# Improving High-Frequency Transit Reliability: A Case Study of the MBTA Green Line Through Simulation and Field Experiments of Real-Time Control Strategies

by

Joshua Javier Fabian

Bachelor of Science in Civil and Environmental Engineering
Massachusetts Institute of Technology, 2015

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 2017

Author . . . . . . **Signature redacted** . . . . . . . . . . . . . . . . . . . . . .
Department of Civil and Environmental Engineering
May 19, 2017

Certified by . **Signature redacted** . . . . . . . . . . . . . . . . . . . . . .
John P. Attanucci
Research Associate of Civil and Environmental Engineering
Thesis Supervisor

Certified by . . . . **Signature redacted** . . . . . . . . . . . . . . . . . . . .
Frederick P. Salvucci
Research Associate and Senior Lecturer of Civil and Environmental Engineering
Thesis Supervisor

Certified by . . . **Signature redacted** . . . . . . . . . . . . . . . . . . . .
Gabriel E. Sánchez-Martínez
Research Associate and Lecturer of Civil and Environmental Engineering
Thesis Supervisor

Accepted by . . . . . **Signature redacted** . . . . . . . . . . . . . . . . . . . .
Jesse Kroll
Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee
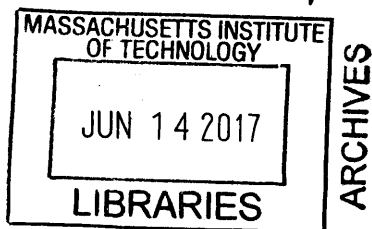
# Improving High-Frequency Transit Reliability: A Case Study of the MBTA Green Line Through Simulation and Field Experiments of Real-Time Control Strategies

by

Joshua Javier Fabian

Submitted to the Department of Civil and Environmental Engineering
on May 19, 2017, in partial fulfillment of the
requirements for the degree of
Master of Science in Transportation

## Abstract

Service reliability is a major concern for public transportation agencies. Transit services experience natural variability in scheduled service, due to factors such as traffic congestion, irregular demand, multi-route and branching corridors, and operator behavior. This variability leads to irregular headways, resulting in longer passenger waits and decreased effective capacity as gaps in service form. Real-time control strategies allow controllers to intervene at terminals and en route to regulate headways and improve performance.

This research tests the effectiveness of holding control strategies on the Massachusetts Bay Transportation Authority (MBTA) Green Line in Boston, a complex, four-branch light rail line. A simulation model is developed to estimate and compare the benefits of different schedule-based and headway-based holding strategies. Dispatching trains at terminals to target headways is found to minimize wait time, and the addition of en route holding improves service further, albeit slightly.

The simulation results inform the design of a field experiment, in which headway-based dispatching is implemented at a Green Line branch terminal. Terminal personnel are provided with tablet computers showing departure times optimized by an even-headway policy. When optimized departure times are adhered to, peak-hour headway variability is reduced by 40%. The average wait is shortened by 15% (30 seconds), and the 90th percentile wait is shortened by 21% (90 seconds). Compliance with the recommended departure times in the experiment was hampered by various human factors and station features. During the experiment, only 49% of trips left within 45 seconds of the departure times recommended by the algorithm. These results show that adopting headway-based dispatching at terminals promises significant benefits to service and passengers if operational changes are accompanied by improved supervision practices. This research fully supports the idea that transit agencies, such as the MBTA, should allocate supervisory resources for high-frequency services to prioritize terminal headway control versus en route and schedule-based strategies.

Thesis Supervisor: John P. Attanucci
Title: Research Associate of Civil and Environmental Engineering

Thesis Supervisor: Frederick P. Salvucci
Title: Research Associate and Senior Lecturer of Civil and Environmental Engineering

Thesis Supervisor: Gabriel E. Sánchez-Martínez
Title: Research Associate and Lecturer of Civil and Environmental Engineering

# Acknowledgments

# Contents

# List of Figures

14

# List of Tables

# Related publications by author

Fabian, J. J. and G. E. Sánchez-Martínez (2016, August). An algorithm for the processing and cleaning of automatic vehicle location data from multiple sources. In *2nd International Workshop on Automated Data Collection Systems*, Boston, Massachusetts.

Fabian, J. J. and G. E. Sánchez-Martínez (2017). Simulation-Based Comparison of Holding Strategies for a Multi-Branch Light-Rail Service. *Transportation Research Record: Journal of the Transportation Research Board* (2648). Forthcoming. doi:10.3141/2648-03

Prokosch, A. and J. J. Fabian (2017, April). Using public real-time data for operational improvements: the Green Line Headway Pilot. In *2017 New England ITS Annual Interchange*, Boston, Massachusetts. http://www.newenglandits.org/annual-meeting

# Chapter 1

# Introduction

The quality of a public transportation service's operations is critical to its ability to meet the demand of passengers traveling to and from places of residence, employment, education, and retail. An ineffective or low quality transit service can be one that does not meet the expectations listed in a schedule, or one that passengers cannot rely on to transport them in a consistent time and manner. This research uses the increasingly common real-time vehicle location data in order to transit improve service from the perspective of both the agency and the passenger. The Massachusetts Bay Transportation Authority (MBTA) is used as a case study for a simulation model testing control strategies and for implementing a decision-support tool for terminal dispatching on its light rail line, the Green Line. The Green Line operates as a complex, high-frequency line over four branches and is projected to experience increases in demand in the future as the Boston region grows and as the line is extended to the north from its current easternmost terminal, Lechmere (Malikova, 2012). Since trains already run at capacity during portions of the peak periods, any growth in demand is likely to further increase crowding and decrease service reliability, unless actions are taken to increase service regularity. While changes to routes and schedules can planned and implemented, service variability will still be experienced due to a number of factors. Traffic congestion, issues with rolling stock, or a spike in demand at a station can all make the best-scheduled transit service unreliable. Ensuring a high quality of service requires strategies for deviating from the scheduled plan and controlling vehicles in real time.

This introduction provides a motivation for the research, describes real-time control actions available to transit agencies, outlines the research objectives, and summarizes the

approaches taken to meet those objectives. This chapter also contains a description of the Green Line and its current operations and a literature review.

## 1.1 Research motivation in transit operations

This research is motivated firstly by an interest of transit authorities to run high quality service, to minimize unexpected operating costs, and to meet established performance standards. Technological advancements now provide the feasible means for agencies to track their rolling stock in granular detail and in real time, record and store the information digitally, and send related information to dispatchers on devices ranging from desktop computers to mobile phones. Finally, there is an incentive to improve transit reliability in order to spur economic growth. This is particularly true in places with relatively high transit utilization, such as Boston.

### 1.1.1 Service reliability

Improving or maintaining reliability is a chief concern for transit agencies. Reliability also ranks very high among the priorities of transit riders, higher than station and vehicle amenities (Iseki and Taylor, 2010). Those who depend on transit for daily commuting, for example, expect to arrive at work on time every morning.

From the perspective of agencies, *service reliability* can be broadly defined as the degree to which the published schedule reflects reality. High-frequency transit lines are typically scheduled at regular intervals, or *headways*, even if the dispatching of the lines is strictly based on the schedule rather than the headway. From the perspective of a commuter, service reliability may be defined as the ability to reach a location—such as a place of employment—in a predictable duration of time.

The reliability of transit lines is a function of the variability of in-vehicle and dwell times, themselves affected by weather, variability in demand, differences in operator behavior, and interference within the right of way (by pedestrians, autos, and other transit vehicles). Poor reliability results in irregular headways, which are manifested in long service gaps and vehicle bunching. Vehicle bunching occurs when one bus or train runs with delays, experiencing long dwell times as it encounters more passengers than typical, and a following vehicle, running early or on time while experiencing short dwell times due to the shorter gap (fewer

passengers are found to be waiting for the bus or train), eventually meets the leader. The two or more vehicles will usually run together as a platoon with a very short headway between them, unless a special action is taken to separate them and control service. The unreliability of a transit line, as a whole, can make for an inefficient use of capacity during periods when headways become uneven and inconsistent, as passengers become unequally distributed among operating vehicles (Fellows, 1990). Passenger wait times and vehicle crowding increase.

### 1.1.2 Technology advancements

The availability of real-time vehicle location data on transit lines has increased, allowing for more precise decision-making and closer monitoring of vehicle and operator activity. Through the use of cellular and Wi-Fi networks, information can be sent to mobile devices in a variety of formats, as simple as showing a map of train locations or as complex as a system recommending service interventions.

In Boston, recent improvements to the Green Line rolling stock and track infrastructure allow the agency—and the public, through the use of mobile applications—to know the location of all trains in real time. Over the past few years, the number of automatic vehicle identification (AVI) and track circuit points within the trunk (the Central Subway) has been expanded so that trains can be detected entering and exiting stations. Above ground, most trains now contain Global Positioning System (GPS) transmitters that send location data as often as every six seconds. Detailed demand information for every origin-destination pair and time of day in the system, even if unavailable in real time, can now be estimated using the Origin, Destination, and Transfer (ODX) inference model by Gordon et al. (2013) and Sánchez-Martínez (2017b). These advances enable improving the work of Malikova (2012) and other researchers, as well as expanding analyses to areas of the line that previously lacked accurate tracking.

### 1.1.3 Economic importance of transit reliability

A high reliability and positive public perception of a transit system also yields economic benefits. Improved public transportation usually results in ridership increases from a higher willingness to utilize it over other travel modes such as private automobiles or ridesharing services. This, in turn, yields environmental and public health benefits, and increases equity

in a metropolitan region, since those who are transit dependent realize improved mobility (Litman, 2017). New urban developments may be built more efficiently if less parking is required, and a reliance on transit can benefit businesses built at a human level (American Public Transportation Association, 2007).

These benefits are likely to be reduced if transit is unreliable. In Fiscal Year 2016, the Washington Metropolitan Area Transit Authority (WMATA) experienced a ridership decrease of 6.0%, largely attributed to safety and reliability concerns following a deadly smoke incident, and service disruptions caused by an ongoing initiative to rebuild much of the system (Anosike and Thomas, 2016). In Boston, a (failed) bid to host the 2024 Summer Olympic Games listed a number of MBTA improvements that would have been required to support the event, including decreased headways on the heavy rail lines and increasing the number of cars on Green Line trains from two to three (Bazelon et al., 2015). After a series of blizzards crippled MBTA services for several weeks, support for the Olympic bid plummeted (Levenson and Arsenault, 2015).

The Green Line will be extended into the cities of Somerville and Medford in the early 2020s (see Figure 1-1, which shows the extension in red). This is a capital project aiming to spur economic growth in the Boston region by creating jobs through improved transportation access, encouraging transit-oriented development, and reducing auto traffic on nearby highways and arterials, which itself has positive air quality impacts (Massachusetts Department of Transportation, 2016). The actual benefits may be limited, however, if the service proves ineffective for its intended passengers due to the unreliability potentially caused by lengthening routes from their current patterns. Ensuring service reliability, including through the use of real-time control strategies, is therefore important for protecting the economic growth potential of new capital investments in transit, such as the Green Line Extension.

## 1.2 Real-time control strategies

A *real-time control strategy* refers to a type of deviation from a scheduled service plan that is implemented in order to improve the service quality of a transit route. A single *control action* is ordered and implemented when service deteriorates and becomes inconsistent due to the factors described in Section 1.1.1. The decision-making process for control can be based on information from sources ranging from previous observations of vehicles at a static
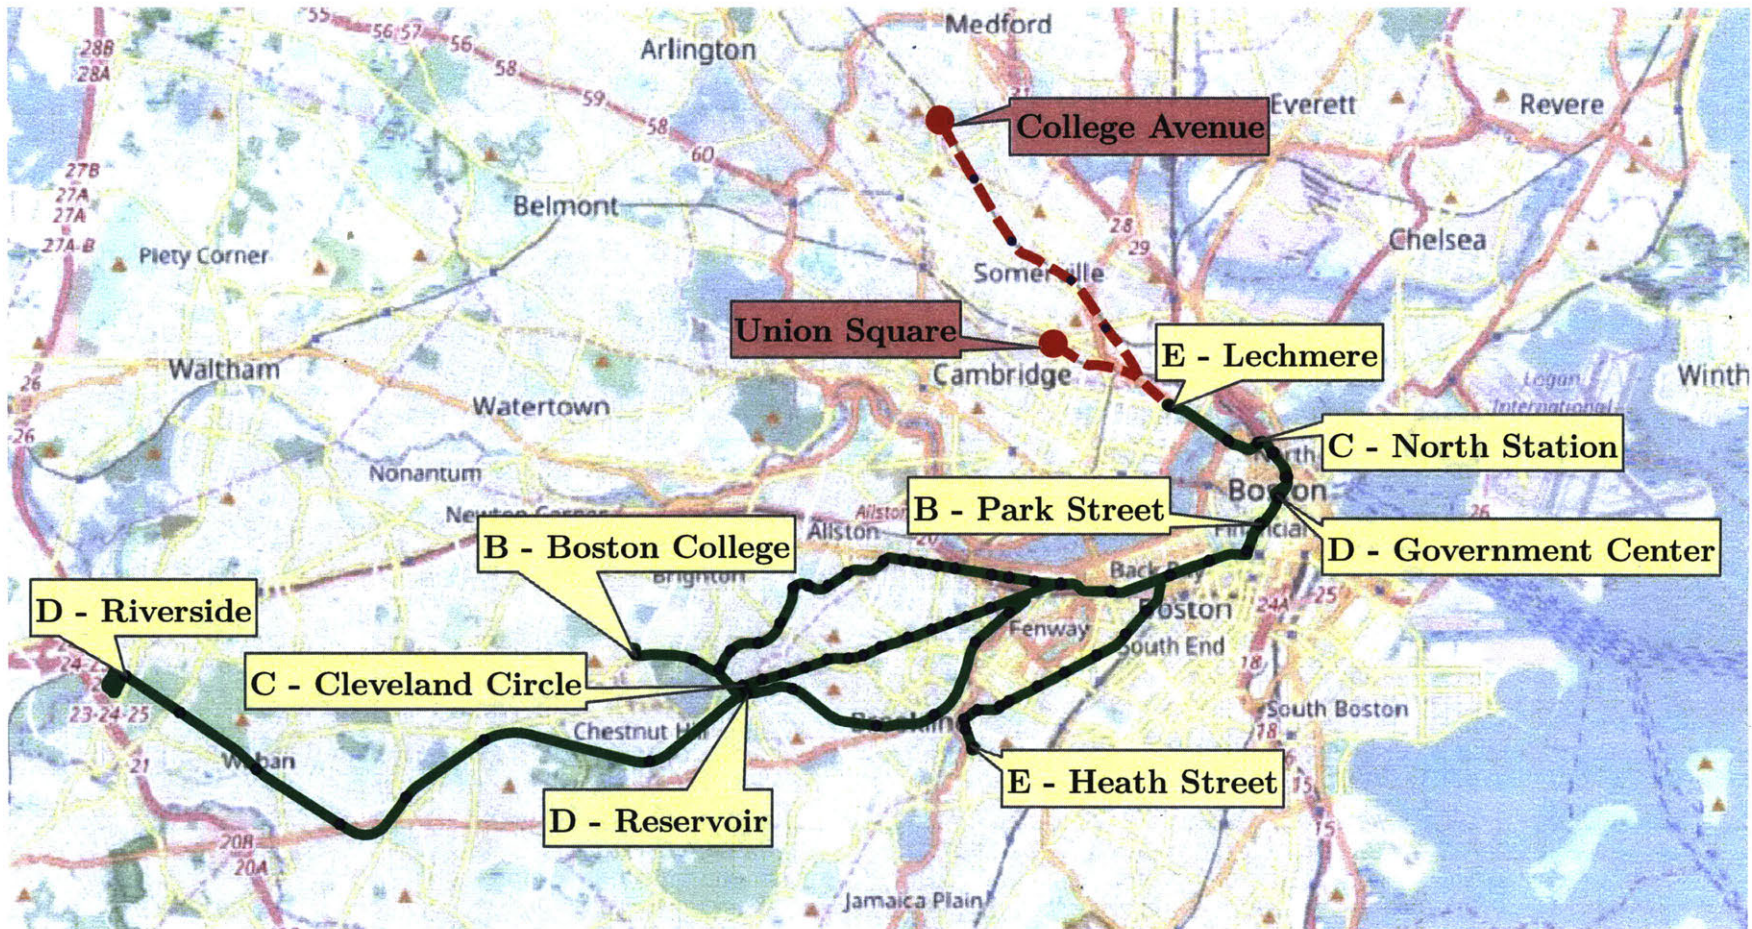
Sources: Office of Geographic Information (MassGIS) (2014a) and OpenStreetMap contributors (2017).

Figure 1-1: Map of MBTA Green Line, including the Green Line Extension to Somerville (Union Square) and Medford (College Avenue).

point to vehicle location data displaying the live locations of vehicles on a digital display.

Several real-time control strategies are briefly described below. This research focuses on holding, including the dispatching of vehicles at terminals.

## 1.2.1 Holding

*Holding* occurs when a bus or train operator is asked to remain stopped at a station for a duration longer than required for passengers to alight and board the vehicle.[1] Holding occurs with the intent of improving service reliability by bringing the vehicle into adherence with either a scheduled time or headway. Holding can implemented at one or more stations along a service, including its terminals, in which case it becomes part of *dispatching*, or the sending of crews and vehicles from terminals. It can be based on schedules, headways, or other criteria. Some passengers may experience shorter wait times downstream of a holding point due to improved reliability, though others may have to sit or stand through the duration of the hold, lengthening their journey times. Much of this thesis will focus on vehicle holding as a control strategy.

In *schedule-based holding*, a vehicle running ahead of its published schedule must stop and wait at assigned timepoints until the scheduled departure time. Vehicles running on or behind schedule are not held. The strategy is convenient to apply because the only information required is the vehicle's schedule—no real-time information about other activity on the line is needed. Operators can execute this strategy on their own with little to no outside guidance and supervision.

*Headway-based holding*, on the other hand, aims to keep vehicles evenly spaced apart. Emphasis is on the regularity of service rather than punctuality, which may or may not be feasible to keep. Headway-based policies are beneficial for high-frequency routes because passengers generally arrive at origin stops at random, basing their experience on the known, scheduled headway. There exist different headway-based holding strategies, such as the *target-headway strategy*, which prevents a departure until the scheduled headway has elapsed since the previous train's departure, and the *even-headway strategy*, which holds a train until it is equally-spaced between its leader and follower, among others (Barnett, 1974, Turnquist, 1981). Headway-based holding requires the use of real-time automatic vehicle location (AVL)

---

[1] Holding can occur at locations away from stations, but this is uncommon due to the inconvenience this poses to passengers without the ability to alight.

data for maximum effectiveness, although target-headway based holding could be executed by direct observation.

Strategies based on a *rolling-horizon optimization* incorporate a cost model with inputs including not only vehicle locations but also demand and downstream arrival predictions (Delgado et al., 2012, Sánchez-Martínez et al., 2016b). The cost model may contain one weight for headway regularity and another for the cost of operator overtime. The goal of these strategies is to optimize holding times based on the specifics of the cost model, although an optimization model could also decide between other control action types. This requires weighting different factors that may influence the decision for or against a control action.

### 1.2.2 Other strategies

A *short-turn* is the termination of a transit service at a station prior to its scheduled terminal. When a vehicle short-turns, it begins its reverse trip at the control point downstream from the scheduled starting location. Short-turning can be implemented for a vehicle when service in its opposite direction experiences a large gap, which may coincide with bunching in the initial direction (Deckoff, 1990). In this situation, passengers from two bunched vehicles are consolidated into one, and service is improved in the opposite direction, with wait times decreased for those passengers at and downstream of the revised origin point. Different criteria can inform decisions as to whether short-turning is worthwhile for a particular situation. Criteria can include schedules, operator costs, headways, passenger load, and estimated demand, depending on what information can be made available to decision-makers in real time.

*Expressing* a vehicle means to have it skip one or more regularly-scheduled station stops in order to decrease its running time. Like short-turning, this is a tactic which can be used for vehicles which are running behind schedule or have both a large gap behind and a short gap in front, depending on the specific criteria used. While it may benefit passengers onboard the train heading for destinations beyond the express area, there is a burden to passengers destined for one of the skipped stops who are forced to alight the train and wait for the next arrival making local stops (Macchi, 1989, Wilson et al., 1992). On certain transit lines, expressing can mean that a vehicle on a circuitous route can take a more direct path to its destination, or that a train can switch to a special track which does not run adjacent

to station platforms, allowing for faster speeds.

*Deadheading* is similar to expressing except that the vehicle operates to a specified location out of revenue service, carrying no passengers. This practice may be experienced near terminals in order to fill a large service gap.

A number of other real-time control strategies can comprise an agency's toolkit for managing service. These can include instructing operators to drive vehicles at reduced speeds, limiting the duration doors are opened for at stations, extending a trip one or more stations past its scheduled terminal, and temporarily switching a train's branch entirely (Delgado et al., 2012). Like the more common strategies described above, these can be used in conjunction with other actions in order to give supervisors a wide range of options for managing service unreliability in real time.

## 1.3   Objectives

This research aims to examine what reliability improvements a transit agency can feasibly expect on a high-frequency service with real-time control technology by reviewing current operations, simulating changes in control policies, and piloting and evaluating a field implementation of a decision-support system. Several goals in this research work toward this overarching objective:

1   Reviewing current service to codify existing real-time control and service monitoring practices.

2   Developing data processing algorithms to account for incomplete, erroneous, and incompatible data collected automatically in real time.

3   Using simulation to evaluate potential impacts of real-time control and to motivate a field experiment.

4   Developing a decision-support system to provide personnel with optimized dispatching times given the real-time state of service.

5   Conducting a field experiment to measure the actual impacts of real-time control and learn how human factors affect the effectiveness of strategies when compared to simulation.

6   Recommending the best steps forward for those interested in implementing a real-time control system.

## 1.4   Research approach

This research begins with the development of a simulation model of a complex light rail transit line for evaluating and comparing the effects of different control strategies. The results of the simulation suggest that even simple control strategies are promising for improving service reliability, and motivate field experiments for testing control strategies in real conditions. A decision-support tool is developed to enable real-time optimization of headways on an actual transit line. The results of the pilot are analyzed. The MBTA Green Line is used as a case study.

### 1.4.1   Simulation modeling

Simulation modeling is used to recreate the MBTA Green Line in an environment resembling current operations, but with added or modified control strategies. Different strategies, and different locations for implementing strategies, are simulated, and their effects on service and passenger experience are compared to a baseline model. The results shed light into how and where the greatest benefits can be achieved, and inform the design of the field experiment.

### 1.4.2   Field experimentation

A field experiment is conducted to test the validity of the simulation results, along with the hopes of tangibly improving service for Green Line passengers. Based on the conclusions of the simulation modeling, an even-headway based dispatching scheme at the Riverside terminal is piloted. Quantitative data on schedule and headway-recommendation adherence are collected and statistics on service characteristics such as headways and running times are calculated. Qualitative data on operator behavior and the reactions to the piloted technology are also collected. The results are analyzed to determine the real impact of the control strategy and to learn about the role of outside factors on the execution and effectiveness of service control.

### 1.4.3 MBTA case study context

This research uses the MBTA Green Line as a case study, with the field experiment focusing on the Riverside branch. The Green Line makes for a suitable case study due in part to its complexity: it comprises four independently-dispatched branches featuring different types of right of way, lengths, and demand patterns. The light rail line is described further in the next section.

The Green Line is also suitable for use as a case study because it has been studied previously. Previous Green Line research has focused both on operations management and infrastructure improvements.

## 1.5   MBTA Green Line

This section describes the case study used throughout this research, the MBTA Green Line. The Green Line is a complex, multi-branch light rail system with a recently upgraded vehicle tracking system allowing for more detailed analysis along segments than was previously possible. In addition to describing the tracking system, an overview of current operations is given.

### 1.5.1   Background

The MBTA Green Line is a four-branch light rail system which covers 37 kilometers (23 miles) of two-way track in the cities and towns of Boston, Cambridge, Brookline, and Newton, and carries over 225,000 passenger trips per weekday (Massachusetts Bay Transportation Authority, 2014). Figure 1-2 shows the official schematic of the MBTA Rapid Transit and Key Bus Routes system, which includes the Green Line. The schematic hides the fact that the four branches of the Green Line vary greatly in track length and operating environments. Figure 1-1 overlays the Green Line on a geographic map of the Boston area for reference. The map shows the upcoming Green Line Extension in red, which will add two branches to the northeast and require a reconfiguration of service patterns (Sindel, 2017). In the Central Subway, each of the four current branches terminates at a different station. The Riverside branch, for instance, runs entirely on a grade-separated, exclusive right of way, whereas a portion of the Heath Street branch operates in the roadway with mixed automobile traffic. The four branches are described by segment in Table 1.1.

28

Source: Central Transportation Planning Staff (2016)

Figure 1-2: Diagram of the MBTA rapid transit system and major connecting bus routes.

Table 1.1: Green Line branches with operating environment by segment.

| Branch | Letter | Branch segment | Right of way operating environment |
|---|---|---|---|
| Boston College | B | Boston College—Blandford St | Protected street median |
| | | Blandford St—Park St | Dedicated subway |
| Cleveland Circle | C | Cleveland Circle—St Mary's St | Protected street median |
| | | St Mary's St—North Station | Dedicated subway |
| Riverside | D | Riverside—Fenway | Dedicated at-grade railway |
| | | Fenway—Government Center | Dedicated subway |
| Heath Street | E | Heath St—Brigham Circle | In-street with mixed traffic |
| | | Brigham Circle—Northeastern | Protected street median |
| | | Northeastern—North Station | Dedicated subway |
| | | North Station—Lechmere | Dedicated elevated railway |

The Green Line typically operates trains with two cars, each featuring about 45 seats, with a policy capacity of approximately 100 passengers (Massachusetts Bay Transportation Authority, 2014). Depending on the branch, service is scheduled every 5-7 minutes during the peak periods and every 7-12 minutes off-peak. The flow of trains is regulated in the Central Subway and along the surface portion of the Riverside branch by a fixed block signaling system, which is not enforced by any automatic train control (ATC) system (Malikova, 2012). The rest of the Green Line, including all stations, is run as line-of-sight operations with minimal signalization. Terminal operations are facilitated using loops and crossover turnbacks. Turnbacks and pocket (or stub) tracks at other locations on the Green Line can be used to switch the direction of trains during emergencies or other disruptions. Figure 1-3 shows a track schematic of the Green Line and depicts loops, turnbacks, and pocket tracks at which service can terminate and turn if needed.

### 1.5.2 Vehicle location tracking

Real-time AVL data with high spatial granularity has recently become available for the Green Line with the installation of GPS units onto the rolling stock. Location data comes from a tracking system with three different sources: AVI, track circuits, and GPS.

AVI transmitters and receivers make up the oldest existing tracking system, originally installed wayside between 1987 and 1990 in the Central Subway and at a handful of locations on the surface level (Fellows, 1990). Every Green Line car contains two AVI transmitters (for bidirectional travel) coded with vehicle information which is set to the route provided
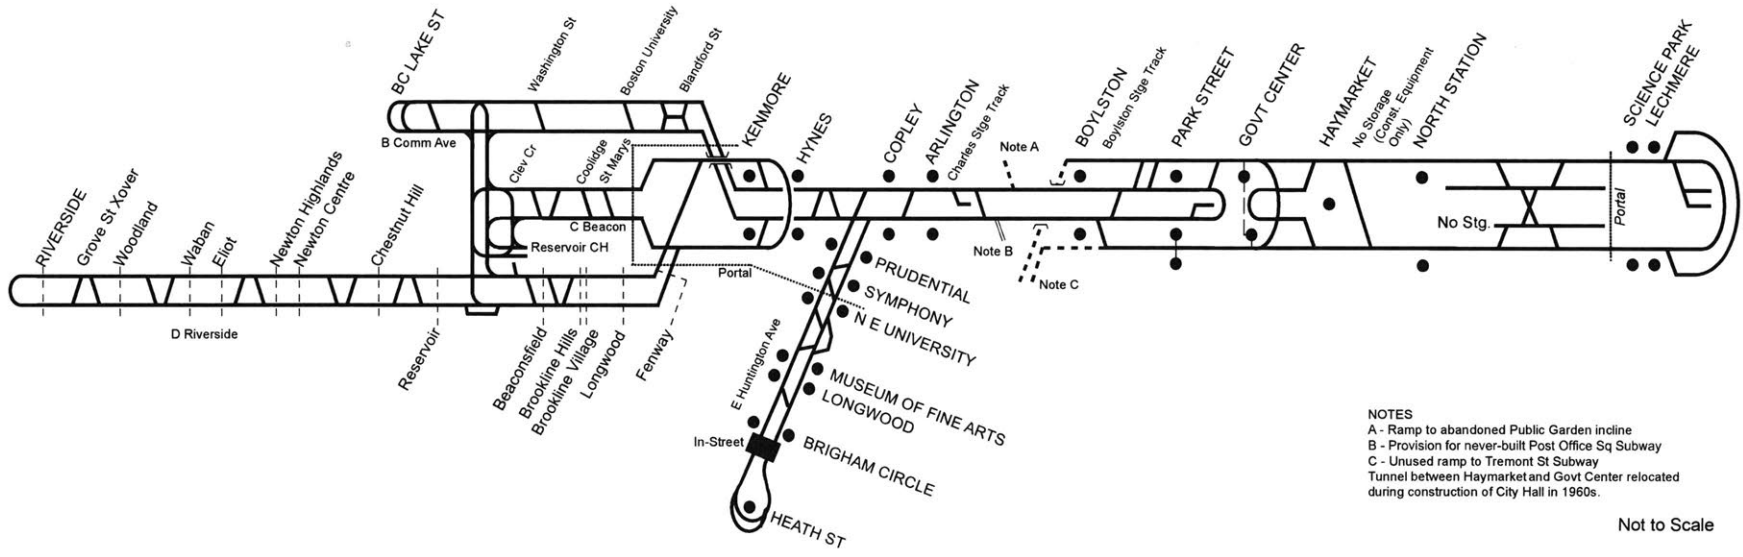
Original source: Massachusetts Bay Transportation Authority (2014). Modified by author.

Figure 1-3: Green Line track schematic.

through manual input by the vehicle operator. Whenever a train passes a wayside AVI receiver, several bits of information are shared between the train and the receiver, containing the identification numbers of all cars and a code corresponding to the train's current routing. AVI receivers are used not only for vehicle tracking but, perhaps more importantly, also for controlling switches at junctions and crossovers and setting the status of wayside signal lights near those switches. Such switches are controlled electronically with a train's route code. The active system of AVI vehicle tracking is susceptible to a variety of data issues, with car numbers misread or extra cars inserted into the reported consist information. Nevertheless, the quality of data from the AVI system has improved in the last five years with the addition of new AVI points. At every Central Subway station, trains are now detected both when arriving at and departing from the stop, allowing dwell times to be analyzed. Data quality issues are discussed further in Chapter 2.

Track occupancy circuits are used on certain stretches of the Green Line trunk in lieu of AVI infrastructure. A track circuit works simply as an electrical circuit which is closed and powered except when a train occupies the track segment. Thus, these circuits detect the presence of a train indirectly, and no information on a train's consist or routing information is transmitted.

The third and newest method of tracking Green Line trains uses GPS. Preparation for the system began in 2013, and real-time location data from GPS-equipped trains began to be published by the MBTA in October 2014 (Barry and Card, 2014). Of the two types of Green Line rolling stock in service today, the newer stock, Type 8 cars, have GPS transponders and cellular routers installed. Similar to the AVI transmitters, the cellular routers are linked to a car number for identification. Routing information input to the AVI system cannot be transmitted by the router, however. The older stock, Type 7 cars, are in the process of having the equipment installed as part of a vehicle overhaul project (Massachusetts Bay Transportation Authority, 2015). At least one car per train needs to have working GPS in order for the entire train to be tracked. When they are above ground west of Kenmore or Symphony stations, trains send their locations via a commercial 3G network every six seconds while in motion. GPS-based locations are also geofenced to predefined locations along the track.

All three sources are unified before vehicle location data are published, internally to the Operations Control Center (OCC) and other applications, and externally through the

real-time General Transit Feed Specification (GTFS-realtime) and MBTA-realtime, an application program interface (API) open for public consumption. These two public sources provide, for MBTA buses and trains, the latest locations of vehicles, as well as predictions of when they will reach downstream stops. The MBTA-realtime API is used for obtaining the current state of the Green Line in subsequent chapters.

### 1.5.3 Current operations

An understanding of current operations on the Green Line is critical for attempting to improve service. Depending on the specific location, the light rail line is controlled either by scheduled departure times or by the scheduled headway. *Dispatchers* and *field inspectors* control operations from the OCC, branch and trunk terminals, and at other stations along the line.[2]

The OCC acts as the central dispatching point for the MBTA Red, Orange, Blue, Green, and Mattapan Lines. Two dispatchers in the OCC are dedicated to the light rail lines (the Green and Mattapan Lines) at all times. They monitor the real-time locations of trains on large-screen monitors. The schematics, as exemplified in Figure 1-4, show the last-recorded location of each train and its routing (coded as the color of the lead car number). Yellow or red-colored flags can also appear at each station, informing dispatchers that a significant period has elapsed since the last train passed on the respective branch. A separate screen, shown in Figure 1-5, lists details from the last twenty trains to pass AVI points as defined by the user. Both of these tools are powered using an internal source of location data. In these ways, dispatchers also keep track of headways at stations and make decisions to express, reroute, or hold certain trains in order to regulate operations. The real-time information available at the OCC lacks operator information, which dispatchers often need in order to determine whether a potential control action may make train drivers late for their scheduled relief. When such information is needed, dispatchers call either the operator directly by radio or the terminal inspector responsible for the operator by telephone. Other than when they check on crew constraints, OCC dispatchers do not keep track of whether trains are running on schedule, as monitoring is generally headway-based.

OCC dispatchers are considered the final authority above field inspectors for implement-

---

[2]The general term "official" is often used interchangeably with "inspector" within the MBTA to refer to any type of supervisory inspector operating in the field. This thesis uses the term "inspector" to avoid confusion when discussing agency management.
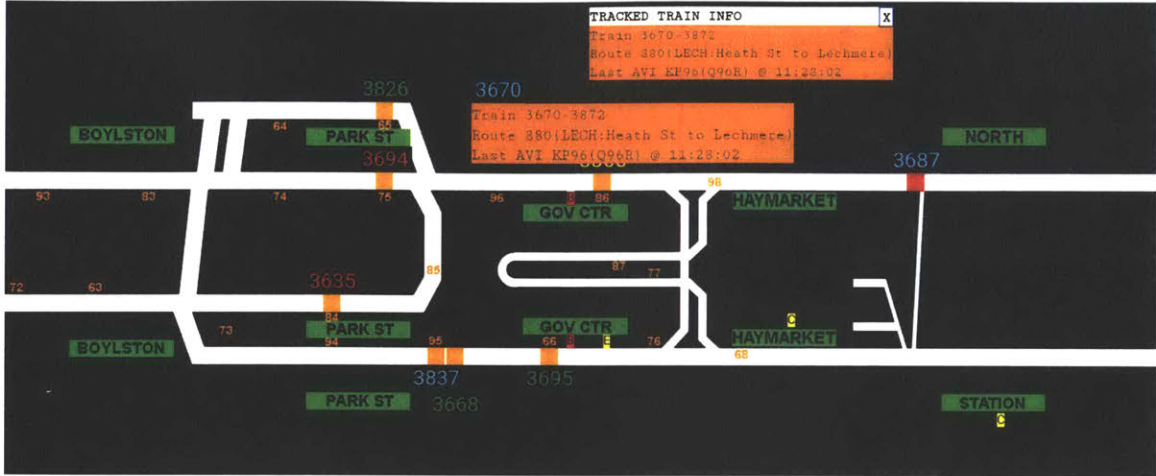
Figure 1-4: Green Line vehicle location tracking interface available to OCC dispatchers and inspectors at select locations.
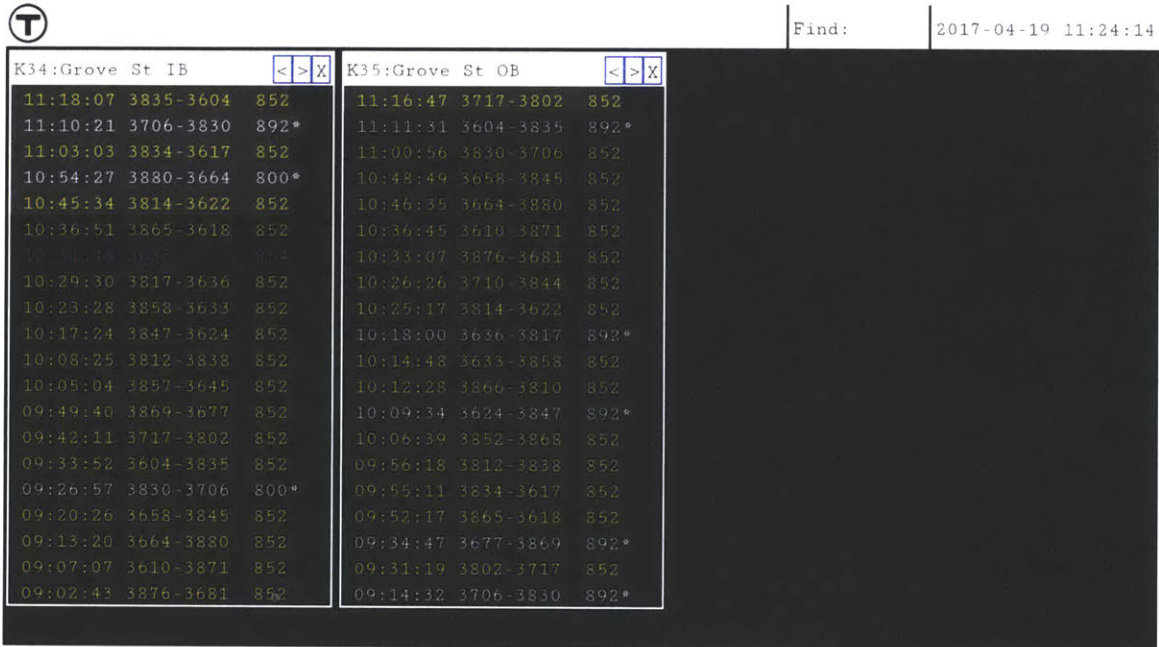


Figure 1-5: Interface allowing OCC dispatchers and select inspectors to view the times and other information of recent vehicles passing selected points.

ing service changes. As will be discussed below, field inspectors can and often do implement adjustments to service, usually minor in nature. Before any such adjustments can take place, however, they must call and obtain permission from the OCC (by radio or telephone).[3] In this sense, the role of dispatchers is akin to that of coordinators: they may, at different times, make calls to ensure that inspectors are acting cooperatively, approve a request to hold a train for several minutes, relay messages about a passenger's lost personal belongings, arrange for medical services to tend to a sick customer, or inform operators of the presence of maintenance workers in the right of way. The presence of two dispatchers on the Green Line, though not allowing for dedicated monitoring of each branch, does allow for multiple issues to be handled simultaneously. In the case of a major disruption, each dispatcher can focus on different aspects of the response.

Terminal inspectors are positioned at the dispatched end of each branch, including Boston College, Cleveland Circle, Riverside, and Lechmere, where there is layover space for trains (Massachusetts Bay Transportation Authority, 2014).[4] These are also the reporting locations for operators. The primary responsibilities of terminal inspectors are to ensure that operators report ready for duty and that trains are dispatched from the terminal.

When operators report for shifts, they check in with the terminal inspector on duty, at least ten minutes before their first trip of the day and at least two minutes before their next trip if they are returning from a break. The inspector confirms that the operator is fit for duty. If an operator reports an absence, the inspector is responsible for finding a substitute driver for the open trips, typically drawing from a pool of extra board operators on hand for this purpose (known as the "cover list").

Terminal inspectors dispatch trains primarily by following the schedule. They refer to printed paper train sheets that include all scheduled departures and arrivals for the day, their scheduled times, and the operators due for each time, and complete it with actual times as they occur. These sheets serve as the only official record matching operators to actual departures. An example of an incomplete train sheet is shown in Figure 1-6. Inspectors aim to have trains depart as close to the scheduled departure time (under the column "trip time" in Figure 1-6), and can take actions such as pulling additional trains from the yard, reordering operators, or drawing from the cover list to ensure on-time performance.

---

[3]Terminal inspectors are an exception, as discussed below.

[4]Terminal inspectors are often referred to as "pull-out officials" or "starters." A pull-out is the act of pulling a train from a rail yard and placing it in revenue service.

These decisions are made while considering the current fleet and crew constraints. In some cases, adjustments to departure times are made, such as by telling operators to leave a minute or two before or after scheduled. Trips can also be canceled due to a vehicle or operator shortage or other disruption, and in these cases, inspectors may adjust the times of adjacent trips so that headways are more even. Dispatching adjustments made by terminal inspectors, except for trip cancellations, do not require the consent or notification of the OCC. Terminal inspectors also decide when trains should be pulled into the yard (resulting in increased headways), and are in charge of validating operators' claims for overtime pay.

Other inspectors are posted at non-terminal stations, or may be assigned to roving shifts along a branch. Common non-terminal stations for inspectors include, but are not limited to, Reservoir, Harvard Avenue, Kenmore, and Boylston. The primary responsibility for these inspectors is to ensure that passing trains are running at the appropriate headway. For example, if the Riverside branch is scheduled to run every five minutes and two Riverside branch trains arrive two minutes from each other, the trailing train may be held at the station for three minutes. To assist in service regulation, inspectors at Kenmore and Boylston have access to a computer workstation in their platform-level booths with access to the same vehicle location screens available to OCC dispatchers.

En route, instructions may be recommended by the mid-line inspector and approved by the OCC by radio, or they may be ordered by the OCC and passed down by telephone to the inspector, who then relays the instruction to the train operator upon his or her arrival. Mid-line inspectors may carry train sheets similar to those used by terminal inspectors, but these are used only to record train arrivals rather than to dispatch them. The roles of mid-line inspectors at surface-level stations may, at times, extend beyond direct service control, with such tasks as off-board fare validation to speed boarding and reduce dwell time.

