Project 2008 | October 2024

SJSU SAN JOSÉ STATE UNIVERSITY

MINETA TRANSPORTATION INSTITUTE

Measurement and Prediction of Transit System Performance Using Probe Data Generated through DSRC and non-DSRC Technologies

Gregory L. Newmark, PhD

MINETA TRANSPORTATION INSTITUTE

Founded in 1991, the Mineta Transportation Institute (MTI), an organized research and training unit in partnership with the Lucas College and Graduate School of Business at San José State University (SJSU), increases mobility for all by improving the safety, efficiency, accessibility, and convenience of our nation's transportation system. Through research, education, workforce development, and technology transfer, we help create a connected world. MTI leads the Mineta Consortium for Equitable, Efficient, and Sustainable Transportation (MCEEST) funded by the U.S. Department of Transportation, the California State University Transportation Consortium (CSUTC) funded by the State of California through Senate Bill 1 and the Climate Change and Extreme Events Training and Research (CCEETR) Program funded by the Federal Railroad Administration. MTI focuses on three primary responsibilities:

Research

MTI conducts multi-disciplinary research focused on surface transportation that contributes to effective decision making. Research areas include: active transportation; planning and policy; security and counterterrorism; sustainable transportation and land use; transit and passenger rail; transportation engineering; transportation finance; transportation technology; and workforce and labor. MTI research publications undergo expert peer review to ensure the quality of the research.

Education and Workforce Development

To ensure the efficient movement of people and products, we must prepare a new cohort of transportation professionals who are ready to lead a more diverse, inclusive, and equitable transportation industry. To help achieve this, MTI sponsors a suite of workforce development and education opportunities. The Institute supports educational programs offered by the Lucas Graduate School of Business: a Master of Science in Transportation Management, plus graduate certificates that include High-Speed and Intercity Rail Management and Transportation Security Management. These flexible programs offer live online classes so that working transportation professionals can pursue an advanced degree regardless of their location.

Information and Technology Transfer

MTI utilizes a diverse array of dissemination methods and media to ensure research results reach those responsible for managing change. These methods include publication, seminars, workshops, websites, social media, webinars, and other technology transfer mechanisms. Additionally, MTI promotes the availability of completed research to professional organizations and works to integrate the research findings into the graduate education program. MTI's extensive collection of transportation-related publications is integrated into San José State University's world-class Martin Luther King, Jr. Library.

Report 24-33

Measurement and Prediction of Transit System Performance Using Probe Data Generated through DSRC and non-DSRC Technologies

Gregory L. Newmark

October 2024

A publication of the
Mineta Transportation Institute
Created by Congress in 1991

College of Business
San José State University
San José, CA 95192-0219

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
24-33		

4. Title and Subtitle
Measurement and Prediction of Transit System Performance Using Probe Data Generated through DSRC and non-DSRC Technologies

5. Report Date
October 2024

6. Performing Organization Code

7. Authors
Gregory L. Newmark ORCID: 0000-0002-8780-8832

8. Performing Organization Report
CA-MTI-2008

9. Performing Organization Name and Address
Mineta Transportation Institute
College of Business
San José State University
San José, CA 95192-0219

10. Work Unit No.

11. Contract or Grant No.
65A0660

12. Sponsoring Agency Name and Address
California Department of Transportation
1120 N Street
Sacramento, CA 95814

13. Type of Report and Period Covered

14. Sponsoring Agency Code

15. Supplemental Notes
10.31979/mti.2024.2008

16. Abstract
This research explores the application of two different probe data standards to transit performance measurement. The first section chronicles the proposed and implemented transit uses of dedicated short-range communication (DSRC) technologies during the two decades between the standard's emergence and its announced sunset. The research finds that, despite proposed applications across safety, operation, and information domains, DSRC never became embedded in transit operations. By contrast, the general transit feed specification (GTFS) standard with its real-time (RT) extension has been widely embraced and offers the potential to use the associated *VehiclePosition* messages as probe data to generate detailed transit performance metrics. The second section presents a method to decompose transit routes into segments (defined as the path between subsequent stops) and to impute three key time points for each segment: arrival time at the segment start, departure time from the segment start, and arrival time at the segment end. This method is applied to several days of GTFS-RT data from the Modesto Area Express to assess the accuracy of the imputation (in comparison to reported times from the *TripUpdate* messages) and, in the third and final section, to generate and visualize a series of performance metrics. These metrics assess the deviation between the designed travel times from the GTFS Schedule data and the actual travel times imputed from the GTFS-RT feeds. The research demonstrates an innovative approach to transform probe data generated by GTFS-RT feeds into valuable measures of transit performance at segment-level granularity. The research emphasizes the importance of segment-level transit performance measurement while recognizing that the methods used to impute segment time points are likely to change over time as the GTFS-RT standard evolves.

17. Key Words
GTFS,
Transit,
Performance Measurement,
Probe-Data,
DSRC

18. Distribution Statement
No restrictions. This document is available to the public through The National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	67	

Form DOT F 1700.7 (8-72)

DOI: 10.31979/mti.2024.2008

Mineta Transportation Institute
College of Business
San José State University
San José, CA 95192-0219

Tel: (408) 924-7560
Fax: (408) 924-7565
Email: mineta-institute@sjsu.edu

transweb.sjsu.edu/research/2008

Acknowledgments

The author is deeply grateful to Bradley Mizuno of Caltrans and Karen Philbrick and Hilary Nixon of the Mineta Transportation Institute for their structured and supportive project management. Their organizational acumen, personal kindness, and generous administrative efforts were essential to the success of this project.

The author thanks Balwinder Tarlok and Chad Riding for shaping the direction of this research at its outset and, along with Bradley Mizuno, for organizing the feedback from Caltrans. It is inspiring to do research for an agency committed to putting that research into action.

Finally, the excellent computer science skills and creative efforts of Alex Brufsky, Andrew "Ollie" Elmgren, and Peter Haas made this project possible. Their work to gather, manage, and explore the data at the heart of this research is very much appreciated.

This project is dedicated to the transit agency planners seeking to optimize the performance of their systems.

Contents

List of Figures

List of Tables

1. Introduction

Background

The wireless communication revolution has enabled the tracking of transit vehicles in real time, which offers new possibilities for measuring (and therefore improving) transit performance. These possibilities are strengthened by the emergence of message standards for describing the behavior of vehicles. While vendors can and have offered bespoke tracking solutions, the use of standardized approaches fosters scalability and more market entrants.

Public agencies particularly benefit when standards are adopted as that embrace enhances market competitiveness for all related products and reduces the risk that an investment in any given product will quickly become obsolete. As stewards of tax revenues, public agencies have heightened concerns over the prudent use of those resources. Standards, especially if supported by the federal government, offer more certainty to public agencies regarding the appropriateness of their investment decisions. Such certainty is important to officials seeking to wisely allocate tax revenues and avoid criticism regarding major infrastructure spending decisions.

Two standards have been particularly notable within the transportation space: dedicated short-range communication (DSRC) and the general transit feed specification (GTFS). The origin stories of these two standards vary substantially. DSRC was designed initially as a way for vehicles to communicate quickly with each other in real time. The standard was developed to support an anticipated profusion of new smart vehicle technologies and advanced by the federal government through the dedication of bandwidth. GTFS was created through a partnership of the Tri-County Metropolitan Transportation District of Oregon (TriMet) the transit operator for the Portland region, and Google, the search engine behemoth. Both parties wanted to solve the discrete problem of representing transit schedules in a standardized machine-readable fashion that would facilitate online trip planning. Both parties sought to directly apply the standard to their respective domains and thus were invested in the success of the standard.

Despite the substantial effort to craft DSRC and significant investment by the federal government and private companies to create products using the DSRC standard, no "killer app" arose to cement DSRC technology into the nation's use patterns. Concomitantly, other communication applications, such as Wi-Fi, competing for the bandwidth reserved for DSRC, did become embedded into American life—with the result that the original DSRC approach was sunset by the federal government. By contrast, the GTFS standard, initially designed to power Google Map's transit trip-planning feature, was widely and voluntarily embraced by the transit industry and, subsequently, the federal government.

Approach

This research was designed to explore both DSRC and non-DSRC probe technologies for their potential for transit performance measurement; however, soon after the project's launch, the federal government announced its decision to eliminate DSRC due to a lack of uptake. Consequently, this work shifted slightly to chronicle the published consideration of DSRC and transit. This accounting is important for understanding the potential and implemented applications and their lessons for the future.

At the same time, GTFS, which had been initially designed to describe static schedule information, was burgeoning into a broader suite of standards. One of these additional standards, GTFS Realtime (GTFS-RT), both functions as a probe technology capable of generating transit performance measures and has enjoyed a wide uptake by transit agencies. For this reason, GTFS-RT was chosen as the non-DSRC probe technology.

GTFS-RT offers an array of message sets with different information. The completeness of these feeds varies substantially despite theoretical adherence to a single standard. This variation complicates the design of an algorithm to transform GTFS-RT data into performance measures since expected data points are inconsistently available. An exploration of feeds from five different transit agencies in California confirmed this challenge. Consequently, the approach used to generate the core data to feed the transit performance measures was derived entirely from the most basic *VehiclePosition* messages that provides geocoded timestamps of transit vehicles.

This lowest common denominator approach to selecting the underlying data ensures that the data required for the analysis is always available. Furthermore, relying on the *VehiclePosition* data alone reduces the amount of data to be processed, an important cost reduction feature. This approach does, however, increase the need to process the data to generate the key fields for analysis. This step is an important contribution of this research, but one that would be unnecessary if all the GTFS-RT *TripUpdate* message sets included vehicle arrival and departure times for all stops along the way. Such feeds would require extensive data collection as they are substantially larger than the *VehiclePosition* feeds but have the benefit of reporting the desired values directly. (Eliciting those values does involve additional programming and processing, however.) It is expected that in the future, newer standards, such as the Transit ITS Data Exchange (TIDES) Data Specification Suite, will provide cleaner data sets for similar analysis. The goal of the current work is to provide a framework of metrics and a means to generate them from the most basic GTFS-RT available. This approach ensures that the GTFS-RT feeds currently being produced yield probe-based transit performance metrics.

An additional innovation of this work is to break transit trips down into segments defined spatially as the path taken by a vehicle between adjacent stops and temporally as the time from the arrival of the bus at the first stop to the arrival of the bus at the second stop. This time period is further subdivided into the dwell time at the stop and the travel time from the stop to the end of the segment. These spatial and temporal components structure the proposed transit performance metrics.

This "segments" approach is designed to produce transit performance metrics sensitive to granular aspects of trip provision—precisely those aspects essential to improving service. The goal of these measures, all of which in some way or another are tied to schedule adherence, is to help transit agencies harmonize the actual provision of transit service with the scheduled provision of those same services. It is also hoped that this clear-cut approach to generating transit performance metrics will democratize access to this information so that all stakeholders in transit policy—from local advocates to city planners to state agencies—will be able to meaningfully interrogate transit provision and therefore meaningfully contribute to its improvement.

2. Dedicated Short-Range Communication and Transit

Introduction

Planners in the late twentieth century imagined applying the burgeoning microcomputer and communication technologies to transportation. In the United States, this coupling became known as intelligent transportation systems (ITS), and its promise was vast. The federal government worked with industry to develop standards and allocate bandwidth to facilitate ITS technologies. One early outcome was the set of standards for dedicated short-range communication (DSRC) between vehicles and infrastructure, for which the United States Federal Communications Commission (FCC) dedicated 75 MHz of bandwidth in October 1999 (*Amendment of Parts 2 and 90 of the Commission's Rules to Allocate the 5.850-5.925 GHz Band to the Mobile Service for Dedicated Short Range Communications of Intelligent Transportation Services*, 1999; Federal Communication Commission, 2019).

DSRC was designed to facilitate a range of activities between vehicles themselves and between vehicles and infrastructure, from safety and emergency vehicle notification to automated tolling and enhanced navigation (Research and Innovative Technology Administration, 2009), all of great interest to transit agencies; however, despite both the potential of the standard and the government's and industry's promotional efforts, DSRC never became entrenched in vehicular operations. In October 2020, the FCC chairman noted this low uptake and announced his intention to phase out the DSRC standard, first by trimming the spectrum allocation by 60 percent and second by authorizing a new vehicular communication standard in the remaining 30 MHz. He argued that the transportation industry was not sufficiently utilizing the valuable bandwidth, which would be reallocated more productively to expand the Wi-Fi services' capacity ("FCC Chairman Signals an End to DSRC as He Prepares to Carve up 5.9GHz Spectrum – but ITS America Hits Back," 2020). In November 2020, the FCC voted to affirm the chairman's proposal, effectively ending the DSRC standard in its current form (Wiquist, 2020).

This decision came as a tremendous shock to the transportation industry, which had long invested in the DSRC standard—often with the direct financial support of the United States Department of Transportation. This decision also impacts the subset of the industry focused on providing public transportation.

It is reasonable to assume that, despite the sunset period allowed by the FCC, there will be few new implementations of DSRC on transit. Consequently, it is both possible and appropriate to fully review this two-decade chapter in applying this ITS technology to transit. The purpose of this review is to explore the proposed uses of DSRC technologies on transit vehicles as well as the actual implementations. This chronicle consolidates the thinking about and experience of DSRC technologies on transit to guide properties, policymakers, and vendors in future decisions about vehicular communication and public transportation.

DSRC Background

Dedicated Short-Range Communication (DSRC) refers to a set of standards designed for the unique needs of vehicular environments. The technology employing these standards are based on the IEEE 802.11 standard that undergirds Wi-Fi (Wu et al., 2013) and is housed in on-board units (OBUs), also known as on-board equipment (OBE), on vehicles themselves (and, more recently, also on devices carried by bicyclists and pedestrians) or at road-side units (RSUs) stationed along the transportation corridor. DSRC enables direct, secure, and high-speed communication between vehicles (V2V); between a vehicle and a fixed infrastructure (V2I), which is also referred to as vehicle-to-roadway (V2R) communication; and between vehicles and non-motorists, such as bicyclists and pedestrians (V2P) (Wu et al., 2013). More recently, DSRC has been presented as offering vehicle-to-everything (V2X) communication, without distinguishing between mobile and fixed transponders.

In 1999, the FCC allocated 75 MHz bandwidth (between 5.850 and 5.925 GHz) for DSRC use and split this allocation among seven channels of 10 MHz each (with the remaining 5MHz held in reserve). Six of these channels are designated as service channels (SCH), which deliver the general data, while the seventh is a control channel (CCH), which carries the high-priority or management data. Two pairs of channels can be combined to form higher-bandwidth 20 MHz channels. DSRC-enabled devices communicate wirelessly, typically within a range of 300 meters (but up to a full kilometer). The standard is designed for very fast communication and therefore has low latency. (Vehicles traveling up to 300 mph can use the standard.) DSRC communications, such as Wi-Fi, are not restricted to line-of-site and thus can function around corners and through inclement weather.

The DSRC standard includes a variety of message types. The basic purpose of the standard was V2V safety; consequently, the Basic Safety Message (BSM) allows mobile devices to anonymously broadcast their location, heading, and speed up to every tenth of a second to all transponders within range. Another common message type, Signal Phase and Timing (SPaT), allows RSUs to broadcast traffic signal status and anticipated change, by lane, to vehicles. This dynamic information is broadcast in conjunction with the MAP message that broadcasts static information on intersection geometry. While these three message types have tended to dominate DSRC considerations, particularly regarding transit, there is also a Personal Safety Message (PSM) broadcast by vulnerable users, such as pedestrians, to vehicles and a Traveler Information Message (TIM) through which RSUs broadcast roadway conditions and attributes, such as a speed warning for a dangerous curve. Less discussed, but highly relevant for transit management, is the Probe Vehicle Data (PVD) message type, which augments the BSM-type of location information with additional vehicle data, such as automated passenger counts, and aggregates these periodic "snapshots" into packets to relay to stationary roadside infrastructure (Perry, 2017).

All these message functions and transmission capabilities are included in a competing vehicular communication standard, C-V2X, that relies on cellular radio chipsets rather than Wi-Fi ones and has been more recently developed. C-V2X is not significantly better at meeting the DSRC goals but may be less costly to implement (given the ubiquity of the chipsets) and more able to multicast (and thus reducing communication traffic in a more connected future roadway environment). The standards are not interoperable given the different specifications (Gettman, 2020). While DSRC bandwidth will remain available for some time and legacy installations should remain functional (Gettman, 2020), it appears that C-V2X has won the standard battle in the vehicle communication space.

Proposed DSRC Applications to Transit

Given its strong concerns for vehicle safety and operations, the transit industry is a natural user for the types of services that DSRC enables. Furthermore, DSRC offers some unique benefits over competing technologies. For example, since transit agencies are often more able to fund capital rather than operational costs, the hardwired connections from the roadside infrastructure to transportation agencies common to DSRC (Johnson, 1998) are favored over reliance on data transmission via external vendors, such as cellphone companies (Regional Transit Authority of Southeast Michigan, 2018). Similarly, for transit agencies operating in dense, urban areas, DSRC's geolocation capabilities, which are based on internal reference points, are likely superior to those of satellite-based Global Positioning System (GPS) technologies (Alam, 2012; Gong, 2017; Liao & Davis, 2008).

DSRC technology's potential has led to many proposed transit applications. This section clusters these proposed applications into three categories: safety, operations, and information.

Table 1. Proposed and Implemented DSRC Technologies on Transit

Categories	**Proposed DSRC Use**	**Actual DSRC Implementations**	**Sources**
Safety	Situational Awareness	Ann Arbor, MI; Cleveland, OH	(Bu & Chan, 2005; Han Su et al., 2012; Turnbull et al., 2017; Valentine et al., 2019; R. Zimmer et al., 2017; R. E. Zimmer et al., 2014)
	Traffic Signal Preemption	--	(Choi et al., 2006; Johnson, 1998)
	Emergency Broadcast	--	(Johnson, 1998)
Operations	Driver Guidance	--	(Han Su et al., 2012; Johnson, 1998)
	Transit Signal Priority	Salt Lake City, UT; Pittsburgh, PA	(Federal Transit Administration et al., 2002; Johnson, 1998; Leonard, 2018; Leonard et al., 2019; Liao & Davis, 2008; Smith et al., 2018)
	Electronic Fare Collection	--	(Johnson, 1998)
Information	Traffic Conditions	--	(Maitipe & Hayee, 2010)
	Real-time information	Japan	(Ahmed et al., 2016; Hirai et al., 2006; Johnson, 1998)

Safety

Proposed safety applications sought to use DSRC to feed traffic information to bus operators (I2V), to allow buses to pre-empt traffic signals during emergency situations (V2I), and to offer a way for operators to signal for help from transit management personnel (V2I).

