Leveraging spatial relationships and visualization to improve public transit performance analysis

by
Nicholas S. Caros
M.S., New York University (2019)
B.A.Sc., University of British Columbia (2014)

Submitted to the Department of Civil and Environmental Engineering
in partial fulfillment of the requirements for the degree of
Master of Science in Transportation
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2021

Author . . . .
Department of Civil and Environmental Engineering
May 14, 2021

Certified by . . . .
John P. Attanucci
Research Associate, Center for Transportation and Logistics
Thesis Supervisor

Certified by . . . .
Anson F. Stewart
Research Scientist, Department of Urban Studies and Planning
Thesis Supervisor

Certified by . . . .
Andrew J. Whittle
Edmund K. Turner Professor of Civil Engineering
Thesis Supervisor

Accepted by . . . .
Colette L. Heald
Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

Leveraging spatial relationships and visualization to improve public transit performance analysis

by

Nicholas S. Caros

Submitted to the Department of Civil and Environmental Engineering
on May 14, 2021, in partial fulfillment of the
requirements for the degree of
Master of Science in Transportation

Abstract

Public transit agencies collect a tremendous amount of data in order to measure bus performance, including vehicle positions and passenger counts. These data are typically organized in the same way that the network is organized: split up by route, with each route further divided into stop-to-stop segments or timepoint-to-timepoint segments. This type of structure suffers from three drawbacks. First, it does not capture the spatial relationships between different routes. Second, route and stop identifiers are arbitrary and change over time, making it difficult to compare performance across time periods. Finally, even the stop-to-stop resolution is insufficient for certain applications, such as planning for transit priority infrastructure at the block or intersection level.

This thesis addresses those three issues by developing new practical methods that incorporate the geography of a transit network into bus performance measurement and analysis. These tools can be used to automate transit planning tasks that have typically involved specialized knowledge and considerable manual effort. Furthermore, each of the methods is intentionally designed to support visualization of performance data as an alternative to tabular representation in order to facilitate the identification of spatial patterns. A total of eight case studies, developed in concert with transit agency staff, are included to demonstrate how spatial analysis and visualization can address real transit planning challenges.

First, a map matching algorithm is described that facilitates the identification and classification of corridors served by multiple bus routes. A framework is then established for systematically aggregating different types of performance metrics across parallel routes, even if those performance metrics are only available at the stop-to-stop segment level. Case studies show how corridor identification and performance aggregation can be used to improve transit priority infrastructure planning, schedule coordination for parallel routes, and balancing service in local-express corridors.

Next, a method for increasing the resolution of performance data to the block-to-block level is proposed. Stop-to-stop segments are split at intersections and bus stops to create a unit of analysis that experiences uniform transit service across its length. Performance measures are then assigned to the block-level segments, eliminating the dependence on arbitrary and fungible identifiers. This geography-based representation enables longitudinal comparison of performance that automatically captures changes in transit service as well as route and stop numbers. Two case studies demonstrate how these methods can be used to map the evolution of bus networks and ridership over many years.

Finally, a process is developed for extending the previous methods to include origin-destination (OD) estimates, enabling the spatial analysis and visualization of passenger journeys throughout the transit network. It also allows OD-based performance metrics that are not available from traditional sources to be assigned to block-to-block segments for longitudinal comparison. Case studies illustrate the strength of these methods in mapping the journeys of passengers whose trips involve a transfer, and for exploring travel pattern changes before and after route modifications.

Future research in this area includes the development of new bus performance metrics that incorporate spatial relationships between bus routes. Richer data sources, such as vehicle positions with a high sampling rate, could be leveraged to visualize travel speeds at the block-to-block level. Another research area is facilitated by the longitudinal analysis methods in this thesis: the ability to test theories related to bus network evolution over time. Finally, these methods lay the foundation for future research into new transit planning tools for strategic restoration of bus service after the COVID-19 pandemic.

Thesis Supervisor: John P. Attanucci
Title: Research Associate, Center for Transportation and Logistics

Thesis Supervisor: Anson F. Stewart
Title: Research Scientist, Department of Urban Studies and Planning

Thesis Supervisor: Andrew J. Whittle
Title: Edmund K. Turner Professor of Civil Engineering

Acknowledgments

This thesis would not have been possible without support and help from a huge number of people. First, thank you to John Attanucci and Anson Stewart, for introducing me to this project and for all of your help and feedback along the way. I am grateful for the guidance and leadership of Professor Jinhua Zhao, who encouraged and supported my research over the past two years. Thank you also to the staff at the MBTA and CTA for generously sharing your time, answering my many questions and for inspiring many of the case studies in this thesis. Finally, thank you to Xiaotong Guo for your collaboration on the visualization dashboard used for this thesis. It has been a pleasure working with you.

Of course there are also many people who contributed indirectly to this project through their support and friendship. Thank you to all of the members of the Transit Lab and Urban Mobility Lab for your thoughtful comments, camaraderie and for building a wonderful community at MIT even when we couldn't meet in person. Thank you also to my friends and family, near and far, for putting up with my non-stop rambling about public transit. A special thank you to my parents, Mark and Mary Caros, who started all of this when they gave me train sets as a kid, and who have encouraged me to pursue my dreams ever since. Lastly, thank you to Nikki, who spent a lot more time with me than she bargained for over the past year, for everything that you do.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and Motivation

Modern public transit agencies collect a tremendous amount of data in order to measure performance, identify problems and provide information to the public. For buses, this typically includes vehicle position data from Automatic Vehicle Location (AVL) systems and passenger boarding and alighting data from Automatic Passenger Counting (APC) systems. Vehicle locations and passenger counts are generally aggregated at the bus stop, segment, or route level.

A typical framework for organizing performance measurement involves the entire network as the highest level, which is divided into many individual routes, with each route further divided into stop-to-stop or timepoint-to-timepoint segments. This conceptual structure is intuitive as it reflects how schedules are organized and how performance data is collected. It is a structure used by many U.S. transit agencies; the U.S. Federal Transit Administration's Transit Co-operative Research Program (TCRP) [3] proposes over 400 performance measures, the vast majority of which are scoped to the system, route and stop level. These scopes differ from the user's experience of the transit system, which is rooted in the geography of the network and how effectively it enables them to travel from one place to another.

Aggregation, partitioning and analysis by bus stop and route has three key drawbacks that limit the use of automated data to conduct certain types of analysis. First, it ignores the spatial relationships between routes that are necessary for a holistic understanding of network performance and service quality, especially where bus routes serve the same corridor. This is particularly critical in regions that are served by several different transit providers with overlapping service areas. Second, it does not permit analysis or comparison at a high-resolution, for example the city block level. Bus performance and service quality can vary significantly across two consecutive blocks. Additionally, routes may overlap for only a portion of a stop-to-stop segment, which should be reflected in performance analysis and visualization. Finally, the route and stop identifiers are entirely arbitrary and often change over time. Changing identifiers limits any systematic comparison of performance and service over long periods of time.

To illustrate these compounding issues, consider the following challenge that is frequently confronted within transit agencies. The transit planning team is interested in identifying and visualizing high impact locations for installing transit signal priority, a type of traffic signal timing that reduces wait time for buses at signalized intersections. They have access to boarding and alighting data at each bus stop ID for each route ID. To determine how many passengers would benefit from transit signal priority at a given traffic signal, they must identify the passenger load for trips traveling through that intersection. If more than one route (or agency) traverses an intersection, the sum of passenger loads should be taken to incorporate all passengers that stand to benefit from transit signal priority. Locating these overlapping routes can be difficult if they do not share stops, such as a local and express route pair. Furthermore, the overlapping routes may branch together or apart between stops, so a smaller unit of analysis than the stop-to-stop segment is necessary in order to accurately represent aggregate passenger flows through traffic signals. If the planners are interested in average flows over a several year period, they would then need to determine whether any route IDs or stop IDs have changed throughout that period.

As a result of these limitations, evaluating the relationships between routes or comparing changes in performance over time are typically undertaken on an ad-hoc basis with considerable manual effort and rely heavily on the experience and knowledge of staff. There is no standard data format that identifies overlapping bus routes, nor is there a typical means for converting between previous stop or route IDs and their new counterparts. To mitigate these limitations, this thesis presents a process for identifying shared corridors and for aggregating performance across these shared corridors depending on the type of performance measure. Then, a method is developed for decomposing the standard bus network representation into small, city-block level segments that are compared based on their geography rather than arbitrary identifiers. The spatial decomposition allows for simple aggregation of transit performance measures where routes overlap, and the geographic representation facilitates comparisons of performance over time. Finally, this method is extended for use with origin-destination data. Case studies based on real transit planning applications are provided throughout to demonstrate how these methods can be used in practice. These case studies are accompanied by maps of transit service developed for this thesis in order to illustrate how visualization can complement transit performance analysis.

It should be noted that some agencies do maintain the tools to mitigate some of the aforementioned drawbacks, although the processes typically involve significant manual effort. The need for improved methods was established through conversations with staff at several of the largest U.S. transit agencies from September 2019 to May 2021, many of whom expressed that these issues are salient and that systematic methods are generally not available. A review of the literature, which is included in Section 1.5, confirmed that the specific methods described in this thesis represent advances in the state of the practice and have not been previously developed.

1.2 Research Objectives

To address the limitations of current bus performance analysis practices, the objectives of this research is to develop generalizable, systematic methods for the following tasks:

1. Identifying shared corridors in a bus network and aggregating different types of

performance metrics across the shared corridors,

2. Decomposing bus networks into block-level units of analysis; and,

3. Converting a network representation based on impermanent identifiers to a consistent representation based on geography.

The benefits of these methods are illustrated through a series of realistic transit planning case studies, many of which were developed in concert with agency staff.

1.3 Research Approach

Data synthesis, geospatial analysis and visualization are the primary methods used in this thesis. Methods proposed in this thesis are described in detail for reproducibility, typically with worked examples. The methods are then tested for different applications using real transit data. The results, benefits and limitations of each method are discussed.

The methods described herein were intentionally developed to be generalizable to any transit agency, big or small, that maintains a compliant GTFS feed. The performance metrics are calculated using data sources that are widely available to transit agencies, with the deliberate exception of the metrics presented in Chapter 4. All of the software packages and libraries referenced herein are open-source and free to use. The programs for map matching, identifying shared corridors, block-level decomposition and edge matching are all publicly available at `https://github.com/nick-caros/mst-thesis`.

This thesis focuses exclusively on bus transit and intentionally omits rail transit. The methods included in this thesis certainly could be applied to rail networks; however, they are likely to be more beneficial for bus transit analysis. Most transit agencies have significantly more bus routes and stops than rail lines and stations, so comprehensive identification of relationships between bus routes is a more complex task. Furthermore, bus routes are much more likely change over time, thus increasing the need for an automated system of identification and aggregation.

1.4 Data Sources

There are five primary data sources used in this research. The first two are automated transit data collection systems used by agencies: Automated Vehicle Location (AVL) data and Automatic Passenger Counting (APC) data. AVL systems provide a location and timestamp at regular intervals or upon the occurrence of an event (e.g. the bus arrives at a bus stop). AVL is used to track the vehicle position over time, facilitate real-time dispatching, and provide estimates of speed, running times, delay relative to the schedule, and so on. APC systems provide records of passenger boardings and alightings at each bus stop, which are used to determine the passenger flow through the bus network. This research was generously sponsored the Chicago Transit Authority (CTA) and Massachusetts Bay Transportation Authority (MBTA); as part of the sponsorship, AVL and APC data from both agencies was shared with the author and are used as inputs for the case studies in this thesis.

The third data source is the Generalized Transit Feed Specification (GTFS). GTFS is a widely adopted data standard originally developed for integrating transit information into web maps [4]. There are static and real-time versions of GTFS; the static version is used as an input for the methods in this thesis. It contains information relating to the transit schedule and geography for a specific transit agency, including scheduled arrival times, stop sequences and bus stop locations [5].

The fourth primary data source is a graph representation of the street network in the transit service area, which is not included in the standard GTFS feed. The geography of the street network is used to determine which streets are traversed by each bus route, enabling the identification of spatial relationships between routes and longitudinal analysis that is not reliant on arbitrary identifiers. This thesis uses open street network data collected from OpenStreetMap.org [6], but other similar sources could be used instead.

An additional data source is included as part of Chapter 4: origin, destination and transfer flows. These records are used to illustrate the how a wide variety of data can be included in the overall analysis framework of this thesis to provide further benefits. Origin-destination flows can be estimated for transit systems with “tap-on” fare collection (see [7]) or collected directly for “tap-on, tap-off” systems. Systems for estimating or collecting origin, destination and transfer data are becoming more common as transit agencies move to provide real-time information to riders [8].

1.5 Literature Review

There is a long history of research into the measurement, visualization and analysis of transit networks. This section reviews relevant literature on creating transit performance measures, spatial analysis of bus networks, transit origin-destination data, visualization for transit planning, the analysis of shared bus corridors and longitudinal analysis of bus performance.

1.5.1 Transit Performance Measurement

Developing and applying TPMs has been the subject of considerable research over many decades [9, 10, 11]. Automated data such as AVL were not initially designed for performance measurement, so additional processing is typically required [12]. Hammerle et al. [13] describes an early application of AVL and APC to measure transit performance for the CTA. The TCRP subsequently published a report describing how archived AVL and APC data can be used to measure performance [14]. Ma and Wang [12] uses these data sources as well as Automated Fare Collection (AFC) to estimate speed, travel time reliability, ridership and headway variance. Pi et al. [15] created a platform for assessing service quality in the Pittsburgh region using archived APC and AVL data.

These data sources can also be used in support of operational decisions. APC and AVL data have been used in several studies to estimate bus arrival and departure times [16, 17, 18]. In another application, APC and AVL data were combined to recommend additional recovery time between inbound and outbound trips in order to reduce bus bunching [15]. Berkow et al. [19] combines historical data with weather information and uses statistical analysis to estimate passenger demand and dwell times. Moreira-

Matias et al. [20] reviews how transit agencies use AVL data to improve control strategies for the purpose of improving reliability. Real-time AVL, which is not used in this thesis, has also been studied and found to have great potential for improving dispatching and control [21].

Recent research has investigated the service planning applications of passive data sources, including APC, AVL, smart card records, and mobile location records [22, 12, 23, 24]. AVL data has been used to measure bus reliability and to recommend interventions for improved reliability such as stop consolidation and schedule changes [25, 22]. It has also been used to design optimal schedules with a high on-time probability [26]. Another research area is the development of methods for automatic identification of the causes of certain service deficiencies, such as bus bunching and arrival delay, using AVL and APC [27, 28]. Route design has also been informed by automated data [23]. These automated data are collected for individual trips and vehicles, so the applications have been focused on service improvements at the stop and route level. Aggregating the data to the corridor level permits a more holistic representation of the network that considers the relationship between different routes where overlapping occurs. A recent review paper provides an overview of big data applications in public transit [29].

It should be noted that bus performance analysis extends beyond operational performance. Other areas of interest, which are beyond the scope of this thesis, include productivity, equity, and accessibility. Interested readers may refer to Carleton and Porter [30] for a good summary of equity in public transit, and to Wei et al. [31] for a discussion of productivity and access.

1.5.2 Spatial Analysis

More recently, research into transit performance measurement has been expanded to include spatial characteristics of the network. One previous study argues that "a network wide and spatial perspective in exploring the operational performance of large transit systems is a worthwhile approach to identifying priorities for transit" [32]. Horner and Murray [33] finds that the spatial representation of transit networks can have a significant impact on the results of transit service assessment.

One of the key spatial analysis methods used in this thesis is a process known as map matching, which estimates the travel path of a bus route using only the stop locations. Prior studies have established methods for matching GTFS to a representation of the road network, although not for the purpose of identifying overlapping routes [34]. Wessel et al. [35], Li [36] and Ordóñez Medina [37] take different approaches to matching bus stop locations to street segments with directionality. Li [36] shows that a shortest-path matching algorithm is accurate for 98% of segments. Perrine et al. [38] proposes an algorithm for matching GTFS routes to a road network for multi-modal modeling. Zhou et al. [39] uses a method similar to Perrine et al., but uses raw GPS traces instead of GTFS as their input. None of these studies attempt to identify relationships between bus routes using the results of the map matching process.

The map matching process and other methods described in this thesis use the GTFS static feed as the source of the geographical and transit network data used to identify shared corridors. A standard GTFS feed contains information about bus routes, schedules and stop locations, but is not required to include information about the road network or the path that each route takes between stops. GTFS has been widely used for static analysis of transit network characteristics such as accessibility [40], equity [41] and competitiveness [42], and for other applications such as route planning [35]. Other studies have developed methods that use GTFS to evaluate the scheduled performance of transit networks [5].

Spatial decomposition of transit networks can also complement new data sources to develop performance metrics at a high spatial resolution. The relatively recent availability of GPS data with a high temporal sampling rate has led to several studies that analyze bus performance between stops [43]. For example, Stoll et al. [44] and Figliozzi and Stoll [8] use high-resolution AVL data to generate speed plots and determine the delay caused by intersections and crosswalks. In an interesting study, Figliozzi and Glick [45] reviews bus speeds before and after infrastructure improvements using high-resolution GPS traces, where routes on shared corridors are matched together in order to calculate average speeds. In each of these studies, the smaller unit of analysis is used to identify sources of delay due to different infrastructure elements, not to identify spatial relationships between overlapping routes that can change between bus stops. The concept of splitting routes at road intersections and stops was inspired by a TransLink project used for aggregating transit performance within shared segments [46].

1.5.3 Origin-Destination Flows in Public Transit

Understanding how passengers move throughout the transit network is another important component of transit performance measurement. This information is typically derived from passenger origin-destination flows, which can collected in a number of different ways. Origin-destination data is used to determine which routes feed one another, if and where passengers make transfers, and how long each journey takes. Many transit systems, including the MBTA and CTA, have “tap-on” fare collection systems whereby users present their fare payment upon entering the system, whether that is while boarding a bus or passing through fare gates at a rapid transit station. Transit agencies with “tap-on” fare collection systems can infer destinations for boarding records and match journeys to vehicle trips. There are many different approaches depending on data availability and desired application.

Destination inference is a well-known problem in transit planning and as such has been the subject of considerable research over decades. Many of the algorithms use trip chaining behavior to identify patterns [47, 48, 7]. Statistical inference is a popular approach for handling unlinked trips [49, 50, 51, 52]. A recent survey paper summarizes past research in this area and provides a comprehensive overview of ongoing challenges [53].

Alternative sources for origin-destination estimates are also an active field of research. Mobile phone records have been used for this purpose [54, 55], as have onboard surveys [56]. Additional studies have also leveraged Bluetooth and WiFi signals to trace passenger flows throughout a transit network [57, 58].

Some transit systems have implemented “tap-on, tap-off” fare collection where passengers also use their fare card when exiting the system, allowing for direct collection of origin-destination flows. Whether estimated or collected, origin-destination flows have been used in many transit planning many applications. These applications include network design and modification [59, 60], coordinating schedules to improve transfers [61, 62], and for equity analysis [63, 64]. Wang et al. [65] discusses how origin-destination flows can be used to generate passenger-based performance metrics, such as average transfer time.

1.5.4 Shared Bus Corridors

The shared bus corridors concept is an ongoing area of research that combines performance measurement and spatial analysis to identify how overlapping bus routes affect passenger decisions and service quality. Chriqui and Robillard [66] developed a method for modelling transit assignment when multiple routes are available to serve a passenger journey. Subsequent papers have adopted or expanded upon their method to solve transit network design problems [67], estimate passenger destinations from smart card data [68] and for real-time bus control [69]. Bie et al. [70] finds that explicitly considering route overlap in scheduling can help to reduce travel time and operating cost. There has been limited research on systematic aggregation of transit measures at the corridor level. Dimond et al. [71] describes a process for map matching to link bus routes to the road network, with some discussion of how ridership could be aggregated and visualized at the road segment level. It does not include any comparison between different routes or any TPMs beyond ridership. This thesis extends their concept to many different types of corridor-level TPMs, including derived TPMs that compare performance between routes in a corridor.

1.5.5 Longitudinal Analysis

Another dimension to transit performance analysis is comparison of performance and coverage over time. There are relatively limited comprehensive efforts to analyze the temporal evolution of bus transit service, despite a formal methodology proposed by the Transportation Research Board [72]. Long-term studies of the growth of transit networks often focus on rail systems, which typically grow more slowly than bus networks [32, 73, 74]. Some previous longitudinal evaluations of bus transit systems do exist, but are based on customer satisfaction surveys rather than network performance [75]. Other research has reviewed the aggregate changes in bus transit supply across different urban areas using network-wide measures such as revenue hours [76, 77].

The studies that have reviewed the performance of individual bus routes or entire networks before and after service changes are almost entirely focused on accessibility rather than performance [78, 79, 80, 81, 82]. One interesting paper used GIS-based analysis to evaluate the changes in accessibility for several different proposed bus network changes [83].

Some examples of longitudinal bus service studies do exist. One recent longitudinal study used GTFS to measure the quantity bus service supplied over the course of five years, which was then included as an explanatory variable in a model of bus ridership at the route level [84]. The authors do not discuss whether changing route or stop identifiers over the course of their study period was a concern. Addition of service was found to have a positive correlation with ridership. Another reviews the travel time changes that arise from a large scale network redesign of the bus network in Helsinki, Finland [85]. The authors note that one issue in comparing schedules before and after the network changes is the differing stop identifiers, which must be manually corrected.

The literature most similar to this research are three studies of transit network evolution in the Toronto area, which compute the changes in transit service in order to evaluate equity [86] or predict future demand [87, 88]. To the authors' knowledge, there has been no previous effort to develop a systematic method for evaluating temporal changes in transit performance measured at a disaggregate level.

1.5.6 Visualization for Transit Planning

Finally, each of the different components of transit performance measurement described above have a spatial context and can therefore benefit from visualization.

Recent studies have explored the benefits of using Geographic Information System (GIS) software for dynamic, interactive visualization of transit performance measures at the stop, route, and network level [12, 89, 90]. Visualization allows the user to integrate data from disparate sources and understand the spatial patterns, relationships and trends [91]. Data visualization serves to enhance communication during the problem solving process, and can also be used for exploratory analysis and development of hypotheses [92]. A 2007 report on the status of transit performance visualization at the transit agency in Portland, OR [91] argues that "efforts to incorporate new data visualization techniques will do much to assist with the identification of operational problems as well as provide insight into potential solutions."

Kurkcu et al. [93] develops an interactive performance analysis tool using AVL to visualize metrics at the stop-to-stop level. Stewart et al. [94] describes how automated data can be visualized at the stop-to-stop and route level to support service planning and communicate with the public. The authors argue that visualizing combined headways when multiple routes serve a single stop can support decision making by transit planners, although they use GTFS `stop_id` property to identify these shared stops, which inherently overlooks stations with multiple GTFS stop ids.

Other transit visualization tools include animated and real-time visualizations. One study used smart card data to map journeys throughout the public transit network [95], allowing users to identify patterns among the load profiles. Tools have also been developed to visualize the movements of vehicles throughout the transit network, for evaluation of operations and communication with the public [90]. One ambitious project has collected and actively displays real time vehicle positions for over 700 different transit agencies [89].

1.6 Thesis Organization

This section describes the contents of the remaining chapters of this thesis.

Chapter 2 presents a method for identifying shared corridors within a bus network, and for aggregating different classes of performance metrics across those corridors.

The "degree of overlap" is compared between the transit agencies of two cities with very different road topology: Chicago, IL and Boston, MA. Case studies are provided to demonstrate how these methods can be used for planning transit priority infrastructure, reducing effective headways for overlapping routes and allocating service between parallel routes to reduce crowding.

Chapter 3 presents a method for decomposing bus networks into a set of city block-level units of analysis. This is a natural extension of the framework Chapter 2, as it enables a simpler aggregation process. These segments are represented by their geography rather than agency identifiers, enabling longitudinal comparison even across periods where stops and routes are renamed. Two case studies are included. The first demonstrates how this method can be used to track changes in service over time for different neighborhoods in Boston. The second case study applies spatial decomposition and geographic representation to visualize changes in ridership across the MBTA network from 2013 to 2019 at the city block level.

Both Chapter 2 and Chapter 3 focus on analyses that can be conducted using widely available sources of transit data, including event-based AVL, APC and GTFS. Chapter 4 shows how the framework can be adapted for a richer data source, origin-destination flow estimates, in order to conduct more advanced analyses. The first case study demonstrates how origin-destination data can be integrated with the methods of previous chapters to identify routes with high percentages of transfer passengers. The second case study investigates spatial trends in passenger destinations before and after a significant service pattern change.

Chapter 5 offers conclusions, discusses the implications of this research and provides possible directions for future research in this area.

Chapter 2

Corridor Identification and Performance Aggregation

2.1 Introduction and Motivation

Understanding bus transit performance is critical for managing operations, planning future service changes and identifying locations for improvement. While much research has been devoted to measuring and benchmarking performance at the stop, route and network level, little attention has been paid to systematic quantification of the spatial and operational relationships between routes. Identifying corridors shared by multiple routes, aggregating performance measures across those corridors and visualizing the results can provide transit planners with a powerful tool for analyzing bus performance and planning infrastructure changes.

Transit performance measures (TPMs) are used by agencies "to provide quantitative information to operating personnel that can be used to diagnose and correct service problems and to improve overall performance" [11]. A transit network can be divided into multiple levels for performance analysis, each with a different resolution and unit of analysis. Such a structure ignores, however, the spatial relationship and overlap between different routes, which is critical for a holistic understanding of the system performance.

Overlapping routes are common in bus networks, such as bus networks that are designed to feed a rapid transit system where many bus routes may converge at a single multi-modal station. In addition, many bus networks include both express routes and local routes that serve the same corridor. An example of the many overlapping bus routes within the Massachusetts Bay Transportation Authority (MBTA) network is shown in Figure 2-1. In some cases, such as the Alameda corridor in Santiago, Chile and the South East Busway in Brisbane, Australia, corridors may be served by more than 20 different routes. Identification of routes that traverse a common street, when required, is typically undertaken by consulting route maps or based on the personal knowledge of the analyst. Systematic identification for the purpose of facilitating TPM calculation can be onerous when there are many overlapping routes within a bus network.

