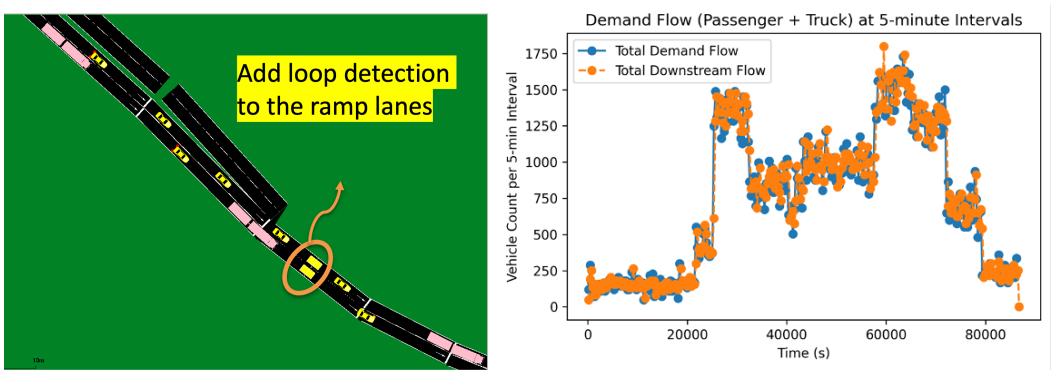


Figure 2.12: I-55 Corridor Geometry Layout in SUMO




(a) Loop Detector Position (b) Demand Flow at 5-minute Intervals

Figure 2.13: Demand Detectors for Traffic Monitoring

- **Downstream of the merge zone**, to track the actual throughput, i.e., the number of vehicles successfully exiting the congested region.

Vehicle counts were aggregated at 5-minute intervals, allowing comparison between demand and throughput over time. This data forms the basis for analyzing capacity drop, bottleneck spillback, and congestion propagation. **Figure 2.13** illustrates a sample time series where demand and throughput curves closely follow each other under free-flow conditions, but diverge during peak periods due to capacity constraints and queue spillback.

By differentiating between total demand and actual downstream flow, the simulation can quantify the extent and persistence of congestion, and evaluate how different conditions (e.g., truck percentage or lane configurations) influence network performance.

**Traffic Demand & Temporal Profile** To accurately reflect the operational characteristics of trucks in the simulation, several key parameters were adjusted to distinguish them from passenger vehicles. These configurations are critical, as heavy vehicles behave differently in acceleration, lane-changing, and car-following scenarios, all of which impact network dynamics.

- **Acceleration and deceleration limits:** Trucks are modeled with lower maximum acceleration (`accel` = 1.0 m/s$^2$) and slightly reduced deceleration

capabilities (`decel` = 3.5 m/s$^2$), reflecting their greater inertia and longer braking distances. Passenger vehicles use higher acceleration (`accel` = 2.0 m/s$^2$) and deceleration values.

- **Car-following behavior:** The `carFollowingModel` is set to IDM (Intelligent Driver Model) for both vehicle types, but with differing sensitivity parameters such as `tau` and `delta`. Trucks are assigned a higher `tau` (headway time), and a greater `delta`, indicating a quicker decay of acceleration response as speed increases.

- **Truck percentage:** The proportion of trucks in the traffic stream is systematically varied across simulation scenarios (0%, 10%, 25%, and 50%) to assess their impact on congestion dynamics and network resilience.

### 2.5.2 Replication of Power-law Congestion Distribution in Simulation

To analyze the congestion clusters, three primary dimensions are considered, which all follow a power-law distribution as discussed in previous sections Ambühl et al. (2023): **width**, the congestion duration measuring how long (in time) congestion persists; **area**, indicating the overall size of the congestion region in the time–space plane; and **height** corresponding to the maximum queue length. By extracting these metrics from each cluster, it becomes possible to construct survival functions for width, area, and height and plot them on a log–log scale. **Figure 2.14** demonstrates how microscopic simulation replicates the power-law distributions often observed in real-world traffic congestion. In the upper portion of the figure, the time-series plot illustrates the total queue length on the I-55 bridge over a given period. Each significant peak exceeding a specified speed or density threshold defines a "congestion cluster."

These clusters can be represented in a space–time diagram by shapes whose horizontal extent corresponds to the duration of the congestion and whose vertical span indicates the queue length. As shown in **Figure 2.14**, the approximately linear shape of these curves in log–log space indicates the presence of power-law tails, signifying that while most congestion events remain small, rare but exceptionally large disruptions can dominate total delay.

From a resilience perspective, these results highlight the "long-tail" nature of urban congestion. Though typical queues dissolve relatively quickly, a small fraction of episodes can persist for long durations or cover significant spatial extents. By systematically adjusting key variables in simulation—such as truck percentages, ramp design, or lane management policies—we can directly observe how these changes impact the distribution of congestion events, particularly in the tail.

This kind of controlled experimentation—exploring "what if" scenarios and measuring their effects on rare, high-impact events—is something that's extremely difficult to achieve with real-world data alone. It demonstrates how simulation serves as a powerful tool for testing resilience-enhancing strategies and making informed infrastructure or policy decisions aimed at preventing system-wide gridlock.

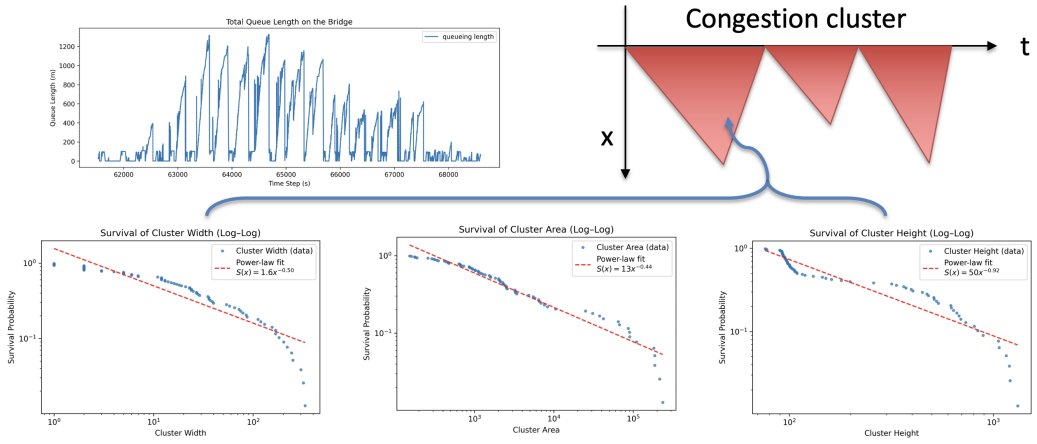



Figure 2.14: Queue length distribution in SUMO with Varying Truck Percentages

### 2.5.3 Simulating the Impacts of Truck Percentage and Ramp Lane

**Figure 2.15** compares congestion distribution under three truck ratios: 0%, 25%, and 50% within a constant upstream flow rate of 2,000 vehicles per hour. The histogram at the top displays queue-length frequencies for each scenario, revealing how higher truck shares shift congestion toward longer queues. Specifically, with 50% trucks (blue bars), the system experiences a disproportionate number of moderate-to-large queues (e.g., above 500,m). In contrast, the 0%-truck case (green bars) rarely encounters extensive queues. This shift arises from trucks' longer start-off times and lower acceleration, which slow the discharging rate of queued traffic.

The lower plots examine power-law distributions of *cluster width* (left) and *cluster area* (right), both plotted on log–log survival axes. Each colored point set corresponds to a distinct truck-percentage scenario, and the dashed lines represent best-fit power-law curves of the form:

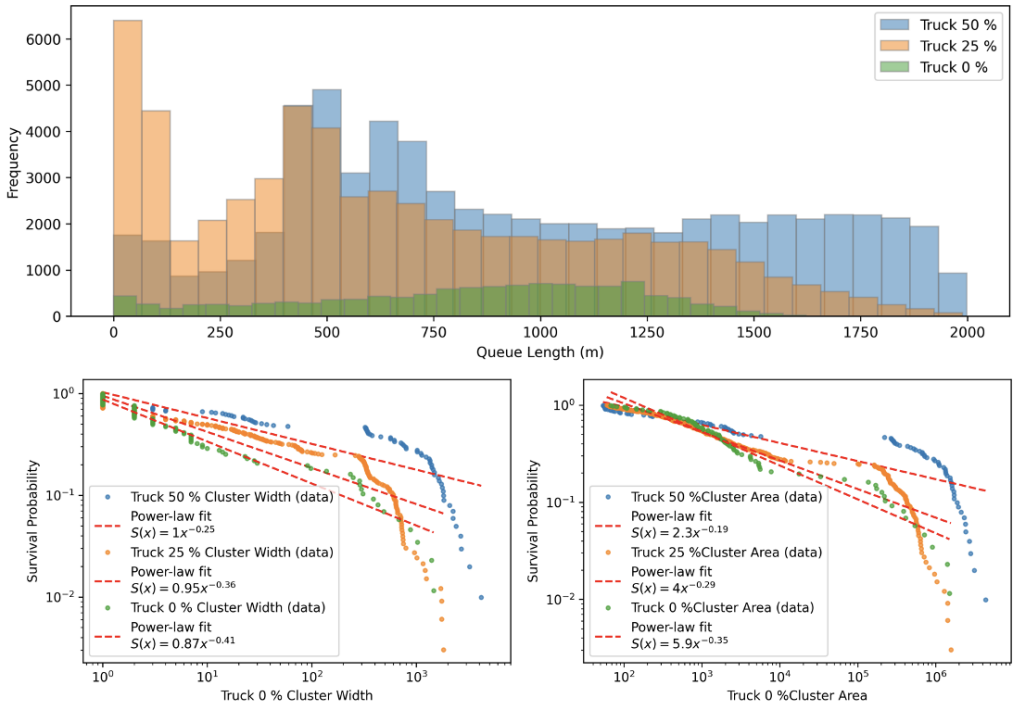



Figure 2.15: Queue length distribution in SUMO with Varying Truck Percentages

$$S(x) = kx^{-\alpha} \tag{2.10}$$

In the 50%-truck case, the slopes ($\alpha$) are notably shallower ($\alpha \approx 0.19$ for area, 0.25 for width), indicating heavier-tail behavior—i.e., a higher probability of extremely large or long-lasting congestion clusters. With fewer trucks (25% or 0%), the curves shift upward and steepen, implying a lower likelihood of severe events. Overall, these findings underscore how an elevated truck share intensifies the formation and prolongs the dissolution of queues, ultimately yielding **more frequent and more extreme** congestion episodes.

A key element of this study involved examining how adding a second ramp lane can mitigate large-scale congestion. **Figure 2.16** (top) shows daily peak congestion areas from May through mid-August 2021, revealing that prolonged, high-magnitude congestion events declined significantly after the mid-June opening of a new ramp lane on I-40. This real-world improvement motivated a series of SUMO simulations to verify whether comparable benefits appear in a controlled environment. In the SUMO network; see **Figure 2.16**; one scenario used a 1-lane ramp merging onto the mainline, while another introduced a 2-lane ramp to handle additional entering traffic. As seen in **Figure 2.16**, the queue-length distribution in the 2-lane scenario shifts notably toward shorter queues—particularly at or above 500,m—compared to the 1-lane case. This outcome indicates that expanding the ramp capacity reduces the likelihood of lengthy backups forming at the merge point.

To further capture extreme congestion events, **Figure 2.16** plots the survival functions (log–log) for congestion-cluster widths in both scenarios. The 2-lane curve shows a steeper power-law slope (approximately -0.43 vs. -0.32), implying that extremely large clusters become significantly less probable. In other words, the two-lane configuration both speeds up discharge at the ramp and lowers the risk of spillback during high-demand periods, reflecting a more resilient network topology.

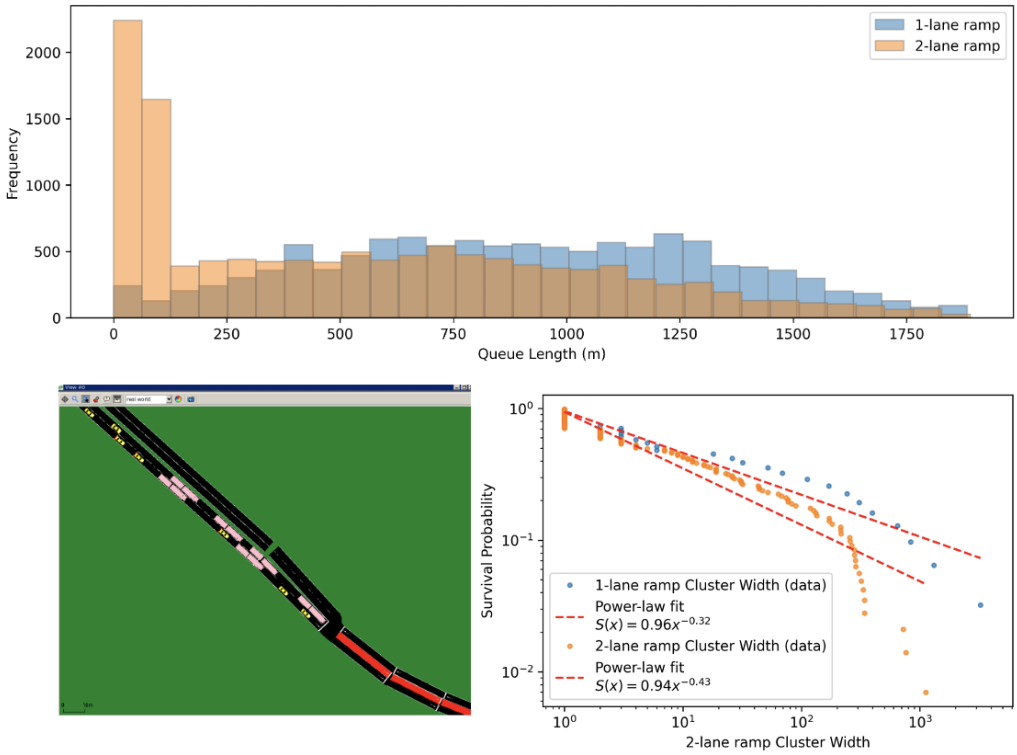



Figure 2.16: Queue length distribution in SUMO in 1-lane and 2-lane Scenarios

Overall, these empirical and simulated findings confirm that adding a ramp lane at a critical merge can substantially reduce the frequency and severity of congestion clusters. From a resilience perspective, removing or alleviating bottlenecks—especially where heavy traffic enters—helps minimize large-scale disruptions, leading to a steeper (more favorable) tail in the congestion-cluster distribution.

### 2.5.4 Understanding the Impact of Truck Percentage on Network Resilience

**Figure 2.17** illustrates how higher truck percentages alter both the formation and dissipation of congestion from a wave-propagation standpoint. In the top schematic, each slanted line represents a shock or rarefaction wave on a time–distance diagram (sometimes called a "jam triangle"). When truck percentages increase (orange lines), the effective discharge rate decreases—i.e., the system's capacity declines—because

trucks accelerate more slowly and take longer to perform lane changes. This reduced outflow causes the queue's dissolution wave to slope more gradually, so congestion clears at a slower rate. Simultaneously, the jam density (vehicles per kilometer) can be lower with trucks because they occupy more space, causing faster deceleration waves when incidents arise. As a result, the queue forms more quickly while also lingering for a longer period.

In the bottom schematic, each colored polygon (orange vs. blue) depicts a congestion cluster in space (horizontal axis) and time (vertical axis). A cluster with more trucks (orange) typically grows rapidly (steep leading edge) and takes longer to subside (flattened trailing edge), leading to a wider, more persistent jam region overall. From a resilience perspective, this dynamic implies that an elevated truck share "hurts both ends" of the congestion process: it speeds up the onset of large queues and slows down their recovery. Consequently, fast-growing and slow-resolving congestion clusters indicate weaker system resilience, whereas corridors with fewer trucks (blue) generally exhibit more robust performance—faster queuing onset can still occur, but the queue tends to dissipate more quickly once capacity is restored.

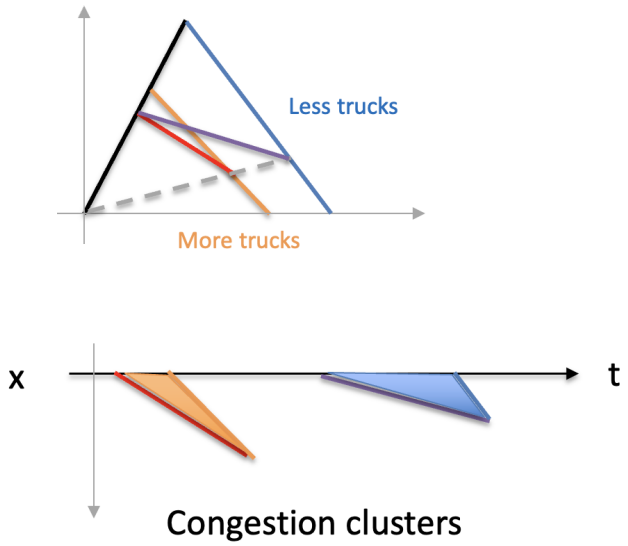



Figure 2.17: Theoretic Visualization Using Fundamental Diagram

This section presented a series of microscopic traffic simulations aimed at evaluating how truck percentages, network topology adjustments, and traffic-control strategies influence overall network resilience. By capturing detailed spatio-temporal data on queue formation and dissipation, the simulations enabled three-dimensional analyses of congestion clusters (duration, spatial extent, and total area), offering a controlled environment in which to assess "what-if" scenarios relevant to high-demand freight corridors.