Inspectors and OCC dispatchers together make up the staff in charge of maintaining Green Line operations. As real-time data first became available for the Green Line, an information asymmetry developed, with maps of real-time train locations available only to central dispatchers. This partly resulted in the strengthening of the role of the OCC dispatchers (Fellows, 1990). Computer workstations have been installed at several locations in recent years, but the only terminal currently equipped with the ability to view real-time train locations is Lechmere. At terminals, knowing when trains will be arriving is important for making the best decisions on running trips and pulling trains from the yard. Traditionally,

03/07/2017   WEATHER _____   INSPECTOR _____

| CAR NUMBER | ROUTE | PULL OUT | PREV TRIP | DUE RIVER | ACTUAL ARRIVAL | NEXT TRIP | TRIP TIME | ACTUAL TIME | IDENT | RUN | SUB | BADGE | COMMENTS | SIGNATURE | ELECTRONIC DEVICE CHECK |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| _____ | 852- | | | | | | 829 | A | 094 | 0545 | | | | | |
| _____ | 852-T | | | | | | 829 | A | 122 | 0546 | | | | | |
| _____ | 852- | | 653 | 832 | _____ | 845 | | | 098 | 0557 | | | | | |
| _____ | 852-T | | 653 | 832 | _____ | 845 | | | 126 | 0558 | | | | | |
| _____ | 852- | | | | | | 835 | A | 096 | 0586 | | | | | |
| _____ | 852-T | | | | | | 835 | A | 124 | 9022 | | | | | |
| _____ | 852- | | 659 | 839 | _____ | | | | 100 | 9023 | | | | | |
| _____ | 852-T | | 659 | 839 | _____ | | | | 128 | 0502 | | | | | |
| _____ | 852- | | | | | | 840 | A | 097 | 1425 | | | | | |
| _____ | 852-T | | | | | | 840 | A | 125 | 1426 | | | | | |
| _____ | 852- | | 704 | 844 | _____ | 856 | | | 102 | 0594 | | | | | |
| _____ | 852-T | | 704 | 844 | _____ | 856 | | | 130 | 0570 | | | | | |
| _____ | 852- | | | | | | 845 | A | 098 | 0557 | | | | | |
| _____ | 852-T | | | | | | 845 | A | 126 | 0558 | | | | | |
| _____ | 852- | | 709 | 849 | _____ | 901 | | | 103 | 9027 | | | | | |
| _____ | 852-T | | 709 | 849 | _____ | 901 | | | 131 | 0521 | | | | | |
| _____ | 852- | | | | | | 851 | A | 100 | 0502 | | | | | |
| _____ | 852-T | | | | | | 851 | A | 128 | 9023 | | | | | |
| _____ | 852- | | 715 | 855 | _____ | 907 | | | 105 | 0535 | | | | | |
| _____ | 852-T | | 715 | 855 | _____ | 907 | | | 133 | 9031 | | | | | |
| _____ | 852- | | | | | | 856 | A | 102 | 0594 | | | | | |
| _____ | 852-T | | | | | | 856 | A | 130 | 0570 | | | | | |
| _____ | 852- | | 720 | 900 | _____ | | | | 106 | 0563 | | | | | |
| _____ | 852-T | | 720 | 900 | _____ | | | | 134 | 0583 | | | | | |
| _____ | 852- | | | | | | 901 | A | 103 | 9027 | | | | | |

Figure 1-6: Sample train sheet for use by terminal inspectors at Riverside terminal

37

train locations have been determined by the terminal inspector using radio communications with train operators, and terminal arrivals are then estimated. With GPS tracking now available for most trains, most terminal inspectors observed during this research regularly use third-party applications on personal smartphones to inform their decision-making. This practice is neither sanctioned nor prohibited by the MBTA.

## 1.6 Literature review

This research builds upon previous work in several areas. This section contains a review of literature on work related to real-time holding strategies, transit simulation, decision-support tools for transit applications, and the MBTA Green Line.

### 1.6.1 Holding as a real-time control strategy

The issue of controlling high-frequency transit in real time has been studied at length, although typically in the context of bus services, and only more recently using agent-based simulation. Barnett (1974) and Turnquist (1981) conducted early work related to holding, creating a model for minimizing wait times on platforms and in vehicles, and proposing the *Prefol* strategy of holding, which equalizes leading and trailing headways. Turnquist (1981) additionally found that passengers view the quality of service in terms of either schedule reliability or headway adherence, with passengers on routes with scheduled headways greater than 10 minutes timing arrivals to stops in an effort to minimize waiting, and riders on high-frequency services arriving to stops at random, suggesting that reducing headway variability could best-serve those latter passengers. Eberlein et al. (2001) formulated the analytical holding problem with the availability of real-time vehicle locations and found that the optimal holding solution depends more on prevailing headways than on passenger demand. Cats et al. (2011) simulated target headway-based, Prefol-based, and schedule-based strategies for controlling a bus rapid transit (BRT) trunk line, finding the Prefol method superior when the objective was to minimize total cost. Focusing on a single-line light rail service, van Oort et al. (2010) showed that service could be improved by limiting holding times to reduce the cascading effect of en route holding, while noting that headway-based holding could increase passenger waits when few people travel through the control point.

## 1.6.2 Transit simulation modeling with holding

Transit simulation modeling in holding research has also progressed over the past decade. Bartholdi and Eisenstein (2012) tested a strategy involving only limited holding at terminals, allowing headways to equalize without manual intervention; the strategy was trialled using a simulation containing passenger demand per half-hour and uniform station-to-station running times. Research by Berrebi et al. (2015) used stochastically-drawn travel times and passenger destinations with headway dynamics to optimize holding for passenger wait time with as little recovery time as possible. Delgado et al. (2012) formulated a rolling horizon optimization model to generate holding and boarding limit policies sensitive to capacity constraints, and showed their strategy outperforming heuristics in a simulation experiment with stationary running times and demand. Sánchez-Martínez et al. (2016b) extended this model to capture time-varying running times and demand, and showed further performance increases in simulated service. Sánchez-Martínez et al. (2015) applied the dynamic model to simulated transit subject to event-driven dynamics. Hernández et al. (2015) simulated control strategies on a multiple-service corridor, the only such study found in this review. The authors found that holding on the scheduled headways of the two lines separately proved worse for wait times than without control, and that holding on combined headways yields the greatest user time savings.

## 1.6.3 Decision-support systems in transit operations

Decision-support systems have been proposed for use in public transport operations for a number of years, with fieldwork focusing either on simply providing dispatchers with increased information or providing recommendations to operators in-vehicle.

Fellows (1990) describes the installation of an AVI system for the MBTA Green Line and conceptualizes a user interface for a decision-support system. The concept allows the use of a desktop computer to click on a displayed train, select a potential control strategy, and see the estimated costs and benefits (including from vehicle delays and crowding) of the action. The attention that would need to be paid to the entire system would be reduced with an exception-based interface that highlights, on a map, anticipated problems with overcrowding, schedule adherence, or headway adherence. Finally, Fellows documents the shift of supervisory roles on the Green Line in the mid-1980s when an analog event recorder

used to measure headways was relocated from a downtown station to the OCC. Carnaghi (2014) provides a theoretical framework for decision-support systems in tram operations and investigates the causes of service unreliability.

The need for developers of decision-support tools to engage with prospective users from the point of project conception is emphasized by Tribone et al. (2014). The research finds a stakeholder-involved development process critical for gaining the acceptance of dispatchers. An increasing number of transit agencies, if not using decision-support tools, are setting the building blocks for future implementation by setting performance standards, calculating the metrics in real time, displaying them on internal dashboards, and releasing the data to the public (Tribone et al., 2016). Carrel et al. (2010) write that tools should be developed to deliver only relevant, real-time information for decision-making, rather than attempt to automate service control. The researchers determine that dispatchers take into account crew and vehicle constraints, which real-time control algorithms have not successfully incorporated.

Pangilinan et al. (2008) trialled the use of a handheld tablet displaying a map with vehicle locations on a lengthy Chicago bus route and recorded a 21% decrease in headway variability. The researchers, like Fellows (1990), found that exception-based triaging is necessary for preventing controllers from needing to filter out unnecessary data. In Stockholm, in-vehicle displays provided bus operators with headway adherence information and instructions based on an even-headway holding strategy (Cats, 2016). Passenger travel time was reduced by 10%, although the experiment also included other improvements, such as bus lanes. In Santiago, the rolling horizon-based simulated model by Delgado et al. (2012) was used in a BRT pilot, also using in-vehicle consoles (Lizana et al., 2014). The average cycle time was reduced by 8% and the standard deviation of cycle time was cut in half. Maltzan (2015) reduced waiting times on a high-frequency bus line by one minute after providing terminal supervisors with a tablet recommending departure times based on a target-headway dispatching strategy. Both Lizana et al. (2014) and Maltzan (2015) found that operator and supervisor compliance requires attention before the full benefits of algorithm-based recommendations can be achieved.

## 1.6.4 MBTA Green Line

This work adds to a large body of research using the Green Line in Boston as a case study. Most of the prior research focuses on one particular aspect of Green Line operations or planning, such as studying the length of trains or evaluating a specific service control option.

Real-time control on the Green Line has been a popular subject. The practice of expressing trains along parts of the route as a control strategy was modeled by Macchi (1989) as the AVI system was being installed. Deckoff (1990) focused on the short-turning of trains at a downtown terminal and determined that, without real-time vehicle location information, many instances of the practice yield few positive or even negative benefits. Soeldner (1993) synthesized the work of the previous authors by documenting several control actions commonly used on the Green Line, created a model to calculate the passenger benefits possible from each strategy. He determined that, at the time of the study, 20% of trips did not depart their terminals on time, and recommended heuristics to be used by Green Line inspectors or dispatchers for each strategy.

A system for monitoring the throughput of trains and passenger waiting times, described by Wilson et al. (1992), is reliant on a pre-AVI system of train circuits at terminals. The potential service control benefits of installing an information-rich AVL system is summarized by Wilson et al. (1992). Fellows (1990) further details decision-support tools that can incorporate the AVI system, as described in Section 1.5.2, while commenting on the changing roles of dispatchers and field inspectors on the Green Line. Wong (2000) proposes an AVL-based transfer coordination system, using the Green and Red Lines at Park Street as the case study.

Several previous works touch or focus on demand-related topics. ming Lin and Wilson (1992) create a dwell time model for Green Line stations within the Central Subway that explains 70% of the variation of dwells using crowding, boardings, and alightings as explanatory variables. Sindel (2017) evaluates demand estimates for the forthcoming Green Line Extension and makes recommendations for the next vehicle procurement process. This follows the work of Malikova (2012), who evaluated service reliability in relation to the inconsistent use of three-car trains on the Green Line. The research also calls for better headway management at terminals. Sánchez-Martínez (2017a) creates a stochastic model

41

for estimating the rates at which Green Line surface-level boardings do not interact with the onboard farebox machines and therefore are not included in fare transaction data. The demand used as an input for the simulation model in Chapter 3 comes from the ODX model as developed and implemented by Gordon et al. (2013), Dumas (2015), and Sánchez-Martínez (2017b).

Prior research enhances this work by providing readily-accessible data, details, histories, such as the dwell time model formulated by ming Lin and Wilson (1992) and the non-interaction factors by Sánchez-Martínez (2017a), which are incorporated into the simulation model specified in Chapter 3.

## 1.7 Thesis organization

This thesis is organized to align with the natural progression of this research as described in Section 1.4. This chapter provides a description of real-time control strategies, the objectives and motivation for the research, a summary of the case study used, and a literature review. Chapter 2 describes the data available from the MBTA and formulates methods to improve the data quality before it is used in further applications. Chapter 3 uses the cleaned data to simulate Green Line service as it currently operates and as it would operate if different control strategies were implemented. Chapters 4 and 5 deal with the field experiment in which trains are dispatched to maximize headway regularity from a Green Line terminal. Chapter 4 details the architecture and features of the decision-support software developed for the experiment, and Chapter 5 presents the experiment and analyzes results. Chapter 6 presents conclusions and recommendations, and discusses opportunities for future research.

# Chapter 2

# AVL data cleaning

This chapter describes an algorithm designed for cleaning archived automated vehicle location (AVL) data for the MBTA Green Line (itself described in Chapter 1). The algorithm is used to clean location data for use in the MBTA version of ODX, the MIT Transit Lab's origin, destination, and transfer inference model, which in turn informs demand rates when simulating Green Line passengers in Chapters 3 and 5. Strategies for cleaning AVL data in real time, necessary in the field experiment described in Chapters 4 and 5, are also discussed. The reading of this chapter is not required for the understanding of the subsequent chapters.

## 2.1   Importance of accurate location information

The accuracy of archived train location data is important for several reasons in the development of a simulation model for testing different control strategies. AVL data are used directly to obtain a distribution of inter-station running times as well as a minimum dwell time for each station. Vehicle locations are used as an input to ODX, which fuses AVL data with AFC data to infer demand patterns by origin, destination, and time of day (Gordon et al., 2013, Sánchez-Martínez, 2017b, Dumas, 2015). These demand rates are then used to simulate passenger activity. If a data problem with a particular vehicle is significant enough, a passenger's destination may not be inferred. Observed running and dwell times on the Green Line are used to validate baseline simulation models, so inaccuracies in the data are likely to be carried over without correction when vehicle movement is simulated.

Real-time data accuracy is crucial for making optimal service recommendations on the fly. In the field experiment discussed later in this thesis, departing headways from a ter-

minal are equalized using knowledge of the previous train to depart and estimation of the following train to arrive. Insufficient granularity in the data—both in spatial and temporal resolution—may lead to incorrect departure times being recorded. This could then erode the effectiveness of any strategy which uses inaccurate information.

## 2.2 Common AVL data issues

This section describes common issues and limitations of the AVL data sources available for the MBTA Green Line.

### 2.2.1 AVL data limitations

Archived vehicle location data are made available via the nightly transfer of AVL records from the MBTA's real-time data server to a data warehouse maintained by the MIT Transit Lab and accessed by MIT and MBTA researchers. The data are at the same level of detail that dispatchers and other Operations personnel at the MBTA can view internally in real time. It is processed lightly. As this data becomes available to MIT researchers on a nightly basis, however, it is not suitable for use in real-time applications.

The real-time data used in the field experiment comes from the publicly-available *MBTA-realtime*, upon which many other vehicle-tracking applications depend. MBTA-realtime is an application program interface (API) that can be called by software developers to output schedule information, real-time vehicle locations, real-time stop arrival predictions, and service advisories. Developers' requests can be for information related to a specific stop or vehicle, or for an entire line. A small amount of processing occurs prior to the publishing of vehicle locations, particularly on the Green Line, and there is typically a delay of about a minute between the time a vehicle passes a location and when the associated record is available through the API. Figure 2-2 depicts the flow of information via MBTA-realtime.

Neither data source contains the vehicle location data in its original, unprocessed form. Vehicle locations detected by GPS above ground are approximated to the centroid of one of a set of predefined *geofences*, or virtual geographic areas. Geofences for the Green Line stretch 60-100 feet in length along the track depending on the branch and specific location. This means that, when a train sends a location to the MBTA server, the coordinates read by downstream sources are those of the center of the geofence regardless of whether the train is
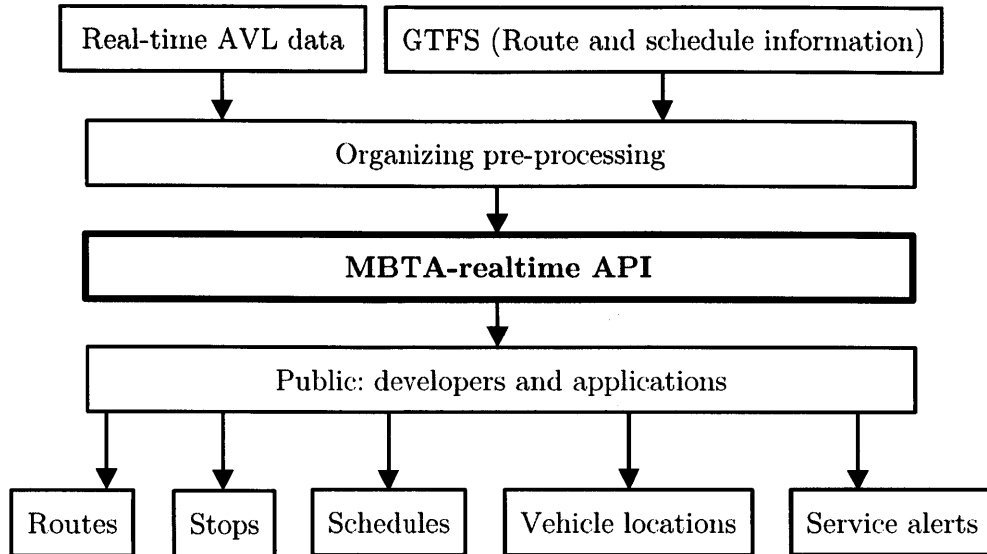
Figure 2-1: Flow of information to and from MBTA-realtime.

at one end of the area or the other. If a train is stopped on the edge between two geofences, there is the possibility of the wrong geofence being recorded, a concern at certain terminals. The main purpose of using geofences rather than raw locations is to have static, periodic locations, similar to existing AVIs and track circuits, for which arrivals can be reviewed using real-time software and analysis tools. Additionally, using geofences reduces bandwidth and storage requirements—a train which makes multiple pings within the same geofence, such as when dwelling at a station or red signal, only has the first timestamp recorded—and to facilitate the analysis of location data by combining areas into readily-identifiable points. For bandwidth reasons also, vehicles send their positions at most often once every six seconds. This means that a train recorded as having been within a particular geofence at 19:50:17 may have entered it as early as 19:50:11 or as late as 19:50:17. Furthermore, while every geofenced GPS record is available in the archived data, only half of the geofences on the Green Line are published to the MBTA-realtime API. While this is not an issue for most uses of the API, which may use stop arrival predictions to inform customers of their next vehicle, it does reduce the quantity and quality of information available to the decision-support tool described in Chapter 4.

The algorithm and techniques for cleaning Green Line AVL data discussed in the following section must work around these limitations.

## 2.2.2 Archived AVL data

A number of issues can reduce the quality of archived AVL data, all stemming from either incomplete or incorrect vehicle location records. A small percentage of trains do not carry GPS units. These are trains which only contain cars from the Green Line's oldest fleet (Type 7) which have not been recently refurbished—as the fleet is refurbished, the problem of trains without GPS units will continue to diminish. For these cases only 1-3 points are recorded on each branch, leaving most surface-level stations without arrival and departure time information.

Trains containing both GPS and AVI units are still susceptible to issues causing incorrect vehicle location records. While duplicate GPS records from a single train at the same location are prevented, the same cannot be said of records coming from track circuits and AVI points. These points have been observed to have a tendency to record some passing trains twice, making stop visit information ambiguous, or misread a bit on one of the cars, resulting in, for example, Car 3734 recorded as Car 3732 for one stop only. On occasion, a third car may be seen as having been added to a two-car train consist for a short duration—a car which, in reality, may have actually passed the location several minutes earlier. These issues can be caused by faulty wayside or onboard equipment, or by the MBTA's pre-processing methods. The onboard GPS units also cause error from time to time. While it remains unclear whether the issue is from a problem with the GPS units themselves or with the routers to which they are attached, certain trains may stop providing location data above ground for a while before sending all of the missing vehicle locations with the same timestamp—as if the vehicle had traveled several miles in several seconds.

The format of the archived train data is helpful in the cleaning algorithm but can itself be susceptible to issues. Each location record contains the vehicle consist as well as a trip identification number. This trip ID is assigned sequentially for each one-way trip on the Green Line from the start of each service day and is not related to trip IDs provided through GTFS. At times, the hardware or software-related issues mentioned above will change the consist information en route but not the trip ID; at other times, the consist information will not be affected while the trip ID changes; sometimes both the consist and the trip ID change en route. As a result, neither the vehicle nor trip ID fields can be used as a reliable key that stays constant throughout the course of a trip.

46

### 2.2.3 Real-time AVL data

Most of the common issues with archived data also apply in the real-time case with data from MBTA-realtime, as that information is derived from the MBTA servers providing internal data (sent as archives to MIT).

Working with data in real time, versus retrospectively, eliminates the possibility of looking at a consist's or trip's activity after an anomaly to determine how to handle the encountered train. For example, when a train previously shown with two cars is suddenly shown as two separate consists (with different trip IDs) with one car each, a decision must be quickly made as to whether or not the alleged splitting of the consist should be trusted.

There are also limitations on data collected from MBTA-realtime not discussed in the preceding subsections. These include the vulnerabilities of the publicly-facing MBTA-realtime API to high user demand. Outages in the API service were experienced several times during the experiment in Chapter 5, typically during the peak periods and lasting for a few minutes each time.

## 2.3 Archived data cleaning algorithm

The algorithm for cleaning archived Green Line AVL data, as described in the six steps below, is implemented in Java. The program is installed on the same server hosting the MIT-MBTA Data Warehouse and is scheduled to run every night after the AVL data are received from the MBTA. Figure 2-2 shows the place of the cleaning program in the flow of automatically-collected data from the transit agency to MIT. On the server, the program takes no more than two minutes to process, clean, and export one day's worth of Green Line data; on a personal computer using an Intel i7 processor, one day's run takes approximately four minutes.

### 2.3.1 Data filtering

The first and simplest step is to pare down the number of movements to be processed to the minimum required for the level of detail desired from the cleaned location data—arrival and departure times at each stop. The filtering process speeds up the algorithm and prevents bad imputation in the final step.

```
┌──────────────────┐  ┌──────────────────┐  ┌─────────┐
│ Active AVI points │  │  Track circuits  │  │   GPS   │
└──────────────────┘  └──────────────────┘  └─────────┘
         │                     │                  │
         ▼                     ▼                  ▼
┌───────────────────────────────────────────────────────┐
│           Organizing pre-processing by MBTA            │
└───────────────────────────────────────────────────────┘
                          │
                          ▼
┌───────────────────────────────────────────────────────┐
│          Cleaning algorithm upon MIT receipt           │
└───────────────────────────────────────────────────────┘
                          │
                          ▼
┌───────────────────────────────────────────────────────┐
│             Insertion in MIT-MBTA database             │
└───────────────────────────────────────────────────────┘
         │                     │                  │
         ▼                     ▼                  ▼
    ┌────────┐          ┌────────────┐     ┌───────────────┐
    │  ODX   │          │ Simulation │     │ Other analyses │
    └────────┘          └────────────┘     └───────────────┘
```
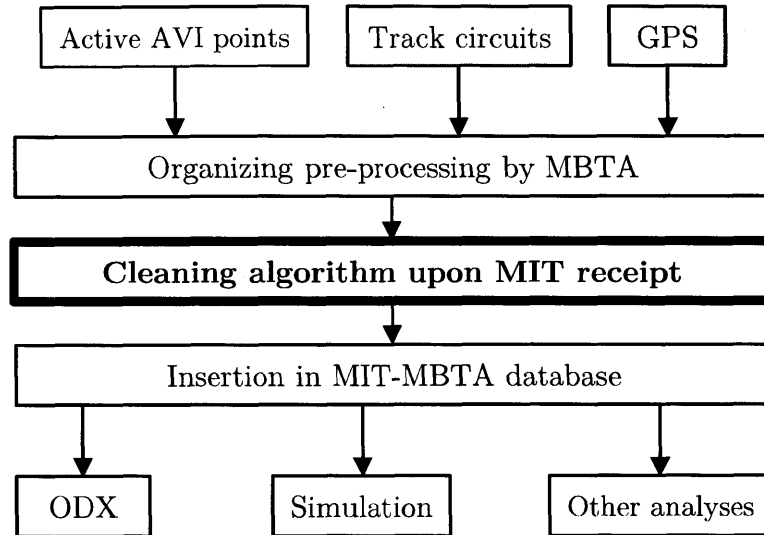
Figure 2-2: Flow of archived vehicle location data from MBTA Green Line trains to MIT.

In the remainder of this chapter, a *movement* is defined as a record of a vehicle's location with a timestamp. The filtering step sorts movements having the same trip identifier chronologically over the service day. Movements are then removed from the program (and not used in further cleaning) if they match any one of the following criteria:

- Movement occurs outside of the service day hours (02:00-04:40).

- Movement is a duplicate of another with the same consist, location, and timestamp.

- GPS geofence has a centroid in a yard or is otherwise away from revenue tracks.

- Movement is located at a AVI point or track circuit with a history of many misreads. Movements from nine such points across the Green Line are removed.

- Additionally, if any car numbers are repeated within a recorded consist, duplicate car numbers are removed.

### 2.3.2 Creation of movement formations

Movements are then grouped by the full consist if they occur within a short time of each other.

This is done by iterating over all valid car numbers $C$ and adding all movements which contain $c$ as one of the cars. Those movements are iterated chronologically. If the full

consist (usually $c$ and a second car) for a movement $m_1$ differs from that of the previous movement $m_0$, then a decision is made to include or delete all of the movements with the consist of the previous movement. The decision is made based on the density of the number of movements with consist $m_0$ to duration of consist $m_0$. If the density is too small, it means that vehicle locations with consist $c$ were not observed frequently and that it is likely that the car numbers are in error. These movements are then deleted from further cleaning unless the consist contains multiple cars all from the Type 7 fleet—potentially operating without GPS units. If the density is too large, it is due to an error on a GPS unit which has resulted in the recording of many different locations in a short period of time. These movements are deleted too.

Those movements, grouped by consist, which have acceptable movement count to duration densities, are made into formations, which are carried over to the next step. *Formations are groups of movements which contain the same consist and have not been broken by a movement containing a different consist with at least one common car number.*

### 2.3.3 Grouping of formations

Formations are ordered chronologically by the timestamp of their first movement, and then grouped by time and by car.

Grouping begins by taking the first formation (for the first train to depart a terminal at the start of the service day), and following it until its end at time $t$. At that point, we consider all other formations which begin shortly before or after $t$ which share at least one car with the first formation. Based on the specific issues observed in the data, the threshold is set to search for formations between $t - 5$ and $t + 10$ (in minutes). If such a train exists, it likely means that a data error occurred changing one of the car numbers temporarily. All movements from the second formation are merged with the first formation to create an enlarged formation. If no such train exists, it likely means that the train entered a yard and was not put back into revenue service until later in the day, if at all. In these cases, the enlarged formation is moved to a new list for further cleaning.

The process of finding the formation with the next earliest first movement is repeated until all formations have been added to an enlarged formation, not allowing any to be added to multiple enlarged formations. Any enlarged formation lasting for less than ten minutes in duration is not added to the collection, as there is not enough information contained to

infer an origin and destination for a one-way trip.

The result of this step is a collection of enlarged formations, each of which contains movements sharing at least one car and occurring temporally adjacent to one another. These enlarged formations are assumed to hold the locations and timestamps of a single consist for the rest of the algorithm.

### 2.3.4 Cleaning consist information

This step entails the editing of train consists for movements suspected to have incorrect car information. Such changes are made between different consists within each enlarged formation, which share at least one but not all cars. The identification and repairing of trains with suspect data uses heuristics based on the issues listed in the previous section and formulated after a close review of data containing quality issues.

First, enlarged formations are broken up into their constituent sub-formations by consist. The cleaning method employed on the enlarged formation depends on the number of sub-formations enclosed, with different methods for formations with one, two, three, or four or more sub-formations. Each method is described briefly below.

**Formation with one sub-formation**

- With only one consist, all movements are assumed to contain the correct vehicle information, and no changes are made.

**Formation with two sub-formations**

- If the earliest-starting sub-formation contains less than a threshold fraction of movements $x_2$ than the second sub-formation, the consist from the latter sub-formation is applied to all movements in the former. In this case, the second sub-formation is inferred to contain the correct consist for the entire formation. If the first sub-formation contains locations for only part of one trip, running time and demand analyses on the train may be impossible without this correction.

  - In this implementation, $x_2 = \frac{1}{8}$. This corrects cases when a consist is misread upon a train's pull-out from the yard and initial record at an AVI point, leading to one missing car in the data. The fraction is set large enough to not over-correct in

situations where a train is scheduled to operate in the morning as one car before a second car is scheduled to be joined to it later in the service day—this occurs on weekends.

- Otherwise, no changes are made to consist information within the different sub-formations.

## Formation with three sub-formations

- If the earliest-starting sub-formation contains the same consist as the third sub-formation and the number of movements in the second sub-formation is less than the product of a threshold fraction $x_3$ of the sum of movements in the first and third sub-formations, then the consist from the first sub-formation is applied to the second sub-formation. The consist of the second sub-formation is inferred to contain a short-lasting inaccuracy.

  - In the Green Line implementation, $x_3 = \frac{1}{10}$.

- Otherwise, if the number of movements in the first sub-formation is less than the product of a threshold fraction $x_3$ of the sum of movements in the second and third sub-formations, then the consist from the second sub-formation is applied to the first sub-formation. The consist of the first sub-formation is inferred to be in error because it appears for only a relatively short time.

- Otherwise, no changes are made to consist information within the different sub-formations.

## Formation with four or more sub-formations

- If the consist of the first sub-formation, $c_0$ contains only one car but the consists of the second and final sub-formations ($c_1$ and $c_{\text{final}}$, respectively) both contain at least two cars, and if the first sub-formation exists for less than thirty minutes, then the consist of the second sub-formation is applied to all movements in the first sub-formation. This is to correct instances where a train's consist is misread upon pull-out from the yard. The first sub-formation is inferred to contain an error in its apparent consist.

51

- The consist of the final sub-formation is not modified—it is assumed that it was recorded correctly.

- For all other sub-formations, the sub-formations are iterated. For each iterated sub-formation $s$, we also consider the sub-formations immediately before, $s - 1$, and after, $s + 1$.

  - For each $s$, if the consist $c_s$ contains only one car but $c_{s-1}$ and $c_{s+1}$ contain more than one car, and $c_{s-1} = c_{s+1}$, then we assume that $c_s$ was recorded in error and replace all the movements in $s$ with consist $c_{s-1}$. This handles cases in which one of a train's cars is not detected for a short time, but is detected before and after.

  - Otherwise, if the number of cars in $c_{s-1}$ equals that of $c_{s+1}$, the number of cars in $c_s$ is greater than one, $s$ exists for less than thirty minutes, the number of movements in sub-formation $s$ is less than the sum of movements in $s - 1$ and $s + 1$, consist $c_s$ is not the same as the consist in the first sub-formation and that of the final sub-formation, and either $c_{s-1}$ or $c_{s+1}$ equal $c_0$, $c_1$, or $c_{final}$, then a fix is applied. This is typically used to handle situations where a third car is inadvertently appended to the recorded consist of a two-car train. The inferred misreading must last for less than half an hour, and confidence in the accuracy of the consist before or after the misreading is assured by comparing it to the first, second, or final consist of the formation. The methods of handling this case are listed:

    * If $s$ exists for more than five minutes in duration, the time between the end of $s - 1$ and the start of $s$ is less than seven minutes, and the time between the end of $s$ and the start of $s + 1$ is also less than seven minutes, then $c_s$ is replaced with $c_{s-1}$.

    * Otherwise, with $s$ existing for such a short duration or the time between sub-formation existing for an excessive duration, the best guess is that the observations were made in error. Sub-formation $s$ and all movements contained within it are deleted. Any stop information removed is imputed in the final step.

For example, consider a formation containing three sub-formations. The first contains consist 3813-3658 for 52 movements. The second contains consist 3813-3658-3880 for 7 move-

ments. The third contains consist 3813-3658 for 145 movements. Since the second sub-formation contains fewer than $\frac{1}{10}(52 + 145) = 19.7$ movements, all movements have the third car 3880 removed. Another example is visualized in Figure 2-3. Part (a) shows several movements in a formation before cleaning; at 09:17:30, many movements were recorded with cars 12 and 38. Part (b) shows the formation during the cleaning phase. Car 38 is filled in at 09:17:18, the multiple movements at 09:17:30 are removed, and the movement containing car 20 at 09:18:12 is removed. The result is shown in part (c). If data are deleted for any stop visit made along the trip, times are reintroduced in the imputation step (Section 2.3.6).

After cleaning each formation of train movements, any formation lasting in duration for less than ten minutes or containing fewer than ten movements is removed.

### 2.3.5 Movement clustering around stops

The goal of this step is to reduce the vehicle location data to only the arrival and departure times at each station.

For each movement $m$ in each cleaned formation, the recorded coordinates $l_m$ are compared to a table of Green Line station platforms and their respective coordinates, and all stations within an $r$ radius of $l_m$ are noted. Based on the regularity of trains' GPS pings and the distances between stops, $r$ is set to 100 meters (or 328 feet). There exists a handful of locations on the Green Line where the distance between two stops is less than 200 meters, and in these cases some movements may have multiple stations recorded. For any single trip on a train, a collection of adjacent movements tagged with a station forms a *stop cluster*—in the cases where stations are located within 200 meters, a stop cluster may encompass all movements containing two or three stations. All movements not in proximity to any station are removed. Additionally, a check is made to ensure that adjacent clusters do not occur simultaneously, which would be indicative of a malfunctioning GPS unit onboard a vehicle; movements in clusters which fail the check are also removed. Stop clusters containing movements from multiple stations are then broken up into multiple clusters, each containing only the movements linked to a single station.

For each stop cluster, the movement located closest to the midpoint of the station platform is considered the arrival time of the train. For surface-level stations, the entire platform in each direction is considered in one geofenced area, so a train will only record locations while it is stopped at or passing a station. The following movement is selected for the

⋮
09:17:06 | Car 12 - Car 38
09:17:18 | Car 12
09:17:24 | Car 12 - Car 38
09:17:30 | Car 12 - Car 38
⋮
09:17:30 | Car 12 - Car 38
09:18:00 | Car 12 - Car 38
09:18:12 | Car 12 - Car 38 - Car 20
09:18:18 | Car 12 - Car 38
⋮

(a) Movements before cleaning.

⋮
09:17:06 | Car 12 - Car 38
09:17:18 | Car 12 *- Car 38*
09:17:24 | Car 12 - Car 38
*09:17:30 | Car 12 - Car 38*
*⋮*
*09:17:30 | Car 12 - Car 38*
09:18:00 | Car 12 - Car 38
*09:18:12 | Car 12 - Car 38 - Car 20*
09:18:18 | Car 12 - Car 38
⋮

(b) During the cleaning phase.

⋮
09:17:06 | Car 12 - Car 38
09:17:18 | Car 12 - Car 38
09:17:24 | Car 12 - Car 38



09:18:00 | Car 12 - Car 38

09:18:18 | Car 12 - Car 38
⋮

(c) Movements after cleaning.

Figure 2-3: Depiction of the cleaning of an example movement formation.

departure timestamp. At this point, all stop clusters contain a best estimate for both an arrival and a departure time. The different actions occurring in this step are depicted in Figure 2-4. In part (b), clusters of westbound movements are formed in 100-meter radii around Symphony and Northeastern University stations. As part (c) shows, for each cluster, the movement closest to the platform (in yellow) becomes the arrival time at the respective station. The following movement (in red) becomes the departure time from the station.

### 2.3.6 Imputation

The final step in the cleaning algorithm is to impute any missing data.

At this point, stop clusters still contain two possible stop locations: one for a westbound platform and another for an eastbound platform, both at the same parent station. For each train, stop clusters are iterated chronologically while consulting a table containing all possible direct routings from each stop. The parent stop for each cluster station $i$ is compared to the parent stop for the following cluster station $i + 1$, and the table is used to compute all possible paths from station $i$ to station $i + 1$. If there are multiple paths, the shortest path by number of stops traversed is selected. Figure 2-5 shows an example where two possible paths exist from station $i$ to station $i + 1$, between which no location data was recorded. In this case, path $b$ is shorter than path $a$, so it is selected. The three intermediate stations along path $b$ are added to a list of stations which the train visited but are missing from the location data—these stop visits need to be imputed.

Before stops are imputed, a check is made to ensure that imputation makes sense given the duration and distance between the two known stops. The distance between the two known stops is computed by summing the straight-line distances between adjacent stations along the imputed path. If the average speed between the known stations is greater than 17 meters/second (or 38 miles/hour)—the top speed on the fastest segment of the Green Line, along the Riverside branch—then imputation is not carried out. If the distance to be imputed is greater than 4 kilometers (or 2.5 miles), approximately the length of the Central Subway, the average speed greater than 13 meters/second (or 29 miles/hour), and the segment is not on the Riverside branch, then imputation is also blocked.

The remaining stops requiring imputation have timestamps filled in proportion to the distances of stations along the imputed section. For example, if the stop Harvard Avenue is to be imputed and Harvard Avenue is 300 meters from the preceding stop, Griggs Street,
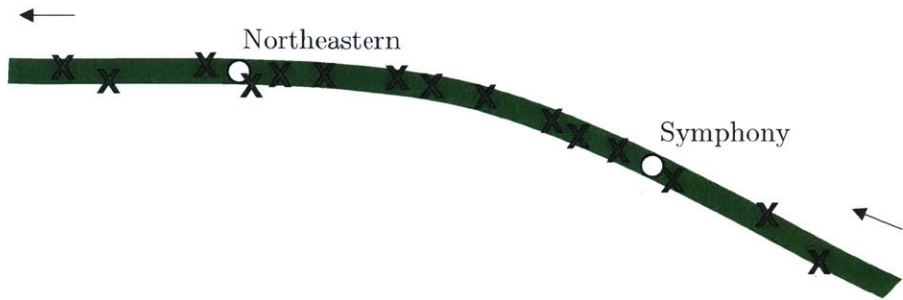
(a) Observed movements in the westbound direction.
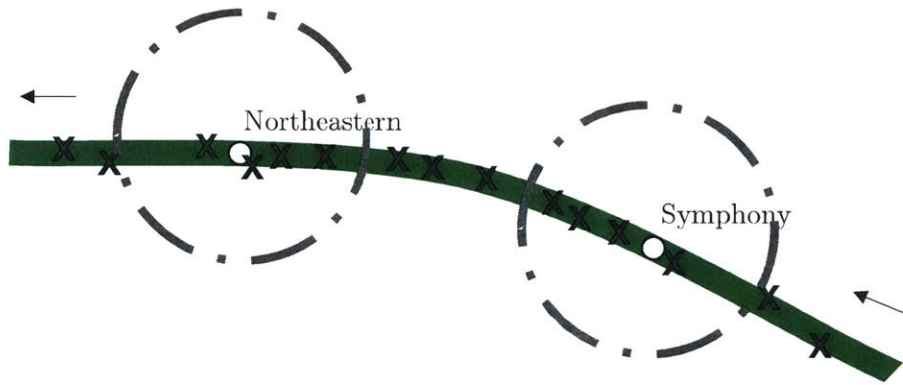


(b) Clusters around two stations.
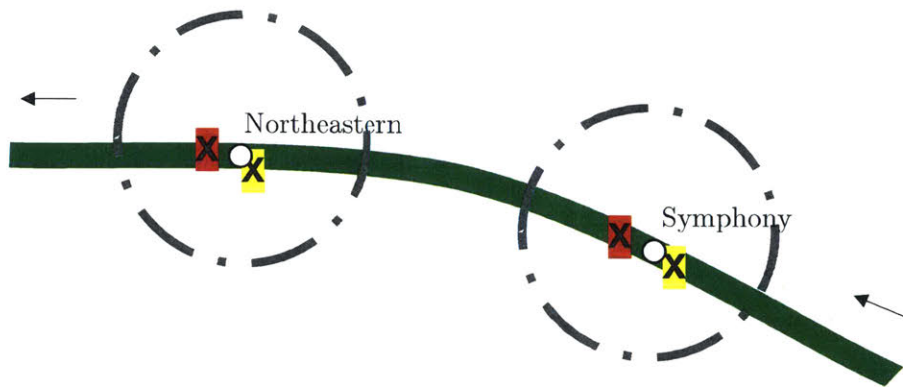


(c) Selection of station arrival and departure movements.

Figure 2-4: Depiction of the selection of movements around two stop clusters on the Heath Street branch.
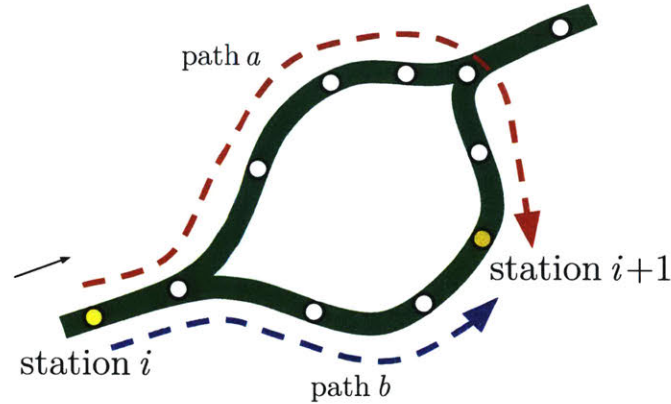
Figure 2-5: Selection of path for imputation between two known points.

recorded at 06:10, and 600 meters from the following station, Packards Corner, recorded at 06:16, then the time recorded for the imputed stop would be $\frac{1}{3}$ of the time between 06:10 and 06:16, or 06:12. Arrival and departure times for an imputed stop are the same, as if no dwell time had occurred. While this is unrealistic, the lack of any dwell time is a useful indicator of time having been interpolated and the stop visit imputed.

Trips within a vehicle's tour of the day are recorded with a unique and sequential trip ID and the appropriate branch and stop pattern. The branch corresponds to the location of each stop visit (the trunk is considered a separate branch for this purpose), and the stop pattern for all stop visits on a trip is determined at the trip's first visit to a non-trunk station. Each stop visit becomes a single record having a vehicle consist, an arrival time, and a departure time. All records are inserted to a table on the MIT-MBTA database for use in various applications.

### 2.3.7  Implementation and future work

The algorithm presented above has been used to clean vehicle location data from the MBTA Green Line. It has improved the quality of data input to the simulation model discussed in Chapter 3, particularly by increasing the origin and destination inference rate of Green Line fare transactions in ODX. It was by using ODX, in aggregate statistics as well as disaggregate examples of journeys, that the algorithm was validated and some parameters calibrated.

The heuristics work best for interpolating short stretches of missing stops or for correcting inaccurate consist information. Improvements can still be made in adjusting conditions for

allowing longer gaps between known AVL points to be imputed. For example, westbound trains without any GPS coverage often will appear, in the uncleaned data, at Kenmore westbound before reappearing, perhaps an hour later, at Kenmore eastbound. While aided somewhat by the route designations included in the original data, the designations are not always accurate, and such a large lack of data makes uncertain which branch was taken, whether a short turn was made, or whether the trip terminated at the end of the Central Subway and laid over (as trains occasionally do when run as directed). Although speed constraints are added to aid in the decision-making, the current implementation errs on the side of having fewer trips imputed. The issue of having entire trips imputed should arise less often as GPS units continue to be installed on the rest of the Green Line fleet.

## 2.4 Real-time data cleaning strategies

As noted previously, the issues related to real-time vehicle location data are largely the same as those related to archived data. The largest challenge is the unavailability of data for the entire service day. Several heuristics are used to help clean the data as it arrives in order to make control decisions in real time.

### 2.4.1 Maintenance of train activity storage

In the decision-support tool described and implemented in Chapters 4 and 5, vehicle locations from MBTA-realtime are matched to a trip appearing in the GTFS schedule and, for each trip, the most recent location coordinates, trip ID, and consist are stored in memory. This provides a comparison against sudden changes to vehicle or trip information.

