SUPPLEMENT C

CASE STUDIES

CHAPTER C1

INTRODUCTION

This supplement contains five case studies that demonstrate approaches to travel time reliability monitoring described in the guidebook. In particular, the case studies illustrate real-world examples of using a travel time reliability monitoring system to quantify the effect of various influencing factors on the reliability of the system.

The goal of each case study is to illustrate how agencies apply best practices for monitoring system deployment, travel time reliability calculation methodology, and agency use and analysis of the system. To accomplish this goal, prototype travel time reliability monitoring systems were implemented at each of the five sites. These systems take in sensor data in real-time from a variety of transportation networks, process this data inside a large data warehouse, and generate reports on travel time reliability for agencies to help them better operate and plan their transportation systems. Each case study chapter consists of the following sections:

- Monitoring System
- Methodological Advancement
- Use Case Analysis
- Lessons Learned

These sections map to the master system components, as shown below in Exhibit C1-1.

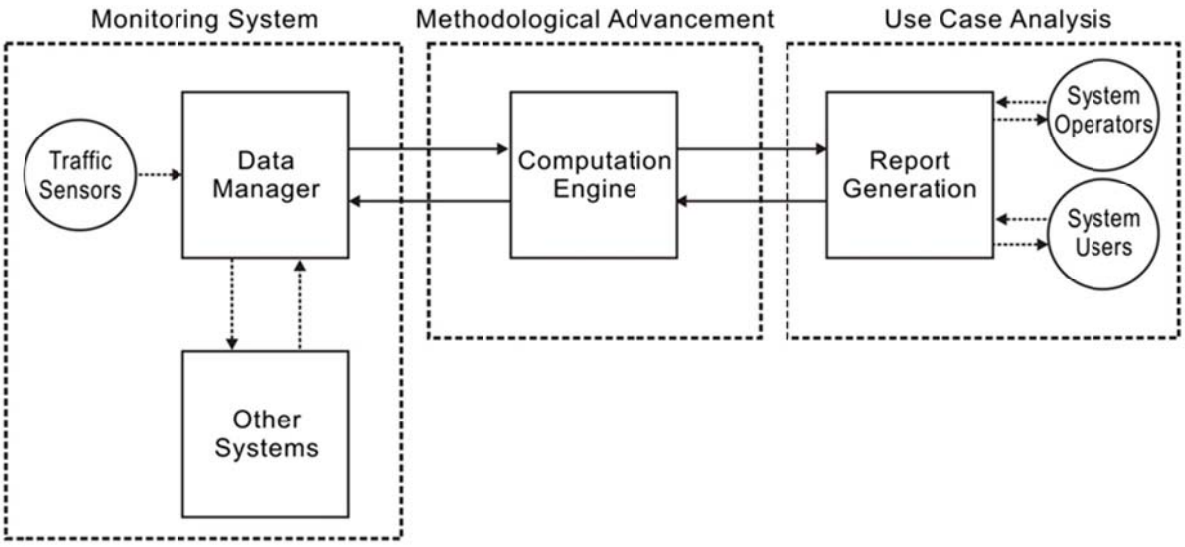

Exhibit C1-1: Reliability Monitoring System Overview

The case studies were performed in San Diego, California; Northern Virginia; Sacramento/Lake Tahoe, California; Atlanta, Georgia; and New York/New Jersey. Exhibit C1-2 shows the case study locations.

Exhibit C1-2: Case Study Locations

This supplement is organized into six chapters. The five chapters following this introductory chapter are titled by the location of the case study demonstration:

- Chapter C2: San Diego
- Chapter C3: Northern Virginia
- Chapter C4: Sacramento/Lake Tahoe
- Chapter C5: Atlanta
- Chapter C6: New York/New Jersey

CHAPTER C2

SAN DIEGO, CALIFORNIA

This case study focused on using a mature reliability monitoring system in San Diego, California to illustrate the state of the art for existing practice. Led by its Metropolitan Planning Organization, the San Diego Association of Governments (SANDAG), and the California Department of Transportation (Caltrans), the San Diego region has developed one of the most sophisticated regional travel time monitoring systems in the United States. This system is based on an extensive network of sensors on freeways, arterials, and transit vehicles. It includes a data warehouse and software system for calculating travel times automatically. Regional agencies use these data in sophisticated ways to make operations and planning decisions.

Because this technical and institutional infrastructure was already in place, the team focused on generating sophisticated reliability use case analysis. The rich, multimodal nature of the San Diego data presented numerous opportunities for state of the art reliability monitoring, as well as challenges in implementing guidebook methodologies on real data.

The purpose of this case study was to:

- Assemble regimes and travel time probability density functions from individual vehicle travel times;
- Explore methods to analyze transit data from Automatic Vehicle Location (AVL) and Automated Passenger Count (APC) equipment;
- Demonstrate high-level use cases encompassing freeways, transit, and freight systems; and
- Relate travel time variability to the seven sources of congestion.

The *monitoring system* section further details the reasons for selecting San Diego as a case study and gives an overview of the region. It briefly summarizes agency monitoring practices, discusses the existing travel time sensor network, and describes the software system that the team used to analyze use cases. The section also details the development of travel time reliability software systems, and their relationships with other systems.

The section on *methodology* is the most experimental and least site specific. It is dedicated to an ongoing investigation, spread across all five case studies, to test, refine, and implement the Bayesian travel time reliability calculation methodology outlined in Chapter 3. For this section, the team is using, as appropriate, site data and other data in order to investigate this approach. The goal of each case study methodology section is to advance the team's understanding of the theoretical framework and practical implementation of the new Bayesian methodology.

Use cases are less theoretical and more site specific. Their basic structure is derived from the user scenarios described in Supplement D, which were derived from the results of a series of interviews with transportation agency staff regarding agency practice with travel time reliability.

Lessons learned summarizes the key findings from this case study, with regards to all aspects of travel time reliability monitoring: sensor systems, software systems, calculation methodology, and use. These lessons learned will be integrated into the final guidebook for practitioners.

MONITORING SYSTEM

Site Overview

The team selected San Diego as an exemplar of the leading edge of the state of the practice for using conventional monitoring systems within an urbanized metropolitan area. Led by its Metropolitan Planning Organization, the San Diego Association of Governments (SANDAG), and the California Department of Transportation, the San Diego region has developed one of the most sophisticated regional travel time monitoring systems in the United States. This system is based on an extensive network of sensors on freeways, arterials, and transit vehicles. It includes a data warehouse and software system for calculating travel times automatically. Regional agencies utilize this data in sophisticated ways to make operations and planning decisions.

In California, the San Diego Metropolitan Area encompasses all of San Diego County, which is approximately 4,200 square miles and the fifth most populous county in the United States. The county, bordered by Orange and Riverside Counties to the north, Imperial County to the east, Mexico to the south, and the Pacific Ocean to the west, contains over 3 million people. Approximately 1.3 million of these people live within the City of San Diego, with the rest concentrated within the southern suburbs of Chula Vista and National City, the beach-side cities of Carlsbad, Oceanside, and Encinitas, the northern, in-land suburbs of Escondido and San Marcos and the eastern suburb of El Cajon. The metropolitan area also includes significant rural areas within and to the east of the Coastal Range Mountains, with the Sonoran Desert and the Cleveland National Forest on the far eastern edge and the Anza-Borrego Desert State Park in the northeast corner of the county. The county has a large military presence, containing numerous Naval, Marine Corps, and Coast Guard stations and bases. Tourism also plays a major role in the regional economy, behind the military and manufacturing, particularly during the summer months.

Over the past several years, transportation agencies operating within the San Diego region have, through partnerships between SANDAG, Caltrans, local jurisdictions, transit agencies, and emergency responders, been updating and integrating their traffic management systems, as well as developing new systems, under the concept of Integrated Corridor Management (ICM). The goal of ICM is to improve system productivity, accessibility, safety, and connectivity by enabling travelers to make convenient and informed shifts between corridors and modes to complete trips. The partnering agencies selected I-15 from SR-52 in San Diego to SR-78 in Escondido as the corridor along which to implement an ICM pilot project using Federal ICM Initiative funding. A Concept of Operations document for this pilot project was completed in March of 2008, and San Diego was selected for the Demonstration Phase of the ICM Initiative early in 2010.

Because of this effort and others, San Diego has a sophisticated travel time monitoring software infrastructure. Among the systems that will share data as part of the planned Integrated Corridor Management System (ICMS) are the Advanced Transportation Management System (ATMS), Performance Measurement System (PeMS), Ramp Meter Information System (RMIS), Lane Closure System (LCS), the managed lane closure and congestion pricing systems on I-15, the Regional Arterial Management System (RAMS), and the Regional Transit Management System (RTMS).

Sensors

Freeway

The California Department of Transportation (Caltrans District 11) manages San Diego's freeway network. District 11 (D11) encompasses San Diego and Imperial County, though only the managed portion of the freeway system in San Diego County will be considered as part of this case study. Within San Diego County, D11 is responsible for 2,000 centerline miles of monitored freeways, 64 lane-miles of which are Managed HOV/HOT lane facilities.

A number of major interstates pass through the district, including Interstate 5, which passes through many major cities on the west coast between Mexico and Canada, Interstate 8, which connects Southern California with Interstate 10 in Arizona, and Interstate 15, which connects San Diego with Las Vegas. Within the county, I-5 connects downtown San Diego with the Mexican Border at Tijuana to the south, and the North County beach-side suburbs and Orange County to the north. I-8 connects the north part of the City of San Diego with El Cajon and the southern California desert. I-15 connects downtown San Diego with the inland suburbs of Rancho Bernardo and Escondido, then travels up through the Los Angeles suburbs in Riverside County. Other major freeways include Interstate 805, which parallels I-5 on the inland side between the Mexican border and its intersection with I-5 between La Jolla and Del Mar. State Route 163 connects I-5 in downtown San Diego with I-15 near the Marine Corps Air Station in Miramar. State Route 94 links I-5 downtown with eastern suburbs, paralleling I-8 to the south. State Route 78 is the major east-west freeway in North County, connecting Oceanside and Carlsbad with Escondido, and traveling further east into the mountainous regions of the county. A map of San Diego's freeway network is shown in Exhibit C2-1.

To monitor its freeways, District 11 has 3,592 ITS traffic sensors deployed at 1,210 locations that collect and transmit data in real-time to a central database. 2,558 of these sensors are in the freeway mainline lanes, 20 are in HOV lanes, and the rest are located at on-ramps, off-ramps, or interchanges. These sensors are a mixture of loop detectors and radar detectors. Approximately 90% of the ITS detection is owned by Caltrans, with the remainder owned by NAVTEQ/Traffic.com. San Diego County has had freeway detection in place since 1999, with the number of detectors steadily increasing over time.

Detectors are spaced relatively frequently on major freeway facilities. Most monitored freeways have an average detector station spacing of between ½ mile and 1 mile. The number and average spacing of detector stations for each monitored mainline facility in the County are indicated in Table C2-1.

Exhibit C2-1: San Diego Freeway Network

Table C2-1: San Diego County Freeway Detection

Freeway	Monitored Lane-miles	Detector Stations	Average Spacing	HOV
I5-N	61.8	98	0.65	X
I5-S	60.8	89	0.70	X
18-E	26.3	45	0.60	
18-W	26.3	46	0.60	
I15-N	39.1	50	0.80	X
I15-S	37.9	45	0.85	X
I805-N	28.7	49	0.60	
I805-S	28.7	46	0.60	
I905-W	3	2	1.50	
SR52-E	14.8	17	0.90	
SR52-W	14.8	16	0.90	
SR54-E	7	3	2.30	
SR54-W	6.8	3	2.30	
SR56-E	5.7	3	1.90	
SR56-W	5.7	3	1.90	
SR78-E	20.2	17	1.20	
SR78-W	20.2	23	0.90	
SR94-E	11.1	14	0.80	
SR94-W	11.6	20	0.60	
SR125-N	10.8	13	0.85	
SR125-S	10.7	13	0.80	
SR163-N	11.1	15	0.75	
SR163-S	11.1	15	0.75	

District 11 also owns and maintains almost 2,000 census count stations. All of these stations report data on traffic volumes and 20 additionally provide vehicle classification and weight information. These stations do not report conditions in real-time, but are obtained and input into the PeMS database via an offline batch process.

In San Diego County, real-time flow, occupancy, and- at some locations- speed data are collected in the field by controller cabinets wired to the individual sensors. Data are transmitted from these controller cabinets to the Caltrans District 11 Traffic Management Center (TMC) via a Front End Processor (FEP). The TMC's Advanced Transportation Management System software (ATMS) parses the raw, binary field data from the field and writes outputs into a TMC database. These values (measured flow and occupancy values for every 30-second time period at every detector) are then transmitted to the PeMS Oracle database in real-time via the Caltrans Wide Area Network (WAN). PeMS then performs a number of database routines on the data, including detector diagnostics, imputation, speed calculations, performance measure computations, and aggregation. These processing steps are fully described in Chapter 3 of the Guidebook.

Arterial

Although San Diego's arterial facilities are managed by the cities in which they are physically located, SANDAG assists these local agencies in implementing the Regional Arterial Management System, a region-wide traffic signal integration system that allows for inter-jurisdictional management and coordination of freeway/arterial interchanges. As part of a project to evaluate technologies for monitoring arterial performance, SANDAG installed an arterial travel time monitoring system along four miles of Telegraph Canyon Road and Otay Lakes Road between I-805 and SR-125 in Chula Vista, a suburb in San Diego's South Bay. The corridor has 18 sensor locations (9 in each direction of travel). The sensors deployed along this corridor are wireless magnetometer dots, which directly measure travel times by re-identifying unique vehicle magnetic signatures across detector locations. In order to read a vehicle's magnetic signature, the dots need to be deployed in series of five at each location. Consequently, a total of 90 wireless magnetometer sensors have been deployed along this corridor.

After a vehicle passes over a sensor location, each set of five sensors wirelessly transmits the vehicle's magnetic signature information to an access point on the side of the roadway. If the sensors are located further than 150 feet from the access point, a battery-operated repeater is needed to transmit the data from the sensor to the access point. The access point collects the sensor data then transmits it via Ethernet or a high-speed cellular modem to a data archive server in the TMC. At the TMC, the magnetic signatures are matched between upstream and downstream sensor stations and travel times are computed.

Transit

The largest share of San Diego County's transit service is operated by the San Diego Metropolitan Transit System (MTS). MTS operates bus and light rail service (through its subsidiary, San Diego Trolley) in 570 square miles of the urbanized area of San Diego, as well as rural parts of the East County, totaling 3,420 square miles of service area. To monitor its transit fleet, MTS has equipped over one-third of its bus fleet with Automatic Vehicle Location (AVL) transponders and over one-half of its fleet with Automated Passenger Count (APC) equipment. The AVL infrastructure allows for the real-time polling of buses to obtain real-time location and schedule adherence data. The APC data are not available in real-time, but can be used for off-line analysis to report on system utilization and efficiency.

Data Management

Freeway

The primary data management software system in the region is PeMS. All Caltrans districts use PeMS for data archiving and performance measure reporting. PeMS integrates with a variety of other systems to obtain traffic, incident, and other types of data. It archives raw data, filters it for quality, computes performance measures, and reports them to users through the web at various levels of spatial and temporal granularity. It reports performance measures such as speed, delay, percentage of time spent in congestion, travel time, and travel time reliability. These performance measures can be obtained for specific freeways and routes, and are also aggregated up to higher spatial levels such as county, district, and state. These flexible reporting options are supported by the PeMS web interface, which allows users to select a date range over

which to view data, as well as the days of the week and times of the day to be processed into performance metrics. Since PeMS has archived data for San Diego County dating back to 1999, it provides a rich and detailed source of both current travel times and historical reliability information.

In Southern California, PeMS obtains volume and occupancy data for every detector every 30 seconds from the Caltrans ATMS, which governs operations at the District TMCs. The ATMS is used for real-time operations such as automated incident detection and for handling special event traffic situations. ATMS data transmitted to the PeMS Oracle database supports the majority of transportation performance measures reported by PeMS and serves as the primary source of data for the travel time system validations discussed in this case study.

PeMS integrates, archives, and reports on incident data collected from two different sources: the California Highway Patrol (CHP) and Caltrans. CHP reports current incidents in real-time on its website. PeMS obtains the text from the website, uses algorithms to parse the accompanying information, and inserts it into the PeMS database for display on a real-time map, as well as for archiving. Additionally, Caltrans maintains an incident database, called the Traffic Accident Surveillance and Analysis System (TASAS), which links to the highway database so that incidents and their locations can be analyzed. PeMS obtains and archives TASAS incident data via a batch process approximately once per year. Incident data contained in PeMS has been leveraged to demonstrate use cases associated with how different sources of congestion impact travel time reliability.

PeMS also integrates data on freeway construction zones from the Caltrans Lane Closure System (LCS), which is used by the Caltrans districts to report all approved closures for the next seven days, plus all current closures, updated every 15 minutes. PeMS obtains this data in real-time from the LCS, displays it on a map, and lets users run reports on lane closures by freeway, county, district, or state. Lane closure data in PeMS was used in the validation of the use cases associated with how different sources of congestion impact travel time reliability.

Arterial

Arterial travel time systems are an emerging concept in San Diego. As described in the Sensors subsection, San Diego currently only has detection for arterial travel time support on one corridor in the suburb of Chula Vista. The system used to evaluate arterial travel times in San Diego is the Arterial Performance Measurement System (A-PeMS), an arterial extension of PeMS that collects and stores arterial data. A-PeMS receives a live feed of travel times and volume data from a server at Sensys Networks (the manufacturer of the arterial sensors deployed on this corridor) and stores them in the PeMS database. Within PeMS, these data are integrated with information on each intersection's signal timing, which allows for the computation of arterial performance measures. As part of the San Diego A-PeMS deployment, cycle-by-cycle timing plan information is parsed from time-of-day signal timing plans. A-PeMS can also integrate real-time signal timing cycle lengths and phase green times from traffic signal controllers. The performance reporting capabilities within A-PeMS are similar to those within PeMS. Users can view arterial-specific performance measures such as control delay and effective green time, as well as general performance measures such as travel times.

Outside of the reliability and performance monitoring aspects of arterial operations, the various agencies operating within San Diego County, led by SANDAG, are working toward development of a Regional Arterial Management System (RAMS). This system has relevance to this project since its signal timing plan data could eventually be used to support the widespread

monitoring of travel time variability on county arterials. This would facilitate a greater understanding of how different arterial facilities interact with one another, with transit service, and with freeway operations.

Transit

District 11 also uses a transit extension of PeMS, the Transit Performance Measurement System (T-PeMS), to obtain schedule, AVL, and APC data from its existing real-time transit management system, compute performance measures from this data, and aggregate and store them for further analysis.

METHODOLOGICAL ADVANCES

Overview

One objective of the case studies is to test and refine the methods developed in Phase 2 for defining and identifying segment and route regimes for freeway and arterial networks. The team's research to date has focused on identifying operational regimes based on individual vehicle travel times and determining how to relate these regimes to system-level information on average travel times. Since individual vehicle trip travel times on freeways are not available in the San Diego metropolitan region, data from the Berkeley Highway Laboratory (BHL) was used in this analysis.

Analysis Setting and Data

The Berkeley Highway Laboratory (BHL) is a 2.7-mile section of Interstate 80 in west Berkeley and Emeryville. The BHL includes fourteen surveillance cameras and sixteen directional dual inductive loop detector stations dedicated to monitoring traffic for research purposes. The sensors are a unique resource because they provide individual vehicle measurements. The system collects individual vehicle actuations from all 164 loops in the BHL every 1/60th of a second and archives both the actuation data and a large set of aggregated data, such as volumes and travel times. The loop data collection system is currently generating approximately 100 megabytes of data per day. A suite of loop diagnostic tests has been developed over the last 2 years, which continuously tests the data stream received from the loops and archives the test results.

The BHL loop data are unique because it provides event data on individual vehicle actuations, accurate to 1/60th of a second. Most other loop detector systems collect only aggregated data over periods of 20 seconds or longer. Collecting the individual loop actuations allows the generation of data sets which are not found elsewhere, such as vehicle stream data, which can be used for headway studies, gap analysis, and merging studies. The BHL loops also provide individual vehicle length measurements, allowing for the classification of freeway traffic. Rich data sets of individual vehicle travel times are also available on the BHL, stemming from research that developed a vehicle re-identification algorithm to calculate travel times between successive loop stations. A final benefit of the BHL data is that the corridor was temporarily instrumented with two Bluetooth reader stations (BTRs) along eastbound I-80. These stations record the timestamps and MAC addresses of Bluetooth devices in passing vehicles.

Travel times can be derived from the matching of MAC addresses between two readers. A map of the BTR locations is shown in Exhibit C2-2.

Analysis was performed on a day's worth of BHL data, collected on Tuesday, 11/16/2010. One data file was obtained for each of the two BTRs, with each file containing every MAC address captured by that sensor on that day. Some MAC address IDs were repeated within the file, due to the fact that passing devices can be sampled multiple times by a single reader. Since the BTRs are located along the eastbound side of the freeway, the majority of MAC address re-identifications were for eastbound traffic, though some westbound vehicles were also captured. There was a one-hour gap in the data between 4:30 AM and 5:30 AM due to a bug in the BHL database. Additionally, some of the initial time-stamps in the file for the midnight hour were negative, possibly due to clock error. Six files of loop detector actuation data were also obtained. Together, these files contain all of the vehicles records at all of the BHL stations on this day.

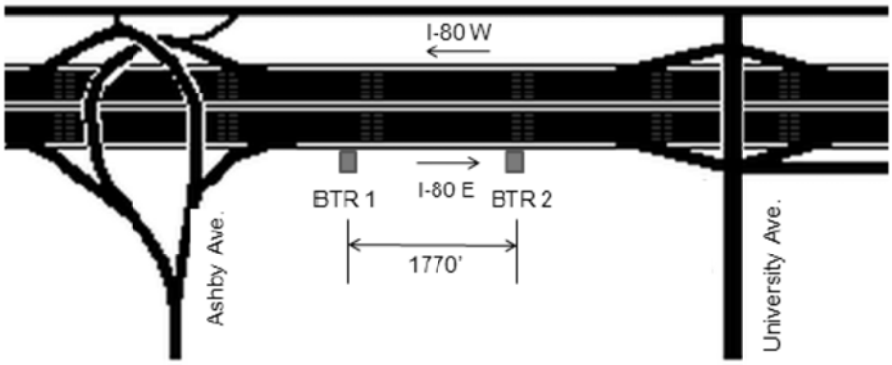

Exhibit C2-2: Bluetooth Reader Locations, I-80E

Methodological Use Cases

Overview

Five concepts are important in this analysis:

- **Concept 1**: Regardless of the data source, the methodology must always generate a full travel time probability density function (PDF). All reliability measures can be generated from the PDF.
- **Concept 2**: We need to distinguish between two types of PDFs:

1) Those pertaining to the distribution of travel times derived from *individual travelers* along a segment or route; this accounts for travel time variability (for a route or a segment) among individual travelers and over time.
2) Those pertaining to the distribution of the mean travel time along a segment or route; this accounts for variations in the mean travel time (for a segment or a route) over time.

- **Concept 3**: It is desirable (and we think possible) to generate individual traveler travel time PDFs directly from some data sources (for example, Bluetooth or GPS) and indirectly from others (for example, loop detectors or video).

- **Concept 4**: The travel time PDFs can be reasonably characterized by a Shifted Gamma Distribution with parameters (α, β, δ) as follows:
  1) α: the shape of the density function, with $\alpha > 1$ implying that it has a "log-normal" type shape
  2) β: the spread in the density function, with larger values implying more spread
  3) δ: the offset of the "zero-point" from the value of zero, or, in this context, the smallest possible travel time
- **Concept 5**: A finite number of traffic states, or regimes, describe all possible travel time PDFs for a route or a segment. Regime PDFs can be continuously updated using real-time data.)

For use cases that serve motorists in need of traveler information, the development of reliability statistics from individual travel time PDFs is ideal. The use cases examined in this chapter are shown in Table C2-2. They are intended to provide information on recommended trip start times (ST) for constrained trips, subject to certain arrival time performance criteria.

Table C2-2: Use Cases MC1, MC2, and MC3

Use Case	Description	What is known?	Desired Deliverable	Metrics
MC1	User wants to know *in advance* what time to leave for a trip and what route to take—planning level analysis	Origin position, Destination position, Day of Week, Desired Arrival Time at Destination	A list of alternative routes, their mean travel time and required start time on each route to ensure meeting arrival time 95% of the time	Average O-D travel time by path, planning time
MC2	User wants to know *immediately* what route to take and time to leave for a trip to arrive on time at destination—real time analysis	Origin position, Destination position, Desired Arrival Time at Destination	A ranked list of alternative routes, their mean travel time based on current conditions and required start time on each route to ensure meeting arrival time 95% of the time	Average O-D travel time by path, planning time
MC3	User wants to know *the extra time needed* for a trip to arrive on time at destination with a certain probability	Origin position, Destination position, Prob. arriving on time, day of week, time of day	Map of the route with lowest travel time meeting the threshold, the route average travel time, selected % travel time and buffer time.	Buffer time, % travel time, average travel time for O-D pair

In this discussion, the analysis is focused on developing the probability density function of travel times for those individual travelers who depart an origin in a pre-specified time interval in order to meet a pre-specified arrival time at the destination within an acceptable – and specified- level of risk. The size of the time interval is selected in such a way as to ensure stationary travel conditions within the interval as well as to capture a sufficient sample of travelers to characterize or update the developed travel time distribution.

It is hypothesized that the route travel time distribution can be "stitched" from the distribution of segment travel times which make up the route. This hypothesis is still subject to testing and validation using field data. Furthermore, it is assumed that there is a finite number of travel time PDFs (or regimes) that can fully characterize the travel time distribution between an origin and a destination on a given route over a full year. Exhibit C2-3 illustrates an example that uses four PDFs and a transition PDF (labeled T) , where each cell color represents a unique travel time regime based on historical travel time data for a given origin-destination pair on a given route.

It is further hypothesized that the individual auto travel times on links or routes can be characterized by a 3-parameter shifted Gamma distribution (α, β, and δ) of the form:

$$g_{\alpha,\beta,\delta}(t) = \frac{\beta^{\alpha}}{\Gamma(\alpha)}(t-\delta)^{\alpha-1}e^{-\beta(t-\delta)} \quad for t \geq \delta, o.e.w$$

For α=1, the Gamma distribution degenerates into the shifted exponential distribution. Exhibit C2-4 shows a diagram of the distribution for α > 1.0. There is a unique set of distribution parameters associated with each origin-destination pair, route, and PDF regime.

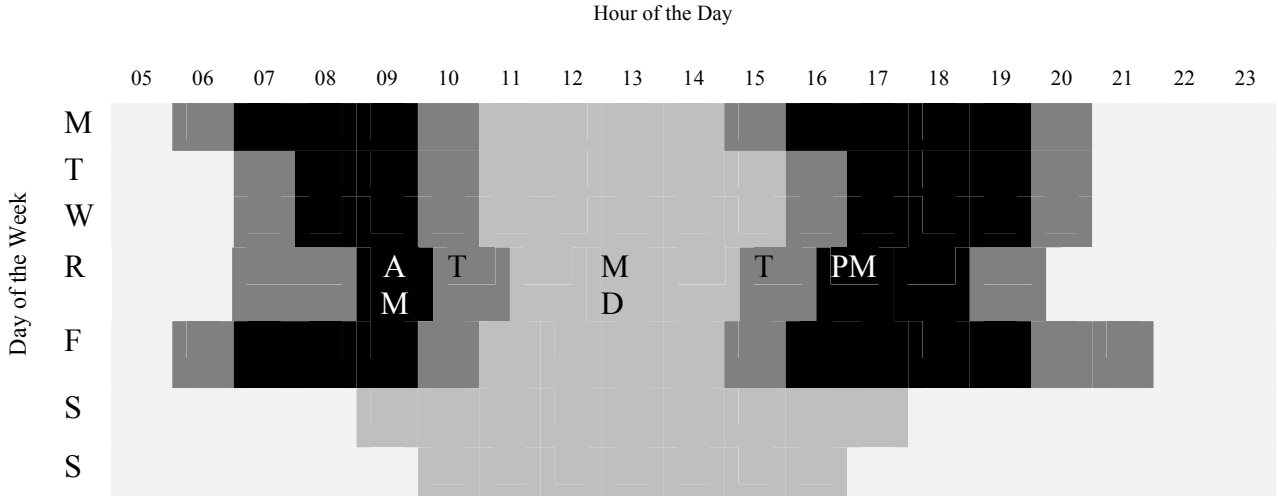

Exhibit C2-3: Historical route travel time PDFs by time of day and day of week

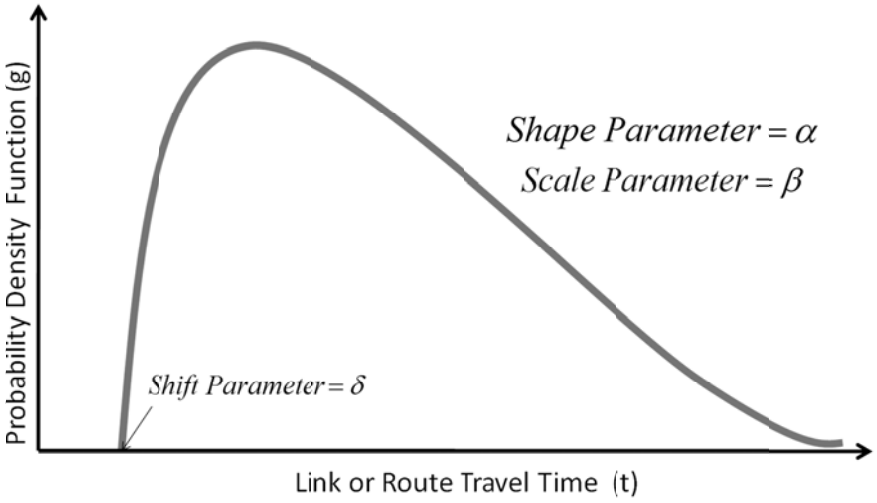

Exhibit C2-4: Shifted gamma distribution of travel times

Use Case MC1: User wants to know in advance what time to leave for a trip and what route to take. The procedure for validating use case MC1 is depicted in Exhibit C2-5. The top-right corner represents user-driven input, such as origin-destination (O-D) selection, desired arrival time at destination, and possible routes to be evaluated. The top-left corner represents field data collection of travel times to develop and update off-line historical travel times PDFs which follow the shifted Gamma distribution described in the previous section. The bottom section represents the actual algorithm to determine the computed user start time (ST) in order to meet the desired arrival time (DAT) criterion.

The outcomes, shown in the table in Exhibit C2-5, match the use case MC1 results requirement specified in Chapter 3 of the Guidebook, which is to generate "… A list of alternative routes that displays the required start time to arrive on-time 95% of the time and the required start time based on the average travel time". Based on this example, the entry time PDF consistent with the desired arrival time (DAT) of 8:40 AM is the 8:00-10:00 entry time.

An example application of the procedure using hypothetical travel time parameter values is shown in Exhibit C2-6. The procedure works as follows:

- User enters origin, destination, and a DAT of 8:40 at the destination on a Thursday.
- The user or the system identifies (or retrieves from a route library) a finite number of routes connecting the input O-D (or nearby locations). Let the first route be labeled Route 1.
- The system identifies the relevant time-dependent PDF (the AM peak) consistent with the user-inputted DAT and DOW. It represents all travel times for entry times between 8:00 AM and 10:00 AM on Thursdays.
- Based on the retrieved PDF, achieving a 95% on-time arrival requires a planned 30 minute travel time, compared to the average travel time of 23 minutes.
- Thus the recommended start time ST is 8:10 AM. Other DAT scenarios and outcomes are also shown in the table in Exhibit C2-5.

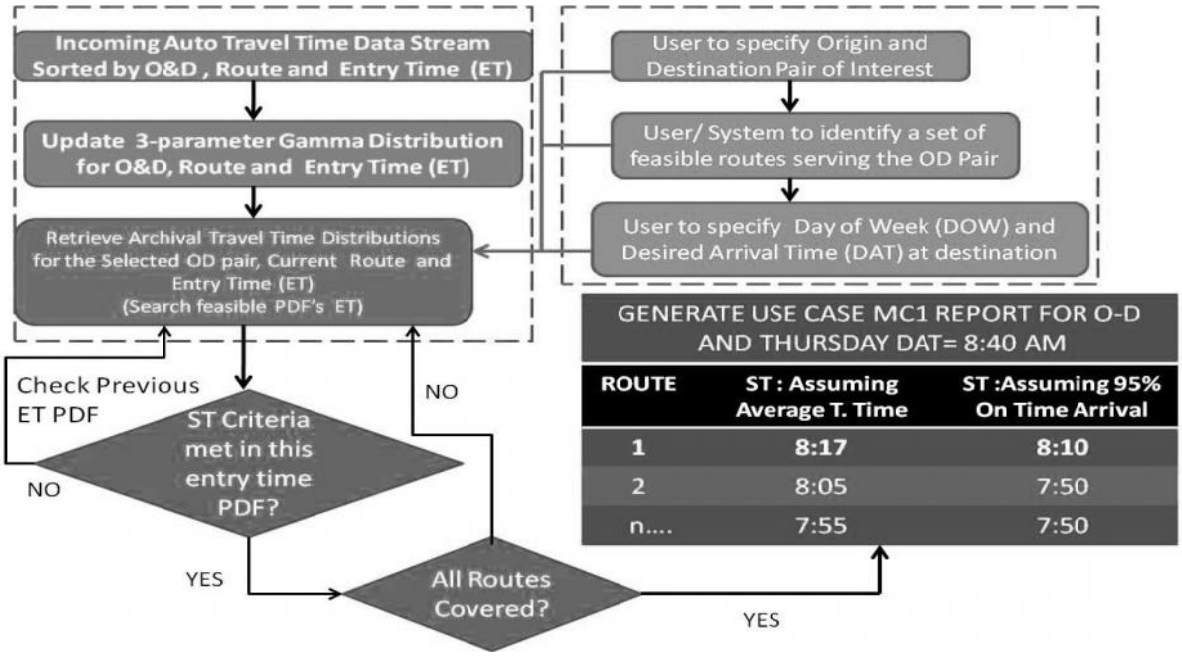

Exhibit C2-5: Validation process for Use Case MC1

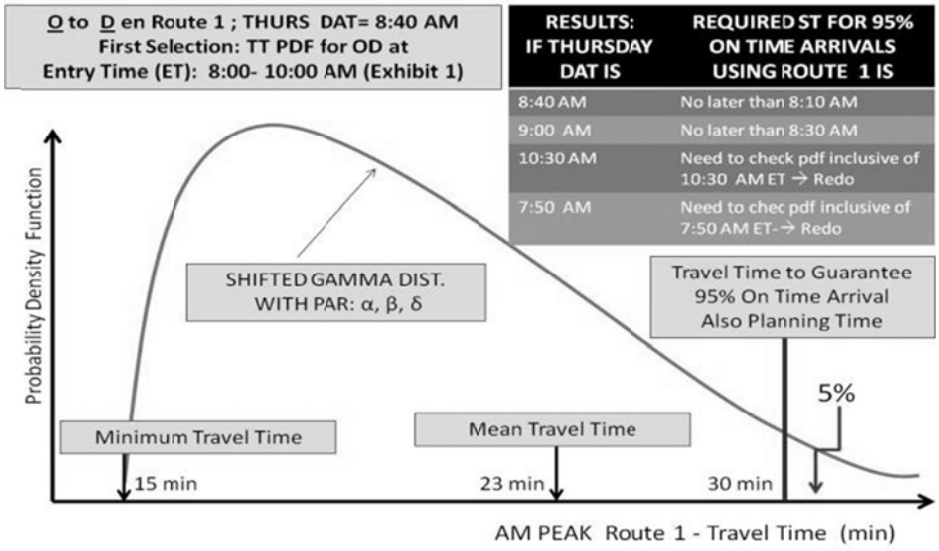

Exhibit C2-6: Example application of Use Case MC1

Use Case MC3: User wants to know the extra time needed for a trip to arrive on time at destination with a certain probability. This use case represents a simple variation of use case MC1 and is therefore discussed before the real-time use case, MC2. Here, the user is interested in identifying, for a known O-D, DAT, and DOW, a route, average travel time (AT) and planned travel time (PT) that will ensure his or her on-time arrival R% of the time. The algorithm for MC1 is adjusted slightly to meet these new requirements, as shown in Exhibit C2-7. The hypothesized PDFs for the two candidate routes are shown in Exhibit C2-8. These are designed to highlight the contrast between a shorter route (Route 2) and a more reliable route (Route 1). In this case, the system would recommend the selection of Route 1 and a departure time of no later than 8:44 AM in order to guarantee arrival at the destination by 8:40 AM with

90% certainty. The user would have to depart 10 minutes earlier on Route 2 to achieve the same probability of on-time arrival. This is confirmed by comparing the buffer times between the two routes.

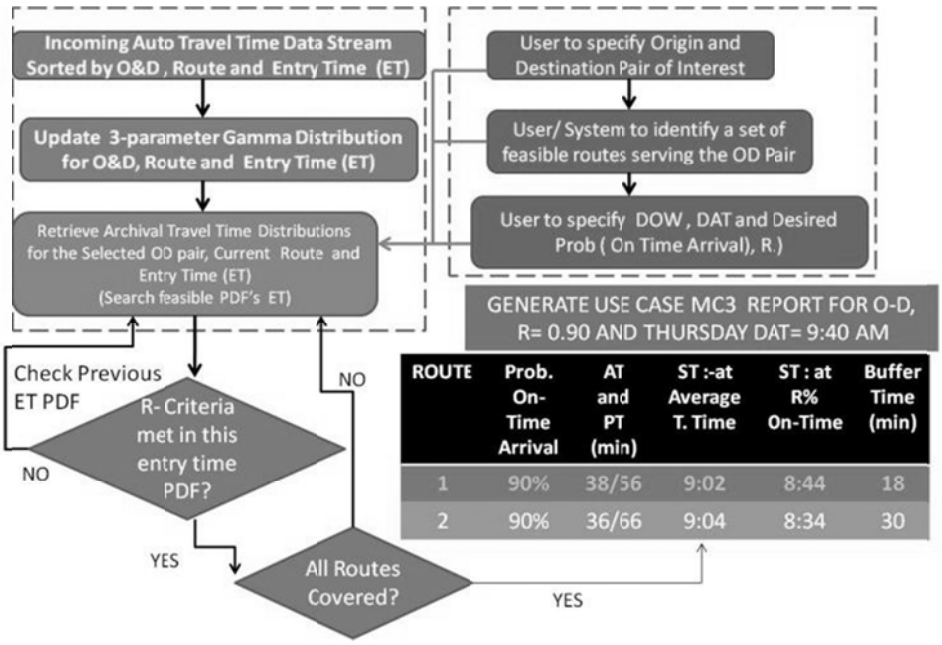

Exhibit C2-7: Validation process for Use Case MC3

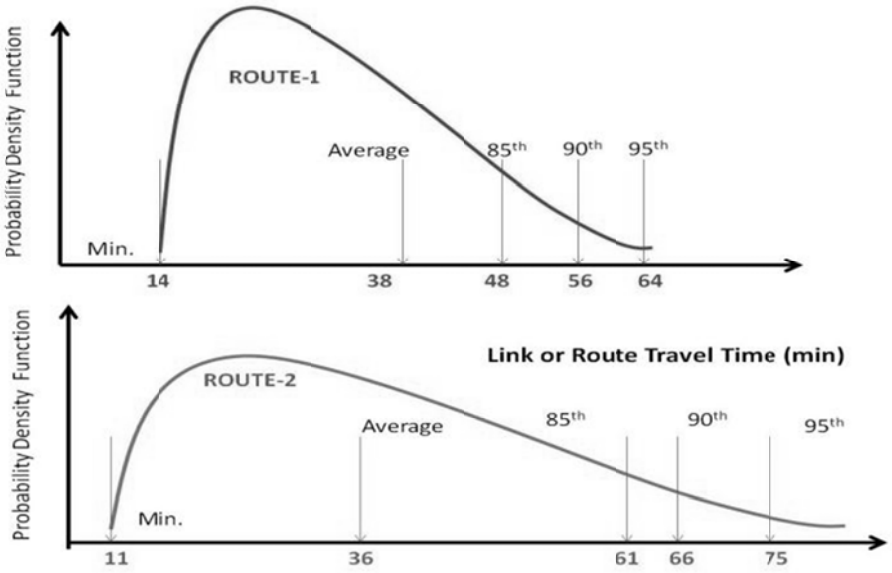

Exhibit C2-8: Illustration of a reliable route PDF (top) and a faster average route PDF (bottom)

Use Case MC2: User wants to know immediately what route to take and what time to leave for a trip to arrive on time at a destination. This use case is different and much more challenging to demonstrate than MC1 or MC3. It also represents the application with the highest utility from the driver's perspective since it will provide real-time information on the

recommended trip start time, including the effects of incidents or other events not explicitly accounted for in historical travel time PDFs. The principal issue, therefore, is how to combine the historical and real-time data streams in order to provide up-to-date travel time estimates and predictions based on current conditions. As an example, during major weekend road construction projects, the more accurate distribution may be the weekday AM peak profile, rather than the historical weekend travel time PDF.

Several stipulations are important to note:

- It is possible that there are no feasible solutions to the current user request. A departure at the earliest departure time may not guarantee the user's DAT at the specified probability R on some or all of the feasible routes.
- While historical PDFs are still important, they are not appropriate for use in a real time context. The system must be able to detect which PDF regime each link or route is operating in, based on the real-time data stream.
- The PDF regime selection process is akin to the "plan selection" algorithm that is used in many urban traffic signal control systems. Those algorithms collect traffic data (typically key link volumes and occupancies) to be matched with the signal plans most appropriate for the collected data patterns.
- In a real time context, where computational speed is of the essence, the number of PDFs to be considered should be kept to a minimum. Each link or route could theoretically be considered to operate in four regimes: uncongested, transition from uncongested to congested, congested, and transition from congested to uncongested.

The procedure for Use Case MC2 is shown in Exhibit C2-9. It assumes that there are three feasible alternate routes, and that the earliest departure time is 8:15 AM, while the DAT is 9:40 AM. The system checks which of the routes is feasible, and determines the required start time assuming average and 95th percentile travel times. In this case, Route 3 is deemed infeasible, while Routes 1 and 2 are both feasible. Work is underway to apply Bayesian techniques to match real-time travel data to historical regime-based PDFs and to develop the simplified four PDF regimes described earlier in this section.

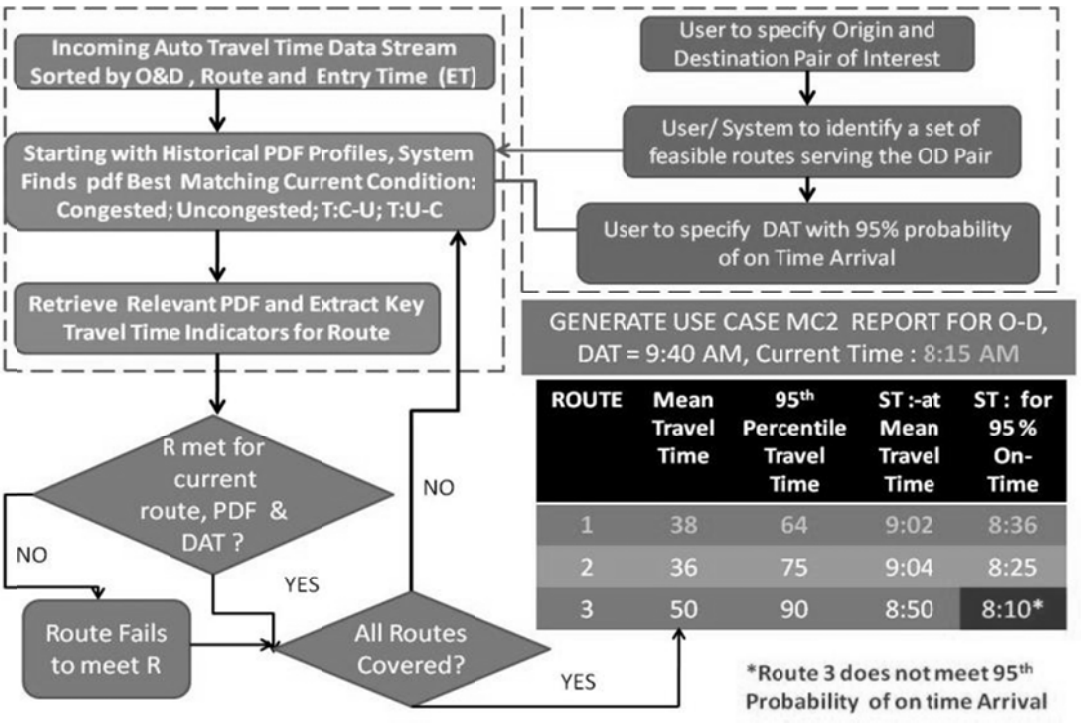

Exhibit C2-9: Validation process for Use Case MC2

Route Selection Criteria

An interesting byproduct of the use case analyses is the possibility of developing additional route selection criteria that can account for the differential utilities of early and late arrivals. Thus far, the selection between routes has been made on the basis of the route yielding the latest trip start time while ensuring a pre-specified on time arrival probability (for example, Route 1 in Exhibit C2-9). Specifying different penalty functions for late and early arrivals could change the selection.

Analysis of Bluetooth Travel Times

To support the methodologies presented in use cases MC1, MC2, and MC3, Bluetooth data from the BHL was analyzed to see what could be learned about individual vehicle travel times and the probability density functions.

The raw data were filtered to remove MAC addresses with six or more timestamps on either reader. Contiguous timestamps from the same vehicle were averaged to obtain an estimate of when the vehicle was adjacent to the sensor. The filtering process resulted in a data set of 5,028 travel time measurements. These were then filtered a second time to remove observations where the speed between the readers was below 5 mph. This resulted in 5,012 final measurements. These travel times are plotted in Exhibit C2-10 and Exhibit C2-11.

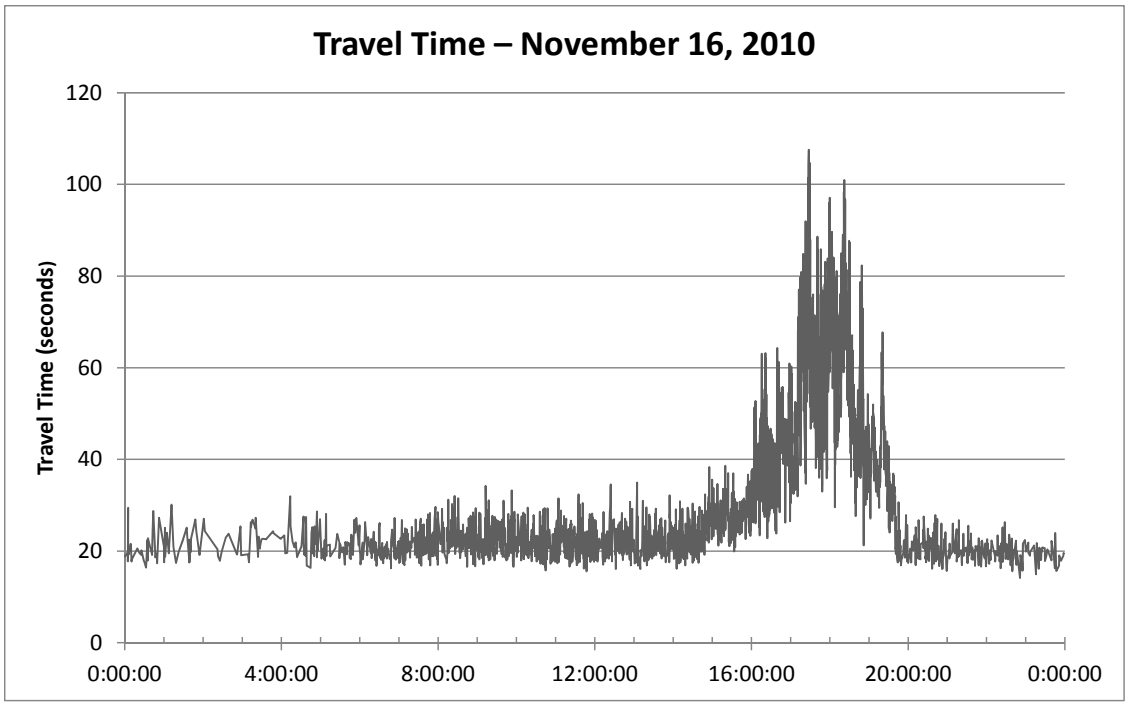

Exhibit C2-10: BHL Bluetooth-measured travel times, 11/16/2010

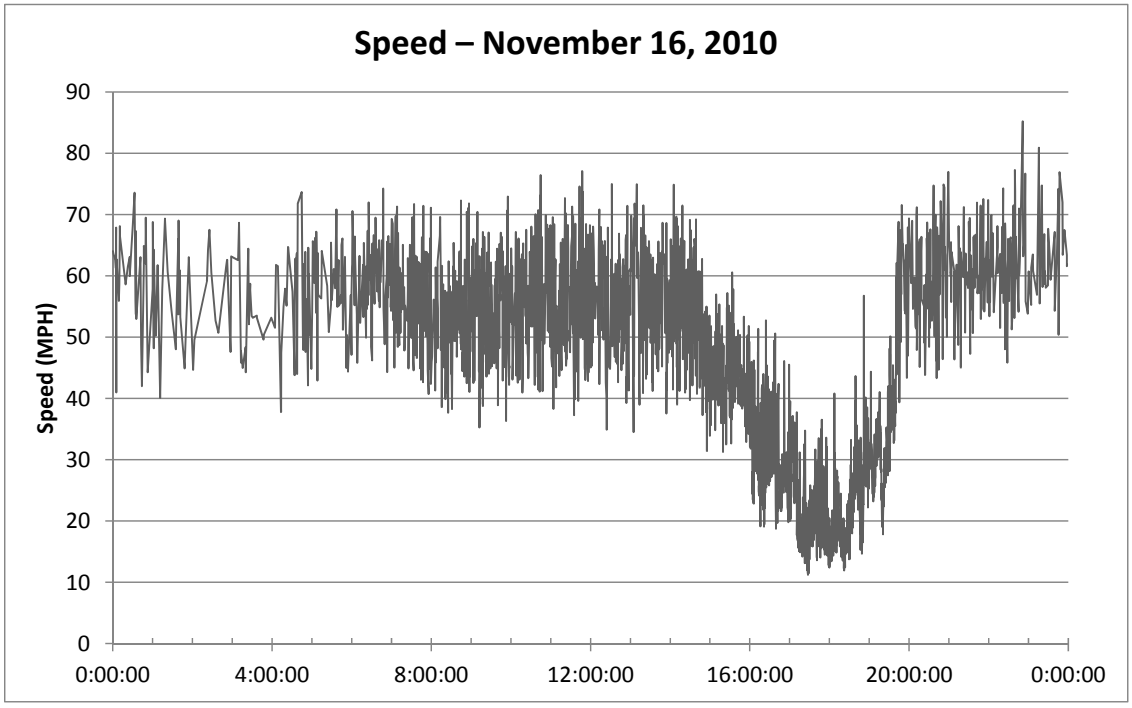

Exhibit C2-11: BHL Bluetooth-measured speeds, 11/16/2010

By inspection, three time periods of operative regimes were identified as follows:

- Free flow: 0:00:00-14:30:00 and 19:45:00-23:59:59
- Transition: 14:30:00-15:45:00 and 19:30:00-19:45:00
- Congested: 15:45:00 -19:30:00

The resulting distribution of the Bluetooth travel time observations is shown in Table C2-3.

The data were then analyzed using EasyFit software to see how different probability density functions fit the data and to estimate the parameters for each density function.

Table C2-4, Table C2-5, and Table C2-6 present the goodness of fit results down to the 3-parameter Gamma distribution (Gamma(3p)), sorted by the Anderson-Darling statistic. Exhibit C2-12, Exhibit C2-13, and Exhibit C2-14 show the resulting plots of the Gamma(3p) density functions. The Gamma(3p) fits relatively well for the Free Flow and Congested conditions. It is likely that there will be multiple transition regimes, and Gamma(3p) fit may be improved for stratified transition regimes.

Table C2-3: Bluetooth data regime classifications

Category	Flag	Observations
Free flow	1	2679
Transition	2	484
Congested	3	1849

Table C2-4: Goodness-of-fit results for the free-flow regime

Distribution	Kolmogorov Smirnov		Anderson Darling		Chi-Squared	
	Statistic	Rank	Statistic	Rank	Statistic	Rank
Pearson 5 (3P)	0.01352	1	1.0731	1	6.7917	1
Pearson 6 (4P)	0.01377	2	1.0875	2	7.2844	2
Dagum	0.01795	4	1.4442	3	13.684	3
Burr (4P)	0.02118	6	2.1326	4	22.291	5
Gen. Logistic	0.01975	5	2.3254	5	23.294	6
Log-Logistic (3P)	0.02309	9	2.4624	6	25.444	8
Frechet (3P)	0.02139	7	2.726	7	23.419	7
Gen. Extreme Value	0.02174	8	2.9185	8	27.343	9
Burr	0.02749	11	3.748	9	30.88	10
Lognormal (3P)	0.0172	3	5.798	10	16.413	4
Frechet	0.03534	15	7.4445	11	44.274	13
Gen. Gamma (4P)	0.02908	12	11.258	12	51.262	14
Inv. Gaussian (3P)	0.03043	13	11.749	13	36.427	11
Fatigue Life (3P)	0.03055	14	11.915	14	38.129	12
Log-Logistic	0.04617	19	12.611	15	117.0	18
Pearson 5	0.03864	17	13.959	16	69.484	15
Gamma (3P)	0.03686	16	18.252	17	84.176	16

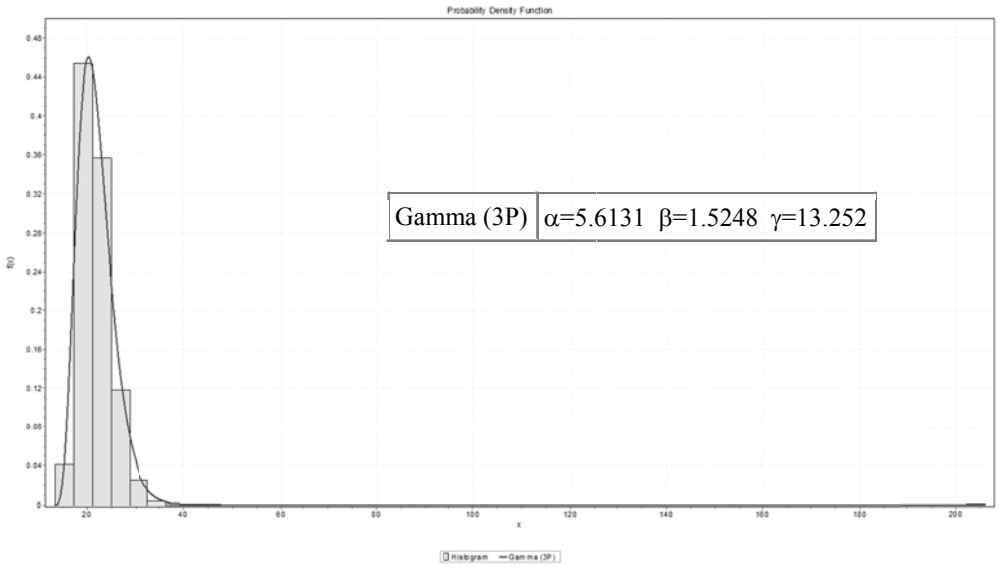

Exhibit C2-12: 3-Parameter Gamma distribution for the free-flow regime

Table C2-5: Goodness-of-fit results for the transition regime

Distribution	Kolmogorov Smirnov		Anderson Darling		Chi-Squared	
	Statistic	Rank	Statistic	Rank	Statistic	Rank
Burr	0.02676	1	0.81555	1	16.645	1
Burr (4P)	0.02709	2	0.82053	2	19.224	2
Johnson SU	0.03208	4	0.95065	3	22.543	10
Dagum (4P)	0.03373	6	0.97157	4	21.446	6
Dagum	0.03512	7	1.0092	5	21.472	7
Gen. Extreme Value	0.03317	5	1.0168	6	22.294	8
Frechet	0.02965	3	1.058	7	20.188	5
Frechet (3P)	0.03808	10	1.1188	8	22.351	9
Log-Logistic (3P)	0.03514	8	1.1732	9	19.289	3
Gen. Logistic	0.03904	12	1.2923	10	19.714	4
Pearson 5 (3P)	0.04297	13	1.3807	11	23.961	11
Pearson 6 (4P)	0.04478	14	1.5089	12	25.479	12
Lognormal (3P)	0.05205	15	2.0513	13	29.111	13
Inv. Gaussian (3P)	0.05956	16	2.5343	14	33.83	14
Fatigue Life (3P)	0.06274	18	2.8342	15	37.124	16
Gen. Gamma (4P)	0.06117	17	3.0654	16	36.792	15
Log-Pearson 3	0.03856	11	5.3555	17	N/A	
Pearson 5	0.08043	21	5.4889	18	47.002	17
Wakeby	0.03568	9	5.4998	19	N/A	
Gamma (3P)	0.08052	22	5.6067	20	53.894	20

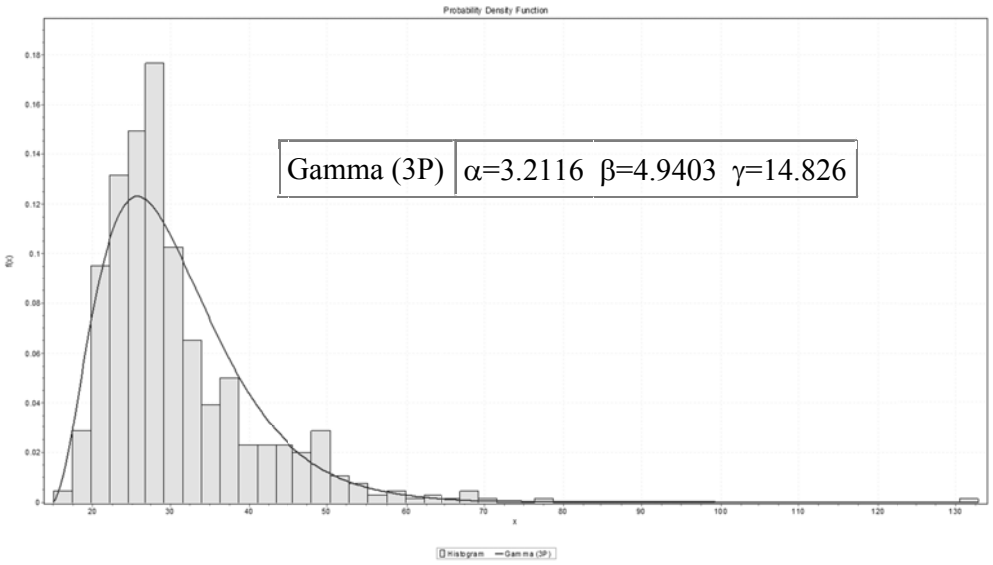

Exhibit C2-13: 3-Parameter Gamma distribution for the transition regime

Table C2-6: Goodness-of-fit results for the congested regime

Distribution	**Kolmogorov Smirnov**		**Anderson Darling**		**Chi-Squared**	
	Statistic	**Rank**	**Statistic**	**Rank**	**Statistic**	**Rank**
Fatigue Life (3P)	0.02266	3	0.9031	1	13.229	4
Inv. Gaussian (3P)	0.02297	4	0.94129	2	13.141	3
Gamma (3P)	0.02734	10	1.0359	3	13.408	5

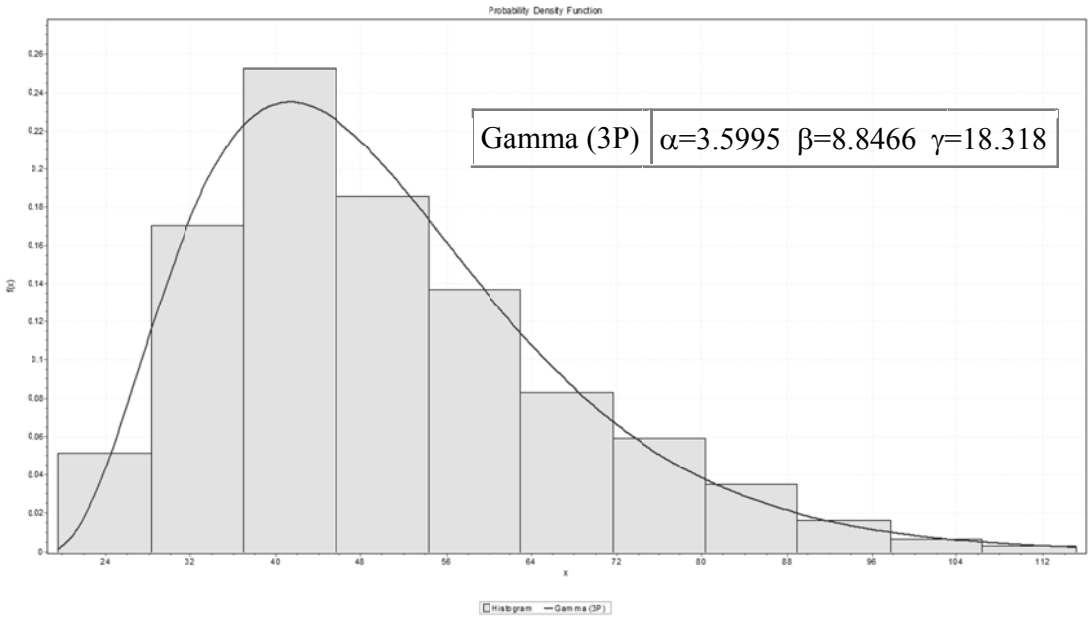

Exhibit C2-14: 3-Parameter Gamma distribution for the congested regime

The three PDFs are superimposed in Exhibit C2-15. It is apparent that the free-flow PDF has a lower mean travel time, a smaller standard deviation, and the lowest 95th percentile value. The congested PDF is at the other end of this extreme, with the largest mean, the largest standard deviation, and the highest 95th percentile value. Not unexpectedly, the PDF for the transition regime lies between these two. The numerical values are presented in Table C2-7.

Table C2-7: 3-Parameter Gamma distribution means, standard deviations, and 95th percentiles

Condition	Mean (sec)	StdDev (sec)	95th Percentile (sec)
Uncongested	21.8	3.57	28.3
Transition	30.4	9.11	47.7
Peak	50.0	17.0	83.5

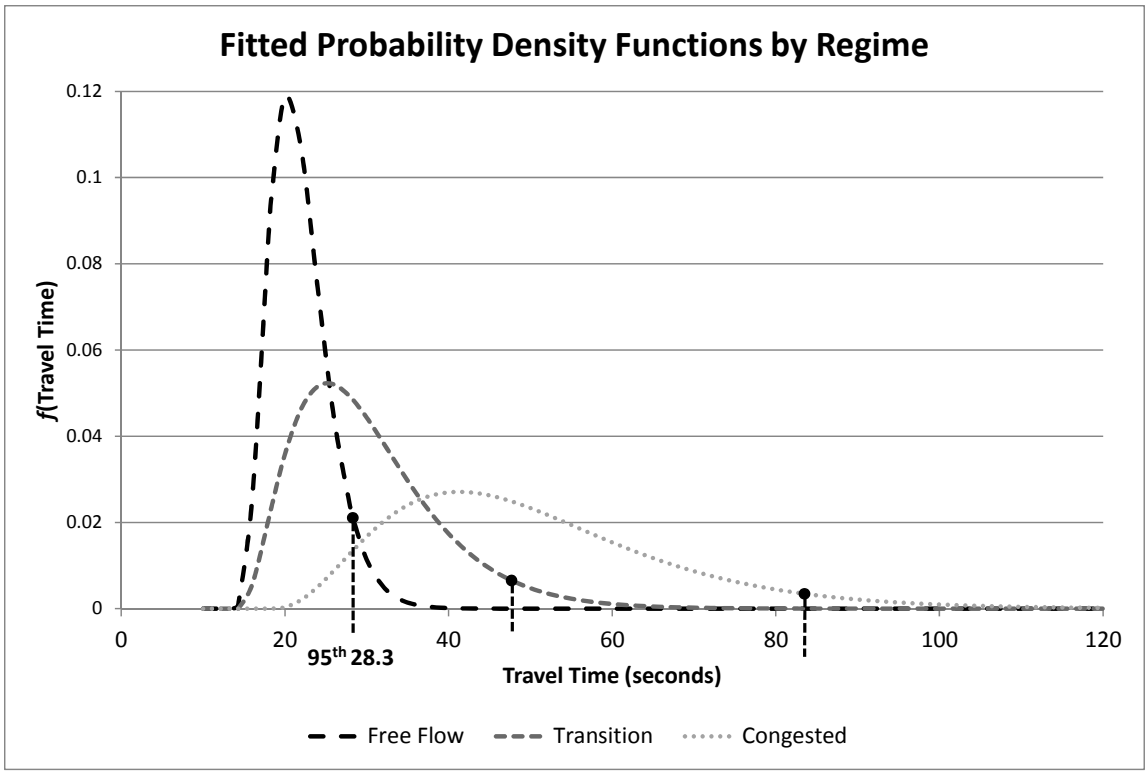

Exhibit C2-15: 3-Parameter Gamma distributions for all three regimes

Conclusions

This analysis examined data from the Berkeley Highway Lab to see if operative regimes for individual vehicle travel times can be identified from Bluetooth data. The research team concluded that this can, indeed, be done. Based on more than 5,000 observations of individual travel times, three different regimes can be identified: (1) off-peak or uncongested; (2) peak or congested; and (3) transition between congested and uncongested. All three can be characterized by 3-parameter Gamma density functions. More specifically, the PDF for the free flow condition has the lowest mean, the smallest standard deviation, and the lowest 95th percentile. The congested PDF is at the other extreme; and the transition PDF is in between.

Further investigation is needed into the individual vehicle PDFs and the parameters that describe them, but the efficacy of the concepts seems sound. Two issues that need to be explored in the very near future are: (1) how the PDFs for individual vehicle travel times relate to mean travel times (for example, those computed from loop detectors) during the same time periods and (2) whether there are ways to retrieve information from loop detectors that would be help to infer the PDFs that describe individual vehicle travel times.

USE CASE ANALYSIS

Overview

Chapter 4 of the guidebook and Supplement D: Use Case Demonstrations present dozens of use cases intended to satisfy the myriad ways that different classes of users can derive value from a reliability monitoring system. For the San Diego case study, a number of these use cases

were combined to form six high-level use cases that broadly encompass the types of reliability information that users are most interested in and that were suited for validation using the San Diego data sources. These six use cases, their primary user groups, and the guidebook use cases that they encompass, are shown in Table C2-8.

Table C2-8: Demonstrated use cases in San Diego

Use Case	Primary users	Guidebook sub-use cases
Freeways		
Conducting offline analysis on the relationship between travel time variability and the seven sources of congestion	Planners and Roadway Managers	MC4, PE1, PE2, PE3, PE4, PE5, PE11, PP1
Using planning-based reliability tools to determine departure time and travel time for a trip	Motorists	MC1, MC2, MC3
Combining real-time and historical data to predict travel times in real-time	Operations Managers	MM1, MM2, MC5
Transit		
Using planning-based reliability tools to determine departure time and travel time for a trip	Transit Riders	TP1, TS2, TO2, TC4
Conducting offline analysis on the relationship between travel time variability and the seven sources of congestion	Transit Planners and Managers	PE1, PE2, PE3, PE4, PE5, PE11, PP1
Freight		
Using historical data to evaluate freight travel time reliability	Drivers and Freight Carriers	FP1, FP3, FP4, FP6

In line with the use case divisions shown in the table, the remainder of this chapter is broken up into three sections: Freeways, Transit, and Freight. Each section presents the analytical results of validating the use cases with reliability monitoring system data and methods.

Freeways

Use Case 1: Conducting offline analysis on the relationship between travel time variability and the seven sources of congestion

Summary. This use case aims to quantify the impacts of the seven sources of congestion: (1) incidents; (2) weather; (3) lane closures; (4) special events; (5) traffic control; (6) fluctuations in demand; and (7) inadequate base capacity, on travel time variability. To perform this analysis, methods were developed to create travel time probability density functions (PDFs) from large data sets of travel times that occurred under each event condition. From these PDFs, summary metrics such as the median travel time and planning travel time were computed to show the variability impacts of each event condition.

Users. This use case has broad applications to a number of different user groups. For planners, knowing the relative contributions of the different sources of congestion toward travel time reliability helps them to better prioritize travel time variability mitigation measures on a facility-specific basis. For example, if unreliability on a particular route is predominantly caused by the frequent occurrence of incidents, planners may want to consider measures such as freeway service patrol tow truck deployments to help clear incidents faster. If unreliability on a route has a high contribution from special event traffic impacts, planners may want to consider providing better traveler information before events to inform travelers of alternate routes.

The outputs of this use case are also of value to operators, providing them with information on the range of operating conditions that can be expected on a route given certain source conditions. Knowing the historical impacts of the different sources of congestion helps operators better manage similar conditions in real-time by, for example, changing ramp metering schemes to mitigate congestion or posting expected travel times on variable message signs. It is important for operators to have outputs from this use case at a time-of-day specific level. For example, on some facilities, incidents may significantly impact reliability during one or more peak hours, but may have little impact during the midday due to lower baseline traffic volumes. On some facilities, weather may have a major impact at all times of the day, since all vehicles may need to slow to safely travel in the conditions. Understanding the time-dependency of variability impacts would help operators more effectively manage events as they occur.

Finally, the outputs of this use case have value to travelers, by providing better predictive travel times under certain event conditions that could be posted in real-time on variable message signs or on traveler information websites. This information would help users better know what to expect during their trip, both during normal operating conditions and when an external event if occurring.

Sites. Two routes were selected for the evaluation of this use case, to highlight the varying contributions of congestion factors to travel time reliability across different facilities, days of the week, and times of the year. These routes are shown in Exhibit C2-16. The first route analyzed is a 10 mile stretch of westbound Interstate-8 beginning at Lake Murray Boulevard in the eastern suburb of La Mesa and ending at Interstate-5 north of the San Diego International Airport. This route was selected because it provides access to Qualcomm Stadium, located at the major interchange of I-8 and I-15, which hosts San Diego Chargers football games as well as college football bowl games, concerts, and other events. Because this route is a major commute route, the impacts of the sources on travel time variability were investigated for weekdays between the months of November and February (when Qualcomm Stadium regularly hosts events and when San Diego experiences the most inclement weather).

The second route is a 27 mile stretch of northbound I-5 beginning just south of the I-805 interchange in San Diego and ending north of SR-78 in the northern suburb of Oceanside. This route was selected because it has a significant amount of congestion and incidents, and it sees special event traffic impacts during the summer months due to the San Diego County Fair and Del Mar horse races. The route also has significant traffic congestion on weekends. For this reason, travel time variability and its relationship with the sources of congestion were evaluated over a year-long period on Saturdays and Sundays.

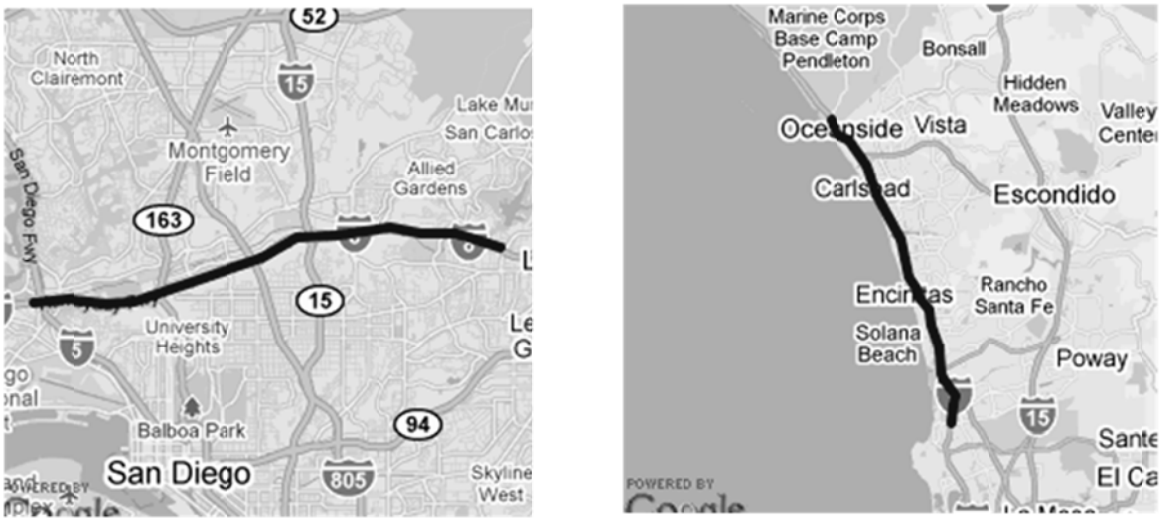

Exhibit C2-16: Freeway Use Case 1 routes

Methods. These routes were analyzed to determine the travel time variability impacts caused by five sources of congestion: (1) incidents; (2) weather; (3) special events; (4) lane closures; and (5) fluctuations in demand. Traffic control contributions were not investigated as ramp metering location and timing data could not be obtained. The impacts of inadequate base capacity were also not considered due to the difficulty of quantifying this factor.

For each route, five-minute travel times were gathered from PeMS for each day in the time period of analysis (four months of weekdays for the westbound I-8 route and one year of weekends for the northbound I-5 route). To ensure data quality, five-minute travel times computed from more than 20% imputed data were discarded from the data set.

To link travel times with the source condition active during their measurement, each 5-minute travel time was tagged with one of the following sources: (1) baseline; (2) incident; (3) weather; (4) special event; (5) lane closure; or (6) high demand. A travel time reliability monitoring system that supports this use case would ideally integrate data on external sources of freeway congestion such as incidents, weather, lane closures, special events, and demand levels. The PeMS system operational in San Diego integrates statewide incident data from Caltrans' Traffic Accident and Surveillance Analysis System (TASAS) and statewide lane closure data from Caltrans' Lane Closure System. PeMS also reports peak-period vehicle-miles-travelled data for freeway routes. This PeMS data was used to evaluate the relationship between travel time variability and incidents, lane closures, and demand. Hourly weather data from the Automated Weather Observing System (AWOS) station at the San Diego International Airport was obtained from the NOAA National Data Center. Special event data was collated manually from various sport and event calendars for venues adjacent to the study routes.

- **Baseline**: A travel time was tagged with "baseline" if none of the factors was active during that five-minute time period.
- **Incident**: Incident data was obtained from the PeMS system operational in San Diego, which integrates statewide incident data from Caltrans' Traffic Accident and Surveillance System (TASAS). A travel time was tagged with "incident" if an incident was active anywhere on the route during that five-minute time period.

Incident start times and durations reported through PeMS were used to determine when incidents were active along the route. Incidents with durations shorter than 15 minutes were not considered.

- **Weather**: A travel time was tagged with "weather" if the weather station used for data collection reported precipitation during that hour.
- **Special Event**: A travel time was tagged with "special event" if a special event was active at a venue along the route during that time period. Special event time periods were determined from the start time of the event and the expected duration of that event type. For example, if a football game at Qualcomm Stadium had a start time of 6:00 PM and was scheduled to end around 9:00 PM, the event was considered active between 4:00 PM and 6:00 PM and between 8:30 PM and 10:00 PM, as this is when the majority of traffic would be accessing the freeways surrounding the venue.
- **Lane Closure**: A travel time was tagged with "lane closure" if a lane closure (scheduled or emergency) was active anywhere along the route during that time period.
- **High Demand**: Finally, a travel time was tagged with "high demand" if the vehicle-miles-travelled measured during that time period were more than 10% higher than the average vehicle-miles travelled for that time period. This approach was adapted from the SHRP2 L03 project, which considered high demand to be any time period where demand was 5% higher than the average for that time period. 10% was selected in this research effort because a 5% increase in demand had no measureable impact on travel times on either of the selected corridors.
- **Multiple Factors**: There were a few time periods within each data set where more than one factor was active during a single 5-minute period; in these cases, the travel time was tagged with the factor that was deemed to have the larger travel time impact (for example, when an incident coincided with light precipitation, the travel time was tagged with "incident").

Tagged travel times were then divided into different categories based on the time of the day, since the impacts of the congestion sources are time-dependent. For the westbound I-8 route, which was analyzed for weekdays, two different time periods were evaluated: (1) AM Peak, 7:00 AM-9:00 AM and (2) PM Peak, 4:00 PM-8:00 PM. For the northbound I-5 route, which was analyzed for weekends, two different time periods were evaluated: (1) Morning, 8:00 AM-12:00 PM and (2) Afternoon, 12:00 PM-9:00 PM.

Finally, within each time period, travel time probability density functions (PDFs) were assembled separately for all travel times and for those occurring during each source condition. The PDFs were plotted and summarized in various ways to give a thorough description of how the sources of congestion impact travel time variability and conditions on a route.

Route 1 (I-8) Results. For the westbound I-8 route, travel time variability and its contributing factors were investigated for weekdays during the four month period between November 2008 and February 2009. Data on incidents, weather, lane closures, special events, and demand fluctuations was collected from PeMS and external sources as described in the Methods section. Due to the preference of scheduling freeway lane closures during overnight, weekend hours, no lane closures were active on the route during the selected hours and date range. As a result, the contribution of lane closures to travel time variability on this route is zero. Analysis of vehicle-miles-travelled for the demand fluctuations component showed that demand is very steady and consistent on this corridor. Only three days were identified as having a

demand level not otherwise attributable to a special event that exceeded 10% of the average weekday demand level. All of these hours of high demand were during the PM period.

AM Peak. Exhibit C2-17 illustrates the distribution of 5-minute travel times in the AM period (7:00 AM-9:00 AM), divided by source condition. The AM period is the peak period for commute traffic on this route, since it begins in the eastern suburbs and terminates near downtown San Diego. As such, it is the time period with the most travel time variability. As evidenced by the plot, there is a wide distribution of travel times during the morning hours, ranging from approximately 8.5 minutes free-flow to 25 minutes at a maximum, a travel time measured when there was an incident. The only source conditions active during the weekday AM period over the four month study period were incidents and precipitation; no special events or hours of high demand were noted. The histogram shows that, almost 25% of the time, the travel time is a near-free-flow 9 minutes. The travel time only falls below 9 minutes when there is no external source of congestion active. The "tail end" of the travel time distribution, however, is dominated by weather and incident events. In particular, travel times ranging between 15 and 20 minutes (or double the free-flow travel time) only occur when either an incident or a weather event is active. Travel times greater than 20 minutes only occur when there is an incident on the route.

Interestingly, it is apparent from this graph that sometimes, even when an incident is active, the travel time falls below 10 minutes. This is likely due to the fact that this analysis does not account for the severity of incidents in the travel time tagging process. The incident travel times shown in this figure that are near the median are likely minor incidents that were promptly moved to the shoulder and then cleared.

Another way of viewing the travel time reliability impacts of different sources is to plot the travel time probability density functions (PDFs) under each source condition. Travel time PDFs for the baseline, incident, and weather conditions are each shown in Exhibit C2-18. The PDFs shown in this use case were assembled using non-parametric kernel density estimation. As the baseline PDF plot shows, the distribution of travel times is very small when there is no external congestion source active on the corridor; there is only a 2 minute difference between the median travel time and the 95th percentile travel time in this case. When an incident is active on the corridor, the distribution of travel times is much wider. An incident increases the median travel time on the facility by 2 minutes over the baseline condition and, with a 95th percentile travel time of 18.7 minutes, requires travelers to add a buffer time of 9.8 minutes, almost doubling their typical commute, to arrive on time during an incident. A weather event increases the median travel time even higher, to 15 minutes, resulting in a buffer time comparable to that caused by an incident.

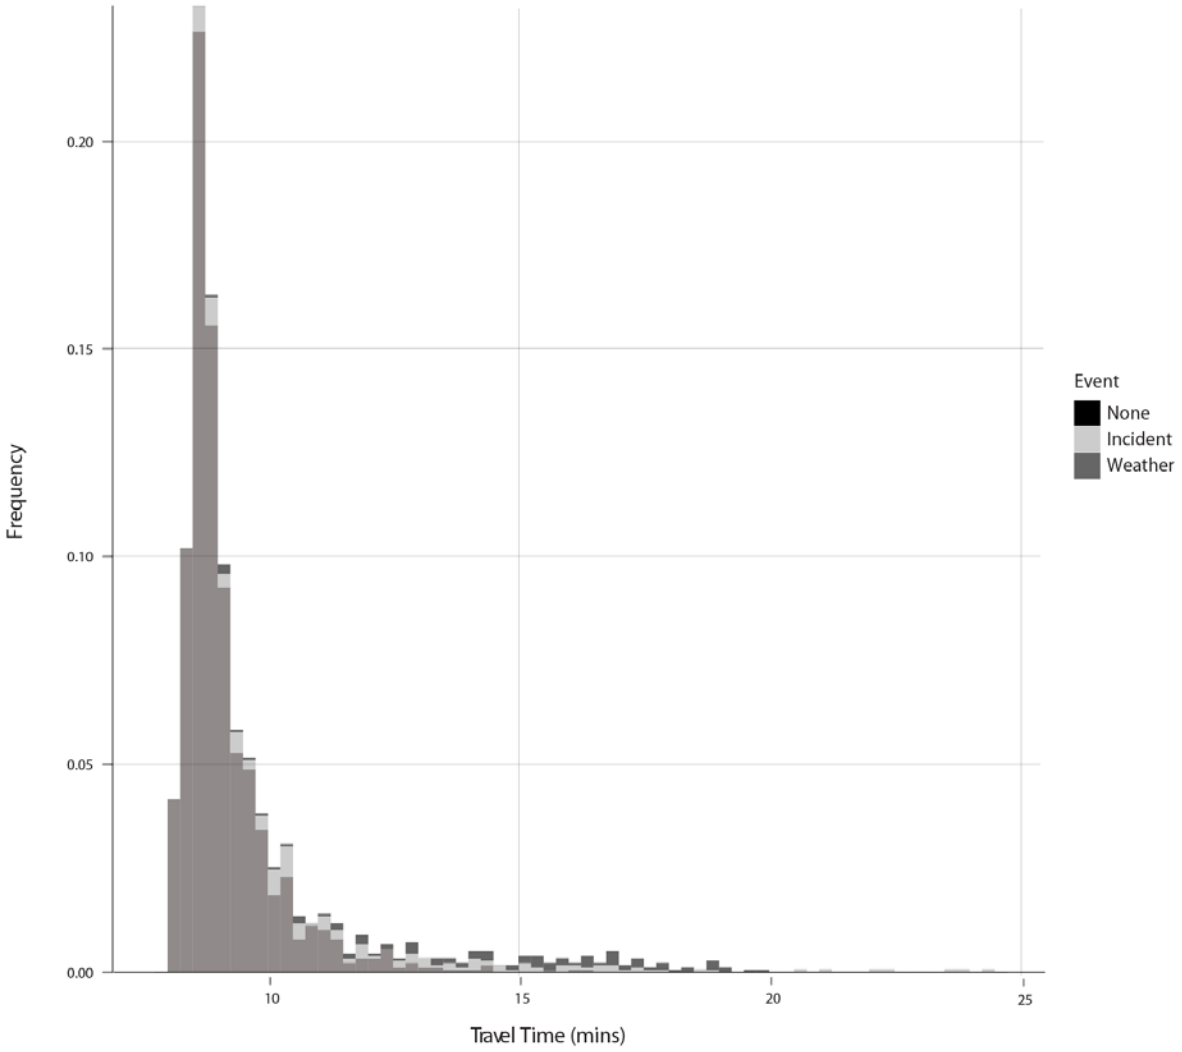

Exhibit C2-17: AM weekday distribution of travel times, WB I-8

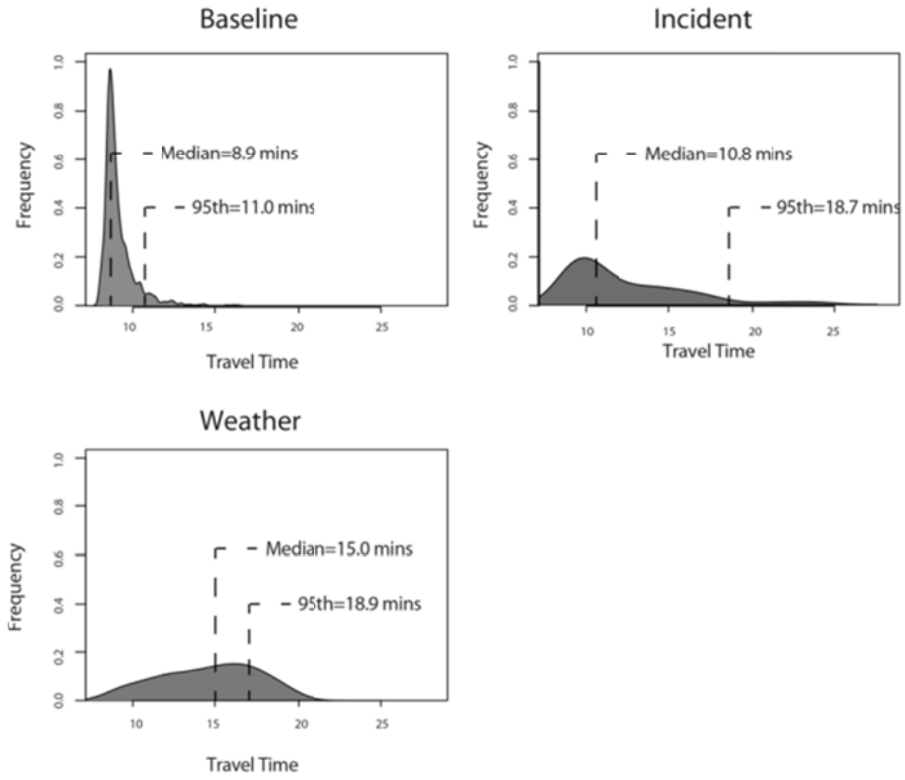

Exhibit C2-18: AM weekday travel time PDFs, WB I-8

A final way of summarizing this analysis is shown in Table C2-9, which lists the percentage of time that each source condition was active when travel times exceeded the 85th percentile travel time (10.6 minutes) and the 95th percentile travel time (15.0 minutes). As shown in the table, each of the three source conditions (none, incidents, and weather) occurred approximately 1/3 of the time that travel times exceeded the 85th percentile. For travel times that exceed the 95th percentile, weather is responsible for the largest share, followed closely by incidents. When the travel time exceeds the 95th percentile during the AM period on this facility, there is almost always some type of causal condition active on the roadway.

Table C2-9: AM weekday travel time variability causality, WB I-8

Source	Active when travel time exceeded 85th percentile	Active when travel time exceeded 95th percentile
Baseline	37.7%	3.3%
Incident	31.2%	41.1%
Weather	30.6%	55.6%

The conclusions that can be made from the AM time period analysis are that weather almost always slows down travel times significantly. Travelers need to plan to more than double their travel time over the typical condition when it is raining on this route. Incidents have a wider range of impacts on the corridor, depending on the severity. At their 95th percentile level, incidents increase travel times by almost 10 minutes over the median condition. Given no incidents or weather on this route, travelers can expect to see a travel time less than the 14.5

minute 95th percentile. Thus, when no non-recurrent sources of congestion are active, travelers need only add a buffer time of 5.5 minutes to arrive at their destination on-time.

PM Peak. The same analysis was also conducted for the PM peak period. The travel time variability source analysis for the PM period includes two factors that were not active during the morning: special events and high demand. There were three special events active on this corridor over the study period: one San Diego Chargers Monday Night Football game and two college football games. All three events took place at Qualcomm Stadium. Additionally, there were three time periods over the study date range that experienced greater than 1.1 times the normal demand level that were unrelated to special events. The breakdown of travel times by source is shown in Exhibit C2-19. Since the majority of traffic on this route commutes during the AM time period, the distribution of travel times during the PM period is small: there is a difference of only 0.7 minutes between the median travel time and the 95th percentile travel time. Travel times exceeding the 95th percentile have contributions from multiple factors. Travel times between 10 minutes and 12 minutes appear to be predominately caused by precipitation. Travel times exceeding 12 minutes appear to be caused by incidents or special events. The travel times measured during high demand time periods do not vary significantly from the median travel time.

Exhibit C2-20 shows the different PDFs for the five source conditions active during the PM period over the four months. At a glance, it is clear that the baseline and high demand event conditions have very tight, similarly shaped distributions, with less than a minute difference between the median and 95th percentile travel times. The lack of variability impacts of high demand is likely because the baseline volume is low enough during this time period that increasing it by 10% has minimal traffic impacts. While special events are rare on weekdays on this route, they can have a significant travel time impact when they do occur. The large difference between the median special event travel time and the 95th percentile special event travel time is likely due to the uncertainty of determining when the special event's travel time impacts would occur during the data tagging process. The 16.2 minute 95th percentile travel time likely represents the short time period when the majority of people are trying to access the special event venue, and the faster special event travel times are likely from the periods further before the event start when attendees are just beginning to trickle in. The impacts of incidents during the PM period are similar to those in the AM period, though the travel time variability impact of incidents is larger during the heavier morning commute. The PDF for the weather condition is of a different shape and a smaller distribution than it is in the other two time periods. This is possibly due to smaller amounts of precipitation in the PM period that were noted during

the data collection process.

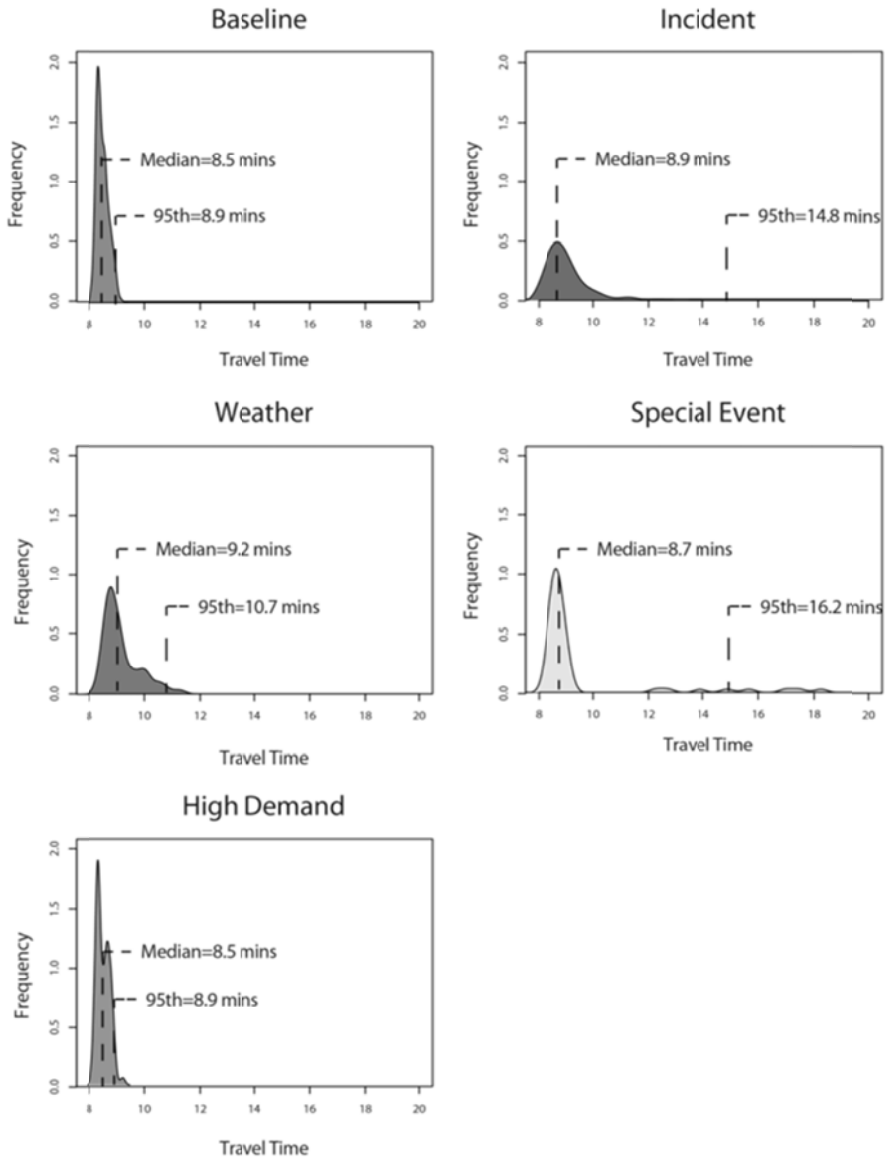

Exhibit C2-20: PM weekday travel time PDFs, WB I-8

Table C2-10 summarizes the contribution of each source condition to travel times exceeding the 85th percentile (8.9 minutes) and the 95th percentile (9.2 minutes). The 85th percentile travel time is very close to the median travel time, so there are many cases when the travel time exceeds the 85th percentile but no causal source is occurring. However, when travel times exceed the 95th percentile, there is a weather event 50% of the time and an incident 30% of the time. The contribution of the other factors to high travel times is low, due to the fact that they are infrequent.

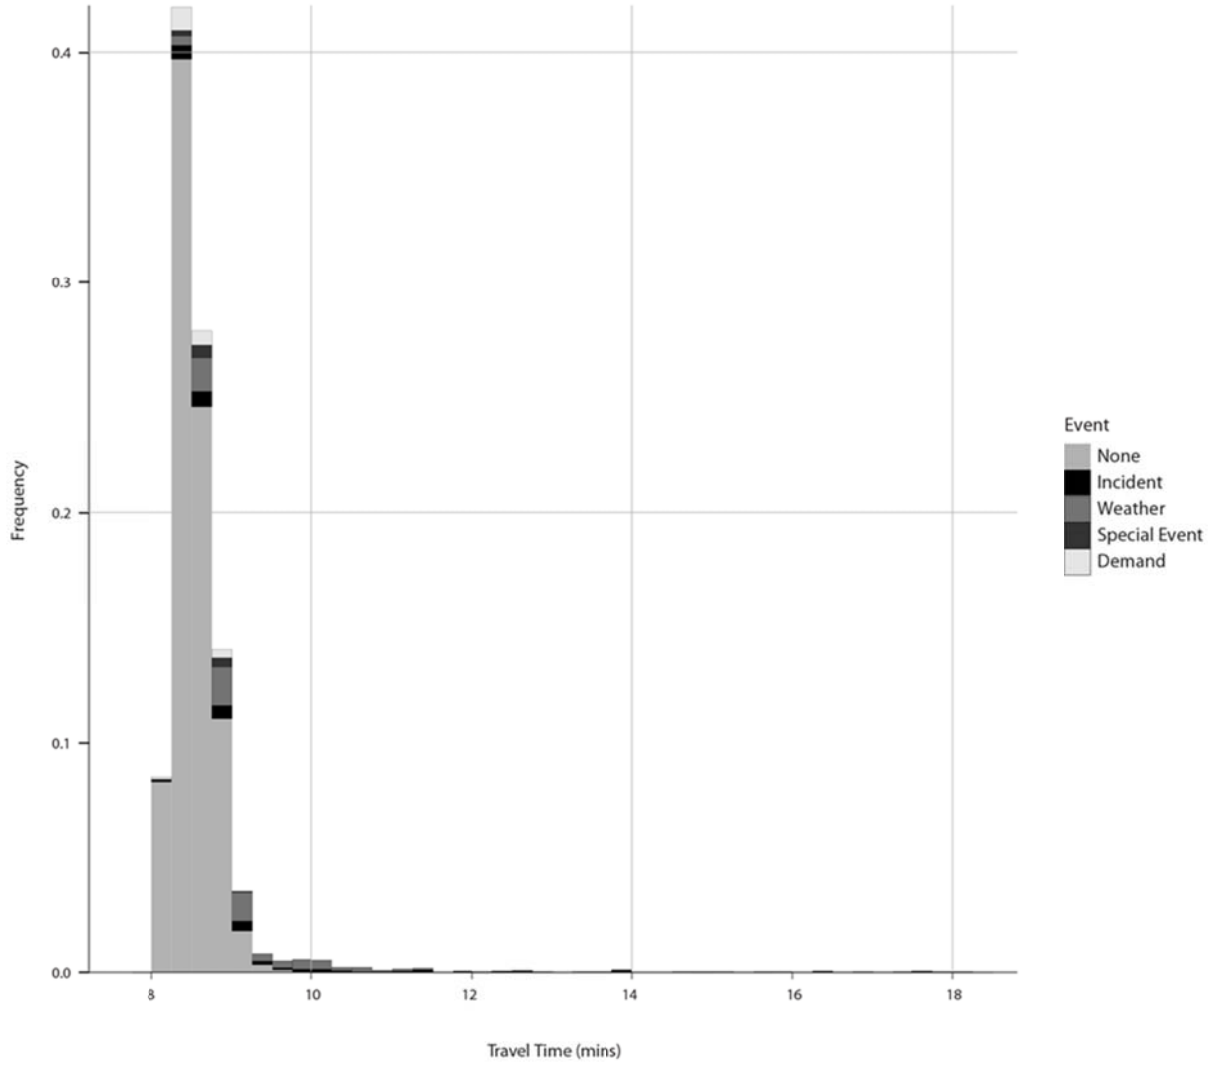

Exhibit C2-19: PM weekday distribution of travel times, WB I-8

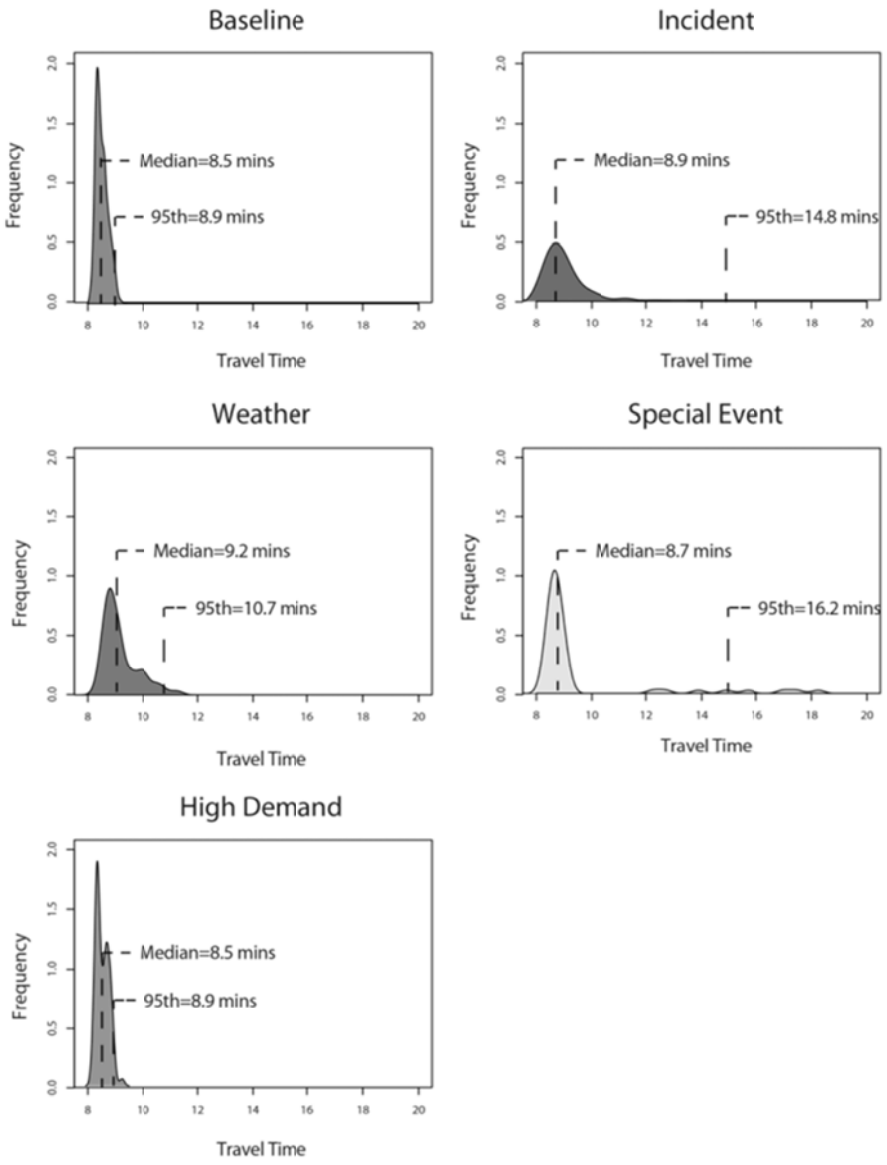

Exhibit C2-20: PM weekday travel time PDFs, WB I-8

Table C2-10: PM travel time variability causality, WB I-8

Source	Active when travel time exceeded 85th percentile	Active when travel time exceeded 95th percentile
Baseline	59.7%	15.2%
Incident	13.4%	29.8%
Weather	22.4%	50%
Special Event	3.6%	4.5%
High Demand	1.0%	0.6%

Synthesis. From a planning and operational standpoint, the only room for reliability improvement on this route exists during the AM period, as this is the only time period where

substantial travel time variability exists. While little can likely be done to reduce the variability caused by weather, focusing on better incident response or incident reduction methods could reduce the overall variability on the facility, which currently requires travelers to add a buffer time of 5.6 minutes (63%) to their AM commute to consistently arrive on time. In the other two time periods, travel time variability is minimal and the travel time impact of incidents is less severe than in the AM.

From a traveler perspective, this analysis provides insight into the range of conditions that can be expected given certain events. For instance, weather appears to slow down travel times across all time periods. It may prove useful to provide information to travelers on the travel times that they can expect to experience during rainy conditions, so that they can appropriately plan for an on-time arrival or defer a trip until conditions improve. Additionally, special events, when they occur, cause travel times to more than double on this route. In these instances, operators may want to consider providing information for alternate routes so that through-travelers can avoid the event-based congestion.