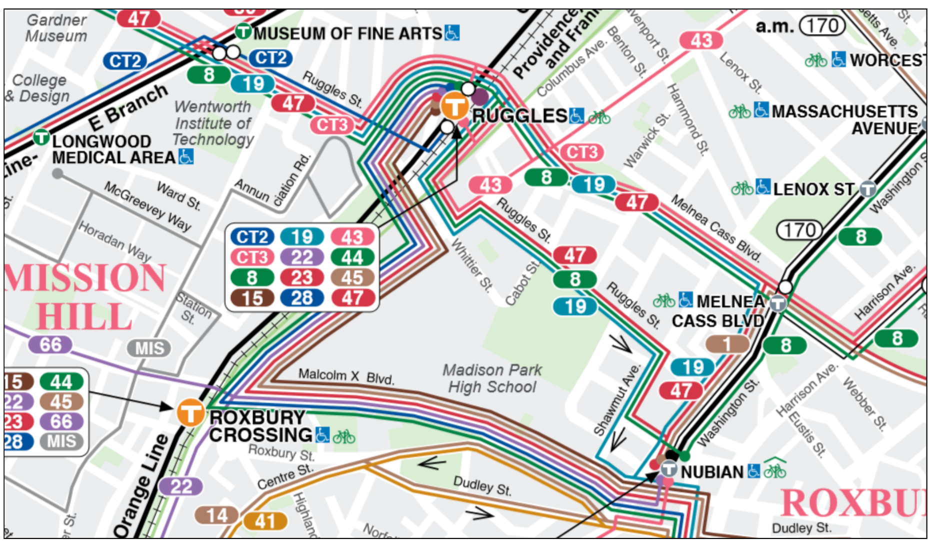

Figure 2-1: Overlapping bus routes near Ruggles Station in Boston, MA [1]

Aggregating TPMs by shared corridor enables three types of analysis that are not available when routes are considered independently:

1. **Infrastructure-level analysis**. By aggregating the TPMs of all routes that traverse each corridor (and therefore use the same lanes, traffic signals, etc.) it is easier to identify locations where infrastructure interventions such as transit priority systems would have the greatest impact.

2. **Journey-level analysis**. Corridor aggregation allows planners to compare the

performance of routes that serve the same origin-destination pairs and thus would be interchangeable to many passengers. For example, if three routes share a common trunk corridor, passengers whose journeys begin and end within the trunk corridor would be concerned with the combined headways of all three routes.

3. **Comparative analysis**. Derived TPMs can be used to illustrate differences in the performance of routes that share a common corridor. For example, visualizing a derived TPM that represents the difference in crowding between an express route and a local route may help to identify locations where service adjustments could improve the passenger experience.

Systematic identification and matching of routes within shared corridors are not trivial exercises. Many transit agencies do not maintain an inventory of the streets traversed by each bus route. The standard General Transit Feed Specification (GTFS) feed is not required to contain any information relating to the street network or the path that a bus follows between stops [96]. Combining routes that only share stops will exclude express routes that run parallel to local routes. In some cases, such as split near/far side stops or situations where multiple agencies serve the same region, comparing stop IDs is not sufficient to generate a list of all shared corridors in the network.

By combining shared corridor-level TPMs with visualization, service planners can quickly locate areas where transit performance is substandard. Past research has shown that visualization of TPMs allows planners to identify trends that might not be apparent when the TPMs are viewed in a tabular format [19]. The capacity for rapid analysis of performance enables "tactical transit" interventions, a recent trend in transit service improvement [97]. Tactical transit interventions involve inexpensive, short-term changes to improve the delivery of transit service as a method of testing and gaining support for more permanent strategies. TPM visualization can facilitate tactical transit modifications by locating the areas where such interventions would have the greatest impact. Visualization can also be used to highlight the spatial correlation of performance metrics and to combine transit performance data with other spatial data.

This chapter establishes a methodology for generating shared bus corridors by matching routes using a standardized GTFS feed. First, a map matching process for snapping bus routes to the road network is proposed. Then, spatial relationships between routes are used to identify shared corridors independent of stop or route identifiers. Performance measures are separated into three classes based on their input data and a procedure is proposed for aggregating each class of performance measures at the shared corridor level. Finally, three extended case studies using real transit data are used to illustrate how corridor-level analysis can be leveraged and combined with visualization to generate insights that are not available from typical performance measures. Each case study demonstrates how corridor-level analysis can be used to solve well-known transit planning issues: identifying opportunities for transit priority intervention, improving schedule coordination, and balancing resources between express and local routes serving the same corridor. This research addresses a gap in the state of the practice by systematically producing performance measures that account for the spatial relationships between routes without relying on user knowledge or conventions. The methods can be adopted by transit agencies immediately to facilitate long term infrastructure planning, optimize routes and schedules, and identify areas for improved performance.

Here the terminology used throughout the remainder of the chapter is defined. “Edge” refers to a street segment between two intersections, which is roughly analogous to one side of a city block. “Segment” refers to the section of an individual bus route between two consecutive stops. Segments are unique for each route. “Corridor” refers to the streets, or set of edges, that constitute the path of a segment. A corridor may be served by multiple routes.

2.2 Map Matching

The first step in developing corridor-level TPMs is to determine which routes share common edges. Map matching is used to identify the list of edges that is traversed by a bus when traveling between two consecutive stops. The minimum inputs required for map matching are the coordinates of each bus stop and a network file representing the road network. Many GTFS feeds include the optional `shapes.txt` file that contains coordinates describing the path of each route. These coordinates can be used to improve the accuracy of map matching by adding coordinates between bus stops. Note that these coordinates are sometimes derived from a GPS trace, which is subject to noise.

If the GTFS `shapes.txt` coordinates are used, additional processing is required to split the set of coordinates into stop-to-stop segments. The optional `shapes.txt` standard only requires the list of coordinates corresponding to a route, and a field indicated the order in which those coordinates are traversed. There is no required field to link a coordinate in the `shapes.txt` to the bus stops. Per GTFS best practices, shapes should follow road center-lines and not deviate to boarding locations, so there are typically no shape coordinates that match the bus stop coordinates specified in the `stops.txt` file [98]. We use a custom program to determine the Euclidean distance between each bus stop and each coordinate in the `shapes.txt` file for each route pattern. The coordinate that is closest to each stop is then designated as the "break point" at which the sequence of coordinates for the full route is divided. The subset of coordinates occurring between two bus stop break points is then used as the input coordinates for map matching the path between those stops.

Open source data downloaded from OpenStreetMap (OSM) [6] is used for the road network information. OSM provides network files for any location across the globe. The open-source Valhalla map matching service converts the input coordinates from GTFS into a set of edges [99]. Valhalla builds a graph from the OSM input data, where the edges of the graph represent the section of street between two intersections. Two-way streets are represented by two separate edges, one for each direction of travel.

Given a pair of input coordinates for routing, Valhalla identifies the set of nearby edges in the road network, then computes the shortest path between them, similar to the map matching procedure described in Li [36]. The output includes the set of edges traversed by the shortest path between the two input coordinates. The output also includes a list of coordinates that define the shortest path, which can be used as the input to a data visualization program, or to replace inaccurate GPS trace coordinates in a GTFS feed.

Valhalla allows the user to include a wide range of input parameters to ensure the desired result is returned. For the purpose of bus corridor matching, the bus mode is specified so that transit-only right-of-ways are included among the candidate edges. Other parameters, such as the search distance for nearby edges and the interpolation distance (a maximum coordinate pair distance threshold) can be tuned to suit the accuracy and resolution of the input coordinates.

Using the `shapes.txt` coordinates can also create issues due to errors in GPS traces. For example, assume that the reported position of a westbound vehicle near an intersection is 30 feet north of the actual position. The map matching algorithm might therefore erroneously impute that that the vehicle makes a brief detour by turning northbound at the intersection, as shown in Figure 2-2. These issues can be largely avoided through configuration of the map matching parameters.

Figure 2-2: Example of how GPS noise can produce a map matching error

The first solution to this issue is to include a `turn_penalty_factor` parameter in the Valhalla input. This parameter enforces a penalty for excessive turns in the shortest path algorithm, reducing the likelihood that routes with small detours are included in the result. It can be calibrated to correspond to the level of accuracy of the input coordinates. To improve the results even further, the entire bus route should be matched as a single continuous path, and then split into segments at the specified input coordinates that correspond to bus stops. This full route map matching method with break points combined with the turning penalty ensures that the result is a direct path without unnecessary detours.

2.3 Corridor Identification

Once the set of unique edge IDs is determined for each segment in the network, they can be compared to identify shared corridors. If two segments have no overlap, or partial overlap where neither segment is a subset of the other, they will not form a shared corridor. Otherwise, there are three possible corridor types when comparing one segment to another. Examples of each corridor type are shown in Figure 2-3.

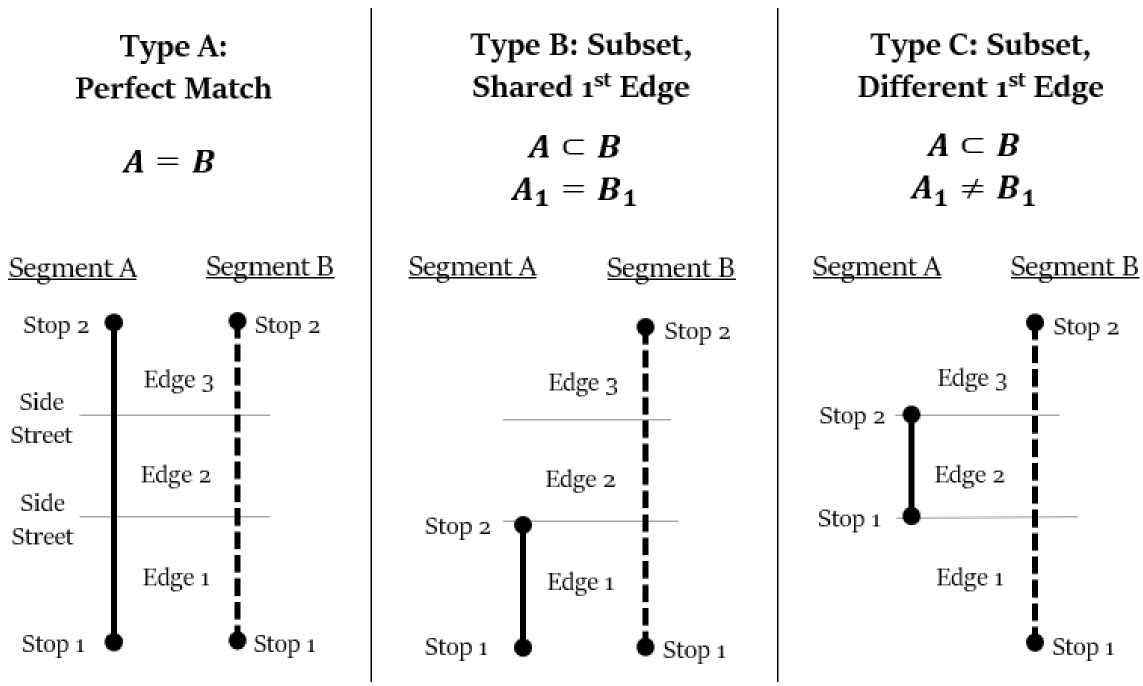

Figure 2-3: Corridor classification scenarios for different spatial relationships between routes

A Type A corridor is created when the same set of edges defines each segment and they share both stops. This is the most common type of corridor, as bus networks often feature stop pairs served by multiple bus routes. Corridors that are created when one segment is a subset of another (Type B and C in Figure 2-3) are critical for uncovering new insights that are not evident from segment level analysis. Type B corridors are those that overlap with a shared initial stop, while Type C corridors overlap with no shared stops. These types of corridors allows for the combined and comparative analysis of local-express route pairs, such as the corridor shown in Figure 2-5, that do not have an explicit relationship in the standard GTFS because they do not share any stops. Counting the number of corridors in a bus network can be used to quantify the degree of route overlap in the network design. The difference in overlap between the Chicago Transit Authority (CTA) and MBTA networks is discussed in the next section. We do not make any normative claims about transit route design with overlapping segments; it may be desirable to have considerable overlap based on the service area topology and demand profile.

Comparing the edges of each segment to identify corridors is almost always sufficient, however there are rare situations where errors may occur due to the relationship between edges and segments. Consider an arterial road with large gaps between intersections. If three or more bus stops are located between successive intersections, then there would necessarily be two consecutive segments of the same route located on a single edge. A real life example of this scenario in Lynn, MA is shown in Figure 2-4, where the Route 435 bus stops three times between intersections. There are two segments shown, one from Stop 7266 to Stop 7267, and Stop 7267 to Stop 7268. The map matching process will assign the same edge to both segments, and a direct comparison would determine that these segments are identical (Case #1). This issue is unlikely but does occur. For the full MBTA network, 3 such cases were identified out of more than 8,600 segments, while 8 cases were identified for the Chicago Transit Agency (CTA) network.

Another similar issue may arise if one segment is contained within a single edge. The preceding and succeeding segments will also traverse that edge, so it would appear that the middle segment is a subset of its adjacent segments. To resolve these issues, we apply a second condition on the corridor identification process: no two segments

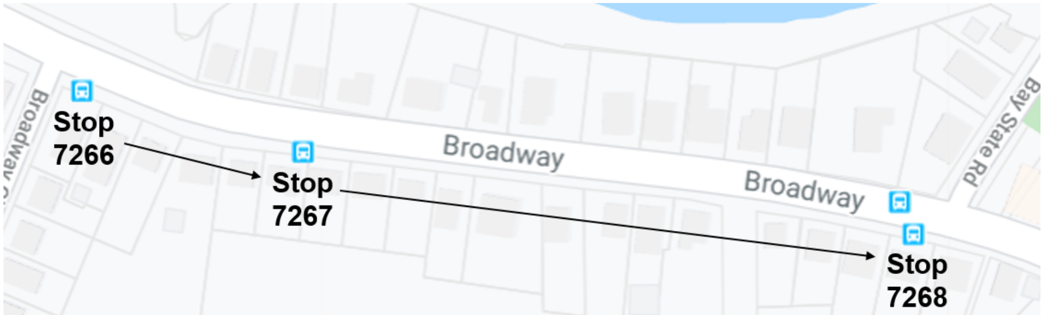

Figure 2-4: Two segments contained within a single edge

may form a corridor if they appear in the same route variant. It may be helpful to implement an initial screening condition that no two segments may form a corridor if they are consecutive (i.e. the first stop of one segment is equal to the last stop of the other, and vice versa). Because stops may have multiple `stop_id` values in the GTFS feed, it is also necessary to check that the "same route" condition is satisfied.

2.4 Performance Measure Aggregation

Careful consideration is needed for accurate aggregation of TPMs across shared corridors. For example, consider the corridor between the two local (Route 86) stops shown in Figure 2-5, an MBTA bus corridor in Somerville, MA. Note that the express route (Route CT2) does not serve either of the stops within the corridor, so any matching method that is based on stop IDs would not identify the overlap. In this case, it would be accurate to express the total passenger flow of the corridor as the sum of the passenger flows of all trips across both routes. Determining the average running time of the corridor, however, is not as straightforward if running time data is only available at the stop-to-stop resolution for the express route. The running time of the express route might be affected by conditions upstream or downstream of the corridor, so it would be incorrect to use the express route running time in the corridor running time calculation. To ensure that corridor-level TPMs are accurate, different rules must be used for aggregating (or not aggregating) different types of TPMs across corridors.

We propose sorting TPMs into three broad classes, each of which is treated ac-

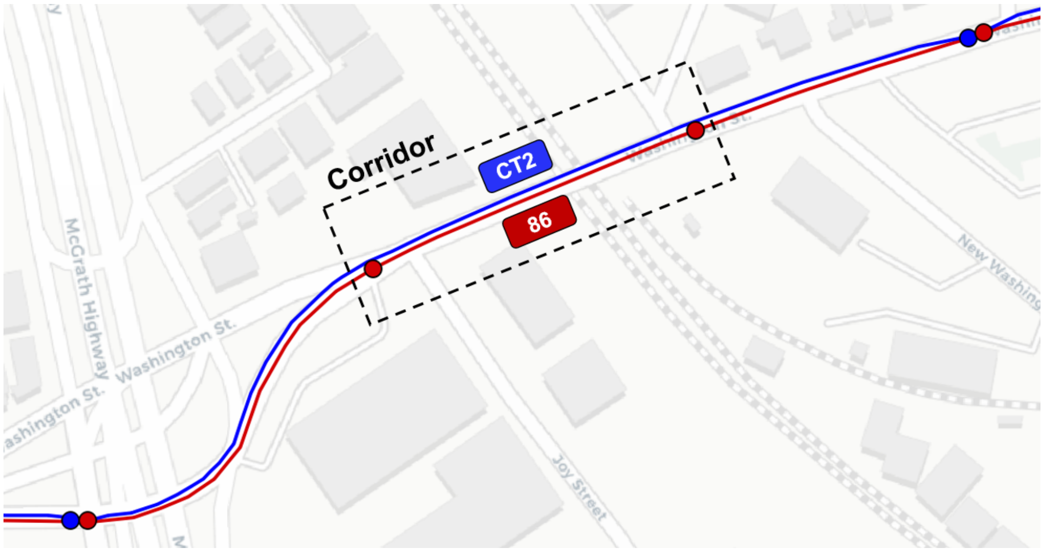

Figure 2-5: MBTA bus corridor in Somerville, MA served by express and local routes

cording to different logic during the aggregation step: “spatial”, “passenger” and “stop-based”. In general, this classification system and the assignment of TPMs to different classes is undertaken according to the principle that segments should be aggregated into corridors wherever possible to enable new analyses that are not available at the segment level. It is intended to generate the most granular corridor data possible while aggregating TPMs in a systematic and intuitive way.

- **Spatial** TPMs are those where the calculation is based on stop-to-stop distance (e.g., speed) or the metric is affected by traffic conditions between stops (e.g., arrival delay). These metrics should therefore only be aggregated across identical (Type A) segments.

- **Stop-based** TPMs are perceived at the stop-level, therefore they should be aggregated across segments that share a common first stop even if the segments are not identical (Type A and B). One example would be dwell time, which is aggregated across all routes serving a stop. Stop-based TPMs also include boardings, frequency/headway, expected wait time and excess wait time. The corridor frequency would include the combined frequency of all routes serving a stop.

Transit Performance Metric	Typical TPM Class	Corridor Type		
		A	B	C
Speed Running Time Delay Reliability factor	Spatial	X		
Dwell time Boardings, Alightings Frequency, Headway Waiting time On-time performance Headway regularity	Stop-based	X	X	
Passenger flow Crowding	Passenger	X	X	X

Table 2.1: Correspondence between corridor types and transit performance metric (TPM) classes

- **Passenger** TPMs are measures that are based on passenger flow and that are not affected by conditions between stops. These metrics can be aggregated across any corridor type. Passenger flow may be used as a passenger TPM for infrastructure-level analysis. Crowding may also be considered a passenger TPM for comparative analysis in some applications.

The different classes of TPMs are aggregated between the corridors depending on the corridor type. All TPMs can be aggregated across Type A corridors, since all routes share the same stops and traverse the same set of edges. Passenger and stop-based TPMs such as passenger flow and boardings can be aggregated across Type B corridors, because the segments overlap and share a common first stop. Only passenger TPMs such as passenger load can be aggregated across Type C corridors. These are typically local-express corridors where some local segments are entirely contained within an express segment and there are no shared stops (see Figure 2-5). The correspondence between TPM classes and corridor types is shown in Table 2.1, along with several examples of commonly used TPMs.

Each of the corridor types and TPM classes can be used in many different applications, including all three analysis types described in Section 2.1 (infrastructure-level, journey-level and comparative analysis). Practical examples are provided in Section 2.5. Case Study #1 uses Type A corridors along with spatial and passenger TPMs for an infrastructure-level analysis. Case Study #2 uses Type A and B corridors with passenger and stop-based TPMs for a journey-level analysis. Finally, Case Study #3 uses Type A, B and C corridors with passenger TPMs for a comparative analysis.

It is critical to note that these general categorizations are application-dependent; it may be necessary to shift TPMs between categories for certain use cases, or to apply additional conditions. For example, crowding might be used as a passenger TPM for comparing supply allocation between express and local routes serving the same corridor. Yet it would not be appropriate to compare crowding between a commuter shuttle and an urban route with one overlapping segment, because the service areas and demand profiles are much different. Data availability may also determine these categorizations; if high resolution vehicle position data are available, allowing interpolation between stops, then spatial metrics could be applied to Type B and C corridors. Professional judgement should be used to determine how TPMs are assigned for a given application and whether additional conditions are needed.

2.5 Case Studies

Three case studies are presented in the following subsections. Real operational data, schedules and route geography of the the CTA and MBTA bus networks are used. Data from non-holiday weekdays in October 2019 are used for the MBTA, while similar data from November 2019 are used for the CTA. These periods were chosen due to data availability, and to represent a typical month during the school year when transit systems are generally operating at peak demand. Each case study was developed through consultation with transit planning professionals to demonstrate the value of corridor analysis in solving problems that arise frequently in practice.

Before the case studies were conducted, the matching procedure was implemented for both networks in order to quantify the degree of overlap within each network.

Network	Type A Corridors	Type B Corridors	Type C Corridors
MBTA	37.6%	2.2%	3.0%
CTA	18.0%	3.2%	5.4%

Table 2.2: Percentage of all bus network segments included in each corridor type by agency

This type of analysis allows planners to measure the degree of route overlap within a network and compare it to other networks or alternative network designs. The regular MBTA bus network featured 13,564 route segments serving 8,461 unique stop pairs. 37.6% of route segments are therefore perfect matches with at least one other segment (Type A corridors). In comparison, only 18% of segments in the regular CTA bus network were combined into Type A corridors. Table 2.2 shows the percentage of segments that could be combined into the different corridor types for each network. Note that none of the Type B and Type C corridors would be identified by a stop ID-based matching procedure as they do not share both stops. Of the MBTA Type A corridors, 66 also would not have been identified if stop IDs were used, because the corridors have different terminal stop IDs at the same location (i.e. depots with multiple bays).

The difference in the degree of matching between networks is primarily caused by the difference in the street layouts between cities as well as the bus network design. Chicago's street network follows a rectilinear grid pattern so the bus routes are designed to operate on separate arterial roads that are parallel or orthogonal to one another. The Boston street network is much more circuitous and thus many routes travel along shared arterial roads. In the MBTA network, many of the matched corridors featured three or more segments, including a local segment in downtown Boston that was found to be a subset of six additional route segments. The CTA network, however has a greater proportion of subset matches (Type B and C corridors) than the MBTA network. This is a result of several local-express route pairs (e.g. Route 9 and X9, Route 49 and X49) in CTA network that travel the same path but stop at different intervals. The MBTA network, on the other hand, does not contain any local-express route pairs that run parallel to each other.

2.5.1 Transit Priority Infrastructure

This case study highlights the ability of corridor-level TPMs to represent bus performance at the infrastructure level. It uses Type A corridors and a combination of both passenger and spatial metrics to identify locations for bus queue jump lanes. The National Association of City Transportation Officials (NACTO) defines queue jump lanes as "... short dedicated transit facilities with either a leading bus interval or active signal priority to allow buses to easily enter traffic flow in a priority position" [100]. NACTO also states that "...queue jump treatments can reduce delay considerably, resulting in run-time savings and increased reliability." Queue jump lanes provide the same benefit to any bus route that passes through the intersection, so a corridor-level analysis is especially suitable for this application. This analysis is not specific to queue-jump lanes; it could be use to plan other bus priority infrastructure like transit signal priority or dedicated right-of-way.

Identification of priority locations for transit interventions involves two components: benefit and impact. Benefit refers to the actual performance improvement that can be expected as a result of the intervention. The benefit of installing a queue jump lane at an intersection with little congestion would be minimal. Impact is the number of passengers who would be affected by the intervention. Combining these two components allows analysts to identify locations where intervention would provide a significant improvement for a large number of riders. While reductions in running time have downstream benefits related to schedule efficiency, downstream reliability and increased supply capacity, here we will consider only the passenger time savings.

To combine the performance benefit and ridership impact, we will use a use a TPM derived from combining delays in bus travel times with the number of passengers experiencing those delays: passenger-weighted congestion delay (PWCD). Congestion delay is defined as the difference in running time between the minimum running time in the analysis period (assumed to be free flow speed) and the actual run time for each trip. AVL records identifying the arrival time at each stop are used to derive the segment running time. Dwell time is excluded from the speed calculation to remove any delays caused due to passenger loading and offloading. This delay is multiplied by the passenger flow, collected from APC data, to weight by the number of passengers affected. Finally, the result is normalized by the length of the segment to highlight locations where a small change could have a significant impact. The result is a PWCD TPM for each trip measured in passenger-minutes per mile. We use historical APC and AVL data for 22 weekdays in October 2019 to generate average daily PWCD results for each segment in the MBTA network.

This analysis benefits substantially from corridor-level aggregation of performance measures. Reviewing the passenger flow of the routes independent from one another does not provide a full account of the number of riders that stand to benefit from the intervention. Passenger flow, which falls into the passenger class of TPMs, is aggregated as long as one segment in the corridor is a subset of another. Corridor-level PWCD therefore provides an estimate of the congestion delay encountered by all passengers traveling through the corridor.

The difference between segment-level and corridor-level PWCD is illustrated using the MBTA example. Table 2.3 presents top five segments and corridors by PWCD for the average non-holiday weekday in October 2019. The ranking of corridors by PWCD includes two locations that were not among the top five segments in PWCD because the passenger flow is distributed among many routes. The two high-PWCD corridors are adjacent to a major transit hub, Nubian Station, which is a terminus station for many bus routes.

This analysis represents a first-order prioritization; other analysis would be needed to assess the feasibility of implementation. Excess running time compared to the free-flow condition is used in this case study to identify delay caused by congestion. Excess running time might occur as a result of conditions such as traffic incidents, that could not be improved by installing queue jump lanes. Additionally, queue-jump lanes may not be feasible at all intersections due to infrastructure limitations, fiscal considerations or political constraints.