### 2.4.2 Checks for consist and trip changes

All real-time vehicle records are cleaned lightly upon download from MBTA-realtime. They are checked for invalid car numbers or situations in which a car number is repeated in a consist (particularly if a train is recorded as having three cars). When a new record is being linked to an ongoing trip, it can be matched on either the trip ID (assigned by MBTA-realtime and not related to the GTFS trip ID) or the vehicle consist. Changes in trip ID are saved in memory if the previous record occurred within fifteen minutes. If the vehicle consist changes, heuristics are used to determine whether or not the change should

be saved in memory. These heuristics are based on the location of the consist change and the number of cars in the old and new consists. Only in the rare event that both values change simultaneously (with the newly-recorded consist not containing any cars from the old consist) does this method fail.

### 2.4.3 Checks for duplicate trips

Trips first appearing in the MBTA-realtime *predictionsByStop* feed—occurring upon departure from a terminal in the eastbound direction and upon passing Park Street station in the westbound direction—are checked to ensure they are not duplicates of preexisting trips. Duplicate trains can appear in the AVL data when a vehicle's GPS or AVI unit is set to the improper car number. In this situation, AVL records from GPS and AVI pings show two different consists operating on two distinct trips. The check functions by comparing the car numbers of a new record to the previous consist to pass the location. If the two consists contain at least one car in common, then the new AVL record is ignored. In the eastbound direction, a departure is prevented from being stored and matched to a GTFS trip if it occurs within a one-minute headway—shorter than the smallest practical headway along the branches due to platform and signal constraints. Trips appearing to depart in shorter succession have been confirmed to be duplicates of the same train, reporting as separate trains due to issues between the different data sources. In the westbound direction, a similar check is used but with a threshold of only four seconds—this has proven to be sufficient to eliminate duplicate trips observed in this direction.

### 2.4.4 Start of trips en route near yards

Lastly, a simple check is made to prevent new trips from being matched to the schedule and saved in memory if a train first appears in the middle of a branch adjacent to a train yard. This is an issue when a train is moved around a yard on a track adjacent to revenue tracks— GPS drift may result in a non-revenue train being detected as moving along a branch. In these cases, if the train is beginning to make a legitimate revenue trip from the yard, it will be properly processed once it travels a sufficient distance away from the yard.

## 2.5 Summary

This chapter begins by describing the limitations and common issues experienced with tracking data for the MBTA Green Line case study. Many of these issues come from the tracking hardware or the software which interprets the raw data in real time. Having the most complete and accurate data possible is critical for making informed service decisions in real time, as well as for post-hoc analyses.

An algorithm for cleaning historic AVL data is described. The algorithm is used on a nightly process before data are stored in a database for further use. It can be summarized in the following steps:

1   Filter completely invalid AVL movements.

2   Form movements into groups lasting without gaps in time and containing the same vehicle consists.

3   Form train blocks, or formations, from movement groups partially sharing consists and matching temporally.

4   Identify and correct inconsistencies in vehicle consists for movement groups within formations.

5   In each formation, form stop clusters and, for each cluster, set stop, direction, arrival time, and departure time.

6   Impute missing stop visits along routes using interpolation.

Heuristics used for cleaning real-time AVL data, applied in Chapter 5, are also described, and include checking for consist and trip changes, checking newly-appearing trains for duplications, and using caution for trains first appearing close to storage yards.

# Chapter 3

# Control strategy simulation experiment

The use of simulation models allows researchers to conduct and analyze a range of experiments quickly and at a low cost. Field experiments require the time and cooperation of a transportation agency, which may be understandably difficult to obtain if the experimental scenarios are unproven and operations might be adversely affected. By using simulation as an early step in the analysis process, reasonable estimates of the costs and benefits of the proposed scenarios can inform the design of and build support for field experimentation. This chapter begins by describing a simulation framework created by Sánchez-Martínez (2013) before describing its use to simulate a series of schedule-based and headway-based holding strategies on the Green Line.

Significant portions of this chapter were published in Fabian and Sánchez-Martínez (2017).

## 3.1 Simulation framework

The simulation experiments conducted are based on a framework developed by Sánchez-Martínez (2013). Simulations conducted within the framework are agent-based and event-driven, with agents including individual passengers and vehicles, and events including passenger or vehicle arrivals at a station. Any type of relevant information from automatically-collected or static sources can be incorporated into model specifications. This means that simulations need not rely on a static distribution of station-to-station running times, for

example, and can instead draw on an archive of observations. Other possible inputs include passenger demand by time and origin-destination pair, signal blocks, and vehicle block pull-out and pull-in times.

In each simulation run, vehicles travel between segments, and segments are linked together to form a route. Where the edge between segments is at a station, passengers, who themselves are simulated arriving at platforms, can board vehicles if the capacity allows. Custom-specified controllers are added to regulate activity at particular locations. Controllers can define how long it takes passengers to board vehicles, for instance, control terminal procedures, or define holding or short-turn strategies. After the specified number of simulation replications are run, two tables are output:

1    One table showing the details of every stop visit from every vehicle. Details can include arrival and departure time, arriving passenger load, alightings, boardings, and amount of time the vehicle was held at the station, if applicable.

2    One table showing the details of every passenger. Details can include arrival time at the origin station, vehicle boarding time, boarding vehicle, route, and direction, alighting time, and details on transfers, if applicable.

The data output to the tables can be easily changed for specific implementations—for example, transfer information was added to the output as part of this research, as previous uses of the framework have focused on single routes. The outputs can then be used in various analyses. The framework and its architecture can be found, in full, in Sánchez-Martínez (2013), Chapter 3.

Two case studies have implemented the framework thus far, both simulating high-frequency bus services. Sánchez-Martínez et al. (2016a) applies the framework for use on nine MBTA bus lines. Service on each line is simulated independently. The objective is to optimize vehicle allocations between routes while minimizing a linear combination of waiting on platforms and crowding in vehicles. The other application, by Maltzan (2015), tests different control action strategies on two MBTA bus routes. This research represents the first use of the framework both for a rail-based service and a network of multiple routes.

## 3.2 Framework implementation for MBTA Green Line

This section describes how the framework was implemented in this research for a complex, multi-branch light rail system. Specifically, this section includes the sources and dates of data used as well as models and distributions assumed.

### 3.2.1 Routes and segments

The Green Line simulation model includes a file that defines the route patterns and track geometry to be simulated. Each route pattern is unidirectional, so following the description of the four branches in Table 1.1, a list of eight route patterns can be formed. Three additional route patterns are added, as many trains operating on the Heath Street branch are based at Reservoir Yard and must travel via the Riverside branch when pulled out or pulled in.

Each route pattern contains a number of segments. A segment, in the simulation framework, can be related to a length of track: it is the unit at which vehicle running times are distributed, and can represent a length of track shorter than the distance between two adjacent stations. Most segments for the Green Line are between stations, but are broken up when a track switch is located between stations. This is done so that, in the example of a switch splitting two branches, travel times can be simulated as accurately as possible. Information on the locations of fixed signal blocks provides the capacities for the input segments (LTK Engineering Services, 2012, IBI Group, 2011). As described by Malikova (2012), the Green Line signaling system is dependent on fixed blocks, each of which can only be occupied by one train.[1] While safety is the primary motivation behind a signaling system, its design has an effect on the congestion of vehicles, as a segment of track between two stations containing three signal blocks can at most hold three trains at any one time, regardless of length. A controller in the simulation enforces the constraint, restricting the number of vehicles in a segment, and no overtaking of trains is allowed except in certain circumstances at Park Street and Kenmore stations, both of which contain two tracks per direction.

For each station or crossover, the routes and segments file contains information such as

---

[1]Only the Central Subway trunk, the Heath Street branch between Copley and Northeastern University, and the Riverside branch are protected with a block signaling system. The surface-level portions of the Boston College, Cleveland Circle, and Heath Street branches are protected by line-of-sight procedures—only major intersections with automobile conflicts are signalized.

the stop identification number, route pattern, direction, capacity (number of signal blocks) ahead of it, traffic signal phases, and fare payment information (such as whether the location contains faregates).

### 3.2.2  Segment running times

Running times for each segment are drawn actual observations that come from AVL data cleaned by the algorithm described in Section 2.3. The dates used are the 25 weekdays from March 22, 2016 to April 26, 2016.[2]

Outliers are removed by truncating lower and upper percentiles. Observations below the 15th percentile and above the 75th percentile for each segment are removed. The range was selected after inspecting the cumulative distributions of running times for every segment. Two examples of running time distributions are shown in Figure 3-1, with lines drawn at the 15th and 75th percentiles, to show that the tails are removed from the running time distributions. The distributions are truncated, particularly for upper values, to prevent the simulator's use of running times which are lengthy due to train blockages. As this implementation aims to simulate those blockages, the inclusion of observations biased by irregularities in actual service may result in a double-counting of blockages. The truncated distribution, therefore, should represent the distribution of running times in normal, free-flow conditions over all times of day.

From this truncated distribution, observations are divided into half-hour periods throughout the service day. Separate probability distributions are generated for each period. Selecting a running time for simulation involves drawing a random observation from the distribution related to the appropriate segment and 30-minute period of day. The total running time simulated may be longer than that selected if the train is blocked from passage into the next station or segment due to congestion, as often happens during the peak hours in the trunk.

### 3.2.3  Passenger demand

Passenger arrivals at origin stations are simulated via a Poisson process with arrival rates grouped by origin-destination pair within 30-minute periods. In a Poisson process for pas-

---

[2]The start date coincides with the opening of Government Center Station following a two-year reconstruction closure, and the reconfiguration of branch terminals in downtown. Additionally, Monday, April 18 is excluded due to the atypical demand during the Boston Marathon.
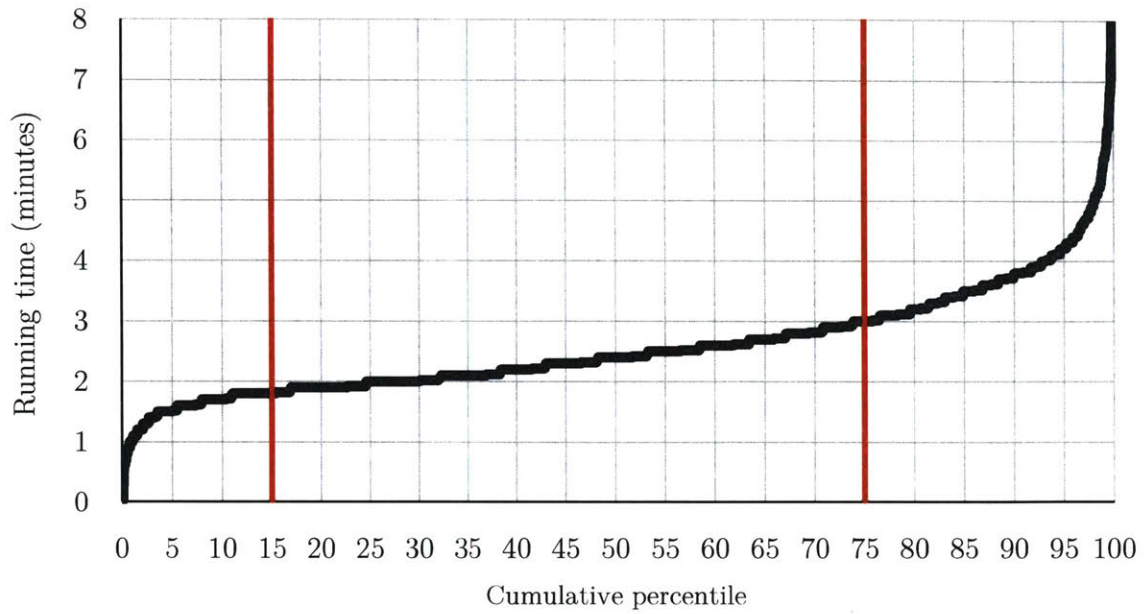
(a) Museum of Fine Arts to Northeastern University, eastbound
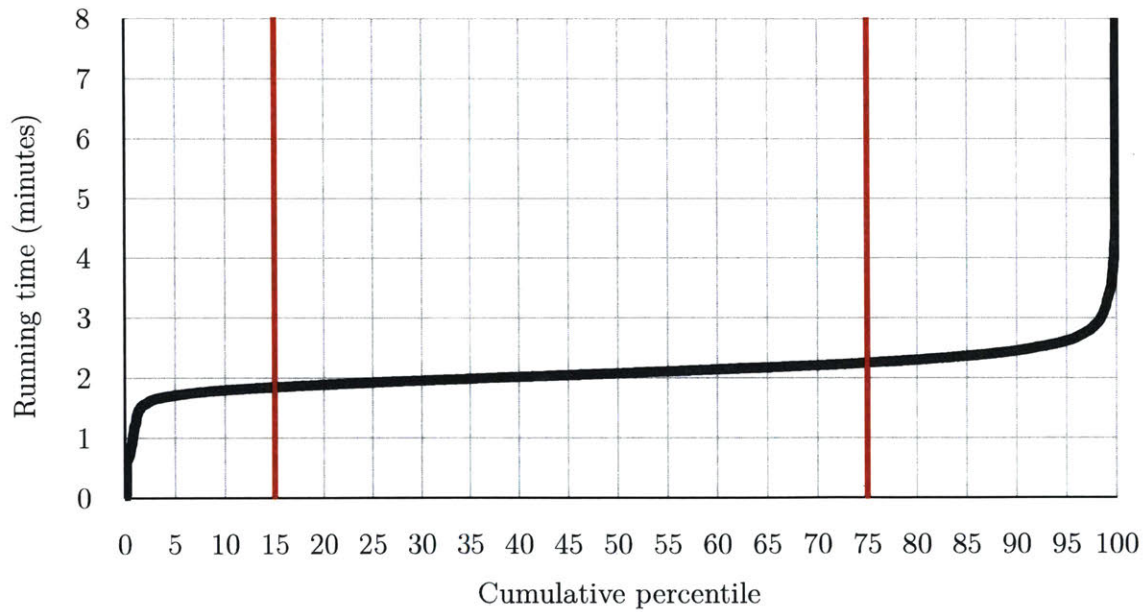


(b) Hynes Convention Center to Kenmore Square, westbound

Figure 3-1: Examples of interstation running time distributions.

senger arrivals, arrival times at a stop are mutually independent. The arrival rate is the mean number of arrivals by unit time that would be observed over a long period. This assumes that passengers arrive randomly at all hours of the day due to the high-frequency of service on the Green Line, and that they do not time their arrivals based on any available real-time information or schedule.

The arrival rates are obtained from ODX, a model that infers origins, destinations, and transfers using disaggregate AFC and AVL data (Gordon et al., 2013, Sánchez-Martínez, 2017b). The passenger load data from which the arrival rates are derived are themselves based on data from the twenty-one weekdays in April 2016. Since the ODX model by itself cannot account for passengers who evade the fare or otherwise do not interact with the AFC system upon boarding trains, non-interaction factors are applied as given by Sánchez-Martínez (2017a), who also conducted the analysis using farecard data from fall 2015.

### 3.2.4 Station dwells

A model is created to calculate the dwell time trains spend at each station. Dwell time, for this implementation, has two components: the constant, or static, portion of dwell time which all trains experience, and the variable dwell time which is a function of boardings and alightings.

A constant (or minimum) dwell time component is required because the granularity of the Green Line AVL system only provides the times of when a vehicle begins entering and finishes exiting a platform, which is longer than the time a vehicle spends stopped. The entire duration from entering through exiting of the platform must be simulated, even if the train does not actually need to stop. The constant component of dwell time is determined separately for each station by examining the cumulative distribution of observed dwell times and choosing the minimum dwell above the lower tail. In the example shown in Figure 3-2, the selected minimum dwell time for two stations is depicted with red lines. The time is then specified into the routes and segments file. In simulation, the constant is assumed to be divided equally between the arrival and departure ends of the dwell period.

Once a train stops, passengers may alight and board. Alightings are simulated before boardings. Passengers will alight once they reach their destination. If the destination is not on the train's route pattern, then passengers will alight once they reach Park Street,
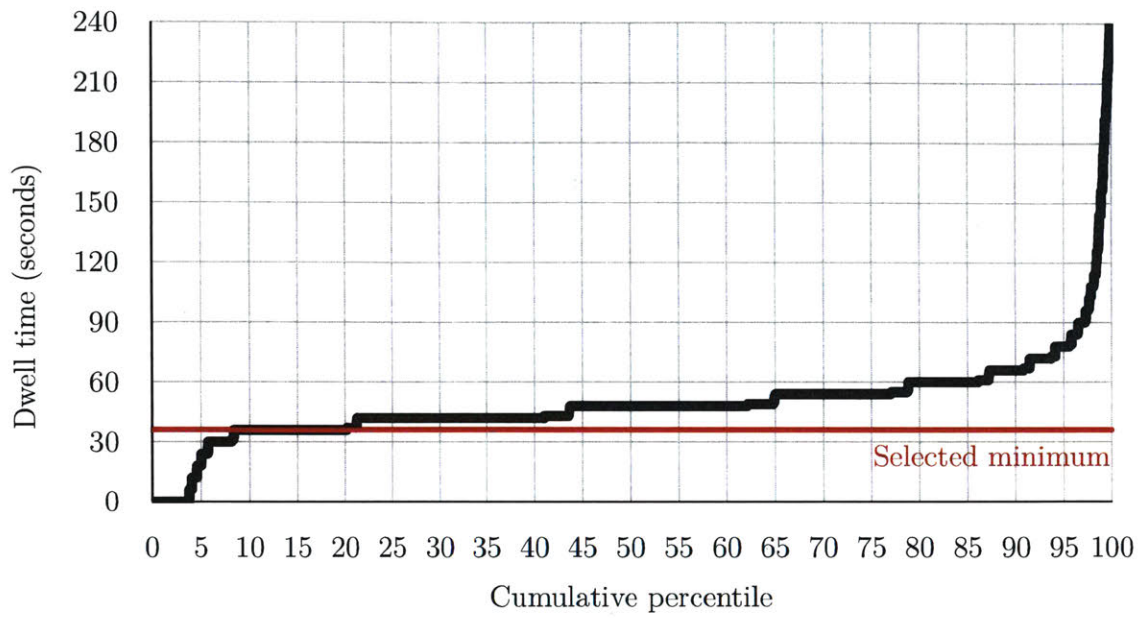
(a) Reservoir, westbound
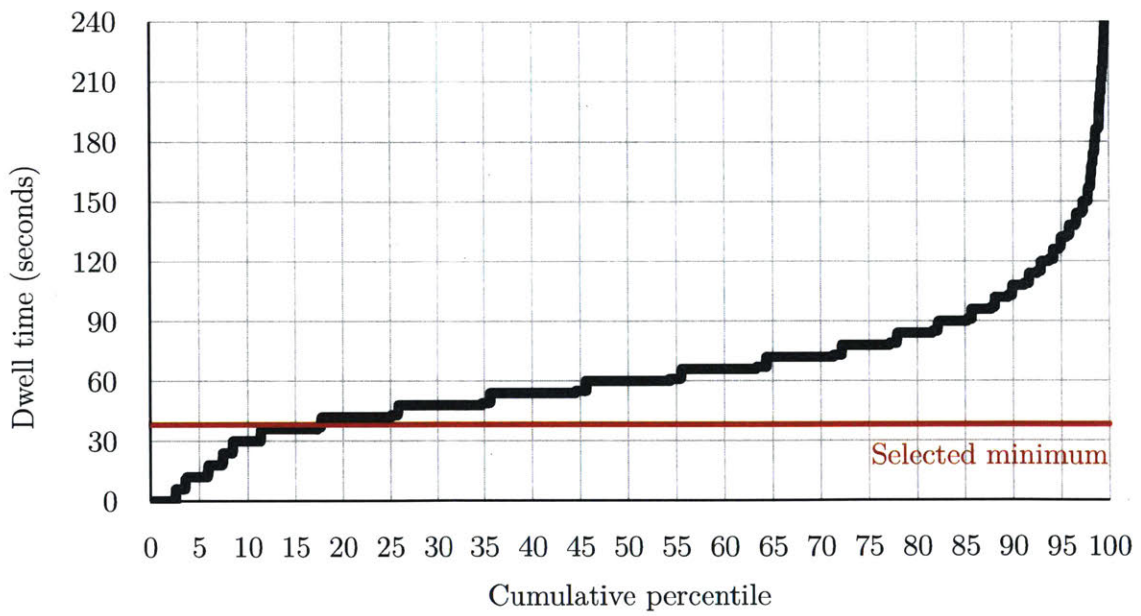


(b) Longwood Medical Area, eastbound

Figure 3-2: Examples of station dwell time distributions.

and then wait for a train that continues to their destination.[3] All four branches share the Central Subway at Park Street, so any required transfers are simulated to occur there. Following the alightings, waiting passengers can board a train if it is scheduled to serve their destination; if the origin station is not scheduled to connect directly with the destination, then they will board the first available train and change at Park Street. An example journey requiring a transfer is between Fenway and Haymarket: since the Riverside branch (which is the only branch to serve Fenway) terminates at Government Center, one would travel from Fenway to Park Street and then transfer to a train from another branch to continue to Haymarket. Passengers cannot board trains which have loads exceeding capacity, which is set to 168 passengers per car, or 336 passengers per train (Massachusetts Bay Transportation Authority, 2014).[4]

The variable dwell time is modeled differently at gated and ungated stations.[5] For gated stations, the model developed by ming Lin and Wilson (1992) for stations along the Central Subway is applied, which makes dwell time a function of alightings, boardings, and the seated capacity of the train. For ungated stations, a model is derived from AFC transaction data for boardings and industry standards for alightings (Kittelson & Associates, Inc., 2007), as follows:

$$t_{\text{alight}} = \frac{(f_{\text{front}} t_{\text{front}}) + (f_{\text{rear}} t_{\text{rear}})}{c} \tag{3.1}$$

where:

$t_{\text{alight}}$ = alighting time per passenger,

$c$ = number of cars operating per train = 2,

$f_{\text{front}}$ = fraction of passengers alighting from the front door = $\frac{1}{4}$,

$f_{\text{rear}}$ = fraction of passengers alighting from the rear doors = $\frac{3}{4}$,

$t_{\text{front}}$ = alighting time per passenger from the front door = 2.6 seconds, and

$t_{\text{rear}}$ = alighting time per passenger from the rear doors = 2.1 seconds.

---

[3]In simulation, passengers alight at a station that is neither Park Street nor their destination, if the train makes an unscheduled short-turn prior to the destination or runs express past the destination. This situation is not encountered in the experiment described in this chapter.

[4]All Green Line trains are simulated as running with two cars, which is scheduled for all hours during weekdays. In reality, a train may occasionally run with only one car due to vehicle or operator constraints. The capacity of 168 passengers is derived from the average of the policy (101) and crush capacities (234, assuming each train has one Type 7 car and one Type 8 car) of Green Line cars, as given in Massachusetts Bay Transportation Authority (2014).

[5]Gated stations include all underground stops, Science Park, Lechmere, and Riverside. All other stations are ungated and require fares to be paid at the farebox adjacent to the operator's seat, in either car.

$$t_{\text{board}} = \frac{(f_{\text{card}}t_{\text{card}}) + (f_{\text{ticket}}t_{\text{ticket}}) + (f_{\text{cash}}t_{\text{cash}}) + (f_{\text{card}}f_{\text{topup}} * t_{\text{topup}}) + (st_s)}{c} \quad (3.2)$$

where:

$t_{\text{board}}$ = boarding time per passenger,

$$s = \begin{cases} 1, & \text{if boarding passenger is a standee} \\ 0, & \text{otherwise} \end{cases},$$

$f_{\text{card}}$ = fraction of passengers validating fare with a smartcard,

$f_{\text{ticket}}$ = fraction of passengers validating fare with a magnetic-stripe ticket,

$f_{\text{cash}}$ = fraction of passengers paying fare in cash,

$f_{\text{topup}}$ = fraction of smartcard-validating passengers topping-up their card,

$$t_{\text{card}} = \text{boarding time per passenger with smartcard} = \begin{cases} 3 \text{ seconds}, & \text{during AM period} \\ 2 \text{ seconds}, & \text{otherwise} \end{cases},$$

$t_{\text{ticket}}$ = boarding time per passenger with magnetic-stripe ticket = 5 seconds,

$t_{\text{cash}}$ = boarding time per passenger with cash = 9 seconds,

$t_{\text{topup}}$ = time spent by a passenger topping-up smartcard onboard = 21 seconds, and

$t_s$ = additional boarding time required if train is above seated capacity = 0.5 seconds.

For the boarding model, $t_{\text{card}}$, $t_{\text{ticket}}$, $t_{\text{cash}}$, and $t_{\text{topup}}$ are the medians of the amount of time passengers with the respective fare transaction take to board a vehicle from a sample of Green Line boardings, where the amount of time is defined as the time between one fare transaction for a farebox and the next transaction, if both occur at the same station. The fractions of passengers using each fare payment media (smartcard, magnetic-stripe ticket, and cash) and the fraction of smartcard-equipped passengers topping-up their cards on the train are provided in Table 3.1 by Green Line branch and time of day.

Beyond the duration prescribed by the constant and dynamic terms, dwell times can be lengthened if a train must be held at a station. This may occur due to a control strategy in place (such as headway-based holding) or due to a train in the segment ahead preventing a clear signal out of the station from being shown. Passengers who arrive to the origin station during a holding period are allowed to board, capacity permitting.

Table 3.1: Fraction of Green Line passengers by branch, time, and fare media.

| Branch | Time period[a] | $f_{card}$ | $f_{ticket}$ | $f_{cash}$ | $f_{topup}$ |
|---|---|---|---|---|---|
| Boston College | AM | 0.94 | 0.04 | 0.02 | 0.019 |
| | Midday | 0.89 | 0.07 | 0.04 | 0.041 |
| | PM | 0.88 | 0.09 | 0.03 | 0.035 |
| Cleveland Circle | AM | 0.93 | 0.04 | 0.03 | 0.01 |
| | Midday | 0.88 | 0.07 | 0.05 | 0.038 |
| | PM | 0.87 | 0.10 | 0.03 | 0.035 |
| Riverside[b] | AM | 0.91 | 0.07 | 0.02 | 0.006 |
| | Midday | 0.85 | 0.12 | 0.03 | 0.017 |
| | PM | 0.83 | 0.14 | 0.03 | 0.014 |
| Heath Street | AM | 0.89 | 0.08 | 0.03 | 0.02 |
| | Midday | 0.84 | 0.11 | 0.05 | 0.033 |
| | PM | 0.86 | 0.11 | 0.03 | 0.025 |

[a] For this table, AM is defined as 04:00-10:00, midday 10:00-14:00, and PM as the rest of the service day, 14:00-26:00.

[b] All ungated stations on the Riverside branch feature unstaffed fare vending machines where passenger can obtain or top-up fare media.

## 3.2.5 Terminal behavior

Terminal controllers define the behavior of trains at the start and end of every route pattern. As current Green Line operations revolve around scheduled trips, the complete schedule for each station is built into the controller for the respective terminal or en route stop. When a train reaches a terminal, either at a train's pull-out time and location or at the end of a revenue trip, the next route pattern is inferred from the current location. For example, an eastbound train terminating at Government Center via the Riverside branch will have, as its next pattern, the westbound route on the Riverside branch toward Riverside. Interlining between branches is not permitted, per current practice on the Green Line.

If the train is at its assigned pull-in location, an estimate is made on the duration of the next route trip from and to that location. This estimate is made based on the method of selecting segment running times presented in Section 3.2.2 and the constant portion of each station's dwell time. (The effects of demand and segment blocking are not simulated for these estimates, so it is likely that the estimated cycle time will be lower than that actually simulated.) If the estimated time of return to the pull-in location is later than the scheduled pull-in time, then the train does not go back into service and is no longer simulated until the next replication. If the train is at a terminal which is not its assigned

pull-in location, the train operates on the next route pattern regardless of time of day. An important exception is for Heath Street branch trains based at Reservoir. Trains pulling out from the yard operate from Reservoir to Government Center via the Riverside branch, loop, and then travel to Heath Street. For such trains approaching the pull-in time, the decision point is at Lechmere: a calculation is made as to whether the train can travel to Heath Street, return to Lechmere, and then travel to Reservoir before the pull-in time. Trains pulling in travel in revenue service from Lechmere to Reservoir via the Riverside branch.

Departing trains are assigned the first available trip for the route pattern departing that terminal. A trip is available for dispatching from a terminal if it has not already been assigned to another train and its scheduled departure time has not been elapsed by more than twenty minutes. Unavailable trips which have not been dispatched are considered canceled from the schedule. In the experiment scenarios simulated, the number of canceled trips during the morning peak period (for which the analysis is conducted) resulting from these rules is below the 6-7% of peak period Green Line trips which were not observed, and likely canceled, by Sindel (2017).

Operators are not explicitly modeled in this simulation, leaving scheduled vehicle blocks, which are publicly-obtainable from an agency's GTFS feed, to act as the closest proxy. While it is true that operators cannot work the 12 to 14 hours for which some blocks are scheduled, the calculations estimating whether a train can operate an additional cycle before the pull-in time are similar to the guesswork terminal inspectors must do for operators in order to avoid situations in which they must clock-out late. The simulation, in its existing form, underestimates the number of time constraints that inspectors and dispatchers have to consider when making service decisions.

Vehicle blocks, including their pull-out and pull-in locations and times, are taken from the Spring 2016 schedule as provided in GTFS. Figure 3-3 depicts the span of all 100 scheduled blocks on the Green Line as well as their respective pull-out and pull-in locations. The availability of spare vehicles at terminal yards is not accounted for in this implementation. In existing operations, the use of spare vehicles (and cover operators), when available, could be considered a control action, as they can fill gaps in service and ensure on-time departures.

Minimum layover times at terminals are also determined from the GTFS schedules as well as from recommendations by Operations staff. For most terminals, the minimum turnaround time is 3 minutes. At Park Street and Government Center, it is 2 minutes because there is
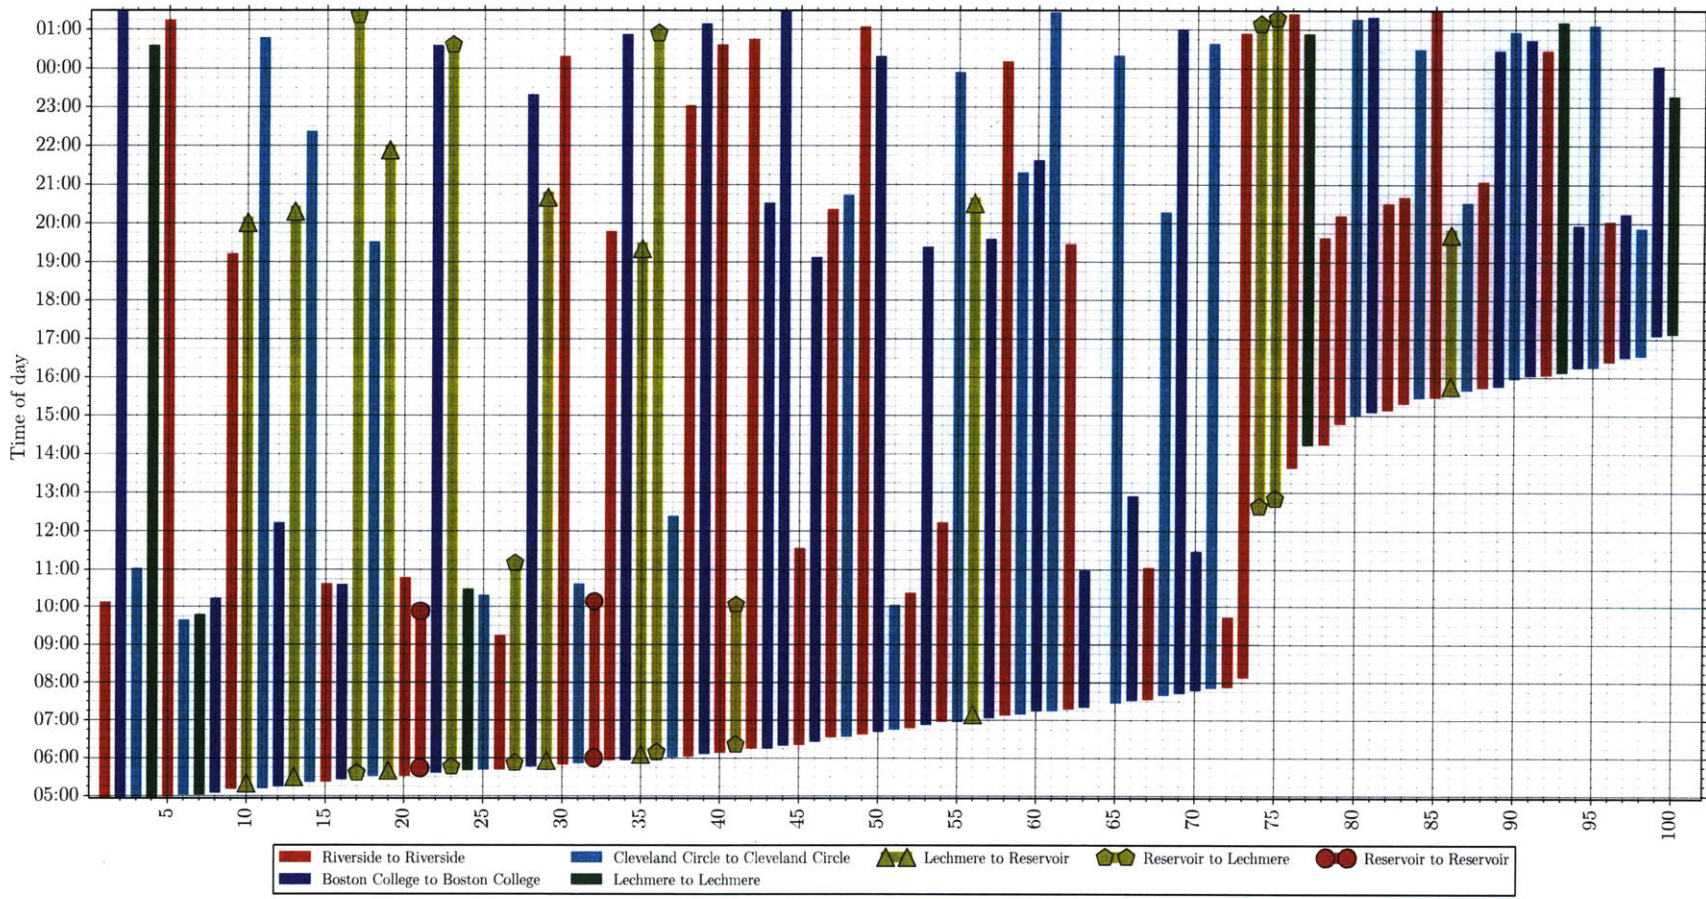
Figure 3-3: Vehicle blocks on MBTA Green Line for spring 2016, from GTFS schedule.

no need to change crews, and similarly, this is lowered to 1 minute at Heath Street. At North Station, the turnaround time is extended to 4 minutes because the crossover is located well past the station, and operators are required to change ends when switching directions.

## 3.3 Simulation experiment overview

Using the Green Line simulation model described above, a number of schedule-based and headway-based holding strategies are tested. The strategies are programmed in the controller classes for stations and terminals, and the controllers are set to a strategy based on flags in the routes and segments file. In order to evaluate only the effects of different holding strategies, other commonly-used control strategies, such as short-turning, are not tested in this experiment. This means that, while the simulation includes many realistic aspects of current operations on the Green Line, the results are unlikely to precisely match a real-world implementation.

In schedule-based holding, a train can depart a terminal or station at or after its scheduled time. If a train arrives at the stop later than the next scheduled departure, then it departs after the minimum layover time has elapsed, if at a terminal, or immediately following any normal dwell time if en route. This is formulated in Equation 3.3.

$$D_{ij} = \max \left( \min(D_{ij}^*, A_{ij} + H_{ijt} + T_j + W_{ij}), A_{ij} + T_j + W_{ij} \right) \tag{3.3}$$

where:

$D_{ij}$ = departure time from stop $j$ for trip $i$,

$D_{ij}^*$ = scheduled departure time (trip $i$) from stop $j$,

$A_{ij}$ = arrival time to stop $j$ for trip $i$,

$H_{ijt}$ = maximum allowed holding duration for trip $i$ at stop $j$ based on time of day $t$,

$T_j$ = minimum layover and turnaround time at $j$ if stop $j$ is a terminus, and

$W_{ij}$ = dwell time for trip $i$ at stop $j$.

In headway-based holding, the main determinants of a train's departure are the headway between it and the preceding vehicle, $A_{ij} - A_{(i-1)j}$, and the scheduled target headway, $F_{ijt}$ at time of day $t$. The strategy is detailed in Equation 3.4.

Table 3.2: Stations used for en route holding in the simulation model

|  | Branch/segment | Station |
|---|---|---|
| Branch | Boston College | Harvard Avenue EB |
|  | Cleveland Circle | Coolidge Corner EB |
|  |  | St Mary's St EB |
|  | Riverside | Reservoir EB |
|  |  | Fenway EB |
|  | Heath Street | Brigham Circle EB |
| Trunk | Central Subway | Park Street WB |
|  |  | Kenmore WB |
|  |  | Kenmore EB |

$$D_{ij} = \max \Big( \min(A_{(i-1)j} + F_{ijt} + T_j + W_{ij}, A_{ij} + H_{ijt} + T_j + W_{ij}), A_{ij} + T_j + W_{ij} \Big) \quad (3.4)$$

Holding at terminals is not limited to a maximum duration in any scenario. For en route holding (along the branch or trunk), some of the scenarios tested allow holding at each station to be as long as necessary, while other scenarios limit holding to two minutes per station.

Holding is tested in branch and trunk terminals, selected branch-line stations, selected trunk-line stations, and combinations thereof. Holding at select stops along the trunk is tested with strategies that hold either to the vehicle's branch's headway or to the combined headway of all branches serving the station. In the latter case, Equation 3.5 determines the combined trunk headway, $F_{jt}$ for stop $j$ at time of day $t$:

$$F_{jt} = \Big( \sum_{k=1}^{K} f_{kt}^{-1} \Big)^{-1} \quad (3.5)$$

where $K$ is the set of branches that stop at $j$, each having a scheduled headway of $f_{kt}$.

On the Heath Street branch, the eastbound Brigham Circle station acts as the terminus point for holding. En route holding is applied only on commonly used control points, which are listed in Table 3.2. There are five eastbound branch control points, and one eastbound and two westbound trunk control points.

Twenty replications are run for each scenario, each spanning an entire service weekday. Analyses are run on only the subset of passengers arriving at origin stations between 05:00

and 10:00. The focus is placed on this morning peak period, as it captures a large share of daily ridership and is generally uniform in scheduled headway—5–6 minutes—reducing natural variability in wait times due to different service levels throughout the day. $T$-tests confirm that twenty replications suffice to obtain statistically significant results. Performance across scenarios is compared based on both passenger-oriented and operator-oriented performance measures calculated from vehicle and passenger movements. In addition to visually comparing cumulative probability distributions, means are compared with $t$-tests and variances with $F$-tests. Passenger-oriented metrics include wait time, wait time reliability, and journey time. Wait time is the difference in time between a passenger's arrival at a station and boarding of a train, including any transfer time, and is given in Equations 3.6 (no transfer required) and 3.7 (transfer required).

$$w_i = b_i - a_i \tag{3.6}$$

$$w_i = (b_i - a_i) + (b_i^{\text{transfer}} - d_i^{\text{transfer}}) \tag{3.7}$$

where:

$w_i$ = wait time of an individual passenger $i$,

$b_i$ = vehicle boarding time for passenger $i$ at the origin station,

$a_i$ = origin arrival time for passenger $i$,

$b_i^{\text{transfer}}$ = vehicle boarding time for passenger $i$ at the transfer station, and

$d_i^{\text{transfer}}$ = vehicle alighting time for passenger $i$ at transfer station.

Wait reliability is a metric used by the MBTA (Tribone et al., 2016). A passenger's experience is considered reliable if the wait time is less than the scheduled headway, and unreliable otherwise. The metric is detailed in Equation 3.8.

$$W = \frac{\sum_{i=1}^{n} r_i}{n} \tag{3.8}$$

where:

$W$ = wait reliability over $n$ passengers,

$h_i$ = scheduled headway for $i$ (a function of origin, destination, and time of day),

$r_i = \begin{cases} 1, & \text{if } h_i - (b_i - a_i) \geq 0 \\ 0, & \text{otherwise} \end{cases}$ , and

$n$ = number of passengers analyzed.

Given AVL data, arrival rates are used to estimate how many people experienced a wait longer than the scheduled headway, and this subset of passengers is counted as having an unreliable wait. This headway-based metric is adopted for analysis because passengers expect service to operate frequently and generally make trips without consulting the schedule (Turnquist, 1981). Tabulated results show reliability as the percentage of passengers with a reliable wait time. Total journey time combines a passenger's wait time and in-vehicle travel time, and is given in Equation 3.9.

$$j_i = d_i - a_i \qquad\qquad (3.9)$$

where:

$j_i$ = total journey time of an individual passenger $i$, and

$d_i$ = vehicle alighting time for passenger $i$ at destination station.

Evaluating journey times is of interest because drawing conclusions based only on wait times neglects the additional journey time of passengers travelling through (and held at) control points. Holding en route might help decrease wait times, but potentially at the detriment of passengers' journey times. One-way running time distributions are used to evaluate performance from the operator's perspective, since they affect operations planning, required fleet and crew, and the ease with which the operations plan can be followed (Sánchez-Martínez, 2013).

Performance is evaluated for vehicles and passengers on all branches during the morning period, separately for each direction because the trunk is concentrated to the east and the branches to the west, so different control points may disproportionately benefit service in one direction. Since the same passenger demand and quantity of service is used for each scenario, the observed differences in performance across scenarios are due to differences in control strategies. While vehicles operating on different branches and passengers with different origins and destinations will naturally experience a wide range of travel times, the distributions of these travel times can be compared across control strategies.

## 3.4 Experiment results

### 3.4.1 Unlimited holding durations

In the first scenarios, holding durations are not limited at any location. In the eastbound direction, the predominant direction of travel along the branches during the morning peak, regulating service for schedule adherence results in only small fluctuations in running time means and variances when changing control points. Controlling for headway adherence, on the other hand, results in either significantly higher or lower mean running times depending on the choice of holding points. Scenarios with holding at terminals or controlling by the combined headways of multiple branches at trunk stations yield faster trains than scenarios with separate control on each branch. This is primarily because trains held in the trunk to their branch headway prevent following trains of other branches from moving forward as fast as they otherwise would. Examining westbound running times supports these findings: the highest variance is found in the scenario where holding based on branch headways occurs in the trunk with no control elsewhere. Headway-based holding results in significant improvements over schedule-based holding in most other scenarios. In general, however, the difference in average running times between control policies is less than one standard deviation. These small differences, while statistically significant, suggest that other performance metrics should first be compared before drawing conclusions.