Route 2 Results. For the northbound I-5 route, travel time variability and its contributing factors were investigated for weekends during the entire year of 2009. Data on incidents, weather, lane closures, special events, and demand fluctuations were collected from PeMS and external sources as described in the Methods section. Due to the preference of scheduling freeway lane closures during overnight, weekend hours, no lane closures were active on the route during the selected hours and date range. As a result, the contribution of lane closures to travel time variability on this route is zero. The contributions of the factors to travel time variability were investigated for two different time periods, which corresponded to observed traffic patterns on the facility: (1) Morning, 8:00 AM-12:00 PM; and (2) Afternoon, 12:00 PM-9:00 PM.

Morning. Exhibit C2-21 shows the distribution of travel times during the weekend morning hours on northbound I-5. There is very little spread in the travel times measured on this corridor during the AM period; there is only a difference of one minute between the median and 95th percentile travel times. The travel times exceeding the 95th percentile predominantly occurred under incident and weather conditions. There were a number of high demand time periods on this corridor, when VMT exceeded 1.1 times the average VMT for weekend mornings, especially during the summer months due to increased beach traffic. However, travel times during high demand time periods never exceeded the 95th percentile, so the demand increases in the morning are typically not significant enough to cause severe congestion. There were no special events recorded during the morning hours of the study period.

Exhibit C2-22 illustrates the travel time PDFs that were assembled for each source condition. The baseline and weather PDFs have a very small distribution. The lack of travel time variability during weather conditions is likely related to the fact that there were only a few weekend days of precipitation over the study year, and the precipitation was relatively light during those days. The high demand PDF has a longer tail, showing that enough demand can cause slower travel times on this facility. Incidents appear to have the biggest impact on travel time variability during the AM hours, requiring motorists to add a buffer time of 8.5 minutes to the typical travel time

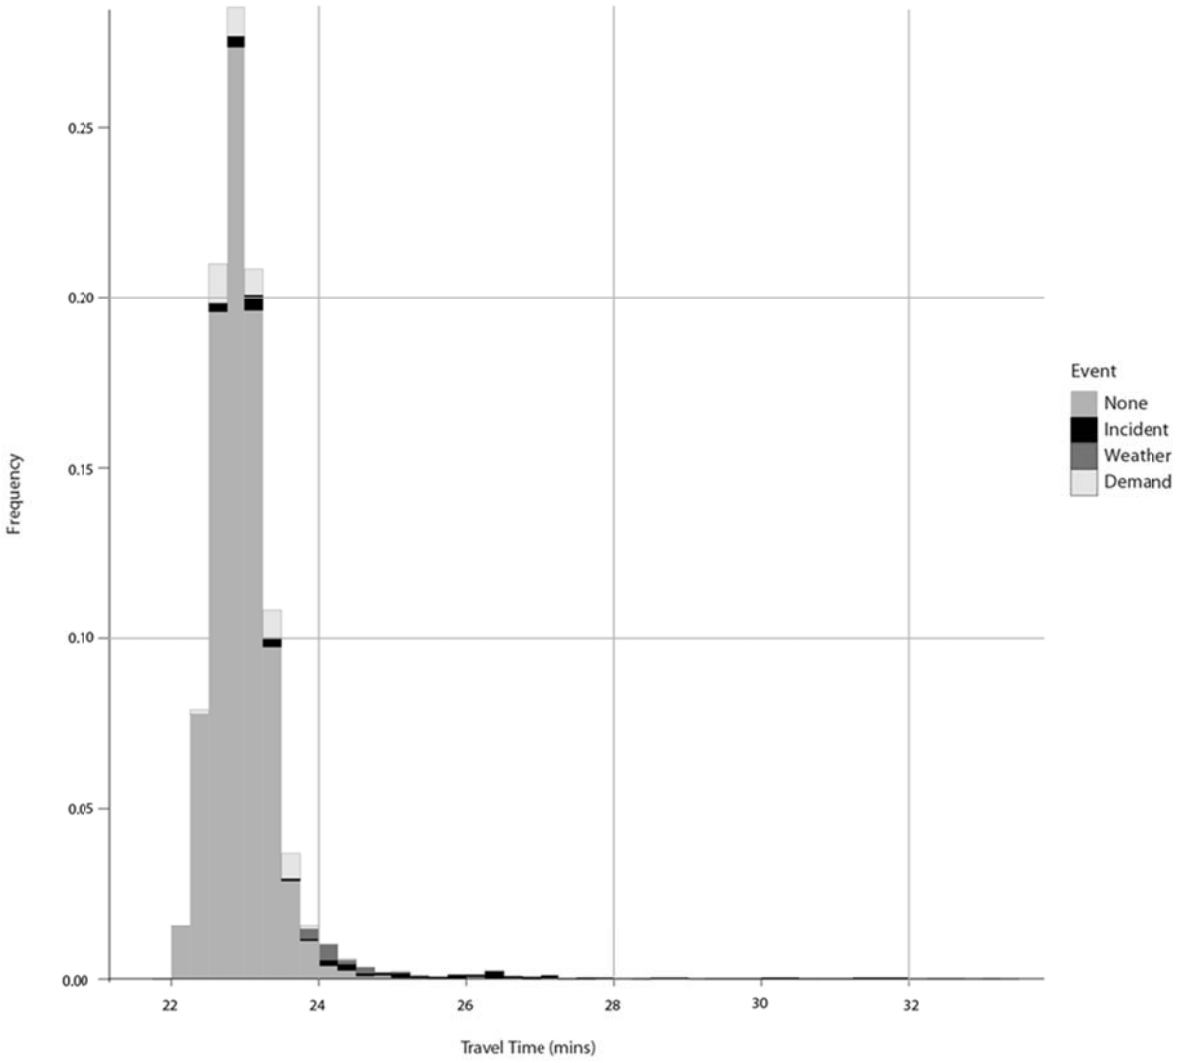

Exhibit C2-21: Weekend morning distribution of travel times, NB I-5

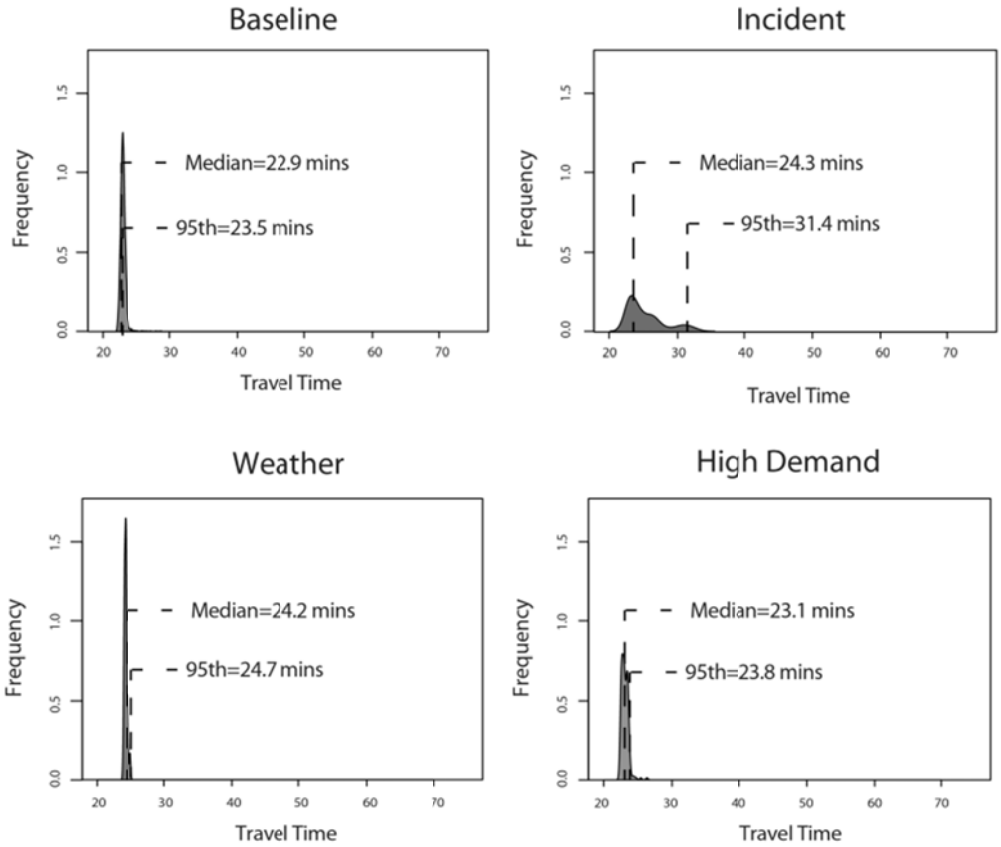

Exhibit C2-22: Weekend morning travel time PDFs, NB I-5

Finally, Table C2-11summarizes which source conditions were active when travel times exceeded the 85th and 95th percentile travel times on this route. While the high percentages for the baseline condition indicate that the sources of congestion cannot explain much of the variability, the variability on this route is very low. As such, it is conceivable that a number of travel times that would be considered typical for the corridor are falling outside of the 95th percentile threshold.

The results of the weekend morning analysis show that travel conditions remain relatively uniform throughout the year, though some variability is caused by incidents and rare levels of high demand.

Table C2-11: Weekend morning travel time variability causality, NB I-5

Source	Active when travel time exceeded 85th percentile	Active when travel time exceeded 95th percentile
Baseline	79.5%	64.2%
Incident	11.3%	20.9%
Weather	0.4%	0.1%
High Demand	6.5%	8.8%

Afternoon. Exhibit C2-23 shows the distribution of travel times by source condition during weekend afternoons and evenings on northbound I-5. As compared with the morning, the

PM travel time distribution has a significantly longer tail, with travel times ranging from 23.5 minutes free-flow to over 70 minutes, which occurred during a special event. Travel times exceed the 95th percentile travel time under various source conditions: in particular, incidents and special events. The special events considered in this analysis were the San Diego County Fair and the Del Mar horse races. Both events are active on multiple days during the summertime and are known to have major impacts on corridor traffic.

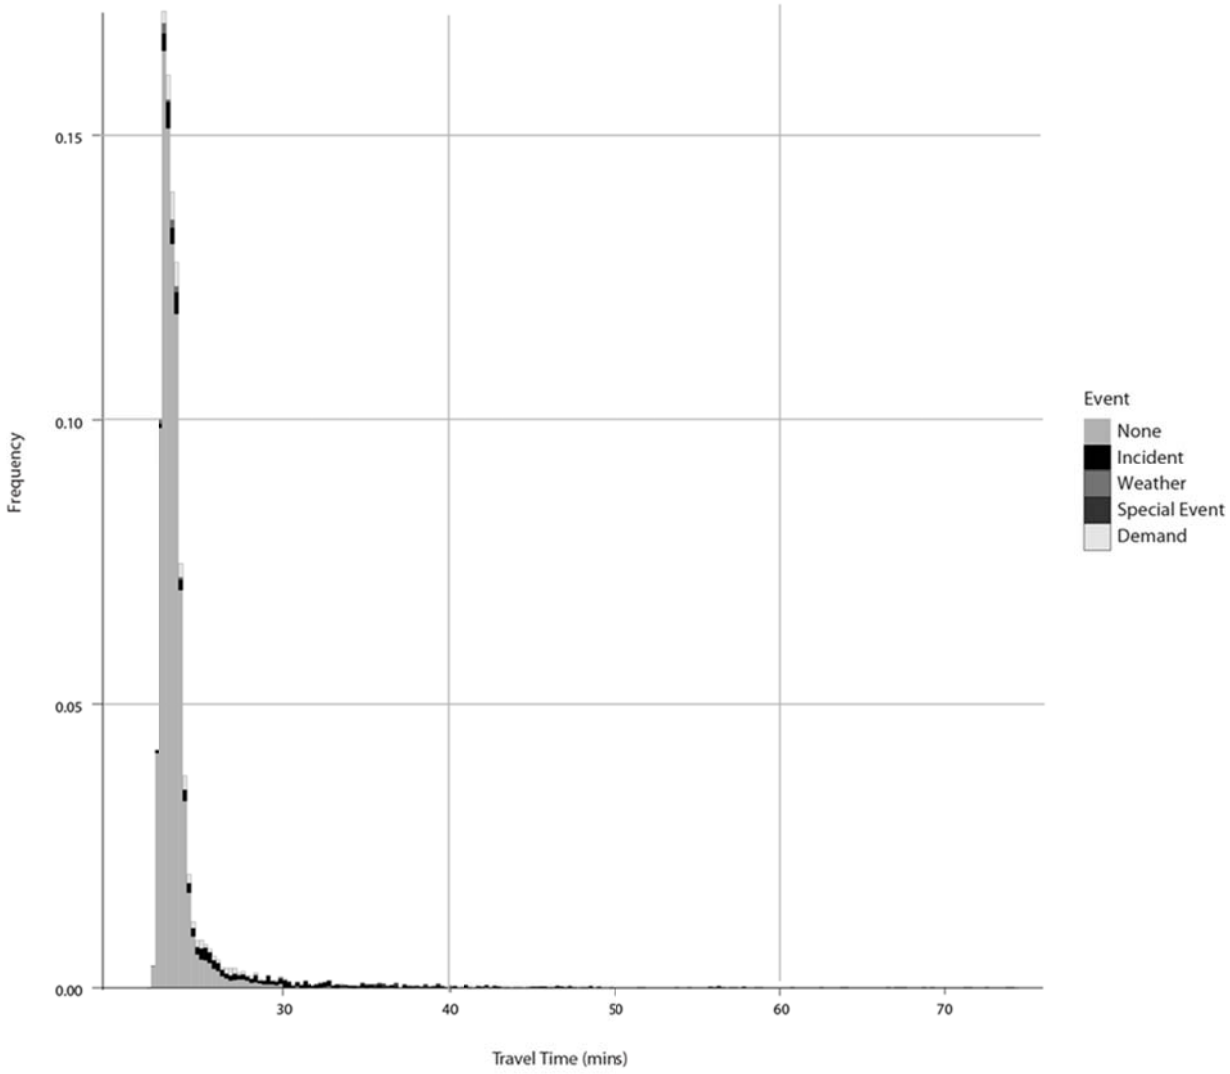

Exhibit C2-23: Weekend afternoon distribution of travel times, NB I-5

Exhibit C2-24 illustrates the different travel time PDFs assembled for the various source conditions that occurred on weekend afternoons on this study corridor. Similar to the morning time period, the PDFs for the baseline condition and the weather condition show little travel time variability. The weather events recorded over the study period were very minor, which might explain the difference in weather variability impacts between this corridor and the westbound I-8 corridor analyzed previously in this use case validation. High demand unrelated to any specific special event has the potential to increase travel times, but only in extreme circumstances; the typical demand fluctuations on the corridor incur only minor variability impacts. The sources

that cause the most travel time variability are incidents and special events. The median travel time during an incident is three minutes higher than the normal median travel time, and can be almost double the free-flow travel time at the 95th percentile level. On this corridor, special events are the source that has the potential to cause the highest travel time variability. Though they are relatively infrequent in that they are concentrated in the summer months, the median travel time during a special event requires an additional travel time of 15 minutes, a 64% increase over the ordinary median travel time. The 95th percentile travel time during a special event requires a buffer time of 45 minutes over the normal median travel time, requiring travelers to almost triple their typical travel time during this time period.

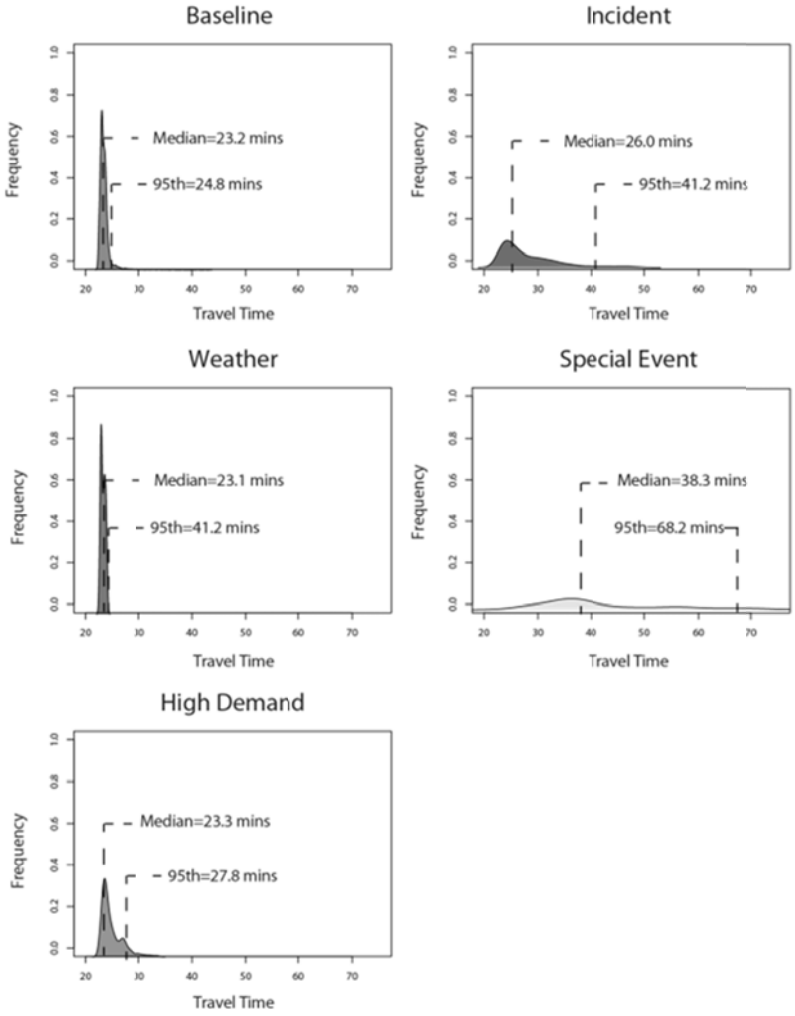

Exhibit C2-24: Weekend afternoon travel time PDFs, NB I-5

Finally, Table C2-12 summarizes which source conditions were active when travel times exceeded the 85th percentile and 95th percentile travel times on the route. Incidents and special events appear to be responsible for the majority of travel times that exceed the 95th percentile.

Table C2-12: Weekend afternoon travel time variability causality, NB I-5

Source	Active when travel time exceeded 85th percentile	Active when travel time exceeded 95th percentile
Baseline	51.4%	20.2%
Incident	29.1%	48.2%
Weather	0.0%	0.0%
Special Event	8.8%	25.3%
High Demand	10.8%	6.3%

Synthesis. The morning weekend travel time variability on the corridor is very minor, leaving little room for improvement from planning or operational interventions. The afternoon period, however, has significant travel time variability. This variability is predominantly caused by incidents throughout the year and by high demand and special events during the summer months. Because special events can cause such extraordinary travel time variability (causing travel times to double or triple the typical travel time on the route), traveler information during these event time periods is key. Diverting through traffic whose destination is not the event to alternate routes, or encouraging them to travel when the event is not active, could help mitigate the variability caused by these events.

Conclusion. This use case analysis illustrates one potential method for linking travel time variability with the sources of congestion. The methods used are relatively simple to perform with data that is generally available, either from the travel time reliability monitoring system or from external sources. The application of the methodology to the two study corridors in San Diego reveals key insights into how this type of analysis should be performed. Firstly, to ensure that sufficient travel time samples within each source category are being captured, this analysis should be performed on no less than three months' worth of data. It also should be performed separately for different days of the week, depending on the local traffic patterns. For example, the magnitude of the contributions of the sources to travel time variability on the northbound I-5 study corridor would likely be very different on weekday afternoons, when the corridor is serving commuters, than on weekend afternoons, when the corridor is serving recreational and event traffic. Additionally, it is important to consider the seasonal dependence of the congestion factors when selecting the time period for analysis, and when reviewing the analyses. For example, weather was shown to be a large contributing factor to travel time variability on the westbound I-8 corridor because the study period was November through February. If the analysis period was over the summer, the contribution of weather to travel time variability on this corridor would be nearly zero, as San Diego receives virtually no precipitation outside of winter. Finally, the contributions of the sources should be analyzed separately by time of the day, in a manner consistent with local traffic patterns. For example, while incidents had a major impact on the median travel time and the planning time during the AM commute period on the westbound I-8 study corridor, they had little variability impact during other parts of the day. Elucidating the time-dependence of the factors is critical to providing outputs that can be used by planners and engineers to improve the reliability of their facilities.

Use Case 2: Using planning-based reliability tools to determine departure time and travel time for a trip.

Summary. The purpose of this use case is to demonstrate how a reliability monitoring system can help travelers better plan for trips of varying levels of time-sensitivity. Currently, most traveler information systems that report travel times to end users focus solely on the average travel time, and give users little insight into the variability of their travel route. While this may be fine for trips with a flexible arrival time, it is less useful for trips for which the traveler must arrive at the destination at or before a specified time (such as a typical morning commute to work). This use case demonstrates how a reliability monitoring system can provide information both on the average expected travel time and the worst-case planning travel time so that the user can choose a departure time commensurate with their need for an on-time arrival. It also helps users choose between alternate routes; whereas one route may offer a faster average travel, it may have more travel time variability than a parallel route that is slower on average but has more consistent travel times.

Users. This use case is of most value to travelers, who are the end consumers of information that informs on the average and planning travel times for alternate routes between selected origins and destinations. The analysis behind this use case is also of value to operators, who can post estimated average and planning travel times throughout the day on variable message signs, to help travelers on the road choose between different routes based on their need for an on-time arrival.

Scope. The use case demonstrated in this section is broad and could provide a wide range of travel time reliability metrics to end users in a number of different formats. To narrow down the scope of this use case for validation purposes, this section will explore the specific use case defined below:

The user wants to view, for alternate routes, the latest departure times needed to arrive to a destination at 5:30 PM on a Friday: (1) on average and (2) to guarantee on-time arrival 95% of the time.

This definition means that the system needs to provide, for each alternate route, the median travel time and planning time for trips traveling between 5:00 PM and 5:30 PM on Fridays. It is envisioned that this use case involves the traveler utilizing the monitoring system for information in advance of a trip, likely from a computer, although other applications and dissemination methods are possible.

Site. Three alternate routes, which travel from just south of the I-5/I-805 diverge near La Jolla and Del Mar to the US Naval Base in National City, south of downtown San Diego, are studied in this use case. The three routes are shown in Exhibit C2-25. Route 1 is approximately 17 miles long and travels only along southbound I-5. Route 2 is approximately 16 miles long and travels along southbound I-805, southbound I-15, and southbound I-5. Route 3 is also 16 miles long and travels along southbound I-805, southbound SR-163, and southbound I-5.

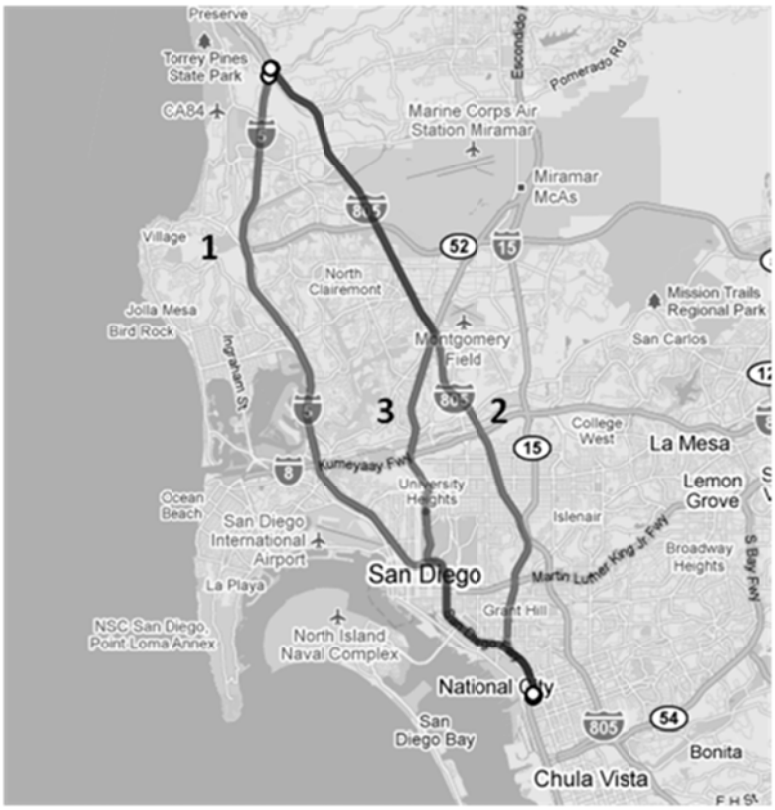

Exhibit C2-25: Freeway Use Case 2 Alternate Routes

Methods. The state of the practice for the few agencies who currently report travel time reliability metrics through their traveler information systems is to compute them from travel time probability density functions (PDFs) assembled based on the time of day and day of the week of the trip for which information is being requested. For example, to give a user the average and 95th percentile travel time for a Wednesday afternoon trip departing at 5:30 PM, the system might obtain all of the travel times for trips that departed between 5:15 PM and 5:45 PM for the past 10 weekdays.

The time-of-day and day-of-week approach to travel time reliability is valid and is used to demonstrate multiple use cases at the San Diego site. However, this use case evaluation incorporates the work that the research team has conducted into categorizing a route's historical and current performance into "regimes", and assembling travel time PDFs based on similar regime designations. Regimes are a way of categorizing travel times based on the prevailing operating condition when the travel time was measured. Regimes can be considered an extension of the time-of-day approach to reliability; on most corridors, regimes typically have a strong relationship with the time of day of travel. For example, a route that travels from a suburb to a downtown area may have four different operating regimes on weekdays: (1) a severely congested regime during the AM peak; (2) a mildly congested regime during the midday period; (3) a moderately congested regime during the PM peak; and (4) a free-flow regime that occurs during the middle of the night. There may also be "transitional" regimes that are observed when a route switches from congested to uncongested. Weekends may only have a free-flow regime and a slightly congested regime. An example regime assignment for a route that has five weekday regimes and two weekend regimes is shown in

Exhibit C2-26. As is evident from this figure, regimes are closely related to the time-of-day, but help capture the variability in operating conditions that occur across different days of the week, as well as to show the similarity in operating conditions across certain hours of the day.

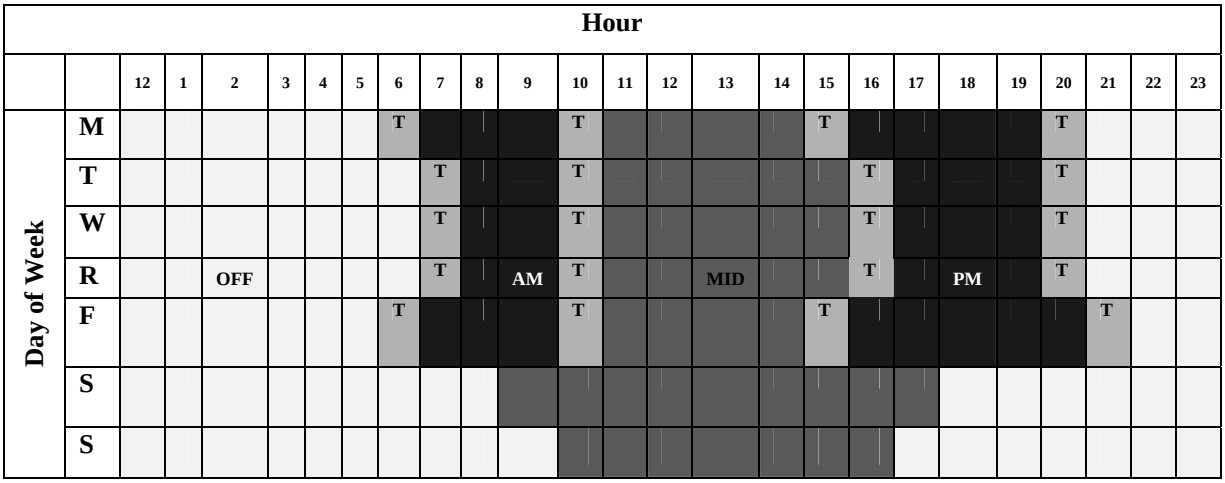

Exhibit C2-26: Example regime assignment for a route

Regime assignment is addressed in the Methodological Advances chapter of this case study, and the team is further refining its regime assignment methodologies. In this use case validation, each route is assigned a regime for each 5-minute time period of each day of the week. Routes are categorized into one of four regimes (free-flow, slightly congested, moderately congested, and severely congested) based on the ratio of the average travel time during the time period to the free-flow travel time, otherwise known as the travel time index (TTI). This metric was selected for regime identification because it is travel time-based and groups sets of travel times based on similar baseline operating conditions and levels of congestion, rather than a strictly time-of-day based categorization.

Following the regime identification process, travel times are assembled into regime-based PDFs based on the time of day and day of week of the traveler's request for trip information. From these PDFs, average travel times and planning times are computed and used to generate required departure times for each route based on the time-sensitivity of the trip.

Validation. The validation consists of three steps: (1) regime identification; (2) PDF generation; and (3) user output.

Regime Identification. In this use case, the travel time PDFs used to calculate reliability metrics for alternate routes are assembled based on regime conditions. In a travel time reliability monitoring system, this regime assignment step would be done prior to the user making the request for travel time information for alternate routes. For the three alternate routes, regime assignments were made for each day of the week type, based on local traffic patterns. The five day of week types selected for separate regime classifications were: (1) Monday; (2) Mid-week days (Tuesday, Wednesday, Thursday); (3) Friday; (4) Saturday; and (5) Sunday. Each 5-minute period of each day of the week type was assigned to a regime based on average TTI during that time-period. Average TTIs for each time period were calculated using 6 months of 5-minute travel time data (excluding holidays). The breakdown of regimes by TTI is shown in Table C2-13. These TTIs were selected by assuming a free-flow speed of 65 mph, then assuming that

average speeds less than 40 mph represent severely congested conditions, speeds between 40 and 50 mph represent moderately congested conditions, and speeds greater than 60 mph represent slightly congested conditions. Other routes in other regions may need different thresholds or numbers of regimes to accurately capture the varying levels of congestion along an individual corridor.

Table C2-13: Regimes by travel time index

Regime	TTI	Color	Route 1 Travel Time	Route 2 Travel Time	Route 3 Travel Time
Free-flow	<1.1		<15.6 mins	<13.4 mins	<15.4 mins
Slightly congested	1.1-1.3		15.6-18.5 mins	13.4-15.8 mins	15.4-18.2 mins
Moderately congested	1.3-1.6		18.5-22.7 mins	15.8-19.5 mins	18.2-22.4 mins
Severely congested	>1.6		>22.7 mins	>19.5 mins	>22.4 mins

The connection between regimes and travel times for each of the three study routes is shown in Table C2-13. The colors in the table correspond with the regime assignments by day of week type for each of the routes, shown in Exhibit C2-27, Exhibit C2-28, and Exhibit C2-29. While the regime assignments in these tables are shown for each 20-minute time period, regimes were actually assigned to each 5-minute time period.

The regime assignment allows for a comparison of the average performance by day of week and time of day on each of the three different routes. The free-flow travel times on each route are fairly comparable. Route 2 is the shortest route and has the fastest free-flow travel time (12.2 minutes). Route 1 and Route 3 are of comparable length; Route 1 has a slightly faster free-flow travel time (14 minutes) than Route 3 (14.2 minutes). Analysis of the regime tables shows that the duration of congestion on Route 1 is much narrower than it is on the other route, and there is only severe congestion right around the 5:00 pm hour during the midweek days. The duration of congestion on Route 2 is very wide; it lasts throughout the midday, is severe during the 5:00 PM hour on the midweek days, and is severe beginning at 4:00 PM on Friday. Route 3 is the only one of the routes to have AM congestion throughout the work week. It is also the only one of the routes to have weekend congestion, possibly because it traverses through San Diego's Balboa Park, a popular tourist destination. Congestion is severe on Route 3 Tuesday-Friday during the 5:00 PM hour.

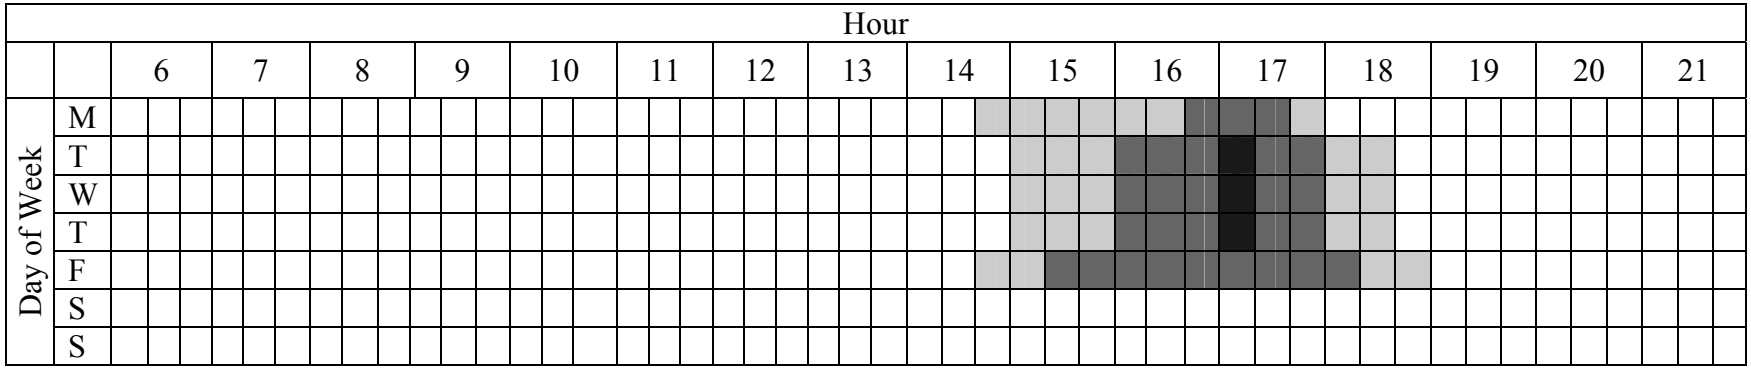

Exhibit C2-27: Route 1 (southbound I-5) regimes

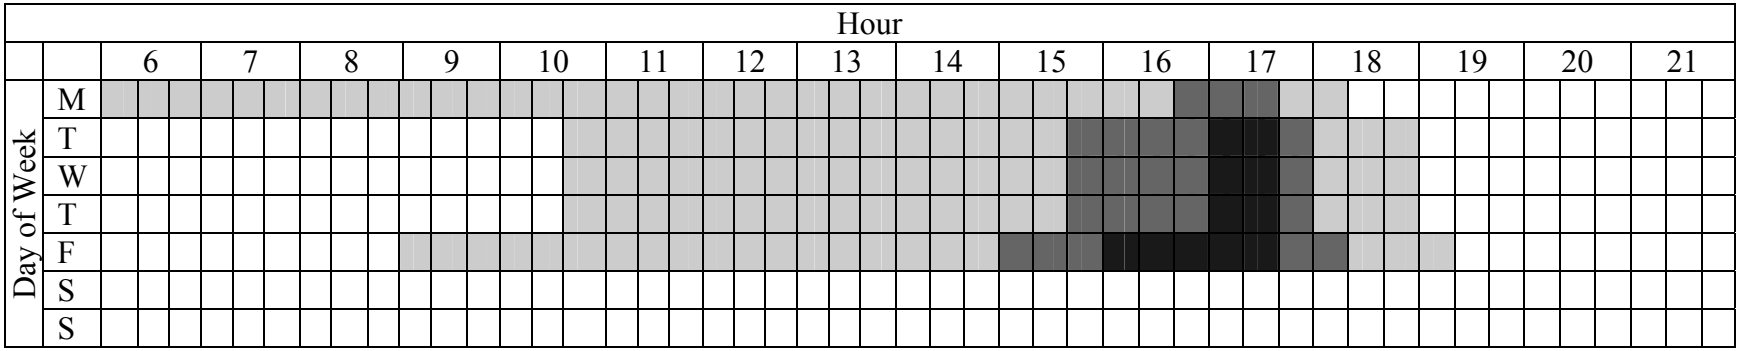

Exhibit C2-28: Route 2 (southbound I-15) regimes

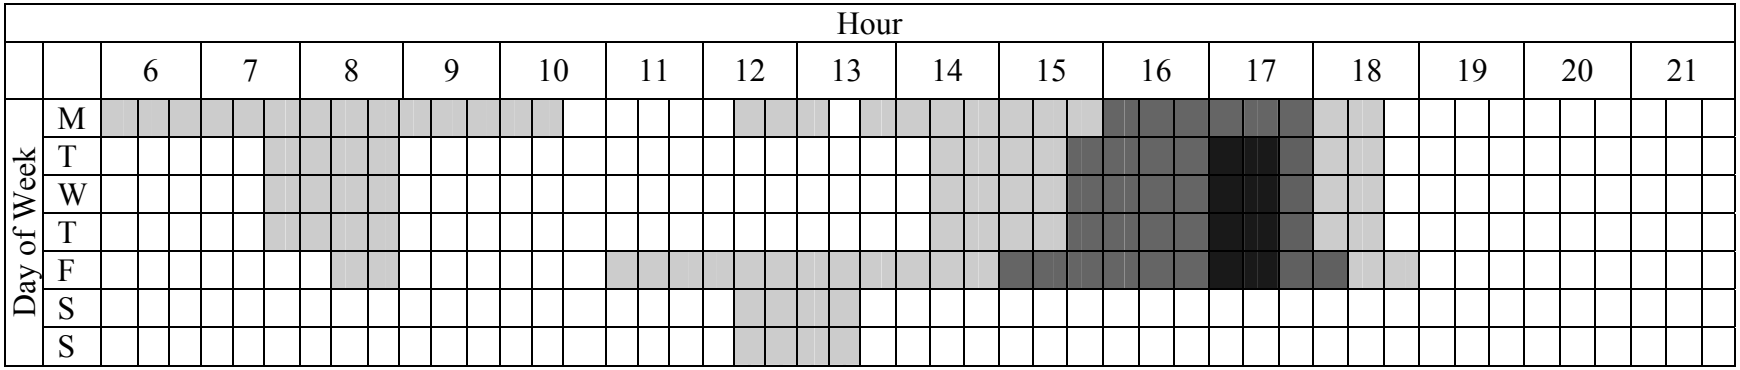

Exhibit C2-29: Route 3 (southbound SR-163) regimes

PDF Generation. While regime assignments are made off-line, this validation assumes that the regime-based PDFs are assembled in real-time, in response to a user's request for information. Future work by the research team will develop methods for creating PDFs off-line and storing them in advance of a user query, to reduce the need for real-time computation.

This validation assumes that the user wants to know the average and planning departure times for three different routes that allow for arrival at 5:30 PM on a Friday to the destination. As such, PDFs are generated for each of the three routes' operating regimes during the Friday 5:00 PM hour. The regime matrices show that Route 1 is in the moderately congested regime and Routes 2 and 3 are in the severely congested regime during this time period. As such, this validation effort generates travel time PDFs for each route using all of the travel times within the same regime category measured on Fridays over the past six months. An alternate method is to generate PDFs based on travel times within the same regime category measured on *any* day. In this case, since 6 months of data were used to form the PDFs, it was determined that Friday data alone would generate sufficient travel time data points to form an accurate PDF.

The plots of each PDF are shown in Exhibit C2-30. Route 1 appears to have the smallest distribution of travel times during this time period; the most frequently occurring travel time is around 20 minutes. Route 2 has significantly more travel time variability during this time period; while the most frequently occurring travel times on Friday during severe congestion are around 18 minutes, the travel time PDF has a long tail end, and travel times upward of 30 minutes can occur. The most frequently occurring travel time on Route 3 is approximately 24 minutes, and the route has significant travel time variability on Fridays.

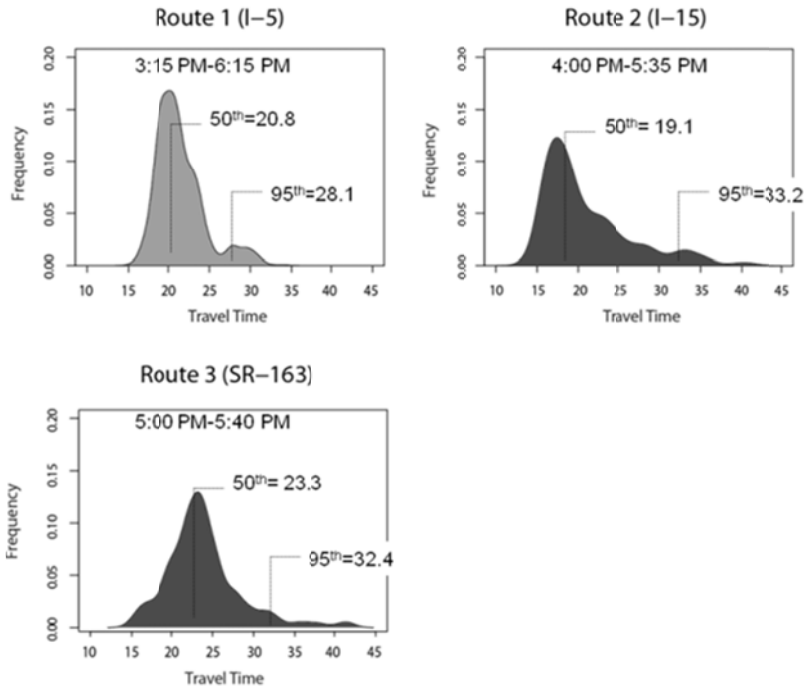

Exhibit C2-30: Alternate route travel time PDFs, 5:30 PM trip

User Outputs. In this step, the travel time PDFs are generated into useful summary metrics to assist the user in itinerary planning. In this case, the goal is to provide the user the departure times needed to arrive on-time on average and with a buffer time along each route. The

median and planning travel times on each route during the user's desired time of travel are summarized in Table C2-14. Route 2 has the fastest median travel time, but this route also has significant travel time variability, requiring a traveler with a non-flexible arrival time to add a buffer time of 14 minutes to the median travel time to ensure on-time arrival 95% of the time. Route 1 is almost 2 minutes slower than Route 2 on average, but offers significant (5 minutes) time savings when variability is included. Even Route 3, which has a much slower median travel time than the other two routes, has a faster planning time than Route 2.

Table C2-14: Median and planning travel times along alternate routes

Route	Median Travel Time (mins)	Planning Travel Time (mins)
Route 1 (I-5)	20.8	28.1
Route 2 (I-15)	19.1	33.2
Route 3 (SR-163)	23.3	32.4

Table C2-15 synthesizes these travel time estimates into information that is of most use to the end user- recommended departure times. These departure time estimates are termed "departure time for 50% on-time arrival" and "departure time for 95% on-time arrival" to help the user plan the trip with consideration of the need for an on-time arrival. Other applications of this use case could provide departure times calculated from other reliability metrics, such as the 85th percentile travel time rather than the 95th, or the 99th percentile travel time for trips where on-time arrival is imperative.

Table C2-15: Alternate route departure time estimates

Route	Departure time for 50% on-time arrival	Departure time for 95% on-time arrival
Route 1 (I-5)	5:09 PM	**5:01 PM**
Route 2 (I-15)	**5:10 PM**	4:56 PM
Route 2(SR-163)	5:06 PM	4:57 PM

Conclusion. This use case validation illustrates the value of incorporating reliability-based travel time estimates into traveler information systems for use in advance of trips, so that travelers can plan itineraries based on their need for on-time arrival. As proven by the San Diego validation, the route that is the fastest on average is not always the route that consistently gets travelers to their destination on-time. Providing buffer time measures for alternate routes conveys this message to the end user, ultimately giving them more confidence in the ability of the transportation system to get them to their destination on-time.

Use Case 3: Combining real-time and historical data to predict travel times in real-time

Summary. The purpose of this use case is to extend the system capabilities described in the freeway planning time use case in order to support the prediction of travel times along a route in real-time, using both historical and real-time data. While a number of methods for performing this data fusion to predict travel times have been implemented in practice, most only generate a single expected travel time estimate. This use case validation extends the methodology to generate, in addition to a single expected travel time, a range of predictive travel times that incorporate the measured historical variability along a route.

Users. This use case if of the most value to travelers, who currently lack quality real-time information on expected travel times while en route to a destination. The analysis behind this use case is also of value to operators, who can use these methodologies to provide better predictive travel times to post on variable message signs or via other dissemination technologies.

Scope. This use case validation describes methodologies for predicting near-term travel time ranges along a route. Specifically, it predicts travel time ranges for a 5:35 PM Thursday trip for two alternate routes.

Site. Two of the same alternate routes used to demonstrate freeway use case 2 (alternate route planning times) were used to demonstrate this predictive travel time use case. Both routes begin just south of the I-5/I-805 diverge and end near the United States Naval Base in National City. The first route, called the I-15 route, travels along southbound I-805, southbound I-15, and southbound I-5. The second route, called the I-5 route, travels solely along I-5. Maps of these two routes are shown in Exhibit C2-31.

Southbound I-15 Route Southbound I-5 Route

Exhibit C2-31: Freeway Use Case 3 alternate routes

Methods. Per the use case requirements, the validation needs to use both data from the historical archive as well as real-time data to generate travel time predictions for trips that are already occurring or are to begin immediately. To meet these requirements, a "nearest neighbors" approach was adopted, which uses the measured real-time conditions along a route to search for similar conditions in the past, then predicts a travel time based on historical travel times measured under similar conditions. Similar approaches have been well-documented in literature, and a nearest neighbors approach is currently used in PeMS to predict travel times along a route for the rest of the day *(1, 2)*. The method used in this validation extends upon traditional techniques to incorporate reliability information; instead of providing one predictive travel time, this use case validation outputs a range of predictive travel times that incorporate the potential variability in travel times that may occur, as gathered from similar historical conditions. The employed methodology is only valid for near-term travel time prediction. As such, this use case assumes that predictions are only made for the next three upcoming five-minute time periods.

To estimate a real-time predictive travel time range for a route, the methodology compares travel time data collected over the past six five-minute time periods with travel time data collected over the same six time periods over the most recent 15 days of the same day of the week. In this use case, which aims to predict travel times for a 5:35 PM Thursday trip, this means that travel times measured between 5:00 PM and 5:30 PM on the current day are compared with travel times measured between 5:00 PM and 5:30 PM over the 15 most recent Thursdays.

The "nearest neighbors" to the current day are selected by comparing the "distance" between the measured five-minute travel time on the historical day with the measured travel time for the same five-minute period on the current day. The distances between travel times for different five-minute periods are weighted differently, such that similarity for the five-minute trip that immediately precedes a trip is weighted more than similarity for the five-minute trip that occurred 30-minutes before the current trip. The weighting factors used for each 5-minute period are shown in Table C2-16.

The following variables are used to explain the methodology:

- T_C = current day travel time
- T_h = historical day travel time
- d_h=distance between current day five-minute travel time and historical day five-minute travel time
- D_h=total distance between current day travel time and historical day travel times for all five-minute periods prior to a trip
- x= time period prior to trip start (ranges from 1 for 5-minutes prior to 6 for 30-minutes prior)
- w= weight factor

The distance d_h between the current day travel time and the historical day travel time is calculated using the following equation:

$$d_h(x) = w(x) * (T_h(x) - T_c(x))^2$$

The total distance D_h between a current day and a historical day is calculated by summing up all the distances d_h using the following equation:

$$D_h = \sqrt{\sum_{x=1}^{6} d_h(x)}$$

Table C2-16: Weighting factors (w) for minutes prior to trip

Minutes Prior to Trip	Weight factor
5 (x=1)	1
10 (x=2)	½
15 (x=3)	¼
20 (x=4)	1/8
25 (x=5)	1/16
30 (x=6)	1/32

The result of the distance calculation step is a measure of travel time closeness between each historical day and the current day. From here, the method of k-nearest neighbors is

followed; rather than selecting the travel time profile of the nearest day as the predicted travel time, the method considers the travel time profiles from the three nearest days in order to make a prediction. The goal of this use case is to predict a travel time range for the next three five-minute time periods. In this validation, the expected travel time for the next three time periods is computed as the median of the travel times from the three nearest neighbor days. The lower bound of the predictive range is computed as the expected travel time minus the variance of the three neighbor travel times. The upper bound of the predictive range is computed as the expected travel time plus the variance of the three neighbor travel times.

Results. The travel time prediction methodology was used to compute predictive travel time ranges for the two example alternate routes between 5:35 PM and 5:45 PM on Thursday, August 12, 2010. Because there is data on what the travel times actually were on this day, this validation has a "ground-truth" data source with which to compare the estimates generated by the selected methodology.

I-15 Route. To predict 5:35 PM to 5:45 PM travel times on Thursday, August 12, 2010, five-minute travel times between 5:05 PM and 5:45 PM were obtained for 15 Thursdays, between April 29, 2010 and August 12, 2010.

The distance calculation method was used to determine the nearest neighbors. Table C2-17 shows the travel times measured for each five-minute time period over the 15 selected days. The first row shows the travel times measured on the "current" day of August 12, and all other rows show the travel times measured on each previous Thursday. The last column shows the total distance measured between the travel times on each day and the travel times on the current day. The three shaded rows indicate the days on which the distance was lowest, which were concluded to be the most similar to the current day.

Exhibit C2-32 compares the travel times measured on the predicted day with those measured on the closest three Thursdays, and extends the x-axis to show the travel times on these three days for the periods of 5:35 PM, 5:40 PM, and 5:45 PM. These are the travel times from which the predictive range for the current day is to be calculated. The thick black line indicates the travel times for the current day up until 5:30 PM.

Exhibit C2-33 shows the results of using the median of the nearest neighbor travel times approach to make a prediction of the expected travel times for the upcoming 15 minutes, and compares these predictions to the travel times that were actually measured on this day. Table C2-18 shows this information in tabular form, and also gives the predictive travel time ranges, which account for travel time variability in the evolving traffic conditions. As shown in the table, each actual measured travel time fell within the predictive range. The expected travel times only varied from the measured travel times by 5%.

Table C2-17: Neighboring Thursday travel times on I-15

Date	5:05 PM	5:10 PM	5:15 PM	5:20 PM	5:25 PM	5:30 PM	Distance
8/12/10	28.3	28.1	27.9	26.1	25.9	24.6	--
5/06/10	17.1	18.1	18.7	18.8	18.5	17.7	10.4
5/13/10	18.6	18.0	18.4	18.6	18.1	18.3	10.2
5/20/10	18.8	19.7	19.8	19.7	18.8	18.4	9.4
5/27/10	15.6	16.5	16.9	17.0	16.9	16.4	12.5
6/03/10	28.0	27.9	28.1	27.9	27.6	26.8	2.7
6/10/10	17.5	19.1	20.1	21.0	21.0	21.0	6.9
6/17/10	18.2	19.2	19.0	19.2	18.5	17.1	10.6
6/24/10	34.8	35.0	36.0	37.4	37.0	37.4	16.5
7/01/10	24.6	25.2	25.9	24.9	24.8	24.3	1.6
7/08/10	17.5	18.2	18.4	17.1	16.1	15.7	13.0
7/15/10	15.8	16.1	16.4	16.5	16.9	16.5	12.6
7/22/10	20.7	22.0	22.4	22.5	22.6	22.9	4.4
7/29/10	20.8	20.4	20.3	20.0	20.2	19.5	8.0
8/05/10	22.9	24.5	26.2	26.4	25.7	25.1	1.6

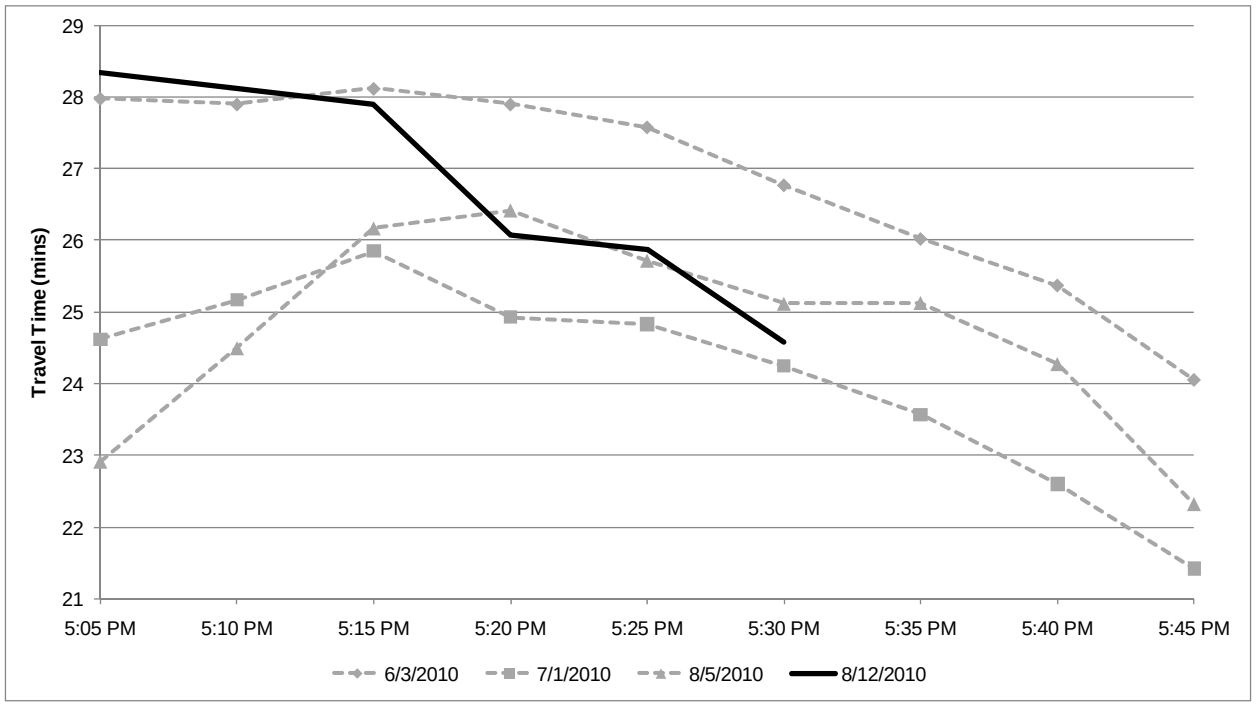

Exhibit C2-32: Travel time profiles of three closest Thursdays, I-15

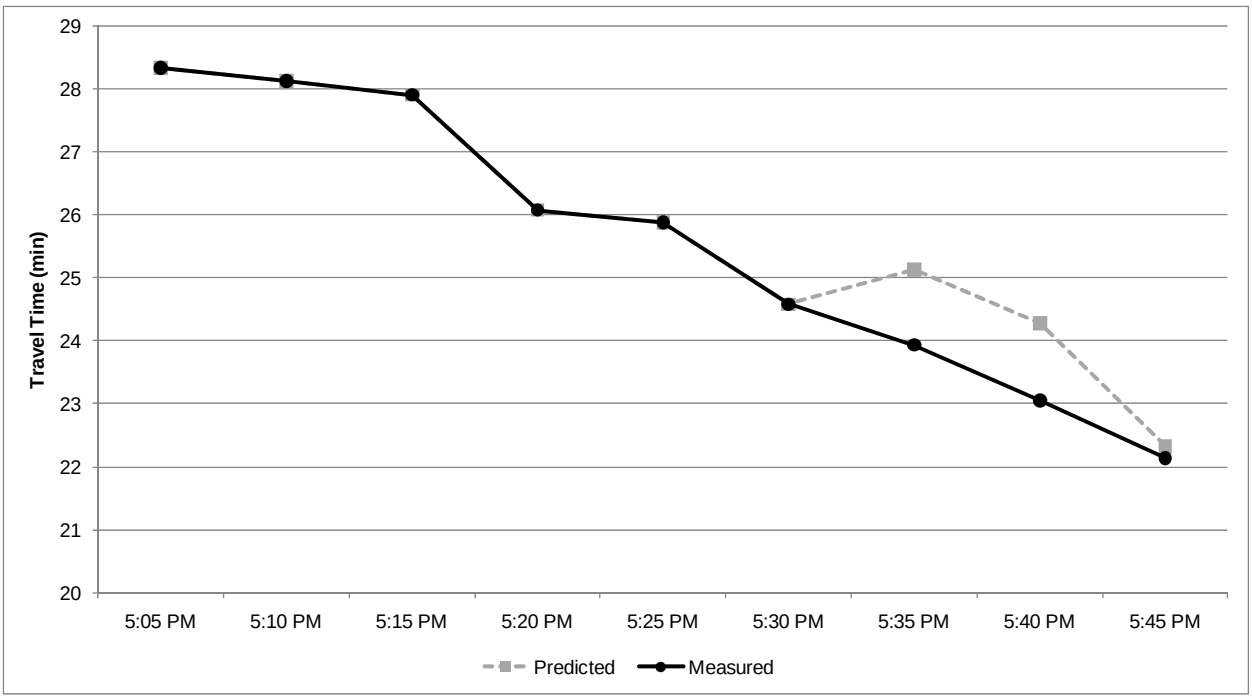

Exhibit C2-33: Measured and Predicted Travel Times, 8/12/2010, I-15

Table C2-18: Predicted vs. actual travel times, 8/12/2010, I-15

	5:35 PM	**5:40 PM**	**5:45 PM**
Predicted Lower Range	23.6 mins	22.3 mins	20.6 mins
Predicted Upper Range	26.7 mins	26.2 mins	24.1 mins
Predicted	25.1 mins	24.3 mins	22.3 mins
Measured	23.9 mins	23.1 mins	22.1 mins
Measured in range of predicted?	Yes	Yes	Yes
% difference between predicted and measured	5.0%	-5.2%	-1.0%

I-5 Route. The same approach was taken to estimate a predictive travel time range for the alternate southbound I-5 route for the same 15 minute time period. Exhibit C2-34 plots the travel times for the three closest Thursdays identified by the distance calculation method. The heavy black line indicates the travel times for the current day up until 5:30 PM.

Exhibit C2-35 compares the median travel time prediction for the upcoming 15 minute period with the actual travel times that were measured on this route and day. Table C2-19 expands this information to show the lower and upper bounds of the predicted travel time ranges, and compares the estimates with the travel times actually measured on this day. Each measured travel time fell within the predictive range, and expected travel times varied from the measured travel times by less than 5%.

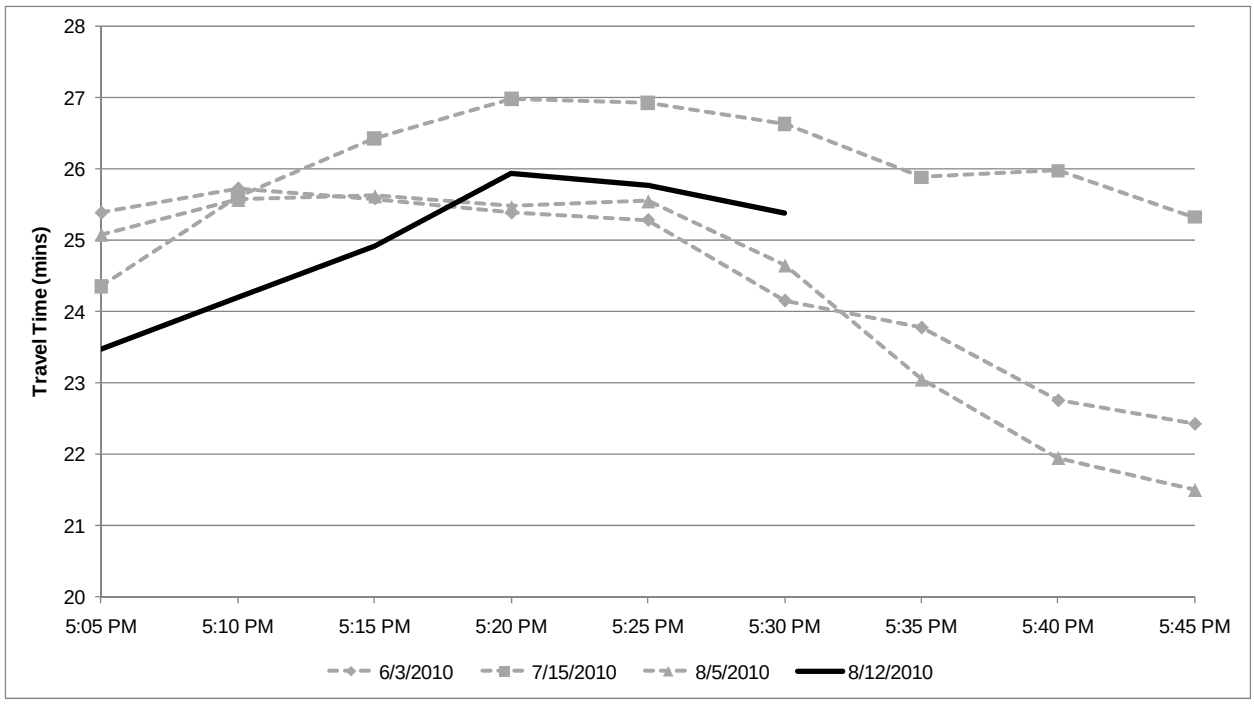

Exhibit C2-34: Travel time profiles from three closest Thursdays, I-5

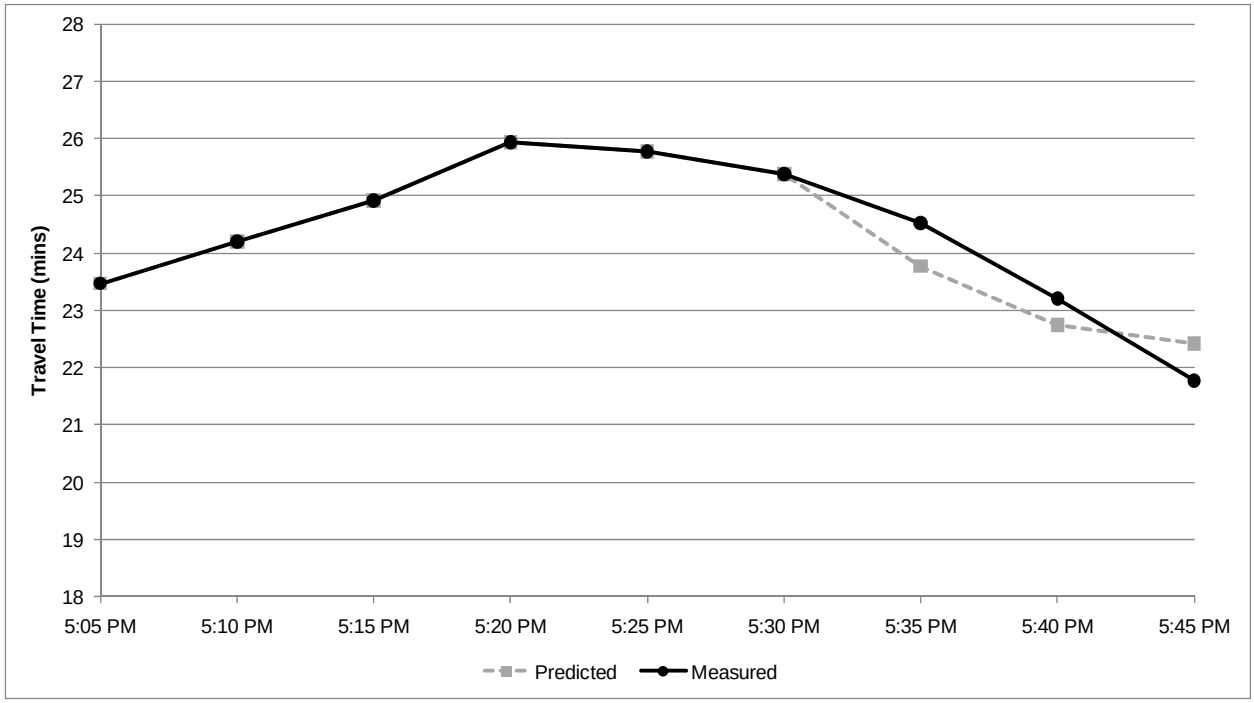

Exhibit C2-35: Measured and predicted travel times, 8/12/2010, I-5

Table C2-19: Predicted vs. actual travel times, 8/12/2010, I-5

	5:35 PM	**5:40 PM**	**5:45 PM**
Predicted Lower Range	21.6 mins	18.2 mins	18.4 mins
Predicted Upper Range	25.9 mins	27.3 mins	26.4 mins
Predicted	23.8 mins	22.8 mins	22.4 mins
Measured	24.5 mins	23.2 mins	21.8 mins
Measured in range of predicted?	Yes	Yes	Yes
% difference between predicted and measured	-2.9%	-1.7%	3.8%

Synthesis. It is envisioned that the results of the travel time prediction methodologies can be used to provide updated travel time information in real-time, to help users select alternate routes based on current traffic conditions as well as historical travel time patterns and reliability. For the example case, the following information could be posted on a variable message sign to provide travelers with current information.

TRAVEL TIMES TO NATIONAL CITY
I-5: 21-26 MIN
I-805/I-15: 23-27 MINS

Conclusion. This use case validation shows that it is possible to provide predictive travel time ranges and expected near-term travel times by combining real-time and archived travel time data. The validation uses a k-nearest neighbors approach to compare recent travel times from the current day with travel times measured on previous days. It then approximates near-term travel times based on the measurements from the most similar days. The travel time predictions for both study routes proved very similar to the actual travel times measured on the sample day. The travel time ranges output by the prediction method provide a way to report travel time reliability information in real-time to give travelers a more realistic idea of the range of conditions they can expect to see during a trip.

Transit

Use Case 1: Conducting offline analysis on the relationship between travel time variability and the seven sources of congestion

Summary. This use case aims to quantify the impacts of the seven sources of congestion: (1) incidents; (2) weather; (3) lane closures; (4) special events; (5) traffic control; (6) fluctuations in demand; and (7) inadequate base capacity, on travel time variability for transit trips. To perform this analysis, methods were developed to extract travel times from Automated Passenger Count (APC) bus data. These travel times were then flagged with the type of event they occurred under (if any) and aggregated into travel time probability density functions (PDFs). From these PDFs, summary metrics such as the median travel time and planning travel time were computed to show the extent of the variability impacts of each event condition.

Users. This use case has broad applications to a number of different user groups. For transit planners, knowing the relative contributions of the different sources of congestion toward travel time reliability would help them to better prioritize travel time variability mitigation

measures on a route-specific basis. The outputs of this use case would also be of value to operators, providing them with information that informs on the range of operating conditions that can be expected on a route given certain event conditions. Finally, the outputs of this use case would have value to travelers, by providing better predictive travel times under certain event conditions that could be posted in real-time on variable message signs at stops or on vehicles, or on traveler information websites. This information would help users better know what to expect during their trip, both during normal operating conditions and when a congestion-inducing event is occurring.

Site. Three routes were selected for the evaluation of this use case, in order to highlight the varying contributions of congestion factors to travel time reliability across different routes, service patterns, and times of day. The first route analyzed is Route #20, Southbound, which travels from the Kearny Mesa area down SR-163 into downtown San Diego. For this analysis, we select a subset of the route spanning 16.4 miles. This study section of Route #20 begins near the intersection of Miramar Road and Kearny Villa Road on the northern edge of the Marine Corps Air Station Miramar and continues South along SR-163 to downtown San Diego. At Balboa Ave. and SR-163, after traveling along SR-163 for 6.6 miles, Route #20 takes a detour to Fashion Valley Transit Center at Friar's Road and SR-163 before reentering SR-163 at I-8. Finally, the route terminates in downtown San Diego at 10th Avenue and Broadway.

The second route analyzed here is Route #20 X, which is identical to Route #20 except that it does not stop at the Fashion Valley Transit Center. Here, we study a 14.7-mile long stretch of Route #20 X beginning near the intersection of Miramar Road and Kearny Villa Road on the northern edge of the Marine Corps Air Station Miramar and continuing South along SR-163 for 12.6 miles to downtown San Diego, terminating at 10th Avenue and Broadway.

The third route analyzed is Route #50, Southbound, which travels along I-5 into downtown San Diego. This route begins near the Clairemont Drive on-ramp to I-5, continues south along I-5 for 6.4 miles, and ends 0.8 miles later at 10th Avenue and Broadway. The route is 7.2 miles long.

Both Routes #20 and #50 were chosen because they travel for significant distances along freeways, meaning that roadway incident data can be obtained for them through PeMS. Secondly, these routes were chosen because they travel towards downtown, which hosts several special events during the period of study, so their travel times can be analyzed for the effect of special events. Finally, these are routes for which a comparatively large amount of APC data is readily available.

A map of all routes is shown in Exhibit C2-36.

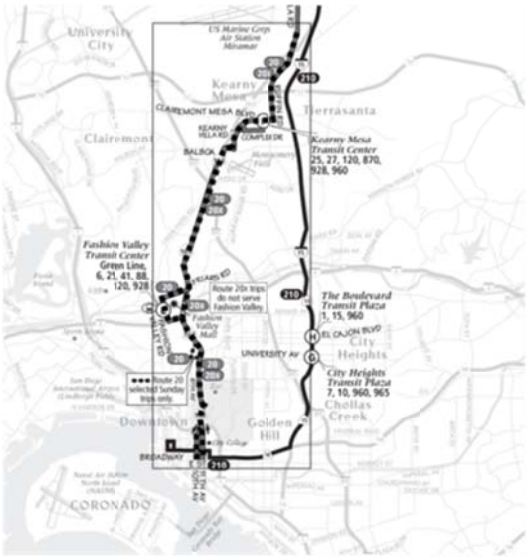

Route #20 and 20x (dashed)

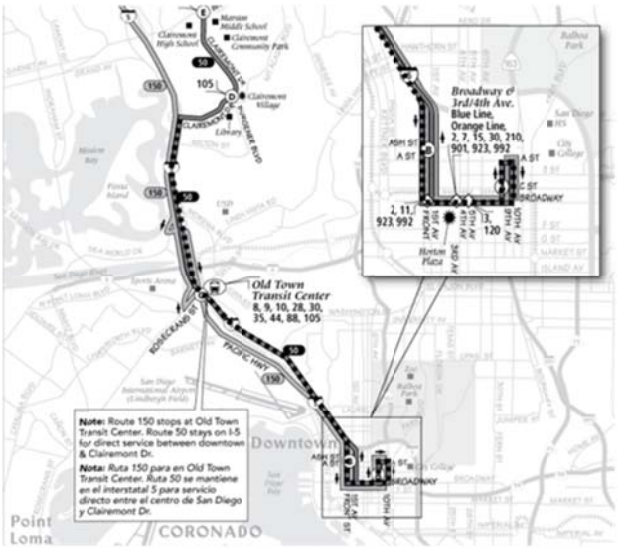

Route 50 (dashed)

Exhibit C2-36: Transit Use Case 2 routes

Methods. These routes were analyzed to determine the travel time variability impacts caused by three sources of congestion: (1) incidents; (2) special events; and (3) fluctuations in transit demand. Traffic control contributions were not investigated as ramp metering location and timing data could not be obtained. Weather contributions were not considered due to the lack of inclement weather in San Diego for the August 2010 study period (the only month for which the APC data could be obtained). Lane closures were also not considered as they are expected to have little impact on transit service, even when the transit route runs along a freeway. The impacts of inadequate base capacity were not considered for the same reason.

For every weekday run for which data was available on each of the three routes described above, APC data was analyzed to determine the in-vehicle travel time from delivered service records. Passenger loadings were also extracted from the APC data.

To link travel times with the event condition that was active during their measurement, each transit run for which a travel time was obtained was tagged with one of the following events: (1) baseline (none); (2) special event; (3) incident; or (4) high demand. A travel time was tagged with "baseline" if none of the factors were active during that run. A travel time was tagged with "incident" if an incident was active anywhere on the route during that run. Incident start times and durations reported through PeMS were used to determine when incidents were active along the route. Incidents with durations shorter than 15 minutes were not considered. A travel time was tagged with "special event" if a special event was active at a venue along the route during that time period. Special event time periods were determined from the start time of the event and the expected duration of that event type. For example, if a football game at Qualcomm Stadium had a start time of 6:00 PM and was scheduled to end around 9:00 PM, the event was considered active between 4:00 PM and 6:00 PM and between 8:30 PM and 10:00 PM, as this is when the majority of traffic would be accessing the venue. Finally, a travel time was tagged with "high demand" if the number of passengers on board the transit vehicle reached or exceeded 50 at any point during the run. For cases in which more than one factor was active,

the travel time was tagged with the factor that was deemed to have the larger travel time impact (for example, when a long-lasting incident coincided with a trip that also ran during the edge of a low-attendance special event window, the travel time was tagged with "incident").

Tagged travel times were then divided into different categories based on the time of the day, since the impacts of the congestion sources are time-dependent. For all three transit routes, three different time periods were evaluated: (1) AM Peak, 7:00 AM-9:00 AM; (2) Midday, 9:00 AM-4:00 PM; and (3) PM Peak, 4:00 PM-8:00 PM.

Finally, within each time period, travel time probability density functions (PDFs) were assembled for all measured travel times.

Results. This section describes the results for the different routes.

Route #20, Southbound. For Route #20, Southbound, travel time variability and its contributing factors were investigated for the 22 weekdays in August 2010. The period of study was limited to a single month due to a shortage of data on other months. Data on incidents, special events, demand fluctuations and travel times were collected from PeMS, external sources, and the in-vehicle APC sensors, as described in the Site Description chapter.

Scheduled travel times for the subset of Route #20 considered here over the period of study range from 39 to 50 minutes, averaging 51.7 minutes. In August 2010, vehicles averaged 8.1 minutes longer to complete this portion of the route than the scheduled time.

The travel time distribution of trips on this route appears to be roughly unimodal with a high standard deviation, greater frequency on the smaller side of the mode, and several outlying trips with long travel times. The mode occurs at 54.2 minutes.

Over the period of study, this route saw 129 transit trips made. Among these 129 total trips, 7 special event, 2 incident, and 16 high demand trips were recorded. Exhibit C2-37 shows the travel time distribution for all trips over the study period, according to the event present (if any) during that trip.

Exhibit C2-38 shows the distribution of travel times during the weekday AM peak of August 2010. Relatively few trips occurred during the AM Peak on Route #20. Those that did appear to be clustered together around 44.2 minutes. This could be due to fluctuations in the transit schedule throughout the day, with trips occurring early in the morning scheduled with shorter travel times than trips occurring later in the day. No events were flagged for trips occurring in this time period.

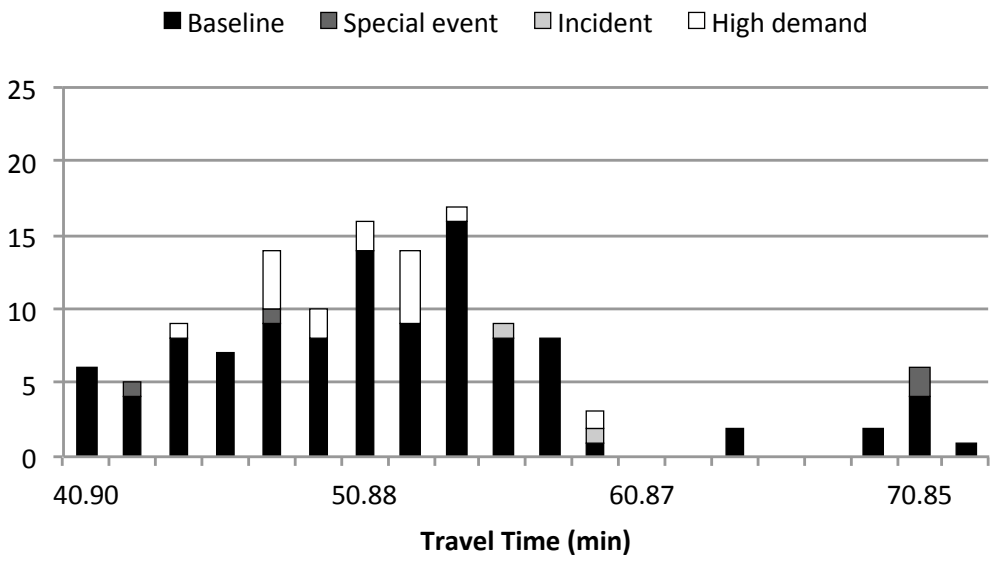

Exhibit C2-37: Total travel time distribution for Route #20, August 2010

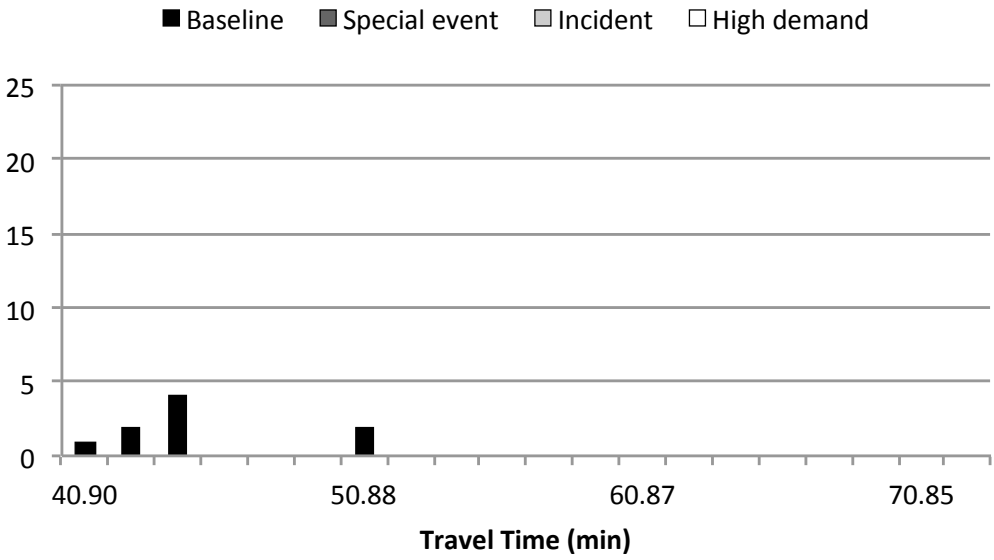

Exhibit C2-38: AM peak travel time distribution for Route #20, August 2010

Exhibit C2-39 shows the travel time distribution over the month for midday trips. Travel times for the midday period, in contrast to those seen in the AM peak, appear clustered around the primary mode seen in Exhibit C2-37 of 54.2 minutes. The distribution of variability-causing events is interesting, with the two recorded incident trips associated with longer-than-average trips, and two of the six longest trips associated with special events. However, all 11 high demand trips had shorter than average travel times, indicating that large passenger loadings do not have much effect on travel times along this route during the middle of the day. This is good as 11 of the 14 high demand events on this route occurred during the middle of the day.

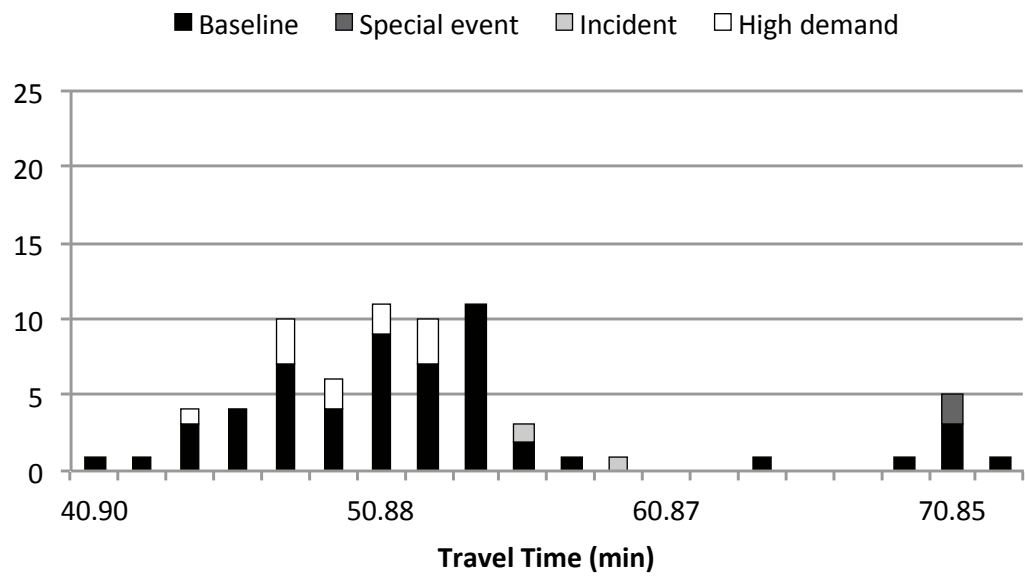

Exhibit C2-39: Midday travel time distribution for Route #20, August 2010

Exhibit C2-40 shows the travel time distribution of trips taken during the PM peak period. There is one special event associated with a relatively low travel time of 42.6 minutes. This event was a San Diego Padres baseball game and occurred late in the evening. High-demand events are also visible throughout the distribution, although they do not appear to be correlated with longer travel times. This is the most highly variable time period for which this route was analyzed.

Table C2-20 summarizes the contribution of each event condition to all travel times, to those exceeding the 85th percentile (57.2 minutes), and to those exceeding the 95th percentile (70.6 minutes). It can be seen that, although just 3.82% of all trips were associated with a special event, 10% of trips where travel times exceeded the 85th percentile were associated with a special event. When limiting the pool to trips that exceeded the 95th percentile travel time, a full 14.29% of that total can be associated with special events. From a planning and operational standpoint, this indicates that special events are associated with long travel times on this route. Thus, there could be some room for reliability improvements by improving signage, adding capacity, or advertising alternative routes during special events.

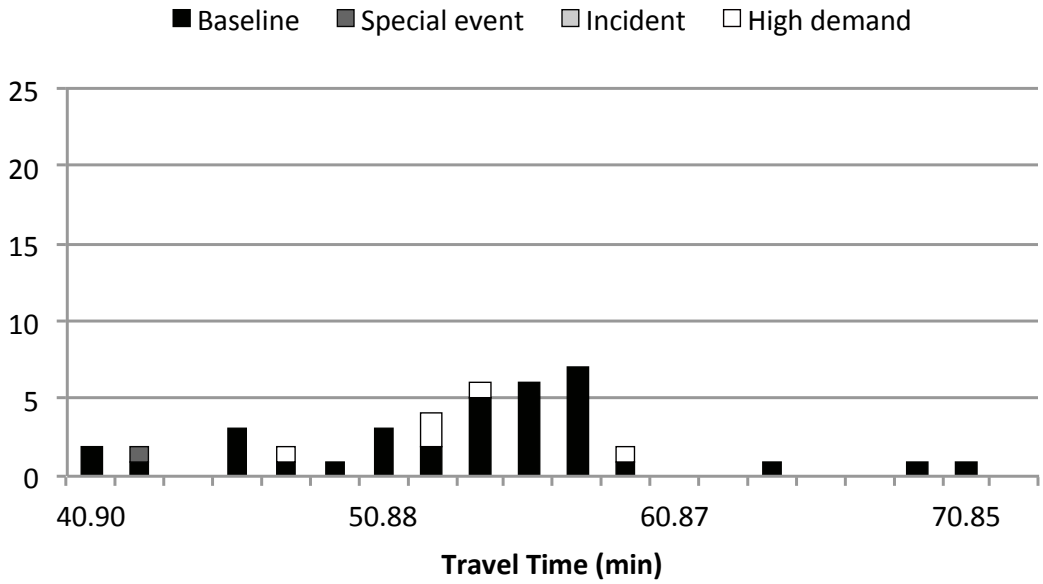

Exhibit C2-40: PM peak travel time distribution for Route #20, August 2010

Table C2-20: Travel time variability causality for Route #20

	Active	**Active when travel time exceeded 85th percentile**	**Active when travel time exceeded 95th percentile**
Baseline	82.4%	80.0%	85.7%
Special Event	3.8%	10.0%	14.3%
Incident	1.5%	5.0%	0.0%
Demand	12.2%	5.0%	0.0%

Route #20X, Southbound. Similarly to Route #20, for Route #20 X, Southbound, travel time variability and its contributing factors were investigated for the 22 weekdays in August 2010. The period of study was limited to a single month due to a shortage of data on other months. Data on incidents, special events, demand fluctuations and travel times were collected from PeMS, external sources, and the in-vehicle APC sensors, as described in the Methods section.

Scheduled travel times for the subset of Route #20 X considered here over the period of study range from 29 to 35 minutes, averaging 32.5 minutes, nearly a full 10 minutes less than Route #20. In August 2010, on average, buses took 10.1 more minutes than they were scheduled to complete this portion of the route.

A bimodal distribution can immediately be seen in Exhibit C2-41, which plots all trip travel times over the month, with most travel times clustered around the higher mode, 42 minutes, and a smaller grouping around 32 minutes. The source of the bimodal distribution is not immediately clear. There is virtually no correlation between the scheduled travel time and actual travel time ($R^2 = 0.043$) on this route; trips belonging to the lower mode do not necessarily have shorter scheduled travel times. However, of the 11 trips with travel times less than 36 minutes, 10 correspond to the 7:13 AM run, and nine were made by the same driver. Of the 11 days when this smaller travel time was not seen, 10 had no 7:13 AM run scheduled. Thus, there seems to be

an unknown factor associated with this particular run and driver which leads to a smaller travel time on this portion of the route.

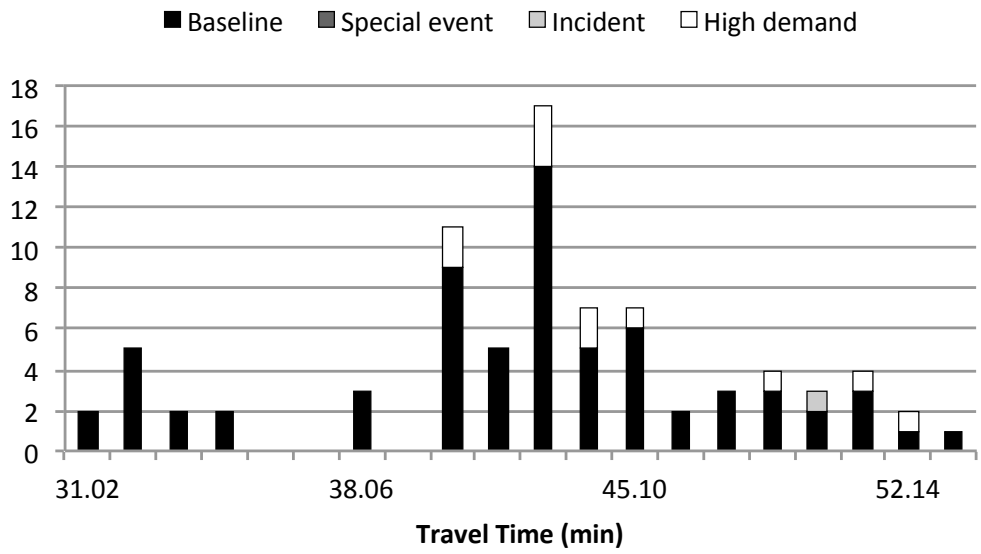

Exhibit C2-41: Complete travel time distribution for Route #20X, August 2010

Exhibit C2-42 depicts the distribution of travel times along Route #20 X during the AM Peak period (7:00 AM to 9:00 AM), labeled by event condition. The bimodal distribution described above can be seen most clearly here as all of the low-travel-time trips occur during the AM Peak, as discussed earlier. This is in stark contrast to the distribution of travel times for the AM Peak period on Route #20. Both modes appear to be tightly bunched. This bimodal distribution makes the AM Peak the period with the largest travel time variability for this route. There was only a single event condition measured during the AM Peak on this route: a high demand event which was associated with a travel time of 48.6 minutes.

Exhibit C2-43 depicts the distribution of travel times along Route #20 X during the midday period (9:00 AM to 4:00 PM), labeled by event condition. Here, a single mode is seen around 42.6 minutes. As with Route #20, the midday period saw the largest number of high passenger loadings on this route, with 8 high demand events. However, also similar to Route #20, these high loadings do not appear to be associated with longer travel times. There was a single incident event which was associated with the highest midday travel time seen on this route, 49.8 minutes.

Exhibit C2-44 depicts the distribution of travel times along Route #20 X during the PM Peak period (4:00 PM to 8:00 PM), labeled by event condition. There were few trips taken on this route during this time span, so no overwhelming travel time trend can be identified other than the high variability of travel times. The largest travel times seen on this route occurred during the PM period. Of the five largest travel times, two were associated with high demand events.

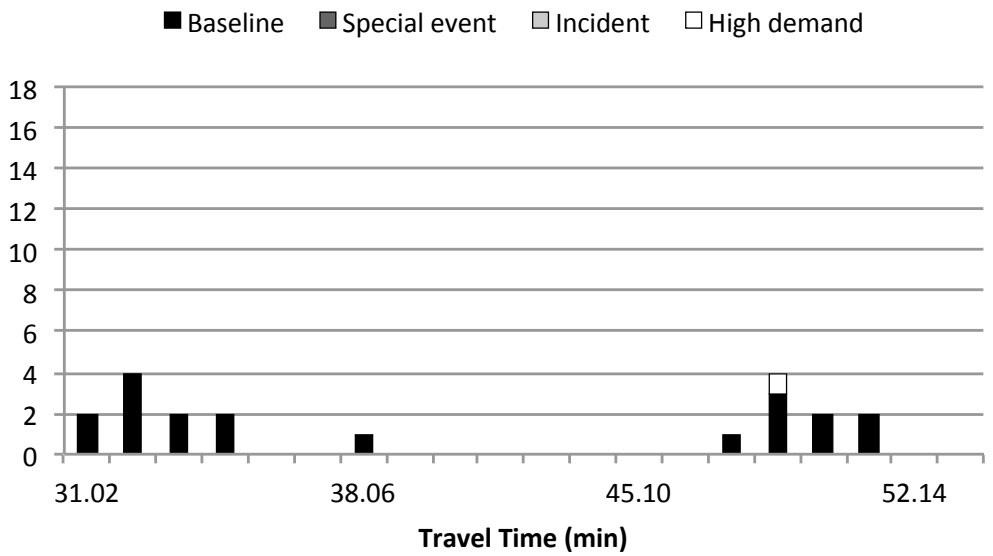

Exhibit C2-42: AM peak travel time distribution for Route #20X, August 2010

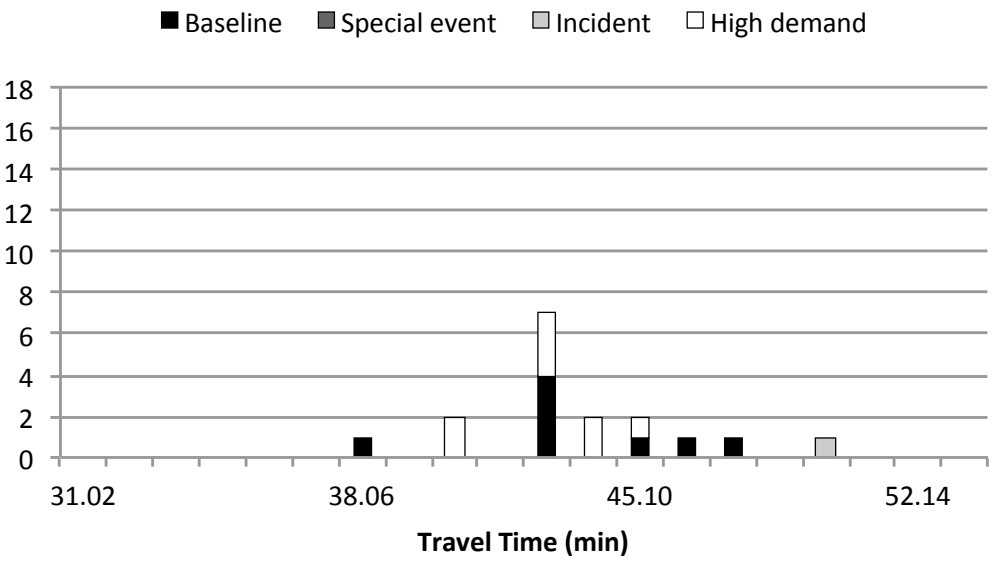

Exhibit C2-43: Midday travel time distribution for Route #20X, August 2010

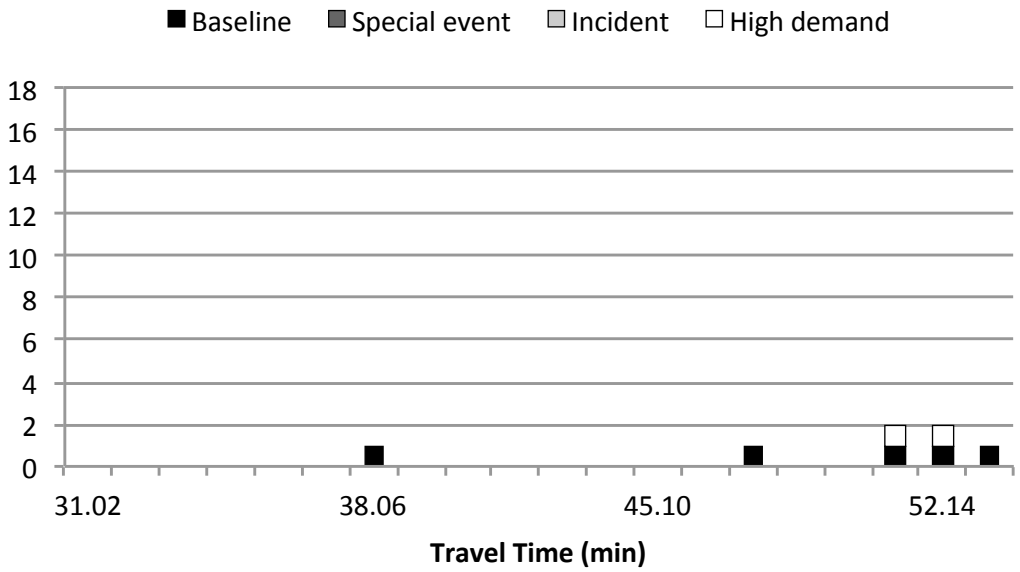

Exhibit C2-44: PM peak travel time distribution for Route #20X, August 2010

Table C2-21 summarizes the contribution of each event condition to all travel times, to those exceeding the 85th percentile (49.1 minutes), and to those exceeding the 95th percentile (51.4 minutes). It can be seen that, although 85.19% of all trips had no associated variability-inducing event, of those trips that exceeded the 85th percentile travel time, 25% were associated with either an incident or high demand, with high demand events occurring more often. Furthermore, all of these high demand trips which exceeded the 85th percentile travel time occurred during the PM Peak period. From a planning and operational standpoint, this indicates that there could be some room for reliability improvements by adding capacity to high-demand trips on this route during the PM peak.

Table C2-21: Travel time variability causality for Route #20X

	Active	**Active when travel time exceeded 85th percentile**	**Active when travel time exceeded 95th percentile**
Baseline	85.2%	75.0%	75.0%
Special Event	0.0%	0.0%	0.0%
Incident	1.2%	8.3%	8.3%
Demand	13.6%	16.7%	16.7%

Route #50, Southbound. For the subset of Route #50, Southbound considered here, travel time variability and its contributing factors were investigated for the 22 weekdays in August 2010. The period of study was limited to a single month due to a shortage of data on other months. Data on incidents, special events, demand fluctuations and travel times were collected from PeMS, external sources, and the in-vehicle APC sensors, as described in the Methods section. Scheduled travel times for the 158 runs analyzed for this route range between 18 and 21

minutes, averaging 19.5 minutes. The average delivered travel time for this route was 28.75 minutes, a full 9.25 minutes more than the average scheduled travel time. Exhibit C2-45 shows the total distribution of trip travel times by event condition over the study period.

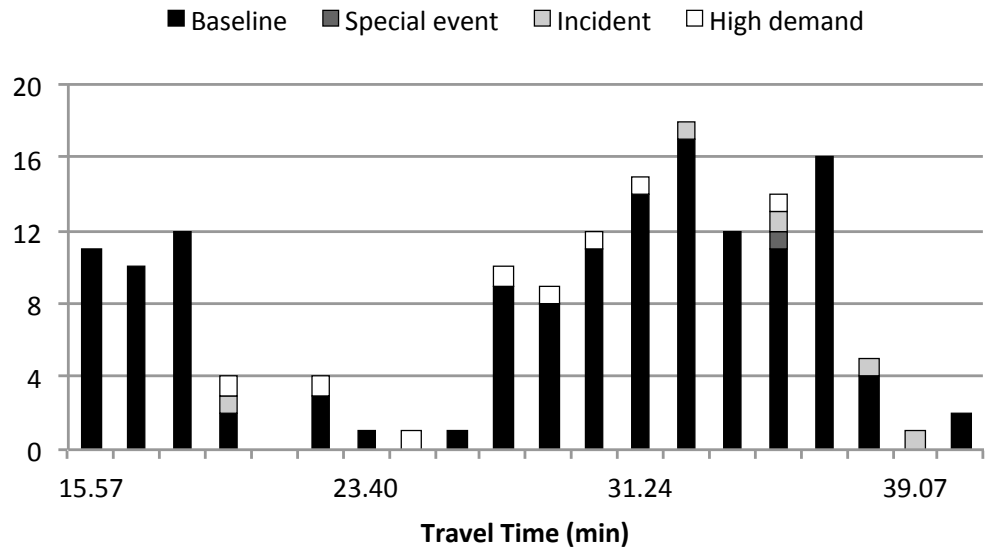

Exhibit C2-45: Complete travel time distribution for Route #50, August 2010

The AM peak distribution for this route, shown in Exhibit C2-46, appears similar to Route #20 X with two widely distributed modes appearing on either side of the distribution. However, unlike Route #20 X, this bimodal distribution was not exclusive to the AM peak period for this route. No events were flagged for trips occurring in this time period.

Similar to the other two routes analyzed here, the midday period, shown in Exhibit C2-47, carried the majority of high demand trips on this route, with four of the five high demand trips occurring here. However, continuing the trend of Routes #20 and #20 X, those high demand trips are not particularly strongly associated with longer travel times. A majority of the trips clustered around the low end of the travel time distribution occurred during the midday period.

Exhibit C2-48 depicts the travel time distribution of trips taken during the PM peak period on Route #50. Immediately visible is the apparent relationship between incident events and long travel times, as two of the three longest travel times seen during this month were associated with incidents (with the third associated with a special event).

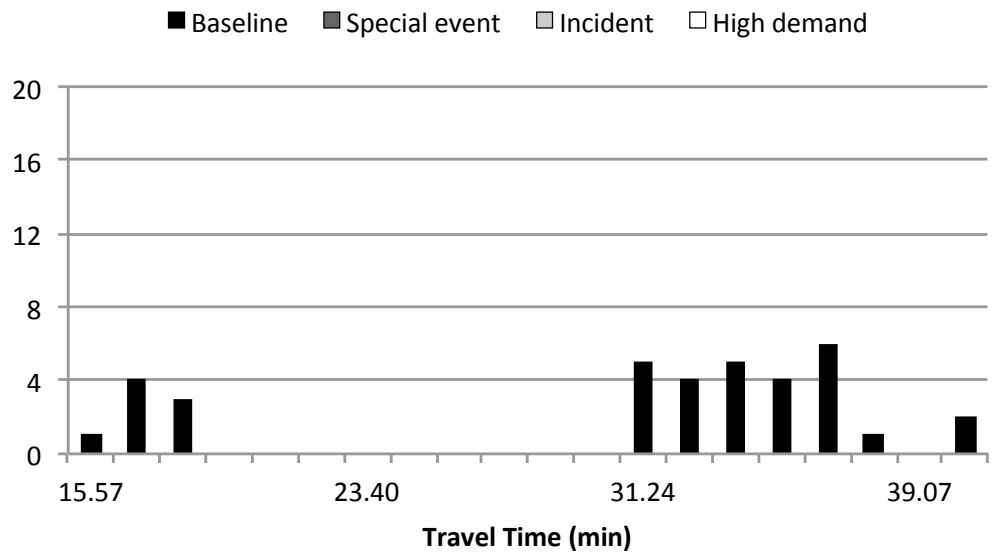

Exhibit C2-46: AM peak travel time distribution for Route #50, August 2010

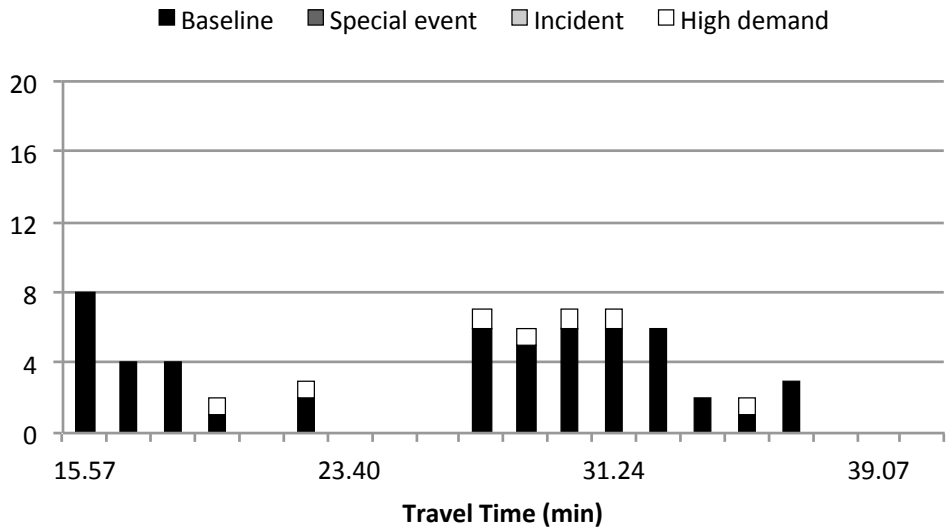

Exhibit C2-47: Midday travel time distribution for Route #50, August 2010

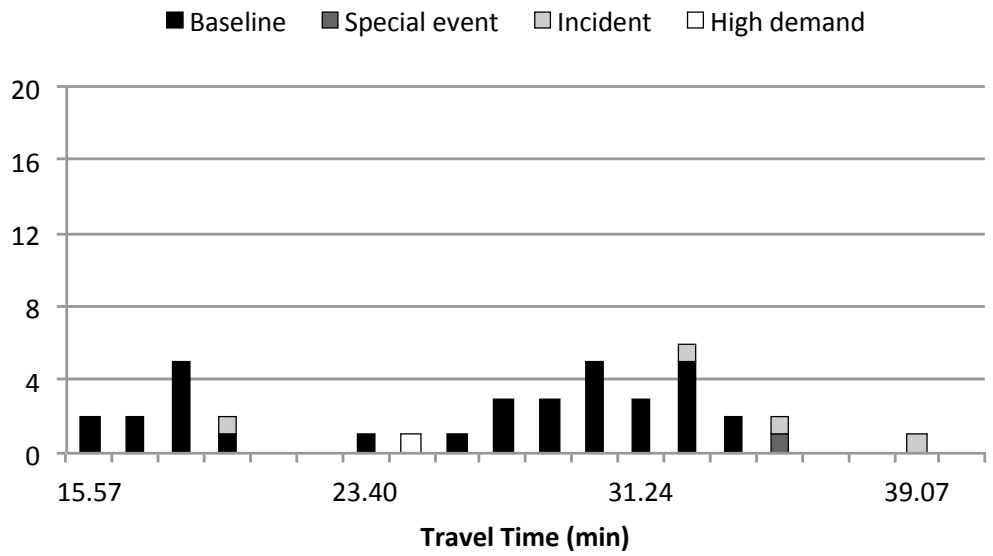

Exhibit C2-48: PM peak travel time distribution for Route #50, August 2010

Table C2-22 summarizes the contribution of each event condition to all travel times, to those exceeding the 85th percentile (35.9 minutes), and to those exceeding the 95th percentile (37.1 minutes). It can be seen that, although 92.36% of all trips had no associated variability-inducing event, of those trips that exceeded the 85th percentile travel time, the percent with no variability-inducing event dropped to 91.67%. When limiting the pool to trips that exceeded the 95th percentile travel time, a full 25% of that total can be associated with incidents (although special events were associated with just 3.18% of all trips). From a planning and operational standpoint, this indicates that there could be some room for reliability improvements by focusing more resources on clearing roadway incidents more quickly along this route to lessen the severity of their impact.