A key outcome was that increasing truck percentages consistently led to larger and more persistent congestion clusters. In particular, higher truck ratios (e.g., 25% or 50%) produced flatter power-law slopes for cluster-size distributions, revealing a heavier-tailed phenomenon. These findings align with theoretical expectations that limited truck acceleration and slower lane changes slow down queue dissipation, exacerbating the likelihood of extreme events.

Additionally, network topology improvements, such as adding an extra ramp lane, significantly reduced the incidence of severe congestion. The simulations showed shorter maximum queues and steeper tail exponents under two-lane ramp configurations, corroborating empirical observations that the mid-June ramp-lane addition during the Hernando de Soto Bridge closure effectively mitigated widespread delays.

Finally, targeted traffic-control measures, including possible "no-truck-in-left-lane" rules or ramp-metering strategies, further demonstrated potential to stabilize flow and lessen the frequency of extreme congestion. While the exact benefits depended on freight volumes and lane geometry, restricting trucks to specific lanes generally helped avoid turbulent merges, supporting higher throughput and resilience. Collectively, these results underscore the crucial interplay between vehicle mix, infrastructure capacity, and operational policies in shaping how transportation networks withstand and recover from significant disruptions.

# Chapter 3

# Percolation-based Evacuation Control Strategies

Evacuations under disastrous conditions are time sensitive, optimizing the use of the existing transportation network capacity during emergencies is crucial to traffic efficiency. In the context of vehicle-based evacuation, evacuation times are dependent on the network's traffic flow size; as traffic density increases, the evacuation times escalate non-linearly, leading to traffic congestion and potentially resulting in network deadlock.

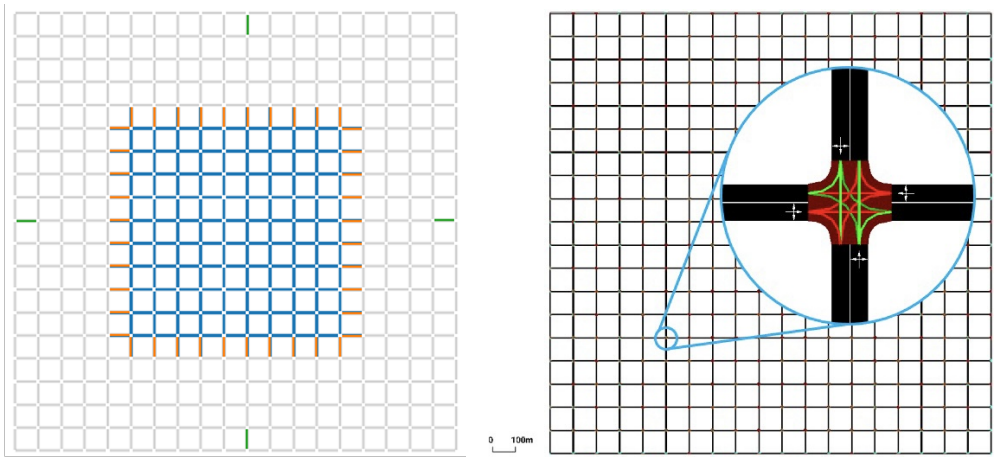
This chapter investigates large-scale vehicle-based evacuations using a simulated 20×20 grid network, where a 10×10 Central Business District (CBD) must be cleared in response to an emergency. The findings reveal that the efficiency of evacuation is highly dependent on the balance between how quickly vehicles are introduced into the network (loading rate) and how effectively they exit the CBD. When this equilibrium is disrupted—especially through excessive loading—it leads to severe congestion spillback, ultimately causing gridlock and slowing down the evacuation process.

## 3.1 Experiment Setup

In this section, we describe the network configurations and OD used for the evacuation simulations in SUMO. We are studying an urgent evacuation scenario where traffic from the central area/CBD needs to evacuate to the outside shelters in hours.

### 3.1.1 Network configuration

The network used is an abstract grid network, which is composed of 400 square cells, laid out in a regular 20 by 20 matrix, as shown in **Figure 3.1**. The grid structure is reminiscent of a checkerboard, where each cell represents a road segment, also referred to as an 'edge', and each point where these edges meet signifies a junction or intersection. The intersections are equally spaced at a distance of 40 meters from each other. Each edge in the grid network serves as a bi-directional connection between two junctions, containing two lanes that accommodate traffic in opposite directions. A Central Business District (CBD), is represented by the central 10*10 grid. This area within the grid symbolizes a concentrated commercial and business hub.




(a) CBD, exiting lanes, and shelters. (b) Road and signal configuration.

Figure 3.1: Network for evacuation simulations

All intersections within the network are equipped with traffic lights that operate on a fixed cycle, with a cycle time of 90 seconds. U-turns are disabled at all intersections. The customized network provides high controllability, and every parameter of the network, from the design features such as the number of lanes and road segment lengths, to operational aspects like the timing of traffic signals and speed limits, can be precisely controlled. This granularity in control enables us to adjust the network to reflect a variety of specific traffic scenarios. Later we will show that the monitoring and control measures developed based on this grid network can be transferred to real city networks easily.

### 3.1.2 Evacuation trip definition

The evacuation process in our grid network commences at time 0, marking the issuance of an evacuation order. As this metaphorical "evacuation clock" starts to tick, every vehicle currently located within the Central Business District (CBD) - the central 10*10 grid - is mandated to evacuate. The orange edges represent exists from the CBD where the evacuation flow must pass. The CBD, under these circumstances, serves as the origin for all evacuation trips. The green edges represent the four designated shelters and serve as the destinations for all evacuation flows. These shelters receive the traffic exiting the CBD. In real-world terms, these green edges could symbolize the entrances to highways or other major exit routes. No traffic exists outside the CBD before the onset of the evacuation; this aspect allows for a clear focus on the evacuation dynamics.

The evacuation demand here is defined as the total number of vehicles to be evacuated from each edge within the CBD. In this paper, the total demand is generated using a Poisson random variable with a mean value of 20, in order to reflect the inherent uncertainty of real-world evacuations.

The loading of evacuees is controlled by the time interval between two successive attempts at vehicle insertion onto a particular edge within the grid. Due to spatial

limitations and safety concerns within the network, not all attempts at vehicle insertion may be successful. This could result in delays, known as 'insertion delays'. However, it's crucial to note that despite these potential delays, no vehicle is discarded from the simulation. Every vehicle is guaranteed eventual insertion into the network. More detailed rules governing vehicle insertion, including handling of insertion delays, will be discussed shortly.

### 3.1.3 Key simulation settings

There are hundreds of simulation settings available in SUMO, here we focus on settings that mostly affect the evacuation simulation. In light of the importance of traffic loading (referred to as insertion into the network in SUMO), we first introduce the **insertion rule by SUMO**: If there is not enough space to insert the vehicle at its designated departure time, that vehicle is placed in an insertion queue and insertion is repeatedly attempted in subsequent simulation steps. If the option `--max-depart-delay <TIME>` is used, vehicles are discarded if they could not be inserted after the specified `<TIME>` in seconds. Apparently, we do not want to abandon any vehicle in the real world even in congestion, hence `--max-depart-delay <-1>`.

By default, SUMO teleports vehicles that have waited too long to avoid deadlock, which is unrealistic. We change this setting by hence `--time-to-teleport <-1>`, which means no vehicles will be teleported.

The driver adaptation level can play a significant role in congestion propagation and deadlock development. After a few trails, all drivers are set to be adaptive with a rerouting period of 60s, considering that non-adaptive drivers tend to deteriorate the simulation too soon with deadlock.

We would like to highlight a few simulation settings that allow us to better track evacuation-specific information. We managed to track the number of pending vehicles that are waiting to be inserted into the network for evacuation. This information

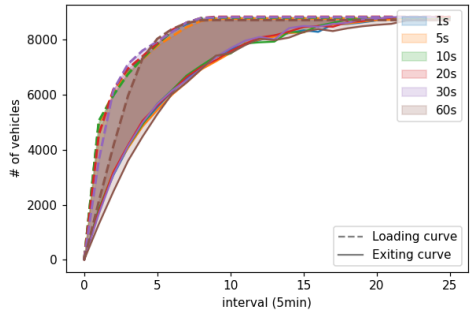
would allow us to better understand the impact of the loading rate on the delay in the insertion of the evacuating traffic. During the simulation, we recorded data from both each signal phase and each 5-min interval for aggregated attributes such as space-mean speed, density, number of arrivals on each link.

## 3.2 Equilibrium between loading and exiting

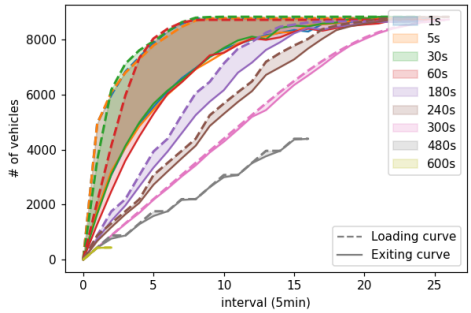
Now we present our simulation results. First, we show the loading, exiting and evacuating histories with varying desired loading rates $L$, defined as the number of vehicles inserted into the road per lane per hour in the evacuation zone. In simulation, the loading rate $L$ is implemented by setting a constant time gap $\tau$ (seconds) between two successive attempts of insertion of the vehicle into the origin of the trip, that is, the departure road in SUMO. The desired loading rate $L$ can be converted directly from $\tau$ to $L = 3600/\tau$ (vehs/hour/lane). The evacuating flow $Q_e$ (veh/h) is defined as the total number of vehicles that have reached shelter destinations per hour, or the sum of the arrival flow at those green links in **Figure 3.1**. Except for the definition of the loading rate and the evacuation flow, we also define the exiting flow, $Q_x$ as the number of vehicles leaving the CBD, which can be considered as the danger zone where vehicles are exposed to high risks.

### 3.2.1 The bounds of loading and exiting rates

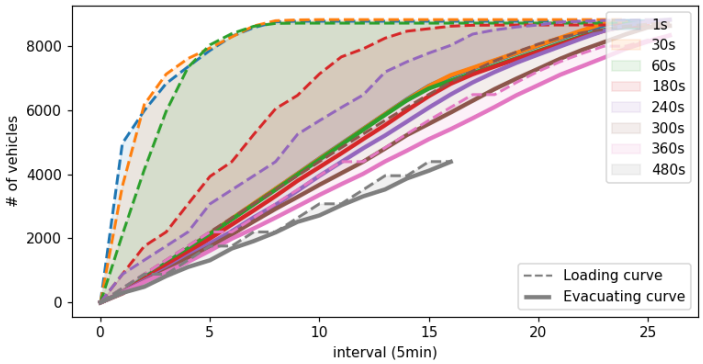
Our first finding is that the loading rate $L$ has its limit. **Figure 3.2** explains the performance of three different loading rates of by monitoring the exiting rate of traffic out of the evacuation zone. When the desired loading gap $\tau_{min} \leq 30$s or equivalently, the desired loading rate $\tilde{L} \geq 120$ veh/hr/lane, the actual loading rate remains the same. It suggests that the CBD, encounters a limit on the loading rate, $L_{max}$, as a result of congestion.




(a) Excessive loading and exiting




(b) CBD exiting rates given varying loading rates




(c) Evacuating curve at different loading rates

Figure 3.2: CBD Evacuating rate and exiting rates at different loading rates

**Figure 3.2a** shows that CBD encounters a limit on exit rate $Q_{x,max}$, despite a possible faster loading from the interior of CBD, all exiting curves with $\tau \leq 60s$ shares almost the same exiting curve, suggesting that this bottleneck effect could originate from the outside region, or the exits are jammed to some extent. **Figure 3.2b** also shows that, when $\tau \geq 60s$, or the loading rate $L \leq 60$ veh / ln / h, the exiting flow of CBD starts to decrease as the loading rate slows down.

### 3.2.2 Equilibrium between loading and exiting

Despite the faster loading and exiting when $\tau \leq 300s$, those flows are never stable, which means that the system has not reached an equilibrium state. In particular, when $\tau = \tau^* = 300$ s, the loading and exit flow are equal, where the evacuating

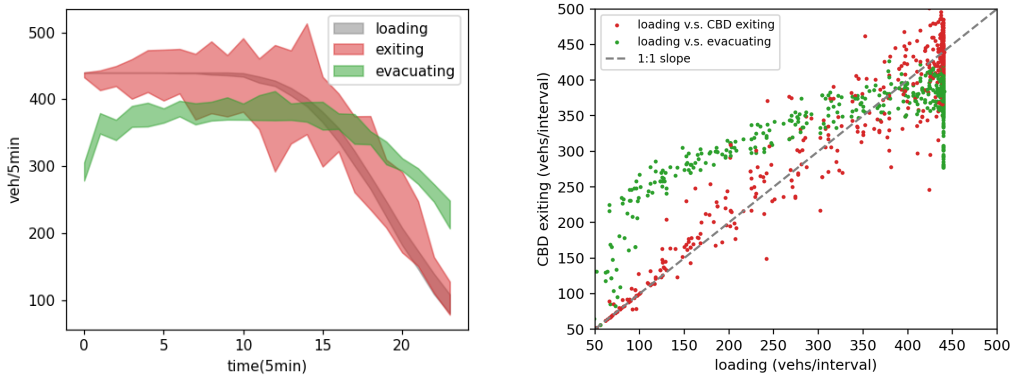
flow is slightly smaller at the beginning but larger in the later. More importantly, we found that all these flows tend to remain stable for a significant long period, as shown in **Figure 3.3**. **Figure 3.3a** shows that in equilibrium, the loading and exiting rates are always approximately equal throughout the entire evacuating process. The equivalence is also shown in **Figure 3.3b**, as the scattering points align with the line $x = y$, suggesting an equilibrium state of the evacuation system.

Therefore, we define the equilibrium between loading and exiting as follows:

$$Q_L^* = N \cdot L^* = N \cdot L(\tau^*) = Q_x^* \approx Q_e^* \tag{3.1}$$

where $N$ is the number of lanes in the CBD network, $L^*$ is the equilibrium loading rate, $Q_L^*$ and $Q_x^*$ are the equilibrium loading and exiting flow respectively.

In **Figure 3.3a**, even at equilibrium, the evacuating flow $Q_e^*$ is not strictly equal to the loading or exit flow, which undergoes a varying process that is initially lower but later higher than the loading curve. This is probably due to the limited space of the outer region of the CBD, as well as the limited duration of our simulation. We believe that with a larger outer region and longer simulation time, all three flows will converge to the same equilibrium value.




(a) histories from 20 rollouts (b) loading/exiting/evacuating scatters

Figure 3.3: Characteristics of the equilibrium loading

In light of the equilibrium between loading and exiting, we classify evacuation loading into three categories, over-loading, equil-loading, and under-loading, which is greater, equal, and smaller than $L^*$, respectively. In addition, we believe that there possibly exists a fourth loading state, namely excessive loading, e.g., $\tau < 60s$, forcing a higher but unstable exiting flow, as shown in **Figure 3.2a**. We will shortly show that in this scenario, the network quickly becomes jammed due to congestion from the inner side, rather than the limited capacity of the congested outside region.

Ideally, we should have the following equality at the equilibrium:

$$\Sigma Q_j^D \approx N_l \cdot 3600/\tau \tag{3.2}$$

$$\tau \approx \frac{3600 N_l}{\Sigma Q_j^D} \tag{3.3}$$

## 3.3 Significance of Traffic Loading for Evacuation Efficiency

Next, we investigate the significance of loading for evacuation efficiency, especially the issues cased by overloading by comparing several evacuation-specific metrics between over-loading and equi-loading. By utilizing simulation tools, the focus will be on the detrimental effects of overloading on evacuation-specific performances that were not shown by conventional macroscopic models or simulations.

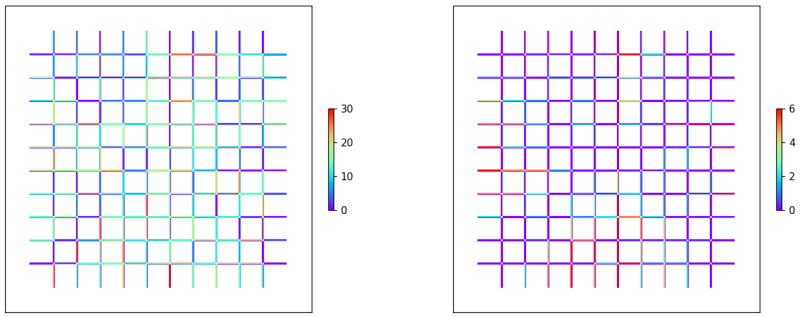
- **Spatial disparity of loading and evacuation** With detailed simulation information, we are also able to track lane-based loading and evacuation, which further allows us to understand where the evacuees are loaded to the network first and from which lane the evacuees reached the shelter first.

  First, we look at the distribution of vehicle loading. As we mentioned earlier, the desired vehicle loading can fail due to congestion or lack of space on the road. Here, we are interested in the distribution of those fail-to-load vehicles. As these data can be difficult to track in the real world, this information is available

in SUMO, known as backlogged vehicles. **Figure 3.4** shows the distribution of the fail-to-load vehicle on the network. Apparently, we can see that the small desired loading headway is infeasible as many vehicles are unable to get on the road. Interestingly, we found that the spatial distribution of fail-to-load vehicles is not even, where the roads closer to the exits face a more severe delay of being loaded.