One safety vision of DSRC is to use the technology to consolidate right-of-way information for transit vehicle operators. (These are presented in terms of bus operations even though they extend to all types of transit vehicles.) The underlying idea is that DSRC can allow the pre-existing transportation agency's investment in road cameras and sensors to be leveraged to provide additional situational awareness to bus drivers. For example, surveillance video from local traffic cameras might be broadcast directly into transit vehicles (Han Su et al., 2012) or analyzed to warn of potential attempts by vehicles to turn in front of idling buses (Zimmer et al., 2014). Alternatively, traffic cameras or other roadside sensors might provide pedestrian or cyclist detection to buses (Bu & Chan, 2005; Turnbull et al., 2017; Zimmer et al., 2014). These technologies could also work in the other direction by sending bus operation information to surrounding users. For example, a bus making a turn might signal that information to surrounding infrastructure which might then warn proximate pedestrians and cyclists indirectly with an audio or visual cue or directly through DSRC chips on phones and bicycles (Turnbull et al., 2017).

A second DSRC safety vision is to allow bus operators to pre-empt traffic signals during emergency operations (Johnson, 1998). This approach would offer the same safety benefits for transit as envisioned for emergency vehicles (Choi et al., 2006). Bus emergency operations are, thankfully, relatively infrequent despite their deceptive prevalence in films (Chan, 1985; de Bont, 1994; Hamilton, 1973; Mann, 2015); nonetheless, automatically pre-empting traffic signals during such situations offers major safety benefits for all involved.

The last DSRC safety vision is to allow buses to signal the "transit management center that something abnormal is occurring" (Johnson, 1998). Such alerts might be initiated by the driver via a panic button or by the vehicle itself based on on-board diagnostics. The signal might further include data identifying the type of emergency (Johnson, 1998). This approach's success would require a sufficient network of RSUs to receive the signal for relay to the transit center.

Operations

Proposed operational improvements sought to use DSRC to guide vehicles (I2V), to adjust traffic signals to vehicles (V2I), and to facilitate fare payment (V2I).

One operational improvement is simply to provide information to bus drivers approaching intersections as to whether to continue through or to stop at a traffic light (Han Su et al., 2012). In this intervention, the bus transmits its position and velocity to the traffic signal, which then alerts the driver as to whether to decelerate or not. This guidance optimizes transit interactions with pre-established traffic phasing. Another proposal seeks to provide real-time information to drivers on their schedule adherence. In this model, the DSRC system could interact with roadside infrastructure to identify the variance from the posted stop arrival and departure times. The driver could simply receive guidance on whether to slow down or speed up (Johnson, 1998).

A more involved operational improvement is to have the traffic signal dynamically adjust its timing to smooth bus operations in non-emergency situations (Federal Transit Administration et al., 2002; Liao & Davis, 2008). This transit signal priority (TSP) might even reflect bus delay or passenger loads (Johnson, 1998). This approach would favor transit operations on a given corridor. (TSP to foster public transportation flows during normal situations is distinct from the safety option of traffic signal pre-emption during emergency situations presented earlier.)

The last proposed DSRC operational improvement for transit is to facilitate electronic fare collection. In this model, the vehicle might transmit electronic ticketing information to the transit agency to expedite boarding and alighting. The fare collection machines might also be updated via DSRC to display relevant information, such changing fare zones, to customers (Johnson, 1998).

Information

Proposed information improvements sought to use DSRC to warn vehicles of upcoming traffic conditions, such as work zones (Maitipe & Hayee, 2010) and to provide near real-time information to transit agencies and potential users on bus locations and speeds.

Several information improvements seek to provide guidance to bus drivers. For example, one proposal seeks to use DSRC-enabled vehicles (not just transit vehicles) as probe devices for understanding the congestion impacts of a work zone. In this model, a roadside unit (RSU) is temporarily installed at the construction site. The RSU gathers location and speed information from the on-coming vehicles to identify the start point for the congested section and then relays that information across a daisy chain of in-vehicle DSRC devices to warn traffic further upstream (Maitipe & Hayee, 2010). This intervention, depending on local rules for bus operations, might allow for effective route deviations and better management of passenger expectations by providing up-to-date information.

A set of proposals seeks to provide real-time service information to the transit agency and, in turn, to the public at large (Cho et al., 2000; Johnson, 1998). This approach allows transit operators to bypass the cost and uncertainty of using GPS as DSRC allows real-time communication between an OBU on the vehicle and appropriately placed RSUs (Cho et al., 2000). Transit agencies can use these data to generate performance metrics to assess their operations including "vehicle location, deviance from schedule, and passenger counting" (Johnson, 1998). The public at large can use these data to make personal scheduling decisions that build ridership and reduce car use

(Hirai et al., 2006). One proposed idea, not specific to transit, is to use the geolocation messages broadcast by OBUs via DSRC through single- and multi-hop relay to view vehicle locations on a map (Ahmed et al., 2016).

These proposed DSRC interventions are somewhat limited given the range of potential activities that DSRC facilitates; however, they far exceed the actual implementations.

Implemented DSRC Applications

In practice, there have been relatively few DSRC implementations on transit documented in the research literature and most of those have been tied to transit signal priority (TSP). This section presents the documented implementations discovered by the author. This discovery was based on discussions with key informants, published research literature, and internet searches. The implementations are organized along the same categories of safety, operations, and information as in the previous section. Implementations that combine more than one of these categories are presented in all relevant places.

Safety

Two safety implementations were identified, one in Ann Arbor, Michigan and one in Cleveland, Ohio.

In Ann Arbor, Michigan, the federal government funded a pilot project, the Transit Safety Retrofit Package (TRP), to test five DSRC collision avoidance systems on three buses from the University of Michigan's transit fleet. Three of these applications were generic vehicle safety applications not tailored to the transit environment. These include two V2V technologies: a forward collision warning to prevent striking the rear-end of a vehicle ahead in traffic and an alert when a nearby vehicle engages an emergency brake in a position that might be of danger to the bus. Both applications require the non-transit vehicle to also be equipped with DSRC technologies. The third safety application was a V2I curve speed warning in which roadside equipment at one notable curve would signal the bus if that vehicle's velocity exceeded a threshold considered safe. The remaining two applications addressed the unique needs of a transit environment. One was a V2I technology that alerted the bus driver while turning at designated intersections if passengers were in the intended path of the right or left turn. The other was a V2V technology that alerted the bus driver if a vehicle also equipped with DSRC was attempting to make a right turn in front of a bus departing from a stop (R. E. Zimmer et al., 2014).

The TRP pilot ran for roughly nine months throughout 2013 and early 2014. The data analysis found a high rate of false alerts for both transit specific applications although the DSRC system itself saw no problems in communicating information (R. E. Zimmer et al., 2014).

In Cleveland, Ohio, the federal government funded a pilot project, the Transit Bus Stop Pedestrian Warning (TSPW) Application, at four bus stops on the Greater Cleveland Regional Transit Authority (GCRTA) system. This project was part of a follow up to the earlier TRP project and was called E-TRP for enhanced TRP. System architecture documentation reports that this project was designed to use roadside equipment to detect pedestrians in the roadway at the four bus stops and relay that information via DSRC to on-board units on 80 to 100 buses "to improve the situational awareness and ultimately safety of pedestrians at transit stops" (R. Zimmer et al., 2017). The pilot that was installed appears to be slightly different. Three roadway sections (a signalized intersection, a non-signalized intersection, and a mid-block crossing) in Cleveland were fitted with pedestrian-sensing technology that alerted bus drivers on 24 vehicles equipped with DSRC on-board units when a pedestrian was in a bus's potential path. The pilot lasted for six months until August 2018 when it was dismantled. The evaluation showed that drivers increased their response to pedestrians in roadways and reduced their reaction times (Valentine et al., 2019).

Operations

Two operational implementations were identified for transit signal priority, one in Salt Lake City, Utah and one in Pittsburgh, Pennsylvania.

In Salt Lake City, Utah, the Utah Department of Transportation (UDOT) deployed a DSRC-activated transit signal priority to improve schedule reliability for the Utah Transit Authority's (UTA) 217 bus line along six-miles of Redwood Road just west of downtown Salt Lake City (Leonard, 2018). UDOT equipped 13 of the corridor's 17 traffic lights with TSP managed by a variant of the Multi-Modal Intelligent Traffic Signal System (MMITSS) that incorporated bus schedule adherence and occupancy criteria. Buses more than five minutes behind their published schedule were considered "behind schedule" while buses with occupancies of at least 20% (at least 9 riders) were considered "minimally occupied" for purposes of requesting signal priority. The "behind schedule" criterion was calculated using GPS-data transmitted by cellular service to the UTA operations software system. The "minimally occupied" criterion was calculated using optical sensors at the bus doorways. As the bus approaches a signalized intersection with TSP, the DSRC system automatically requests priority if both the schedule adherence and occupancy standards are met. The MMITSS-Utah system determines whether or not to alter the signal timing without offering any indication to the driver (Leonard et al., 2019).

UDOT conducted a four-month test between April and July 2018. A relatively small portion of the buses on the corridor (9% northbound in the morning peak and 14% southbound in the afternoon peak) were equipped with OBUs to transmit DSRC to the traffic signals. This allowed UDOT to compare schedule adherence between the two groups to find that, on average, the DSRC-equipped buses were on-time 91% of the time compared to 86% for buses without DSRC technology—considered a clear improvement in reliability. UDOT is installing the same

MMITSS-Utah TSP on a new bus rapid transit link between Provo and Orem, Utah (Leonard et al., 2019).

In Pittsburgh, Pennsylvania, researchers from Carnegie Mellon University worked with the Port Authority of Allegheny County (PATCO) to pilot one DSRC OBU to improve transit signal priority in Pittsburgh's East End along signals fitted with the SurTrac adaptive signal control system. The single bus was tracked along the corridor using the Basic Safety Message (BSM) location information. A future stage was planned to integrate the DSRC OBU with the bus's onboard computer to integrate schedule adherence and load, as in Utah, as well as bus door status into the TSP calculation (Smith et al., 2018).

In addition to these two documented implementations, other regions have explored TSP using DSRC. The Regional Transit Authority (RTA) of Southeast Michigan has explored building off the existing DSRC infrastructure (noted earlier) in Ann Arbor, Michigan to introduce a DSRC-based TSP near the University of Michigan (Regional Transit Authority of Southeast Michigan, 2018). This example demonstrates how DSRC implementations can build on each other. The City of Columbus, Ohio is rolling out a smart city program with support from federal and other sources and is planning to add DSRC-equipped buses to its bus rapid transit corridor that already has a TSP with a proprietary Wi-Fi system to promote interoperability (Bollo et al., 2018).

Information

Only one documented project implemented DSRC for information.

In Japan, as part of the Bus Transit Revitalization Project, the Ministry of Land, Infrastructure, and Transport developed a pilot bus location information system for express buses using DSRC. The buses identified their location using GPS antennae and transmitted that information via DSRC to RSUs placed every 5 to 10 kilometers along the expressway—an approach that eliminated the costs associated with cellular transmissions. The probe data collected location information every second or fifteen meters of travel and made it "possible to grasp traffic conditions in more detail than ever before" (Hirai et al., 2006), a key goal of the project. However, in addition to this general traffic tracking objective, the data were "also used by bus companies for their management of operations" and shared with the public for bus arrival information (Hirai et al., 2006). This project demonstrates the multiple benefiting audiences from transit probe data made available via DSRC.

Discussion

It appears that the attempts by the federal government to anoint DSRC as the standard for vehicular communication were somewhat thwarted in practice by wary transit properties and competing vendors. For an example of the former, King County Metro Transit, which serves the

Seattle region, deployed an ITS system along its high use transit corridors in 2006 using the 4.9 GHz Public Safety standard—not DSRC. The agency felt that the Public Safety standard, which emerged after the 9/11 terrorist events, matured more quickly regarding security and control technologies and intended to migrate to DSRC only as the technologies matured (Nace & Toone, 2011). For an example of the latter, UDOT planned to implement TSP across eleven miles in Salt Lake City using an array of different DSRC vendors; however, it became clear that unforeseen interoperability issues would restrict the study area to a six-mile section using DSRC equipment from a single vendor (Leonard et al., 2019).

The reduction of federal support for DSRC does not reflect any inherent deficiency in the technology, but rather the stronger competing desire for the available bandwidth for communication beyond that tied to vehicle movements. DSRC remains the vehicle communication standard of choice in Europe, where it is likely to see more development. The replacement standard in the United States has been the cellular vehicle-to-everything (C-V2X) approach, which replicates the main DSRC capabilities while surmounting its limited range via additional connectivity to the cellular network. This additional feature does allow for more communication options (such as moving less time-sensitive data to the slower cellular network or providing connectivity to infrastructure in less dense environments). Those options will only increase with the development of increasingly fast cellular technologies, such as 5G, and the expanding coverage of those services.

Conclusions

While this paper focuses on the limited embrace of DSRC within the transit world, that resistance might be contextualized within the larger transportation environment. As a networked good, DSRC required a public uptake that it was unable to achieve. Transit agencies are wary of adopting a technology that does not demonstrate its lasting power. Despite many good faith efforts and federally sponsored pilot implementations, the proverbial "killer app" that would make DSRC indispensable for transit never materialized. It would also appear that the suite of benefits offered by DSRC, such as low latency communication, were less relevant to the transit realm that typically operates in lower speed, lower risk roadway environments. Without the need for ultra-speedy communication, many of the benefits of DSRC could be sufficiently met by other, more familiar technologies, such as cellular communications. Finally, the internal incompatibilities among DSRC vendors proved harmful to the success of installations at transit properties. The experience in Utah, in which, despite having broadly purchased equipment with the expectation of interoperability, the state had to contract its implementation to that stretch of the road that could be covered with purchases from a single vendor, demonstrates a failure of industry players to work together to foster adoption of a new technology. Those companies likely assumed that with federal support the standard would be permanently in place and their role was to maximize market share at the outset. It appears as if such individually maximizing but collectively undermining behaviors hampered the development of the type of off-the-shelf systems most appealing for transit agencies. In short, DSRC failed to uncover its business case for transit in time to stave off its demise.

It may also be the case that the DSRC standard was simply too much technology too early. There is a case to be made for simpler technologies that can be adopted one piece at a time and not all at once. The irony perhaps is that vehicles are currently smarter than ever but largely use their processing power to react to the actions of surrounding infrastructure and vehicles rather than communicate directly with them.

3. Imputing Stop Arrival and Departure Times from GTFS Realtime Data

Introduction

Emerging transit agency practice complements the public provision of static schedule information in the General Transit Feed Specification (GTFS) format with dynamic updates on vehicle locations in the GTFS-Realtime (GTFS-RT) extension. While the goal of this real-time extension is to improve trip planning by adjusting transit schedules for actual road conditions, the availability of standardized vehicle position data offers an exciting off-label use to generate a suite of performance metrics based simply on the differences between scheduled and actual stop arrival and departure times.

In theory, the knowledge of the actual times when transit vehicles arrive and depart from stops (along with the scheduled times and spatial paths) enable the measuring and mapping of an array of metrics tied to schedule deviation—from bus bunching to departure delays—by anyone, not just transit agency staff with internal access to automatic vehicle location (AVL) feeds. A metropolitan planning organization might examine schedule adherence to evaluate future busways or transit signal prioritization programs. A state department of transportation might track performance across the state to evaluate operational subsidy policies. A local community-based organization might explore service equity on certain lines to structure an advocacy campaign. Furthermore, transit agencies themselves may want equipment-agnostic access to consistent metrics rather than being locked into a bespoke solution tied to the feeds from a specific AVL vendor.

In practice, however, there is a significant barrier to directly deriving performance metrics from GTFS-RT feeds: the GTFS-RT standard currently allows for but does not require the explicit coding of all stop arrival and departure times. These timepoints, critical to the consistent calculation of performance metrics, are unfortunately not essential for the GTFS-RT purpose of dynamic trip planning and are therefore inconsistently reported. For example, among the five California systems in this analysis, four provide only arrival time information in their GTFS feeds, and the one system that provides both arrival and departure time information does not do so for the terminal stop—a key timepoint for understanding route-level performance. This research overcomes this barrier by proposing and testing a robust method to impute transit stop arrival and departure times from GTFS-RT feeds.

Such imputation is not straightforward. GTFS feeds (both schedule and real time) come in many flavors with distinct qualities. At times, the standard feels more like a suggestion. Helpful data fields present in one feed may not be present in another. Geocoding conventions closely adhered to in one feed may be more casually implemented in another. These variations may have little impact on the intended use of GTFS-RT for trip planning but can confound algorithms to consistently identify stop times for the off-label use of performance measurement.

This research embraces this variation by exploring GTFS feeds from five California transit systems to develop a robust method to impute stop arrival and departure times. Specifically, this method first identifies unique route variants from GTFS Schedule data and then cuts those variants into consistent segments between consecutive stops. The vehicle location data from the GTFS-RT *VehiclePosition* message set is then matched to each segment by day and transit trip to estimate stop arrival and departure times. These imputed times are then compared to the more limited arrival and departure time data available in the GTFS-RT *TripUpdate* message set to assess the accuracy of the approach. While it is hoped that the need for such imputation is relatively short-lived as transit agencies implement more fulsome GTFS-RT feeds, this method provides a self-contained means for estimating the actual stop arrival and departure times necessary for generating consistent and comparable transit performance metrics.

Methodology

This study explores five full days of GTFS-RT *VehiclePosition* data from five different California transit systems to generate and test a protocol for imputing bus arrival and departure times. The imputed times are tested against the limited stop time data available from the GTFS-RT *TripUpdate* data using several assessment metrics. An innovation of this evaluation approach is that it relies entirely on GTFS data that are available to the public—and no other data source.

The imputation requires two key steps. In the first step, the scheduled routes are cut into segments between each successive stop. In the second step, the vehicle location data are applied to each segment and combined with the derived vehicle speed data associated with the stop locations to impute arrival and departure times.

To assess the validity of the imputation method, the estimated values are compared to the stop arrival and departure times available in the *TripUpdate* message sets of Modesto Area Express (MAX), the only one of the five studied systems that includes both stop arrival and departure times in its message sets. The distribution of these reported times is compared to the imputed ones from this research. Since one goal of this work is to link specific arrival and departure times to measured dwell times, those values from MAX are also compared to the imputed values.

Imputation is an imperfect craft. The approach used in this research to forecast arrival times and backcast departure times occasionally results in the impossible situation in which the departure occurs before the arrival. Mathematically subtracting the latter from the former yields a negative dwell time. As part of the quality control work, the number, magnitude, and distribution of these events is also explored.