Table 3.3 presents statistics on passenger wait times. On average, wait times are cut by 27% when adopting headway-based over schedule-based holding policies. Holding only at terminals decreases average wait time by 32% and the standard deviation by 44%. Adding en route holding along the branches further improves performance, albeit marginally, with a decrease in average wait time of 34% and a decrease in standard deviation of wait time of 47%. Headway-based holding in the trunk only slightly improves performance over schedule-based control. Figures 3-4a and 3-4b show the percentage of simulated passengers served by a train within any duration of wait. By percentile of total ridership, wait times in headway-based holding scenarios are either the same or lower than in schedule-based holding. The no-control scenario underperforms all other scenarios in both directions.

Table 3.4 shows wait time reliability in both directions; numbers give the percentages of passengers who waited less than the scheduled headway. Schedule adherence results in lower wait time reliability than headway-based holding, and holding at terminals (alone or

Table 3.3: Passenger wait times with unlimted holding strategies.

| Locations | | | Eastbound passengers | | | | | | Westbound passengers | | | | | |
| --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |
| | | | Hold to keep branch headways | | Hold to keep applicable combined headway at each point | | Hold to keep published vehicle schedules | | Hold to keep branch headways | | Hold to keep applicable combined headway at each point | | Hold to keep published vehicle schedules | |
| Hold at terminals | Hold at selected branch points | Hold at selected trunk points | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) |
| ✓ | | | 3.1 | 2.9 | —c | — | 4.9 | 6.2 | 3.6 | 4.1 | — | — | 4.8 | 5.9 |
| | ✓ | | 3.4 | 3.8 | — | — | 6.0 | 7.9 | 3.9 | 4.3 | — | — | 5.5 | 6.8 |
| | | ✓ | 5.8 | 7.4 | 5.8 | 7.2 | 6.0 | 7.2 | 5.1 | 5.9 | 5.1 | 6.2 | 5.3 | 6.2 |
| ✓ | ✓ | | 2.9 | 2.8 | — | — | 4.9 | 6.4 | 3.6 | 4.0 | — | — | 4.7 | 6.0 |
| ✓ | | ✓ | 3.9 | 4.2 | 3.3 | 3.3 | 5.1 | 6.1 | 3.8 | 4.3 | 3.4 | 4.0 | 4.8 | 5.5 |
| ✓ | ✓ | ✓ | 2.9 | 2.9 | 2.8 | 2.8 | 5.0 | 6.0 | 4.2 | 5.0 | 3.4 | 4.0 | 4.7 | 5.6 |

| No control | Mean = 6.1 | Standard deviation = 7.5 | Mean = 5.4 | Standard deviation = 6.3 |
| --- | --- | --- | --- | --- |

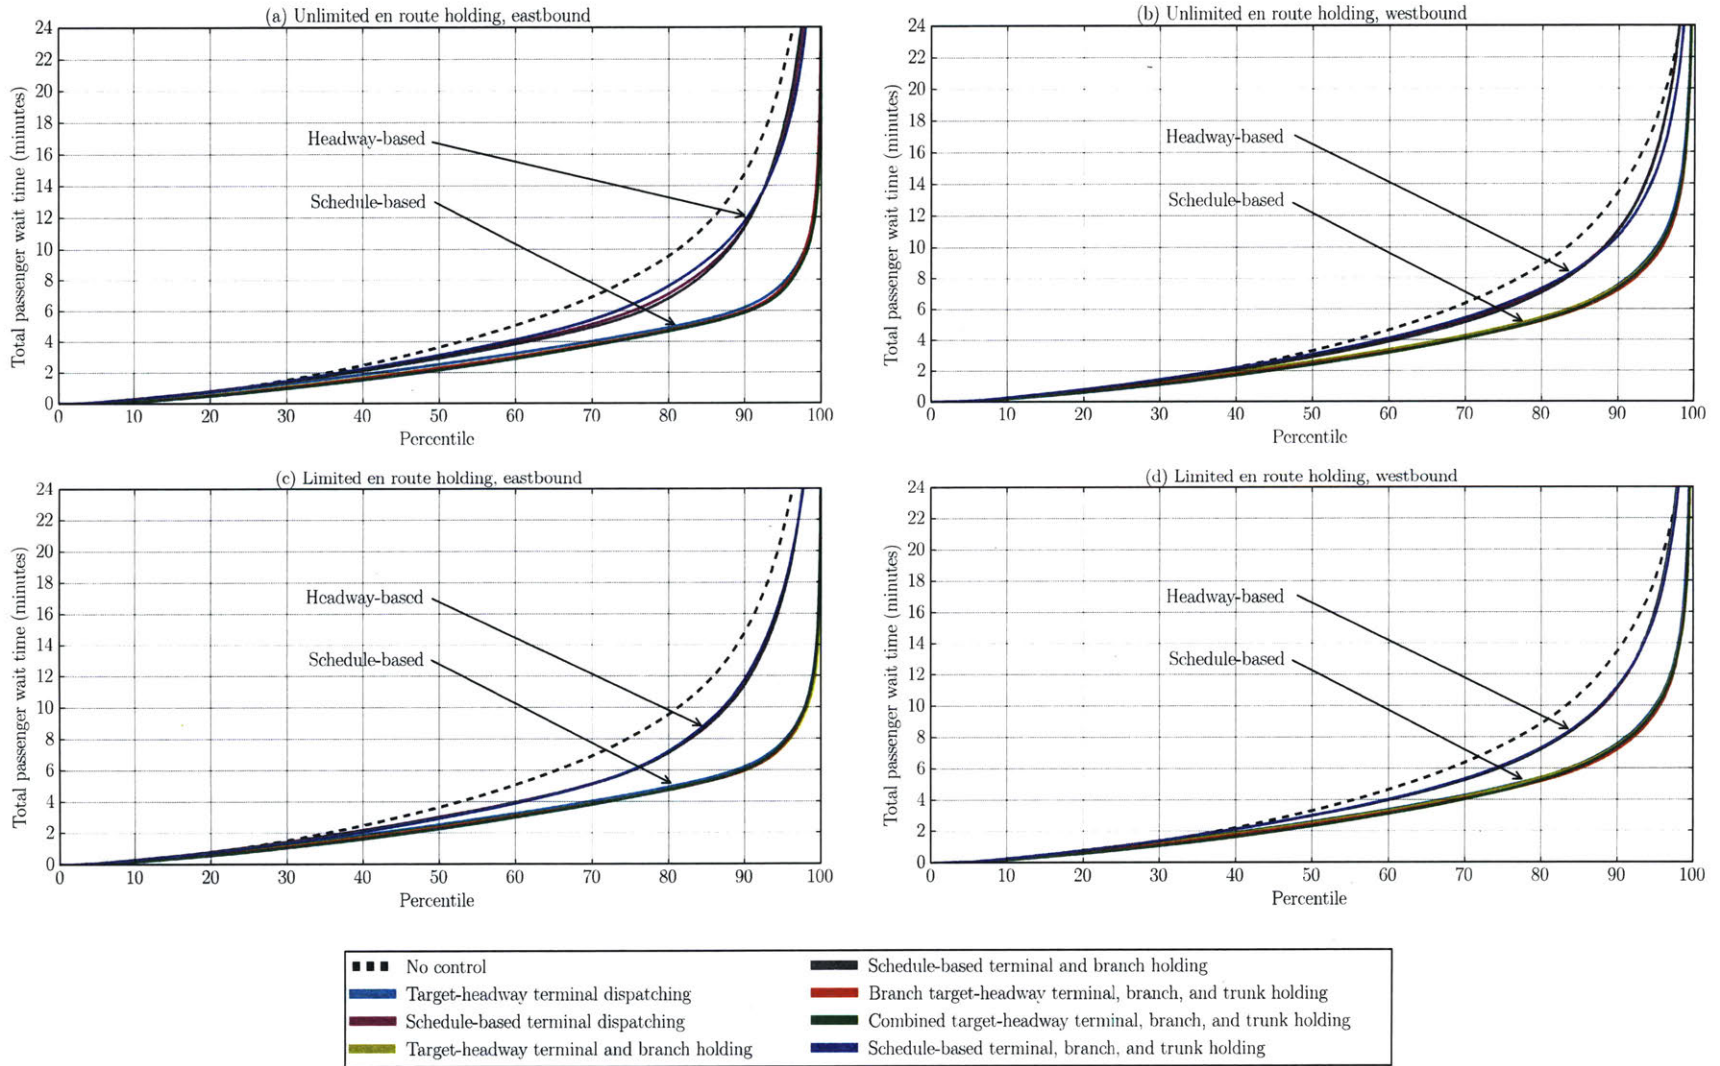c In all tables, "—" denotes inapplicable information.

Figure 3-4: Cumulative distributions of simulated passenger wait times.

Table 3.4: Wait time reliability with unlimited en route holding strategies.

| Locations | | | Eastbound passengers | | | Westbound passengers | | |
| --- | --- | --- | --- | --- | --- | --- | --- | --- |
| Hold at terminals | Hold at selected branch points | Hold at selected trunk points | Hold to keep branch headways | Hold to keep applicable combined headway at each point | Hold to keep published vehicle schedules | Hold to keep branch headways | Hold to keep applicable combined headway at each point | Hold to keep published vehicle schedules |
| ✓ | | | 82% | — | 70% | 76% | — | 62% |
| | ✓ | | 76% | — | 63% | 72% | — | 57% |
| | | ✓ | 60% | 62% | 61% | 63% | 64% | 57% |
| ✓ | ✓ | | 83% | — | 71% | 76% | — | 63% |
| ✓ | | ✓ | 74% | 81% | 67% | 74% | 78% | 60% |
| ✓ | ✓ | ✓ | 84% | 86% | 67% | 78% | 78% | 61% |
| No control | | | 60% | | | 62% | | |

in addition to en route) benefits the greatest proportion of peak-direction passengers. Table 3.5 gives a summary of passenger journey times across each unlimited holding scenario. Headway-based holding significantly decreases average journey times, except in certain scenarios implementing holding within the trunk.

Considering both directions, the strategies minimizing journey times are (a) terminal-only holding on headways and (b) terminal and trunk holding on combined headways. In the eastbound direction, significant increases in journey time means and variances for scenarios including trunk holding by branch headways can be explained by upstream congestion: while vehicles can overtake at two stations, multiple branches may need to be controlled, causing blockages. This is mitigated when trains are held by combined headways, since trains are scheduled to pass very frequently. The journey time distributions in Figures 3-5a and 3-5b reveal that headway-based holding decreases journey times for nearly all passengers, and that in-vehicle times also decrease slightly. That most of the journey time improvements come from wait time should be preferred: revealed preferences studies show that waiting is more onerous to passengers than in-vehicle time (Maltzan, 2015, Sánchez-Martínez, 2015).

Vehicle trajectories help explain the high wait and journey times resulting from holding based on vehicle schedules. It only takes several junction blockages, service gaps, or ridership spikes to delay vehicles and cause bunching. Once several trains fall behind schedule, it becomes difficult to recover with only holding, and thus schedule-based holding strategies eventually produce similar service as without any control. This convergence to bunching speaks to the lack of robustness of the schedule used. It is likely that, without additional control actions, increasing recovery time would decrease passenger wait times, but also increase scheduled headways or fleet requirements.

Table 3.5: Passenger journey times with unlimted holding strategies.

| Locations | | | Eastbound passengers | | | | | | Westbound passengers | | | | | |
| --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |
| | | | Hold to keep branch headways | | Hold to keep applicable combined headway at each point | | Hold to keep published vehicle schedules | | Hold to keep branch headways | | Hold to keep applicable combined headway at each point | | Hold to keep published vehicle schedules | |
| Hold at terminals | Hold at selected branch points | Hold at selected trunk points | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) |
| ✓ | | | 22.2 | 12.4 | — | — | 25.2 | 14.9 | 17.5 | 9.9 | — | — | 19.1 | 11.4 |
| | ✓ | | 26.7 | 17.7 | — | — | 27.0 | 16.5 | 17.9 | 10.2 | — | — | 20.1 | 12.3 |
| | | ✓ | 38.8 | 27.9 | 27.2 | 16.5 | 29.0 | 17.6 | 20.1 | 11.7 | 20.1 | 12.0 | 20.2 | 12.4 |
| ✓ | ✓ | | 23.3 | 13.5 | — | — | 24.9 | 14.7 | 17.4 | 9.8 | — | — | 19.0 | 11.4 |
| ✓ | | ✓ | 32.1 | 22.3 | 22.0 | 12.2 | 27.6 | 16.8 | 17.9 | 10.1 | 17.5 | 10.0 | 19.1 | 11.2 |
| ✓ | ✓ | ✓ | 24.2 | 14.3 | 22.9 | 13.4 | 27.5 | 16.5 | 18.0 | 10.1 | 17.4 | 10.0 | 19.0 | 11.2 |

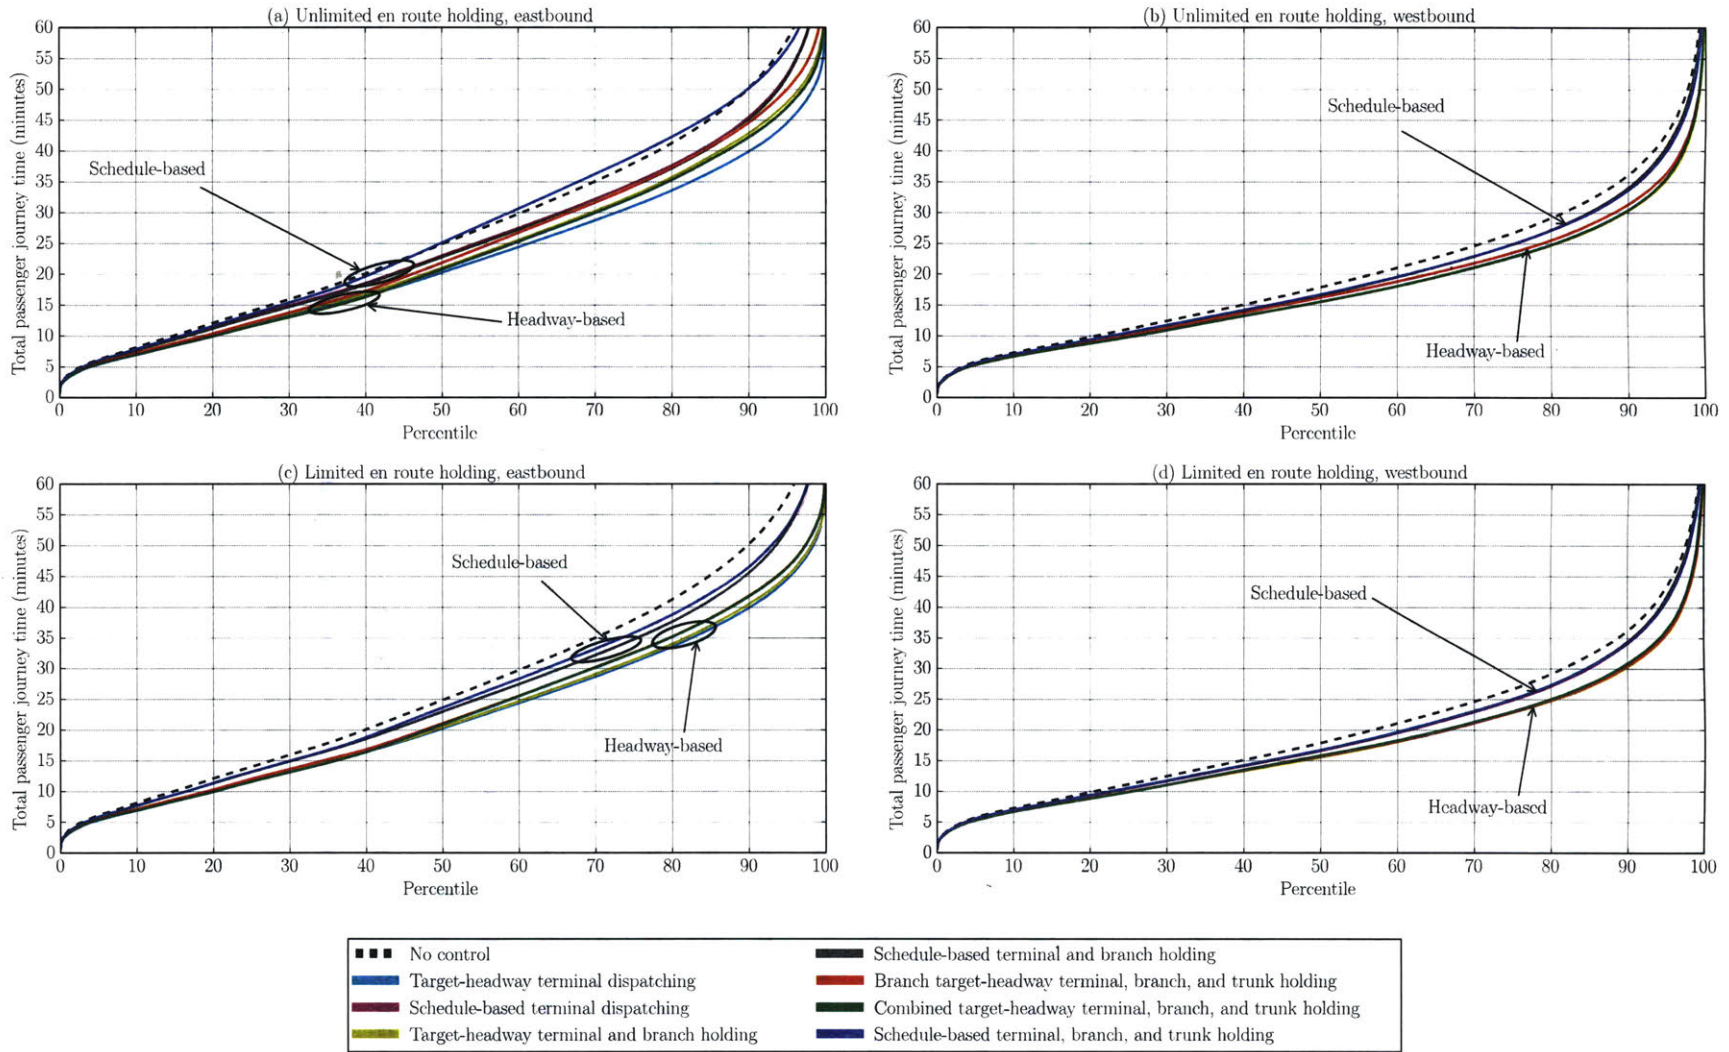No control     Mean = 27.4     Standard deviation = 16.5     Mean = 20.3     Standard deviation = 12.0

Figure 3-5: Cumulative distributions of simulated passenger journey times.

### 3.4.2 Limited en route holding durations

The findings above are compared against those of scenarios that constrain en route holding times, balancing the discomfort of those riding through the control point with service regulation (van Oort et al., 2010). Aside from limiting holding-induced congestion in the trunk, holding limits make the control strategy more palatable to passengers by decreasing their in-vehicle delays, an important consideration in practice. In previous research, maximum holding limits of 1-3 minutes have been used (van Oort et al., 2010, Maltzan, 2015, Sánchez-Martínez, 2015), though other implementations are possible, such as using a specific fraction of the current headway. This research uses a two-minute holding limit, except at terminals, where no limit is applied for dispatching.

Results with limited holding are similar to those without limits, but with several differences in optimal control locations. Average running times decrease significantly (to the 3% significance level) in six of thirteen tested scenarios, with another four experiencing no significant change. In the remaining three scenarios, a lack of sufficient holding en route may increase time spent in queues approaching stations, resulting in longer running times overall. The wait time means and standard deviations, summarized in Table 3.6, are generally the same or slightly lower when compared against the corresponding unlimited holding scenarios. As shown on the cumulative distribution plots, Figures 3-4c and 3-4d, headway-based scenarios result in waits longer than 10 minutes for 4% of passengers, whereas this increases to about 13% of passengers in schedule-based scenarios. The 90th percentile wait time under headway-based scenarios is 44% shorter than under schedule-based scenarios. Table 3.7 reports similar trends in wait time reliability: fewer passengers are well-served when terminals are not holding points. Limiting holding times at some control points lessens opportunities to recover the schedule or headway. Figures 3-5c and 3-5d show that this decreases journey times under branch-based trunk holding scenarios. All other headway-based strategies have similar journey time distributions with the limit. Especially when used in conjunction with headway-based terminal holding, limiting hold durations en route provides similar or increased passenger benefits as without limits, while making the strategy more practical.

Table 3.6: Passenger wait times with strategies limiting en route holding to two minutes per station.

| Locations | | | Eastbound passengers | | | | | | Westbound passengers | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | Hold to keep branch headways | | Hold to keep applicable combined headway at each point | | Hold to keep published vehicle schedules | | Hold to keep branch headways | | Hold to keep applicable combined headway at each point | | Hold to keep published vehicle schedules | |
| Hold at terminals | Hold at selected branch points | Hold at selected trunk points | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) | Mean (min) | Standard deviation (min) |
| ✓ | | | 3.1 | 2.9 | — | — | 4.9 | 6.2 | 3.6 | 4.1 | — | — | 4.8 | 5.9 |
| | ✓ | | 4.0 | 5.5 | — | — | 6.1 | 7.7 | 4.2 | 5.0 | — | — | 5.7 | 6.9 |
| | | ✓ | 5.3 | 6.6 | 5.8 | 7.5 | 5.9 | 7.5 | 4.6 | 5.8 | 5.0 | 6.3 | 5.3 | 6.5 |
| ✓ | ✓ | | 2.9 | 2.7 | — | — | 4.9 | 6.5 | 3.5 | 4.0 | — | — | 4.8 | 6.2 |
| ✓ | | ✓ | 3.3 | 3.3 | 3.4 | 3.5 | 4.7 | 5.9 | 3.4 | 3.9 | 3.4 | 4.1 | 4.7 | 5.6 |
| ✓ | ✓ | ✓ | 2.9 | 2.8 | 2.9 | 2.8 | 4.9 | 6.4 | 3.4 | 4.0 | 3.4 | 3.4 | 4.8 | 6.0 |
| No control | | | Mean = 6.1 | | Standard deviation = 7.5 | | | | Mean = 5.4 | | Standard deviation = 6.3 | | | |

Table 3.7: Wait time reliability with strategies limiting en route holding to two minutes per station.

| Locations | | | Eastbound passengers | | | Westbound passengers | | |
|---|---|---|---|---|---|---|---|---|
| Hold at terminals | Hold at selected branch points | Hold at selected trunk points | Hold to keep branch headways | Hold to keep applicable combined headway at each point | Hold to keep published vehicle schedules | Hold to keep branch headways | Hold to keep applicable combined headway at each point | Hold to keep published vehicle schedules |
| ✓ | | | 82% | — | 70% | 76% | — | 62% |
| | ✓ | | 70% | — | 61% | 70% | — | 55% |
| | | ✓ | 66% | 63% | 62% | 69% | 66% | 57% |
| ✓ | ✓ | | 84% | — | 70% | 76% | — | 62% |
| ✓ | | ✓ | 80% | 81% | 71% | 79% | 78% | 62% |
| ✓ | ✓ | ✓ | 85% | 86% | 71% | 78% | 79% | 62% |
| No control | | | 60% | | | 62% | | |

## 3.5 Limitations and future work

This simulation-based study is not without its limitations, many of which could be overcome with further modeling and experimentation. Future research may optimize the placement and number of control points at terminals, branches, and the trunk, as well as the maximum holding duration allowed at en route stations. As mentioned previously, passenger demand rates are given for 30-minute periods during the morning peak, which may not precisely capture demand variations during peak-of-the-peak periods. A more applicable area of further research for other transit systems is the integration of terminal capacity constraints, spare vehicle counts, and operator schedules. Capacity is currently considered between stops and within non-terminal stations, but in practice, constraints at terminal and short-turn points within the trunk often results in the early dispatching of trains. Such infrastructure-based limitations may reduce the effectiveness of headway-based dispatching, but further work is needed to quantify this. Passenger choice research may also be conducted to better allocate the preferred locations of transfers along the Green Line trunk or determine whether passengers prefer waiting for a one-seat ride versus boarding the first train to arrive, regardless of destination, potentially requiring a transfer.

Analyses of services or passengers can also be broken down by branch, origin, and destination, which may result in different policy recommendations for different branches or during off-peak hours. The simulation modeling approach could be used to test additional holding strategies such as the even-headway strategy and those based on rolling horizon optimization, or stop-skipping control strategies such as short-turning and expressing. The framework used in this research could be applied, with minor modification, to multi-route BRT corridors with exclusive rights-of-way.

## 3.6 Conclusions

This chapter focuses on determining effective holding strategies for a multi-branch light rail service experiencing congestion where branches merge into a high-frequency trunk, as well as track capacity constraints seldom present in bus operations. A simulation of the MBTA Green Line was built to model operations based on historic demand and running time data. Various forms of schedule-based and target headway-based control were tested, including control at terminals, on branches, and in the trunk, as well as limited holding durations.

The analyses were limited to the morning period to limit the scheduled variability in service and isolate the effects of control strategies on performance.

Based on the case study, holding based on target headway prevents cascading delays and bunching more effectively than schedule adherence strategies. The performance resulting from both types of holding strategies significantly depends on the choice of holding points. Terminals are the single most effective locations to control service, and holding at terminals is essential even when service can also be regulated en route. This finding confirms what others have found in previous research (Eberlein et al., 2001). On transit services with branches merging into a trunk where there is not much opportunity to overtake, holding vehicles of each branch to branch-specific headways can be counterproductive due to the interaction of vehicles serving different branches. This performance-decreasing congestion can be avoided to a large extent by limiting holding times in the trunk or by holding to the combined trunk headway. Limiting en route holding also makes the control strategy more palatable to passengers, an important consideration in practice. Out of the strategies considered in this research, holding at terminals and branches based on headway yielded the best performance. Compared to schedule adherence, this strategy decreased mean wait time by 35%, standard deviation of wait time by 49%, 90th percentile wait time by 43%, and shortened journey times and running times.

While adding recovery time to the vehicle schedule could solve some performance problems observed in schedule-based holding scenarios, it would also require either longer headways (and decreased passenger capacity) or higher operating costs. The latter can be a difficult option for agencies, including the MBTA, who face constrained vehicle capacities at terminals and throughout the trunk. Light rail lines often do not rely solely on schedule-based control, but may dispense with the schedule and use other control strategies, often based on headways, when significant delays accrue. Consistent headway-based terminal and en route holding may reduce the need for more disruptive strategies such as short-turning, and increase the effectiveness of available resources. Adopting a headway-based holding strategy may require shifting supervisory resources to staff terminals and en route holding points. Schedule-based control is relatively simple to conduct, but headway-based control requires field officials to track multiple vehicles, requiring technological improvements to communicate real-time headways, and perhaps a decision-support system that recommends control actions.

These results guide further portions of this research. Chapter 4 creates a decision-support tool for controlling headways at terminals, which the simulation results support as the method of obtaining the highest benefits. Chapter 5 uses the tool in a real-world experiment, trialling headway-based dispatching at one Green Line terminal (Riverside).

# Chapter 4

# Decision-support tool for transit performance improvement and management

## 4.1 Introduction

The creation of tools for supporting decision-making in transit operations is not a new movement. The tools themselves, however, have evolved with technology. This chapter describes an application created for the purpose of improving operations performance and management on the MBTA Green Line. The chapter begins by listing several transit applications of decision-support tools and systems before outlining a framework containing the software architecture needed to operate such technology. An implementation of the framework for the Green Line is then described. The experiment run with the decision-support tool described in this chapter is presented in Chapter 5.

## 4.2 Applications of decision-support tools

Decision-support tools and systems can be used to manage different aspects of transit operations. These aspects include real-time control strategy implementation, fleet management, crew management, and the integration of the three.

## 4.2.1 Real-time control strategy implementation

The use of decision-support tools is integral for the effective implementation of consistent control strategies. Currently, service actions on the Green Line are generally chosen by a human controller based on his or her own set of heuristics—through practice and experience—with only basic rules (such as locations from which express segments can begin). This has the advantage of often assigning dispatchers or field inspectors who have had to previously address an expansive list of service situations. Experience-based knowledge—and the ability to successfully draw upon it—is unlikely to be uniform among personnel, however. For instance, different terminal inspectors have different strategies for how far ahead they plan for a departing trip with a late-arriving crew, or at what point a scheduled trip should be canceled.

Tools providing information on vehicle locations and predictions of when trains will reach stops can assist inspectors in determining whether a train will arrive in time to serve a scheduled departing trip, for example, with greater accuracy and precision than by calling for an operator's location via radio, which might not be returned promptly. Many field inspectors are already doing this using a variety of mobile applications on personal smartphones, but this is still far from a uniform or official practice.

Train locations, vehicle identification numbers, destinations, and stop arrival predictions make up the information currently available for the Green Line in real time internally and externally (via the MBTA-realtime API). With future vehicle fleets it may become possible to transmit data on passenger load or fare transactions in real time, although for now it is possible to estimate these based on historical data and ODX. All this information can be used to evaluate the full impacts of potential control strategies.

An intelligent decision-support system could, in theory, automate the simpler aspects of decision-making and management for events of lesser significance, thereby increasing the productivity of operations support personnel. Automated events may include holding at stations for very short durations or informing operators of their terminal departure times. This automation could be the norm at all times or only when more complex situations arise, allowing inspectors to focus on the less common issues requiring human intelligence and understanding. In situations requiring the input or approval of multiple inspectors before enactment, such as a short turn, the parties, each with their own decision-support tool,

could remotely view possible control plans simultaneously and discuss them on private audio channels. Paperwork can also be reduced. OCC dispatchers currently keep a log of major events, and terminal starters must submit daily logs of scheduled trips that were missed, information that would be easy to log automatically, with notifications asking inspectors to provide reasoning or approval as required. The information would also be available for supervisors to review and reference as events occur instead of making phone calls or waiting until the next day.

Finally, feedback can be made instant if a device showed adherence to instructions as well as performance metrics—such as average wait times or headways—in real time and as projected under proposed adjustments. The day's departures and other events could be reviewed as an annotated log by inspectors on different shifts or by supervisors.

### 4.2.2 Fleet management

A decision-support tool can be particularly useful for managing fleets, both in and out service.

Inspectors managing terminal dispatching are often conscious of the number of spare vehicles in the adjacent yard. At Riverside, for example, the number is usually above 40 at the start of the day before all the fleet is pulled out by the height of the peak periods. The *yard count*, as it is known, helps determine if a trip can depart on time when there are no trains available on the platform. An electronic device could aid, or even make suggestions, in this process by keeping track of trains in yards and adjusting the count as trains enter and leave service throughout the day. A display could show which tracks in the yard are occupied.

Such a system could help in the work of the *yardmaster*, the inspector in charge of organizing yard tracks each day, relaying yard occupancy to the terminal inspector, and notifying maintenance crews elsewhere in the yard of required repairs when a disabled train comes out of revenue service. A decision-support system could rely on distributing mobile devices to yardmasters and installing large-screen displays in maintenance facilities. If a train in service experiences a malfunction, an OCC dispatcher could enter the details of the mechanical issue upon receipt of the operator's radio message. The train would be directed to proceed to a specific terminal with maintenance facilities (often Riverside). The relevant yardmaster may miss the radio transmissions, but a notification may appear on his or her

tablet. He or she will then know to expect the faulty train and be able to place it on the proper track into the yard. Maintenance workers will see information beforehand, be able to prioritize the situation and, if urgent, make preparations so that work can begin as soon as the train arrives. Currently, such a situation would require multiple radio transmissions and telephone calls. The system described would reduce irrelevant radio traffic that inspectors have to filter out. Even more possibilities could be envisioned if digital terminals were installed in each vehicle, similar to the TransitMaster units currently installed on MBTA buses, as messages could be transmitted to and from vehicles without the use of radio.

### 4.2.3 Crew management

The management of vehicle operators is another area in which decision-support tools can aid terminal inspectors. Terminal starters act as the operators' first-line supervisors. An operator, for half or all of his or her run, is assigned to a single branch and is based at that branch's dispatching terminal. The terminal starter ensures that operators are checked in and fit to work prior to their first scheduled trip. They must also assign personnel coverage to trips when operators call in sick or when they cannot be at the terminal for a trip due to a disruption elsewhere on the Green Line. The decision-making process could be improved if inspectors needed to only enter the name or identification number of an operator or run to see a complete schedule for the day. Currently, a large book inside the inspector's booth or office must be consulted to definitively determine when an operator is off for the day. Scheduled trips listed on a tablet could include the number of trips or hours remaining for operators until pay moves to overtime rates.

Similar to fleet management, there is the potential to significantly reduce paperwork when operator information can be stored and displayed digitally to inspectors and supervisors on demand. Delay time (or operator overtime pay) is perhaps the single-largest source of paperwork: whenever an operator finishes a run later than scheduled, he or she writes the late time and reason on a small paper slip and turns it in to the terminal inspector, who in turn verifies the slip's validity and logs the details onto a separate log that is sent to the human resources (HR) department nightly. It is not difficult to imagine a system where overtime pay is handled automatically and out of the hands of inspectors whose primary responsibility is to run service. Abnormal patterns in delay time (or other attributes, such as running time or average speed) could be detected automatically and flagged for

further review by a human auditor—maintaining and perhaps enhancing accountability on operators.

Logging crew activity digitally can also make it readily available to agency departments which require it outside of Operations, such as HR (when conducting payroll) or the Transit Police (when investigating vehicle accidents or criminal activity onboard trains). The potential is here to reduce bureaucratic delays and human effort spent archiving and retrieving information when offices legitimately attempt to access crew data.

### 4.2.4 Integration of system management

The full potential of digital decision-support tools is realized only when the pieces described above—real-time tracking, fleet management, and crew management—are integrated as one product. Carrel et al. (2010) discuss how dispatchers make service decisions based on a number of factors beyond headways, including crew and fleet constraints. The objective function, in reality, is not simply to reduce passenger wait time but also to increase throughput, decrease operator overtime, and ensure future trips can be covered. Indeed, decisions in the MBTA OCC are seldom made without consideration to crew constraints: before a train's destination in the trunk is extended, for example, it is verified that the impacted operators are not on their final trips. Likewise, a train carrying operators running significantly behind schedule is more likely to be short-turned in an effort to reduce or eliminate delay time.

Much of the process of juggling several service-related factors could be automated in a device used by terminal inspectors to help with real-time control. As an example, consider an operator who has just arrived to a terminal 20 minutes late for his or her final round trip before a meal break. If the inspector assigns the operator to the next trip due to leave, a red warning sign may appear next to the operator's name—clicking on it would warn that the projected return time to the terminal would put him or her late for the meal break before asking the inspector if the plan is worth implementing. It would not be completely inconceivable for an agency-defined cost function to be used to recommend short-term service plans, balancing the monetary costs of overtime with the passenger benefits of running even headways. Implementing such an optimization would require monetizing all agency and passenger costs, which presents a difficult value proposition for agency managers.
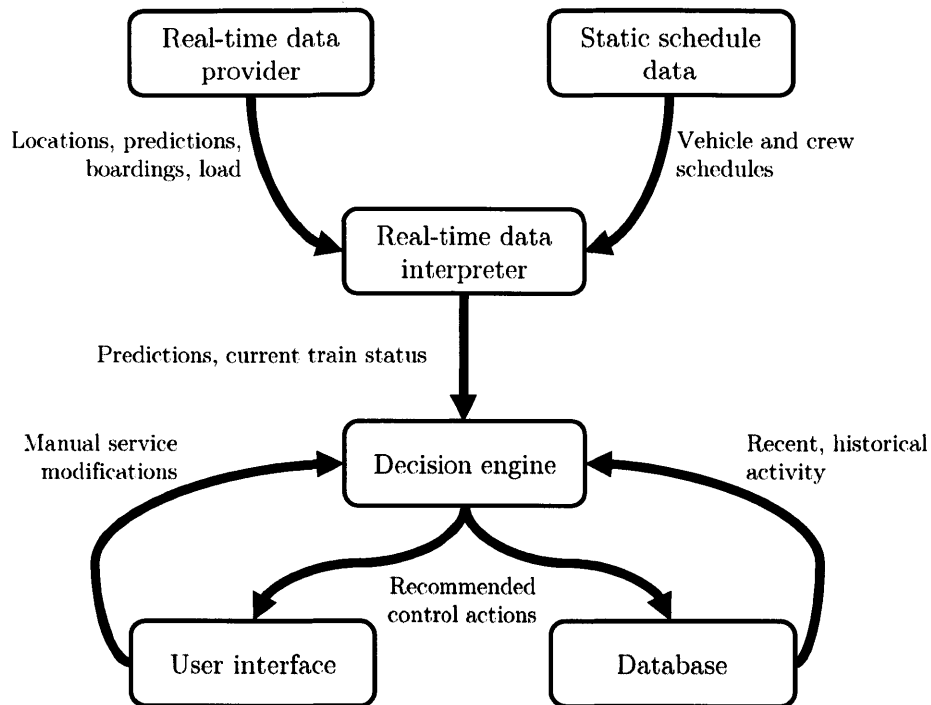
Figure 4-1: Software architecture framework for a transit decision-support system.

## 4.3  Decision-support tool software architecture

This section describes a framework containing the different software components required for a decision-support system capable of managing a transit line's fleet and crew while recommending control actions. The system software recommends service decisions which personnel can choose to follow or ignore. Figure 4-1, adapted from Maltzan (2015), summarizes the software components in a decision-support tool. Data flows both from real-time services and static schedules to a data interpreter, which in turn feeds train statuses and arrival predictions to the decision engine. The decision engine contains all possible service control actions and makes recommendations that are stored in a database and displayed by a user interface.

### 4.3.1  Information requirements for fleet management

Any system or tool fully able to assist in fleet management must be able to determine, store, or retrieve various information about each vehicle. Not every piece of information which follows is required for a tool that supports decisions on control actions, but better recommendations can be made with less input from terminal starters, yardmasters, and

others, with each addition.

- Real-time location of every vehicle in and out of service. In rail yards, this can be accomplished by adding AVI points or ensuring GPS is installed on each car.

- Counts of cars in yards which are available to put into service. This would preferably include accurate track allocations so that a terminal dispatcher does not have to call for the information (which often helps to determine switch placement).

- Counts of how many cars should be in service based on the schedule. This can be displayed as a percentage of the service scheduled to be operating at the time of day.

- Ability for tool users to consult a vehicle's location history for the current day.

- Ability for maintenance crews to look up a vehicle's inspection and maintenance record to determine patterns and schedule preventive maintenance. In the event that a vehicle must be inspected after a regular interval of miles in service, users can also be notified. Information about the vehicle can easily be documented digitally.

### 4.3.2   Information requirements for crew management

The following is information or features that would assist in digitally managing operators. Once again, not everything listed is necessary for a simple decision-support tool. A tool for controlling service can have a very simple interface such as that in Maltzan (2015). In the Green Line case study, however, Operations management suggested that an application resembling the current paper sheets might realize more buy-in from inspectors and supervisors. Possible requirements and features include:

- The ability for users to view an operator's past locations for the current day.

- Digitalization of the run schedule book. Each run and operator can be searched for clock-in, clock-out, and meal break times.

- Complete roster of all operators, including extra board operators.

- Coverage lists which can be accessed and modified remotely.

- Ability for HR to update the system when operators report or schedule absences. Full automation may entail an automatic assignment of scheduled coverage to trips when absences become known.

97

Table 4.1: Required and optional real-time information by control strategy.

| | Static schedule | Previous departure time | Predicted arrival of trailing vehicle | Passenger activity on control vehicle[d] | Downstream demand |
|---|---|---|---|---|---|
| Schedule holding | ✓✓ | | | ✓ | |
| Target headway holding | | ✓✓ | | ✓ | |
| Even headway holding | | ✓✓ | ✓✓ | ✓ | |
| Deadheading | | ✓✓ | ✓✓ | | ✓ |
| Expressing | | ✓✓ | ✓✓ | ✓ | ✓ |
| Short-turning | | ✓✓ | ✓✓ | ✓ | ✓ |

Legend: ✓ optional information, ✓✓ required information.

[d] Passenger activity includes not only the load but also the destinations of passengers onboard the vehicle. Passenger information could be inferred from the combination of static sources, such as ODX, and real-time headways.

### 4.3.3 Information requirements for real-time control

The information requirements for generating recommendations for real-time service controls is dependent on the strategy to be implemented. Maltzan (2015) details a number of strategies and the information which is required to be known in real time in order for a system to recommend them. Table 4.1 summarizes the minimum requirements for each strategy marked with a double checkmark. The most common information requirements are the times of the previous departure from a stop and the next arrival at the stop. Additional information which is not commonly quantified by transit agencies but would be useful in providing controllers with a more complete picture of operating service is denoted with a single checkmark. The lack of a checkmark does not mean that the particular piece of information is irrelevant to the respective control action but rather that it is likely to be less important. Maltzan (2015) walks through each case and gives the requirements in text.

### 4.3.4 Real-time data interpreter

A computer program must be included somewhere to interpret the incoming stream of real-time information. This program must include a direct call to one or more APIs which return vehicle locations, arrival predictions, and vehicle loads, typically in JavaScript Object Notation (JSON) or the Extensible Markup Language (XML) format. A prediction interpreter for transit data must take the following steps:

98

1   Call real-time APIs at a regular interval.

2   Store the API response as a set of objects, one for each record returned (for example, a distinct train with a location).

3   In the case of information related to a train (location, predictions, load, and boardings), link the returned object to a previously-observed train, and update fields.

4   If vehicle locations are available but downstream predictions are unavailable directly from available APIs, predict downstream location times based on historical running times.

5   Impute times, particularly terminal departures, if the closest location-recording point is more than a negligible distance downstream.

6   Update archival database entries with revised information if it does not first need to be processed through the decision engine. After this step, updated information will be shown via the user interface at its next refresh.

### 4.3.5   Decision engine

The decision engine for a real-time decision-support system is its most critical piece: the engine uses the information passed through the data interpreter and makes service, maintenance, or other types of recommendations to the end user. The engine is programmed with the decision rules for all permissible control actions. When updated information (not only related train locations, but also crew availability, etc.) is passed to the decision engine, decision rules related to each possible control action are applied, and updated recommendations are passed along to the database from which they will be shown in the front-end user interface.

### 4.3.6   Database

A database is necessary for storing the data passed from the real-time data interpreter and the decision engine. In terms of the system architecture, the front and back ends of the decision-support software are likely to be two separate applications. Bridging the gap between these two components is a database that the data interpreter, decision engine, and user interface (UI) can all read from and write to.