Visualization of corridor-level PWCD allows analysts to make quick feasibility judgements by providing additional geographic context. Figure 2-6 shows corridor-

Segment	Route(s)	PWCD
BOS Terminal E to Congress @ World Trade Center	SL1	19,478
Third @ Chestnut to N. Washington @ Medford	111	17,750
N. Washington @ Medford to Beacon @ Broadway	111	17,340
Airport Station to Congress @ World Trade Center	SL3	15,189
Silver Line Way to BOS Terminal A	SL1	12,296
Corridor	**Route(s)**	**PWCD**
BOS Terminal E to Congress @ World Trade Center	SL1	19,478
Malcolm X @ Shawmut to Nubian Station	15, 23, 28, 44, 45, 66	18,729
Nubian Station to Malcolm X @ Shawmut	14, 15, 23, 28, 41, 44, 45, 56	18,281
Third @ Chestnut to N. Washington @ Medford	111	17,750
N. Washington @ Medford to Beacon @ Broadway	111	17,340

Table 2.3: Highest passenger-weighted congestion delay in the MBTA bus network, by segment and by corridor

level PWCD visualized across the MBTA network. Note that the color bins shown in Figure 2-6 (as well as Figures 2-8 and 2-9) represent the quintiles of the range. Corridors with the greatest PWCD are concentrated in high-ridership areas surrounding downtown Boston, primarily to the south and west in neighborhoods like Roxbury, Dorchester and the South End. The average daily PWCD per corridor across the full MBTA network is 1,098 passenger-minute per mile, but the average for corridors within the City of Boston is more than double at 2,430 passenger-minutes per mile. Along some corridors, each passenger experiences an average of more than 3.5 minutes of congestion delay per mile.

The visualization allows quick identification of locations where queue-jump lanes would be most impactful *and* feasible. These methods could also be used for other transit improvements such as transit priority implementation, which has been identified as a method of reducing bus delays in Boston [101]. In addition, visualization offers an advantage over tabulation because it provides the spatial relationship between different high-priority locations, allowing installation work to be planned efficiently.

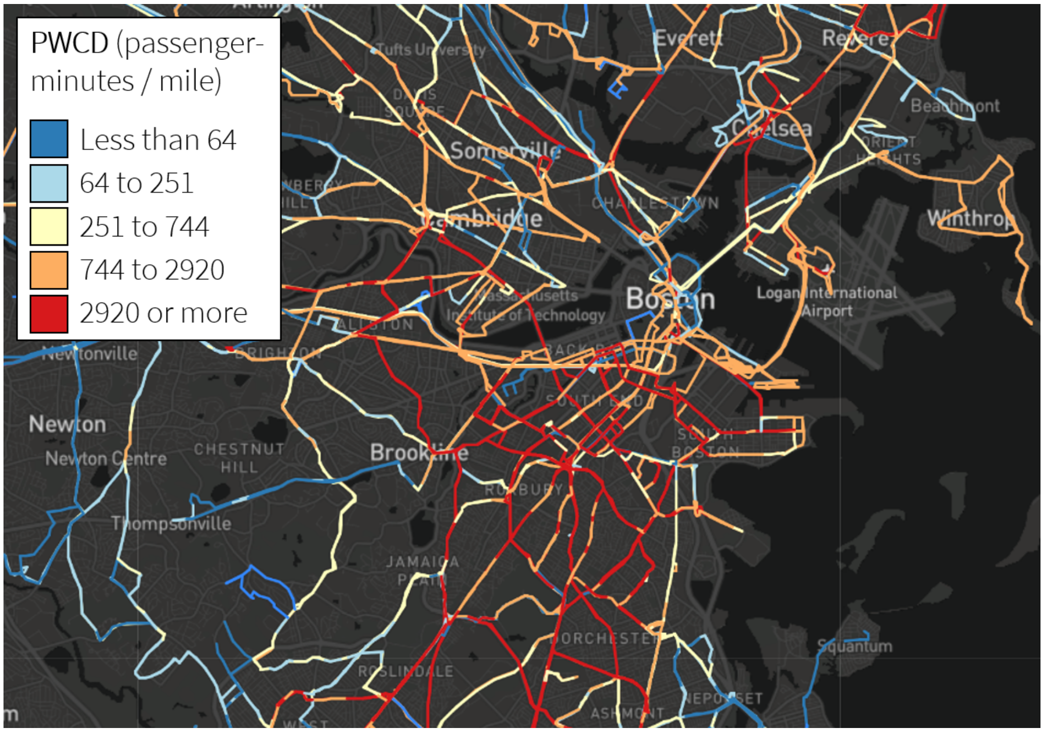

Figure 2-6: Daily average passenger-weighted congestion delay (PWCD) for MBTA Type A corridors in October 2019

2.5.2 Effective Headway Analysis

This case study demonstrates how corridor identification can be used in journey-level analysis. Type A and B corridors are identified and used as the basis for aggregating passenger and stop-based TPMs.

Some transit systems feature routes which serve many of the same stops. These routes may be functionally equivalent to passengers whose journeys start and end within the shared section of the routes. In that case, corridor-level TPMs are more representative of the rider experience than TPMs for individual routes. Coordinating schedules at shared stops may help to improve the passenger experience by reducing expected wait time [102]; however, it can be difficult or impossible to coordinate stop events in both directions and multiple route combinations. The corridor-level analysis can help planners identify the highest priority corridors for schedule coordination based on their potential for improving passenger wait time.

In this case study, we will use the example of the MBTA's 116 and 117 routes that serve communities to the north of downtown Boston. The routes share terminal stops, Wonderland Station and Maverick Station, as well as the majority of their intermediate stops. Inferred origin-destination pairs [103] for all weekday journeys in October 2019 reveal that 81% of Route 116 and 117 journeys could have been served by either route. Using corridor-level TPM aggregation, we can evaluate the combined performance of these two routes. The segments within the shared portion of the routes are matched for both passenger and stop-based TPMs using the methodology described in Section 2.3.

For this example, we will focus on the period between 6:30 AM and 8:00 AM as it is a period of high demand for bus service from the suburbs north of Boston into the downtown core. Scheduled inbound departures (towards downtown Boston) from the northern terminus, Wonderland Station, are shown in Figure 2-7 for the October 2019 weekday schedule. These scheduled departures are determined using the publicly-available GTFS feed. It is estimated (using origin-destination inference) that 53% of the 14,038 weekday journeys on Route 116 and Route 117 originating at Wonderland Station could be served by either route. The scheduled headways for each route are evenly spaced, but when the two routes are viewed as a single combined route, significant bunching is observed. The two routes are often scheduled to depart in close succession, followed by gaps as large as 16 minutes.

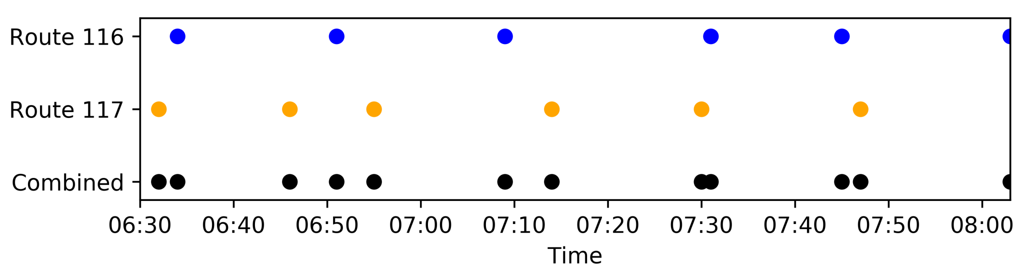

Figure 2-7: Scheduled weekday bus departures from Wonderland Station, October 2019

There are 11 departures after the first departure at 6:32 AM, with the last departing at 8:03 AM. Departures distributed evenly across this 91-minute period would produce 8.3-minute headways. Assuming that passenger arrivals are random, the average wait time for a passenger whose trip may be served by either route is then half of the headway, or 4.1 minutes [104]. The expected waiting time for the actual schedule, however, is 6.2 minutes (49% greater), due to the uneven departures. We will refer to the difference between the actual expected waiting time and the minimum possible expected waiting time as the "excess wait time" in this context. Excess wait time can be included as a stop-based TPM for systematic identification of corridors where schedule coordination between routes could reduce expected wait time. As in Case Study #1, excess wait time is combined with boarding information to generate a passenger-weighted TPM. APC records report 1,861 weekday boardings at Wonderland Station for these routes during October 2019, so the excess wait time for this example is $(6.2 \text{ min} - 4.1 \text{ min}) \times 1,861$ passengers $= 3,908$ passenger-minutes.

This is an example of a stop-based TPM, since it is derived from the departures of multiple routes serving the single stop. To extend this analysis from a single example to the entire network, we calculate excess wait time for all high frequency Type A and B corridors using the typical weekday schedule with arrivals between 6:30 AM and 8:00 AM. For these purposes, we consider high frequency segments to be any segments where the average combined headway of all buses serving the segment is 10 minutes or less within the study period. Spatial patterns can be identified by visualizing excess wait time for the MBTA network as shown in Figure 2-8. The 116-117 corridor is highlighted in the upper right portion of the map. Service planners could use such a map to review opportunities for coordinating schedules across routes in future schedule adjustments.

Not all corridors are comprised of routes that share a significant number of stops, so the corridor-level analysis is necessary but not sufficient for identifying opportunities to improve schedule coordination. Transit agencies with access to reliable origin-destination estimates or smart card data can combine this information with passenger-weighted excess wait time to create a new TPM that considers the proportion of passengers who stand to benefit from schedule coordination. The results are described above simply to illustrate the utility of excess wait time analyses regardless of the availability of journey origin-destination information.

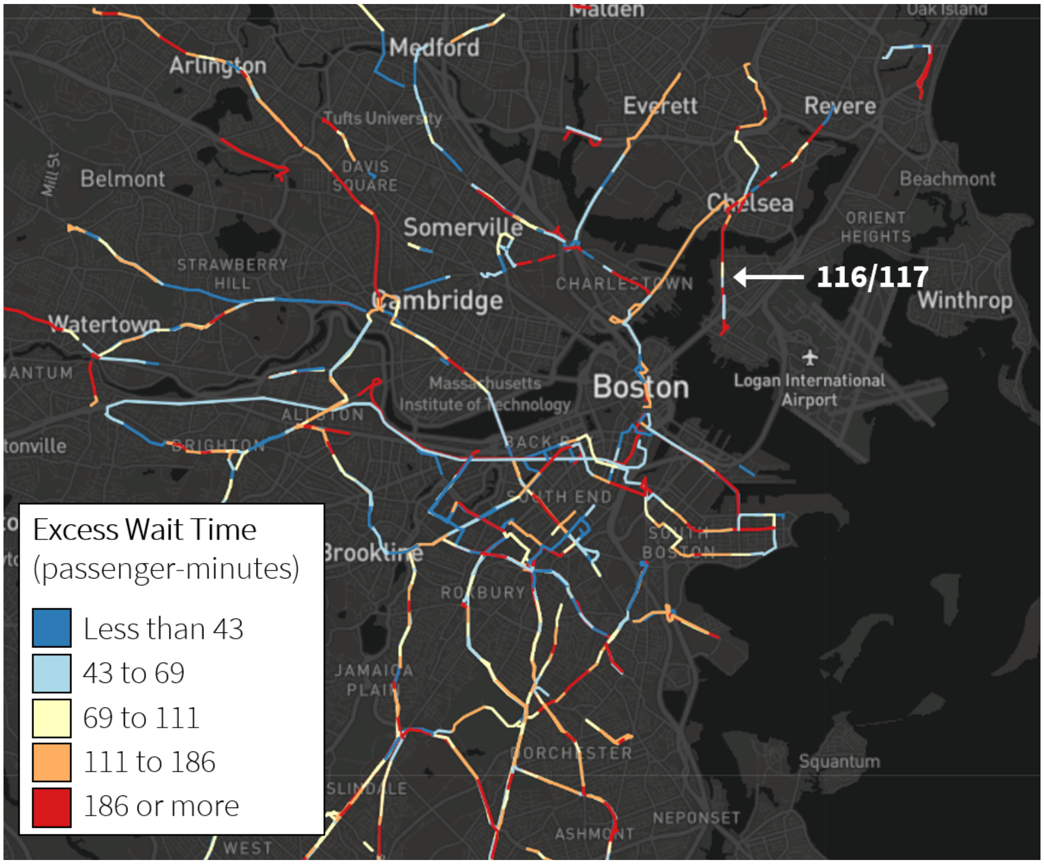

Figure 2-8: Sum of passenger-weighted excess wait time for the period of 6:30 AM - 8:00 AM across all weekdays in October 2019

Local transit planners are well aware that many passengers treat the 116 and 117 routes as interchangeable. The strength of the corridor-level analysis is that it does not rely on prior knowledge to identify locations where coordinated scheduling could improve the passenger experience. Furthermore, it automates the computation of the combined TPMs wherever routes are found to overlap, replacing ad hoc data queries and significantly reducing the overall effort.

This case study should not be interpreted as a criticism of current bus scheduling practices. The routes and time period used in this example were chosen in part for contrast between the actual and ideal schedule, but generally the schedules of these routes and others are coordinated wherever feasible. For example, the 116 and 117 scheduled outbound departures during the PM peak hour are evenly spaced and therefore have minimal excess wait time. Preparing bus schedules is an extremely complex process and there are many trade-offs involved. Our intent is simply to demonstrate the ability of corridor aggregation to facilitate new analyses from the passenger perspective and quickly identify locations where performance improvements could be targeted.

2.5.3 Passenger Crowding Comparison

The third benefit of corridor-level analysis is the ability to conduct comparative analysis of routes serving the same corridors. In this case study, we aggregate crowding, which is classified as a passenger TPM, across Type A, B and C corridors. Crowding has become an important concern during the COVID-19 outbreak, given that maintaining physical distance between passengers to limit spread of the virus is a priority for transit agencies. In this context, crowding is defined as the passenger load divided by the seated capacity of the bus. A comparison of bus crowding across several routes along the same corridor can be used by service planners to identify opportunities to re-allocate supply between different routes. It could also be used to inform route design changes.

The corridor-level "crowding variation" TPM derivation is straightforward: for each corridor where passenger TPMs of multiple segments are matched (Type A, B and C corridors), we take the difference between the minimum and maximum average crowding from the set of segments that make up the corridor across the same time period. APC data and a table of vehicle capacities are used to determine crowding for each segment. As an example, we use the 49 and X49 bus routes in Chicago. Route X49 is an express (limited stop) version of Route 49, both running along Ashland Avenue just west of downtown Chicago. During the PM peak hour (5:00 PM to 6:00 PM), both routes experience significant crowding in the northbound direction as shown in Figure 2-9.

While both routes experience higher crowding at the northern end of the route (consistent with commuter travel away from the central business district), crowding on the local route is highest. For eleven consecutive stop-to-stop segments, the average weekday crowding on Route 49 is above 75% of seated capacity with an average

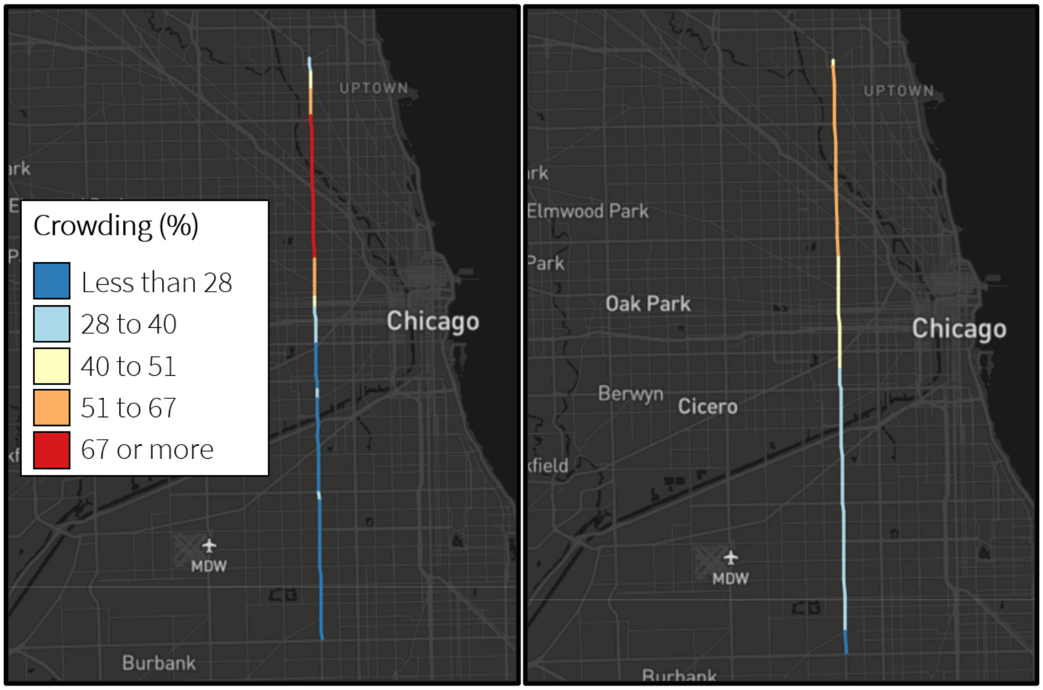

Figure 2-9: PM peak hour crowding for northbound Route 49 and Route X49 buses

crowding level of 85% per segment. The average crowding level on the express route while traveling across these same 11 segments is 62%. This result suggests that more supply could be allocated to the local route during the PM rush hour, or that the express route could make additional stops at local stops within the most crowded portion to attract some of the ridership away from the local route.

Crowding variation is calculated and visualized for all express-local corridors in the CTA bus network during the PM peak hour as shown in Figure 2-10 below. Express-local corridors were identified as any two separate routes that are matched for at least 75% of their segments. It is evident that while the NB 49/X49 corridor reaches above 30% crowding variation near the northern end of the route, several other routes have even higher crowding variation. Several corridors along Lakeshore Drive and others south and west of downtown feature routes with a difference in crowding levels greater than 40%.

As mentioned earlier, professional judgement is required in the interpretation and presentation of these results. It would not be appropriate to compare crowding between a limited-stop commuter shuttle and a local urban route even if they share

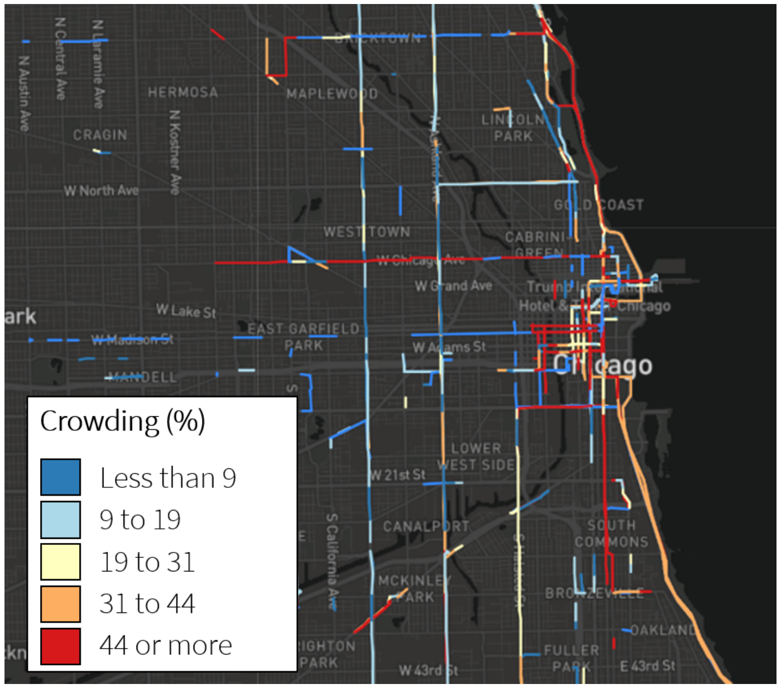

Figure 2-10: PM peak hour crowding variation for multi-route bus corridors in Chicago

many of the same segments, because the service areas and demand profiles are much different. Additional conditions can be applied to this analysis to filter out corridors where the allocation of supply between routes is not relevant.

This corridor-level comparative analysis is not limited to crowding. Other TPMs, such as reliability, dwell time or boardings could also be computed and visualized to locate disparities in performance between routes serving the same corridors. One of the challenges of improving transit service is disentangling the wide variety of conditions that can contribute to poor performance. Diagnosing performance issues by comparing routes within the same corridor allows planners to eliminate the conditions related to infrastructure, thus highlighting other contributing factors that may be easier to resolve.

2.6 Discussion

This chapter describes a systematic method for matching overlapping routes into corridors for different classes of transit performance measures. Typical transit performance measures are calculated at the resolution of the segment or route, which fail to consider the network topology and relationships between different routes. The corridor matching method enables transit planners to identify and prioritize opportunities for performance improvements that cannot be identified through segment, route and network-level analysis alone. Several case studies are described to demonstrate the potential for corridor analysis in planning infrastructure changes, schedule coordination and service allocation.

The corridor matching procedure is an example of how transit networks can be decoupled from their somewhat arbitrary route, route variant and stop designations in order to identify fundamental topological relationships. This approach also shifts the focus to performance measures that consider passenger travel patterns. As origin-destination data become more widely available, transit agencies will be able to better tailor service design to the needs of passengers.

There are limitations to this method. Steps are included to limit exposure to incorrect GTFS information, but it ultimately relies on reasonably accurate stop coordinates. The use of additional coordinates from the optional `shapes.txt` file is encouraged when possible to ensure that the map matching result traces the actual path of the bus route. If no additional shapes are available, there is no way to determine the actual path between two stops using GTFS, so the map matching process assumes that the shortest path is followed. In most cases this assumption is valid, but it may not hold for all paths [37]. Additionally, we have proposed using OSM as the source of road network information to ensure that the method remains open source and generalizable. Past research has found OSM to be reasonably accurate [105], but ultimately it is populated and maintained by volunteer contributors, so its accuracy cannot be guaranteed. Other sources of road network data may be substituted if available.

The case studies presented in this chapter are intended to be illustrative examples of the high-level analyses enabled by these methods. There are limitations and shortcomings to each. The case study presented in Section 2.5.1 involves ranking transit segments by delay in order to identify locations for queue jump lanes. It does not, however, differentiate between delay that could be ameliorated through queue jump lanes and delay that is a result of other problems, such as poor signal timing. Additionally, some locations may not be suitable for queue jump lanes due to infrastructure limitations. While this first order analysis can be done systematically using the methods described, additional investigation will be required before an installation plan can be developed.

The second case study, described in Section 2.5.2, finds opportunity to coordinate departures between overlapping routes. Schedules may have been designed for other reasons, however, that provide benefit to a greater number of passengers. For example, each of the overlapping routes could be timed to allow for passenger transfers at separate downstream transit hubs, making it difficult to coordinate schedules between the two routes. Planners should consider the downstream implications of coordinating departures before making any schedule changes.

The third case study shows how aggregation across shared corridors can help to balance service between express and local route pairs based on peak hour crowding. Looking at average crowding across the peak hour does not capture variations within the hour, where there can be significant difference. Furthermore, APC data is subject to error, especially during peak periods, as it relies on sensors that can become blocked when buses are crowded. The actual results may underestimate or overestimate peak crowding levels.

Chapter 3

High-Resolution Geographic Decomposition

3.1 Introduction and Motivation

Transit planners, analysts and advocates often seek to understand how the supply and quality of bus transit service changes over time. In most applications, performance metrics such as reliability and crowding are measured and reported at the route or stop level, and tabulated by route or stop identifiers. Bus networks are not stable over time, however; the pattern of stops that constitute a bus route are revised periodically and arbitrary stop numbers are sometimes reassigned to entirely different locations. As an example, consider the route changes implemented by the Massachusetts Bay Transportation Authority (MBTA) in December 2019.

This rather complex route redesign has varying impacts depending on the stop. For some locations, the service will not change, but the bus route number is different. In other cases, the service is eliminated entirely. Such changes make it difficult to conduct an accurate long-term comparison based only on route or stop identifiers. In order to review the evolution of bus service near Waltham Center, an agency analyst would have to have prior knowledge (or track down the information) that Route 70A was replaced by Route 61, and then compare the performance of both routes.

Tracking service changes has become especially important during the COVID-19

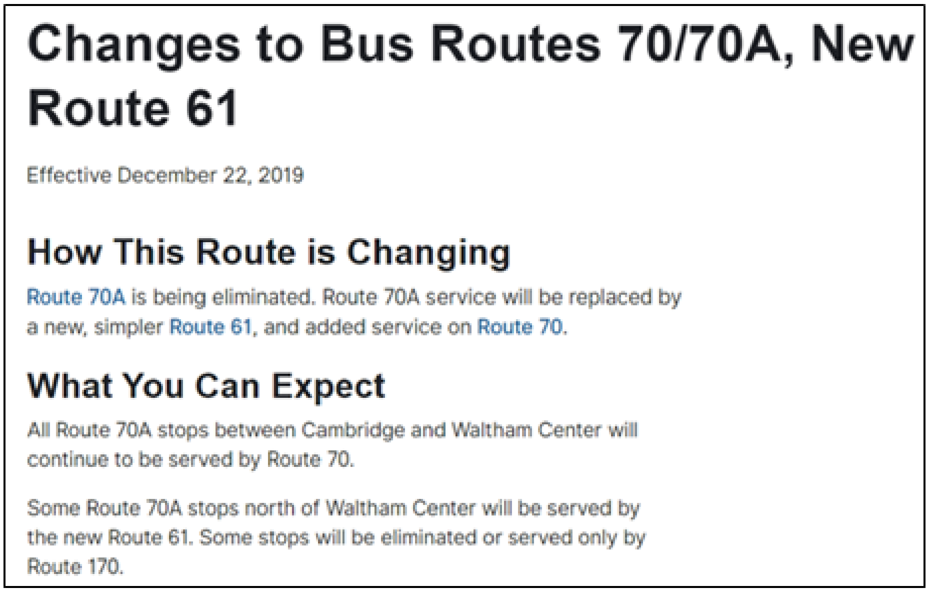

Figure 3-1: Massachusetts Bay Transportation Authority (MBTA) route changes in December 2019 [2]

pandemic, as many transit agencies have implemented substantial service reductions in response to the decline in ridership. Identifying these service changes systematically is critical to ensure accuracy, given that transit networks are often large and complex. Fortunately, many North American transit agencies maintains accurate records of transit service in a consistent format: the General Transit Feed Specification (GTFS).

A systematic method of comparing bus service and performance that is based in geography, rather than identification numbers, can be used for consistent representation over time. Location-based comparison also aligns planning more closely with the way that customers interact with the transit system. However, shifting towards a better geographic context for the transit network presents challenges in terms of scope and interpretation. If the city block is taken as the unit of measurement, there may be tens of thousands of individual data points for a large transit network. Spatial visualization can solve these issues, however, by representing large sets of disaggregate transit service data in a familiar setting.

In this chapter, a new method is proposed to decompose a bus transit network into city block length segments using a standard GTFS feed and open road network data.