Data Collection

The primary data for this research include five full days (March 1-5, 2022) of GTFS-RT *VehiclePosition* and *TripUpdate* feeds from five California transit agencies: Modesto Area Express (MAX), Alameda-Contra Costa Transit District (AC Transit), Monterey-Salinas Transit (MST), Big Blue Bus (BBB), and Orange County Transportation Authority (OCTA). (These five-day data sets include all trips originating during the study period, even if they continued into the following day. Therefore, some trips that began on February 28 and extended into March 1 were excluded, while other trips that began on March 5 and extended into March 6 were included.) Finally, the contemporaneous static GTFS data for the five systems were gathered to provide route and scheduling information. Figure 1 presents the route footprint of studied transit agencies as well as their relative locations within California. (In March 2023, a year after these data were collected, MAX, which was rebranded as StanRTA, reorganized its bus lines. Therefore, the current routes do not match those in Figure 1. This research uses the MAX name and the associated data to reflect conditions when the data were collected.)

Figure 1. Map of Studied Transit Agencies

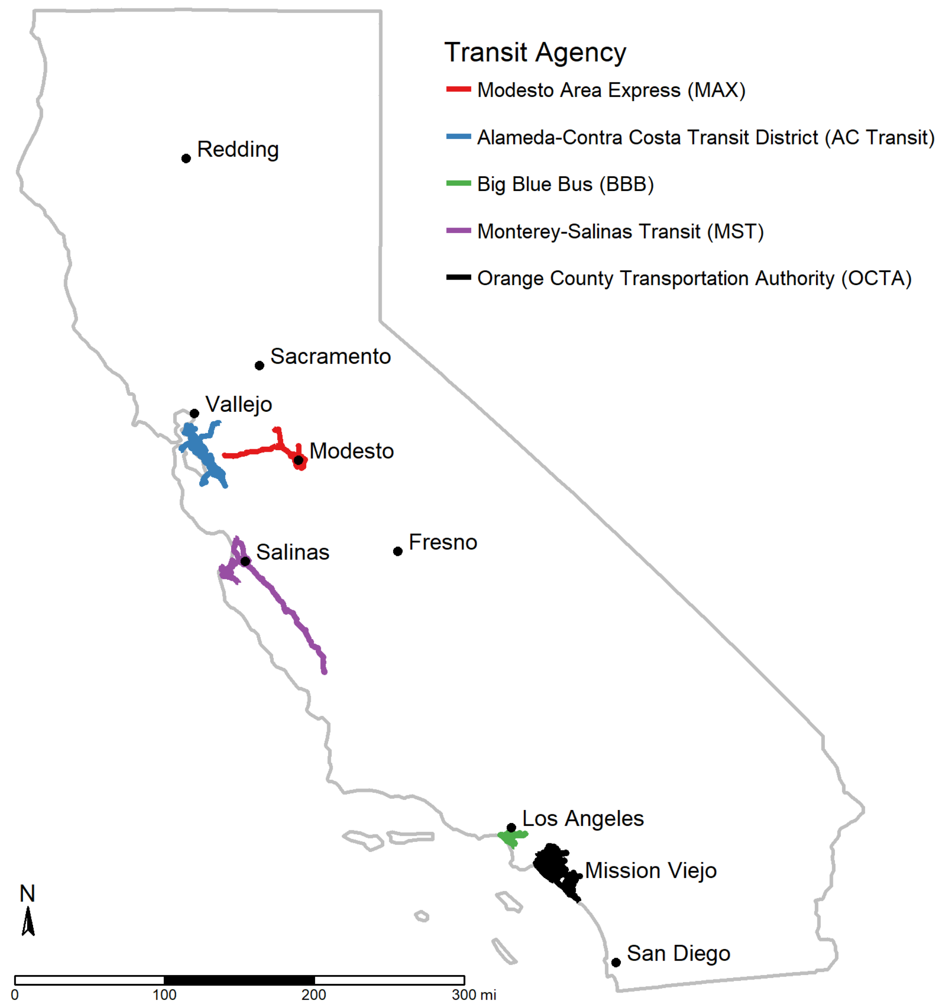

These agencies were chosen to represent a range of different operations in terms of size and location across the state. Furthermore, these agencies have different GTFS-RT vendors, resulting in different feed designs. In practice, while all the systems' feeds informed the analysis, the bulk of the comparison focuses on MAX.

Bus Route Segmentation

To expedite computing time and yield output data for detailed sub-route analysis, this research first parsed the GTFS static information to identify unique "variants" for each route and component segments for each variant. In this nomenclature, a "variant" refers to a set of trips on a given route that have the same set of stops in the same order and follow an identical path as defined by the associated vertices in the GTFS Schedule data. A "segment" refers to the portion of that path that connects two consecutive transit stops.

Variant Identification

Variant identification is a two-part process. One part identifies unique paths from the vertices provided in the *shapes.txt* file of the GTFS Static data. The other part identifies unique sets of ordered stops. A single variant is then defined as the unique combination of a unique path and a unique set of ordered stops for a unique route.

GTFS Schedule data associates an ordered set of vertices, defined in the shapes file, with each defined *trip_id*. The polyline connecting these vertices represents the path of the transit route for a given trip, and the order of the vertices defines the direction of that travel. This directionality component is an important part of the GTFS system. A shuttle service that goes back and forth along the exact same alignment represented by the exact same vertices would still be associated with two distinct paths: one path with the vertices ordered in the outbound direction and a second path with the same exact vertices ordered in the inbound (and therefore opposite) direction.

Most transit routes are associated with more than one unique path. (The few routes that are associated with a single path are typically loop routes in which transit vehicles always follow the same course in the same direction.) Most commonly, routes have two unique paths—an outbound one and an inbound one—but, not infrequently, routes may be associated with more than two paths. For example, a bus route might make a detour during a limited time period to serve a significant destination with peaking characteristics, such as a major employer on a shift schedule or a community center. An example of the latter type of detour can be seen clearly in Figure 2, which maps the four paths associated with MAX Route 29. This route has distinct inbound and outbound patterns which offers some additional coverage in the corridor's southern neighborhoods. Twice a day in each direction, the bus line makes a detour to the east to serve Howard Prep, a non-profit center that provides daytime opportunities for disabled adults.

Figure 2. Different Paths Associated with MAX Route 29

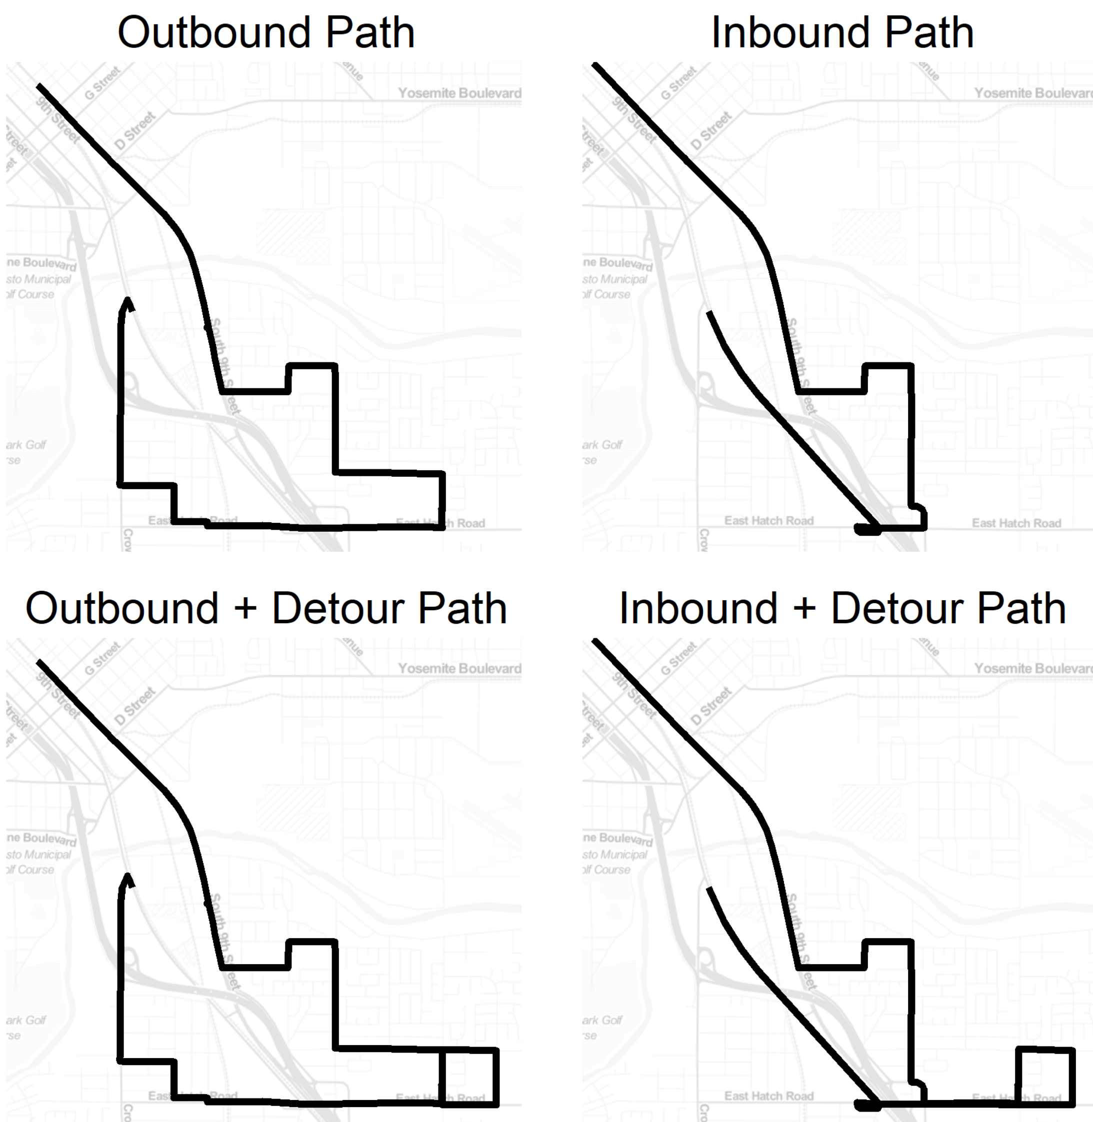

Surprisingly, GTFS Schedule data will sometimes include two or more differently labeled shapes that are defined by the same ordered set of vertices. This redundancy has no impact on the intended use of GTFS Schedule for trip planning purposes, but it does mean that unique shape names cannot be used to identify unique paths—since two or more identical sets of ordered vertices could have different labels. To avoid any possibility of mislabeling two identical paths as two unique paths, this research instead examined the location of each vertex for each shape to identify unique permutations. All shapes sharing the same permutation of vertices were considered as sharing the same unique path for the purposes of defining variants.

Across the five transit systems studied, 713 unique paths were identified. Roughly a sixth of these were defined by more than one shape, as shown in Figure 3. Several unique paths were associated with three, four, or even five shapes within the GTFS Schedule data. This finding underscores the importance of determining paths by their vertices and not their labels. This finding also suggests an area for improved data cleaning processes as GTFS files are updated and altered over time.

Figure 3. Shares of Shapes Associated with a Single Path Across Studied Systems

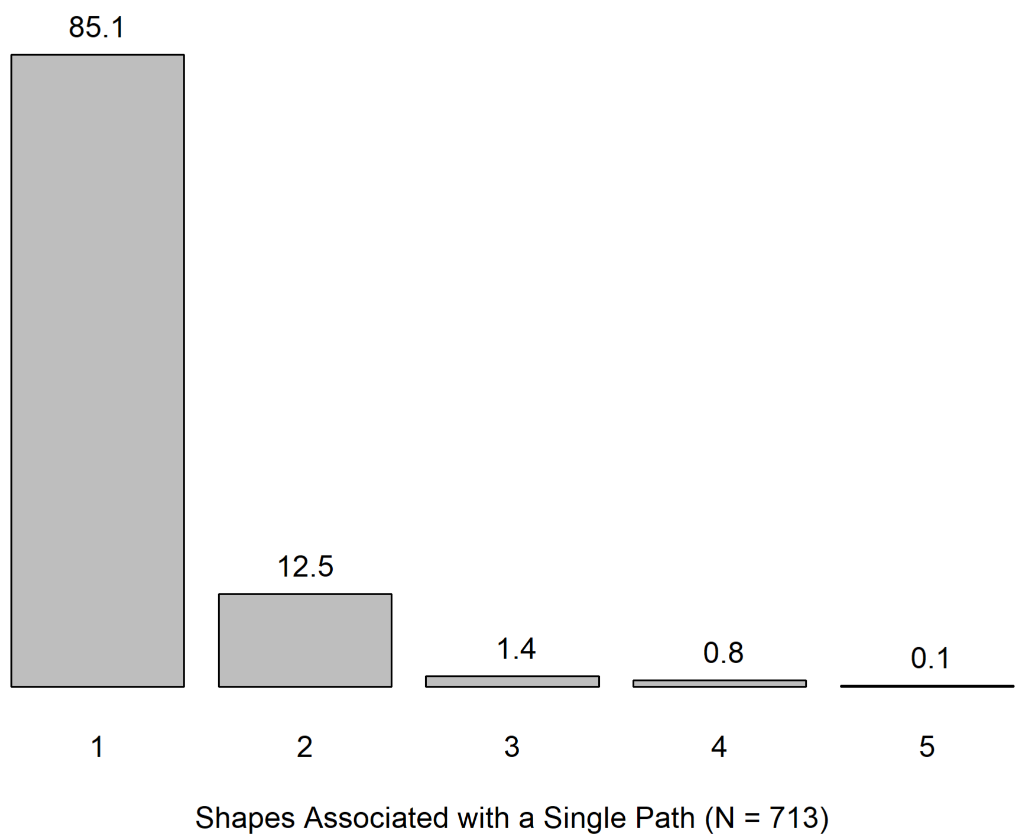

GTFS Schedule data also associate an ordered set of stops with each defined *trip_id*. The GTFS standard defines stops and provides the sequence in which they occur. The sequencing is determined by an integer aligned with each *stop_id*. Those integers are sorted in ascending order to determine the stop sequence. While it is tempting to define a unique order of stops based on a unique set of *stop_ids* and a unique set of ranked integers, that approach does not yield an appropriate solution since any set of integers that retain the same ranking would describe the same stop order. Furthermore, a review of the GTFS data found several cases of identical stop orders defined by slightly different sets of sequencing integers. The differences are usually for just a single stop and do not affect the overall organization of the stops but make directly using the sequencing number as part of the unique variant identifier problematic. Such use would erroneously yield additional variants.

To solve this problem, this research represents a unique set of ordered stops by simply creating an array of *stop_id*s in the appropriate order. This approach yields a parsimonious sorting variable unaffected by inconsistencies in the sequencing numbers that do not affect stop order. Furthermore, it reduces the amount of data that needs to be managed in defining a unique set of stops since only the stop values are stored (as their order reflects the stop sequence).

A variant is defined as a unique combination of path and stop structure. There are cases in which several stop structures share the same path—for example if a route offers local and express service along the same corridor. Such cases are defined as different variants since their stop structures are distinct. There are also cases in which several different paths share the same stop structure, typically for minimal differences in the path geometry. Some of these instances reflect legitimate variation, such as alternate paths to serve the same stops around the Alameda Point and Ferry Terminal at different times of day shown in Figure 4. Interestingly, the AC Transit brochure for this route does not show a specific path around the port. The map demonstrates the stops and notes “Stop in Alameda Point, showing the direction of the bus. (Exact routing in Alameda Point not shown.)” (AC Transit, 2022).

Figure 4. AC Transit Route 78 – Alternative Paths to the Same Stops

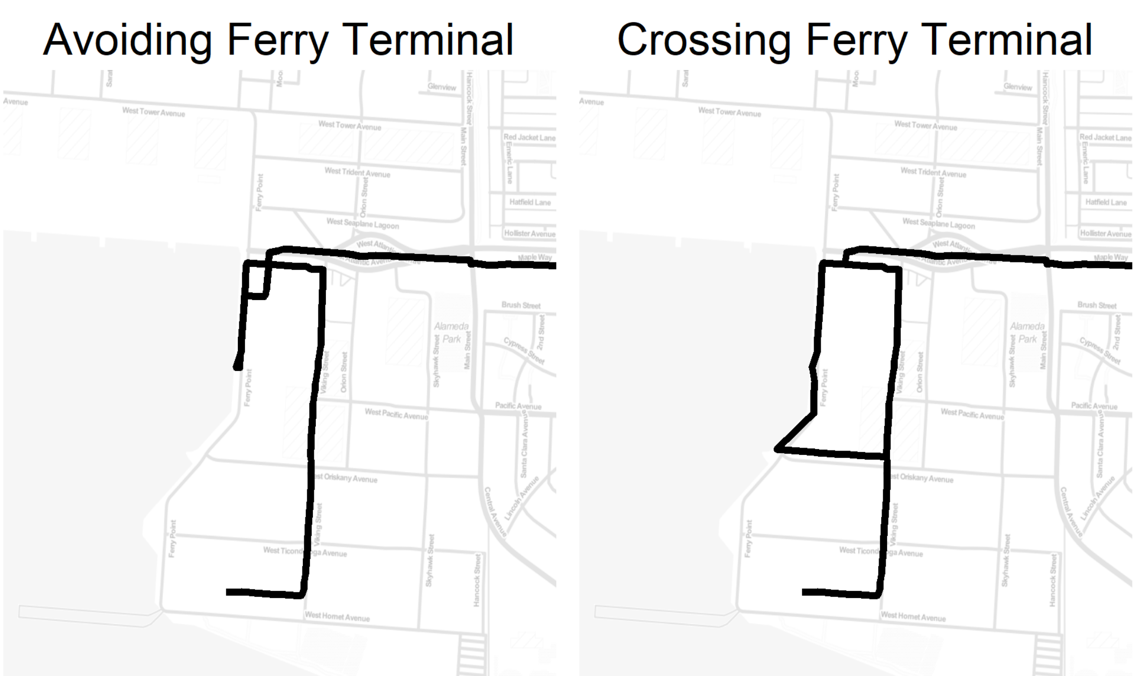

Some of these instances reflect illegitimate variation, here defined as two shapes whose vertices code slightly different paths—even though those differences have no bearing on the course of the transit vehicle. Once again these are cases with two shapes within the GTFS Schedule data meant to code the same path, but in this case the shapes are not identical. Despite this distinction, the prescription for transit agencies to tidy up the shape files to avoid these near duplications is recommended. To be true to the data provided to the public and to ensure the smooth functioning of all GTFS processes, this research uses the paths provided in the GTFS Schedule data, which slightly (and artificially) increases the number of variants.

Finally, there are some variants that are identically coded in terms of paths and stops but are tied to distinct routes and therefore qualify as separate variants. This feature is particularly prevalent in the Big Blue Bus data.