Any database used must meet basic requirements of security and reliability. Since the types of data available in the database may be sensitive or confidential in nature—particularly information related to individual operators—a system of permissions must be put in place so that personnel can access only the information necessary for their jobs, and so that misuse can be detected. Reliability must be ensured by placing the database (and all other software components) on a stable server secured from outside vulnerabilities. All data should be backed up regularly so that, in the event of a failure or security breach, records can be restored with no permanent loss.

### 4.3.7 User interface

The user interface (UI) is the only component visible to end users of the system. As Maltzan (2015) describes, there are three main roles to which the UI can be adapted: centralized dispatchers using desktop computers, field inspectors and supervisors using tablets or smartphones, and operators consulting basic data terminals installed inside train cars. A fourth interface can be added: a portal for analysts and supervisors to access archived information from previous days—a user-friendly alternative to querying the archival database directly. Separate UIs can be created to support multiple role types, or a one-size-fits-all approach can use only one UI for different groups of users. Regardless of which roles the application is adapted for, there are several basic requirements for all UIs:

1  They must show information relevant to the role and location of the specific user.

2  They must hide extraneous information that will be unimportant for the role and location of the end user.

3  They must allow for intuitive modification of the service plan (with few clicks needed to make common changes).

4  They must be responsive. End users must be able to access and interact with the application without significant lag.

## 4.4 Implementation for MBTA experiment

The following section describes an implementation of the decision-support tool framework in Section 4.3. The framework is implemented for an experiment that tests controlling service

by headway at the Riverside terminal of the MBTA Green Line in Newton, Massachusetts. The experiment itself is described in Chapter 5.

For this implementation, the framework components were developed as three applications and one database, the relationships of which are shown in Figure 4-2. The three applications are a logic program, an API, and a user interface. The rest of this chapter describes the applications in the context of the framework.

### 4.4.1 Vehicle location data

Vehicle location data comes from MBTA-realtime, a publicly-accessible API which developers can call to view, among other information, the locations of and predictions for trains on a given line. As noted in Maltzan (2015), data can also be downloaded using GTFS-realtime, which uses a standard format shared by transit agencies around the world. MBTA-realtime is chosen for use in this implementation for its ease of use (no special permissions are required), documentation, and general reliability (although uptime is less than 100%). Unlike GTFS-realtime, using the MBTA-realtime API does not require the system to download location data for all lines and stations, allowing users to select the ones of interest, which decreases download and processing time.

Data from MBTA-realtime pertaining to the Riverside branch of the Green Line is downloaded by the real-time data interpreter once every 10 seconds (the most frequent interval allowed by the API use policy). The API returns a JSON document with the location of each train on the branch, the time at which the train was recorded at the location, its destination, and predicted arrival times at all downstream stops.[1] Each train is then categorized as one of the following:

1 **Terminal-bound train** A train which is en route toward the controlling terminal and for which a prediction for arrival at the terminal is available.

2 **Terminal-bound but short-turning train** A train heading in the direction of and along the branch of the controlling terminal but which will be terminating early. On the Riverside branch, this is typically a train with a destination of Reservoir.

---

[1] Predicted arrival times for downstream stops are not currently provided for westbound trains east of Park Street. In-station signage lists the number of stops away a train is rather the time until arrival, because for trains at Park Street and east, it is difficult to predict movements with accuracy, since many trains are held.
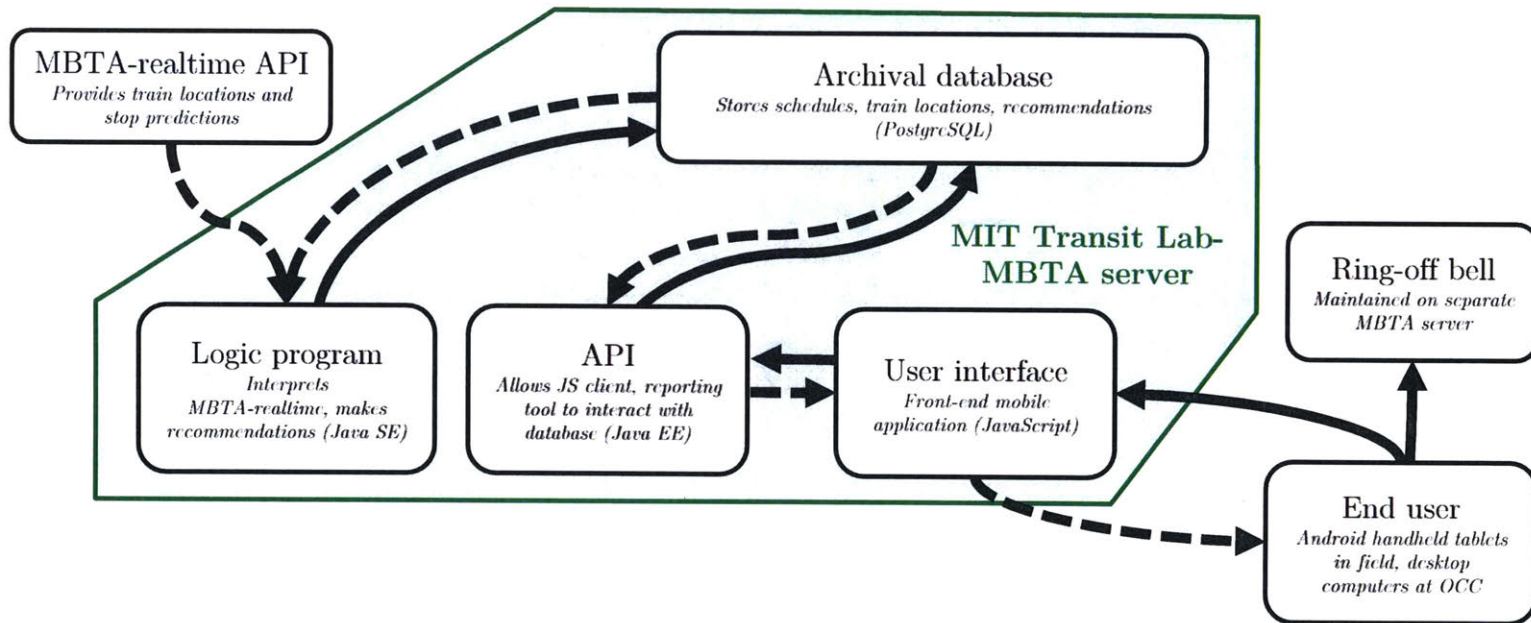
MBTA-realtime API
*Provides train locations and stop predictions*

Archival database
*Stores schedules, train locations, recommendations (PostgreSQL)*

MIT Transit Lab-MBTA server

Ring-off bell
*Maintained on separate MBTA server*

Logic program
*Interprets MBTA-realtime, makes recommendations (Java SE)*

API
*Allows JS client, reporting tool to interact with database (Java EE)*

User interface
*Front-end mobile application (JavaScript)*

End user
*Android handheld tablets in field, desktop computers at OCC*

Figure 4-2: Software architecture for the MBTA Green Line decision-support tool.

3 **Terminal-departing train close to the terminal** A train which is within two
   stops of the controlling terminal and is moving away from it.

4 **Terminal-departing train far from the terminal** A train more than two stops
   away from the controlling terminal and moving away from it.

Once this categorization is applied for each train, the location data are processed and
cleaned with the methods described in Section 2.4. If a train's AVL record is still believed
to be valid after cleaning and is not a duplicate of another record, an attempt is made to
link the train to an existing trip previously detected at another location by comparing train
consists and trip IDs. Depending on the categorization of the train, one of the following
processes or checks occur:

1 **Terminal-bound train** The predicted arrival time at the terminal, if changed, is
   updated. If the current time is more than 15 seconds after the predicted arrival time
   at the terminal, the train is assumed to have arrived. If the train has not yet been
   linked to a trip, the matching algorithm described in Section 4.4.3 is run, which infers
   the train's crew and scheduled trip based on its last departure from a control terminal.

2 **Terminal-bound but short-turning train** If the train was previously designated
   to go to the control terminal but has had its destination changed to short-turn, then
   the object relating to the train is updated so that it is no longer expected at the control
   terminal. If the trip was previously scheduled on paper, then the arrival at the control
   terminal is shown as canceled. Otherwise, the arrival is removed entirely. In either
   case, if a departing trip from the control terminal was recommended to use this train,
   the trip and the train are unpaired, and the trip is reassigned to the next available
   train arriving at the terminal.

3 **Terminal-departing train close to the terminal** If the train matches an existing
   trip, the algorithm checks whether the destination of the trip has changed. If it has, it
   updates the database. An unmatched train suggests that this is the first time the train
   has been tracked since departure. The algorithm estimates the terminal departure time
   from the current location and timestamp. This departure time is given by $t_g - r_g$, where
   $t_g$ is the time the vehicle is reported at location $g$, and $r_g$ is the average running time
   from the terminal to point $g$ (typically a geofence). The average running time from

the first geofence departing Riverside to each possible location $g$ is calculated from one month of AVL data. The remaining running time, from when a train begins to move upon departure from Riverside until it reaches the first possible geofence past the station is obtained from pre-experiment observations. After a departure time is estimated, the trip-matching algorithm is run and the matched trip is marked as having departed at the appropriate time. As an example, consider area $a$ to be the first geofence departing Riverside, and a train is observed traveling away from the terminal, detected at area $c$ at 15:30:25. If the average running time from the platform to area $a$ is 19 seconds, and that from area $a$ to area $c$ is 35 seconds, then it is estimated that the train departed Riverside 54 seconds before its detection at area $c$, at 15:29:31.

4   **Terminal-departing train far from the terminal** If the train matches an existing trip, the algorithm ensures that the destination of the trip has not changed. If it has, the database is updated.

The real-time data interpreter is part of a Java SE program installed on a server in the Massachusetts Green High Performance Computing Center (MGHPCC).

### 4.4.2   Crew and vehicle block data

Several data sources are incorporated in the decision-support implementation in order to provide vehicle and crew information. One of them is open to the public, but others are internal to the agency.

The schedule is published by the MBTA in the GTFS format. This includes all scheduled trips and their scheduled times at all stations, along with trip and vehicle block identifiers. From this information we can determine when vehicle blocks are due to be pulled out of and into yard as well as how many vehicles should be in service at any given time. The schedule does not specify which trains should be pulled from the yard but only uses block numbers as placeholders for available rolling stock.

Digitized crew schedule information was sent to MIT weekly. This was a digital version of the paper train sheets used by terminal inspectors, in the same tabular format, showing operators scheduled to depart from or arrive at a terminal at a scheduled time. The files were parsed by a Python script that summarized the information into a table containing all scheduled runs and operators for each day of the week. After parsing, the table is imported

104

to the database. A separate table, also obtained from the MBTA, correlates scheduled trips in GTFS to crew run numbers. These three sources are combined to provide end users with a schedule including all information from GTFS while adding flags indicating, for instance, if a scheduled trip is the first or last for a run. Any real-time updates to crew or vehicle availability are input directly via the user interface to the dispatcher or terminal inspector.

### 4.4.3 Scheduled trip matching

Several situations require the matching of a train to a scheduled trip. Even though departure times are being modified in the experiment, every departing train is still linked to a scheduled trip, unless it is explicitly added to supplement scheduled service. The situations requiring trip matching and the algorithms developed to link newly-seen trains to a scheduled trip and crew are as follows:

1    A new train is seen heading toward control terminal.

  a    All completed departures from the control terminal in the previous two hours are iterated in reverse chronological order. The most recent departure containing at least one of the cars which is now heading toward the terminal is selected. The two-hour window is sufficiently longer than the one-way running time on the Green Line but shorter than the time typically needed to make three one-way journeys. This ensures that the wrong trip is not selected if, for instance, the train switches branches for a trip or is not tracked for one or more trips.

  b    If such a departure exists, the runs of the crew operating the train are noted. The earliest available trip in the direction of the terminal containing ether of the runs is selected for the train, as long as it has not been marked as canceled. The scheduled arrival time must occur after the actual departure time of the previous trip.

  c    If no previous departure contains either of the cars, the train is likely being transferred from another branch (perhaps to undergo a mechanical repair). A new trip ID is generated for the train and the trip is considered a supplement to the regular schedule.

2    A train is heading toward the control terminal, has already been linked to an arriving

trip, but is no longer the furthest train away from the terminal traveling in the same direction.

a   This situation usually arises simultaneously with that in Item 1: once a train is no longer the furthest from the terminal heading toward it, the trailing headway can be calculated, and an initial recommendation can be made on the train's eventual departure time from the terminal. On the Riverside branch, this usually occurs more than half an hour before its arrival at the terminal.

b   The schedule of departures from the terminal is consulted, and the earliest scheduled departure without a recommendation (and which has not been canceled) is selected. This is trip $a$.

c   The previous step is repeated with the restriction that the selected trip cannot be scheduled as a pull-out from a yard during the early morning hours (before 07:00). This is trip $b$.[2]

d   The train heading toward the terminal is assigned to depart from the terminal as trip $a$. If trip $a$ and $b$ are not the same (i.e., if trip $a$ is scheduled as a pull-out), then trip $b$ is reassigned to now be covered with a pull-out from the yard, to ensure that the number of trains in service matches the schedule.

3   A new train is observed departing from the control terminal via manual observation input by an assistant supervising the pilot.

a   The scheduled trip which is selected by the assistant is matched to the departing train. During the experiment, trains were manually input by assistants because of the current shortcomings in the Green Line tracking system, as described in Chapter 2.

b   Unless the train is being pulled out of the yard, the train will have been predicted to depart on some trip. The algorithm looks at which trip the train was previously matched to and sees if this is the trip that the train is actually serving. If the trips match, no further work is needed. If they do not match, the trains of upcoming departures need to be rematched to trips, as explained in the next section. An

---

[2]Only scheduled pull-outs ahead of the morning peak period are considered, since pull-outs ahead of the afternoon peak period are currently dispatched as needed to fill scheduled trips without regard to the pull-out trip designations.

example of when this may occur is when two trains are dwelling at a terminal station and the first to depart is not first to have arrived.

   c   Since this departure was marked manually without accurate consist information, four minutes are allowed after the input to wait for AVL data (which contains the consist) to be reported by the train-tracking system. During this four-minute window, the actual departure time can still be used for the purpose of optimizing service.

4   A new train is observed departing from the control terminal using AVL data (because it was not manually input by an assistant).

   a   Scheduled departure trips (and any extra service added) for the terminal are iterated chronologically. The trip with the earliest recommended departure time that has not yet departed is selected to be matched to the observed train.

   b   If no trains to depart have recommended departure times, the first scheduled trip to depart is selected.

   c   As with manually-matched departures, trains are rematched to trips if the departing train was not previously matched to a trip. Since the observation is made using AVL data, the consist is known.

### 4.4.4 Rematching of upcoming departures

As mentioned in the previous subsection, there are several cases that require rematching trains to departure trips. These cases include the following:

1   A train departed from the terminal out of the order anticipated by the decision-support system.

2   An inspector indicates that an arriving vehicle previously marked to continue in service will now be pulled into the yard, or vice versa.

3   An inspector indicates that an upcoming departure will be served by a train coming from the yard instead of a train already in service.

4   An upcoming departure trip is canceled.

Table 4.2: Example effects of a trip cancellation.

| (a) Before cancellation. | | | (b) After cancellation. | | |
|---|---|---|---|---|---|
| Trip ID | Scheduled time | Train serving trip | Trip ID | Scheduled time | Train serving trip |
| A | 09:10 | 3642-3888 | A | 09:10 | 3642-3888 |
| B | 09:15 | 3615-3805 | B | 09:15 | |
| C | 09:20 | 3868-3653 | C | 09:20 | 3615-3805 |
| D | 09:25 | 3808-3710 | D | 09:25 | 3868-3653 |

5   A scheduled trip shown as departed in the decision-support tool does not match with the crew which actually departed. The associated train must be moved forward or backward one or more trips so that the train is correctly aligned with the operators.

How departing trips are reorganized is best explained through the use of examples. First we consider an example of a trip needing to be canceled shortly before its recommended departure time from Riverside due to a crew shortage, as shown in Table 4.2. Before the cancellation, the decision-support software anticipates that train 3615-3805 will serve trip $B$, originally scheduled for 09:15, as it will arrive several minutes beforehand. The following trip $C$, scheduled for 09:20, is to be served by train 3868-3653. Trip $D$, scheduled for 09:25, is anticipated to be served by train 3808-3710 based on real-time AVL data. After trip $B$ is canceled, trains are rematched to trips such that train 3615-3805 serves trip $C$, train 3868-3653 serves trip $D$, and train 3808-3710 serves the subsequent trip.

The second example uses the same trains and scheduled trips, and is shown in Table 4.3. The terminal inspector at Riverside learns that the approaching train 3615-3805 has a door malfunction and cannot be put back into revenue service. Using the decision-support tool, the status of train is manually changed to *Pull-in*, designating that the train will not be used for a departing trip upon arrival. As a result, trip $B$ is reassigned to train 3868-3653, trip $C$ reassigned to train 3808-3710, and trip $D$ is reassigned to some other train.

The routines for reorganizing trips are located in multiple locations across the application so that updates to the user interface can occur as quickly as possible. Changes initiated by an end user implement any required reorganizing within the API interfacing with the front-end (see Section 4.4.6).

Table 4.3: Example effects of a train pulled into the yard.

(a) Before train status modification.

| Trip ID | Scheduled time | Train serving trip |
|---------|----------------|--------------------|
| A | 09:10 | 3642-3888 |
| B | 09:15 | 3615-3805 |
| C | 09:20 | 3868-3653 |
| D | 09:25 | 3808-3710 |

(b) After train status modification.

| Trip ID | Scheduled time | Train serving trip |
|---------|----------------|--------------------|
| A | 09:10 | 3642-3888 |
| B | 09:15 | 3868-3653 |
| C | 09:20 | 3808-3710 |
| D | 09:25 | TBD |

## 4.4.5 Database

Data for the experiment was stored in a PostgreSQL relational database configured to allow only authorized developers, analysts, and applications to access and modify its contents. The database was used in real time to power the decision-support tool and to archive results and conduct analyses. A number of tables were created in the database. The following is a list of the tables and a brief description of each:

1 **stop_visits** Contains all terminal stop visits, along with assigned crew and vehicle information, for past service days. The table columns are similar to those in the paper train sheets, but with added fields to store information such as whether the operator has signed in and which inspector was on duty when the stop visit took place.

2 **stop_visits_today** and **stop_visits_tomorrow** Identical to *stop_visits*, except that the contents are only the stop visits which have occurred or are scheduled to occur during the current or next service days, respectively. These tables are kept separate from the archive to reduce retrieval time from the database by the real-time decision-support system.

3 **actions_log**. Logs the user, timestamp, location, and type of every modification made using the front-end UI. Inspector sign-in and sign-out times are also recorded. The logs were used for debugging the application before and during the experiment. In practice, they could be used to audit usage of the tools.

4 **comments, comments_today, and comments_tomorrow** Contains all comments input manually by end users as well as automatically when special actions are taken. Comments can be linked to a trip, vehicle, operator, or single stop visit, and may follow the object for the rest of the service day.

109

5   *operators*   A correlation table of operator identification numbers (or badge numbers) and operator names. This table is joined with *stop_visits_today* when train information is shown in the UI, since the former contains only badge numbers.

6   *run_assignments*   See Section 4.4.2.

7   *signature_images*   If operators sign in using the mobile device, an image file of the signature is stored into this table along with a timestamp and location. The sign-in feature was not used during the experiment.

8   *users*   A table of all authorized application users and hashed credentials.

## 4.4.6   Client-side application

For most users, the client-side application and its user interface is the method of interaction to the decision-support tool. The front-end is built as a browser-accessible web application, built primarily using the JavaScript language supported with the jQuery, jQuery Mobile, and Knockout libraries. Responsive web design is used to ensure that the user interface is easily-viewed and used on both tablet and personal computer devices. Backwards compatibility is ensured so that all features can be used on the OCC terminals using legacy browser versions. The entire application is hosted on an Apache Tomcat web server.

The default user interface, shown in Figure 4-3, is the view that inspectors and other users of the tool see the majority of the time. The layout can be broken into three sections:

1   An upper bar with the current time and administrative features. From left to right, the administrative features are as follows:

   a   The location at which the decision-support tool is being used.

   b   The identification badge number and name of the inspector signed into the tool.

   c   A button to scroll to the next upcoming trip.

   d   A button to activate the ring-off bell on the terminal platform.

   e   A switch between the schedule of the current and following service days (for use when making next-day crew substitutions).

   f   A logout button to switch between users or terminal locations.

110
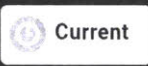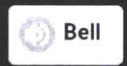
g    A notification showing service alerts for the branch. The button appears only when alerts are issued to the public and can be clicked to show more details.

2    A lower bar with the following figures, from left to right:

a    A count of cars in Riverside Yard. This is only a crude approximation based on pull-outs and pull-ins, as a demonstration, since there is no access to real-time train locations in yards.

b    The number of cars in service on the Riverside branch in each direction based on real-time train locations.

c    The previous departure headway.

d    The current scheduled headway.

e    The average of the recommended headways of the next two departures.

3    The main section showing scheduled and actual service. The section contains a table which shows service to and from the control location for the entire day. Scrolling the section shows either earlier or later service, and a click of the *Current* button in the upper bar returns to the current time of day. The left half of the section shows trains arriving to the user's location, while the right half shows departing trains. Each row in the table represents a single car in service. A row with information only on the left side depicts a train being pulled into the yard upon arrival, while a row with only the right half filled is a train being pulled out. A train will almost always occupy two rows in the display, since trains usually run with two coupled cars. Trains are separated in the table with a thin, horizontal line across the screen.

a    The car column shows the identification number of each car in a train. (The leading digit "3" is omitted, following current practice and for space considerations.) The second car of a train is labeled with a "T" after the car number, and the third car of a train (if applicable) is labeled with a "K", following current practice on the paper sheets.

b    The *Trip* column in the departures half of the table shows the originally-scheduled time for a trip, which may differ from the recommended or actual departure time as shown in the *Time* column.

c   Some trips may have a button within the *Note* column. Such a button means that a comment has been added to the trip, stop visit, operator, or car corresponding to the row (in the arrivals or departures half). The button can be clicked to view the comments.

d   The button next to every car-trip under the *Edit* column opens a dialog window from which modifications can be made, as described later in this section.

Separating arrivals and departures to the left and right sides of the display is reminiscent of the paper sheet format, except that a car's arrival and departure are shown together in a single row instead of two. The decision was made for each row to represent a single unit of rolling stock rather than an operator, since the top priority for terminal inspectors is to dispatch vehicles. The display for a train is likely also to be simpler across the left and right halves because train cars are seldom coupled or uncoupled on the platform, whereas one crew member on a train may be assigned to a meal break upon arrival at the terminal and the other may immediately serve another trip.

An effort is made through the use of symbology to distill the displayed information to only what is necessary. Arrival and departure times in the main section of the user interface can be displayed in parentheses and italics without seconds to indicate they are scheduled times, in large and bold text with seconds to indicate they are times for upcoming arrivals or recommended departures, or in the standard script and size with seconds to indicate they are actual times of completed arrivals or departures. Row background colors also provide context, as exemplified in Figures 4-3, 4-4, and 4-5. The different colors used are listed and defined as follows:

1   **Dark gray (default background color)** The trip has already arrived or departed and thus no action needs to be taken.

2   **Green** The trip has yet to occur. A recommended time is displayed under the *Time* column indicates that the arrival or departure will occur in the near future.

3   **Light gray** An arriving train to the control terminal is being sent to the yard rather than than immediately returning to service.

4   **Black** The car will not run in revenue service for the related trip. A button appears under the *Note* column to view further information. If the text in the row is crossed

**9:34:55** am  **Riverside** NA  Current  Bell  ⚠ Delay

### Arrivals

| Car | Time | Run | Badge | Name | Note Edit |
|---|---|---|---|---|---|
| 829 | 9:19:20 | 543 | | | |
| 675 T | 9:19:20 | 544 | | | |
| | | | | | |
| | | | | | |
| 861 | 9:33:49 | 1420 | | | |
| 821 T | 9:33:49 | 1417 | | | |
| 642 | 9:41:05 | 522 | | | |
| 841 T | 9:41:05 | 523 | | | |
| 843 | 9:42:37 | 530 | | | |
| 624 T | 9:42:37 | 531 | | | |
| 801 | 9:45:23 | 1424 | | | |
| 615 T | 9:45:23 | 1423 | | | |

### Departures

| Car | Time | Trip | Run | Badge | Name | Note Edit |
|---|---|---|---|---|---|---|
| 675 | 9:29:10 | (9:27) | 517 | | | |
| 829 T | 9:29:10 | (9:27) | 518 | | | |
| | 9:36:40 | (9:34^) | 1420 | | | |
| T | 9:36:40 | (9:34^) | 1417 | | | |
| | 9:44:40 | (9:41) | 522 | | | |
| T | 9:44:40 | (9:41) | 523 | | | |
| | 9:52:40 | (9:49) | 1424 | | | |
| T | 9:52:40 | (9:49) | 1423 | | | |
| | 10:00:40 | (9:57) | 537 | | | |
| T | 10:00:40 | (9:57) | 593 | | | |
| | 10:08:40 | (10:05) | 569 | | | |
| T | 10:08:40 | (10:05) | 607 | | | |

Yard **11**  WB cars **22**  EB cars **15**  Prev **9.6m**  Sched **7.5m**  Recommend **7.8m**

Figure 4-3: Standard view of decision-support tool in use for the Green Line experiment.

Figure 4-4: Black and light gray background colors for canceled (top rows) and pull-in (bottom rows) trips.

out, the trip or car has been canceled by a user and will not run at all.[3] If *NR* appears next to the time in the row, the trip is run as a non-revenue train.

5 **Yellow under the *Run, Badge,* or *Name* columns** Crew information cannot be inferred or has not been assigned. The color is also paired with *?* symbols under the relevant rows.

6 **Blue under the *Car* column** A trip is en route to or from an unscheduled location opposite the control terminal. This color typically indicates that a short turn has taken place. This is dependent on the branch: for the Riverside branch, any trip to or from Riverside that does not originate or terminate at Government Center is considered unscheduled.

7 **Red** A recommended departure time exceeds the allowable headway due to a lack of equipment or personnel. This color acts as a warning for the inspector or dispatcher to take a corrective action, usually a pull-out. Red also highlights undesirable performance or yard availability in the lower bar (see Figure 4-5).

8 **Yellow for the entire departing half of a row** The trip is recommended to depart within 30 seconds. The button for the ring-off bell should be pushed to ensure an on-time departure.

Manual changes to the status of trains, crews, or trips are made through pop-up dialog windows accessed by clicking on the button in the *Edit* column for the desired train car. The edit dialogs differ for arrivals (see example, Figure 4-6) and departures (see examples, Figures 4-7 and 4-8). Along the top of the window is a summary of the trip being modified:

---

[3]A car can be considered canceled if, for example, a trip scheduled to be served by a two-car train is served by a one-car train.

**12:58:48** pm                    Riverside  NA   Joshua Fabian   [Current]  [Bell]

| Arrivals | | | | | | | Departures | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Car | Time | Run | Badge | Name | Note | Edit | Car | Time | Trip | Run | Badge | Name | Note | Edit |
| 642 | 12:29:53 | 544 | | | | ✎ | 841 | 12:34:48 | (12:35) | 531 | | | | ✎ |
| 841 T | 12:29:53 | 543 | | | | ✎ | 642 T | 12:34:48 | (12:35) | 530 | | | | ✎ |
| | | | | | | | 618 | 12:41:35 | (12:42 X) | 537 | | | | ✎ |
| | | | | | | | 817 T | 12:41:35 | (12:42 X) | 535 | | | | ✎ |
| 878 | 12:45:43 | 570 | | | | ✎ | 652 | 12:50:08 | (12:50) | 544 | | | | ✎ |
| 652 T | 12:45:43 | 501 | | | | ✎ | 878 T | 12:50:08 | (12:50) | 543 | | | | ✎ |
| 840 | 12:56:05 | 583 | | | | ✎ | | 12:59:05 | (12:57) | 546 | | | | ✎ |
| 641 T | 12:56:05 | 556 | | | | ✎ | T | 12:59:05 | (12:57) | 545 | | | | ✎ |
| 829 | 12:56:39 | 564 | | | | ✎ | | 1:06:35 | (1:05) | 626 | | | | ✎ |
| 675 T | 12:56:39 | 502 | | | | ✎ | T | 1:06:35 | (1:05) | 627 | | | | ✎ |
| 848 | 1:07:10 | 1420 | | | | ✎ | | 1:14:05 | (1:12) | 593 | | | | ✎ |
| 716 T | 1:07:10 | 1417 | | | | ✎ | T | 1:14:05 | (1:12) | 551 | | | | ✎ |
| 877 | 1:10:27 | 1424 | | | | ✎ | | 1:21:35 | (1:20) | 556 | | | | ✎ |

**Yard 12**  WB cars **19**  EB cars **10**  Prev **8.6m**  Sched **7.5m**  **Recommend 8.0m**

Figure 4-5: Instant feedback concept and indicators recommending that actions be taken.

colored banners provide explanations for the colored rows in the main table. Toggles and buttons below the summary allow the user to make the most common changes, including adding a comment or pulling a train into or out of the yard. Some of these actions open secondary dialog windows such as the shown in Figure 4-9 for adding a comment. The main edit dialog contains text boxes for modifying car numbers, runs, or operator identifiers. In the case of an operator, a change can applied to either the upcoming cycle (to and from the terminal) only or to the rest of the original operator's scheduled trips for the day. The button labeled *Special actions* opens a drop-down menu with options used less frequently.

Refinement of the interface occurred after periods of user testing at Riverside. Before the final iteration of the UI, two columns appeared on the left and right halves of the main section labeled *Code*. The columns contained, for certain trips, symbols indicating trips that were scheduled as pulls to and from the yard or trips which started or completed a run or led to a meal break. The symbols came from printed train sheets, but initial feedback found that few inspectors used the information on a regular basis, and many did not even know the meaning of every symbol, so the column was removed. The column was considered redundant because users could click on a car number, run, or operator in the main table
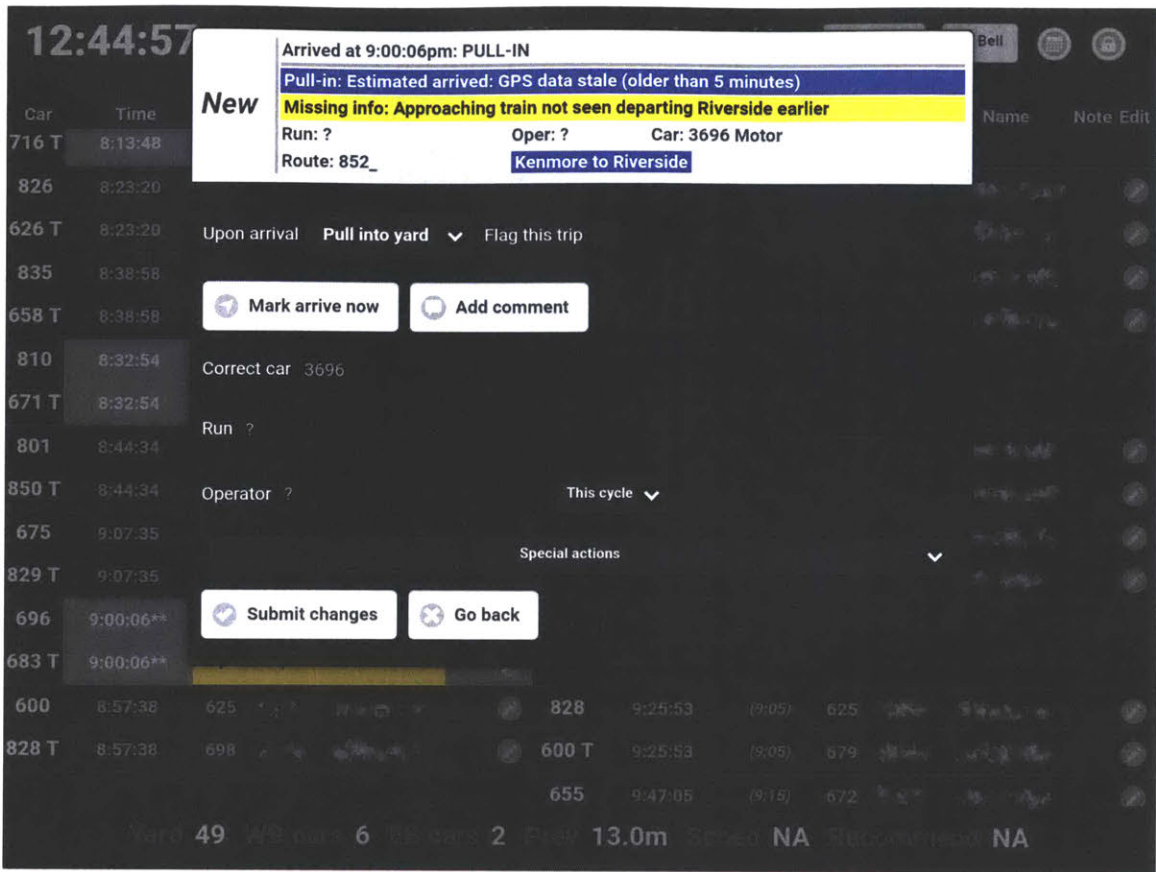
Figure 4-6: Main dialog window for an arriving train originating from another branch.
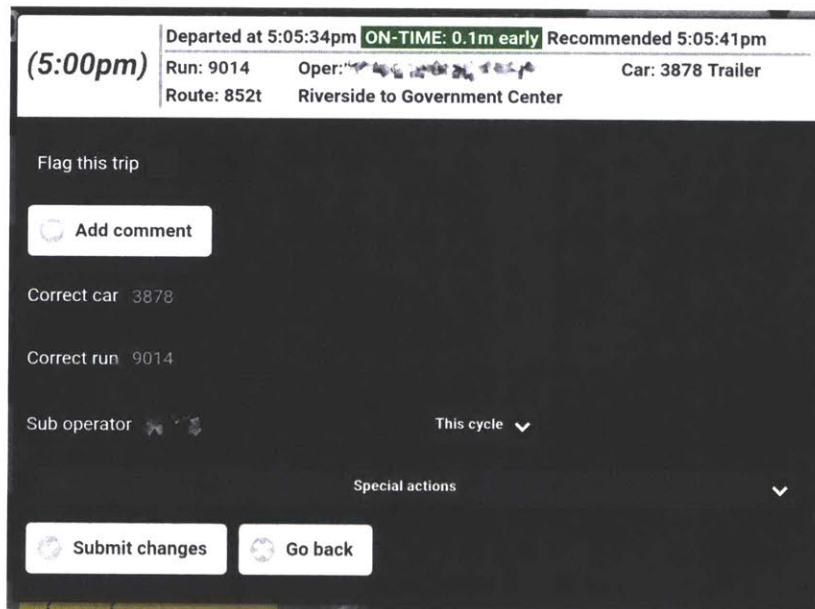


Figure 4-7: Main dialog window for a departing train.
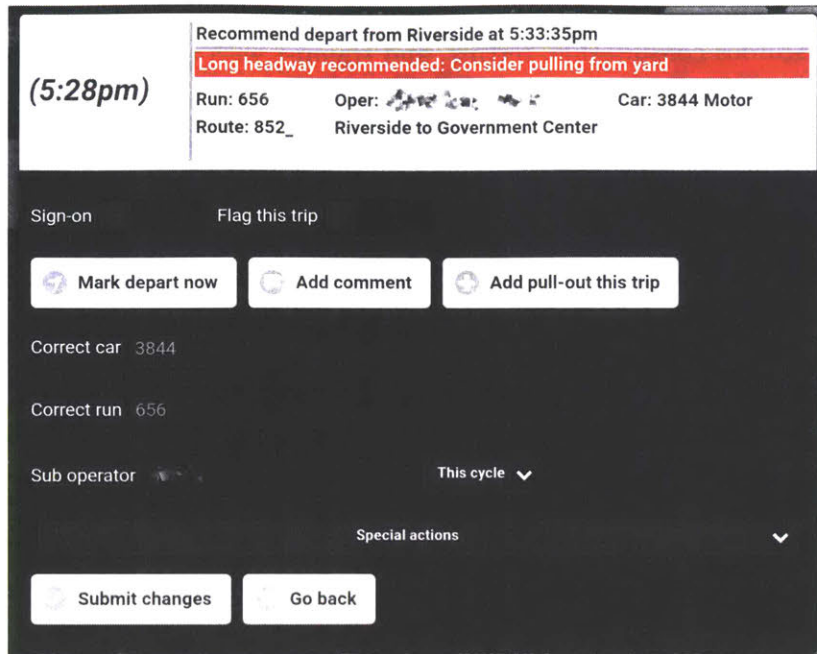
Figure 4-8: A banner explains why a recommended departure time is colored red.
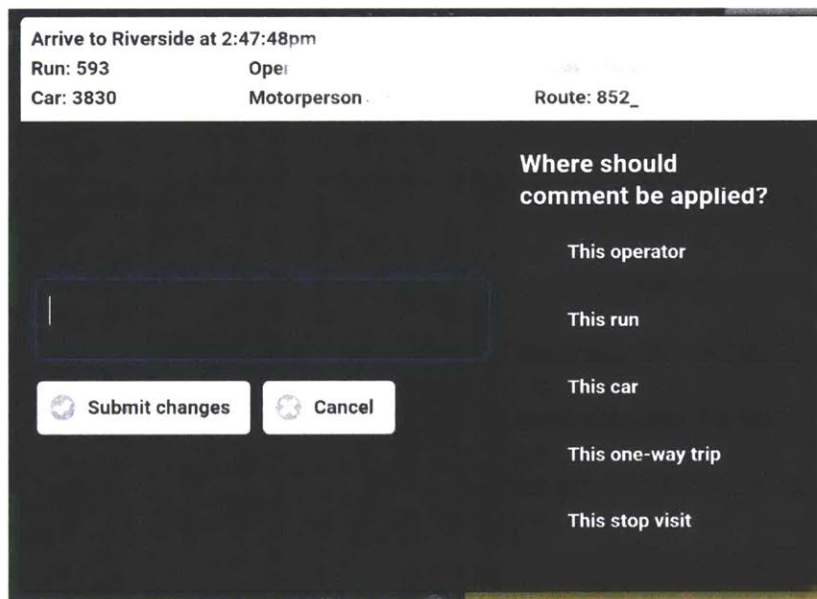


Figure 4-9: Dialog opened upon clicking the *Add comment* button.

to view only desired trips during the service day. Nevertheless, at least one complaint was made by a user after the removal of the column, suggesting that a permanent system must undergo more rigorous consultation, testing, and training with prospective users.

The removal of the *Code* columns also made way for the *Trip* column for departures. Before this addition, scheduled departure times were only shown in the editing dialog window when the trip had a recommended time. This feature was designed to discourage inspectors from referring to the scheduled time, but resulted in confusion as inspectors would often refer to scheduled times when directing trains to depart early or late. Using an example from Figure 4-3, a terminal inspector may announce to waiting operators that "the 10:05" is instead leaving at 10:08:40. In order to smooth the transition from dispatching by schedule to dispatching by headway, the *Trip* column was added so that it could be used as a reference point. The addition was well-received by users who could better compare what the algorithm was recommending with the paper sheets.

There are more features which a decision-support device could have but which were not implemented for the described Green Line experiment. One example is shown Figure 4-3, where small red squares are drawn next to the 12:41 departure. These squares, a prototype that was not used in the pilot, were meant to give instant feedback on recommendation adherence to inspectors. Many red squares would suggest that the official was doing a poor job of ensuring that operators left on time. Since the squares are visible on the main screen, the theory is that inspectors may become more motivated to ensure on-time departures. The green banner in Figure 4-7 shows how recommendation adherence feedback was incorporated into the dialog window for the final version of the UI.

When the main screen is left open without dialog windows, the train data are updated on the front-end application every seven seconds; while dialog windows are open, updating is suppressed. While the refresh rate of the interface is shorter than that of the back-end application, they are not synced together, meaning that in some cases, the effects of a UI-initiated change may not be seen by the user for seventeen seconds. This would commonly result in a significant delay between an action and its effects, which annoyed several end users. Future implementations must either reduce this delay or clearly indicate that modifications are being processed and prevent further user actions in the meantime.

Acting as an intermediary between the database and the user interface is an API. This API is written as a Java EE web application and is hosted on the same web server as the

front-end application. The API is accessed using HTTP POST and GET methods sent with a body containing the changes required or a request for train information along with an authentication token in JSON. For calls requesting information, portions of the database tables are returned to the front-end client as a JSON array.

### 4.4.7   Departure ring-off bell

An important logistical problem is the communication of recommended departure times to operators. Currently, departures are scheduled to the minute, and adjustments made by inspectors are never at sub-minute intervals. At Riverside, the terminal inspector is usually situated at the end of the platform opposite of departing trains, making it difficult for him or her to speak to operators once they have boarded the train, except by way of radio, which can be heard by all operators and supervisors on the Green Line. This leads to departures which are not made on time, as operators use different methods to determine when to depart. Any success at implementing a control strategy using sub-minute accuracy is dependent on ensuring that operators know precisely when to depart. It is unreasonable, under current conditions, to expect an operator to depart at 40 seconds after the minute, for example.

This decision-support implementation included the installation of an audible bell that notifies operators when they should depart. A similar system of "ringing-off" an operator exists in many metro systems, including at heavy rail terminals on the MBTA network. The departure bell at Riverside was configured as a 15-second audio file that could be played through the existing public address system on the platform. This is the same system which announces the arrival of trains using speakers and LED displays at downstream stations. On the main user interface, the *Bell* button sends a message to an MBTA server instructing the public address system to play the audio file. The audio levels on the platform were adjusted by a technician to ensure that the bell could be heard from within the train cab. A standalone application was also made to activate the bell with a single tap of a large-sized button, as shown in Figure 4-10. At most times the terminal inspector had one tablet open with the main dispatching application next to another tablet with the companion bell-ringing application open. The size of the bell was designed on purpose to encourage inspectors to utilize it on a routine basis.