These segments are identified based on geography rather than identifiers, allowing them to be compared over time. Performance metrics are then can be assigned to each segment to enable long-term comparison of service quality along each block. Two case studies of the Massachusetts Bay Transportation Authority (MBTA) bus network in Boston with 10 years of GTFS feeds illustrate the advantages of this method over identifier-based comparisons.

3.2 Block-Level Decomposition

To enable consistent representation of a bus transit network over time, a unit of analysis must be chosen. The goal is to select the largest possible unit of analysis for across the entire length of which the aggregate transit service is consistent. As an example, consider a segment of the network between two bus stops as the unit of analysis. This would not be appropriate, because two bus routes could run parallel to one another for the first half of the stop-to-stop segment, and then one bus could branch off at an intersection between the two stops. There would not be consistent transit service across the entire segment. Instead, the distance between two intersections is chosen as the unit of analysis, with additional boundaries at bus stop locations. Choosing this unit of analysis, which will now be referred to as an "edge", provides a stable unit of analysis across which bus service and performance can be measured.

The decomposition method requires two sets of input data. The first is the GTFS feed, which includes information relating to the transit network geometry and schedules. A standard GTFS feed contains the `stop_times.txt` table, from which all of the unique stop sequences (or "patterns") that constitute a bus trip can be extracted. Each pattern is defined as the sequence of stops for one bus route in one direction. In some cases, bus routes will have multiple variants under the same route ID. For example, a bus may deviate from the standard route by one stop to serve school trips during specific weekday hours. The MBTA has 172 bus routes and 833 unique patterns, meaning that on average there are 2.4 variants per route-direction. Each variant is retained as a separate pattern in the decomposition process.

The latitude and longitude coordinates of each stop are collected from the `stops.txt` table in the GTFS feed. Patterns are then converted to a sequence of stop coordinates that represent the path of the bus. The stop coordinates are used as inputs for the open-source Valhalla map matching library, which determines the path of the bus between stops within the road network. The output is the coordinates of a line representing that path. OpenStreetMap provides the second set of input data for the decomposition procedure, used as source of road network information for the map matching process. Additional discussion of map matching and some important considerations is included in Section 2.2.

The lines representing the route paths are then split at street intersections to create a series of discrete directional "edges" that are approximately equivalent to a city block. Finally, the lines are split again at mid-block bus stops, which allows us to differentiate performance metrics like crowding and passenger load before and after a stop event. A distance threshold can be used for determining which stops are mid-block. This additional step improves the visualization by avoiding the creation of very small edges between bus stops and nearby intersections. A threshold of 75 feet was found to produce reasonable results and is used to generate the visualizations shown in this chapter. An illustration of the edge decomposition is shown in Figure 3-2 below, where a bi-directional bus route is split into fourteen discrete edges. The diagram on the right shows how stops within 75 feet of an intersection can be ignored to avoid creating very small edges.

Figure 3-3 provides an example of the edges generated from the map matching process applied to the MBTA bus network just west of Nubian Square. Note that there are separate edges on either side of each road intersection and each bus stop. The MBTA operates one of the largest bus networks in the United States, with 172 routes traversing 39,544 unique edges.

A keen observer will note that the edges are also split at a parking lot driveway on the north side of Dudley Street, east of Dudley Place (this is represented in Figure 3-3 by a thin white line with no name or arrows). This raises the question of which street intersections should be considered significant enough to warrant splitting edges. The

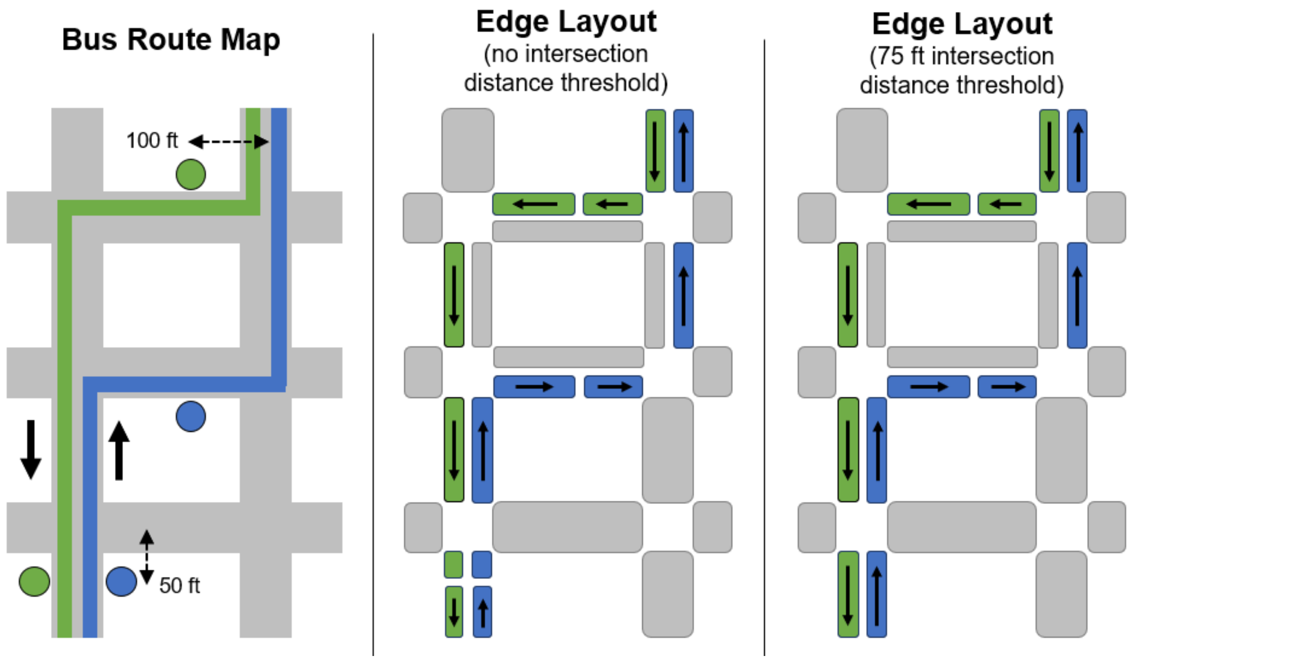

Figure 3-2: Example of bus network discretization into time-invariant geographic edges, with and without an intersection distance threshold applied

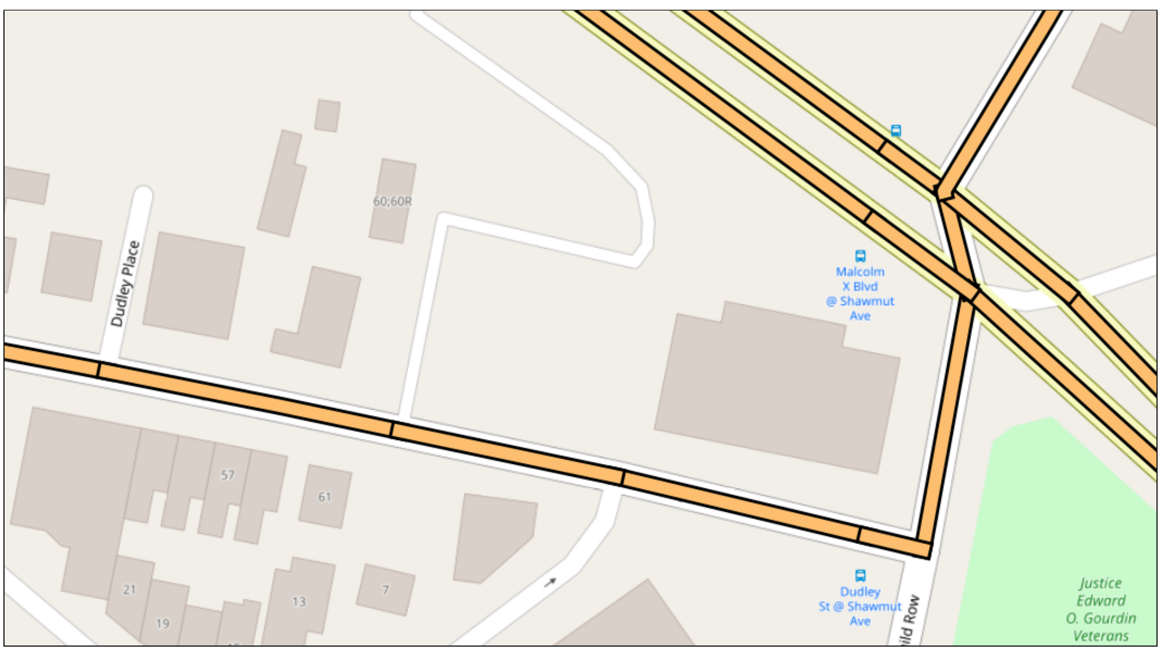

Figure 3-3: An example of transit network decomposition into block-length edges using map matching

overall purpose of decomposing the transit network is to create a unit of analysis that is small enough so that transit service is consistent across the full length of the edge. In this case, it is unlikely that transit service would stop or turn at the parking lot driveway, so it would reasonable to ignore the driveway and create one continuous edge. This can be accomplished by filtering out roads of type “service” from the OpenStreetMap network.

For each analysis period, the stop pair corresponding to each edge are stored to allow route or stop-level performance metrics to be assigned to the edges. In the case that multiple routes serve a single edge, the performance metrics are aggregated across routes to generate a edge-average or edge-total measure. Schedule-based performance metrics, such as scheduled frequency or span of service, can be derived directly from GTFS. Metrics based on passenger loading or observed arrivals would require other data sources such as APC and AVL. A thorough discussion of metric aggregation across overlapping bus segments is provided in Chapter 2.

One advantage of the edge-level decomposition is that it limits the complexity of aggregation across shared stop-to-stop corridors described in Chapter 2. Instead of three different corridor types, each with different processes for aggregation, there is only one type of corridor that can be created when comparing edges. Whenever multiple routes share an edge, they can be matched as a Type A corridor, because edges are defined such that transit service is consistent across the entire edge. Spatial and passenger metrics can be aggregated as desired by the analyst without concern about the type of the metric. Many edges will not start or end at a bus stop. For those edges, stop-based metrics such as boardings or dwell would not be defined. When shared edges do abut a bus stop, however, the stop-based metrics can be aggregated across all routes that traverse the edge.

3.3 Geographic Representation

The edges are more stable than route or stop IDs over time, and are therefore used as the basis for longitudinal comparison. The comparison could be conducted directly using geometry to check for equivalence, but such operations are typically too computationally intensive to be used for comparison at the network scale. To enable a more efficient method of comparison, the coordinates representing each edge can be encoded as a unique text string. For this application, the Google Encoded Polyline compression algorithm is used [106]. There can be some loss associated with the compression due to rounding after the fifth decimal place of the latitude and longitude decimal degrees, which creates a maximum error of approximately 0.5 meters or 1.5 feet. The potential loss is inconsequential for this application given that the edges are typically hundreds of feet in length. This text encoding step improves computation time for comparison between two time periods by two orders of magnitude.

3.4 Longitudinal Comparison

To match edges across two time periods, the first step is to screen for exact matches between the encoded test strings described in the previous section. This step generally identifies most of the common edges between periods. If the comparison is being conducted for the same service area across two time periods, the vast majority of edges in both periods will have a match. The rate of matching depends on the number of changes to the bus network between the two comparison periods. An edge match indicates that a block with bus transit service in the first period continues to be served in the second period. If an edge from the first period does not have a match, that indicates that the edge was served and is no longer served (i.e. cancelled service). If an edge from the second period does not have a match, that indicates that the edge has new service that did not exist during the first period. To provide a sense of the matching rate, a comparison of the MBTA bus service between January 2011 and January 2021 found that 86% (31,230 / 36,296 total) of the edges in the January 2011 network were also included in the 2021 network. The remaining 14% (5,066 edges) were no longer being served, but an additional 8,314 new edges were included in the January 2021 network. In many cases the matching rate will be higher, as this comparison covers a significant period of time during which there were several network changes.

In some cases, the encoded polylines of two segments may not match exactly, but the segments could be considered equivalent for this analysis. One example is a segment that begins at a bus depot with multiple bays. Another scenario where this issue could arise is when an agency revises the coordinates of a bus stop in GTFS between successive feeds for greater accuracy. This is common, especially when comparing early GTFS feeds to more recent ones, as GPS sensors and tools for automating GTFS feed creation have become more sophisticated. Even if the position was revised by only a few feet, the encoded polylines representing the edges on either side of the stop would be different between the two feeds. A third and final scenario can occur when multiple agencies serve the same physical bus stop. If the agencies use different methods for generating GTFS feeds, they may have minor differences in the reported coordinates for the same bus stop. In each of these scenarios, the encoded polylines do not match, yet the edges should be matched for comparison.

An efficient method for comparing edges geographically is therefore needed. Assume there are two periods with different service patterns, Period A and Period B. When comparing edges from Period A to Period B, any edges from Period A that are not matched during the first screening are stored in a new list. As noted above, this is a typically a small fraction of the total number of edges. Then we check for spatial intersections between the unmatched Period A edges and the full set of Period B edges. The intersection operation is relatively efficient; checking for intersection between one edge and full set of edges in the January 2021 MBTA network takes less than 0.2 seconds. To confirm a match based on the intersection of two segments, it is necessary to then check that the edges have the same direction of travel, as it is possible that a segment could intersect another segment traveling in the opposite direction. The direction of travel is determined by computing the angle of a vector between the first and last coordinates of the segment. If the two directions are within a certain tolerance, the segments can be considered a match.

There is one more important consideration specific to longitudinal comparison. Over time, new bus stops may be introduced and old ones may be removed. This has the effect of creating two completely different edges along the same block for the two comparison periods. As an example, consider a block that does not have a mid-block bus stop during the baseline period. The block would be decomposed into a single edge for that network. If a new mid-block bus stop is introduced in the comparison period, then the block would be split into two separate edges for that network. To avoid this scenario, the individual edges for each network should be created by splitting at any mid-block bus stop that exists in *either* bus network.

3.5 Case Study: MBTA Network Evolution, 2011 to 2021

The MBTA maintains a public archive of GTFS feeds dating back to 2009 [107]. This period covers several sweeping changes to the bus network, including the Bus Network Redesign that began in 2018, and the significant reduction in bus service that began during the summer of 2020 in response to low ridership during the COVID-19 pandemic. These feeds are used to visualize service changes over the course of a decade, even under changing route and stop identifiers. This case study demonstrates how the geography-based representation, rather than identifier-based representation, enables a simple yet thorough comparison of transit service across long periods of time. Such a comparison could be used by transit planners or advocacy groups to identify spatial disparities in transit service provision over time. The segments displayed in Figure 3-3 are used as the unit of analysis for comparing the change in transit service from 2011 to 2021.

Figure 3-4 shows how service has evolved by highlighting segments that were served in 2011 but no longer served in 2021, regardless of route ID. This includes elimination of several routes in the northwestern part of the region, as well as consolidation of routes along fewer corridors near downtown Boston. One interesting piece of Boston history can be observed, although it almost appears to be a mapping error: the elimination of bus service on Long Island in Boston Harbor, which is currently abandoned and is not connected to the mainland. There was, however, a bridge to Long Island until 2014, across which several MBTA bus routes traveled in order to provide access to social services located on the island. Other changes can also be observed, such as discontinuation of Route 90 service to Wellington Station after the opening of the Assembly Square Station on the Orange Line. Several new routes can be easily identified, such as Route 714 serving the Hull / Hingham peninsula in the southeastern part of the region. A comprehensive list of all MBTA transit network changes over many decades is compiled by Belcher [108].

The longitudinal comparison enables further analysis beyond simple binary indicators of service. Figures 3-5 and 3-6 show how the average number of inbound and outbound weekday trips have changed at the edge resolution from January 2011 to January 2021 across the entire MBTA network. The trends for each direction are fairly similar, as service changes are often applied to both directions of a route, but there are some significant differences. Fewer inbound trips from the communities west of downtown Boston have been eliminated than outbound trips serving the same area. Part of this difference is due to route changes presented in Fig. 3-1. Route 70A was cancelled in both directions, while the new Route 61 only replaced inbound trips for certain parts of the former Route 70A service area. Figure 3-7 shows the route maps for Route 70A and Route 61. At the network-scale, it can be observed that trips have been added within some inner ring suburbs like Chelsea, Medford and Somerville, but reduced in more distant suburbs to the west, northwest and southeast.

While the network-scale results are informative, the strengths of this decomposition and geographic representation are more evident at the neighborhood level. One way that the geographic representation can be useful is that it contextualizes service pattern changes that do not involve a route identifier change. Consider the case of MBTA Route 52, a north-south bus route that connected Watertown and Dedham via Newton. As part of the MBTA's "Better Bus Project", a full bus network redesign that began in 2018, Route 52 was consolidated along a single path and service was eliminated for a significant part of the original coverage area. The changes to Route 52 were made in order to improve the match between supply and demand. Com-

Figure 3-4: Locations where bus service was added, removed, or continued between January 2011 and January 2021 for the MBTA Bus Network

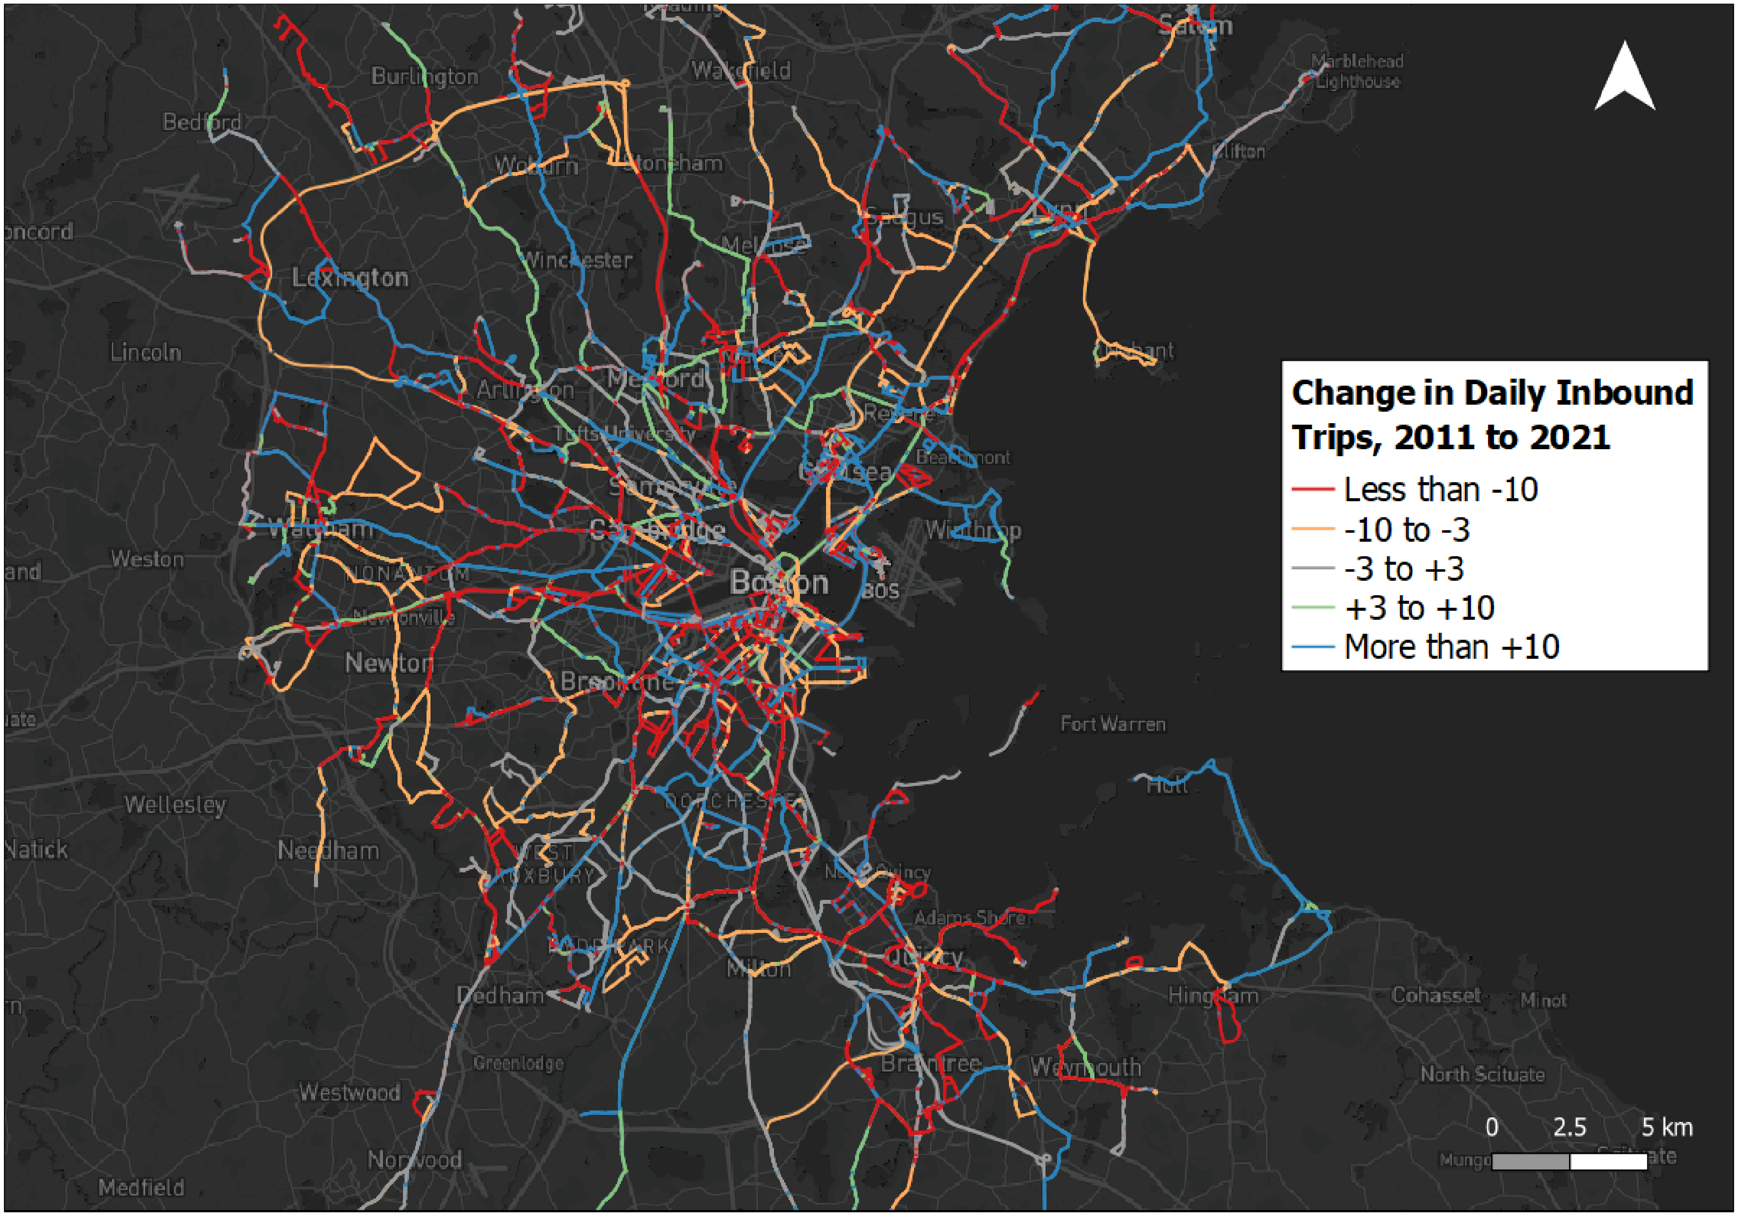

Figure 3-5: Change in average daily inbound weekday trips by edge, January 2011 to January 2021 across the MBTA bus network