Table 2. Route Variants by Transit Agency

Agency	Routes	Variants	Variants to Routes	Routes with a Given Number of Variants							
				1	2	3	4	5	6	7	8
MAX	22	44	2.0	3	17	1	1	0	0	0	0
AC Transit	129	333	2.6	5	81	20	14	4	5	0	0
BBB	39	138	3.5	0	20	2	7	0	6	4	0
OCTA	58	174	3.0	2	33	6	9	3	3	0	1
MST	28	99	3.5	0	11	7	7	0	1	0	1
Totals	**276**	**788**	**2.9**	**10**	**162**	**36**	**38**	**7**	**15**	**4**	**2**
Shares (%)				**3.6**	**58.7**	**13.0**	**13.8**	**2.5**	**5.4**	**1.4**	**0.7**

Note: OCTA also has one route with 11 variants and MST has one route with 14 variants.

On a systemwide basis, among the five agencies explored here, the ratio of variants to routes ranged from 2.0 to 3.5 with an average of 2.9 variants per route. The mode number of variants was 2, which accounted for almost three-fifths of the routes considered. This finding is very reasonable since most routes are coded with a single outbound variant and a single inbound one. What is perhaps more interesting is the high number of routes that have more than 2 variants: 13 percent of routes had 3 variants, almost 14 percent of routes had 4 variants, and more than a twentieth of all routes had 6 variants. This high number of variants might reflect scheduling nuances that transit agencies undertake to tailor their routes for different purposes and times and might also indicate the accumulation of GTFS Schedule coding errors that have artificially inflated the number of variants. A prudent policy for a transit agency would be to review routes with high numbers of variants to ensure they are due to the former and not the latter explanation. Since an outside user can only rely on the data provided, each variant needs to be considered separately to accurately assess the arrival and departure times.

The special role of variants as a defining unit of transit performance analysis offers some additional ways to consider differences in operating characteristics. Table 3 presents route-by-route variant information for those MAX lines numbered in the twenties, a sample of the full network. These data show routes with one, two, and four variants. Both the loops of Route 21 and Route 26 have only a single variant. Routes 22, 23, 24, 25, and 28, all traditional linear routes, have two variants. Route 29, whose paths were shown earlier, has four variants to accommodate an additional loop (in both directions) to serve the Howard Prep center.

Table 3. Sample of Variant Information for MAX Routes

Route	Variant	Stops	Weekly Frequency				
			Weekday	Sat	Sun	Total	Share
Route 21	A	30	280	23	11	314	1.00
Route 22	A	35	150	24	11	185	0.51
	B	34	145	24	11	180	0.49
Route 23	A	33	140	21	10	171	0.50
	B	25	140	21	10	171	0.50
Route 24	A	27	65	0	0	65	0.50
	B	25	65	0	0	65	0.50
Route 25	A	71	160	13	11	184	0.51
	B	73	150	13	11	174	0.49
Route 26	A	28	150	25	10	185	1.00
Route 28	A	21	65	0	0	65	0.50
	B	11	65	0	0	65	0.50
Route 29	A	24	145	25	11	181	0.47
	B	19	145	25	11	181	0.47
	C	27	10	0	0	10	0.03
	D	24	10	0	0	10	0.03

The data in Table 3 also demonstrate the surprising differences in the number of stops on the outbound and inbound portions of the five routes with two variants. Three of these routes see variations of two stops or less between the outbound and the inbound variant; however, Route 23 has an eight-stop difference and Route 28 has a ten-stop difference.

The data in Table 3 also demonstrate that variants of the same route may have different trip frequencies during the week. This might occur if the route network is structured to facilitate interlining within blocks assigned to drivers. Route 22 has a variant that makes an additional trip every weekday, and Route 25 has a variant that makes two additional trips every weekday. Among the sample of MAX variants shown in Table 3, there are no such distinctions in the number of trips made on Saturdays or Sundays.

Route Segmentation

Once the variants are identified, they are broken down into segments bound by consecutive stops. While this process is straightforward in theory, in practice it is somewhat more involved to accommodate the nuances and issues with the GTFS Schedule structure. This section describes the computational process of subdividing variants into segments, which is accomplished through a spatial database, in this case PostGIS, to incorporate each segment's geometries and attributes.

The first step is to simply create a table with rows for each segment. Since the stops bound the segments, there is always one less segment than there are stops in the variant. The segments and, therefore, the rows are defined by consecutive pairs of stops with the first stop of any pair designated the segment's start and the second stop designated the segment's end. These rows are structured in the stop order so that the first stop of the variant is also the start stop of the first row and the second stop of the variant is both the end stop of the first row and the start stop of the second row and so on until the penultimate stop is the start stop of the final row and the last stop is the end stop of the final row.

The second step is to use the end stop locations to progressively excise segments from the variant path and then append those geometries to the appropriate row in the segment table. The primary challenge to accomplishing this segmentation is placing the end stops at the correct location on the variant path. Several common GTFS coding practices complicate this placement. First, there is a tendency to code route shapes along road centerlines. This practice often results in shapes in which the outbound and inbound portion of a lasso-shaped route are coded along the same exact path. The common geographic information systems (GIS) technique of snapping a point to the closest place on a line does not guarantee that an early outbound stop will not be coded to the late inbound portion of the path. While this problem could be addressed with more careful shape coding, there are many situations involving looping portions of transit routes where the path legitimately overlaps itself—raising the same troubling issues for route segmentation. Second, sometimes legitimately overlapping portions of a route are not coded directly on top of each other but slightly separated. In this case, simply snapping the stop to the path will naturally go to the closest point which might be the second time that the path passes the stop and not the first time. Finally, while the GTFS guidelines promote the addition of a vertex along a route shape near each stop location, these vertices are not always present. Therefore, algorithms that are based on finding the nearest vertex to place the end stop on the path—a reasonable approach based on GTFS guidance—will occasionally miscode. This issue is particularly problematic on long straight paths for which vertices are quite distant from each other. Searching for the nearest vertex can place the end stop cut point on the path rather distant from the actual end stop's location.

After encountering each of these problems, this research embraced a sequential algorithm to place the end stop locations on the variant path and extract the segment shapes. The algorithm loops over each row in the segment table and first places a vertex onto the path at the nearest point to the end stop of that row. This placement ensures that there is a vertex near the stop. Second, a line is drawn from the stop through the newly placed vertex. This line extends beyond that vertex for a length equal to half the distance between the stop and the vertex. Additional vertices are added to any part of the path that is crossed by this line. All the vertices that define the variant path, including any newly added vertices, are renumbered to maintain the correct vertex order to define the path line. The algorithm places the end stop at the lowest numbered of these vertices. This second step increases the likelihood that even in areas with poorly coded parallel or overlapping paths, the stop location is placed on the earliest instance. Then the path preceding that end stop is trimmed off and its geometry added to the segment row. This geometry defines the shape of the segment.

Once a segment is excised, its shape is dropped from further segment generation. This procedure ensures that, for routes that contain overlapping or parallel segments, an end stop that occurs the second time a vehicle passes a given stop point is assigned to the correct section of the path. This process continues until the penultimate row in the segment table has been allocated its segment, at which point the variant's remaining path is allocated as the final segment.

This method addressed the problems encountered in the five systems studied to correctly segment the path of each variant; however, this does not leave the algorithm immune to potential challenges. For example, this approach could be confused if a very poorly coded shape had a long straight overlapping section where the outbound path was incorrectly coded on the inbound side of the street and the inbound path was incorrectly coded on the outbound side of the street and the distance between the two paths was more than half the distance between the stop and the outbound path. In such a theoretical case, the algorithm would place the vertex for the outbound stop at the inbound path and not find an earlier vertex in the area to correct this mistake. This misplacement would remove a large section of the path's geometry for subsequent segmentation. Such a situation would reflect several coding errors in the initial GTFS Static feed, which, in practice, was never witnessed. Another potential problem with this algorithm could occur if a transit route needed to go beyond a stop on the far side of the road to make a U-turn to come back on the near side of the road where the stop was located. This algorithm might incorrectly place the end stop prior to the U-turn and add length to the following segment. To avoid such inadvertent miscodings, the research algorithm identifies segments for which more than one vertex was added to facilitate manual checking.

Stop Time Imputation

The imputation algorithm places the timestamped and geocoded *VehiclePosition* pings from the GTFS-RT feed along these segments to estimate vehicle arrival and departure times for each stop. This estimation takes several steps. First, vehicular speeds near the segment's beginning and end are calculated. Second, pings closest to the segment's start and also its end (but beyond a buffer around the stops) are identified. Finally, the estimated speeds (including an accelerating or decelerating factor) are applied to the distance from the scheduled stop to the identified ping to determine likely stop departure and arrival times.

Ping-Segment Placement

The pings that represent the location of vehicles in time and space within the GTFS-RT feed were joined to the appropriate variant's appropriate segment for the given trip on the given day using the *start_date* and *trip_id* from the GTFS Static data. This step is done sequentially segment-by-segment to avoid allocating a ping to the incorrect segment, a particular concern for lasso-shaped routes. To accomplish allocation, the pings are ordered by their timestamp (which has been put into local time from Universal Metric Time), and then their share of distance along each route was calculated.

The closest distance between these pings and the path that defines the segment are calculated. If this distance exceeds 10 meters, the associated pings are removed to ensure that only those pings that could be reasonably considered on the segment are included in future calculations. This removal reduces the understanding of potential route deviations but also reduces the potential for the stop imputation algorithm to be confused by such deviations.

Near-Stop Speed Calculation

The stop time imputation relies on an estimate of traffic speeds near the scheduled stops that define the beginning and end of each segment.

The "near" criterion is important as road conditions on bus routes can vary as stop spacing increases. This feature is particularly pronounced on the long segments that characterize commuter routes that run on limited-access highways. Such routes report average segment speeds that far exceed the "near-stop" speeds exhibited when the vehicle exits a highway to board and alight passengers. It is critical that the traffic speeds used to impute stop times be representative of the portions of the segment to which they are applied. This research defines the "near-stop" regions as the first and last half-kilometer of any given segment (which often overlap). This arbitrary threshold was selected to provide consistent estimates regardless of segment length or bus service type. In practice, this approach often ignores speeds along substantial portions of large segment routes (while incorporating all speeds on small segment routes) to achieve more accurate imputations. Dropping data can feel strange since comprehensiveness is intuitively associated with accuracy; however, the goal of data analysis is to pick out the signal that this “near-stop” approach accomplishes.

Identifying “traffic speeds” rather than posted speeds is the second critical criterion. This work initially explored using posted traffic speeds to facilitate the stop time imputation; however, posted speeds are, for most transit environments, much faster than actual vehicular flow which is impeded by traffic control devices (as well as the impact of other vehicles). While in theory it might be possible to apply a reduction factor to the posted speeds to reflect average speeds, such an approach would remove the speed variation that naturally occurs given changing road conditions. The same speed would be applied during the peak period in a thunderstorm as in the middle of a clear night removing the nuance that makes probe-based performance measure meaningful. Finally, use of posted speeds forces users to bring in an additional data set outside the GTFS products and to introduce an additional algorithm to join the posted speeds to the route segments—all for an inferior product. A more elegant and accurate approach is to use the GTFS products themselves to determine the near-stop speeds.

This approach generates a unique traffic speed for each end of each segment for each trip for each day. The departing speed (i.e., the speed at the beginning of the segment) is calculated by taking an array of ping times and distances from the segment’s origin to create a set of time-distant pairs one less than the length of the array. For each of these pairs, the difference in time and distance are calculated to generate a pair specific speed. For the departing speed, the inclusion of time-distance pairs is constrained to those whose initial distance is less than 500 meters from the segment’s beginning (even if the final distance is more than 500 meters from the segment’s beginning). The departing speed is then calculated as a weighted average of these time-distance pair speeds with the weight based on the share of the total distance traveled attributable to each time-distance pair. The weighting approach is designed to minimize the impact of pings taken while the vehicle is stopped—particularly while allowing passengers to board or alight at the segment’s origin. A consequence of this approach is to slightly increase estimated speeds by eliminating time-distance pairs in which the distance is zero, such as sitting at a particularly long traffic light. The arriving speed (i.e., the speed at the end of the segment) is calculated the same way but from the end of the segment—so, for example, it includes all time-distance pairs whose ends are within 500 meters of the end of the segment (even if the initial distance is more than 500 meters from the segment’s end).

To ensure quality, only pings that are within 10 meters of the segment are considered, and the total distance for which consecutive pings are available must exceed 100 meters. The first condition excludes deviations from the route, and the second condition ensures sufficient data to reflect vehicular speeds along the segment. Both thresholds are arbitrary. Since departing or arriving speeds are not calculated when the relevant distance traveled as measured by pings is less than 100 meters, a trip average speed is instead assigned in its place. This trip average is calculated as the arithmetic mean of all departing and arriving speeds that do meet the criteria for inclusion for a given trip on a given day. This workaround is not sensitive to the segment's conditions but incorporates information from the other segments on that given trip on that day.

Stop Buffering

The stop time imputation approach also involves determining a zone in which the stop is likely to have occurred.

Two buffer points are identified along each segment—one a set distance downstream from the origin stop (the segment's start) and the other the same set distance upstream from the destination stop (the segment's end). The space between the scheduled stop locations and the associated buffer points determines the portion of the segment where the bus is reasonably expected to physically board and alight passengers—the stop zone. This buffer is designed to accommodate the reality that buses do not always perfectly stop at the designated stop location (as well as the fact that these locations are not always perfectly geocoded). A bus stop that serves multiple bus lines or is prone to bus bunching will often result in multiple vehicles servicing the same stop at the same time and therefore parking in tandem. This reality expands the reasonable area that a bus is likely to be considered "at a stop."

A challenge is determining the smallest buffer that reasonably circumscribes the area within which buses are likely to stop. This research places that buffer 30 meters (roughly 100 feet) from either side of the geocoded stop location. Such liberal spacing allows for three articulated buses of roughly 60 feet each to all be at a given bus stop.

Departure/Arrival Time Estimation

The last step of the imputation is to identify a ping outside the stop buffer but within 500 meters from the relevant segment end and apply the near-stop speed (while accounting for typical bus acceleration/deceleration of 1.7) to determine the departure time from the beginning stop and the arrival time at the end stop.

The algorithm looks for the ping closest to the stop zones to minimize the potential impact of traffic controls on the stop times' estimate. As noted earlier, a consequence of this methodology is that time spent entirely still by buses is not counted in the near-stop speed calculation; therefore, minimizing such possibilities is preferred for accurate imputations. Most bus routes are designed not to encounter a need to stop at a traffic control device immediately after the scheduled stop (although many scheduled stops are themselves at traffic control devices). If no ping is found within 500 meters of the scheduled stop but outside the stop zone, the associated stop time for that end of that segment cannot be computed and is left empty.

To impute the departing stop time, the distance from that stop to the selected ping along the path of the segment is divided by the speed and adjusted for acceleration from still to the steady state speed. That calculation yields the number of seconds that is subtracted from the ping's timestamp to determine the departing stop time. To impute the arrival stop time, the same calculation occurs (although the adjustment here is for deceleration from the steady state speed to still), and the resulting number of seconds is added to the ping's timestamp to determine the arrival stop time. To calculate the dwell time at the stop, the arrival stop time from the previous segment is subtracted from the departure stop time of the subsequent segment.

Comparing Imputed and Reported Data

To assess the success of the imputation, those data are compared to the times reported in the *TripUpdate* feed for MAX. Among the GTFS files collected for this report, only MAX provided departure and arrival time data for each stop, as this is an optional component of the GTFS standard. This comparison first examines the arrival and departure times independently before exploring the related pairs that constitute dwell time.

Stop Arrival and Departure Times

The cumulative density plots shown in Figure 5 show that the imputation method appears to predict both the arrival and departure times of bus stop events slightly later than the information provided within the *TripUpdate* messages from the MAX GTFS-RT feed. This feature is evident with both curves primarily located to the right of the origin.

The median arrival time difference is 9.0 seconds while the median departure time difference is 13.7 seconds. The arrival time differences are slightly more varied, as seen in a less vertical curve, with roughly a fifth of these values negative, i.e., the imputed stop time precedes the reported stop time. By contrast, the departure time differences are far more consistent, as seen in the almost vertical nature of most of the curve, with very few negative values.

These findings suggest that the algorithm presented here works similarly to the algorithm used by MAX's vendor. The typical difference for both arrival and departure times is less than 14 seconds, and the offset from the reported values is largely positive for both curves. The consistency of this offset is important since a critical objective of the imputation algorithm is to provide stop dwell times; however, just because median arrival and departure time differences vary by roughly five seconds does not mean that the specific dwell times do as well. Therefore, an additional analysis is needed to compare the imputed and reported dwell times.

Figure 5. Comparing Imputed and Reported Stop Times on MAX

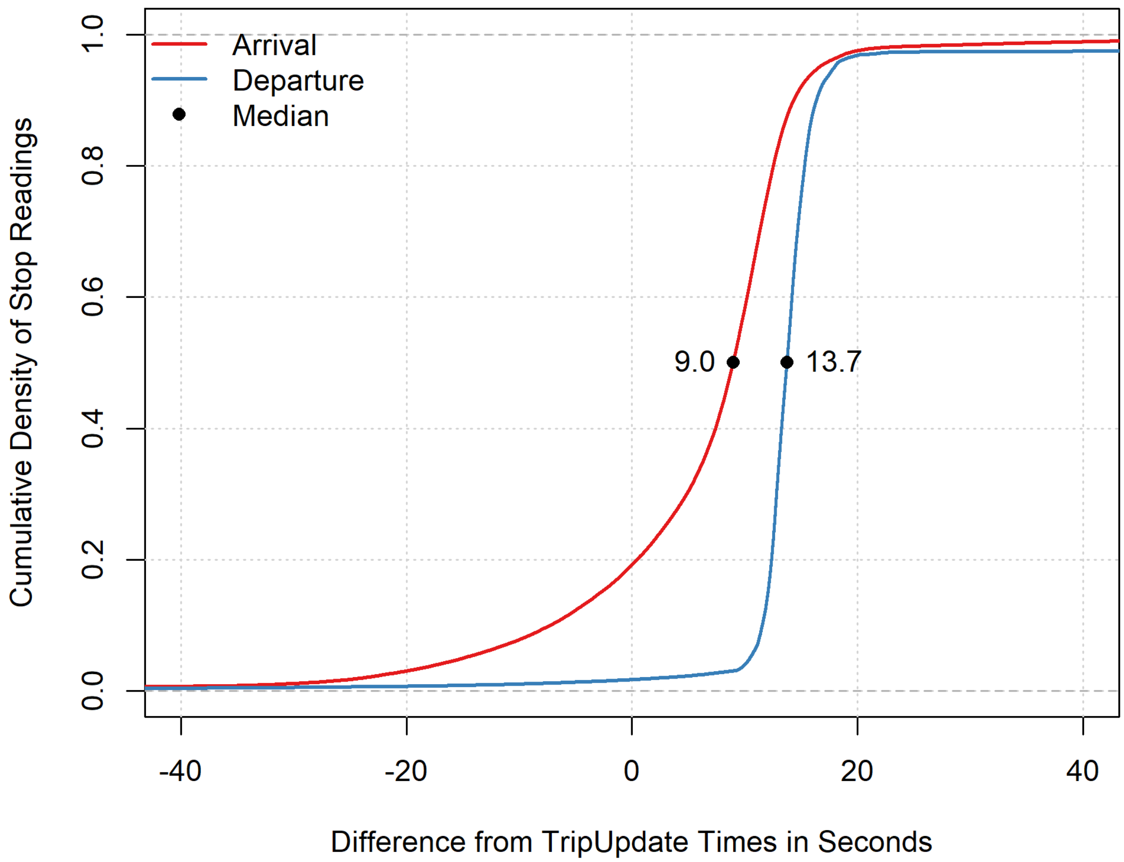

Dwell Times

The arrival and departure times matter on their own because they reflect on-time performance and jointly because dwell time is a feature of transit service that contributes to that on-time performance. Stop dwell times represent the time elapsed at a stop. These values are easily calculated as the difference between the bus arrival and departure times as noted above.