Table C2-22: Travel time variability causality for Route #50

	Active	**Active when travel time exceeded 85th percentile**	**Active when travel time exceeded 95th percentile**
Baseline	92.4%	91.7%	75.0%
Special Event	0.6%	0.0%	0.0%
Incident	3.2%	8.3%	25.0%
Demand	5.1%	0.0%	0.0%

Conclusion. This use case analysis illustrates one method for exploring the relationship between travel time variability and the sources of congestion. The methods used are relatively simple to perform provided that the transit APC data can be obtained and sufficiently cleaned. The application of the methodology to the three San Diego routes revealed key insights into how this type of analysis should be performed.

Of note is the limited sample size used in this analysis. To ensure statistical significance and meaningful analysis, ideally no less than three months' worth of data should be used to avoid invalid conclusions due to anomalies. Breaking the travel times down by time of day according

to local traffic patterns is valuable as it isolates the effects of sources of congestion by time of day. For example, on Route #20 high passenger loadings are associated with longer trip times during the PM peak period, but not at other times of day.

Use Case 2: Using planning-based reliability tools to determine departure times and travel times for a trip.

Overview. Perhaps the most commonly occurring use case related to transit data is the case of the transit user seeking information about the system for trip planning purposes. This happens thousands of times each day in cities across the country, and with good reason. The dissemination of traveler information such as real-time arrivals, in-trip guidance, and routing can lead to a more satisfactory transit experience for the user and potentially increase ridership.

Conversely, uncertainty can also have a significant effect on the traveler experience. The agony associated with waiting for transit service has been well documented; research suggests that passengers overestimate the time they spend waiting by a factor of 2 to 3 compared to in-vehicle time (3). Meanwhile, driving is often perceived as offering travelers a greater sense of control when compared to other modes. Offering transit users accurate and easily accessible information on the transit system, while certainly stopping short of providing direct control over the trip, can give peace of mind to transit riders, reducing uncertainty along with the discomfort of waiting for service. As the reliability of this information improves, so will the experience of transit users.

The use of planning-based reliability tools to determine departure times and/or travel times for a trip therefore has the potential to improve passenger understanding of the state of the transit network, leading to less uncertainty and greater ease of use of the transit system.

Site Characteristics. Transit agencies are rarely able to equip their entire fleets with Automated Passenger Count (APC) or Automated Vehicle Location (AVL) sensors, making it difficult to conduct a thorough analysis of the entire network. For San Diego's bus network, approximately 40% of buses have APC/AVL sensors installed, though not all of these sensors are fully operational. Due to malfunctioning sensors and limitations in the distribution of APC/AVL equipped vehicles, in reality only 30% of San Diego routes are covered by transit vehicles equipped with functioning APC/AVL sensors.

The San Diego #30 North bus route was chosen for this study primarily because it is the route for which the largest quantity of APC data was available for the period of study (August 2010). A subset of the route spanning from the Grand Avenue exit on Highway 5 along the coast to the intersection of Torrey Pines Road and La Jolla Shores Drive (8.13 miles) was chosen for this study.

For comparative purposes, the San Diego #11 North bus route was also examined. This route also contains a comparatively large amount of APC data for August 2010. It travels through the Southcrest neighborhood at 40th Street and National Avenue West on National Avenue, through downtown and north on 1st Avenue to University Avenue and Park Boulevard. The total length for the portion of the route analyzed here is 11.68 miles. Both routes are shown in Exhibit C2-49.

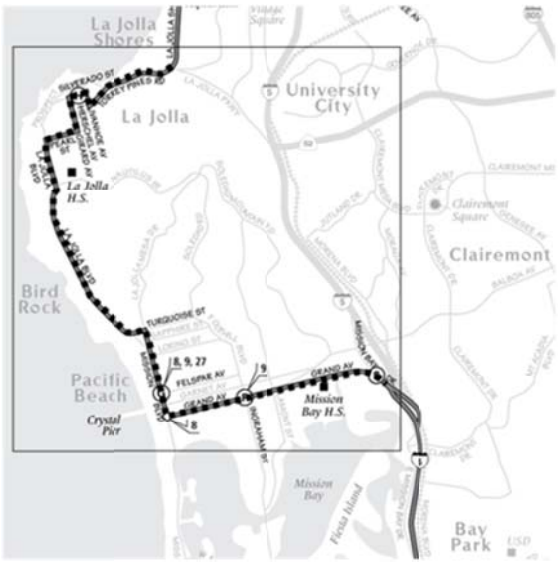

Route #30

Route #11

Exhibit C2-49: Analyzed portions of #30 and #11 bus routes

Data. The data used in this analysis was obtained from SANDAG. It is APC data collected from August 1 to August 31, 2010, and it consists of measurements taken every time the vehicle opens its doors. Each data point contains the following variables, among others:

- Operator ID
- Vehicle ID
- Trip ID
- Route ID
- Door open time
- Door close time
- Number of passengers boarding
- Number of passengers alighting
- Passenger load

Notably absent from this data is any kind of service pattern designation, which is necessary to group similar trips together for comparison purposes. Route ID is not a sufficient level at which to group trips, since a single route often consists of multiple service patterns (e.g., express patterns and alternate termination patterns). This means that the APC data must be preprocessed in order to identify which trip measurements can be grouped into the same service pattern.

The APC passenger count data are collected by detecting disturbances of dual light beams positioned at the doors of the transit vehicle. Boardings and alightings are detected based on the order in which the beams are broken by a passenger entering or exiting the vehicle. This data can be unreliable as some preprocessing of the data occurs on the sensor itself; specifically, the passenger load is never allowed to drop below zero.

For the subset of the #30 route considered here, scheduled trip times range from 32 minutes to 38 minutes, and scheduled headways range between 13 and 46 minutes (the mean scheduled headway is 21.6 minutes). Approximately 700 vehicle trips over 20 weekdays in August 2010 were analyzed. Of the APC data for this entire route, 50% is imputed. It is necessary to impute data for points where the measured data are missing or does not make physical sense. For example, if a given transit stop has no passengers waiting at it, and no riding passengers have requested a stop there, it is common for the transit vehicle to skip this stop. This results in a missing APC data point for that stop that must be imputed. This practice is particularly common at the beginnings and ends of runs, thus for this subset of the route it is expected that the percent of data imputed is lower than 50%.

For the subset of the #11 route considered here, scheduled trip times range from 40 minutes to 56 minutes, and scheduled headways range between 15 and 76 minutes (the mean scheduled headway is 30 minutes). Approximately 850 vehicle trips over 20 weekdays in August 2010 were analyzed. Of the APC data for this route, 53.20% is imputed.

Approach. Most other analyses of AVL and APC data consider transit trip components (e.g., run time, dwell time, and headways) separately (4, 5, 6, 7). This can be considered an agency-centric approach as it attempts to answer questions that a transit system operator may be interested in such as "How are dwell times affecting on-time performance?" and "What is an appropriate layover time?".

In this analysis, we combine headways and in-vehicle travel times in order to view transit performance measurement from a more passenger-centric perspective. The service experienced by the passenger is studied by focusing the analysis on answering the fundamental passenger question "If I were to go to the bus stop at a certain time, when would I arrive at my destination?".

This study assumes that passengers do not plan their transit trips according to real-time or scheduled data, but rather follow a uniform arrival pattern throughout the day, beginning their transit trips independently of the state of the system.

Methods. To begin this validation, the literature was surveyed to determine the recommended planning-based means for calculating the best departure time for a trip in a general way. An appropriate departure time will take into account the variability within the transit system, while being calculated in a way that is intuitive and useful to users.

The SHRP2 L02 Task 2-3 report presents the results of focus group interviews conducted with passenger travelers which attempted to uncover the most meaningful travel time metrics for different trip scenarios. The results show that for daily, unconstrained trips, planning time is the most appropriate metric for passengers. Planning time is a travel time metric that accounts for variability within the system, representing a percentile (often the 85^{th} or 95^{th}) travel time for a trip. That is to say, the planning time for a trip is the travel time that should be accounted for in order for the traveler to be on time a certain percentage of the time. "Trip" here is taken to mean a pattern of movement between two points at a certain time of day, thus planning time is always computed based on travel times for a single trip over a range of dates.

In order to satisfy this use case and determine the planning time for a transit trip, we must find the travel times for a single trip over a range of days. It is possible to calculate such a table based on APC data alone. To do this:

1) We choose 8.13 miles of the #30 North route (from the Grand Avenue exit on Highway 5 along the coast to the intersection of Torrey Pines Road and La Jolla Shores Drive) to analyze for this use case.

2) We use the APC data to measure actual travel times for trips along this route beginning every two minutes throughout the day. These trips begin independently of the bus schedule.
3) We repeat the previous step for each of the dates in the study range.
4) We now have a table whose columns are dates, rows are times of day, and values are travel times along this transit route. We compute the PDF distribution of travel time for each of the trips in this table.

The notion of computing such a table of travel times is common in highway performance measurement, but less common for transit performance measurement, which tends to focus on travel time in relation to a schedule (schedule adherence) rather than absolute travel time.

The results of this analysis for August 31, 2010 can be seen in Exhibit C2-50. The troughs correspond to trips that begin immediately before the departure of a bus. The peaks represent trips that began just after the departure of a bus. The steadily downward sloping lines following peaks indicate trips that begin between bus departures; the trips within a single downward sloping section are related in that they all go on to travel on the same bus, whose arrival is indicated by the following trough. The travel times are complemented by a Marey graph of the trips for this day. It can be seen that the troughs correspond to bus departures.

A similar Marey graph and travel time plot, also for August 31, 2010, are shown below in Exhibit C2-51for Route #11, North.

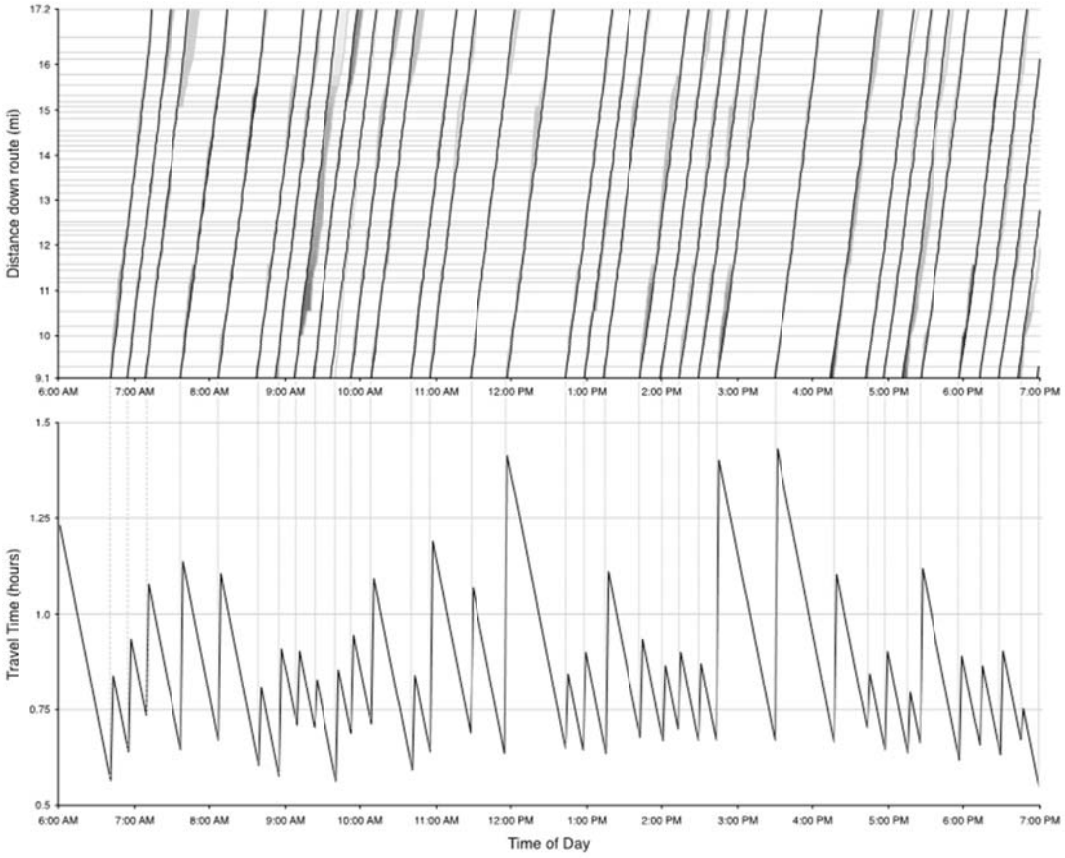

Exhibit C2-50: Marey graph (top) and passenger travel times (bottom) by time of day for Route #30 on 8/31/2010

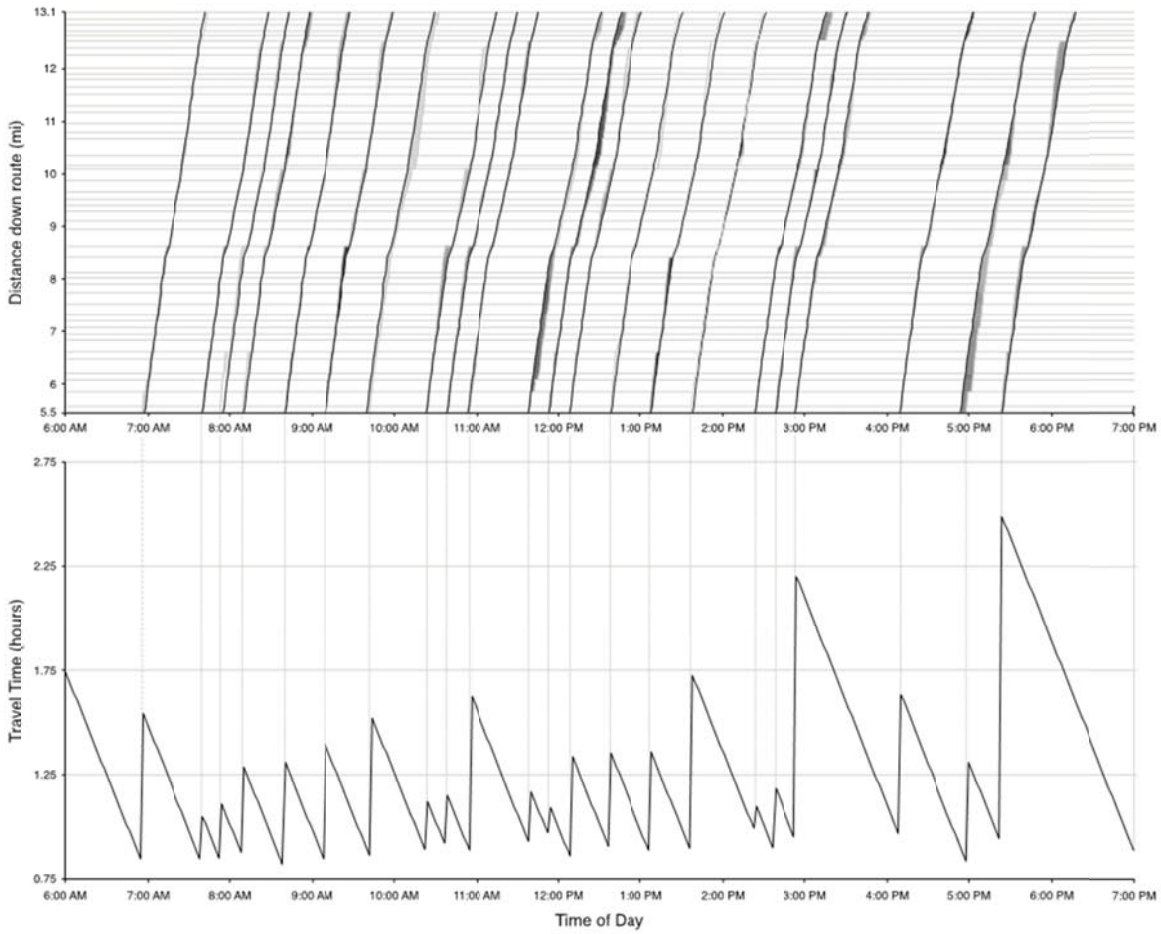

Exhibit C2-51: Marey graph (top) and passenger travel times (bottom) by time of day for Route #11 on 8/31/2010

Results. Analyzing multiple days yields statistical measures of travel time variability. Here, 22 weekdays in August 2010 are analyzed following the preceding methodology to obtain Exhibit C2-52, which depicts average travel time as well as the distribution of travel times along the vertical axis, with darker shading corresponding to higher frequency.

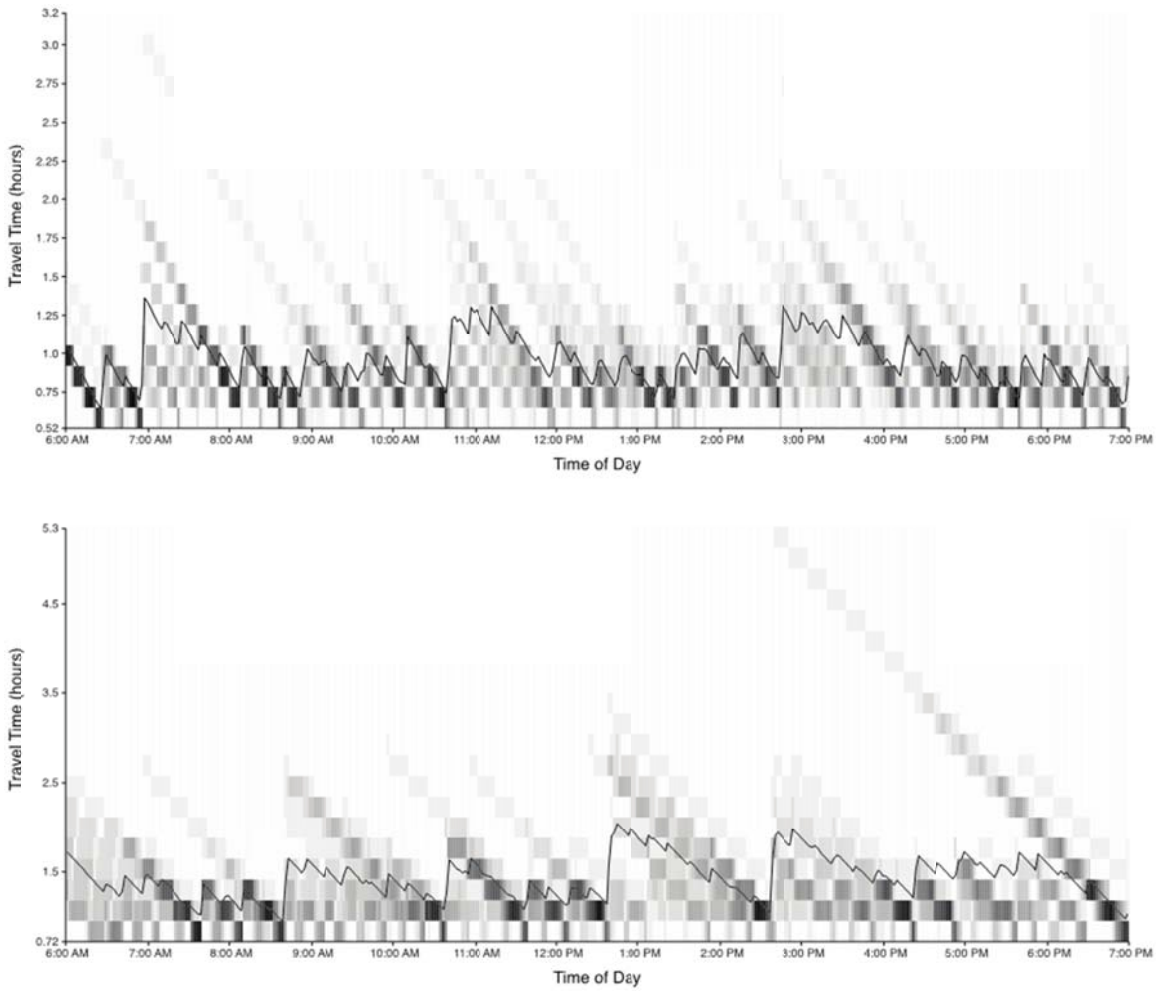

Exhibit C2-52: Planning time for trips on Route #30 North (top) and Route #11 North (bottom)

All that remains to complete the validation of this use case is to select a desired arrival time and subtract the expected travel time from it. The expected travel time can be extracted from the distributions presented in Exhibit C2-52, and a range of expected travel times are given. Interpolation may be necessary to obtain precise arrival times depending on the sample size. Table C2-23 and Table C2-24present departure times and travel times resulting from this analysis. Because bus departures are discrete and not continuous events, it is possible that a range of departure times can correspond to a single arrival time. This effect goes away with larger sample sizes.

Table C2-23: Departure times and planning times on Route #30 North

75th Percentile departure time	**75th Percentile travel time**	**85th Percentile departure time**	**85th Percentile travel time**	**95th Percentile departure time**	**95th Percentile travel time**	**Arrival time**
6:54 AM	1h 6m	6:53 AM	1h 7m	6:52 AM	1h 8m	8:00 AM
10:06 AM	54m	9:53 AM	1h 7m	9:45 AM	1h 15m	11:00 AM
2:02 PM	58m	2:00 PM	1h	1:58 PM	1h 2m	3:00 PM
4:03 PM	57m	4:00 PM	1h	3:27 PM	1h 33m	5:00 PM

Table C2-24: Departure times and planning times on Route #11 North

75th Percentile departure time	**75th Percentile travel time**	**85th Percentile departure time**	**85th Percentile travel time**	**95th Percentile departure time**	**95th Percentile travel time**	**Arrival time**
6:26 AM	1h 34m	--	--	--	--	8:00 AM
8:55 AM	2h 5m	8:41 AM	2h 19m	8:40 AM	2h 20m	11:00 AM
12:40 PM	2h 20m	12:38 PM	2h 22m	12:38 PM	2h 22m	3:00 PM
2:54 PM	2h 6m	2:52 PM	2h 8m	2:37 PM	2h 23m	5:00 PM

Conclusion. The most direct analysis would be achieved by restricting the date range to dates with identical schedules, however, in practice it can be rare to find days with the exact same schedule. Regardless, for routes with headways smaller than 10 minutes it is common for passengers to arrive at bus stops independently of the schedule, thus the constant arrival pattern used in this simulation may be more meaningful.

Agencies should strive to either reduce transit travel times across the day, or establish reliable times of day when the transit travel time can be expected to be low. As seen in the transition between Exhibit C2-50 and Exhibit C2-51, as more days are added to the analysis, the strong peaks correlating to regular bus departures can become obscured if the transit schedule is not regular day to day. This results in the slightly blurry look of the distributions in Exhibit C2-52. However, if a period of study is selected in which the transit schedule is fixed, the troughs will always appear in the same locations indicating good reliability across days from the transit user's perspective.

Use Case 3: Analyzing the effects of transfers on the travel time reliability of transit trips

Summary. The goal of this use case is to demonstrate a methodology for quantifying the effects of missed transfers on travel time (and travel time reliability) for a particular transit trip. The likelihood of a transfer being missed is predicted based on three factors: the measured

performance of the vehicles on the route, the schedule, and an assumed passenger arrival distribution. In this use case, two transfer trips in San Diego are simulated and the resulting passenger travel time histograms (accounting for the effects of missed transfers) for each route are presented. The delay applied when a transfer is missed is based on the vehicles' measured performance as well as the schedule. Practically, this methodology could aid in the identification of a pair of buses whose chronic schedule deviations at a particular location are likely to cause missed transfers.

Missed transfers in a transit system are rarely monitored, despite the problems they cause for passengers. In practice, transit systems are most often evaluated according to the performance of individual vehicles, stops, and routes, not the interactions between them. In contrast, the likelihood of a missed transfer occurring depends on combinations of several factors, making it hard to estimate. This use case takes a systems approach to quantify the effects of three distributions: passenger arrival rate, on-time vehicle performance, and schedule-based transfer time on passenger travel time distributions. Additionally, a sensitivity analysis is used to isolate the effects of changes in each of these three distributions on the percentage of transfers predicted to be missed and the total passenger travel time histogram for the route.

The simulation techniques found in this use case are made possible by the increasing availability of data from APC and AVL systems. These data are typically rich, containing vehicle arrival times and passenger loading information at stops along a route often accompanied by contextual geographic information to relate records from multiple vehicles. All simulations carried out in this use case are based entirely on APC data from the San Diego bus system, the bus schedule, and an assumed passenger arrival distribution.

Users. The anticipated users of this case study are transit agency operators with an interest in minimizing missed transfers and their negative effects on passenger travel time. Operators of transit agencies with APC data collection systems in place will find guidance on how to use their observed schedule adherence data to identify the predicted rates of transfers missed between a given pair of vehicles. Techniques such as schedule or route adjustments can then be used to reduce the rate of missed transfers and decrease passenger travel times.

Transit passengers are expected to be the prime beneficiaries of this use case. For the passenger, missing a transfer that should have been available according to the schedule is costly in terms of increased travel time and stress. Computer-based trip planners almost exclusively route passengers across transfers based on the transit schedule, not real-time data. Furthermore, trip planners can often recommend routes that transfer at unofficial transfer points. This means that any time a transfer is missed (i.e., the scheduled arrival order of two buses at a stop is reversed due to schedule deviations), passengers may be affected, even if the transfer was officially untimed. Any efforts to reduce passenger travel times across the system must consider the effects that missed transfers can have on overall system travel times and travel time reliability.

Site. San Diego's transit network is extensive and well connected, containing many transfer points. This makes it an ideal test setting. It includes 88 bus routes and several light rail lines. Most importantly, many buses in this system are equipped with APC equipment to monitor on-time performance. Two routes containing transfers through San Diego were selected for this analysis and are described in Table C2-25. These routes were chosen for their popularity with riders as well as their high data coverage rates. Maps of the routes are shown in Exhibit C2-53. Route A is from the Gaslight District to the San Diego Zoo and has a predicted travel time of 39

minutes. Route B is from Sea World to the Birch Aquarium and has a predicted travel time of 55 minutes.

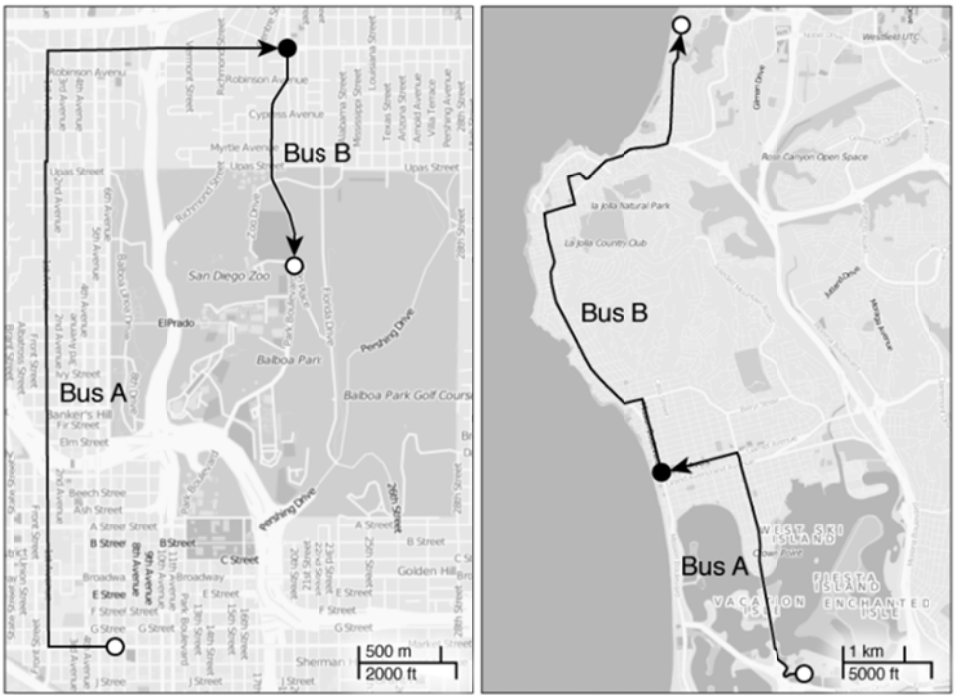

Exhibit C2-53: Routes A (left) and B (right). Route endpoints are hollow. The transfer point is filled in.

Trip times are simulated from APC data collected on these buses. This data originates from location-tracking devices installed directly in the buses themselves, and is based on GPS technology. Each APC device keeps a detailed event-based record of the vehicle's performance as it drives along the route. A data point is created every time the vehicle makes a stop. For this use case, the relevant elements in each data point are:

- The name of the route that the bus was on (e.g., Route 30).
- A unique ID corresponding to the individual run being made.
- A unique ID corresponding to the stop at which the record was made, enabling stops across routes to be cross-referenced.
- The time when the doors opened at the stop.
- The time when the doors closed at the stop.
- The scheduled time when the stop was supposed to be made.

Table C2-25: Route characteristics

Route	Bus A Distance	Bus B Distance	Total Distance	Transfer Location	Estimated Time*
A: Gaslight to the Zoo	3.7 miles	1.0 miles	4.7 miles	Park Blvd. and University Ave.	39 minutes
B: Sea World to Birch Aquarium	3.7 miles	6.4 miles	10.1 miles	Mission Blvd. and Felspar St.	55 minutes

*Estimated Time is from the San Diego MTS trip planner for a trip departing at 10:00AM on a weekday.

Methods. This section describes the data preparation and trip time methodologies.

Data Preparation. In order to relate trip times on these transfer routes to the on-time performance of the buses serving them, several issues with the raw APC data must first be addressed. Most critically, the data are not a complete record of all vehicle activity throughout the system. Only a portion of the vehicle fleet is instrumented with APC equipment, and certain routes have higher coverage rates than others. With data available on only a fraction of the runs, gaps in data coverage become problematic, particularly when exploring missed transfers. Because of the missing data, the number of directly observed transfers between two buses at a given stop and time is relatively low, as either the arriving or departing bus will often be uninstrumented. This means that (in this setting) it is impossible to simply observe the missed transfers and total trip times directly.

To circumvent this problem of incomplete instrumentation, a simulation-based method is used. This method works on the assumption that the on-time performance of the runs for which APC data exists is representative of the on-time performance of all trips. Rather than directly observing on-time performance that would result in a missed transfer, a large number of virtual trips on Routes A and B are simulated based on APC data, an empirical passenger arrival distribution [7], and the schedule.

The APC data contributes distributions of arrival schedule adherence, departure schedule adherence, and travel times for the relevant buses and stops. In order to construct these distributions accurately, only data from runs that serve both the origin and transfer (or transfer and destination) stops should be included. Grouping the data into service patterns facilitates this. A service pattern is a finer unit of organization than a route and represents a grouping of trips that share the same stops in the same order. Route variations with alternate termination points or express service are examples of distinct service patterns within the same route. To detect service patterns in the data, repeating patterns of stops made by different vehicles within a single route were identified. Runs were then labeled according to the service pattern they represent. Considering APC traces at the service pattern level instead of the route level allows data from trips that do not serve the desired stops to be discarded.

The inclusion of the passenger's arrival time at the origin in the simulated trips means that there are actually two transfers on each route (from walking to Bus A and from Bus A to Bus B). Thus, the simulated passenger can either catch both buses, miss only Bus A, miss only Bus B, or miss both buses. The passenger arrival time distribution is based on a distribution empirically derived by Bowman and Turnquist, scaled to the 15-minute headway of Bus A (on both Routes A and B) [7].

The distribution of schedule-based transfer times was constructed based on the daytime weekday schedule for Buses A and B on each route. The transfer times for both routes are irregular as they are untimed. However, despite their irregularity, in each there was some correlation between consecutive transfer times. For example, if one transfer time was short, the following transfer time was scheduled to be longer. Because of this, missed connections at the transfer point were assessed a travel time penalty corresponding to the transfer time immediately following the one that was missed (without another independent sample). This additional travel time is the same as Bus B's headway at that time of day. The relevant distributions used to simulate travel times on Route B are shown in Exhibit C2-54.

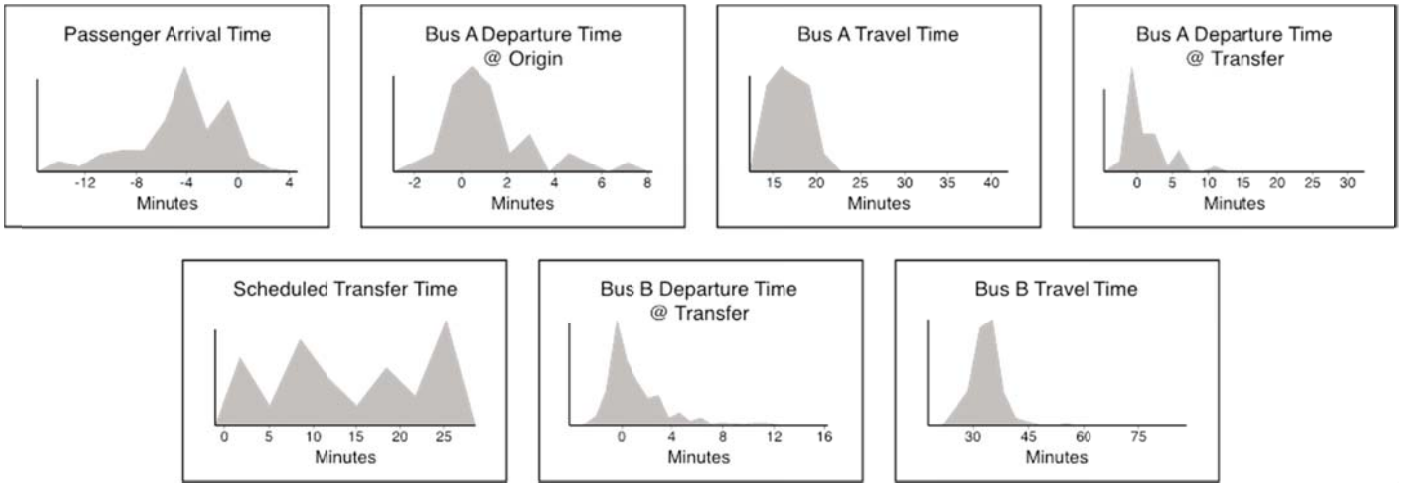

Exhibit C2-54: Distributions used to simulate travel times on Route B.

Several of these distributions correlate with each other, affecting how the samples are drawn in each simulation (see Exhibit C2-55, for example). On both routes, there was found to be some correlation between Bus A's departure time at the origin, Bus A's travel time, and Bus A's departure time at the transfer point. That is to say, a bus that departed late from the origin was more likely to be late when it left the transfer point. Correlation between Bus B's departure time at the transfer point and Bus B's travel time was also found. Because of these relationships between the distributions, simulated trips must not sample values from these related distributions independently. In a single travel time simulation, the values sampled from correlated distributions must come from the same APC trip record because they are related.

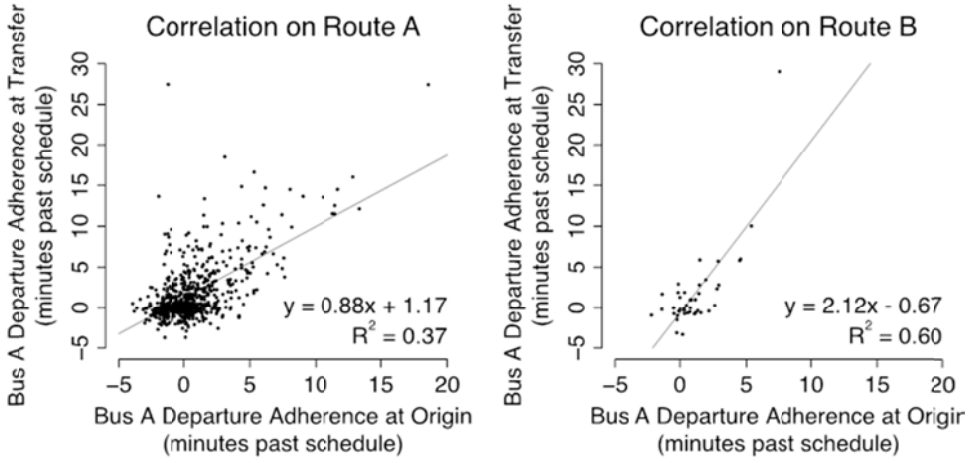

Exhibit C2-55: Positive correlation between Bus B's departure time at the transfer stop and its arrival time at the destination on Routes A and B.

Approach to Obtain Trip Times. The procedure for determining a single trip time can be seen in Exhibit C2-56. To begin, values are randomly sampled from the Bus A departure and passenger arrival distributions. These values (both relative to Bus A's scheduled departure at the origin) are then compared to determine whether or not Bus A is caught. If the departure time for Bus A is greater than the passenger's arrival time, the first bus is caught. Otherwise, the first bus is missed (as a result of the passenger's late arrival, the bus departing early, or some combination of the two). If the bus is missed, a single Bus A headway is added to the total trip time to represent the time spent waiting for the next bus. For both Routes A and B, Bus A maintained regular 15-minute headways during the daytime on weekdays.

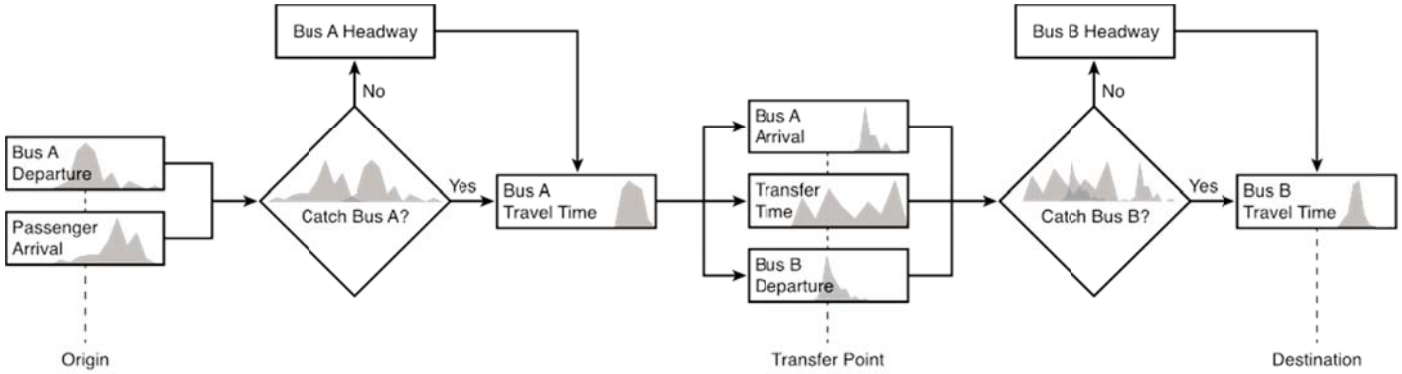

Exhibit C2-56: Procedure followed to generate travel times

Once Bus A is caught, the Bus A travel time value from the same data record as Bus A's departure from the origin is added to the total trip time, bringing the virtual passenger to the transfer point. Whether or not the transfer is made depends on three things: Bus A's departure time from the transfer point relative to the schedule, the scheduled transfer time, and Bus B's departure from the transfer point relative to the schedule. In order to be conservative and to acknowledge the time required by the passenger to move between buses, the measured time that

Bus A departs from the transfer point is actually used to construct the distribution of Bus A's arrival time at the transfer point. This represents a worst-case scenario. If Bus A's departure adherence is earlier than the sum of Bus B's departure adherence and the scheduled transfer time, Bus B is caught. Otherwise, Bus B is missed. Because of their correlation, the value used to represent Bus A's arrival at the transfer point is taken from the same run in the APC data as Bus A's schedule adherence at the origin and Bus A's travel time.

If Bus B is missed, a penalty of one Bus B headway is assessed to the trip time. For both routes A and B, Bus B's headways were irregular. Because of this, the time until the arrival of the next consecutive Bus B is taken, as opposed to simply sampling another transfer time value, or an independent headway from Bus B. After the transfer, the Bus B travel time value from the same run as the sampled Bus B transfer point departure is applied to the total trip time.

This completes the simulation, and the total travel time is computed as the sum of its components. The arrival and departure adherence distributions (passenger arrival time, Bus A's departure time at the origin, Bus A's departure time at the transfer point, and Bus B's departure time at the transfer point) are all in terms of schedule adherence: *actual time – scheduled time*. The other travel time and transfer time distributions are magnitudes of time. This process was repeated 10,000 times for each route in order to obtain travel time histograms that accurately reflect the sample distributions.

Results. This section describes the results for the different routes.

Route A: Gaslight to the San Diego Zoo. A simulation of 10,000 trips on Route A produces the probability density function for travel time shown in Exhibit C2-57. The shortest travel time is 22 minutes and the longest is 96 minutes. The 50th percentile is reached at 47 minutes and the 95th percentile is reached at 62 minutes. The average is 47 minutes. The longest travel time is 104% longer than the mean and 336% as long as the shortest time. Guidance to potential passengers might be that they should expect the trip to take 47 minutes but one out of every 20 trips will take longer than 62 minutes.

The histogram of travel times appears normally distributed with a portion of the simulated travel times skewed to the right. These trips represent times when a very long in-vehicle travel time was sampled for one of the legs of the trip, not necessarily trips where a connection was missed. Further insight into travel times on this route can be gained by dividing the simulated trips into those that made or missed each bus.

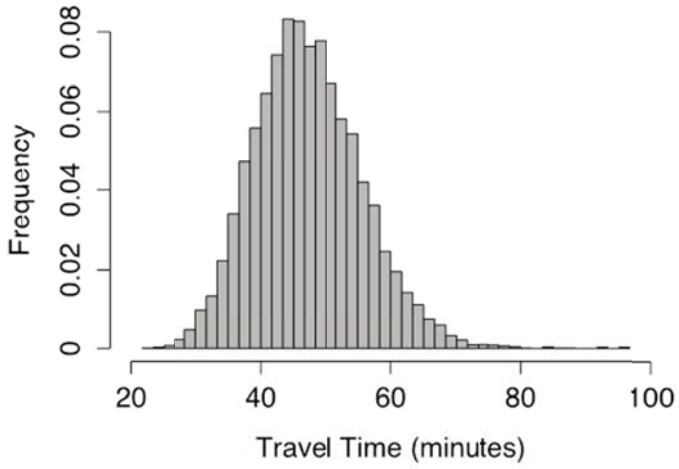

Exhibit C2-57: Histogram of 10,000 simulated trips on Route A

Table C2-26 presents a breakdown of the simulations by scenario. Four out of five simulated trips were able to catch both buses and enjoyed shorter average travel times. The median travel time increased by roughly 9 minutes for each bus that was missed. Surprisingly, the trip time histogram for trips that missed buses were more tightly grouped (and thus had better travel time reliability) than those that made both buses. This is discussed in further detail in the following section.

Table C2-26: Travel time distributions under different trip scenarios on Route A

	Percentage	Minimum (min)	Median (min)	95th Percentile (min)	Maximum (min)	Mean (min)	Standard Deviation (min)
Make Both	81.71%	22	45	59	96	46	7.78
Miss A, Make B	7.69%	36	54	66	96	54	7.05
Make A, Miss B	9.76%	34	52	65	79	53	6.84
Miss A, Miss B	0.84%	42	63	69	71	62	5.25
Total	100%	22	47	62	96	47	8.25

Route B: Sea World to the Birch Aquarium. A simulation of 10,000 trips on Route B produces the probability density function shown in Exhibit C2-58. The shortest travel time is 42 minutes and the longest is 138 minutes. The 50th percentile is reached at 66 minutes and the 95th percentile is reached at 85 minutes. The average travel time is 67 minutes. Thus, the longest travel time is 109% longer than the mean and 229% as long as the shortest time. Guidance to potential passengers might be that they should expect the trip to take 66 minutes but one out of every 20 trips will take longer than 85 minutes. The histogram of travel times appears approximately normal with a longer tail of high travel times.

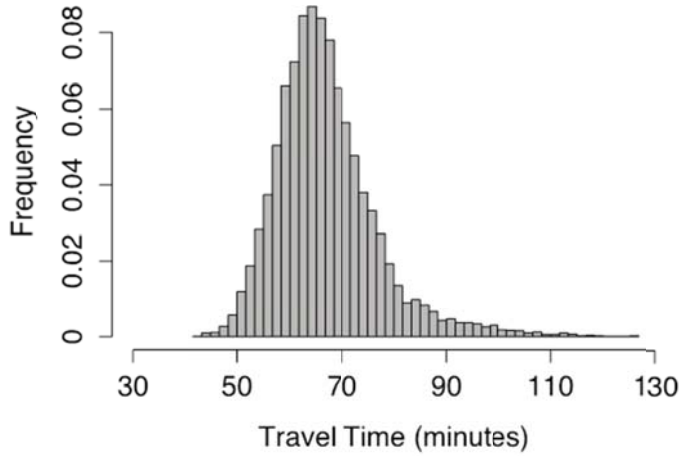

Exhibit C2-58: Histogram of 10,000 simulated trips on Route B

On Route B, whether or not Bus A and/or Bus B were missed was tracked for the purposes of exploring the effects of missed transfers on travel time. Travel time histograms corresponding to each scenario are plotted in

Exhibit C2-59 and described in Table C2-27. Clearly, missing one or more buses leads to increased travel times on this route, although (as with Route A) the travel time reliability actually improves as well.

This apparent improvement in reliability may be unexpected but according to Exhibit C2-56, simulated trips that miss Bus A or Bus B are subjected to no or little additional randomness. If Bus A is missed, a predetermined 15-minute headway is added to the trip time. If Bus B is missed, a Bus B headway (ranging between 13 and 16 minutes) is added to the trip time (note that the standard deviation is greater when missing Bus B than when missing Bus A). Thus, the smaller standard deviations when missing buses are attributed to the smaller sample sizes and the presence of outliers in the "make both" case.

The presence of a few extremely long travel times for Bus A and Bus B on each route contributed to these patterns. With a greater number of simulations catching both buses on each route, more "make both" simulated trips had the opportunity to experience an extremely long in vehicle travel time. Thus, the rare occurrence of an extremely long travel time (roughly twice as long as the average travel time in this data) can have a greater effect than the occasional missed bus. However, it is important to note that trips that miss one or more buses do so unexpectedly, so even though the reliability in those scenarios is improved, the passenger cannot plan for them, and their existence diminishes the reliability of the trip as a whole.

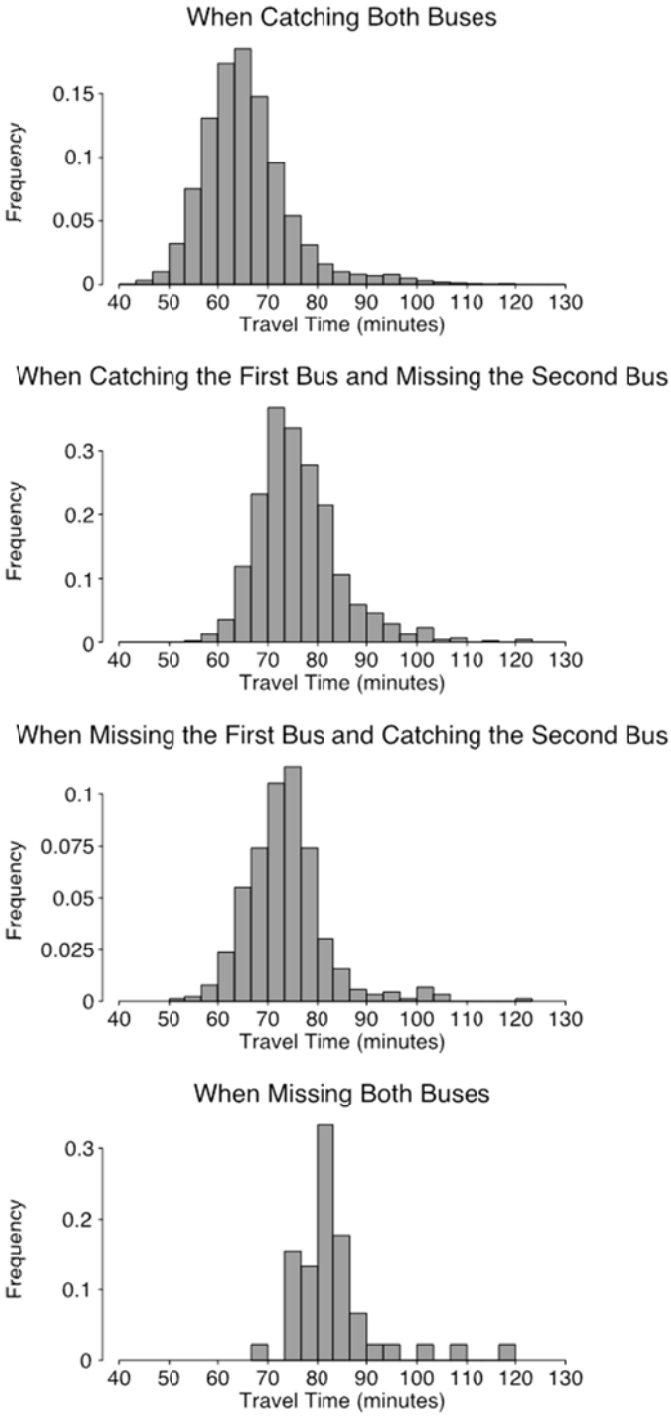

Exhibit C2-59: Travel times when catching or missing buses on Route B.

Table C2-27: Travel time distributions under different trip scenarios on Route B

	Percentage	Minimum (min)	Median (min)	95th Percentile (min)	Maximum (min)	Mean (min)	Standard Deviation (min)
Make Both	85.93%	42	65	82	138	66	9.22
Miss A, Make B	4.71%	51	73	86	122	74	8.10
Make A, Miss B	8.91%	56	75	92	122	76	8.53
Miss A, Miss B	0.45%	69	82	100	117	83	8.67
Total	100%	42	66	85	138	67	9.75

Discussion. A sensitivity analysis comparing the effects on various measures of travel time (as well as the percentages of simulated passengers who miss one or more buses) on Route B is presented in Table C2-28. The baseline case represents the results of the simulation with all distributions unaltered. The passenger arrival distribution, Bus B's departure adherence at the transfer stop, and the scheduled transfer time are then each incrementally shifted or scaled and 10,000 trips with the adjusted distributions are simulated. The scheduled transfer time was held at zero instead of allowing it to go negative.

Table C2-28: Sensitivity analysis on Route B

	Mean (min)	Median (min)	95th Percentile (min)	Standard Deviation (min)	Make Both	Miss A, Make B	Make A, Miss B	Miss A, Miss B
Baseline	67	66	85	9.75	85.93%	4.71%	8.91%	0.45%
Pax Arrival + 1 min	67	65	85	9.93	82.54%	7.33%	9.29%	0.84%
Pax Arrival + 2 min	67	66	85	9.84	76.28%	14.57	7.65%	1.50%
Pax Arrival + 3 min	67	66	86	10.17	66.49%	23.90%	6.95%	2.66%
Pax Arrival * 1.2	68	67	87	10.03	86.86%	3.38%	9.35%	0.41%
Pax Arrival * 1.4	69	68	88	10.46	87.05%	3.24%	9.30%	0.41%
Bus B Departure – 1 min	68	67	87	10.25	83.70%	3.35%	12.47%	0.48%
Bus B Departure – 2 min	68	67	87	10.40	79.23%	3.24%	17.04%	0.49%
Bus B Departure – 3 min	68	67	86	10.21	72.53%	2.87%	23.64%	0.96%
Bus B Departure * 1.2	69	68	88	10.35	86.65%	3.77%	9.09%	0.49%
Bus B Departure * 1.4	69	68	89	10.62	86.11%	3.70%	9.83%	0.36%
Scheduled Transfer Time – 1 min	68	67	87	10.23	83.41%	3.36%	12.63%	0.60%
Scheduled Transfer Time – 2 min	68	67	86	10.21	80.05%	3.27%	16.02%	0.66%
Scheduled Transfer Time – 3 min	68	66	86	10.25	76.51%	3.05%	19.54%	0.90%
Scheduled Transfer Time * 0.8	67	66	85	9.77	85.08%	3.37%	11.02%	0.53%
Scheduled Transfer Time * 0.6	66	65	84	9.65	80.09%	3.64%	15.61%	0.66%

Each of these fifteen alternative scenarios is designed to disrupt transfers. However, missed transfers do not directly affect in-vehicle travel time, which makes up the bulk of the total travel time. For example, when Bus B's departure was shifted 3 minutes earlier, 15.24% more passengers missed the second bus, with each of those passengers experiencing a delay of one Bus B headway. However, the mean travel time in this scenario only increased by one minute. This could be because the duration of the transfer is a relatively small part of the total trip time on Route B due to its length. Also, shifting departures earlier makes all trips in which the bus is not missed start sooner, decreasing wait times overall and offsetting increases in the mean and median due to missed connections. This suggests that traditional performance metrics (even reliability-based metrics) may not be capable of capturing the full effects of missed transfers.

When the scheduled transfer time is confined to a tighter range, travel time reliability (as measured by standard deviation) increases. This is because the distribution of transfer times has such a wide range on this route (from 1 to 26 minutes) that when those long transfer times are cut nearly in half (as in the *Scheduled Transfer Time * 0.6 case*), each simulation benefits equally from reduced transfer times, even though the percentage of passengers who miss Bus B rises.

Conclusion. This use case has leveraged a simulation-based approach to demonstrate the possibility of simulating the percentages missed transfers on a route based on APC data. These missed transfers could be due to late passenger arrivals, mistimed vehicle arrivals at the transfer point, or a transfer time that is too short as scheduled. The impacts of missed transfers on travel time and travel time reliability are explored through a sensitivity analysis. It is concluded that unusually long in vehicle travel times can have a larger effect on traditional reliability measures than missed transfers, potentially hiding the existence of missed transfers on a route.

Freight

Use Case: Using freight-specific data to study travel times and travel time variability across an international border crossing.

Overview. Calculating travel time reliability for freight poses unique data challenges and begs the question: How does travel time reliability for freight transportation systems differ from the question of reliability in the overall surface transportation system? From the research performed in Tasks 2/3 of this project, the team determined that two primary factors differentiate freight systems and the overall surface transportation system: traveler context and trip characteristics.

Traveler context is a primary differentiator between freight trips and all other surface modes: rather than delivering travelers to a destination, a freight trip delivers goods. Because freight drivers are being paid to perform a freight trip, the commercial ecosystem surrounding this concept means that the entire program of scheduling and executing freight trips is much more organized than a typical passenger trip. Thus, freight drivers acquire and utilize travel time reliability information in a fundamentally different manner than other travelers. They also have different concerns. Freight movers were part of the stakeholder interview process conducted by the project team, and these differences have been previously outlined in Tasks 2/3 in this project.

In terms of trip characteristics, freight and overall travel have spatial differences, temporal differences, and facility differences. Spatial differences refer to the fact that origins and destinations with the heaviest freight traffic do not necessarily also have the highest overall traffic volumes. Numerous origin-destination surveys have been employed to identify high-priority freight corridors, and these can be used to focus freight reliability monitoring efforts. In terms of temporal differences, freight traffic generally does not follow the same temporal AM and PM peak pattern of passenger travel. In fact, many freight trips are made during off-peak hours to avoid recurrent congestion. Finally, facility differences refer to the existence in some locations of freight-only lanes or corridors, which would need to be monitored separately from general purpose travel lanes.

Given these differences, the project team decided to take a different approach than that taken for the freeway and transit data, and focus analysis on a very specific freight reliability concern: travel times and reliability across international border crossings.

Data Challenges. This freight use case validation presented a number of data challenges, mostly due to the fact that it is difficult to distinguish freight traffic within an overall traffic

stream using conventional data sources. The project team considered estimating freight traffic volumes from single loop detectors, and then computing freight reliability statistics using the same methodologies employed in the freeway use case validations. However, these estimates, which rely on algorithms that compare lane-by-lane speeds in order to estimate truck traffic percentages, were deemed too unreliable to support accurate travel time variability computations. The team also considered using data from the handful of specialized weigh-in-motion sensors in the region that report vehicle classification data and truck weights, but these were too sparsely located to prove useful for travel time analysis. Because of the unsuitability of traditional traffic monitoring infrastructure for freight reliability calculations, the team's preference was to base analysis on freight-specific data.

There are troves of data on freight vehicle movements, including data on route reliability, available from one stakeholder group: freight movers themselves. Companies such as Qualcomm and Novacom have developed data systems for freight mover operations. They rely on global positioning systems (GPS) outfitted on individual trucks, tracking position and speed, generally on a sub-hour basis. While these data are frequently not fine grained enough to calculate some of the detailed urban reliability information that has been demonstrated elsewhere in this case study, it is adequate for freight movers to understand their travel time reliability environment and to schedule departures appropriately for the just-in-time-delivery windows demanded by their customers.

However, these data are not generally available for studies such as L02, because it is proprietary, competitive information that freight movers gather on their own operations. While these companies have begun to share this data with some partners (such as third party data providers), these deals are struck under terms of strict confidentiality and anonymity. There are some ongoing efforts to leverage this information for public sector agency analysis, such as the border crossing work at Otay Mesa currently underway by the Federal Highway Administration (FHWA) – described in the following section – but these efforts are still in the research phase and are not feasible for public sector agencies to put into operational practice.

In terms of the data required to understand reliability in freight systems, there is strong overlap with freeway and arterial data systems, as freight vehicles are generally part of the overall traffic stream. Because of this, they share the same overall reliability characteristics of the freeway and arterial systems as a whole. However, in many cases, the data required to understand freight movements is scarcer than data needed to understand the overall transportation system, simply because it is data that only pertains to a few percent of overall trips in a given region. The project team was fortunate enough to be given access by the FHWA to freight-specific GPS data collected at the Otay Mesa truck-only border crossing facility from Mexico into the United States. Because of this, this use case validation has a narrow geographic scope, but explores a major issue in freight travel.

Site. The Otay Mesa Crossing has a truck-only facility that, during peak season (which is from October to December), provides access to the US to approximately 2,000 trucks per day. The crossing is equipped to handle trucks that participate in the Free and Secure Trade (FAST) expedited customs processing program, as well as those required to undergo standard processing. US-bound trucks pass through Mexican Export processing prior to entering the US, and are required to be screened at a California Highway Patrol (CHP) commercial vehicle weighing and inspection station before accessing US roadways.

For travel time analysis, the Otay Mesa crossing was broken up into 10 districts, as shown in Exhibit C2-60. These districts are:

1) Export Approach
2) Departure East
3) Departure West
4) Mexico Customs
5) USA Primary
6) USA Secondary Gate
7) USA Secondary
8) Secondary Departure
9) CHP Approach
10) CHP Inspection

Exhibit C2-60: Otay Mesa district map (8)

Data. As part of an FHWA project, data was collected at Otay Mesa from 175 trucks passing through the crossing repeatedly over December 2008 through March 2009. The total number of crossings for GPS-equipped trucks ranged from five percent to twelve percent of the total population of trucks passing through the Otay Mesa crossing. The resulting data set contained 900,000 individual points. A number of these data points were outside of the crossing analysis zone, and thus were discarded prior to analysis. Additionally, almost 30% of trip records contained no travel times, making them unusable for freight reliability analysis. As a result, analysis was performed on the remaining 300,000 individual points, or a third of the total data set.

The Otay Mesa data was used to do two types of reliability analysis: (1) to evaluate the reliability within and across different districts; and (2) to evaluate the reliability associated with different types of inspections. For the district-level analysis, one data complication is that the quantity of reported travel times varies by district. Most individual districts have tens of

thousands of travel time records, as shown in Exhibit C2-61. However, very few trip records (0.07%) contain travel times for all districts. The sparseness of this data makes it challenging to monitor travel times across groups of districts. For example, analyzing the travel time reliability between districts 4 and 7 requires a large set of trips with data points within both districts 4 and 7. Exhibit C2-62 shows the number of trips that spanned multiple districts. Those with zero districts indicate trips where data points were all outside of the geographical analysis range.

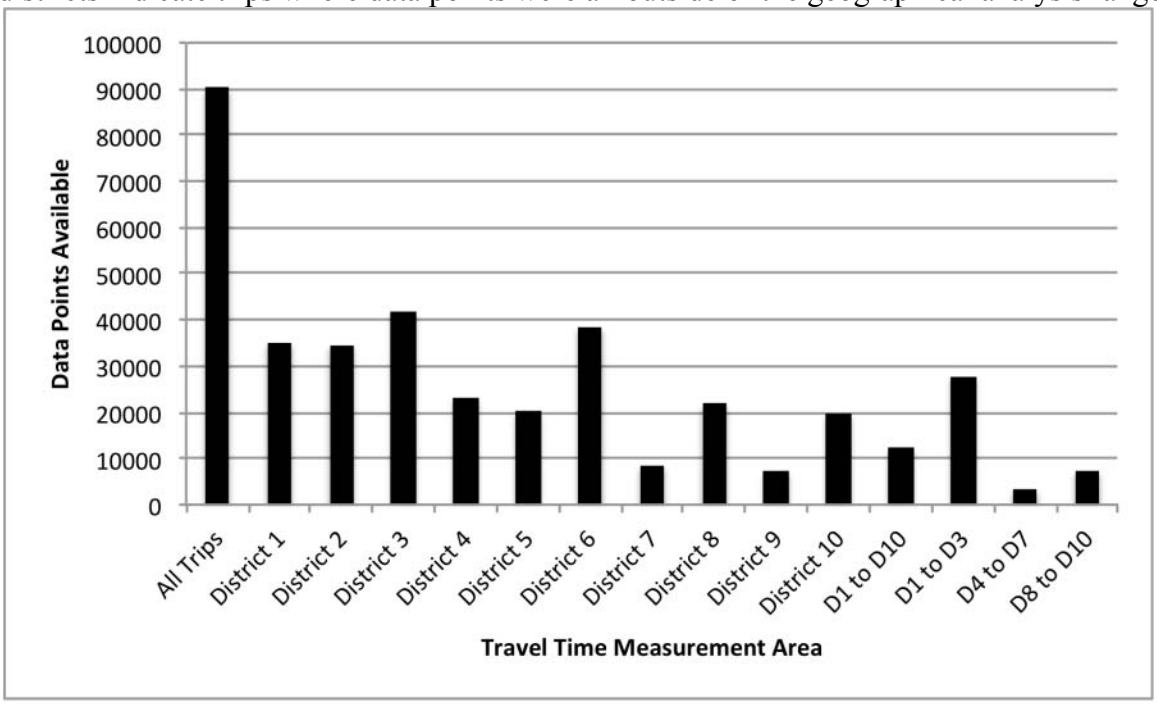

Exhibit C2-61: Otay Mesa GPS points by district

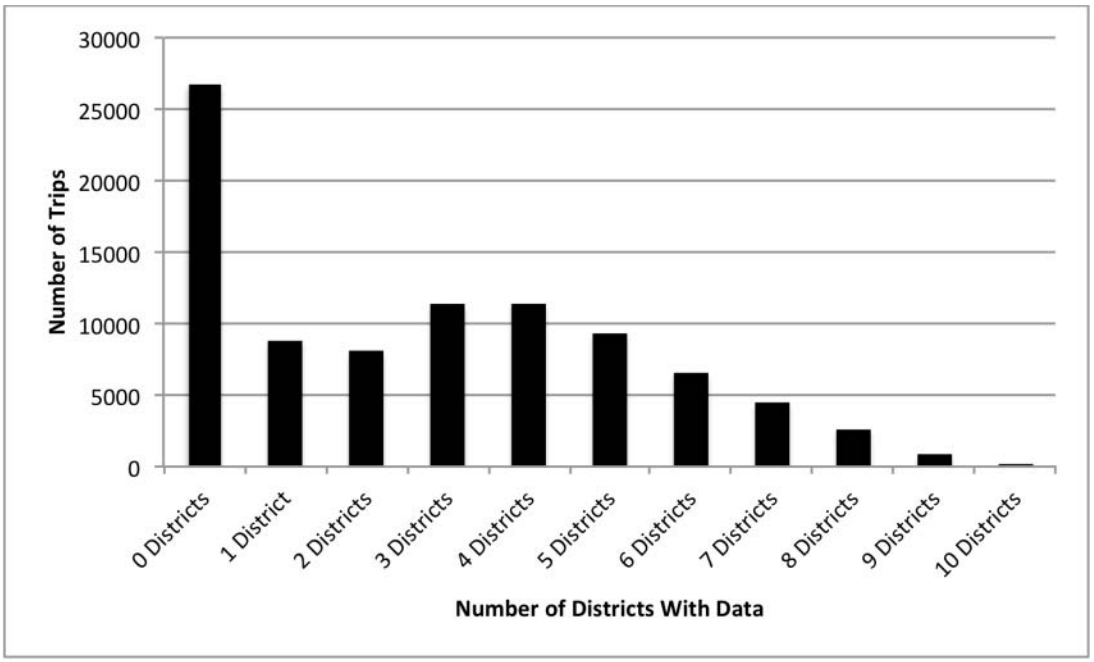

Exhibit C2-62: Otay Mesa trips spanning multiple districts

Results. As outlined in the data section, analysis focused on investigating reliability across Otay Mesa districts and for vehicles subjected to different inspection types. The results of each type of analysis are detailed in the following subsections.

District Reliability. To understand which geographical segments of the border crossing have the most travel time variability, the research team assembled the travel time PDFs for trips within each of the 10 individual districts, and for two trips spanning multiple districts.

The PDFs for districts 1 through 6 are shown in Exhibit C2-63 and the PDFs for districts 7 through 10 are shown in Exhibit C2-64. All of the PDFs are plotted on the same x-axis scale, to facilitate comparison. These data are also summarized into median, standard deviation, and 95th percentile travel times by district in Table C2-29. From the plots, the district that notably stands out as having the most travel time variability is district 7 (USA Secondary Inspection). From the distribution, it appears that the most frequently occurring travel time through district 7 is about 15 minutes, but the trip regularly can take longer than an hour. The median travel time through this district is only 20 minutes, but the 95th percentile travel time is 90 minutes. Districts 1, 2, 3, 4, 8, and 10 also all have 95th percentile travel times at or greater than one hour, which are significantly higher than their median travel times of less than 10 minutes. The district with the most reliability is district 9 (CHP Inspection Approach). Here, the median travel time is only 12 seconds, with a 95th percentile travel time of 2 minutes.

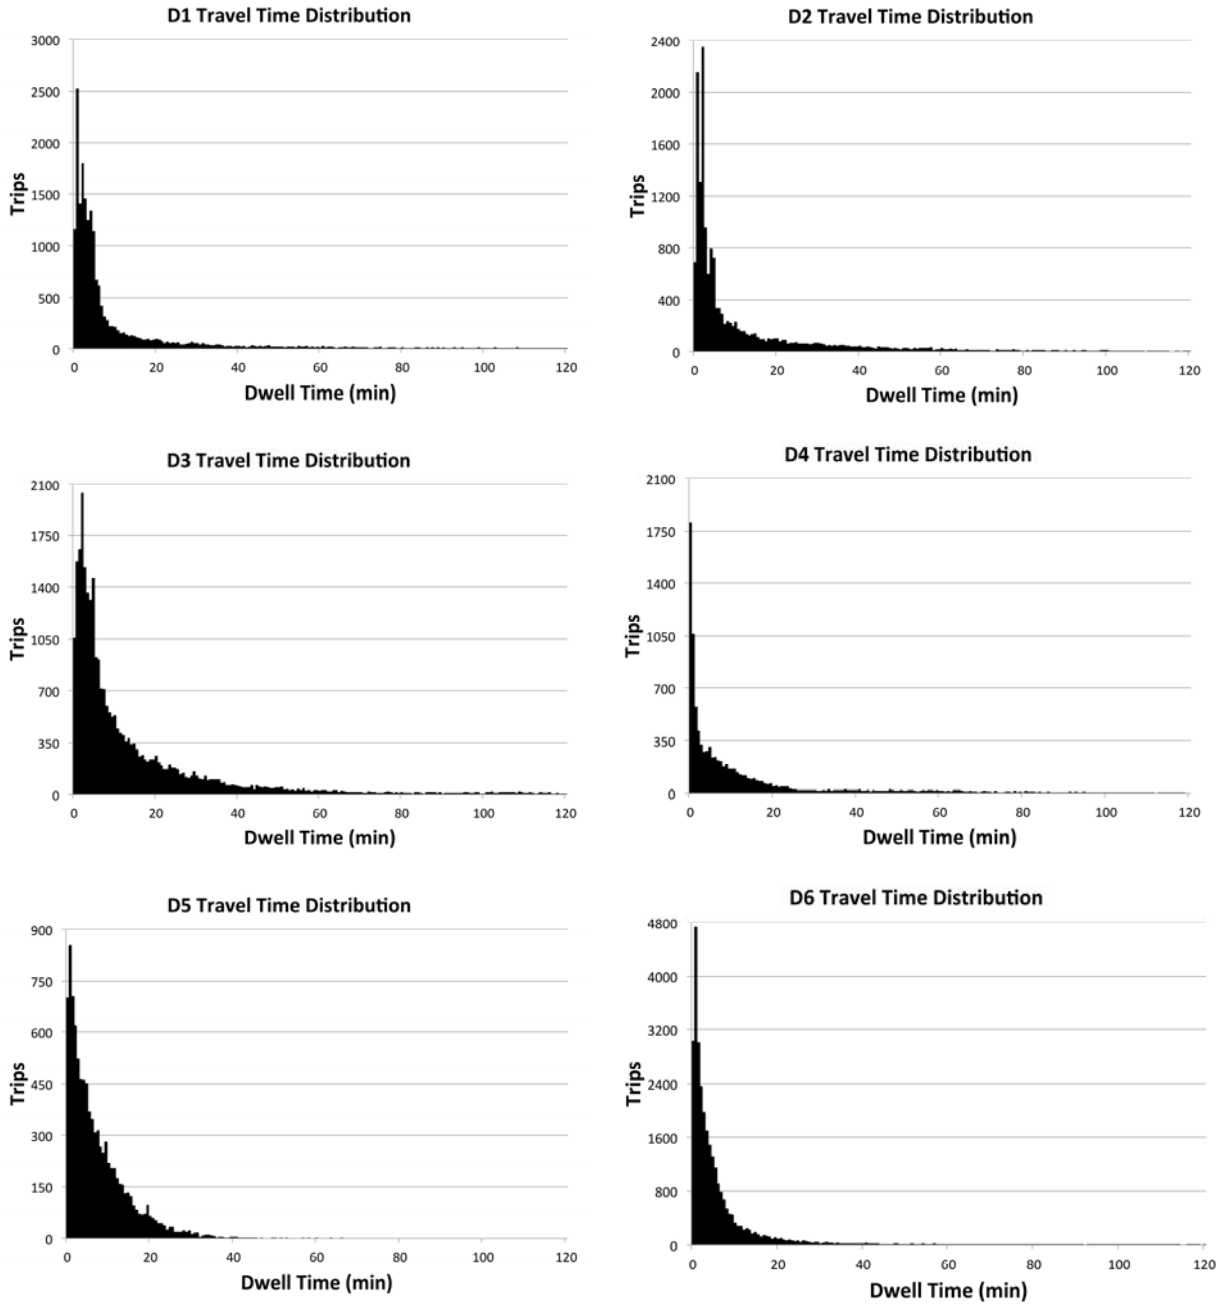

Exhibit C2-63: Districts 1 through 6 Travel Time PDFs

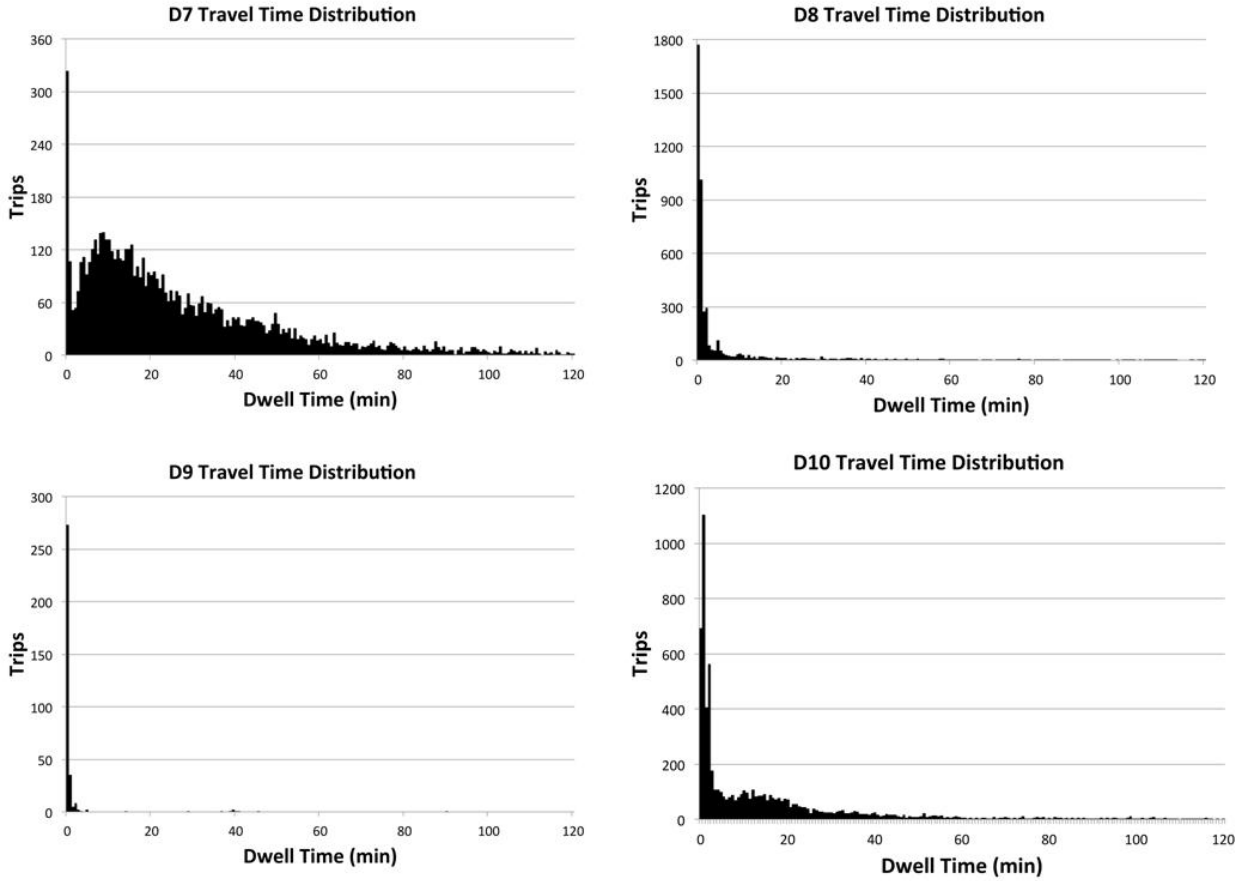

Exhibit C2-64: Districts 7 through 10 travel time PDFs

Table C2-29: District-by-district travel times and variability

District	Median Travel Time (mins)	Standard Deviation (mins)	95th Percentile Travel Time (mins)
D1	4	27	65
D2	4	21	59
D3	7	20	56
D4	5	24	68
D5	5	7	22
D6	3	16	29
D7	21	32	90
D8	1	83	65
D9	0.2	8	2
D10	7	36	87

The research team also looked at the travel times for trucks to get from district 1 to district 6 (the gate to the US Secondary Inspection) and to travel from district 1 to district 10. The PDFs for these two trips are shown in Exhibit C2-65, and the results are summarized into the median, standard deviation, and 95th percentile travel times in Table C2-30.

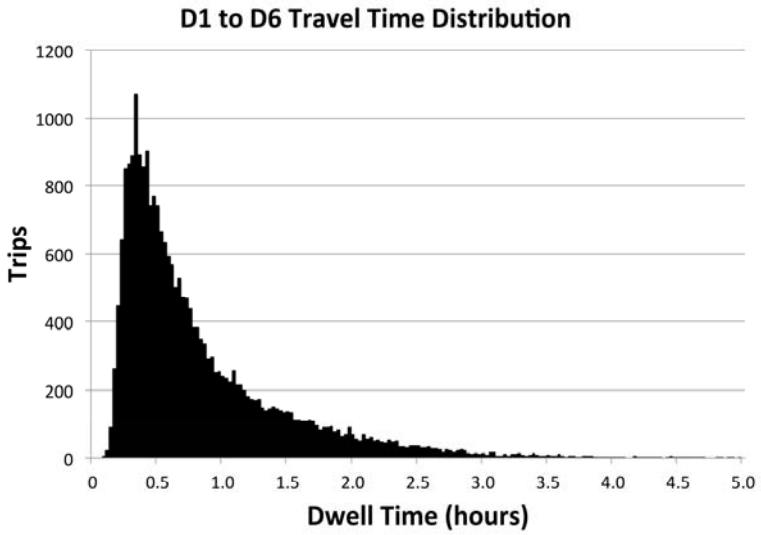

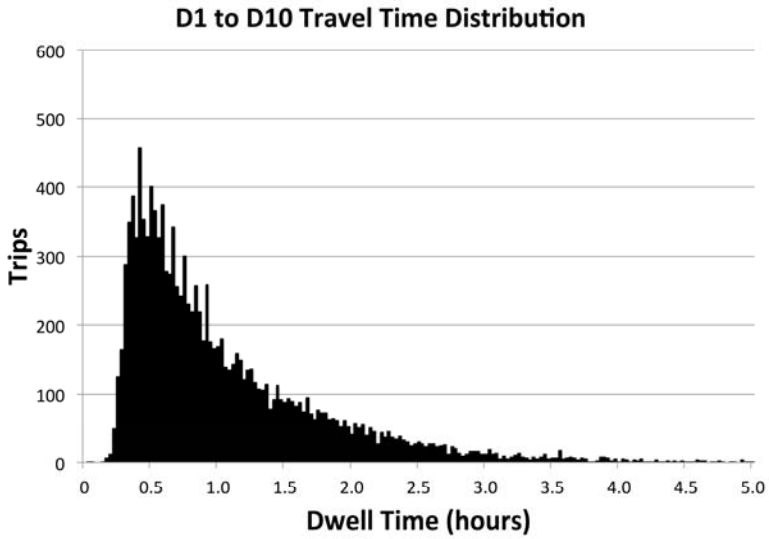

Exhibit C2-65: Cross-district travel times PDFs

Table C2-30: Cross-district travel times and variability

Trip	Median Travel Time (mins)	Standard Deviation (mins)	95th Percentile Travel Time (mins
D1 to D6	37	40	132
D1 to D10	50	48	157

The most commonly occurring travel time between district 1 and district 6 is slightly less than half an hour, though a significant number of trips can take upwards of one or two hours. The median travel time for this trip is 37 minutes, but the 95th percentile travel time is 2 hours and 12 minutes. The median travel time to pass through the Otay Mesa crossing (as represented by the

district 1 to 10 travel time samples) is only 50 minutes, but 5% of trips experience travel times exceeding 2.5 hours.

Checkpoint Reliability. The team also considered the average travel times and travel time variability of trucks passing through certain combinations of checkpoints at different times of the day. As described in the freight data section, many of the freight GPS data records included information on which checkpoints a truck had to pass through while making its trip. While all trucks have to go through certain checkpoints (Mexican Exports, US Inspection, and CHP inspection), some trucks are subjected to additional inspections (Mexico Secondary Inspection and/or US Secondary Inspection). These were used to calculate travel times and reliability for each hour of the day for different checkpoint combinations.

Approximately 15% of trucks that use the crossing qualify for FAST status, which means that, while they have to pass through all the required checkpoints, they can do so in designated FAST lanes *(1)*. Exhibit C2-66 below shows, for all days over which data was received, the total number of sampled FAST lane trucks that traveled during each hour and did not have to stop for any secondary inspections, the average travel time they experienced, and the standard deviation in the travel times they experienced. The data represents over 3,500 records of vehicles that made FAST lane trips. As is evident from the plot, the travel times and travel time variability are actually the highest in the early morning hours, when the fewest sampled trucks were traveling. This may be because drivers are resting or because there is less staff available to perform inspections. The peak number of trucks use the FAST lanes at around noon and between 4:00 PM and 6:00 PM. Average travel times are fairly steady throughout the day, hovering at or slightly above one hour. The standard deviation of the travel times also remains steady at 40 to 50 minutes, meaning that it is fairly frequent for FAST lane border crossings to take almost 2 hours.

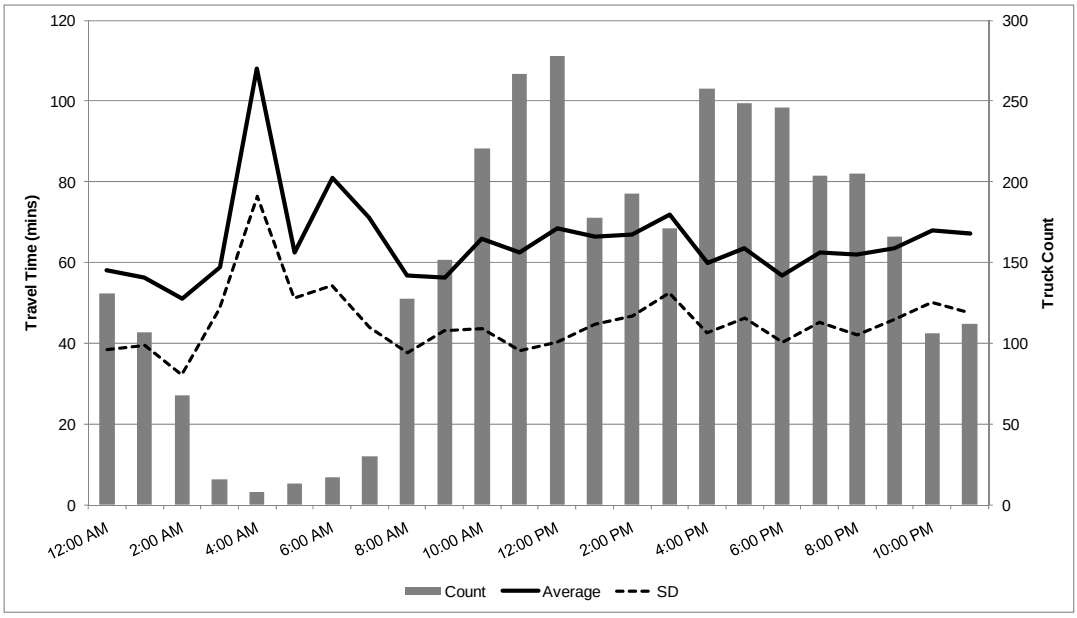

Exhibit C2-66: FAST truck counts, average travel times, and standard deviation travel times by hour

Exhibit C2-67 shows the same plot for 7,400 non-FAST trucks that were selected for US Secondary Inspections. As in the FAST lanes, travel times are the highest during the early morning hours. Throughout the rest of the day, travel times are steady, but are 20 to 30 minutes higher on average than the FAST travel times.

Exhibit C2-67: US Secondary truck counts, average travel times, and standard deviation travel times by hour

Exhibit C2-68 shows the hourly vehicle counts and travel times for FAST trucks who were selected for a US Secondary Inspection. Interestingly, average travel times for FAST vehicles going through a US Secondary Inspection are actually slower (between 90 and 100 minutes) during most hours than they are for non-FAST vehicles (between 80 and 90 minutes) going through a US Secondary Inspection. The standard deviation of travel times for both types of trips are approximately the same.

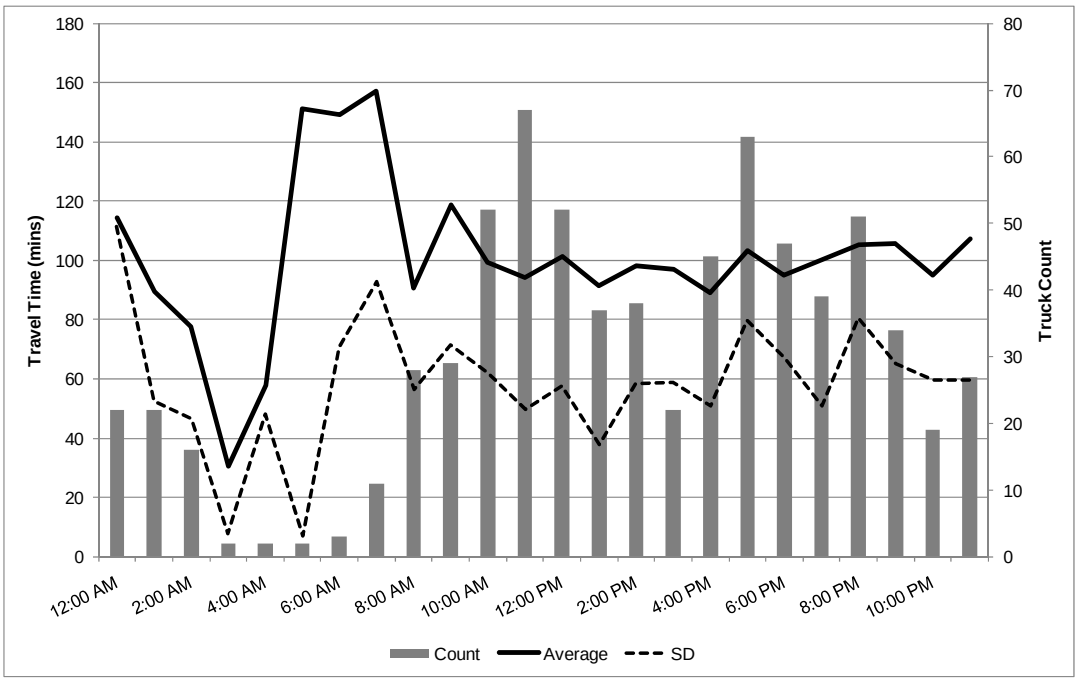

Exhibit C2-68: FAST US Secondary truck counts, average travel times, and standard deviation travel times by hour

Conclusions. This freight use case validation represents an initial use of the Otay Mesa truck travel time data to evaluate travel time reliability for different aspects of a border crossing. The research analyzed and compared travel time reliability across different physical sections of a freight-only border crossing, as well as for different combinations of inspection points passed through by individual trucks. By understanding where the bottlenecks are in the border crossing process and how they are impacting travel times and reliability, managers can begin to take steps to improve operations: for example, adding lanes to capacity-restricted locations or adding staff to checkpoints that are impacting reliability during peak hours of the day.

Extensions of the district-level analysis would group travel times by hour of the day to explain not just where travel time reliability is high, but when it is high as well. Extensions of the checkpoint-based analysis would look at travel time reliability for different days of the week, and for different seasons, because truck border crossings have strong temporal patterns that impact the underlying reliability analysis.

References

1) Kwon, J., B. Coifman, and P. Bickel. Day-to-Day Travel Time Trends and Travel Time Prediction from Loop Detector Data. In *Transportation Research Record, Journal of the Transportation Research Board*, No 1717, Transportation Research Board of the National Academies, Washington, D.C., 2000.
2) Rice, J and E. van Zwet. *A Simple and Effective Method for Predicting Travel Times on Freeways*. In Intelligent Transportations Systems Journal of the IEEE Intelligent Transportation Systems Society, Volume 5, Issue 3. September 2004.

3) Mohring, H.,J. Schroeter, and P. Wiboonchutikula. The Value of Waiting Time, Travel Time, and a Seat on a Bus. *Rand Journal of Economics*, Vol. 18, No. 1, 1987, pp. 40–56.
4) Berkow, M., El-Geneidy, A., Bertini, R., & Crout, D. (2009). Beyond generating transit performance measures: Visualizations and statistical analysis using historical data. Transportation Research Record, (2111), 158- 168.
5) Robert L. Bertini and Ahmed M. El-Geneidy - Modeling Transit Trip Time Using Archived Bus Dispatch System Data (2004)
6) Robert L. Bertini and Ahmed M. El-Geneidy - Generating Transit Performance Measures with Archived Data (2003)
7) Christopher Pangilinan, Wai-Sinn Chan, Angela Moore, Nigel Wilson - Bus supervision deployment strategies and the use of real-time AVL for improved bus service reliability (2007)
8) Delcan Corporation. Measuring Cross-Border Travel Times for Freight: Otay Mesa International Border Crossing Final Report. FHWA-HOP-10-051. September 2010.

LESSONS LEARNED

Overview

During this case study, we focused on fully utilizing a mature reliability monitoring system. We did this to illustrate the state of the art for existing practice. This was possible because of many years of coordinated efforts by transportation agencies in the region, led by the San Diego Association of Governments and Caltrans. These efforts put in a large sensor network, developed the software to process the data from these sensors, and created the institutional processes to utilize this information. Because this technical and institutional infrastructure was already in place, the team focused on generating sophisticated reliability use case analysis. The rich, multi-modal nature of the San Diego data presented numerous opportunities for state of the art reliability monitoring, as well as challenges in implementing guidebook methodologies on real data.

Methodological Advancement

In terms of methodological advancement, the team used data from the Berkeley Highway Laboratory section of Interstate 80. This section is valuable because it has co-located dual loop detectors and Bluetooth sensors. This dataset provided an opportunity for the team to begin to assemble regimes and travel time probability density functions from *individual vehicle travel times.* These travel time PDFs are needed to support motorist and traveler information use cases. Since the majority of the upcoming case study sites will not provide data on individual traveler variability, it was important for the research team to study the connection between individual travel time variability and aggregated travel times, and whether the former can be estimated from the latter. In general, the team found that it was possible to divide the system into specific travel regimes, but has not yet harmonized these two different types of data. Work on answering these questions is ongoing, and will be used to refine the methodologies used at the next four case study sites.

Transit Data

The biggest data challenge in this case study validation was processing the transit data, which is stored in a newly developed performance measurement system. This case study represents the first research effort to use this data and this system. The team found that data quality is a major issue when processing transit data to compute travel times. Many of the records reported by equipped buses had errors, which had to be programmatically filtered out. Errors were due to a variety of reasons. Some buses reported that they were one route, but were serving a completely different set of stops. GPS malfunctions resulted in erroneous locations. Passenger count sensors failed and left holes in the data.

Following the identification and the removal of these data points, assembling route-based reliability statistics using a drastically reduced subset of good data presented the next challenge. This limited the number of routes that the research team could consider, since not all trips on all routes are made by equipped buses, and trips made by equipped buses contain a number of holes due to erroneous data records. From this experience, the research team concluded that transit travel time reliability monitoring requires a robust data processing engine that can programmatically filter data to ensure that archived travel times are accurate. Additionally, transit reliability analysis requires a long timeline of historical data, due to the fact that, typically, a subset of buses is monitored and a large percentage of obtained data points will prove invalid.

Seven Sources Analysis

From a use case standpoint, the research team was challenged to find the best ways to leverage the unique data available in San Diego to demonstrate use cases that might not be possible to explore at other sites. On the freeway side, the research team focused on relating travel time variability with the seven sources, since this dataset was unique to San Diego and the results have high value to planners and operators. In the past, the research team developed a sophisticated statistical model that can estimate the percentage of a route's buffer time attributable to each source of congestion. This model is documented in Chapter 11 of the guidebook. In this case study, the team opted to pursue a less sophisticated but more accessible approach that develops travel time PDFs for each source using a simple data tagging process. This approach was selected because it provides meaningful and actionable results without requiring agency staff to have advanced statistical knowledge.

Conclusions

The San Diego case study validation provided the first opportunity for the team to test guidebook recommendations, implement advanced methodologies, and formally respond to use cases. The research team plans to take the lessons learned during this process to modify the guidebook and better inform the future validation efforts.

CHAPTER C3

NORTHERN VIRGINIA

This case study provides an example of a more traditional transportation data collection network operating in a mixture of urban and suburban environments. Northern Virginia was selected as a case study site because it provided an opportunity to integrate a reliability monitoring system into a pre-existing, extensive data collection network. The focus of this case study was to describe the required steps and considerations for integrating a travel time reliability monitoring system into existing data collection systems.

The purpose of this case study was to:

- Describe the data acquisition and processing steps needed to transfer information between the existing system and the PeMS reliability monitoring system
- Demonstrate methods to ensure data quality of infrastructure-based sensors by comparing probe vehicle travel times using the procedures described in Chapter 3
- Develop multi-state travel time reliability distributions from traffic data

The *monitoring system* section details the reasons for selecting Northern Virginia as a case study and gives an overview of the region. It briefly summarizes agency monitoring practices, discusses the existing sensor network, and describes the software system that the team used to analyze use cases. The section also details the development of travel time reliability software systems, and their relationships with other systems. Specifically, it describes the steps and tasks that the research team completed in order to transfer data from a pre-existing collection system into a travel time reliability monitoring system.

The section on *methodology* describes the implementation of a multi-state travel time reliability model, developed by the SHRP 2 L10 research team, using the Northern Virginia freeway data. It is intended to showcase a tractable method for assembling travel time probability density functions from historical travel time data, as well as highlight the tie-ins of this project with others under the SHRP 2 umbrella. It was selected for emphasis in this case study because the original work was performed using model-generated travel times from the same I-66 corridor being monitored as part of this case study. Work on refining the Bayesian travel time reliability calculation methodology outlined in Chapter 3 and introduced in the San Diego case study will resume as part of the final three case study sites.

Use cases are less theoretical, and more site specific. Their basic structure is derived from the user scenarios described in Supplement D, which are derived from the results of a series of interviews with transportation agency staff regarding agency practice with travel time reliability. Since the focus of this case study is to describe the required steps and considerations for integrating a travel time reliability monitoring system into existing data collection systems, only one use case is described in this case study.

Lessons learned summarizes the lessons learned during this case study, with regard to all aspects of travel time reliability monitoring: sensor systems, software systems, calculation methodology, and use. These lessons learned will be integrated into the final guidebook for practitioners.

MONITORING SYSTEM

Site Overview

The team selected Northern Virginia to provide an example of a more traditional transportation data collection network operating in a mixture of urban and suburban environments. The Northern Virginia (NOVA) District of the Virginia Department of Transportation (VDOT) includes over 4,000 miles of urban, suburban, and rural roadway in Fairfax, Arlington, Loudoun, and Prince William counties. Exhibit C3-1 shows a map of the Northern Virginia District.

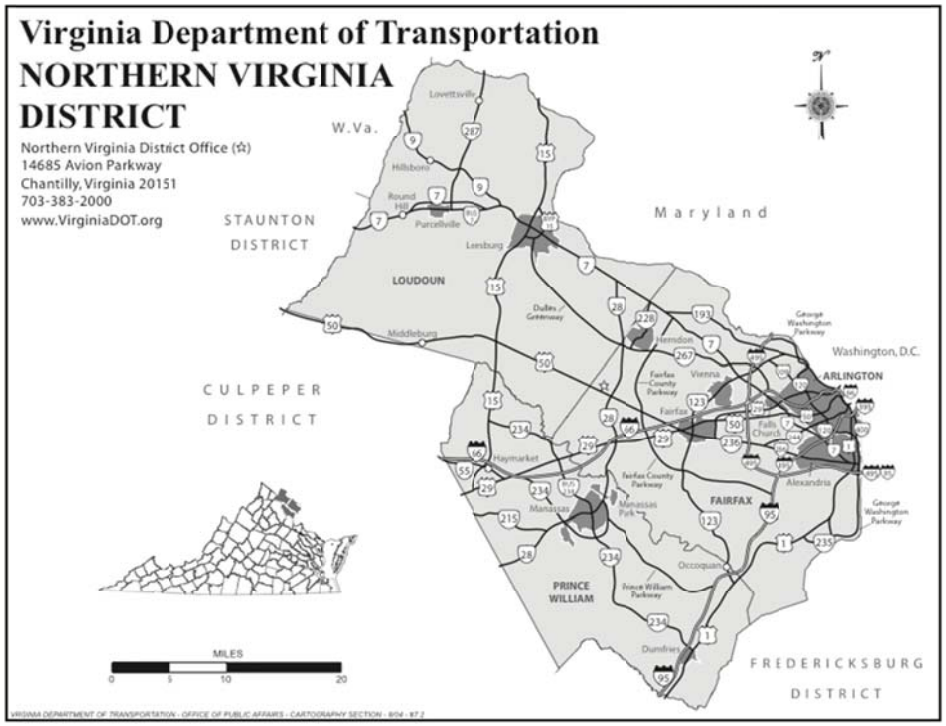

Exhibit C3-1: Map of the NOVA District

Traffic operations in the District are overseen from the NOVA Traffic Operations Center (TOC), which manages more than 100 miles of instrumented roadways, including HOV facilities on Interstates 95/395, 295, 66, and the Dulles Toll Road. To support these activities, the TOC has deployed a wide range of intelligent transportation system (ITS) technologies, including:

- 109 cameras
- 222 dynamic message signs
- 24 gates on I-66 HOV lanes for use during peak travel hours
- 21 gates on I-95/I-395 for reversible HOV lanes
- 25 ramp meters on I-66 and I-395
- 30 lane control signals
- 23 vehicle classification stations

- ~250 traffic sensors (see Exhibit C3-2 for deployment locations)

Overall, the NOVA TOC is a high-tech communications hub that manages some of the nation's busiest roadways. Its systems collect, archive, manage, and distribute data and video generated by these resources for use in transportation administration, policy evaluation, safety, planning, performance monitoring, program assessment, operations, and research applications. Moreover, an Archived Data Management Systems (ADMS) has been developed by the University of Virginia (UVA) Smart Travel Lab (STL) to support VDOT in conducting these activities. TOC staff use dynamic message signs (DMS) and Highway Advisory Radio (HAR) sites to alert commuters about changing traffic conditions. Commuters and other travelers can also tune to AM 1620, call the Highway Helpline at 1-800-367-ROAD (7623) for real-time traffic information, or view the road conditions map on 511 Virginia.

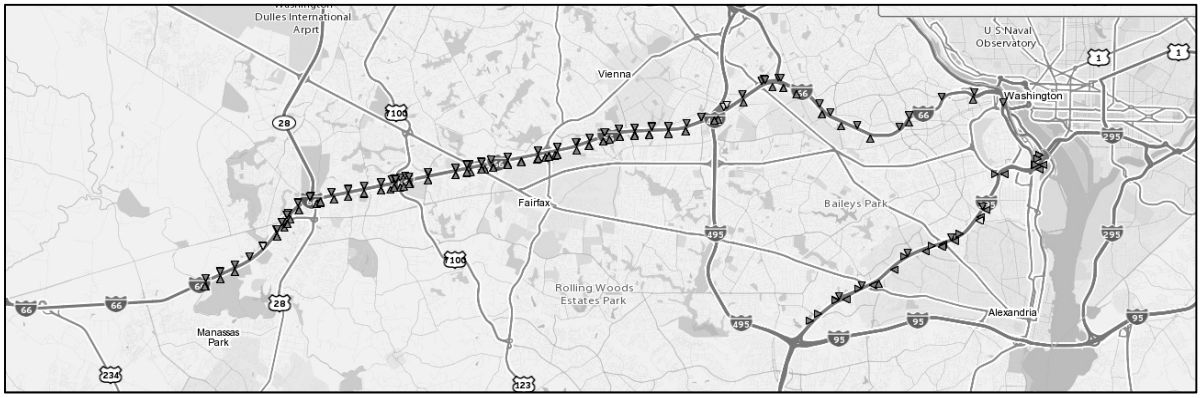

Exhibit C3-2: Locations of Nova District Freeway-based Traffic Sensors

VDOT's management strategy has undergone a dramatic change in the last few years, transitioning from a two-pronged "build-maintain" regime, to a three-pronged "build-operate-maintain" scheme. As such, VDOT is evolving into a customer-driven organization with a focus on outcomes and a "24/7" performance orientation. As part of these efforts, VDOT has developed four "Smart Travel" goals:

1) Enhance public safety
2) Enhance mobility
3) Make the transportation system user-friendly
4) Enable cross-cutting activities to support goals 1-3

These goals are geared toward providing better services to NOVA District customers by improving the quality of their travel and responding promptly to their issues. The focus is on attaining greater operating efficiencies from existing roadway infrastructure as an alternative to building additional capacity. The NOVA Smart Travel Vision is as follows:

"Integrated deployment of Intelligent Transportation Systems will help NOVA optimize its services, supporting a secure multimodal transportation system that improves quality of life and customer satisfaction by ensuring a safer and less congested transportation network."