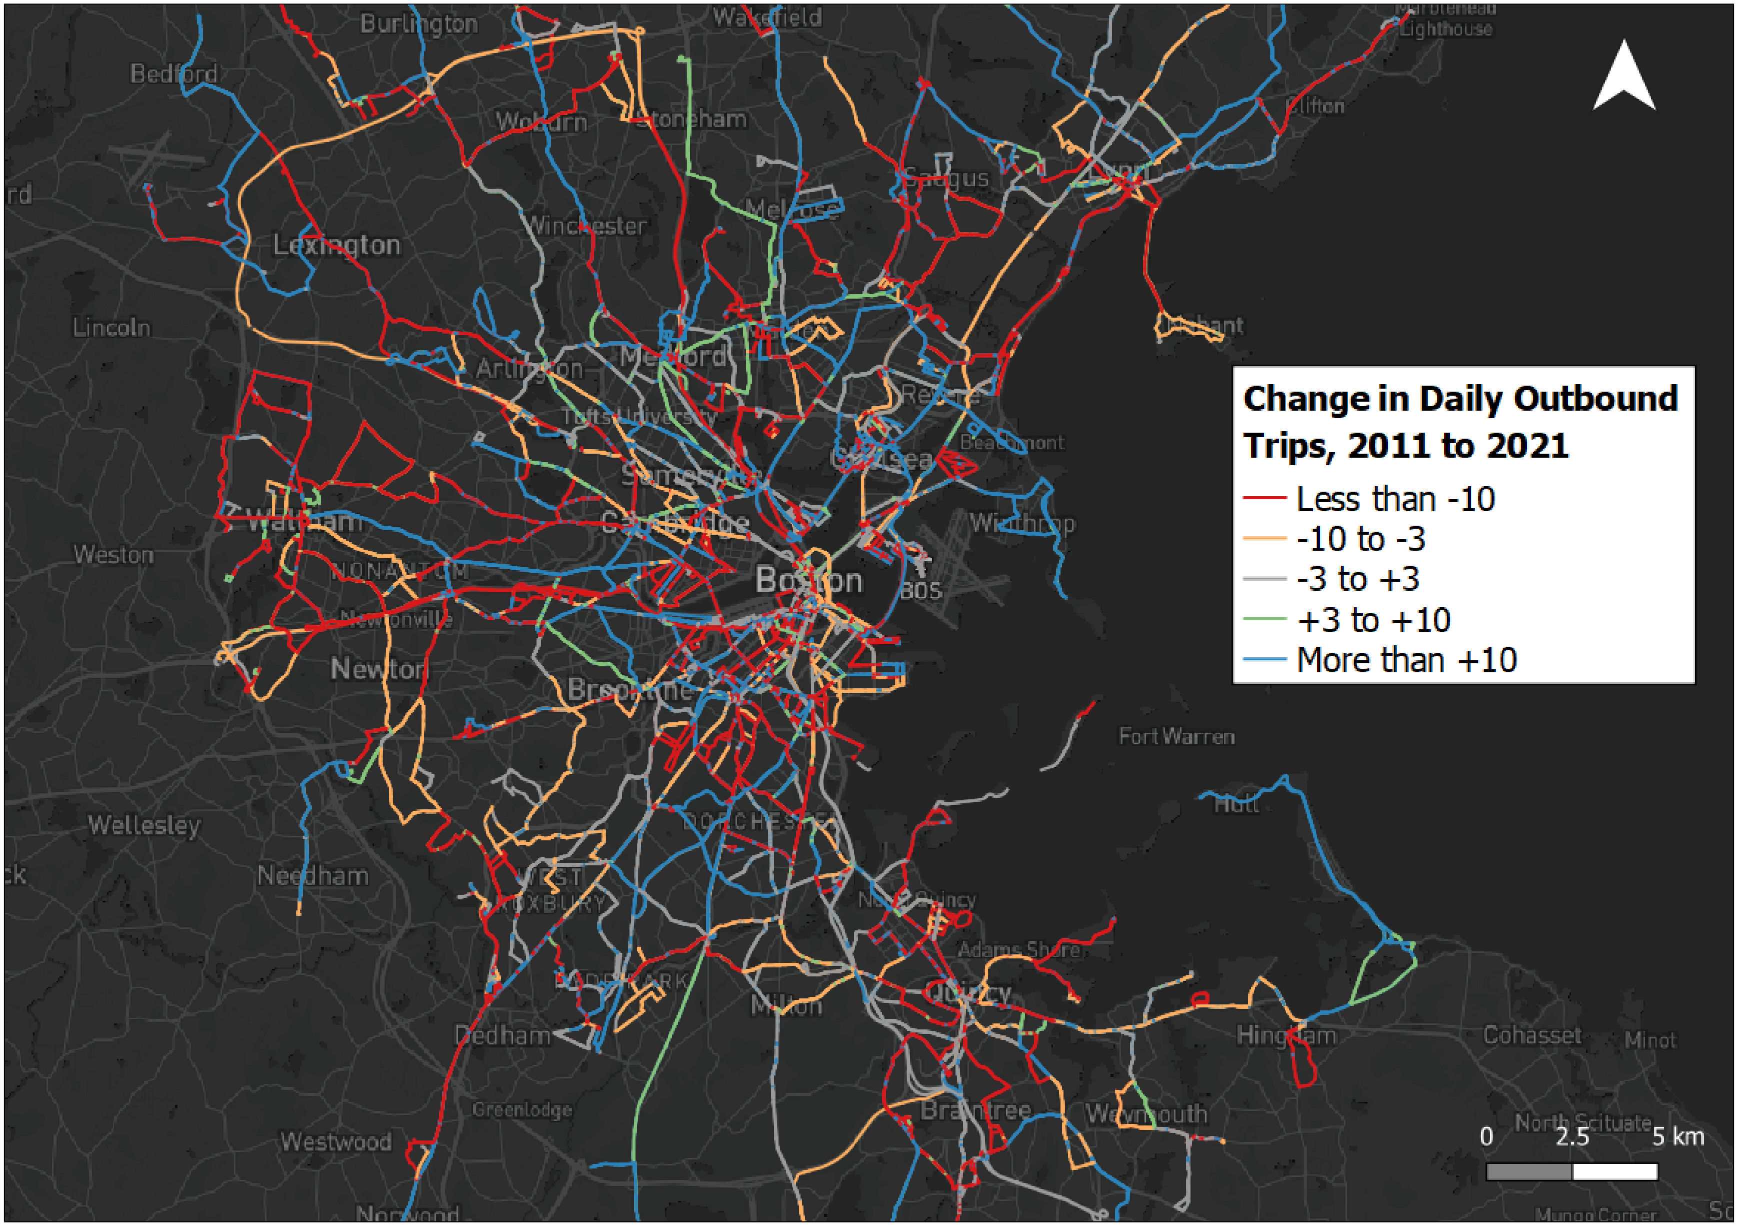

Figure 3-6: Change in average daily outbound weekday trips by edge, January 2011 to January 2021 across the MBTA bus network

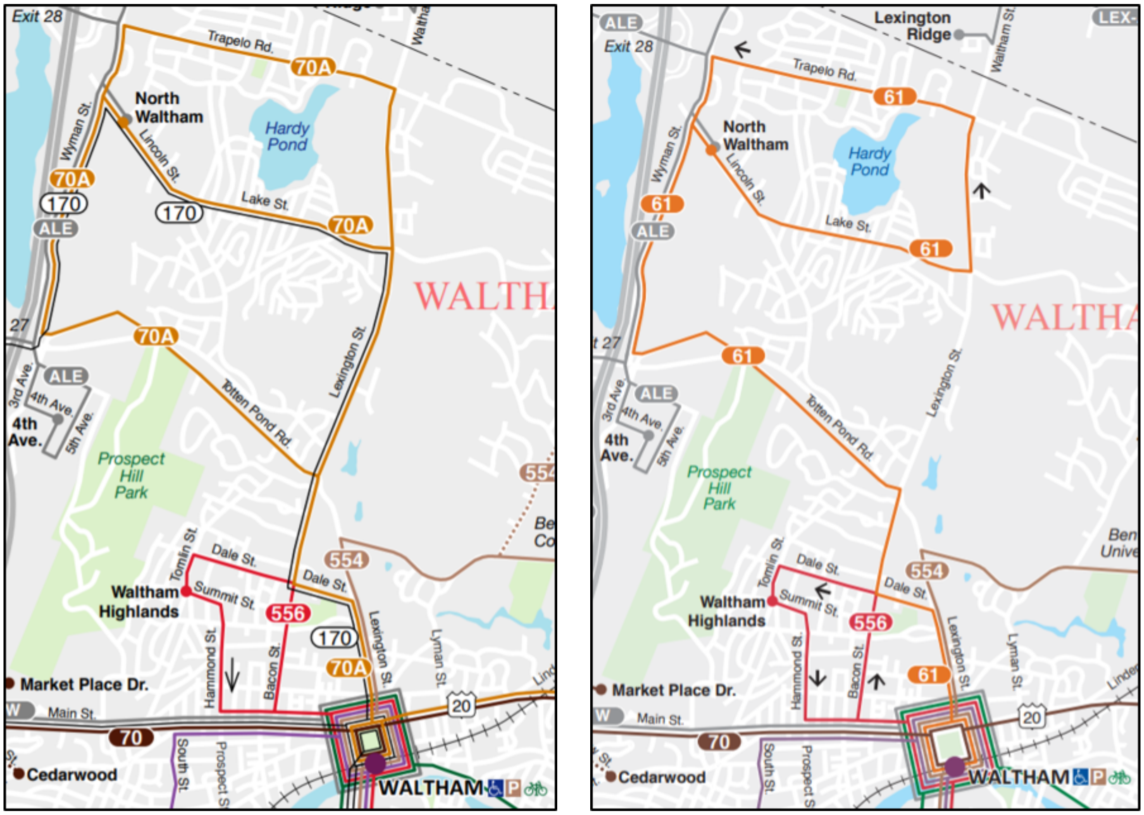

(a) Route 70A Service Pattern

(b) Route 61 Service Pattern

Figure 3-7: Comparison of North Waltham bus routes before (a) and after (b) the Bus Network Redesign in December 2019

paring the performance of Route 52 in terms of frequency, delay, and so on before and after the service pattern changes ignores the fact the Route 52 serves fewer bus stops. The methods described in this chapter, however, provide a clear and holistic visualization of the outcome of the network changes: a decrease in frequency along the blocks where service was cancelled, and an increase in frequency where service was retained. Figure 3-8 presents the changes in frequency along Route 52 before and after the implementation of the bus network redesign. All highlighted sections of Route 52 had fewer trips in 2021 than in 2011, except for a short section where trip counts have remained relatively static. The edges with cancelled service are easily identified as they have lost more trips than the edges where service was retained.

These changes to Route 52 are not being criticized; travel demand patterns change over time and bus networks should be updated periodically in response. This example

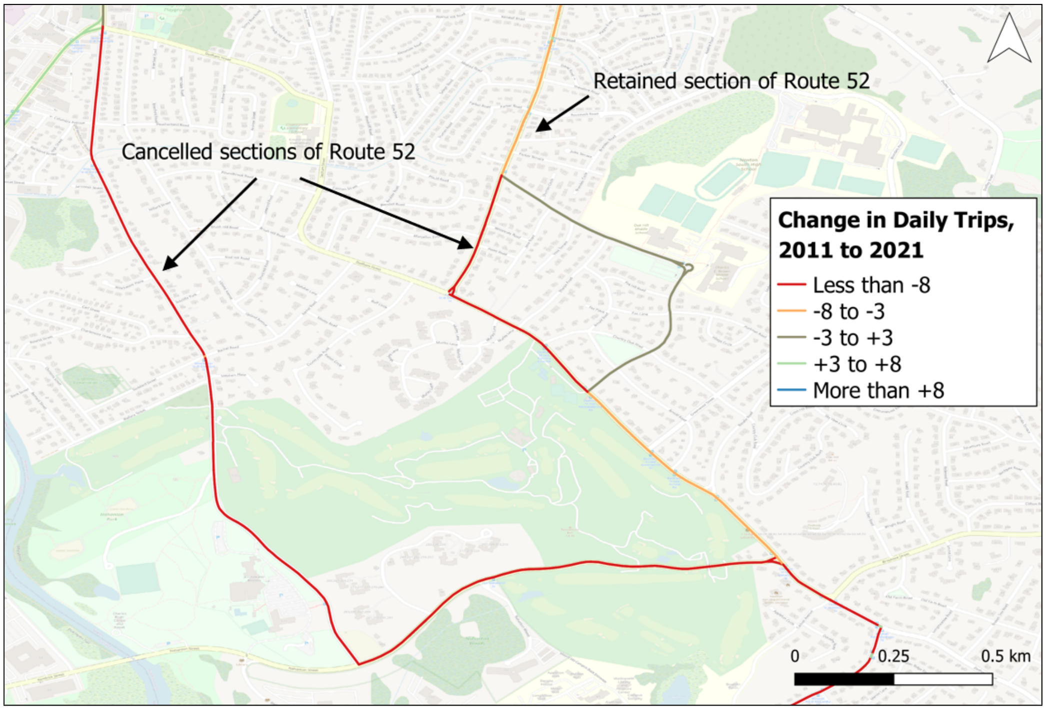

Figure 3-8: Route 52 change in average daily scheduled weekday trips by segment, January 2011 to January 2021

simply demonstrates how performance measurement and analysis at the route level does not capture the reality that routes change over time. A systematic method for computing and visualizing these changes in performance based on geography, using only GTFS feeds, is a significant improvement over existing analysis tools.

One additional application of this method is in the communication of service changes, even temporary detours, to the public. The current practice at the MBTA is to present service changes as a set of stop markers on a web map. In many cases service changes are not shown visually at all. A visualization that shows how the route patterns change at the block resolution, such as Figure 3-8, is more intuitive than showing how each route would change, and would help the public to understand the implications of service changes in their neighborhoods.

3.6 Case Study: Mapping Ridership Trends

Leveraging the block-level decomposition and geographic representation for analysing longitudinal service changes, as shown in the previous section, is a powerful tool because it relies only on (typically) publicly available data sources. However, if additional data is available, these techniques can be used to produce further insights regarding bus performance and ridership in addition to service delivery. In this case study, APC data is combined with the block level decomposition to map ridership changes at the edge level across the MBTA network for a period of several years. This type of analysis and visualization could be used by transit planners to identify medium- and long-term spatial trends in ridership in order to estimate future growth or decline. Much like the previous case study, the advantages of spatial decomposition and geographic representation over existing techniques are that any changes to route or stop IDs in the 6 year period are automatically incorporated into the results.

The earliest APC data available for this case study dates back to January 2013. The intent of this case study is to review the general trends in ridership, so a pre-COVID-19 comparison period, January 2019, was chosen. The ridership changes shown in this case study are therefore the differences across the six year period from January 2013 to January 2019. The APC coverage, or the number of trips with functioning APC devices, has changed over time. For this reason, the average ridership per weekday trip recorded by APC is used as opposed to the total ridership, which is not available from existing data sources. This case study reviews the average ridership per trip for the AM peak period, defined in MBTA's Service Delivery Policy as 7:00 AM - 9:00 AM, which is typically the time period during which ridership is highest. As in the previous case study, ridership trends are reviewed for both directions, inbound and outbound. The inbound direction is busiest during the AM peak period on a typical weekday, as more people travel towards the central business districts.

The results are shown for the full MBTA network in Figures 3-9 and 3-10. Many edges have not seen a significant change in ridership over the course of the 6 year study period. This is fairly consistent with overall ridership trends; the MBTA reported a

2.6% decline in total unlinked bus trips from 2013 to 2019 [109].

Generally speaking, there has been some decline in occupancy during both periods in the western portion of the service area (i.e. Lexington, Waltham and Newton). The corridor served by Route 70 and 70A between Cambridge and Waltham stands out as having a larger than average occupancy decline, especially in the outbound direction (between 10 and 12 passengers per trip depending on the segment). Schedules were very similar between the two periods, so this occupancy decline appears to represent a total ridership decline over time, consistent with the MBTA's public ridership counts [110]. Ridership around the City of Lynn, northeast of downtown Boston, has also declined. Both of these areas experienced decreases in the number of daily trips as shown in the previous case study, although the study periods are somewhat different. Interestingly, many of the routes in or immediately adjacent to the Central Business District have gained ridership. Neighborhoods such as Charlestown, South Boston, the South End and East Boston all have many edges which have seen an increase of more than 3 riders per trip. Note that the bus network is only a one component of the transit network; some of these changes may be a result of transit riders switching to or from the rail network.

Leveraging the high-resolution spatial representation and geographic representation, it is possible to directly observe changes in passenger load patterns due to service changes at the edge level. A new high-capacity bus route, the SL3, was introduced in 2018, connecting Chelsea to downtown Boston. For a very short part of its route in East Boston, it runs parallel to Route 112. Along this part of its route, the SL3 averages a passenger load of 12.5 passengers per inbound weekday AM peak trip. Figures 3-11a and 3-11b show the change in average weekday passenger load during the AM peak from January 2013 to January 2019 for Routes 112 and SL3 separately. Looking at the changes passenger load for Route 112 alone might give the impression that corridor ridership has declined. Figure 3-11 shows the combined changes in passenger load in the area where the SL3 and 112 converge. The edge-level spatial resolution enables differentiation of passenger load patterns before and after the two routes merge, even though the merge points do not coincide with a bus stop. Stop-to-

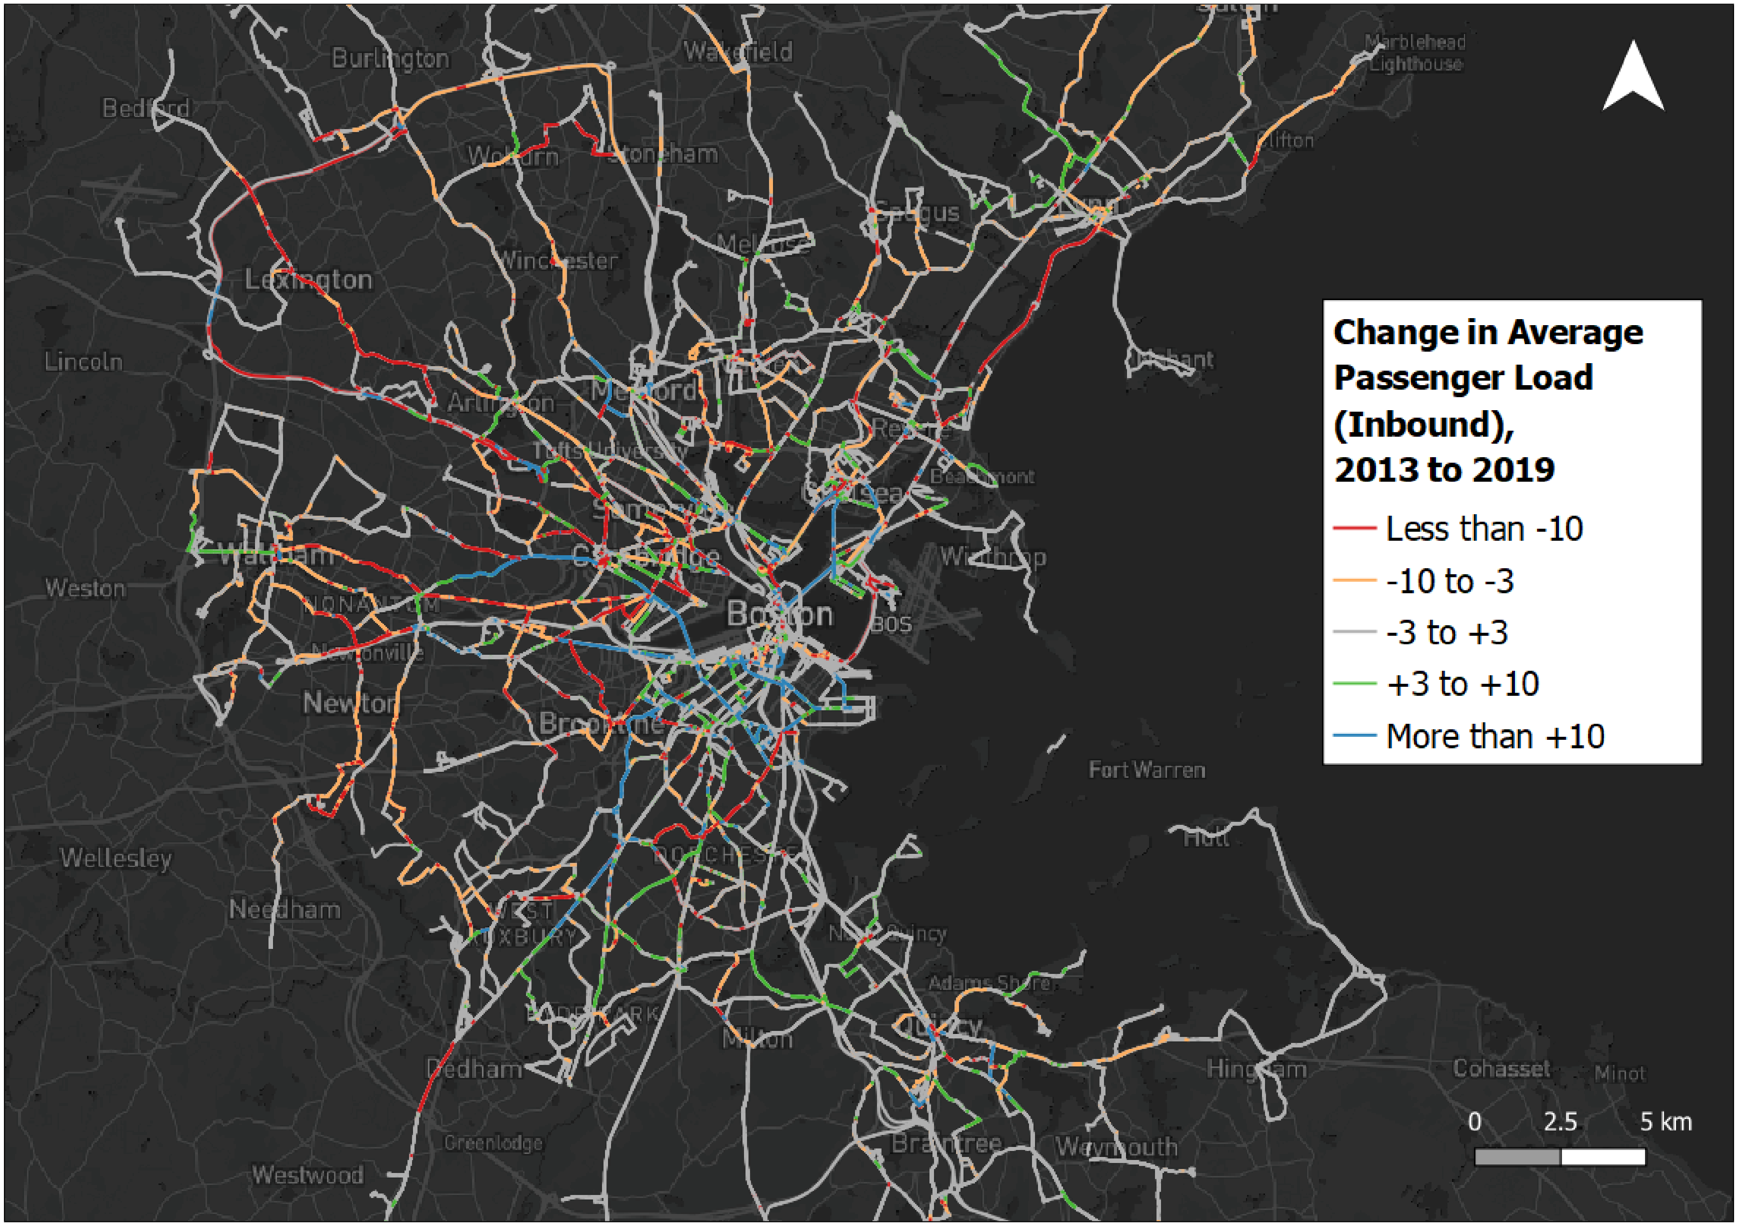

Figure 3-9: Change in average passenger load for **inbound** weekday AM peak hour trips by edge, January 2013 to January 2019 across the MBTA Bus Network

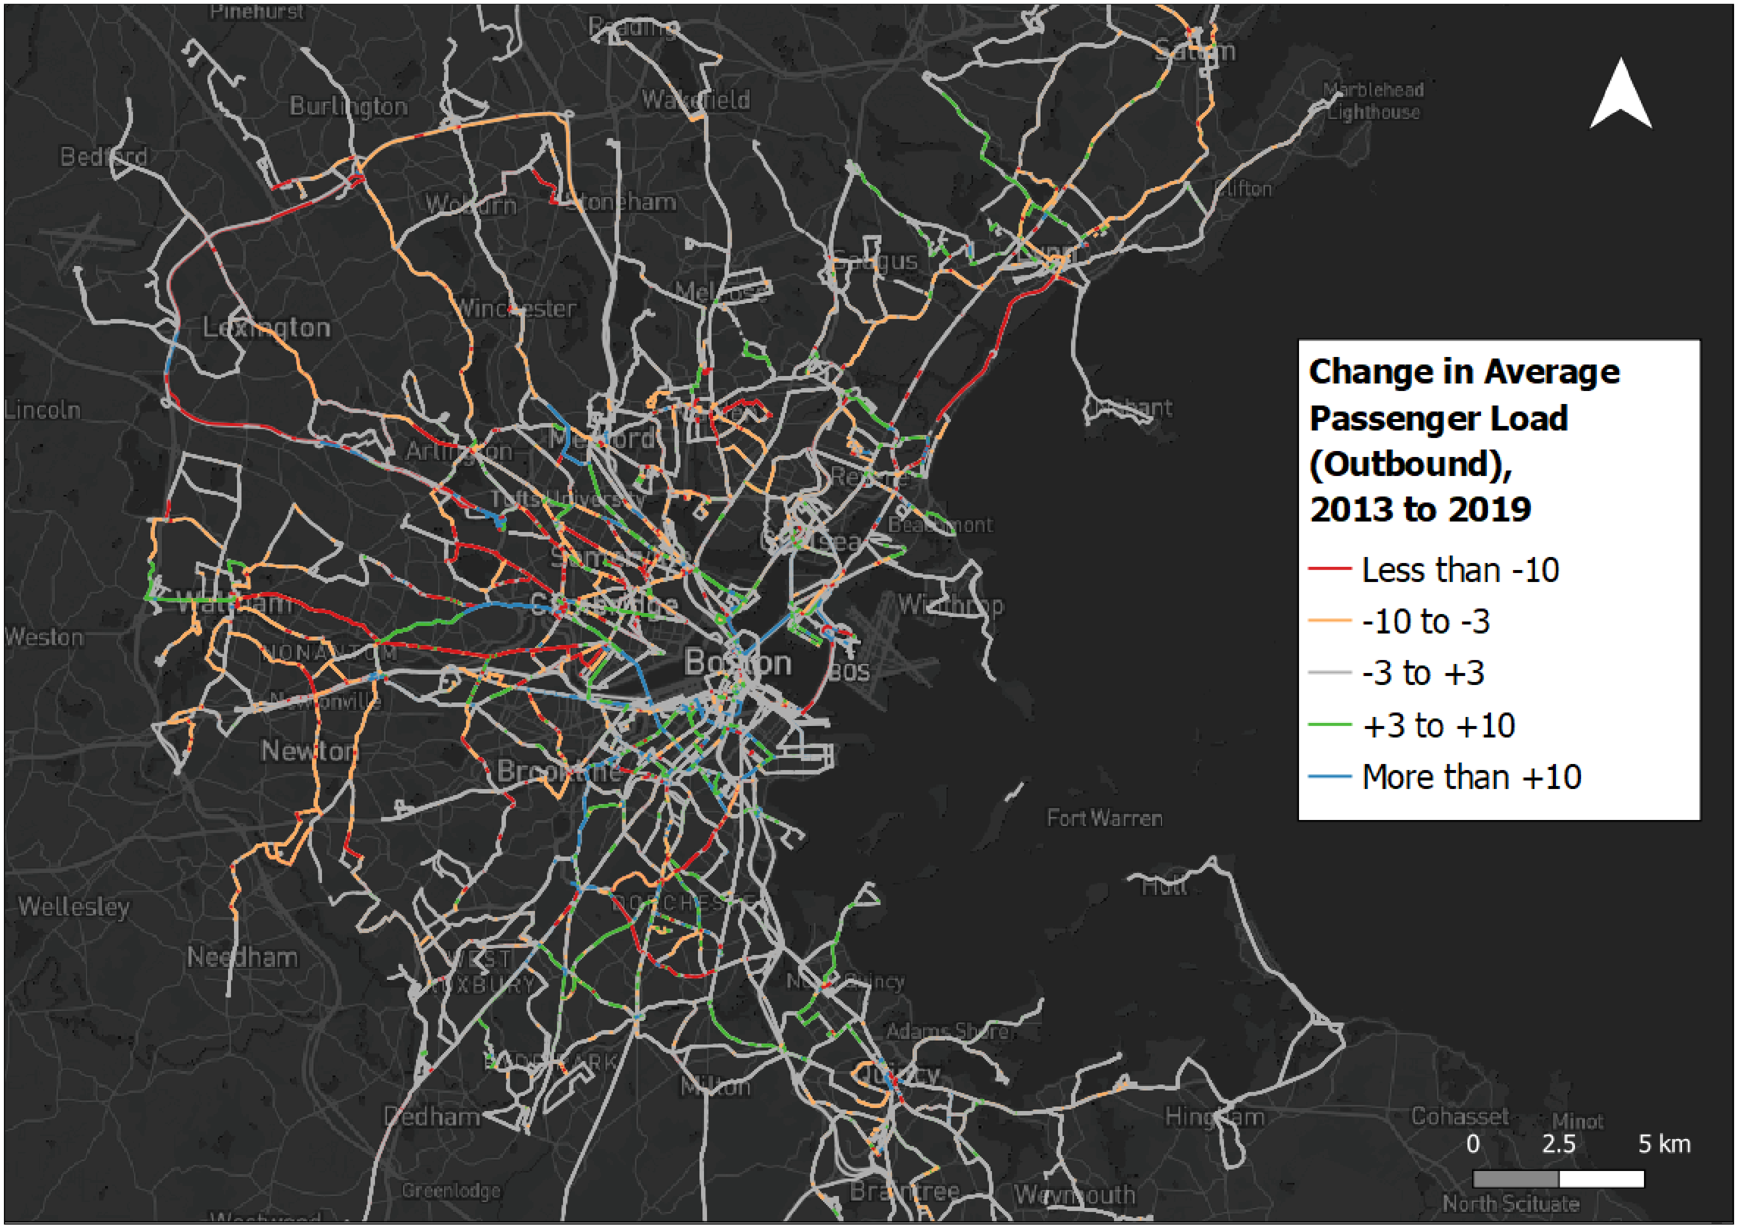

Figure 3-10: Change in average passenger load for **outbound** weekday AM peak hour trips by edge, January 2013 to January 2019 across the MBTA Bus Network

stop or timepoint-to-timepoint level visualization would not be able to capture such patterns.