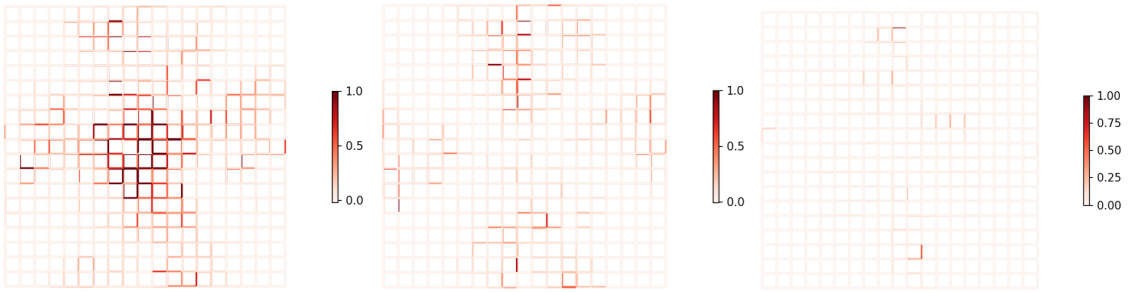
Similarly, we can show the spatial disparity of evacuated vehicles across their origins, or the departure links.

- **Efficiency performance under varying loading rates** The **Figure 3.5** shows the distribution of waiting time (divided by total time) of all links on the network, under different evacuation loading rates. We observed three different patterns. For the case of excessive loading, we can see that the highest waiting time occurs in the central part of the CBD. In the event of overloading, the longest waiting time is in the shelters. In the equilibrium loading case, the waiting time percentage is negligible. Other metrics such as the queue length are found to similar patterns.




(a) $\tau = 2$s (b) $\tau = 30$s

Figure 3.4: Spatial distribution of # of fail-to-load vehicles: vehicles cannot get onto road as at desired times




(a) excessive loading (b) over loading (c) equilibrium loading

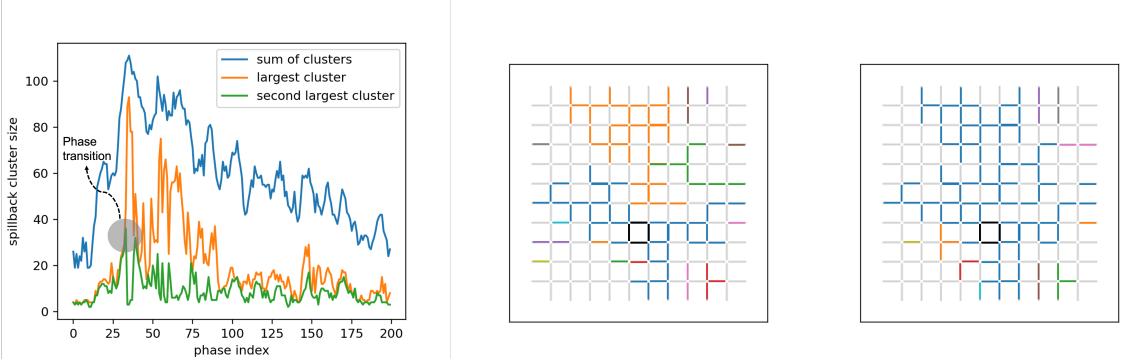
Figure 3.5: Distribution of average waiting/halting time of vehicles on each road of the network: time divided by interval

### 3.3.1 Percolation of spillback jams

Through simulation, we also found that spillover jams propagate extensively in the case of over-loading. Interestingly, the affected links tend to form connected patterns—one or several clusters of spillback roads—which exhibit self-affine characteristics, commonly referred to as fractals in physics.

We now investigate the fractal properties of these spillback jams under over-loading conditions. For example, **Figure 3.6** illustrates the percolation process observed in one simulation rollout with a fixed loading interval of $\tau = 120s/veh/ln$. Phase transition occurs right at the point when the second largest cluster reaches its maximum, as marked in **Figure 3.6a**. After the phase transition, the two largest fractals in **Figure 3.6b** are aggregated into one large jammed cluster in **Figure 3.6c**.

It is also important to note the presence of a deadloop, marked in black, which represents a closed routing path that can sustain the jam indefinitely if drivers do not reroute. In this particular case, the full driver adaptation setting allowed vehicles to reroute dynamically, which helped dissolve the large jammed cluster after the transition.




(a) Percolation process and phase transition (b) Right before phase transition (c) Right after phase transition

Figure 3.6: Percolation of spillback and deadlock formation with over-loading

## 3.4 Implications for Evacuation Traffic Monitoring and control

Congested clusters can be defined not only as spillback, the extreme congestion. We can also extend the percolation observation method for spillback to other congested clusters defined by certain density thresholds. In this section, we will share more percolation-based traffic monitoring and evacuation control measures.

### 3.4.1 Percolation of congestion with varying loading rates

The fractal properties of congestion clusters suggest that to control its growth, localized signal design, or other control measures may not be sufficient, where regional control methods can be more efficient in dealing with clusters of congested links. Before proposing those control methods, we first present the percolation-based observations below, at, and beyond the equilibrium between loading and exiting in evacuation.

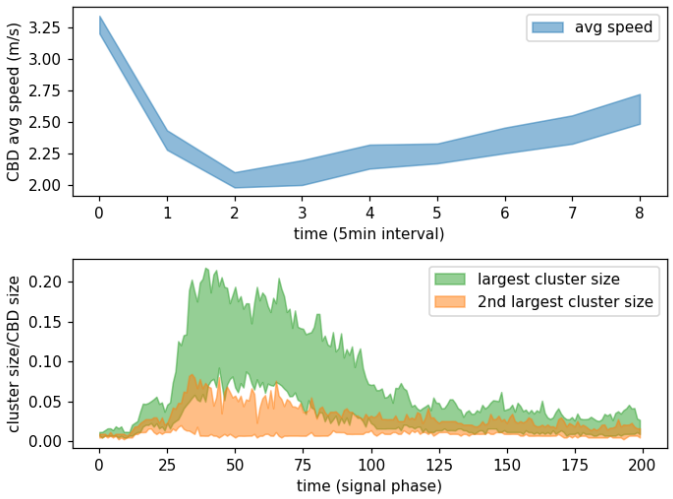



Figure 3.7: Spillback development in the CBD: results from 20 simulation roll outs

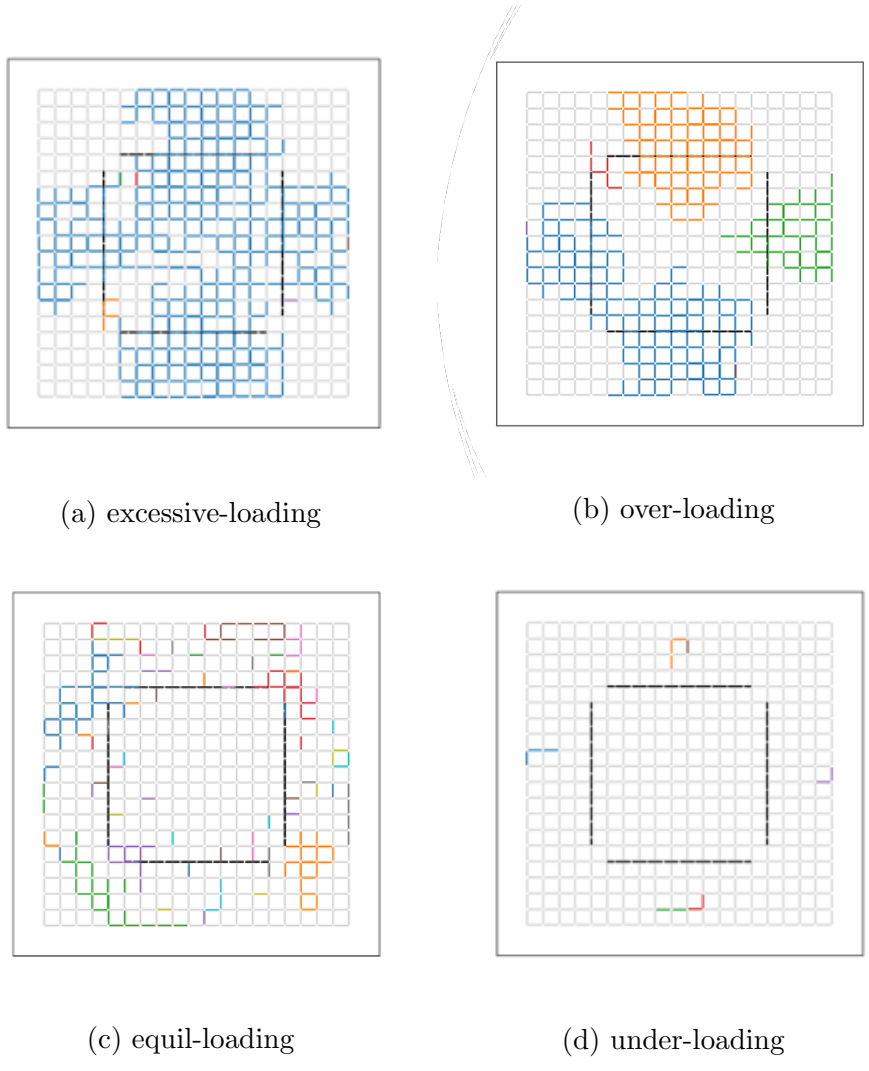
As shown in **Figure 3.8**, we show the percolation status of congested ($k > k_c$) clusters over the network given the four different loading rates mentioned earlier. In the case of excessive loading, **Figure 3.8a** shows that a giant congested cluster percolates almost the entire network. Noticeably, it includes and actually originates from the CBD area, which is significantly different from the other 3 loading cases. In the overloading case, see **Figure 3.8b**, the CBD is no longer congested, but the congestion backs up from the shelters and spreads over the outside region. The congestion further degrades in the outside region in the case of equil-loading, see **Figure 3.8c**. Interestingly, it almost keeps the CBD free from congestion, where no back-up queue encroaches into the boundary (black dashed line). In the case of underloading, **Figure 3.8d** shows that the congestion group becomes further away from the CBD boundary, whose impact tends to be negligible. Another observation is that at equilibrium loading, the number of clusters outside the CBD region reaches the maximum.

We also found that the size $S$ of congested clusters tend to follow a nice power-law distribution as follows:

$$P(S) \approx e^{-\alpha} \tag{3.4}$$

where $\alpha$ is the power-law exponent. Note that the power law exponent $\alpha$ has important implications, as the mean ($\mu$) of a power-law distribution is given by $\mu = \frac{\alpha x_{min}}{\alpha - 1}$, for $\alpha > 1$, because otherwise the mean does not exist.

The variance ($\sigma^2$) of a power-law distribution is given by $\sigma^2 = \left(\frac{x_{min}}{\alpha - 1}\right)^2 \frac{\alpha}{\alpha - 2}$, where the variance does not exist for $\alpha < 2$.



(a) excessive-loading

(b) over-loading

(c) equil-loading

(d) under-loading

Figure 3.8: Percolation of congested clusters under varying loading rates, with congestion defined based on the critical density of the Fundamental Diagram. $k > k_c$

**Figure 3.9** illustrates the empirical distributions of cluster sizes under different loading conditions, where the linear decay in the log-log plots confirms the presence of power-law behavior and highlights how the slope (i.e., exponent $\alpha$) becomes shallower under overloading, indicating reduced network resilience.

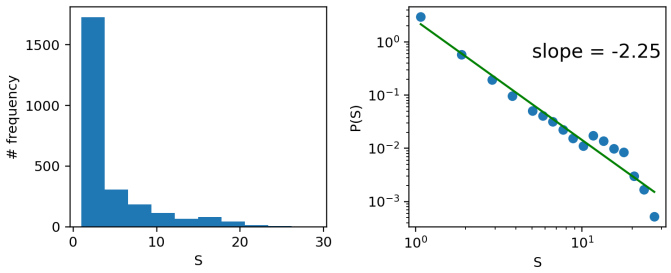
### 3.4.2 Percolation-based evacuation control strategies

Now we further investigate the relationship between the percolation pattern and evacuation efficiency. **Figure 3.10** plots the size of congestion encroachment into the CBD versus the evacuation exit flow under different loading rates. We observe that the maximum exit flow occurs when the congestion encroachment is minimal, highlighting that efficiency is highest when congestion remains outside the CBD.

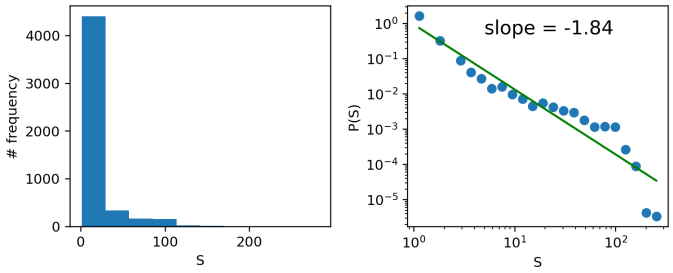
The hysteresis observed in **Figure 3.10b** results from the dynamic nature of loading and unloading phases in the simulation. As traffic builds up and dissipates, the system exhibits delayed responses. This loop further confirms that once congestion breaches the CBD boundary, flow deteriorates and is slow to recover. Therefore, an effective control strategy should aim to prevent congestion from encroaching into the CBD in order to maintain optimal evacuation performance.

In light of these observations, we propose the following percolation-based control rules for evacuation management:

**R-1** The percolation-based rule for managing the loading is to ensure that the external, spreading congestion cluster barely reaches the periphery of the CBD. This approach facilitates an equilibrium state that also delivers the highest exiting rate for evacuating traffic. If traffic engineers have knowledge of the destination of evacuation traffic, they can further refine the loading control, which can be achieved by first observing the congestion pattern in that direction and then determining whether additional traffic toward that direction can be accommodated.

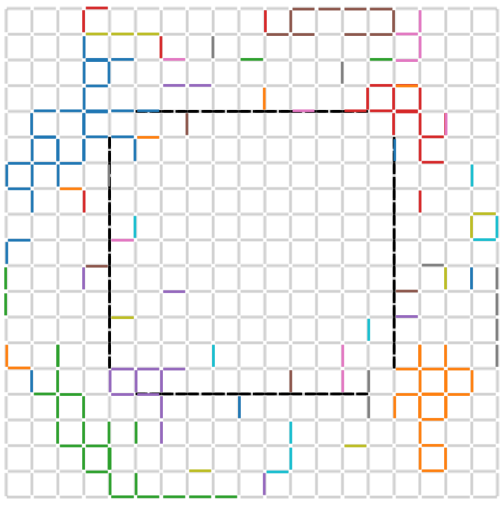



(a) equilibrium loading

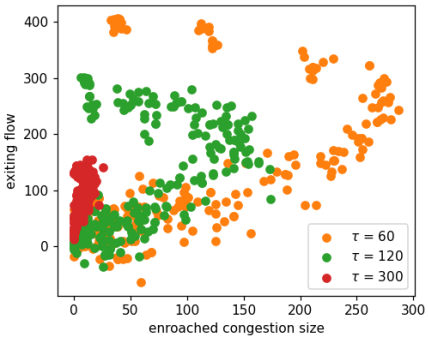



(b) overloading

Figure 3.9: Distribution of congested ($k > k_c$) clusters in the network



(a) maximum CBD exiting requires minimum congestion encroachment from the outside




(b) congestion encroachment v.s. CBD exiting

Figure 3.10: Maximum CBD exiting happens at minimum congestion encroachment at the perimeter

**R-2** The size of the percolating congestion cluster emanating from a particular shelter can serve as an effective indicator for traffic engineers to assess the immediacy of increasing the service capacity of that shelter or tackling incidents that have reduced its capacity. Given that multiple shelters often operate concurrently, these indicators from various shelters can be leveraged for comparative purposes. This comparison can assist in setting priorities and distributing finite resources, as well as guiding the direction of any remaining traffic that has yet to be evacuated.

**R-3** The percolation pattern outside the CBD can provide guidance on where cars can be rerouted and stored before finally being served at the shelter. For example, those uncongested (uncolored) links in **Figure 3.8c** can accommodate more outgoing traffic from the CBD, as shelters are not able to receive them immediately. To see this delay, recall that maximum exit and maximum evacuation flow do not occur at the same time; see **Figure 3.3a**. Therefore, we suggest a better use of the outside region as temporary storage for vehicles that have already left the risk zone but are still waiting for evacuation.

**R-4** This percolation-based observation approach also gives us a dynamic region that allows loading with maximum efficiency. When the outside congestion encroaches into the predefined CBD, the fixed region for traffic loading is no longer suitable, which should be replaced by a dynamic zone surrounded by the outside congestion clusters.

**R-5** Monitoring the percolation of congested cluster in the outside region is more meaningful than only monitoring the boundary exits of the CBD. It is because the outside region can still flow well even when some small congested clusters happen to be close to the CBD outside boundary, which should not concern us since those small congested clusters will resolve soon more downstream traffic reaches its destination.

Now we provide some intuition for the percolation-based control strategies **R-1**. Intuition can be found from the analogy in the 1d evacuation task, e.g., a highway, where the optimal loading rate is very straightforward. As the cell transmission model (Daganzo, 1994) suggests, the upstream demand (loading) should not exceed the downstream supply; see the supply-demand curve in **Figure 3.11**.