Since the proposed imputation algorithm calculates arrival and departure times using different speed estimates along different segments, it is possible that the imputed times could yield a departure time that occurs before the arrival time. Such a situation would result in a negative dwell time, a physical impossibility. The first dwell time assessment identifies the incidence of negative dwell times and their magnitude.

The incidence of negative dwell times within the imputed data is quite low. Only 0.11% of imputed dwell times are less than zero. This low share further affirms the efficacy of the imputation algorithm. In addition, it is expected that many dwell times within the data set would be zero since most bus drivers continue driving if there is neither anyone waiting to board or to alight at a given stop. Nonetheless, it is important to explore the distribution of those negative dwell times. Figure 6 presents the cumulative distribution of these imputations to show that the magnitude of most of these negative dwell times is quite small. The median is 1.6 seconds, and all of the data are less than 10 seconds. Nonetheless, since such times are impossible, negative dwell times are ultimately recoded within the data to be zero by resetting the departing stop time to equal the arrival time whenever that departing time precedes the arrival time.

Figure 6. Distribution of Imputed Dwell Times that are Less than Zero

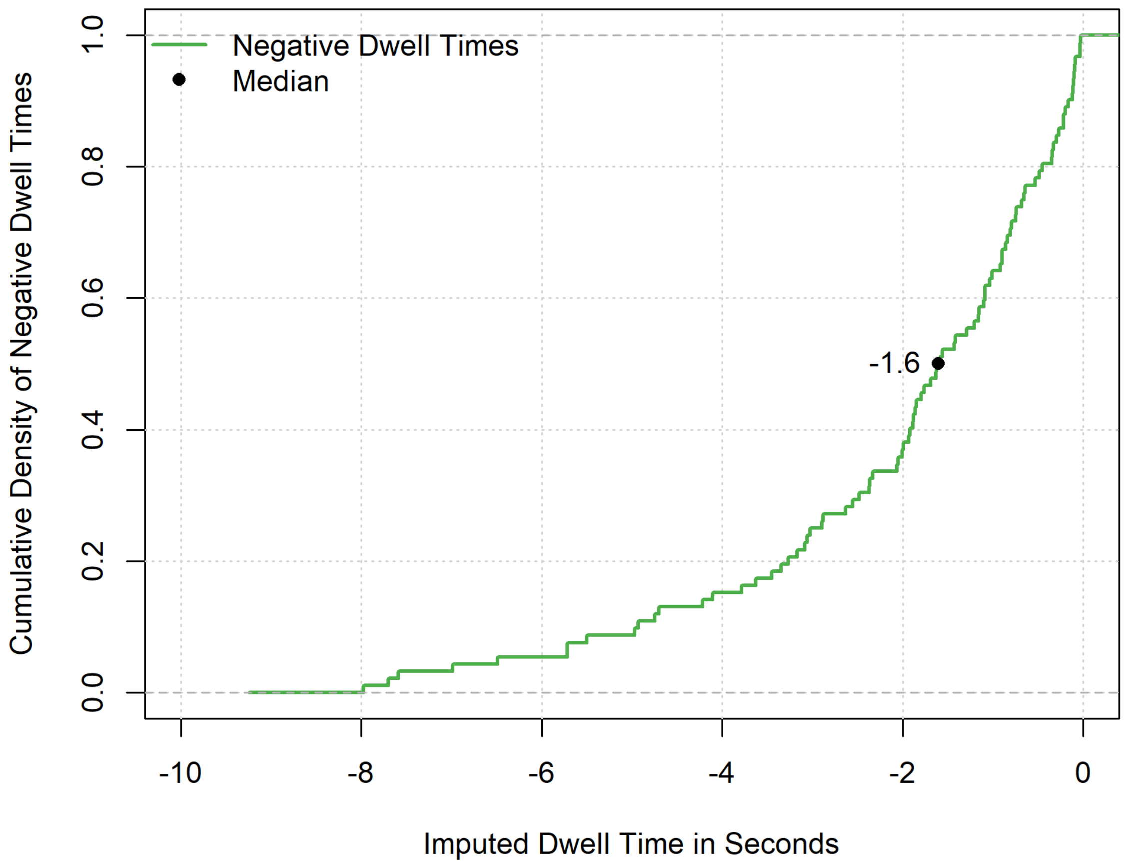

The recoded data allows for a comparison of reported and imputed values. Figure 7 presents a scattergram that places the reported values on the y-axis and the imputed values on the x-axis. These values are highly correlated (Pearson Correlation r = 0.87). A diagonal line of equivalence, where associated dwell times are identical, is superimposed on the plot in green. Many of the values do fall along the line of equivalence, as can be seen most prominently with many of the longer dwell times—such as the value that includes an absurdly long dwell time of 1,200 seconds (or twenty minutes); however, many values fall slightly below this line, likely reflecting the difference between the arrival and departure time offsets from the reported data noted in Figure 5 earlier. This relationship is also evident in the regression line superimposed on the plot in purple, which has a more moderate slope than the line of equivalence. This positioning reflects slightly longer dwell times in the imputed data.

Figure 7. Comparison of Reported and Imputed Dwell Time on MAX

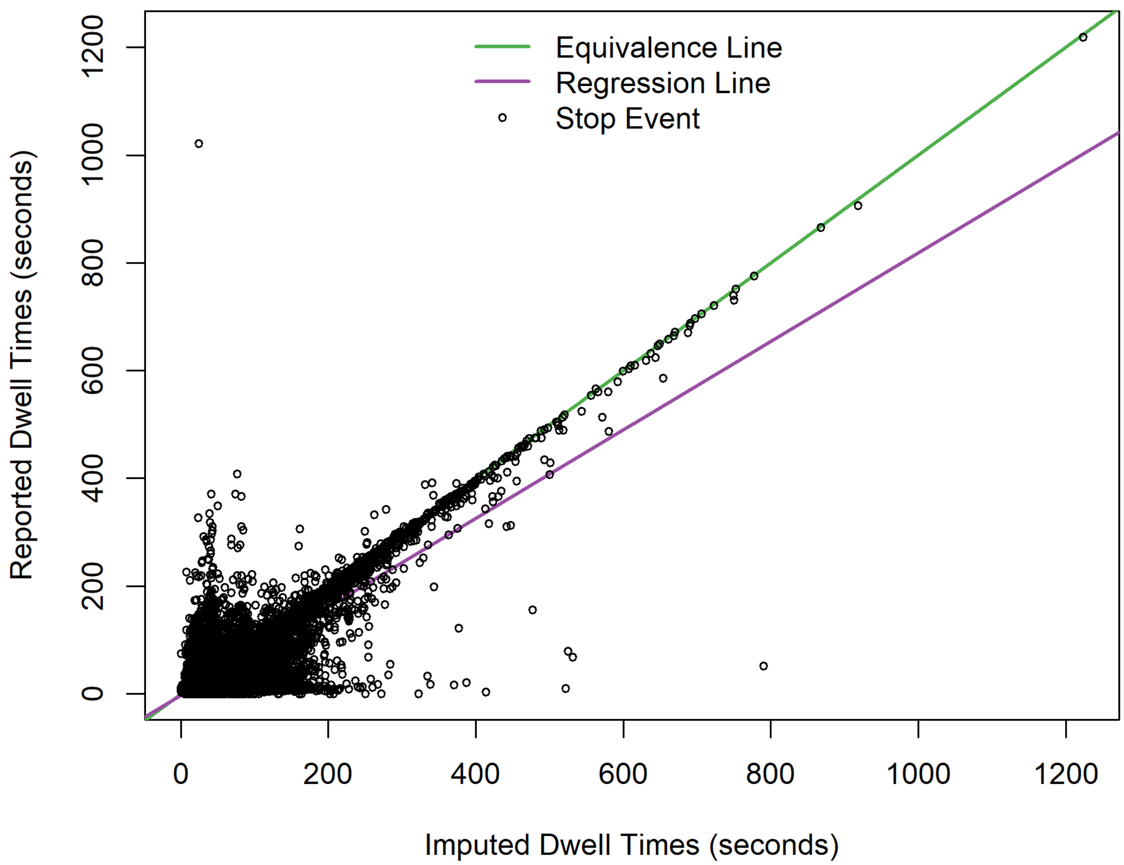

The concentration of values for the lower dwell times makes the lower left-hand portion of the chart difficult to discern. Figure 8 zooms in on that subsection of Figure 7 to present dwell times of three minutes or less. These data show a preponderance of values along the line of equivalence, but many values that are at some distance from the line. Those distances reflect divergent outcomes between the reported and imputed data.

Some spot testing has suggested that the algorithm presented here handles divergences from scheduled routes more conservatively than the MAX GTFS-RT algorithm. The proposed algorithm only considers pings within 10 meters of the scheduled path and ignores route divergences from that path. The MAX GTFS-RT feed reports stop times even when the vehicle is not on the scheduled path.

Figure 8. Comparison of Reported and Imputed Dwell Time on MAX (<180 sec)

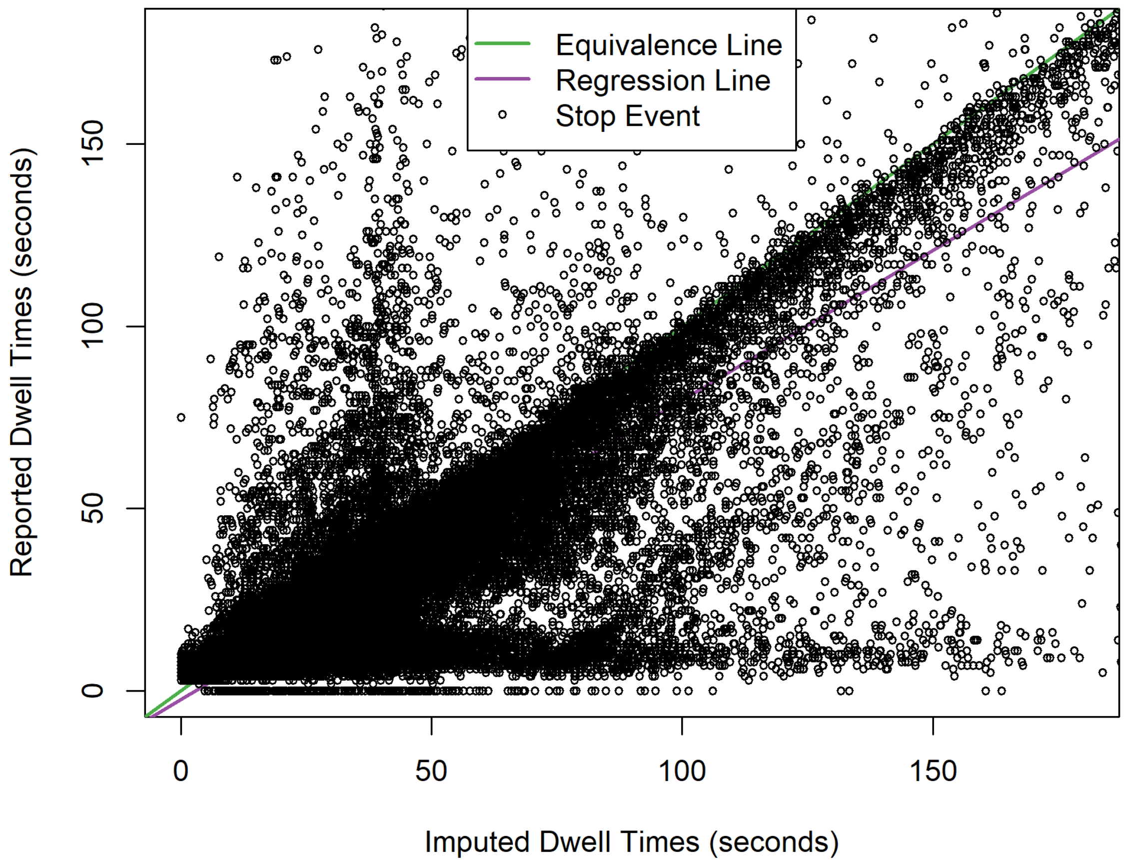

The bottom of Figure 8 shows many dwell times where the reported values are zero, but this research has imputed non-zero values. (There are examples of the opposite, but they are less prevalent.) This feature, as well as the concentration of readings along the x-axis, suggests that the imputation algorithm is estimating longer dwell times than those reported in the MAX GTFS-RT feed.

These dwell time distinctions can be seen more clearly in Figure 9 which shows the distribution of dwell times between the two methodologies. The reported dwell times demonstrate more kurtosis with a taller, slenderer peak than the imputed dwell times. The median value of the imputed values are almost ten seconds longer than those of the reported values, an 81% increase.

Figure 9. Distribution of Reported and Imputed Dwell Time on MAX

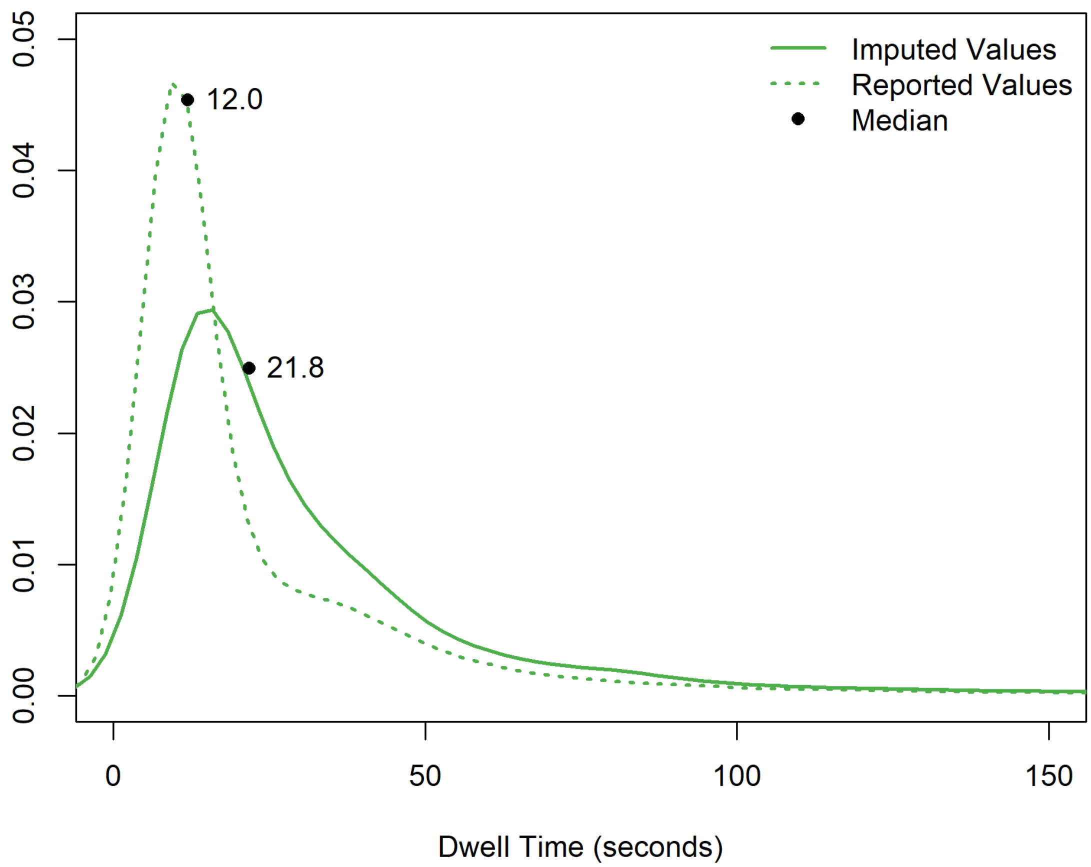

Conclusions

This section describes a challenging but essential process of combining the most basic *VehiclePosition* messages of the GTFS-RT feeds with the GTFS Schedule information to impute actual stop arrival and departure times. This process is designed to prepare a data set composed of spatial segments between stops that include all the stop arrival and departure times. This work is challenging because GTFS schedule coding is messy and inconsistent and the real-world conditions in which transit operates that creates the GTFS-RT data are messy and inconsistent. Nonetheless, this demonstrates that the core GTFS products are capable of being transformed into probe data that can be structured to generate transit performance metrics.

This work provides a highly detailed discussion of the challenges of characterizing the paths that define transit lines in GTFS Static so that they can be cut into segments between stops. This work then provides a discussion of estimating travel speeds along those segments to predict the stop arrival and departure times at the ends of those segments.

The stop times that are imputed are compared to those provided within the MAX GTFS-RT feed. The imputed stop times appear to be slightly later than those reported and to yield longer dwell times, in part because departure times seem to be predicted later than arrival times. The algorithm also resulted in a very small number of negative dwell times which were recoded to equal zero. The consistency of the imputed values suggest they can be appropriately used for generating transit performance measures as is explored in the next section; however, their divergence in key ways from reported numbers, such as dwell time, does raise concerns about fidelity to facts on the ground.

An unusual feature of the MAX GTFS-RT feed is a very sparse set of pings that appear to be triggered more by location than time. This sparseness likely reduces the accuracy of the imputation algorithm. It is expected that in other regions with more pings and more spatially dispersed pings, the accuracy of the algorithm would be improved; however, the fact that quality of the stop time imputations might vary across systems could complicate the cross-system comparison of metrics generated from these probe data. The consistency of those metrics within a single system, however, does suggest their utility for generating performance metrics which is useful for improving transit service, as discussed in the next section of this report.