It can be observed that the passenger load on Route 112 upstream of the SL3 merge point has also increased. The passenger load on the edge between the last stop on Route 112 prior to the connection with the SL3, however, has decreased. This is likely due to passengers alighting from Route 112 and board the SL3 for an express connection to downtown Boston. Transit planners could use this information to adjust schedules or stop locations for improved connections between the 112 and SL3. Ridership downstream of the merge point is higher in 2019, due to the introduction of the SL3. After the two routes split further downstream, ridership on Route 112 is similar to 2013. This example demonstrates how passenger load trends can be inferred from longitudinal analysis enabled by decomposition and geographic representation.

3.7 Discussion

This chapter presents a systematic method for decomposing transit networks into small units of analysis ("edges") that can be used to compare the performance of bus transit networks at a disaggregate spatial level over long periods of time. To enable longitudinal comparison, the edges are represented by their geography, rather than the route or stop identifiers, so that they remain stable in the case of identifier revisions. These methods allow transit planners to study how supply, performance and ridership, have changed over time within a given locality, even if their historical data is only organized by the route and stop identifiers of the era. Two case studies are included to illustrate the strength of these methods. The first uses only GTFS to create maps of service changes for every city block in the service area. The second combines these spatial methods with ridership data to examine trends in bus ridership based on geography rather than only routes.

This approach, like the methods described in Chapter 2, align the measurement of bus performance more closely to the way in which passengers interact with the transit system. Ultimately, changes to route IDs do little to affect the passenger experience,

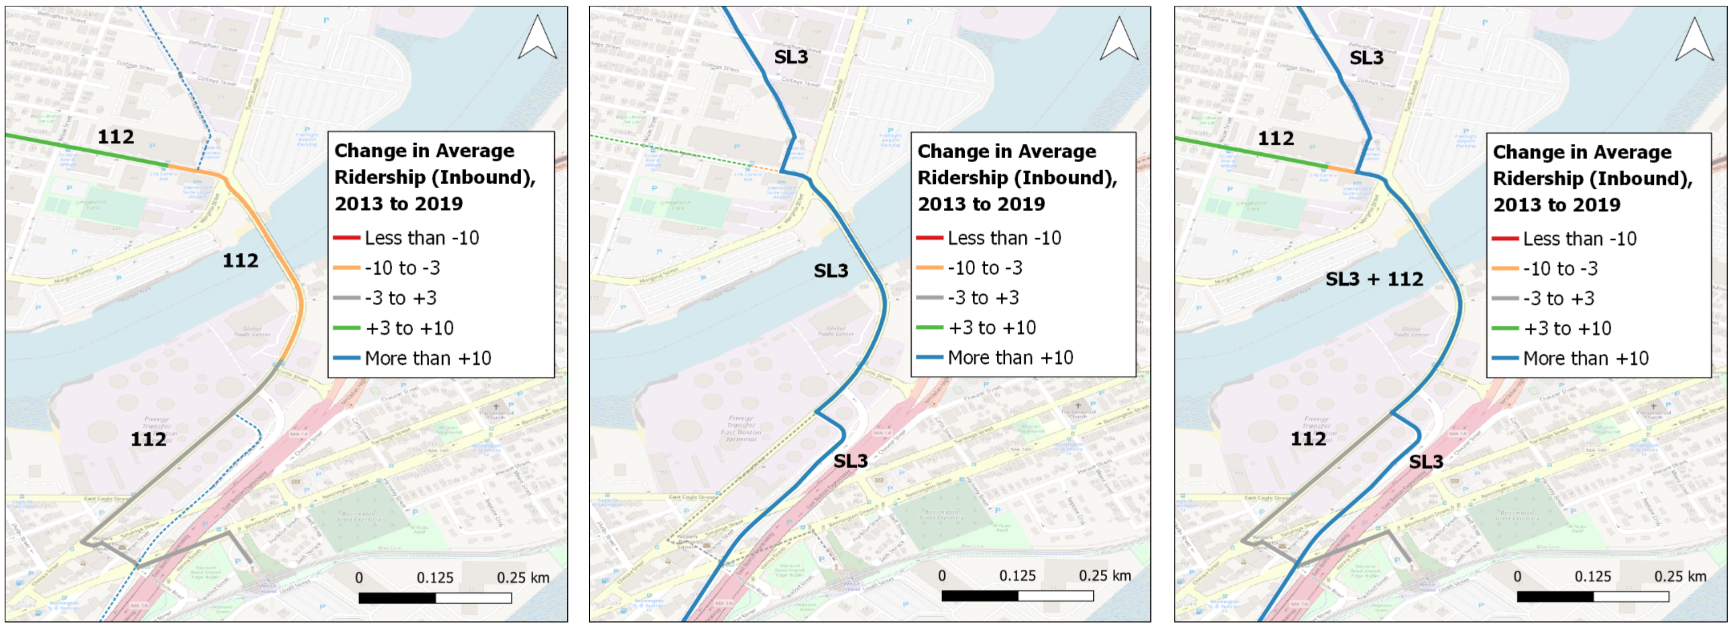

(a) Route 112 only

(b) SL3 only

(c) Both routes combined on shared corridor

Figure 3-11: Change in average passenger load on Route 112 (a), SL3 (b), and combined Route 112 and SL3 (c), for inbound weekday AM peak hour trips, January 2013 to January 2019 so measuring performance by geography can provide a more passenger-oriented perspective. This approach also lowers the barrier to conducting comprehensive reviews of service changes over the course of several years. Curious parties no longer need to parse service modification announcements or have local knowledge in order to compare the headways in a given neighborhood from one decade to the next. If transit performance data are available at the route or stop-to-stop level, but the names or stop patterns of routes have changed over time, it can be onerous to compute how performance has evolved in a given area. This approach provides a method of mapping those data to the appropriate set of edges. Anyone with access to the necessary GTFS feeds (which are often stored in public archives) can apply these methods to understand how their local transit agency has provided service in their community. Advocacy groups can use this information to generate independent visualizations and explore how proposed service changes affect individual neighborhoods.

Some limitations to this method are important to consider. First, given that map matching is used in this chapter, the limitations described in Chapter 2 regarding map matching and OpenStreetMap accuracy guarantees also apply here. Second, GTFS feeds are not always perfectly accurate, for many reasons. They may not be updated frequently, so changes to the schedule might occur without being reflected in the feed. Finally, a more general point about this method, but an important one: quantifying the amount of service provided along a city block (or an edge, as they are called in this Chapter) does not provide any indication of accessibility, i.e. where the transit service enables users to travel. This was a deliberate exclusion, as quantifying accessibility generally requires further analysis and, for some measures, additional data sources beyond GTFS and OpenStreetMap. Future extensions could include weighting by accessibility measures.

Chapter 4

Origin-Destination Applications

4.1 Introduction and Motivation

The methods proposed in the two previous chapters are effective tools for overcoming many common challenges in transit planning. Practicality and generalizability were emphasized in order to make these tools available to transit agencies of all sizes. As a result, input data sources were limited to those that are widely available, such as GTFS, AVL and APC. There are, however, additional applications that leverage rich data sources to generate further insights beyond simple measures of transit supply and performance. This chapter demonstrates how origin, destination and transfer (OD) data can be integrated into the framework proposed in the two previous chapters. Moreover, it demonstrates how the aggregation, decomposition and geographic representation concepts can be used to support a broader set of transit planning tasks and inquiries.

Automated collection and processing of OD data has been enabled by recent advances in technology and academic research. So-called "smart cards" are a type of fare media that are used repeatedly, typically be the same individual, allowing transit agencies to identify patterns in user behavior. As discussed in Section 1.5.3, estimating destinations for transit trips is a very active field of research, developing methods that have been adopted in practice. Simultaneously, more and more transit agencies are installing tap-on tap-off fare collection systems that obviate the need for estimating destinations, although transfers within the fare paid area must still be inferred. Some pilot projects have even used mobile phone records and other digital traces to estimate OD flows. The availability of OD data is mixed but trending upwards. Many larger transit agencies, such as the MBTA and CTA, have access to estimates of OD flows derived from AFC and AVL records.

OD sources can be used in place of APC counts, which are subject to some error, or for completely new applications that are not described in earlier chapters. OD flows are a valuable source of information for understanding spatio-temporal demand patterns throughout a transit network. These data also facilitate the computation of information about journey characteristics, such as the number of transfers or travel times, that is not available from traditional APC or AVL sources. Transit planners may be interested in reviewing the distribution of destinations for passengers boarding at a specific location in order to support long-term service changes. Like other historical data, however, OD flows are often represented by stop identifiers, or trips tagged to a route identifier. As discussed in previous chapters, this can make temporal comparison quite challenging if stop and route identifiers have been revised.

Identifying overlapping routes, spatial decomposition and geographic representation can all be applied to OD data, just as these methods were applied to APC and AVL data in previous chapters. The same practical benefits are accrued; spatial decomposition and aggregation across routes allows performance measures to be analysed at the block level, while geographic representation facilitates long-term temporal comparison. Additional steps are required for pre-processing OD data in order for it to be represented by paths through the transit network rather than straight lines between origin and destination. Representing the data as a path through the transit network has the advantage of informing planners about how the transit network is navigated and which infrastructure elements (i.e. bus lanes, traffic signals) are used.

There are many different planning applications for OD data. Understanding how demand patterns change before and after service changes provides valuable insight into travel behavior. It can be used to evaluate whether service changes produced the desired effect on travel behavior, or if the proffered alternatives to cancelled service were actually adopted. By extending systematic geographic representation and aggregation to the analysis of OD flows, these types of questions can be answered quickly and easily. Furthermore, these methods can be used to support direct comparisons between OD matrices, which is an emerging area of research [111].

The remainder of this chapter is organized as follows. Section 4.2 describes a process for converting identifier-based representation of OD flows into a geographic representation. Sections 4.3 and 4.4 apply these methods and data sources using real data from the MBTA network, providing two practical use cases. Sections 4.5 summarizes the results and implications.

4.2 Processing Origin-Destination Data

The purpose of this section is to develop a method for extracting OD information into a quantity or quantities that can be represented at the edge level. Just as AVL data (vehicle positions) or APC data (passenger counts) can be translated into performance measures like speed and passenger flow, OD data can be used to compute measures relating to travel demand. For example, it can be used to determine the percentage of passengers traveling along each edge that make a transfer to another bus or rail at some point during their trip. Unlike AVL or APC data, which is typically available at the stop-to-stop level, OD flows are represented as a pair of stops. In order to fit OD flows into the framework developed in the two previous chapters, some pre-processing is required. The first step will be to infer the path taken through the transit network for each journey if it is not already inferred during the destination inference process. This translation will allow OD information to be represented geographically, aggregated across routes serving the same edge, and visualized.

OD data, like AVL and APC data, do not yet have a standardized format or structure. In this case, we will consider OD data that consists of a separate record for each vehicle trip within each full passenger journey. For example, if a passenger journey involves a bus to rail transfer, there would be two records, one for the bus trip and one for the rail trip. Each record contains an origin boarding stop ID

Journey ID	Leg	Origin Stop ID	Destination Stop ID
journey_00001	1	875	523
journey_00002	1	16364	6366
journey_00002	2	1722	1726

Table 4.1: Initial origin-destination records

and time, destination alighting stop ID and time. Records are indexed by a unique journey ID, where multiple records for a given journey ID represent a transfer to another vehicle for completion of a one-way passenger journey. This is similar to the structure described by Gordon et al. [7]. From this information, it is possible to infer the path taken through the transit network from origin to destination as a series of stop-to-stop segments as shown below.

Consider the sample MBTA OD records provided in Table 4.1. These sample records will be used as an extended example to illustrate the process. Note that time-stamps for each origin and destination are omitted for clarity, but would be included in any OD table.

The first record represents a single bus trip, from Stop 875 (Forest Hills Station) to Stop 1720 (Morton Street at Canterbury Street) with no transfers. Since both stops are served by Route 31, it can be inferred that Route 31 was taken directly from origin to destination. The second journey, represented by the second and third rows in Table 4.1, involves a transfer from one bus route to another. The first leg of the journey is from Stop 16364 (River Street at Holmfield Avenue) to Stop 6366 (River Street at Blue Hill Avenue), while the second leg is from Stop 1722 (1624 Blue Hill Avenue at Mattapan Square) to 1726 (Blue Hill Avenue at Mattapan Library). Comparing these stops to service patterns from GTFS, it can be inferred that the first leg was served by Route 24, while the second leg was served by Route 28. In brief, the first step involves determining the route taken for each leg of the journey, based on the boarding and alighting stop. If multiple routes serve the boarding and alighting stops, then the boarding and alighting times can be matched with AVL records to identify the correct route.

The schedule of each route (available from a GTFS feed) can then be used to

Journey ID	Leg	Route ID	Origin Sequence	Destination Sequence
journey_00001	1	31	1	5
journey_00002	1	24	17	20
journey_00002	2	28	2	6

Table 4.2: Sample origin-destination records converted to stop sequences

Journey ID	Leg	Traversed Segments
journey_00001	1	(875, 520) (520, 11521) (11521, 5232) (5232, 523)
journey_00002	1	(16364, 6365) (6365, 16365) (16365, 6366)
journey_00002	2	(1722, 1723) (1723, 1724) (1724, 1725) (1725, 1726)

Table 4.3: Origin-destination records converted to sequence of segments, represented by pairs of Stop IDs

derive the set of stop-to-stop segments traversed by each leg of the journey. Table 4.2 summarizes this step by converting the stop IDs from Table 4.1 into a pair of stop sequences along a route.

Stops 875 and 523 represent the first and fifth stops in the service pattern of Route 31, therefore the first journey in Table 4.1 traversed the initial four stop-to-stop segments of Route 31. The segments traversed by the second journey in Table 4.1 will include segments from both Route 24 and Route 28. The list of segments for each of the sample records is shown in Table 4.3.

Once the full list of segments traversed by each journey is determined, the next step is to convert the list of segments to a list of edges. This can be done using the segment-edge correspondence computed in Section 3.2. The final step is to store the journey information (such as whether it includes a transfer, or the location of the final journey destination) as a property of each edge in the journey. Once this process is complete, each edge will contain information for each of the journeys that traverse it, which can be aggregated if desired. Because the edges are represented geographically and are not related to any identifiers, the travel demand measures derived from OD data can be easily compared over time.

To summarize, this section outlines the series of steps that can be used to generate measures related to travel demand from an OD data and translate it into a series of edge properties, or properties of any spatial element of the transit network. The steps are as follows, starting with a set of OD journey records:

1. Determine the travel demand measure of interest
2. For each leg of the journey:
   (a) Use the stop numbers to identify the route taken
   (b) Determine the position of the boarding and alighting stops along the route
   (c) Add each stop-to-stop segment between the boarding and alighting stops to the list of traversed segments
3. For each traversed segment:
   (a) Identify the set of edges that make up the segment
   (b) Add travel demand measure to the properties of each edge

Exploring the fraction of passengers on each edge who make a downstream transfer could help planners to prioritize reliability improvements, as missed transfers can add significant delays to passenger journeys (see Section 4.3). Time-invariant units of analysis are important to track changes over time. Another application would include identifying and visualizing the destinations for all passengers who traverse a given edge, as shown in Section 4.4. Such an analysis could support long-term planning by identifying opportunities for more direct service.

4.3 Case Study: Visualizing Transfer Rates

Understanding where and when passengers make transfers within the transit network can be extremely helpful for transit planners. Schedules of intersecting bus routes and rail lines can be coordinated to reduce transfer times at common transfer points. Efforts to improve running time and running time reliability are especially beneficial when applied to routes with a large number of passengers who transfer downstream. These efforts can help to limit missed transfers, which have an outsized impact on passenger journey time. The location and times of transfers, as well as information about whether certain passengers will make a transfer, are not available directly from AVL or APC systems, but can be inferred from OD estimates.

This case study demonstrates how the methods developed in Chapter 2, aggregation across shared corridors, can be applied to OD data in order to visualize journey-related performance measures for the purpose of transit priority infrastructure planning. In a sense, this case study is an reframing of the first case study in Chapter 2 to focus on the passenger transfer benefits of transit priority infrastructure. Two related metrics are developed: the first to identify the segments with a high absolute number of passengers who make a downstream transfer, and the second to combine the number of downstream transfer passengers with a measure of running time reliability. These metrics are then visualized separately for the entire bus network to identify and prioritize locations where infrastructure-based efforts to improve running time reliability could be implemented. As described in Chapter 2, systematic corridor aggregation benefits infrastructure-level analyses such as these by demonstrating the full impact of any infrastructure improvements.

Real OD data from the MBTA for weekdays in January 2020 is used in this case study. The selected period represents typical (i.e. pre-COVID) ridership patterns. Note that the destinations and transfer locations for this OD data are inferred using AFC and AVL records. Transfers within the rail rapid transit network and from rail to bus are not considered, as this case study focuses on infrastructure improvements that can be applied to the bus network. Bus-to-rail transfers are considered, however, as these benefit from bus reliability improvements. Like the case study in Section 2.5.1, this represents a first order approximation used for illustrative purposes. Transit planners may include additional information and steps as needed to fit the characteristics of their network.

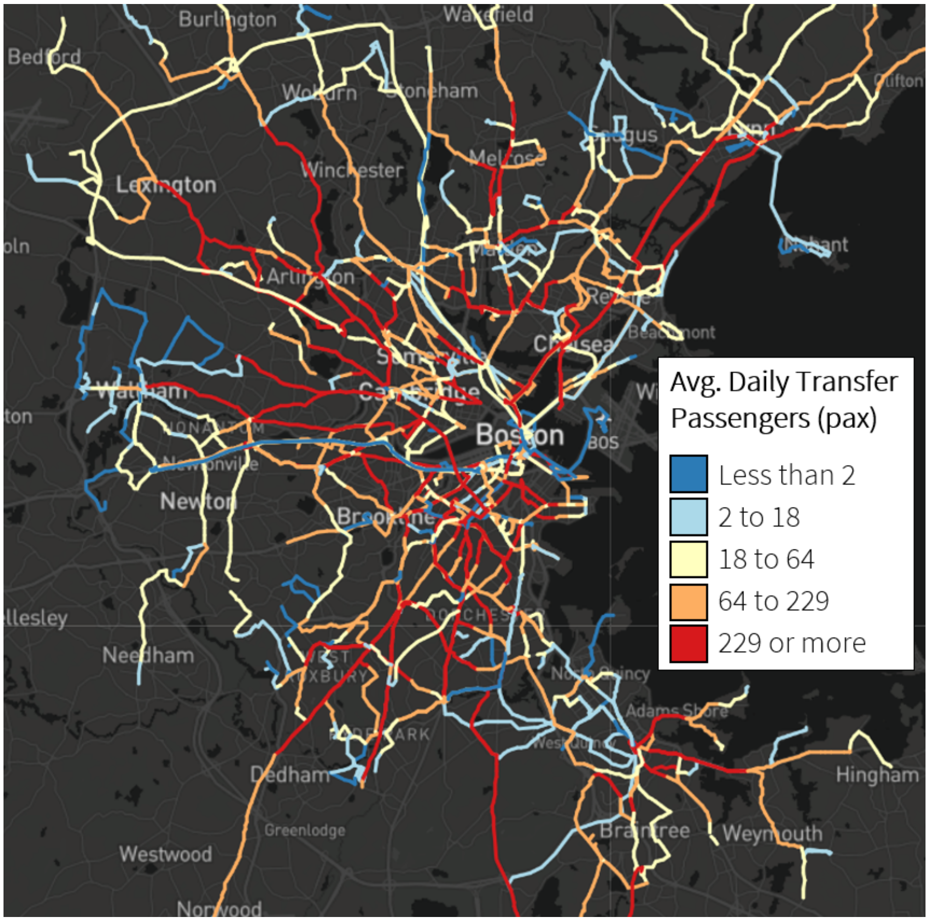

Figure 4-1: Average daily downstream transfer passengers by corridor in the inbound direction for the MBTA bus network, weekdays in January 2020

The first step in evaluating the potential for installing TSP to reduce total journey times is to identify the travel patterns of passengers prior to making a transfer. Using the methods from Section 4.2, the number of total passengers making a downstream transfer is computed for each segment in the bus network. These counts are a passenger-level metric, and are aggregated across shared corridors according to the process defined in Section 2. Downstream transfer passenger counts are visualized for the inbound and outbound directions in Figures 4-1 and 4-2, respectively.

The legend entries reflect quintiles of the distribution. The median number of average daily downstream transfer trips for all corridors is 34. The spatial trends are not entirely surprising. Routes with higher overall ridership generally tend to have a high number of transfer passengers, although that is not always the case. The express

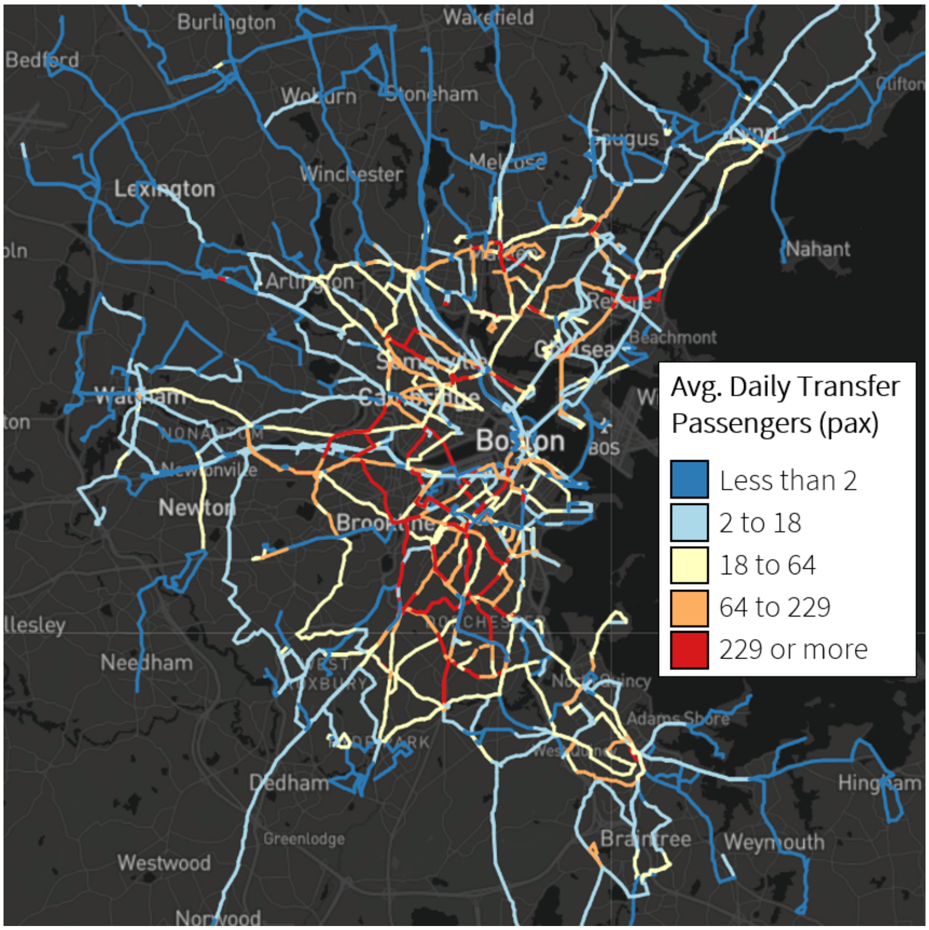

Figure 4-2: Average daily downstream transfer passengers by corridor in the outbound direction for the MBTA bus network, weekdays in January 2020

routes that use I-90, for example, have relatively high ridership but few transfer passengers, likely because they are designed to serve commuters and terminate at major employment centers. In the inbound direction, downstream transfers are more common on segments upstream of rail stations (e.g. Quincy Center, Nubian Square). Some routes on the periphery of the service area simply have fewer passengers overall, and thus a lower number of downstream transfer passengers as well. For the outbound direction, passengers with downstream transfers are much more concentrated around the downtown area and nearby neighborhoods.

Keen observers may note that the SL1, traveling between South Station and Boston Logan Airport, shows very few transfer passengers, despite the fact that it is a high ridership route that connects to the rail network. The SL1 is free when boarding at the airport, so there is no need for the passenger to interact with the fare payment system. As the MBTA's OD estimation process uses AFC as its primary source of data, the OD journey records do not capture the inbound passenger load on the SL1.