To achieve the maximum evacuating rate, downstream congestion must be cleared. In the worst case, the spillback jam, or $k = k_j$, percolates through from the shelter all the way back to the CBD boundary, causing the most catastrophic effects since it could take indefinite long time for the spillback percolating jam to resolve, such that the CBD traffic will also needs indefinite time get out of the danger zone.

## 3.5 Implications for MFD-based Traffic Control

This paper has investigated the significance of loading rates for evacuation. It identifies the equivalence between loading and exiting, and finds that over-loading causes detrimental impacts on evacuation efficiency, where devastating spillback jams can happen and deteriorate the network quickly. The paper has also proposed a few percolation-based traffic monitoring and control strategies for evacuation management. Now, we share some concluding remarks and discuss some potential avenues for future research.

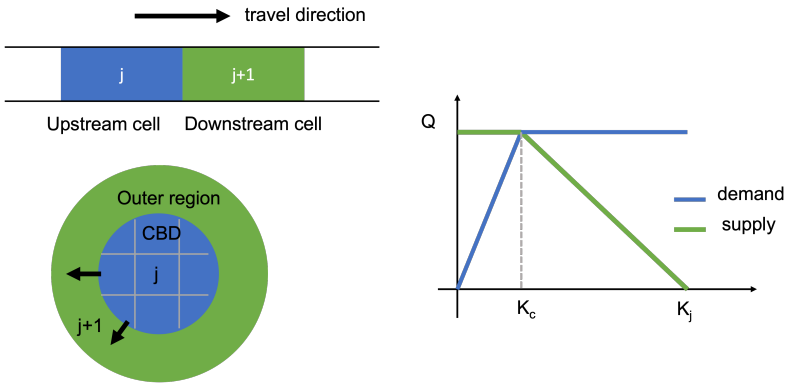
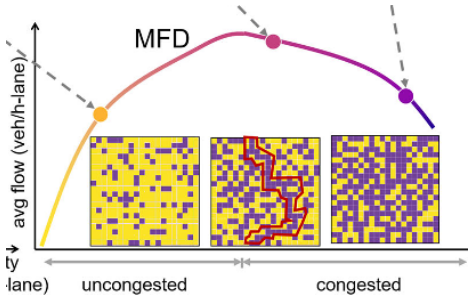


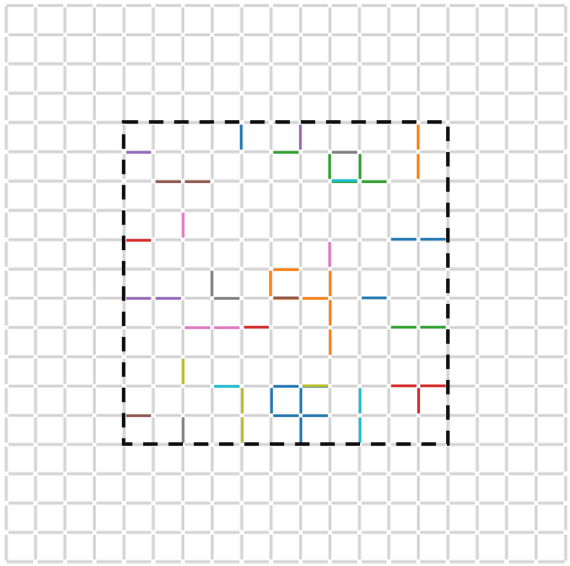

Figure 3.11: Intuition behind the percolation-based control: from 1d to 2d evacuating control

### 3.5.1 Connections with MFD control

One of the main contributions of this paper is the percolation-based control strategy for managing evacuation loading rates, which is to ensure that the external spreading congestion cluster barely encroaches into the periphery of the CBD. To the best of our knowledge, this is the first percolation-based control strategy to move traffic out of the CBD. Interestingly, the dual of this problem, where vehicles head towards the CBD, has been studied more extensively by the field, which is known as MFD control Daganzo and Geroliminis (2008); Geroliminis et al. (2012). A recent study Ambühl et al. (2023) has revealed that MFD hits the highest flow when the number of congested ($k > k_c$) clusters reaches the maximum, or the network percolates. This insight has provided a new pathway for the application of percolation theory in managing incoming arrival traffic. By reversing the OD in this paper, we could easily simulate the incoming arrival traffic, and implement the percolation-based MFD control as shown in **Figure 3.12b**.




(a) MFD flow is highest right before percolation: sub-figure source Ambühl et al. (2023)



(b) CBD has highest MFD flow with max # of cong. clusters

Figure 3.12: The dual of the percolation-based evacuation control: the percolation-based MFD control, where the underlying rule is that MFD flow is highest with maximum number of congested clusters

To our surprise, we find that the widely anticipated signal control policy, the LQF rule, may not be a good option in evacuation scenarios due to its deterministic nature, which may delay the dissolve of large spillback clusters. In contrast, traditional signal control, where the green lights change in turns, provides more incentives for driver adaptation.

### 3.5.2 Implementation for city networks

The method proposed in this paper can easily be transferred from abstract network settings to more realistic urban networks. To see this, we applied our percolation analysis method to the Nashville network, TN, as shown in **Figure 3.13**, where the dashed circle is a CBD zone that needs to be evacuated, and the four highway entrances serve as destinations. The proposed method can be easily transferred to other city networks. As traffic managers, when those congestion clusters infiltrate the CBD, such as Shelters 1 and 2 in **Figure 3.13**, the control should be taken to increase the service capacity at 1 and 2, and reroute evacuees to other exits, to improve exiting and evacuating efficiency. Similarly, one can build a spillback cluster to monitor the heavy congestion in the network and concentrate limited resources to resolve those most devastating traffic jams that delay the evacuation.

It should be noted that the primary focus of this paper is on a symmetric network with symmetric OD pairs, implying that no individual links hold a higher priority than others. However, real-world scenarios often present bottleneck links where certain roads are prone to congestion. These links can be analogized to bond activation in percolation theory, leading to the aggregation of smaller clusters into larger ones, as depicted in **Figure 3.6**. Identifying the location of such bottleneck links and preventing their bond activation pose intriguing questions worthy of exploration in future research endeavors.

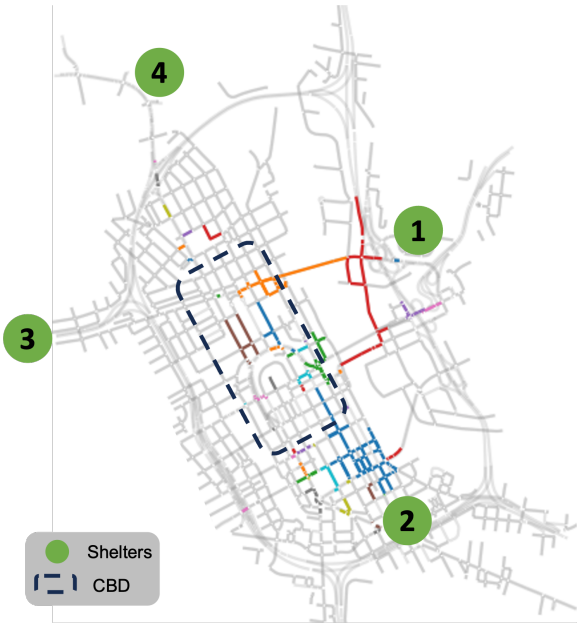



Figure 3.13: Implementation of percolation-based evacuation control in city networks: A case study of Nashville, TN ($k > k_c$).

# Chapter 4

# Potential of electric autonomous vehicles for improving traffic efficiency under extreme events

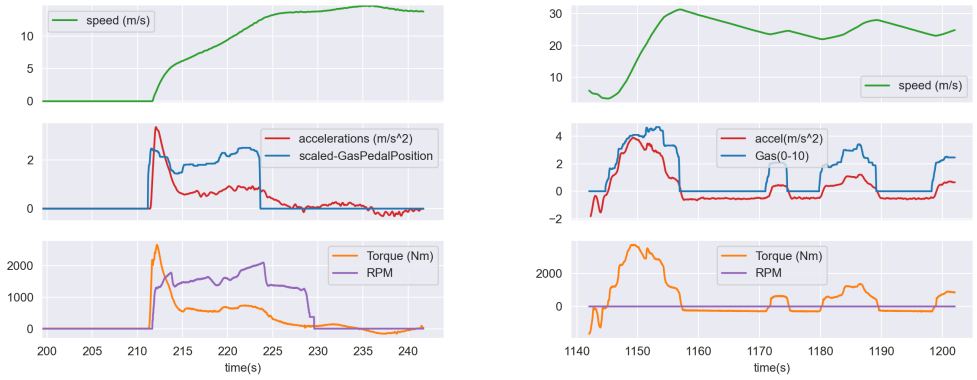
This chapter presents an empirical study to characterize electric vehicles' (EVs) physical performance compared with internal combustion engine (ICE) vehicles and discuss their implications for traffic flow. Through field experiments with rental EVs and CAN-logger, we find that: (i) EVs exhibit more powerful accelerations than ICE vehicles, characterized by a plateau-like acceleration pattern, in contrast to the well-known linearly decreasing accelerations of ICE vehicles. (ii) Regenerative braking in EVs provides quick and efficient deceleration, allowing for tighter following distances with lead vehicles. (iii) EVs demonstrate better responsiveness to human drivers' pedal commands compared to ICE vehicles, resulting in improved predictability and controllability. (iv) Despite their superior acceleration and deceleration capabilities, our EV platoon experiment indicates that current ACC systems do not fully utilize EVs' potential, leading to a waste of capacity. The empirical evidence and insights presented in this study are expected to draw more attention to EVs' capabilities and the integration of automation to fully leverage their potential to improve traffic flow.

## 4.1 Background

### 4.1.1 Torque, power and accelerations

Torque ($T$) is a measure of the rotational force applied to an object, defined as the product of the force and the distance from the pivot point, and is measured in Newton-meters (Nm). Power ($P$) is the rate at which work is done or energy is transferred, and in rotational systems, it is the product of torque and angular speed ($\omega$), expressed as $P = T \cdot \omega$.

For vehicles, power determines how quickly a vehicle can perform work to accelerate. The relationship between power and vehicle acceleration ($a$) can be described by Newton's Second Law ($F = m \cdot a$) and the equation $P = F \cdot v$, where $F$ is force, $m$ is mass, and $v$ is velocity. By combining these equations, acceleration can be expressed as $a = P/(m \cdot v)$, indicating that for a given power, acceleration is inversely proportional to both vehicle speed and mass. Thus, higher power results in greater acceleration, particularly at lower speeds, as shown in **Figure 4.1a** and **Figure 4.1b**.




(a) ICE's power decay example (b) EV's sustained torque example

Figure 4.1: Torque response and accelerations in EV and ICE vehicles

EVs exhibit unique acceleration characteristics due to the properties of electric motors. These characteristics are typically divided into two distinct regimes: a **torque-limited regime** and a **power-limited regime**. Understanding these regimes provides insight into the performance of EVs and their ability to accelerate under various conditions.

### 4.1.2 EV Two-Regime Accelerating Model

The acceleration of an EV can be described using a two-regime physics model:

1. **Torque-Limited Regime:** At low speeds, the motor operates in a torque-limited regime where the maximum torque ($\tau_{\max}$) is constant. Acceleration in this regime is determined by the relationship:

$$F = m \cdot a = \frac{\tau_{\max}}{r},$$

   where:

   - $F$: Force applied to the wheels (in newtons),
   - $m$: Mass of the vehicle (in kilograms),
   - $a$: Acceleration (in $\mathrm{m/s^2}$),
   - $\tau_{\max}$: Maximum motor torque (in Nm),
   - $r$: Effective wheel radius (in meters).

   From this equation, acceleration in the torque-limited regime is constant and independent of speed.

2. **Power-Limited Regime:** At higher speeds, the motor enters a power-limited regime where the total power output ($P_{\max}$) is capped. In this regime, the available torque decreases as speed increases. The relationship is governed by:

$$P_{\max} = F \cdot v = m \cdot a \cdot v.$$

Rearranging for acceleration ($a$):

$$a = \frac{P_{\max}}{m \cdot v}.$$

In the power-limited regime, acceleration decreases inversely with speed.

The *corner speed* ($v_c$) marks the transition point between the torque-limited and power-limited regimes. It is the speed at which the motor's maximum torque and maximum power constraints are equal. Mathematically, this can be derived by equating the torque-limited power output to the motor's maximum power:

$$P_{\max} = \tau_{\max} \cdot \omega_c,$$

where $\omega_c$ is the rotational speed at $v_c$ (in radians per second). The relationship between vehicle speed ($v_c$) and rotational speed is given by:

$$\omega_c = \frac{v_c}{r}.$$

Substituting $\omega_c$ into the power equation:

$$P_{\max} = \tau_{\max} \cdot \frac{v_c}{r}.$$

Rearranging for $v_c$:

$$v_c = \frac{P_{\max} \cdot r}{\tau_{\max}}.$$

The corner speed depends on the following key parameters:

- $P_{\max}$: Maximum motor power (in watts),
- $\tau_{\max}$: Maximum motor torque (in Nm),
- $r$: Effective wheel radius (in meters).

To summarize EVs' accelerating capabilities,

- In the **torque-limited regime**, acceleration is nearly constant and depends on the maximum torque and wheel radius.
- In the **power-limited regime**, acceleration decreases inversely with speed, as the motor operates at its maximum power output.
- The **corner speed** ($v_c$) serves as the dividing line between these two regimes and can be calculated using motor specifications.

### 4.1.3 Limitations of ICE Engines and Advantages of EVs in Sustaining Power

One of the main limitations of ICE vehicles is their inability to sustain high power output beyond a certain RPM. This limitation arises from several factors, such as reduced volumetric efficiency, increased frictional losses, and challenges in maintaining optimal fuel-air mixture quality at higher engine speeds. As a result, after reaching a peak RPM, the engine's power output plateaus or decreases. Consequently, the torque must decrease to prevent exceeding the engine's power capacity, leading to a drop in acceleration. This behavior is evident in the empirical data, as shown in the third plot in **Figure 4.1a**, where ICE vehicles show an initial surge in torque and acceleration, followed by a rapid decline as the RPM increases.

In contrast, EVs can sustain their power output over a broader range of speeds; see the example in **Figure 4.1b**. Electric motors deliver consistent torque across a wide range of RPMs, enabling EVs to maintain high acceleration without the need for power reduction. This characteristic allows EVs to provide immediate and sustained acceleration, making them superior in performance compared to ICE vehicles, especially in scenarios requiring prolonged high acceleration. The ability of EVs to sustain power output and torque without significant drops underscores their advantage in efficiency and performance.

## 4.2 Trapezoidal Modeling of Accelerator Pedal Behaviors

Accurate modeling of accelerator pedal behaviors during electric vehicle (EV) acceleration is essential for realistic vehicle dynamics simulation and driver behavior analysis. Typically, pedal position profiles during acceleration exhibit characteristic trapezoidal shapes, transitioning through clearly defined ramp-up, plateau (hold), and ramp-down stages. However, traditional trapezoidal models using piecewise-linear functions often lack smoothness and differentiability, complicating computational analyses.

### 4.2.1 Pedal profile modeling

To address these limitations, a sigmoid-based trapezoidal function is proposed. This model employs the difference of two sigmoid functions to smoothly represent pedal positions:

$$\text{Pedal}(v) = \sigma(k_1, v_1; v) - \sigma(k_2, v_2; v) \tag{4.1}$$

where the sigmoid function $\sigma$ is defined as:

$$\sigma(k, x_0; v) = \frac{1}{1 + e^{-k(v - x_0)}} \tag{4.2}$$

Here, the parameter definitions are as follows:

(i) $v$: Vehicle velocity (m/s).

(ii) $k_1, k_2$: Steepness parameters controlling the smoothness of the pedal's ramp-up and ramp-down transitions.

(iii) $v_1, v_2$: Velocity midpoints at which the pedal's engagement (rise) and disengagement (fall) transitions occur.

This parameterization offers significant flexibility, enabling the modeling of a wide range of realistic driving behaviors through systematic adjustments:

(i) **Aggressive Driving:** Higher steepness parameters ($k_1, k_2 \gg 1$) yield sharp, clearly defined trapezoidal pedal behaviors, with extended plateau regions representing sustained full-throttle scenarios.

(ii) **Gentle Driving:** Lower steepness parameters ($k_1, k_2 < 1$) result in smoother transitions and shorter plateau phases, resembling bell-shaped or parabolic profiles.

(iii) **Balanced Driving:** Intermediate parameter values capture typical, balanced driving conditions observed in common urban acceleration scenarios.

Note that Eqn.(4.1) is also able to capture unsymmetrical pedaling profiles. For vehicles accelerating from $v_0$ to its desired speed $v_c$, an unsymmetrical pedaling behaviors $Pd(v)$ can be observed if $v_0 - v_1$ is not symmetric with $v_2 - v_c$, which is quite common especially when the driver tapers off the pedal to gradually accelerate to its desired speed.

This approach offers notable advantages, including differentiability for ease of computational implementation, intuitive parameter interpretation, and adaptability to driver-specific or context-specific pedal behavior profiles. Consequently, the sigmoid-based trapezoidal modeling technique provides a robust foundation for enhanced vehicle dynamics simulation and comprehensive driver behavior studies within EV acceleration contexts.