119
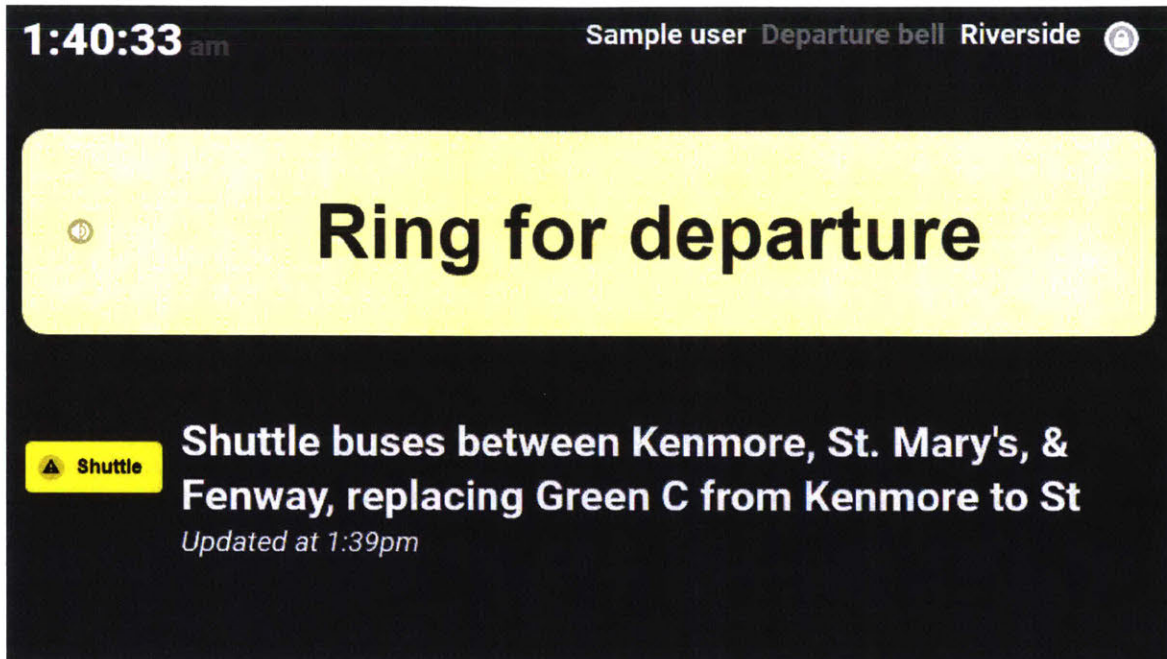
Figure 4-10: Screenshot of companion application for the platform ring-off bell.

## 4.5 Summary

This chapter began by presenting several applications of decision-support tools and systems. These applications include implementing real-time control actions, vehicle management, and crew management. Additionally, the three applications can be combined to manage transit at all levels, since real-time service decisions, for example, often require updated information on spare vehicles and operator coverage.

Descriptions of the applications are followed by an outline of the software architecture needed to implement a decision-support tool. Real-time and static schedule information are fed into a data interpreter, which updates the current state of the transit line or system. The state is read by the decision engine, which makes recommendations on service adjustments in real time. The recommendations are displayed in the user interface and stored in a database.

The framework was implemented for an experiment on the MBTA Green Line. The MBTA-realtime API is used to obtain vehicle locations and predictions, and each train in the feed is matched to a trip and crew using internal schedule data. Trains moving toward a terminal are assigned a trip, crew, and time after arrival at which they should depart again, unless the end user indicates otherwise. A user interface is created to provide dispatchers and inspectors with recommendations and to allow them to adjust the service plan. A database

stores all of the data necessary to feed the decision-making components and display service plan details to authenticated users of the decision-support tool.

# Chapter 5

# Field experiment and results

An experiment to control departing headways and modernize record-keeping using the decision-support tool described in the previous chapter was implemented on the Riverside branch of the Green Line. The pilot was conducted at Riverside terminal over ten weekdays in early 2017: Monday, February 20 through Friday, February 24, and Wednesday, March 1 through Tuesday, March 7. Weekends were excluded, and the experiment lasted for the entirety of each of the included service days. This chapter describes the experiment's implementation and discusses the results.

## 5.1 MBTA Green Line Riverside branch

The Riverside, or Highland, branch of the MBTA Green Line stretches 19 kilometers (12 miles) between Government Center station in downtown Boston and Riverside station in Newton. It is by far the longest of the Green Line's four branches and serves 20 stops per direction. The route's length makes it attractive to commuters from outer suburbs west of Boston. Combined, the two western-most stations feature about 1500 parking spaces (Massachusetts Bay Transportation Authority, 2014). Its surface-level, branch portion, west of Kenmore station, is also unique to the Green Line because it operates at higher speeds along a dedicated right-of-way, without the interference of automobiles. Figure 5-1 shows a diagram of the Riverside branch and its stations, with service in downtown typically terminating at Government Center.

Original source: Central Transportation Planning Staff (2016). Modified by author.

Figure 5-1: Schematic of Riverside branch of the MBTA Green Line.

### 5.1.1 Headway-based control strategy

In consultation with MBTA Light Rail operations officials, a hybrid departure policy combining the target headway (Turnquist, 1981, Daganzo and Pilachowski, 2011) and even headway (Turnquist, 1981) strategies was adopted at Riverside for this experiment. The base of the experimental policy is the Prefol even headway strategy, which equalizes the leading and trailing headways around a controlled train as shown in Equation 5.1.

$$d_{ij} = \frac{d_{i-1,\, j} + (a_{i+1,\, j} + l_j)}{2} \tag{5.1}$$

where:

$d_{ij}$ = departure time for trip $i$ from terminal $j$,

$a_{i+1,\, j}$ = predicted arrival time for trip $i + 1$ to terminal $j$, and

$l_j$ = minimum required layover time at terminal $j$.

Passenger-related dwell time at terminals is not considered because it is assumed that passengers can alight and board during the minimum layover time. $a_{i+1,\, j} + l_j$ is the earliest

estimated departure time for trip $i + 1$. The adopted policy adds the following constraints: any train cannot depart until both the minimum layover, $l_j$, and the target headway at time of day $t$, $h_j t$, have elapsed. This is shown in the following:

$$d_{ij} = \max\left(\frac{d_{i-1,\,j} + (a_{i+1,\,j} + l_j)}{2}, a_i + l_j, a_i + h_{jt}\right) \tag{5.2}$$

At most times, the target headway is equal to the scheduled headway, $h_{jt} = h_{\text{schedule},\,jt}$. Following a large gap in service, defined by the MBTA as 1.5 times the scheduled headway (Massachusetts Bay Transportation Authority and IBI Group, 2016), $h_{jt} = 0.65 h_{\text{schedule},\,jt}$ for the following trip. Following a very large gap longer than twice the scheduled headway, the adjustment is made for the following two trips. These adjustments are made to alleviate concerns of reduced capacity after large gaps if the first train following the gap leaves passengers behind on the platform. In the experiment, these situations were seldom encountered, as long gaps were often filled by pulling extra trains out of the yard and into service.

### 5.1.2  Context within agency organization

The experiment was conducted with the assistance of the MassDOT/MBTA Office of Performance Management and Innovation (OPMI), which conducts data analysis and initiates projects not only in Operations but in all facets of the transit agency. Staff from OPMI acted as assistants during the experiment, as described in Section 5.1.4. Supervisors from the Light Rail Operations department, which oversees the Green Line, as well as agency-level Operations officials, were consulted on and gave approval to pilot specifics. Light Rail officials issued a memorandum during the second week of the pilot asking operators and inspectors to adhere to departure times calculated on the experimental tablet. The memo was posted in the terminal break room and dispatching booth but was not issued to every operator in the way a more binding special order would.

### 5.1.3  Operations personnel

The operations personnel involved in the experiment included operators, inspectors, and supervisors. OCC dispatchers observed the user interface (UI) throughout the pilot, often using it to determine which operators were on which train, though they did not use it directly

to make control decisions.

Field supervisors oversee inspectors and operators. For the first two days of the pilot, one supervisor monitored operations from within the official's booth at Riverside and coordinated the training of officials who were not regularly scheduled to work at Riverside.

Officials regularly scheduled to work at Riverside terminal were the primary users of the decision-support tool. On a normal shift, these officials are responsible for signing in operators, ensuring that they are fit to operate a train, dispatching vehicles and crews toward Boston, and marking departures, arrivals, and schedule modifications on the pre-printed paper train sheets. During the experiment, they had the added responsibility of using the support tool, and were instructed to dispatch trips at the time given by the headway-based algorithm rather than on schedule. During the second week only, the ring-off bell (see Section 4.4.7) was operable via the support tool, and officials were told to ring the bell when indicated in the user interface (30 seconds before the recommended departure). In general, the recommended departure time was communicated to operators orally, even when the ring-off bell was in use.

Most Green Line inspectors and supervisors received a short training session by assistants February 20–21. These trainings were held during their normal shifts at other locations and lasted 15–20 minutes; other officials filled in for their regular duties while they were at Riverside for training. Since officials' schedules often fluctuate due to vacations and sick leave, an effort was made to train even those officials who were not scheduled to work at Riverside. Nine different officials had shifts at Riverside during the two weeks of the experiment.

### 5.1.4  Support personnel

Assistants from OPMI and the MIT Transit Lab provided on-site support at all times of the experiment. After an initial, hands-on training at the start of an official's first shift with the tablet, assistants monitored the official and ensured that they made all necessary adjustments through the user interface. If the official encountered difficulties using the application, the assistant would step in to provide technical support. Three tablets were present at the terminal, one for inspectors, one for assistants, and a third as a spare but also used as a stand-alone activation point for the ring-off bell. Assistants took notes on algorithm adherence, official and operator actions, feedback from all parties involved, and other

126

service-related events. If the official was not telling operators to leave at the recommended time, assistants would determine why and attempt to convince the official to adhere to the algorithm departure times, noting such situations. Finally, assistants were responsible for obtaining and returning tablets to a secure location at the station at the start and close of service each day.

In total, 13 assistants, 11 from OPMI and 2 from MIT, had shifts during the two-week pilot. The shifts ranged between 4 and 7.5 hours in duration. During peak periods, two assistants were often present in order to provide support during busiest times—for example, one assistant can monitor trains on the platform while the other follows the official's planning and actions. The assistants themselves attended one overview training session and one hands-on training session the week before the pilot began. And a debriefing was held after the pilot to obtain their feedback.

## 5.1.5 Pre-experiment baselines

The baseline period consists of weekdays within the three weeks before the experiment, Monday, January 30 through Friday, February 17, 2017. Three weekdays (Tuesday, January 31, Thursday February, 9, and Monday, February 13) were excluded from the baseline due to significant snowstorms, as service and ridership on those days were not representative of normal weekdays. Tuesday, February 7, was also excluded due to a sports parade in Boston which led to added service and increased ridership. For the remaining eleven weekdays, periods with severe delays or during which service was suspended on the Riverside branch (as described in public-facing service alerts) were also excluded.

For all baseline and experimental data presented later in this chapter, the time-of-day periods and associated scheduled headways used in the analyses are shown in Table 5.1. The non-integer scheduled headways listed come from patterns in the existing schedule, since the schedule is rounded to the nearest minute. For example, a 5.33-minute headway means that trains are scheduled to leave with headways in a repeating pattern of 5 minutes, 5 minutes, and then 6 minutes—averaging to 5.33 minutes. These periods are different from those defined by the MBTA and are instead based on periods of uniform departure headways from Riverside.

Table 5.1: Time periods analyzed and scheduled headways on the Riverside branch.

| Time period | Time | Scheduled headway |
|---|---|---|
| Morning peak | 06:00-09:00 | 5.33 min |
| Midday | 10:00-15:00 | 7.5 min |
| Afternoon peak | 15:00-18:00 | 5.66 min |
| Evening | 20:00-24:00 | 10.5 min |

### 5.1.6 AVL data periods of experiment data

Analyzed experimental data during the pilot's ten days excludes periods of time with suspensions or severe delays on the Riverside branch, which occurred only once, during the midday of Thursday, March 2, when a tree branch fell on catenary wires.

In all analyses, data from only the second week (weekdays of Wednesday, March 1 through Tuesday, March 7) of the experiment are considered, for several reasons. The first week of the pilot occurred during the Massachusetts-wide February school vacation period—ridership periods differ from normal during this week, as some parents stay home with their children and/or take off-peak trips from suburbs and into Boston. Much of the first week also consisted of officials and assistants becoming comfortable with the devices and dispatching algorithm, resulting in more periods of non-adherence to recommendations. The first week also did not run with a functioning ring-off bell, decreasing the number of methods officials had to communicate dispatching times precisely and encourage operator compliance.

## 5.2 Compliance with dispatch instructions

The decision-support tool presented each scheduled departing trip from Riverside with the scheduled departure time and, more prominently, the algorithm's recommended time. This section describes the degree to which departures complied with recommendations and differed from the schedule.

### 5.2.1 Definition of compliant departure

Departing trips from the experimental period are labeled as either compliant or non-compliant to the recommended departure time from Riverside, based on a threshold of $\pm 45$ seconds. This is done to separate the effects of using the holding strategy from those of using normal dispatching procedures, as the policy was not followed consistently during the experimental

period. Since a train's running time and headway are often impacted by the vehicles ahead (leading to greater or fewer than normal boardings) and behind (leading to en route control actions), a trip is only considered compliant if the adjacent trains had also departed within the 45-second threshold. In other words, a compliant departure is one with even headways surrounding it per the dispatching algorithm. For passenger-based metrics, the definition of compliance was relaxed in that the following train does not need to have departed Riverside within the threshold, because the simulation used to generate the metrics assumes that passengers board the first available train serving their destination or the furthest station reachable from the origin (whichever is closer to the origin).

The 45-second threshold provides a large enough sample size to compare against the baseline for all times of day except the evening (which sees low eastbound ridership), and it can be reasonably expected that, under a full implementation of headway-based dispatching, trains will be able to depart within 45 seconds of the recommended time. A smaller threshold was not chosen due to the lack of trips matching all criteria and to inaccuracies in determining a train's true departure time from real-time AVL data. During the second week of the experiment, 15.6% of eastbound trips met all the criteria. When broken by time of day, the share of fully compliant trips varied between 12.8% (midday) and 20.3% (morning peak).

### 5.2.2 Deviation of departure times from recommendations and schedule

A total of 782 departures were observed at Riverside during the second week of the experiment. Figure 5-2 summarizes the differences between actual departure times and recommended departure times for all times of day. Negative lateness values signify early departures. Only 19% of departures left within 15 seconds of the recommendation, and 36% departed within 30 seconds. Table 5.2 shows that just fewer than half of all trips left within the 45-second standard of punctuality established above, suggesting that significant difficulties were experienced ensuring that officials and/or operators adhere to the decision-support tool's recommendations. Comparing the punctuality to the departure recommendations to schedule adherence, both in Table 5.2 and in Figure 5-3, it becomes clear that trains were more likely to leave with the scheduled time than to the recommendation.

Figure 5-4 shows adherence to the experimental recommendations for compliant trips by time of day. For all periods except the evening, trips tend to depart before the recommended time—this is also reflected in the negative means and medians of adherence for the respective
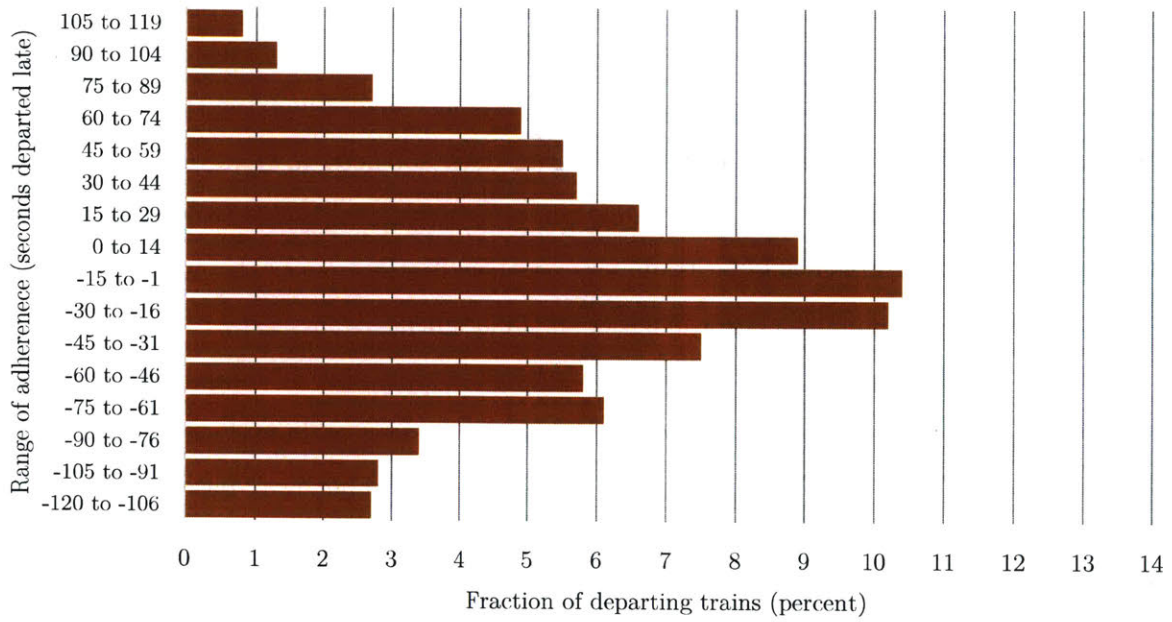
Figure 5-2: Adherence of all departures from Riverside to recommended times.
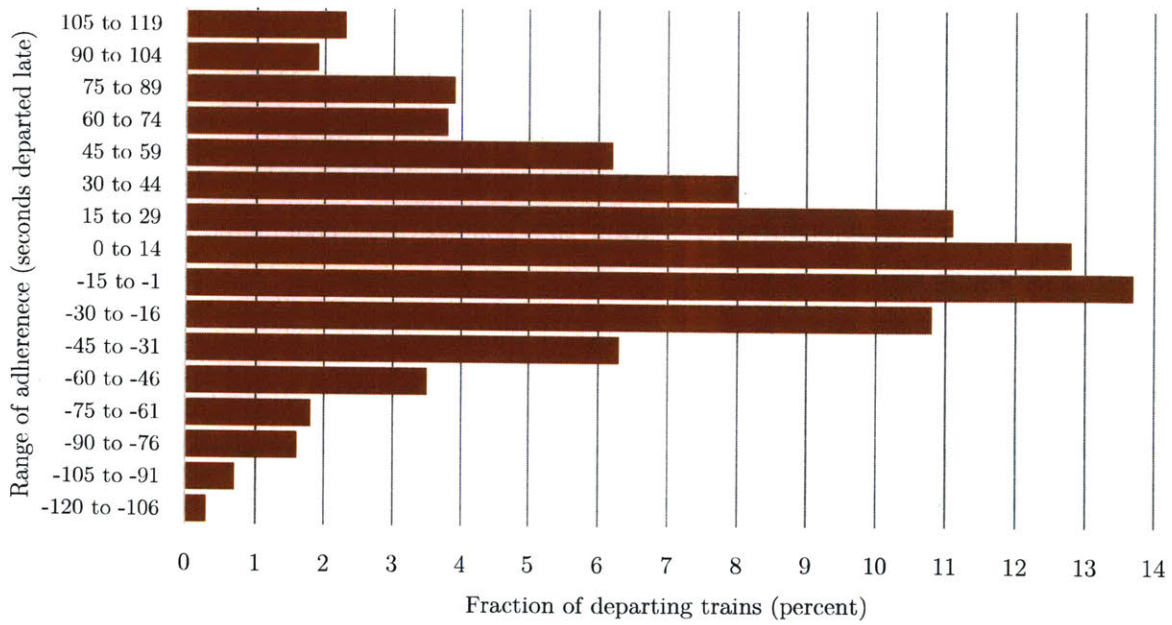


Figure 5-3: Adherence of all departures from Riverside to scheduled times.

Table 5.2: Adherence to recommended, scheduled departure times by compliance threshold.

| Threshold | Trips compliant to recommendation | Trips compliant to schedule |
|---|---|---|
| Within 30 seconds | 36.1% | 48.4% |
| Within 45 seconds | 49.3% | 62.7% |
| Within 60 seconds | 60.6% | 72.4% |
| Within 90 seconds | 77.6% | 83.4% |

times of day.

Based on observations written by assistants, officials occasionally ignored the recommended times if they strayed too far late past the scheduled time. In these cases, an early departure compared to the algorithm may also be an on-time departure compared to the paper schedule. This can be seen in Figure 5-5: adherence of departures against scheduled times, among compliant trips, is centered around the on-time mark and is not skewed early or late. The link between the schedule, recommended departures, and actual adherence to recommendations is explored further in Section 5.3. There is also evidence that some operators departed ahead of the recommendation in order to avoid hearing the ring-off bell; several even negotiated with the official at what time the bell would be rung if he or she had not yet departed.

### 5.2.3 Deviations of recommended departure times from schedule.

Figure 5-6 shows, for compliant trips by time of day, the difference between the time recommended for departure and the scheduled departure time. To a greater degree than in the day-spanning Figure 5-3, compliant trips show a tendency to have a recommended departure time later than scheduled. In order to determine more precisely how the schedule impacts adherence to the recommended departure times, we segregate compliance by the difference in time between the recommended and the scheduled departure of each trip. Table 5.3 summarizes the share of early, on-time, and late departures from Riverside by the difference between scheduled and recommended departure times. When recommendations are made for 90 seconds or more after the scheduled time, trains leave early more often than they do on-time. For recommendations 30 seconds or more ahead of schedule, trains will leave late as often or more often than they depart on-time.

These results suggest that the problem of compliance was widespread. Trips recommended to leave ahead of schedule tend to leave late, and those due to leave behind schedule

(a) Morning peak.



(b) Midday.



(c) Afternoon peak.



(d) Evening.
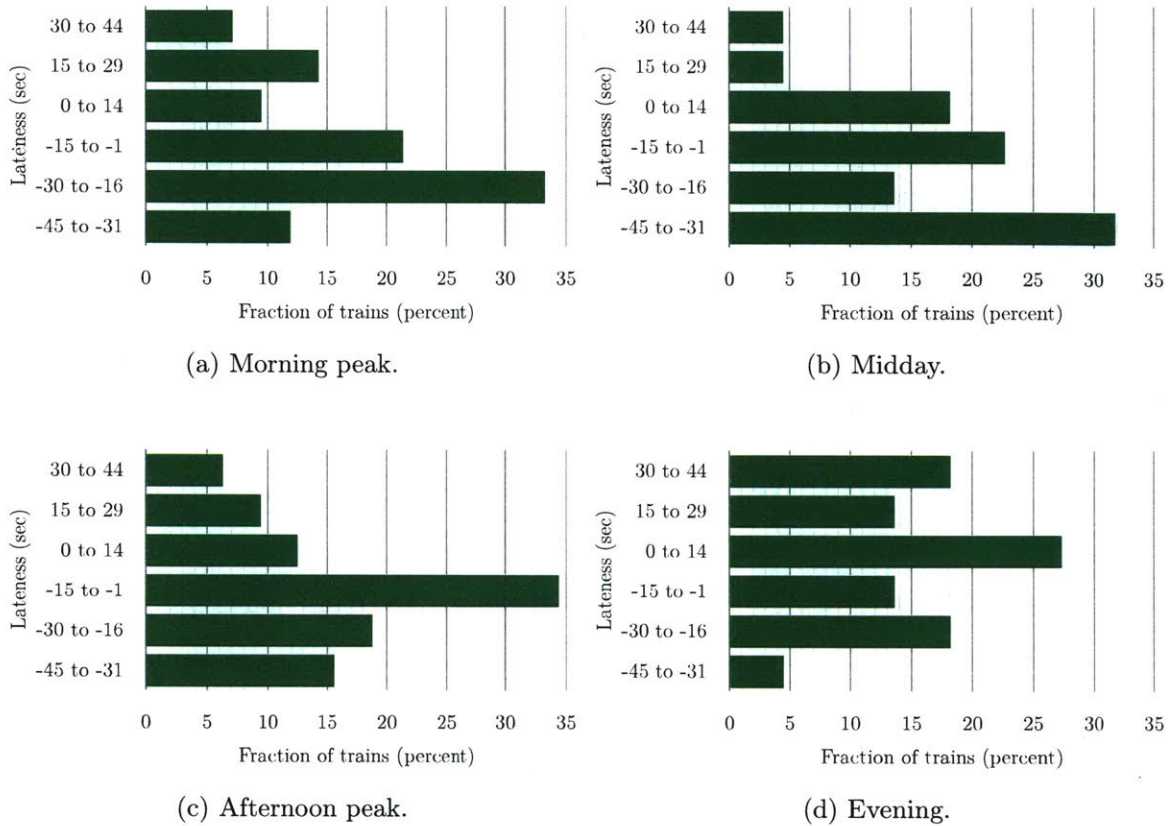
Figure 5-4: Adherence of compliant trips to recommended departure times from Riverside.

Table 5.3: Shares of trips departing early, on time, and late (to recommended).

| | Difference between recommended and scheduled departure times | Left 45 sec+ early | Left on time (±45 sec) | Left 45 sec+ late |
|---|---|---|---|---|
| Recommended behind schedule | 2.5–3.0 min | 68.2% | 18.2% | 13.6% |
| | 2.0–2.5 min | 77.8% | 22.2% | 0.0% |
| | 1.5–2.0 min | 69.2% | 25.6% | 5.1% |
| | 1.0–1.5 min | 46.4% | 44.9% | 8.7% |
| | 45–60 sec | 42.1% | 49.1% | 8.8% |
| | 30–45 sec | 29.6% | 54.9% | 15.5% |
| | 15–30 sec | 19.7% | 62.1% | 18.2% |
| | 0–15 sec | 6.0% | 82.1% | 11.9% |
| Recommended ahead of schedule | 0–15 sec | 17.5% | 60.5% | 21.9% |
| | 15–30 sec | 7.0% | 58.1% | 34.9% |
| | 30–45 sec | 3.2% | 48.4% | 48.4% |
| | 45–60 sec | 4.0% | 48.0% | 48.0% |
| | 1.0–1.5 min | 8.7% | 30.4% | 60.9% |
| | 1.5–2.0 min | 0.0% | 20.0% | 80.0% |
| | 2.0–2.5 min | 0.0% | 25.0% | 75.0% |
| | 2.5–3.0 min | 0.0% | 0.0% | 100.0% |
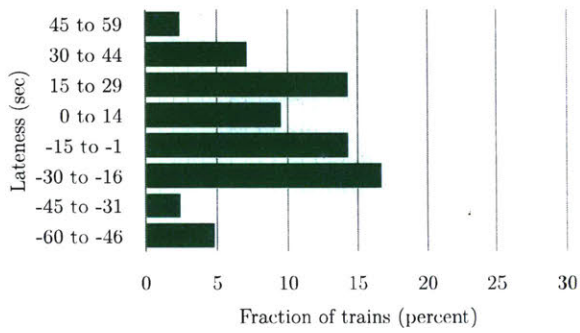
(a) Morning peak
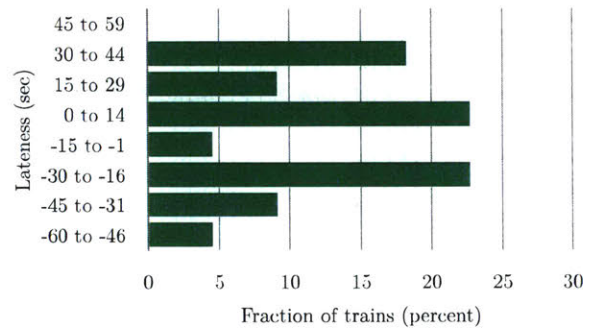
(b) Midday
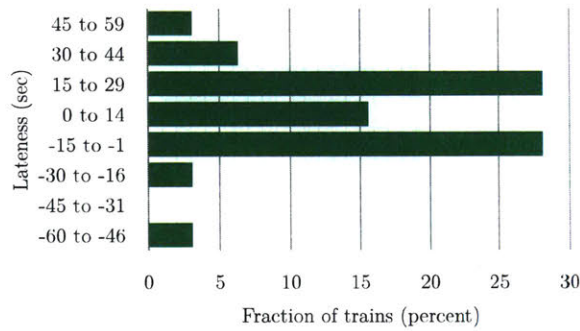
(c) Afternoon peak
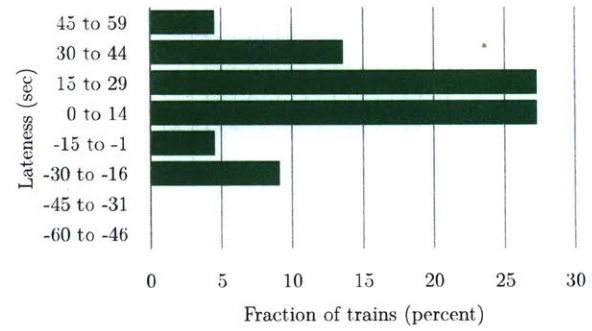
(d) Evening

Figure 5-5: Adherence of compliant trips to scheduled departure times from Riverside.

(a) Morning peak

(b) Midday

(c) Afternoon peak

(d) Evening
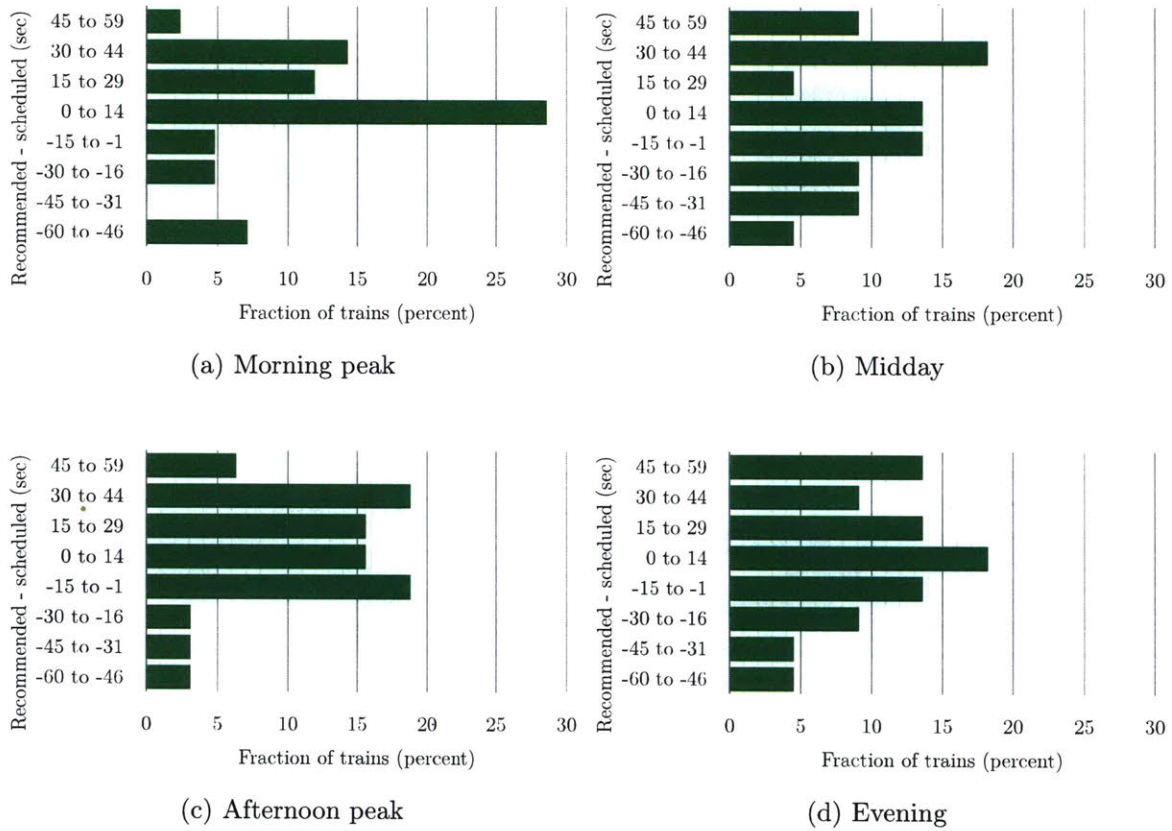
Figure 5-6: Difference between recommended and scheduled departure times for compliant trips.

tend to leave early. The problem is slightly worse for recommendations ahead of schedule than those behind schedule. Observations by assistants show that there were difficulties in having operators walk to their vehicles (which, at Riverside, are usually left on the platform with doors open during the layover period) in advance of the departure times. Even when the recommended times were within one-half minute of the schedule, however, only 65.5% of trips left on time. In most of these cases, the official simply told operators to leave at the "regular time". This relatively low adherence rate may be reflective of typical on-time performance at the terminal. Issues related to compliance are discussed further in Section 5.3.

Given the above observations on recommendation adherence, the analyses on service and passenger-oriented impacts (Sections 5.4 and 5.5) are based on the small proportion of compliant trips, but are representative of improvements which can be experienced if service was operated to the times recommended by the decision-support tool on a more consistent basis.

## 5.3 Factors contributing to recommendation non-compliance

This section explores the contributing factors to operator departure non-compliance during the experiment at Riverside. These factors include poor operator and inspector behavior, the design and reliability of the software application, and specific features of the Riverside terminal. Section 5.6 will then provide several suggestions for addressing these factors for future experiments and implementations.

### 5.3.1 Operator and inspector behavior and culture

The chief reason for the large proportion of non-compliant trips in Section 5.2 is the behavior of operators, who drive the trains, and inspectors, who are supposed to inform operators of their departure times. Some of the issues suggest a culture among personnel which does not prioritize on-time performance, while others relate to inconsistencies in practice by terminal inspectors, or inspectors' workloads.

The ring-off departure bell at Riverside was conceptualized as a way to let operators know when to close the doors of the train and depart. In practice, however, many operators at the terminal do not board their vehicles until the departure time or very shortly beforehand. This

135
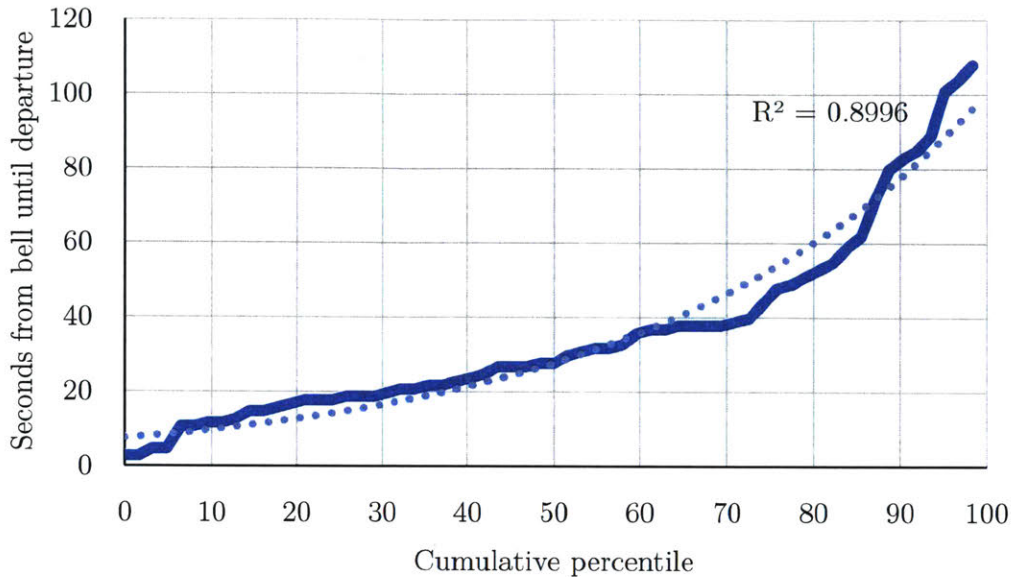
Figure 5-7: Distribution of reaction time from departure bell activation until train departure.

often made the bell impractical. Often during the pilot, assistants would observe operators in or near the inspector's booth, chatting with fellow colleagues, with only a minute to spare until the instructed departure time. More often than not, the bell would be rung while the operator was not yet onboard the train, so the sound acted as a admonishment that he or she should have already departed. Some operators would begin running if they heard the bell ringing for them while more than one was observed to have slowed his or her walking pace, perhaps out of contempt. Figure 5-7 shows that the duration of time between when the departure bell began ringing and when the vehicle began to move was inconsistent, suggesting that the bell was often ineffective at ensuring compliant trips during the pilot.

The bell itself was generally not well-received by operators. Many complained that it was excessive in duration and volume, valid complaints which are addressed in Section 5.3.3, while others likened the bell to being herded around as if they were cattle and expressed fears of being tracked while on the job. The bell was inoperable until the second week of the pilot, meaning that some operators may have considered the bell to be separate from the rest of the tablet application and held different opinions of the two.[1] Separately, this experiment was never run without the presence of an assistant, who came from MIT or the OMPI office at the MBTA headquarters. Therefore, there is the possibility that some

---

[1]All of the analyses of the experiment earlier in this chapter exclusively use data from the second week, as explained in Section 5.1.5.

operators were more diligent in their behavior during the pilot than on typical weeks.

Instructions to inspectors about the usage of the decision-support tool, including the departure bell, were unclear and not consistently followed. The written memorandum to inspectors and operators said that "a bell will sound on the platform; this is to alert Operators to safely close their doors and depart the station on signal", but the memos were not posted until two days before the end of the pilot (Michaud and Thibodeaux, 2017). A memorandum, within the MBTA, does not have the force of a special order, which has an implied disciplinary threat. The use of a memo was deemed appropriate because the pilot was only in effect for a short time and because special orders must be approved and signed by additional supervisory and management personnel. Some inspectors used the bell for every departure, while others did so only if the operator had not yet left after the instructed time, and several stopped using the bell entirely after a short time because it was deemed too bothersome. From at least one inspector there was an indication that, due to the complaints of operators, continuing to use the bell may cause operators to stop following his instructions, worsening service performance.

Ensuring on-time departures begins before operators reach the platform. Operators are required to report to the terminal inspector on duty ten minutes before the start of all shifts and two minutes after breaks (including the start of the second half of a split shift). Some operators were observed reporting much closer to their first trips than the ten-minute buffer required, often without reprimand from the inspector. A two-minute buffer before the scheduled departure time after breaks was insufficient to ensure that operators depart early when that was the recommendation given by the even-headway policy. At present, inspectors are permitted to send operators out on trips earlier than two (or ten) minutes before scheduled if they are present, but operators are allowed to refuse the change. More than one inspector expressed a sentiment of frustrated defeat on this subject: coordinating the pulling out and in of trains and ensuring that trains departed on time left very little time and energy to force operators to sign-in early.

Indeed, the terminal inspectors were observed as having a large workload. While all inspectors manage to run service and record necessary information on their train sheets at Riverside, it was often clear that there is room for increased efficiency. Terminal inspectors play multiple roles: managing and approving delay time requests, fielding inquiries from OCC dispatchers about operators or canceled trips, instructing operators to depart, deter-

mining whether a trip will go out on time, and if not, formulating a plan (by assigning cover, having the trip depart late, or cancelling the trip). At least several tasks (such as delay time processing, reporting train cancellations, and filling in substitute operators for the next day's service) could be streamlined or eliminated with a decision-support system, allowing inspectors to focus on dispatching the next several trips from the terminal.

## 5.3.2   Software application design and reliability

The software application had undergone many hours of field testing on multiple days before the start of the pilot, but the software was still not without issues which either disrupted the pilot or made the interface difficult to follow on occasion.

MBTA-realtime, the public-facing API from which vehicle locations and predictions are obtained, would occasionally go offline for several minutes during the height of rush hour, right as it is most critical for trains to be properly dispatched. During some of these instances, trains were dispatched by the scheduled times, and departures were not recorded electronically.

Often, during the pilot and in normal operations, inspectors would orally give instructions on departure times to operators several minutes in advance. If, between the current and recommended departure times, changes were made using the user interface, such as pulling in the next arriving train, the recommended time would also change. Users of the application found it frustrating when they told an operator one time before having it change on the screen soon thereafter. The reliability and other technical issues which remained may have shown some inspectors that the software was not yet perfected, reducing their trust in the devices.

Some inspectors found maintaining the tablet application to be a difficult task during the busier times of their shifts, especially since they were still required to fill out the paper train sheets. On occasion, this meant that the decision-support system was not kept updated, particularly when the inspector had to leave the booth to investigate a mechanical issue or find an operator. In these cases, the recommended time, using the information known by the algorithm, may have differed from the optimal departure time, had more up-to-date, accurate information been given. The inspector would, at times, ignore the recommended time, or, as a concession, give the operator a departure time between that scheduled and recommended by the tablet. Operations officials assured researchers that inspectors would

138

be easily able to manage both paper sheets and tablets simutaneously, though observations during the experiment suggested this was often not the case. While some exception-based ideas were incorporated in the development of the tool (by highlighting trips with long recommended headways in red or highlighting unknown crews in yellow, as shown in Figure 4-5), the displayed information can be pared down and tasks simplified. The dissemination of information could also be better-managed with the use of additional electronic displays, as described in Section 5.6.1.

### 5.3.3  Riverside terminal infrastructure

The physical layout of the Riverside terminal should bear some blame for the observed non-compliance of departures. The distance between the platform and the operators' lobby is addressed here, as is the lack of an effective, internal communications system.

Before shifts and during breaks, train drivers are usually found in the operators' lobby, the break and locker room at Riverside. The operators' lobby is located 100 meters (325 feet) from the inspector's booth, a non-trivial walk. The path between the booth and the lobby can be seen in Figure 5-8. If an operator is late reporting for a shift, he or she may be in the lobby, but it can be difficult for the inspector to communicate with operators there—an installed telephone often goes unanswered. On a handful of occasions, a frustrated inspector left his or her post in the booth and walked to the lobby to find the tardy operator.

The station is not designed for easy communication between the inspector and the operators once they are onboard a train. The easternmost end of the train, where the departing operator sits, is about 53 meters (175 feet) from the inspector's booth. During testing for the experiment, a working microphone, connected to the public address system on the platforms, was discovered, but the inspector on duty reported not having seen or used it for many months, if not years. The options available to the inspector are to walk from one end of the platform to the other (consuming valuable time), using a flashlight as a signal (which must be pre-coordinated by a different means), or using the open radio (which may be occupied and can be heard by all inspectors on the Green Line).[2] Additionally, the ring-off bell, when used, would ring loudly throughout the entire platform area. The station is split into two public address zones by platform, but both were needed to be used for the bell

---

[2]A special radio channel does exist at Riverside, but this is generally used for communicating with the yard crews. Train operators normally would not be tuned to the appropriate frequency while on the Riverside platform.

Figure 5-8: Map showing the locations of several facilities at the Riverside terminal.