Once the number of transfer passengers is known, determining segments with poor reliability is the next step. There are many different measures of reliability used for bus transit; in this case a simple but reasonable measure that can be weighted by the number of passengers is appropriate. Standard deviation in running time is used to develop measures of reliability [112]. Typically, the inverse of the standard deviation is used to represent reliability, as lower standard deviation then produces a higher reliability score. For this case study, however, the intent is to highlight areas with poor reliability, so the standard deviation will be used directly. Alternative measures of reliability could be substituted here for more detailed analyses, such as headway variability or on-time performance.

The standard deviation of running times for each segment can be computed by collecting the distribution of segment travel times from AVL records. The study period used for running times is the same as for OD flows, weekdays in January 2020. Running times were collected for all times of day to provide a first order approximation. There can be significant variation in running time standard deviation between the peak periods and off-peak periods, so period-specific running times could be used for a closer investigation.

The standard deviation of running times (in minutes) is multiplied by the number of transferring passengers to develop a metric which will be called "Transfer Reliability Improvement Potential", or TRIP, measured in passenger-minutes. Note that while this metric has units of passenger-minutes, it is based on the standard deviation and therefore has no relationship to the total amount of time that could be saved by reliability improvements. The results are visualized by corridor for the inbound and outbound directions in Figures 4-3 and 4-4.

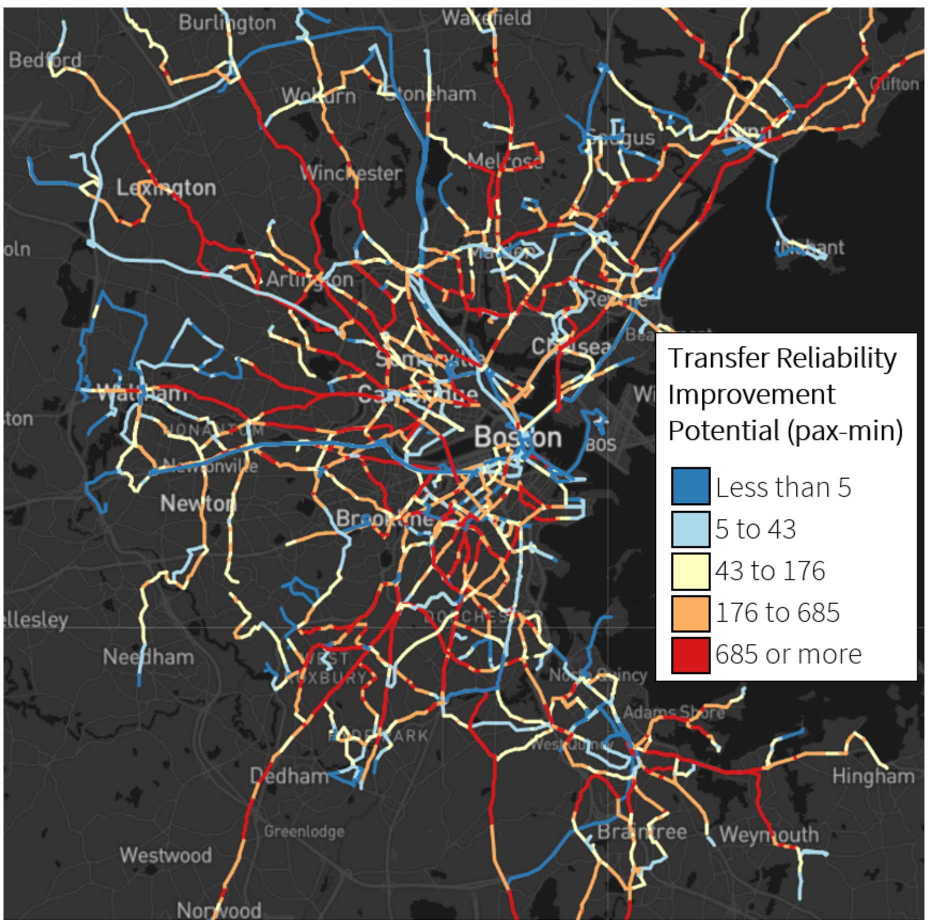

Figure 4-3: Average transfer reliability improvement score by corridor in the inbound direction for the MBTA bus network, weekdays in January 2021

The opportunities for improving bus running time reliability that would affect a significant number of passengers are spread across many routes. The median value overall is 91 passenger-minutes. The transfer reliability improvement score is based in

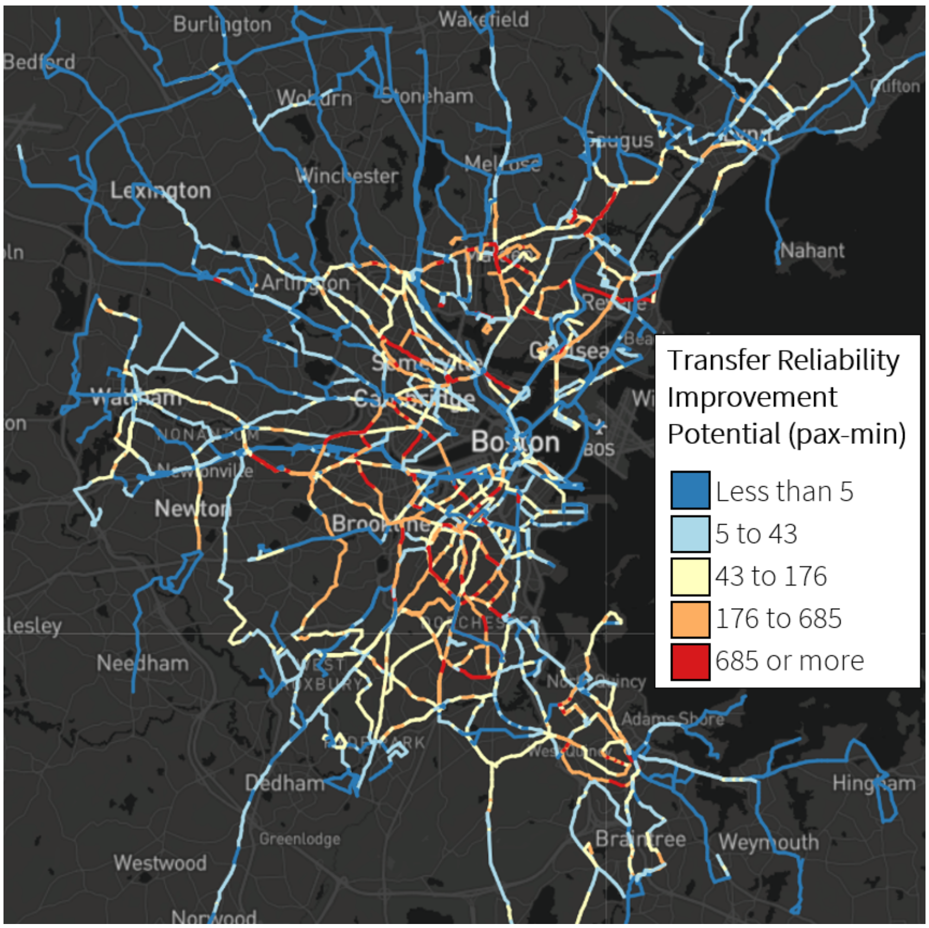

Figure 4-4: Average transfer reliability improvement score by corridor in the outbound direction for the MBTA bus network, weekdays in January 2021

part on the number of downstream passenger trends, so the spatial trends are similar to Figures 4-3 and 4-4. However, introducing reliability does change the results; for example, many of the routes traveling inbound towards Lynn have transfer passenger counts in the third and fourth quintiles, but transfer reliability improvement scores in the fifth quintile, indicating that the travel time reliability of those corridors is below average. Opportunities for improving transfer reliability in the outbound direction are generally concentrated around the inner ring suburbs and Dorchester.

In summary, this case study demonstrates how origin, destination and transfer records can be combined with the methods of Chapter 2 for more advanced use cases. The corridors with high transfer reliability improvement potential could be targeted with infrastructure improvements to increase reliability, such as transit signal priority or queue jump lanes. Visualization also helps planners to identify spatial trends that can be used for implementation planning.

4.4 Case Study: Mapping Destination Changes Over Time

This case study demonstrates how the methods described in Section 4.2 can be used to visualize changes in travel patterns over time when OD data is available. When service changes are introduced, passengers whose routes are disrupted must respond by altering their travel behavior. In many cases, this will include choosing a different path through the transit network in order to get to the same destination. In other cases, the traveler may select an alternative destination that is more convenient, or choose to switch from transit to another travel mode. Examining the distribution of downstream passenger flows from a given edge, and how they adjust in response to a service pattern change, can be helpful in adjusting schedules and planning future changes.

These types of analysis and visualization can be created using traditional OD data. The same issues that were identified in previous case studies also apply to OD data, however. Identifiers are not stable over time, especially in the context of service changes, so longitudinal comparisons may require manual input to correct for changing stop or route IDs. This case study demonstrates how geographic representation facilitates a systematic analysis without any prior knowledge of the service changes.

The MBTA implemented a major bus network change in the fall of 2019 [113, 114]. Five routes were eliminated and dozens of additional routes were modified. To examine the impacts that these changes had on overall passenger demand as well as route choice, this case study compares OD flows between January 2019 to January 2020. The methods from Section 4.2 are used to compute the downstream passenger flows for each edge in the network. The difference between the baseline period (January 2019) and the comparison period (January 2020) are calculated. The last 15 non-holiday weekdays of the month were used to determine passenger flows for both January 2019 and January 2020. System-wide ridership was 3.7% greater for the January 2019 period than the January 2020 period. The difference in passenger flows can then be visualized for any edge in the network, allowing for a comprehensive review of the passenger impacts of specific service changes.

One of the key network changes that was implemented in the fall of 2019 was the replacement of Route 70A with a simpler Route 61 and added service on Route 70, which overlapped with much of Route 70A [114]. These route changes are shown in Figure 3-7. Some passengers who previously rode Route 70A would have their service replaced by Route 70, others by Route 61, and some may have to board or alight at different stops. In January 2019, a typical weekday had 58 scheduled outbound trips on Route 70, which included 10 minute headways during peak hours and 20 - 30 minute headways during off-peak hours. Route 70A was scheduled to run 20 weekday outbound trips in January 2019 between 6 AM and 7 PM, with approximately 30 minute peak headways. The cancellation of Route 70A was partially offset by added service on Route 70. In January 2020, Route 70 had 76 scheduled weekday trips (a slight decrease of 2 trips per day compared to the prior combined service) with 10 minute headways throughout much of the mornings and evenings. But Route 70 does not provide access north of Waltham Town Center, forcing former Route 70A passengers with destinations in that area to transfer to Route 61, which serves many of the same stops as Route 70A and operates on a similar frequency.

By visualizing the destinations of passengers passing through corridors served by Route 70A before and after the route changes, it is possible to estimate the effect that the service change had on passenger route and destination choice. First, OD data is used to determine downstream passenger flows for each stop-to-stop segment using the methods described in Section 4.2. Then, the segment-level values are translated into edge-level using the correspondence between segments and edges developed in Chapter 3. The conversion to edge-level representation allows for stable comparison over time. The result is a set of data that can be used to visualize the set of downstream passenger flow for any edge in the network.

Figure 4-5 shows outbound passenger flows on a typical weekday for all journeys passing through an edge that was served by Route 70A and Route 70 for January 2019: Westbound U.S. Route 20 (Main Street) between Lafayette Street and Willow Street. These are then compared to the passenger flows downstream of the same edge in January 2020, shown in Figure 4-6. As a result of the network redesign, the edge is now served by Route 70 only with enhanced headway reliability on the single route 70 as compared to the two separate routes. Passengers who previously used Route 70A to travel north of Waltham Town Center must now transfer to Route 61. The difference in passenger flows between the two periods is shown in Figure 4-7.

On a typical weekday in January 2019, there were 438 bus passengers traveling outbound through the selected edge on Main Street, which is represented by the dark red color west of the selected edge in Figure 4-5. This passenger flow includes passengers on both Routes 70 and 70A. In January 2020, the typical weekday passenger flow increased to 550 passengers, represented by the dark red color west of the selected edge in Figure 4-6, potentially suggesting that the service changes attracted more ridership in the area. The spatial distribution of destinations, however, changed substantially in response to the service changes. In January 2019, Route 70A carried 48 passengers north of Waltham Town Square, with moderate ridership for many stops afterwards. In January 2020, on the other hand, a daily average of just 6 passengers

Figure 4-5: Average daily passenger flow for passengers traveling outbound on MBTA Routes 70 and 70A in Waltham in January 2019

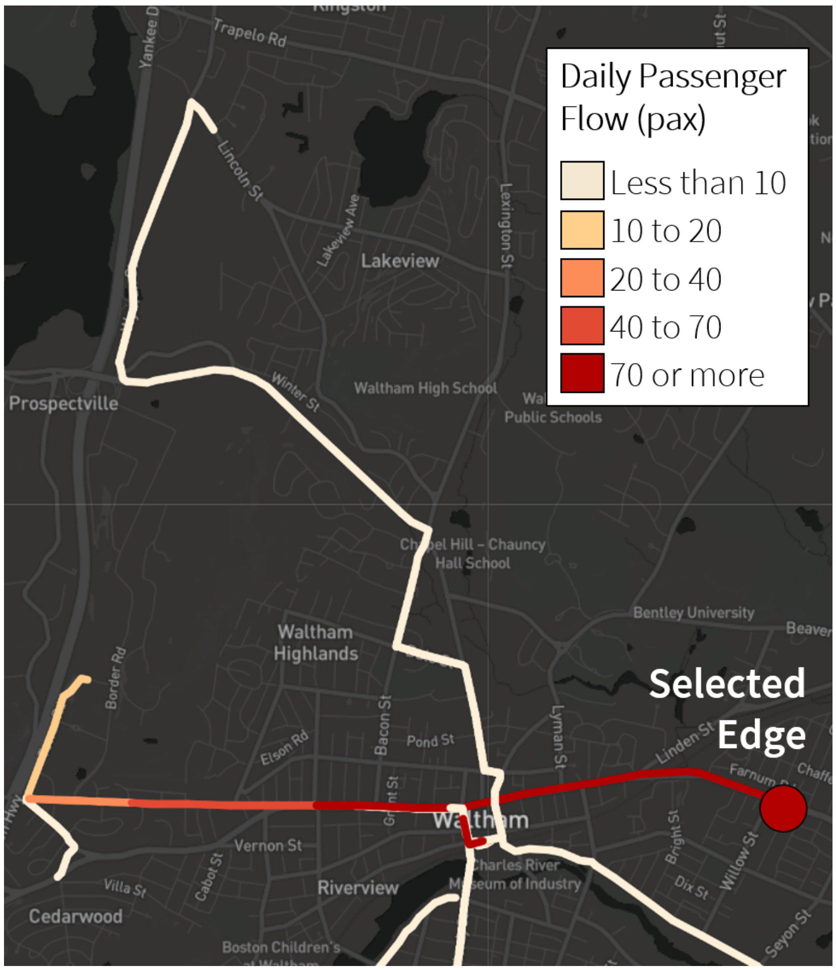

Figure 4-6: Average daily passenger flow for passengers traveling outbound on MBTA Routes 61 and 70 in Waltham in January 2020

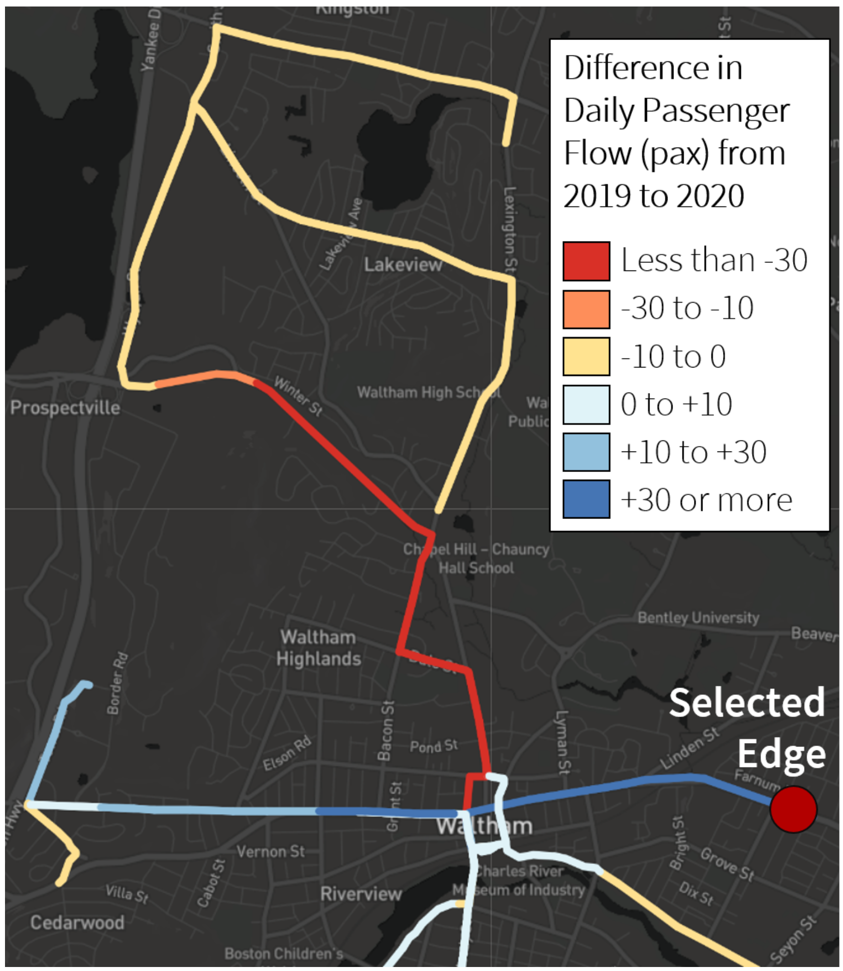

Figure 4-7: Change in average daily passenger flow in North Waltham, January 2019 to January 2020

have an inferred transfer from Route 70 to Route 61 in order to reach destinations north of Waltham Town Square. This suggests that the majority of passengers opted not to follow the recommendations presented in the route change notice (transfer from Route 70 to Route 61), likely due to the disutility of making a transfer to another bus route. Based on this analysis, it appears that the service change was successful overall in attracting significantly more passengers to the primary corridor service while losing some passengers in the outer suburban edges of the the corridor that now require a transfer.

This case study summarizes how OD data can be used to visualize passenger flows throughout the transit network, and provides an example of how that visualization can inform future service planning. The comparison of ridership over time under significant network changes allows planners to understand how passengers respond to new service patterns. The systematic nature of this process makes it simple to conduct these analyses any time a network is changed or disrupted.

4.5 Discussion

This Chapter proposes a systematic method for processing OD records in order to assign journey statistics to spatial elements of a transit network, such as the corridor or edge. This additional step prepares the OD data for the spatial analysis methods described in the two previous chapters. These records can then be aggregated across corridors for infrastructure-level analysis, or decomposed and represented geographically for longitudinal comparison. The proposed process and new data source enable a new set of analyses beyond what are described in Chapters 2 and 3. Two case studies with real OD data are included to demonstrate how the methods can be used in practice for data-driven transit infrastructure planning and for visualizing the effects of service changes on travel behavior.

OD data is just one additional data source that could be combined with the methods developed in this thesis. This chapter develops a framework for including performance measures based on data sources beyond APC, AVL and GTFS. For example, targeted surveys that have trip information could be used to create metrics related to customer satisfaction that are then compared over time. Another example would be operating cost data, which could be used for developing productivity-related metrics. The pre-processing step outlined in Section 4.2 can be adapted for use with these data sources and others.

There are some limitations to the methods described in this chapter. First, the destinations of OD records are typically inferred unless collected directly from "tap-on, tap-off" fare systems. While these inferences can be quite accurate, they are typically subject to some error; the lack of fare system interaction for the SL1 described in the pre-processing step in Section 4.3 is one such example. Data quality issues for any OD estimation inputs could be propagated to the OD results.

In addition, Section 4.2 assumes that the OD data has a specific structure: one record for each vehicle trip by each passenger, with a unique identifier that can be used to identify linked trips. If the OD data is organized differently, then additional data processing may be required to achieve the same results. Furthermore, the algorithms used to determine whether two vehicle trips by the same passenger are part of the same journey or separate journeys will affect the data structure and impact the transfer passenger analysis shown in Section 4.3.

There are additional limitations unrelated to data quality. Section 4.4 describes how changes to passenger destinations before and after network modifications can be determined systematically, without knowing anything about the network. However, interpreting the change in destination choices requires additional context that is not shown in a passenger flow visualization, such as any changes to frequency or forced transfers. Similarly, additional context is needed for thorough review of opportunities to reduce transfer waiting times beyond the methods shown in Section 4.3. For example, the scheduled transfer duration would help to prioritize reliability improvements, because reliability is more important when there is little transfer time. Future work could include additional performance measures that incorporate this information.

Chapter 5

Conclusions and Recommendations

5.1 Summary

This thesis develops three practical methods for identifying spatial relationships between elements of a public transit network and using those relationships to generate new insights:

1. Identification and aggregation of shared bus corridors

2. Spatial decomposition to block-level resolution

3. Geographic representation of transit network elements

These three methods avoid dependence on identifiers such as route names or numbers, which vary over time and thereby obfuscate changes in transit performance. When combined, the methods create a powerful framework for transit analysis at the network scale and at the resolution of individual city blocks.

This analysis framework, with its systematic and robust identification of spatial relationships, removes the need for personal knowledge of where different elements of the transit network intersect with one another. It also drastically reduces the effort needed to conduct comprehensive longitudinal studies by creating a geographic representation of the transit network. The disaggregate unit of analysis enables transit performance to be reviewed for individual components of the road network, such as traffic signals or queue jump lanes. Finally, this framework is oriented around visualization and generates outputs that can be use to enhance communication and identify spatial trends in transit performance.

Multiple case studies are included in each chapter to demonstrate the applications of these methods, individually and combined. Chapter 2 applies shared corridor identification and aggregation to solve common transit planning challenges such as scheduling in shared corridors and balancing service between express and local routes. Chapter 3 uses the spatial decomposition and geographic representation to conduct longitudinal analyses of transit service and ridership over many years at a high spatial resolution. Chapter 4 extends the methods of the previous chapter to include measures of demand generated from OD flows, including passenger transfers and destination distributions.

These case studies present a small fraction of the possible use cases for these methods. It is hoped that these methods will be adopted in practice and that the list of use cases continues to grow. Throughout the development of this research, emphasis was placed on ease of implementation, generalizability and open access. These tools rely on open source software and there are several applications included herein based entirely on open data sources. They can be generalized for use by any transit agency or advocacy group whose transit network is represented by a standard GTFS feed.

5.2 Limitations

The specific limitations of each of the individual methods described in Chapters 2, 3 and 4 are included in the discussion section of the respective chapters. This section provides a broader set of limitations regarding the overall approach to transit performance measurement and analysis taken in this thesis. Also included included are potential barriers to implementation, which are addressed by the recommendations in Section 5.3.

The first set of limitations are related to data. These methods are predicated on the availability of GTFS and OpenStreetMap (or other road network data) to perform map matching and spatial analysis of the transit network. While GTFS is produced by nearly all transit operators in the United States, it has seen lower adoption in other regions. Additionally, reliance on GTFS limits the scope of temporal comparisons as the data standard was not introduced until 2005 and was not widely adopted for another several years. Furthermore, many of the studies shown in this thesis rely on automated transit data including APC, AVL and origin-destination estimation systems, which are subject to error and lack of coverage. Analysts should be aware of the limitation of their data sources and take steps to avoid producing analyses that are affected by data quality results. Some of these steps could include using longer analysis periods or scaling the results to account for data gaps.

The case studies included in this thesis demonstrate how transit-specific data sources such as APC and AVL can be used to aggregate, compare and visualize transit performance. There are other important measures of transit performance that are not included, as these were intended to be illustrative rather than comprehensive examples. Changes in fare or fare structure is an important component of transit service that was not included in the longitudinal comparisons despite their impact on passenger utility. In a similar vein, measures of equity are also omitted. Equity in public transit is an important and complex topic with considerable research into measurement and comparative analysis [30], but is beyond the scope of this thesis. As mentioned in the limitations in Chapter 3, accessibility measures are also important but outside the scope of this thesis.

These methods are also based on analysis of traditional fixed route transit networks. While fixed-route bus transit is the primary subject of this thesis, the methods could be applied to rail transit without significant changes. On the other hand, these methods are not well suited for transit which does not operate on fixed routes or schedules, such paratransit or demand-responsive transit. Paratransit, while a critical component of public transit service that should not be overlooked, has a different service paradigm and performance measures that would benefit from other methods than those included in this thesis. Demand-responsive transit has become more popular in the wake of COVID-19 [115], but ultimately represents a small fraction of total public transit ridership. If the mode share of demand-responsive transit continues to grow, the need for dedicated analysis and visualization tools that capture the inherent flexibility of these modes will increase. Some of these methods, such as the block-level geographic decomposition, might be quite helpful in demand-responsive transit analysis and visualization where traditional route and stop identifiers are not relevant. The tools described in this thesis are intended to make analysis of the spatial relationships in transit networks easier by eliminating the need for multiple data sources, expensive software subscriptions and personal knowledge. That said, implementing these tools will require some technical capacity on the part of transit agencies or advocacy groups. These barriers have been removed to the extent possible by consolidating the scripts for corridor matching and aggregation (Chapter 2) and those for spatial decomposition and edge matching (Chapter 3). Any user, however, must have the software and knowledge for opening, modifying and running Python programs.

5.3 Recommendations

It is recommended that transit planning and performance management staff within transit agencies adopt the methods described herein to streamline and improve the long-term planning process. The timeframes for conducting these analysis are application-dependent. For example, methods presented in Sections 2.5.2, 2.5.3 and 4.3 could be adopted as part of a periodic review process to ensure efficient resource allocation. The analyses described in Sections 3.6 and 4.4, on the other hand, could be required as a post-implementation evaluation of network design changes to determine whether the changes had the desired effect. On the whole, these analyses will improve the information available to transit planners, limit "information silos" within agencies, and provide a new perspective on the spatial relationships within transit networks.

As shown in this thesis, these methods have many practical applications for transit planning, from identifying high-potential locations for transit signal priority to investigating long-term trends in ridership at a high spatial resolution. Agencies can also use these tools to improve the communication of service changes with the public. It is likely that continued engagement with these tools by transit professionals will lead to the identification of further use cases.