To illustrate the practical applicability and flexibility of the sigmoid-based trapezoidal model, **Figure 4.2** demonstrates three different pedal behavior profiles derived by selecting specific parameter sets:

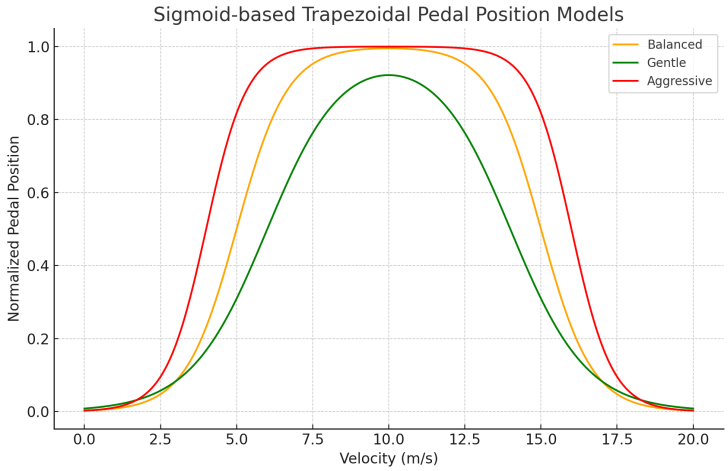



Figure 4.2: Sigmoid-based trapezoidal pedal models capturing various driving styles.

The specific parameter values corresponding to each driving style are summarized in **Table 4.1** below.

As shown in **Figure 4.2**, varying the parameters $k_1, k_2$, and midpoints $v_1, v_2$ significantly alters the pedal profile shape. Specifically:

(1) **Balanced model** displays a moderately trapezoidal configuration that effectively captures the essence of conventional urban driving patterns. This configuration is characterized by its smooth transitions between phases and a prolonged period at the plateau stage. Within this model's balanced scheme, the peak pedal position is frequently set at a level that prioritizes driver comfort, avoiding the utilization of the vehicle's engine or motor to its full capacity.

(2) **Gentle model** is distinguished by its diminished steepness, which culminates in a profile akin to a parabolic or bell-shaped curve. This design underlines the importance of gradual pedal movements and is closely aligned with environmentally conscious driving practices.

Table 4.1: Sigmoid Model Parameters for Different Driving Styles

| **Model** | $k_1$ | $v_1$ | $k_2$ | $v_2$ |
|---|---|---|---|---|
| Balanced (Typical) Driving | 1.2 | 5 | 1.2 | 15 |
| Eco-Friendly Driving | 0.8 | 6 | 0.8 | 14 |
| Aggressive Driving | 1.5 | 4 | 1.5 | 16 |

(3) **Aggressive model** incorporates an increased level of steepness and elongates the plateau regions. This effectively captures the dynamic and vigorous acceleration patterns typically associated with sport or performance driving modes. Such a mode often maintains the pedal position at a persistently high level, although not necessarily nearing or reaching the maximum peak value of one.

Through such flexible parameterization, the sigmoid-based trapezoidal model efficiently represents diverse realistic acceleration scenarios, offering a versatile approach for integrating driver-specific behavior into EV acceleration analyses and simulations.

### 4.2.2 Average Pedal Position During Acceleration

To model the average pedal position during acceleration from an initial speed $v_0$ to a final speed $v_c$, we analyze the pedal position function:

$$\text{Pedal}(v) = P_{\max} \left( \frac{1}{1 + e^{-k_1(v - v_1)}} - \frac{1}{1 + e^{-k_2(v - v_2)}} \right),$$

where $P_{\max}$ is the maximum pedal position, $v_1$ and $v_2$ are the speed thresholds for pedal engagement and disengagement, and $k_1$ and $k_2$ control the steepness of the transitions. The average pedal position over the interval $[v_0, v_c]$ is given by:

$$\overline{\text{Pedal}} = \frac{P_{\max}}{v_c - v_0} \int_{v_0}^{v_c} \left( \frac{1}{1 + e^{-k_1(v - v_1)}} - \frac{1}{1 + e^{-k_2(v - v_2)}} \right) dv.$$

This integral can be evaluated exactly, but its complexity depends on the values of $k_1$ and $k_2$. Below, we analyze two cases: (1) sharp transitions ($k_1, k_2 \gg 1$) and (2) moderate transitions ($k_1, k_2$ finite).

**Case 1: Aggressive Driving with Sharp Transitions ($k_1, k_2 \gg 1$)**

For sharp transitions, the sigmoid functions approximate step functions, and the pedal position behaves like a trapezoidal pulse. Specifically: - The pedal is fully applied ($P_{\max}$) between $v_1$ and $v_2$. - The pedal is zero outside this range.

The average pedal position simplifies to:

$$\overline{\text{Pedal}} \approx P_{\max} \cdot \frac{v_2 - v_1}{v_c - v_0}.$$

This result assumes $v_0 \leq v_1 \leq v_2 \leq v_c$, and the transitions at $v_1$ and $v_2$ are sufficiently sharp ($k_1(v_1 - v_0) \gg 1$ and $k_2(v_c - v_2) \gg 1$).

**Case 2: Balanced Driving with Moderate Transitions ($k_1, k_2$ Finite)**

For moderate steepness parameters, the transitions are smooth, and the pedal position exhibits gradual ramping. The average pedal position can be approximated by dividing the integral into three stages corresponding with the trapezoid: 1. **Ramp-up** ($v_0$ to $v_1 + \Delta v_1/2$): The pedal position transitions from 0 to $P_{\max}$. 2. **Plateau** ($v_1 + \Delta v_1/2$ to $v_2 - \Delta v_2/2$): The pedal position is approximately $P_{\max}$. 3. **Ramp-off** ($v_2 - \Delta v_2/2$ to $v_c$): The pedal position transitions from $P_{\max}$ to 0.

Here, $\Delta v_1 = \frac{4}{k_1}$ and $\Delta v_2 = \frac{4}{k_2}$ represent the width of the rising and falling transitions, respectively. The average pedal position is:

$$\overline{\text{Pedal}} \approx P_{\max} \cdot \frac{v_2 - v_1 + \frac{2}{k_1} + \frac{2}{k_2}}{v_c - v_0}.$$

The two cases illustrate the impact of transition steepness on the average pedal position i) for sharp transitions, the average pedal position depends only on the

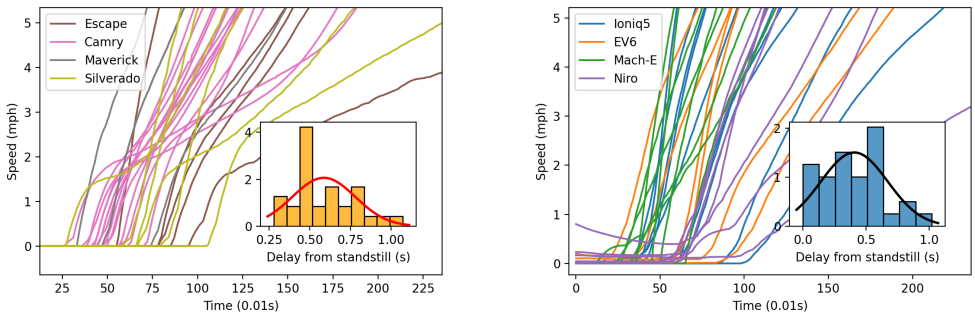
plateau width ($v_2 - v_1$). and ii) for moderate transitions, the rising and falling edges contribute additional terms proportional to $\frac{1}{k_1}$ and $\frac{1}{k_2}$, respectively.

## 4.3 Empirical characteristics of EV accel/decelerations

In this section, we will present empirical data from our EV experiments on their acceleration/deceleration capabilities. We will first investigate the response delay, especially from a standstill state, and then present the acceleration and deceleration evidence from market EVs.

### 4.3.1 Delay off standstill

**Figure 4.3** compares the delay in acceleration responses to gas pedal inputs from a standstill state for both EVs (right) and Non-EVs (left). Several key observations can be drawn from the graphs: The delay in acceleration response at standstill is highly dependent on the specific vehicle rather than the type (EV or Non-EV). Within each category, there are significant differences. For example, the Mach-E and Camry exhibit much smaller delays compared to the Silverado and Ioniq 5, which have longer delays. When viewed from a broader category perspective, the delay times at the standstill state are similar between EVs and Non-EVs. This indicates that, on average, the acceleration delay from a standstill is not solely dictated by whether a vehicle is an EV or a Non-EV. As shown by the histogram, while there are notable differences between specific vehicles, the overall delay in acceleration from a standstill is similar when comparing EVs to Non-EVs. Although EVs are expected to have more instant torque than ICE vehicles, our empirical data suggest that EVs do not have a significant advantage in response delay from standstill, which might contradict common expectations.



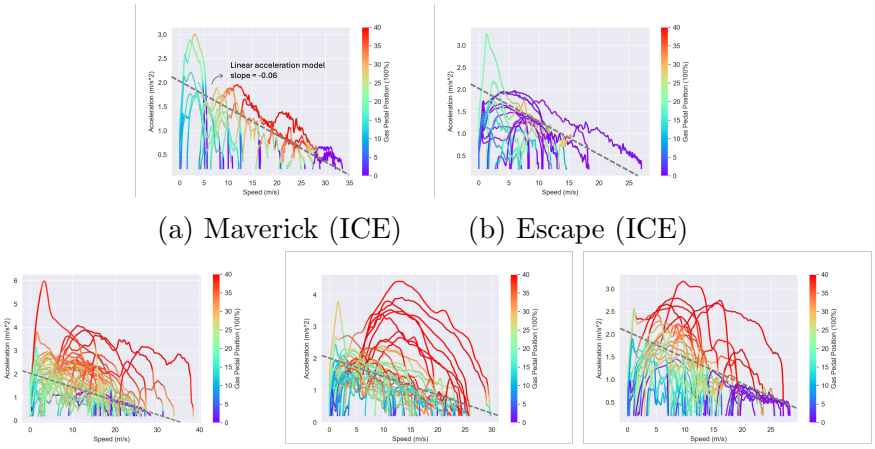

(a) ICE vehicles (b) EV vehicles

Figure 4.3: Delay of EV and Non-EV acceleration responses from the standstill state

### 4.3.2 Acceleration power

To show EVs' acceleration capabilities, we have conducted a good amount of human-driven experiments covering different speed ranges and various traffic conditions. Notably, to see the impact of EVs' accelerations on traffic flow, we have done a few 'catching-up' experiments, where the leader vehicle drives away or changes lane and the following EV needs to fill the gap. In comparison, we also tested a few traditional ICE vehicles. The comparison results are shown in **Figure 4.4**. Note that the dashed line in these subplots is the linear acceleration model in the traffic literature Laval et al. (2014); Zhou et al. (2022a), which is also found to be the main contributing factor to capacity loss.

In **Figure 4.4**, our data from ICE vehicles, such as the Maverick and Escape, align well with the existing linear acceleration theory. This theory, represented by the dashed line in the subplots, suggests that acceleration decreases linearly with increasing speed. This pattern is evident in the acceleration profiles of the ICE vehicles, where the curves follow a linear decaying trend. However, our findings indicate that EVs do not adhere to this linear acceleration model.




(a) Maverick (ICE) (b) Escape (ICE)

(c) Mach-E (EV) (d) Ioniq5 (EV) (e) EV6 (EV)

Figure 4.4: Extended high accelerations in EVs vs. power decline in ICE vehicles. The dashed line is the linear acceleration model for ICE vehicles.

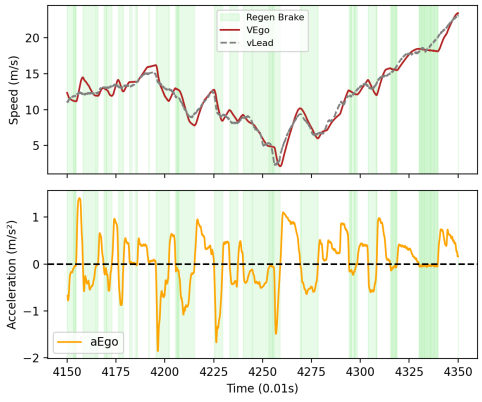
Instead, EVs exhibit extended high accelerations across a broader range of speeds, from low to very high. The acceleration curves for EVs, such as the Mach-E, Ioniq 5, and EV6, display a parabolic shape, maintaining higher acceleration levels even at higher speeds. This divergence from the linear model highlights the superior acceleration capabilities of EVs, which can sustain higher performance across various speed ranges, unlike the gradually declining acceleration observed in ICE vehicles. This characteristic of EVs could have significant implications for traffic dynamics, as their ability to maintain higher accelerations may contribute to more efficient and responsive traffic flow.

The acceleration curves for EVs, such as the Mach-E, Ioniq 5, and EV6, do not show the linearly decaying trend observed in ICE vehicles. Instead, they display a plateau-like shape, maintaining higher acceleration levels over a broad range of speeds. This suggests that EVs can sustain consistent and robust acceleration even at higher speeds, which is a marked contrast to the gradual decline seen in ICE vehicles. This characteristic of EVs could have significant implications for traffic dynamics, as the linearly-decreasing acceleration in ICE vehicles is a major contributing factor to capacity loss. With their ability to maintain high and stable acceleration, EVs have

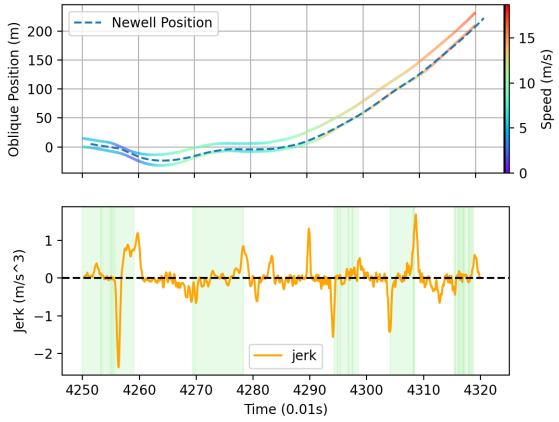
the potential to fill in voids in traffic more effectively and hence improve overall traffic capacity. Additionally, we present data on gas pedal positions to justify that these superior EV performances are not merely a result of aggressive driving but rather their inherent physical capabilities. The gas pedal data shows that high acceleration values are achieved without extreme inputs, underscoring the efficiency and power of EVs.

### 4.3.3 Regenerative braking

Regenerative braking is a technology used in EVs that allows the vehicle to recover and store energy during deceleration. When the driver lifts their foot off the gas pedal, the electric motor switches to generator mode, converting kinetic energy into electrical energy and storing it in the battery. One-pedal driving leverages regenerative braking to enable the driver to control the vehicle's acceleration and deceleration using just the gas pedal. By modulating the pressure on the pedal, the driver can smoothly transition between speeding up and slowing down without needing to use the brake pedal frequently.




(a) one-pedal driving




(b) following in stop-and-go traffic

Figure 4.5: Regenerative braking provides responsive deceleration. Data is collected using Kia Niro's One-pedal driving following an oscillatory lead vehicle

**Figure 4.5a** illustrates the effect of regenerative braking in an EV, showcasing the Kia Niro's one-pedal driving mode in oscillatory traffic conditions. The top plot shows the vehicle speed (vEgo) in red, overlaid with the speed of the leading vehicle (vLead) in gray, while the green shaded areas indicate periods of regenerative braking. The bottom plot displays the vehicle's acceleration (aEgo) in orange. Notice the frequent spikes in the speed profile during the green shaded areas, which correspond to rapid decelerations when the gas pedal is lifted. These spikes demonstrate how Regen provides quick and efficient deceleration, allowing the vehicle to slow down responsively and recover energy simultaneously. This capability enhances driving efficiency and improves overall control, making EVs particularly well-suited for stop-and-go traffic scenarios.

The quick and strong decelerations provided by regenerative braking not only enhance driving efficiency, but also has the potential to allows EVs to maintain shorter time headways safely, which can increase traffic capacity. As shown in **Figure 4.5b**, one-pedal driving helps catch up quickly, leaving very small gap compared to the Newell trajectory. By reducing the spacing between vehicles, EVs can effectively utilize road space more efficiently, leading to smoother traffic flow and potentially reducing congestion. This advantage highlights the potential of EVs to contribute to better traffic management and higher road capacity, especially in stop-and-go traffic scenarios where responsive deceleration is critical.

### 4.3.4 Responsiveness of EVs

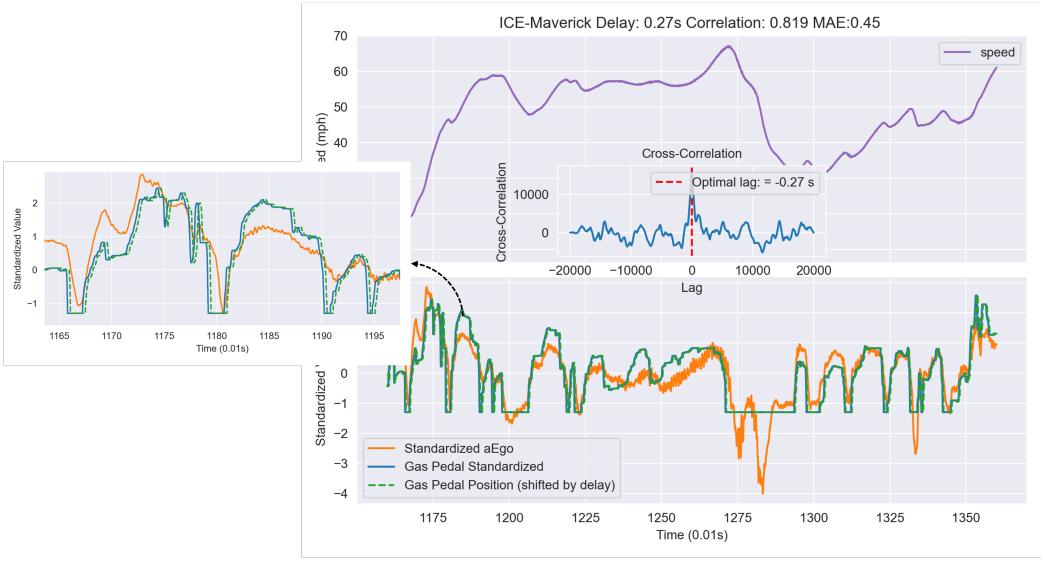
Thanks to the superior accelerating and decelerating capabilities, we also found that EVs tend to be more responsive, in both human-driven and AV modes.