Labels on the map:
- Operators' lounge
- Boarding and alighting berths (both tracks)
- Fare boundary
- Inspector's booth
- Publicly-accessible are within fare control
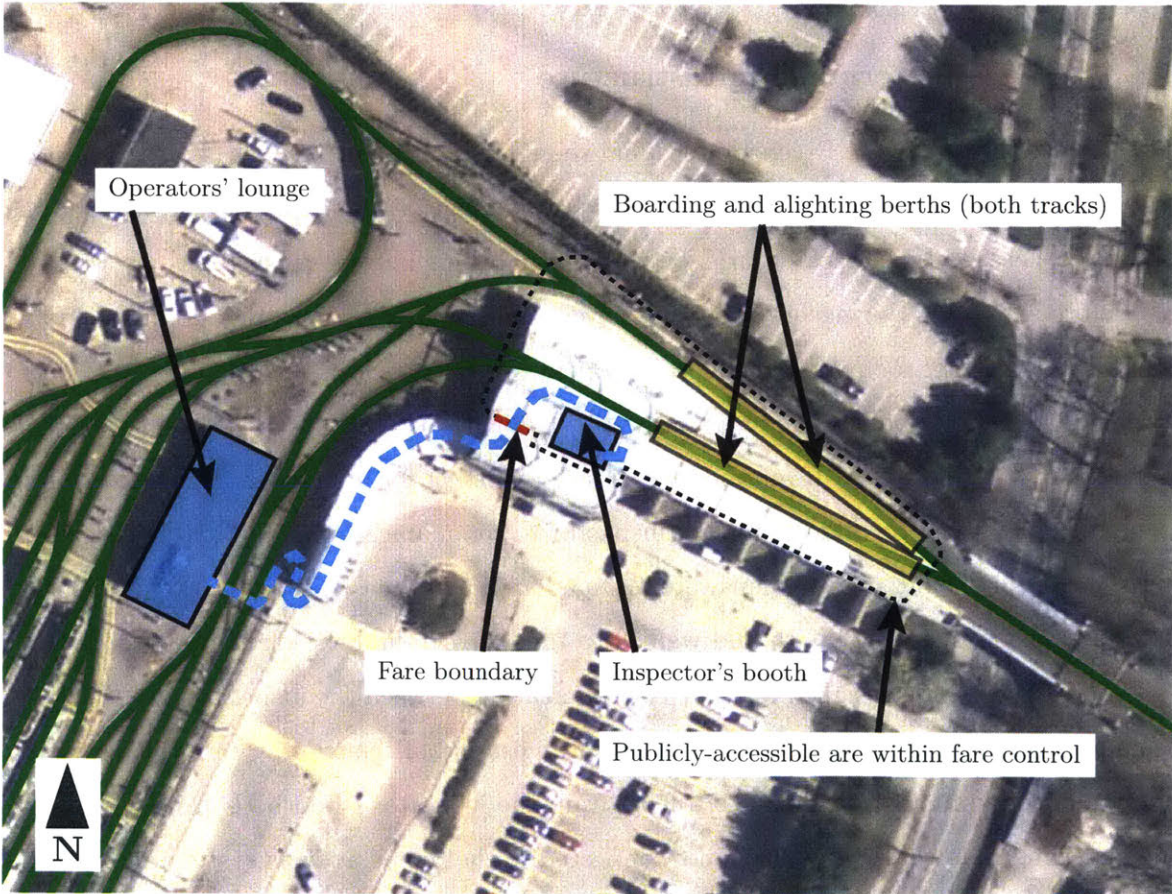
Original sources: Office of Geographic Information (MassGIS) (2014a,b). Modified by author.

to be audible from the operator's cab due to the location of properly working speakers.[3] This meant that people far away from the front of the train, including the inspector and operators waiting around the booth, experienced a quite loud bell ring, perhaps reducing the inspector's willingness to use it.

## 5.4 Experiment impacts on service

This sections considers the impacts of the terminal control experiment on Green Line service. The metrics used to compare changes in service are running and cycle times, headway variability, and the additional number of trains that would be required to obtain the service quality produced by a perfectly run schedule.

### 5.4.1 Running times

Significant improvements to running times were not anticipated based on the results of the simulation experiment in Chapter 3. Table 5.4 lists summary statistics for running times by time of day in the eastbound direction. Running time is calculated as the time between departure at Riverside and arrival at Government Center. The improvement in mean eastbound running time from the baseline to the set of compliant trips is small but statistically significant to the 1% level during the morning peak and to the 5% level during the afternoon peak. The decrease in running time when following the headway-based dispatching policy comes from fewer longer dwells following long gaps and fewer en route holds (to adjust the headway mid-trip).

In the westbound direction (returning to Riverside), similar improvements are observed. While it may appear counterintuitive for running times to have decreased in the direction that did not have additional terminal control, bunched trains are still often held in the westbound direction. This happens most often at Park Street and Kenmore Square stations, where multiple platforms allow overtaking.

In practice, however, vehicle schedules are set based on the longest of observed trips over a survey period, typically between the 85th and 95th percentile of cycle times (Muller and Furth, 2000). This means that upper percentile observed running times are more important

---

[3]A public address zone is a set of network-connected speakers which are linked together and can be controlled as a unit. The zones are defined using software. The MBTA uses public address zones to send different audio and visual messages to different platforms within the same station.

Table 5.4: Statistics for eastbound running times.

| Time period | Service scenario | Median (min) | Mean (min) | Std dev (min) | 90th percentile (min) |
|---|---|---|---|---|---|
| Morning peak | Baseline | 46.7 | 47.2 | 4.7 | 52.7 |
| | Non compliant trips | 45.1 | 45.4 | 3.7 | 50.3 |
| | Compliant trips | 45.3 | 45.1 | 3.7 | 49.6 |
| Midday | Baseline | 44.0 | 44.2 | 3.9 | 48.4 |
| | Non compliant trips | 43.4 | 43.9 | 4.8 | 47.4 |
| | Compliant trips | 42.7 | 43.0 | 3.4 | 46.9 |
| Afternoon peak | Baseline | 46.3 | 46.9 | 3.6 | 52.0 |
| | Non compliant trips | 44.5 | 45.1 | 3.0 | 49.3 |
| | Compliant trips | 45.2 | 45.4 | 2.3 | 48.1 |

Table 5.5: Statistics for cycle times from Riverside EB to Riverside WB.

| Time period | Service scenario | Median (min) | Mean (min) | Std dev (min) | 90th perc (min) | Sched (min) |
|---|---|---|---|---|---|---|
| Morning peak | Baseline | 98.1 | 97.8 | 9.7 | 107.4 | 100.0 |
| | Non compliant trips | 96.9 | 97.9 | 7.1 | 107.0 | |
| | Compliant trips | 94.2 | 94.4 | 5.4 | 101.0 | |
| Midday | Baseline | 91.3 | 90.8 | 8.7 | 99.6 | 100.5 |
| | Non compliant trips | 90.7 | 91.0 | 5.9 | 98.4 | |
| | Compliant trips | 89.2 | 89.7 | 4.0 | 94.5 | |
| Afternoon peak | Baseline | 98.6 | 98.7 | 11.6 | 111.4 | 107.0 |
| | Non compliant trips | 94.9 | 95.2 | 7.3 | 103.3 | |
| | Compliant trips | 93.3 | 90.9 | 12.2 | 101.2 | |

Table 5.6: Reductions in 90th percentile cycle time from baseline to compliant trips.

| Time period | Reduction in cycle time from baseline to compliant trips | Reduction in cycle time from non-compliant to compliant trips | Reduction in cycle time from current schedule to compliant trips |
|---|---|---|---|
| Morning peak | 5.9% (6.4 min) | 5.6% (6.0 min) | — |
| Midday | 5.1% (5.0 min) | 3.9% (3.8 min) | 5.2% (5.2 min) |
| Afternoon peak | 9.2% (10.2 min) | 2.0% (2.1 min) | 3.8% (4.1 min) |

than the mean for potential schedule changes. Table 5.5 lists the statistics for observed cycle times on the branch excluding the layover time at Riverside. As with running times, the difference in means of cycle times between compliant trips and the baseline is significant and large during the peak periods. The differences between the 90th percentile cycle times observed for compliant experimental trips, non-compliant experimental trips, the baseline period, and the current schedule are summarized in Table 5.6:

Compared to the baseline, consistently using the times recommended by the headway-based algorithm results in at least 5% shorter cycle times through the afternoon peak period. By adjusting the scheduled cycle time likewise, the current number of trains could be scheduled to serve additional trips, or fewer trains could serve the current number of scheduled trips. In the example of the morning peak, a 6.4-minute cut from the scheduled cycle time would require one fewer train in service. The scheduled cycle time in the morning, however, is shorter than what was observed during the experiment: this means that trains are still arriving back to Riverside late. Nevertheless, the algorithm used in the pilot would still reduce the number of late trains, which in turn reduces the operating costs by reducing the amount of overtime worked.

### 5.4.2 Headway variability

The effect of the decision-support tool on headway throughout the route is examined. Figures 5-9, 5-10, and 5-11 show the coefficient of variation of headway along the Riverside branch for the morning peak, midday, and afternoon peak periods, respectively. The coefficient of variation of headway is the standard deviation of the headway divided by the mean headway at each station. In each figure, the left half shows the headway variability in the eastbound direction departing Riverside and the right half shows the westbound direction heading
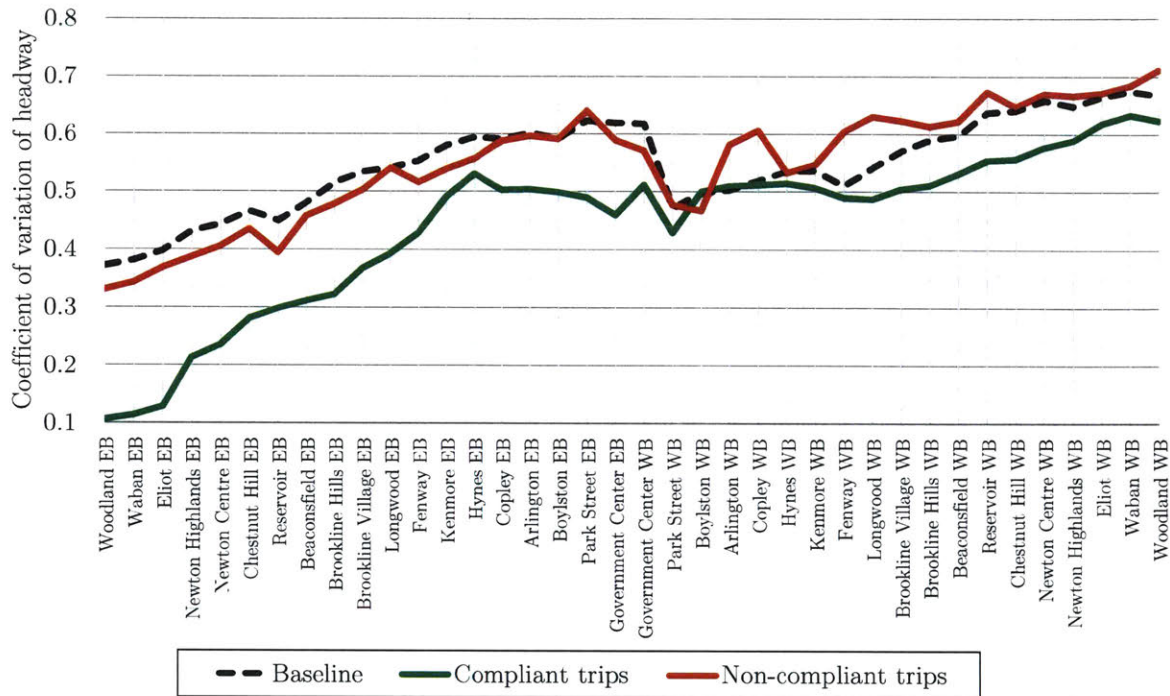
Figure 5-9: Coefficients of variation of headways during the morning peak (06:00-09:00).

toward Riverside.

Abrupt changes in the coefficient of variation between stops, such as at Reservoir eastbound or Park Street westbound, likely indicate a holding control point. At all times of day, eastbound headways generally increase in variability moving toward Boston. During the peak hours, the headway variability of compliant eastbound trips begins significantly lower than the baseline, and the difference decreases as trains approach and merge with other branches at Kenmore. At midday, the magnitude of improvement from the baseline is small leaving the terminal but increases before reaching the eastern terminus. This indicates either that terminal dispatching is already effective during this time period or that compliant trips and their surrounding trains are not as punctual at midday as during the peak periods (see Figure 5-2).

In all cases, westbound service from Government Center is more variable than eastbound service, largely due to layover space constraints within the Central Subway: only so much adjustment can be made before a held train blocks another, and such control is out of the scope of this experiment. Improvements from baseline coefficients of variation are smaller in magnitude westbound but tend to increase along the route.

Across both directions, the morning peak headway variability of compliant trips de-

144
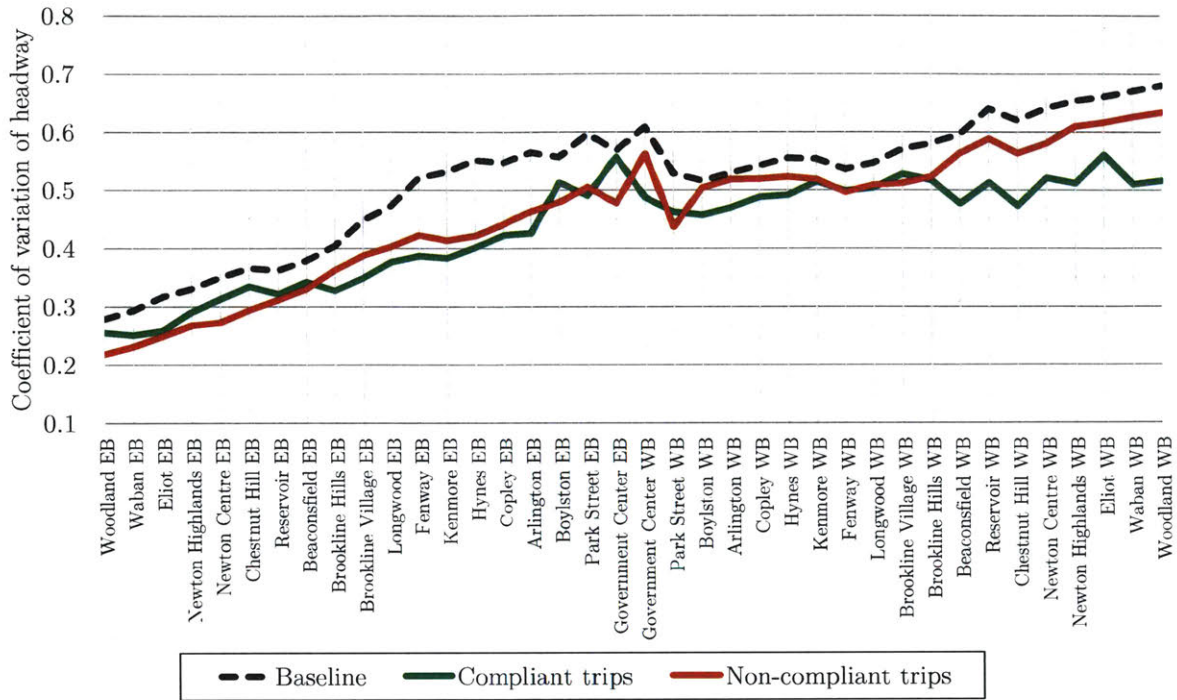
Figure 5-10: Coefficients of variation of headways during midday (10:00-15:00).
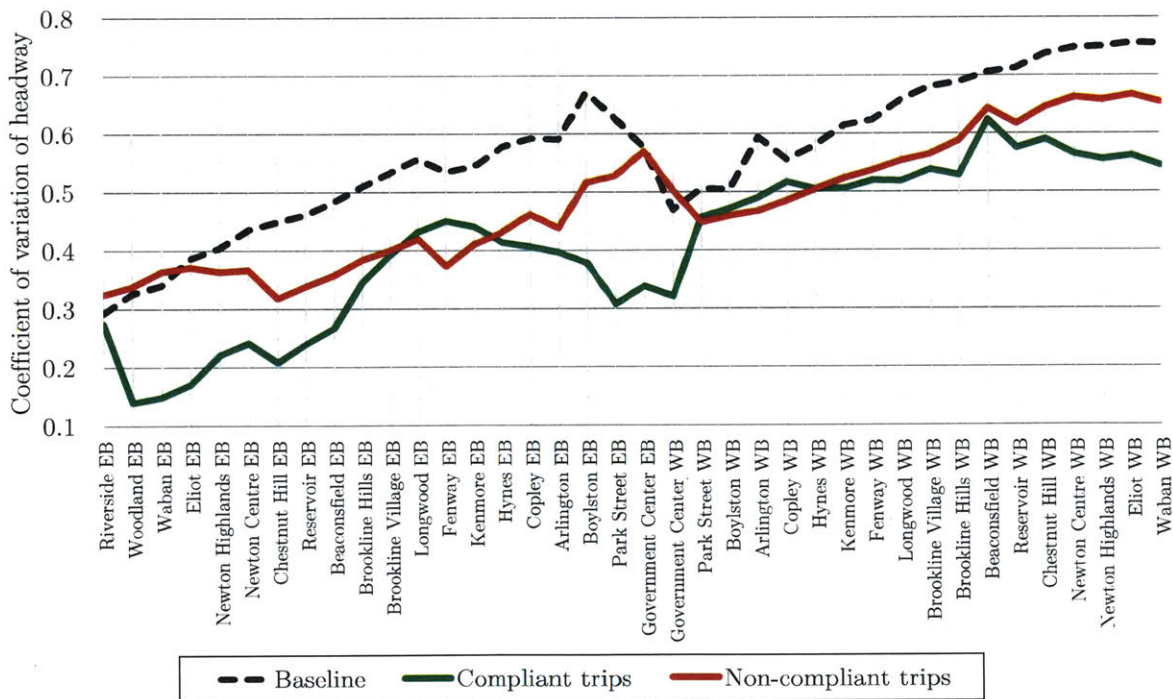


Figure 5-11: Coefficients of variation of headways during the afternoon peak (15:00-18:00).

creased by 41% compared to the baseline when the coefficients of variation are weighted by boardings.

### 5.4.3 Additional vehicles required

The additional vehicles, or capacity, required metric, formulated by Maltzan (2015), is an operations-oriented indicator of the amount of resources required to operate at the level of service scheduled and promised to passengers. It is useful for agencies when allocating vehicles and crews and estimating additional or reduced costs from implementing new service delivery strategies. The metric is based on the relationship of a transit route's scheduled cycle time at a given time of day, $c_t$ and its headway, $H_t$, to the number of vehicles required to run the service. This requires the assumption that the scheduled cycle time, scheduled headway, and thus scheduled number of vehicles in service are constant across a time period. In this analysis the assumption is true for scheduled headway and is true for cycle time and number of vehicles after an initial ramp-up phase in service for the morning and afternoon peak periods.

As given by Maltzan, the headway used for $H_t$ is the route's effective headway, $H_t^*$, a function of its unreliability:

$$H_t^* = H_{\text{scheduled, }t}(1 + \widehat{CV_t}^2) \tag{5.3}$$

where $\widehat{CV_t}$ is the coefficient of variation of the headway at time of day $t$ weighted across all stations on the line. The weighted coefficient of variation is the weighted mean $\widehat{\mu}$ divided by the weighted standard deviation $\widehat{\sigma}$ over all stations $i$, which are given in Equations 5.4 and 5.5:

$$\widehat{\mu} = \sum_I w_i \mu_i \tag{5.4}$$

$$\widehat{\sigma} = \sqrt{\sum_I w_i(\mu_i - \widehat{\mu})^2} \tag{5.5}$$

where $w_i$ is the passenger weight assigned to the station and $\mu_i$ is the average headway at stop $i$ during the studied time of day.

Passenger weights are obtained by running the passenger simulation model from Chap-

Table 5.7: Improvements in addition vehicles required to run scheduled level of service.

| Service scenario | Morning peak | Midday | Afternoon peak |
|---|---|---|---|
| Baseline | 22.7% (3.9 additional trains) | 2.3% (0.3 additional trains) | 24.2% (4.1 additional trains) |
| Non-compliant trips | 21.5% (3.7 additional trains) | — (0 additional trains) | 18.8% (3.3 additional trains) |
| Compliant trips | 11.7% (2.2 additional trains) | — (0 additional trains) | 12.8% (2.4 additional trains) |
| Improvement with experiment strategy | 11.0% (1.7 trains added) | 2.3% (0.3 trains added) | 11.4% (1.7 trains added) |

ter 3 with baseline period AVL data (simulating only demand and not service) and determining each station's share of boardings out of the total boardings at all stations in both directions of the line. The weights are calculated separately for each time period and are normalized to sum to 1. The number of additional trains required to run the scheduled service is given by $V_{required}$ in Equation 5.6:

$$V_{\text{required}} = V_{\text{scheduled, }t} - \frac{H_t^*}{c_{\text{scheduled, }t}} \tag{5.6}$$

where $V_{\text{scheduled, }t}$ is the scheduled number of vehicles to operate on the branch and $c_{\text{scheduled, }t}$ is the scheduled cycle time at time of day $t$.

Summarized in Table 5.7 is the amount of service, given baseline or experimental conditions, that would have to be added to the Riverside branch in order to operate at the scheduled level of service (at which the effective headway is equal to the scheduled headway), both as a percentage and the number of vehicles. For example, we consider the afternoon peak, which experienced the greatest improvement: the added vehicles required decreased from 24.2% to 12.8% with the proper compliance of the experimental dispatching strategy. This is an improvement of 11.4%, which implies that controlling headways at Riverside produced an effect equivalent to increasing the amount of service provided on the branch by 11.4%. Translated into vehicles, the strategy results in only 2.4 trains needing to be added to provide the scheduled service during afternoon peak instead of 4.1 trains, the equivalent of obtaining the benefit of 1.7 additional trains. The benefits are similar in the morning peak as the afternoon peak, with smaller benefits realized at midday.

These improvements mean that an agency can improve reliability with fewer capital and operating expenses spent on adding resources, or that it can run current levels of service

147
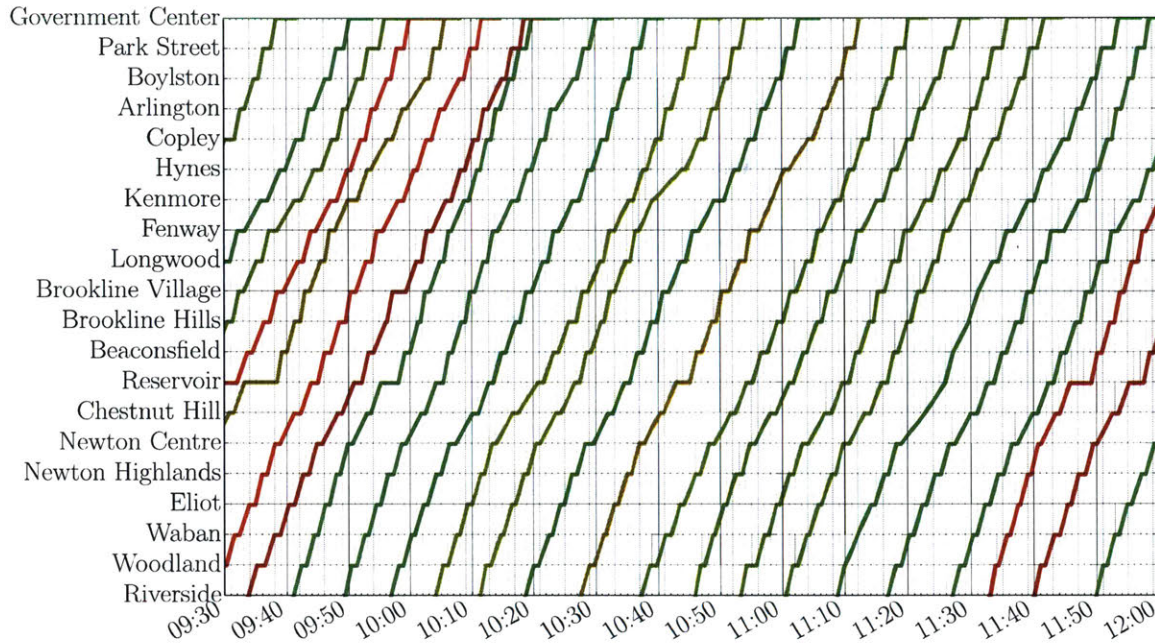
Figure 5-12: Time-space diagram of eastbound service on Tuesday, March 7, 09:30–12:00.

by utilizing fewer resources, lowering operating costs. These improvements can be increased across the Green Line if the strategy was applied at all terminals and in both directions, and to a lesser extent if fine-tuned control actions consistent and coordinated with terminal dispatching were applied sparingly at select en route control points, as shown in the results from Chapter 3.

### 5.4.4 Examples of effects on holding

Shown below are several examples of time-space diagrams showing eastbound service on the Riverside branch when the recommended departures were and were not followed. In the three following figures (5-12, 5-13, 5-14), each line represents the trajectory of a train. Lines in bright red signify trains which left more than two minutes before or after the recommendation; bright green lines signify trips departing at the recommended time; colors in the gradient between green and red signify a varying degree of adherence.

Figure 5-12 shows a midday stretch of trains which departed Riverside relatively close to the recommended time displayed. The departures running along the bottom of the figure are spaced relatively equally apart, and trains generally remained evenly spaced downstream until reaching Government Center.
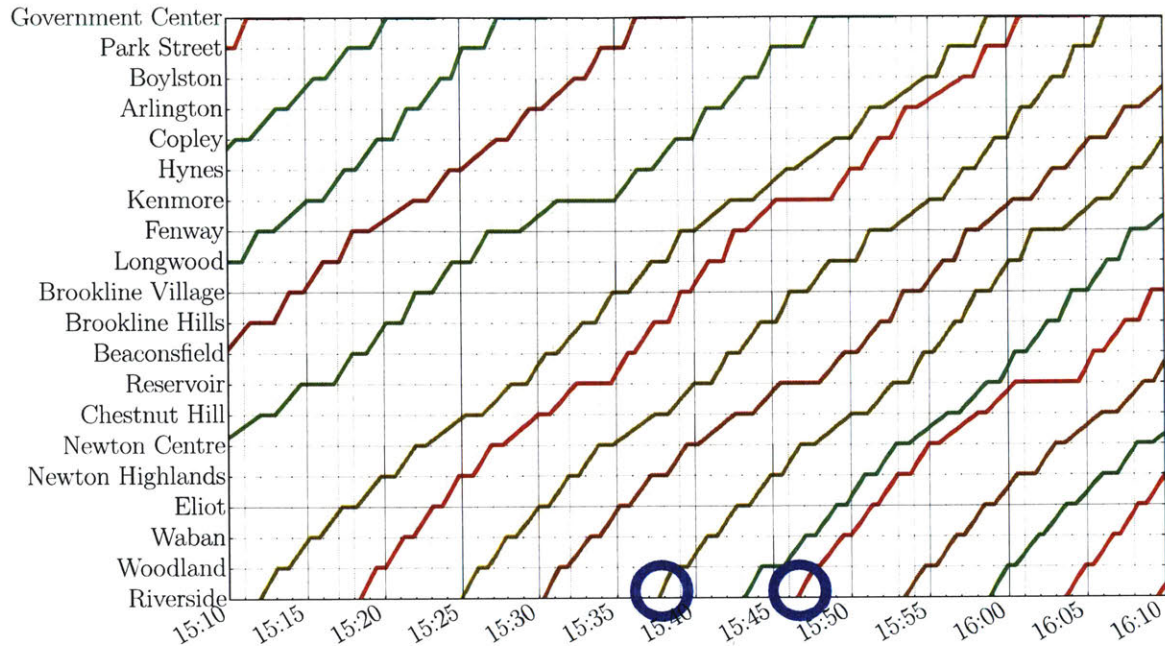
Figure 5-13: Time-space diagram of eastbound service on Tuesday, March 7, 15:10–16:10.

In contrast, Figure 5-13 shows uneven departures at the start of the afternoon peak on March 7. Both of the circled departures occurred at their regularly-scheduled time but did not follow the recommendation. Both experienced long dwells at Reservoir station—likely the respective operators were told by an official at Reservoir or the OCC via radio to hold several minutes to adjust headway. The train departing at 15:46 left less than 4 minutes after the previous, likely experiencing fewer passengers and thus shorter dwell and running times. The result is that two trains came within one minute of each other by Newton Centre, and the second was held for about four minutes at Reservoir, inconveniencing the 27 passengers estimated to be onboard. There is also the consideration of passengers waiting downstream of Reservoir who had no knowledge of the hold from the real-time information available to them (via station signage or mobile devices). If the train had departed at the recommended time, such an en route hold at Reservoir likely would not have been needed.

In the final example, Figure 5-14, the focus is placed on the train departing just before 09:20. The previous train left at its originally-scheduled time and was then held for three minutes at Reservoir from 09:29. The circled train itself was scheduled to depart at 09:20:00 and was recommended to depart very close to schedule at 09:19:46. The train departed more than a minute early at 09:18:45 and experienced a faster running time, catching up to the previous train by Reservoir. At Reservoir, the early train was held for five minutes, and a

149
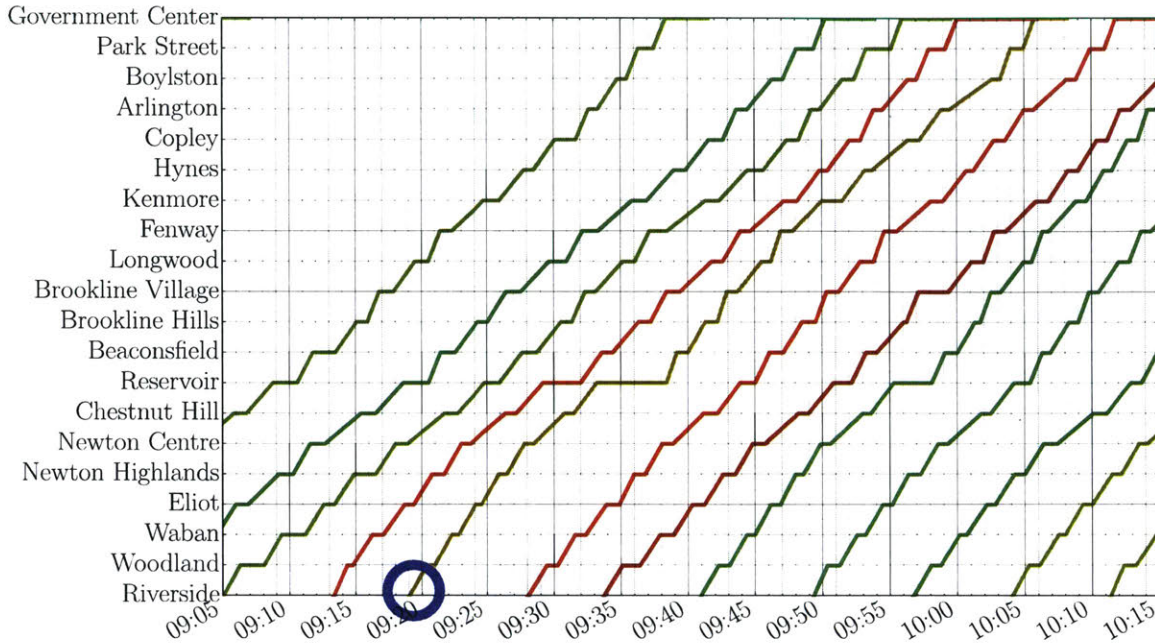
Figure 5-14: Time-space diagram of eastbound service on Tuesday, March 7, 09:05–10:15.

shorter hold may have additionally occurred at Kenmore station. 53 passengers rode through Reservoir station and 90 through Kenmore, resulting in a total of 6.7 passenger-hours spent onboard the particular vehicle during the holds.

The preceding examples show how the control algorithm can impact service at the level of the individual trip—and how a lack of adherence to dispatching instructions can quickly affect commuters' journeys.

## 5.5    Experiment impacts on passenger experience

This section looks at the impacts of the terminal dispatching experiment to passengers. In particular, we look at changes in passenger wait and journey times as well as in a wait reliability metric. Values for metrics are obtained by running a simulation model similar to that used in Chapter 3, adapted to assign passengers to vehicles defined by archived AVL data rather than from randomly-drawn running times. For each scenario, at least 50 replications were run.

Table 5.8: Passenger wait time statistics during the morning peak.

| Service scenario | Eastbound | | Westbound | |
|---|---|---|---|---|
| | Mean (min) | Std dev (min) | Mean (min) | Std dev (min) |
| Baseline | 3.26 | 3.07 | 4.34 | 3.35 |
| Non-compliant trips | 3.03 | 2.60 | 4.32 | 3.46 |
| Compliant trips | 2.75 | 2.07 | 4.20 | 3.01 |

Table 5.9: Passenger wait time statistics during the afternoon peak.

| Service scenario | Eastbound | | Westbound | |
|---|---|---|---|---|
| | Mean (min) | Std dev (min) | Mean (min) | Std dev (min) |
| Baseline | 3.65 | 2.93 | 4.36 | 3.73 |
| Non-compliant trips | 3.44 | 2.68 | 3.98 | 3.40 |
| Compliant trips | 3.22 | 2.35 | 3.77 | 2.83 |

### 5.5.1 Passenger wait times

A passenger's wait time is defined as the time between his or her arrival at the origin station and that person's boarding time. If the journey requires a transfer between trains within the Green Line, the transfer time is added to the wait time.

Tables 5.8 and 5.9 show the means and standard deviations of wait times for the morning and afternoon peak periods, respectively. Analyzed eastbound passengers include all those who began trips along the branch portion and went to a destination within the Central Subway trunk; westbound passengers include all who traveled from the trunk to a surface-level stop along the branch.[4] Assuming that passengers arrive at their origins randomly through a stochastic process, the mean wait time for the ideal transit route running with regular headways should equal one-half the scheduled headway.

Passengers in each of the four time-direction pairs shown experienced significant decreases in the mean headway as well as headway variability. The greatest improvements from baseline wait times are realized in the two patterns with the highest demand: morning peak eastbound and afternoon peak westbound. The average passenger, in those patterns, waits 30 seconds less for a train following the headway-based strategy. More importantly, the change in standard deviation means that passengers experience a smaller range of waiting.

We also examine the ranges of experiences of eastbound passengers' waits as presented

---

[4]Passengers traveling only within the branch are not analyzed, as their journey times tend to be much shorter compared to those riding between the trunk and branch, and they represent a small share of all travelers, especially during the peak periods.

Table 5.10: Riverside-branch wait time reliability by direction and time of day.

| Direction | Service scenario | Morning peak (%) | Midday (%) | Afternoon peak (%) |
|-----------|------------------|------------------|------------|--------------------|
| Eastbound | Baseline | 81.2 | 85.0 | 80.2 |
| | Non compliant trips | 84.9 | 84.8 | 82.8 |
| | Compliant trips | 89.2 | 83.8 | 86.9 |
| Westbound | Baseline | 77.6 | 75.0 | 74.3 |
| | Non compliant trips | 78.6 | 75.5 | 78.7 |
| | Compliant trips | 80.4 | 75.4 | 79.5 |

in box plots in Figures 5-15 and 5-16. The ends of each plot's box represent the first and third quartiles, and the range between the end of the box and the closest whisker is 1.5 times the interquartile range. In the eastbound direction, the third quartile (75th percentile) is reduced by more than half a minute during the peak periods. The 90th percentile wait time, representing a segment of customers more likely than others to complain publicly about their unreliable journeys, decreases by 90 seconds (21%) in the morning peak and by 75 seconds (15%) in the afternoon peak. The midday period shows no decrease in the interquartile range for complaint trips but rather a marginal shift toward longer waits. As mentioned in Section 5.4, service already runs with greater regularity at midday than during the peak periods. This means that recommended departure times are not likely to significantly differ from the schedule, leading to a lesser sense of need to adhere to the tablet recommendation.

Figure 5-16 gives the wait time distribution for the afternoon peak in the westbound direction. The improvement in customer experience is similar to that in the morning peak eastbound, with a 75-second (or 14%) decrease in wait time at the 90th percentile.

## 5.5.2  Passenger wait time reliability

Wait time reliability, as defined in Chapter 3 and by Tribone et al. (2016), is a measure of customer experience. It provides the proportion of passengers with wait times deemed acceptable by the transit agency, with a duration shorter than the scheduled headway. The MBTA uses this metric for measuring performance on its bus and rail services. Table 5.10 shows the wait time reliabilities of passengers riding between the trunk and branch portions of the Riverside branch by time of day.

The greatest improvements are realized in the peak directions of travel. In the morning, wait reliability in the eastbound direction increases from 81.2% to 89.2% for those passengers
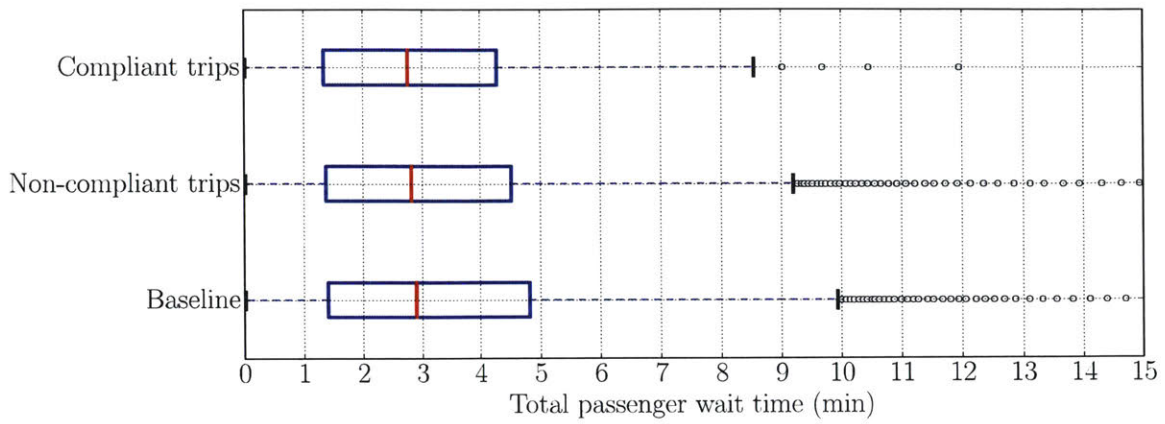
(a) Morning peak.
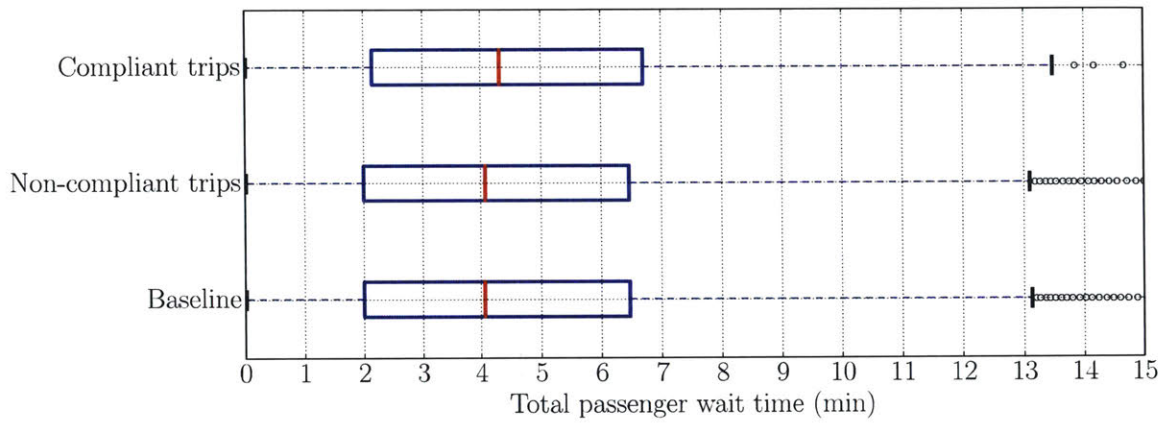
(b) Midday.
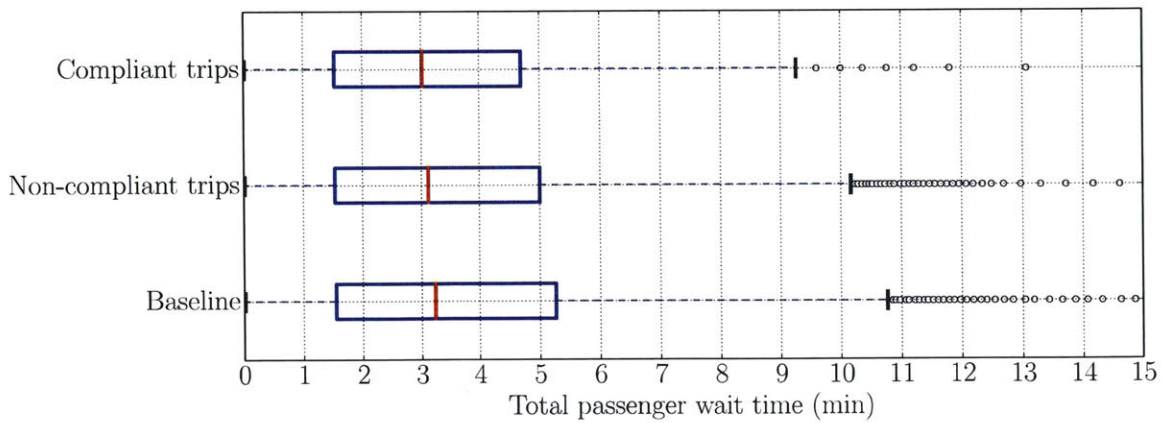
(c) Afternoon peak.

Figure 5-15: Distributions of wait times for eastbound passengers.
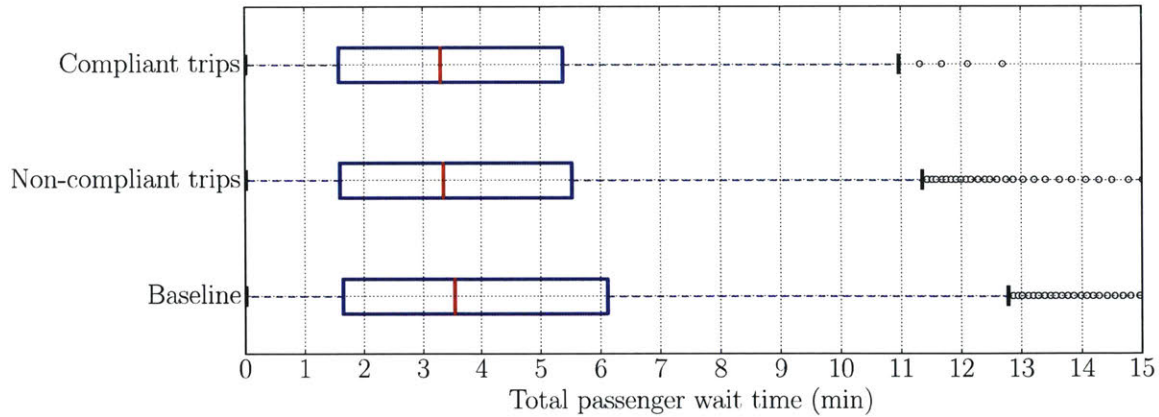
Figure 5-16: Distributions of westbound wait times during the afternoon peak.

boarding compliant trains, and for westbound afternoon peak passengers, wait reliability increases from 74.3% to 79.5%. This suggests that, in the peak directions of demand, 5–8% of passengers, who wait longer than the acceptable duration under current operations, would experience acceptable wait times with terminal headway-based control.