As part of its activities, the NOVA District has significant interaction with agencies in the District of Columbia and Maryland (in particular in Montgomery and Prince Georges Counties). A number of Federal, state, and local transportation stakeholders, including transit, police, emergency, medical, and other agencies, also play important roles in operating and managing area roadways and other regional transportation systems. Recently, there has been a push within

the region to strive towards increased regional coordination and interoperability. To that end, a regional coordinating entity called CapCOM (Capitol Region Communications and Coordination) has been created to focus on collecting data from a variety of sources to facilitate the creation of a "big picture" of regional traffic.

Due to the major transportation-related construction that began in the region during 2008 and which is anticipated to continue through 2011, mitigation of construction-related congestion is a major focus for the district. Major projects are concurrently occurring, including:

- Construction of 14 miles of HOT lanes on I-495;
- Construction of 56 miles of HOT lanes on I-395/95;
- Widening of I-95 between Newington and Dumfries;
- Widening of I-495, and;
- Roadway improvements at the I-495/Telegraph Road interchange

Sensors

Northern Virginia suffers from severe road congestion, and is generally considered one of the most congested regions in the nation. To help alleviate gridlock, VDOT encourages use of Metrorail, carpooling, slugging, and other forms of mass transportation. Major limited-access highways include Interstates 495 (the Capital Beltway), 95, 395, and 66, the Fairfax County Parkway and Franconia-Springfield Parkway, the George Washington Memorial Parkway, and the Dulles Toll Road. High-occupancy vehicle (HOV) lanes are available for use by commuters and buses on I-66, I-95/395, and the Dulles Toll Road. A portion of the region's HOV lanes have been designed to be reversible, accommodating traffic flow heading north and east in the morning and south and west in the afternoon.

VDOT operates five (5) regional TOCs located in NOVA, Hampton Roads, Richmond, Staunton, and Salem. At the core of each VDOT TOC is an Advanced Transportation Management System (ATMS), which controls each region's field devices and manages information associated with the operation of the roadway network. Operators at each TOC monitor traffic and road conditions on a continuous basis via closed circuit television (CCTV) cameras, vehicle detection infrastructure, and road weather information sensors. In Northern Virginia, VDOT has deployed an extensive network of point-based detectors (primarily inductive loops and radar-based detectors) to facilitate real-time data collection on freeways. Volume, occupancy, and (limited) speed data are collected from these detectors and used by NOVA TOC staff to manage traffic and incidents and provide information to motorists regarding current conditions. The breakdown of NOVA data sources is as follows:

- Multiple types of traffic sensors along I-95, I-495, I-395, and I-66. The mix of sensors deployed along these roadways includes: inductive loop detectors, RTMS radar, magnetometers, SmartSensor digital radar, and SAS-1 sensors.
- Trichord – has deployed acoustic sensors on I-95, I-395, I-495, and I-66.
- Traffic.com – has deployed sensors on I-495, I-395, I-66, and the Dulles Toll Road.

TOC operators also enter incident data, planned events/work zones, and weather events into a web-based application called the Virginia Traffic Information Management System (VaTraffic). VaTraffic information is shared with the public, VDOT management and other key state and local emergency response agencies.

Although a number of major Interstate roadways pass through the NOVA region, including I-95, I-495, I-395, and I-66, for the purposes of this study we conducted analyses

exclusively on I-395 and I-66, the two primary entry/egress interstates southwest of Washington, D.C.

On I-66 and I-395, point detectors are placed at approximately ½ mile intervals. Due to accuracy and maintainability issues with inductive loop detectors and other older sensors, there are no plans to replace failed units which have been deployed on the mainline lanes. Instead, plans are in motion to transition to the use of non-intrusive radar-based detection technologies. These sensors are being deployed both as replacements for older failed units, as well as at all locations where detection infrastructure is being deployed for the first time. As a result of a combination of older loop detector station failures, ongoing roadway construction, and the need to configure many of the newer radar-based units, data is currently available for only about 75 of the detectors. Exhibit C3-3 provides a visual indication of the availability of data on I-66 and I-395; lighter colored icons indicate working stations, darker icons indicate non-working stations.

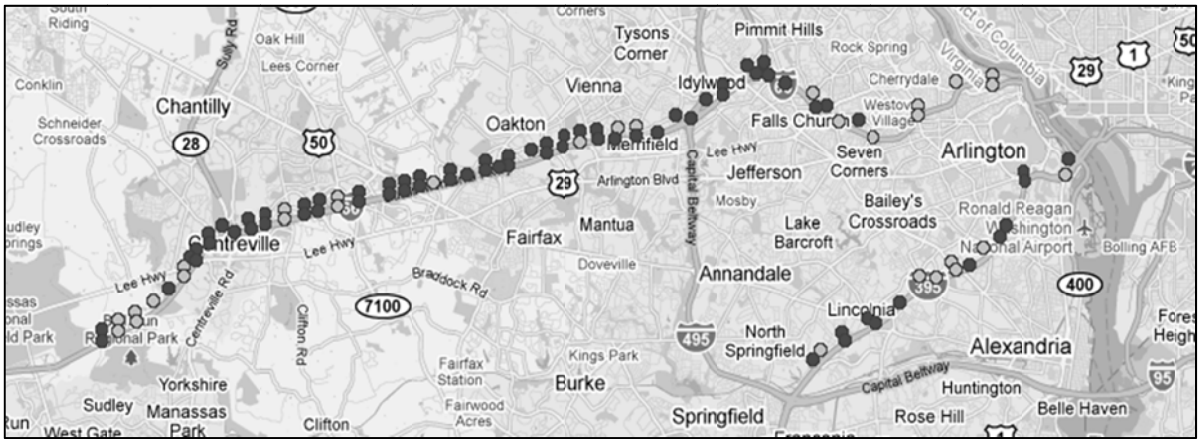

Exhibit C3-3: Map of Working vs. Non-Working Sensor Stations

Data Management

NOVA TOC staff use a regional Freeway Management System (FMS) to monitor and manage traffic data from the ATMS, respond to incidents, and disseminate traveler information. The FMS is linked to the Virginia Traffic Information Management System (VaTraffic), a statewide traffic information management and conditions reporting system developed by VDOT to provide an efficient, integrated platform for managing activities that affect the quality of travel experienced by motorists. It comprises a suite of applications that VDOT staff use to manage planned events such as roadway maintenance, unplanned events such as traffic accidents and heavy congestion, and to provide information for use by other VDOT systems. These data are made available via a Data Gateway.

The Data Gateway was first deployed in VDOT in 2004 as an interconnection between the Virginia State Police (VSP) and the Richmond Traffic TOC. Since that time, it has grown into a statewide network that is used to exchange critical information. The Data Gateway is an XML Publish and Subscribe network fully compliant with the Emergency Data Exchange Language (EDXL) standard, providing the maximum degree of interoperability between systems. The Data Gateway presently allows a number of diverse systems to share data, including:

- VaTraffic - uses the Data Gateway to exchange information with nearly 1500 statewide users, the 511 Interactive Voice Response (IVR) and Web applications, and

other VDOT systems. VaTraffic publishes information for incidents, planned events, road conditions, snow conditions, and bridge schedules.

- OpenTMS - is deployed in the Northern, Central, Northwest, and Southwest TOCs, and publishes information concerning incidents and DMS messages. In the future, OpenTMS is planned to provide information on weather sensors, work zones, HOV gate control, and other lane control data.
- Virginia State Police - the Data Gateway has been used to share VSP data since 2004. Data entered in to the VSP CAD system is shared in real-time with all participating TOCs.

VDOT currently reports on roadway conditions via a number of performance-related products, including its Quarterly Report, Web-based Performance Dashboard, and bi-monthly performance reports to the VDOT Commissioner (internal).

- The VDOT Performance Dashboard (http://dashboard.virginiadot.org/) provides a wide range of transportation performance-related data, including:
- Travel Times on Key Commuter Routes
- Congestion along Interstates
- HOV Travel Speeds
- Incident Duration
- Annual Hours of Delay

Performance measurement has become an important function within VDOT and serves to enable TOC engineers and operators to identify, measure, and report the status of the both the freeway system and individual facilities at different geographic (spatial) and temporal scales.

System Integration

Overview

For purposes of this case study, data from NOVA's data collection network and system were integrated into a developed archived data user service and travel time reliability monitoring system. The steps and challenges encountered in enabling the information and data exchange between these two large and complex systems are described in detail in this section. The goal of this section is to provide agencies with a real-world example of the resources needed to accomplish data collection to monitoring system integration, and the likely challenges that will be encountered when procuring a monitoring system.

This section first describes the source system (VDOT's data collection system) and the reliability monitoring system (PeMS). It then describes the data acquisition and processing steps need to transfer information between the two systems. Finally, it summarizes findings and lessons learned.

Source System

VDOT's Northern Region Operations site receives detector data from two different systems; one that collects data along part of I-66, and one that collects data for the rest of I-66 and I-395. These two data streams are integrated into a standardized format in a single text file that is generated every minute. This text file is passed in real-time to the Regional Integrated Transportation Information System (RITIS), developed and maintained by the CATT Laboratory

at the University of Maryland (UMD). RITIS, without doing any further processing of the data, parses the text file and puts it into an XML document that is updated every minute on a page of the RITIS web site. Access to this webpage is limited to pre-approved IP addresses. These real-time detector data XML documents were the primary traffic data source for NOVA PeMS. When data quality, largely due to recent construction on monitored roadways, proved to be a major issue impeding the study of reliability on the 2011 data, the research team also acquired a database dump of detector data along I-66 and I-395 for the entire year of 2009 from the UMD CATT lab.

Reliability Monitoring System

PeMS is a traffic data collection, processing, and analysis tool that extracts information from real-time intelligent transportation systems (ITS) data, saves it permanently in a data warehouse, and presents it in various forms to users via the web. PeMS can calculate many different performance measures; and as such, the requirements for linking PeMS with an existing system depend on the features being used. Since the function of PeMS in this case study is to collect traffic data from point detectors, quality control it, generate and store travel times, and report reliability statistics, the following describes what PeMS uses from the source system to support these functions:

- Metadata on the roadway linework of facilities being monitored
- Metadata on the detection infrastructure, including the types of data collected and the locations of equipment
- Real-time traffic data in a constant format at a constant frequency (such as every 30-seconds or every minute)

The foundation of PeMS is the traffic detector, which reports at least two of the three fundamental parameters that describe traffic on a roadway: flow, occupancy, and speed. Detectors report or are polled for data in real-time at a pre-defined time interval. In PeMS, detectors have a location denoted by a freeway number, direction of travel, latitude and longitude, and a milepost that marks the distance of a detector down a freeway. Each detector is assigned a unique ID which remains with it throughout time, and can never be assigned to another detector, even if the original detector is removed. Every detector belongs to a station, which is a logical grouping of detectors that monitor the same type of lane (for example, mainline versus HOV) along the same direction of freeway at the same location. Each station has a unique ID, a type (such as mainline, HOV, ramp, etc.), a number of lanes, and a corresponding set of detectors. The final pieces of equipment in the PeMS framework are controllers, which are located along the roadside and collect data from one or more stations. They have a latitude/longitude and mile marker location, as well as a set of corresponding stations. This hierarchy- a controller collecting data from stations composed of detectors- gives structure to the roadway instrumentation configuration, making it easy to spatially aggregate data and diagnose problems in the data collection chain, such as a broken detector or controller, or a failed communication line.

PeMS collects detector data- either by directly polling each detector or obtaining it from an existing data collection system- in real-time and stores it in an Oracle database. The raw data is permanently stored in a raw database table, and is also aggregated up to the five-minute level, at which point PeMS computes the average five-minute speed for detectors that transmit flow and occupancy, and the average five-minute occupancy for detectors that transmit flow and speed. This data is stored in a five-minute detector database table. At the five-minute level,

PeMS also aggregates the lane-by-lane detector data up to the station level, which represents the total flow, average occupancy, and average speed across all the lanes at that location during that five-minute period. This data is stored in a five-minute station database table. The station data is further aggregated up to the hourly and daily levels, and stored in corresponding database tables.

PeMS computes travel times on routes, which can traverse more than one freeway, and which are defined by a starting on-ramp, freeway-to-freeway connectors (if any), and an ending off-ramp. It computes travel times for routes at the five-minute and hourly levels from the data in the detector and station database tables, using the infrastructure-based sensor calculation method described in Chapter 11 of the Guidebook. It stores these travel times permanently in five-minute and hourly travel time database tables. These travel times can then be queried to assemble the historical distribution of travel times along a route for different times of the day and days of the week, as well as compute reliability metrics such as the buffer time index and percentile travel times.

Data Acquisition

This section describes, in general, the transfer of data between the source system and the monitoring system in order to monitor travel time reliability. It also details the specific data exchanges occurring between the source system and PeMS in this case study.

General. Typically, reliability monitoring systems must acquire two categories of information from the source system in order to produce accurate performance metrics: (1) metadata on the roadway network and detection infrastructure; and (2) traffic data. The traffic data are unusable for travel time calculation purposes if not accompanied by a detailed description of the configuration of the system. Configuration information provides the contextual and spatial information on the sensor network needed to make sense of the real-time data. Ideally, these two types of information should be transmitted separately (i.e., not in the same file or data feed). Roadway and equipment configuration information is more static than traffic data, as it only needs to be updated with changes to the roadway or the detection infrastructure. Keeping the reporting structure for these two types of information separate reduces the size of the traffic data files, allowing for faster data processing, better readability, and lower bandwidth cost for external parties who may be accessing the data through a feed.

Additionally, the data acquisition step often involves reconciliation between the framework of the source system and the monitoring system. For example, different terminology can lead to incorrect interpretations of the data. As such, this step often requires significant communication between the system contractor and the agency staff who have familiarity with the data collection system, in order to resolve open questions and make sure that accurate assumptions are being made.

Metadata. PeMS needs to acquire two types of metadata before traffic data can be stored in the database: roadway network information and equipment configuration data. To represent the monitored roadway network and draw it on maps, PeMS needs to have GIS-type roadway polylines defined by latitudes and longitudes. To help the agency link PeMS data and performance metrics with their own linear referencing system, PeMS also associates these polylines with state roadway mileposts. In most state agencies, mileposts are a reference system used to track highway mileage and denote the locations of landmarks. Typically, these mileposts reset at county boundaries. In cases where freeway alignments have changed over time, it is likely that the difference between two milepost markers no longer represents the true physical distance down the roadway. For this reason, PeMS adds in a third representation of the roadway

network, called an absolute postmile. These are akin to mileposts, but they represent the true linear distance down a roadway, as computed from the polylines. They do not reset at county boundaries, in order to facilitate the computation of performance metrics across long sections of freeway. In PeMS, this information is ultimately stored in a freeway configuration database table that contains a record for every 10^{th} of a mile on every freeway. Each record contains the freeway number, direction of travel, latitude and longitude, state milepost, and absolute postmile.

The research team was not able to obtain any GIS data for the NOVA network within the project time frame. Since the monitored network consisted of only two corridors, roadway linework was obtained by entering the starting and ending points of each corridor into Google Maps and exporting the results into a KML file. From these data, polylines and their latitudes and longitudes were parsed and placed in a PeMS database. The next step was to add state milepost markers to these latitude/longitude freeway locations. Since both of the monitored freeway segments fell into only one county, this was done by researching the mileposts at the county boundary, and then interpolating the mileposts in at least 0.10 mile increments along the rest of the freeway segment. In the NOVA case, state mileposts and PeMS absolute postmiles are the same.

The second type of metadata required is information about the detection equipment from which the source system is collecting data. PeMS has a very strict equipment configuration framework which is described in the Reliability Monitoring System subsection. All source information must conform. The rigidity of this framework is due to the need to standardize data collection and processing across all agencies, regardless of their source system structures. Configuration information ultimately populates detector, station, and controller configuration database tables in PeMS, and is used to correctly aggregate data and run equipment diagnostic algorithms.

NOVA equipment configuration information was obtained from an XML file posted on the RITIS website that is updated periodically (typically, not more than every few days). A representative section of this file is shown in Exhibit C3-4

Exhibit C3-4. The file is composed of <detector> elements, which each have a unique ID, a textual name that includes a mile marker, a latitude and longitude, a type (such as inductive loop), and one of more <detection-zone> elements. Each <detection-zone> element has a unique ID, a number of lanes, a latitude and longitude, a direction, and, sometimes, a type (such as shoulder or lane).

Exhibit C3-4: NOVA RITIS Detector Configuration XML Format

Once the file was obtained, the next step was to fit the data into the PeMS configuration framework. The third step was to parse the XML file, insert relevant fields into the PeMS database, and write a program to automatically download the configuration file from the RITIS website and populate relevant information into the database whenever the file is updates. Since the XML file was not accompanied by an explanatory text file, the second step took considerable time and effort, as a number of issues were uncovered that made it challenging to map the NOVA information into the PeMS database. The issues, described below, related to conflicting

terminologies, information required by PeMS that was missing from the configuration file, and equipment types not supported by PeMS.

The first challenge was to determine how the NOVA <detector> and <detection-zone> elements should map to the PeMS equipment framework of detectors, stations, and controllers. From the properties of the NOVA <detectors>, it was clear that they did not refer to the same entity as a PeMS detector. NOVA detectors contain multiple zones, and each zone has a lane count, a location, a direction, and a type. From these attributes, it was concluded that the NOVA detection zone was conceptual equivalent to the PeMS station, and that the NOVA detector was the conceptual equivalent of the PeMS controller. This was confirmed by looking at samples of the RITIS traffic data XML files, which report flow, occupancy, and speed data for each <detection-zone>. After performing this matching and reviewing the traffic data, the team concluded that, despite the terminology used, the NOVA configuration information had no notion of a detector in the PeMS or the conventional sense, i.e., a sensor that monitors traffic in a single lane at a single location. Since PeMS is built around the collection of lane-specific data from detectors, which enables the capability to report lane-by-lane flows, volumes, and occupancies at point locations and lane-by-lane travel times along routes, this presented a challenge. The problem was ultimately solved by using the number of lanes reported for each NOVA detection zone to assign artificially constructed PeMS detectors to monitor each lane. Each detector was given an ID, assigned by appending to the detection zone ID an integer representing the lane number. Then, during the real-time data integration, the flows, volumes, and occupancies reported by each detection zone were divided by the number of lanes and assigned to each detector.

Another challenge was matching the NOVA detection zone types with the station types supported by PeMS. Every station in PeMS is assigned a type to denote the lane type that it monitors. Station types must be one of the following: mainline, HOV, collector/distributor, freeway-freeway connector, off-ramp, or on-ramp. In the NOVA configuration XML file, not every detection zone is assigned a type, and the types that are assigned (shoulder, lane, exit ramp, rhov, and hov) do not align with those defined in PeMS. The NOVA “shoulder” zone type is a reflection of the fact that, during peak hours, the shoulder lanes on I-66 are open to traffic. The “rhov” zone type is assigned to HOV lanes that are reversible based on the time of day. These two operational characteristics added significant complexity into the monitoring process. The operation of shoulder lanes meant that the number of lanes at a given location changed by time of day, a characteristic that PeMS could not accurately represent. Similarly, the reversible HOV lane operation meant that sensors monitored different directions of travel based on the time of day, which PeMS also could not accurately configure. For this reason, “shoulder” and “rhov” stations were not stored in the PeMS database. A related problem was that many detection zones were not assigned types in the configuration file. To solve this, the latitude and longitude of each NOVA detection zone was mapped in Google Earth and manually inspected to determine which PeMS station category it belonged to. The end product of this step was a csv file that listed each detection zone ID and its corresponding PeMS station type.

A third issue was that, through the metadata, PeMS needed to learn what types of data it would be receiving from each station. Typically, detectors can report up to three values: flow, occupancy, and speed. Some detectors, such as on- and off-ramp loop detectors, only report flows. Single inductive loop detectors report flow and occupancy. Radar detectors report flow and speed. Double loop detectors report flow, occupancy, and speed. PeMS needs to know which detectors report which values, so that, for detectors reporting two of the three values, the third is

calculated via an algorithm. This information is not directly present in the NOVA configuration XML file. NOVA detectors (PeMS controllers) are assigned types (either inductive loop or microwave radar) in the configuration file. Since VDOT staff confirmed that the inductive loops are single loop detectors, we expected that zones made up of inductive loops would report flow and occupancy, and zones made up of microwave radar sensors would report flow and speed. However, in the traffic data XML file, all zones, regardless of their detector type in the configuration file, reported all three values or only flow. The implications of this finding are further described in the Traffic Data section that follows. From a metadata perspective, there was no sure way of tagging NOVA zones with the types of data expected to be received. For this reason, PeMS ultimately stored whatever values each zone transmitted via the XML file. This meant that, for detectors reporting only flow, their speeds and occupancies were entered as zero, even though this clearly did not reflect the actual field conditions.

The metadata quality control steps described above were the bulk of the work to insert NOVA configuration information into PeMS. Following this, a custom program was written to parse the PeMS-required fields from the XML configuration file, supplement them with the zone type information in the csv file, throw away metadata for elements that PeMS could not support, and insert information into the required database tables in PeMS. Ultimately, PeMS consumed configuration information for a total of 260 mainline zones and 69 HOV zones, which became the equivalent of PeMS mainline and HOV stations, respectively.

Traffic Data. Following the metadata acquisition, the next step was to acquire traffic data and archiving it. Real-time traffic data was acquired via an XML file posted every minute onto the RITIS web page, in the same location as the configuration XML file. The end goal of the traffic data acquisition process was to take one-minute traffic data from the XML file and insert it into the appropriate tables in the PeMS database. Before this could be done, the research team had to develop a full and accurate understanding of the NOVA real-time data. Because the generation of accurate reliability information requires a large set of historical travel times, the team wanted to minimize the delay in acquiring traffic data. For this reason, as soon as the metadata were inserted into PeMS, the team implemented a program to download the traffic data XML file from the RITIS website every minute and save it, so that data could be parsed from the files and placed into the PeMS database as soon as the file format was thoroughly understood.

A sample of the real-time traffic data XML file is shown in Exhibit C3-5. It is composed of <collection-period-item> elements each defined by a timestamp and a 60 second measurement duration. This element contains the most recent measurements for each NOVA <detector> that most recently sent data during that timestamp. Working controllers are reported in the <collection-period-item> element marked by the most recent timestamp. If a controller is not currently transmitting data, its most recent data transmission is included in a <collection-period-item> marked by the timestamp for which the system last received data from it. Each <collection-period-item> element contains a <zone-report> element for each controller that last reported data during that timestamp. Each <zone-report> element then contains a <zone-data-item> with each zone's most recent flow, occupancy, and speed values. For many zones, the flows are non-zero while the occupancies and speeds are zero. For others, all three values are non-zero.

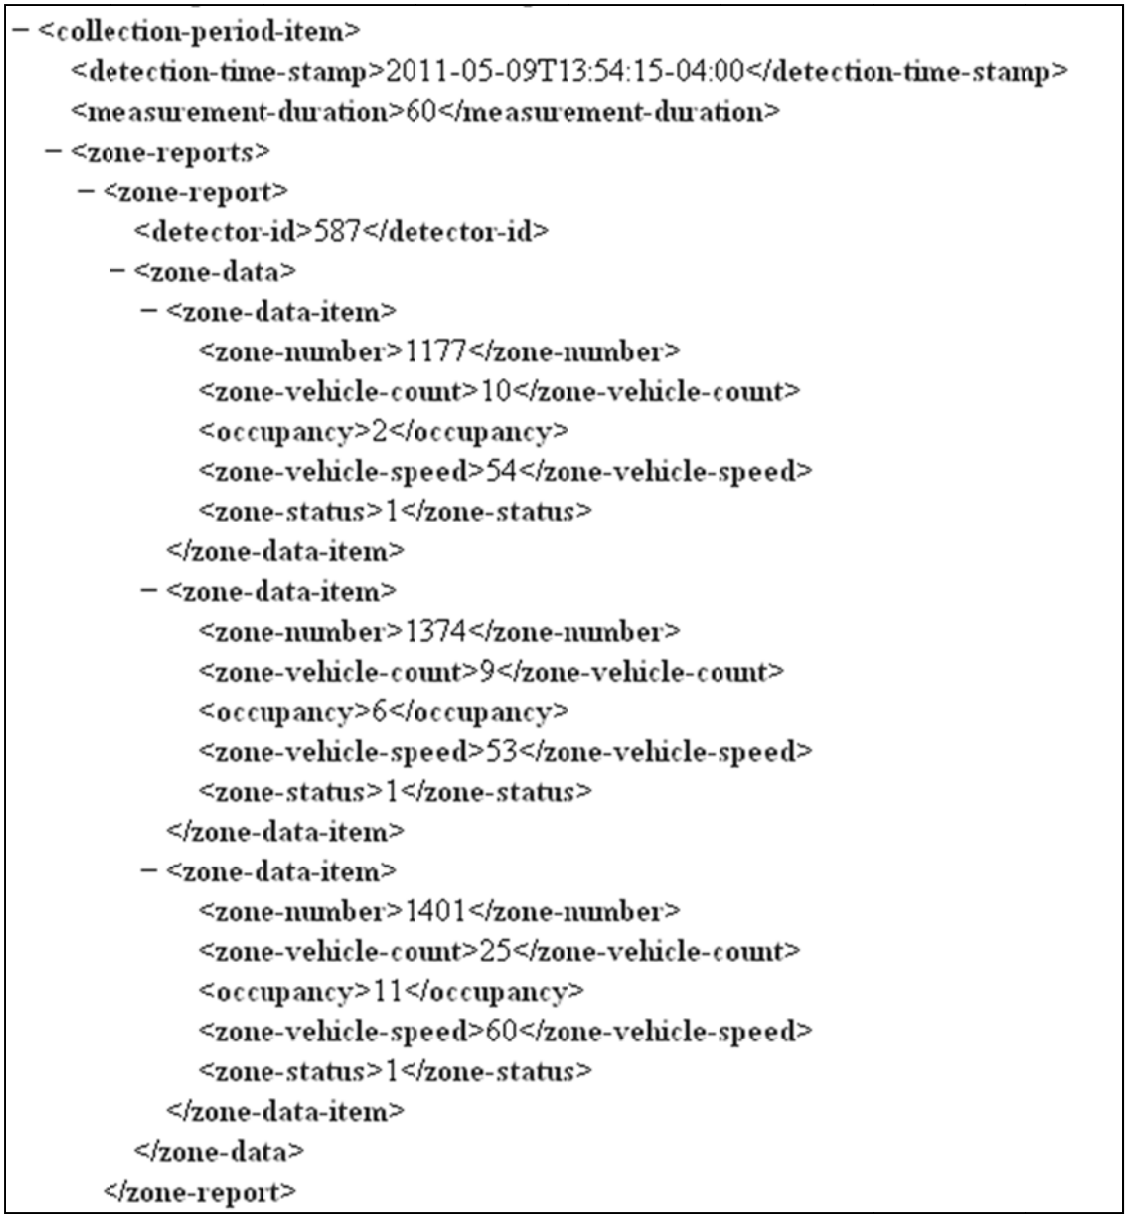

Exhibit C3-5: NOVA RITIS XML Traffic Data Format

Review of the data led to a number of questions. We first wanted to know what processing was done on the data to generate the values in the XML files. This relates to a fundamental issue that agencies collecting data should consider. Many agencies encounter external parties that have an interest in obtaining a traffic data feed generated from public-sector detection infrastructure. The level of interest in raw versus processed data differs depending on the intended use. Maintaining even one data feed can be a challenge; maintaining multiple data feeds is likely to be infeasible for many agencies. As such, if agencies want to provide a feed of processed data, all steps should be documented in as much detail as possible to inform data users on what is being reported and how values are being generated. Optimally, a mature reliability monitoring system would collect raw data, then apply quality controls, and aggregate and report

it using robust methods. This would ensure uniformity in the data at the lowest temporal and spatial level possible while accurately evaluating and reporting the quality of data.

We concluded that the NOVA data was heavily pre-processed before being placed in the data feed. Firstly, the XML file contains no lane-by-lane data, despite the fact that a number of the NOVA "detector types" are single inductive loop detectors, which monitor individual lanes. This means that, at some point, the source system aggregates values from individual lanes into total flow and average occupancy and speed across all lanes at a given location. Since the foundation of PeMS is lane-by-lane data, this issue was addressed by dividing the flows by the number of lanes and assuming that the reported average speeds and occupancies applied to all lanes. While this allowed the data to be transformed into the PeMS framework, it showed that a loss of information can occur when an agency pre-processes the data. In this case, the reliability monitoring system no longer has the ability to report on the differences in travel times along different lanes on the same route. Another sign that the NOVA data was preprocessed lay in the fact that many zones reported flow, occupancy, and speed. Since the corridor detectors were all single loop detectors or microwave radar detectors, they only directly transmitted two of the three values. In these cases, PeMS would normally calculate the third value from the known two using a lane, location, and time-specific assumption about the average vehicle length, called a g-factor *(1)*. When receiving all three values, PeMS does not have to perform this calculation step, but this comes at the expense of not knowing how the third value was computed. The team contacted UMD and VDOT staff to determine what is being done, but both organizations stated that they do no processing on the data ultimately posted in the RITIS XML file. From this, it was concluded that the data collection system in the field is doing the pre-processing, but we were not able to ascertain exactly what was being done. Without being able to evaluate the methods, we decided to have PeMS store whatever data it received via the XML files. In all cases, whether PeMS received all three values, or whether it received a non-zero flow and zero occupancy and speed, it stored these values in the database.

The second issue that had to be addressed was determining the units of the occupancy values being reported. Typically, occupancy is reported as the percentage of the reporting period that a vehicle was directly above the detector. Reasonable values range from 0 to 15. When reviewing the traffic data XML file, we noted that many of the occupancy values were high, with some even consistently exceeding 100. We surmised that perhaps occupancy was being reported in tenths of a percent. The issue was ultimately discussed with VDOT staff, who confirmed that the units of occupancy are whole percentage points, and that zones reporting high occupancies are broken, largely due to construction projects in the vicinity.

The third issue related to the discrepancy between the NOVA data being reported at the "zone", or "station" level, and the PeMS requirement for lane-by-lane data. From a metadata perspective, as previously described, this was resolved by assigning detectors to each zone within PeMS. For the real-time data, the team decided to simply divide the zone flows by the number of lanes at the zone and assign them to each lane. To keep flows as whole numbers, any remainders following the division were assigned starting at lane 1 (the left-most lane), resulting in an overall upward bias of vehicle counts in the left-hand lanes. For occupancy and speed, the team assigned the values reported by the zone to each of its detectors.

By the time the above-described issues had been resolved, we had been downloading and saving the one-minute traffic data XML files from the RITIS website for three weeks. The next step was to write a program to parse out the zone values, assign them to the PeMS detectors, and store them in the PeMS database according to the <detection-time-stamp> element. The team did

this for the three weeks of archived traffic data XML files, and also developed a program to download the XML files every minute from the RITIS website and perform the same steps to place data in the database in real-time.

Data Processing

The data acquisition phase resolved all discrepancies between the NOVA framework and PeMS framework, and successfully mapped over all of the relevant fields in the XML files to the PeMS database. It also resulted in an automated, real-time acquisition chain between the RITIS web page and PeMS, with PeMS obtaining data from the web page every minute and inserting it into the PeMS database. From this point forward, PeMS could perform its standard data processing to assess the health of NOVA detection infrastructure, throw out bad data and impute values, aggregate data across lanes and over time, and calculate performance metrics such as travel times.

In its detector health assessment step, PeMS looks at the data transmitted by each detector over a single day and makes a determination as to whether the data is good or problematic. PeMS makes this assessment based on the flow and occupancy values for each detector. There are a few common problems with detection infrastructure, and they manifest themselves in distinct ways in transmitted data, allowing for an automated quality control process. One example is the situation where PeMS receives no data or few data samples from a detector, a station, or a controller over a day. This is most likely evidence that a communication line is down or that there is a hardware malfunction in the device. Another example is a detector repeating the same flow and/or occupancy values across multiple time periods. Other examples include detectors reporting high occupancy values, indicating that the detector is stuck on, or reporting mismatched flow and occupancy values (for example zero flow and non-zero occupancy, or vice versa), indicating that the detector is hanging on. If PeMS detects any of these scenarios over a day, it discards the detector's data and imputes replacement values. In this imputation process, PeMS makes estimates of what the detector's data might have been based on developed statistical relationships with nearby detectors, or based on historical averages observed at the broken detector. PeMS then aggregates the full set of observed and imputed detector data across all lanes to the station level and computes spatial performance metrics. To inform the user about the quality of the data or performance measure that they are viewing, PeMS reports the "percent observed" of every metric, which represents the percentage of data points used to compute the metric that were directly observed from a detector, as opposed to imputed. For example, if the percent observed for a 5-minute travel time along a route is 75%, then 75% of the detectors on the route were reporting good data, and 25% were reporting bad data.

When the detector health algorithms were run on the NOVA data, we realized that the majority of the detectors on the selected corridors were reporting no data or bad data. Exhibit C3-6 plots the daily percentage of good detectors between March 1, 2011 and May 9, 2011, as well as the percentage of bad detectors attributable to the two leading causes: no data and stuck. The number of good detectors never exceeds 30%, and generally hovers around 25%. The highest percentage of detectors are in the "stuck" category, meaning that they are reporting constant flow and/or occupancy values. VDOT staff attributed this high percentage to the need to calibrate new detectors following large-scale construction projects. Additionally, a significant percentage of the detectors (around 30%) never sent PeMS any data. The days where there were

drops in every category represented times when internet outages prevented the research team from acquiring the XML files from the RITIS website.

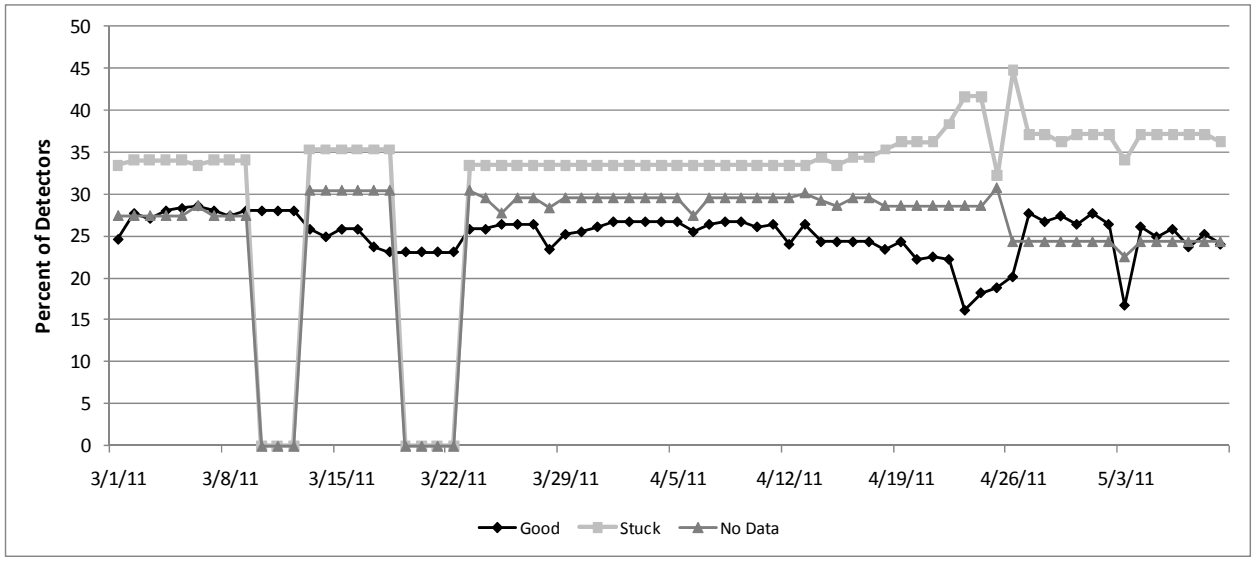

Exhibit C3-6: Daily Detector Health Status, NOVA PeMS Deployment, 2011

The low percentage of usable data available over the 2011 study time frame greatly concerned the research team, as the quality of computed travel times would be poorer given the missing data. Additionally, because the majority of detectors in the network never sent PeMS any good data, it was not possible to develop historical statistical relationships with the data at nearby working detectors needed to drive the most accurate imputation algorithms. Without these statistical relationships, PeMS had to use less robust imputation algorithms, which further decreased the accuracy of computed travel times. To show the effect this has on the detector data recorded, consider a detector on WB I-66 that fell into the "stuck" category for two weeks (Monday-Sunday). This resulted in imputed flows across that entire time period as shown in Exhibit C3-7. Because this detector never sent PeMS any good data, PeMS could only crude estimates of its flow values based on flows observed at nearby detectors. This meant that PeMS repeated the same flow, occupancy, and speed data for a given hour from week to week. In the sample plot, the hourly flows imputed for the first week are identical to those imputed for the second week. This constancy in imputed data is not ideal for computing travel time reliability, which relies on the ability of the traffic network to capture the real variability in conditions over time. Since data had to be imputed for such a large percentage of the detectors in 2011, the research team decided to seek additional data for 2009 hoping that the data quality would be sufficient to support methodological advancement and use case analysis. This effort is described in the following section.

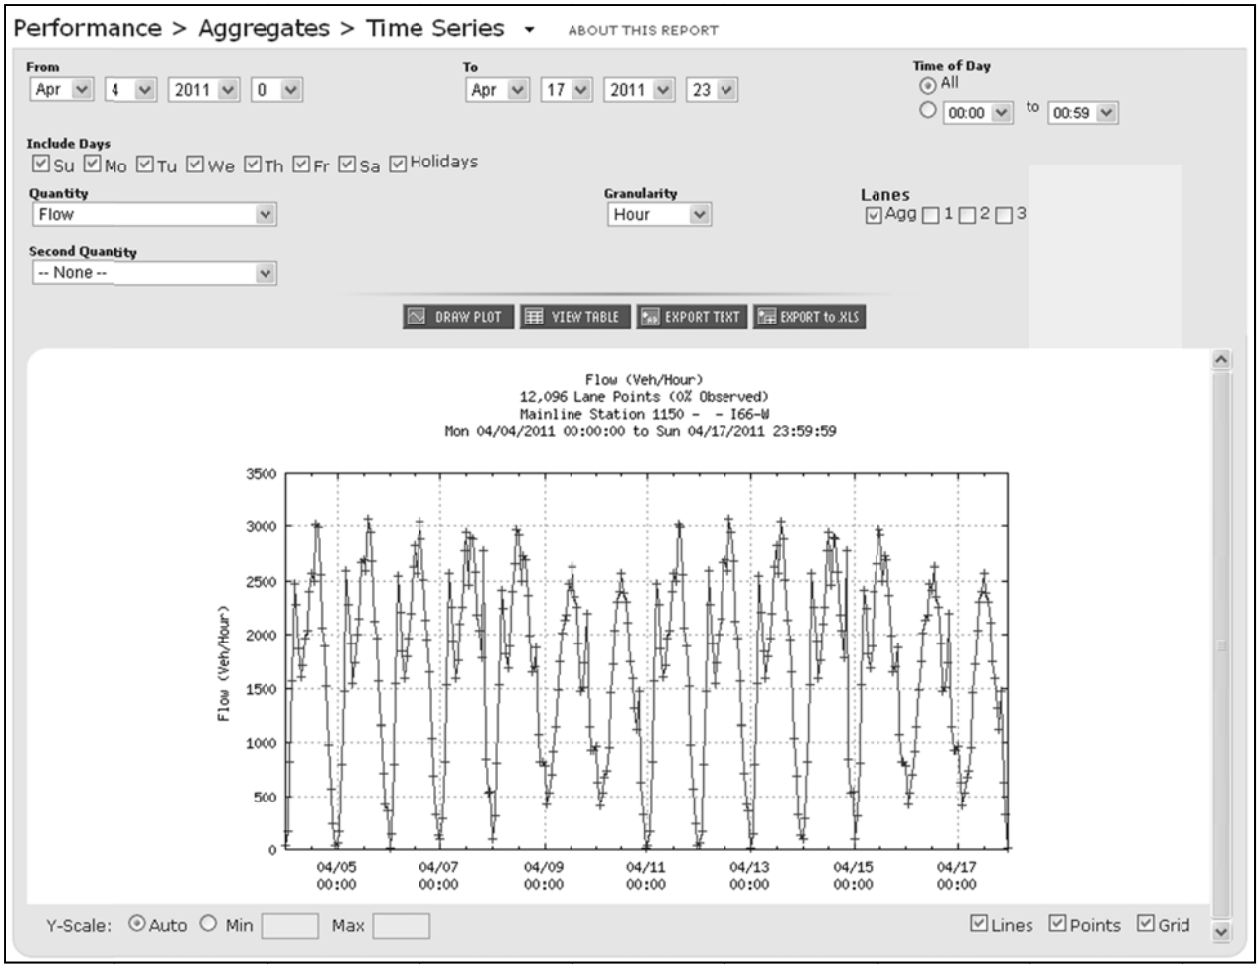

Exhibit C3-7: Imputed Flow Values at a Broken Detector

Historical Data

The research team worked with the University of Maryland CATT lab to obtain traffic data for 2009. The historical data was delivered in 12 zipped csv files, each about 45 MB in size. The csv files contained the same information as the traffic data XML files, so it was straightforward to write a program to parse information from the csv files and put it in the correct database tables in PeMS, with an associated timestamp corresponding to when the data was collected. The one issue that was encountered was that no historical configuration data was available. We manually compared the IDs of detectors and zones reported in the archived data with those present in the 2011 configuration XML file, and determined that the 2011 configuration data would suffice to represent the 2009 detector locations.

After the historical data was entered into the PeMS database, and processed, its health was then investigated to see if the 2009 data was better than the 2011 data. Exhibit C3-8 plots the weekly percentage of good detectors over 2009, as well as the percentage of detectors falling into the leading two error categories- No Data and Stuck. During 2009, the number of working detectors was significantly higher than in 2011, generally hovering above 70% for most of the year. The percentage of detectors that were stuck and transmitting constant data was much lower, being less than 5% across the whole year. The biggest error category remained the No Data

condition, which likely represented detectors that were listed in the configuration file but were not yet calibrated to send data.

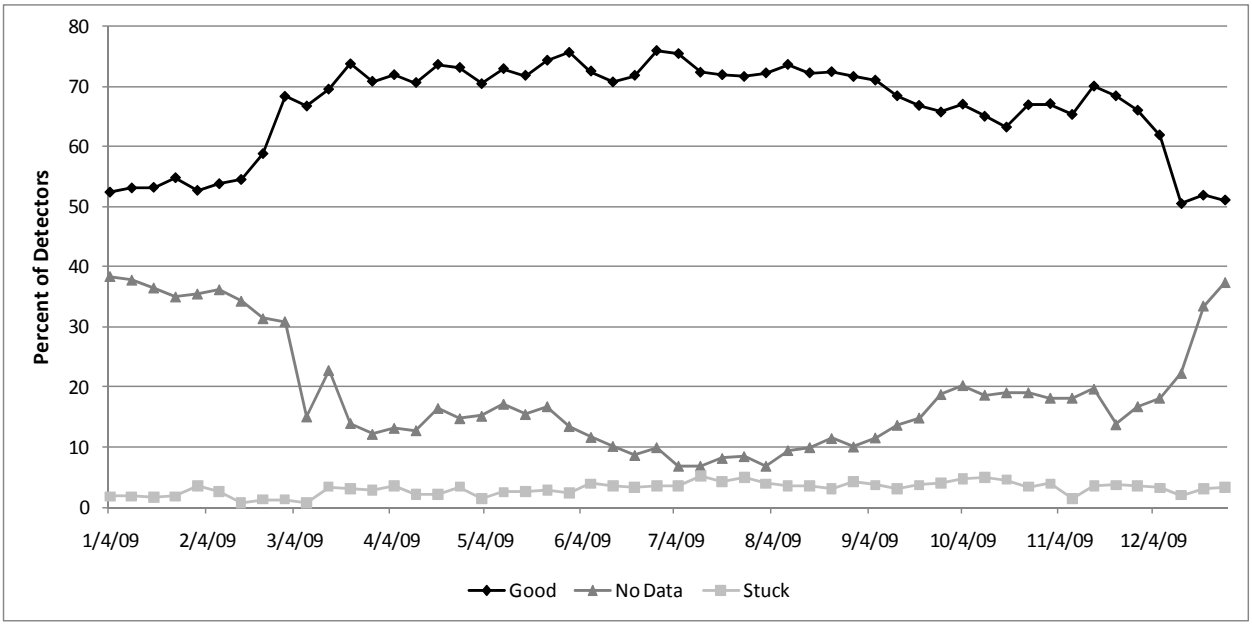

Exhibit C3-8: Weekly Detector Health Status, NOVA PeMS Deployment, 2009

Travel Time Results

Following acquisition of the real-time and 2009 data, eight different routes were constructed in PeMS to monitor reliability across different segments of the four study freeway-directions. Five-minute and hourly travel times were created for these eight routes for the entire year of 2009 and March through May in 2011. To evaluate the impact of individual detector data quality on route-level travel times, the team compared the route travel times for 2009 with those for 2011. Exhibit C3-9 plots the hourly travel times calculated on a 26 mile stretch of westbound I-66 for March through April of 2009. Overall, the PeMS percent observed for this travel time data was 79%. Exhibit C3-10 plots the same data for the same months of 2011; in this case, only 22% of the data were observed. Overall, the 2009 data follows the pattern expected of a highly congested facility with a peak hour commute; travel times are high on every weekday, but the peak value varies from day to day. Due to the high percentage of imputed detector data, the travel time patterns for the month of April 2011 the weekly travel time patterns look almost identical. It is doubtful that such consistency exists. These patterns are more likely to be caused by the high percentage of imputed data. For this reason, we chose to base the methodological advancements of this case study on the 2009 data, and use the 2011 data only to compare with probe travel time runs, to further evaluate the data quality.

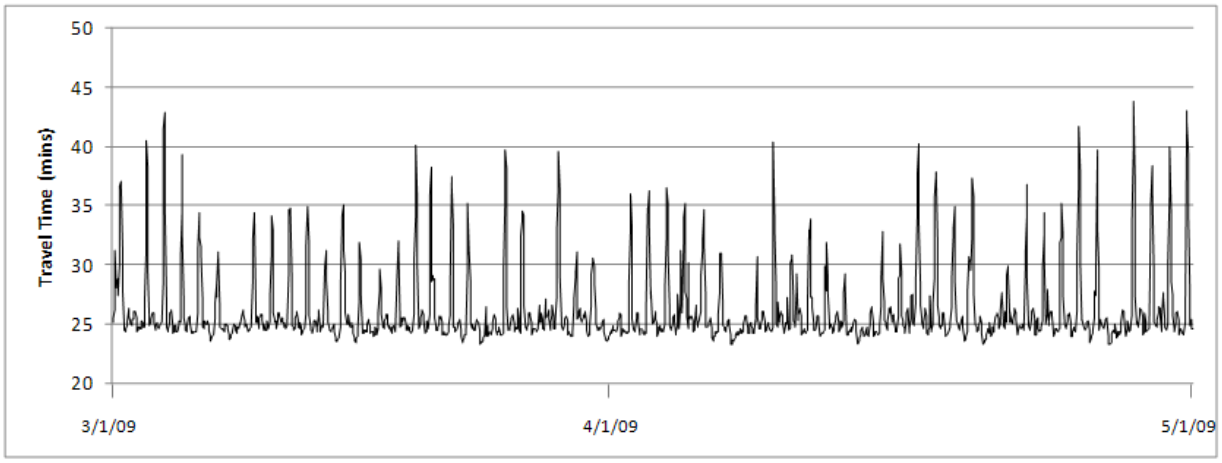

Exhibit C3-9: Travel Time, WB I-66, 3/1/09-4/30/09

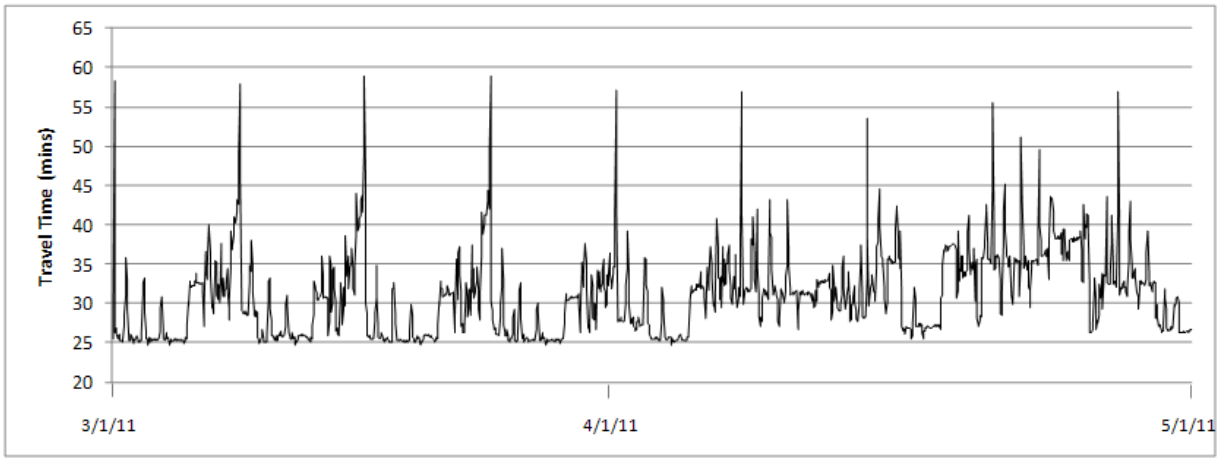

Exhibit C3-10: Hourly Travel Times, WB I-66, 3/1/11-4/30/11

Summary

Data collection is an essential part of any transportation planning or operations activity. Today, transportation agencies are increasingly turning to sophisticated sensor arrays to monitor the performance of their infrastructure, which allow for the use of advanced traffic management techniques and traveler information services. External systems, such as a reliability monitoring system, can leverage this data to further maximize the value of installing and maintaining detection. As evidenced by this case study, data collection for a travel time reliability monitoring system communication can be automated, but it requires significant time and resources to get it started.

References

1) Zhanfeng Jia, Chao Chen, Ben Coifman and Pravin Varaiy. The PeMS algorithms for accurate, real-time estimates of g-factors and speeds from single-loop detectors.

METHODOLOGY EXPERIMENTS

Overview

Because of the type of data available in this case study, and investigations done previously in the I-66 corridor, the research team elected to experiment with travel time reliability monitoring ideas that are being developed in SHRP2 project L10, Feasibility of Using In-Vehicle Video Data to Explore How to Modify Driver Behavior that Causes Non-Recurrent Congestion. In the SHRP2 project L10, researchers are experimenting with a multi-state travel time reliability modeling framework using mixed mode normal distributions to represent the PDFs of travel time data from a simulation model of eastbound I-66 in Northern Virginia. (They are also using this same technique to analyze travel times from toll tag data collected on I-35 in San Antonio *(1)*). This case study adopted that technique and applied it to the travel times calculated from the freeway loop detectors on eastbound I-66.

According to the SHRP2 L10 research, multi-state models are appropriate for modeling travel time distributions because most freeways operate in multiple states across the year (or some other timeframe): for example, an uncongested state, a congested state, and a state caused by non-recurrent events, such as incidents, construction, weather, or fluctuations in demand. This concept is illustrated in Exhibit C3-11, which shows the distribution of weekday travel times on a corridor in Northern Virginia. Three travel time "modes" are evident, which may be interpreted as the most frequently occurring travel times for the uncongested state, the congested state, and the non-recurring congestion state. Multi-state models also provide a helpful framework for delivering understandable information to the end consumer of travel time reliability information: the driver. They provide two pieces of information: (1) the probability that a particular state will be extant during a given time period; and (2) the travel time distribution for that state during that time period. This provides a way of creating reliability information that is similar to how people are accustomed to receiving weather forecasts, for example: "there is a 60% chance that it will rain tomorrow, and if it does rain, the expected precipitation will be 1 inch". The reliability analog to this is, for example: "The percent chance of encountering an incident-based congestion state during the AM peak period is 20%. If one does occur, the expected average travel time is 45 minutes and the 95th percentile travel time is one hour".

Beyond its suitability for modeling travel time distributions and providing useful metrics, this methodology was also fits well with the work that the SHRP2 L02 team is doing to develop travel time distributions for different operating regimes. The different states of the normal mixture models are the conceptual equivalent of the regimes that L02 is working on to classify the operating conditions of routes and facilities. It also provides an opportunity to test a methodology that was developed for modeling the distribution of *individual vehicle travel times* on *aggregated travel times* calculated from loop detectors.

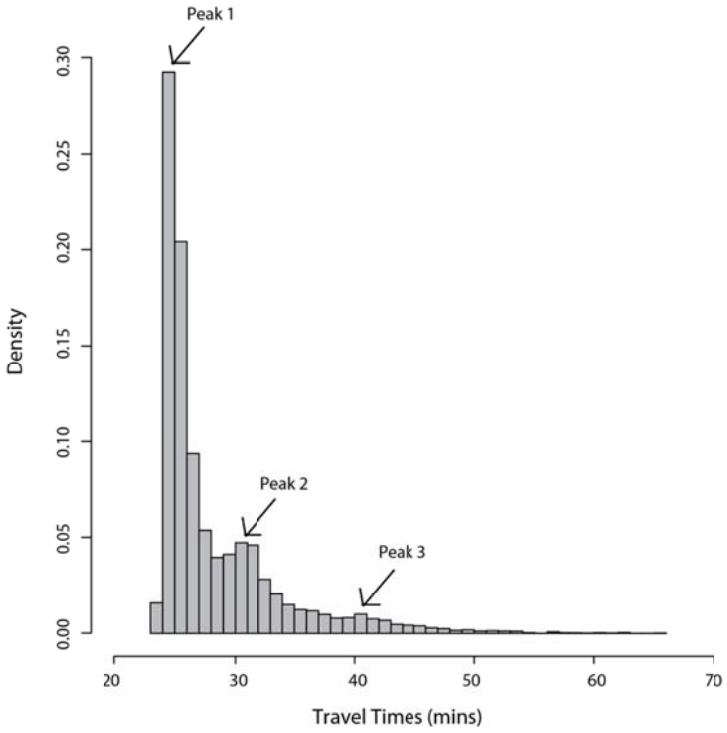

Exhibit C3-11: Distribution of Travel Times on EB I-66, 5:00 AM-9:00 PM

Site Description

A multi-state model was developed for a 26 mile stretch of eastbound I-66 from Manassas to Arlington, Virginia. A map of the corridor is shown in Exhibit C3-12. This segment of freeway is monitored by 96 sensors, which are a mix of radar detectors and loop detectors. The selected dataset consists of 17,568 travel times at the five-minute level aggregation, which represent the average travel time experienced by vehicles departing the route origin during that five-minute time period. This dataset covers the travel times for departures every five minutes during the weekdays between January 1, 2009 and March 30, 2009.

The route is a major commute path from the suburbs of Northern Virginia into Washington D.C. As such, it sees the highest demand levels during the AM peak period, as well as a smaller increase in demand during the PM peak. A PeMS-generated plot of the minimum, average, and maximum travel times by hour of the day measured over the study time frame is shown in Exhibit C3-13.

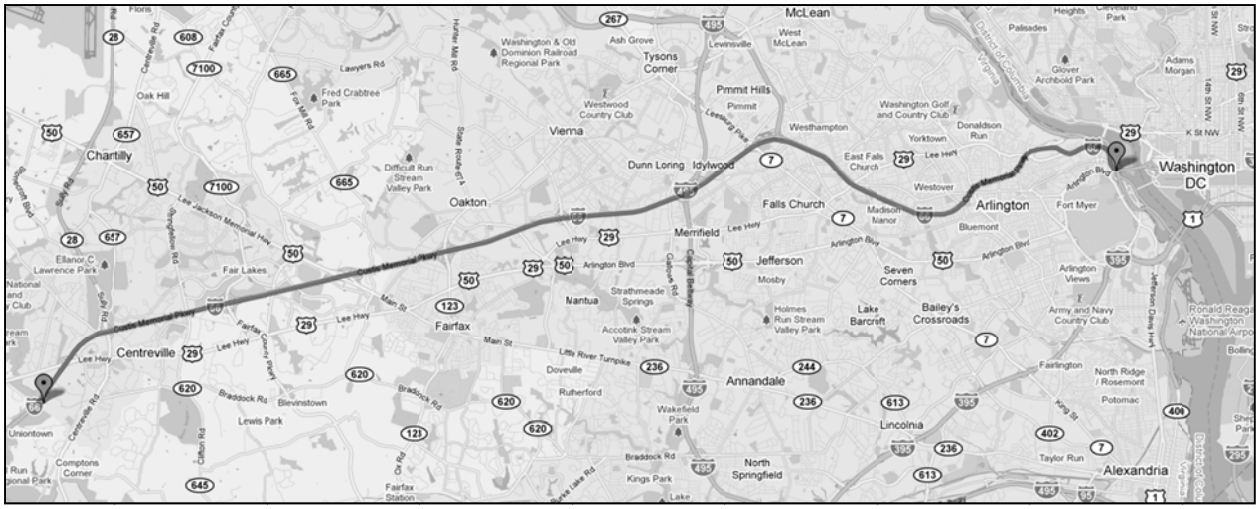

Exhibit C3-12: Map of Eastbound I-66 Study Corridor

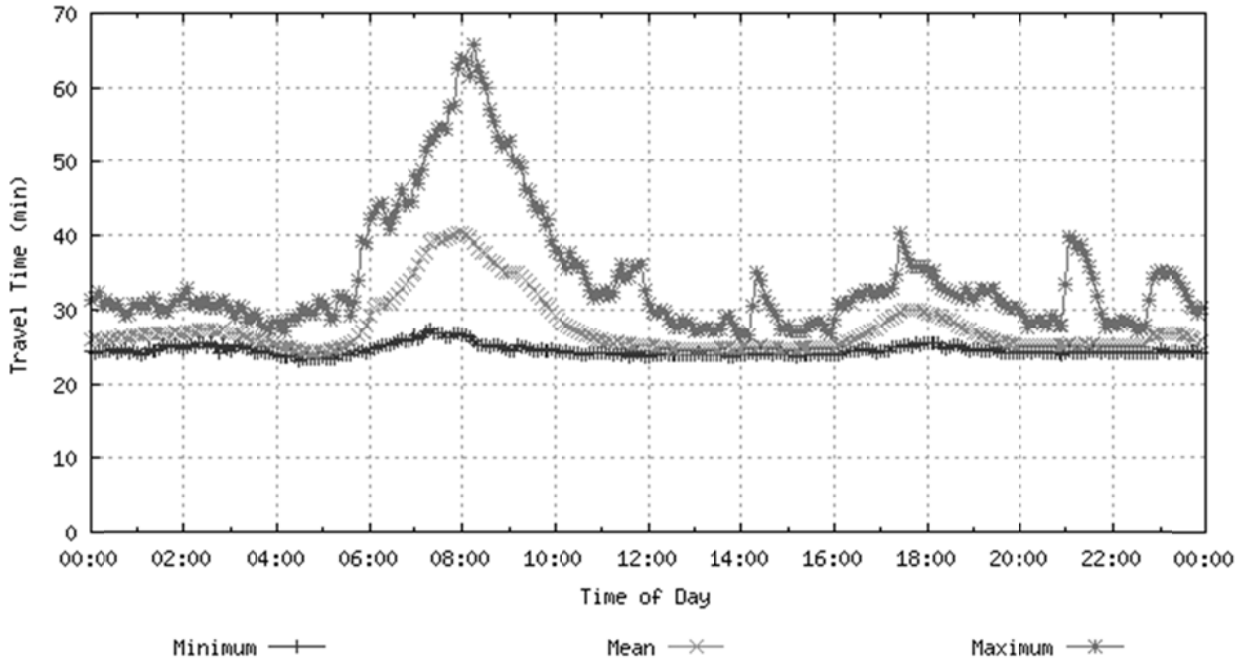

Exhibit C3-13: Minimum, average, and maximum corridor travel times, weekdays, 1/1/2009-3/30/2009

Method

The goal of this study was to generate, for each hour of the day, two outputs: (1) the percent chance that the traveler would encounter a certain condition; and (2) for each condition, the average and 95th percentile travel time. The mathematical details of these steps are explained in the referenced paper by Rakha et al *(1)*. Under this framework, three questions are answered for each time period:

1) How many states are needed to model the travel time distribution?

2) What is the probability of each state occurring?
3) What parameters describe the normal distribution for each state?

Analysis was performed using the Mclust package in R, which provides functions to support normal mixture modeling and model-based clustering *(2)*. Normal mixture models were developed to represent travel times for each hour of the day between 4:30 AM and 12:30 AM. The early morning hours were not considered due to the lack of any congestion. The first question above was answered by putting the data set for each hour into a function that initially clusters the data into the number of states that provide a best fit (in this paper, the "optimal" number of states). The best-fit was determined using the Bayesian Information Criterion (BIC), defined as $-2\log(L) + k\log(n)$, where L is the likelihood function of the model parameters, k is the number of parameters, and n is the sample size of the data. This function considers the fit of the model while also penalizing for an increased number of parameters, to prevent against over-fitting. The model with the number of states that produces the lowest BIC is selected as the optimal model, and each data point is given an initial probability of belonging to each state.

The outputs of this step (the model type, number of states, and initial probabilities of a data point belonging to each state) are then put into an Expectation-Maximization (EM) algorithm, which is an iterative method, appropriate for mixture models, that is used to find maximum likelihood estimates of parameters. The EM algorithm outputs the mixture component for each state (its probability of occurrence), the mean and standard deviation of each state, and the final estimates of the probability that a data point belongs to each state. These estimates are used to form the user-centric output of, for example, "If you depart on a trip at noon, you will have a 20% chance of experiencing congestion. During congestion, the average travel time is 30 minutes and the 95^{th} percentile travel time is 45 minutes."

During the analysis process, complications arose that required the research team to balance the desire for a best-fitting model with the need to provide useful and clear information to the end user. The initial clustering step suggested that either three or four states were needed to optimally model the travel times for each hour. The "optimal" number of states for each hour's model is summarized in Table C3-1along with the associated BICs. However, in the practical realm, a historical set of travel times from a given hour can be conceptualized as consisting of only up to three states. Early morning time periods may only have one state, non-congested, and can thus be described by a single distribution. Time periods where demand fluctuates may have two states: non-congested and congested, with congestion being triggered either by high demand or a non-recurrent condition. Finally, the peak periods may have three states: a non-congested state, likely rare, when demand is low, a congested state, which is common, and a very congested state, which may be triggered by an incident or special event. The fourth state has no clear physical explanation that can be effectively conveyed to the end user. As such, each hour's data set was run through the clustering algorithm again, this time with a constraint of three maximum states. The "constrained" best-fit state for each hour and its associated BIC is shown in Table C3-1. Three states provided the best-fit for all but two hours (12:30 PM and 2:30 PM), when two states provided the best-fit.

Following the EM step using the "constrained" number of states, the mean travel time estimates for each state were evaluated. These mean travel times are summarized in Table C3-2. For the majority of hours (all hours outside of the AM peak), the mean travel times for state 1 (S1) and state 2 (S2) were very similar (within 3 minutes of each other). These are denoted in the table by gray shading. Because such small differences in average travel times are not meaningful enough to the end user to be considered different states, any hour where three states were

suggested but mean travel times between consecutive states differed by less than three minutes were reduced to two states. The model parameters were then re-estimated for this final number of states. In the end, the models for each hour were composed of two states, with the exception of the AM peak hours (6:30 AM-10:30 AM) which remained composed of three states. The final number of states and associated BICs for each hour are shown in the final column in Table C3-1.

Table C3-1: Selection of States

Hour	**Optimal**		**Constrained**		**Final**	
	States	**BIC**	**States**	**BIC**	**States**	**BIC**
4:30-5:30 AM	3	1387	3	1387	2	1443
5:30-6:30 AM	4	3580	3	3595	2	3620
6:30-7:30 AM	3	4322	3	4322	3	4322
7:30-8:30 AM	3	5017	3	5017	3	5017
8:30-9:30 AM	4	4854	3	4855	3	4855
9:30-10:30 AM	3	3876	3	3876	3	3876
10:30-11:30 AM	4	2561	3	2567	2	2622
11:30-12:30 PM	3	1578	3	1578	2	1640
12:30-1:30 PM	4	960	2	968	2	968
1:30-2:30 PM	3	1081	3	1081	2	1132
2:30-3:30 PM	3	1118	3	1118	2	1153
3:30-4:30 PM	3	1675	3	1675	2	1725
4:30-5:30 PM	2	3074	2	3074	2	3074
5:30-6:30 PM	3	3160	3	3160	2	3170
6:30-7:30 PM	3	2793	3	2793	2	2812
7:30-8:30 PM	4	1459	3	1464	2	1477
8:30-9:30 PM	3	1283	3	1283	2	1291
9:30-10:30 PM	4	1220	3	1220	2	1233
10:30-11:30 PM	3	2398	3	2398	2	2488
11:30-12:30 AM	3	2162	3	2162	2	2178

Table C3-2: Mean Travel Times by State for Constrained Parameters

Hour	S1	S2	S3
4:30 AM-5:30 AM	24	25	29
5:30 AM-6:30 AM	25	26	31
6:30 AM-7:30 AM	28	33	37
7:30 AM-8:30 AM	28	39	46
8:30 AM-9:30 AM	28	34	42
9:30 AM-10:30 AM	26	29	34
10:30 AM-11:30 AM	25	26	28
11:30 AM-12:30 PM	25	25	29
12:30 PM-1:30 PM	23	25	--
1:30 PM-2:30 PM	24	26	30
2:30 PM-3:30 PM	24	25	27
3:30 PM-4:30 PM	24	26	29
4:30 PM-5:30 PM	26	27	--
5:30 PM-6:30 PM	25	31	31
6:30 PM-7:30 PM	27	27	31
7:30 PM-8:30 PM	25	26	28
8:30 PM-9:30 PM	25	26	31
9:30 PM-10:30 PM	25	26	28
10:30 PM-11:30 PM	25	26	31
11:30 PM-12:30 PM	25	26	30

Results

This section first summarizes the travel time reliability findings for each weekday hour. It then provides an in-depth analysis of model results for the AM peak hours.

Overall

Exhibit C3-14 presents, for each hour of the day and for each state, the probability of the state's occurrence (top-left), the mean travel time (top-right), the standard deviation of its travel times (bottom-left), and the 95th percentile travel time (bottom-right). Estimates for state 1 are shown in the dashed line, state 2 in the solid line, and state 3 (where applicable) in the bold line. Values are also summarized in Table C3-3.

As can be seen in the plot of each state's probability, state 1 is by far the most common state encountered during the early morning, the midday, and the late night hours. When this state is active during these hours, the mean travel time tends to be near free-flow, at around 25 minutes, the standard deviation is low, and the 95th percentile is close to the mean. During these off-peak periods, the percent chance of congestion (state 2), generally stays between 10% and 20%. Even when the congested state is active during these hours, the mean travel time is still generally less than 30 minutes, and the 95th percentile travel time generally less than 35 minutes.

At the beginning of the PM peak (4:30 PM-5:30 PM), state 1 and state 2 each have a 50% chance of occurring. During the PM peak hour (5:30 PM-6:30 PM), the probability of congestion increases to 67%. At the end of the PM peak hour (6:30 PM-7:30 PM), the probability of state 1

and state 2 effectively swap; state 1 has a 64% change of occurring and state 2 a 36% change of occurring. Throughout the PM peak, the mean and 95th percentile travel times of each state are consistent. State 1 has a mean travel time of 26 to 27 minutes and a 95th percentile travel time of 27 to 28 minutes, and state 2 has a mean travel time of 30 to 31 minutes and a 95th percentile travel time of 33 to 34 minutes.

The four hours of the AM peak (6:30 AM-10:30 AM) have three active states, as they have both the most congestion and travel time variability. Within these four hours, however, both the relative probabilities of each state and the parameters of each state differ significantly. State 3 (conceptualized as the non-recurrent congestion state) has the greatest chance of occurring at the beginning of the AM peak, between 6:30 AM and 7:30 AM, and between 8:30 AM and 9:30 AM (41% and 39%, respectively). Its likelihood is around 25% during the other two hours. The severity of congestion in this state differs across each hour. It has the highest mean travel time (46 minutes) and 95th percentile travel time (58 minutes) during the 7:30 AM hour, indicating that this is the true AM peak hour. At 8:30 AM, the mean travel time of this state is reduced to 42 minutes, and the 95th percentile travel time to 51 minutes. On the shoulders of the AM peak, the mean travel times of state 3 are 34 and 37 minutes and the 95th percentile travel times are 40 and 44 minutes. State 2 occurs with varying probabilities during the AM peak, ranging from a low of 32% at 8:30 AM to a high of 58% at 7:30 AM. The mean and 95th percentile travel times of state 2 are significantly higher during the AM peak than at any other time period of the day. Even though this time period usually experiences congestion and some travel time variability, there are days (approximately one out of five) when the corridor operates in the uncongested state, and mean travel times are around 28 minutes.

The information gained from the example plot and accompanying table can be used to provide intuitive and useful information to the traveling public, in ways illustrated in the following section, which focuses on interpreting the results for the AM peak hours.

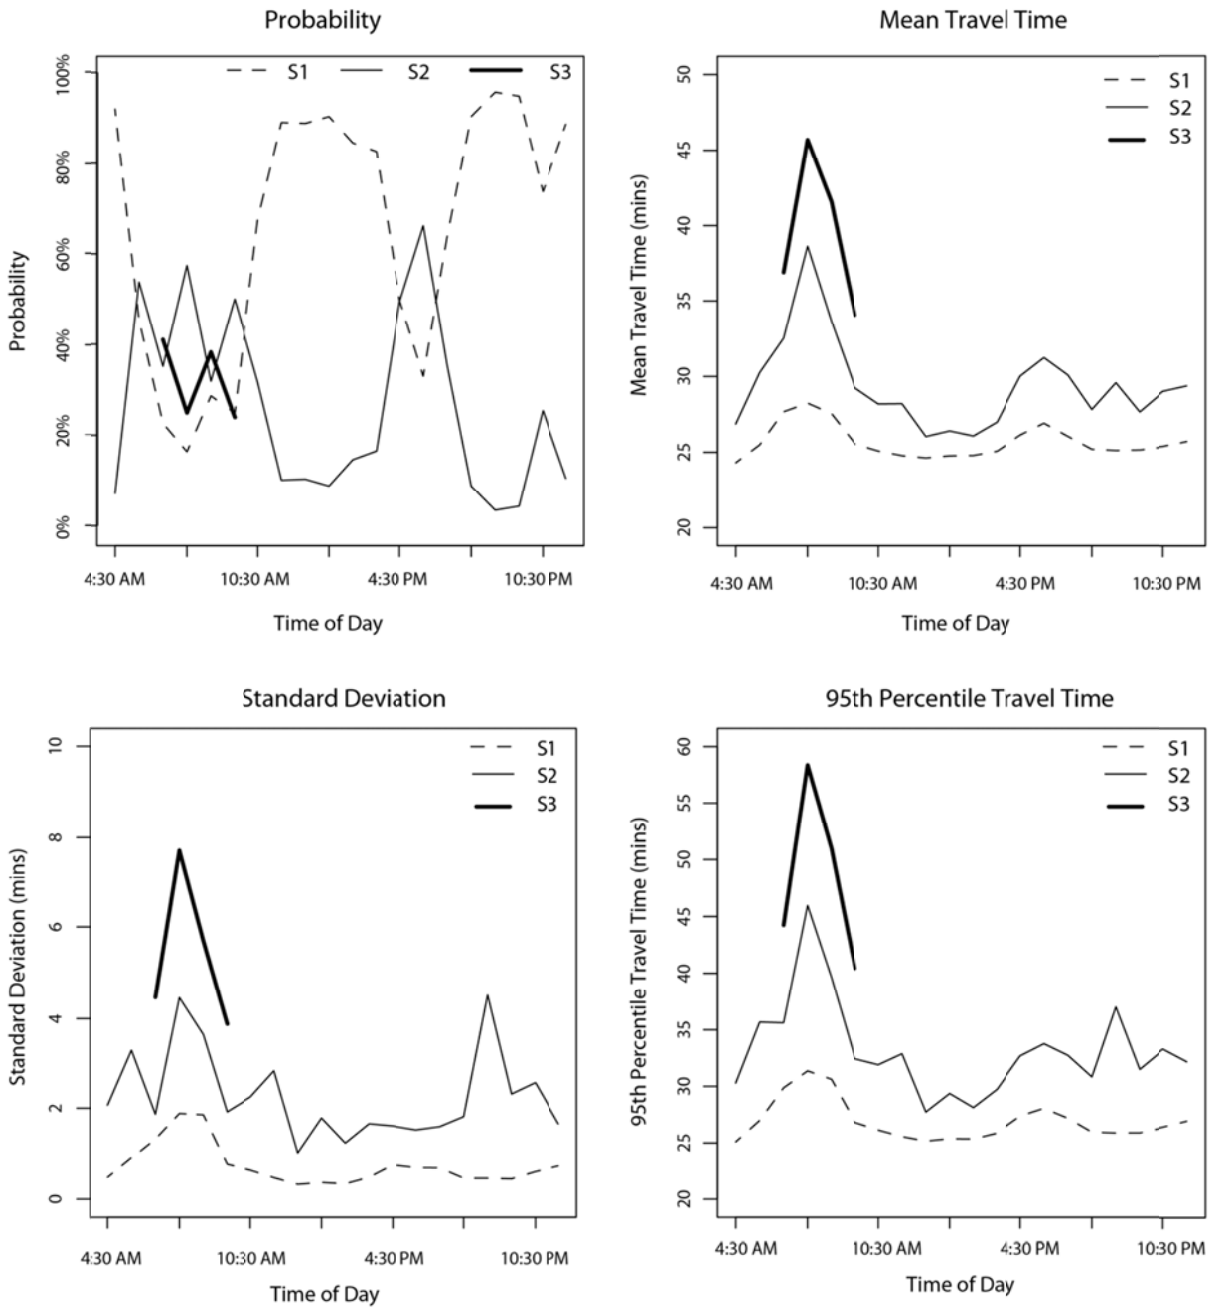

Exhibit C3-14: State Probabilities, Mean Travel Times, Standard Deviation, and 95th Percentile Travel Times by Time of Day

Table C3-3: Probability, Mean Travel Time, Standard Deviation, and 95th Percentile Travel Time by State

Time	Probability			Mean			Std. Dev.			95th percentile		
	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3
4:30 AM	92%	8%	--	24	27	--	0.5	2.1	--	25	30	--
5:30 AM	46%	54%	--	25	30	--	0.9	3.3	--	27	36	--
6:30 AM	23%	36%	41%	28	33	37	1.3	1.9	4.5	30	36	44
7:30 AM	17%	58%	25%	28	39	46	1.9	4.5	7.7	31	46	58
8:30 AM	29%	32%	39%	28	34	42	1.9	3.6	5.7	31	40	51
9:30 AM	25%	50%	24%	26	29	34	0.8	1.9	3.9	27	32	40
10:30 AM	68%	32%	--	25	28	--	0.6	2.3	--	26	32	--
11:30 AM	89%	11%	--	25	28	--	0.5	2.8	--	25	28	--
12:30 PM	89%	11%	--	25	26	--	0.3	1.0	--	25	28	--
1:30 PM	91%	9%	--	25	26	--	0.4	1.8	--	25	29	--
2:30 PM	85%	15%	--	25	26	--	0.3	1.2	--	25	28	--
3:30 PM	83%	17%	--	25	27	--	0.5	1.7	--	26	30	--
4:30 PM	50%	50%	--	26	30	--	0.8	1.6	--	27	33	--
5:30 PM	33%	67%	--	27	31	--	0.7	1.5	--	28	34	--
6:30 PM	64%	36%	--	26	30	--	0.7	1.6	--	27	33	--
7:30 PM	91%	9%	--	25	28	--	0.5	1.8	--	26	31	--
8:30 PM	96%	4%	--	25	28	--	0.4	4.5	--	26	37	--
9:30 PM	95%	5%	--	25	28	--	0.4	2.3	--	26	31	--
10:30 PM	74%	26%	--	25	29	--	0.6	2.6	--	26	33	--
11:30 PM	89%	11%	--	26	29	--	0.7	1.7	--	27	32	--

AM Peak

As discussed in previous sections, a three-state normal mixture model was selected to measure reliability statistics for the four AM peak hours. Exhibit C3-15 provides a visual comparison of the relative model fits of the three-state normal mixture model, a two-state normal mixture model, and a lognormal distribution model. These fits are also quantitatively summarized in Table C3-4. It compares the BICs for each model for each hour. Visually, it is clear that for every hour except 8:30 AM, the three-state normal model approximates the data the most closely. This is also reflected in the BIC values, which are the lowest for the three-state normal mixture model. During the 8:30 AM hour, the fits between the three-state and two-state mixture models appear comparable, and their BICs are essentially equivalent.

Exhibit C3-16 provides a clearer visual comparison of the different travel time distributions within each morning hour by plotting them on the same x- and y-axis scales. It is evident that the two middle peak hours (7:30 AM and 8:30 AM) have the most travel time variability, while the distributions for the shoulder hours are more tightly packed. In particular, there is a large spike in the travel time distribution for the 9:30 AM hour at 25 minutes, which is essentially free-flow for this corridor. In this figure, each bar of the travel time histogram is shaded according to which state the model determined it was the *most likely* to fall into. It is important to make clear that there are no clearly defined boundaries for each state; rather, for each observed travel time, the model provides the percentage chance that the data point belongs

to each state. For some values (for example, 24 minutes), there is a near 100% likelihood that the travel time belongs in state 1. For others, such as a 46 minute travel time during the 7:30 AM hour, there is a near 50% chance that the data point belongs to state 2 and a near 50% chance it belongs to state 3. As such, these shadings are meant only to be a rough visualization of the component travel times of each state.

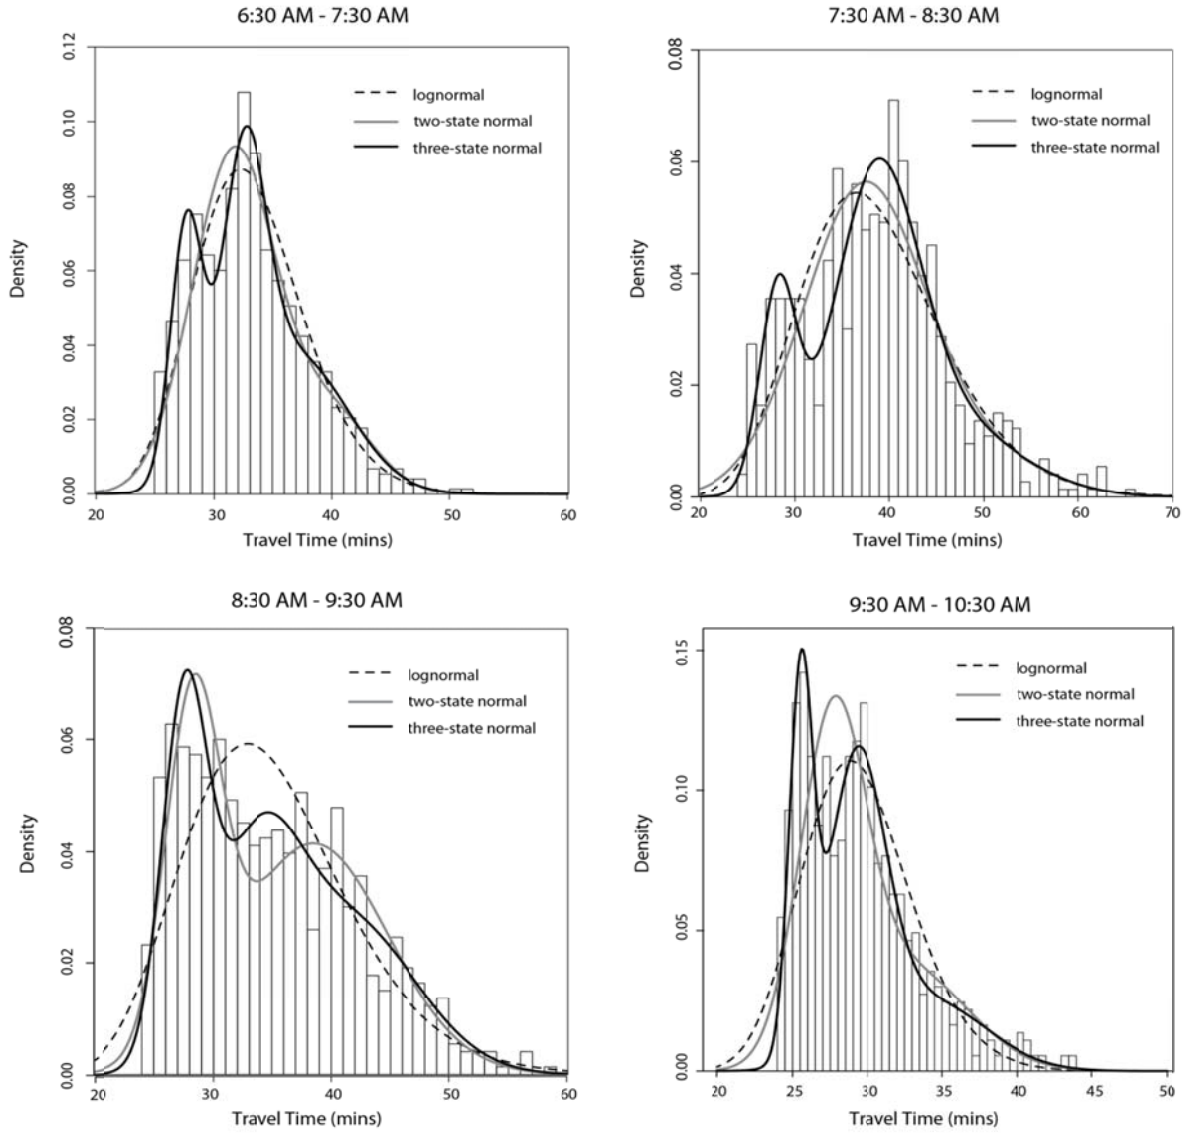

Exhibit C3-15: Lognormal and two- and three-state normal mixture models for AM peak hours

Table C3-4: BICs by distribution model

	3-state normal	**2-state normal**	**Log-normal**
6:30 AM-7:30 AM	4322	4346	4330
7:30 AM-8:30 AM	5017	5053	5034
8:30 AM-9:30 AM	4856	4856	4910

9:30 AM-10:00 AM	3876	3954	3981

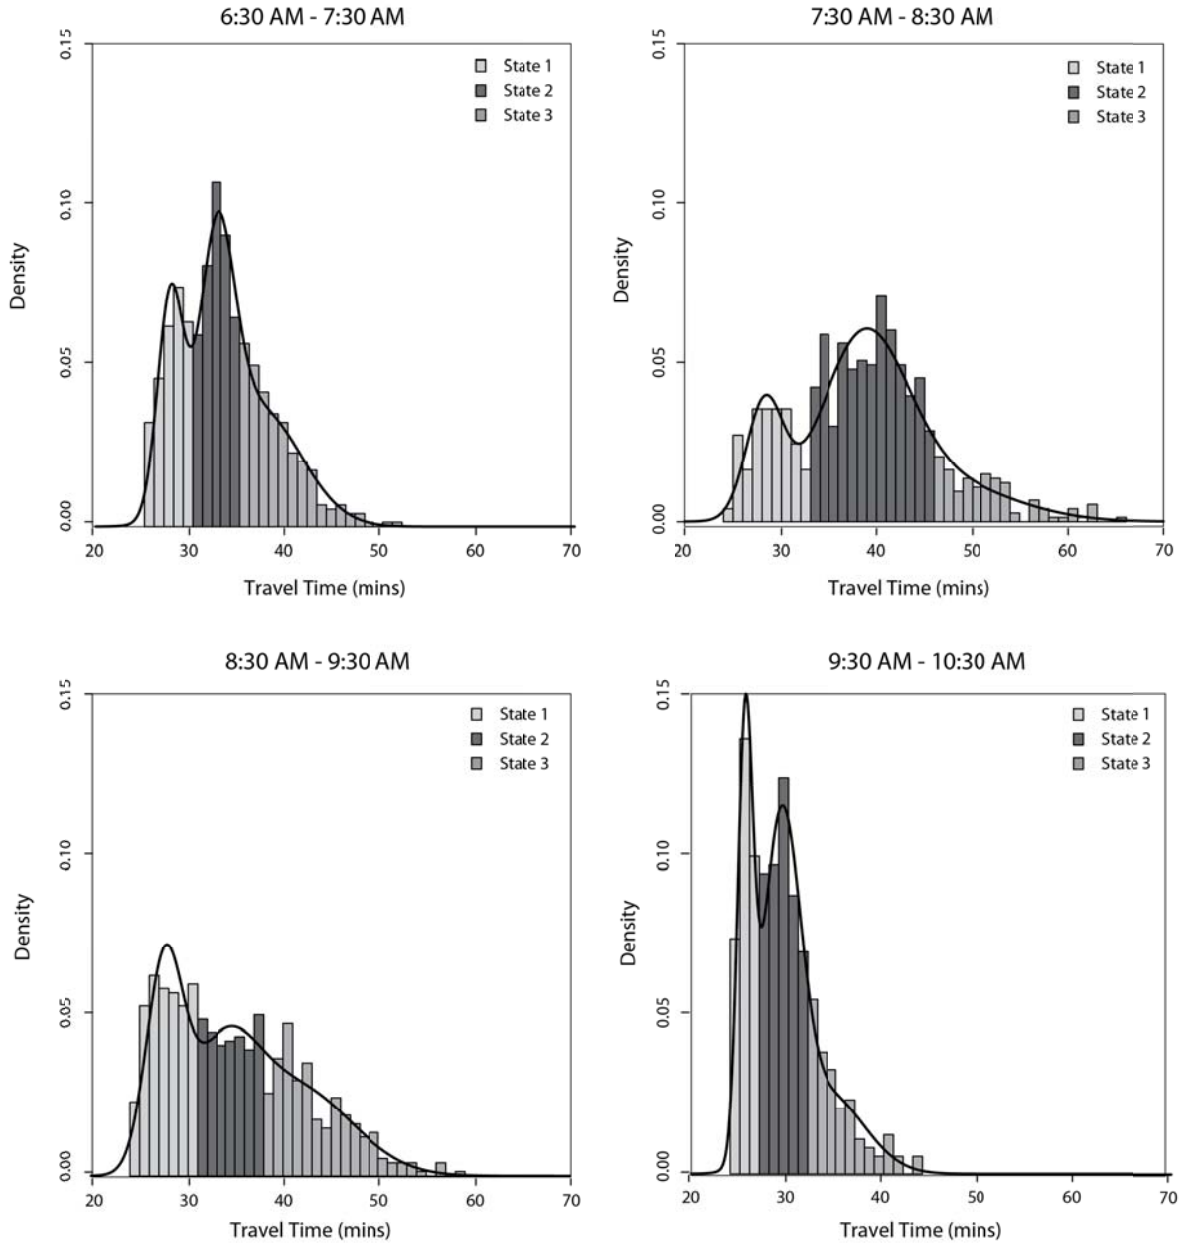

Exhibit C3-16: Travel time distributions and states, AM peak

The desired final output of these analyses is reliability information that can be readily interpreted and consumed by corridor drivers who are planning to make a trip at a certain time. From the information presented above, the following examples convey information that could be provided to drivers on a pre-trip basis, to aid them in their planning process:

- For trips made between 7:30 AM and 8:30 AM, there is a 60% chance of experiencing congestion. If congestion occurs, the expected travel time is 39 minutes

and the 95th percentile travel time is 46 minutes. There is also a 25% chance of experiencing severe, incident-based congestion. If this occurs, the expected travel time is 46 minutes and the 95th percentile travel time is 58 minutes.

- For trips made between 9:30 AM and 10:30 AM, there is a 50% chance of experiencing congestion. If congestion occurs, the expected travel time is 29 minutes and the 95th percentile travel time is 32 minutes. There is also a 25% chance of experiencing severe, incident-based congestion. If this occurs, the expected travel time is 34 minutes and the 95th percentile travel time is 40 minutes.

Summary

This case study leverages the methodologies developed by the SHRP2 L10 research team and applies them to three months of five-minute aggregated loop detector data collected on a 26 mile corridor of eastbound I-66 in northern Virginia. The results indicate that normal mixture models reasonably approximate travel time data observed within a given time period. Two-state models seem sufficient to accurately model off-peak hours, while three-state models are needed to capture the variability during the peak hours. Beyond providing a good fit to travel time data, mixture models also output data in a form that can be easily conveyed to help end users better plan for trips.

References

1) Rakha, H., F. Guo, S. Park. A Multistate Model for Travel Time Reliability. In Transportation Research Record: Journal of the Transportation Research Board, No. 2188, Transportation Research Board of the National Academies, Washington, D.C., 2010, pp. 46-54.
2) C. Fraley and A.E. Raftery. MCLUST Version 3 for R: Normal Mixture Modeling and Model-Based Clustering. Technical Report No. 504. Department of Statistics, University of Washington, 2009. http://www.stat.washington.edu/fraley/mclust/tr504.pdf.

PROBE VEHICLE COMPARISONS

Introduction

To better understand the implications of the data quality issues on travel times, the team performed a quality control procedure. Probe vehicle runs were conducted along I-66 to amass "ground-truth" data that could be compared with the sensor data. A GPS-based data collection device was used capable of collecting data at 1-second intervals. The sections of roadway along which probe runs were conducted, and details concerning the sensor data collected as part of this effort are described in Table C3-5:

Table C3-5: Overview of Probe Runs

Segment	Route	Time Period	Runs	Start and End Mileposts	# Sensors	Date
A > B	I-66 EB	PM Peak	1, 2, & 3	68.5 – 74.3	4	April 19, 2011
C > D	I-66 WB	PM Peak	4, 5, & 6	74.2 – 69.9	3	April 19, 2011
E > F	I-66 EB	AM Off-Peak	7, 8, & 9	54.4 – 56.3	4	April 20, 2011
G > H	I-66 WB	AM Off-Peak	10, 11, & 12	56.3 - 54.4	4	April 20, 2011

Along this corridor, as elsewhere in the study region, point detectors are placed at approximately ½ mile intervals. Due to accuracy and maintainability issues with inductive loop detectors and other older sensors, there are no plans to replace failed units which have been deployed on the mainline lanes of NOVA region freeways. Instead, plans are in motion to transition to the use of non-intrusive radar-based detection technologies along the freeways. These sensors are being deployed both as replacements for older failed units, as well as new installations. As a result of a combination of the failure of some older loop detector stations, ongoing roadway construction, and the need to configure many of the newer radar-based units, data is currently available for only about 75 of NOVA's freeway detectors. Exhibit C3-17, below, provides a visual indication of the availability of data on I-66 and I-395; darker colored icons indicate working stations, lighter colored icons indicate non-working stations.

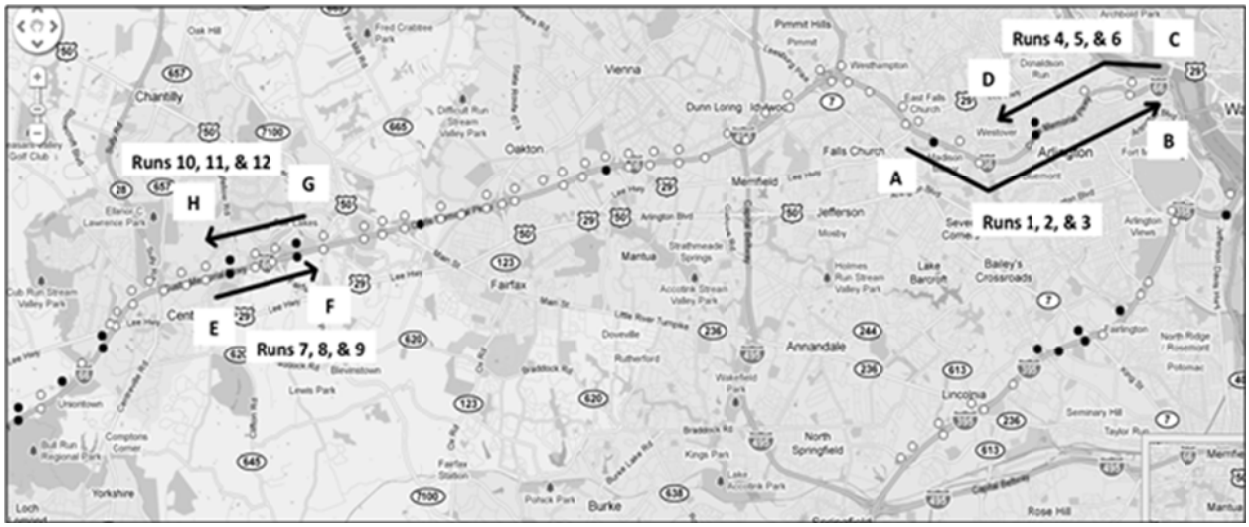

Exhibit C3-17: Display of Functioning vs. Non-Functioning Sensor Stations

Data-Related Issues Associated with NOVA Sensors

As discussed, construction and maintenance-related issues have resulted in a limited number of operational sensors from which data is available for use. In addition, a number of sensors that at first appear to be in working order are actually transmitting speed and/or flow data of questionable quality. For example, Exhibit C3-17 indicates that there are five (5) working

sensors operating in close proximity to one another along I-395. However, a closer analysis of the data output by several of these sensors indicates conditions that are either decidedly irregular, or are simply inaccurate. Examples are shown in Exhibit C3-18:

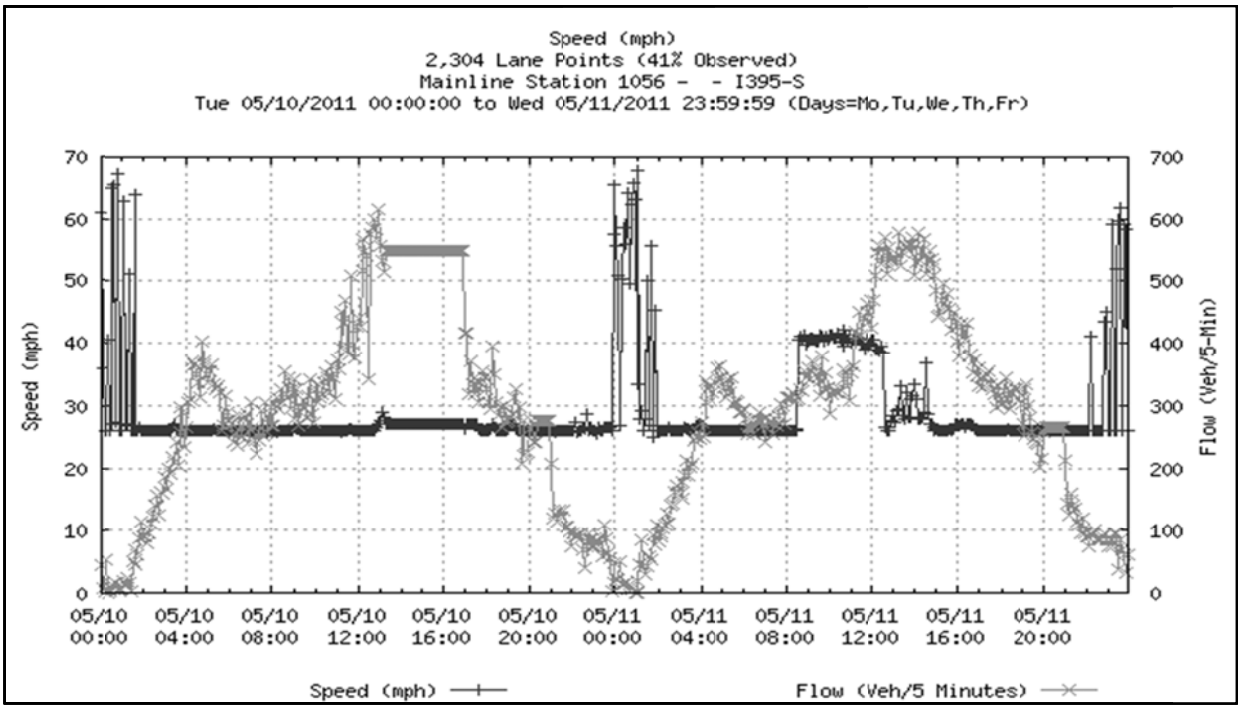

Exhibit C3-18: Speed/Flow Data from Suspect Sensor Along I-395 (From NOVA PeMS System)

Although the sensor providing the speed and flow data in Exhibit C3-18 appears to be functioning properly (as reported by the automated system used by the team to collect and analyze data as part of this project), a review of the speed data (Y-axis) and flow data (Z-axis) indicates the following:

- Speeds reported by this sensor are approximately 27 mph at all times of day except during the middle of the night, when traffic speeds increase significantly.
- The reported traffic flows appear fairly normal (with the exception of an apparent issue occurring between approximately 1pm and 5pm on May 10th), except that the peak traffic volume is reported as occurring between noon and 3pm, rather than the typical 4 to 7 pm. A field review of conditions by team members at this location and during this time period does not support this suggested condition.

A review of data collected from other sensors along I-395 (southbound) adjacent to this detector show similar conditions in that the peak traffic flow is reported as occurring between noon and 3pm, resulting in a concomitant drop in speeds to between 30 and 40 mph. As indicated above, a field review of conditions did not support this reported condition.

Error! Reference source not found., below, shows similar issues for sensors along I-66.

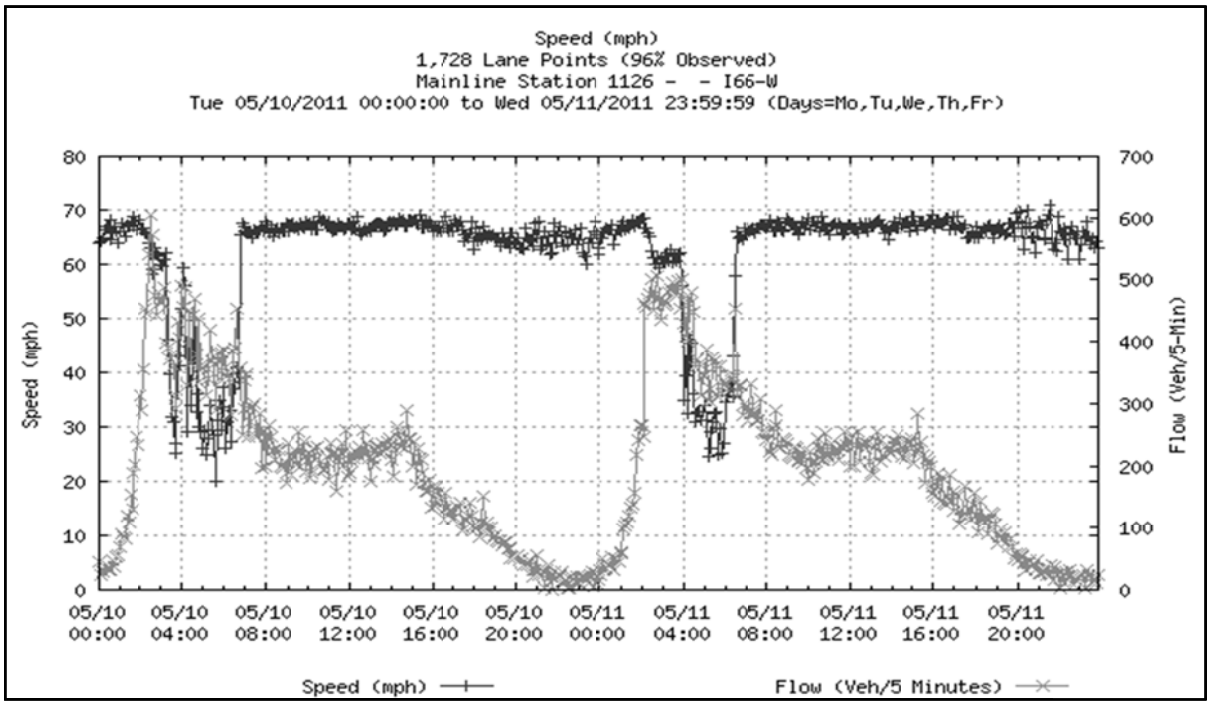

Exhibit C3-19: Speed/Flow Data from Suspect Sensor Along I-66 (From NOVA PeMS System)

As with the sensor data reported in Exhibit C3-18, data from the sensor displayed in **Error! Reference source not found.** indicates the existence of conditions along I-66 that diverge from conventional wisdom concerning the time of day at which the peak travel condition occurs. As per this data, peak volumes and the lowest speeds regularly occur at this location between approximately 2:00 am and 6:30 am, with speeds near 70 mph present during the remainder of the day. Again, a field review conducted by team members indicated that these data do not accurately represent the conditions that really exist.

It is likely that some portions of these data-related issues are the result of the high percentage of imputed detector data being used to represent conditions at many detector stations (e.g., 59% of data used generate the contents of Exhibit C3-18are imputed rather than observed). However, an even more significant issue is related to the need for these detectors to be fully calibrated on a system-wide basis so as to ensure they accurately represent real world conditions. Although VDOT is currently in the process of doing this, the team recommends that speed, flow, and estimated travel time data derived from these quantities be used sparingly until this process is complete. Failure to do so may result in decisions based on largely erroneous data, potentially resulting in a significant waste of resources and labor.