#### Human-driven mode

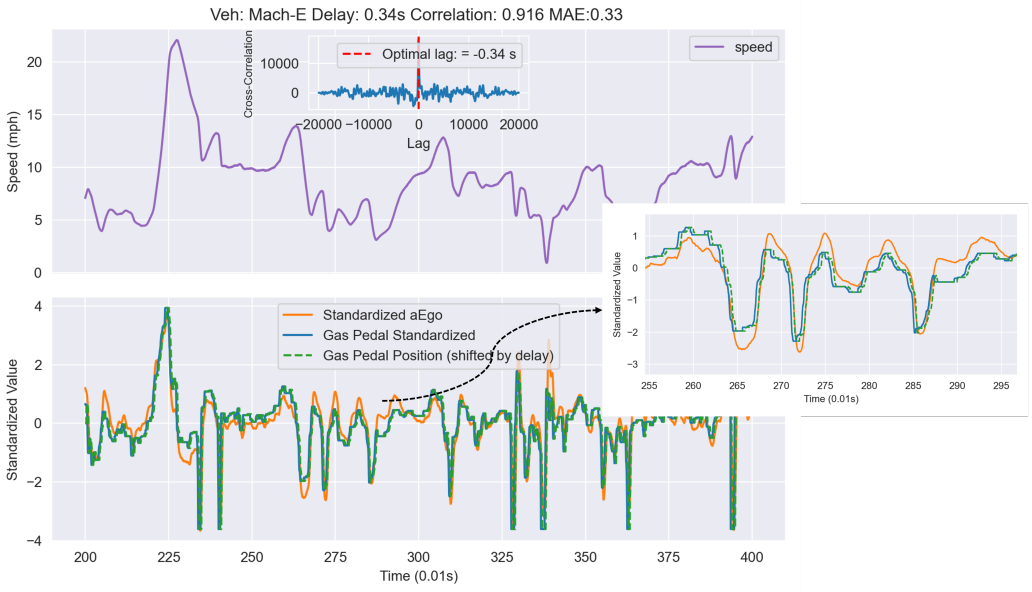
To see the impact of EVs on human driving, we mainly focus on the predictability of the vehicle, i.e., whether EVs respond to human drivers' pedal pressing in a

more anticipated way. Thanks to access to vehicle CAN messages, we were able to investigate the correlation between gas pedal commands and the acceleration response. The results are summarized in **Figure 4.6**, including one example from ICE vehicles and one for EVs. We used the cross-correlation method to estimate the delay in acceleration response to pedal commands. Then, we added the shifted gas pedal history to see if there exists a good correlation between the pedal position and the vehicle acceleration. As shown in **Figure 4.6b**, the good correlation between gas pedal input and EV acceleration responses suggests that EVs have simpler and more linear dynamics, making them easier to predict for human drivers and control for self-driving systems. This linearity ensures that changes in gas pedal input result in proportional changes in acceleration, leading to a more intuitive and responsive driving experience. In contrast; see **Figure 4.6a**; ICE vehicles often exhibit frequent overshooting and undershooting in their acceleration responses, indicating more complex and non-linear dynamics. This variability can cause difficulties for human drivers in making precise speed adjustments and poses challenges for autonomous systems in maintaining consistent and safe control, as the control algorithms must account for the unpredictable behavior of the vehicle.

The correlation between gas pedal input and acceleration after applying the optimal delay shift provides significant insights into vehicle controllability and predictability. For human drivers, a high correlation in EVs suggests a more intuitive and responsive driving experience. The immediate and predictable acceleration response allows drivers to make precise adjustments to speed, which enhances overall vehicle control and safety. This responsiveness can be particularly beneficial in dynamic driving conditions, such as urban environments or during emergency maneuvers, where timely and accurate vehicle reactions are critical.




(a) ICE response delay to pedal input




(b) EV response delay to pedal input

Figure 4.6: ICE/EV response delay to human driver's gas pedal commands

### AV mode

For self-driving, the pedal commands are canceled, instead the vehicle receives and tries to execute a commanded/requested/desired acceleration from the ACC system, which is referred to as aReq in this paper. Similarly, we are looking at how EVs react

to aReq and if its true acceleration response (aEgo) has a good tracking/correlation with the commanded values.

Apparently, a higher correlation observed should imply a more reliable and predictable control environment. Autonomous systems rely heavily on the predictability of vehicle responses to ensure smooth and safe operation. High correlation facilitates the development of accurate control algorithms and models, improving the system's ability to anticipate and react to various driving scenarios. This predictability reduces the complexity of the control logic and enhances the overall robustness of the autonomous system. In contrast, the lower correlation would indicate more variability and less predictable behavior, posing challenges for both human drivers and self-driving systems in maintaining consistent and safe control.

We found that, although ICE vehicles have more constrained accelerations and can be more difficult to control for human drivers, they still display good tracking capabilities, meaning that those ICE vehicles are still able to closely track their desired accelerations in their self-driving/ACC systems. By using the reversed vehicle CAN messages, we are able to obtain the desired acceleration (aReq) and compare it with true acceleration output (aEgo).

To ensure good tracking while having limited acceleration capabilities, ICE vehicles often plan very conservative accelerations. As illustrated in **Figure 4.7**, the acceleration values in ACC systems for a 2024 Camry and Rav4 are generally modest. The Rav4's acceleration values typically remain below 0.5 $m/s^2$, and its jerk values are usually under 0.25 $m/s^3$. This conservative approach helps maintain smooth driving dynamics and prevents abrupt changes that could compromise passenger comfort and vehicle stability. Although EVs possess high acceleration capabilities, not all utilize this advantage in self-driving mode. Their acceleration and jerk values often resemble those of internal combustion engine (ICE) vehicles, as shown in **Figure 4.7** and **Figure 4.8**. The acceleration values of these EVs rarely exceed 1.0 $m/s^2$, and their jerk values are typically less than 0.5 $m/s^3$. However, the Ioniq 5 stands out with more aggressive acceleration behavior; its ACC system includes a unique setting

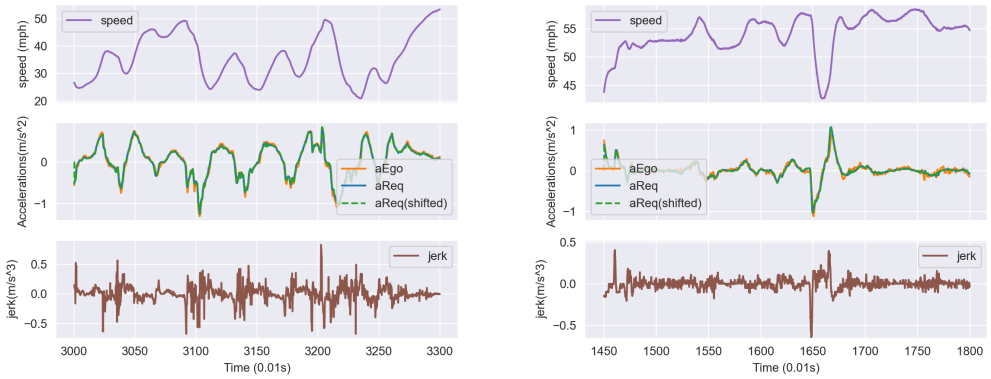
that allows users to toggle higher speeds and acceleration rates, resulting in higher acceleration and jerk values. Unfortunately, such features are not common across the EV market, leaving the great potential of EVs' acceleration and jerk capabilities largely untapped.

## 4.4 Impact of EVs' Agility on Lane-Changing

Lane changing is a critical component of highway driving, directly influencing traffic flow efficiency, safety outcomes, and the overall driving experience. The decision-making process behind lane changes has been the subject of extensive research, with the MOBIL model emerging as one of the most widely accepted frameworks.

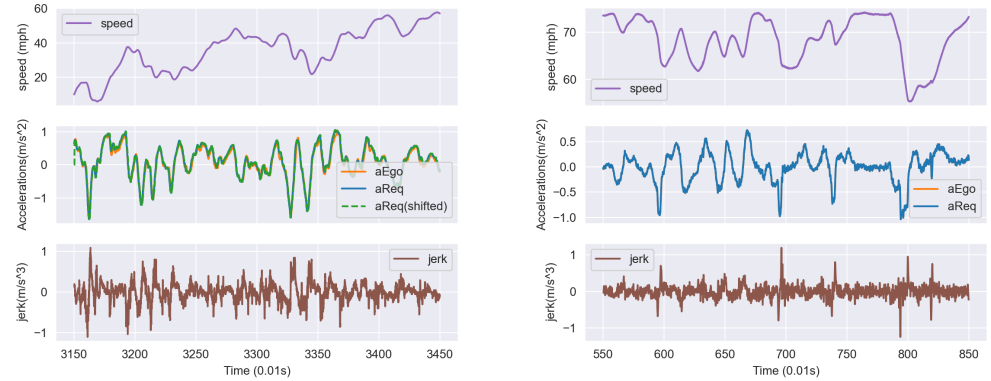
The MOBIL model posits that a driver's decision to initiate a lane change is based on a cost-benefit evaluation of acceleration gains. It assesses not only the potential improvement for the ego vehicle but also considers the impact on surrounding vehicles — both in the current lane and the target lane. Central to this evaluation is the need to estimate acceleration responses, which MOBIL computes using an underlying car-following (CF) model.

However, traditional CF models — such as Newell's model, IDM, and Gipps — were primarily designed for longitudinal scenarios in steady-state or car-following conditions. While effective in capturing normal acceleration and deceleration behavior, these models are not well-suited for the dynamic, multi-vehicle interactions that occur during lane changes. Specifically, they often assume instantaneous reactions to lead vehicles, lacking the ability to model forward-looking or anticipatory behavior. This presents a fundamental limitation: human drivers do not make lane-change decisions based solely on the present moment. Instead, they engage in a predictive process — anticipating how other vehicles will move over the next few seconds, whether gaps will remain open, and whether the maneuver can be completed safely. Instantaneous CF models fail to account for this predictive reasoning, which is particularly critical when merging into high-speed or dense traffic.




(a) Camry (b) Rav4

Figure 4.7: Acceleration and jerk values from ACC in EVs




(a) Ioniq5 (b) Mach-E

Figure 4.8: Acceleration and jerk values from ACC in EVs

### 4.4.1 MPC in the spirit of MOBIL

To bridge this research gap and more realistically represent human anticipation behavior during lane changes, we propose the use of Model Predictive Control (MPC) as an advanced hypothetical car-following model within the MOBIL framework. MPC naturally accommodates prediction over a future horizon, capturing anticipatory driving behavior by considering not only immediate but also future states of both ego and surrounding vehicles.

This MPC design introduces a novel approach by explicitly considering both the target lane leader and follower vehicles during a lane change maneuver. Our primary innovation lies in extending the traditional state representation to include the follower vehicle's dynamics. This extension is crucial since the ego vehicle's actions directly influence the follower's response, thus affecting overall system safety and performance across the entire MPC horizon.

#### Dynamic Model

The system models both the ego vehicle and a follower vehicle. The state vector is defined as:

$$\mathbf{x} = \begin{bmatrix} x_{\text{ego}} \\ v_{\text{ego}} \\ a_{\text{ego}} \\ x_{\text{follow}} \\ v_{\text{follow}} \end{bmatrix}, \tag{4.3}$$

where:

- $x_{\text{ego}}$ and $x_{\text{follow}}$ are the longitudinal positions of the ego and follower vehicles, respectively,
- $v_{\text{ego}}$ and $v_{\text{follow}}$ are their velocities,
- $a_{\text{ego}}$ is the ego vehicle's acceleration.

The control input is the jerk of the ego vehicle, denoted as $j_{\text{ego}}$. The system dynamics are described by the following equations:

$$\dot{x}_{\text{ego}} = v_{\text{ego}}, \tag{4.4}$$

$$\dot{v}_{\text{ego}} = a_{\text{ego}}, \tag{4.5}$$

$$\dot{a}_{\text{ego}} = j_{\text{ego}}, \tag{4.6}$$

$$\dot{x}_{\text{follow}} = v_{\text{follow}}, \tag{4.7}$$

$$\dot{v}_{\text{follow}} = \text{sat}\left(K_S \Delta x + K_V \Delta v, \text{accel}_{\min}, \text{accel}_{\max}\right), \tag{4.8}$$

where:

- $\Delta x = x_{\text{ego}} - x_{\text{follow}} - \tau_{\text{relax}} v_{\text{follow}} - d_{\text{stop}}$ represents the relative distance between the ego and follower vehicles, adjusted for a relaxation time constant $\tau_{\text{relax}}$ and a stopping distance $d_{\text{stop}}$,
- $\Delta v = v_{\text{follow}} - v_{\text{ego}}$ is the relative velocity,
- $\text{sat}(\cdot)$ ensures the follower's acceleration remains within the bounds $[\text{accel}_{\min}, \text{accel}_{\max}]$.

**Cost Function**

The cost function is designed to minimize tracking errors, control effort, and unsafe gaps. It is formulated as:

$$J = \sum_{k=0}^{N} \left(\mathbf{w}^\top \mathbf{e}_k^2\right) + \mathbf{w}_e^\top \mathbf{e}_N^2, \tag{4.9}$$

where:

- $\mathbf{e}_k$ is the error vector at time step $k$, defined as:

$$\mathbf{e}_k = \begin{bmatrix} x_{\text{obstacle},k} - x_{\text{ego},k} - v_{\text{ego},k}\tau_{\text{follow}} - d_{\text{stop}} \\ x_{\text{ego},k} \\ v_{\text{ego},k} - v_{\text{lead},k} \\ a_{\text{ego},k} \\ a_{\text{ego},k} - a_{\text{ego},k-1} \\ x_{\text{ego},k} - x_{\text{follow},k} - \tau_{\text{relax}} v_{\text{follow},k} - d_{\text{stop}} \\ j_{\text{ego},k} \end{bmatrix}, \tag{4.10}$$

where each component is described below:

1. Lead Spacing Error ($e_{\text{obstacle},k}$): This term penalizes deviations from a safe and efficient spacing with the leader:

$$e_{\text{obstacle},k} = x_{\text{obstacle},k} - x_{\text{ego},k} - v_{\text{ego},k}\tau_{\text{follow}} - d_{\text{stop}}. \tag{4.11}$$

Here, $x_{\text{obstacle},k}$ is the position of the leading obstacle, $x_{\text{ego},k}$ is the ego vehicle's position, $\tau_{\text{follow}}$ is the desired time headway, and $d_{\text{stop}}$ is the stopping distance.

2. Ego Position Error ($e_{\text{ego},k}$): This term penalizes deviations from a desired reference position:

$$e_{\text{ego},k} = x_{\text{ego},k} - x_{\text{ref},k}. \tag{4.12}$$

In most cases, $x_{\text{ref},k}$ is set to zero unless a specific reference trajectory is provided by neural networks.

3. Velocity Error ($e_{\text{velocity},k}$): This term penalizes deviations from the desired velocity:

$$e_{\text{velocity},k} = v_{\text{ego},k} - v_{\text{ref},k}, \tag{4.13}$$

where $v_{\text{ref},k}$ is the reference velocity, often set to the speed of the lead vehicle or a predefined cruise speed.

4. Acceleration Error ($e_{\text{acceleration},k}$): This term penalizes deviations from a desired acceleration profile:

$$e_{\text{acceleration},k} = a_{\text{ego},k} - a_{\text{ref},k}. \tag{4.14}$$

The reference acceleration $a_{\text{ref},k}$ is typically set to zero for comfort.

5. Acceleration Change Error ($e_{\text{accel\_change},k}$): This term penalizes abrupt changes in acceleration to ensure smooth driving:

$$e_{\text{accel\_change},k} = a_{\text{ego},k} - a_{\text{ego},k-1}. \tag{4.15}$$

6. Follower Risk ($e_{\text{follower},k}$): This term penalizes deviations from a safe distance to the lane change follower vehicle:

$$e_{\text{follower},k} = x_{\text{ego},k} - x_{\text{follow},k} - \tau_{\text{relax}} v_{\text{follow},k} - d_{\text{stop}}. \tag{4.16}$$

Here, $x_{\text{follow},k}$ and $v_{\text{follow},k}$ are the position and velocity of the follower vehicle, respectively, and $\tau_{\text{relax}}$ is a relaxation time constant.

7. Jerk Error ($e_{\text{jerk},k}$): This term penalizes high jerk values to ensure passenger comfort:

$$e_{\text{jerk},k} = j_{\text{ego},k}, \tag{4.17}$$

where $j_{\text{ego},k}$ is the jerk of the ego vehicle.