4. Performance Measurement of Transit Using Probe Data

Introduction

The ultimate goal of this project is to transform probe data into useful performance metrics for transit. While other efforts, such as TIDES, have looked at larger units, this work innovatively focuses on sub-trip components and then aggregates these to higher levels for analysis. This approach leverages the strengths of probe data's granularity. Specifically, this work breaks down transit trips into segments, which are defined spatially as the scheduled path of the vehicle between two adjacent stops and temporally as the time from the arrival at the initial stop to the arrival at the subsequent stop. By creatively combining the special and temporal information from segments, it is possible to generate an array of useful performance measures.

Approach

While performance measurement within the context of transit agencies can take many forms, this report focuses on the aspect that is most pertinent to the riding public, namely schedule adherence. This admittedly dryly named concept cuts to the heart of public transportation. Humans value predictability and fear uncertainty. Much of transit's perception, both positive and negative, flows from its reliability. The schedule published by a transit agency, like a syllabus for a college class, is a contract that sets expectations (although GTFS-RT does provide some wiggle room by offering users updates to schedule deviations). The extent to which those expectations are upheld often determines customer satisfaction with transit service. Poor satisfaction undermines transit by disappointing users, driving away riders, and depressing public support from the allocation of space to the allocation of subsidies.

The goal of the proposed probe-based performance measures is to provide detailed information to stakeholders on where and when schedule deviations are occurring to guide policies that will result in more reliable service delivery. This work is agnostic to the policies implemented to ameliorate transit service. Rather these metrics are designed to provide transit agencies with constant feedback as to the performance of their systems to enable constant refinement. While broadly aimed at identifying schedule adherence, the individual metrics focus on different aspects of this concept to allow planners and policy makers to best diagnose issues and formulate mitigation strategies.

Data Source

The data used in this section of the report were generated as part of the process described in the previous section. Since the algorithm defined earlier uses arrival times at the end of a previous segment to determine the start time of the next segment, it does not generate a start time for the first segment of a trip (although it does calculate a departure time from that start point). To enable a full consideration of trip delivery, the scheduled start times for a given trip are taken from the GTFS Schedule data and applied as the start time for the first segment of the imputed data. Since this represents the official start time of the route, it is reasonable to include this time point in an analysis of transit performance measurement.

As noted in the previous section, if a ping was not available within an acceptable distance envelope from the end of a segment, no stop time could be imputed for those sections. In the MAX data used in this report, 11% of segment readings were dropped because of either a missing departure or arrival time. The structure of the proposed transit performance metrics easily sidesteps this missing data as it is entirely based on the segments that do appear; however, if there is a systematic reason for the missing data, such as a faulty transponder along a low frequency route, the absence of complete data might result in a misrepresentation of that route among the remaining data.

The proposed metrics are clustered thematically. This arrangement is purely for presentation. Users of these metrics should view this presentation as a menu of options and not hesitate to pick and choose amongst them à la carte to best meet their needs.

Route-Level Metrics

The route remains a fundamental unit of transit planning, provision, and perception. Transit agencies are often structured so that planners are responsible for a discrete set of bus lines; garages are often organized to maintain and drive associated vehicles for those lines; and riders are often using only a fraction of a system's available routes for their transit trips. Route level measures remain a primary mechanism for triaging problems with the network.

Travel Time Mismatch

The Travel Time Mismatch metric identifies the extent to which actual route travel times exceed the scheduled travel times expressed as the percentage difference as shown in Figure 10. Ideally, the data points would be gathered close to the origin with actual times matching the scheduled ones. This chart includes two 10-percentage point bands around the origin to propose a triaging mechanism of addressing concerns from the outer bands inward. The bands might be used to prioritize improvement efforts. This approach could be fine-tuned to consider the ridership affected by the routes, an aspect that is beyond the scope of the current project but would be possible if aligned with APC data.

Figure 10. Travel Time Mismatch

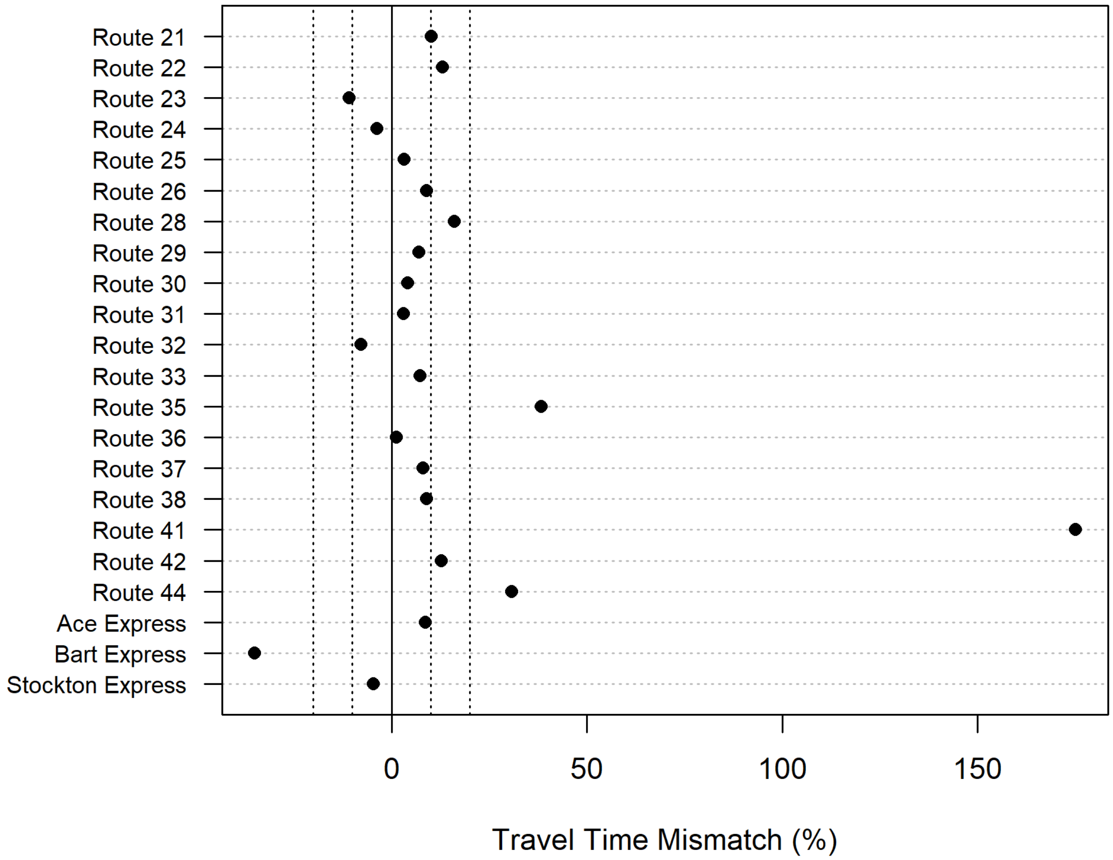

Three routes exceed these bands, with Route 41 appearing to typically take more than twice the scheduled time to complete the route. These routes might be examined first. Several routes have scheduled travel times that substantially exceed their actual travel times. This type of mismatch is less problematic for riders but may be a concern for transit agencies that must schedule drivers. The substantial negative value for the BART Express may reflect scheduled buffering to account for potential traffic that has yet to materialize on this long route (or possibly disappeared after the Covid-19 pandemic.

Late Start Propensity

Transit agencies commonly report their systemwide and route-level on-time performance. These numbers are typically determined by establishing a threshold of what constitutes "on-time" (e.g., between one minute early and five minutes late), tracking those values at specified time points, and then calculating the share of values that are within the threshold. This traditional approach has been criticized for its binary determination of being "on time" and its insensitivity to the number of people affected by any given delay. A rarely mentioned criticism is that this approach obscures the cascading impact of starting a trip late.

The Late Start Propensity metric calculates the share of trips on a given route that, as its name implies, start late. While this is a route-level measure, its calculation is accomplished using the first segment of a given trip. As with the traditional measure, a threshold determining “lateness” is established, and the share of first segments meeting this standard determines the Late State Propensity. For this report, a threshold of five minutes was used, but this could similarly vary by system.

These scores are represented as a bar chart in Figure 11. A horizontal line is superimposed on the chart to denote a breakpoint separating routes with more than a quarter of starts being late and those with less than a fifth. Once again, a transit agency might want to work systematically to reduce the Late Start Propensity among the most frequently late routes.

Interestingly, there appears to be a limited association between the Late Start Propensity and the Travel Time Mismatch metrics. One might assume that a late start would consistently extend the travel time; however, these measures capture different attributes of transit trips. This is strongly seen in Route 41 which, while having the highest Travel Time Mismatch, has the lowest Late Start Propensity.

Figure 11. Late Start Propensity by Route

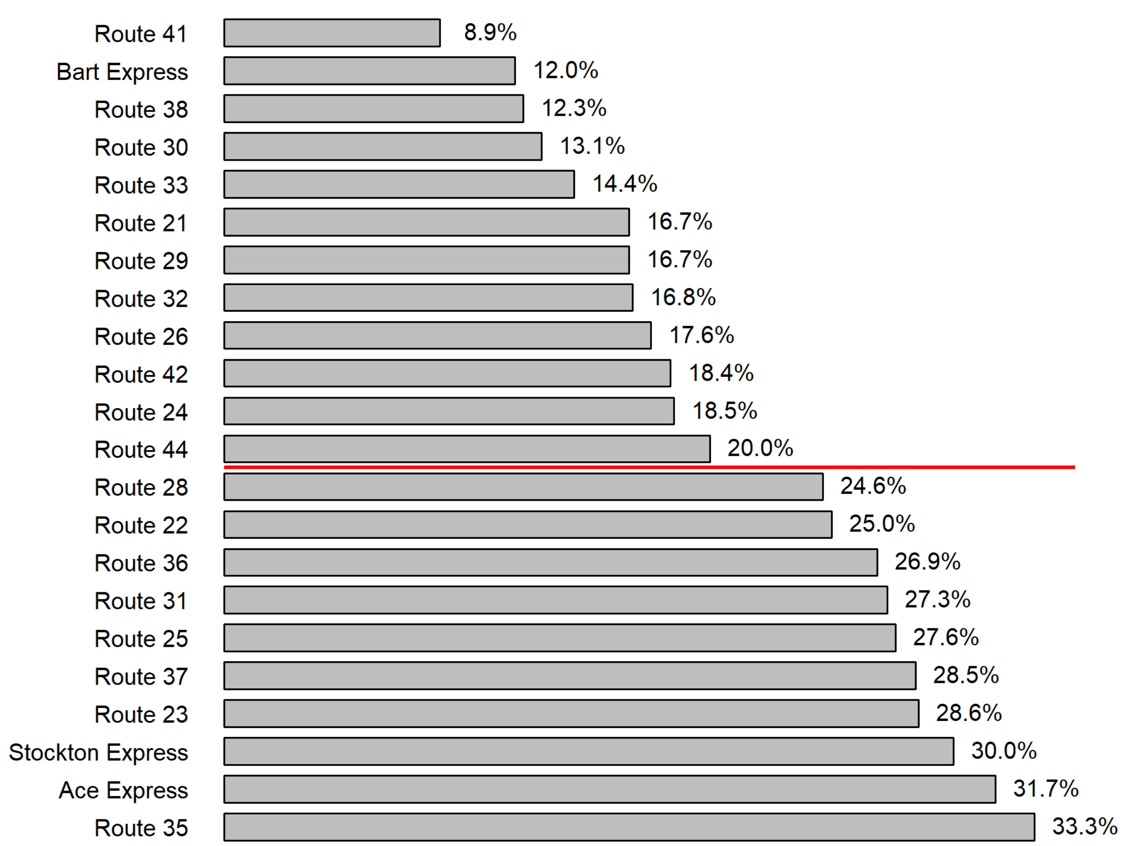

Late Start Distribution

The Late Start Propensity suffers from the threshold effect of classical on-time performance measures. An alternative that might be more useful for transit agencies is the Late Start Distribution. This metric provides box plots to demonstrate the distribution of late starts, again using the dwell time at the first segment as the indicator. This information augments the simple propensity indicator to identify those routes whose variation might be more problematic for passengers. For example, Routes 25 and 31 have similar Late Start Propensities at 27.6 and 27.3, respectively; however, the median start delay for Route 25 is about three minutes—a minute less than Route 31. This finding suggests that, despite a slightly higher Late Start Propensity, typical trips on Route 25 will start substantially more on time than those on Route 31 (and that MAX might focus on addressing Route 31 before Route 25).

Figure 12. Late Start Distribution by Route

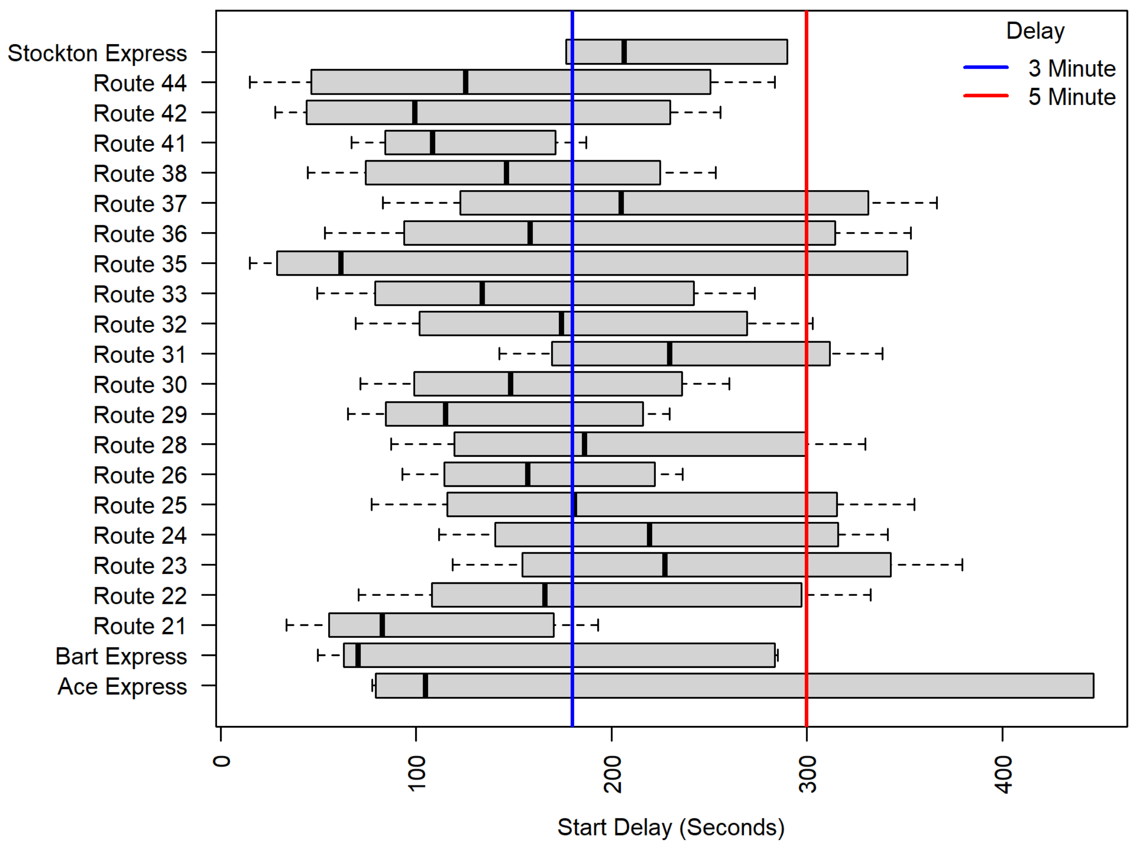

Dwell-Travel Ratio

A route-level metric that entirely emerges from the granular focus on the component segments is the Dwell-Travel Ratio. This metric exploits the structure of the segment data that breaks down the total segment time into the dwell time for boarding and alighting passengers and the travel time for transporting them. All things being equal, passengers prefer to minimize the former and maximize the latter. This preferred outcome results in a Dwell-Travel Ratio less than one. Despite this preference, the data in Figure 13 show the opposite is typical on all MAX routes (excluding the express routes that run long distances on highways). Routes 23, 32, and 41, report twice as much time spent waiting at stops than moving between them.

While these values are likely somewhat inflated because of the relatively sparse ping availability increasing the estimation of dwell time at the expense of travel time, they can be used comparatively within a single transit agency to identify routes of concern. Figure 13 includes a dotted line at a Dwell-Travel Ratio of two to set a possible threshold for prioritizing exploration among MAX routes. A different agency with a denser array of pings might see generally lower Dwell-Travel Ratios and adjust the thresholds of concern accordingly.

Figure 13. Dwell Time Ratio by Route

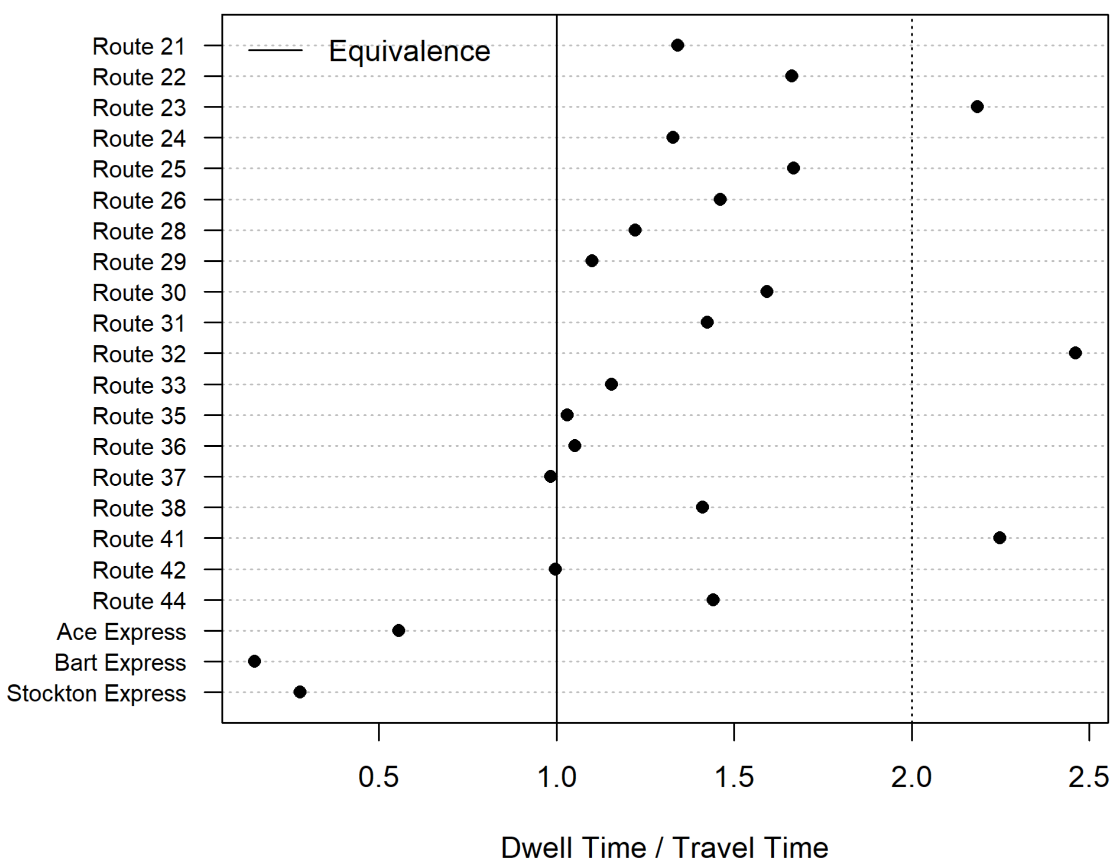

In any case, it is important to note that this ratio is strongly impacted by late starts. Starting routes on time is an easy way to reduce these ratios by removing all of the initial dwell time from the calculation. Conversely, there might be some justification for excluding initial segments from the calculation on the logic that the late start is a distinct performance aspect from the riding experience following the trip start. Some agencies might choose to calculate this ratio without those initial segments (or excluding the initial dwell) to focus their policy analysis and the associated interventions on the trip portion rather than the pre-embarkation portion. Such calculations are reasonable but need to be documented to avoid comparisons to the more complete approach presented here.