Methodology

The primary question the team wanted to answer in this probe-based experiment was: how well do the probe data align with the traffic speed and travel time estimates provided by the sparsely deployed point-based detectors? The primary method for answering this question was to compare data collected at 1-second intervals from a GPS-based data collection device against speed estimates generated based on data from Virginia DOT sensors deployed along each of the

four sections of I-66 described above. As part of this effort, the following analytical approach was used:

For each segment of roadway, graphs were used to compare the speed of the probe vehicle with speeds reported by the sensors. Speeds were displayed on the vertical axis and milepost on the horizontal axis. The solid line represented the speed estimates generated by the sensors (based on aggregate data collected from all lanes of travel), and the dotted line represented the probe vehicle speeds. In cases where data from the sensors was of suspected quality, the line representing the speed estimate provided by that sensor was dashed rather than completely solid. The locations of all the sensors from which data were collected along each roadway was indicated by a solid circle at the mid-point of each segment, accompanied by the sensor's identification number. We subsequently provided analysis of the differences between these two data sets along each segment.

In addition to analyzing the speed data as described above, the team conducted an analysis of the differences between the travel times experienced by the probe vehicle during each trip versus the estimated travel times generated from the sensor speeds. In situations where unreliable sensor data was present, a combination of observed sensor speeds and imputed speeds was used to fill in the gaps. Results of each analysis were then compared to calculate the average (absolute) error for each segment of roadway, as well as for the complete set of runs as a whole.

Data Analysis

The speed data from the probe-based runs was compared with the speed estimates generated using the spot speed sensors located along the same sections of roadway.

Data Analysis Along I-66 Inside of I-495 (Eastbound)

Exhibit C3-20, Exhibit C3-21, and Exhibit C3-22 show plots of the instantaneous speeds recorded by the vehicle probe as it traversed I-66 eastbound inside of I-495 at three times on Tuesday April 19th, 2011 plotted against the speeds reported by the detectors along that stretch of roadway (804, 822, 808, and 817) at that those same times.

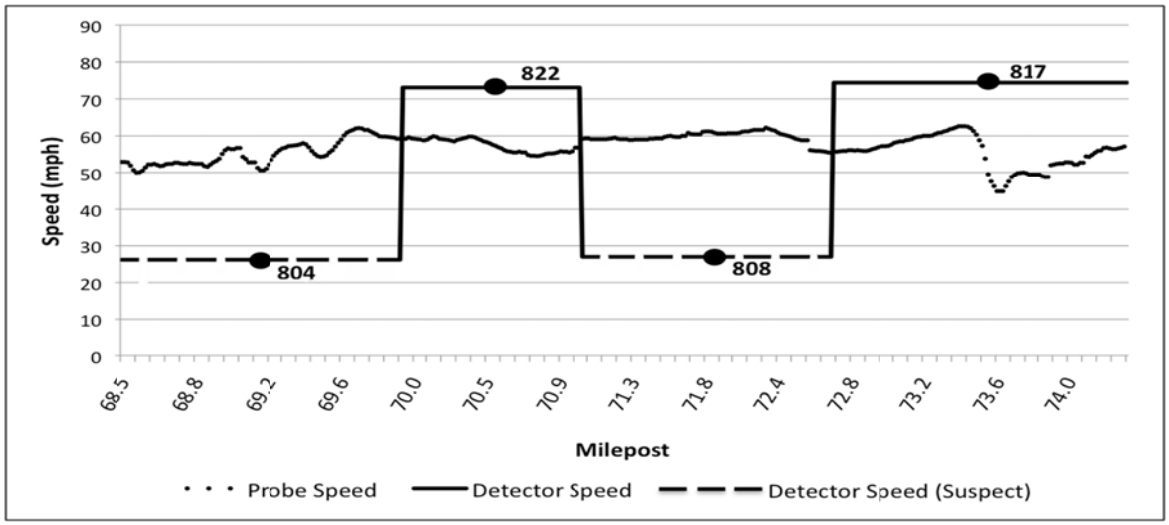

Exhibit C3-20: Segment A > B, Run 1 (I-66 Eastbound - 3:40 PM on Tuesday, April 19th)

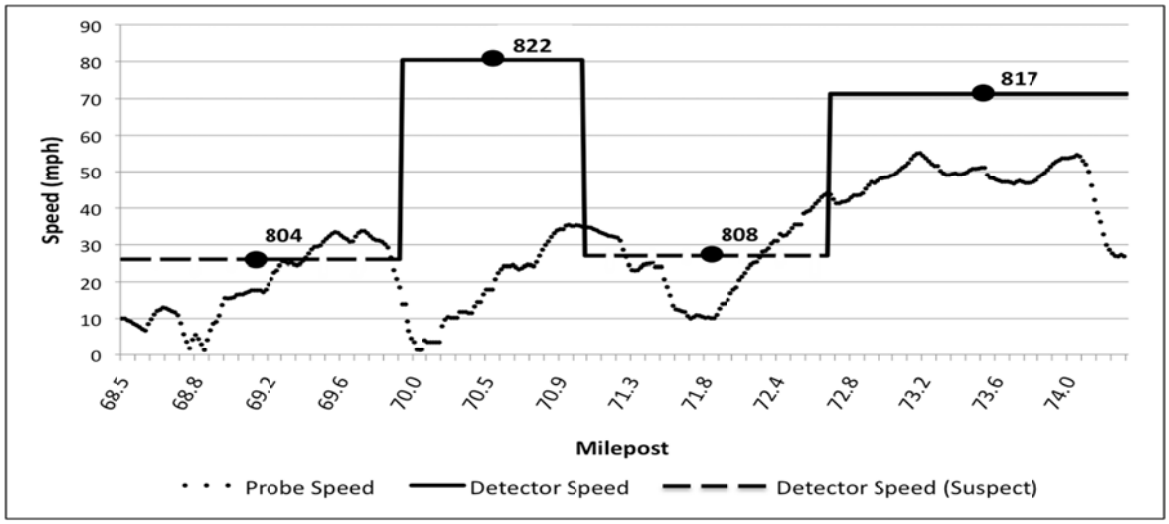

Exhibit C3-21: Segment A > B, Run 2 (I-66 Eastbound - 5:23 PM on Tuesday, April 19th)

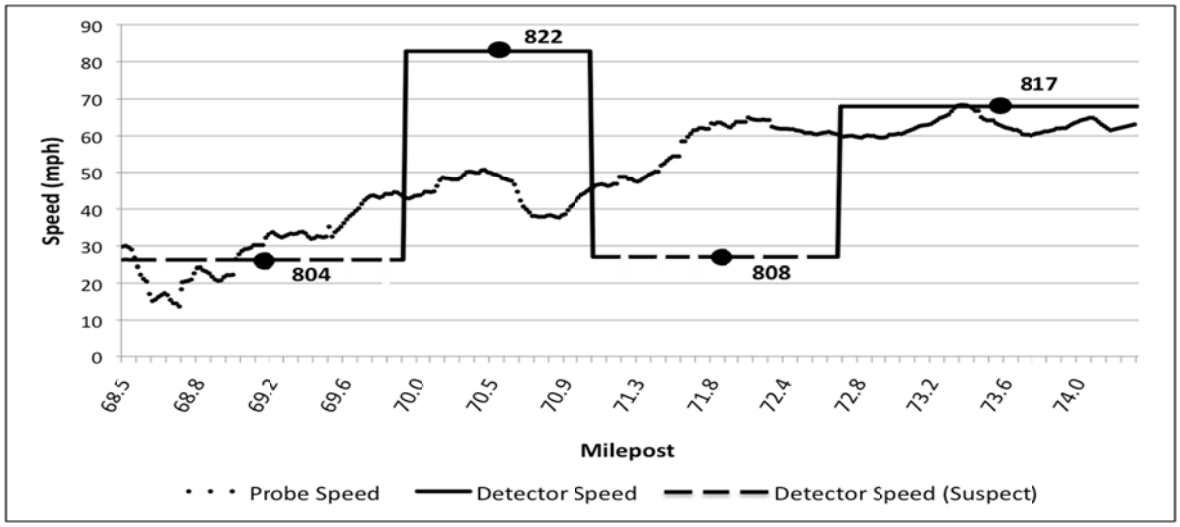

Exhibit C3-22: Segment A > B, Run 3 (I-66 Eastbound - 6:18 PM on Tuesday, April 19th)

Comparison of the probe speeds with the sensor-based speeds suggests the following:

- Sensor 804 (Milepost 68.5 – 70) – Data generated by this sensor are not consistent with the probe data collected along this roadway segment. The most likely explanation is data quality issues with the sensor. The speed reported by this sensor for most of the day is about 27mph.
- Sensor 822 (Milepost 70 – 71.05) – All the data (100%) for this sensor were imputed. The imputed data suggest a sustained free-flow speed which is clearly inaccurate based on the speeds observed by the probe vehicle.
- Sensor 808 (Milepost 71.05 – 72.7) – Data generated by this sensor are not consistent with the probe data. Again the explanation is likely to be data quality issues with the sensor. The speed reported by this sensor is about 28mph for most of the day.
- Sensor 817 (Milepost 72.7 – 74.3) – This is the one sensor which appears to be providing reliable speed data for the time periods during which the probe runs were conducted. Even so, the probe vehicle speeds are lower, and significantly so for probe runs 1 and 2.

Data Analysis Along I-66 Inside of I-495 (Westbound)

Exhibit C3-23, Exhibit C3-24, and Exhibit C3-25 show plots of the instantaneous speeds recorded by the vehicle probe as it traversed I-66 westbound inside of I-495 at three times on Tuesday April 19th, 2011 plotted against the speeds reported by the detectors along that stretch of roadway (819, 1422, and 806) at that those same times.

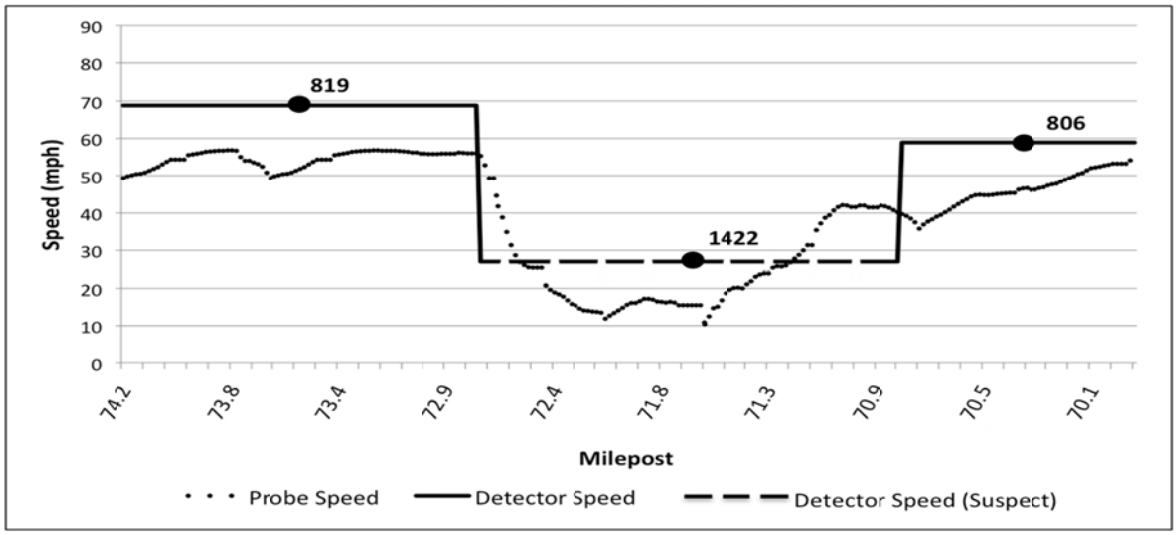

Exhibit C3-23: Segment C > D, Run 4 (I-66 Westbound - 3:27 PM on Tuesday, April 19th)

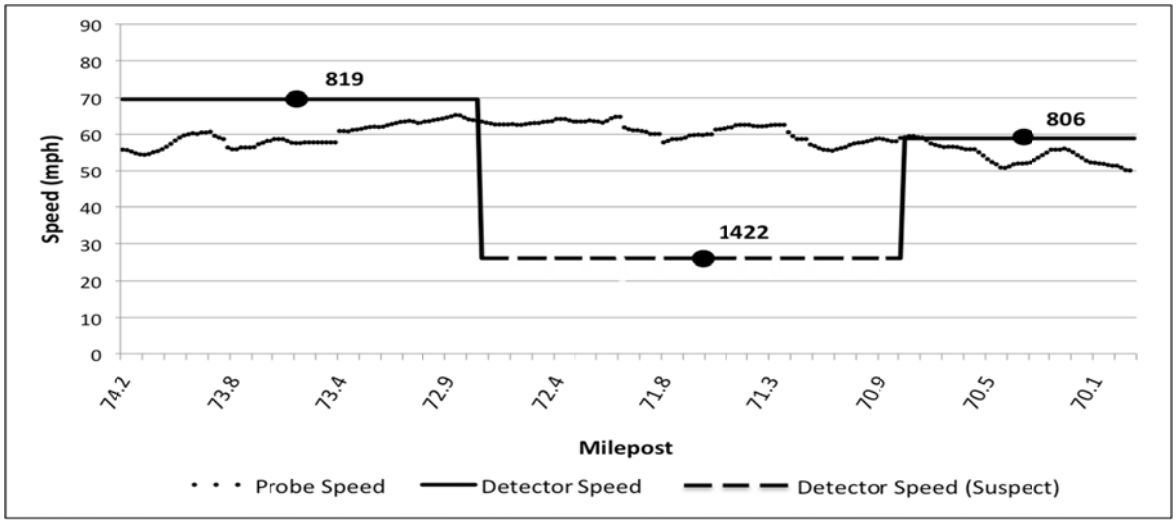

Exhibit C3-24: Segment C > D, Run 5 (I-66 Westbound - 4:05 PM on Tuesday, April 19th)

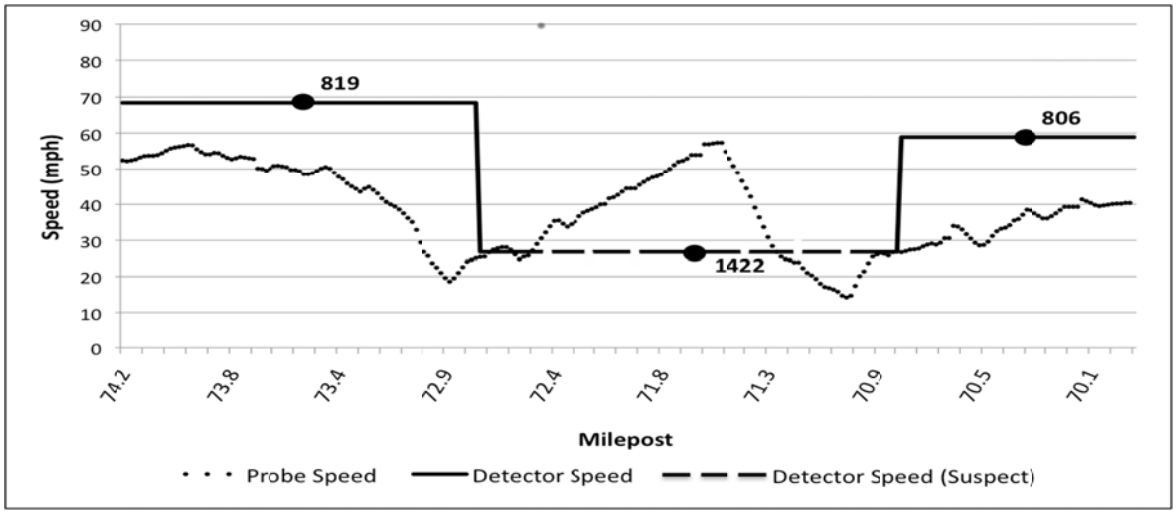

Exhibit C3-25: Segment C > D, Run 6 (I-66 Westbound - 6:38 PM on Tuesday, April 19th)

Comparison of these probe data with the sensor-based speeds suggests the following:

- Sensor 819 (Milepost 74.2 - 72.7) – This sensor was reporting fairly reliable speeds for the time periods during which the probe runs were conducted. Even so, the probe speeds are lower than those reported by the sensor, especially during the latter portion of probe run #3, during which significant congestion was encountered.
- Sensor 1422 (Milepost 72.7 – 70.8) – Data generated by this sensor was not consistent with the probe data due to data quality issues with the sensor. The speed reported by this sensor was about 28mph for most of the day.
- Sensor 806 (Milepost 70.8 - 69.9) – All of the data (100%) for this sensor was imputed (estimated). No field observations were generated by the sensor during any of the probe runs. Imputed data for this section of roadway indicates near free-flow speeds which were demonstrated to be inaccurate by the probe vehicle.

Data Analysis Along I-66 Outside of I-495 (Eastbound)

Exhibit C3-26, Exhibit C3-27, and Exhibit C3-28 show plots of the instantaneous speeds recorded by the vehicle probe as it traversed I-66 eastbound outside of I-495 at three times on Wednesday April 20th, 2011 plotted against the speeds reported by the detectors along that stretch of roadway (1139, 1157, 1141, and 1142) at that those same times.

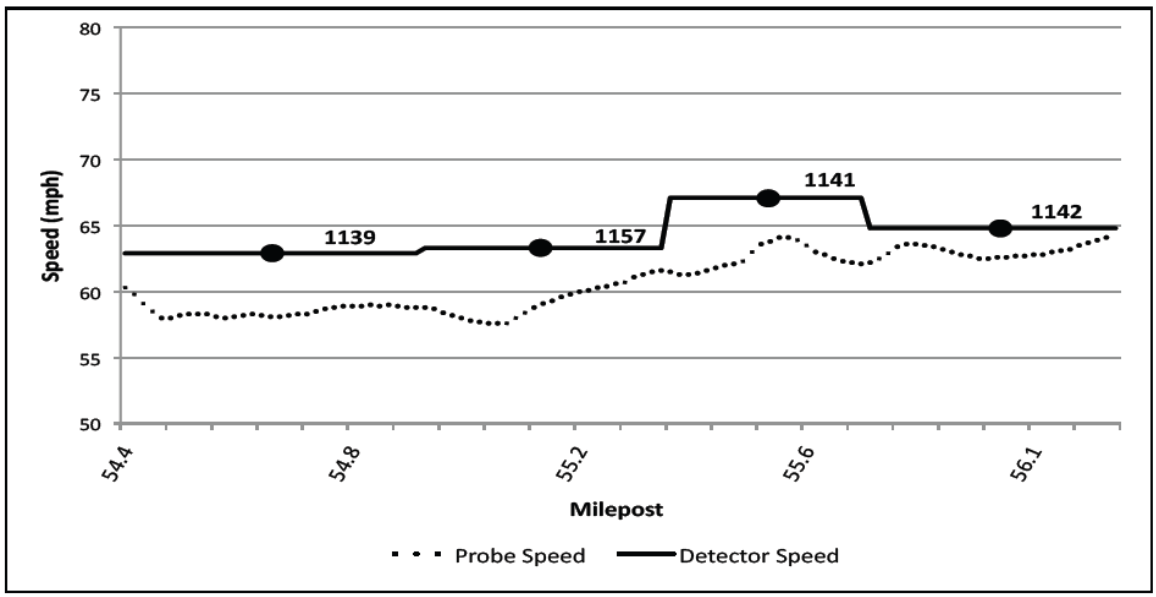

Exhibit C3-26: Segment E > F, Run 7 (I-66 Eastbound – 9:43 AM on April 20, 2011)

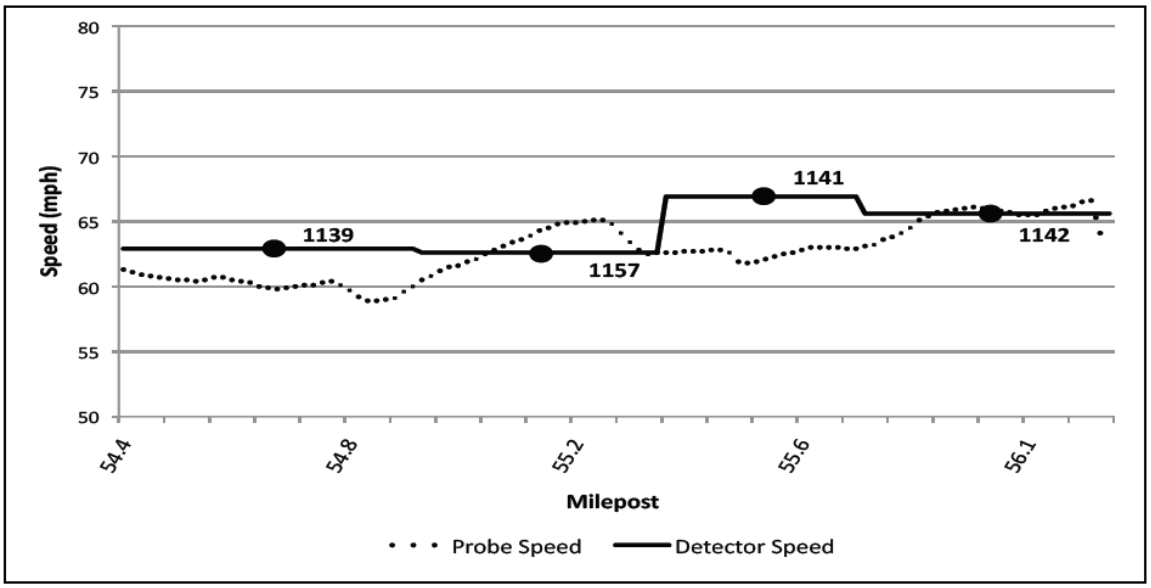

Exhibit C3-27: Segment E > F, Run 8 (I-66 Eastbound – 10:20 AM on April 20, 2011)

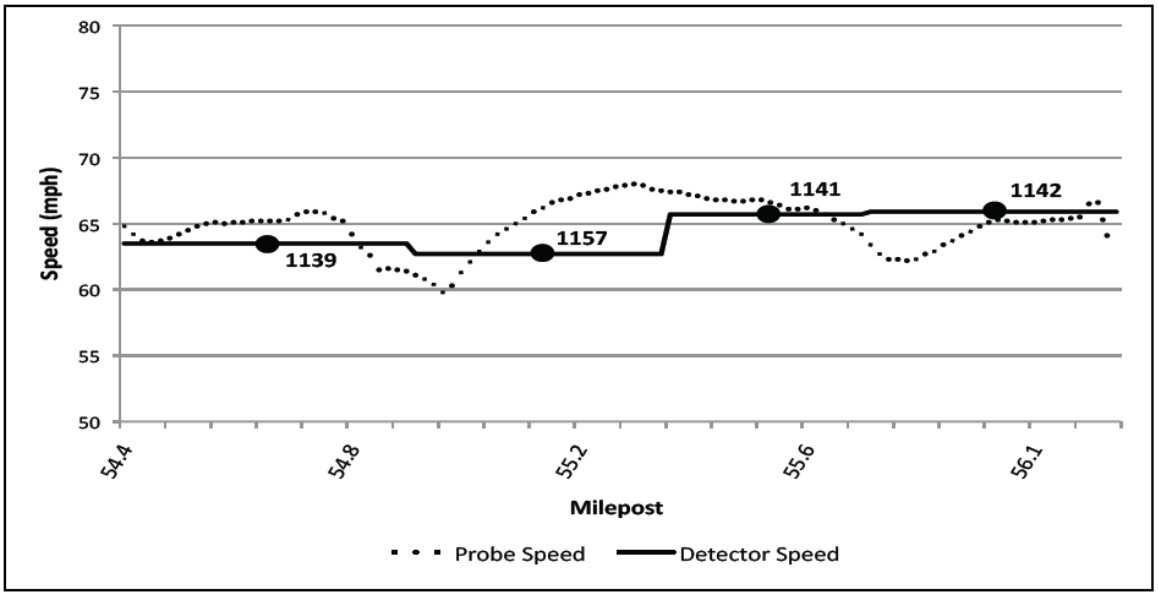

Exhibit C3-28: Segment E > F, Run 9 (I-66 Eastbound – 10:36 AM on April 20, 2011)

Comparison of these probe data with the sensor-based speeds suggests the following:

- Sensor 1139 (Milepost 54.4 – 54.9) – Only 15% of the speeds reported by this sensor were actually observed. Consequently, although those speeds are reasonably consistent with the conditions observed by the probe vehicle, it is unclear whether this sensor would provide accurate data under other conditions.
- Sensor 1157 (Milepost 54.9 – 55.4) – All of the speeds (100%) reported by this sensor were imputed. Those imputed data suggested sustained free-flow speeds, which is consistent with the conditions encountered by the probe vehicle.
- Sensor 1141 (Milepost 55.4 – 55.8) – All of the speeds (100%) reported by this sensor were imputed. Those imputed speeds suggest sustained free-flow conditions, which is consistent with the conditions encountered by the probe vehicle (although the sensor shows slightly higher speeds during 2 of the 3 probe runs).
- Sensor 1142 (Milepost 55.8 – 56.3) - As with sensor 1139, only 15% of the speeds reported by this sensor were actually observed. As such, although the sensor suggests the conditions encountered by the probe vehicle, it is unclear whether this sensor would provide accurate data under other conditions.

Data Analysis Along I-66 Outside of I-495 (Westbound)

Exhibit C3-29, Exhibit C3-30, and Exhibit C3-31 show plots of the instantaneous speeds recorded by the vehicle probe as it traversed I-66 westbound outside of I-495 at three times on Wednesday April 20th, 2011 plotted against the speeds reported by the detectors along that stretch of roadway (1143, 1156, 1158, and 1140) at that those same times.

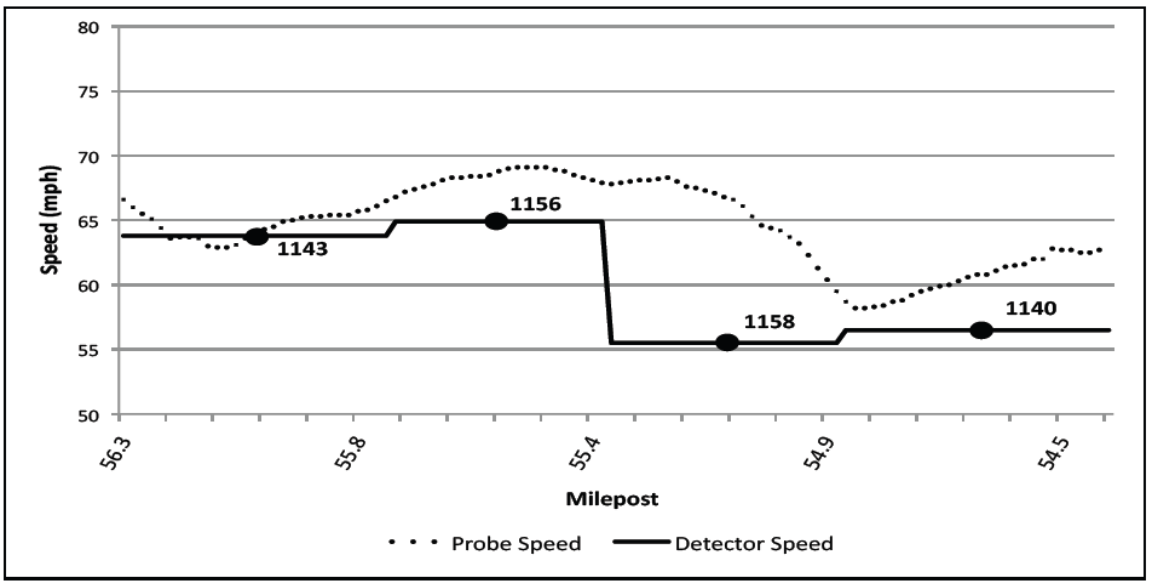

Exhibit C3-29: Segment G > H, Run 10 (I-66 Westbound – 9:34 AM on April 20, 2011)

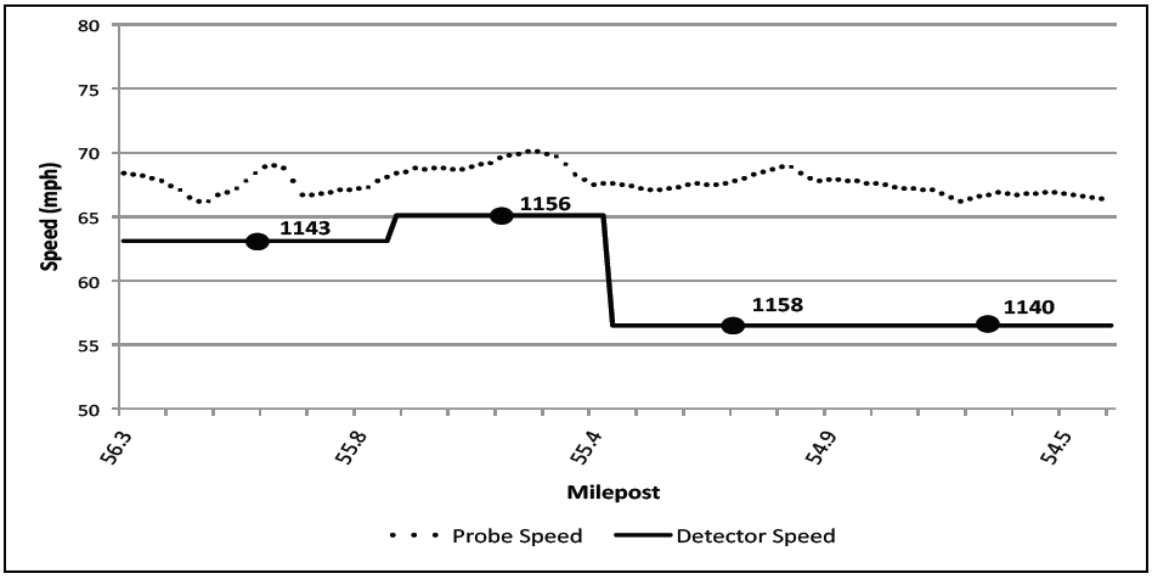

Exhibit C3-30: Segment G > H, Run 11 (I-66 Westbound – 9:53 AM on April 20, 2011)

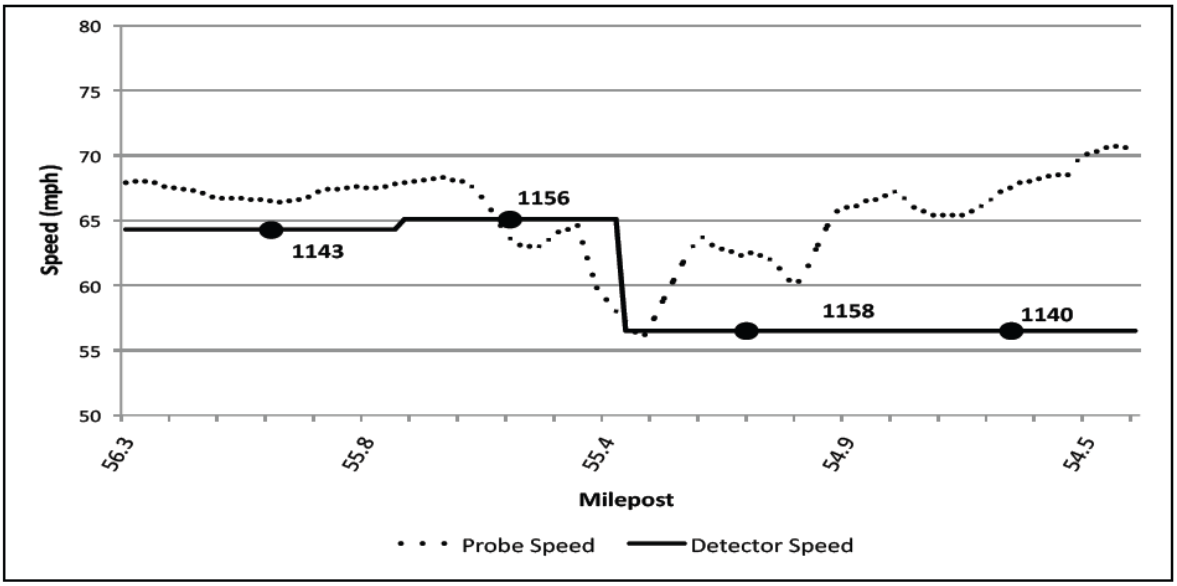

Exhibit C3-31: Segment G > H, Run 12 (I-66 Westbound – 10:27 AM on April 20, 2011)

Comparison of these probe data with the sensor-based speeds suggests the following:

- Sensor 1143 (Milepost 56.3 – 55.7) – Only 15% of the speeds generated by this sensor were actually observed. Consequently, although the speeds reported by this sensor are close to those observed by the probe, it is unclear whether this sensor would provide accurate data under other conditions.
- Sensor 1156 (Milepost 55.7 – 55.3) – All (100%) of the speeds reported by this sensor were imputed. The Imputed speeds suggest sustained free-flow speeds along this portion of the freeway mainline, which is consistent with the conditions encountered by the probe vehicle.
- Sensor 1158 (Milepost 55.3 – 54.85) – All (100%) of the speeds for this sensor were imputed. Those imputed speeds suggest sustained near free-flow speeds, which is somewhat lower than speed data generated by the probe vehicle.
- Sensor 1140 (Milepost 54.85 – 54.4) - As with sensor 1143, only 15% of the data generated by this sensor were observed. This lack of observed data helps to explain the lower speeds generated by this sensor versus those reported by the probe vehicle.

Comparison of Travel Times - Probe (Measured) vs. Sensor (Estimated)

Based on the speed data from the probe vehicle runs and speed estimates provided by the sensors, segment travel times were generated for each of the 12 probe runs described above. Two approaches were used to calculate roadway travel times based on the sensor data.

- Approach 1 –*ALL* of the speed data received by the team from the sensors was used regardless of whether the data was good, imputed, or suspect.
- Approach 2 –data from nearby sensors were used in place of the data from the sensors that were flagged (manually) as likely generating suspect data – based on the reporting of very low speeds over significant periods of time:
  - Runs 1, 2, and 3 – substituted data for sensors #804 and #808

- Runs 4, 5, and 6 – substituted data for sensor #1422
- Runs 7, 8, and 9 – no substitution of data
- Runs 10, 11, and 12 – no substitution of data

As no substitution of sensor data occurred for runs 7 – 12, Approach 2 was not employed as part of the travel time estimation process along those segments of roadway.

Travel Times for Runs 1, 2, and 3 (A > B) – April 19th

Start Time	Road	Start MP	End MP	Probe Vehicle Travel Time (Measured)	VDOT Sensor Travel Time (Estimated) Approach 1	VDOT Sensor Travel Time (Estimated) Approach2	Percent Error App. 1	Percent Error App. 2
3:40 PM	I-66 EB	68.5	74.3	6.3 minutes	7.0 minutes	4.7 minutes	+ 11%	- 25%
5:23 PM	I-66 EB	68.5	74.3	10.1 minutes	7.0 minutes	4.6 minutes	- 31%	- 54%
6:18 PM	I-66 EB	68.5	74.3	7.4 minutes	7.1 minutes	4.6 minutes	- 4 %	- 37%

Travel Times for Runs 4, 5, and 6 (C > D) – April 19th

Start Time	Road	Start MP	End MP	Probe Vehicle Travel Time (Measured)	VDOT Sensor Travel Time (Estimated) Approach 1	VDOT Sensor Travel Time (Estimated) Approach2	Percent Error App. 1	Percent Error App. 2
3:27 PM	I-66 WB	74.2	69.9	7.2 minutes	6.3 minutes	4.0 minutes	- 12%	- 44%
4:05 PM	I-66 WB	74.2	69.9	4.6 minutes	6.3 minutes	4.0 minutes	+ 37%	- 13%
6:38 PM	I-66 WB	74.2	69.9	12.2 minutes	6.1 minutes	4.1 minutes	- 50%	- 66%

Travel Times for Runs 7, 8, and 9 (E > F) – April 20th

Start Time	Road	Start MP	End MP	Probe Vehicle Travel Time (Measured)	VDOT Sensor Travel Time (Estimated) Approach 1	VDOT Sensor Travel Time (Estimated) Approach2	Percent Error App. 1	Percent Error App. 2
9:43 AM	I-66 EB	54.4	56.3	1.8 minutes	1.7 minutes	N/A	- 6%	N/A
10:20 AM	I-66 EB	54.4	56.3	1.8 minutes	1.7 minutes	N/A	- 6%	N/A
10:36 AM	I-66 EB	54.4	56.3	1.8 minutes	1.7 minutes	N/A	- 6%	N/A

Travel Times for Runs 10, 11, and 12 (G > H) – April 20th

Start Time	Road	Start MP	End MP	Probe Vehicle Travel Time (Measured)	VDOT Sensor Travel Time (Estimated) Approach 1	VDOT Sensor Travel Time (Estimated) Approach2	Percent Error App. 1	Percent Error App. 2
9:34 AM	I-66 WB	56.3	54.4	1.7 minutes	1.8 minutes	N/A	+ 6%	N/A
9:53 AM	I-66 WB	56.3	54.4	1.7 minutes	1.8 minutes	N/A	+ 6%	N/A
10:27 AM	I-66 WB	56.3	54.4	1.7 minutes	1.8 minutes	N/A	+ 6%	N/A

Travel times collected during the first day of probe data collection (April 19) differed significantly from the estimated travel times generated by the sensor data (using either Approach 1 or 2). For example, for runs 1, 2, and 3 there was an overall absolute average error of 15% for Approach 1 and 39% for Approach 2. Although this might result in a perception that the sensor data along this segment are useful for calculating travel times, it must be remembered that two of the sensors generated suspect speed data – in this case, very low freeway speeds. Incorporating these speeds into the travel time estimation appears to have offset the higher roadway speeds generated by the other two roadway sensors, speeds that were generally much higher than those reported by the probe vehicle. Consequently, incorporation of the likely erroneous slow speeds resulted in travel times closer to those experienced by the probe vehicle – an unintended consequence of the use of this data. Moreover, the nearly identical travel time estimates generated using both approaches over the course of several hours speaks to the likely impact of the considerable amount of data imputation which occurred. The steadiness of these travel time estimates is not ideal for computing reliability, which relies on the ability of the system to detect variability in traffic conditions over time. Reviewing the content of the histogram found in Exhibit C3-32(below), which provides a breakdown of PM Peak Period (3 – 7 pm) travel times along the roadway segment used as part of runs 1, 2, and 3 (A > B) for a two month period (March 15th – May 15th) demonstrates a fairly low amount of travel time variability over the 2000+ 5-minute data collection periods for which data was collected.

Travel times collected during the second day of probe runs conform much more closely to the estimates from the sensors, with an average error of 6% in each direction of travel. However, it must be pointed out that nearly all of these data were imputed (only 15% observed data provided by 4 of the 8 sensors from which data was collected). As a result, it is highly unlikely that these sensors would provide accurate travel times under most congested conditions. The full extent of this problem is made clear by the histogram contained in Exhibit C3-33 (below), which demonstrates that over the course of two months, a total of only 44 (of 2156 total) 5-minute time slices along segment E > F (runs 7, 8, and 9) were reported as having travel times in excess of 2 minutes during the AM peak period. It should be noted than a nearly identical travel time distribution exists for westbound travel times along this segment of I-66 during the AM peak period.

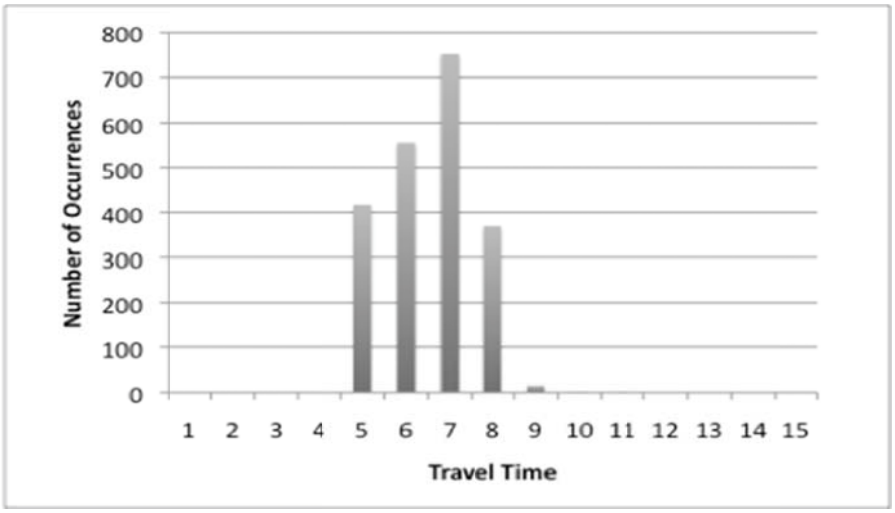

Exhibit C3-32: I-66 EB PM Peak Travel Times between MP 68.5-74.3 (3/15/11 - 5/15/11)

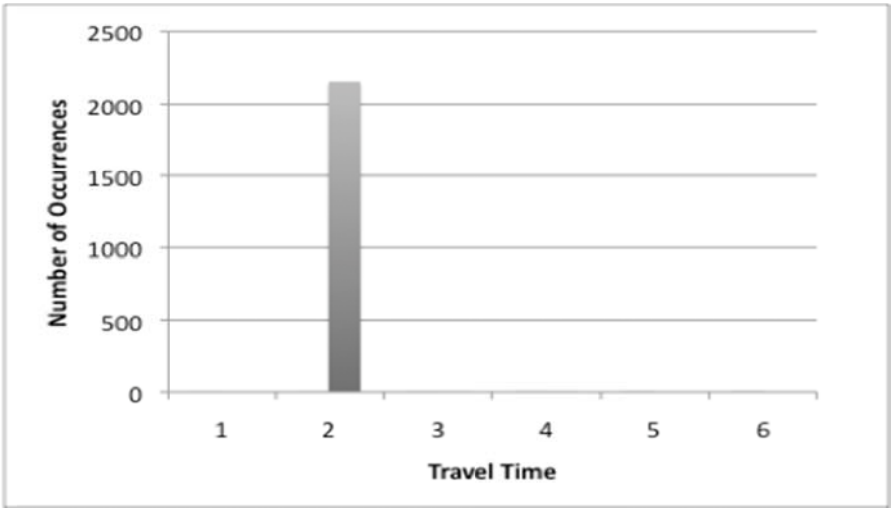

Exhibit C3-33: I-66 EB AM Peak Travel Times between MP 54.4 - 56.3 (3/15/11 - 5/15/11)

LESSONS LEARNED

Overview

The team selected Northern Virginia as a case study site because it provided an opportunity to integrate a reliability monitoring system into a pre-existing, extensive data collection network. The data collected on NOVA roadways is already passed to a number of external systems, including RITIS at the University of Maryland, the ADMS at the University of Virginia, and the statewide 511 system. Configuring PeMS to receive NOVA data helped define the requirements for complex traffic systems integration and illustrate what agencies can do to facilitate the process of implementing reliability monitoring.

Systems Integration

The process of fully integrating the NOVA data with PeMS took several weeks. While this amount of effort is standard when integrating archived data user systems with traffic data collection systems, there are a number of steps that agencies can take to make this integration go more smoothly and quickly.

For one, it is important that the implementation and maintenance of a traffic data collection system be carried out with a broad audience in mind. Efforts such as the Federal Government's 2009 "Open Government Initiative" underscore the value of providing public access to government data. Often, increasing access to data outside of an organization can help to further agency goals; for example, providing data to mobile application developers can help agencies distribute information in a way that increases the efficiency of the transportation network. It will also help the agency support contractor's efforts to implement procured systems, such as travel time reliability monitoring systems.

One of the ways that agencies can facilitate the distribution of data from their data collection system is by establishing one or more data feeds. As discussed in the first chapter, different parties will want to acquire data processed to different levels, depending on the intended use. For example, a mobile application developer may only be interested in heavily processed data, such as route-level travel times. A third-party data aggregator may be interested in obtaining speeds computed from loop detectors, to be fused with other travel time data sources. A traffic engineering firm may prefer raw detector flow and occupancy data that they can quality-check using their own established methods and use to calculate performance measures. Since maintaining multiple data feeds can be a challenge, if agencies want to provide a feed of processed data, it will save resources in the long run to document the processing steps performed on the data. This will allow implementers of external systems to evaluate them and undo them, if needed.

Aside from the processing documentation, maintaining clear documentation on the format of data files and units of data will greatly facilitate the use of data outside of the agency. Additionally, documentation on the path of data from a detector through the agency's internal systems can be of value to contractors and other external data users. Clearly explaining this information in a text file minimizes the back-and-forth communication between agency staff and contractors and prevents inaccurate assumptions from being made.

Methodological Advancement

From a methodological standpoint, this case study focused on implementing a multi-state travel time reliability model developed by the SHRP2 L10 project. The original research developed this model on AVI travel time measurements in San Antonio, as well as travel times generated by a micro-simulation traffic model on a section of I-66 in Northern Virginia. This effort extended the research by applying it to point speeds generated by multiple loop detectors along a freeway segment.

The methodological findings of this case study are that multi-state normal distribution models can approximate travel time distributions generated from loop detectors better than normal or log-normal distributions. During the peak hours on a congested facility, three states are generally sufficient to balance a good model (distribution) fit with the need to generate information that can be easily communicated to interested parties. During off-peak hours, two states typically provide a reasonable model (distribution) fit. The outputs of this method can

inform travelers of the percent change that they will encounter moderate or severe congestions and, if they do, what their expected and 95th percentile travel times will be.

Probe Data Comparison

Most public agency managed data processing systems currently rely on fixed sensor infrastructure to support the calculation of roadway travel times and subsequent generation of travel time reliability metrics. Although this state of affairs may change over time as more private sector sources of data become available, this will not happen overnight. To that end, agencies need to consider how to make the best use of the data currently available to them. As part of this use case, we have examined the data available from a network of fixed infrastructure sensors (a combination of single loops and radar-based sensors) going through the process of being modernized by the Virginia DOT. The team's analysis of the data available from these sensors has yielded a number of findings of potential interest to a wide variety of agencies, particularly those facing maintenance and calibration issues associated with older sensor systems, as well as those agencies with more sparsely spaced spot sensors. Overall, we found that there were five (5) primary factors that accounted for differences between the probe vehicle data and speed / estimated travel times generated based on VDOT sensor data; these factors are detailed below.

Likely one of the most significant, and at the same time most difficult to measure, impacts on sensor based speeds is associated with research that suggests that fixed roadway sensors may not always accurately measure very low speeds during highly congested conditions. Although impossible to definitively evaluate here, it is something that should be taken into consideration as part of all such analysis.

As the Virginia DOT is in the process of modernizing its sensor network in NOVA, the vast majority of sensors are not fully calibrated and/or fully configured so as to properly communicate with back-office data analysis systems. This resulted in the types of data quality issues discussed earlier in the case study. This issue makes clear the need for public agencies to conduct regular sensor maintenance programs in order to ensure that their detection networks are generating the most accurate data possible.

Beyond any issues that spot sensors may have accurately assessing low-speed, stop-and-go traffic conditions, another issue that sensor users must contend with is the problem associated with extrapolating speeds (and subsequently travel times) for a segment of roadway based solely on conditions within the sensor's field of detection. As such, all speeds and travel times for a segment are based on the assumption that conditions along the segment are identical to those experienced within the sensor's field of view. As a result, it is likely that any data generated by spot sensors will fail to detect congestion, incidents, etc., that occur outside of the sensor's immediate vicinity, with the impact becoming more pronounced the longer the segment.

Related to the problem associated with extrapolating spot sensor data to cover entire segments of roadway is that related to the need to impute data from adjacent sensors or segments of roadway to fill in gaps in sensor coverage. Although not necessarily an enormous problem in cases where data for a single lane of travel is "filled in" based on conditions experienced by adjacent sensor stations, the types of imputation required as part of this case study resulted in speeds being generated for segments of roadway based largely on historical data for a sensor or macroscopic speed and flow data for a section of the roadway network. Although a necessity for computing speed and estimated travel time for the given segment, use of this replacement data further aggravated the data-related issues described above.

Another dynamic impacting the comparison of sensor data with probe vehicle data stems from a basic difference between these data sets:

- Sensor Data – Represents five minute, average conditions across all lanes of travel observed at the sensor location.
- Probe Data – Represents the movement of a single vehicle through one lane of travel across the segment being evaluated.

These differences have the potential to result in significant differences in speed/estimated travel time between the two data sources if one lane of travel experiences significant congestion, while the other(s) do not. This is especially true in cases where the probe vehicle is slowed by congestion outside of a sensor's detection zone, while other lanes of travel are moving at higher, less congested (or even free-flow) rates of speed.

Each of the factors described above almost certainly had some degree of impact on the differences between the probe vehicle speeds we collected and speed / estimated travel times generated based on VDOT sensor data. Moreover, with the exception of the final factor (basic differences between probe and sensor data sets), each of these has the potential to impact the quality of data collected by spot sensor-based fixed data collection infrastructure. As such, public agency staff should take each of these into consideration when making decisions concerning both the deployment of new data collection infrastructure, as well as the maintenance and/or expansion of existing systems.

CHAPTER 4

SACRAMENTO–LAKE TAHOE, CALIFORNIA

The *monitoring system* description section details the reasons for selecting the Sacramento–Lake Tahoe region in northern California as a case study and provides an overview of the region. It briefly summarizes agency monitoring practices, discusses the existing sensor network, and describes the software system that the team used to analyze the use cases. Specifically, it describes the steps and tasks that the research team completed in order to transfer data from the data collection systems into a travel time reliability monitoring system.

The section concerning *methodological experiments* describes the manner in which different types of filtering techniques might be applied at different stages of the analytical process to further refine the travel times estimates generated from Bluetooth-based datasets.

Use cases are less theoretical, and more site specific. The first two use cases assess the impact of detector network configuration on the data ultimately available for use by travel time reliability monitoring systems. The third use case attempts to quantify the impact of adverse weather and demand–related conditions on travel time reliability using data derived from the Bluetooth and electronic toll collection-based systems deployed in rural areas as part of this case study.

The section on *privacy considerations* addresses the challenges associated with collecting data using toll tag and Bluetooth-based technologies in a manner that respects the privacy of the individuals from whom the data is being collected.

Lessons learned summarizes the lessons learned during this case study, with regard to all aspects of travel time reliability monitoring: sensor systems, software systems, calculation methodology, and use. These lessons learned will be integrated into the final guidebook for practitioners.

MONITORING SYSTEM

Site Overview

The team selected the Lake Tahoe region located in Caltrans District 3 in order to provide an example of a rural transportation network with fairly sparse data collection infrastructure. Caltrans District 3 encompasses the Sacramento Valley and Northern Sierra regions of California. Its only metropolitan area is Sacramento. The District DOT is responsible for maintaining and operating 1,410 centerline miles and 4,700 lane-miles of freeway in eleven counties. District 3 includes urban, suburban, and rural areas, including areas near Lake Tahoe where weather is a serious travel time reliability concern and there is heavy recreational traffic. The District also contains 64 lane-miles of HOV lanes, with more than 140 more lane-miles proposed, all within the greater Sacramento region. Two major interstates pass through the District, Interstate-80, which travels from east to west, and Interstate 5, which travels from north to south. Other major freeway facilities include US-50, which connects Sacramento and South Lake Tahoe, and SR-99.

Built in 2000, the District 3 Regional Traffic Management Center (RTMC) is located in Rancho Cordova, 15 miles east of Sacramento. The RTMC serves as the focal point for traffic information within District 3. RTMC staff are responsible for managing (1):

- Regional network of sensors, cameras, CMS, HARs, and RWIS
- Delivery of traveler information
- Dispatch of other Caltrans resources

As mentioned above, weather-related conditions contribute to serious travel time reliability concerns in District 3, including (1):

- Fog/Visibility – The region is prone to thick 'tule' fog during periods after heavy rain;
- High Winds - Several bridges in the District are exposed to high winds;
- Frost/Ice – Freezing can occur on longer viaduct sections during cold weather; and,
- Snow in Sierras - High winds combined with snow accumulation create white out conditions over mountain roadways.

Caltrans and its regional partners are pursuing the creation of Corridor System Management Plans (CSMPs) for the most heavily congested transportation corridors in the region, "aimed at increasing transportation options, reducing congestion, and improving travel times. A CSMP is a comprehensive, integrated management plan for increasing transportation options, decreasing congestion, and improving travel times in a transportation corridor. A CSMP includes all travel modes in a defined corridor – highways and freeways, parallel and connecting roadways, public transit (bus, bus rapid transit, light rail, intercity rail) and bikeways, along with intelligent transportation technologies. CSMP success is based on the premise of managing a selected set of transportation components within a designated corridor as a system rather than as independent units. Each CSMP identifies current management strategies, existing travel conditions and mobility challenges, corridor performance management, planning management strategies, and capital improvements. In District 3, six CSMPs have been developed along I-80, I-5/SR-99, US-50, SR 99 North, SR 49, and SR 65." (2)

Sensors

Caltrans District 3 currently only collects traffic data along freeway facilities. It operates a total of 2,251 point detectors (either radar or loop detectors) located in over 1,000 roadway locations across the District. Point detection infrastructure in the mountainous regions of the District is sparser, with detectors often miles apart. To supplement the point detection network in rural portions of the Sierra Nevada Mountains near Lake Tahoe, the District has installed electronic toll collection (ETC) readers on I-80 and Bluetooth-based data collection readers along I-5 and US-50 (see Exhibit C4-1). These readers register the movement of vehicles equipped with FasTrak tags (Northern California's ETC system) and Bluetooth-based devices (e.g., Smart Phones) for the purpose of generating roadway travel times. Table C4-1 provides details about the ETC readers deployed in this case study, and Table C4-2 shows the Bluetooth readers deployed in this case study.

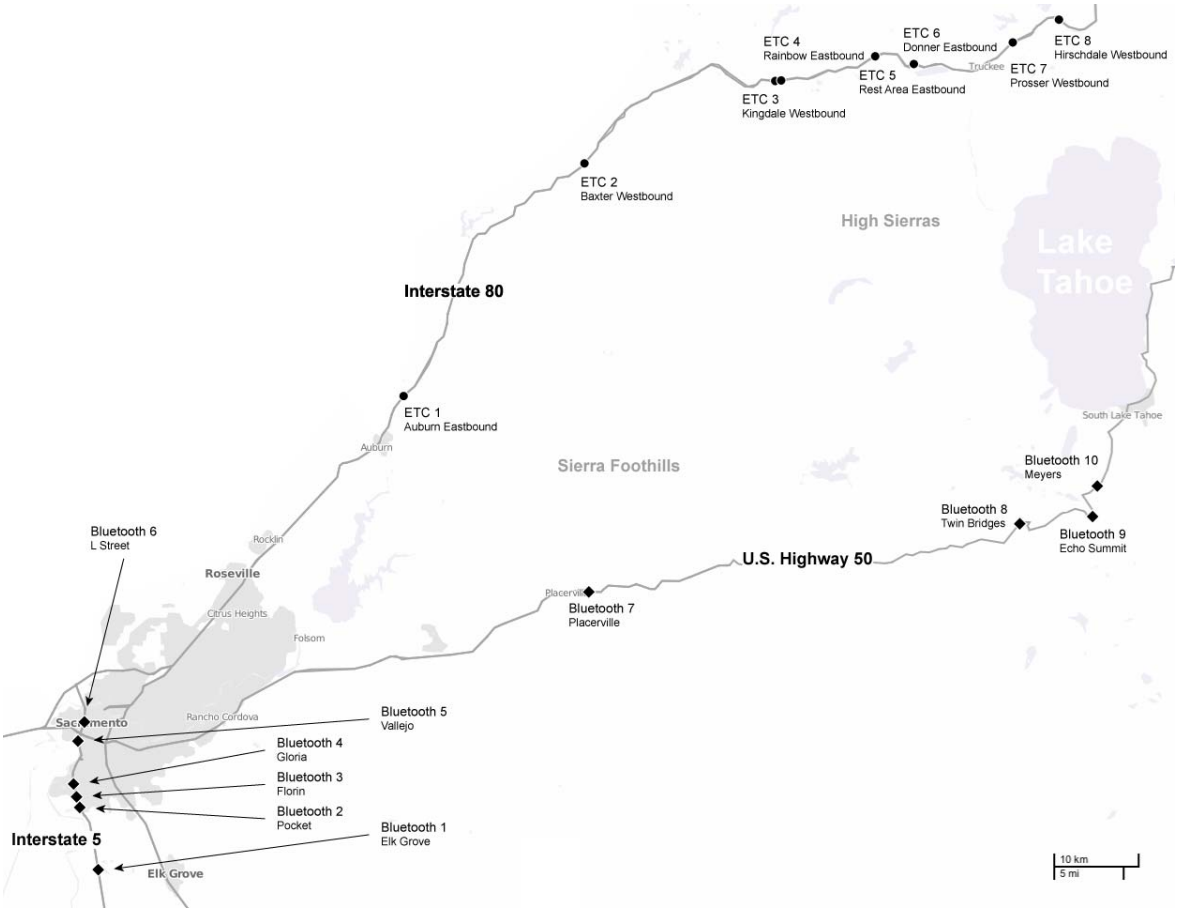

Exhibit C4-1: Map of ETC and Bluetooth Readers Deployed in Caltrans District 3

Both ETC and Bluetooth-based data collection technologies utilize vehicle identification technologies to record the presence of vehicles as they pass instrumented points along a roadway. Field controllers typically record location, time, and vehicle identification information for each vehicle to support the calculation of travel times. By knowing the length of the road segment between two instrumented points, and the starting and ending times at which travel between those points took place, the travel time for that section of roadway can be determined.

Table C4-1: Breakdown of Deployed ETC Readers

Exhibit C4-1 ID	Roadway / Direction of Travel	ETC Reader ID	Nearest Crossroad	Postmile
ETC 1	I-80 E	42003	Auburn	123.1
ETC 2	I-80 W	42035	Baxter	148.5
ETC 3	I-80 W	42041	Kingale	168.0
ETC 4	I-80 E	42036	Rainbow	168.1
ETC 5	I-80 E	42042	Rest Area	176.2
ETC 6	I-80 E	42044	Donner Lake	179.9
ETC 7	I-80 W	42006	Prosser Village	189.0
ETC 8	I-80 W	42015	Hirschdale	193.4

Table C4-2: Breakdown of Deployed Bluetooth Readers

Exhibit C4-1 ID	Roadway / Direction of Travel	Bluetooth Reader ID	Nearest Crossroad	Postmile
Bluetooth 1	I-5 N	1005	Elk Grove	506.4
Bluetooth 2	I-5 N	1011	Pocket	511.5
Bluetooth 3	I-5 S	2101	Florin	512.4
Bluetooth 4	I-5 S	2009	Gloria	513.5
Bluetooth 5	I-5 N	1039	Vallejo	517.2
Bluetooth 6	I-5 N	1004	L St.	518.9
Bluetooth 7	US-50 E	1054	Placerville	48.4
Bluetooth 8	US-50 E	2055	Twin Bridges	87.1
Bluetooth 9	US-50 W	2058	Echo Summit	94.9
Bluetooth 10	US-50 E	2056	Meyers	98.7

ETC-based Data Collection in District 3

The ETC-based data collection infrastructure deployed along I-80 consists of eight (8) Fastrak toll tag reader stations, installed and operated by Caltrans District 3. The readers were initially installed to provide the Bay Area's 511 system with travel times to Lake Tahoe, but have not yet been used for that purpose.

According to Caltrans, each reader is either mounted on an overhead Changeable Message Sign (CMS) or other fixed overhead sign. Each reader station consists of a cabinet mounted to the sign pole, which is connected to antennas mounted on the edge of the sign closest to the roadway; directed such that they monitor traffic in each lane of travel. All of the readers are deployed at roadway sections that have two lanes of travel in each direction, with the exception of one location, where there are three lanes of travel in each direction. ETC transponders passing these readers are each encoded with a unique identification number. Data from these transponders is collected via Dedicated Short-Range Communication (DSRC) radio by the reads and assigned time/date stamps, as well as an antenna identification stamp for use in calculating travel time.

Bluetooth-based Data Collection in District 3

This case study also leverages data from Bluetooth readers (BTRs) deployed on I-5 in Sacramento and along US-50 between Placerville and Lake Tahoe; these BTRs were installed by Caltrans' research division.

From a travel time data collection standpoint, Bluetooth readers are typically placed on the side of a roadway, ideally at a vehicle windshield height or higher to minimize the obstructions between the reader and the in-vehicle Bluetooth-enabled devices. In Caltrans' case, each BTR was mounted inside an equipment cabinet strapped to poles along the freeway.

The BTRs deployed by Caltrans used the standard Bluetooth device inquiry algorithm, scanning all 32 available channels every 5.12 seconds (split into two 2.56 phases of 16 channels each). Each Bluetooth reader records the unique Media Access Control (MAC) address generated by every Bluetooth device it detects during each scan cycle for use in calculating travel time.

Data Management

The primary data management software system in the District 3 region is Caltrans' Performance Measurement System (PeMS). All Caltrans districts use PeMS for data archiving and performance measure reporting. PeMS integrates with a variety of other systems to obtain traffic, incident, and other types of data. It archives raw data, filters it for quality, computes performance measures, and reports them to users through the web at various levels of spatial and temporal granularity. It reports performance measures such as speed, delay, percentage of time spent in congestion, travel time, and travel time reliability. These performance measures can be obtained for specific freeways and routes, and are also aggregated up to higher spatial levels such as county, district, and state. These flexible reporting options are supported by the PeMS web interface, which allows users to select a date range over which to view data, as well as the days of the week and times of the day to be processed into performance metrics. Since PeMS has archived data for Caltrans dating back to 1999, it provides a rich and detailed source of both current travel times and historical reliability information.

PeMS integrates, archives, and reports on incident data collected from two different sources: the California Highway Patrol (CHP) and Caltrans. CHP reports current incidents in real-time on its website. PeMS obtains the text from the website, uses algorithms to parse the accompanying information, and inserts it into the PeMS database for display on a real-time map, as well as for archiving. Additionally, Caltrans maintains an incident database, called the Traffic Accident Surveillance and Analysis System (TASAS), which links to the highway database so that incidents and their locations can be analyzed. PeMS obtains and archives TASAS incident data via a batch process approximately once per year. Incident data contained in PeMS has been leveraged to validate use cases associated with how different sources of congestion impact travel time reliability.

PeMS also integrates data on freeway construction zones from the Caltrans Lane Closure System (LCS), which is used by the Caltrans districts to report all approved closures for the next seven days, plus all current closures, updated every 15 minutes. PeMS obtains this data in real-time from the LCS, displays it on a map, and lets users run reports on lane closures by freeway, county, district, or state. Lane closure data in PeMS was used in the validation of the use cases associated with how different sources of congestion impact travel time reliability.

Systems Integration

Data Acquisition Prior in Support of Travel Time Reliability Analysis

PeMS can calculate many different types of performance measures; and as such, the requirements for linking PeMS with an existing system depend on the features being used. The following bullet points describe the basic data that PeMS requires from the source system to support these functions:

- Metadata on the roadway linework of facilities being monitored
- Metadata on the detection infrastructure, including the types of data collected and the locations of equipment (configuration)
- Real-time traffic data in a constant format at a constant frequency (such as every 30-seconds or every minute)

Traffic data are generally unusable for travel time calculation purposes if not accompanied by a detailed description of the configuration of the system. Configuration

information provides the contextual and spatial information on the sensor network needed to make sense of the real-time data. Ideally, these two types of information should be transmitted separately (i.e., not in the same file or data feed). Roadway and equipment configuration information is more static than traffic data, as it only needs to be updated with changes to the roadway or the detection infrastructure. Keeping the reporting structure for these two types of information separate reduces the size of the traffic data files, allowing for faster data processing, better readability, and lower bandwidth cost for external parties who may be accessing the data through a feed.

To represent the monitored roadway network and draw it on maps, PeMS requires Geographic Information System (GIS) type roadway polylines defined by latitudes and longitudes. To help the agency link PeMS data and performance metrics with their own linear referencing system, PeMS also associates these polylines with state roadway mileposts. In most state agencies, mileposts are a reference system used to track highway mileage and denote the locations of landmarks. Typically, these mileposts reset at county boundaries. In cases where freeway alignments have changed over time, it is likely that the difference between two milepost markers no longer represents the true physical distance down the roadway. For this reason, PeMS adds in a third representation of the roadway network, called an absolute postmile. These are akin to mileposts, but they represent the true linear distance down a roadway, as computed from the polylines. They do not reset at county boundaries, in order to facilitate the computation of performance metrics across long sections of freeway. In PeMS, this information is ultimately stored in a freeway configuration database table that contains a record for every 10^{th} of a mile on every freeway. Each record contains the freeway number, direction of travel, latitude and longitude, state milepost, and absolute postmile.

PeMS also requires metadata concerning the detection equipment from which the source system is collecting data. This is due to the need to standardize data collection and processing across all agencies, regardless of their source system structures. Configuration information ultimately populates detector, station, and controller configuration database tables in PeMS, and is used to correctly aggregate data and run equipment diagnostic algorithms.

Finally, the data acquisition step often involves reconciliation between the framework of the source system and the monitoring system. For example, different terminology can lead to incorrect interpretations of the data. As such, this step often requires significant communication between the system contractor and the agency staff who have familiarity with the data collection system, in order to resolve open questions and make sure that accurate assumptions are being made.

Integration of District 3 Case Study Data Sources into PeMS

Keeping the above in mind, the two sources of data utilized in support of this case study, based on the movement of vehicles equipped with Electronic Toll Collection (ETC) and Bluetooth devices, are extremely new and not currently integrated into Caltrans District 3's existing PeMS data feed. Consequently, it was necessary to ingest these data sets into project–specific instances of PeMS for analysis as part of this project. This section provides an overview of the resources needed to conduct the pre-requisite data collection through monitoring system integration-related activities, as well as discusses some of the challenges likely to be encountered when developing such a monitoring system. Such activities included:

- **ETC Data** - With the Tahoe area ETC data, the goal was to use pre-existing PeMS ETC processing and equipment configuration software, as well as the road network

definitions in use by PeMS for Caltrans. This effort proved to be fairly straightforward and no special accommodations were required, other than dealing with detectors that would occasionally go off-line during real-time collection. As per public agency policy, all individual toll tag identifiers from the ETC readers were deleted every 24 hours.

- For each ETC reader station, the research team was provided with information regarding the county in which it was deployed, freeway on which it was located, a direction of travel for which it was collecting data, milepost, textual location, and the Internet Protocol (IP) address used to communicate with it to obtain data. To integrate each reader into PeMS so that data could be collected in real-time, the research team assigned each reader a unique ID and determined its latitude and longitude. Software was then developed to communicate with each reader's IP address, obtain its data, and incorporate that data into the PeMS database.
- **Bluetooth Data** – With the Lake Tahoe area Bluetooth data, the goal was to configure PeMS so that the Bluetooth readers and data they produced could be utilized as if it was from standard ETC reader stations. For each BTR, the research team received configuration data in a text file, with fields for the node (reader) ID, a textual location, and a latitude/longitude. Configuration data was provided for a total of 26 Bluetooth readers. Caltrans also provided the research team with a 2 gigabyte SQL file containing all of the Bluetooth data collected by the BTRs between December 25, 2010 and April 21, 2011. The research team subsequently integrated this data into PeMS and processed it to compute travel times between each pair of BTRs.

Analyzing ETC and Bluetooth Data

PeMS collects sensor data, either by directly polling each detector, obtaining it from an existing data collection system, or via integration of data from another archival resource, and stores it in an Oracle database. Reliability measures available based on this data will depend on the type of detector from which it has been collected – e.g., loop detectors will provide different raw data for analysis than ETC or Bluetooth-based data collection systems. Reliability metrics available in PeMS based on data from the ETC and Bluetooth systems are as follows:

- **Min** - The fastest vehicles that traveled across a roadway segment during a given period of time.
- **25th**: The 25th percentile travel time during a given period of time.
- **Mean**: The mean travel time during a given period of time.
- **Median**: The median travel time during a given period of time.
- **75th**: The 75th percentile travel time during a given period of time.
- **Max**: The slowest moving vehicles that traveled across a roadway segment during a given period of time. It is likely that much (if not all) of this data is composed of outliers that made at least one stop between two consecutive readers before completing their trip.

Each of the reliability measures described above is available for analysis based on five-minute and hourly time periods.

As stated above, the research team utilized pre-existing PeMS ETC processing and equipment configuration software to support the development and deployment of ETC and BTR

instances of PeMS. Existing PeMS analysis tools create reports of travel time versus starting time. For a given starting (or source) tag reader, the travel time to a destination tag reader is defined as the amount of time it takes for a specific tag to be seen at the destination tag reader. Due to public agency policy, PeMS does not store travel times for individual ETC tag reads, only recording summary statistics for all of the tags that traversed the distance between each consecutive pair of readers during a given period of time. That said, similar regulations do not currently exist regarding the use of data collected from Bluetooth devices. As such, the research team had access to a much wider variety of raw and summary data concerning the movement of Bluetooth-enabled vehicles for use as part of this case study.

It is important to note that the algorithm currently used by PeMS to calculate travel times based on ETC and Bluetooth data is fairly simple, with its only purpose being to identify travel times for vehicles that pass between consecutive readers regardless of whether the resultant travel time makes logical sense. For example, there is no way of knowing if a given vehicle got off the freeway in between reader stations. We only know when they were seen at each station. As a result, the travel times produced by PeMS based on this data have the potential to be significantly influenced by outliers and can at times be quite "noisy."

Lastly, there are two key differences between the ETC and Bluetooth technologies that needed to be accounted for as part of the research team's efforts to utilize the BTR data available as part of this project:

- **Directionality** – ETC detectors are aimed in such a way as to sense traffic flowing in a particular direction. In most cases, well over 95% of data collected by an ETC device is from traffic flowing in the direction that the detector is anticipated to be measuring. The Bluetooth readers do not have this directional bias. Both ETC and Bluetooth readers are capable of recording the presence of a single vehicle multiple times as it passes through the reader's detection zone. In the case of ETC readers, a vehicle is seldom detected more than twice, due to the limited range and directionality (aimed down on a spot on the road, not parallel to the ground) of the ETC antenna. However, Bluetooth readers can record any device generating a Bluetooth signal within its sensing radius, sometimes from 100 meters away. This can result in a single Bluetooth device being detected many times as it passes through the reader's detection zone, especially in cases where it is traveling slowly or is stopped.
- Keeping the above in mind, PeMS expects data to come from devices that have a directional bias. To accommodate this issue, the research team configured PeMS to view each Bluetooth reader as generating data for two directions of travel and fed the data into PeMS twice, assigning it first to one detector in one direction of travel and then assigning a copy of that data to the other direction of travel as well.
- **Background noise** - Several Bluetooth readers deployed as part of this project are located within a few dozen meters of office buildings, homes, or parking lots. Consequently, there are many stationary (or nearly, so) Bluetooth devices residing within these locations that produce a reading every few seconds for hours on end. This data has the potential to overwhelm legitimate vehicular data, sometimes by as much as a factor of 10 times or greater. The research team's initial solution for dealing with this issue in order to generate roadway travel times for analysis was to eliminate all subsequent reports of unique Bluetooth media access control (MAC) addresses collected within one hour of its initial reporting.

Additional information concerning activities undertaken by the project team to optimize the usefulness of these data sets is contained in the section entitled "Methodological Experiments" and the first two use cases.

METHODOLOGICAL EXPERIMENTS

Overview

Due to the significant amounts of Bluetooth-based travel time data available for analysis as part of this case study, the research team elected to focus its methodological efforts on this dataset rather than on data generated by the ETC-based system. This stems from an awareness that Bluetooth-based systems, while new, have been rapidly embraced by a wide range of transportation agencies interested in identifying low-cost, easy to deploy solutions for collecting roadway travel times. A great deal remains largely unknown regarding the underlying nature of this data, including how filtering techniques might be applied at different stages of the analytical process to further refine generated travel times. As such, this section focuses on the evaluation of methods for identifying individual vehicle trips between Bluetooth readers, followed by a statistical analysis of procedurally generated vehicle travel times. Filtering techniques at both the procedural and statistical levels are also explored as methods for improving the quality of travel time estimates.

The primary output of this section is a methodology for obtaining filtered travel time histograms that depict the distribution of travel times within a sample of Bluetooth data. It is possible to generate parameterized probability distribution functions (PDFs) from these histograms as was done in the San Diego case study, however this step is omitted here in favor of analysis of the underlying data issues.

Bluetooth Device Data

Impact of Bluetooth Reader Hardware on Data Available for Analysis

The characteristics of Bluetooth device data available for analysis are determined largely by the capabilities of the Bluetooth reader (BTR) deployed at the roadside. For example, only five of the 10 BTRs deployed by Caltrans had the ability to read and store signal strength measurements for each observation. Signal strength measurements are important because they provide the ability to determine the relative distance of each Bluetooth enabled mobile device from the reader. Whether or not a specific BTR has the ability to read and report signal strength values for each mobile device depends on the nature of the BTR's Host Controller Interface (HCI). The HCI is an interface between the Bluetooth protocol stack and the device's controller hardware. BTRs that reported signal strength values were based on Linux boards using the BlueZ protocol stack, while units not reporting signal strength values used a microcontroller-based implementation.

In the case of the Bluetooth Class I devices deployed by Caltrans, which have a range (radius) of detection of approximately 100m (see Exhibit C4-2), knowing the signal strength of each mobile device observation can be important to accurately calculate the travel times of those devices to the next BTR. To clarify, if a vehicle is traveling at 40 MPH, it will pass through the device's full 200m detection zone in approximately 10 seconds. However, if heavy congestion is present and the BTR zone traversal speed is only 5 MPH, it will take approximately 82 seconds.

In cases where BTRs are fairly close together, the accurate calculation of travel time can be significantly affected by whether or not the travel time analysis system has the ability to determine the time at which each Bluetooth device is closest to each BTR; the impact is even greater during periods of congestion when vehicles are moving slowly and generating many more observations. This issue is underscored by the fact that within the Caltrans dataset, Bluetooth-enabled mobile devices each generated approximately 1 observation (on average) per second (see Table C4-3), resulting in a mean number of observations per mobile device, per visit to each BTR of between 1.06 and 21.30.

Exhibit C4-2 provides a graphical depiction of the nature of the detection zones generated by Bluetooth and ETC-based data collection technologies.

Table C4-3: BTR Detection Zone Traversal Times and Observations

Bluetooth Reader ID	Mean Zone Traversal Time (Secs.)	St. dev. of Zone Traversal Time (Secs.)	Mean Observations per Visit
BTR 1	1.43	12.02	1.06
BTR 2	0.88	9.88	1.10
BTR 3	1.27	25.45	1.16
BTR 4	4.20	48.56	1.44
BTR 5	0.58	10.91	1.09
BTR 6	7.93	48.01	1.38
BTR 7	8.77	37.49	11.49
BTR 8	8.73	65.38	11.44
BTR 9	5.77	36.78	8.19
BTR 10	23.10	116.65	21.30

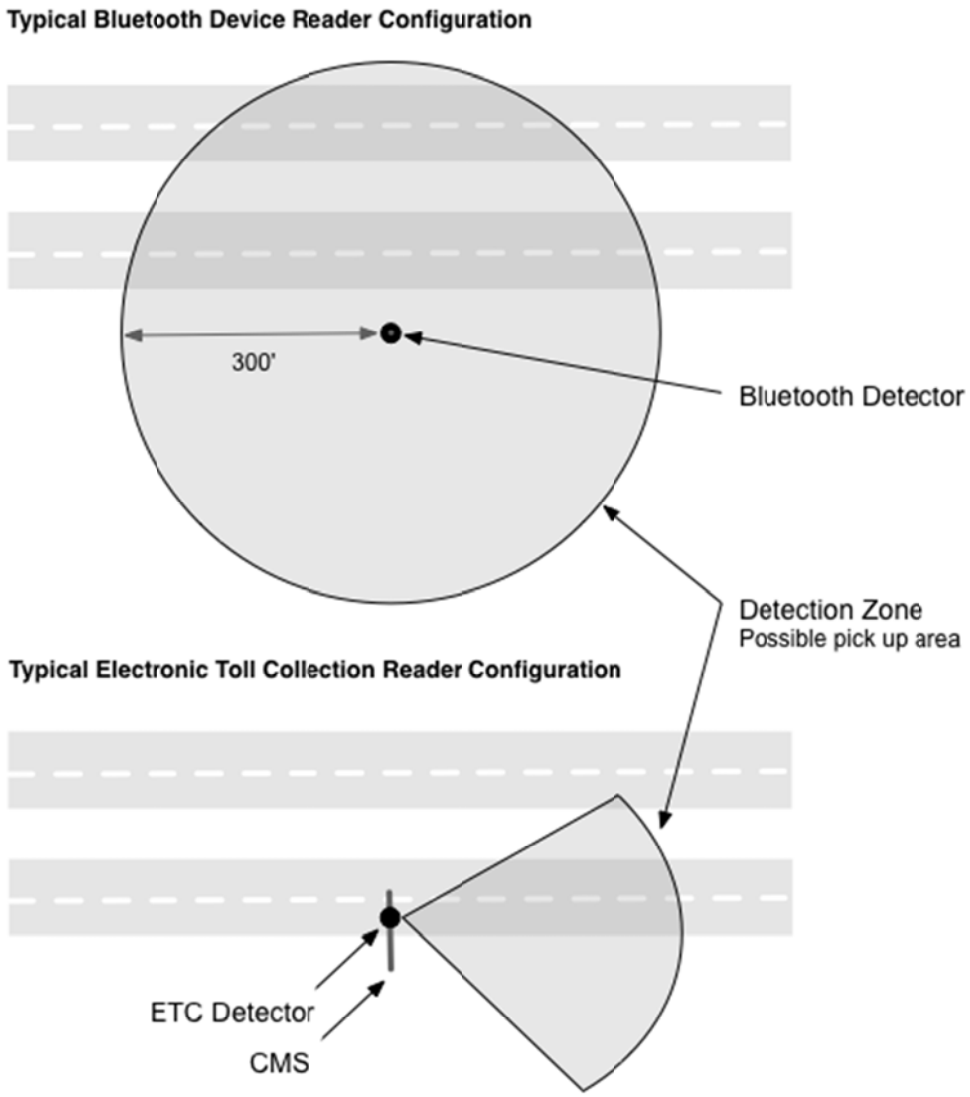

Exhibit C4-2: Bluetooth and ETC Reader Detection Zones

Table C4-4 provides examples of mean, maximum, and standard deviation of mobile device signal strengths collected by the various BTRs involved in this study. BTRs with signal strength characteristics noted as "N/A" did not have the capability to collect signal strength data. Signal strength readings are proportional to the distance between a BTR and each mobile device. A BTR's mean signal strength is therefore a function of the location of the BTR relative to the roadway. In addition, BTR antenna gain varies as a function of manufacturer and type which affects mean signal strength (3).

Exhibit C4-3 compares observed signal strengths over time for 3 vehicles traveling through BTR detection zones; each plot is centered (from a temporal perspective) on the time at which the peak signal strength was detected for each vehicle. The first vehicle arrives in the detection zone, travels past the reader and stops for approximately 11 minutes within the

detection zone. The second vehicle passes through the detection zone in approximately 17 seconds, traveling at 24 MPH. The third vehicle enters the BTR's detection zone, pauses for approximately 18 seconds, passes the BTR, and then departs the detection zone.

Table C4-4: BTR Signal Strength Characteristics

Bluetooth Reader ID	Number of Observations	Mean Signal Strength	Maximum Signal Strength	Signal Strength St. Dev.
BTR 1	319	**N/A**	**N/A**	**N/A**
BTR 2	430,679	**N/A**	**N/A**	**N/A**
BTR 3	442,739	**N/A**	**N/A**	**N/A**
BTR 4	1,055,037	27.24	55	16.04
BTR 5	870,362	**N/A**	**N/A**	**N/A**
BTR 6	401	**N/A**	**N/A**	**N/A**
BTR 7	1,507,667	77.66	96	3.53
BTR 8	893,232	77.68	94	3.38
BTR 9	403,628	77.32	93	3.01
BTR 10	2,178,002	77.18	95	3.38

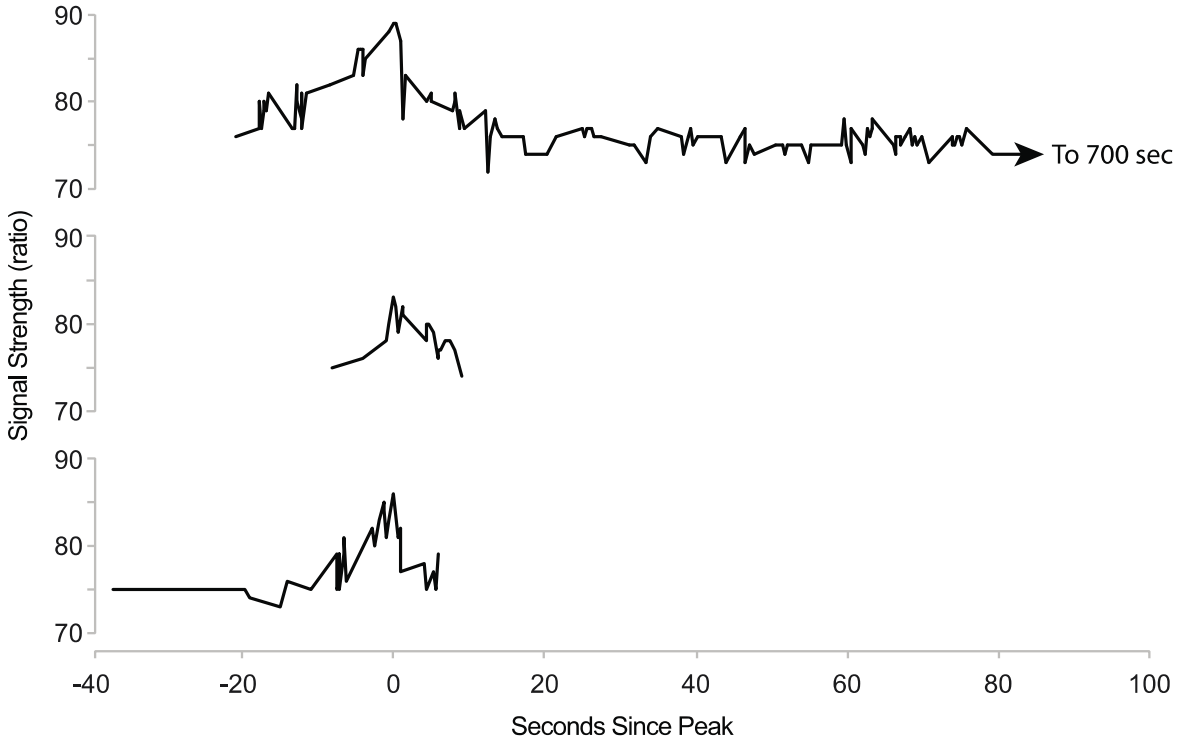

Exhibit C4-3: Comparison of Observed Signal Strengths versus Time for 3 Vehicles

Mobile Device Data Characteristics

Bluetooth device data collected as part of this case study exhibited a number of characteristics that should be understood prior to attempting the calculation of roadway travel times; these are discussed below.

Devices Visiting Only One BTR. One way to classify mobile devices is by the total number of unique BTRs they visit. For the purposes of calculating segment (BTR to BTR) travel times, observations generated by devices that visit only a single BTR can be ignored. Based on the team's analysis, approximately 29% of all mobile devices represented in the Caltrans dataset visited only a single BTR during a given trip; these devices contributed 12.5% of all mobile device observations (Table C4-5).

Table C4-5: Observations Generated by Devices – By Number of BTRs Visited

	Visited 1 BTR	Visited > 1 BTR	Total
Number of Devices	146,075 (29%)	356,408 (71%)	502,483 (100%)
Number of Observations	2,315,389 (13%)	16,176,143 (88%)	18,491,532 (100%)

Variable BTR Detection Zone Traversal Times. As discussed above, mobile devices take varying amounts of time and generate unpredictable numbers of observations each time they pass through a given BTR's detection zone. Generally, the number of observations generated by a device is proportional to the amount of time the vehicle is present within the detection zone; which is proportional to the vehicle's speed. Based on analysis conducted as part of this case study, the research team believes that the "Mean Zone Traversal Time" (see Table C4-3) is affected by a combination of the physical location of the reader relative to the roadway and other roadway characteristics. For example, BTR #2 (see Table C4-2 for the location of each BTR) is located at the end of an entrance ramp and is isolated from nearby arterials and buildings. It has a mean detection zone traversal time of .88 seconds with approximately 1.10 observations per visit. This can be seen in the zone traversal time frequency distribution (top distribution in Exhibit C4-4) with no delay time for vehicles passing through the detection zone. This reader contrasts with BTR #10, which has a mean detection zone traversal time of 23.1 seconds and 21.3 observations per visit (bottom distribution Exhibit C4-4). This reader is located on one leg of a T-intersection with a single stop sign. Consequently, cars queuing at the stop sign may be contributing significantly to the long zone traversal times.

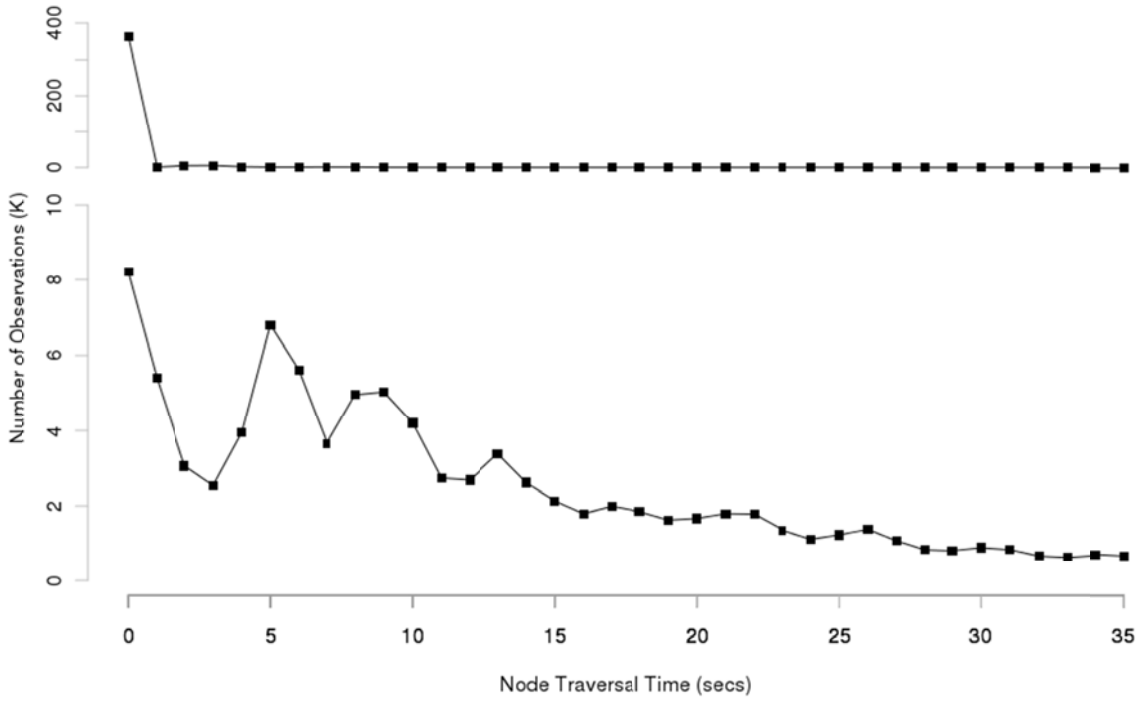

Exhibit C4-4: Node Traversal Time Frequency Distribution for BTRs #2 (top) & #10 (bottom)

Multiple Mobile Device Observations per BTR. Individual mobile devices can enter and exit a single BTR's detection zone multiple times during a sufficiently lengthy period of time. Depending on the size of the window of time, these individual observations have the potential to be matched with a significant number of observations from other BTRs. Table C4-6 displays the results of one vehicle visiting BTR #10 four times during one day. The final 2 visits are separated by just 10 minutes (the 3rd and 4th visit are shown in Exhibit C4-5. This demonstrates that a travel time algorithm that processes device observation data must have the ability to aggregate and differentiate between clouds of such observations separated in time as a step in the process of calculating travel times between BTRs.

Table C4-6: Multiple Device Observations for One Device at BTR #10

Visit Number	Time	Number of Observation	Time Delta (Seconds)
1	09:50 am	1	0.00
2	15:00 pm	5	2.54
3	16:33 pm	2	18.05
4	16:43 pm	39	42.98

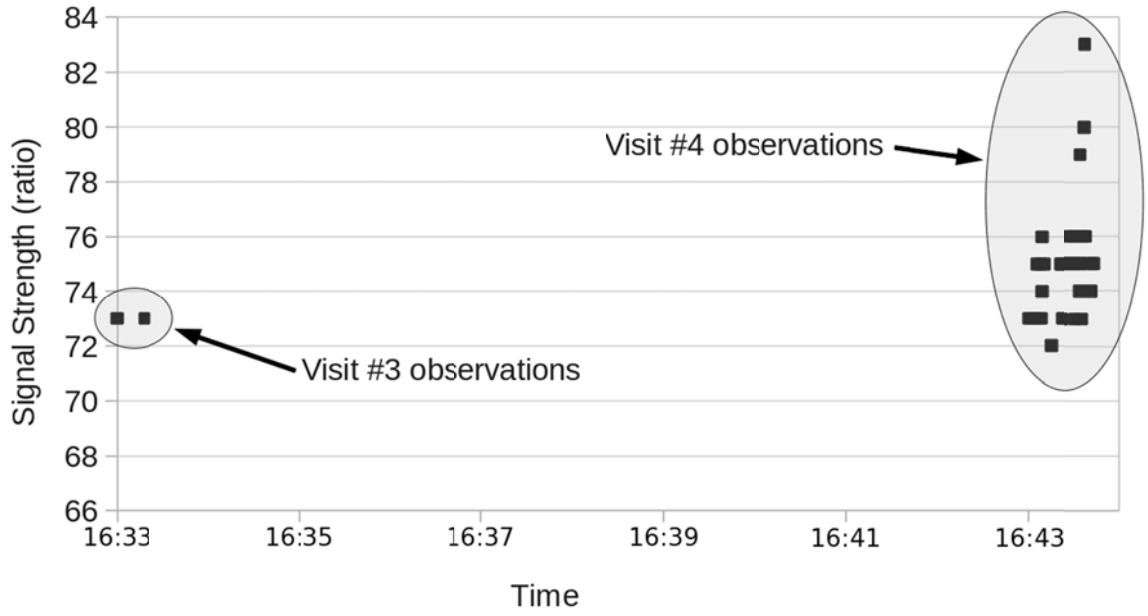

Exhibit C4-5: Details of Visits 3 and 4 for a Single Mobile Device at BTR #10

Calculating Travel Times Based on Bluetooth Device Data

The primary goal of BTR-based data analysis is to characterize segment travel times between BTRs based on the re-identification (re-id) of observations derived from unique mobile devices. Generally, the data processing procedures associated with the calculation of BTR to BTR travel times can be broadly broken down into 3 processes, as shown below. The first two processes are procedural. The third is statistical:

1. Identification of Passage Times
    A. Aggregating device observations into visits
    B. Selecting BTR passage time
2. Generation of Passage Time Pairs
    A. Method 1: Maximum origin and destination permutations
    B. Method 2: Use of all origin visits
    C. Method 3: Aggregation of visits
3. Generation of Segment Travel Time Histograms
    A. Filter outliers across days
    B. Filter outliers across time intervals
    C. Remove intervals with few observations
    D. Remove highly variable intervals

The steps involved in these 3 processes are discussed below.

Process 1: Identification of Passage Times

The first step in the process of calculating segment travel time PDFs for a roadway is the calculation of segment travel times for individual vehicles. A vehicle segment travel time is calculated as the difference between the vehicle's *Passage Times* at both the origin and destination BTRs. Passage Time is defined as the single point in time selected to represent when a vehicle passed through a BTR's detection zone. As previously discussed, mobile devices typically generate multiple observations as they pass through a BTR's detection zone. Consequently, selection of appropriate passage times is an important step in maximizing the accuracy of calculated segment travel times for individual vehicles.

Aggregating Device Observations into Visits. The goal here is to identify clusters of observations that represent a vehicle's continuous presence in the detection zone. Each group of observations is referred to as a *visit*. For example, Exhibit C4-6 displays numerous observations during a single vehicle's visit over the course of several minutes. Exhibit C4-7 displays two separate visits by a single vehicle (each with multiple observations) separated in time by a stop outside of the detection zone. Identifying unique visits is an important step in increasing the accuracy of segment travel time calculations. Associating multiple observations clustered in time as part of a single visit rationalizes the selection of a single passage time for calculating the vehicle's travel time to a destination BTR. The alternative, which makes little sense, would be to calculate a travel time for each origin observation to the destination BTR. Identifying visits also enables the assessment of arrival (first mobile device observation) and departure (last mobile device observation) times for each mobile device, which can, depending on the circumstances, be used as the passage time. Exhibit C4-6 depicts clusters of observations associated with two distinct visits by a single vehicle to the same BTR.

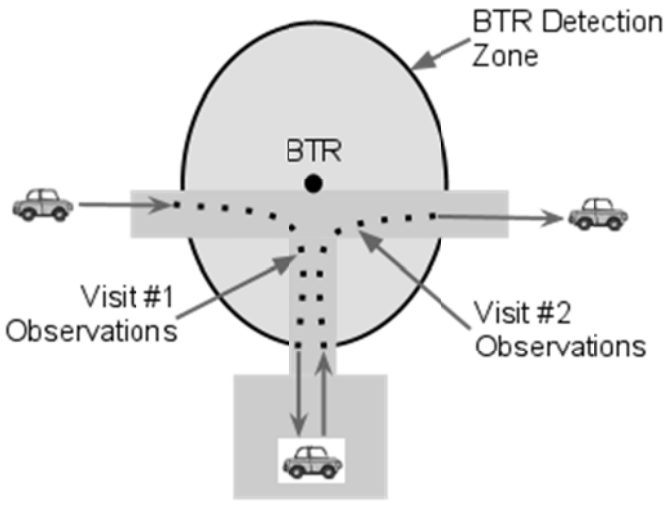

Exhibit C4-6: Visits as Clusters of Observations in Time

The method used to aggregate visits is a causal sliding time window filter. It is a filter in that it removes unnecessary observations during the aggregation process. It is causal in that it uses only past and present observations to support its decision-making. This filter discards all subsequent observations that are within a fixed time span from the time of a prior observation.

However, when the time between an observation and the prior observation with which it is being compared exceeds this time span, it is considered to be part of a new visit. This has the effect of aggregating observations into visits by arrival time (or departure time, depending on how it is implemented) and discarding all other observations. Exhibit C4-7 displays 11 observations that have been aggregated into 2 visits due to a sufficiently large time gap between the 5th (part of visit #1) and 6th (part of visit #2) observations. This filter is an efficient method of processing real-time observations and compressing large quantities of observation data for efficient storage.

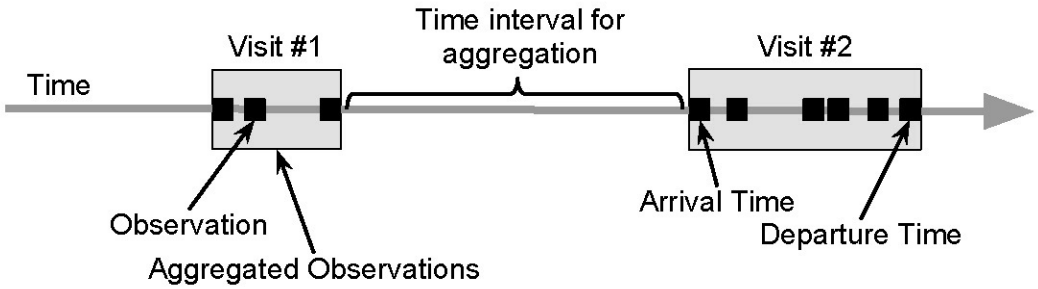

Exhibit C4-7: Aggregation of Observations into Visits Using Time Intervals

The size of the filter interval time depicted in Exhibit C4-7 determines the granularity of identified visits. The effect of different sized interval times is shown in Exhibit C4-8. In general, selecting the largest reasonable interval time is desirable because it results in more accurate estimates for arrival and departure times (and hence, passage times, depending on the method used). However, over-aggregating visits is potentially problematic. The research team has identified the following error types to consider when selecting an interval time:

- **Observation over-aggregation**: When observations belonging to multiple visits are incorrectly aggregated as a single visit, the arrival passage time and departure passage time may be calculated as too early or too late, depending on the method used. This may also result in the classification of stopped non-delay time as stopped delay time because the vehicle is incorrectly identified as being continuously in the detection zone. For example, if a filter time interval of 20 minutes is used and the vehicle leaves the detection area and returns 10 minutes later, this 10-minute absence would be classified as having been spent within the zone. If the distance to adjacent BTRs is close, over-aggregation risks subsuming valid origin visits, resulting in the deletion of valid trips. For these reasons, under-aggregation is preferred to over-aggregation.
- **Observation under-aggregation**: Incorrectly sub-dividing observations from a single visit into multiple visits may result in the incorrect calculation of passage time, depending on the method used. Under aggregation is less problematic because multiple sequential visits that are not interwoven with visits to other BTRs can be aggregated and considered to be a single visit (discussed below).

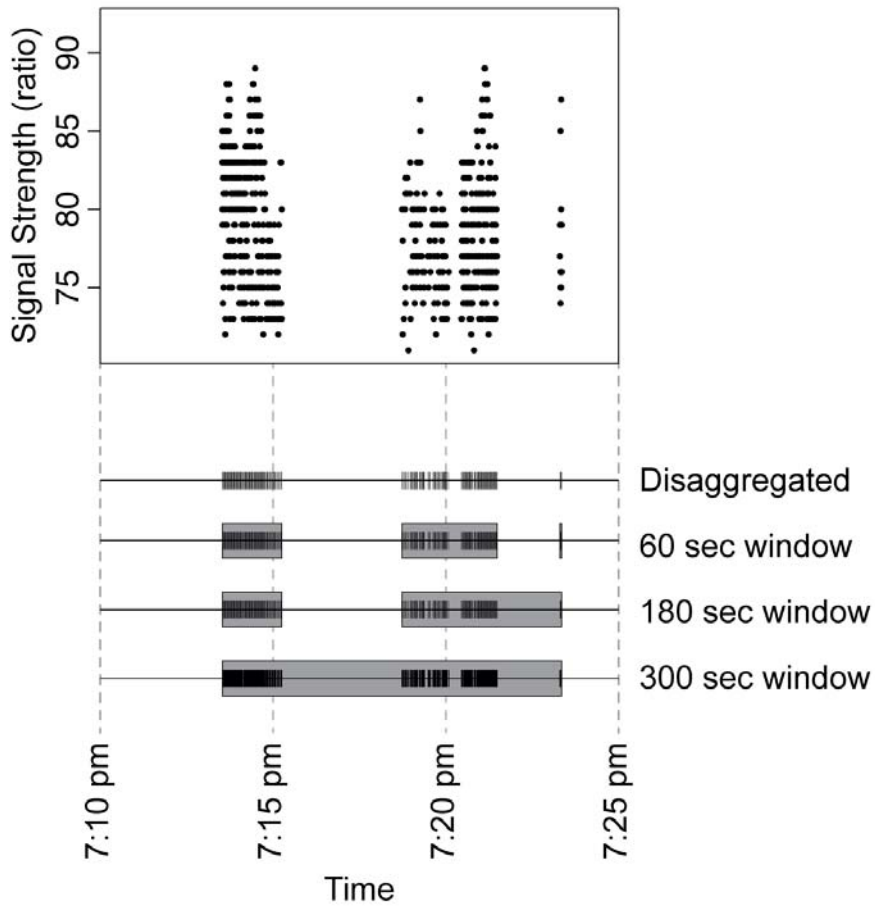

Exhibit C4-8: Influence of Window Sizes on Observation Aggregation

The Caltrans Bluetooth data was processed by storing the arrival, departure, and maximum signal strength (where available) for each identified visit. Observations were aggregated using a 120-second time window. The 120-second window size was selected due to the small distance between BTRs #9 and #10 (about 3.8 miles) and the preference for under-aggregating visits. Moreover, using a smaller, 60-second time interval was found to be under-aggregating observations in too large a number of cases. Other researchers appear to be using a 5-minute interval (4), which may be appropriate for large BTR to BTR distances. When deploying permanent travel time data collection systems based on BTR (or related) technologies, it is likely that the filter interval should be adjusted for each BTR as a function of its location and the characteristics of the surrounding region. For example, if a snow chain fitting area is nearby, longer interval times may be optimal.

Vehicles that are continuously within a BTR's detection zone (and generating observations) are either in travel mode (e.g. driving, in congestion, at a stop light) or trip mode (e.g. stopped at a fuel station, parked at the side of the road). Without more information, distinguishing between trip and travel behavior within a single visit is difficult. In contrast, distinguishing between trip and travel behavior across multiple visits is possible. Repeat visits to the same BTR (without visiting any other BTR) can be assumed to be non-travel time oriented and therefore eliminated. For example, if the vehicle in Table C4-6 did not visit other BTRs

between 09:50 am and 16:43 pm, then these visits can be eliminated from travel time calculations.

Selecting BTR Passage Time. The precise methodology used to determine a vehicle's passage time depends on the availability of signal strength data, the distance to adjacent BTRs and traffic flow patterns in the area surrounding a BTR. When signal strength data is available, passage time can be considered as corresponding to the mobile device observation with the greatest signal strength. In cases where signal strength data is not available and the distance to adjacent BTRs is large, the arrival, mean, or departure time may be used as the passage time without introducing a significant bias. However, if traffic through the detection zone is subject to stop delay time (e.g., traffic signals, stop signs, congestion, etc.) then use of arrival or departure times may either introduce or eliminate significant bias. This is illustrated with BTRs #9 and #10, below.