As noted in the discussion in Chapter 3, the MBTA currently has set as a standard a wait reliability of 90%. Table 5.10 shows that this level may be achievable across all passengers on the Riverside branch when controlling service only at the outer terminal. Figure 5-17 compares morning peak wait time reliability by station. Under headway-based control from Riverside, the 90% standard is met at six out of thirteen surface-level stations (at two more stations, 89% of passengers have reliable waits). Under existing operations, only one station achieves a reliability of 89%. Wait reliability decreases eastbound of Riverside until the final station along the branch segment, Fenway. Even with proper, headway-based dispatching at Riverside, headways become more variable and wait times more unreliable due to factors such as the variability of operator speeds and of dwell times (because of unequal demand within a short period of time, an unequal distribution of passengers boarding the two cars of a train, or even passengers interacting with the onboard farebox with cash).

The results in this section reinforce the conclusion from Chapter 3 that implementing proper terminal dispatching is a major—but only one—tool for agencies to improve service reliability.
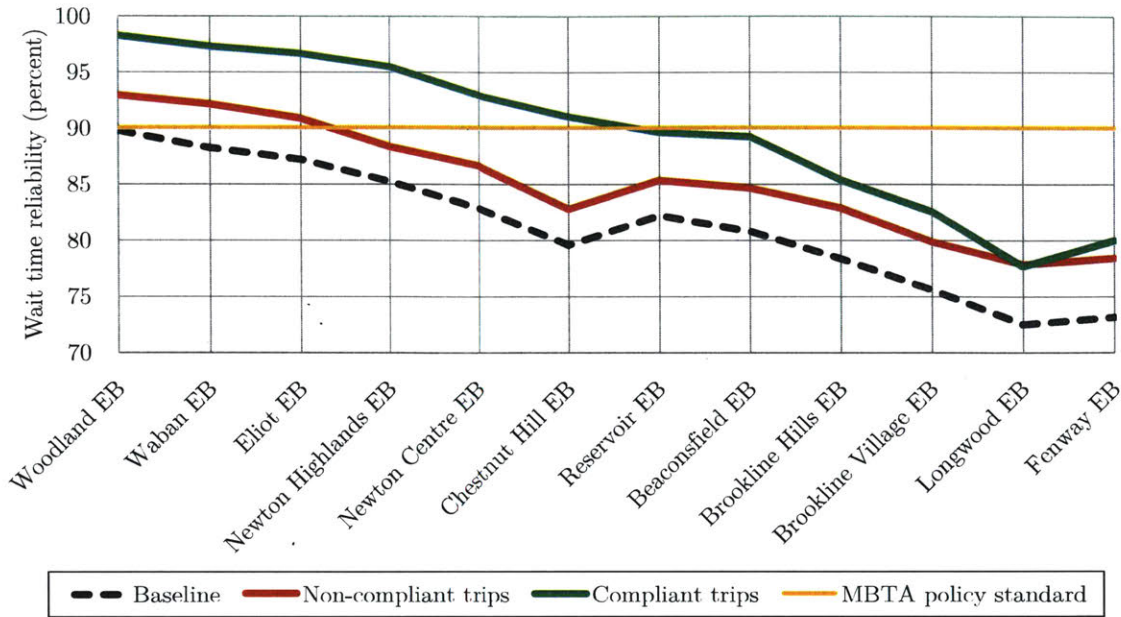
Figure 5-17: Stop-level wait time reliability along the surface-level of the Riverside branch during morning peak period.

### 5.5.3 Passenger journey times

We compare passenger journey times between existing service and compliant service during the experiment. Journey time is defined as the time from a passenger's arrival at the origin station until his or her arrival at the destination station.

Statistics for journey times between the Riverside branch and the trunk are shown in Tables 5.11 and 5.12. The variation of journey times is naturally high because different passengers have different origins and destinations, but the samples of passengers are large enough to compare between different operating environments. During the dominant directions of travel during peak periods (eastbound in the morning and westbound in the afternoon), average journey time decreases by more than 90 seconds for passengers riding compliant trips. In all other directions and times, except the morning peak traveling away from Boston, the mean journey time decreases by smaller but still statistically significant durations. In both directions of travel, the standard deviation of journey times decreases when recommended departure times are followed.

The decrease in average journey time is greater than that of average wait time, meaning that the journey time improvements come from both improvements in wait times and faster running times.

Table 5.11: Passenger journey time statistics during the morning peak.

| Service scenario | Eastbound | | Westbound | |
| --- | --- | --- | --- | --- |
| | Mean (min) | Std dev (min) | Mean (min) | Std dev (min) |
| Baseline | 30.6 | 17.1 | 21.6 | 11.8 |
| Non-compliant trips | 29.2 | 11.0 | 22.0 | 10.5 |
| Compliant trips | 28.9 | 10.9 | 22.6 | 10.4 |

Table 5.12: Passenger journey time statistics during the afternoon peak.

| Service scenario | Eastbound | | Westbound | |
| --- | --- | --- | --- | --- |
| | Mean (min) | Std dev (min) | Mean (min) | Std dev (min) |
| Baseline | 26.8 | 11.5 | 24.7 | 13.6 |
| Non-compliant trips | 25.6 | 11.0 | 23.3 | 12.8 |
| Compliant trips | 25.9 | 10.9 | 23.0 | 9.9 |

## 5.6 Ideas to improve recommended departure compliance

This section lists several ideas to improve the compliance of departures for future implementations. The proposed solutions are based on the possible reasons for non-compliance as suggested in Section 5.3.

### 5.6.1 Real-time terminal displays

As recommended by Maltzan (2015), the LED displays on the platform should display the two upcoming recommended departure times, adjusted automatically as needed. This is similar to the current practice at the MBTA's heavy rail terminals, except the Green Line departure times would be based on the recommendations made by the decision-support system. An example of an LED sign displaying real-time train information at a mid-route station is shown in Figure 5-18; the existing signs at Riverside or other terminals would show similar information. Since the display is visible to the public, operators may be incentivized to depart on time in order to avoid complaints from passengers who rightly believe that the train should be departing. On the Green Line, LED displays are installed at Riverside and Lechmere and require only software changes; at the other two dispatching terminals, Boston College and Cleveland Circle, network upgrades would be necessary in order to install such displays.

This research proposes that additional displays be installed at terminals. These large-

Figure 5-18: LED display at Newton Centre showing the next two westbound arrivals.

format LCD displays, similar to those already used by the MBTA to show real-time information at several bus terminals, would list upcoming, recommended departure times to the second, the operators assigned to those trips, and the state of the train to be covering the trip (whether it is being pulled from the yard, is on the platform, or will be arriving from service). A mockup of the interface is proposed in Figure 5-19. Each large-format display would be located in a portion of the station inaccessible to the public—likely the inspector's booth—in order to protect the privacy of operators. If the operators' break room is located away from the platform, as is the case at Riverside, additional displays should be installed in those non-public areas so that operators can confirm their departures before walking to the platform.

## 5.6.2 In-vehicle electronic communication devices

The recent research by Lizana et al. (2014) in Santiago and Cats (2016) in Stockholm used in-vehicle communication consoles in order to send holding and speed-adjustment messages directly to en route operators. This is another avenue for the MBTA or other agencies to consider when aiming to improve departure adherence.

A network-connected console can provide various types of messages to the operator, including a readout of the train's current headway, or instructions to hold, express, or terminate at an upcoming stop. Naturally, any electronic devices in the operator's cab must include the simplest-possible interface to avoid dangerous distractions while operating a train (Maltzan, 2015). An example interface may have a black screen display only the words "At Reservoir: Hold for 2 minutes" or "At Reservoir: Hold until 15:25:30" next to a icon of a stop sign. At terminals or other locations, they could count down the seconds until the departure time. Use of such a tool could reduce the amount of radio traffic required. In addition to receiving messages, operators could press a button to request to speak to dispatchers, using either public or private audio channels; canned text messages could also be sent to the OCC with the touch of a button. Instructions sent to consoles could be connected directly to a decision-support system which automates the dissemination of messages with effects below a specified threshold, but this is not necessary for a functioning system of in-vehicle consoles, as inspectors and dispatchers may instead send all messages manually.

The MBTA currently uses in-vehicle consoles from the company Trapeze for its bus fleet. The units are linked to the radio network and display a bus operator's adherence to

Figure 5-19: Mock interface of a display for operators at a terminal.

scheduled timepoints but they do not display headway-based instructions or metrics, so a potential exists to also improve service for high-frequency bus routes if capable hardware and software were installed.

### 5.6.3 Adherence reporting

All parties involved in the operating the Green Line—operators, inspectors, and OCC dispatchers—need to have the same understanding of which service objectives and standards will be evaluated. Terminal inspectors primarily work toward schedule adherence, whereas mid-line inspectors and OCC dispatchers enforce headway adherence and are typically not concerned with the schedule. This can make for inefficient operations if, given a state in Green Line service, schedule and headway-based controls lead to different outcomes. As an example, a terminal inspector may release two trains within three minutes of each other, only to have the second train held for two or three minutes by another inspector several stops downstream in order to increase the spacing.

Until 2016, the MBTA Service Delivery Policy defined light rail reliability, along the branches, in terms of the percentage of trips which "operated within 1.5 scheduled headways" over a service day, with a goal of reaching this for 85% of trips (Massachusetts Bay Transportation Authority, 2010). The running time was also evaluated: trips should have been run within five minutes of the scheduled duration from one terminal to the other. A train running with a headway of 8.5 minutes in a period with a 6-minute scheduled headway would have been considered to be on time under previous policy. The newest revision is tied more directly to the customer experience by using the passenger wait reliability metric, defined in Section 3.3 (Massachusetts Bay Transportation Authority, 2017). The emphasis on headways and passenger waits must be conveyed to inspectors by ensuring that the tools produced for use in controlling service reflect the objectives set by the agency upon which these employees will be evaluated.

Previous researchers have suggested posting the on-time performance of operators at each terminal on a regular basis (Malikova, 2012), or providing incentives (in the form of extra pay or vacation hours) to the top percentiles of operators complying with instructions Maltzan (2015). The decision-support system used in the experiment generates compliance data. Reporting these to management and even providing operators with personalized compliance scorecards could motivate improved compliance. These ideas of softly incentivizing on-time

performance should be pursued, to improve morale among some operators. Any requirements to negotiate such programs with unions should not preclude their consideration. Indeed, union leaders may express interest in the idea of providing better service to passengers through improved behavior, thus improving their public perception.

## 5.7 Recommendations

Several recommendations are made to the MBTA for future work in implementing a headway-based dispatching policy relying on automated decision-support tools. The recommendations are based on the results of the experiment, a review of issues leading to non-compliance, and ideas to improve compliance. While the recommendations are specific to the Green Line case study, they can be adapted for similar projects on other transit modes and in other cities.

### 5.7.1 Operator culture and departure discipline

Requiring operators to check in with a terminal inspector two minutes before the first trip following a break is not enough to ensure that early departures can be made when necessary. A change in operator policies to allow for a five-minute buffer would allow for increased flexibility in operations, likely at a small monetary cost.

Separate rules requiring operators to be onboard their trains ahead of the instructed departure time should also be proposed. One simple implementation may involve having operators begin walking toward the train (and moving the vehicle onto the platform, if necessary) at the "boarding time," two or three minutes before the recommended departure time. The large-format displays depicted in Figure 5-19 should show the boarding time for each operator rather than the departure time.

Ultimately, the implementation of any major changes to terminal departures requires the retraining of inspectors and operators in order to realize the intended behavioral changes. Training sessions must emphasize the importance of service reliability as well as how it relates to other key goals, including safety. Supervisors should receive regular reports on departure compliance, and operators should be able to track their compliance, either through messaging in break rooms or by distributing personalized scorecards. Incentives can be offered for good compliance, and tips can be offered to improve poor compliance.

161

### 5.7.2 Terminal infrastructure

Improvements in communication systems at terminals are also crucial for implementing a dispatching system which uses sub-minute intervals. At Riverside, much of the infrastructure is already in place and only needs modification. A microphone connected to a public address system should be installed at every station with an inspector's booth and adjusted so that it can be clearly heard.

Despite the negative reactions to the ring-off bell during the Riverside pilot, the use of some signal to cue departures, audio or visual, is strongly recommended. Examples of a visual signal include an LCD sign posted just past the front of the train, or an indicator at the front of the cab. If the signal remains an audible bell, those implementing it must work closely with those in charge of the public address system to ensure that speakers are relocated as needed, and to make sure that the public address zones are configured so that only necessary speakers ring the bell. At Riverside, only the speakers closest to the eastern ends of the platforms should sound the departure bell (a practice done on some heavy rail terminals on the MBTA).

LED and LCD screens should be used as discussed in Section 5.6.1, with the recommended departure times appearing on the countdown clocks as well as in the publicly-accessible MBTA-realtime API. An additional, less intrusive chime should be considered to sound at the boarding time before departure. This short chime would be heard throughout the platform as an instruction for the next operators; it could be replaced with a recorded voice announcing that the "next Green Line train departs in two minutes".

Vehicle tracking should be made more accurate and more precise around terminals and in yards, potentially by adding AVI points on each track within a yard, or sending trains' GPS pings within yards to internal data feeds. This knowledge of vehicle availability would reduce the need for manual data entry and improve the data the headway-based algorithm uses to recommend departure times.

### 5.7.3 Terminal and layover procedures

This research follows the suggestion of (Malikova, 2012) to have trains arriving at terminals move into the departing position, leaving the doors open for passengers to board before the departure time. This already occurs at most times at Riverside and generally also at

Lechmere, both of which feature faregates. The practice can be followed at other locations with the MBTA's forthcoming movement to a proof-of-purchase system for the Green Line as part of the AFC 2.0 project (Block-Schachter, 2016). At Boston College, a terminal with a loop, trains which are arriving and will be reused on the next trip should be relieved of passengers upon arrival to the station and immediately looped to the boarding platform prior to transferring the train to another operator. This may add a minute or two to the duties of an operator about to be relieved, but any associated costs should be worth the increased reliability in departures.

Operators should sign in electronically so that check-in times can be recorded and the associated rules enforced. Requiring operators to sign in electronically may reduce late check-ins if they are aware that they may easily be held accountable for their tardiness. Radio-frequency identification (RFID) targets can be installed near or at the entrances of the booths at terminal stations. An inspector is still required with any individual's sign-in, however, due to the need to confirm an his or her fitness to work. One way this check can be ensured while still easing the responsibilities of the inspector is to introduce a multi-tier system. Upon arrival to the booth, following ths example, an operator sees her name next to a red dot on a large screen, indicating that she has not yet signed-in. She taps her employee badge at the entrance of the booth and the dot turns yellow; the screen lists her boarding time. She approaches the inspector on duty, who confirms that she is fit for duty and clicks her name on a handheld tablet, making the dot next to the name turn from yellow to green. At this point, the operator would be authorized to board and operate the train when signaled.

### 5.7.4  Future vehicle procurement

Procurements for the Type 10 and future light rail fleets should incorporate elements which will transmit more data or allow service control instructions to be transmitted more easily. At this time, the existing fleet can only transmit real-time data related to location. The existing routers, which are connected to the GPS units, could also be connected to fare validators and automatic passenger counting (APC) systems to gain real-time information on boardings and load, which could in turn be incorporated into the control strategy algorithms. Alternatively, passenger counts may be estimated using computer vision techniques if surveillance cameras are installed on Green Line cars, just as they are on many MBTA

buses. There may also be an ability to stream information from the electronic headsign, which would provide more frequent and detailed updates on a train's destination. The current destinations in the AVI system are limited in quantity (with more terminals to be added with the Green Line Extension), and updates can only be recorded when the train passes an AVI point. There may also be the potential for this information to be made public through in-station signage and APIs: passengers may be able to see, before leaving their homes, for example, that the next train is near capacity and choose to postpone their travel for several minutes in order to obtain a seat.

There should be a consideration into adding in-vehicle consoles to future fleets, not unlike those on the existing bus fleets. In addition to the benefits of sending and receiving instructions directly in the cab, operators would be able to sign into trains with their run and operator numbers (possibly with a tap of their RFID-enabled badges), reducing the need to infer which crew a train is carrying and which scheduled trip it is serving. Ideally, the AVI route code, headsign, and public address announcements would be controlled from the same console to reduce the amount of dashboard clutter.

### 5.7.5 Decision-support tool integration in operations

The MBTA should pursue the creation and use of a unified decision-support system for all Green Line supervisory personnel. The system, during the development phase, must have the full support of everyone to be impacted, including inspectors and operators and dispatchers. Tribone et al. (2014) concluded that, when creating new tools and metrics, the process involved in the creation is as important as the innovation itself. If a group outside of the Light Rail Operations department is in charge of the project, a liaison within Operations should be established with regular updates and meetings scheduled with Operations staff.

The dispatching control algorithm should be improved upon to include vehicular and crew constraints, and algorithms to determine the feasibility of different actions (such as short turns) should be formulated and tested. Algorithms and performance metrics should be aligned to be based on similar inputs, and inspectors and operators must be retrained on what they will be evaluated in addition to the functionality of the new system.

Opportunities for using decision-support tools beyond recommending real-time controls should also be explored. Yard masters and maintenance crews, for instance, may carry devices which show the real-time positions of trains within a yard, or warn them when an

in-service train requires repairs. Administrative tasks, such as the requesting and approval of overtime delay pay, can also be simplified if not automated, increasing productivity while reducing the likelihood of fraud.

Any decision-support system should rely on internal data sources and APIs. An internal API running on a server separate from public-facing API will be more resilient by not experiencing periodic downtime from increased demand during the peak periods, as the researchers experienced. An internal API will also be secure from external threats, such as denial of service (DoS) attacks. Security and redundancy is important for any electronic system upon which service depends.

# Chapter 6

# Conclusions

This chapter begins with a summary of the thesis, including findings from the simulation of and field experiment with real-time control strategies on a complex light rail network. The key findings and contributions of the research are then provided, followed by recommendations for the MBTA or another transit organization to follow when implementing a headway-based service control policy. Finally, avenues for future research are offered.

## 6.1 Summary and findings

This research has expanded the study of headway-based control actions to a complex, multi-branch light rail service, using simulation and an experiment implementing a decision-support system. The research uses the MBTA Green Line, which has recently expanded vehicle tracking to cover the entire system, as a case study. Chapter 1 of the thesis begins by outlining the motivation for this research in terms of service reliability, technology advancements, and social and economic growth, particularly as the Green Line prepares for an expansion which will add two branches to its eastern end. An overview of the Green Line and its existing operations is provided, followed by a literature review.

### 6.1.1 AVL data cleaning

Chapter 2 begins by presenting common issues and limitations experienced with the vehicle location data available for the MBTA Green Line, in both its archived and real-time forms. Some of these issues relate to the pre-processing of the different raw data sources, while others involve the reliability of the tracking infrastructure itself. The chapter continues

by detailing an algorithm for cleaning archived AVL data using information from an entire service day. Use of this algorithm results in improved continuity within trips when analyzing characteristics such as headway and running times, since errors in train information are corrected and missing data are imputed. The algorithm is run each morning to clean the previous day's vehicle location data before they are stored on a server for later use by researchers (such as in Chapters 3 and 5). Strategies for cleaning real-time data are also described.

### 6.1.2 Service control simulation

Chapter 3 presents a simulation model of the Green Line based on the framework created by Sánchez-Martínez (2013). The model takes as inputs running times from AVL data cleaned with the algorithm in Chapter 2, demand from ODX (which also uses the AVL data), track and signal infrastructure, and schedule data from GTFS. The model allows for scenarios implementing different service control strategies to be tested and compared.

An experiment is run to determine effective holding strategies for a light rail service as complex as the Green Line. The current practice of schedule-based dispatching is compared against a target headway-based strategy. Combinations of holding points at terminals and en route were tested for each control strategy to determine where supervisory resources would be most efficiently utilized. Holding trains at terminals until the scheduled headway from the previous trip has elapsed significantly improves passenger wait time reliability during the peak hours and, as a result, decreases the 90th percentile of passenger journey times by 12%. Holding trains en route yields only marginal additional benefits to wait and journey times, and should be limited in duration if conducted along the trunk of the line. These results inform the strategy implemented in the experiment described in Chapter 5, that is to focus resources on terminal headway control rather than en route holding.

### 6.1.3 Decision-support system and tool

Chapter 4 begins by specifying a framework for implementing a decision-support system for improving reliability on a transit line. The framework builds upon what Maltzan (2015) conceptualized by incorporating information that inspectors and dispatchers consider when making control action decisions, such as vehicle and crew constraints. The framework is generalized so that it can be adapted and used outside of the case study for this research.

The second part of the chapter describes the creation of the decision-support tool for use in the headway control experiment in Chapter 5. The tool's software takes as inputs real-time AVL data, crew and vehicle schedules, and changes entered via the user interface. Trains en route toward and departing from the Riverside branch terminal, the experiment site, are detected and matched to a schedule trip. The user interface allows the user to switch the order of departing operators, designate trains as coming in and out of service, cancel trips, and make other common adjustments. Based on these inputs, the decision-support tool, in the form of a web application accessible from a browser, applies the service control policies specified in Section 5.1.1 and displays optimal departure times. A bell, activated within the same web application, was installed at the testing terminal in order to notify operators on the platform of revised departure times.

### 6.1.4   Field experiment

Chapter 5 covers the implementation of the headway control experiment at Riverside terminal on the MBTA Green Line. The chapter begins by describing the experimental even-headway dispatching policy on the Riverside branch before outlining the logistics involved in running the pilot. During the two-week experiment, terminal inspectors simultaneously used both the preexisting paper train sheets and the new decision-support tool, implemented on a commercial tablet computer. Assistants from outside of the Operations department supervised the experiment at all times, providing technical support as necessary.

The results from the experiment show that, during the times that the recommendations were properly followed, the variability of headways during the peak hours decrease by 40%. Passenger wait times, on average, decrease by 30 seconds during the peak period. The improvement in service quality is equivalent to adding 1.7 trains to service. Problems with ensuring that operators departed the terminal at the recommended time prevented these benefits from being recorded at all times. Recommendations for improving terminal operations and increasing compliance in future implementations are provided.

## 6.2   Contributions

This research contributes to the body of research into the holding of transit vehicles as a control strategy, the implementation of real-time control strategies using a decision-support

system, and AVL data cleaning.

Many of the techniques for cleaning archived and real-time AVL data, as described in Chapter 2, while formed specifically to address the issues found on the Green Line, could be generally applied to other transit agencies or technology companies wishing to improve the quality of their vehicle location data. Before the start of this research, the quality of the Green Line data was sufficient to provide station predictions in real time, but inadequate for aggregate service analyses by time of day or for a decision-support system.

The analyses in Chapter 3 extends the simulation of holding strategies to a complex, multi-branch light rail service. The framework used for the simulation is agent-based and event-driven and incorporates various automatically-collected data. The results from the simulation experiment are similar to previous, simpler applications on light rail and BRT trunks: headway-based holding is more effective than schedule-based holding, and regular dispatching at terminals is critical to maintaining service quality throughout the entire line.

Finally, this research represents the first known use of a decision-support system which recommends headway-based holding while incorporating crew data. Compared to the works of Lizana et al. (2014) and Cats (2016), this system consists of a tool designed for decision-makers potentially deciding between many alternative service options, rather than automating the process and sending instructions directly to operators. The results and feedback show that a support system similar to the one designed is feasible for use as long as considerable retraining is conducted. The holding algorithm, pilot application, and the broader framework on which the application is based can be used for future implementations, on the Green Line and other complex transit lines.

## 6.3   Recommendations

The key recommendation following this research is for the MBTA to adopt a policy of headway-based dispatching at the outer terminals of the Green Line. Consistent, headway-based dispatching at terminals should be prioritized over en route control, and supervisory resources should be shifted accordingly. These recommendations can be adapted and applied to other high-frequency light rail and bus services of varying complexity, and are mostly taken from the recommendations previously made in this thesis.

1   Dispatching at terminals should be based on maintaining the scheduled headway rather

than attempt to follow the scheduled times precisely. If ensuring the success of the headway-based dispatching requires additional supervisory posts at terminals, mid-line inspectors along the branches (such as at Coolidge Corner or Harvard Avenue) should be reassigned to the terminals. Headway-based holding at terminals will reduce the need for en route holding, which then could be directed by OCC dispatchers.

2    Any implementation of headway-based dispatching at terminals must involve the consistent use of real-time AVL information. This requires having inspectors use a single website or application in order to standardize the decision-making process. This need can be filled using a new decision-support system.

3    The MBTA should consider deploying a decision-support system that is adaptable to operate within a variety of situations and can detect and flag exceptions in normal service patterns. This research envisions supervisory personnel being provided with handheld tablets, but other configurations are possible, such as installing a console in each vehicle for the operator to receive direct messages from inspectors or the OCC. The implementation of any support system requires an agreed upon control policy to be programmed in the recommendation engine. The opportunity may also warrant a reevaluation of the roles and relationships between field inspectors and control center dispatchers.

4    Dispatchers, inspectors, and operators must understand how performance of their own jobs and of service as a whole are evaluated. Agency performance metrics should align with the goals of the agency, and in turn, any decision-support tools developed must be geared to work toward those metrics.

5    Beyond using headway-based dispatching, procedures at Green Line terminals must change in order to encourage on-time departures. At every terminal, arriving trains should be pulled into the departure berth, with all doors left open for passengers to board without delaying the departure. The adjustment can coincide with a change to a proof-of-payment fare enforcement policy. Operator policy must also be changed to ensure that those returning from breaks check in with the inspector more than two minutes before the next scheduled departure. The use of digital displays, departure signals, and potentially financial rewards for compliance may also encourage a culture

of on-time performance.

6   Future light rail vehicle procurements should consider adding features to allow passenger load, operator, run, and destination information to be wirelessly transmitted from the train in real time. Information on vehicle crowding and stops to be skipped could be published through public APIs and added to digital station signage.

## 6.4   Future research

This research concludes by listing several possible areas for further researchers to investigate. They are generally in regard to simulation modeling but are also relevant for informing the development of decision-support tools and reducing operator behavior variability.

The simulation experiment presented in Chapter 3 tested headway-based holding, but controllers were also specified for short-turning and expressing. Using the framework by Sánchez-Martínez (2013), controllers for other control actions can be easily specified and added to stations or inter-station links. Additional experiments should be run to determine the optimal use of these additional strategies on the Green Line or other transit services. The results can be applied in the field in various ways, from providing pamphlets listing guidelines and typical effects of specific actions to full incorporation into a decision-support system. More sophisticated holding strategies, such as the one by Sánchez-Martínez et al. (2016b), could also be tested on the Green Line.

The algorithms in the controllers should take into account practical considerations when deciding on control actions. These include multiple costs and constraints, such as passenger comfort, the availability of cover personnel and extra vehicles, and scheduled operator relief times. Passenger comfort can already be estimated from passengers' in-vehicle travel times and vehicle load, both of which are provided by the simulation engine and can be used within the model. Another avenue for future study is the variability of operator behavior. The simulation experiment in this research included a factor to simulate the current variability in schedule adherence, but additional research may include the creation of an operator speed factor. Such a factor would simulate the difference between conservative and aggressive operators, or slow and fast operators. If used in simulation, this would require changes in how the simulated running times are drawn from historic observations. In real-time operations, an operator speed factor could be assigned to a trip in progress and used to improve downstream

arrival predictions or to suggest slowing, holding, or expressing interventions to regulate a particularly fast or slow driver.

The assumption of stochastic passenger arrivals at stations, used for simulating demand in Section 3.3, should be examined when real-time arrival predictions are publicly available. This is likely to be a factor at some stations on the Green Line more than others. At Riverside, for instance, many passengers drive to the station from suburbs and thus may be unable to time their journeys as precisely than if by foot. Additional research should develop a new model for passenger arrival rates as a function of the scheduled headway and the mode shares of station arrivals. If a stochastic distribution is found to not be a good approximation of passenger arrivals, changes in wait time-based performance metrics may be necessary.

The work of Wong (2000), who evaluated the benefits of coordinating transfers between the MBTA Green and Red Lines in Boston, should be reevaluated in light of more accurate vehicle location information. The previous research found transfer coordination to be only effective during off-peak hours, when scheduled headways are long. An advanced decision-support system, coordinated across different lines and modes within a public transportation system, may be able to continuously evaluate transfer opportunities between metro lines and bus services, regardless of scheduled headways. Based on estimated or real-time vehicle load information, it may determine whether holding a vehicle for a short duration would result in a large passenger benefit, in which case instructions can be sent directly to the bus or train.

Finally, more can be learned from labor research and best practices in the transit industry on operator and supervisor behavior and culture. A plan or toolkit should be developed for fostering a positive and passenger-focused work culture which makes service reliability a top priority. Recommendations for retraining employees should be included in such a plan.

# Bibliography

American Public Transportation Association (2007). Public Transportation: Benefits for the 21st Century. http://www.apta.com/resources/reportsandpublications/Documents/twenty_first_century.pdf

Anosike, D. and L. T. Thomas (2016, November). Comprehensive Annual Financial Report for the Fiscal Years Ended June 30, 2016 and 2015. Technical report, Washington Metropolitan Area Transit Authority. https://www.wmata.com/about/records/public_docs/upload/CAFR-Final-12-15-16.pdf

Barnett, A. (1974). On Controlling Randomness in Transit Operations. *Transportation Science 8*(2), 102–116. doi:10.1287/trsc.8.2.102

Barry, M. and B. Card (2014, October). MBTA Releases Green Line Data. http://mbtaviz.github.io/green-line-release/

Bartholdi, J. J. and D. D. Eisenstein (2012). A self-coördinating bus route to resist bus bunching. *Transportation Research Part B: Methodological 46*(4), 481–491. doi:10.1016/j.trb.2011.11.001

Bazelon, C., P. Seth, S. Herscovici, M. Berkman, A. R. Sanderson, B. Humphreys, J. J. Floyd, and M. P. Abasciano (2015, August). Analysis of the Boston 2024 Proposed Summer Olympic Plans. Technical report, The Brattle Group, Inc. http://www.mass.gov/governor/docs/news/final-brattle-report-08-17-2015.pdf

Berrebi, S. J., K. E. Watkins, and J. A. Laval (2015). A real-time bus dispatching policy to minimize passenger wait on a high frequency route. *Transportation Research Part B: Methodological 81, Part 2*, 377–389. doi:10.1016/j.trb.2015.05.012

Block-Schachter, D. (2016, October). AFC 2.0: The next generation of MBTA fare collection. Discussion presentation. http://www.mbta.com/uploadedfiles/About_the_T/Board_Meetings/I.%20AFC2%20Policy%20Board%20Oct2016%20-TO%20POST.pdf

Carnaghi, P. (2014). Decision support systems in transportation: The case of surface urban public transport. Master's thesis, Politecnico di Milano. http://hdl.handle.net/10589/94981

Carrel, A., R. Mishalani, N. Wilson, J. Attanucci, and A. Rahbee (2010). Decision Factors in Service Control on High-Frequency Metro Line. *Transportation Research Record: Journal of the Transportation Research Board* (2146), 52–59. doi:10.3141/2146-07

Cats, O. (2016, August). A Brief Summary of Real-time Bus Regularity Control: From a Simulation Study to Full-scale implementation. `https://odedcats.weblog.tudelft.nl/files/2016/08/Real-time-control-weblog-summary.pdf`

Cats, O., A. Larijani, H. N. Koutsopoulos, and W. Burghout (2011). Impacts of Holding Control Strategies on Transit Performance. *Transportation Research Record: Journal of the Transportation Research Board* (2216), 51–58. doi:10.3141/2216-06

Central Transportation Planning Staff (2016, February). Massachusetts Bay Transportation Authority Rapid Transit/Key Bus Routes Map.

Daganzo, C. F. and J. Pilachowski (2011). Reducing bunching with bus-to-bus cooperation. *Transportation Research Part B: Methodological 45*(1), 267–277. doi:10.1016/j.trb.2010.06.005

Deckoff, A. A. (1990). The short-turn as a real time transit operating strategy. Master's thesis, Massachusetts Institute of Technology. `http://hdl.handle.net/1721.1/41774`

Delgado, F., J. C. Muñoz, and R. Giesen (2012). How much can holding and/or limiting boarding improve transit performance? *Transportation Research Part B: Methodological 46*(9), 1202–1217. doi:10.1016/j.trb.2012.04.005

Dumas, R. A. (2015). Analyzing transit equity using automatically collected data. Master's thesis, Massachusetts Institute of Technology. `http://hdl.handle.net/1721.1/103650`

Eberlein, X. J., N. H. M. Wilson, and D. Bernstein (2001). The Holding Problem with RealâĂŞTime Information Available. *Transportation Science 35*(1), 1–18. doi:10.1287/trsc.35.1.1.10143

Fabian, J. J. and G. E. Sánchez-Martínez (2017). Simulation-Based Comparison of Holding Strategies for a Multi-Branch Light-Rail Service. *Transportation Research Record: Journal of the Transportation Research Board* (2648). Forthcoming. doi:10.3141/2648-03

Fellows, R. E. (1990). Using and Enhancing Automatic Vehicle Identification to Improve Service Control on the MBTA Green Line. Master's thesis, Massachusetts Institute of Technology.

Gordon, J. B., H. Koutsopoulos, N. Wilson, and J. Attanucci (2013). Automated Inference of Linked Transit Journeys in London Using Fare-Transaction and Vehicle Location Data. *Transportation Research Record: Journal of the Transportation Research Board* (2343), 17–24. doi:10.3141/2343-03

Hernández, D., J. C. Muñoz, R. Giesen, and F. Delgado (2015). Analysis of real-time control strategies in a corridor with multiple bus services. *Transportation Research Part B: Methodological 78*, 83–105. doi:10.1016/j.trb.2015.04.011

IBI Group (2011). Appendix B - Proposed Green Line tracking system - AVI and signal system elements.

Iseki, H. and B. D. Taylor (2010). Style versus Service? An Analysis of User Perceptions of Transit Stops and Stations. *Journal of Public Transportation 13*(3), 39–63. doi:10.5038/2375-0901.13.3.2

Kittelson & Associates, Inc. (2007). *TCRP Report 118: Bus Rapid Transit PractitionerâĂŹs Guide*, Chapter 5. Transportation Research Board. Research sponsored by the Federal Transit Administration in cooperation with the Transit Development Corporation.

Levenson, M. and M. Arsenault (2015, July). In gloomy time, poll finds drop in Games support. *The Boston Globe*. https://www.bostonglobe.com/metro/2015/02/19/ poll-finds-riders-snow-slammed-are-less-likely-support-boston-bid-for-olympics/ moUSAH3qfPlEA4abZQz1pO/story.html

Litman, T. (2017, February). Understanding Transport Demands and Elasticities: How Prices and Other Factors Affect Travel Behavior. Policy paper, Victoria Transport Policy Institute. http://www.vtpi.org/elasticities.pdf

Lizana, P., J. C. Muñoz, R. Giesen, and F. Delgado (2014). Bus Control Strategy Application: Case Study of Santiago Transit System. *Procedia Computer Science 32*, 397–404. doi:10.1016/j.procs.2014.05.440

LTK Engineering Services (2012). LTK OnTrack Positive train control simulations: Green Line base case.

Macchi, R. A. (1989). Expressing Trains on the MBTA Green Line. Master's thesis, Massachusetts Institute of Technology.

Malikova, A. A. (2012). MBTA Green Line 3-Car Train Operating Plans to Enhance Capacity and Reliability. Master's thesis, Massachusetts Institute of Technology. http: //hdl.handle.net/1721.1/74271

Maltzan, D. (2015). Using Real-time Data to Improve Reliability on High-Frequency Transit Services. Master's thesis, Massachusetts Institute of Technology. http://hdl.handle. net/1721.1/99541

Massachusetts Bay Transportation Authority (2010, June). Service Delivery Policy: 2010 Update. Technical report. http://mbta.com/uploadedfiles/About_the_T/T_ Projects/T_Projects_List/2010ServiceDeliveryPolicy.pdf

Massachusetts Bay Transportation Authority (2014, July). MBTA Ridership and Service Statistics: Fourteenth Edition 2014. Technical report.

Massachusetts Bay Transportation Authority (2015, November). Green Line Type 7 Fleet Selective Systems Overhaul. Discussion presentation. http://www.mbta.com/ uploadedfiles/About_the_T/Board_Meetings/Type7-SSO%20updated.pdf

Massachusetts Bay Transportation Authority (2017, January). Service Delivery Policy: 2017 Update. Technical report. http://www.mbta.com/uploadedfiles/About_the_T/Board_ Meetings/MBTA%20Service%20Delivery%20Policy%202017%20FINAL.pdf

Massachusetts Bay Transportation Authority and IBI Group (2016, April). MBTA-realtime API for Performance Data Documentation (v0.9.1). Technical report. http://realtime.mbta.com/Portal/Content/Documents/MBTA-realtime_ Performance_APIDocumentation_v0_9_1_2016-04-14_public.pdf

Massachusetts Department of Transportation (2016, May). Green Line Extension Review Interim Project Management Team: Final Report. http://www.apta.com/resources/reportsandpublications/Documents/twenty_first_century.pdf

Michaud, N. and T. Thibodeaux (2017, March). O.T.P. Tablet - Riverside Station Pilot program. Memorandum.

ming Lin, T. and N. H. M. Wilson (1992). Dwell Time Relationships for Light Rail Systems. *Transportation Research Record: Journal of the Transportation Research Board* (1361), 287–295.

Muller, T. H. J. and P. G. Furth (2000). Integrating Bus Service Planning with Analysis, Operational Control, and Performance Monitoring. In *ITS America 10th Annual Meeting and Exposition: Revolutionary Thinking, Real Results.*

Office of Geographic Information (MassGIS) (2014a, September). MBTA Rapid Transit. http://www.mass.gov/anf/research-and-tech/it-serv-and-support/application-serv/office-of-geographic-information-massgis/datalayers/mbta.html

Office of Geographic Information (MassGIS) (2014b, September). USGS Color Ortho Imagery (2013/2014). http://www.mass.gov/anf/research-and-tech/it-serv-and-support/application-serv/office-of-geographic-information-massgis/datalayers/colororthos2013.html

OpenStreetMap contributors (2017). OpenStreetMap, Standard Map. https://www.openstreetmap.org

Pangilinan, C., N. Wilson, and A. Moore (2008). Bus Supervision Deployment Strategies and Use of Real-Time Automatic Vehicle Location for Improved Bus Service Reliability. *Transportation Research Record: Journal of the Transportation Research Board* (2063), 28–33. doi:10.3141/2063-04

Sánchez-Martínez, G. E. (2013). Running Time Variability and Resource Allocation: A Data-Driven Analysis of High-Frequency Bus Operations. Master's thesis, Massachusetts Institute of Technology. http://hdl.handle.net/1721.1/79498

Sánchez-Martínez, G. E. (2015). *Real-Time Operations Planning and Control of High-Frequency Transit.* Ph. D. thesis, Massachusetts Institute of Technology. http://hdl.handle.net/1721.1/99550

Sánchez-Martínez, G. E. (2017a). Estimating Fare Noninteraction and Evasion Using Disaggregate Fare Transaction Data. *Transportation Research Record: Journal of the Transportation Research Board* (2652). Forthcoming.

Sánchez-Martínez, G. E. (2017b). Inference of Public Transportation Trip Destinations by Using Fare Transaction and Vehicle Location Data: Dynamic Programming Approach. *Transportation Research Record: Journal of the Transportation Research Board* (2652), 1–7. doi:10.3141/2652-01

Sánchez-Martínez, G. E., H. N. Koutsopoulos, and N. H. Wilson (2016a). Optimal allocation of vehicles to bus routes using automatically collected data and simulation modelling. *Research in Transportation Economics 59*, 268–276. doi:10.1016/j.retrec.2016.06.003

Sánchez-Martínez, G. E., H. N. Koutsopoulos, and N. H. M. Wilson (2015). Event-Driven Holding Control for High-Frequency Transit. *Transportation Research Record: Journal of the Transportation Research Board* (2535), 65–72. doi:10.3141/2535-07

Sánchez-Martínez, G. E., H. N. Koutsopoulos, and N. H. M. Wilson (2016b). Real-time holding control for high-frequency transit with dynamics. *Transportation Research Part B: Methodological 83*, 1–19. doi:10.1016/j.trb.2015.11.013

Sindel, D. (2017). Strategies for Meeting Future Capacity Needs on the Light Rail MBTA Green Line. Master's thesis, Massachusetts Institute of Technology.

Soeldner, D. W. (1993). A Comparison of Control Options on the MBTA Green Line. Master's thesis, Massachusetts Institute of Technology.

Tribone, D., D. Block-Schachter, F. Salvucci, J. Attanucci, and N. Wilson (2014). Automated, Data-Driven Performance Regime for Operations Management, Planning, and Control. *Transportation Research Record: Journal of the Transportation Research Board* (2415), 72–79. doi:10.3141/2415-08

Tribone, D., L. Riegal, R. Warade, and D. Barker (2016, January). Measuring transit agency performance using open realtime data. In *TRB 95th Annual Meeting Compendium of Papers*, Transportation Research Board 95th Annual Meeting, Washington, District of Columbia.

Turnquist, M. A. (1981). Strategies for Improving Reliability of Bus Transit Service. *Transportation Research Record: Journal of the Transportation Research Board* (818). http://onlinepubs.trb.org/Onlinepubs/trr/1981/818/818-002.pdf

van Oort, N., N. Wilson, and R. van Nes (2010). Reliability Improvement in Short Headway Transit Services. *Transportation Research Record: Journal of the Transportation Research Board* (2143), 67–76. doi:10.3141/2143-09

Wilson, N. H. M., R. A. Macchi, R. E. Fellows, and A. A. Deckoff (1992). Improving Service on the MBTA Green Line Through Better Operations Control. *Transportation Research Record: Journal of the Transportation Research Board* (1361), 296–304.

Wilson, N. H. M., D. Nelson, A. Palmere, T. H. Grayson, and C. Cederquist (1992). Service-Quality Monitoring for High-Frequency Transit Lines. *Transportation Research Record: Journal of the Transportation Research Board* (1349), 3–11.

Wong, C. J. (2000). Improving the network performance of urban transit systems. Master's thesis, Massachusetts Institute of Technology. http://hdl.handle.net/1721.1/81551