Finally, the Dwell-Travel Ratio is unique among the measures presented in this report in that it does not directly address schedule adherence but rather focuses on rider experience. One can readily imagine that the same interventions that facilitate schedule adherence, such as transit signal priority and dedicated transit rights-of-way, also reduce the Dwell-Travel Ratio.

Segment-Level Metrics

A focus of this work is to introduce the idea of segments as the core unit of transit trips and the building block for transit performance measures. As demonstrated in the previous section, segments can generate route-level metrics. Segments can also complement those higher-level measures by offering a means to more deeply explore and diagnose issues identified at that higher level. In addition to this complementary role, segments can also, on their own, provide a different path to analyzing performance across a network.

Decoupling segments from the routes they serve might involve slightly recoding them. To enable fidelity to the full provision of transit service, segments are defined uniquely for each route variant. While this feature allows transit service to be studied fully at its most granular level, it results in identical segments being defined as distinct entities. This problem occurs in two ways: multiple variants on a single route share segments along the common portions of those variants, and multiple routes share segments where routes offer overlapping service, such as along a downtown corridor. Despite sharing the same physical definition, the service along those segments can vary with the actual route and variant. This variation is not due to the traffic conditions and, therefore, is unlikely to affect the travel time; rather this variation is due to the nature of the transit route and is likely to affect dwell times. For example, one might expect that segments shared by an express and local route will see more time spent boarding and alighting passengers on the express rather than on the local.

Defining segments as unique to each route-variant combination allows for these nuances to be explored but can complicate both presentation and interpretation. Since geographic information system software will map segments on top of each other, it is possible that traits of one segment will be obscured by those of a shared segment. If segments are recoded to be unique to their physical description—which can be useful for many performance analyses—the unique aspects of those segments tied to their route might be averaged away in an aggregated analysis. For example, the long dwell times of an express that runs only a few times a day might not show as problematic once the times are offset by the shorter dwell times of more frequent local service.

Transit analysts need not be dissuaded from segment analysis because of these issues; rather, they should be mindful of how they define and visualize segments and these approaches' potential consequences. This segment-level section incorporates a discussion of these issues.

Slow Segments

This approach can be used to identify Slow Segments, whose average actual travel time exceeds its scheduled time by two (somewhat analogous to the Transit Time Mismatch metric presented at the route level).

In this presentation, the initial trip segments are excluded from Slow Segment consideration to remove concerns about the role of late starts. This exclusion is entirely arbitrary and made here given the extensive previous consideration of the Late Start Propensity and Late Start Distribution metrics. This exclusion also is made to demonstrate the flexibility of the segment approach to transit performance measurement, which leverages its granularity to enable all sorts of nuanced consideration. (It is perhaps worth mentioning that decisions regarding what segments to include or exclude should always be documented in the presentation of the resulting performance measures. This documentation is particularly important when measures are used to track transit performance over time to ensure apples-to-apples comparisons.)

Here the Slow Segment calculation is simply the ratio of the actual travel time to the scheduled travel time rounded to the nearest integer. Any segment with a score of two or higher is designated a Slow Segment as shown in grey in Figure 14. This rounding approach shaves off some subtlety but alerts the analyst to the segments of greatest concern.

Figure 14. Slow Segment Shares (%)

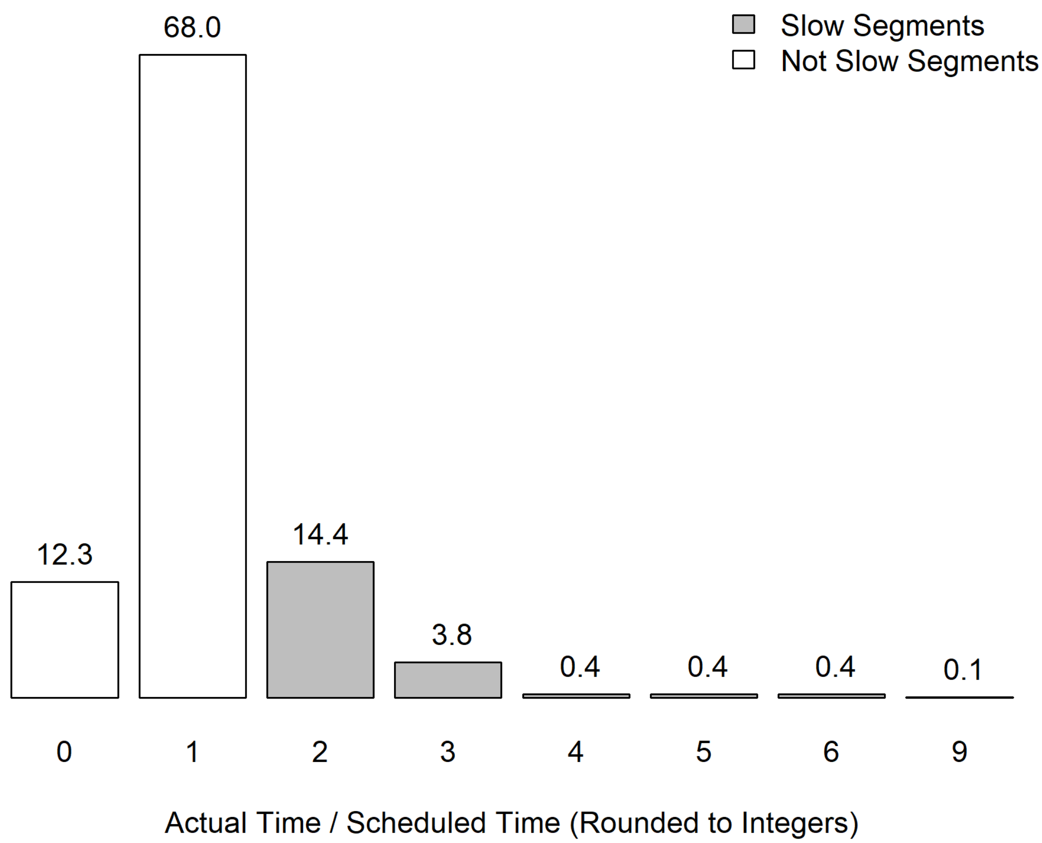

The shares shown in Figure 14 demonstrate that almost a fifth of segments meet the threshold to qualify as Slow Segments. Visualizing these segments spatially while symbolizing them by their scores translates this transit performance metric into a diagnostic tool.

Before mapping Slow Segments, it is useful to see if any of them overlap, which might confuse the interpretation of the score visualized on a map. In the MAX data presented above, 180 segments were characterized as slow (i.e., with Slow Segment scores of two or more). Adding an additional identifier to these segments based on their unique first stop / second stop pairing reveals 17 stop dyads that are repeated. These unique pairings account for 35 segments or 19.4% of the Slow Segments. Table 4 presents these segments defined by their unique stop dyads, the routes traversing those segments, and their Slow Segment Scores.

Table 4. Overlapping Slow Segments and their Associated Scores

Start Stop	End Stop	Routes	Scores
9th St & D St	9th St & River Rd	29, 29	2,2
Amtrak Station	Held Briggsmore	25, 25	2,2
Dale Rd & Veneman Ave	Standiford Ave & Dale Rd	37, 22	2,2
Hatch Rd & Herndon Rd	Herndon Rd & Hatch Rd	44, 29	3,3
Hatch Rd & Richland Ave	Hatch Rd & Herndon Rd	44, 29, 29	2,6,5
Herndon Rd & Hatch Rd	Herndon Rd & Eugene Ave	29, 29	6,5
Herndon Rd & Pecos Ave	Nadine Ave & Herndon Rd	29, 29	2,2
Herndon Rd & Pecos Ave	Herndon Rd & Latimer Ave	29, 29	2,2
I St & 7th St	H St & 4th St	36, 26	2,2
K St & 15th St	McHenry Ave & Jones St	23, 22	2,2
McHenry Ave & Bowen Ave	McHenry Ave & Bowen Ave	23, 22	2,2
McHenry Ave & Coralwood Rd	McHenry Ave & Grecian Ave	23, 35	3,2
McHenry Ave & Morris Ave	McHenry Ave & Stoddard Ave	23, 22	3,2
McHenry Ave & Norwegian Ave	McHenry Ave & Tokay Ave	23, 22	2,2
McHenry Ave & Orangeburg Ave	McHenry Ave & Coolidge Ave	23, 22	3,2
Sisk Rd & Plaza Pkw	Sisk Rd & Earl St	36, 25	2,2
Sonora Ave & Herndon Rd	Herndon Rd & Pecos Ave	29, 29	2,2

These data are interesting on several levels. First, they demonstrate how segments can be shared by different variants of the same bus route (e.g., Routes 25 and 29) as well as by different routes entirely. Second, they show general consistency between the services on any given stop dyad. With the exception of the one section between Richland Avenue and Herndon Road along Hatch Road that has three overlapping segments, all the other overlaps have only two segments, and the difference between the Slow Segment scores is never more than one unit. It is unsurprising that there would be consistency along the same section of roadway. What is a bit surprising is that the one section with three overlapping segments reports very different scores between the Route 44 (Score = 2) and the Route 29 (Scores = 6 & 5) segments.

The segment approach presented here allows for a deeper investigation into this finding. The stop dyad of interest is roughly 430 meters from end to end, a little more than a quarter mile, along a major road with two traffic lights. Within the GTFS products, Route 44 defines only one segment for this section and schedules exactly one minute (60 seconds) for both dwell and travel time; by contrast, Route 29 defines three segments for this section and schedules two-and-a-half minutes (150 seconds) for one but only 10 and 20 seconds, respectively, for the other two. These absurdly low scheduled times result in a very high Slow Segment score. By contrast, the segment that has a somewhat high schedule time does not remotely qualify for being considered a Slow Segment. The average time to traverse this section across all four segments is 87 seconds. This finding suggests that the minute allocated for Route 44 is too low and the two-and-a-half minutes for most of Route 29 segments is too high. The transit performance metrics reveal disparities between the schedule and travel information, and the probe data that informs those metrics can offer guidance for removing those disparities.

Figure 15. Slow Segments

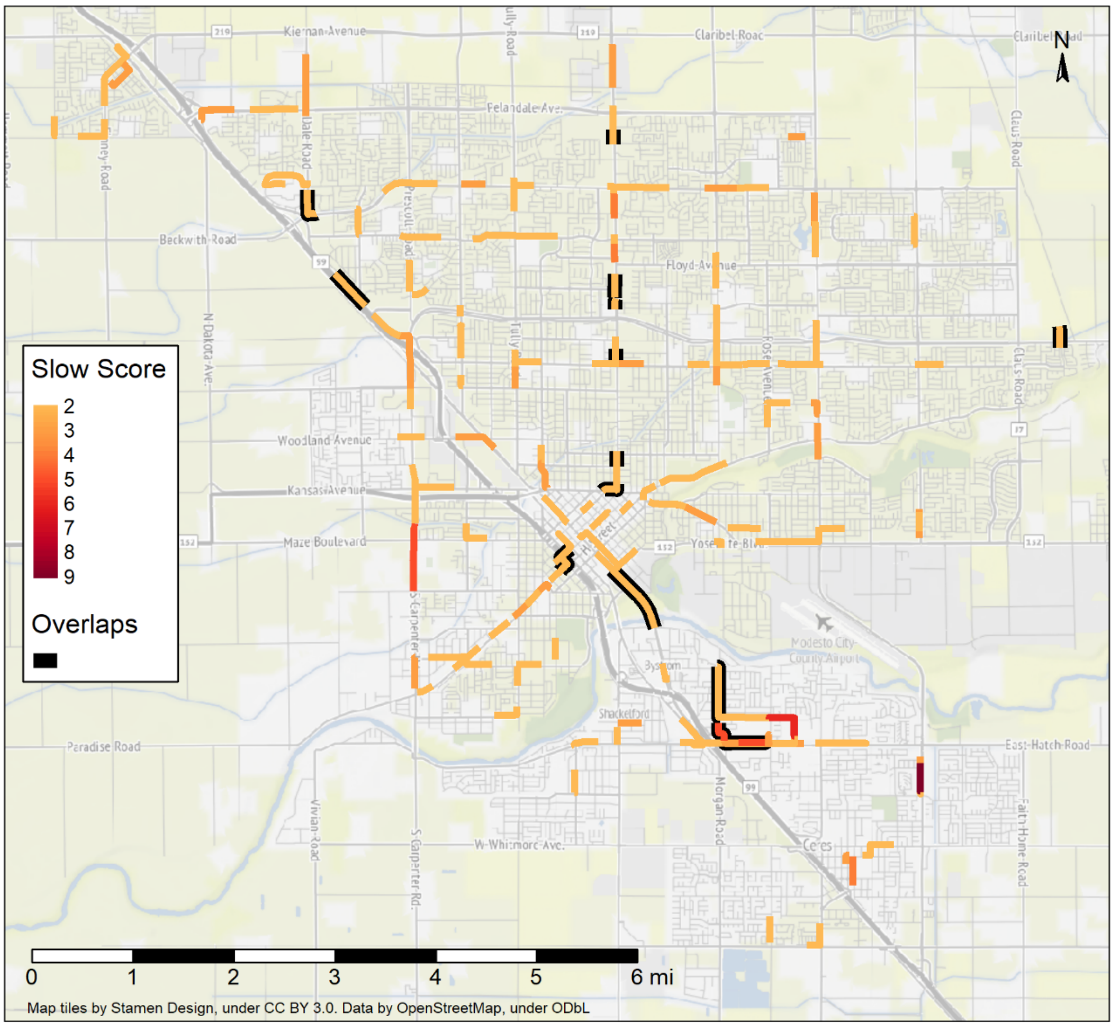

Figure 15 maps these Slow Segments and identifies overlaps by outlining them in black. This type of visualization of transit performance metrics helps transit agencies understand where their schedule deviates from their provision of transit service. The visualization does not directly reveal the source of those deviations but can provide a starting place for planners knowledgeable about the system to explore further. Similar segment-based visualizations can be done for a specific route or for a given time window. Furthermore, the information provided can be aggregated into all sorts of reportable metrics. For example, a transit agency might keep a monthly record of the share of segments traveled that are Slow Segments.

Disproportionate Dwell

By splitting total segment time into boarding / alighting and travel components, it is hoped that this approach will enable a better understanding of at-stop activities. Transit vehicles already collect an array of information during the boarding and alighting period from door opening to passenger counting to ramp usage. While the GTFS standard includes fields to capture some of this useful information, these fields are inconsistently populated. The substantial effort of this project to impute stop arrival and departure times is an attempt to approximate information that is simply not made readily available to the public.

While the dwell times imputed as part of this research do not distinguish in detail what activities are occurring at the transit stop, they do provide a basic accounting of the time spent there. That information allows for the identification of Disproportionate Dwell Time, a performance measure that captures unusually long stop durations.

Figure 16 shows the distribution of dwell times that are heavily right skewed (so much so that the chart has been censored for readability). The data appear to be log-normal as shown by the curve superimposed on the histogram in Figure 16. (Please note that the log-normal curve, which can only incorporate values greater than zero, was generated without 82 zeros constituting 0.10% of the total sample.)

Figure 16. Histogram of Dwell Times

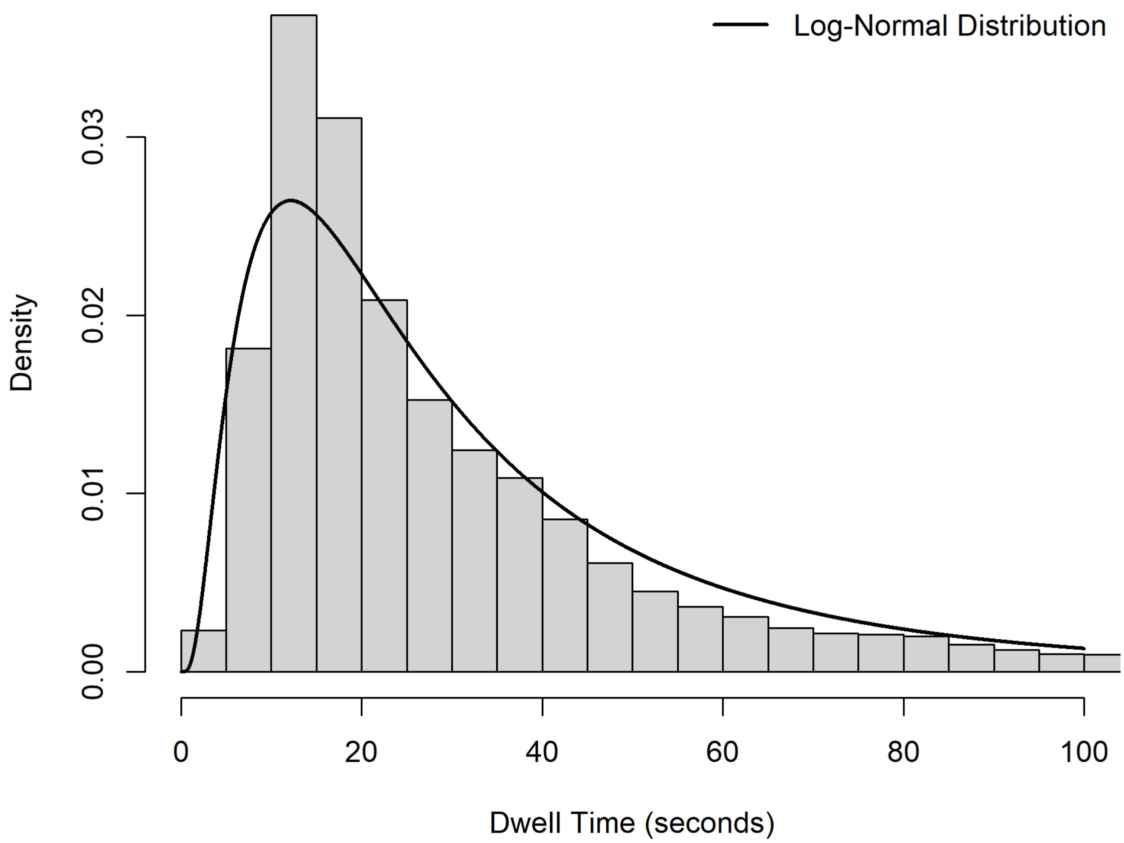

A log-normal distribution occurs when the logarithm of the underlying values is normally distributed. This property enables the dwell times to be log-transformed. (To incorporate the 82 zero values, a one was added to each dwell time—both zero and non-zero—before taking the logarithm.) These logged values were then converted to standardized values.