BTR #10 provides an example of how the use of arrival vs. departure times as a proxy for passage time (in cases where no signal strength data is available) can influence the calculation of segment travel times. BTR #10 is located on one leg of a T-intersection with a single stop sign, as shown in Exhibit C4-9. Its nearest neighboring BTR is 3.8 miles away. The mean detection zone traversal time for BTR #10 is 23.1 seconds, which is likely due at least in part to vehicles queuing at the nearby stop sign. Vehicles queued at the stop sign either turn right, away from BTR #10, or turn left and pass it. The free-flow speed of traffic passing the BTR is approximately 45 MPH. At this speed, traffic passes through the detection zone in 9 seconds. As such, for vehicles not queued at the stop sign, arrival times are (on average) 4.5 seconds earlier and departure times 4.5 seconds later than when the mobile device passes the BTR. If BTR #10 is used as the point of origin for generating a segment travel time, e.g., with BTR #9 as the destination, left-turning vehicles proceeding from the stop sign will pass the reader and generate an arrival time (and consequently a passage time) that is approximately 23.1 – 4.5 = 18.6 seconds early, or nearly 7% of the travel time to the next BTR (3.8 miles away with a free flow speed of 45 mph). This error may be further compounded by heavy traffic, causing longer queues at the stop sign within the detection zone of BTR #10. In contrast, basing vehicle passage time on departure time, thereby removing the delay associated with the presence of the stop sign, would introduce only about 4.5 seconds of error, representing a substantial improvement over use of arrival time. This example demonstrates why significant attention needs to be paid to the process used to calculate passage time in situations where signal strength data is not available.

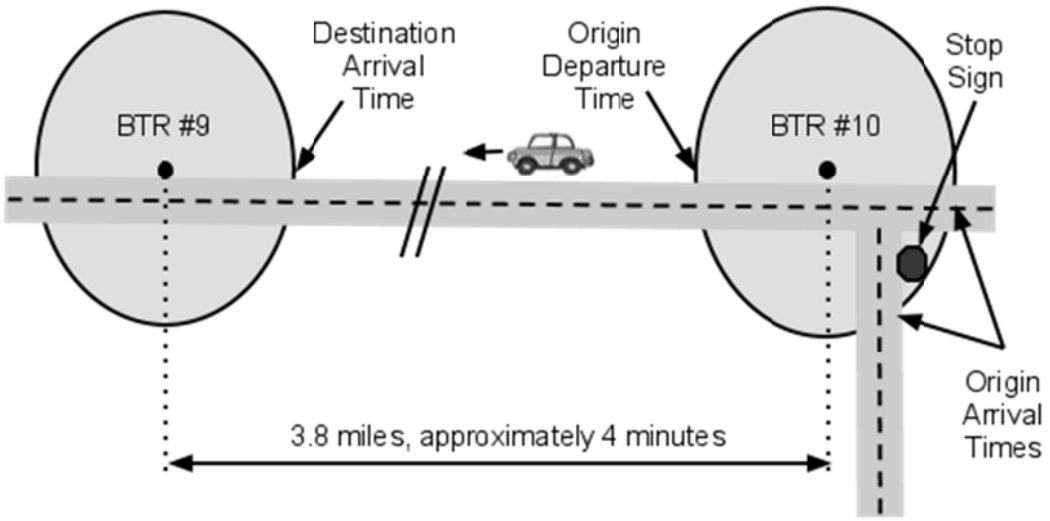

Exhibit C4-9: BTR #10 Geometry in Relation to Adjacent BTR #9

In addition to considering the impact of BTR passage times, users of Bluetooth data must also consider that the accurate calculation of segment travel time is a function of the relationship between BTR-to-BTR distance and the maximum speed error. Following on the analysis performed by Haghani, et. al. (4), this relationship is depicted in Exhibit C4-10 and Exhibit C4-11as the maximum error in segment speed versus BTR-to-BTR distance ("Distance between Nodes") for 4 speeds. For this analysis, BTR-to-BTR distance is L, vehicle speed is S, and the travel time between adjacent BTRs is T, such that:

1) $L = S * T$

2) $L + \Delta L = (S + \Delta S)(T + \Delta T)$

3) $\Delta S_{max} <= (\Delta L_{max} - \Delta T * S) / (L / S + \Delta T)$

Equation #2 introduces error terms for L, S, and T. As per Equation #3, the maximum error in distance, ΔL_{max}, is assumed to be 600 ft. (the diameter of each BTR's detection zone).

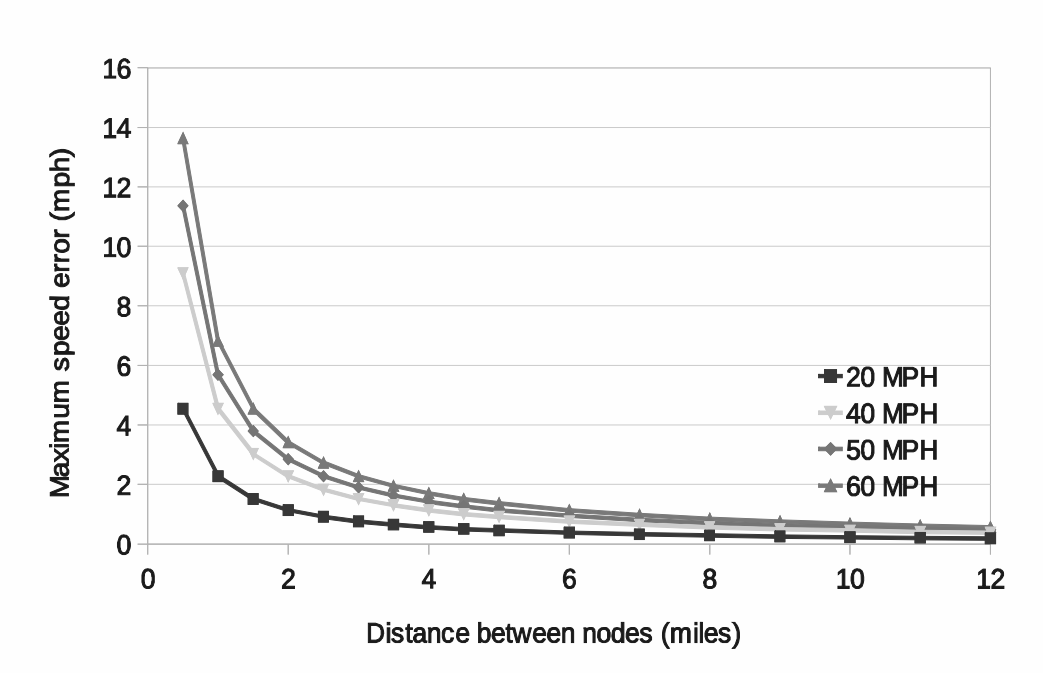

Exhibit C4-10: Relationship Between Maximum Speed Error & BTR-to BTR Distance with ΔT=0

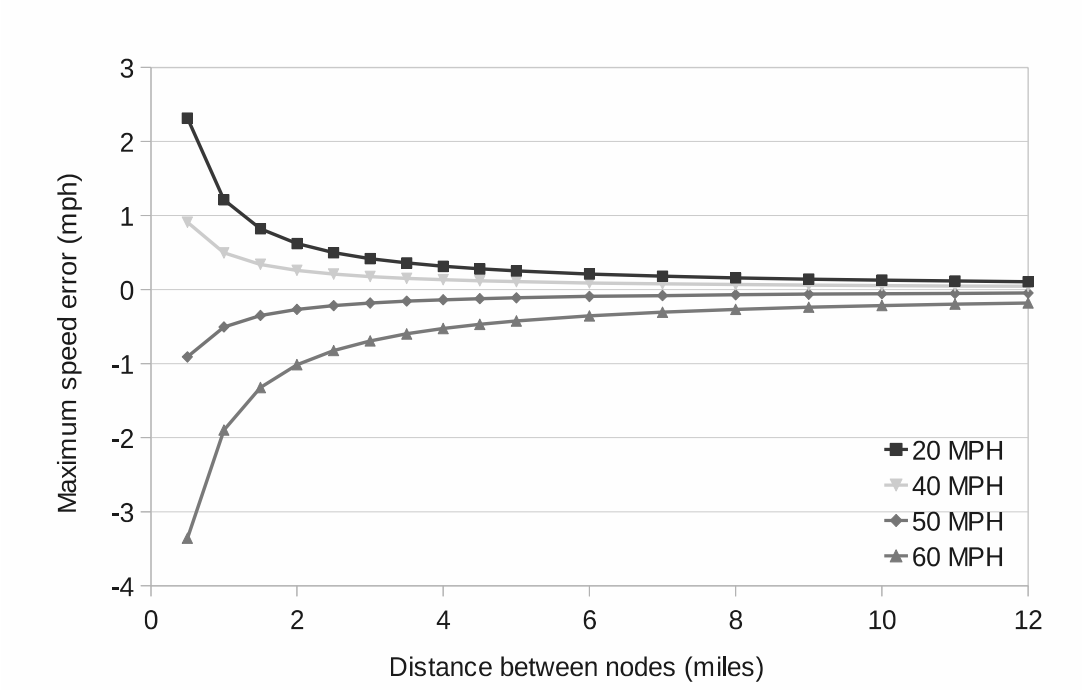

Exhibit C4-11: Relationship Between Maximum Speed Error & BTR-to BTR Distance with ΔT>0

As per Exhibit C4-10, if time error (ΔT) is 0, then speed error is maximized as vehicle speed increases. As a result, for BTRs spaced less than 2 miles apart collecting Bluetooth data from vehicles traveling at high rates of speed, the maximum speed error becomes quite significant. However, due to a combination of clock synchronization error and/or Bluetooth time stamp inaccuracies, it is highly unlikely that ΔT will often (if ever) equal 0.

As per Exhibit C4-11, if time error (ΔT) is greater than zero, then both slower, as well as faster vehicle speeds have the potential to maximize speed errors. Within the context of this graph, the influence of time errors has a tendency to negate the effect of distance errors. A time error of 4 seconds was used based on clock synchronization error; associated with Caltrans' method for synchronizing BTRs when local time differed from network time by more than 2 seconds.

Process 2: Generation of Passage Time Pairs

It is common for vehicles to generate multiple sequential visits per BTR, which may be interwoven in time with visits at other BTRs (see Table C4-7). For BTRs with a significant mean zone traversal time, it is common for vehicles to generate multiple visits close in time. The motivation for grouping visits is evident in Table C4-7, where the vehicle was at the origin BTR multiple times (see rows 1-3) before traveling to the destination BTR (see row 4). Based on this data, 3 different travel times could be calculated: 1 to 4; 2 to 4; or 3 to 4. Which pair or pairs represent the most likely trip? The benefit of performing more complex analysis of visits is that many likely false trips can be eliminated, increasing the quality of the calculated travel time. Three methods of identifying segment trips are discussed below.

Table C4-7: Visits for a Single Vehicle Between Two BTRs

Row	Origin BTR Visit	Destination BTR Visit	Time	Observations Per Visit
1	1		Fri Jan 28 13:07:19 2011	7
2	2		Fri Jan 28 16:24:06 2011	13
3	3		Fri Jan 28 17:41:50 2011	25
4		1	Fri Jan 28 22:07:36 2011	3
5	4		Sat Jan 29 15:07:40 2011	4
6		2	Sun Jan 30 10:49:03 2011	129
7	5		Sun Jan 30 12:15:33 2011	3
8	6		Mon Jan 31 13:05:54 2011	10

Method 1: The first potential method for identifying segment trips is simple: create an origin and destination pair for every possible permutation of visits, except those generating negative travel times (Exhibit C4-12). For example, the visits in Table C4-7 show 6 origin and 2 destination visits, resulting in 12 possible pairs. Five pairs can be discarded because they generate negative travel times. Even so, this approach will generate many passage time pairs that do not represent actual trips. Using this method, 243,777 travel times were generated between one pair of BTRs over a three-month period.

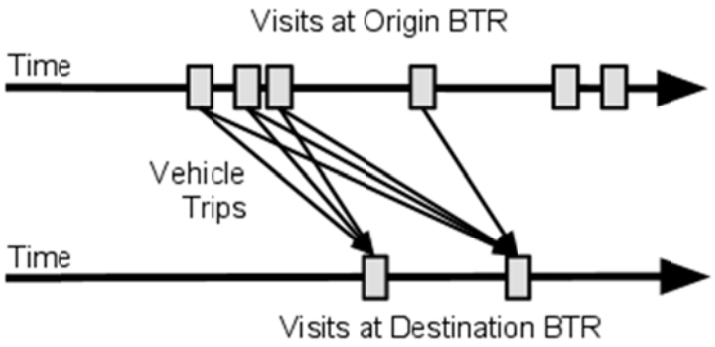

Exhibit C4-12: Trips Generated from All Visit Permutations

Method 2: The second potential method for identifying segment trips is also simple, but represents an improvement from the first method. It creates an origin and destination pair for every origin visit and the closest (in time) destination visit, as shown in Exhibit C4-13. Multiple origin visits would therefore potentially be paired with a single destination visit. Using this method with the data in Table C4-7 would generate 4 pairs: 1-4, 2-4, 3-4, 5-6. This method generated 60,537 travel times between the origin and destination BTRs, eliminating 183,240 (75%) potential trips compared with the first method.

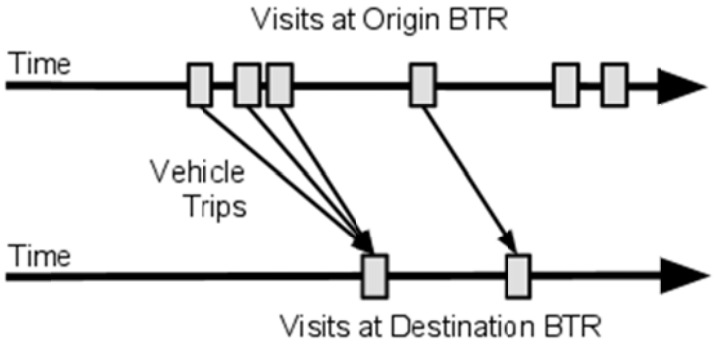

Exhibit C4-13: Trips Generated from All Origin Visits to First Destination Visit

Method 3: Vehicles frequently make multiple visits to an origin BTR before traveling to a destination BTR. The third method of eliminating invalid segment trips aggregates those origin visits that would otherwise be interpreted incorrectly as multiple trips between the origin and

destination readers. This method can be described as aggregating visits at the BTR network level. This is an additional level of aggregation beyond aggregating individual observations into visits, as discussed in the previous section. Logically, a single visit represents a vehicle's continuous presence within a BTR's detection zone. In contrast, multiple visits aggregated into a single grouping represent a vehicle's continuous presence within the geographic region around the BTR, as determined by the distance to adjacent BTRs. This method is an example of using knowledge of network topology to identify valid trips.

This method can be applied to the data displayed in Table C4-7, which shows 3 origin visits in rows 1, 2, and 3. The question is whether any of these visits can be aggregated or should each be considered a valid origin departure? The distance from the origin (BTR #7) to the destination (BTR #10) is 50 miles (or 100 miles for the round-trip). Driving at some maximum reasonable speed (for that road segment, anything over 80 MPH is unreasonable) a vehicle would take 76-minutes for the round-trip. Therefore, if the time between visits at the origin is less than 76-minutes, they can be aggregated and considered as a single visit. In Table C4-7, visits 2 and 3 (rows 2 and 3) meet this criterion and can therefore be aggregated (Exhibit C4-14). This eliminates 1 of 3 potential origin visits that could potentially be paired with the destination visit in row 4. Again, the idea is to identify when the vehicle was continuously within the geographic region around the origin BTR and eliminate departure visits wherever possible. When this method was applied to the data set, it generated 39,836 travel times, eliminating 20,701 (34%) potential trips compared with the second method discussed above.

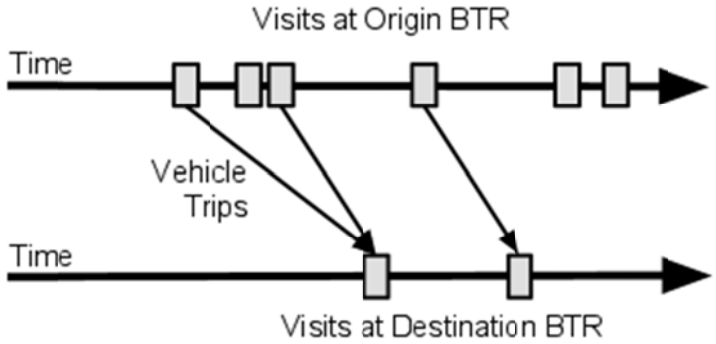

Exhibit C4-14: Trips Generated from Aggregating Origin Visits

Additional filters could be used to identify and eliminate greater numbers of trips. For example, an algorithm could take advantage of graph topology and interspersed trips to other BTRs to aggregate larger numbers of visits. In addition, the algorithm could potentially track which destination visits had previously been paired with origin visits, eliminating unlikely trips. If PDFs are developed based on historical data, selection among multiple competing origin visits paired with a single destination visit could be probabilistic. These are potential topics for future research.

Process 3: Generation of Segment Travel Time Histograms

Previous sections described methods for determining travel times based on Bluetooth data. This was done by first identifying vehicle passage times at each of the Bluetooth readers,

then pairing those passage times from the same vehicle at origin and destination locations. These techniques were developed with the goal of maximizing the validity of the travel times. However, because of Bluetooth data's susceptibility to erroneous travel time measurements, even the most careful pairing methodology will still result in trip times (which could include stops and/or detours) that need to be filtered in order to obtain accurate ground truth travel times (the actual driving time).

This section of the methodology describes a four-step technique for filtering travel times, presents travel time histograms before and after filtering, and compares the effects of two passage time pairing techniques ("Method 2" and "Method 3" from *Generation of Passage Time Pairs*) on the resulting travel time histograms. The underlying parameterized travel time PDFs could be approximated from the filtered travel time histograms presented here. However, this step is omitted from this methodology section in order to more closely focus on the low-level issues associated with obtaining travel time distributions from Bluetooth data.

To begin, the distribution of raw travel times obtained from two different passage time pairing methods can be seen in Exhibit C4-15. The data presented here as "Method 2" was developed using the second passage time pairing method described in Generation of Passage Time Pairs. Data labeled as belonging to "Method 3" was built according to the third passage time pairing method in that same section. No "Method 1" analysis is included here due to that method's lack of sophistication. In Exhibit C4-15, the unfiltered travel time distributions appear similar apart from the quantity of data present. Both distributions have extremely long tails, with most trips lasting an hour or less and many taking months. It is clear from these figures that even the carefully constructed "Method 3" data is unusable before filtering.

Several plans have been developed to filter Bluetooth data. Here, we adopt a four-step method proposed by Haghani, et. al. (4). In Haghani's filtering plan, points are discarded based on their statistical characteristics, such as coefficient of variation and distance from the mean. The four data filtering steps are:

1) **Filter outliers across days.** This step is intended to remove measurements that do not represent an actual trip but rather a data artifact (i.e., the case above of a vehicle being missed one day and detected the next). Here, we group the travel times by day and plot PDFs of the speeds observed in each day (rounded to the nearest integer). To filter the data, thresholds are defined based on the moving average of the distribution of the speeds (with a recommended radius of 4 miles per hour). The low and high thresholds are defined as the minima of the moving average on either side of the modal speed (i.e., the first speed on either side of the mode in which the moving average increases with distance from the mode). All speeds above/below these values are discarded (see Exhibit C4-16).
2) **Filter outliers across time intervals.** For the remaining steps 2-4, time intervals smaller than one full day are considered (we use both 5-minute and 30-minute intervals). In this step, speed observations beyond the range $mean \pm 1.5\sigma$ within an interval are thrown out. The mean and standard deviation are based on the measurements within the interval.
3) **Remove intervals with few observations.** Haghani determines the minimum number of observations in a time interval required to effectively estimate ground truth speeds. This is based on the minimum detectible traffic volume and the length of the interval. Based on this, intervals with fewer than 3 measurements per 5-minutes (or 18 measurements per 30-minutes) were discarded.

4) **Remove highly variable intervals.** In Step 4, the variability among speed observations is kept to a reasonable level by throwing out all measurements from time intervals whose coefficient of variation (COV) is greater than 1.

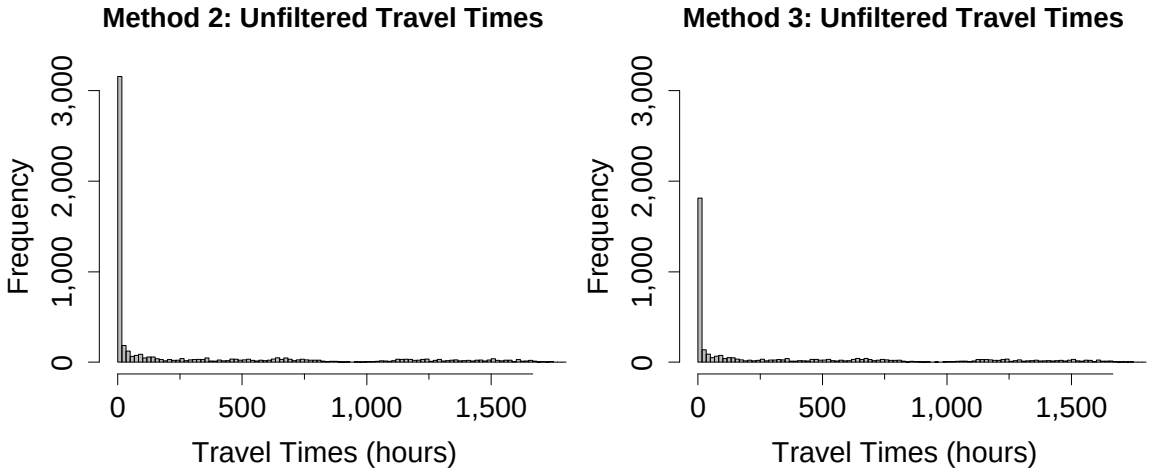

Exhibit C4-15: Unfiltered Travel Times Between One Pair of Bluetooth Readers, February 6, 2011 to February 12, 2011

To carry out Step 1, the moving average (with radius of 4 mph) is computed over the speed distribution for each day (note that speeds are found by simply dividing route length by travel time). The moving average and distribution of speeds from a single day can be seen in Exhibit C4-16. To exclude unreasonably low speeds, the modal speed is defined as the speed corresponding to the peak of the moving average above 20 mph (53 mph in this case). On this day, as a result of filtering Step 1, the upper threshold was set to the maximum observed speed (the minimum of the moving average above the modal speed), and the lower threshold was set to 25 mph (the minimum of the moving average below the modal speed). Thus, on this day, all data points representing speeds below 25 mph or above 62 mph were discarded as a result of Step 1.

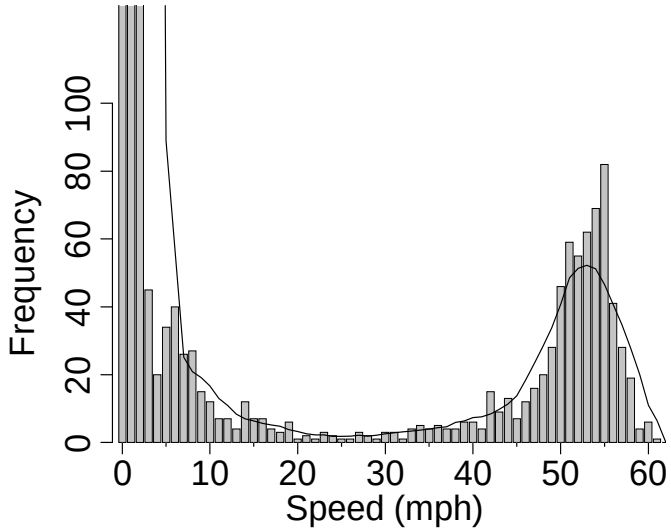

Exhibit C4-16: Distribution of Speeds

While Step 1 is carried out across days, Steps 2-4 are carried out across 5-minute and 30-minute intervals. The 5-minute interval was chosen to match what was done by Haghani (4) and represents a standard, baseline filter. Filtering results based on a 30-minute interval are also included to compare the effects of a wider filtering interval. A wider filtering interval may be more appropriate for sparse data sets like that available in the rural Lake Tahoe area where many 5-minute intervals contain no measurements at all. The particular details of steps 2-4 are more straightforward and are omitted from further discussion.

The results of Haghani's four-step filtering method on the data obtained using passage time pairing method 2 are presented in Exhibit C4-17 for the week beginning on February 6, 2011. The white points are points identified to be thrown out in that step, and the black points are points to be kept following that step. Note that the steps are performed sequentially, so that points discarded after Step 1 are not considered in Step 2, and so on. Higher traffic volumes due to weekend traffic in the area can clearly be seen on Friday and Saturday.

After the data has been filtered, the travel time distributions over the week appear much more meaningful, as can be seen by comparing Exhibit C4-18 with Exhibit C4-15. The filtered travel times have lost their unreasonably long values and in both cases, a nicely shaped distribution is visible. Note that the earlier comparison of the data sets remains true: both have similarly shaped distributions, but the data prepared using passage time pairing method 2 contains a greater quantity of data. This is because that data set was larger initially, and also a larger percentage of points from it survived filtering.

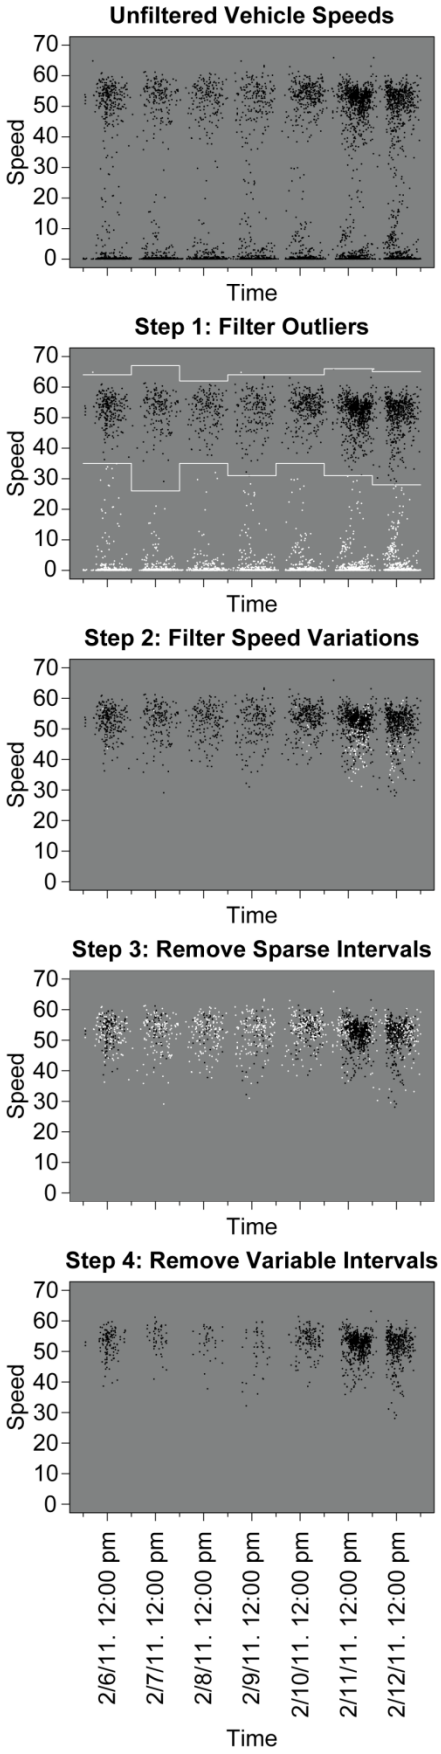

Exhibit C4-17: Four-step filtering on passage time pairing method 2 with 5-minute intervals.

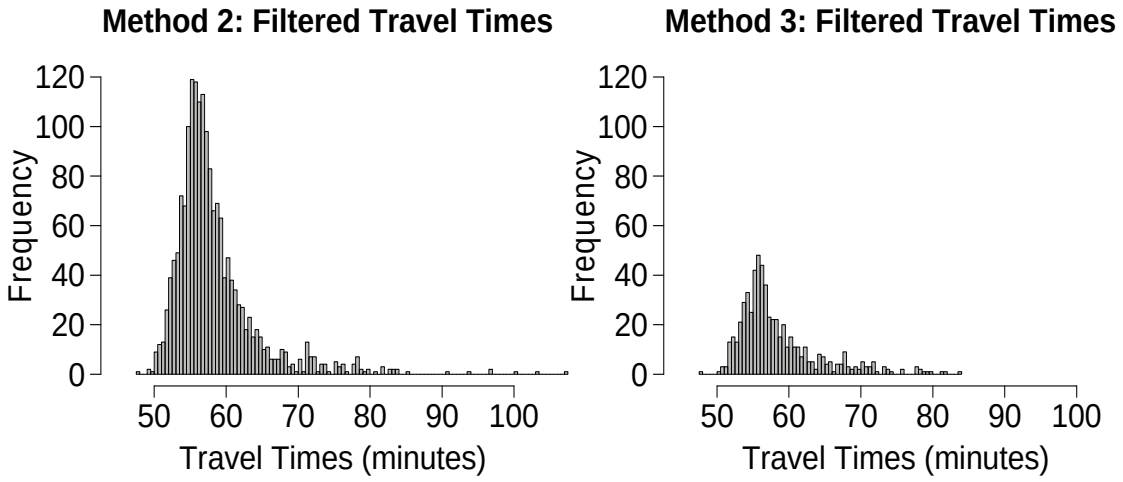

Exhibit C4-18: Filtered Travel Times (5-minute Intervals)

Table C4-8 presents a summary of the filtering results on both data sets using 5-minute and 30-minute intervals. It can be seen that Step 1, which removes outliers by day, takes out a much smaller percentage of the data from Method 2. This is because the data in Method 2 was prepared in a way such that the resulting data is grouped more closely, even though it was not prepared as carefully. For example, if a particular O-D pair contained three vehicle passage times at the origin and one at the destination, Method 3 would report the single travel time from the latest origin timestamp to the destination timestamp. Method 2, on the other hand, would report this as three separate travel times, all likely with similar magnitudes. Thus, the data in Method 3 is of higher quality, but is not as closely grouped, and it is penalized for this in filtering Step 1.

Additionally, Method 3, which has fewer points, is much more vulnerable to overaggressive filtering in Step 3 (which removes sparse intervals). This can be seen in the larger bands of dark gray in the method 3 columns of Exhibit C4-19. This is because the data sets prepared using Method 3 were much sparser initially. As a result, filtering routines that discard intervals with sparse detection may be overaggressive for sparse data sets such as those prepared by Method 3, even if the data itself is more meaningful.

Overall, data sets constructed with passage time pairing Method 2 had a higher percentage of points survive the filtering process when using both 5-minute and 30-minute intervals (see Table C4-8 and Exhibit C4-19), although both data sets performed more poorly when 30-minute intervals were used. This could be because the longer time intervals do not allow for quickly changing conditions such as weekend traffic congestion or adverse weather events.

Table C4-8: Comparison of Passage Time Pairing Methods

	5-Minute Intervals		30-Minute Intervals	
	Method 2	**Method 3**	**Method 2**	**Method 3**
Total points	5,886	4,185	5.886	4,185
Removed at step 1	3,118 (53%)	2,687 (64%)	3,118 (53%)	2,687 (64%)
Removed at step 2	117 (2%)	20 (0%)	273 (5%)	119 (3%)
Removed at step 3	915 (16%)	883 (21%)	1,272 (22%)	1,185 (28%)
Removed at step 4	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Total points removed	4,150 (71%)	3,590 (86%)	4,663 (79%)	3,991 (95%)
Remaining points	1,736 (29%)	595 (14%)	1,223 (21%)	194 (5%)
Mean after filtering	58.3 min	58.4 min	57.9 min	58.2 min
Standard deviation after filtering	5.9 min	5.7 min	3.8 min	4.2 min

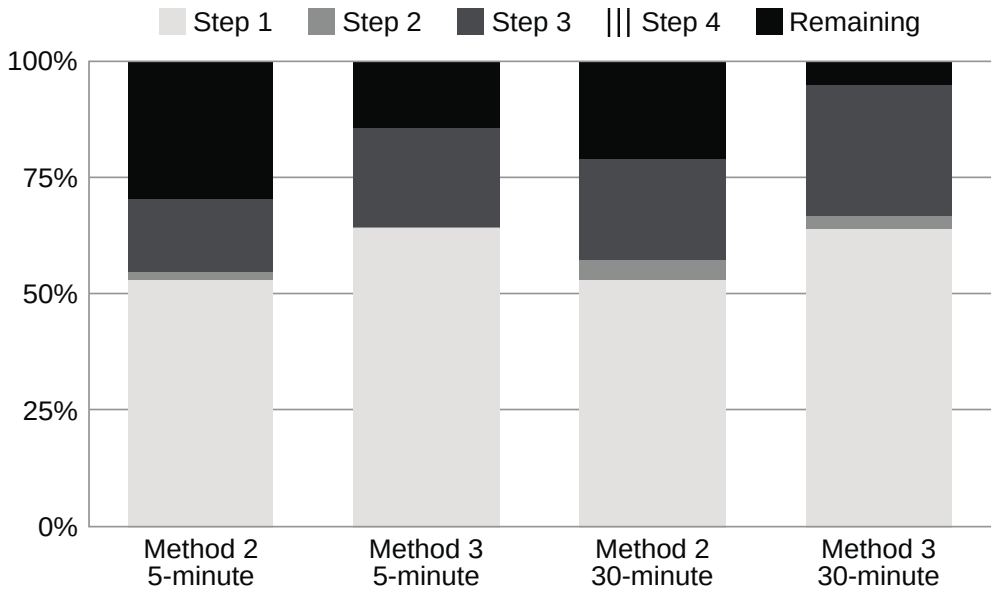

Exhibit C4-19: Proportions of Data Points Discarded

Summary

This section has evaluated various methodological approaches and processes for estimating ground-truth segment travel times based on Bluetooth data. The characteristics of Bluetooth data at each node were found to vary significantly, as a function of the surrounding roadway configuration. In cases where inter-node distances were small, the availability of signal strength was determined to be an important factor in increasing the accuracy of calculated travel times. Methods were also explored for identifying invalid segment trips, most especially via the analysis of network topology. In turn, this facilitated the generation of fewer and higher quality segment trips for use in statistical analysis.

The generation of travel time histograms used filters proposed by Haghani, et al. A comparison of two passage time pairing methods was made through histograms of filtered and unfiltered data. Potential pitfalls of using standard filtering procedures on Bluetooth data (such as discarding sparsely populated intervals) were also identified. The filtering methodology demonstrated herein was statistical in nature, in the sense that data points were discarded based on their statistical characteristics, such as coefficient of variation and distance from the mean. By comparison, passage time pairing strategies were based on the physical characteristics of the network. This exercise showed that to obtain valid travel times, knowledge of the characteristics of the network should be leveraged to the greatest extent possible, although there will still likely be a need for statistics-based filtering due to the nature of Bluetooth data.

USE CASE ANALYSIS

This case study explores the use of two vehicle re-identification technologies in support of travel time reliability monitoring within a rural setting. These data collection technologies (ETC and Bluetooth-based) work by sampling the population of vehicles along the roadway, subsequently matching unique toll tag ids or Bluetooth MAC addresses between contiguous reader stations. Their effectiveness in accurately calculating roadway travel times is dependent on a number of factors, including:

- The percentage of the total traffic steam sampled at individual readers
- The re-identification rate between pairs of readers

In general, the percentage of the vehicle population sampled by individual readers depends on the penetration rate of the technology within the vehicular population, the positioning and mounting of the reader, and the roadway configuration at the reader's location. The re-identification rate between pairs of readers can depend on all of the above factors, as well as the distance between readers and the likelihood of a trip diverting between the origin and destination reader. Since little can be done to increase the technology's penetration rate when deploying a reliability monitoring system, locating, positioning, and configuring readers to maximize their collection of quality data is crucial to the success of the system.

As this case study leveraged data generated by networks of existing data collection devices, the research team could not evaluate the process for installing and configuring detection infrastructure. However, this case study did provide the opportunity to analyze the impacts of the configuration of existing ETC and Bluetooth reader networks on the nature of the data ultimately collected for use by a travel time reliability monitoring system. Based on this concept, the team developed two network configuration-related use cases. The first use case details the findings of the research team's investigation into the configuration of the Lake Tahoe ETC network, and discusses issues including: time-of-day dependency of the toll tag penetration rate, the number of

lanes that can be monitored using ETC infrastructure, and re-identification of toll tags between readers separated by different distances. The second use case details the team's investigation into configuration-related issues associated with the Bluetooth reader network, including the relationship between reader location and the number of lanes monitored and the sample sizes measured between readers on different freeways over varying distances.

A third use case seeks to quantify the impact of adverse weather and demand–related conditions on travel time reliability using data derived from the case study's Bluetooth and ETC-based systems deployed in rural areas. To examine travel time reliability within the context of this use case, methods were developed to generate probability density functions (PDFs) from large quantities of travel time data representing different operating conditions. To facilitate this analysis, travel time and flow data from ETC readers deployed on I-80W and Bluetooth readers deployed on I-50E and I-50W were obtained from PeMS and compared with weather data from local surface observation stations. PDFs were subsequently constructed to reflect reliability conditions along these routes during adverse weather conditions, as well as according to time-of-day and day-of-week. Practical data quality issues specific to Bluetooth and ETC data were also explored.

Impact of ETC Reader Deployment Configuration on Data Quality

This first use case details the findings of the research team's investigation into the impact of the configuration of the Lake Tahoe ETC network on the quality of travel time data collected.

Introduction

In this case study, the ETC detection network consisted of eight Fastrak readers located on Interstate-80 between the eastern outskirts of Sacramento and North Lake Tahoe. For each reader, Caltrans provided us with its county, freeway, a single direction of travel, mile post, a textual location, and the IP address that could be used to communicate with it and obtain its data. To place data from this network into PeMS, the research team assigned each reader a unique ID and determined its latitude and longitude from the provided mile post. Code was then written to connect with each reader's IP address, obtaining its data once per minute for storage onto the PeMS database.

Methodology

The configuration data obtained from Caltrans was sufficient to place each reader at a location alongside the roadway. Based on that information, the team sought to answer the following questions in order to more fully understand the impacts of the network's configuration on the characteristics of the reported travel times:

1) Are the readers where they are reported to be?
2) Are any of the readers monitoring multiple directions of travel?
3) What percentage of total traffic is being detected?
4) What percentage of toll tags is matched between pairs of readers?

The first question, which addresses where the readers are located, appears straightforward, but agencies often struggle to track detection equipment in the field. This is especially problematic with vehicle re-identification technologies, which can be easily moved from location to location. While one solution to this problem is to equip readers with GPS units,

this is not common practice. The issue is compounded when multiple departments within a single agency, or multiple agencies, are using the data from these readers for their own purposes, and are not informed in a timely manner of configuration changes. To verify reader locations, the team evaluated the travel time data reported between each pair of readers to make sure that the travel times and the number of samples reported within a given time period were reasonable given the distance between the readers and the direction of travel for which they were supposed to be collecting data.

Answering the second question is important because, in some cases, ETCs can be deployed such that they monitor two directions of travel. This question was addressed by examining the roadway configuration of each reader deployment and evaluating the ETC data collected to determine whether a significant number of toll tag matches occurred between that reader and the neighboring reader in each direction of travel.

The third question addresses the "hit rate" occurring at each reader. The team calculated this by comparing hourly ETC tag reads against hourly volumes collected from nearby loop detectors. In an effort to relate mounting configuration to the percentage of traffic sampled, hit rates were subsequently compared between readers.

The final question relates to the quality of travel times being reported. As the higher the percentage of matches, the more accurate a travel time estimate is likely to be, the research team assessed the percentage of tags matched between all possible combinations of upstream and downstream ETC readers. Results from each combination of ETC readers were then compared to determine how the percentage of matches is impacted by the hit rates of each individual reader as well as the distance between readers.

Analysis

This subsection documents the process used and analysis conducted by the team to develop answers to the aforementioned four questions.

Are the readers where they are reported to be? According to Caltrans, each reader is either mounted to an overhead Changeable Message Sign (CMS) or an overhead fixed sign. Each reader consists of a cabinet mounted to the sign pole, which is connected to two antennae mounted on the edge of the sign closest to the roadway. Exhibit C4-20 shows a photograph (courtesy of Caltrans District 3), taken during installation, of the reader at the Donner Lake exit on eastbound I-80.

Exhibit C4-20: ETC Installation

Using the information provided by Caltrans, the team verified that there was a CMS or overhead sign at the latitude/longitude reported for each ETC station. Photographs of each deployment were obtained to determine each ETC's mounting configuration, its positioning over the roadway, and the roadway geometry at that location. Photographs of each reader's mounting structure, as indicated by Caltrans, are displayed in Exhibit C4-21.

The team next evaluated the minimum travel times reported between each pair of readers in order to ensure they were reasonable given the distances involved. All travel times were determined to be reasonable with the exception of trips that originated or ended at the Kingvale reader, stated by Caltrans as being located on I80-W, adjacent to the Rainbow reader on I80-E. Results of the team's analysis indicated that the Kingvale reader was actually located on I80-E, approximately 3 minutes downstream of the Donner Lake reader, which was later confirmed by Caltrans (see Exhibit C4-22).

Exhibit C4-21: ETC Locations

Exhibit C4-22: Updated Kingvale Reader Location, I80-E

Are any of the readers monitoring multiple directions of travel? The next step in understanding the impact of the various ETC reader configurations on the nature of the data collected was to determine whether any readers were capturing traffic in both directions of travel. The photographs in Exhibit C4-21 indicated that the eastbound and westbound directions of travel at the Rainbow, Rest Area, and Donner Lake reader deployments were completely separated from one another. As a result, it was not possible for these readers to monitor the opposite direction of travel. For the other readers, their ability to capture bi-directional traffic depended on the size and orientation of the detection zone generated by their antennae.

To conduct this analysis, the research team calculated the minimum and median travel times and the number of matches reported between each pair of adjacent readers monitoring opposite directions of travel along I-80. In cases where the minimum travel times reported between two readers approximated the free-flow speed given their geographic distance, and significant numbers of matches were generated that approximated that speed, the research team determined that the destination reader was likely capable of monitoring bi-directional traffic. Alternatively, if the minimum travel times were high and the number of matches low, then the matches likely represented vehicles making a round-trip (see Exhibit C4-23 for a graphical depiction of a one-way trip vs. a round trip).

For example, Exhibit C4-24 display the hourly travel times and tag matches from Friday May 27th through Saturday May 28th, 2011 between Auburn/Bell on I80-E (origin) and Prosser Village on I80-W (destination). The Prosser Village reader is 66 miles east of the Auburn/Bell reader, and is deployed in the freeway median. As indicated in Exhibit C4-23, the minimum travel times between Auburn/Bell on 80-E and Prosser Village on 80-W ranged between 60-80 minutes, which is reasonable given the 60 mile distance between them. The 75th percentile travel times are higher, likely reflecting the travel times of vehicles detected passing the Prosser Village reader on 80-W as part of a round trip after having first passed both the Auburn/Bell reader, as well as the Prosser Village reader, but not being detected by it, while traveling east. Finally, the median travel times more closely reflect the minimum travel times, indicating that the Prosser Village reader was matching more toll tags traveling past it along 80-E than were being generated based on 80-W round trips as reflected in the 75th percentile travel time. Overall, the research team's analysis indicated that only the Prosser Village reader was capable of monitoring bi-directional traffic; Caltrans later indicated that the Prosser Village reader had been deployed with antennae facing in both directions of travel.

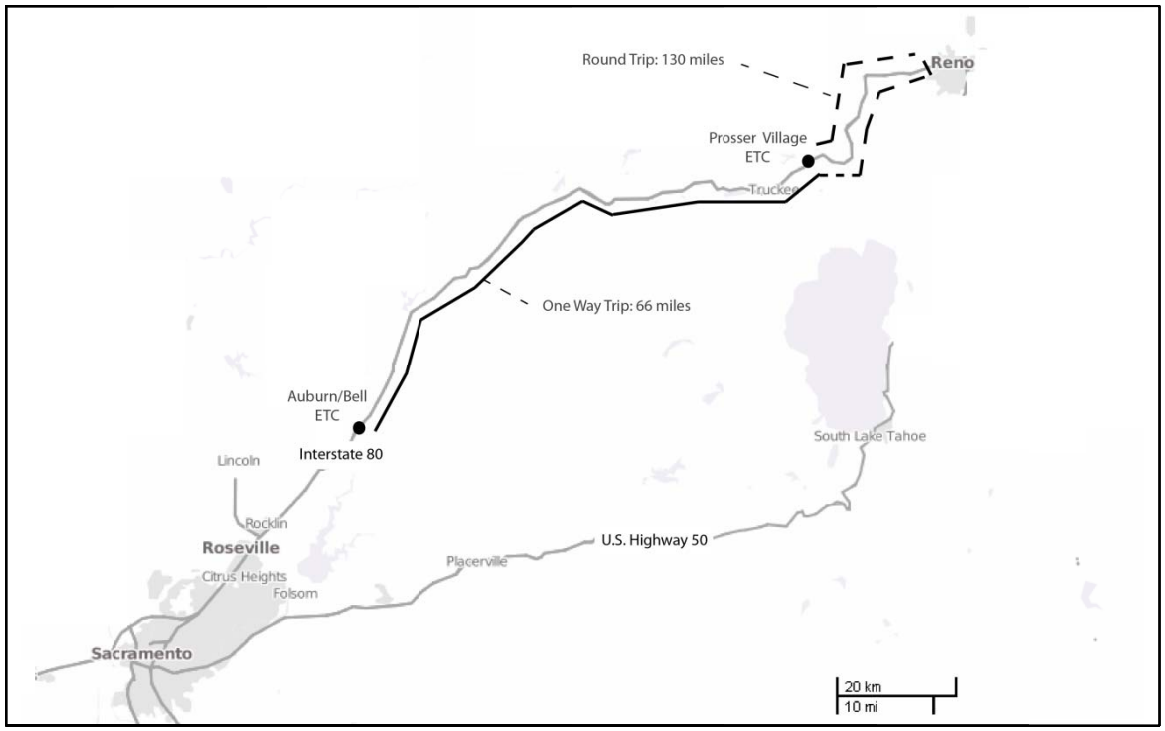

Exhibit C4-23: Graphical Depiction of a One-way Trip vs. a Round-Trip

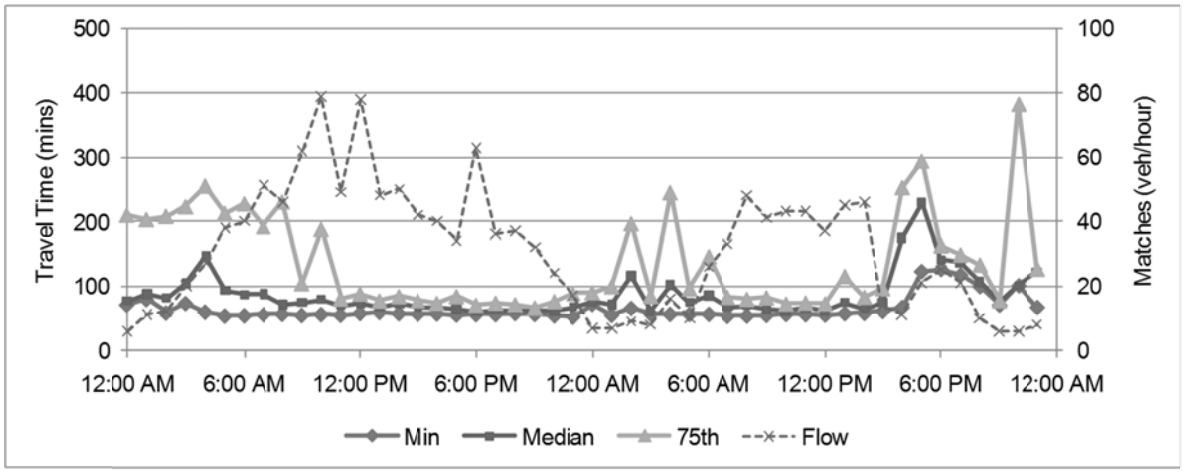

Exhibit C4-24: Travel Times Between Origin I80-E at Auburn/Bell and Destination I80-W at Prosser Village

What percentage of total traffic is being detected? As mentioned previously, the percentage of the vehicle population sampled by individual readers depends on a number of factors, including the penetration rate of the technology within the vehicular population and the positioning and mounting of the reader.

The Bay Area Toll Authority (BATA) reported in January 2011 that 53% of drivers passing through its toll plazas were equipped with FasTrak tags, with that percentage increasing to 65% during weekday peak periods (5). Even so, the ETCs in the Tahoe area are more than 100

miles from the nearest toll plaza. Consequently, the percentage of vehicles equipped with FasTrak tags depends, to a great extent, on traffic patterns between the Bay Area and Lake Tahoe.

With respect to the mounting configuration of the readers, previous ETC-based travel time data collection deployments noted that a number of configuration-related factors have the potential to impact the quantity and quality of tag reads (6). For example, when readers are positioned directly overhead, such as at tolling facilities, they reliably capture data from almost all toll tags. That said, in many real-world traffic monitoring deployments, such as in Lake Tahoe, ETC readers are placed at the side of the road, increasing their distance from vehicles and reducing the efficiency of their tag reads. Such configurations also make it more difficult for readers to capture traffic across all lanes of travel, particularly when there are multiple lanes of traffic.

To calculate the percentage of vehicles sampled at each reader, the research team compared ETC tag reads with the traffic flows measured at nearby loop detectors; the result is referred to as the *hit rate*. The hit rate at Prosser Village was not analyzed since there were no working loop detectors nearby.

Table C4-9 displays the average daily hit rates, by day of the week, for each of the ETCs along I-80, collected during the week of May 9th to May 15th, 2011. Low hit rates on Sunday and Monday were common to all of the eastbound readers, especially on the eastern end of the monitored corridor. Another trend common across all readers, though especially marked at the Auburn/Bell Road reader, was the spike in the hit rate during the overnight hours (see Exhibit C4-25). This could be due to the higher percentage of freight traffic during these hours, which may be more likely to be equipped with FasTrak tags.

Table C4-9: Average ETC Hit Rates by Day of Week (7:00 AM-8:00 PM)

Reader	Average	Sunday	Monday	Tuesday-Thursday	Friday	Saturday
Auburn/Bell Road (80-E)	3.4%	2.9%	2.6%	4.0%	4.0%	3.6%
Rainbow (80-E)	5.4%	2.9%	3.3%	6.9%	7.3%	6.7%
Rest Area (80-E)	3.2%	1.6%	1.8%	5.0%	4.0%	3.4%
Donner Lake (80-E)	6.3%	3.6%	3.7%	8.9%	8.5%	6.8%
Kingvale (80-E)	6.4%	3.9%	4.4%	9.6%	7.2%	7.0%
Hirschdale (80-W)	4.5%	4.1%	3.7%	5.6%	5.0%	4.1%
Baxter (80-W)	6.7%	6.2%	6.0%	8.1%	7.1%	5.9%

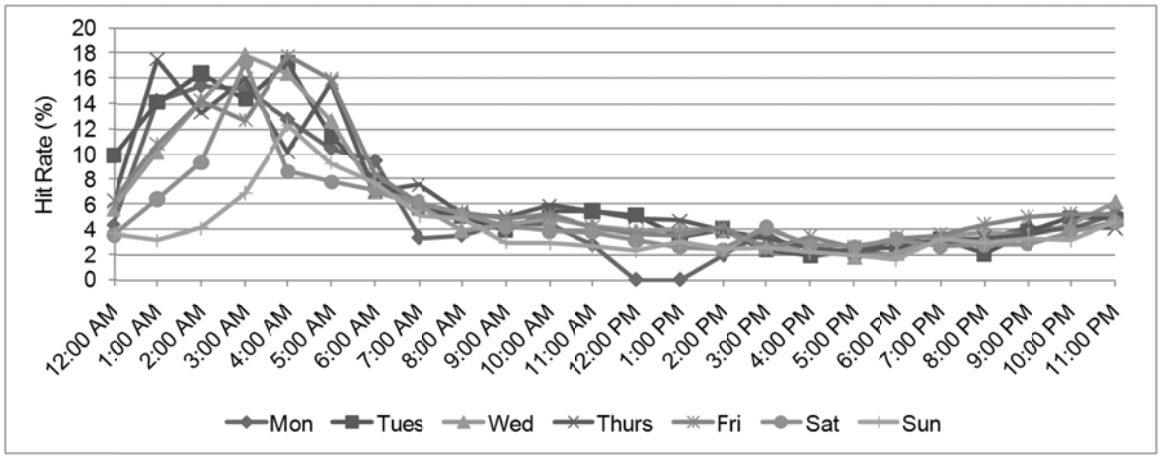

Exhibit C4-25: Hourly Hit Rates, Auburn/Bell Road Reader

Comparing average hourly hit rates for all of the readers on I-80-E from Tuesday through Friday makes it clear that some readers are sampling a significantly higher percentage of traffic than others. An examination of photographs of the signs onto which each reader was mounted provides no clear explanation for why the hit rates at some readers are approximately double those at other readers. The hit rate at Auburn/Bell Road may be lower because it is the only location with three lanes of travel (all other readers only monitor two) being monitored by two antennae. The Rest Area reader, though appearing to be optimally positioned above of the roadway, also has a low hit rate. Another possible reason for the low hit rate is that the antennae here are not properly aligned with the two lanes of travel, resulting in a reduced number of toll tag reads.

To gauge the sampling rate in another way, we also looked at the raw number of hourly tag reads reported by each reader, the results of which are displayed in Exhibit C4-26 for the eastbound direction of travel. Despite its low percentage of tag reads, the reader at Auburn/Bell Rd. still records a large number of reads, simply due to the fact that traffic volumes are higher here than at any other reader. The highest number of reads recorded across readers is on Friday, due to the recreational pattern of weekend trips to Lake Tahoe.

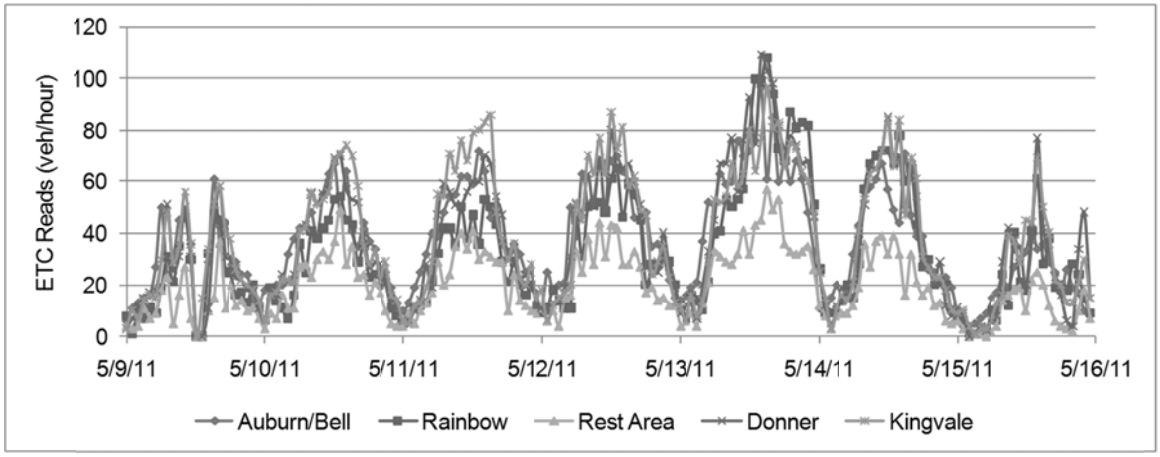

Exhibit C4-26: Hourly Tag Reads, EB I-80 Readers

What percentage of toll tags is matched between pairs of readers? For the purpose of calculating travel time reliability-related metrics it is most important to have the ability to quantify the typical percentages and volumes of toll tags matched between multiple readers, as this directly impacts the quality of aggregated travel times. To quantify typical tag match rates between readers, the team looked at the percentage of vehicles being matched between the furthest upstream readers (Auburn/Bell Road in the eastbound direction and Hirschdale Road in the westbound direction) and all subsequent downstream readers between May 9, 2011 and May 15, 2011 (see Exhibit C4-1 for the deployment layout).

Exhibit C4-27 shows the percentage of toll tags detected at the Auburn/Bell reader that are re-identified at each downstream 80-E reader (ordered from left to right by distance from origin). If each reader's data collection capabilities, and therefore their hit rates, were identical, we would expect to see the percentage of matched tag reads decrease with distance from the origin reader as vehicles detected at Auburn/Bell Road deviated from 80-E. However, this trend does not hold for these readers. Instead, the highest matching percentage (91%) is seen between Auburn/Bell Rd. and the Kingvale reader, which are separated by a distance of 50 miles. At the same time, matches between Auburn/Bell Rd. and the Rainbow and Rest Area readers are much lower, which is consistent with the low hit rates measured at these three stations.

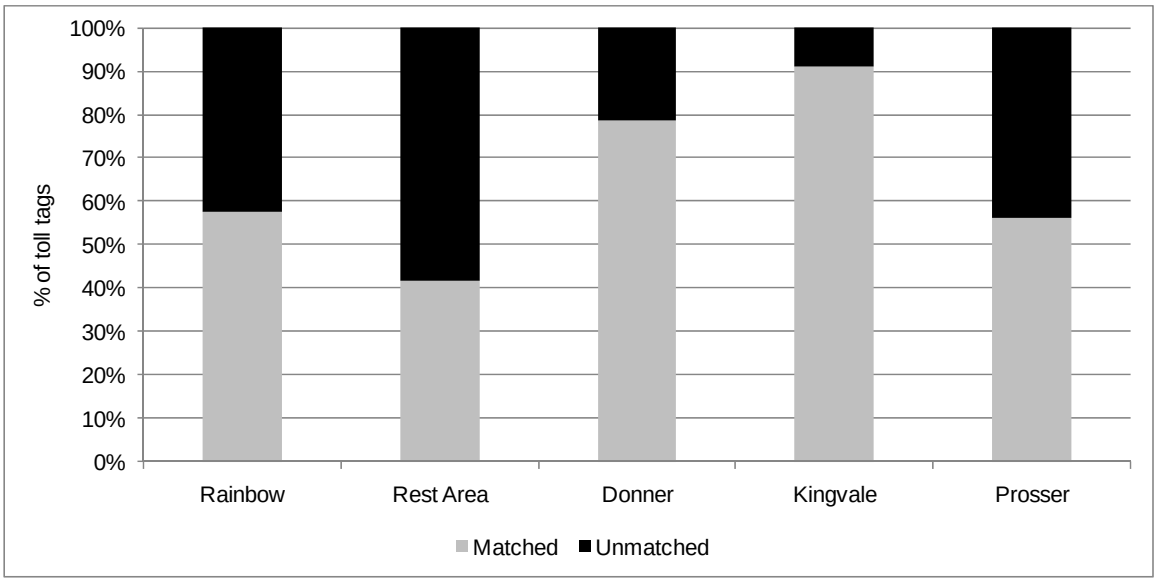

Exhibit C4-27: Average Percentage of Toll Tags Matched on EB I-80 From Origin Auburn/Bell

Exhibit C4-28 displays the total number of hourly matches measured on eastbound I-80 between the origin reader at Auburn/Bell Rd. and all downstream readers. At all readers, matches are the highest on Fridays, which is supported by local traffic patterns between the Bay Area and Lake Tahoe. Again, despite being the second furthest reader from the origin, the Kingvale reader often sees the most matches on eastbound I-80. Daytime matches between Auburn/Bell and Kingvale generally exceed 30 per hour (3 per five minutes or eight per fifteen minutes). While this number of samples is likely too low to compute accurate five-minute average travel times,

this data has the potential to be used to generate average fifteen-minute or hourly statistics throughout the week.

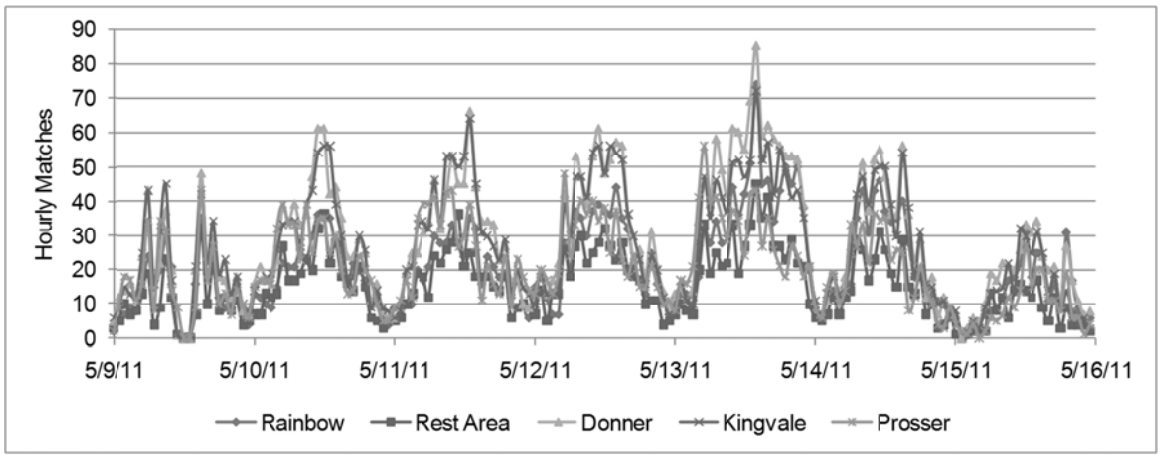

Exhibit C4-28: Hourly Toll Tag Matches, I80-E; Origin Reader at Auburn/Bell Rd.

Findings

There are two primary variables that impact hit rate: (1) the total number of ETC tags in the population of vehicles that pass a reader; and, (2) the number of tags actually read by a specific reader. The product of these factors has a significant influence on the accuracy of travel time data generated between any two ETC readers. With that in mind, this use case has sought to demonstrate how the hit rates and matching percentages generated by the Tahoe region ETC network may be impacted by the configuration of individual ETC readers. Table summarizes the configuration and average data collection results for each of the single-directional ETC readers deployed along I80-E during Friday afternoons/evenings (12:00 – 8:00 PM); considered the peak period for this roadway due to weekend traffic between the Bay Area and Lake Tahoe. All of the readers included in this table were deployed along I-80, are mounted on overhead signs, and are monitor two or three (in one case) lanes of traffic. Despite these readers being deployed under what would appear to be such similar conditions, this table's content indicates that there are a number of differences in the percentage of the traffic stream sampled at the different locations that are worth noting.

To begin with, hit rates (% of Traffic Sampled) for the Donner Lake and Rainbow readers are more than twice those generated by the Auburn/Bell Road and Rest Area readers. Although the team was not able to investigate the underlying reason for these differences, we believe they are most likely the result of:

- Reader antennae being misaligned at some locations.
- The reader at Auburn/Bell Road attempting to collect data from three lanes of traffic using only two ETC readers.

As seen in the table, differences in hit rate of only 2-3% can make a significant difference in the number of tag reads collected, which is crucial for ensuring that a sufficient number of samples are re-identified downstream to generate accurate travel times and travel time distributions.

As expected, the hit rate for an individual reader has a profound impact on that reader's ability to re-identify vehicles initially detected at upstream readers. For example, as shown in Table C4-10, even though the Auburn/Bell Road reader is 45 miles and 24 exits from the downstream reader at Rainbow, the high hit rate at this downstream reader enables it to re-identify 83% of vehicles initially detected at Auburn/Bell Road. Given the number of opportunities to exit the freeway, this likely represents nearly all of the ETC-equipped vehicles that pass between the readers. Conversely, despite the fact that the Rest Area reader is only 8 miles from the Rainbow reader with only one exit ramp in between, it is only able to match 42% of vehicles initially identified at Rainbow. Overall, at least on rural roads experiencing fairly significant through traffic, the readers' hit rates appear to impact the percentage of matched vehicles to a greater extent than the distance between the readers.

However, even with ideal reader placement and configuration, the primary constraint on the percentage of traffic sampled will always be the penetration rate of toll tags in the population. In rural areas, it is uncommon to have electronic tolling infrastructure, so deploying ETCs in these locations requires that at least some portion of the traffic stream be composed of vehicles equipped with toll tags used in nearby urban areas. The results of this use case show that this penetration rate can vary by time of day and day of week; for example, on I80-E, far fewer Fastrak-equipped vehicles travel the corridor on Sundays and Mondays than during the rest of the week.

Table C4-10: I-80 E ETC Reader Summary, Fridays, 12:00 PM-8:00 PM

Reader	Mounting Type	Lanes	Tag Reads	% of Traffic Sampled	Distance to Next Reader (mi.)	Exits Between Readers	% Hits Re-identified Downstream
Auburn/ Bell Rd	EB Roadside VMS	3	648	3.5%	45	24	83%
Rainbow	EB Roadside VMS	2	789	7.4%	8	1	42%
Rest Area	EB Roadside Sign	2	380	3.6%	4	1	99%
Donner	EB Roadside Sign	2	785	7.4%	3	2	96%
Kingvale	EB Roadside Sign	2	696	6.6%	6	--	61%

Impact of Bluetooth Reader Deployment Configuration on Data Quality

This use case details the findings of the research team's investigation into the impact of the configuration of the Lake Tahoe Bluetooth network on the quality of travel time data collected.

Introduction

The Bluetooth reader network leveraged in this case study was deployed along Interstate 5 (I-5) in Sacramento and US 50 between Placerville and Lake Tahoe. For each BTR, the research team received configuration data in a .CSV file, with fields for the node ID, a textual location, and a latitude/longitude. The research team was also provided with a 2-gigabyte SQL file containing all of the Bluetooth data collected at the readers between December 25, 2010 and April 21, 2011; this use case only utilizes the eight BTRs that provided more than a week's worth of data. This data was downloaded into PeMS for use in computing travel times between each BTR pair.

Methodology

In evaluating the impact of the network's configuration on the characteristics of the reported travel times, the team sought to answer questions of a similar nature to those explored as part of the ETC use case, including:

1) Are the readers where they are reported to be?
2) Which facilities is each reader monitoring?
3) What percentage of total traffic is being detected?
4) What percentage of Bluetooth devices is matched between pairs of readers?

The first question was particularly important for this use case as the BTRs were deployed as part of a test, and not as permanent data collection infrastructure. As a result, each BTR changed locations multiple times over a span of several months. To compute accurate travel times, the team had to ensure that the locations provided in the configuration file matched the data delivered in the SQL file. This was achieved by mapping the latitude and longitude provided by Caltrans to determine whether the data matched the textual locations provided.

Answering the second question is critical to all Bluetooth studies. Class I Bluetooth devices, like the ones used in this case study, have a detection radius of 300'. As a result, the potential exists for the BTRs to monitor bi-directional traffic along a roadway, as well as traffic along parallel facilities, which presents challenges when trying to compute accurate travel times. As such, the research team evaluated the reader locations and data to approximate the lanes of travel they each monitored, whether they were monitoring traffic bi-directionally, whether they might also be detecting vehicles on on-ramps, off-ramps, or frontage roads, and whether the potential existed for them to capture data concerning the movement of other modes of travel, such as from bicyclists.

The third question addresses the "hit rate" occurring at each reader. The team calculated this by comparing hourly BTR reads against hourly volumes collected from nearby loop detectors. In an effort to relate mounting configuration to the percentage of traffic sampled, hit rates were subsequently compared between readers.

The final question relates to the quality of travel times being reported. As the higher the percentage of matches, the more accurate a travel time estimate is likely to be, the research team assessed the percentage of Bluetooth devices matched between all possible combinations of upstream and downstream BTRs. Results from each combination of BTRs were then compared to determine how the percentage of matches is impacted by the hit rates of each individual reader, as well as the distance between readers.

Analysis

This subsection documents the process used and analysis conducted by the team to develop answers to the aforementioned four questions.

Are the readers where they are reported to be? Using the information provided by Caltrans, the team verified the location of each BTR according to both its latitude/longitude and textual description. While a number of the readers represented in the configuration file were erroneously located (for example, the latitude/longitude of one placed it in a lake), the eight readers used as part of this use case all appeared to be in roughly the correct location. Photographs of each BTR station used as part of this use case are displayed in Exhibit C4-29. One BTR, deployed on US-50 at Echo Summit is not visible as a result of being buried in snow. Despite this, the team was able to use the data collected from this station as part of its analysis.

As a final location confirmation step, the team evaluated at the minimum travel times computed between each BTR to ensure they were reasonable given the distances involved. All minimum travel times were subsequently determined to be reasonable, and the BTR locations deemed accurate.

Which facilities is each reader monitoring? The next step in understanding the impact of each BTR's configuration on the nature of the data collected was to determine which readers might be capturing traffic data for multiple directions of travel. The BTRs evaluated as part of this use case were deployed as follows:

- Three of the readers on US-50 were deployed in locations where there is one lane of travel in each direction.
- The reader at US-50 and Placerville monitored two lanes in each direction.
- The reader at US-50 and Meyers was located near an intersection that might result in it picking up MAC addresses from vehicles turning onto US-50 from a cross street.
- The reader at I-5 and Vallejo potentially monitored up to five lanes in each direction.
- The reader at I-5 and Gloria (the only BTR along I-5 located on southbound side of the roadway) had the potential to monitor up to four lanes of travel in each direction.
- The reader at I-5 and Florin was located in the middle of the clover-leaf on-ramp of Florin Road to I-5 North. It was adjacent to four mainline northbound and southbound lanes. Given the reader's location, it was likely detecting significant numbers of vehicles entering and exiting I-5, both traveling at slower speeds and being detected earlier (for on-ramp vehicles) or later (for off-ramp) than if they were actually traveling on I-5.
- The reader at I-5 and Pocket was located some distance from the northbound side of the roadway. It had the potential to monitor two on-ramp lanes to I5-N, three mainline lanes in each direction, and a clover-leaf off-ramp from I5-S.

Based on this analysis, the research team concluded that all BTRs were likely monitoring at least bi-directional traffic, a conclusion that was confirmed by the data analysis conducted in support of the following subsection. This effort also provided some insight into how the reader's locations have the potential to impact the sampling of non-representative trips.

What percentage of total traffic is being detected? As with the ETCs, the percentage of the traffic stream monitored by a Bluetooth reader depends on the penetration of Bluetooth-enabled devices within the vehicle population. Although it is estimated that 20% of travelers now have Bluetooth devices with them in their vehicles, at least a quarter of them do not have the device set to discoverable mode.

The detection rate also depends on the reader's configuration. Class I Bluetooth readers have a 300' detection radius. Based on this, a single BTR could easily monitor all lanes of a freeway that has four lanes of traffic in each direction of travel and is barrier-separated. That said, it might also collect a number of undesired samples, such as Bluetooth devices on parallel facilities or within office buildings. All readers used in this case study had approximately the same average signal strength, so this variable was not a factor.

To calculate the percentage of vehicles sampled at each BTR, the research team compared Bluetooth mobile device reads with the traffic flows measured at nearby loop detectors; the result is referred to as the *hit rate*. Hit rates were computed for the four readers on I-5 (there were no working loop detectors near the US-50 readers). Because all readers were presumed to monitor both directions of travel along I-5, and as it is impossible to assign a direction of travel to unmatched Bluetooth reads, hit rates were calculated by comparing hourly detections at each reader with hourly volumes summed up from nearby northbound and southbound loop detectors over a week-long period (Monday, February 28th to Sunday, March 6th, 2011). In addition, because the Florin and Pocket readers clearly detected traffic on roadway on-ramps, hit rates at these readers were computed by comparing the hourly reader detections with the hourly volumes summed up from the mainline and on-ramp loop detectors (so as not to upwardly bias the hit rates at these readers).

Bluetooth hit rates were first evaluated to determine if they exhibited any time of day or day of week patterns. As with the ETC readers, hit rates were lowest during the early morning hours. There were no other discernible patterns. Exhibit C4-30 compares the hourly hit rates measured over three days (Tuesday-Thursday) across the four readers. The hit rates at all readers generally ranged between 6% and 10%. Hit rates were usually highest at the Gloria reader, which was directly adjacent to the southbound lanes; hit rates between 8% and 10%. The reader at Pocket typically had the lowest hit rates, between 6% and 8%, possibly due to its setback from the roadway.

raph

1

Exhibit C4-29: Bluetooth BTR cabinet locations

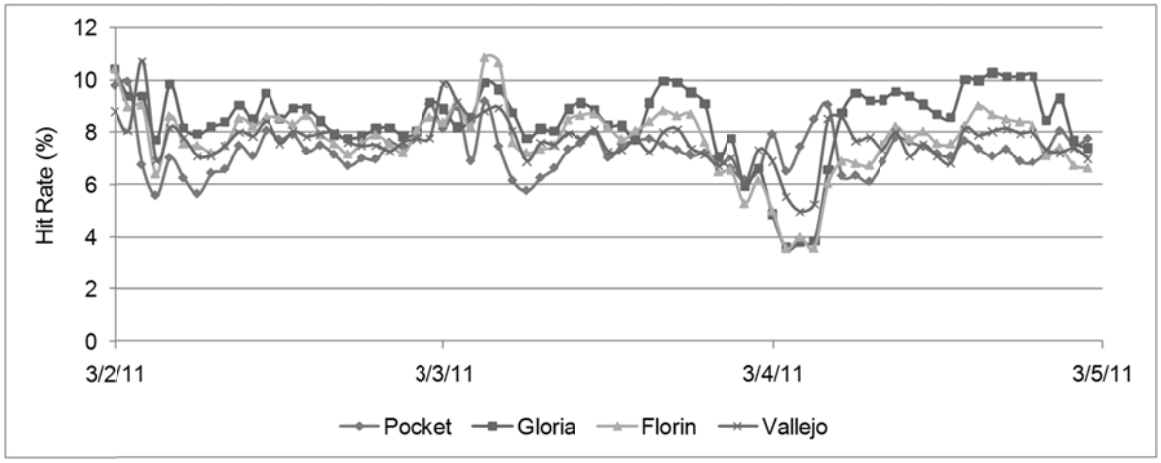

Exhibit C4-30: Hourly Hit Rates, I-5 Bluetooth Readers

The data displayed in Exhibit C4-31 presents the raw number of hourly MAC address reads at each reader listed. The reads shown in this plot are based on the number of MAC address reads remaining following filtering to remove duplicate IDs at the same reader during the same hour. While the reader at I-5 and Vallejo does not have the highest hit rate, it generally records the largest number of MAC addresses per hour, reaching nearly 1,000 reads per hour during the weekday PM peak. In contrast, the reader at I-5 and Pocket has both the lowest hit rate and the lowest number of reads, with between 500 and 600 MAC address reads per hour during the peak hours and only 300 to 400 per hour during the midday.

While the hit rate could not be computed for the readers on US-50, the research team did evaluate the raw number of hourly MAC address reads at each reader on this road (see Exhibit C4-32). The pattern of reads on US-50 differs from that on I-5, which follows the more typical AM/PM peak commute pattern. On US-50, each reader detects the most reads on Fridays, Saturdays, and Sundays, due to recreational traffic patterns near Lake Tahoe. At the Meyers, Echo Summit, and Twin Bridges readers, which are all relatively closely spaced (within 12 miles of one another) near South Lake Tahoe, the number of hourly reads are fairly similar, and are quite low (30-50 per hour, or 2 to 4 per five minutes) from Monday through Thursday. The number of reads at the Placerville reader, which is closer to Sacramento, are higher, especially during the work week, when this location has higher traffic volumes.

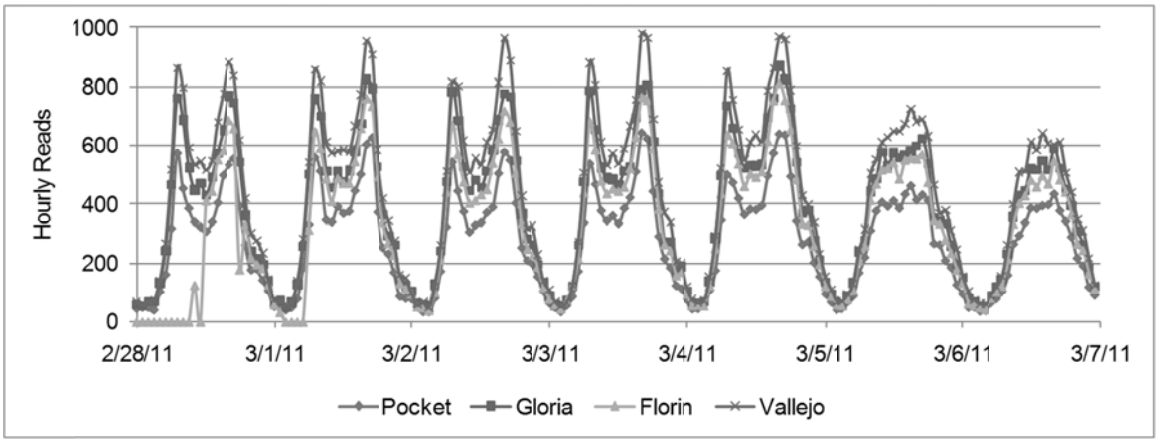

Exhibit C4-31: Hourly Reads, I-5 Bluetooth Readers

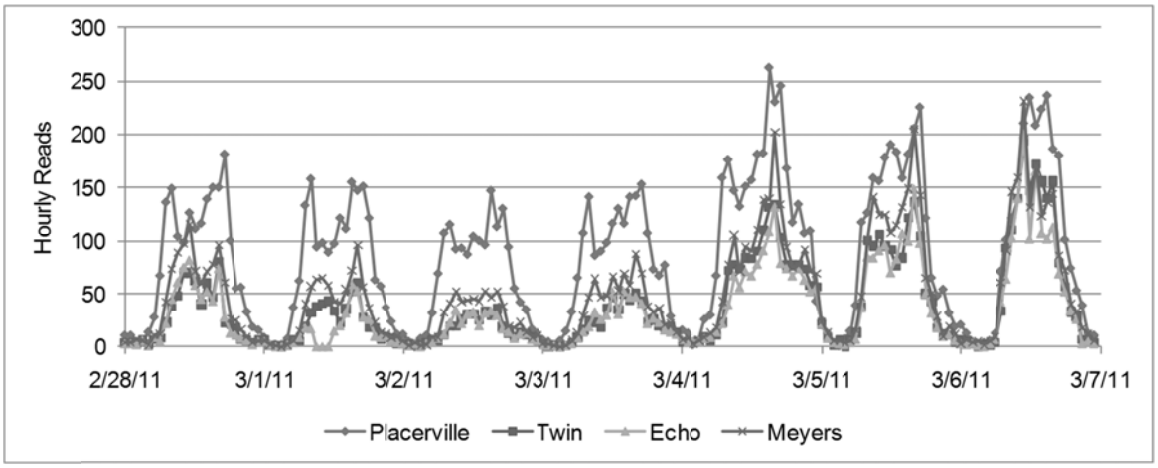

Exhibit C4-32: Hourly Reads, US-50 Bluetooth Readers

What percentage of Bluetooth devices is matched between pairs of readers? For the purpose of supporting the calculation of travel time reliability-related metrics it is most important to have the ability to quantify the typical percentages and volumes of Bluetooth devices matched between multiple readers, as this directly impacts the quality of aggregated travel times. The first step in performing this analysis was to evaluate the percentage of each reader's Bluetooth MAC address reads re-identified at downstream readers. Results for the readers along I-5 are shown in Exhibit C4-33 and for the US-50 readers in Exhibit C4-34.

On I-5, the Vallejo (northern-most) and Pocket (southern-most) readers only have downstream readers in one direction of travel (see Exhibit C4-1 for the deployment layout). Re-identification of devices between these readers occurred as follows:

- For the Vallejo reader, approximately 42% of its MAC address reads were re-identified at the Gloria reader located about 4 miles to the south.
- For the Gloria reader, about 48% of its reads were re-identified in the north-bound direction at Vallejo, while 50% of its reads were re-identified in the south-bound direction at Florin; 2% were not re-identified at all.
- For the Florin reader, 53% of reads were re-identified in the northbound direction at Gloria, while 39% were re-identified in the southbound direction at Pocket; 8% were not re-identified in either direction.
- For the Pocket reader, 48% of reads were re-identified in the northbound direction at Florin.

Overall, the rate of matching between readers was very high, with the vast majority of Bluetooth devices matched at another sensor for use in generating travel times.

Re-identified rates were also high between the readers along US-50, particularly the three deployments closest to Lake Tahoe. Re-identification of devices between these readers occurred as follows:

- For the Meyers (eastern-most) reader, for which there is no downstream reader in the eastbound direction, 50% of reads were re-identified at the Echo Summit reader, four miles to the west.
- For the Echo Summit reader, 57% of reads were re-identified at Meyers and 43% were re-identified at Twin Bridges, 8 miles to the west. Virtually none of the reads captured at Echo Summit went unmatched, likely due to its location at a point on the

roadway that has no parallel facilities, and fact that there were few possible exits between Echo Summit and Meyers or Echo Summit and Twin Bridges.

- For the Twin Bridges reader, 40% of reads were re-identified to the east at Echo Summit and 47% were re-identified at Placerville, 39 miles to the west. 13% of reads captured at Twin Bridges are not re-identified.
- For the Placerville (western-most) reader, 22% of reads were re-identified downstream at Twin Bridges.

Based on these high re-identification rates, the team concluded that the readers on US-50 were capable of detecting and re-identifying a very high proportion of the Bluetooth devices that pass through their detection zones, likely due to the narrow roadway width at these locations and the limited options available to exit the roadway.

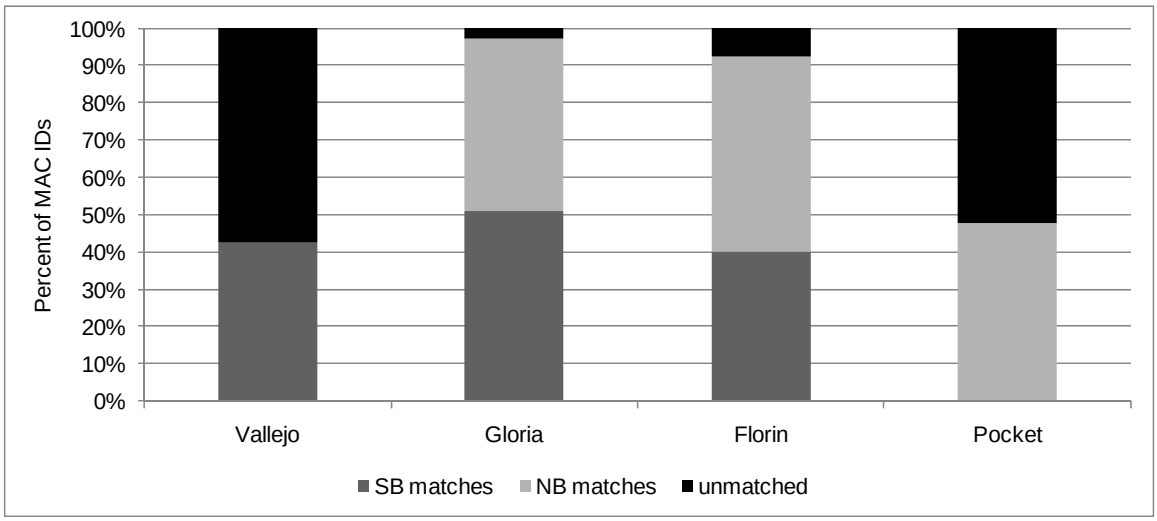

Exhibit C4-33: MAC address matching rates, I-5 readers

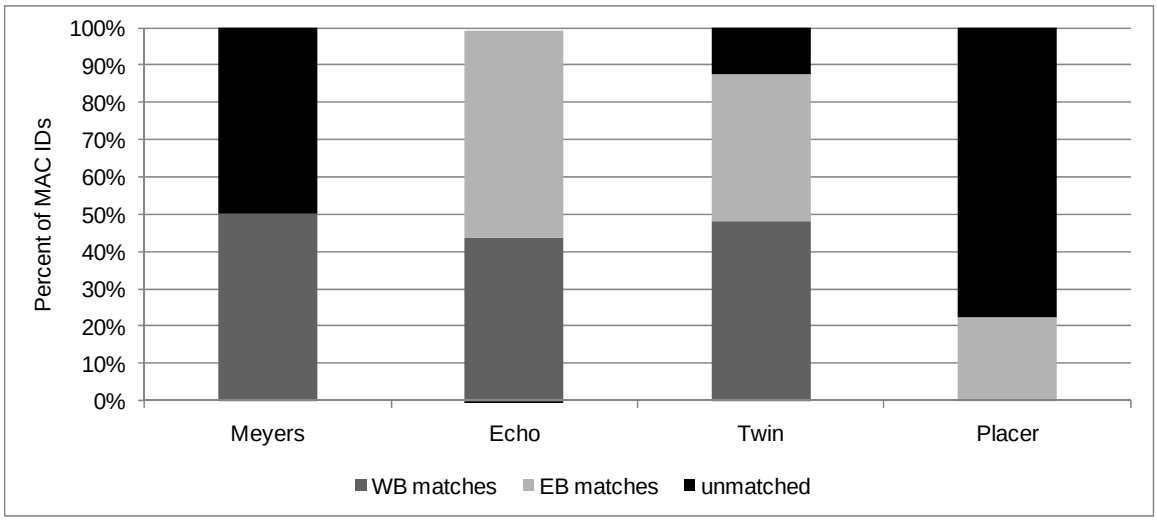

Exhibit C4-34: MAC address matching rates, US-50 readers

The next technique for evaluating Bluetooth device reidentification between readers was to examine the raw number of matches between readers in order to assess whether the match volumes were sufficient to yield accurate average travel times. This was carried out by selecting an origin reader and computing the hourly matches to a series of destination readers. Exhibit C4-35 displays the results of this analysis in the southbound direction along I-5 for the Gloria, Florin, and Pocket readers from the origin reader at Vallejo. Highlights of this analysis included:

- As the team expected, the greatest number of matches occurred with the closest downstream reader, Gloria, and the fewest matches with the reader furthest away, Pocket. These matches differed by about 100 during each PM peak hour, representing a difference of 25%.
- Matches between Vallejo and all downstream readers averaged about 16 per five-minutes during daytime hours, likely sufficient for obtaining five-minute travel times.
- The number of matches peaked during the PM period, at around 350 to 500 per hour, when travelers were departing Sacramento for its southern suburbs.
- Volumes were lower on weekends, but still sufficient to support average travel time computations at a fine granularity.

Exhibit C4-36 displays results of this analysis for hourly northbound matches at the Florin, Gloria, and Vallejo readers from the origin reader at Pocket. Highlights of this analysis included:

- As Pocket had the lowest hit rate of the readers on I-5, a smaller percentage of vehicles were available for re-identification when using this reader as an origin.
- The numbers of matches at each of the three destination readers were very similar, and generally differed by less than 25 per hour, representing a difference of about 10%.
- The number of matches peaked during the AM period, at around 350 to 400 per hour, when the majority of traffic was commuting north to Sacramento.
- As in the southbound direction, matches were lower on Saturdays and Sundays, but still likely sufficient to calculate fine-grained average travel times.

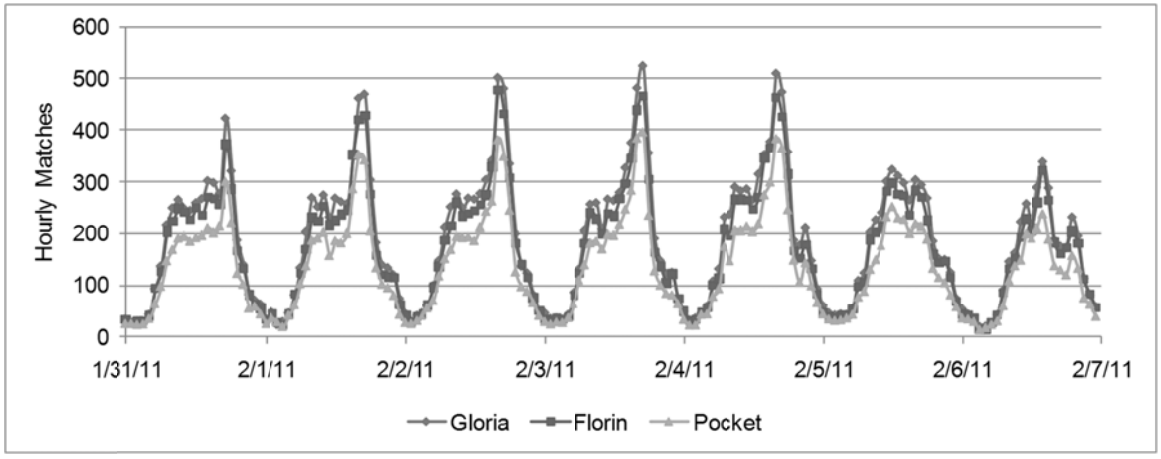

Exhibit C4-35: MAC address matches, I-5 South from Vallejo reader

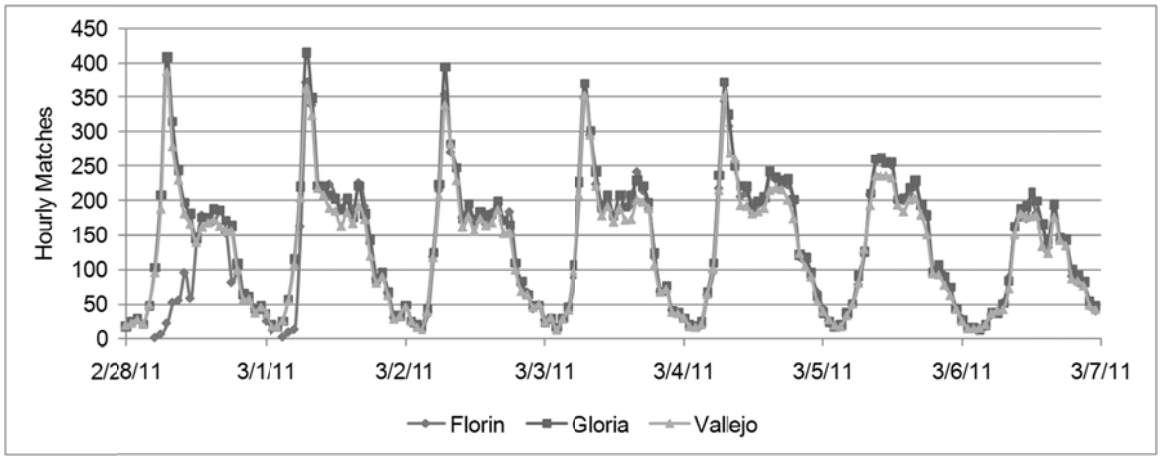

Exhibit C4-36: MAC address matches, I-5 North from Pocket reader

Results for the number of hourly matches between the Meyers reader (eastern-most) and downstream readers on US-50 are displayed in Exhibit C4-37. Highlights include:

- Although the number of matches decreased with distance from the origin reader, the number of matches was similar between the three destinations due to the fact that a significant amount of traffic on US-50 travels its entire length from Lake Tahoe to Sacramento.
- The number of matches was much lower than along I-5, likely due to its rural characteristics and lower traffic volumes.
- The number of matches was highest on Saturdays and Sundays, due to recreational traffic, peaking on Sunday afternoons, when travelers are returning from Lake Tahoe to Sacramento and the Bay Area.
- During the peak hours on Sunday, there are 100 to 140 hourly matches (8 to 12 per 5 minutes or 25 to 35 per 15 minutes) between the Meyers reader and the Placerville reader, 50 miles away, likely sufficient to calculate 15-minute travel times, and possibly 5-minute travel times, for this facility's peak hour.
- During the rest of the week, the number of hourly matches ranged from around 20 to 50 during the daylight hours (2 to 4 per five minutes or 5 to 12 per fifteen minutes). This number of matches is not likely sufficient to compute average travel times every five-minutes, though it might be used to compute fifteen-minute or hourly average travel times.

The number of hourly matches for traffic between the origin reader at Placerville (western-most) and the destination readers at Twin, Echo, and Meyers, is shown in Exhibit C4-38.

- Matched peaked on Fridays and Saturdays as vehicles traveled from the Bay Area and Sacramento to Lake Tahoe.
- During the peak hours on Friday and Saturday afternoons, matches between the Placerville reader and the Meyers reader, near South Lake Tahoe, were around 100 per hour (8 per five minutes or 25 per 15 minutes), likely enough to compute average travel times at a five-minute or a 15-minute granularity. During the rest of the week, however, there are only about 20 matches per hour.

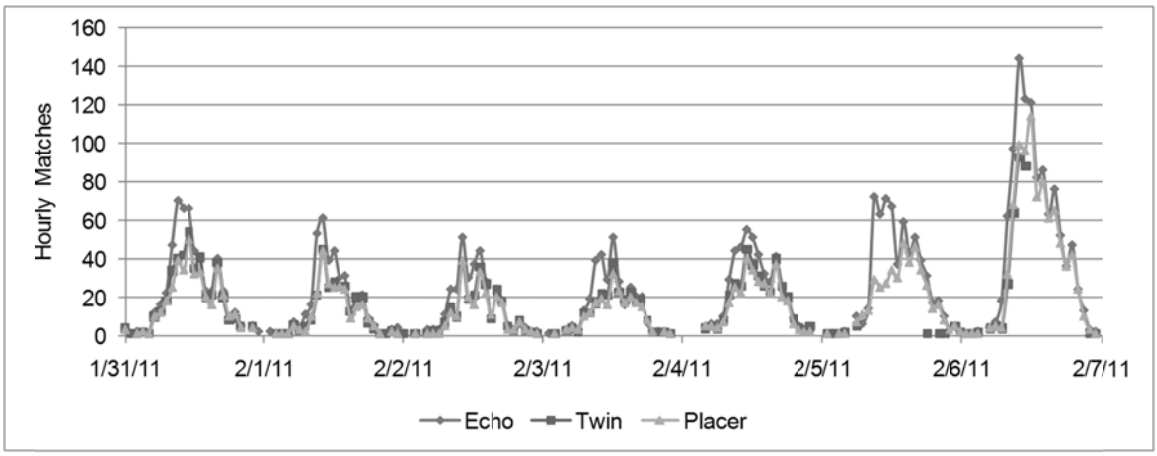

Exhibit C4-37: MAC address matches, US50-W from Meyers reader

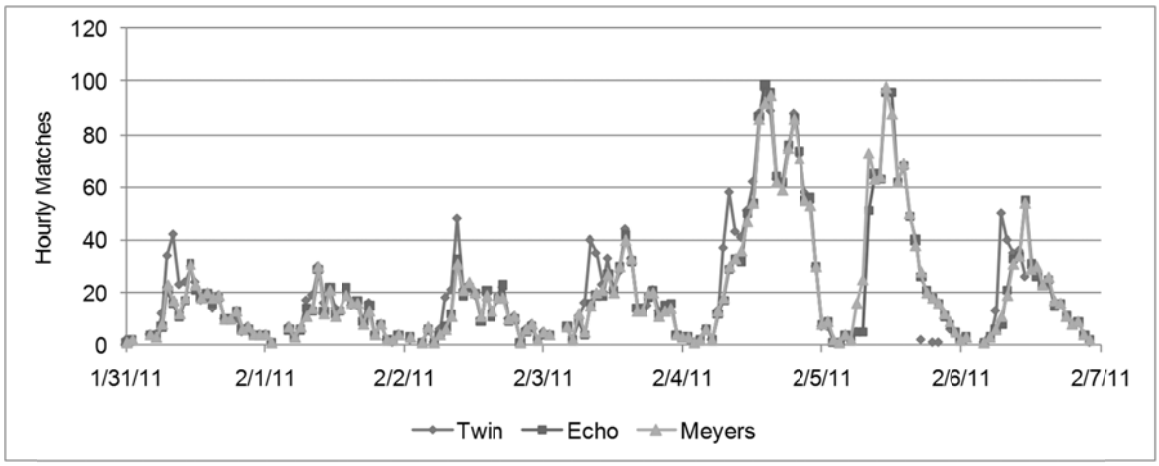

Exhibit C4-38: MAC address matches, US50-E from Placerville reader

As a number of travelers make trips between Sacramento and Lake Tahoe, there is potentially value in knowing the travel times between the two. For this reason, the team also examined the number of hourly matches between the readers along I-5 and the readers on US-50. Exhibit C4-39 shows the number of matches between the reader at Meyers (closest to South Lake Tahoe) and other readers along I-5. As these readers are on different freeways and are about 100 miles apart, the key question is whether there are sufficient matches to compute travel times at any level of granularity. Exhibit C4-39 displays the results of this analysis, representing trips along US-50W, exiting onto I-5S.

- The peak number of matches occurred on Sunday afternoon, when some hours had up to 16 - 18 matches. However, even during this peak, there were some hours when the number of matches dipped to only 5 per hour. Consequently, it does not appear possible to consistently calculate travel times at a fine granularity even on, Sunday afternoons. However, there are sufficient matches to compute hourly travel times, which could provide a reasonable indication of travel time reliability for who want to make this return trip from Lake Tahoe.

- During the rest of the week, there were insufficient matches to compute accurate average travel times by time of day and day of week, though travel times could be collected over a period of many weeks to compute average travel times and travel time variability.

Exhibit C4-40 displays the hourly matches between the Pocket reader (southern-most) on I-5 and each of the readers along US-50. These matches most likely represent vehicles traveling north on I-5 towards Sacramento, and then exiting onto US-50 E in the direction of Lake Tahoe. The number of matches peaked at between six to 12 per hour during Friday afternoon, and was also higher on Saturday morning, at around 8 per hour. Matches during the rest of the week were lower, but could potentially be studied over time to better understand the variability of travel times by time period.

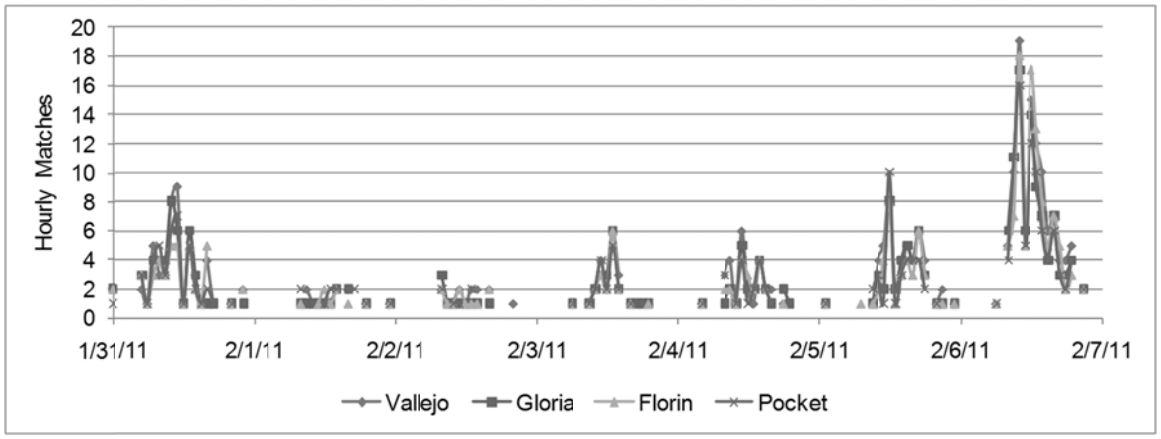

Exhibit C4-39: MAC address matches, I5-S from Meyers on US-50

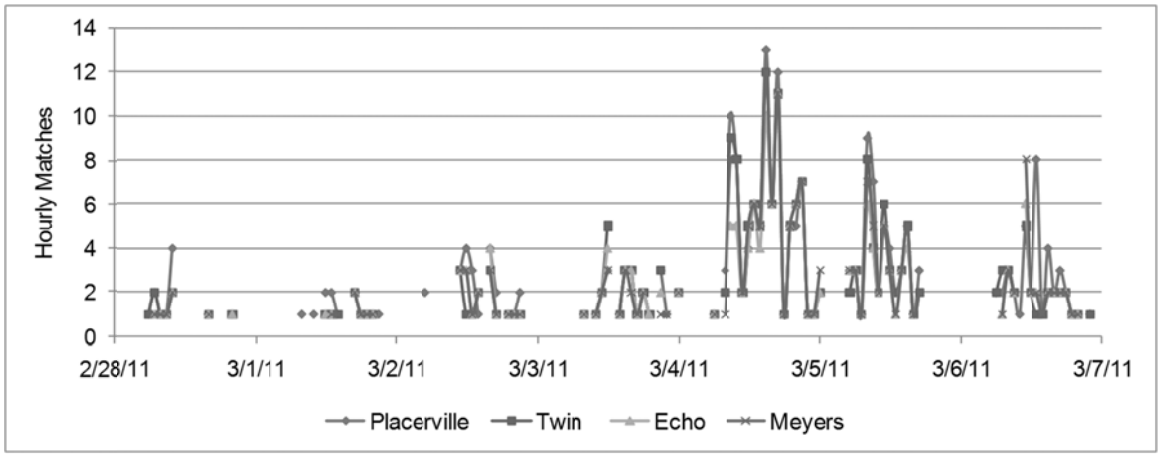

Exhibit C4-40: MAC address matches, US50-E from Pocket reader on I-5

Findings

This use case provided the opportunity to assess the performance of a Bluetooth-based travel time monitoring system deployed in an urban environment with that deployed in a rural

environment, while simultaneously demonstrating how sensor configuration impacts both the amount and quality of data collected.

One overarching finding of this use case is that the potential exists to use Bluetooth readers to generate travel times over long distances between urban and rural settings based on travel along adjoining roadways; though this is heavily dependent on the presence of the right set of conditions. For example, as indicated in Table C4-11, during an average Friday afternoon/evening, 132 vehicles detected at the Vallejo reader on I-5 (2% of the Bluetooth reads at this location) are later re-identified at the Placerville reader on US 50, more than 46 miles away. For this origin-destination pair, this degree of mobile device re-identification is sustained only on Fridays and Saturdays; on other days of the week, far fewer matches are registered.

Within urban environments, Bluetooth readers placed along the same freeway have the capacity to produce sufficient numbers of matches to continuously compute fine-grained five-minute travel times. In contrast, due to lower overall traffic volumes in rural areas, fewer travel time matches are generated and this capacity is therefore reduced. Even so, at least in the area around Lake Tahoe, sufficient matches were generated to compute fine-grained travel times during peak days.