- $\mathbf{w}$ and $\mathbf{w}_e$ are weight vectors for the stage and terminal costs, respectively.

**Constraints**

The system is subject to the following constraints:

$$v_{\text{ego}} \geq 0, \tag{4.18}$$

$$a_{\min,k} \leq a_{\text{ego},k} \leq a_{\max,k}, \tag{4.19}$$

$$x_{\text{obstacle},k} \geq x_{\text{ego},k}, \tag{4.20}$$

$$x_{\text{ego},k} \geq x_{\text{follow},k} + \tau_{\text{relax}} v_{\text{follow},k} + d_{\text{stop}}. \tag{4.21}$$

These constraints ensure:

i. Non-negative ego vehicle velocity,

ii. Acceleration within dynamically adjustable bounds $[a_{\min,k}, a_{\max,k}]$, which is significant here to capture the dynamic constraints changing with the speed,

iii. Safe following distance from the leading obstacle,

iv. Safe separation from the follower vehicle.

The MPC design incorporates several key features to ensure robust and efficient performance:

**Non-Uniform Time Grid** A quadratic time grid is used for the prediction horizon:

$$t_k = T_{\max} \left( \frac{k}{N} \right)^2, \quad k = 0, 1, \ldots, N.$$

This prioritizes near-term predictions, reducing computational load while maintaining accuracy.

**Coupled Ego-Follower Dynamics** The follower's acceleration is modeled as a function of the ego vehicle's state, ensuring interactive safety and realistic behavior in multi-vehicle scenarios.

**Time-Varying Acceleration Limits** The acceleration bounds $a_{\min,k}$ and $a_{\max,k}$ are adjusted dynamically over the horizon to account for actuation delays and other real-world constraints.

**Time-Varying Weight Design** The MPC formulation incorporates time-varying weights in the cost function to dynamically prioritize control objectives over the prediction horizon. This design ensures that the controller responds aggressively to immediate safety concerns while gradually relaxing constraints on comfort-related metrics, such as jerk and acceleration changes, over time. The time-varying weights mimic human driver behavior, where near-term hazards are addressed with urgency, and long-term goals are pursued more smoothly.

The cost function at each time step $k$ is defined as:

$$J_k = \mathbf{w}_k^\top \mathbf{e}_k^2,$$

where $\mathbf{w}_k$ is the time-varying weight vector and $\mathbf{e}_k$ is the error vector. The weights $\mathbf{w}_k$ are adjusted by a time-varying scaling factor defiend as follows:

1. Acceleration Change Cost ($w_{a,k}$): The weight on acceleration changes is reduced over the horizon to allow smoother transitions in later stages:

$$w_{a,k} = w_{a,0} \cdot \text{interp}\left(t_k, [0, t_1, t_2], [1.0, 1.0, 0.5]\right),$$

where interp$(\cdot)$ denotes linear interpolation. This reduction reflects the decreasing importance of abrupt acceleration changes as the prediction horizon extends.

2. Follower Separation Cost ($w_{f,k}$): The weight on maintaining a safe distance from the follower vehicle is increased initially and then gradually relaxed:

$$w_{f,k} = w_{f,0} \cdot \text{interp}\left(t_k, [0, t_1, t_2], [10.0, 5.0, 1.0]\right).$$

This design ensures that immediate separation is prioritized, while the urgency diminishes over time.

3. Jerk Cost ($w_{j,k}$): The weight on jerk is kept constant or adjusted based on the driving scenario. For example, in high-speed scenarios, the jerk weight may be

reduced to prioritize smoothness:

$$w_{j,k} = w_{j,0} \cdot \text{scaling factor}.$$

The time-varying weights are computed at each time step $k$ based on the current time $t_k$ and predefined interpolation profiles. These weights are applied to the corresponding terms in the cost function, ensuring that the optimization problem reflects the desired priorities at each stage of the horizon. The interpolation profiles are designed to be flexible, allowing for scenario-specific tuning (e.g., urban driving vs. highway driving).

The time-varying weight design offers several advantages:

(i) Behavioral Realism: The weight profiles mimic human driving behavior, making the controller's actions more intuitive and predictable.

(ii) Improved Safety: By prioritizing near-term safety constraints, the controller can react aggressively to immediate hazards, reducing the risk of collisions.

(iii) Enhanced Comfort: The gradual relaxation of comfort-related weights ensures smoother trajectories over the long term, improving passenger comfort.

(iv) Computational Efficiency: The design reduces the need for overly conservative constraints, leading to faster convergence of the optimization solver.

For example, consider a scenario where the ego vehicle is approaching a slower-moving lead vehicle. The time-varying weights ensure that, in the near term, the controller prioritizes maintaining a safe distance from the lead vehicle (high $w_{f,k}$) and minimizing abrupt acceleration changes (high $w_{a,k}$). Over the longer term, the controller focuses on achieving the desired speed smoothly, with reduced emphasis on immediate separation and acceleration changes.

The time-varying weight design is a key feature of the MPC formulation, enabling the controller to balance competing objectives such as safety, comfort, and efficiency.

By dynamically adjusting the weights over the prediction horizon, the controller can respond effectively to immediate hazards while maintaining smooth and realistic trajectories.

### 4.4.2 MPC Results and Impact Analysis

To demonstrate the agility advantages of EVs and how they facilitate lane changes (LC) compared to ICEVs, we conducted a two-lane evacuation traffic simulation, where a lane drop occurs — such as one caused by a fallen tree, crash, or other obstruction, as illustrated in **Figure 4.9**., The road is divided into:

- **Lane 0 (Congested Lane)** downstream congestion and a physical bottleneck is simulated.
- **Lane 1 (Free-flow Lane)** an open alternative route, available to vehicles capable of making lane changes.

Simulation results indicate that EVs perform lane changes more effectively near the bottleneck area. In contrast, ICEVs frequently become trapped behind the bottleneck, experiencing difficulty in maneuvering into shorter gaps due to their limited acceleration capabilities. The enhanced acceleration and superior LC agility of EVs enable them to navigate these challenging conditions more effortlessly, efficiently bypassing the bottleneck.

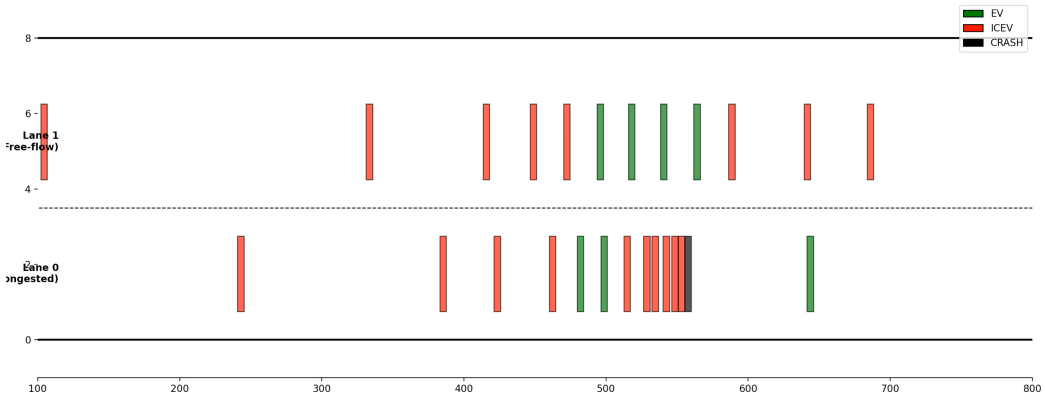



Figure 4.9: Simulating EV&ICEV under Evacuation.

### 4.4.3 Comparing EV and ICEV in lane change performance

The designed MPC model explicitly captures the enhanced capabilities of EVs in executing lane changes. **Figure 4.10** compares EV and ICEV behaviors within an identical LC scenario, featuring a substantial gap ahead and a tighter gap with the following vehicle in the target lane. The subplots clearly demonstrate EVs' superior acceleration, allowing quicker gap closure with the leader vehicle, consequently enhancing overall traffic efficiency by minimizing road spacing. Additionally, the third subplot highlights the accurate representation of EVs' trapezoidal acceleration pattern and distinctly illustrates ICEVs' delayed acceleration responses. Furthermore, ICEVs exhibit notable velocity reductions, amplifying oscillations in traffic flow, which may propagate and exacerbate congestion upstream.

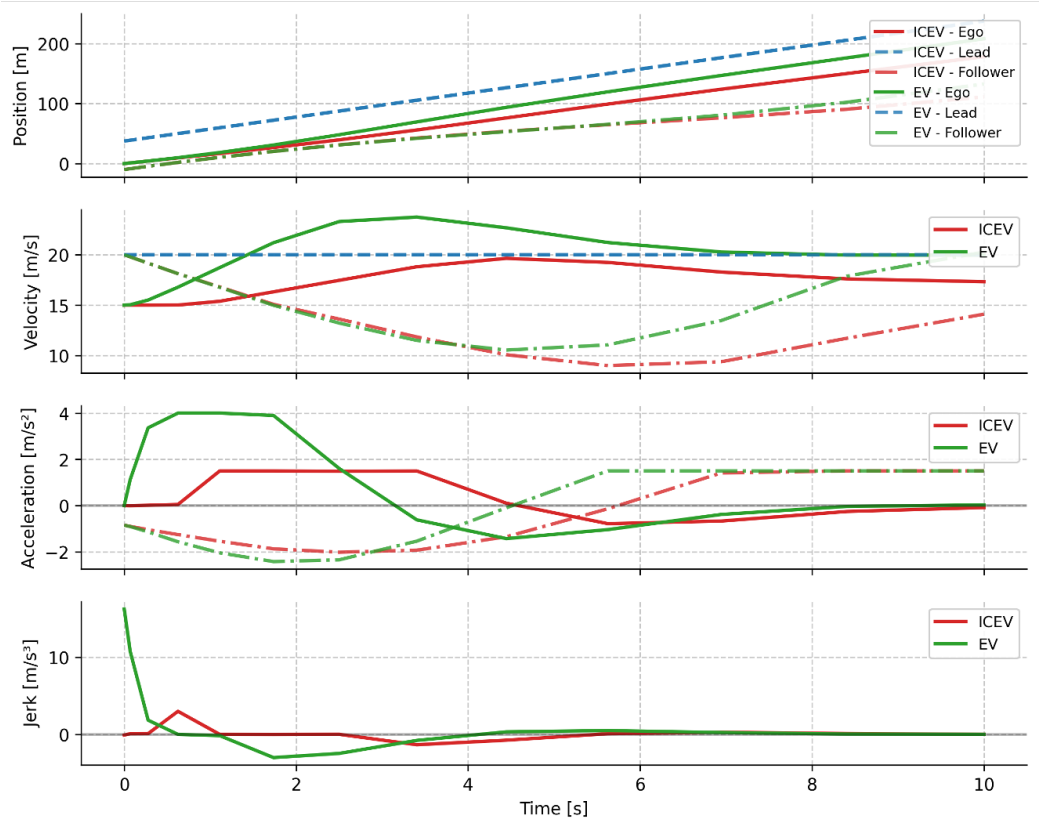


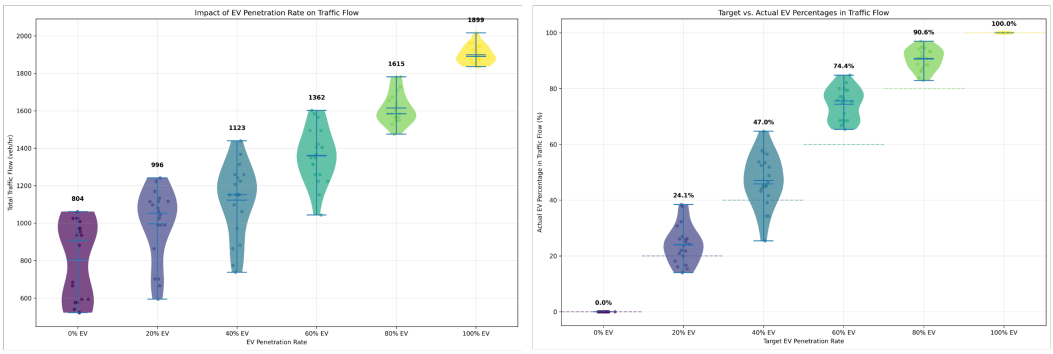

Figure 4.10: Comparing EV & ICEV's response in the same LC scenario.

### 4.4.4 Impact of EV penetration on evacuating traffic throughput

**Figure 4.11** visualizes the distribution of simulation outcomes across 20 randomized runs for each EV penetration level, capturing both the variability and central tendency of traffic throughput and EV flow share. **Figure 4.11a** illustrates the impact of EV penetration on overall traffic throughput under an evacuation scenario. The results demonstrate a clear upward trend: as the proportion of EVs in the traffic stream increases, total evacuation throughput (measured in vehicles per hour) improves significantly. This increase is nonlinear, with throughput rising from 804veh/hr at 0% EV to 1899veh/hr at 100% EV. The improvement is attributed to the superior acceleration and lane-changing capabilities of EVs, which allow them to navigate bottlenecks and lane-drop conditions more efficiently than ICEVs.

Meanwhile, **Figure 4.11b** compares the intended (target) EV penetration rate to the actual EV percentage observed in the post-bottleneck traffic flow. Interestingly, the actual EV share consistently exceeds the initial target. This deviation is a result of ICEVs being more likely to get stuck behind the bottleneck due to their limited agility and slower lane-changing behavior. In contrast, EVs are more capable of escaping congestion, thereby disproportionately contributing to downstream flow. This result underscores a form of "natural selection" under high-stress traffic conditions: the system effectively self-sorts in favor of more agile vehicles, amplifying the mobility benefits of EVs and enhancing evacuation efficiency even beyond their numerical presence in the initial fleet composition.

In summary, EVs significantly outperform ICEVs in traffic conditions that require frequent lane changes, such as evacuation scenarios. The MPC model effectively simulates the decision-making process of human drivers during lane changes, demonstrating EVs' superior capability in maneuvering into smaller gaps, a critical advantage for efficient and safe traffic evacuation.




(a) Experiment view from the driver (b) Dash-cam view from the 3rd EV

Figure 4.11: Traffic throughput from 20 simulation rollouts for each EV penetration level)

## 4.5 Outlook for EAV's Potential in Evacuation

According to existing traffic flow theories, the superior acceleration and deceleration capabilities should prevent capacity loss, and fewer response delays in EVs should offer better string stability, especially when combined with automation to form electric autonomous vehicles (EVAs). To understand the potential impact of EVAs on mobility, we conducted an EV platoon experiment consisting of three EVs: Mach-E, Ioniq 5, and EV6, as shown in **Figure 4.12**. We used the Comma 3X device to obtain CAN data from these EVs, allowing us to measure their radar-measured spacings and construct accurate platoon trajectories.

Current automation technologies in EVs have yet to fully capitalize on the performance advantages of electrification. Despite their stronger acceleration and deceleration capabilities, EVs under current Adaptive Cruise Control (ACC) systems often behave conservatively, leaving excessive spacing and reacting sluggishly—particularly when the lead vehicle begins to accelerate. This leads to overshooting behavior and inefficient use of available capacity, a pattern previously observed in ICE vehicle platoons.




(a) Experiment view from the driver (b) Dash-cam view from the 3rd EV

Figure 4.12: EV platoon experiments in Tampa, Florida.

To investigate this further, we conducted an EV platoon experiment using three popular models: the Ford Mach-E, Hyundai Ioniq 5, and Kia EV6. With CAN data captured via the Comma 3X device, we were able to reconstruct radar-based spacing and acceleration trajectories for each vehicle. These trajectories allowed us to assess how real-world EVs behave under automated following conditions, especially during acceleration transitions.

Although EVs are mechanically capable of quick, smooth acceleration, our observations show that current automation systems do not fully exploit this potential. Instead, the vehicles leave large gaps when the lead vehicle pulls away, failing to close them effectively or regain cruising speed in a timely manner. This inefficiency results in a measurable loss of traffic capacity.

**Fig. 4.13** below illustrates this capacity loss. As the lead EV accelerates, the follower vehicles respond with delay, leaving persistent gaps and disrupting flow continuity. This result highlights a key limitation in current EV automation: while the physical capabilities of EVs offer clear potential to improve traffic flow, they are underutilized by existing control logic. To unlock the benefits of Electric Autonomous Vehicles (EAVs), automation systems must be redesigned to align with EV-specific performance characteristics. Only then can we reduce capacity loss, improve traffic efficiency, and fully realize the advantages of electrification in dynamic traffic environments.

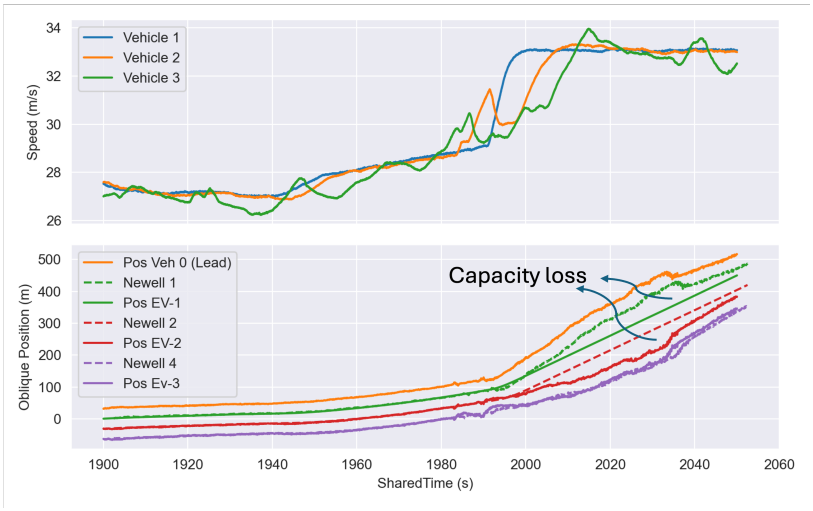



Figure 4.13: Capacity loss in the EV platoon experiment.