Standardized values offer a way for a transit agency to identify areas of concern and address them. Standardized values do not represent a comparable transit performance metric since they are unique to a given agency, but they do nicely convert the skewed values of the raw data into a more analytically approachable form. Figure 17 presents those section-starting locations whose averaged standardized values for dwell time exceed 0.75. These values were calculated by aggregating all segments that shared the same start and end point so that there are no overlapping representations of the data. This analysis also includes initial segments so that late starts are incorporated.

Figure 17. Standardized Dwell Times (Logged) Above 0.75

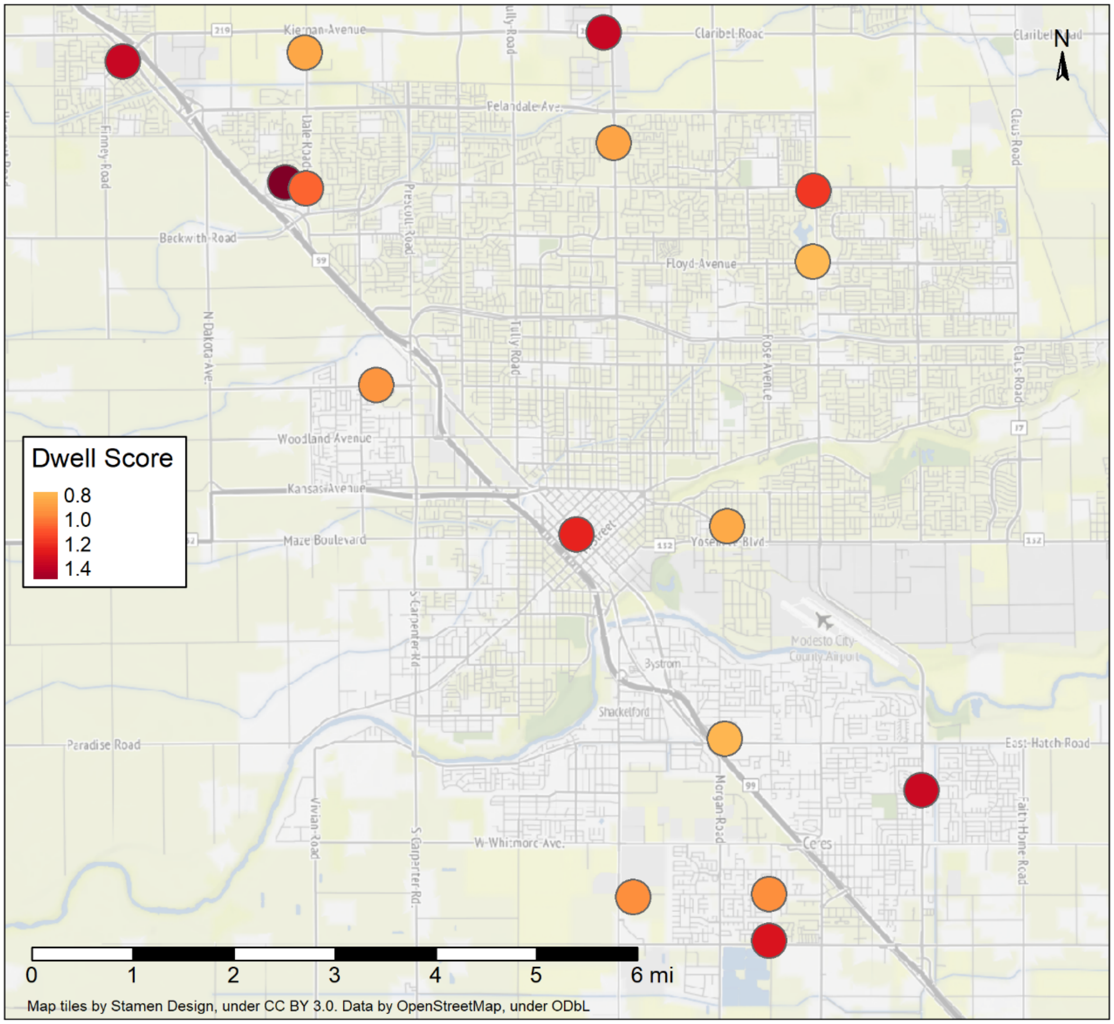

This approach to transit performance assessment is aimed at internal analysis to quickly determine the locations of the most outlying dwell times. Another approach that is comparable is to set a threshold for excessive dwell time and identify all the stops that cross that threshold.

More Than a Minute Dwell

This research proposes an easily understood cut point for excessive dwell time—one minute. Any segment that begins with an average dwell time of more than a minute warrants exploration. Figure 18 presents this metric spatially and symbolizes the stop locations by their average dwell time—some of which are over four minutes. This metric can easily be turned into a systemwide indicator that calculates the share of all segments that have More Than a Minute Dwell Times.

Figure 18. More than a Minute Dwell Times

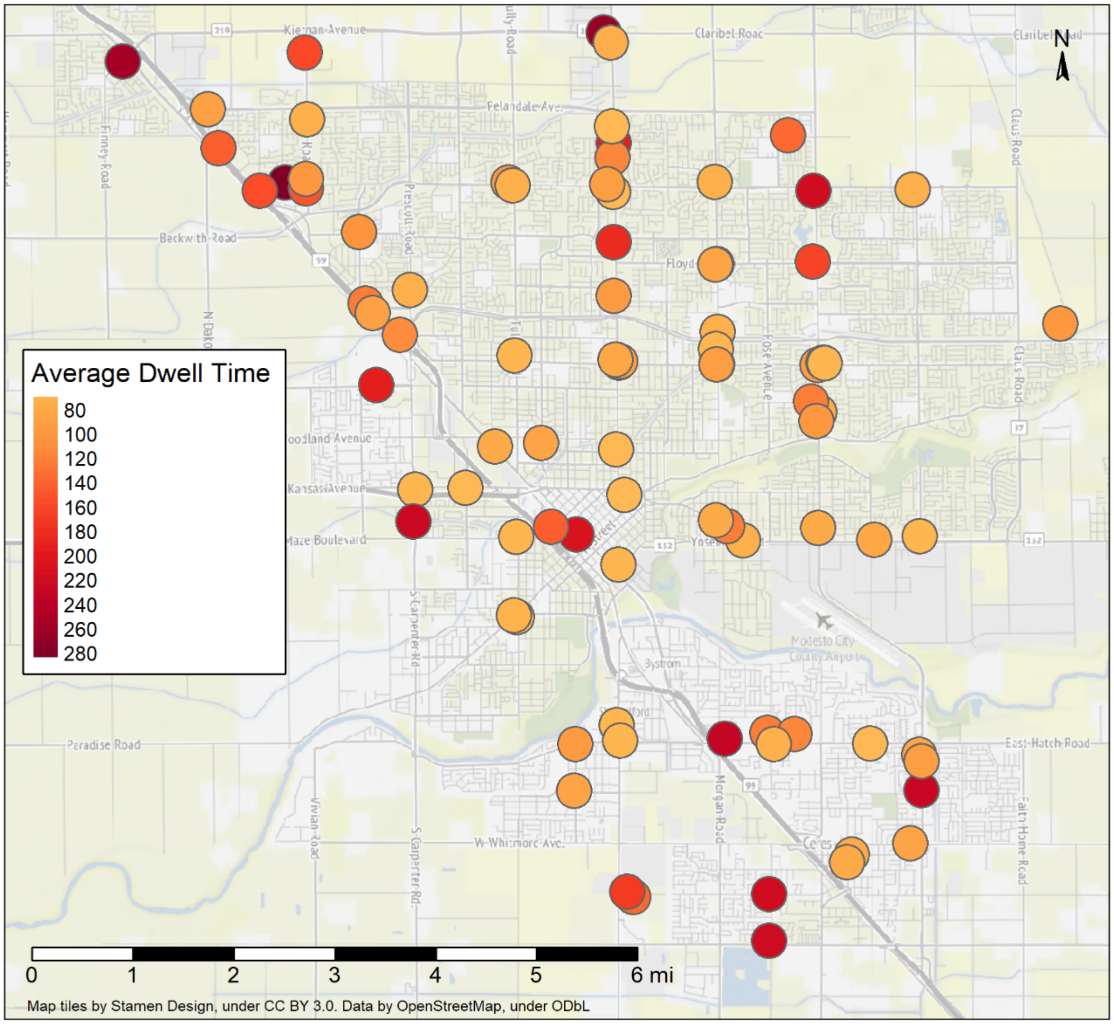

Conclusion

This section applied the segment data with imputed stop times to turn probe data collected by buses and reported to the public into different transit performance measures. The primary goal of this effort was to demonstrate the added potential of probe data for enabling a granular understanding of transit service. The metrics presented emphasize tools for transit agencies to assess their own networks as well as comparable metrics for transit agencies to assess their systems over time and compare them to other systems that employ the same metrics.

This work relied entirely on publicly available GTFS Static and Realtime feeds. A secondary goal of this effort is to democratize access to probe data and therefore place transit performance measurement into the hands of a larger suite of stakeholders from state and local governments to community activists and transit advocates. The goal of any performance metric is to improve performance. Increasing access to such measures increases the likelihood that the service will improve as well.

Bibliography

AC Transit. (2022). 78 AC transit schedule. AC Transit.

Ahmed, Md. S., Hoque, M. A., & Khattak, A. J. (2016). Demo: Real-time vehicle movement tracking on Android devices through Bluetooth communication with DSRC devices. *2016 IEEE Vehicular Networking Conference (VNC)*, 1–2. https://doi.org/10.1109/VNC.2016.7835956

Alam, N. (2012). *Vehicular positioning enhancement using DSRC.* [PhD Thesis]. University of New South Wales, Sydney, Australia.

Amendment of Parts 2 and 90 of the Commission's rules to allocate the 5.850-5.925 GHz Band to the mobile service for dedicated short range communications of Intelligent Transportation Services, ET Docket No. 98-95 RM-9096 (October 22, 1999).

Bollo, R., Wolpert, A., Toth, C., Timcho, T., Smith, T. A., McLaughlin, K., Toussant, E., Wisor, N., James, R., Baker, J., & Justis, M. (2018). *Concept of operations for the Connected Vehicle Environment for the Smart Columbus Demonstration Program* (pp. 1–215) [Final Report]. City of Columbus.

Bu, F., & Chan, C.-Y. (2005). Pedestrian detection in transit bus application: Sensing technologies and safety solutions. *IEEE proceedings. Intelligent Vehicles Symposium, 2005.*, 100–105. https://doi.org/10.1109/IVS.2005.1505085

Chan, J. (Director). (1985, December 14). *Police story*. Golden Harvest.

Cho, H., Chung, T., & Yim, C. (2000). Bus Information Service and Internet Service Plan using active DSRC System. *Proceedings of the 7th World Congress on Intelligent Systems*. 7th World Congress on Intelligent Systems, Turin, Italy.

Choi, K.-J., Kim, D.-H., Yoon, D.-W., & Park, S.-K. (2006). The traffic management system for emergency vehicles based on DSRC System. *Journal of the Institute of Electronics Engineers of Korea TC*, *43*(9), 40–48.

de Bont, J. (Director). (1994). *Speed* [Action Movie]. 20th Century Fox.

FCC chairman signals an end to DSRC as he prepares to carve up 5.9GHz spectrum – but ITS America hits back. (2020, October 27). *Traffic Technology Today*. https://www.traffictechnologytoday.com/news/connected-vehicles-infrastructure/fcc-chairman-signals-an-end-to-dsrc-as-he-prepares-to-carve-up-5-9ghz-spectrum-but-its-america-hits-back.html

Federal Communication Commission. (2019, April 22). *Dedicated Short Range Communications (DSRC) Service*. Federal Communications Commission. https://www.fcc.gov/wireless/bureau-divisions/mobility-division/dedicated-short-range-communications-dsrc-service

Federal Transit Administration, Institute of Transportation Engineers, & American Public Transportation Association. (2002). *Prioritizing transit in a connected vehicle world—A Transit Signal Priority (TSP) Workshop* (pp. 1–12) [Workshop Summary]. US Department of Transportation.

Gettman, D. (2020, June 3). *DSRC and C-V2X: The future of connected vehicles*. Kimley-Horn: News and Insights - Perspectives. https://www.kimley-horn.com/dsrc-cv2x-comparison-future-connected-vehicles/

Gong, Y. L. (2017). *Use of DSRC and other sensor signals for vehicular localization*.

Hamilton, G. (Director). (1973, June 27). *Live and let die* [Action Movie]. United Artists.

Han Su, Cheng-Hsuan Cho, Yi-Hong Chu, & Wen-Yao Chang. (2012). A traffic information dissemination mechanism based on DSRC/WAVE and its applications. *2012 12th International Conference on ITS Telecommunications*, 635–638. https://doi.org/10.1109/ITST.2012.6425258

Hirai, S., Ogawa, T., Yoshinaga, H., Imamura, T., Yamada, H., Iwata, T., & Tanaka, K. (2006, October). A field trial of a highway bus location system utilizing DSRC. *Proceedings of the ITS World Congress, London*. ITS World Congress, London.

Johnson, S. A. (1998). Extension of the Dedicated Short Range Communications standard for application to transit management systems. In *Compendium: Graduate student papers on Advanced Surface Transportation Systems, 1998* (pp. 193–224). Texas Transportation Institute.

Leonard, B. D. (2018, February 14). *The Utah DSRC MMITSS Project*. National Operations Center of Excellence 2018 SPaT Challenge Peer Exchange, Salt Lake City, UT.

Leonard, B. D., Mackey, J., Sheffield, M., Bassett, D., Larson, S., & Hooper, I. (2019). Demonstrating transit schedule benefits with a Dedicated Short-Range Communication-Based Connected Vehicle System. *Transportation Research Record: Journal of the Transportation Research Board*, *2673*(12), 215–224. https://doi.org/10.1177/0361198119859321

Liao, C.-F., & Davis, G. A. (2008). An adaptive transit signal priority strategy using GPS/AVL and wireless communication systems. *15th World Congress on Intelligent Transport Systems and ITS America Annual Meeting 2008*, 688–699.

Maitipe, B., & Hayee, M. I. (2010). *Development and field demonstration of dsrc-based v2i traffic information system for the work zone*.

Mann, S. (Director). (2015, November 13). *Heist* [Crime Movie]. Lionsgate Premiere.

Nace, B., & Toone, J. (2011). Connected Vehicle at 4.9 GHz for Transit ITS-Unified Communication Architecture for an ITS enterprise and pathway to DSRC. *18th ITS World Congress, Orlando, 2011. Proceedings*, 1–13.

Perry, F. (2017, May 3). *Overview of DSRC messages and performance requirements*. University of Florida Transportation Institute DSRC and Other Communication Options for Transportation Connectivity Workshop, Gainesville, Florida.

Regional Transit Authority of Southeast Michigan. (2018). *Washtenaw BRT refinement study Transit Signal Priority (TSP) Evaluation & Recommendations* (pp. 1–35) [Draft Report]. Regional Transit Authority of Southeast Michigan.

Research and Innovative Technology Administration. (2009). *SAE J2735—Dedicated Short Range Communications (DSRC) message set dictionary*. U.S. Department of Transportation. https://www.standards.its.dot.gov/Factsheets/PrintFactsheet/71

Smith, S. F., Isukapati, I., Bronstein, E., & Igoe, C. (2018). *Integrating transit signal priority with adaptive signal control in a connected vehicle environment* (pp. 1–37) [Phase 1 Final Report]. Mobility 21 University Transportation Center Carnegie Mellon University.

Turnbull, K. F., Cherrington, L., Elgart, Z., Zmud, J., Baker, T., & Hudson, J. (2017). *Automated and Connected Vehicle (AV/CV) test bed to improve transit, bicycle, and pedestrian safety: Technical report* (Technical Report FHWA/TX-17/0-6875-1; p. 133). Texas (A&M) Transportation institute.

Valentine, D., Polinori, A., Windholtz, M., Gibbs, W., Neumeister, D., Baumgardner, G., & Paselsky, B. (2019). *Connected Vehicle (CV) infrastructure – Urban bus operational safety platform* (FTA Report No. 0133; pp. 1–109). USDOT Federal Transit Administration.

Wiquist, W. (2020). *FCC Modernizes 5.9 GHz band for Wi-Fi and auto safety: New plan makes additional spectrum available immediately for improved indoor Wi-Fi connectivity while authorizing market-driven transportation safety services*. Federal Communications Commission.

Wu, X., Subramanian, S., Guha, R., White, R. G., Junyi Li, Lu, K. W., Bucceri, A., & Tao Zhang. (2013). Vehicular communications using DSRC: Challenges, enhancements, and evolution. *IEEE Journal on Selected Areas in Communications*, *31*(9), 399–408. https://doi.org/10.1109/JSAC.2013.SUP.0513036

Zimmer, R. E., Burt, M., Zink, G. J., Valentine, D. A., & Knox, W. J. J. (2014). *Transit safety retrofit package development final report* (Final Report FHWA-JPO-14-142; pp. 1–123). US Department of Transportation.

Zimmer, R., Pierce, B., Burns, M., Polinori, A., Guspan, A., & Sedda, A. (2017). *Transit bus stop pedestrian warning application – architecture and design* (FHWA-JPO-16-401; pp. 1–107). Federal Highway Administration.

About the Author

Gregory L Newmark, PhD

Dr. Newmark is an Associate Professor in the Department of Design and Planning in the School of Architecture and Planning at Morgan State University. He is a research associate of the Mineta Transportation Institute and a faculty researcher at the National Transportation Center. His research focuses on fostering the sustainability of travel, particularly through transit provision and use.

MTI FOUNDER

Hon. Norman Y. Mineta

MTI BOARD OF TRUSTEES

Directors

MINETA TRANSPORTATION INSTITUTE