In order to do adopt these tools, transit agencies should ensure that they have the technical capacity to automate the processes, produce visualizations and make them widely available to staff. A limited version of the methods described in Chapter 2 have already been deployed at two large U.S. transit agencies and have been used to support decision making, and there are plans to deploy the methods of Chapter 3 in the future. Lessons learned from these deployments could be shared with other agencies to reduce implementation challenges. Standardization of data inputs from APC and AVL systems would also reduce the amount of agency-specific customization needed to include performance measures beyond those related to scheduled service. Creating data standards for these systems is the subject of an ongoing project sponsored by the Transportation Research Board [116].

In addition to agencies, this tool was developed for use by transit advocacy groups. The primary application that could benefit advocates is the analysis presented in Section 3.5, where changes in transit service delivery can be reviewed using only GTFS feeds. It is recommended that these methods be adopted by such groups. If technical capacity presents a barrier to adoption, local governments should provide small grants to promote technical development for advocacy groups. Transit agencies and their riders benefit from a well-informed public.

5.4 Future Work

There is considerable potential for future work related to this thesis. The work in this thesis provides the practical tools that could enable new and interesting studies in a variety of directions. Further refinement of these methods would also benefit the state of the practice.

Future research in this area could develop new TPMs to identify high-priority locations for different infrastructure improvements. Additional data sources, such as high-frequency GPS data, could be incorporated to improve the spatial resolution of speed and delay data. While this thesis focuses on the bus transit mode almost exclusively, the methods could also be applied to rail transit where there is often significant overlap among different branch lines.

Applications of this research could include using these methods to study the growth of bus transit networks over time. GTFS feeds have been adopted by most American transit agencies for over a decade, providing a wealth of data on network topology and service delivery. Transit network data could be combined with socioeconomic indicators to examine how service changes have affected different groups. Future work could also conduct comparative analyses based on the methods described herein. Statistical analysis of the bus network growth patterns combined with ridership data across different urban areas could be used to identify causal relationships between network evolution and ridership.

Finally, these methods could be applied to a very topical challenge for transit agencies: COVID-19 pandemic recovery planning. Chapter 4 demonstrates how spatial decomposition and identifier-free representation enables fast, comprehensive studies of changes in ridership patterns. While the given example was a before and after analysis of a specific route design change, the methods could be especially helpful in reviewing ridership changes throughout the COVID-19 pandemic and into the recovery period. Systematic methods, like those presented in this thesis, would be helpful as COVID-19 pandemic ridership changes affected the entire network. Most transit agencies also modified service patterns in response to the pandemic, so these tools also could be used to distinguish ridership changes that resulted from COVID-19 and changes that resulted from longer headways or cancelled routes.

Bibliography

[1] Massachusetts Bay Transportation Authority. MBTA 2021 System Map (March 2021). `https://mbta-massdot.opendata.arcgis.com/datasets/eec03d901d2e470ebd5758c60d793e8e_0`, 2021. Online; accessed on 2021-05-08.

[2] MBTA. Changes to bus routes 70/70A, new route 61, 2019. URL `https://www.mbta.com/projects/better-bus-project/61-70-70A`. [Online; accessed January 26, 2021].

[3] Kittelson & Associates and United States. Federal Transit Administration and Transit Cooperative Research Program and Transit Development Corporation. *A guidebook for developing a transit performance-measurement system*, volume 88. Transportation Research Board, 2003.

[4] Bibiana McHugh. Pioneering open data standards: the GTFS story. *Beyond Transparency: Open Data and the Future of Civic Innovation*, pages 125–135, 2013.

[5] Philippe Fortin, Catherine Morency, and Martin Trépanier. Innovative GTFS data application for transit network analysis using a graph-oriented method. *Journal of Public Transportation*, 19(4):2, 2016.

[6] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org . `https://www.openstreetmap.org`, 2020. Online; accessed on 2020-11-22.

[7] Jason B Gordon, Harilaos N Koutsopoulos, Nigel HM Wilson, and John P Attanucci. Automated inference of linked transit journeys in London using fare-transaction and vehicle location data. *Transportation Research Record*, 2343(1):17–24, 2013.

[8] Miguel A Figliozzi and Nicholas B Stoll. A study of bus high-resolution GPS speed data accuracy. *Transportation Research Record*, 2672(8):187–198, 2018.

[9] Gordon J Fielding, Timlynn T Babitsky, and Mary E Brenner. Performance evaluation for bus transit. *Transportation Research Part A: General*, 19(1):73–82, 1985.

[10] Paola Prioni and David A Hensher. Measuring service quality in scheduled bus services. *Journal of Public Transportation*, 3(2):4, 2000.

[11] Anthony Cramer, John Cucarese, Minh Tran, Alex Lu, and Alla Reddy. Performance measurements on mass transit: Case study of New York City Transit Authority. *Transportation Research Record*, 2111(1):125–138, 2009.

[12] Xiaolei Ma and Yinhai Wang. Development of a data-driven platform for transit performance measures using smart card and GPS data. *Journal of Transportation Engineering*, 140(12):04014063, 2014.

[13] Meghan Hammerle, Michael Haynes, and Sue McNeil. Use of automatic vehicle location and passenger count data to evaluate bus operations: experience of the Chicago Transit Authority, Illinois. *Transportation Research Record*, 1903(1): 27–34, 2005.

[14] Peter G Furth, Brendon Hemily, Theo HJ Muller, and James G Strathman. Using archived AVL-APC data to improve transit performance and management. *Transit Cooperative Research Program, No. Project H-28*, 2006.

[15] Xidong Pi, Mark Egge, Jackson Whitmore, Amy Silbermann, and Zhen Sean Qian. Understanding transit system performance using AVL-APC data: An analytics platform with case studies for the Pittsburgh region. *Journal of Public Transportation*, 21(2):2, 2018.

[16] Mei Chen, Xiaobo Liu, Jingxin Xia, and Steven I Chien. A dynamic bus-arrival time prediction model based on APC data. *Computer-Aided Civil and Infrastructure Engineering*, 19(5):364–376, 2004.

[17] Jayakrishna Patnaik, Steven Chien, and Athanassios Bladikas. Estimation of bus arrival times using APC data. *Journal of Public Transportation*, 7(1):1, 2004.

[18] Amer Shalaby and Ali Farhan. Prediction model of bus arrival and departure times using AVL and APC data. *Journal of Public Transportation*, 7(1):3, 2004.

[19] Mathew Berkow, Ahmed M El-Geneidy, Robert L Bertini, and David Crout. Beyond generating transit performance measures: visualizations and statistical analysis with historical data. *Transportation Research Record*, 2111(1):158–168, 2009.

[20] Luis Moreira-Matias, João Mendes-Moreira, Jorge Freire De Sousa, and João Gama. Improving mass transit operations by using AVL-based systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 16(4):1636–1653, 2015.

[21] Christopher Pangilinan, Nigel Wilson, and Angela Moore. Bus supervision deployment strategies and use of real-time automatic vehicle location for improved bus service reliability. *Transportation Research Record*, 2063(1):28–33, 2008.

[22] Ahmed M El-Geneidy, Jessica Horning, and Kevin J Krizek. Analyzing transit service reliability using detailed data from automatic vehicular locator systems. *Journal of Advanced Transportation*, 45(1):66–79, 2011.

[23] Jeffrey Hanft, Shrisan Iyer, Brian Levine, and Alla Reddy. Transforming bus service planning using integrated electronic data sources at NYC Transit. *Journal of Public Transportation*, 19(2):6, 2016.

[24] Gregory D Erhardt, Oliver Lock, Elsa Arcaute, and Michael Batty. A big data mashing tool for measuring transit system performance. In *Seeing Cities Through Big Data*, pages 257–278. Springer, 2017.

[25] Laura Cecilia Cham. *Understanding bus service reliability: a practical framework using AVL/APC data*. PhD thesis, Massachusetts Institute of Technology, 2006.

[26] Yindong Shen, Jia Xu, and Zhongyi Zeng. Public transit planning and scheduling based on AVL data in China. *International Transactions in Operational Research*, 23(6):1089–1111, 2016.

[27] Jayakrihsna Patnaik, Steven Chien, and Athanassios Bladikas. Using data mining techniques on APC data to develop effective bus scheduling plans. *Journal of Systemics, Cybernetics and Informatics*, 4(1):86–90, 2006.

[28] Michael Mandelzys and Bruce Hellinga. Automatically identifying the causes of bus transit schedule adherence performance issues using AVL/APC archived data. *Transportation Research Board of the National Academies*, 2010.

[29] Timothy F Welch and Alyas Widita. Big data in public transportation: a review of sources and methods. *Transport Reviews*, 39(6):795–818, 2019.

[30] Phillip R Carleton and J David Porter. A comparative analysis of the challenges in measuring transit equity: definitions, interpretations, and limitations. *Journal of Transport Geography*, 72:64–75, 2018.

[31] Ran Wei, Xiaoyue Liu, Yongjian Mu, Liming Wang, Aaron Golub, and Steven Farber. Evaluating public transit services for operational efficiency and access equity. *Journal of Transport Geography*, 65:70–79, 2017.

[32] Mahmoud Mesbah, Graham Currie, Claudia Lennon, and Trevor Northcott. Spatial and temporal visualization of transit operations performance data at a network level. *Journal of Transport Geography*, 25:15–26, 2012.

[33] Mark W Horner and Alan T Murray. Spatial representation and scale impacts in transit service assessment. *Environment and Planning B: Planning and Design*, 31(5):785–797, 2004.

[34] Mark Dimond, Neil Taylor, and Robert Houghton. Estimating and editing transit topology over the road graph using supply data feeds. *Proceedings of the Association of Geospatial Information Laboratories Europe (AGILE)*, 2016.

[35] Nate Wessel, Jeff Allen, and Steven Farber. Constructing a routable retrospective transit timetable from a real-time vehicle location feed and GTFS. *Journal of Transport Geography*, 62:92–97, 2017.

[36] Jing-Quan Li. Match bus stops to a digital road network by the shortest path model. *Transportation Research Part C: Emerging Technologies*, 22:119 – 131, 2012.

[37] Sergio A Ordóñez Medina. Semi-automatic tool for bus route map matching. In *The Multi-Agent Transport Simulation MATSim*, pages 115–121. Ubiquity Press, 2016.

[38] Kenneth Perrine, Alireza Khani, and Natalia Ruiz-Juri. Map-matching algorithm for applications in multimodal transportation network modeling. *Transportation Research Record*, 2537(1):62–70, 2015.

[39] Ying Zhou, Guiyuan Jiang, and Guifeng Jiang. Discover the road sequences of bus lines using bus stop information and historical bus locations. *International Journal of Distributed Sensor Networks*, 15(2):1550147719830552, 2019.

[40] Nate Wessel and Steven Farber. On the accuracy of schedule-based GTFS for measuring accessibility. *Journal of Transport and Land Use*, 12(1), 2019.

[41] Alex Karner. Assessing public transit service equity using route-level accessibility measures and public data. *Journal of Transport Geography*, 67(1):24–32, 2018.

[42] Yuan Liao, Jorge Gil, Rafael HM Pereira, Sonia Yeh, and Vilhelm Verendel. Disparities in travel times between car and transit: Spatiotemporal patterns in cities. *Scientific Reports*, 10(1):1–12, 2020.

[43] Travis B Glick, Wei Feng, Robert L Bertini, and Miguel A Figliozzi. Exploring applications of second-generation archived transit data for estimating performance measures and arterial travel speeds. *Transportation Research Record*, 2538(1):44–53, 2015.

[44] Nicholas B Stoll, Travis Glick, and Miguel A Figliozzi. Using high-resolution bus GPS data to visualize and identify congestion hot spots in urban arterials. *Transportation Research Record*, 2539(1):20–29, 2016.

[45] Miguel A Figliozzi and Travis B Glick. Evaluation of roadway reallocation projects: Analysis of before-and-after travel speeds and congestion utilizing high-resolution bus transit data. Technical Report NITC-RR-887, Transportation Research and Education Center (TREC), Portland, OR, 2017.

[46] Graeme Brown. Corridor-based bus performance analysis. In *TransitData 6th International Symposium, Online, August 2020.*, 2021.

[47] Jinhua Zhao, Adam Rahbee, and Nigel HM Wilson. Estimating a rail passenger trip origin-destination matrix using automatic data collection systems. *Computer-Aided Civil and Infrastructure Engineering*, 22(5):376–387, 2007.

[48] Neema Nassir, Alireza Khani, Sang Gu Lee, Hyunsoo Noh, and Mark Hickman. Transit stop-level origin–destination estimation through use of transit schedule and automated data collection system. *Transportation Research Record*, 2263 (1):140–150, 2011.

[49] Baibing Li. Markov models for Bayesian analysis about transit route origin–destination matrices. *Transportation Research Part B: Methodological*, 43(3): 301–310, 2009.

[50] Martin L Hazelton. Statistical inference for transit system origin-destination matrices. *Technometrics*, 52(2):221–230, 2010.

[51] Li He and Martin Trépanier. Estimating the destination of unlinked trips in transit smart card fare data. *Transportation Research Record*, 2535(1):97–104, 2015.

[52] Ding Luo, Oded Cats, and Hans van Lint. Constructing transit origin–destination matrices with spatial clustering. *Transportation Research Record*, 2652(1):39–49, 2017.

[53] Etikaf Hussain, Ashish Bhaskar, and Edward Chung. Transit OD matrix estimation using smartcard data: Recent developments and future research challenges. *Transportation Research Part C: Emerging Technologies*, 125:103044, 2021.

[54] Lauren Alexander, Shan Jiang, Mikel Murga, and Marta C González. Origin–destination trips by purpose and time of day inferred from mobile phone data. *Transportation Research Part C: Emerging Technologies*, 58:240–250, 2015.

[55] Danya Bachir, Ghazaleh Khodabandelou, Vincent Gauthier, Mounim El Yacoubi, and Jakob Puchinger. Inferring dynamic origin-destination flows by transport mode using mobile phone data. *Transportation Research Part C: Emerging Technologies*, 101:254–275, 2019.

[56] Yuxiong Ji, Rabi G Mishalani, and Mark R McCord. Transit passenger origin–destination flow estimation: Efficiently combining onboard survey and large automatic passenger count datasets. *Transportation Research Part C: Emerging Technologies*, 58:178–192, 2015.

[57] Matthew Dunlap, Zhibin Li, Kristian Henrickson, and Yinhai Wang. Estimation of origin and destination information from Bluetooth and Wi-Fi sensing for transit. *Transportation Research Record*, 2595(1):11–17, 2016.

[58] Neveen Shlayan, Abdullah Kurkcu, and Kaan Ozbay. Exploring pedestrian Bluetooth and WiFi detection at public transportation terminals. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pages 229–234. IEEE, 2016.

[59] Bin Yu, Zhong-Zhen Yang, Peng-Huan Jin, Shan-Hua Wu, and Bao-Zhen Yao. Transit route network design-maximizing direct and transfer demand density. *Transportation Research Part C: Emerging Technologies*, 22:58–75, 2012.

[60] Pierre-Léo Bourbonnais, Catherine Morency, Martin Trépanier, and Éric Martel-Poliquin. Transit network design using a genetic algorithm with integrated road network and disaggregated O-D demand data. *Transportation*, pages 1–36, 2019.

[61] Yinhi Wang, B De Schutter, T van den Boom, B Ning, and Tao Tang. Origin-destination dependent train scheduling problem with stop-skipping for urban rail transit systems. In *Transportation Research Board 93rd Annual Meeting*, 2014.

[62] Yu Zhou, Yun Wang, Hai Yang, and Xuedong Yan. Last train scheduling for maximizing passenger destination reachability in urban rail transit networks. *Transportation Research Part B: Methodological*, 129:79–95, 2019.

[63] Sigal Kaplan, Dmitrijs Popoks, Carlo Giacomo Prato, and Avishai Avi Ceder. Using connectivity for measuring equity in transit provision. *Journal of Transport Geography*, 37:82–92, 2014.

[64] Malvika Dixit, Subeh Chowdhury, Oded Cats, Ties Brands, Niels van Oort, and Serge Hoogendoorn. Examining circuity of urban transit networks from an equity perspective. *Journal of Transport Geography*, 91:102980, 2021.

[65] Wei Wang, John P Attanucci, and Nigel HM Wilson. Bus passenger origin-destination estimation and related analyses using automated data collection systems. *Journal of Public Transportation*, 14(4):7, 2011.

[66] Claude Chriqui and Pierre Robillard. Common bus lines. *Transportation Science*, 9(2):115–121, 1975.

[67] Héctor Cancela, Antonio Mauttone, and María E Urquhart. Mathematical programming formulations for transit network design. *Transportation Research Part B: Methodological*, 77:17–37, 2015.

[68] Marcela A Munizaga and Carolina Palma. Estimation of a disaggregate multimodal public transport origin-destination matrix from passive smartcard data from Santiago, Chile. *Transportation Research Part C: Emerging Technologies*, 24:9–18, 2012.

[69] Georgios Laskaris, Oded Cats, Erik Jenelius, Marco Rinaldi, and Francesco Viti. Multiline holding based control for lines merging to a shared transit corridor. *Transportmetrica B: Transport Dynamics*, 7(1):1062–1095, 2019.

[70] Yiming Bie, Ruru Tang, and Linhong Wang. Bus scheduling of overlapping routes with multi-vehicle types based on passenger OD data. *IEEE Access*, 8: 1406–1415, 2019.

[71] Mark Dimond, David Brenig-Jones, and Neil Taylor. Exploiting OpenStreetMap topology to aggregate and visualise public transport demand. In *Proceedings of the 25th Annual Conference of GISRUK, 18th-21st April 2017*, 2017.

[72] Philippe Fortin, Catherine Morency, and Martin Trépanier. A methodology to evaluate supply changes in public transit networks using GTFS data. Technical Report 17-03004, Transportation Research Board, 2017.

[73] Hyun Kim and Yena Song. Examining accessibility and reliability in the evolution of subway systems. *Journal of Public Transportation*, 18(3):6, 2015.

[74] Oded Cats. Topological evolution of a metropolitan rail transport network: The case of Stockholm. *Journal of Transport Geography*, 62:172–183, 2017.

[75] Juan de Oña, Rocío de Oña, Laura Eboli, and Gabriella Mazzulla. Index numbers for monitoring transit service quality. *Transportation Research Part A: Policy and Practice*, 84:18–30, 2016.

[76] Corinne Mulley and Chinh Ho. Evaluating the impact of bus network planning changes in Sydney, Australia. *Transport Policy*, 30:13–25, 2013.

[77] Michael Graehler, Richard Alexander Mucci, and Gregory D Erhardt. Understanding the recent transit ridership decline in major US cities: Service cuts or emerging modes. In *98th Annual Meeting of the Transportation Research Board, Washington, DC*, 2019.

[78] Joe Grengs. Measuring change in small-scale transit accessibility with Geographic Information Systems: Buffalo and Rochester, New York. *Transportation Research Record*, 1887(1):10–17, 2004.

[79] Ahmed El-Geneidy and David Levinson. Mapping accessibility over time. *Journal of Maps*, 3(1):76–87, 2007.

[80] Hyun Kim, Keumsook Lee, Jong Soo Park, and Yena Song. Transit network expansion and accessibility implications: A case study of Gwangju metropolitan area, South Korea. *Research in Transportation Economics*, 69:544–553, 2018.

[81] Kelly Bertolaccini. An analysis of changes to transit accessibility and equity after the opening of a bus rapid transit system in Hartford, Connecticut. *Journal of Transport and Land Use*, 11(1):1163–1171, 2018.

[82] John C Handley, Lina Fu, and Laura L Tupper. A case study in spatial-temporal accessibility for a transit system. *Journal of Transport Geography*, 75:25–36, 2019.

[83] Belinda M Wu and Julian P Hine. A PTAL approach to measuring changes in bus service accessibility. *Transport Policy*, 10(4):307–320, 2003.

[84] Ehab Diab, Jamie DeWeese, Nick Chaloux, and Ahmed El-Geneidy. Adjusting the service? Understanding the factors affecting bus ridership over time at the route level in Montréal, Canada. *Transportation*, pages 1–22, 2020.

[85] Christoffer Weckström, Rainer Kujala, Miloš N Mladenović, and Jari Saramäki. Assessment of large-scale transitions in public transport networks using open timetable data: case of Helsinki metro extension. *Journal of Transport Geography*, 79:102470, 2019.

[86] Nicole Foth, Kevin Manaugh, and Ahmed M El-Geneidy. Towards equitable transit: examining transit accessibility and social need in Toronto, Canada, 1996–2006. *Journal of Transport Geography*, 29:1–10, 2013.

[87] Amr Mohammed, Amer Shalaby, and Eric J Miller. Empirical analysis of transit network evolution: case study of Mississauga, Ontario, Canada, bus network. *Transportation Research Record*, 1971(1):51–58, 2006.

[88] Nicole Foth, Kevin Manaugh, and Ahmed M El-Geneidy. Determinants of mode share over time: how changing transport system affects transit use in Toronto, Ontario, Canada. *Transportation Research Record*, 2417(1):67–77, 2014.

[89] Hannah Bast, Patrick Brosi, and Sabine Storandt. Real-time movement visualization of public transit data. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 331–340, 2014.

[90] Postsavee Prommaharaj, Santi Phithakkitnukoon, Merkebe Getachew Demissie, Lina Kattan, and Carlo Ratti. Visualizing public transit system operation with GTFS data: A case study of Calgary, Canada. *Heliyon*, 6(4):e03729, 2020.

[91] Thomas Jeffrey Kimpel. Data visualization as a tool for improved decision making within transit agencies. Technical Report TNW2006-14, Transportation Northwest, University of Washington, 2007.

[92] Harvey J Miller, Shih-Lung Shaw, et al. *Geographic information systems for transportation: principles and applications*. Oxford University Press on Demand, 2001.

[93] Abdullah Kurkcu, Fabio Miranda, Kaan Ozbay, and Claudio T Silva. Data visualization tool for monitoring transit operation and performance. In *2017 5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, pages 598–603. IEEE, 2017.

[94] Colin Stewart, Ehab Diab, Robert Bertini, and Ahmed El-Geneidy. Perspectives on transit: Potential benefits of visualizing transit data. *Transportation Research Record*, 2544(1):90–101, 2016.

[95] Antoine Giraud, Martin Trépanier, Catherine Morency, and Félix Légaré. Data fusion of APC, smart card and GTFS to visualize public transit use. Technical report, CIRRELT, Centre interuniversitaire de recherche sur les réseaux d'entreprise, 2016.

[96] Google. General Transit Feed Specification Reference. `https://developers.google.com/transit/gtfs/reference/`, 2020. Accessed on July 9, 2020.

[97] Transportation Research Board. *Fast-Tracked: A Tactical Transit Study*. The National Academies Press, Washington, DC, 2019. doi: 10.17226/25571. URL `https://www.nap.edu/catalog/25571/fast-tracked-a-tactical-transit-study`.

[98] Mobility Data. Best Practices for GTFS. `http://gtfs.org/best-practices/`, 2020. Accessed on July 9, 2020.

[99] Valhalla. Valhalla - Open Source Routing Engine for OpenStreetMap. `github.com/valhalla`, 2020. Online; accessed on 2020-10-08.

[100] National Association of City Transportation Officials. *Transit Street Design Guide*. Island Press, 2016.

[101] Peter G Furth, Burak Cesme, and Tarannum Rima. Signal priority near major bus terminal: Case study of Ruggles Station, Boston, Massachusetts. *Transportation Research Record*, 2192(1):89–96, 2010.

[102] Anthony F Han and Nigel HM Wilson. The allocation of buses in heavily utilized networks with overlapping routes. *Transportation Research Part B: Methodological*, 16(3):221–232, 1982.

[103] Gabriel E Sánchez-Martínez, Laurel Paget-Seekins, Christopher W Southwick, and John P Attanucci. Bus load inference and crowding performance evaluation through disaggregate analysis of fare transaction, vehicle location, and passenger count data. *Transportation Research Record*, 2672(8):464–474, 2018.

[104] Avishai Ceder. *Public transit planning and operation: Modeling, practice and behavior*. CRC Press, 2016.

[105] MA Brovelli, M Minghini, and ME Molinari. An automated GRASS-based procedure to assess the geometrical accuracy of the OpenStreetMap Paris road network. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 7, 2016.

[106] Google Developers. Encoded polyline algorithm format. https://developers.google.com/maps/documentation/utilities/polylinealgorithm, 2021. Online; accessed on 2021-03-24.

[107] Massachusetts Bay Transportation Authority. MBTA GTFS Archive. https://github.com/mbta/gtfs-documentation, 2021. Online; accessed on 2021-03-25.

[108] Johnathan Belcher. Changes to Transit Service in the MBTA district, 1964-2021. http://roster.transithistory.org/MBTARouteHistory.pdf, 2021. Online; accessed on 2021-04-30.

[109] Federal Transit Administration. The National Transit Database (NTD). https://www.transit.dot.gov/ntd, 2021. Online; accessed on 2021-03-29.

[110] Massachusetts Bay Transportation Authority. MBTA Bus Ridership by Trip, Season, Route/Line, and Stop. https://mbta-massdot.opendata.arcgis.com/datasets/eec03d901d2e470ebd5758c60d793e8e_0, 2021. Online; accessed on 2021-04-30.

[111] Krishna N. S. Behara, Ashish Bhaskar, and Edward Chung. A novel approach for the structural comparison of origin-destination matrices: Levenshtein distance. *Transportation Research Part C: Emerging Technologies*, 111:513–530, Feb 2020.

[112] Avishai Polus. Modeling and measurements of bus service reliability. *Transportation Research*, 12(4):253–256, 1978.

[113] Massachusetts Bay Transportation Authority. See how your bus route changed on September 1. https://www.mbta.com/projects/better-bus-project/update/, 2021. Online; accessed on 2021-04-01.

[114] Massachusetts Bay Transportation Authority. See How Some Bus Routes Changed on December 22, 2019. https://www.mbta.com/projects/better-bus-project/update/see-how-some-bus-routes-changed-december-22-2019, 2021. Online; accessed on 2021-04-04.

[115] Christina Kakderi, Nicos Komninos, Anastasia Panori, and Eleni Oikonomaki. Next City: Learning from cities during COVID-19 to tackle climate change. *Sustainability*, 13(6):3158, 2021.

[116] Transportation Research Board. TCRP G-18: Improving access and management of transit ITS data. https://apps.trb.org/cmsfeed/TRBNetProjectDisplay.asp?ProjectID=4687, 2021. Online; accessed on 2021-03-29.