The research team's results also indicate that a single Bluetooth reader can typically be used to monitor bi-directional traffic. Although the number of lanes at each reader used as part of this case study ranged from two to ten, data indicate that each reader was able to capture traffic in most, if not all, of the lanes at its location.

Finally, this use case enabled the research team to compare hit rates and matching percentages for readers located in both urban and rural environments. In this study, as is typical for urban versus rural settings, the biggest differences between the readers deployed on I-5 and US 50 included:

- The number of lanes at each reader;
- The distance between readers; and,
- The traffic volumes at each reader.

The I-5 readers were all configured to monitor traffic across six or more lanes of bi-directional traffic. Although three of the four I-5 readers were placed adjacent to the northbound lanes, the content of Table C4-11 and Table C4-12, demonstrates that they generate significant hit rates in both directions of travel; a similar situation exists for the Gloria reader deployed adjacent to the southbound lanes. This demonstrates that Bluetooth readers have the potential to monitor wide bi-directional freeway segments.

The content of these tables also indicates that directional traffic patterns have a significant degree of influence on Bluetooth device matching patterns. For example, on I-5, where northbound and southbound traffic volumes are comparable throughout the day, none of the readers re-identify more than 50% of the hits from upstream readers. In contrast, 68% of the hits from the Twin Bridges reader on rural US 50 are re-identified at the Placerville reader (39 miles away); see Table C4-12. These higher rural matching percentages, despite longer distances between readers, are in part due to US 50 exhibiting much stronger directional trends (e.g., eastbound US 50 carrying the majority of traffic on Friday afternoons/evenings). Despite this, volumes of Bluetooth reads along I-5 are several times greater than those along US 50, facilitating the calculation of more granular travel time reliability metrics.

Finally, each of the eight Bluetooth readers from which data was collected as part of this use case were mounted on roadside controller cabinets, and used directional antennae to focus signal strength toward the roadway. The fact that each of the readers had high hit rates and

produced significant matching percentages with downstream readers demonstrates that this is an effective configuration for capturing multi-lane, bi-directional traffic. However, the team also found that the readers are most effectively used when deployed in locations where they only monitor traffic in the mainline lanes. This is particularly a problem with readers placed adjacent to on- or off-ramps, such as the readers on I-5 at Florin and Pocket, as the travel time re-identification between the vehicle's timestamp at the on-ramp and its timestamp at the next downstream reader will be higher than the true travel time on the freeway; this is especially true if the ramp is congested or has ramp metering. For agencies using Bluetooth networks already in the field, it is important to determine which readers may be monitoring ramp traffic so that these travel time biases can be understood and mitigated.

Table C4-11: BTR Reader Summary, I-5N to US 50E, Friday 12:00 PM-9:00 PM

Reader	Mounting Type	Lanes	MAC ID Reads	% of Traffic Sampled	Distance to Next Reader (mi.)	Exits Between Readers	% Hits Re-identified Downstream
Pocket (I-5)	NB roadside controller cabinet	3	4208	7.2%	0.9	1	43%
Florin (I-5)	NB roadside controller cabinet	4	5402	8.1%	1.1	0	47%
Gloria (I-5)	SB roadside controller cabinet	4	5843	9.6%	4	2	45%
Vallejo (I-5)	NB roadside controller cabinet	5	6642	7.7%	46	27	2%*
Placerville (US 50)	EB roadside controller cabinet	1	1676	Not available	39	7	34%
Twin Bridges (US 50)	EB roadside controller cabinet	1	882	Not available	8	3	55%
Echo Summit (US 50)	WB roadside controller cabinet	1	771	Not available	4	2	74%
Meyers (US 50)	EB roadside controller cabinet	1	1059	Not available	--		--

Table C4-12: BTR Reader Summary, US 50W to I-5S, Sunday 12:00 PM-9:00 PM

Reader	Mounting Type	Lanes	MAC ID Reads	% of Traffic Sampled	Distance to Next Reader (mi.)	Exits Between Readers	% Hits Re-identified Downstream
Meyers (US 50)	EB roadside controller cabinet	1	936	Not available	4	2	52%
Echo Summit (US 50)	WB roadside controller cabinet	1	771	Not available	8	3	66%
Twin Bridges (US 50)	EB roadside controller cabinet	1	968	Not available	39	7	68%
Placerville (US 50)	EB roadside controller cabinet	1	1495	Not available	46	27	6%*
Vallejo (I-5)	NB roadside controller cabinet	5	4940	8.1%	4	2	45%
Gloria (I-5)	SB roadside controller cabinet	4	4352	8.7%	1	1	48%
Florin (I-5)	NB roadside controller cabinet	4	4003	8.5%	1	1	42%
Pocket (I-5)	NB roadside controller cabinet	3	3208	7.5%	--	--	--

* Readers should note that the low re-identification rates for Vallejo on Table C4-11 and Placerville on Table C4-12 are primarily the result of the next downstream reader for each being 46 miles away and on an adjoining roadway; I-5 to US 50 and US 50 to I-5.

Using Bluetooth and Electronic Toll Collection Data to Analyze Travel Time Reliability in a Rural Setting

This use case aims to quantify the impact of adverse weather and demand–related conditions on travel time reliability using data derived from Bluetooth and electronic toll collection-based systems deployed in rural areas.

Introduction

At present, loop detectors provide the majority of transportation data used for highway analysis. These detectors must be embedded in the roadway and require regular quality checking and often-costly maintenance. Bluetooth and electronic toll collection-based systems, on the other hand, can be mounted onto existing infrastructure either overhanging or adjacent to the roadway, thereby reducing the costs of deployment, reconfiguration, repair, and replacement. These systems work by scanning compatible devices deployed inside passing vehicles for unique identification information (i.e., Media Access Control (MAC) ids for Bluetooth devices and tag id numbers for toll transponders). If multiple readers detect identification information for a uniquely identifiable device, a record of that vehicle's travel time can be constructed. Because these devices do not need to be permanently fixed on the roadway, they offer a more flexible and often more cost effective method for detection, especially in rural locations.

To examine travel time reliability within the context of this use case, methods were developed to generate probability density functions (PDFs) from large quantities of travel time data representing different operating conditions. To facilitate this analysis, travel time and flow data from ETC readers deployed on I-80W and Bluetooth readers deployed on I-50E and I-50W were obtained from PeMS and compared with weather data from local surface observation stations. PDFs were subsequently constructed to reflect reliability conditions along these routes during adverse weather conditions, as well as according to time-of-day and day-of-week. Practical data quality issues specific to Bluetooth and ETC data were also explored.

This use case has value to a broad range of user groups. Transportation agencies with data collection needs in rural areas will benefit from seeing a travel time reliability analysis of real world data obtained from Bluetooth and ETC devices. This type of data is still fairly uncommon in practice and this use case should help to demystify it, demonstrating how such data sets compare to more commonly available types of traffic data. Operators and analysts will benefit from a discussion of the quality and typical characteristics of this type of data. Transportation agencies with specific data needs and cost constraints seeking a flexible sensor deployment may find Bluetooth or ETC-based systems more attractive based on the results of this analysis.

This use case also has value to operators who are interested in the effects of varying weather conditions and weekend travel on travel time reliability within a rural setting. Understanding the historical effects of different weather and demand conditions on the performance of a given roadway enables operators to respond to similar conditions as they occur, for example, by posting the expected range of travel times on dynamic message signs located at key decision points.

Use Case Analysis Sites

Two sites were used in the validation of this use case to compare similar phenomena in different locations, as well as to highlight the different types of data available in this region (see Exhibit C4-41). Site 1 is a 45.2-mile stretch of primarily 4-lane divided highway along I-80W with an estimated free flow travel time of 46 minutes. It begins east of the Truckee-Tahoe Airport weather station and ends just after I-80 exits the western border of the Tahoe National Forest. This roadway is instrumented with ETC readers mounted on sign structures overhanging the roadway. Exhibit C4-42 shows an example of Site 1.

Site 2 is a shorter 10.8-mile stretch of 2-lane highway along US 50 with an estimated free flow travel time of 14 minutes. This site was examined in both the Eastbound and Westbound directions of travel. It approaches the South Lake Tahoe airport on its eastern end and terminates in the West just outside of Twin Bridges. Site 2 is instrumented with Bluetooth readers deployed along the side of the roadway. Exhibit C4-43 shows an example of Site 2.

Table C4-13 provides details about the two sites.

Table C4-13: Site Characteristics

	Highway	Distance	Estimated Travel Time	Type
Site 1	I-80 W	45.2 miles	46 minutes	4 lane, divided
Site 2	US 50 E & US 50 W	10.8 miles	14 minutes	2 lane

These two sites were selected due to their strong weekend traffic patterns, as well as their proximity to local weather observation stations. They were made as short as possible (within the constraints of the detection infrastructure) in order to enable the research team to closely tie its analysis to the data generated by the weather stations, thereby maximizing the relevance of the weather data. Both Sites 1 and 2 are rural and receive relatively little commute or intercity traffic during the week. However, Lake Tahoe is a popular weekend and holiday destination for residents of the Bay Area, which is just a 3.5-hour drive away. I-80 and US 50 are both popular routes to take to get to Lake Tahoe from the Bay Area, and they are known for their heavy weekend traffic as large numbers of people enter and leave the area at nearly the same time.

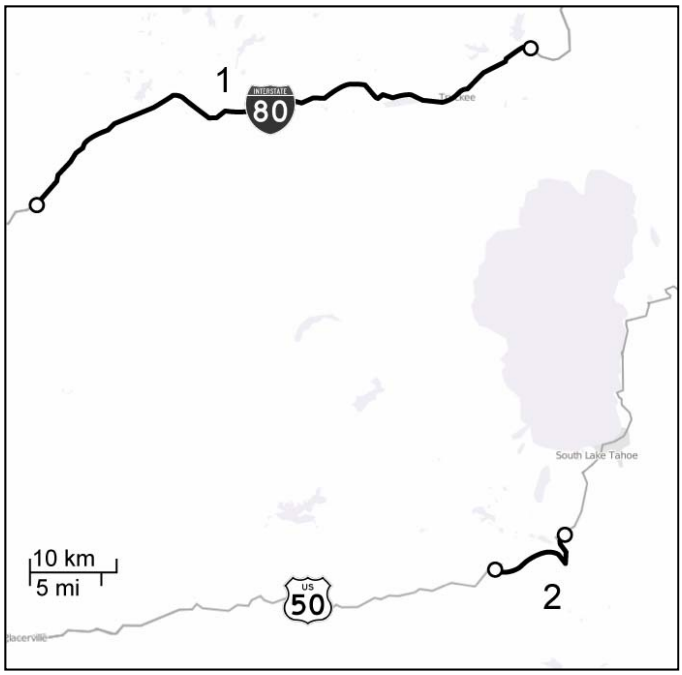

Exhibit C4-41: Use Case Site Map

Exhibit C4-42: Example of Site 1, I-80

Exhibit C4-43: Example of Site 2, US 50

Analysis Methodology

The routes included as part of this use case were analyzed to determine the effects of weather and weekend travel conditions on travel time reliability. To do this, travel time PDFs that isolate certain operating scenarios (e.g., snow on a weekday) were constructed. Time-of-day, day-of-week, and weather conditions based PDFs were generated for Site 1, and time-of-day and day-of-week conditions based PDFs were generated for Site 2.

To begin the analysis, travel time statistics for 5-minute windows at both sites were obtained from PeMS. For Site 1, where weather conditions were considered, travel time data was matched with weather data from the nearby AWOS-III surface observation station. Each 5-minute time interval was marked with its corresponding weather event if any (rain, snow, fog, or thunderstorm), visibility distance (0 to 10 miles), and precipitation (in inches). For site 2, where only weekend travel effects were considered, intervals were grouped into three categories. Travel times were labeled as belonging to a weekday (Monday through Thursday), a Friday, a Saturday, a Sunday, or a holiday.

With the travel time data collected and labeled, an effort was made to determine which data points, if any, should be thrown out. As was discussed previously and will be explored further in the *Data* section of this use case, travel time data obtained from Bluetooth and ETC readers can, depending on a number of variables, contain artificially long travel vehicle times.

The travel time data, labeled with weather condition and day-of-week, was used to construct PDFs of travel times under varying operating conditions. The effects of weather and weekend travel can be seen in the differences in travel time variability as indicated in the PDFs reflecting different conditions. Finally, aggregate travel time reliability statistics such as the 95th percentile travel time were computed for different conditions.

Data Collected. To complete this use case, Bluetooth and ETC data was retrieved from PeMS for the two sites described above. ETC data was obtained at Site 1 and Bluetooth data at Site 2 (see Table C4-14). To benefit from the availability of this rich data set, all available data was used in both cases. This was particularly desirable, as the available data does not span seasonal changes. The data obtained from PeMS included:

- Minimum travel time,
- Average travel time,
- Maximum travel time,
- 25th, 50th, and 75th percentile travel times, and
- Flow (number of vehicles observed during the window).

Each of these metrics was collected for a series of consecutive 5-minute windows. It should be noted that not all 5-minute windows during the periods of observation contained usable data, so some gaps exist in the data do exist.

Table C4-14: Dataset Descriptions by Site

Roadway	Data Type	Date Range	Data Completeness	Quantity of Data
I-80 West	ETC	4/25/2011 to 6/29/2011	59.5%	11,071 points
US 50 West	Bluetooth	1/28/2011 to 4/21/2011	35.9%	8,576 points
US 50 East	Bluetooth	1/28/2011 to 4/21/2011	38.9%	9,376 points

To examine the effect of weather on travel times across Site 1, weather data was obtained from the nearby AWOS-III surface observation station located at the Truckee-Tahoe Airport. This data was available in windows ranging between 5 and 20 minutes, fine grained enough to match well with the 5-minute travel times. Here, the research team focused on optional event tags (fog, rain, snow, or thunderstorm), visibility (0 to 10 miles), and precipitation (in inches).

After the weather and travel time data sets were obtained, the travel time data was next quality checked to ensure no erroneous data points were included. As mentioned previously, Bluetooth and ETC-based data collection systems are susceptible to data errors due to the way they measure travel times. These detectors work by recording the MAC address or toll tag id of vehicles that pass them on the roadway, along with a timestamp. This identification data is matched between detectors such that a vehicle passing multiple BTRs produces a travel time for that link. However, if that vehicle stops somewhere along the roadway after passing the first BTR before it continues on to the second, an artificially large travel time will be seen. Similarly,

if a vehicle visits the first BTR, then travels to an adjacent, but unmonitored roadway prior to returning to the monitored roadway and passing the second BTR, the travel time for that trip will be artificially large. Additionally, vehicles traveling past the same BTR more than once in different directions can also cause data errors when readers are capable of measuring multiple directions of travel simultaneously.

To prepare the raw data for analysis, these inaccurate travel times should typically be removed individually. The data set for this use case was composed of aggregate statistics that had already been computed based on all available travel time data; including data that is potentially inaccurate. To prevent the analysis conducted as part of this use case from being skewed by those values, the research team used the median travel time for each 5-minute interval. In this case, working with the median as opposed to the mean has a significant effect on the analysis, reducing the appearance of implausible extreme values. This works well for periods of time with significant traffic flow because unreasonably long travel times are muted as the sample size increases. However, the problem remains when the sample size is small, as a time interval containing a single extreme value will still result in an unreasonable median. As can be seen in Exhibit C4-44, below, representing conditions for Site 1, virtually all "long" travel times in the data occur during low volume time periods. It should be noted that the flows shown in Exhibit C4-44 are not sustained, but rather 5-minute aggregates).

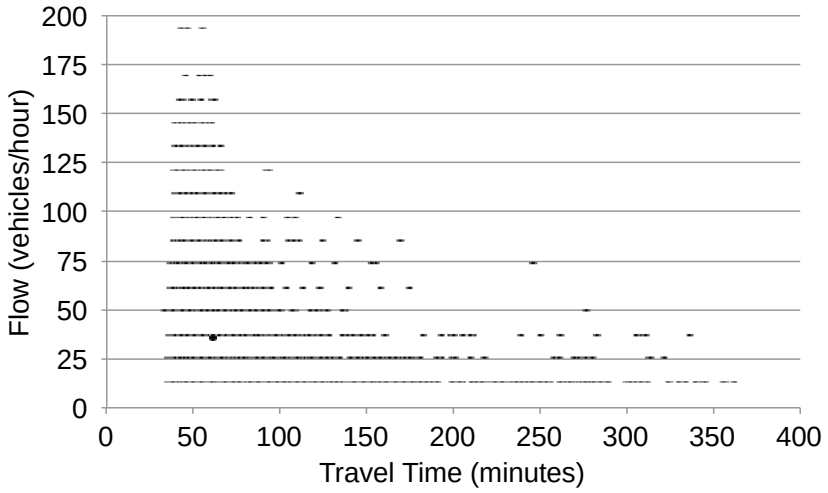

Exhibit C4-44: Travel Time vs. Flow on I-80

However, this does not mean that poorly represented time intervals should be discarded. While it is true that median travel times from sample intervals with larger numbers of vehicles should better conform to the expected value, and medians of smaller samples are more likely to contain outliers, this phenomenon is also representative of the fundamental behavior of traffic: both high (uncongested) and low (congested) speeds are seen at low flows. Thus, points from sparsely populated time intervals should not necessarily be discarded on those grounds alone (as long as the points can be assumed to be valid). Plotting the data for Site 1 from Exhibit C4-44 another way yields an empirical fundamental diagram for speed and flow (see Exhibit C4-45). The expected triangle shape can be seen with congested conditions represented by the points sloping down and to the left from the peak flow (seen around 60 mph) and uncongested

conditions represented by the points sloping down and to the right from the peak flow (with speeds above 60 mph). When viewed like this, all points appear to be valid as they are behaving according to basic traffic flow theory. Because longer travel times can be reasonably expected for time periods with few observations (i.e., during congested flow), it is determined that for the purposes of this use case no data points will be excluded.

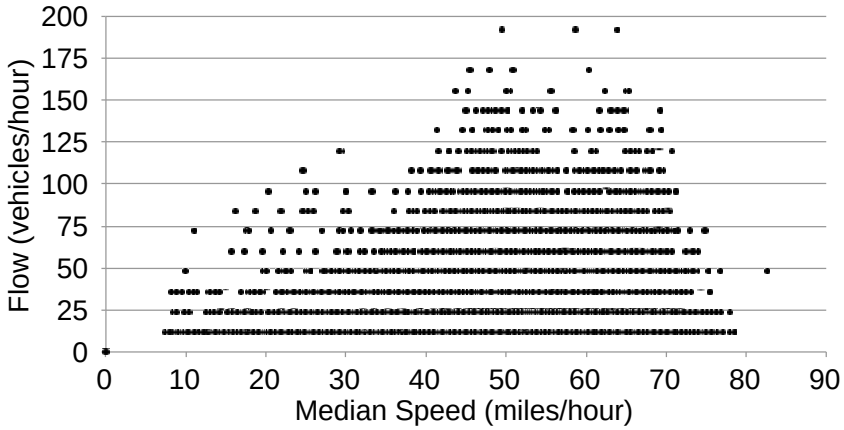

Exhibit C4-45: Speed vs. Flow on I-80

Travel Time Analysis: Site 1. Site 1, which lies on I-80W and begins just North of Lake Tahoe, is known to receive heavy traffic from vehicles returning to the Bay Area from weekend trips on Sunday evenings. As such, the breakdown of travel times by day-of-week from April 25 to June 29, 2011 shown Exhibit C4-46 indicates that the Sunday 95th percentile travel time exceeds that of a normal weekday by ~34%. This difference indicates increased travel time unreliability on Sundays, whereas the rest of the week appears fairly consistent. Since Sundays exhibit a significantly different pattern of traffic, they were considered separately as part of the research team's weather analysis.

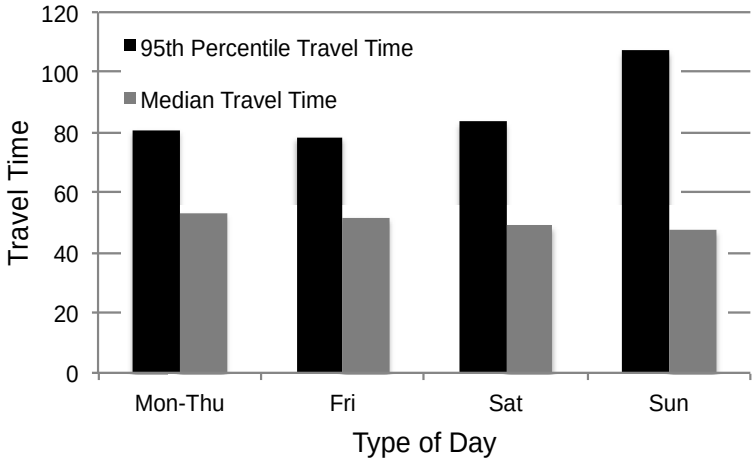

Exhibit C4-46: 95th Percentile and Median Travel Times for Site 1

Having assessed travel time reliability trends over the entire week, we next examined travel times within individual days to determine if it was necessary to handle AM and PM peak conditions separately during the analysis. To facilitate this, the distribution of travel times for each 5-minute interval over the course of a full day was plotted (see Exhibit C4-47). This figure demonstrates that no significant time-of-day trends exist on this route. If the typical day had shown some periodicity, it would have been necessary to examine weather effects during peak and off-peak hours separately. However, as travel times appear to be consistently between 30 and 45 minutes throughout the day the research team was able to conduct its weather-related travel time reliability analysis without accounting for differences between for daily peak conditions.

Exhibit C4-47 also indicates the presence of significantly longer travel times (hovering near the top of the chart). As these travel times do not appear to follow any time-of-day trends, it was surmised that they might be the result of adverse weather conditions. The team explored this idea further by generating PDFs of travel times collected during varying weather conditions. These PDFs were built by placing each 5-minute median travel time into a bin, each of which is 5 minutes wide. To define discrete weather conditions, we adopted five (5) labeled event categories (baseline, snow, rain, fog, and thunderstorm). We then broke the quantitative measures precipitation and visibility down into categories. For precipitation, we created "no precipitation" and "some precipitation" cases, and for visibility, we defined "low visibility", "medium visibility," and "high visibility" cases which corresponded to 0-3, 3-7, and 7-10 miles of visibility, respectively. The event conditions were all mutually exclusive, as were the visibility and precipitation categories. Note that the "baseline" event condition does not necessarily mean that driving conditions were ideal, but that no event was associated with that time (there may have been precipitation or low visibility).

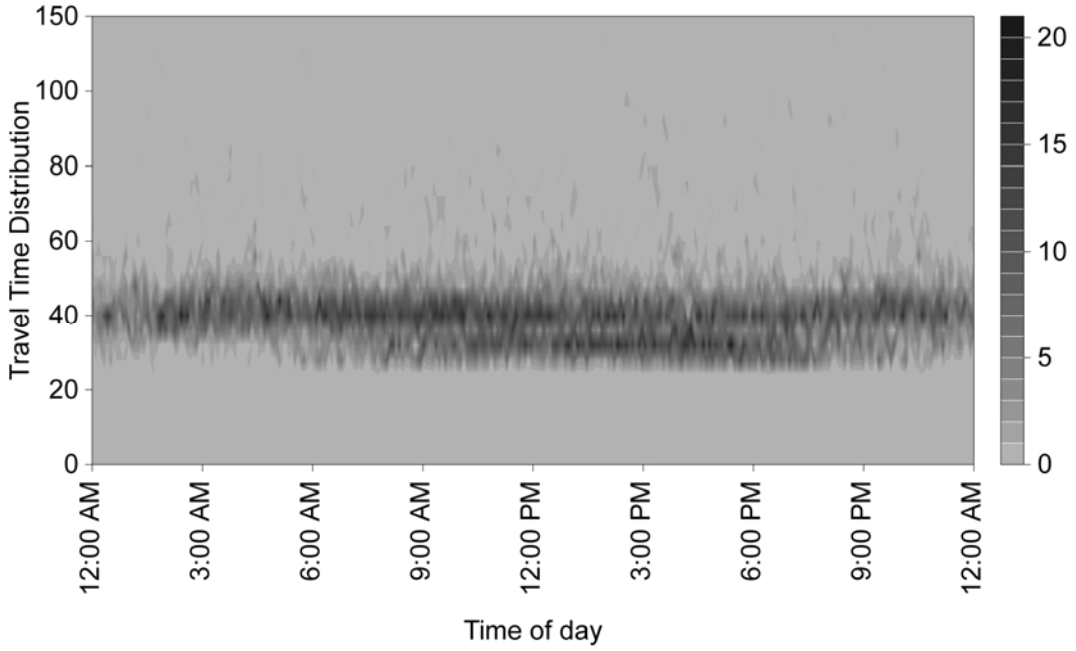

Exhibit C4-47: Site 1 Time of Day Travel Time Distribution

The resulting weather PDFs can be seen in Exhibit C4-48 and their effects on travel time are summarized in Exhibit C4-49. Note that the scale of the vertical axis in Exhibit C4-48 is not

consistent across each of the graphs. This is due to the variable quantities of data available for each condition.

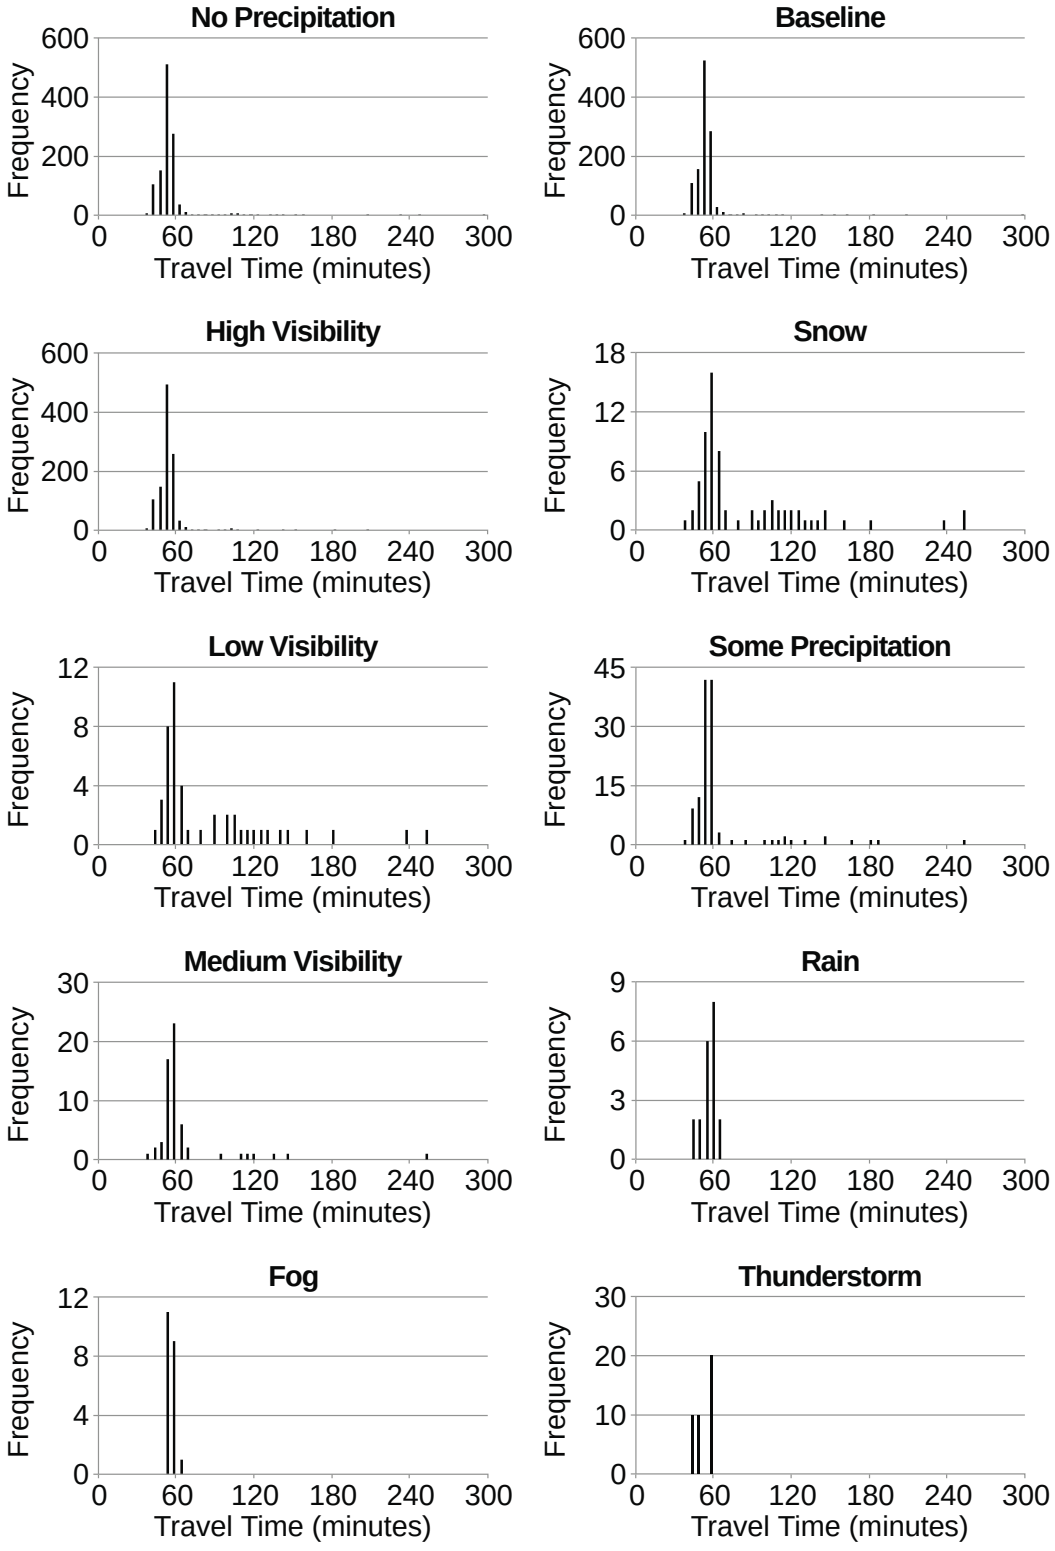

Exhibit C4-48: Site 1 Travel Time PDFs During Various Weather Conditions

It is clear from Exhibit C4-49 that snow, low to moderate visibility, and precipitation have a measurable effect on travel time reliability. The 95th percentile travel times during those weather conditions are significantly higher than their median travel times, indicating that the distribution of travel times is skewed toward the high end.

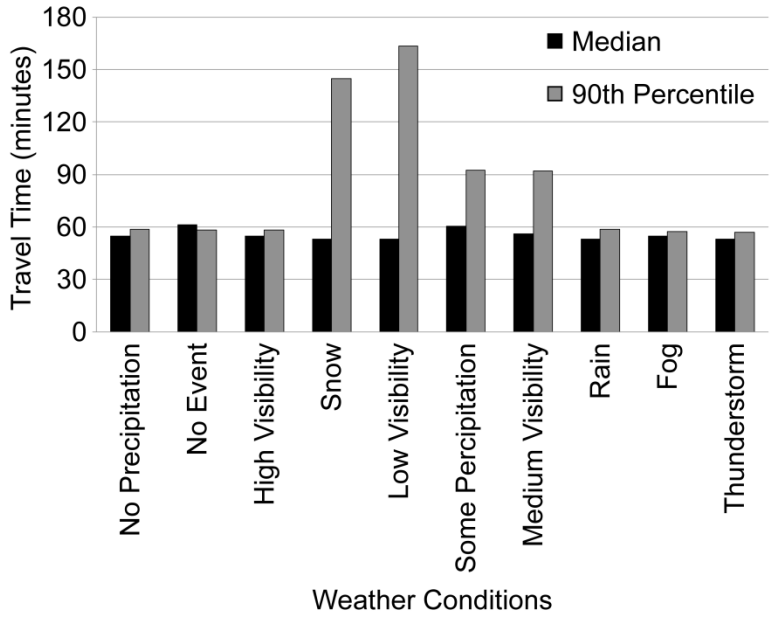

Exhibit C4-49: Site 1 Summary of Weather Effects

Another way to explore this data is to assess which conditions were present during the longest travel times occurring on this route. The results of this analysis are presented in Table C4-15. This perspective complements that of the PDFs displayed above by revealing that adverse weather events are present during many more long travel times than short travel times. In fact, the research team's analysis indicated that if a travel time exceeded the 95th percentile for this route, there was nearly a 50% chance that it was snowing, despite snow accounting for only 5% of all trips.

Table C4-15: Weather Conditions Active During Long Travel Times

Conditions	Active	Active when travel time exceeded 85th percentile	Active when travel time exceeded 95h percentile
No Precipitation	90.3%	84.3%	76.6%
Precipitation	9.7%	15.7%	23.4%
Baseline	90.7%	74.9%	54.7%
Snow Event	5.8%	23.0%	45.3%
Rain Event	1.6%	1.6%	0.00%
Fog Event	1.7%	0.5%	0.00%
Thunderstorm Event	0.3%	0.00%	0.00%
High Visibility	84.7%	67.5%	48.4%
Medium Visibility	4.8%	11.0%	11.0%
Low Visibility	3.8%	14.1%	31.3%

Travel Time Analysis: Site 2. Site 2 was similar to Site 1 in that it is subject to periodic spikes in demand due to weekend travel. However, whereas Site 1 is a 4-lane divided highway, Site 2 is a 2-lane highway (with only intermittent passing opportunities) and thus not as well equipped to handle the additional demand. Site 2 was equipped with Bluetooth detectors that were used to construct travel times in a similar manner to the ETC readers used for Site 1. The goal of the Site 2 travel time analysis was to determine the effects of the weekend travel on this site.

We began by examining a typical day on US 50 to check for the presence of AM or PM peak conditions, which would have to be controlled for as part of day-of-week analysis. Similarly to Site 1, 5-minute median travel times were obtained from PeMS. The time-of-day average of these median travel times is presented in Exhibit C4-50, which does not appear to show any true peak conditions. While the maximum daily travel time appears to occur at around 5:00 AM, this does not appear to be a true AM peak, likely being attributable to artificially high travel times occurring during low volume periods as discussed in the *Data Collection* section.

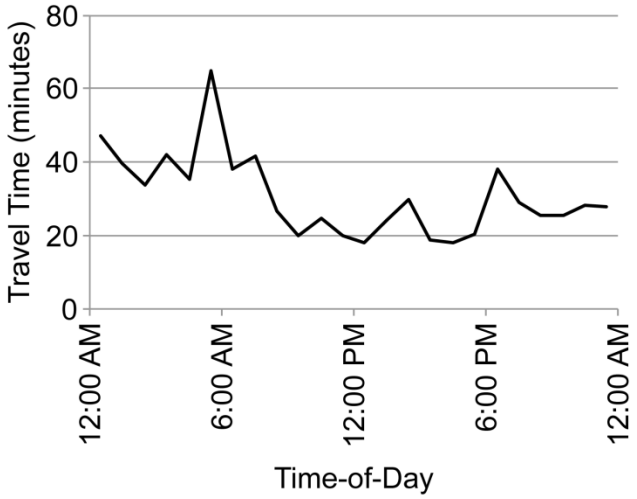

Exhibit C4-50: I-50W Average Travel Time by Time-of-Day

This assessment is supported by the average daily flow data displayed in Exhibit C4-51. Due to the absence of daily peak conditions at this site, the research team decided to consider each day as a whole. If strong peak conditions had been observed, it would have been necessary to develop travel time distributions for peak and off-peak conditions separately.

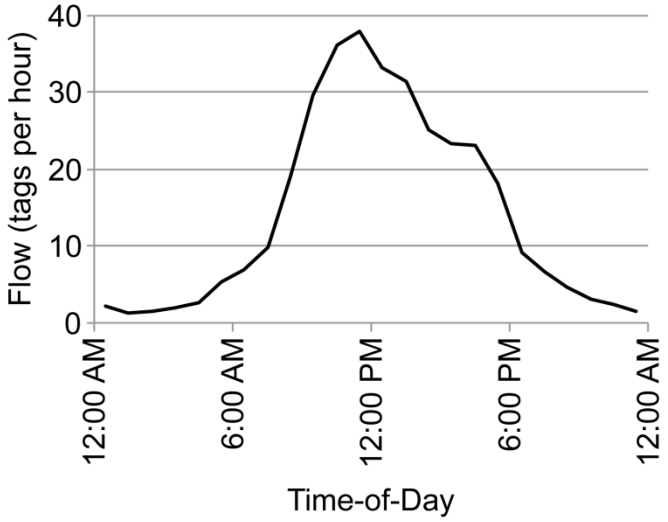

Exhibit C4-51: I-50W Average Flow by Time-of-Day

If this site were indeed subject to heavy weekend demand, it would be expected that travel times would be less reliable during the weekend. In order to explore whether the data supported this, the team first plotted the average vehicle flow over the course of the week for each direction of traffic for this site. It can be seen in Exhibit C4-52 and Exhibit C4-53 that weekend demand dominates the traffic profile for this section of roadway. As a result, we would expect travel time unreliability to follow a similar pattern.

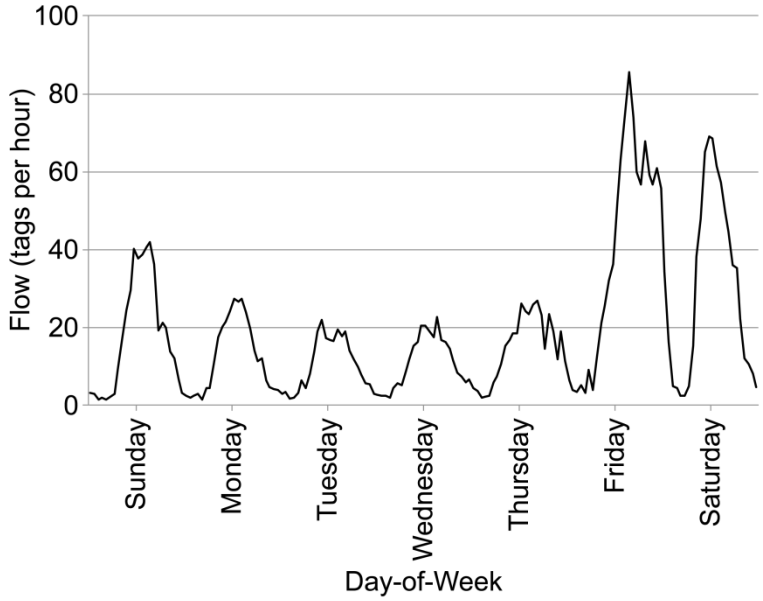

Exhibit C4-52: Weekly Flow on US 50E

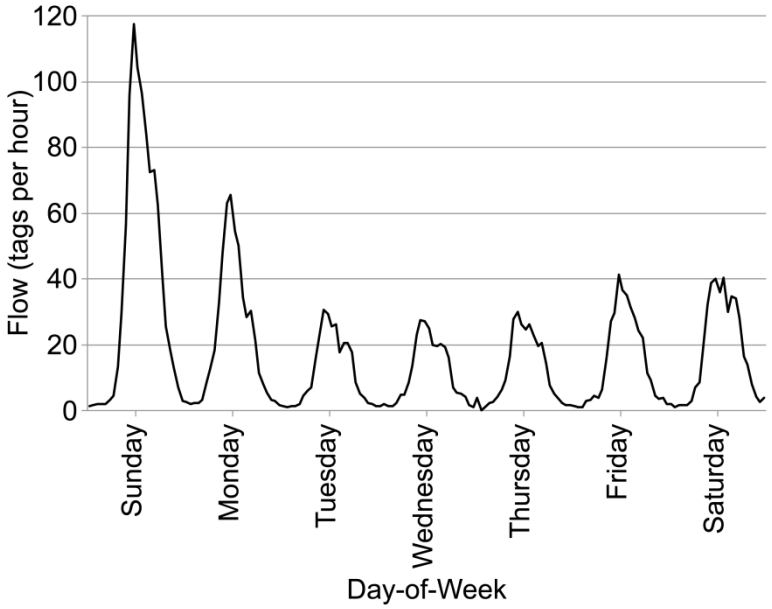

Exhibit C4-53: Weekly Flow on US 50W

To visualize travel time unreliability for this site, the research team constructed a travel time density plot representing a full week for US 50 East (see Exhibit C4-54). This figure is a collection of PDFs for each 5-minute period over the course of the entire week. Since this figure represents travel times in the Eastbound direction, we would expect more unreliability on Friday and Saturday as weekend travelers are making their way to Lake Tahoe from the Bay Area. The PDF appears to confirm this, as it can be seen at a glance that Friday is the day with the most

severe unreliability, with Sunday through Thursday exhibiting much more consistent travel times in comparison.

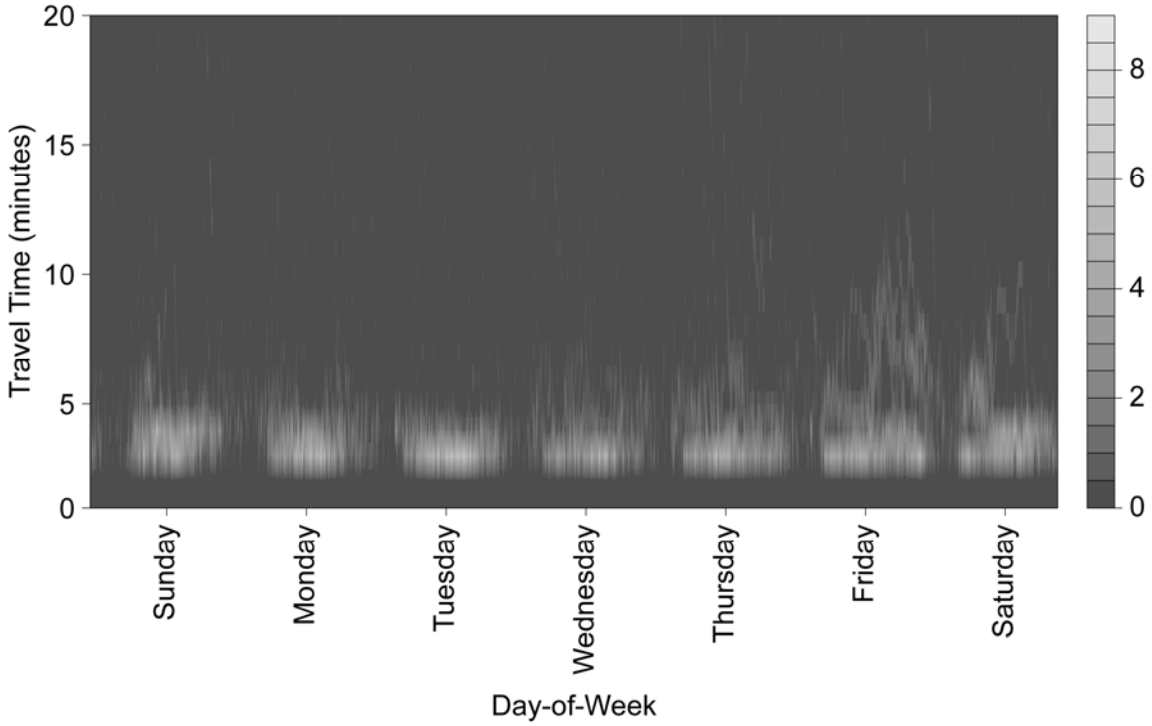

Exhibit C4-54: Week-long Distribution of Travel Times on US 50E

If this variation in travel time reliability over the course of the week is the result of weekend travel patterns and not adverse weather or some other factor, we would expect to see a complementary trend on US 50 West. Sunday should have been the least reliable day in this direction of travel as heavy traffic caused unreliability for travelers returning to the Bay Area from Lake Tahoe at the end of the weekend. After constructing PDFs for the opposite direction of travel, we see that this is in fact supported by the data (Exhibit C4-55).

The travel time variability by weekday on US 50 can be expressed in terms of the 95th percentile of travel time. This is presented for both directions of travel along with the mean by day in Table C4-16 below. Weekend travel patterns appear in the longer 95th Percentile travel times seen on Friday in the Eastbound direction and Sunday in the Westbound direction.

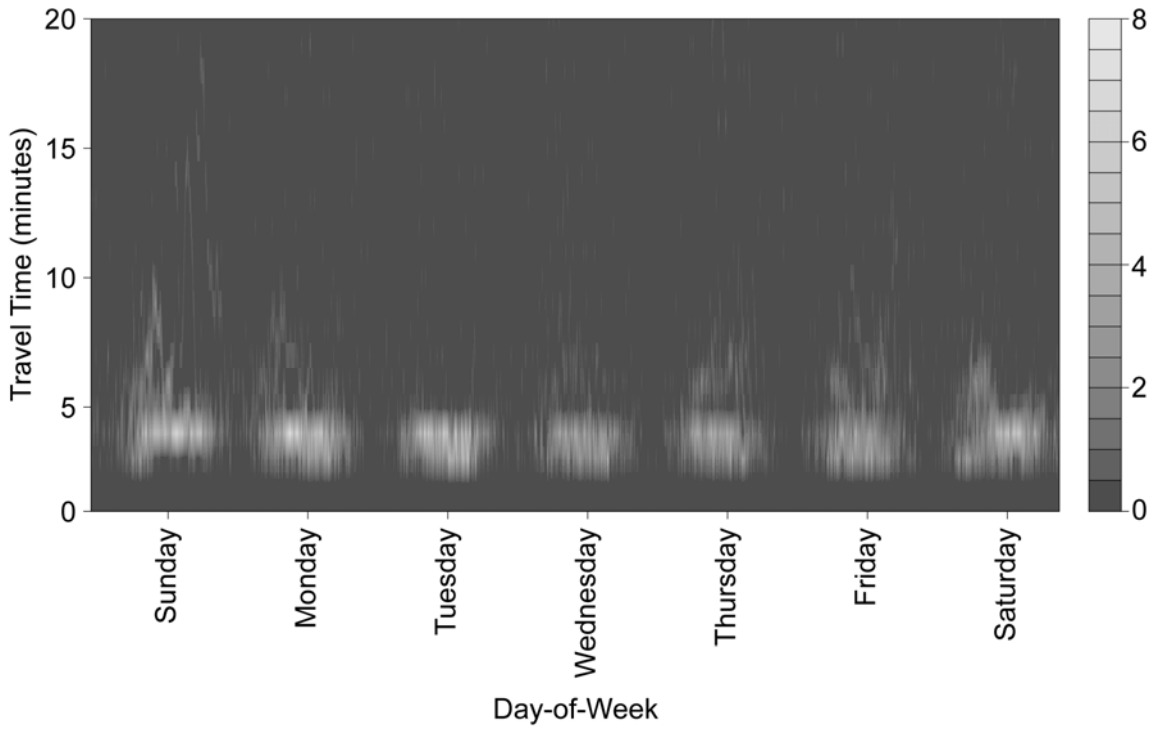

Exhibit C4-55: Week-long Distribution of Travel Times on US 50W

Table C4-16: Travel Time Reliability By Weekday

	Sun.	Mon.	Tues.	Wed.	Thurs.	Fri.	Sat.
US 50 E - Mean Travel Time	9.6 min	8.1 min	8.9 min	8.3 min	12.4 min	14.3 min	10.4 min
US 50 E - 95th Percentile Travel Time	25.4 min	22.0 min	18.2 min	23.7 min	36.0 min	41.1 min	32.2 min
US 50 W - Mean Travel Time	12.7 min	11.5 min	10.7 min	11.3 min	14.4 min	12.3 min	13.1 min
US 50 W - 95th Percentile Travel Time	40.4 min	30.7 min	20.8 min	30.8 min	41.0 min	37.7 min	31.0 min

PRIVACY CONSIDERATIONS

Introduction

As discussed in previous sections of this case study, innovations in data collection technology are providing exciting opportunities in the area of roadway travel time measurement. At the same time, use of these technologies is not without challenges, some technical, others related to protecting the confidentiality of personal information contained in ETC toll tag and Bluetooth mobile device datasets. As individual drivers' privacy has the potential to be compromised when others have the ability to track their movements across the public roadway network, users of this data, both public and private, have developed a variety of plans and programs to ensure that data gathered in support of the generation of roadway travel times cannot

be linked back to individuals. Recognizing that the data collection technologies described in this case study have the potential to raise public concerns over privacy, this section provides examples of the types of privacy protection policies and procedures currently in use by both public agencies and private sector companies to guard against the misuse of drivers' personal information.

Electronic Toll Tag-Based Data Collection

Overview of Personal Privacy Concerns

When used for toll collection purposes, toll transponders are automatically identified whenever they pass within the detection zone of a compatible ETC reader. Every time this occurs, the ETC reader prompts the tolling system to deduct a pre-determined amount of money from the prepaid debit account associated with that transponder's unique ID number. Recognizing that this technology would make it possible to track the path of each transponder-enabled vehicle between successive ETC readers, a number of agencies have deployed supplemental (non-revenue generating) ETC readers and back-office data analysis systems to facilitate the calculation of point-to-point travel times based on this data.

Although not instantaneous, direct connection exists between a toll transponder's unique ID and the personal information of the transponder user, such data does exist within agency databases. As a result, this creates concerns for some users stemming from the potential loss of anonymity associated with their travel behavior.

Policies and Procedures in Place to Protect the Privacy of ETC Transponder Data

Two of the agencies best known for making use of anonymous ETC transponder data in support of travel time data collection are:

- Houston TranStar (Houston, TX)
- Metropolitan Transportation Commission (San Francisco Bay Area, CA)

Whereas both agencies have made significant efforts to protect the personal information of ETC toll tag users, only MTC has developed detailed guidelines concerning the use, archiving, and dissemination of this data.

Houston TranStar. Houston was the first city in the United States to apply ETC-based tolling technology to the collection of data concerning travel times and average speeds. The toll tag data on which this system is based is collected from ETC reader stations deployed at one to five mile intervals along over 700 miles of Houston area roads. Traffic Management Center (TMC) staff use this system to detect congestion along area freeways and high occupancy vehicle (HOV) lanes; this data is also provided to the pubic via media reports, travel times posted to roadside changeable message sign (CMS), and the Houston TranStar website. In an effort to protect the privacy of the driver's from which travel time data is being collected, TranStar has configured its ETC readers to only store the last four digits of each toll tag's ID number. Truncating ID numbers in this way creates an environment where the agency's automated systems can track, but not identify, individual vehicles as they move across the data collection network. TranStar staff are acutely aware of drivers' concerns regarding the protection of their personal information and have made efforts to inform the public that not only do they collect just a portion of each toll tag's ID number, but also that none of the information concerning the movement of individual transponders is available for use by agency staff or law enforcement.

Metropolitan Transportation Commission. In support of its 511-traveler information service, the Metropolitan Transportation Commission (MTC) operates a travel time data collection system based on information collected from the region's FasTrak toll system. As part of this effort, MTC takes the following steps to ensure the protection of toll tag users' personal information (7):

- Encryption software in the central software system encrypts each toll tag ID before any other processing is carried out to ensure that the toll tags are treated anonymously.
- Encrypted toll tag IDs are retained for no longer than twenty-four hours before being discarded. No historical database of encrypted IDs is maintained beyond that time period.

In addition to establishing the guidelines described above concerning the management of toll tag data, MTC has also developed the following principles regarding the protection of personal privacy (8):

1) All traffic data collection activities will be implemented in a manner consistent with Federal and California laws governing an individual's right to privacy;
2) The tag users' consent will be secured before the operation of any data collection system based on toll tags;
3) No information about, or that is traceable to, any individual person will be collected, stored, or manipulated;
4) Information on the data collection, aggregation and storage practices will be available at the 511.org website, which will include traffic data collection methods, privacy policy, and full disclosure on the use of the data;
5) Members of the public will be given the ability to contact the program to discuss any privacy questions or concerns;
6) All recipients of the data shall comply with these privacy principles; and,
7) An annual evaluation will be conducted to assure that individual privacy is protected.

Although MTC provides the third-party contractors who operate its 511 and related services with access to the toll tag data collected as part of this system, as is indicated in items #6 and #7, above, these firms are required to observe all of MTC's privacy principles and are subject to an annual evaluation to verify their compliance.

Bluetooth-Based Data Collection

Overview of Personal Privacy Concerns

Bluetooth-based travel time data collection systems operate, similarly to ETC-based systems, via the re-identification of mobile device ID data at successive locations along a roadway. However, whereas other technologies used to calculate roadway travel times based on the movement of probe vehicles (e.g., toll tag and license plate reader-based systems) have the potential, if abused, to directly link a specific user to the movement of their vehicle, identification of an individual based on their Bluetooth signature (i.e., MAC address) is much less straightforward. In theory, if the MAC address of the mobile device has been set by its manufacturer, the possibility exists, however remote, for a link to be established between the product part number and its owner via a product registration database or product warranty. Even

so, the MAC addresses of mobile devices, though unique, are not linked to specific individuals or vehicles via any type of central database or user account.

Despite these facts, public perception regarding this method of data collection varies widely and has the potential interfere with its implementation. As a result, users of this technology have implemented a range of procedures to minimize the possibility of infringing on users' privacy.

Policies and Procedures in Place to Protect the Privacy of Bluetooth ID Data

Two of the entities currently deploying Bluetooth-based data collection technologies for the purpose of calculating roadway travel times are:

- Post Oak Traffic Systems (Company utilized technology developed at the Texas Transportation Institute)
- Traffax (Company utilizing technology developed at the University of Maryland)

Users of Bluetooth-based data collection technologies stress that the MAC addresses collected by their systems are not directly associated with a specific user and do not contain any personal data or information that could easily be used to identify or "track" an individual person's whereabouts. That said, all recommend taking additional steps to further ensure that the information collected from individual Bluetooth devices is kept as anonymous as possible.

Post Oak Traffic Systems. Techniques used by this firm to help protect the personal privacy of drivers include:

- Only polling the Bluetooth device information necessary to facilitate the calculation of travel times, including:
  - MAC address;
  - Device reader location; and,
  - Timestamp.
- Although other data can be accessed as part of the Bluetooth device polling process (e.g., device name and packets of information concerning data exchanged between a mobile phone and its associated Bluetooth headset), Post Oak staff recommend only collecting the data absolutely necessary to calculate segment travel times.
- To further address potential privacy concerns, Post Oak field processors are programmed to encrypt all Bluetooth ID data immediately upon receipt. Doing so ensures that the actual device ID is not sent or stored anywhere.

Traffax. Company staff recommend implementing the following additional measures to ensure that no unauthorized use of data occurs. This includes (9):

- Implementation of policies concerning the retention and dissemination of Bluetooth MAC address data, including:
  - Destroy or encrypt any base level MAC address information after processing.
  - Use industry standard encryption and network security. Proper security protocols, passwords, encryption and other methods should be incorporated into the data systems that store and process the MAC address data.
- Establishment of data processing safeguards (encryption and randomization) to prevent the recovery of unique MAC addresses:
  - Encryption methods transform MAC address data (at the sensor level) into an output form that requires special knowledge (such as an encryption key) to recover the original information. This activity preserves the uniqueness of the ID

so that matching can still be performed without risking exposure of actual device ID data.

- Randomization methods deliberatively degrade the data such that individual observations are no longer globally unique and the ability to track individuals based on their MAC addresses becomes theoretically impossible. A simple example of this would be to truncate the final 3 characters of the MAC ID.

All of the privacy protection methods recommended by Traffax are implemented at the sensor level, not the central processing station. Doing so makes is virtually impossible to obtain the complete and globally unique MAC address of any particular device.

Application of Privacy Principles

It has been amply demonstrated that travel time data collection technologies based on device re-identification (e.g., ETC toll tags and Bluetooth devices) have the potential to be abused in such a way as to cause significant privacy-related concerns. Although this section of the case study has reviewed a number of techniques currently being utilized to further ensure that drivers' anonymity is preserved, long-term acceptance of these technologies will ultimately rely on maintenance of the public's trust. To that end, the Intelligent Transportation Society of America has established a set of *Fair Information and Privacy Principles* aimed at safeguarding individual privacy within the context of the deployment and operation of Intelligent Transportation Systems. Although advisory in nature, these principles are intended to act as guidelines for use by public agencies and private entities to protect drivers' right to privacy. Principles include (10):

- Individual Centered: Intelligent Transportation Systems must recognize and respect the individual's interests in privacy and information use;
- Visible: Intelligent Transportation Information Systems will be built in a manner "visible" to individuals;
- Comply: Intelligent Transportation Systems will comply with applicable state and federal laws governing privacy and information use;
- Secure: Intelligent Transportation Systems will be secure;
- Law Enforcement: Intelligent Transportation Systems have an appropriate role in enhancing travelers' safety and security interests, but absent consent, statutory authority, appropriate legal process, or emergency circumstances as defined by law, information identifying individuals will not be disclosed to law enforcement;
- Relevant: Intelligent Transportation Systems will only collect personal information that is relevant for ITS purposes;
- Anonymity: Where practicable, individuals should have the ability to utilize Intelligent Transportation Systems on an anonymous basis;
- Commercial or Secondary Use: Intelligent Transportation Systems information stripped of personal identifiers may be used for non-ITS applications;
- Federal and State Freedom of Information Act (FOIA): FOIA obligations require disclosure of information from government maintained databases. Database arrangements should balance the individual's interest in privacy and the public's right to know; and,
- Oversight: Jurisdictions and companies deploying and operating Intelligent Transportation Systems should have an oversight mechanism to ensure that such

deployment and operation complies with their Fair Information and Privacy Principles.

LESSONS LEARNED

Overview

The team selected the Lake Tahoe region located in Caltrans District 3 in order to provide an example of a rural transportation network with fairly sparse data collection infrastructure. The data used as part of this case study was generated by electronic toll collection (ETC) readers on I-80 and Bluetooth-based data collection readers along I-5 and US 50 (see Exhibit C4-1). These readers register the movement of vehicles equipped with FasTrak tags (Northern California's ETC system) and Bluetooth-based devices (e.g., Smart Phones) for the purpose of generating roadway travel times.

Methodological Experiments

This case study examined vehicle travel time calculation and reliability using Bluetooth and RFID re-id systems. A number of factors were identified that influence travel time reliability and guided the development of methods for processing re-id observations and calculating segment travel times. The results show that smart filtering and processing of Bluetooth data to better identify likely segment trips increases the quality of calculated segment travel time data. This approach helps preserve the integrity of the data set by retaining as many points as possible, and basing decisions to discard points on the physical characteristics of the system, rather than their statistical qualities.

It is important to only filter out unlikely trips, so that the correctly measured variability of the data is not lost. The benefit of a more careful accounting procedure during the vehicle-identification stage allows for later statistical filtering of the data to be milder, preserving more meaning. Filtering trips based on statistical properties is less desirable because criteria for eliminating points are not based on the physical system. If all of the data points in an interval are valid, it does not make sense to discard that entire interval simply because it does not contain very many points. It is important to be aware of the interactions between preprocessing procedures. Future research may explore other smarter methods for filtering out unlikely segment trips. For example, considering observations across the entire BTR network would be useful for identifying unlikely segment trips.

A number of factors were found to influence vehicle segment travel times. For example, if the distance between BTRs was small, errors in calculated travel times may be significant and methods for determining passage time must be carefully considered. Signal strength availability enables easy and accurate determination of passage times. Without signal strengths, using arrival and departure times for passage times may improve travel time accuracy. This was found to be likely for BTR #10 based on the location of the reader relative to an intersection, the intersection configuration, and the short distance to the nearest BTR. Aggregating observations into visits was also found to be useful for distinguishing between trip and travel time for individual vehicles at a BTR.

Use Case Analysis

This case study explored four aspects of the ETC and Bluetooth reader networks used in the Lake Tahoe case study: (1) detailed locations and mounting structures; (2) lanes and facilities monitored; (3) percentage of traffic sampled; and (4) percentage and number of vehicles re-identified between readers. As a whole, it showed that vehicle re-identification technologies are suitable for monitoring reliability in rural environments, provided that traffic volumes are high enough to generate a sufficient number of samples. For rural areas that have heavy recreational or event traffic, vehicle re-identification technologies such as ETC and Bluetooth can provide sufficient samples to calculate accurate average travel times at a fine granularity during high-traffic time periods. During these high-volume periods, vehicle re-identification technologies can be used to monitor travel times and reliability over long distances, such as between the rural region and nearby urban areas.

For agencies deploying vehicle re-identification monitoring networks, it is necessary to understand that the quality of the collected data is highly dependent on the decisions made during the design and installation process. The mounting position and antennae configuration of ETC readers impacts the number of lanes sampled at a given location. The positioning of Bluetooth readers, which have a large detection radius, dictates whether ramp, parallel facility, or multi-modal traffic are also sampled, which can introduce large errors into travel time and reliability computations. In addition to choosing an optimal positioning of readers, it is also important to place and space them appropriately. Readers should be placed where they can provide travel time information for heavily traveled origins and destinations. Because vehicle re-identification readers can be easily moved, there are opportunities to do pilot tests to evaluate the quality, quantity, and value of collected data, so that the final deployment robustly supports the desired measures.

For agencies leveraging existing networks, it is important to fully understand the configuration of the network before using its data. At a minimum, this should include taking steps to verify that reader locations are correct and that the computed travel times and number of matches are reasonable given the distance and known traffic patterns between reader pairs. In locations where readers are closely spaced, computing reader hit rates and comparing between readers can help identify the reader most suited for monitoring travel times at a given location. Finally, evaluating percentage and volume of matched reads between each reader pair by time of day and day of week can indicate which time periods typically have sufficient matches to support average travel time computations at different granularities.

This case study also explored an approach for isolating and exploring the effects of weather and weekend travel on travel time reliability. As implemented here, the analysis should be fairly straightforward to replicate with data from a travel time reliability monitoring system such as PeMS and the appropriate weather data. The PDFs of travel times under different operating conditions consistently demonstrated the unreliability associated with low visibility, rain, and travel under high-demand conditions. This use case also described the travel time unreliability associated with such events in terms of 95th percentile travel time. Taken together, these tools should be valuable to planners, operators, and engineers interested in analyzing and communicating the travel time reliability of a section of roadway, especially one of a rural nature. Finally, application of the research team's approach has revealed several insights into the nature of working with Bluetooth and ETC-based sources of data. Specifically, that due to the nature of how these data collection technologies calculate travel time, it is necessary to account for artificially long travel times likely contained in the data set prior to conducting any analysis.

Despite this shortcoming, these technologies both provide users with the potential to effectively assess roadway travel times and consequently, reliability of travel, in rural areas where the cost of deploying and maintaining spot-based sensors (e.g., loop detectors) makes their use impracticable.

Privacy Considerations

For either of the data collection technologies described in this report to be successful over the long-term, safeguards must be put into place to ensure that the privacy of individual drivers being samples is protected. With this in mind, we recommend that any probe data collection program implemented by public agencies or by private sector companies on their behalf adhere to a pre-determined set of privacy principles (e.g., ITS America's Fair Information and Privacy Principles) aimed at maintaining the anonymity of specific users. Additionally, any third party data provider working for a public agency to implement a travel time data collection solution based on either of the technologies described in this Case Study should be required to submit an affidavit indicating that they will not use data collected on the agency's behalf in an inappropriate manner, including:

- Renting, leasing, selling, or otherwise providing data to any entity without explicit written permission of the agency;
- Using data for any purpose(s) other than those described as part of the project-specific requirements; and,
- Attempting to identify the ownership of individual vehicles or devices whose personal information is collected as part of the system's data collection infrastructure.

REFERENCES

1) http://www.ntoctalks.com/webcast_archive/to_aug_6_09/to_aug_6_09_bs.pdf
2) http://www.dot.ca.gov/dist3/departments/planning/corridorplanning.html
3) Porter, J.D., Kim, D.S., Magana, M.E., Poocharoen, P., and C.A. Gutierrez Arriaga. Antenna Characterization for Bluetooth-based Travel Time Data Collection. In Transportation Research Board 90th Annual Meeting Proceeding, Washington, D.C., 2011.
4) Haghani, A., M. Hamedi, K.F. Sadabadi, S. Young, and P. Tarnoff. Data Collection of Freeway Travel Time Ground Truth with Bluetooth Sensors. In Transportation Research Record: Journal of the Transportation Research Board, No. 2160, Transportation Research Board of the National Academies, Washington, D.D., 2010, pp. 60-68.
5) http://bata.mtc.ca.gov/tolls/fastrak.htm
6) Haas, R., M. Carter, E. Perry, J. Trombly, E. Bedsole, and R. Margiotta. iFlorida Model Deployment Evaluation Report. Prepared for the USDOT. Report No FHWA-HOP-08-050. January 2009.
7) http://traffic.511.org/privacy.asp
8) http://traffic.511.org/privacy.asp
9) http://www.traffaxinc.com/content/privacy-concerns
10) http://www.e-squared.org/privacy.htm

CHAPTER C5

ATLANTA, GEORGIA

The team selected the Atlanta, Georgia, metropolitan region to provide an example of a mixed urban and suburban site that primarily relies on video detection cameras for real-time travel information. The main objectives of the Atlanta case study were to:

- Demonstrate methods to resolve integration issues by using real-time data from Atlanta's traffic management system for travel time reliability monitoring
- Compare probe data from a third-party provider with data reported by agency-owned infrastructure
- Fuse the regime-estimation and non-recurrent congestion analysis methodologies to inform on the reliability impacts of non-recurrent congestion

The *monitoring system* section details the reasons for selecting the Atlanta region as a case study and provides an overview of the region. It briefly summarizes agency monitoring practices, discusses the existing sensor network, and describes the software system that the team used to analyze the use cases. Specifically, it describes the steps and tasks that the research team completed in order to transfer data from the data collection systems into a travel time reliability monitoring system.

The section on *methodological advancement* leverages methods developed in previous case studies to propose a framework for analyzing the impacts of non-recurrent congestion on a given facility's operating travel time regimes.

Use cases are less theoretical, and more site specific. The first use case details the challenges of leveraging ATMS data to drive a travel time reliability monitoring system. The second use case compares the results of analyzing congestion with agency-owned infrastructure-based sensors and third-party provider speed and travel time data.

Lessons Learned summarizes the lessons learned during this case study, with regard to all aspects of travel time reliability monitoring: sensor systems, software systems, calculation methodology, and use. These lessons learned will be integrated into the final guidebook for practitioners.

MONITORING SYSTEM

Site Overview

With a population of five and half million people, Atlanta is the 9th largest metropolitan area in the U.S. The layout of the freeway network follows a radial pattern. The core of the city is encircled by a ring road (I-285, known locally as "the Perimeter"), which is intersected by a number of interstates and state routes that radiate from downtown Atlanta into its outlying suburbs. Major radial highways include I-75 and I-85, which merge together to form a section of freeway called the "Downtown Connector" within the I-285 loop, I-20, which is the major east-to-west freeway in the region, and GA 400, which travels from north of downtown toward Alpharetta. A map of the major freeway facilities in the region is shown in Exhibit C5-1. The metropolitan freeway network also contains 90 miles of HOV lanes that operate 24 hours a day, 7 days a week on the following facilities:

- I-75 inside the I-285 loop
- The Downtown Connector
- I-20 east of the Downtown Connector
- I-85 between Brookwood and SR 20

Exhibit C5-1: Map of Atlanta Freeways

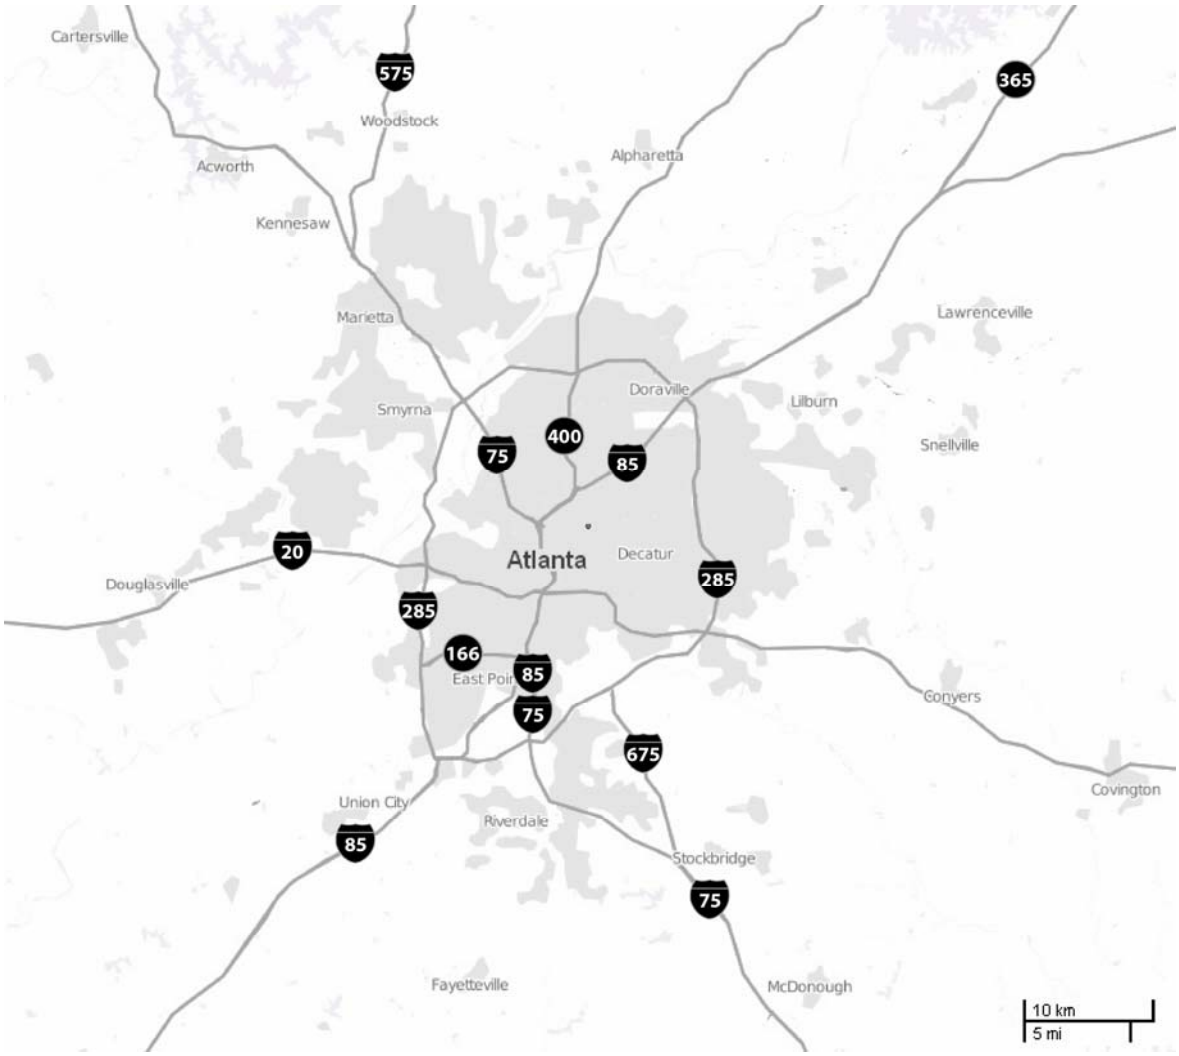

Additionally, on October 1, 2011, GDOT opened its first express lanes in the state of Georgia, which are operational on I-85 from I-285 to just south of the GA 365 split. The agency is also planning to deploy express lanes on I-75 north of Atlanta in 2015.

Atlanta's growing congestion is a major concern to GDOT and other agencies in the region. In 2008, the Atlanta region was granted $110 million by the USDOT for a Congestion Reduction Demonstration Program (CRD). Under this agreement, GDOT is partnering with the Georgia Regional Transportation Authority (GRTA) and the State Road and Tollway Authority (SRTA) to implement innovative strategies to alleviate congestion. The first phase of this program involved the conversion of HOV lanes to HOT lanes on I-85, mentioned above. Future

phases will add additional express lanes to major freeway facilities, enhance commuter bus service, and construct new Park and Ride lots. Aside from this program, GDOT is also undertaking a Radial Freeway Strategic Improvement Plan (RFSIP) to investigate the implementation of operational improvements, managed lanes, and capacity expansion on congested freeways, as well as to study how to increase transit mode-share.

GDOT monitors traffic in the Atlanta Metropolitan Area in real-time through its Advanced Traffic Management System (ATMS), called Navigator. The Transportation Management Center (TMC), located in Atlanta, is the headquarters and information clearinghouse for Navigator. TMC staff support regional congestion and incident management through a three-phase process:

- Phase 1: Collect Information- TMC operators monitor the roadways and review real-time condition information from sensors deployed along regional interstates. Operators also gather information provided by 511 users regarding traffic congestion and roadway incidents.
- Phase 2: Confirm and Analyze Information- TMC operators confirm all incidents by identifying the problem, the cause, and the effect it is anticipated to have on the roadway. Based on their analysis, proper authorities, such as police or fire responders, are notified.
- Phase 3: Communicate Information- TMC operators communicate information regarding congestion and incidents to travelers by posting relevant messages to regional CMS and updating the Navigator website and 511 telephone service.

GDOT's traffic management system integrates with traffic sensors, CCTVs, changeable message signs (CMS), ramp meters, weather stations, and Highway Advisory Radio (HAR). At the TMC, staff use the real-time data and CCTV feed to detect congestion and incidents. To minimize the disruption of traffic caused by lane-blocking incidents, TMC staff can dispatch Highway Emergency Response Operator (HERO) patrols. GDOT estimates that the implementation of HERO patrols through the TMC has reduced the average incident duration by 23 minutes and reduced yearly delay time by 3.2 million hours during the peak commute *(1)*. To facilitate information sharing and coordinated responses, the central TMC in downtown Atlanta is also linked to seven regional Transportation Control Centers, as well as the City of Atlanta and the Metropolitan Atlanta Rapid Transit Authority (MARTA).

Sensors

In the Atlanta region, GDOT collects data from over 2,100 roadway sensors, which include a mix of video detection sensors and radar detectors. Both of these types of sensors consist of single devices that monitor traffic across multiple lanes. The majority of active sensors are monitoring freeway lanes, with some limited coverage of conventional highways. Sensors in the active network are manufactured by four different vendors, as shown in Table C5-1.

Table C5-1: GDOT Sensor Network Summary

Vendor	Sensor Type	Percentage of GDOT Network
Traficon	Video	80%
Autoscope	Video	8%
NavTeq	Radar	8%
EIS	Radar	4%

The make and model of the sensor dictates the type of data that it collects and the frequency at which data is retrieved from the device (and thus, the level of aggregation of the data). Traficon video detection cameras make up approximately 80% of GDOT's active detection network. In Georgia, these sensors monitor flow, occupancy, and speed, and report data to a centralized location every 20 seconds. Autoscope video detection sensors make up another 8% of the GDOT detection network. These cameras also monitor flow, occupancy, and speed but, in the Atlanta region, report it to a centralized location every 75 seconds. The remainder of the detection network is composed of radar detectors, which also report aggregated flows, occupancies, and speeds. NavTeq radar detectors make up 8% of GDOT's active detection network and report data every 1 minute. Finally, EIS's RTMS radar detectors make up 4% of GDOT's active detection network and report data every 20 seconds. In addition to the aggregated flow, occupancy, and speed data, these sensors also report on the percentage of passenger cars versus truck traffic.

In general, the different types of sensors are divided up by freeway. Exhibit C5-2 shows the location of active mainline sensors in the GDOT network, broken down by manufacturer. The predominant sensors, the video detector manufactured by Traficon, exclusively cover the I-285 ring road, I-75, the I-75/I-85 Downtown Connector, and I-575. Traficon sensors also monitor GA-400 north of the ring road and the majority of I-85, and share coverage of I-20 with NavTeq radar detectors. In most of the network, Traficon sensors are placed with a very dense spacing of about one-third of a mile. Autoscope cameras monitor a small portion of I-85 near the Hartsfield-Jackson Atlanta International Airport with a spacing comparable to that of the Traficon cameras. In addition to sharing coverage of I-20 within the ring road with the Traficon sensors, NavTeq radar detectors exclusively monitor I-20 outside of the ring road, I-675, GA-400 inside of the ring road, and GA-316. NavTeq detectors are spaced approximately 1 mile apart. Finally, RTMS radar detectors exclusively monitor US-78, GA-141, and GA-166.

All sensors in the network are capable of monitoring multiple lanes. For this reason, the same sensors that monitor mainline lanes can be configured to also monitor HOV lanes. Exhibit C5-3 shows the sensors that monitor HOV lanes. The monitored HOV lanes are I-75 inside of the ring road (Traficon), the I-75/I-85 Downtown Connector (Traficon), I-85 north of the I-75 split (Traficon), and I-20 from east of downtown Atlanta to east of the ring road. Along each of these freeway segments, HOV lanes are operational seven days a week, 24 hours a day along both directions of travel.

In addition to the real-time detection network, GDOT staff use approximately 500 CCTV cameras positioned at approximately 1-mile intervals on most major interstates around Atlanta to monitor conditions.

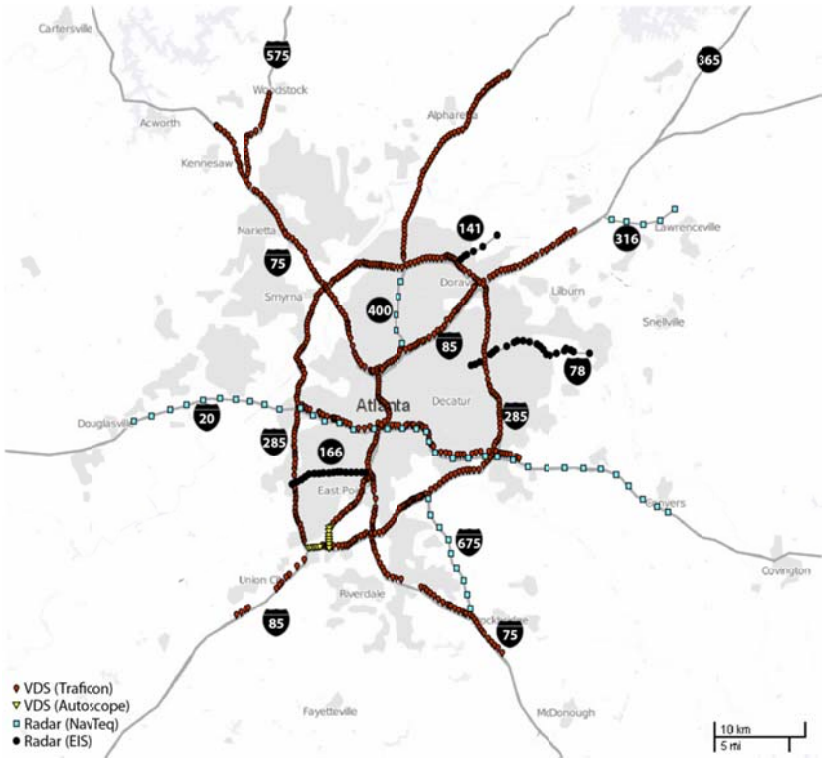

Exhibit C5-2: GDOT Traffic Detector Network

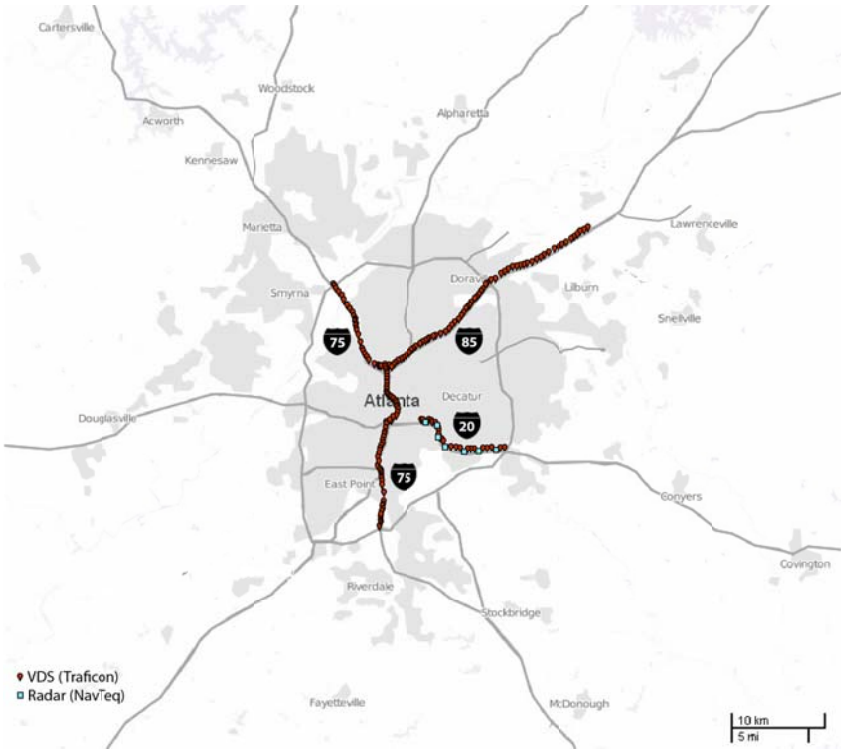

Exhibit C5-3: GDOT Managed Lane Detector Network

Data Management

The primary data management system used in the Atlanta region is the Georgia DOT's Navigator System. Navigator is an Advanced Traffic Management System (ATMS) that was initially deployed in metropolitan Atlanta in preparation for the 1996 Summer Olympic Games. Navigator collects traffic data from video and radar detectors in the field, automatically updates CMSs with travel time information, and controls ramp metering. It also pushes information to the public through a variety of outlets, including a traveler information website and a 511 telephone information service. In addition, Navigator data is used by several private sector companies who enhance and package the data for distribution to media outlets.

The Navigator system is broken up into six subsystems *(2)*:

1) Field Data Acquisition Services
2) Management Services
3) Audio/Video Services
4) System Services
5) Geographical Information Services
6) System Security Services

The Field Data Acquisition subsystem is responsible for device communication and management, and consumes data from CMS, detector stations, ramp meters, a parking management system, and Highway Advisory Radio. The Management Services system helps TMC staff analyze data to determine conditions and develop response plans, and includes the Navigator Graphical User Interface, congestion and incident detection and management services, response plan management, and the historical logging of detector data. The Audio/Video subsystems lets TMC staff control CCTVs in the field as well as the display of information within the TMC. The System Services subsystem communicates speed information with GDOT's Advanced Traveler Information System (ATIS) and logs system alarms. The GIS subsystem provides a graphical view of the roadway network and real-time data. The final subsystem provides system security.

The primary functions of Navigator are the monitoring of and the response to real-time traffic conditions. As such, Navigator collects lane-specific volume, speed, and occupancy data in real-time from the disparate detector types at their respective sampling frequencies (for example, every 20 seconds for the Traficon cameras), and then stores the raw data in a database table for 30 minutes. This database table always contains the most recent 30-minute subset of collected data. An associated table contains configuration data (such as locations and detector types) for all of the devices that sent data within the past 30 minutes. Besides being accessible at the TMC, this raw data is also used to compute travel times on key routes, which are then automatically displayed on regional CMS as well as distributed through traveler information systems. The raw data is not processed or quality-controlled prior to being stored in the real-time data table.

Every fifteen minutes, the raw Navigator traffic data samples are aggregated up to lane-specific 15-minute volumes, average speeds, and average occupancies, and archived for each detector station. The data is not filtered or quality-controlled prior to being archived. Many agencies and research institutions use this data set for performance measurement purposes; for example, the Georgia Regional Transportation Authority (GRTA), the Metropolitan Planning Organization for the Atlanta region, uses it to develop its yearly Transportation AP Report, which tracks the performance of the region's transportation system.

Aside from the traffic data, Navigator also maintains a historical log of incidents. When the TMC receives a call about a incident, TMC staff log it as a "potential" incident in Navigator, until it can be confirmed through a camera or multiple calls. Once the incident has been confirmed, its information is updated in Navigator to include the county, type of incident, and estimated duration. This incident information is archived and stored.

Systems Integration

For the purposes of this case study, data from GDOT's Navigator system was integrated into PeMS, a developed archived data user service and travel time reliability monitoring system. This section briefly describes the steps involved in integrating the two systems. A more detailed account of the integration process and associated challenges is presented in the Use Case chapter of this document.

PeMS is a traffic data collection, processing, and analysis tool that extracts information from real-time intelligent transportation systems (ITS), saves it permanently in a data warehouse, and presents it in various forms to users via the web. PeMS requires three types of information from the data source system (in this case study, Navigator), in order to report performance measures such as travel time reliability:

- Metadata on the roadway linework of facilities being monitored
- Metadata on the detection infrastructure, including the types of data collected and the locations of equipment
- Real-time traffic data in a constant format at a constant frequency (such as every 30-seconds or every minute)

PeMS acquired the first piece of required information- roadway linework and mile marker information- from OpenStreetMap, an open-source, user-generated mapping service.

PeMS acquired the second piece of required information- detection infrastructure metadata- directly from GDOT database tables at the beginning of the integration process. The Navigator data framework is based around two components: devices and detectors. Devices are the physical unit in the field (either the VDS or the radar detector) that collect the data. Detectors represent the specific lanes from which data is being collected. Since all GDOT detectors are VDS or radar, detectors in the GDOT network are virtual, rather than physical, entities. To define devices and detectors, GDOT has database tables that are modified each time that field equipment is added, removed, or modified. The PeMS framework consists of two similar entities: stations (parallel to devices) and detectors. Because of this similarity, the mapping of GDOT infrastructure into PeMS was relatively straightforward. Challenges related to consuming metadata from GDOT's disparate detector types are described in the use case chapter.

PeMS continuously acquires the final piece of required information—real-time data—from GDOT database tables. As described in the Data Management section of this chapter, Navigator stores all of the raw data for the most recent 30-minute period in a database table. To obtain data, PeMS consumes and stores the entirety of this database table every five-minutes, and throws out any duplicate records. The Navigator raw data table is copied into PeMS every five-minutes rather than every thirty-minutes to support the near-real time computation of travel times.

Two aspects of the Navigator framework presented major challenges for incorporating the traffic data into PeMS:

1) The frequency of data reporting differs for different device types; and
2) Many VDS device data samples are missing

These challenges are further discussed in the Use Case chapter of this document.

Other Data Sources

To deepen the case study analysis and explore alternative data sources, the project team acquired a parallel, probe traffic data set, provided by NavTeq. The data set covers the entirety of the I-285 ring road, and is reported by Traffic Message Channel (TMC) ID. The following data is reported every minute for each TMC ID:

- Current travel time
- Free-flow travel time
- Current speed
- Free-flow speed
- Jam factor
- Jam factor trend
- Confidence

The lengths of the TMC segments vary but are generally between 0.3 and 2 miles long. PeMS consumes the NavTeq data through a real-time data feed. While the computational methods and sources of the data are proprietary, the data is generally computed from a mixture of probe and radar data. When there is not sufficient real-time data to generate the reported measures, the data is also based on historical averages. The confidence interval reflects the amount of real-time data used in the computation. This data set is addressed in more detail in the use case section of this document.

To enable investigation into the impact of the seven sources of congestion on travel time reliability, the research team also acquired event data (consisting of incident and lane closure data) collected by Navigator. The issues involved in preparing this dataset for use in analysis are detailed in the first use case. The results of the analysis into the impact of the sources of congestion on unreliability are discussed in the second use case.

Summary

The Atlanta Metropolitan area offers the densest network of fixed point sensors of any of the five sites studied in this project, while presenting the challenges of adapting operational ATMS data for reliability monitoring. The site also provides the opportunity to analyze a third-party probe-based data set.

METHODOLOGICAL ADVANCES

Overview

The methodological advancement of this case study builds upon methods established and validated in previous case studies. Two of the main themes of the case study validations are: (1) estimating the quantity and characteristics of the operating travel time regimes experienced by different facilities; and (2) calculating the impacts of the seven sources of non-recurrent congestion on travel time reliability.

To estimate regimes, the San Diego case study grouped time periods with similar average travel time indices, within which travel time probability density functions were assembled. To refine the regime-estimation process, the Northern Virginia case study validated the use of multi-

state normal density functions to model the multi-modal nature of travel time distributions for a particular facility and time of day. This approach has the advantage of providing a useful, traveler-centric output of the likelihood of congestion and the travel time variability under different congestion scenarios.

With respect to non-recurrent congestion analysis, the San Diego and Lake Tahoe case studies focused on estimating probability density functions for travel times measured during instances of non-recurrent congestion. These distributions help distinguish between the natural travel time variability of a facility due to the complex interactions between demand and capacity, and the travel time variability during specific events.

The methodological goal of this case study is to fuse the previously-developed regime-estimation and non-recurrent congestion analysis methodologies by using multi-state models to inform on the reliability impacts of non-recurrent congestion. Providing a way for agencies to link the travel time regimes that their facilities experience with the factors that cause them, such as incidents or special events, would allow them to better predict travel times when these events occur in real-time, as well as develop targeted projects to improve reliability over the long-term. The background and steps of this analysis are described in this chapter, with detailed results presented in Use Case 2.

Site Description

The methodology was applied to the segment of southbound I-75 starting just north of the interchange with I-85 and ending just north of the I-20 interchange in downtown Atlanta. A map of this corridor is shown below in

Exhibit C5-4. This corridor was selected for the following reasons:

- Significant recurrent congestion during the AM and PM weekday peak periods
- A high frequency of incidents
- Proximity to special event venues, such as the Georgia Dome and Phillips Arena

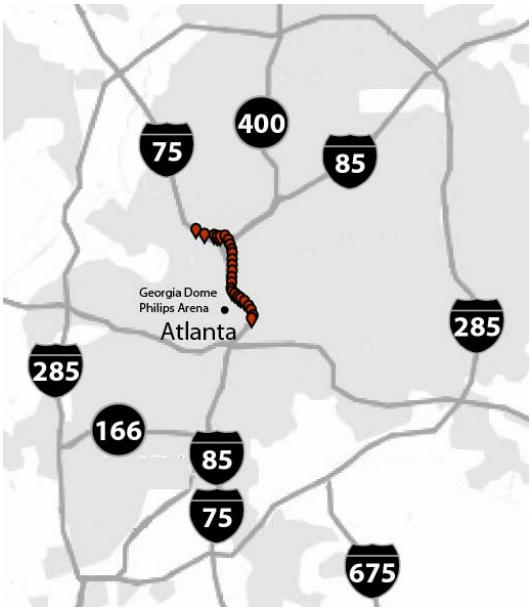

Exhibit C5-4: Downtown Connector Study Route

Method

The method to develop the regimes and estimate the impacts of non-recurrent congestion events consists of three steps:

1) **Regime Characterization**, to estimate the number and characteristics of each travel time regime measured along the facility;
2) **Data Fusion**, to link travel times with the source active during their measurement, and;
3) **Seven Sources Analysis,** to calculate the contributions of each source on each travel time regime.

Regime Characterization

The details of how to implement multi-state normal models for approximating travel time density functions are thoroughly described in the Methodology section of the Northern Virginia case study. With multistate models, the data set is modeled as a function of the probability of each state occurring and the parameters of each state. In generalized form, multistate models take the form of Equation 1,

$$f(T|\lambda,\theta) = \sum_{K=1}^{K} \lambda_K f_K(T|\theta_K) \qquad (1)$$

where T is a travel time, $f(T|\lambda,\theta)$ is the travel time density function for the data set, K is the state number, $f_K(T|\theta_K)$ is the density function for travel times in the Kth state, λ_K is the probability of the Kth state occurring, and θ_K is the distribution parameters for the Kth state. For the multistate normal distribution, θ_K is composed of the mean (μ) and the standard deviation (σ) of the state's travel times.

More practically, if a three-state normal model provides the best fit to a set of travel times collected at the same time of day over multiple days, the first state can be considered the least congested state, the second state a more congested state, and the third state the most congested state. Each state is defined by a mean travel time and a standard deviation travel time, with the first state having the fastest mean travel time and the third state having the slowest mean travel time.

The development of a multi-state model consists of two steps: (1) identifying the optimal number of states to fit the data; and (2) calculating the parameters (probability of occurrence and mean and standard deviation travel times) to define each state. The methods for performing these tasks are described in the Northern Virginia case study.

In addition to providing the number of operating states and their parameters, the model also outputs, for each measured travel time, the percentage chance that it belongs within each state. By assigning each travel time to the state it is most likely to belong to, it is possible to derive a set of travel times that belong within each state. This output is used to drive the non-recurrent congestion reliability analysis, described in the following subsection.

Data Fusion

To test the methodology, the research team downloaded five-minute travel times measured on all non-holiday weekdays between September 9th, 2011 (the first day that PeMS was set up for data collection) and December 31st, 2011 from the reliability monitoring system.

Due to drops in the data feed, there were many days of missing data during the months of November and December. Each travel time was then manually tagged with the source active during its measurement, following the methodology used and described in the San Diego case study, and briefly summarized below. The following sources were included in the fusion process:

1) **Baseline**. No source was active during the five-minute time period.
2) **Incident**. Incident data was acquired from Georgia Tech's Navigator event data archive. The challenges of quality-controlling the incident data set are described in the first use case of this document. The research team ultimately associated incident travel times with the following types of events that were marked as blocking at least one lane in the incident data set:
    a) Accident/Crash
    b) Debris (all types)
    c) Fire/Vehicle
    d) Stall/Lane(s) Blocked

    In previous case studies, the research team assumed that incident impacts began at the start time of the incident and ended fifteen minutes after the incident closed, to allow for queue discharge. However, because the incident durations seemed unusually long in this data set, for this study, it was assumed that incident impacts ended at the incident closure time.
3) **Weather**. Hourly weather data was downloaded from the NOAA National Data Center and was measured at a weather station housed at Atlanta Hartsfield-Jackson International Airport (located approximately 10 miles southwest of the study corridor). The research team assumed that weather impacts were incurred when greater than $1/10^{th}$ of an inch of precipitation was measured during the hour. The Navigator event data set also documented instances of roadway flooding (through the incident type "Weather/Road Flooding"). Travel times measured during these events were also associated with this source.
4) **Special Events**. Special event data from the Georgia Dome and Philips Arena was collated manually from sport and event calendars. Determining when special events impact traffic is challenging, as the impact of the event depends on the type of event. Typically, event traffic impacts begin prior to the start time, and end after the event is over. However, while event start times are typically available, event end times are rarely explicit and have to be assumed. In this study, a travel time was tagged with "special event" if it occurred up to one hour before the event start time and in the hour following the estimated end time.
5) **Lane Closures.** Lane closures were gathered from the Georgia Tech's Navigator event data archive, which contained events marked as "Planned/Maintenance Activity", "Planned/Construction", and "Planned/Rolling Closure". The research team tagged travel times with the lane closure source if a closure affecting at least one lane was active during the five-minute time period.

In the San Diego case study, fluctuations in demand were also measured. In Atlanta, fluctuations in demand were not able to be analyzed due to the high quantity of missing data samples, which impacted the ability of the system to monitor traffic volumes (as explained in Use Case 1).

Seven Sources Analysis

The model development process described above results in a set of travel times, each tagged with the non-recurrent congestion source active during their measurement, that are categorized according to the state that they belong to. From this it is possible to calculate two key measures to inform on the relationships between non-recurrent congestion and the travel time regimes:

1) Within each state, the percentage of travel times measured during each source
2) For each source, the percentage of its travel times that belong in each state

The use case section presents the results of these two measures for a freeway corridor in downtown Atlanta. It also visualizes the results through travel time histograms divided into states and color-coded according to the source active during the travel time's measurement.

Results

Results are presented in Use Case 2.

USE CASE ANALYSIS

Use Case 1: Integrating ATMS Data into Travel Time Reliability Monitoring System

Summary

For this case study, data from GDOT's Navigator ATMS system was brought into a travel time reliability monitoring system (PeMS) and archived to support the computation of historical and real-time travel times and reliability metrics. This case study was the project team's first opportunity to use ATMS data, which is focused on real-time congestion and incident detection, for monitoring travel time reliability. To contrast with the previous case studies, the San Diego and Lake Tahoe sites relied primarily on data within PeMS that had already been quality-controlled and processed, and the Northern Virginia site leveraged data collected from an archived data user service at the University of Maryland. In each of these cases, the data leveraged by the project team had already been processed to fill in any data holes and aggregated to ensure a consistent granularity across all of the raw data samples. Because ATMS data is conventionally used only for real-time operations, the acceptable level of data quality is much lower than it is for the analysis of archived data. Conceptually, it is easier for TMC staff to identify gaps and errors in the real-time data, since they have access to other data sources such as CCTV cameras and reports from the field, than it is for analysts who are evaluating historical travel times and performance measures without the benefit of any other contextual information. Given the nature of the Atlanta data, initial case study efforts focused on the integration issues with consuming unprocessed, incomplete data from disparate sensor types and using it to compute travel time reliability. Encountered issues fell into two categories: (1) metadata integration, where GDOT device and detector information is transferred into PeMS; and (2) data integration, where real-time traffic data is consumed by PeMS, processed, cleaned, and stored, and ultimately used to measure travel times and reliability. The project team acquired metadata and traffic data through direct access to the relevant Navigator database tables. This use case describes the challenges of interpreting the information in the database tables and inputting

it into PeMS. It also describes the process for interpreting the event data acquired from Georgia Tech from Navigator.

Metadata Integration

As described in the Monitoring System chapter, the data model for Navigator detection devices (devices containing multiple detectors) is very similar to the PeMS data model (stations containing multiple detectors). As such, the mapping between the two system models was trivial, and the primary metadata integration challenge was interpreting the fields and formats of the Navigator metadata database tables, and filtering out non-active infrastructure. Navigator defines devices and detectors in two separate database tables. The project team acquired complete copies of these database tables at the beginning of the integration project, and used them to generate the detection network for PeMS.

The device database table contained 14,581 rows, with nearly all device IDs having multiple records corresponding to different version numbers (up to 14 for some devices). The version number appeared to be driven by "modified date" column, with the highest version numbers corresponding to the most recent modified date. As such, the set of devices was reduced to a single record for each device ID with the highest version number. This step reduced the number of devices to 4,633. After excluding those missing latitude and longitude information, which PeMS requires, 3,406 unique devices remained.

The detector database table contained 40,496 records, which was filtered down to 34,135 after excluding detectors associates with devices that had missing locations. Each detector was assigned a "lane_type". PeMS assigns detectors to one of six possible lane types: (1) mainline; (2) HOV; (3) on-ramp; (4) off-ramp; (5) collector/distributor; and (6) freeway to freeway connector. When assessing the Navigator detector lane types, the project team noted a total of 21 possible categories. This high number is because Navigator, because of its operational nature, allows for the same type of lane to be identified in different ways. For example, in the detector database table, the lane types "Entrance Ramp", "Entrance_ramp", "Left_entrance_ramp", "Right_entrance_lane", "Right_entrance_ramp, are all used to denote on-ramp detectors. This required the development of a mapping structure to appropriate categorize Navigator detectors in PeMS, as shown in Table C5-2. In doing this, the project team noted that a large percentage of the devices that had no locations monitored "arterial" detectors. The research team hypothesizes that these devices were planned for deployment, but were not yet configured to report data into the system.

Table C5-2: Mapping of Lane Types from Navigator to PeMS

PeMS Lane Type	Navigator Lane Type
Mainline	Mainline Through_lane Through_lanes Through-lanes THRU/THRU THRU/OFF-RAMP (THRU) THRU/ON-RAMP (THRU)
HOV	High Occupancy Vehicle Hov_lanes THRU/HOV
On-Ramp	Entrance Ramp Entrance_ramp Left_entrance_ramp Right_entrance_lane Right_entrance_ramp
Off-ramp	Exit Ramp Right_exit_lane Right_exit_ramp
Collector/Distributor	Collector/Distributor
Freeway to Freeway Connector	Connecting Lanes
N/A	Arterial

Using the above-structure, Navigator devices and detectors were mapped as stations and detectors in PeMS. This allowed for the step of the real-time data integration, described in the next subsection, to begin.

Agency Data Integration

As described in the Monitoring System chapter, two characteristics of the GDOT detection network presented major data integration challenges for the case study: (1) variable sample rates across detectors; and (2) missing data samples for detectors and devices.

Varying data sampling rates are problematic because PeMS assumes that all detectors within the same data feed report data at a constant, known frequency (for example, in the San Diego case study, this frequency is every 30 seconds). This assumption enables the accurate aggregation of raw data up to the five-minute level, from which travel times and other measures are then calculated. While all GDOT detectors report flow, occupancy, and speed, the frequency at which they report it varies. GDOT stores the most recent 30 minutes of data from each active detector in a database table. PeMS obtains real-time data from GDOT by copying over the GDOT raw database table every five minutes then eliminating duplicate records already acquired in previous five-minute periods. An initial manual review of the database table showed a data reporting frequency of every 20 seconds, so this was the basis for aggregation up to the five-minute level. Through inspection of the aggregated data, however, it became evident that the

frequency of data reporting varies by the vendor type. Table C5-3 shows the observed reporting frequencies by vendor type.

Table C5-3: Data Reporting Frequencies by Device Type

Vendor	Reporting Frequency
Traficon	20 seconds
Autoscope	75 seconds
NavTeq	60 seconds
EIS	20 seconds

As such, while the majority of GDOT detectors report data every 20 seconds, a significant number do not, and thus were not being aggregated correctly in PeMS. The research team decided that the best way the handle this issue was to change the process for extracting data from the GDOT raw database table. Instead of extracting data from al detectors in a single feed, the problem could be solved by establishing three data feeds, each with their own aggregation routines, to obtain data from all detectors that report at the same frequency (20 seconds, 60 seconds, and 75 seconds).

The second issue identified by the research team was that a significant number of expected data samples were missing. For example, since Traficon detectors are configured to send data every 20 seconds, and GDOT stores the most recent 30 minutes of data from each detector, the research team expected to see 90 samples for each Traficon detector in each copy of the database table. Instead, many 20 second time periods were missing data for one or more detectors. For many of the VDS detectors, almost no samples were reported during the nighttime hours. From this, the research team concluded that some of the detectors were not able to monitor traffic in the dark. Many samples were also missing during the daytime hours. This, combined with the fact that none of the data samples ever reported zero volume, made it clear that the detectors send no data sample if they detect no vehicles during the time interval. This data reporting scheme is problematic because monitoring systems need to be able to distinguish between when the detector or data feed is broken (requiring data imputation to fill in the hole) and when no vehicles traveled past the location during the time interval (requiring a recording of zero volume in the database). With PeMS, the GDOT detector reporting framework causes two main problems.