# Chapter 5

# Conclusion

## 5.1 Summary of Key Findings from Work 1: Network Resilience to Infrastructure Failure

Work 1 examined freight network resilience in response to a critical infrastructure disruption—the temporary closure of the Hernando De Soto (I-40) Bridge in Memphis in 2021. This event rerouted significant volumes of truck and passenger traffic onto alternative corridors, primarily I-55, creating an ideal natural experiment to assess the system's response and recovery.

Two complementary approaches were used: traditional resilience metrics (resilience triangle) and a novel distribution-based analysis using congestion clusters. The resilience triangle quantified performance loss and recovery time using indicators such as truck miles traveled and average travel time. It provided a high-level view of service degradation and rebound but was limited in capturing spatial-temporal traffic dynamics.

To address this, we introduced a distributional analysis of congestion events. We defined congestion clusters as contiguous space-time regions below a speed threshold and measured their duration, spatial extent, and area (delay). These metrics were

found to follow long-tail, power-law-like distributions, allowing deeper insight into the system's behavior under stress.

By examining the slopes of these distributions over time, we captured changes in system resilience. Steeper slopes indicate that large congestion clusters are rare—suggesting resilience—while shallower slopes imply the system struggles to contain disruptions. During the bridge closure, the slope flattened significantly, reflecting widespread, prolonged congestion. As traffic adapted, particularly after the opening of a new ramp lane, the slope steepened again, indicating recovery and improved resilience.

Key insights include:

The closure caused a surge in total delay and congestion cluster size, with the system operating near a critical threshold of percolation.

The reopening of a ramp lane mid-closure helped restore system function by absorbing excess demand, visible in a measurable shift in the slope of cluster-size distributions.

After reopening, although traffic volume returned to normal, the slope of the power-law distribution did not fully revert to pre-closure values, suggesting that even temporary infrastructure failures can leave lasting impacts on network behavior and adaptation.

Empirical data lacked direct queue length observations, so SUMO-based simulations were used to explore theoretical limits of resilience under varying truck percentages. These simulations confirmed that higher truck proportions lower both jam dissolution rates and overall capacity, thus reducing resilience.

Altogether, Chapter 1 establishes that resilience is not simply a static performance metric, but an emergent property visible through the accumulation of many small and large congestion events over time. The slope of these distributions becomes a

meaningful proxy for system-wide resilience, capable of capturing both deterioration and adaptation in response to disruption.

One key observation from the power-law slope analysis is that the slope following the reopening of the I-40 bridge became steeper than it was prior to the closure. Intuitively, the power-law slope, often derived from the distribution of congestion size or duration, serves as a proxy for a system's resilience. A steeper slope (e.g., $-1.84$) indicates that large congestion events are significantly less frequent, suggesting that the network is better able to absorb and suppress disruptions before they propagate. Since the slope is not a momentary measure but rather an emergent property based on system behavior over time, it reflects the underlying network's capacity to withstand, absorb, and recover from disturbances.

From a resilience standpoint, one might expect the post-closure slope to return to its pre-closure value, assuming the system merely recovered to its prior state. However, the fact that the post-closure slope is steeper indicates that the system did not just recover—it adapted and improved. A major contributing factor could be the opening of a new ramp lane during the repair period, which constitutes a topological change to the network. Such an infrastructure enhancement increases system redundancy and flow flexibility, reducing the likelihood of bottlenecks and large-scale congestion clusters. Additionally, behavioral adaptations by drivers and freight operators during the closure (e.g., time-of-day shifts or alternate routing) may have had lasting effects that continued to suppress congestion even after the bridge reopened. Combined with potential institutional improvements—such as optimized signal timing or operational adjustments—these factors point to a form of transformative resilience, where the system evolves into a more robust and adaptive state rather than simply returning to baseline conditions.

## 5.2 Summary of Key Findings from Work 2: Evacuation Traffic Monitoring and Control Strategy

This chapter explored a percolation-based framework for managing evacuation traffic, focusing on how to dynamically regulate the rate at which vehicles are loaded into a congested network. Building on the previous chapter's emphasis on structural resilience, this chapter addresses demand-driven disruptions, such as those seen during emergency evacuations, where abnormal surges in traffic can quickly overwhelm a network.

Through microscopic simulation in SUMO, we modeled an abstract city network with a Central Business District (CBD) serving as the evacuation zone. Vehicles were loaded into the network at varying time intervals to assess how loading rates affected system-wide performance. A key finding is the emergence of a critical equilibrium point: a narrow range of loading rates at which vehicle inflow into the CBD matches the outflow to shelters. At this point, the system achieves optimal evacuation efficiency—maximizing throughput while avoiding system collapse.

Importantly, this equilibrium point was not imposed but emerged consistently across multiple randomized simulation runs. It represents a tipping point: just beyond it, localized congestion clusters begin to grow, merge, and form large, connected jams—indicating that the network has reached its percolation threshold. At this stage, mobility collapses and evacuation performance deteriorates rapidly. The percolation transition and the breakdown of flow equilibrium are tightly aligned, reinforcing the validity of this threshold as a target for real-time traffic control.

To manage this risk in operational settings, we introduced a percolation-based evacuation control strategy. The core principle is to maintain congestion clusters in a scattered, contained state—ensuring they do not reach the edge of the CBD. This preserves the equilibrium condition and sustains high evacuation throughput.

The strategy is implemented through four rules (R1–R4), which provide practical guidance for monitoring congestion growth, adjusting demand, and allocating traffic across shelters in real time.

Overall, this chapter highlights that in evacuation scenarios—unlike routine operations—the goal is not to avoid congestion entirely, but to push the system right up to the edge of its capacity without crossing into collapse. The percolation-based framework offers both a theoretical lens and a practical toolkit to achieve that balance.

## 5.3 Summary of Key Findings from Work 3: EV Potential in Improving Evacuation Traffic

This chapter investigates the potential of electric vehicles (EVs) to enhance evacuation traffic efficiency. It addresses critical gaps in EV road testing and empirical data collection, as well as the development of theoretical frameworks and models necessary to capture the impact of EVs on evacuation traffic flow.

Work 3 begins by presenting road experiments conducted on several of the most popular EV models currently available in the U.S. market, encompassing major brands such as Tesla, Ford, Hyundai, and Kia. Using a novel data collection method, we obtained high-resolution kinematic data from the vehicle's Controller Area Network (CAN), including precise speed profiles and pedal position information. This granular data allowed us to characterize EVs' agility, highlighting their superior acceleration capabilities, which exhibit a distinct trapezoidal pattern.

To quantify the impact of EVs' agility on traffic flow, this chapter develops a model predictive control (MPC) framework that replicates the lane-changing (LC) decision process of human drivers. The MPC model incorporates both leader and follower dynamics, while accounting for the linear acceleration constraints typically associated with internal combustion engine vehicles (ICEVs). Model evaluation under varying

acceleration capabilities demonstrates that EVs' superior acceleration incentivizes more frequent lane changes compared to ICEVs, enabling vehicles to maneuver more effectively through bottlenecks and congestion, which are common in evacuation scenarios.

Further, Work 3 introduces two ordinary differential equation (ODE) models to analyze the effects of EV-induced lane changes on target lane traffic disruption and safety, specifically assessing the collision risk between the lane-changing vehicle and the target lane follower. These ODE models provide analytical insights into how EVs' agility can mitigate delays and enhance safety by reducing the risk of collisions during lane changes.

The chapter concludes by presenting a two-lane traffic simulation of an evacuation scenario involving a fallen tree and subsequent lane drop. The simulation results demonstrate that higher EV penetration significantly improves roadway throughput, reinforcing the potential of EVs to enhance evacuation traffic efficiency.

# Bibliography


Alawadi, R., Murray-Tuite, P., and Bian, R. (2022). Determinants of departure timing for hurricane matthew and anticipated consistency in future evacuation departures. *Natural Hazards Review*, 23(2):04022005. 3

Alvarez Lopez, P., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., and Wießner, E. (2018). Microscopic traffic simulation using sumo. In *IEEE Intelligent Transportation Systems Conference (ITSC)*. 4, 32

Ambühl, L., Menendez, M., and González, M. C. (2023). Understanding congestion propagation by combining percolation theory with the macroscopic fundamental diagram. *Communications Physics*, 6(1):26. 4, 36, 60

Ayyub, B. M. (2014). Systems resilience for multihazard environments: Definition, metrics, and valuation for decision making. *Risk analysis*, 34(2):340–355. 2

Barrett, B., Ran, B., and Pillai, R. (2000). Developing a dynamic traffic management modeling framework for hurricane evacuation. *Transportation Research Record*, 1733(1):115–121. 3

Brown, C., White, W., van Slyke, C., and Benson, J. D. (2009). Development of a strategic hurricane evacuation–dynamic traffic assignment model for the houston, texas, region. *Transportation research record*, 2137(1):46–53. 3

Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O'Rourke, T. D., Reinhorn, A. M., Shinozuka, M., Tierney, K., Wallace, W. A., and Von Winterfeldt, D. (2003). A framework to quantitatively assess and enhance the seismic resilience of communities. *Earthquake spectra*, 19(4):733–752. 2

Cascetta, E. (2009). *Transportation systems analysis: models and applications*, volume 29. Springer Science & Business Media. 3

Chen, X. and Zhan, F. B. (2008). Agent-based modelling and simulation of urban evacuation: relative effectiveness of simultaneous and staged evacuation strategies. *Journal of the Operational Research Society*, 59(1):25–33. 3

Chiu, Y.-C., Zheng, H., Villalobos, J. A., Peacock, W., and Henk, R. (2008). Evaluating regional contra-flow and phased evacuation strategies for texas using a large-scale dynamic traffic simulation and assignment approach. *Journal of Homeland Security and Emergency Management*, 5(1). 3

commaai (2024). opendbc. https://github.com/commaai/opendbc. Accessed: 2024-08-02. 5

Daganzo, C. F. (1994). The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory. *Transportation research part B: methodological*, 28(4):269–287. 59

Daganzo, C. F. and Geroliminis, N. (2008). An analytical approximation for the macroscopic fundamental diagram of urban traffic. *Transportation Research Part B: Methodological*, 42(9):771–781. 60

Faturechi, R. and Miller-Hooks, E. (2015). Measuring the performance of transportation infrastructure systems in disasters: A comprehensive review. *Journal of infrastructure systems*, 21(1):04014025. 2

Feng, K. and Lin, N. (2021). Reconstructing and analyzing the traffic flow during evacuation in hurricane irma (2017). *Transportation research part D: transport and environment*, 94:102788. 3

Feng, K. and Lin, N. (2022). Modeling and analyzing the traffic flow during evacuation in hurricane irma (2017). *Transportation research part D: transport and environment*, 110:103412. 3

Ganin, A. A., Kitsak, M., Marchese, D., Keisler, J. M., Seager, T., and Linkov, I. (2017). Resilience and efficiency in transportation networks. *Science advances*, 3(12):e1701079. 3

Geroliminis, N., Haddad, J., and Ramezani, M. (2012). Optimal perimeter control for two urban regions with macroscopic fundamental diagrams: A model predictive approach. *IEEE Transactions on Intelligent Transportation Systems*, 14(1):348–359. 60

Gu, Y., Liu, D., Stanford, J., and Han, L. D. (2022). Innovative method for estimating large truck volume using aggregate volume and occupancy data incorporating empirical knowledge into linear programming. *Transportation Research Record*, 2676(11):648–663. 3

Hsu, Y.-T. and Peeta, S. (2013). An aggregate approach to model evacuee behavior for no-notice evacuation operations. *Transportation*, 40:671–696. 3

Hsu, Y.-T. and Peeta, S. (2014). Risk-based spatial zone determination problem for stage-based evacuation operations. *Transportation research part C: emerging technologies*, 41:73–89. 3

Hsu, Y.-T. and Peeta, S. (2015). Clearance time estimation for incorporating evacuation risk in routing strategies for evacuation operations. *Networks and Spatial Economics*, 15:743–764. 3

Intini, P., Ronchi, E., Gwynne, S., and Pel, A. (2019). Traffic modeling for wildland–urban interface fire evacuation. *Journal of Transportation Engineering, Part A: Systems*, 145(3):04019002. 3

Lapardhaja, S., Yagantekin, K. U., Yang, M., Majumder, T. A., Kan, X., and Badhrudeen, M. (2023). Unlocking potential capacity benefits of electric vehicles (evs) with adaptive cruise control (acc). *Transportmetrica B: Transport Dynamics*, 11(1):1894–1911. 4

Laval, J. and Menendez, M. (2023). Urban network percolation and traffic dynamics: Causality relationships in numerical experiments. 4

Laval, J. A. (2023). Self-organized criticality of traffic flow: Implications for congestion management technologies. *Transportation Research Part C: Emerging Technologies*, 149:104056. 4

Laval, J. A., Toth, C. S., and Zhou, Y. (2014). A parsimonious model for the formation of oscillations in car-following models. *Transportation Research Part B: Methodological*, 70:228–238. 74

Li, D., Fu, B., Wang, Y., Lu, G., Berezin, Y., Stanley, H. E., and Havlin, S. (2015). Percolation transition in dynamical traffic network with evolving critical bottlenecks. *Proceedings of the National Academy of Sciences*, 112(3):669–672. 4, 27

Li, M., Liu, R.-R., Lü, L., Hu, M.-B., Xu, S., and Zhang, Y.-C. (2021a). Percolation on complex networks: Theory and application. *Physics Reports*, 907:1–68. 2

Li, T., Chen, D., Zhou, H., Laval, J., and Xie, Y. (2021b). Car-following behavior characteristics of adaptive cruise control vehicles based on empirical experiments. *Transportation research part B: methodological*, 147:67–91. 4

Makridis, M., Mattas, K., Anesiadou, A., and 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. 4

Makridis, M., Mattas, K., Ciuffo, B., Re, F., Kriston, A., Minarini, F., and Rognelund, G. (2020). Empirical study on the properties of adaptive cruise control systems and their impact on traffic flow and string stability. *Transportation research record*, 2674(4):471–484. 4

Martinez-Pastor, B., Nogal, M., O'Connor, A., and Caulfield, B. (2016). A sensitivity analysis of a dynamic restricted equilibrium model to evaluate the traffic network resilience. Technical report. 16

Mattsson, L.-G. and Jenelius, E. (2015). Vulnerability and resilience of transport systems–a discussion of recent research. *Transportation research part A: policy and practice*, 81:16–34. 2

Murray-Tuite, P. and Wolshon, B. (2013). Evacuation transportation modeling: An overview of research, development, and practice. *Transportation Research Part C: Emerging Technologies*, 27:25–45. 3

Newman, M. (2018). *Networks*. Oxford university press. 2

Pan, S., Yan, H., He, J., and He, Z. (2021). Vulnerability and resilience of transportation systems: A recent literature review. *Physica A: Statistical Mechanics and its Applications*, 581:126235. 21

Pel, A. J., Hoogendoorn, S. P., and Bliemer, M. C. (2010). Impact of variations in travel demand and network supply factors for evacuation studies. *Transportation Research Record*, 2196(1):45–55. 3

Sbayti, H. and Mahmassani, H. S. (2006). Optimal scheduling of evacuation operations. *Transportation Research Record*, 1964(1):238–246. 3

Serulle, N. U., Heaslip, K., Brady, B., Louisell, W. C., and Collura, J. (2011). Resiliency of transportation network of santo domingo, dominican republic: case study. *Transportation research record*, 2234(1):22–30. 16

Sharma, N., Tabandeh, A., and Gardoni, P. (2018). Resilience analysis: A mathematical formulation to model resilience of engineering systems. *Sustainable and Resilient Infrastructure*, 3(2):49–67. 2

Wu, D., Liu, H., Yuan, C., and Wei, D. (2016). A quantitative assessment framework for post-earthquake transportation network resilience using fuzzy logic. Technical report. 16

Zhou, H., Toth, C., Guensler, R., and Laval, J. (2022a). Hybrid modeling of lane changes near freeway diverges. *Transportation research part B: methodological*, 165:1–14. 74

Zhou, H., Zhou, A., Ding, Z., Laval, J., and Peeta, S. (2022b). Empirical study on the acceleration/deceleration constraints under commercial adaptive cruise control. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 3069–3076. IEEE. 4

# Vita

Ruqing Huang was born in Hunan, China. She received her Bachelor's degree in Transportation Engineering from Central South University and her Master's degree in the same field from the Georgia Institute of Technology. In 2020, she began pursuing her Ph.D. in Civil Engineering at the University of Tennessee, Knoxville, focusing on transportation system resilience, evacuation control strategies, and vehicle-level traffic modeling under extreme events.

She is a fan of basketball and works as an on-site sports journalist, covering NBA and UFC events alongside her academic research. Her work bridges data, mobility, and real-world storytelling.