1) PeMS performs detector diagnostics at the end of every day. If more fewer than 60% of expected data samples are received, then the detector is deemed to be broken and all of its data is imputed;
2) PeMS performs imputation for missing data samples in real-time. If the cause of the missing sample is that there were no vehicles at the location over the time period, then the imputation results in an over-counting of volumes.

In the Atlanta site, the first issue was deemed minimal because PeMS only runs the detector diagnostics on samples collected between the hours of 5:00 AM and 9:00 PM. Since the majority of missing samples occur outside of these hours (in the middle of the night), very few detectors sent less than 60% of expected samples during the diagnostic hours. The second issue, however, was deemed more serious, because it means that volumes are over-estimated and speeds are estimated from unnecessary amounts of imputed data. The ideal, permanent solution to mitigate both issues would be to change the way that the field equipment interacts with the data collection system, to ensure that data samples are sent even when no traffic is measured.

This change would need to be made at the device level. However, because this was a case study validation effort and not a procured monitoring system for GDOT, the team decided that the following solution would be more practical:

1) Turn off real-time imputation to allow missing data samples
2) Calculate five-minute volumes by summing up the non-missing raw data samples
3) Calculate five-minutes speeds by taking the flow-weighted average of the non-missing raw data samples
4) Compute travel times from all detectors with non-missing five-minute travel times samples along a route.

The end result of this solution is that the volume-based performance measures (such as vehicle-miles-travelled and vehicle-hours-of-delay) may be under-reported, but speed-based measures are more accurate than they would be under the PeMS traditional real-time imputation regime.

Event Data Integration

To enable seven sources analysis, the research team acquired a database dump of all Navigator events (primarily incidents and lane closures) from September through December 2011 from Georgia Tech. The data was delivered in an excel spreadsheet in a format summarized in Table C5-4. It contained 21,540 event records summarizing Navigator events within the Atlanta metropolitan region.

Table C5-4: Event Data Format

Column	Name	Description	Example
1	ID	Unique ID	244835
2	Primary Road	Freeway number	I-75
3	Dir	Direction of travel	N
4	MM	Mile marker	228
5	Cross	Cross-street	Jonesboro Rd
6	County	County	Clayton
7	Start	Event start date and time	09/01/2011 01:00
8	End	Event end date and time	09/02/2011 06:15
9	Type	Type of event	Accident/Crash
10	Status	Status of event	Terminated
11	Blockage	2	Number of lanes blocked

The breakdown of events by type in the data set is shown in Table C5-5 (grouped and summed into event types in similar categories).

Table C5-5: Event Data Set by Event Type

Type	Number
Accident (Crash, Haz Mat Spill, Other)	3,311
Debris (Animal, Mattress, Tire, Tree, Other)	1,896
Fire (Structural, Vehicle, Other)	237
Infrastructure (Bridge Closure, Downed Utility Lines, Gas/Water Main Break, Road Failure)	120
Planned (Accident Investigation, Construction, Emergency Roadwork, Maintenance Activity, Rolling Closure, Special Event)	4,499
Signals (Bulb Out, Flashing, Not Cycling)	638
Stall (Lane(s) Blocked, No Lanes Blocked)	10,690
Unplanned (Live Animal, Policy Activity, Presence Detection , Rolling Closure)	55
Weather (Dense Fog, Icy Condition, Road Flooding)	99

The data was assessed with an eye towards its ability to detail incidents and lane closures on a ten-mile segment of southbound I-75, for use in analyzing the impacts of the seven sources of congestion on travel time variability on this corridor (see the Methodological Advancement Chapter for more details). In doing this, the team noted the following data set characteristics that complicated the assignment of incidents and lane closures to measured travel times:

1) The same freeway was given different names in the "Primary Road" column;
2) Mileposts were missing from some events;
3) There were inconsistencies between the number of lanes blocked in the event type column and the blockage column; and
4) Durations for many of the events were longer than expected for the event type

With respect to the first issue, the segment of I-75 studied in the document was given the following different names in the data set: 75/85, I-75, 75/85 SB, I-75/85, and 75. As such, the research team had to ensure that all of the possible freeway names were evaluated and narrowed down by milepost so as not to miss any events on the study route. The second issue was dealt with by manually mapping the given cross-street to determine if the location was on the study segment. The third issue related to the numerous events of type "Stall, Lane(s) Blocked" and "Stall, No Lanes Blocked" where the degree of lane blockage was contradicted by the number in the "Blockage" column. In these cases, the research team used the event type description to determine if there was lane blockage. The fourth issue regards event durations; in many cases, the event duration computed from the start and end times seemed longer than would be expected for an event of that type. For example, it was common to see events of type "Stall, No Lanes Blocked" last for longer than 3 hours. Without any other source of data to reference, the research team simply had to accept the reported durations, and note it as a potential inaccuracy in the analysis.

Conclusions

Because most metropolitan areas are already equipped with ATMS detection and software systems, ATMS data is a likely source of information for urban travel time reliability monitoring systems. The integration of ATMS data into a travel time reliability monitoring

system presents challenges in ensuring data quality and quantity. Practitioners may encounter the following issues when acquiring and integrating ATMS data for reliability monitoring purposes:

1) Sensor metadata and event data with missing required attributes, such as location
2) Sensor metadata and event data with unstandardized naming classification
3) Data at miscellaneous sampling rates
4) Missing data samples

When required sensor information is missing, the only alternative to obtaining the information from the field is to discard the sensor from the reliability monitoring system. For unstandardized classifications, the best alternative is to manually translate ATMS terminology into the monitoring system framework, prioritizing the translation of mainline and managed lane detectors. The data variability issues are more challenging to deal with, and are best solved on a permanent level by changing the way that the field equipment communicates with the ATMS system, to ensure that all the information needed for historical travel time monitoring is required.

Use Case 2: Determining Travel Time Regimes and the Impact of the Seven Sources of Congestion

Summary

The Northern Virginia case study analyses developed methodologies for modeling the multi-modal nature of travel time distributions to determine the operating regimes of a facility. The San Diego case study analyses validated ways to evaluate the impact of the seven sources of congestion on travel time variability. This use case seeks to combine these two methods to identify the impacts of the seven sources of congestion on the different travel time regimes that a facility experiences. The methodology that drives this analysis and a description of the study route is presented in the Methodological Advancements chapter of this document. This use case write-up documents the results of performing the regime characterization, data fusion, and seven sources analysis steps on a ten-mile study route through Downtown Atlanta during the weekday AM, midday, and PM periods.

Results

Regime Characterization. The first step in the analysis is to identify the number of modes, or regimes, in the travel time distribution. In this study, the data set consisted of five-minute travel times measured on non-holiday weekdays between September 9th, 2011 and December 31st, 2011. To appropriately identify the number of operating regimes along the study route, the travel time data set was grouped by similar typical operating conditions (defined by the mean travel time) and time of day into the following categories:

- AM Peak, 7:20 AM – 9:20 AM, (mean travel times exceeding 14 minutes)
- Midday, 9:30 AM – 4:00 PM, (mean travel times less than 13 minutes)
- PM Peak, 5:00 PM – 6:20 PM (mean travel times exceeding 18 minutes)

An algorithm in R was used to identify the optimal number of multi-modal normal regimes to model each of the three travel time datasets. Results showed that the AM and PM peak time periods were best modeled with two normal distributions and that the midday period was best modeled with three normal distributions. Exhibit C5-5, Exhibit C5-6, and Exhibit C5-7 show a histogram of the travel time distribution for each time period, as well as the probability

density functions for each of the regimes (the dashed lines) and the overall mixed-normal density function (the solid line).

Table C5-6 summarizes the regime parameters (probability of occurrence and mean travel time) by time period.

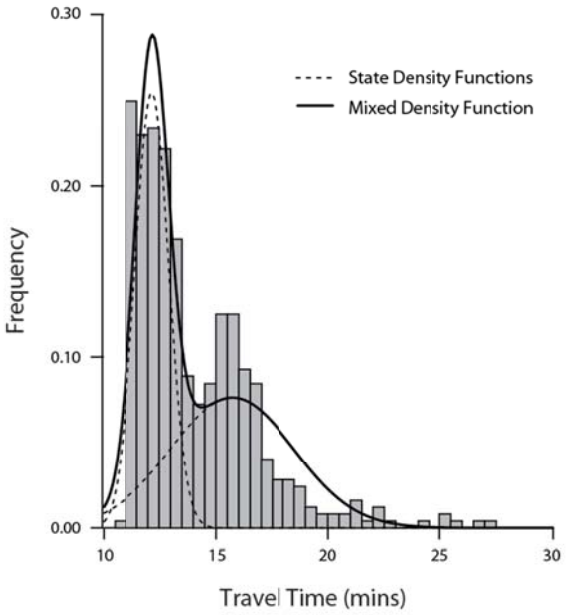

Exhibit C5-5: AM Multi-state Normal PDFs

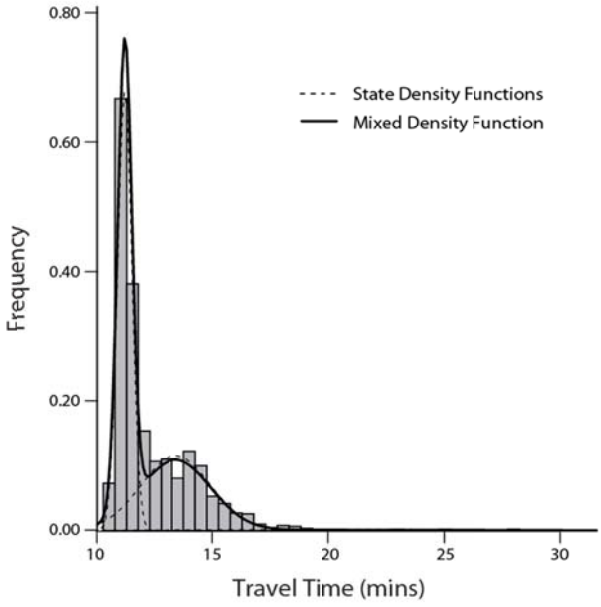

Exhibit C5-6: Midday Multi-state PDFs

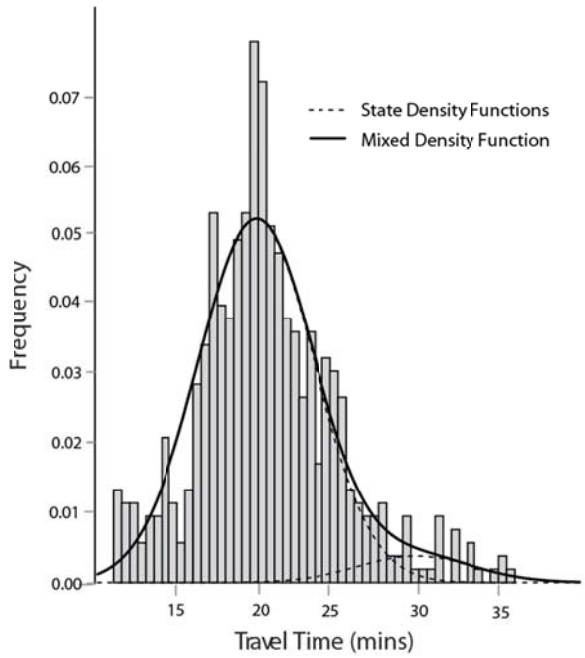

Exhibit C5-7: PM Multi-state PDFs

Table C5-6: Regime Parameters by Time Period

	Probability (%)			**Mean Travel Time (mins)**		
	State 1	**State 2**	**State 3**	**State 1**	**State 2**	**State 3**
AM	47%	53%	--	12	16	--
Midday	52%	44%	4%	11	14	18
PM	92%	7%	--	20	30	

In the AM peak, each regime (uncongested and congested) occurs about half of the time. The mean of the first, uncongested regime is 12 minutes, with little travel time variability in the distribution. The mean of the congested regime is 16 minutes, and the distribution of travel times is wider.

The midday period has three regimes. The uncongested regime happens 52% of the time, the slightly congested regime happens 44% of the time, and the congested regime happens only 4% of the time (this small percentage makes the regime invisible in Exhibit C5-6). The mean of the uncongested regime is 11 minutes (free-flow), the mean of the slightly congested regime is 14 minutes, and the mean of the most congested regime is 18 minutes.

The PM period is characterized by two regimes. The congested regime happens 92% of the time, with a mean travel time of 20 minutes (almost double the free-flow travel time). The

very congested regime happens only 7% of the time, but has a mean travel time of 30 minutes (almost three times the free-flow travel time).

Data Fusion

In the data fusion step, the seven sources data described in the Methodological Advancements chapter was fused with the five-minute travel times. Table C5-7 summarizes the number and percentage of travel time samples by source within each time period. Special events only occurred during the PM time period. Conversely, lane closures only occurred during the AM and midday time periods. Incidents made up a similar percentage of the data set in all three time periods.

Table C5-7: Five-minute Travel Time Samples by Time Period and Source

	AM	**Midday**	**PM**
Baseline	297 (60%)	1,254 (71%)	413 (78%)
Incident	77 (16%)	286 (16%)	73 (14%)
Weather	115 (23%)	119 (7%)	36 (9%)
Special Event	0 (0%)	0 (0%)	10 (2%)
Lane Closure	7 (2%)	115 (6%)	0 (0%)
Total	**496**	**1774**	**532**

Seven Sources Analysis. The final step in the analysis is to assess the contributions of the sources of congestion to each travel time regime. Exhibit C5-8, Exhibit C5-9, and

Exhibit C5-10 illustrate the breakdown of travel times by source within each state. Table C5-8, Table C5-9, and Table C5-10 respectively summarize each state's parameters, the percentages of each state's travel times tagged with each source, and the percentage of each source's travel times that occur within each state.

During the AM peak, state 2 has a four-minute higher mean travel times than state 1, and also contains more variability (a standard deviation of 3 minutes versus less than a minute). Incident travel times are seen in both states, but incidents are three times more likely to result in the most congested state. Weather events, in contrast, are found more frequently in the uncongested state (58%) than the congested state (42%). There were not very many lane closure samples to evaluate, so lane closures do not appear to be a driving factor of AM peak congestion and travel time variability on this route. State 2 contains a significant number of baseline travel times (51%), indicating that something other than incidents, weather, and lane closures is causing delay and unreliability on this corridor during the morning commute.

The midday peak has three states. The most congested state, which occurs only 4% of the time, is composed of around one-third weather-influenced travel times, one-fifth incident-influenced travel times, and one-tenth lane-closure travel times, and the remainder baseline travel times. The fact that the less congested states contain a significant proportion of the congestion-influenced travel times indicates that only the most severe instances of the sources result in a reduction in capacity below the midday demand levels.

During the PM peak, the congested state that happens 93% of the time (state 1) contains nearly all of the congestion source travel times. However, this state has a wide distribution of travel times, and

Exhibit C5-10 shows that many of these incident- and weather-influenced travel times occupy the right-most part of the state 1 travel time distribution. The very congested second state during the PM peak is composed of one-third weather-influenced travel times, one-tenth incident influenced travel times, and the rest baseline travel times, indicating that this most unreliable state is caused by some other influence.

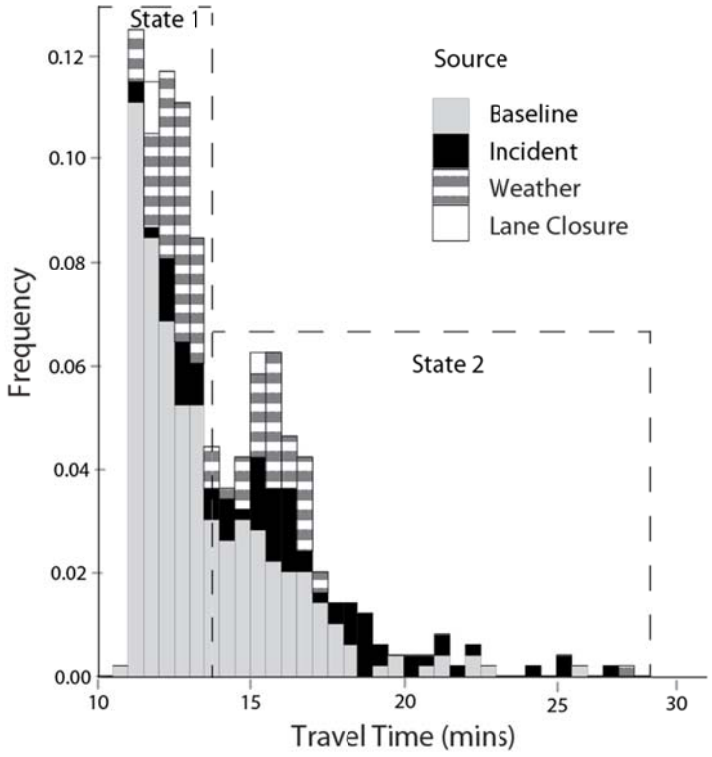

Exhibit C5-8: AM Peak Travel Times by Source

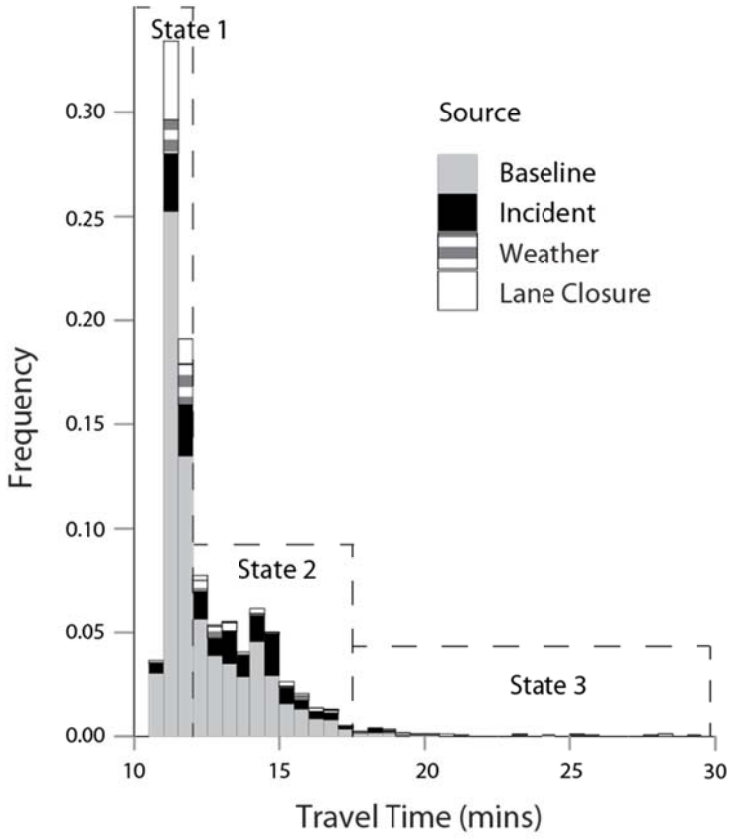

Exhibit C5-9: Midday Peak Travel Times by Source

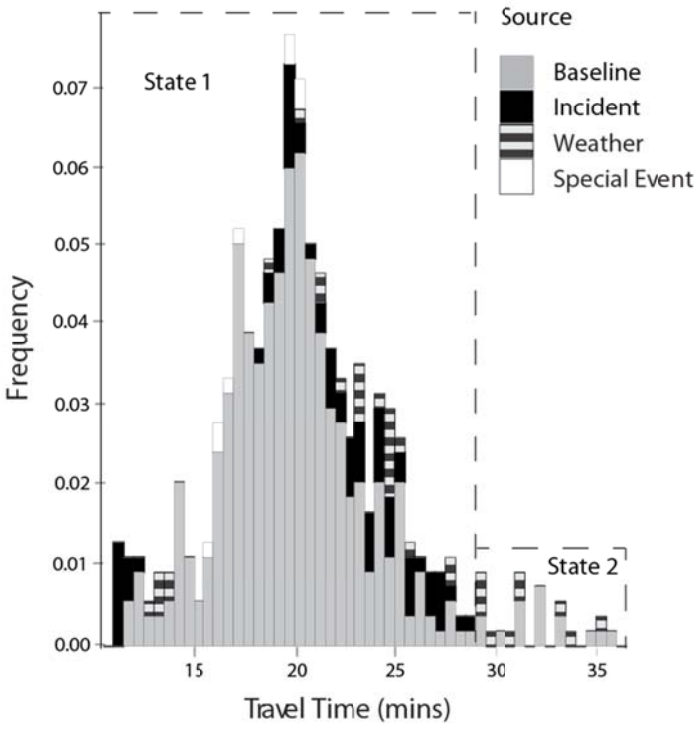

Exhibit C5-10: PM Peak Travel Times by Source

Table C5-8: Source Contributions to AM Peak Regimes

	State 1	State 2
Parameters		
Probability	47%	53%
Mean	12 minutes	16 minutes
Standard Deviation	0.7 minutes	3 minutes
Percentage of State Travel Times by Source		
Baseline	67%	51%
Incident	7%	26%
Weather	24%	22%
Special Event	0%	0%
Lane Closure	2%	1%
Percentage of Source Travel Times by State		
Baseline	62%	38%
Incident	25%	75%
Weather	58%	42%
Special Event	0%	0%
Lane Closure	71%	29%

Table C5-9: Source Contributions to Midday Regimes

	State 1	State 2	State 3
Parameters			
Probability	52%	44%	4%
Mean	11 minutes	14 minutes	18 minutes
Standard Deviation	0.2 minutes	3 minutes	4 minutes
Percentage of State Travel Times by Source			
Baseline	75%	67%	32%
Incident	10%	24%	20%
Weather	6%	6%	35%
Special Event	0%	0%	0%
Lane Closure	9%	3%	13%
Percentage of Source Travel Times by State			
Baseline	59%	40%	1%
Incident	36%	62%	2%
Weather	54%	34%	2%
Special Event	0%	0%	0%
Lane Closure	78%	17%	4%

Table C5-10: Source Contributions to PM Peak Regimes

	State 1	State 2
Parameters		
Probability	93%	7%
Mean	20 minutes	30 minutes
Standard Deviation	4 minutes	4 minutes
Percentage of State Travel Times by Source		
Baseline	79%	59%
Incident	14%	7%
Weather	5%	34%
Special Event	2%	0%
Lane Closure	0%	0%
Percentage of Source Travel Times by State		
Baseline	96%	4%
Incident	97%	3%
Weather	72%	28%
Special Event	100%	0%
Lane Closure	0%	0%

Conclusions

By combining the regime-estimation and seven sources analysis methodologies used in previous case studies, this application showed that it is possible to characterize the impact of the sources of non-recurrent congestion on the different travel time states that a facility experiences.

On the study route of I-75 into Downtown Atlanta, the analysis showed that a driving factor other than weather, incidents, lane closures, and special events is a leading factor of the high and unreliable travel times that make up the right-most portion of the travel time distribution. This factor may be fluctuations in demand and capacity due to a bottleneck; these factors were not measurable at this case study site. On this route, weather is the source that, when it occurs, most frequently drives the travel time regime into the most congested state.

Use Case 3: Quantifying and Explaining the Statistical Difference Between Multiple Sources of Vehicle Speed Data.

Summary

This use case identifies issues associated with the integration of data feeds from multiple sources. Speed measurements from Traficon video detectors and Navteq probe vehicle runs are compared. For each of these technologies, the data comes from a 10-mile segment of I-285 in Atlanta, Georgia where peak period congestion is observed on weekdays. Some preprocessing was necessary to translate the data sets into a common format which could be easily compared. At that point, correlations between pairs of detectors of each type at the same location were computed. A possible source of difference in the measurements, the distance between each pair of compared detectors, was analyzed and found to be moderately significant.

Data from multiple sources, if properly understood, can be aggregated to provide a rich set of performance monitoring information. Multiple data sources add redundancy to the system, preventing a data blackout in the event that one of the data feeds goes down. Multiple data sources also facilitate the cross-validation of detectors, providing an additional way to identify malfunctioning equipment. However, if the additional data sources are integrated incorrectly, they can conflict with each other, decreasing the accuracy of the monitoring system in unpredictable ways.

The observed traffic data is the fundamental driver of the performance measures computed by a travel time reliability monitoring system. While the underlying traffic model also influences the performance measures, its influence is typically static. For example, a particular methodology for computing travel times may be consistently biased towards overestimating travel times. A systematic bias like this can be recognized and accounted for. On the other hand, the effects of misconfigured data sources can change as the incoming data changes. Understanding the peculiarities of data from different sources is critical since the observed data feeds directly into the measures computed by the monitoring system.

Users

This use case is applicable to all users of travel time reliability monitoring systems, particularly those systems that integrate data from multiple sources or technologies. It provides practical guidance on how to properly compare traffic measurements from multiple data sources. The data comparison techniques presented here are the necessary first steps to transform raw detector data from multiple sources into aggregated traffic information. This information will give important context to users of travel time reliability monitoring systems, improving their understanding of the performance measures they compute.

Information technology professionals responsible for the data integration and preprocessing tasks necessary to build and maintain a travel time reliability monitoring system

will also benefit directly from this use case. This use case provides guidance on the steps necessary to compare data from two different sources, a necessary initial step in data integration. Understanding these issues can also help system managers more easily troubleshoot systems whose computed performance measures are suspect. For example, data feeds that are aggregated incorrectly can be compared using the techniques presented in this use case as part of a troubleshooting routine.

This use case is also valuable to transportation professionals interested in exploring new data sources. GPS-based probe data is increasing in availability and offers a roadway monitoring solution that is rich, with speed and position measurements taken from actual vehicles throughout their trip. Probe data is also appealing because it does not require any ongoing maintenance of detection equipment. With this technology, there is no roadway-based detection hardware; the data collection infrastructure resides entirely within the vehicles themselves. When compared with conventional infrastructure-based sensors, which only record roadway information at discrete locations and must be regularly maintained, probe data can be very appealing. This use case provides guidance on how probe data compares with more traditional infrastructure-based data sources.

Data Characteristics

This use case compares two types of traffic data: (1) speed data from vehicle probes, provided by Navteq, and (2) speed data from Traficon video detectors. The vehicle probe data comes from GPS chips residing within individual vehicles, directly measuring their speed and location. In contrast, the Traficon data comes from video cameras installed at fixed locations along the roadway, measuring speed, volume, and density. Data from infrastructure-based sensors such as these (and loop detectors) is currently much more common than probe data. For this reason, many users of travel time reliability monitoring systems conceptualize the data they see primarily in terms of fixed-infrastructure sensors. The rising availability of probe data for transportation system monitoring makes the Navteq probe data a desirable data set to compare with fixed-infrastructure data.

Because the video data comes from fixed-infrastructure sensors and the probe data comes from in-vehicle sensors, they require different types of network configurations to relate them to the roadway. The video data is organized by device, with each device applying to a single location on the roadway. Data from each device then corresponds to traffic at that point. The probe data, on the other hand, is organized directly by location through Traffic Message Channel (TMC) paths. Each TMC path represents a stretch of roadway in a single direction, and is explicitly defined by a starting and ending milepost. The lengths and locations of the TMC paths are irregular, and there are gaps between TMC paths.

The Navteq probe data differentiates between mainline speeds and speeds on managed lanes such as HOV or HOT lanes, although it does not provide mainline speeds disaggregated by lane. A data point is calculated for each TMC path roughly every two minutes (0.5 Hz). This is a lower sampling rate than many other types of detectors, however since the measurements are taken directly from actual vehicles (representing ground truth conditions), they are generally considered more accurate, making sampling frequency less important.

The Traficon video detector data closely resembles traditional infrastructure-based data such as that from loop detectors. Each video detector is assigned to a specific milepost and lane on the roadway, and its measurements apply directly to that point location. Each video detector

directly reports occupancy, speed, and flow at a maximum frequency of once every 20 seconds (3 Hz). This frequency is comparable to that of most loop detectors.

Sites

A 10-mile stretch of I-285 around Atlanta (known locally as "The Perimeter") was chosen for this study for several reasons. As discussed previously, I-285 is covered by both Traficon video detectors and Navteq probe data, and this location has good data availability for both. The heavy commute traffic on I-285 leads to strong peak period congestion and a range of congestion levels, another reason this site was chosen. I-285 carries the largest volume of traffic of any Atlanta freeway, providing the metropolitan area access to major interstates I-20, I-75, and I-85, which lead to several residential suburbs.

Data covering both the Northbound and Southbound directions of travel was examined. The study area spanned mileposts 25 to 35 in the northbound direction, and 45 to 55 in the southbound direction. Although these milepost ranges differ, they represent the same stretch of roadway (see Exhibit C5-11). The study area extends from the Belvedere Park area at its southern end to the I-85 interchange at its northern end. During the time period studied, free-flow speed was measured around 70 mph. The typical weekday flow was 80,000 to 90,000 veh/day northbound and approximately 100,000 veh/day southbound.

In the Northbound direction, 3.9 of the 10 miles in the study area are covered by 8 TMC paths, with an average TMC path length of 0.5 miles. Also in the Northbound direction are 24 working Traficon detectors, seven of which lie within a TMC path. In the Southbound direction, 5.3 of the 10 miles in the study area are covered by eight TMC paths, with an average TMC path length of 0.7 miles. Also in the Southbound direction are 19 working Traficon detectors, 12 of which lie within a TMC path (see Exhibit C5-12).

One reason this site was chosen is its congestion patterns. AM peak period congestion was seen in the northbound direction between 6 and 9 AM. PM period congestion was seen in the southbound direction between 4 and 7 PM. In both directions, the congestion was most pronounced on Tuesdays, Wednesdays, and Thursdays. 5-minute speed measurements were commonly observed in both directions as low as 15 mph.

Methods

The comparison of the probe and video speed data begins with the procurement of that data. PeMS began collecting live Traficon video detector data in the Atlanta region on September 9, 2011. Data from this initial date through December 23, 2011 (the beginning of a gap in availability) was obtained for the 51 total video detectors in the study area from PeMS. All available data for each detector was included, weekends in addition to weekdays, in order to compare the data sets across a range of conditions. PeMS stores Traficon video detector data at 5-minute resolution at the finest, which is the level of aggregation used in the comparison. It was immediately observed that 2 northbound and 6 southbound video detectors were not reporting any data and they were discarded.

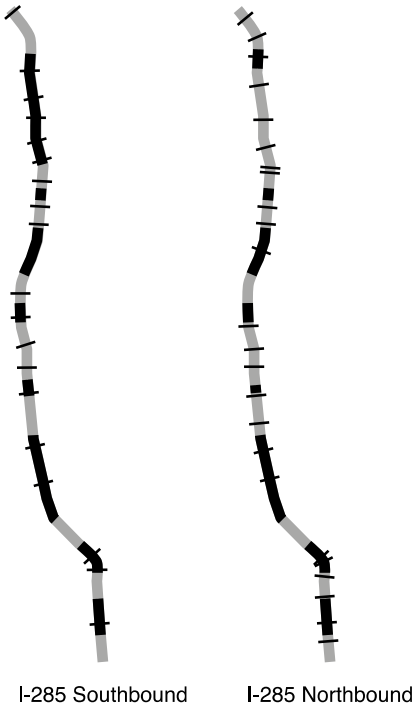

Exhibit C5-11: Locations of Navteq TMC paths (longitudinal black lines) and Traficon video detectors (perpendicular black lines)

Exhibit C5-12: Study Area on I-285

PeMS began archiving Navteq probe data in the Atlanta region on September 18, 2011. All available data from this date through December 23, 2011 was obtained from all 17 TMC paths in the study area. Each probe data point is the result of Navteq's aggregation of many GPS measurements from multiple vehicles into a single speed value for a particular TMC path. PeMS stores these aggregated speed measurements at their finest provided resolution, which is one data point roughly every two minutes (0.5 Hz).

In order to properly compare the two data sets it is immediately necessary to convert them to a common time standard. As obtained from PeMS, the video data and probe data have different time ranges and different sampling frequencies. A perl script was written to fix the time range of all data sets to extend between September 9, 2011 and December 23, 2011, with empty cells for any time points without data. This same script fixed the probe data to the same 5-minute resolution of the video data, the coarser of the two data resolutions. This was done by dividing the predefined time range into 5-minute windows and averaging all probe data points that fell inside each window (see Exhibit C5-13). As discussed in Chapter 1: Data Management GDOT's Navigator system also aggregates Traficon data into 15-minute periods.

Each 5-minute Traficon video speed measurement is also accompanied by a value representing the degree to which that data point represents an actual roadway measurement, called "percent observed". Certain time periods might have a low percent observed due to errors in the detector or feed. In those cases, PeMS fills in the missing data according to certain estimation algorithms. To keep the comparison focused solely on the data generated by the sensors, only 100% observed data points were included. After this filtering, between 40% and 50% of 5-minute periods contained data for most Traficon video detectors. By comparison, the Navteq probe data sets all contained data for 20% of all 5-minute periods, and all TMC paths followed the same pattern of data availability. This indicates the few probe data outages were caused by system issues.

At this point, the video and probe data is all in the same temporal frame of reference. The comparison begins by identifying the pairs of video detectors and TMC paths that apply to the same stretch of roadway. Since video detectors are fixed to a point and TMC paths span a length of roadway, each video detector can have no more than one associated TMC path while each TMC path can have many matching video detectors (see Exhibit C5-11). There were 7 pairs of video detectors and TMC paths in the northbound direction and 12 in the southbound direction.

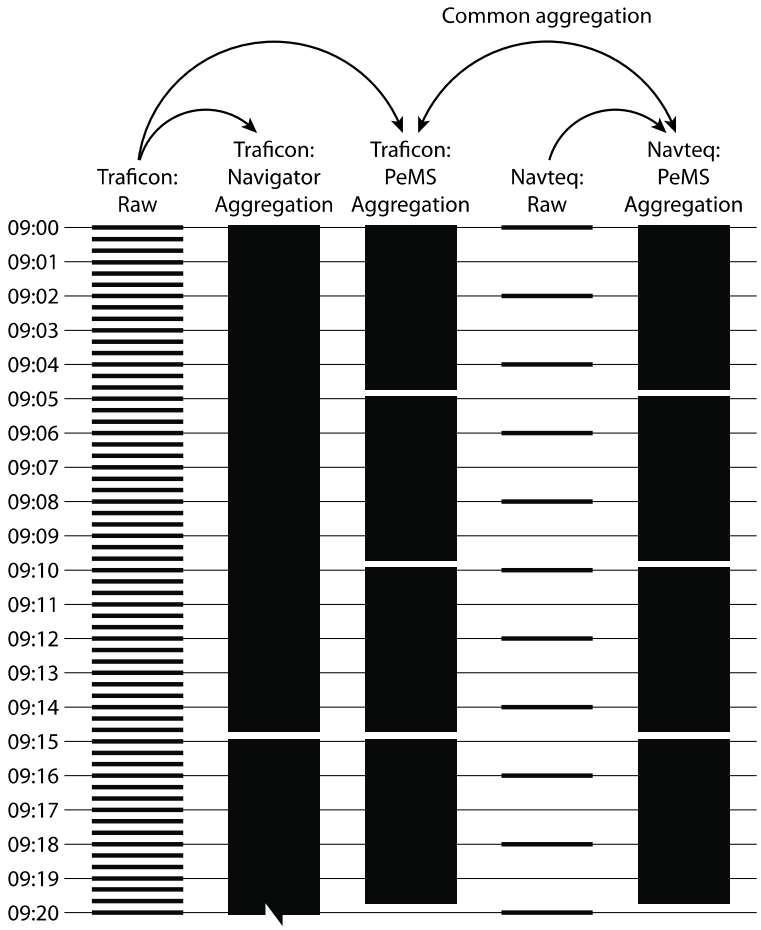

Exhibit C5-13: Common temporal aggregation of comparison data

With video and probe detectors paired by location, their speed measurements can be plotted and compared visually. Exhibit C5-14 shows video detector and probe speeds at the same location on I-285 in the Northbound direction over three consecutive weekdays. Both data sets seem to agree closely on the speed profile during the congested period. However, the Navteq probe data is clearly capped at an artificial ceiling around 55 mph. This means that the probe data is only valid for times when speeds were below 55 mph.

To maintain the integrity of the comparison, all 5-minute periods during which any TMC path had a reported speed of 55 mph were identified as artificial and discarded. Critically, the corresponding time period in the paired video detector was also discarded in order to maintain the same temporal reference in both data sets. Exhibit C5-15 plots the results of this filtering on the time range and data from Exhibit C5-14, showing all of the time points from Exhibit C5-14 during which both data sets contained directly observed data. The removal of data from certain time periods creates discontinuities in the time basis of the data, so each point is now identified by its index in the data set. This procedure effectively removes all non-congested time periods from each comparison. This means that the fundamental basis of comparison of these data sets is the observed speeds during congested periods.

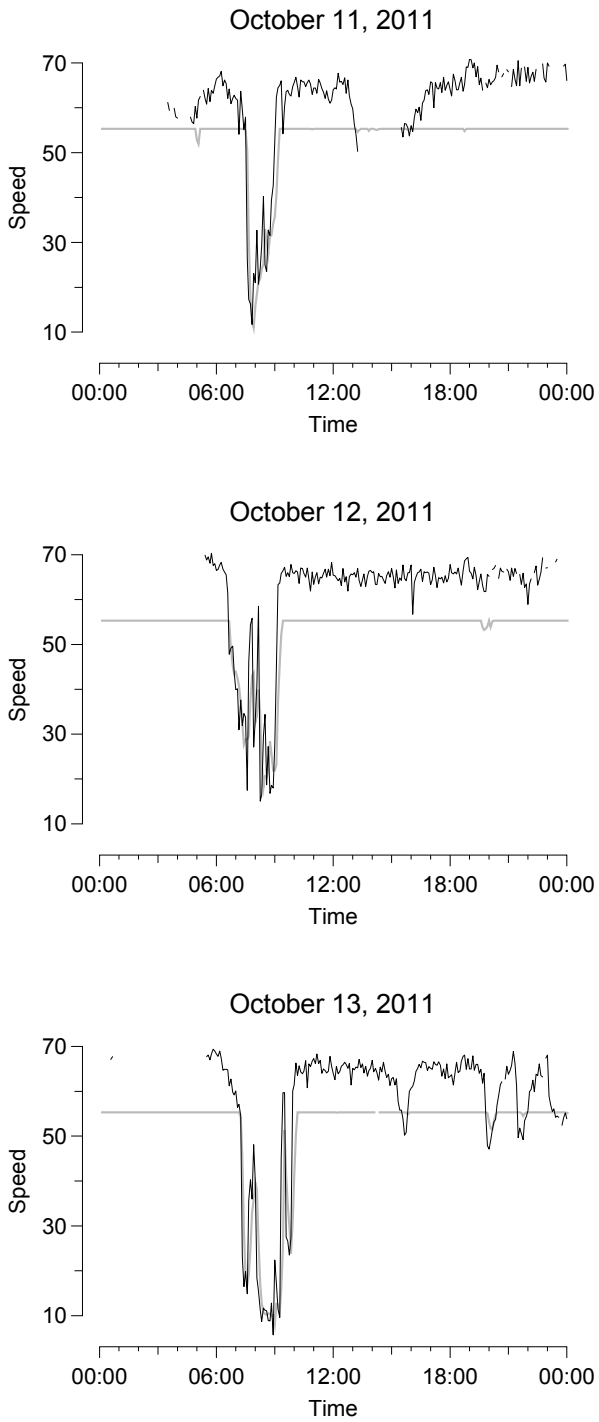

Exhibit C5-14: Comparison of speeds from video (black) and probe (gray) sources

Many techniques are available for numerically computing the similarity of two data sets. In this case, the Pearson correlation coefficient was computed between each pair of processed data sets. The correlation coefficient is defined as the covariance of the two data sets (a measure of their linear dependence) normalized by the product of their standard deviations. Covariance is a useful measure of the degree to which two data sets increase and decrease together, but its magnitude is difficult to interpret. Normalizing the covariance by the product of the standard

deviations allows correlations to be compared across pairs of data sets. Correlation coefficients were computed between each pair of processed data sets in R to determine the degree to which the speed measurements from each source agree.

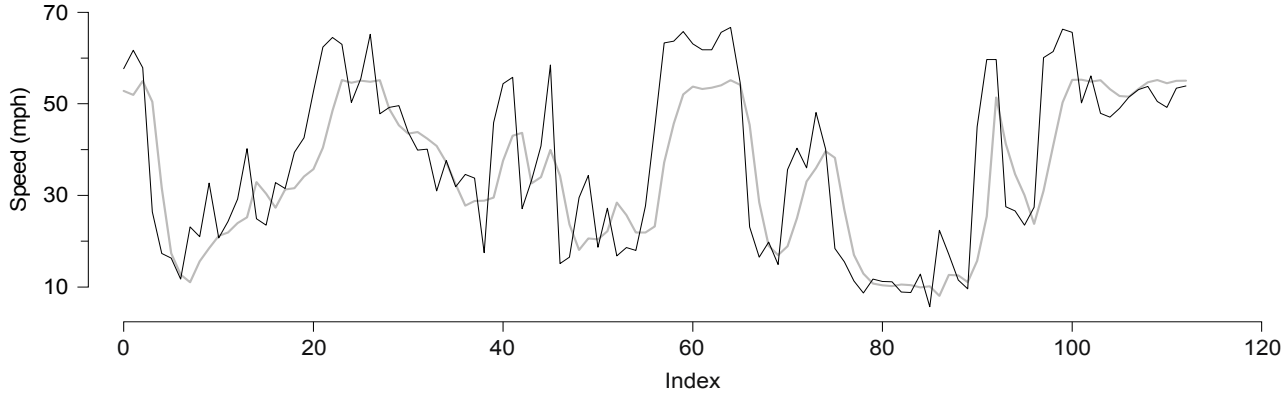

Exhibit C5-15: Comparison of speeds from video (black) and probe (gray) sources

Upon inspection of Exhibit C5-15, the probe data at this location appears to lag slightly behind the video detector data. This lag can be quantified by computing the cross-correlation of the two data sets. To demonstrate this, the cross-correlation for the data shown in Exhibit C5-15 was computed. It can be seen in Exhibit C5-16 that the peak correlation occurs at a lag of -1. The unshifted data, as shown in Exhibit C5-16, has a correlation of 0.80. When the probe data is shifted earlier by one index position, as recommended by the cross-correlation function, the correlation of the two data sets improves to 0.93 (see Exhibit C5-17). This technique can be used to calibrate sensor measurements.

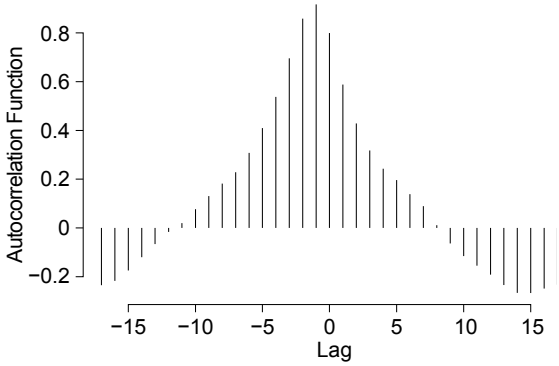

Exhibit C5-16: Cross-correlation of data from Exhibit C5-15

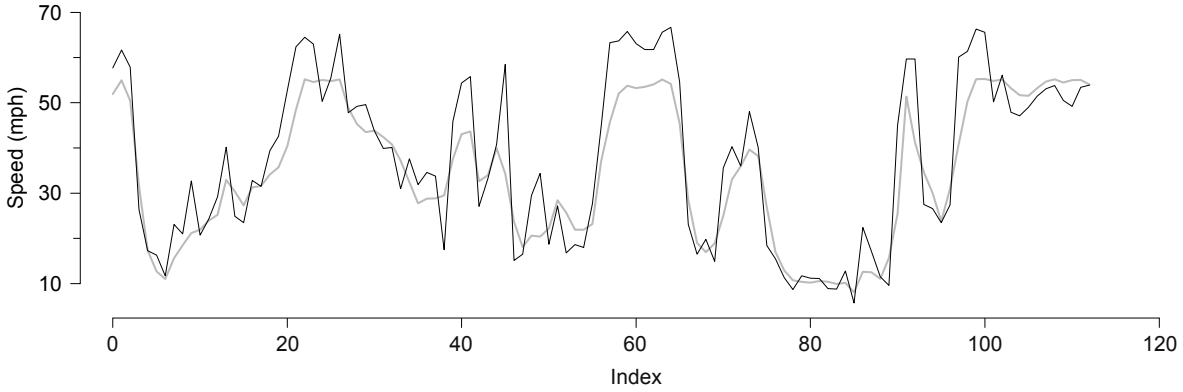

Exhibit C5-17: Data from Exhibit C5-15 after shifting probe data

I-285 Northbound Results

Correlations in speed measurements from the northbound direction of travel were strong, ranging from 0.75 to 0.87. Of the 7 video detector–TMC path pairs, 5 (71%) had correlations exceeding 0.8. The poorest correlated pair was located at the northern end of the study segment, near the North Hills Shopping Center. The best correlation was seen between the longest TMC path and the detector located near its middle, close to the Decatur Road exit.

I-285 Southbound Results

Correlations in speed measurements from the southbound direction of travel were slightly weaker than in the northbound direction, ranging from 0.69 to 0.87. The range of correlations was greater in this direction of travel, perhaps because of the larger number of pairs. Of the 12 video detector–TMC path pairs, only one (8%) had a correlation exceeding 0.8, although 10 (83%) exceeded 0.75, a good correlation. The poorest correlated pair was located on the southern edge of the longest TMC path, near Midvale Road. The best correlation was seen between the TMC path and detector located near the U.S. 78 and I-285 junction.

Discussion

Although the video detector speeds and the probe speeds correlate well with each other, a better understanding of the source of the differences in the measurements was sought. Some part of the difference is likely due to random error, but another part could be related to the locations of the video detectors and TMC paths. Since each detector that sat along any part of a TMC path was paired with that TMC path, one source of difference could be related to the location of the video detector within its paired TMC path. It seems reasonable to assume that a TMC path paired with a video detector located at its midpoint would correlate better than a TMC path paired with a video detector near the TMC path's edge.

To investigate this, the distance between each video detector and the midpoint of its paired TMC path was calculated. These distances ranged from 0.02 to 0.27 miles in the northbound direction and from 0.01 to 0.72 miles in the southbound direction. Scatterplots were made between these distances and the correlation of the corresponding video detector and TMC path for each freeway direction (see Exhibit C5-18). We would expect each pair's correlation to

increase as the distance decreases, and we indeed appear to see this negative relationship in the southbound direction ($R^2 = 0.55$). No linear relationship between correlation and distance is apparent in the northbound direction. When plotting distances and correlations from both directions of traffic together, the same approximate linear relationship that was seen in the southbound direction reemerges, with a slightly lower correlation coefficient ($R^2 = 0.43$). This indicates that part of the difference in the video detector and probe data speed measurements may be due to the distance between the video detector and the midpoint of the TMC path.

Another way to compare two sets of speed measurements would be to simply compute the difference between them at each time point. Exhibit C5-19 shows the difference in speed measurements for the same pair of detectors and time range as in Exhibit C5-14 and Exhibit C5-15. Speed measurements from this pair of detectors matched well, with a correlation coefficient of 0.85. Exhibit C5-15 shows both speed profiles in general agreement. However, when the difference in speed measurements is plotted in Exhibit C5-19, we see that the measurements often differ by as much as 20 mph during individual 5-minute time periods. This indicates that measurements from two types of detectors may not agree at fine time resolutions, even if the detectors are properly configured and in good working order. That the speed difference appears to fluctuate around zero indicates further that this pair is still a good match. Since the detectors agree on the general duration and speed profile of congestion and their difference is centered around zero, their correlation will likely improve as the data is rolled up into coarser levels of temporal aggregation.

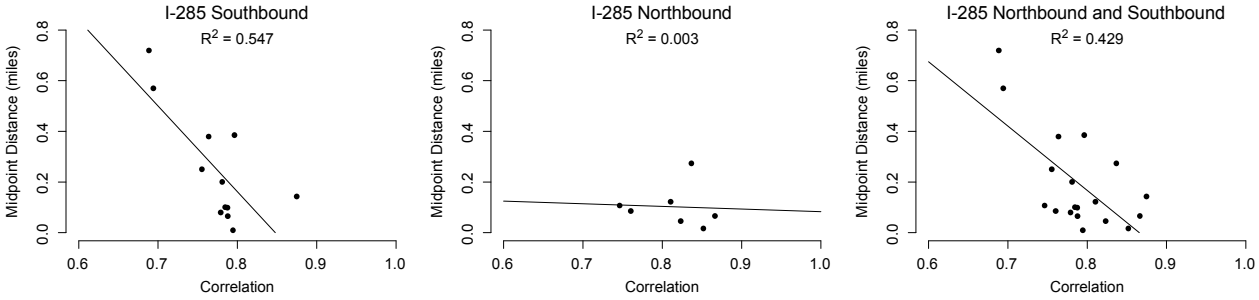

Exhibit C5-18: Scatter plots comparing correlation of speed measurements with distance between detectors

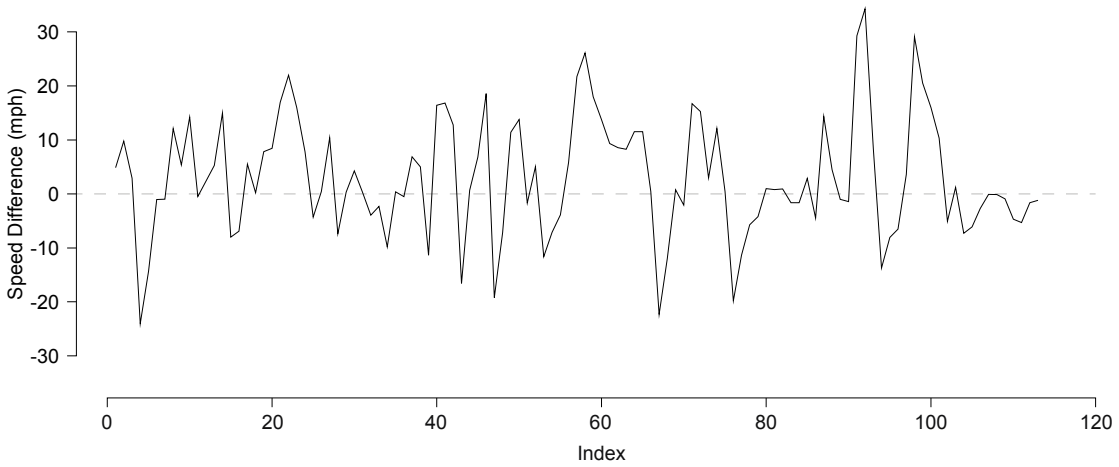

Exhibit C5-19: Difference in Speed Measurements (video – probe)

Conclusion

This use case explored the steps necessary to compare speed measurements from two different types of detectors. Differences in sampling rate (3 Hz vs. 0.5 Hz), configuration basis (detector-based vs. TMC path-based), and data availability range were addressed by aggregating speed measurements at the finest available grain to 5-minute windows. Time points during which a video detector was less than 100% observed, or a TMC path reported the 55 mph speed ceiling were discarded. With this preprocessing carried out, the speed values of detectors from the same roadway segment were compared by computing their correlations. It was seen that the video detector speeds correlate well with probe-based speeds at the same location, particularly in terms of the magnitude of speed drops and their profile. Thus, these disparate detector types can be used together to determine the time, duration, and extent of congestion. Additional analysis revealed that some part of the differences between the two types of measurements may be due to the distance of the video detector from the midpoint of its matched TMC path. Finally, plotting the difference between two data sets reveals the hazards of comparing data from individual 5-minute periods.

LESSONS LEARNED

Overview

This case study showed that, with proper quality control and integration measures, ATMS data can be used for travel time reliability monitoring, including the linking of travel time variability with the sources of non-recurrent congestion. It showed that ATMS systems can be a source of traffic data, as well as a source of information for informing on the relationship between travel time reliability and the seven sources of congestion. In evaluating the similarity between ATMS and third-party probe data, it also sheds light into points of consideration for integrating different data sources into a travel time reliability monitoring systems. The remainder of this chapter describes lessons learned within each of these areas.

Systems Integration

The key systems integration finding from this case study is that ATMS data requires significant evaluation and quality-control processing before it can be used to compute travel times and inform on the causes of unreliability. Four major issues were noted with ATMS data and metadata:

1) Sensor metadata and event data may not contain locational information at the accuracy required for travel time computation and analysis;
2) Descriptive information for sensor metadata and event data can be free-form and non-standardized;
3) Traffic data may not be received at constant sampling rates; and
4) Expected data samples may be missing

Due to the short-term nature of this case study, these issues were handled internally by the research team by changing the properties of the data collection feeds and discarding sensors and events that did not have sufficient information to allow for interpretation. For staff executing a long-term deployment of a reliability monitoring system, these issues highlight the need for a thorough understanding of the ATMS data model and processing steps, as well as a good relationship with ATMS staff so that needed information can be acquired and problems resolved.

Methodological Advancement

The methodology work of this case study linked the regime-estimation work developed in the Northern Virginia case study site with the seven sources analysis developed for the San Diego site. At the San Diego study site, analysis showed incidents and weather events to be leading drivers of travel time variability. On the Atlanta corridor, however, while incidents, weather, lane closures, and special events all contributed to the slowest and most variable travel time regimes, a large portion of travel time variability was not attributable to any of the measured seven sources. This indicates that, particularly for urban corridors that experience a lot of recurrent congestion, the harder-to-measure sources of fluctuations in demand and inadequate base capacity are likely leading drivers of travel time variability.

Probe Data Comparison

This case study provided the first opportunity to compare speed data reported by infrastructure-based sensors with speeds obtained from a third-party data provider. It showed that there are three main points of consideration for integrating different data sources into a reliability monitoring system: (1) standardizing the data sampling rate (in this case study, 3 Hz vs. 0.5 Hz); (2) standardizing the spatial aggregation of the data (in this case study, detector-based vs. TMC path-based); and (3) handling instances of missing or low quality data samples among the sources. These issues must be deal with before disparate data sources can be fused together for reliability monitoring. Following the necessary integration steps and the discarding of any artificial speed bounds in the third-party data set (in this case study, third-party speed were capped at 55 mph) the comparison analysis showed that the agency-owned video detection speeds correlated well with the corresponding probe-based speeds. However, results showed that speed differences between data sources may increase with the distance between the mid-point of the TMC path and the infrastructure detector.

REFERENCES

1) CMAQ: Advancing Mobility and Air Quality. FHWA Office of Planning, Environment, and Realty. http://www.fhwa.dot.gov/environment/air_quality/cmaq/research/advancing_mobility/ 03cmaq07.cfm.
2) Brunswich MPO 2030 Long Range Transportation Plan. Section 12. Glynn County, Georgia. 2005. http://www.glynncounty.org/index.aspx?NID=1024

CHAPTER C6

NEW YORK/NEW JERSEY

The New York City site was chosen to provide insight into travel time monitoring in a high-density urban location. The 2010 United States census revealed New York City's population to be in excess of 8 million residents, at a density near 28,000 people per square mile. While New York City has a low rate of auto ownership compared to other United States cities, more than half of all commute trips are still made in single-occupancy vehicles. In 2010, these factors contributed to New York City having the longest average commute time of any United States city, at 31.3 minutes.

The main objectives of the New York/New Jersey case study included:

- Obtaining time-of-day travel time distributions for a study route based on probe data;
- Identifying the cause of bi-modal travel time distributions on certain links; and
- Exploring the causal factors for travel times that vary significantly from the mean conditions.

The route analyzed in this case study begins in the Boerum Hill neighborhood of Brooklyn and ends at JFK Airport, traversing three major freeways: the Brooklyn-Queens Expressway (I-278), the Queens-Midtown Expressway (I-495), and the Van Wyck Expressway (I-678) and is illustrated later in this chapter.

The *monitoring system* section details the reasons for selecting New York as a case study site and gives an overview of the setting. It briefly summarizes the archived probe vehicle data source and the underlying road network to which it corresponds, and gives an overview of approach that the team took to analyze that data.

The section on *methodology* describes the steps necessary to obtain Probability Density Functions (PDFs) of travel time distributions route based entirely on probe data along a New York City. Critically, this probe data is sparse and no probe vehicle runs traverse the entire route. Techniques are presented to preserve the correlation in speed measurements on consecutive links while synthesizing the aggregate route travel time PDF from segments of multiple probe vehicle runs.

The *use case analysis* section is less theoretical, and more site specific. It is motivated by the user scenarios described in Supplement D, which are the results of a series of interviews with transportation agency staff regarding agency practice with travel time reliability. While the methodology section of this case study describes the steps necessary to process and interpret probe vehicle data, the use case section focus on a specific application of this methodology. This case study contains a single use case that focuses on three alternative methodologies for constructing travel time probability density functions (PDFs) from probe data.

The *lessons learned* section summarizes the lessons learned during this case study, with regards to all aspects of travel time reliability monitoring: sensor systems, software systems, calculation methodology, and use. These lessons will be integrated into the final guidebook for practitioners.

MONITORING SYSTEM

Site Overview

The New York City site was chosen to provide insight into travel time monitoring in a high-density urban location. The 2000 US census estimated New York City's population to be in excess of 8 million residents, at a density near 26,500 people per square mile (1). Kings County (Brooklyn) is the second-most densely populated county in the United States after New York County (Manhattan) (1). New York City has a low rate of auto ownership; only 55% of households had access to an automobile in 2010 (2). For drivers of single occupancy vehicles, 53% of all commute trips take 30 minutes or more, with an average commute travel time of 31 minutes (2).

This site is covered by a probe vehicle data set, provided to the research team by ALK Technologies, Inc. This data is collected from mobile devices inside of vehicles, and consists of two types of data: (1) individual vehicle trajectories defined by timestamps and locations, and (2) link-based speeds calculated from each vehicle's trajectory. Probe vehicle detection technology provides high-density information about the vehicle's entire path, allowing travel times to be directly monitored at the individual vehicle level. In contrast, infrastructure-based sensors such as loop detectors measure traffic only at discrete points along the roadway, and do not keep track of individual vehicles as they travel. Probe data relies on the roadway's users to generate performance data, greatly reducing detection maintenance costs to the agency. These features make probe vehicles an attractive roadway data source to agencies.

The research team obtained this probe data for a region of New York City defined by a rectangular bounding box, 25 miles long east-to-west and 40 miles long north-to-south. This bounding box covers Manhattan, The Bronx, and Brooklyn in their entirety, along with most of Queens (see Exhibit C6-1). Data from all roadway segments within this bounding box was obtained. Probe runs that crossed the boundary of the bounding box were truncated such that only the segments within the bounding box were included in the data set. Segments that had been truncated in this way were treated as unique trips.

The data obtained for this site was a static collection of raw traces and processed speed measurements, as collected by probe vehicles between May 19, 2000 and December 29, 2011. No real-time data was acquired or analyzed for this case study because it was not available. Unlike in other case study sites, an Archived Data User Service (ADUS) was not deployed. All data processing and visualization were carried out through custom routines, run offline.

As with roadway data from other case studies, this probe data set was accompanied by a network configuration. The network configuration connects the traffic data to the physical roadway network through a referencing system. This configuration data is necessary for proper interpretation and analysis of the traffic data, such as computing route travel times from point speeds. For this probe data, the network configuration is made up of links defined by ALK. Links are unique to a roadway segment and direction and are less than 0.1 miles long on average. Due to limitations in GPS location accuracy, these links are not lane-specific; link data is interpreted as the mean speed across all lanes. The full data set obtained for this case study contains 180,061 links representing 14,402 roadway miles over the 1000 mi^2 area enclosed by the bounding box (see Exhibit C6-1).

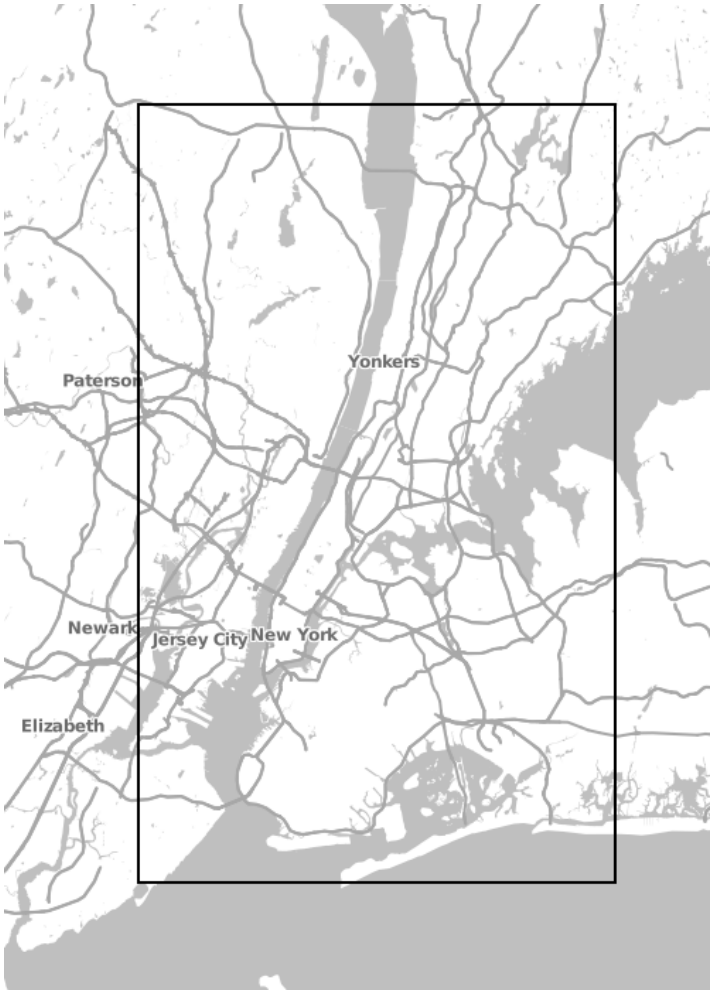

Exhibit C6-1: Site map with data bounding box

Data

Three probe vehicle based data sets contribute to this case study. Each of these three data sets is based on the same original collection of probe vehicle runs, collected in the *Raw* dataset. The raw data contains unaltered GPS sentences, as originally recorded by the probe vehicles, and was not obtained by the research team. The second data set, called *Gridded GPS Track Data (GGD)*, contains most or all raw GPS points, matched to ALK's link-based network configuration. The third data set, called *One Monument*, is an aggregation of the GGD data set. The One Monument data contains a more manageable number of speed measurements that correspond with the vehicle's speed and timestamp at the midpoint of each ALK link.

Table C6-1: Probe vehicle data sets

Data set	Description	Number of data points	Uncompressed size
Raw	Untouched NMEA sentences	36,683,340 (or more)	4.19 GB (or more)
GGD	Data points reformatted and identified by ALK link	36,683,340	4.19 GB
One Monument	One vehicle measurement per link midpoint	4,282,136	0.48 GB

The Raw GPS data set is stored in the standard NMEA sentences originally recorded by the GPS device in the probe vehicle. A different file is typically created for each vehicle trip. The primary GPS data elements of interest for traffic analysis are location (latitude and longitude), speed, heading, and timestamp. GPS sampling frequency affects the temporal resolution of all three probe data sets. The data analyzed by the research team was based on GPS data recorded every three seconds.

The GGD data set is produced through the cleaning and map-matching routines carried out on the raw GPS data; it contains speeds on links and travel times between links, which are organized into trips. This data set is contained in a single file with entries that include timestamp, link ID, position along the link, speed, trip ID, and sequence within the trip. The organization into trips follows that of the GPS files. A gap greater than 4 minutes in a single GPS file is interpreted as the boundary between two trips made by the same vehicle. This preserves continuity in the data and ensures that only travel times (and not trip times) are represented. In this data set, each point is also map-matched to a single ALK link, and includes a value indicating how far along that link the point is located.

The One Monument data set aggregates each trip's data points into single time-stamped speed values for each ALK link that the trip traverses. This is a subset of the GGD data set. When there are multiple observations for the same link within a single trip, only the data point closest to the midpoint of the link is retained and its timestamp is interpolated to the time the vehicle likely passed the link's center point. The speed values in this data set are computed based on the total travel time along the link and the link's length, which effectively evens out the instantaneous speeds over the link. This data set aids travel time analysis by greatly reducing the number of data points required to compute travel times over road segments for a single trip.

The ALK links themselves are defined in three configuration files, referred to as *Links*, *Nodes*, and *Shapes*. Each link lies within a cell of a rectangular grid, and is uniquely identified by the combination of its Grid ID and Link ID. Links are bounded on either end by nodes whose coordinates are defined in the *Nodes* file. The geometry of each link can be drawn from coordinates found in the *Shapes* file. Additionally, links are labeled with a class identifier, which corresponds to one of the following road types: interstate, interstate without ramps, divided road, primary road, ferry, secondary road, ramps, and local road. Local roads make up the vast majority of the links in the network configuration.

Data Management

Analysis of this probe data set was primarily carried out on the aggregated trip-link speeds present in the One Monument data. The aggregated speeds in this data set are similar in

format to the TMC path-based data analyzed in the Atlanta case study. The Atlanta case study compared GPS trace data with video detector data, but only after it had been aggregated into link-based speed measurements. The complete GPS trace data in the GGD data set is the only data from any of the five case studies that traces the entire path of vehicle trips. Even though it is not analyzed directly in this case study, it deepens the analysis done on the *One Monument* data to enable sophisticated computations, as described in the use case section of this document

The data was provided by ALK in flat files and managed by the project team manually through custom processing routines run offline. To focus on issues related to probe vehicle data processing, no additional data sources were considered in this case study.

METHODOLOGY

Overview

The central goal of this use case is to advance the understanding of practical techniques for working with probe vehicle data in travel time reliability monitoring applications. To accomplish this, the research team analyzed a collection of probe vehicle data. This section first describes the study route, illustrating how the probe data set was assembled and processed for the route and explaining the implications of data density on the resulting analysis. The section then describes methods for identifying and visualizing congestion and travel time reliability from sparse probe data. Finally, it lays the groundwork for computing route-level travel time probability density functions, a methodological issue that is explored in depth in the Use Case chapter.

Site Description

The methodological steps in this section are conducted on a 17.4 mile route in New York City that travels from the densely residential Boerum Hill neighborhood of Brooklyn to JFK International Airport. This route was chosen because it lies within a well-connected roadway network, over which several alternate routes can be taken. This makes for a more interesting analysis, as drivers in the area likely base some of their travel decisions on the travel time and travel time reliability of this particular route. The route is also varied, traversing a series of arterials and three major freeways between Boerum Hill and JFK International Airport. The route begins at Atlantic Ave and Flatbush Ave., then travels over the Brooklyn-Queens Expressway (I-278 E), the Queens-Midtown Expressway (I-495 E), and the Van Wyck Expressway (I-678 S), ending near JFK International Airport's cell phone parking lot. This route is shown in Exhibit C6-2, with the origin identified in white and the destination in black.

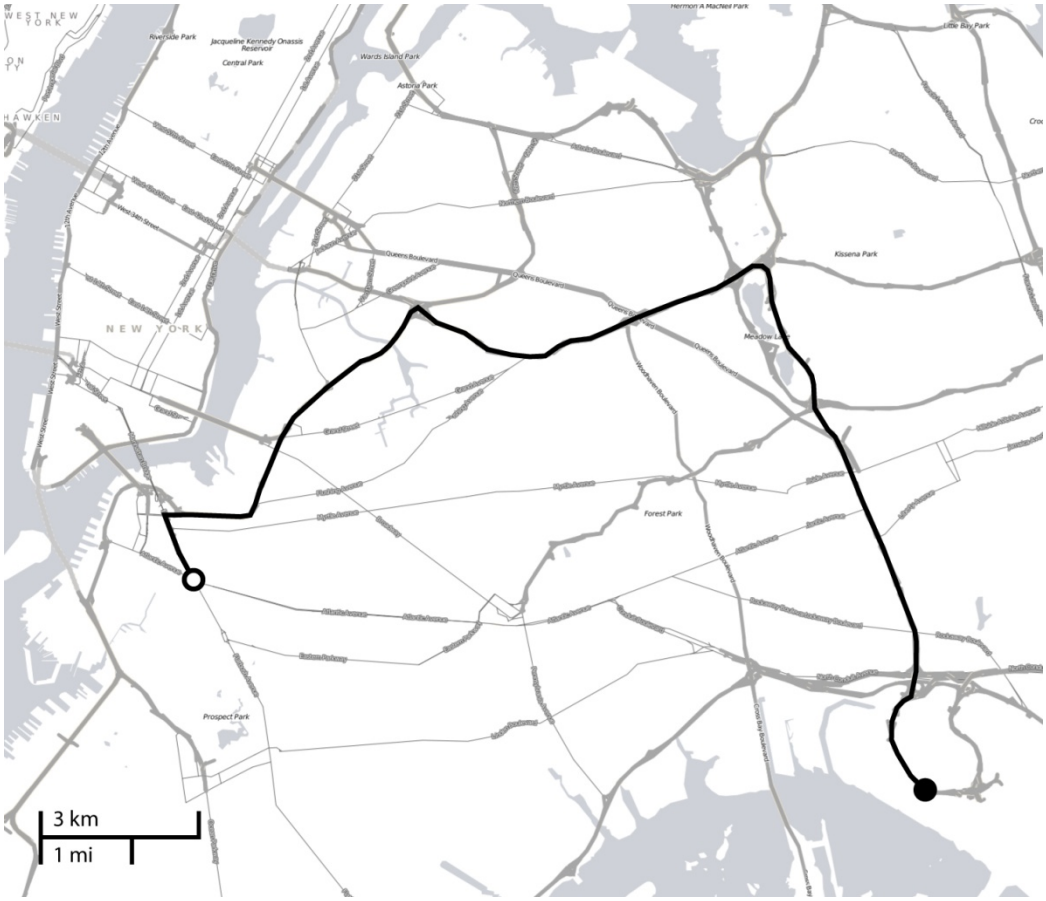

Exhibit C6-2: Study route

The first step was to determine which ALK links make up the study route. This was done by visually identifying the ALK grids that the route travels through. From there, it is possible to map all interstate-class links contained in the relevant grids and visually identify the links which make up the route. Upon the completion of this process, we found the 17.4-mile long route to be made up of 102 ALK links. The Grid IDs and Link IDs of these links are labeled with their order within the route and stored.

After the route links have been identified, it is possible to calculate the number of data points recorded for each link. Probe data is sparser during times when fewer vehicles are traveling (i.e., at night), making certain types of time-of-day analysis more difficult. Since each data point contains a timestamp, counts of data points by link and time of day can be obtained directly from the data. The timestamps must be converted from UTC time to local time (EST) (with adjustments made for daylight savings time) before the counts can be interpreted. Data availability on this route during the 11-year period of coverage is displayed in Exhibit C6-3.

As shown in Exhibit C6-3, data coverage over the route is generally quite sparse, with the most densely covered link-hour containing 71 points. As such, analysis requiring data partitioning, such as comparing weekday and weekend speeds, will likely not yield rich results. The three freeway segments have the best data coverage, while coverage is sparser on the arterials near the origin, the freeway connectors, and the airport roads at the destination. Data coverage is highest in the evenings and around midday. Due to the sparseness of the data, no individual vehicle trips traversed the entire route from beginning to end.

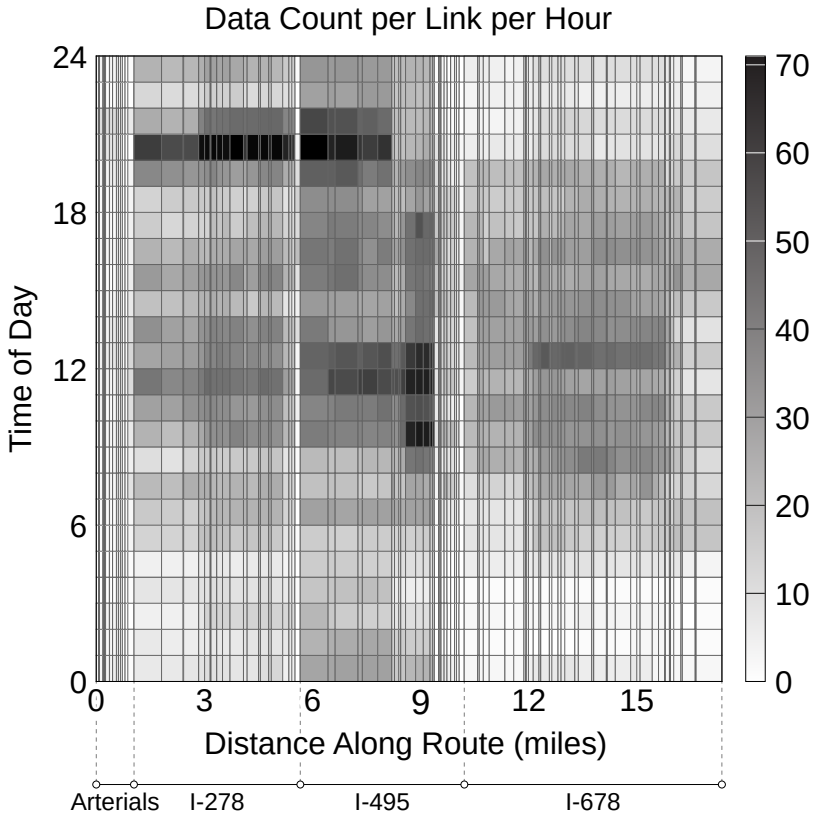

Exhibit C6-3: Quantity of data analyzed

Methods

Since there were no travel time records for the entire route, methodologies had to be developed to construct the route travel time distribution piecemeal from the individual link data. The advantage of this approach is that it utilizes the entirety of the dataset, rather than a subset of long trips. Obtaining composite travel time distributions from vehicles that only traveled on a portion of the route is a complex process, primarily because, as this project has shown, travel times on consecutive links often have a strong linear dependence. This linear dependence must be accounted for when combining individual link travel times into an overall route travel time distribution. This is the core methodological challenge of this case study, fully explored in the Use Case chapter. The research team first approached this complex topic by examining

probability density functions of speeds on an individual link, the results of which are presented in this section.

To understand the traffic conditions represented in the data set, we can plot time-of-day based speed distributions on a single link. Exhibit C6-4 depicts hourly probability density functions of speeds observed on the 38th link in the route (near the I-278 / I-495 interchange). From this visual, it is clear that most speeds fall between 45 and 65 mph, with the exception of the PM peak. From 2:00 p.m. to 7:00 p.m., the speeds appear to be bimodally distributed, with a lower modal speed around 10 mph.

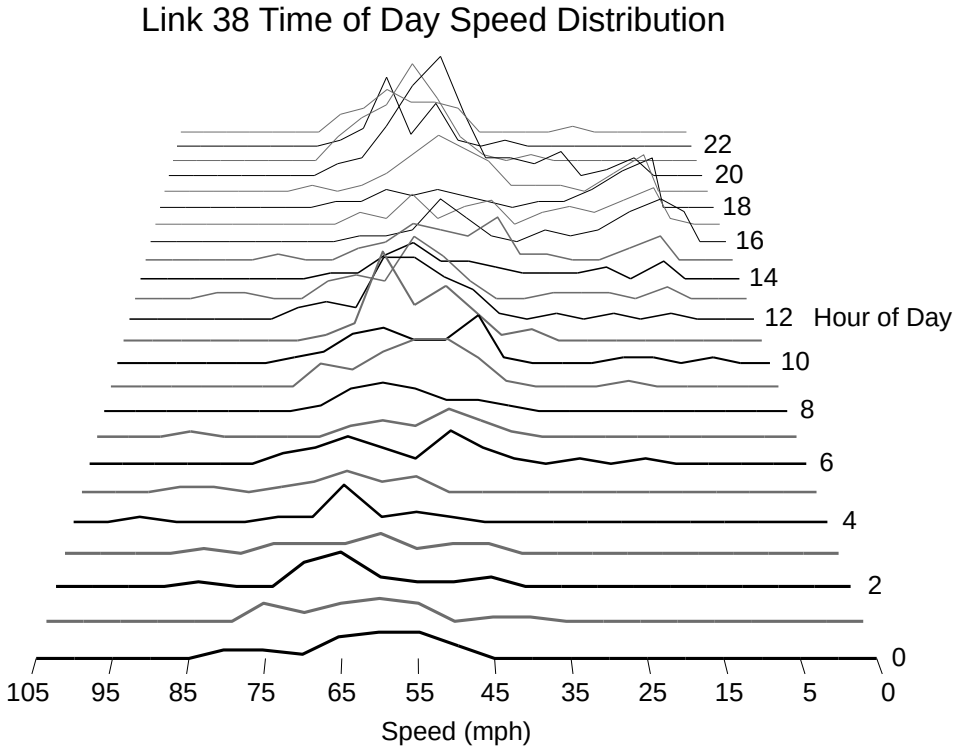

Exhibit C6-4: Time of day speed distribution on a link

With the knowledge that mixed traffic conditions occur during the PM period on the 38th link in the route, we can analyze PM speeds along the entire route. Speed measurements on each link during the 3pm to 8pm commute period were obtained from the *One Monument* data set. To illustrate speed changes along the route in the PM period, the median PM speed for each link is plotted (see Exhibit C6-5). Each link has multiple speed measurements over the 11-year study period during these hours, so speeds between the 25th and 75th percentile for each link are shaded in gray to indicate the rough extent of each link's PM speed distribution. Speeds appear to dip in the middle of the freeway segments. Median speeds along the route outside of the PM period remain relatively high throughout the freeway segments, indicating PM period congestion.

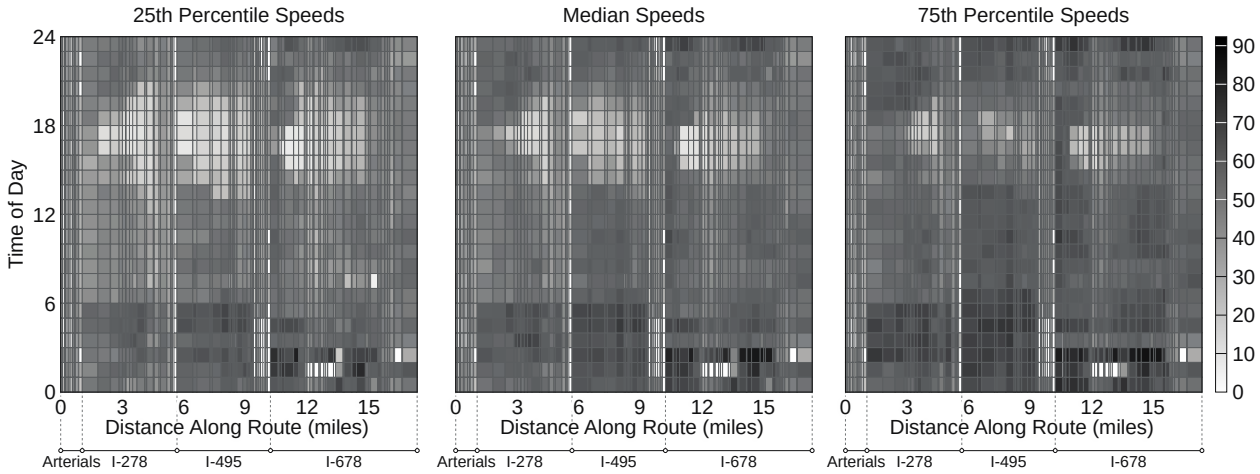

Exhibit C6-5: Quartile speeds along route by time of day

Next we look at how speeds vary across the route throughout the whole day, again considering the entire speed distribution on each link-hour. Speed measurements on each link during each hour of the day are extracted from the *One Monument* data set and the 25th percentile, median, and 75th percentile speeds for each link-hour are computed. The variation of speeds along the route throughout the day is presented in Exhibit C6-6. Link-hours with no data (mostly at freeway interchanges and toward the end of the route at night) were marked with a speed of zero.

The speed data appears to show three triangular regions in the PM period of each freeway segment. These triangular regions indicate bottleneck regions of low speeds during the PM commute period.

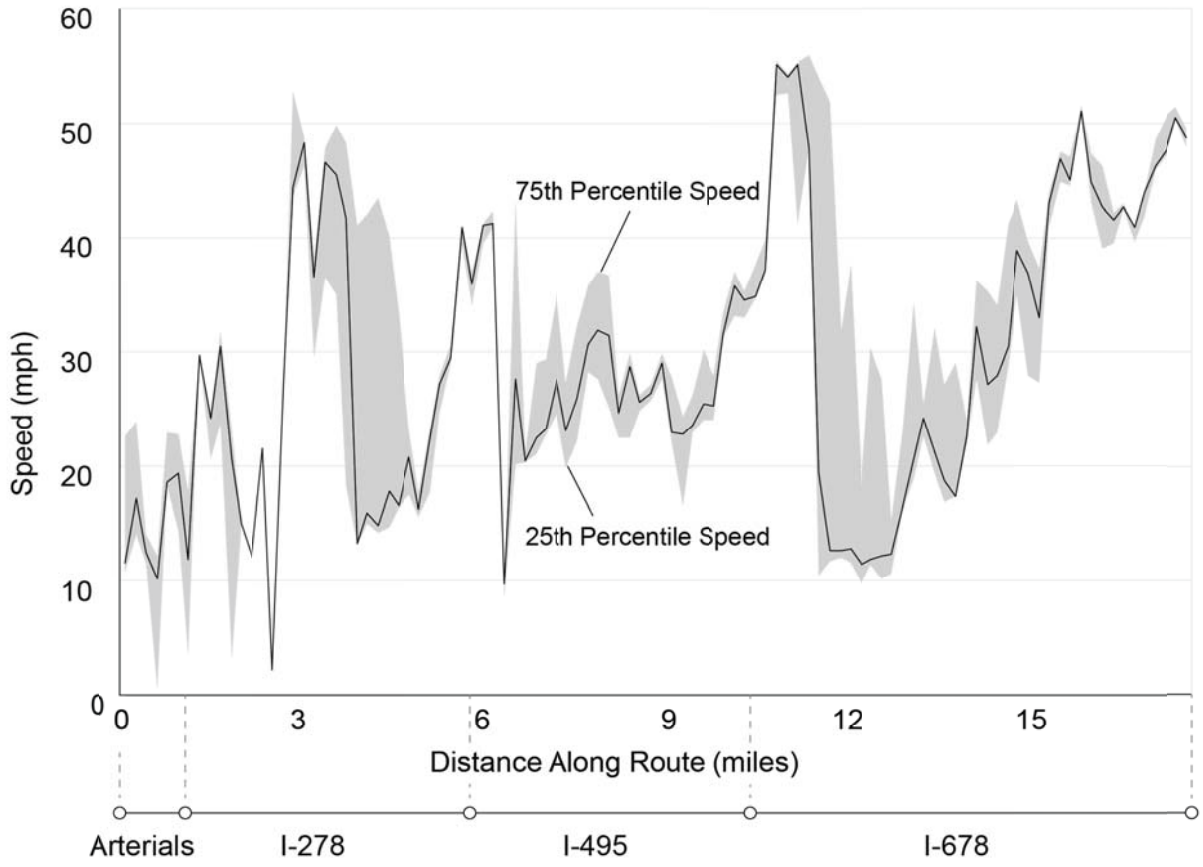

Exhibit C6-6: Route speed profile

Results

Using the quartile speeds for each link throughout the day, it is possible to simulate trip trajectories along the route for any slice of the speed distribution. We do this by first choosing a virtual trip start time, and then moving along the route link-by-link, simulating the arrival time at the next link based on the speed and length of the current link. The link speeds used to advance this simulation must correspond to the time of day in the virtual vehicle's trip.

Exhibit C6-7 shows the trajectory of trips simulated using PM period link mean speeds at 30-minute intervals. This type of time-space contour plot is practical in helping to identify locations or times that experience long travel times and to view how unreliable conditions affect trips at different times of day. For example, the virtual trip departing at 5pm appears to experience more congestion at the beginning of the I-678 segment than later trips do. This gives it a longer travel time than it would have experienced had it departed 30 minutes later.

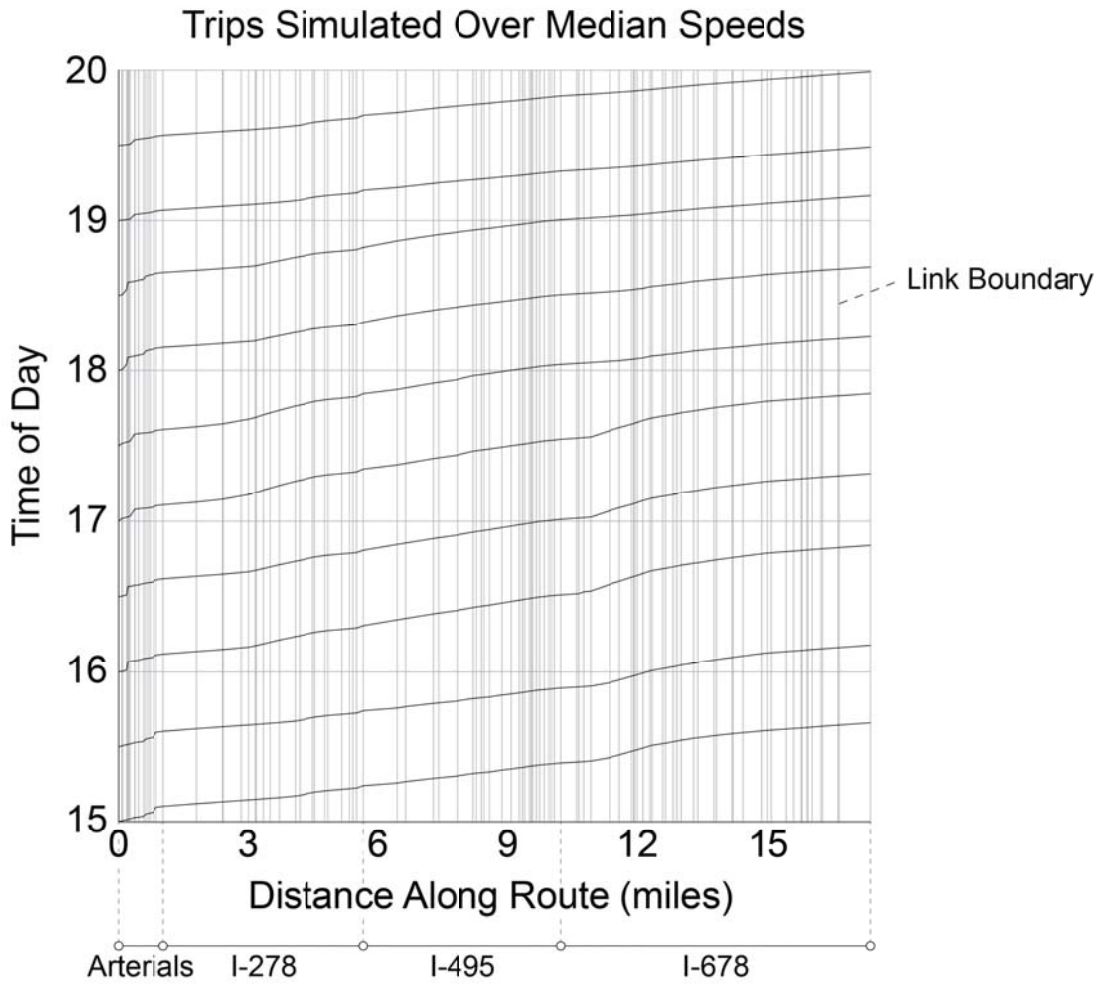

Exhibit C6-7: Virtual trips simulated over median link speeds

USE CASE ANALYSIS

A single use case was evaluated in this case study. This use case is a site-specific application of the probe data processing and analysis techniques described in the Methodology chapter. The motivation for this use case is to generate and compare travel time distributions along a route at different times of day, using only probe data. The methodology chapter of this document describes a technique for simulating trips based on probe speed measurements; however, these simulated trips only apply to a particular slice of the speed distribution (such as the median speed). A more complex approach is needed to measure and illustrate the variation in speeds and travel times on a route at a given time.

This use case demonstrates three methods for obtaining route travel time distributions from probe-based speed data. For continuity, analysis is performed on the route describe in the Methodology chapter. The analysis in each of the three methods is performed on the One Monument data set. For this analysis, the most important variables in the data set are timestamp, speed, trip ID, and indexed position within a trip, if any (many trips are made up of a single point on the route). This yields two types of information useful for travel time analysis: (1) individual

vehicle time-stamped link speeds, and (2) individual vehicle link travel times, as derived from the differences in the timestamps of consecutive trip points (for trips with more than one point). The methods differ in how they use these features of the data set to construct the travel time PDFs.

Method 1

The first method is the only method to use all available data elements in the One Monument probe data set to construct the route travel time PDF. It uses discrete link speeds as well as trip-based travel times to construct the travel time distribution in different time periods of the day.

Since the data coverage on the arterial links at the beginning of the route is so sparse, analysis is focused on the route beginning with link #17 (and continuing to JFK International Airport). The method is divided into two stages: a preparatory stage and a distribution construction stage.

Preparatory Stage

In the preparatory stage, we consider each link in the route, and identify trips that began on that link and traveled at least one link downstream on the route. The goal of this step is to calculate a link-startpoint to link-endpoint travel time for each multi-link trip in the dataset. Each One Monument data point contains a *LinkOffset* value that indicates the distance along the link that the speed value was taken (for, example 0.5 indicates that the data point was taken at the link's midpoint). This trip travel time calculation method uses the data point timestamps to determine the travel time in between each trip's first and last link, and the link speed, length, and offset to extend that travel time to the start point of the first link and the end point of the last link. For a trip that travels from link 1 to link n, the trip travel time equation is:

$$TripTT = \frac{Length_1 * LinkOffset_1}{Speed_1} + \left(Timestamp_n - Timestamp_1\right) + \frac{Length_n * (1 - LinkOffset_n)}{Speed_n}$$

This step results in a set of travel times for each link that measure trips from that link to some downstream link. The travel times were divided up by time period (AM, midday, PM, and nighttime), and were then assembled into trip travel time distributions for each link and time period.

Distribution Construction Stage

The distribution construction stage builds up the full travel time distribution along the route link by link in four steps. Each iteration of the steps adds the subsequent downstream link into the route travel time distribution. The route travel time distribution is initialized as the travel time distribution on the first link on the route, as computed from all data points on the first link. The following four steps are then carried out sequentially down the route for all links:

1) Compute the travel time distribution for the current link using all data points measured on the link.
2) Add the travel time distribution for the current link to the route travel time distribution computed in Step 4 for the upstream link, assuming independence. To

add two independent distributions of data, each point of the first data set must be summed with each point of the second data set. If the size of one dataset is *m* and the size of the other is *n*, the size of the dataset resulting from their sum is the product of the two sizes: *mn*. This is equivalent to convolving the probability density functions of the two independent distributions.

3) Obtain the set of travel times computed in the preparatory stage that end at the current link and merge their adjusted datasets to the dataset of the route travel time distribution computed in Step 2. The adjusted dataset will have been computed in Step 4 for a previous link.
4) For all trips that start at the downstream link, add the route travel time distribution computed in Step 3 to their travel time. This adjusts these travel times such that they represent the travel time distribution between the beginning of the route and the end of the trip.

The resulting travel time probability density functions computed using this method are shown for four time periods in Exhibit C6-8. The odd multimodal distribution of the 10pm to 12am travel times is due to a proportionally larger number of trip-based speeds than discrete link speeds at night. At other times of the day, the number of link speeds overwhelms the number of trip-based speeds, smoothing out the effects of individual trips.

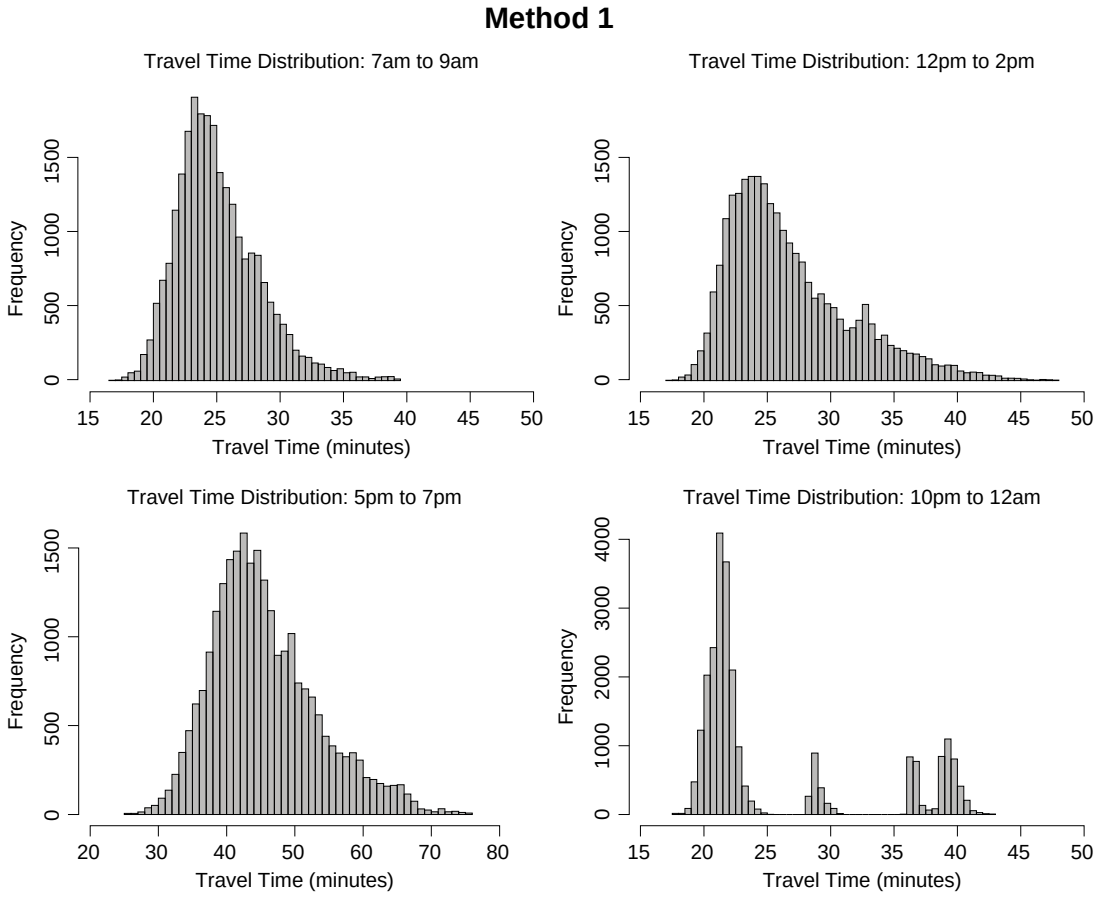

Exhibit C6-8: Route PDF generation method 1

Method 2

The second method for computing route travel time PDFs ignores the linear dependence between consecutive links and directly computes the route travel time distribution as if all link travel times were independent. This method is based entirely on directly-observed link speeds, discarding the timestamp differences between points in the same trip. It works by simply convolving the distributions of travel times on consecutive links down the route. For example, the frequency distribution of travel times on the first link is added to the frequency distribution of travel times on the second link, and so on until a full travel time distribution for the entire route is obtained. The resulting travel time probability density functions computed using this method are shown for four time periods in Exhibit C6-9.

This is the simplest route travel time PDF creation method considered in this case study. Here we treat every single measurement as independent of all others, ignoring all trip relationships between points. As in method 1, we compute travel time distributions for four time periods during the day. With the trip-based travel times discarded, the outlying spikes in the 10pm to 12m travel time distribution are no longer seen. The speeds between 5pm and 7pm appear to be shifted by roughly the same amount as seen in method 1. The 7am to 9am and 12n to 2pm time periods appear to have very similar bimodality to that generated by method 1.

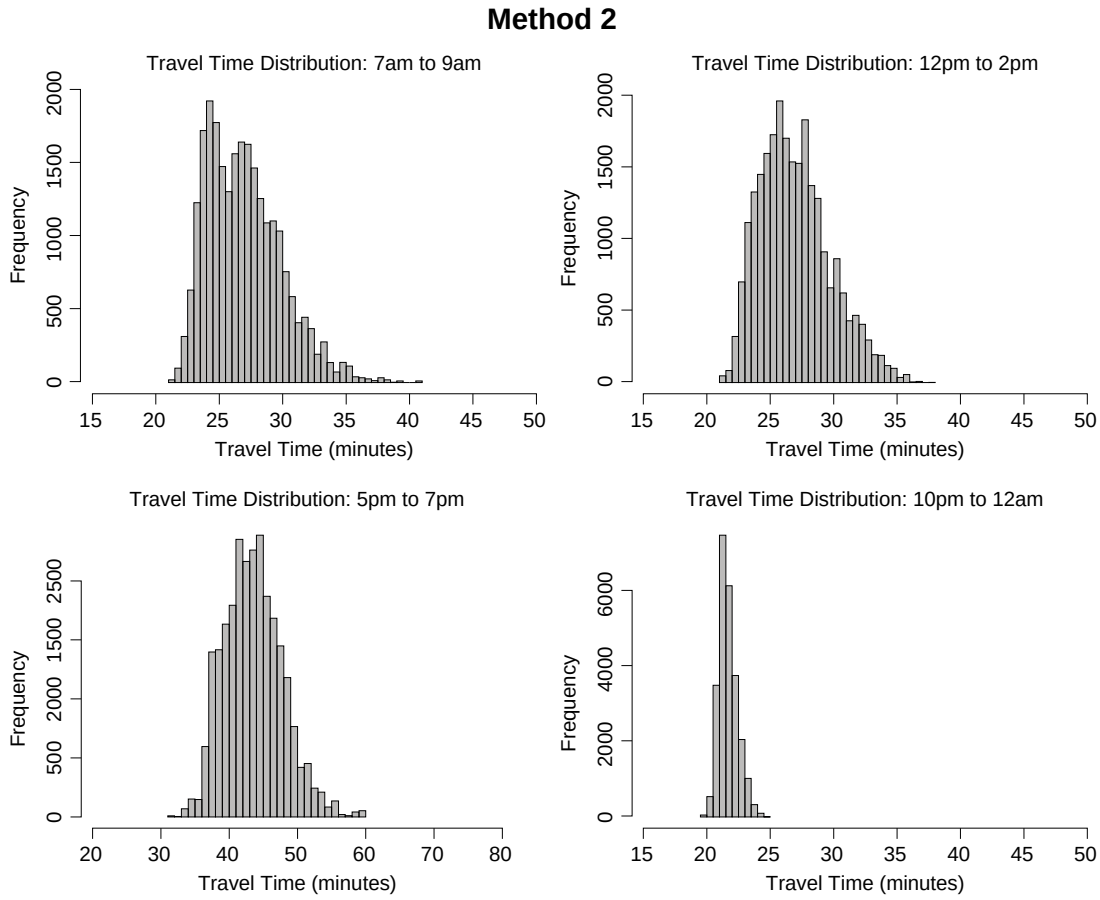

Exhibit C6-9: Route PDF generation method 2

Method 3

The third and final method developed for constructing route travel time PDFs computes and leverages the correlation between speeds on consecutive links within a trip. This method, which only requires speeds measured from trips that traveled on multiple links, uses the fewest One Monument data elements. It builds route travel time PDFs by simulating trips along a route, taking into account the measured data on each link as well as synthesized trips based on observed data and computed incident matrices. It builds up travel times link by link. As with the previous two methods, due to the lack of data on the arterials near the beginning of the route, we begin the route on link #17.

The method begins by computing incidence matrices for each pair of consecutive links. These incidence matrices describe the correlation in speeds between the two links. To construct the incidence matrices, evenly spaced bins are defined to group the speed data for each link. In this use case, 10 bins are used between 0 mph and 80 mph (each bin is 8 mph wide). A 2-D incidence matrix is created for each pair of consecutive links to capture the nature of the speed relationship between the two links within different bins. Speed bins on link #1 are represented in the incidence matrix's rows, and speed bins on link #2 are represented in its columns. Because 10 bins were used in this use case, all incidence matrices are 10 x 10.

Consider an incidence matrix for two consecutive links: link #1 and link #2. The incidence matrix describes the likelihood of a speed on link #2 occurring given a speed on link 1. The entry in the (*m*, *n*) cell of this incidence matrix contains the quantity of link #2 speed measurements that fell into the *n*th bin when the link 1 speed came from the *m*th bin. The counts in the cells of the incidence matrix become synthesized trip points for each observed data point on link #1.

For example, suppose a single link #17 speed observation falls within the 4^{th} speed bin, and the incidence matrix for links #17 and #18 lists two speeds in the 5^{th} bin and three speeds in the 4^{th} bin on link #18 following a 4^{th} bin speed on link #17. This single observed speed on link #17 has resulted in five pairs of speeds across links #17 and #18 (two between it and the 5^{th} speed bin, and three between it and the 4^{th} speed bin). These five speed pairs can be thought of as synthesized trips between the two links since they capture the correlation between speeds on the two consecutive links while using the observed data. This process is repeated for each observed speed on link #17 and all synthesized trips over the first two links are recorded.

To continue the process on the next pair of links, #18 and #19, the speeds on link #18 resulting from the incidence matrix technique described above (there were 5 such speeds in the example) are combined with the directly observed speeds on link #18. This collection of speeds is then subjected to the same incidence matrix procedure to obtain synthesized link #19 speeds for each speed on link #18 that was either directly observed or synthesized from link #17's directly observed speeds. When the final link in the route is reached in this way, the speed on each link in each synthesized trip can be used to obtain its travel time. The distributions of these travel times calculated at different times of day are shown in Exhibit C6-10.

Since each preceding link speed generates multiple speeds for the following link, this method generates a large amount of data very quickly. To keep the travel time data set manageable, the growing data set of synthesized speeds was periodically reduced to a random sample whenever it grew too large to efficiently process.

The multimodal pattern seen in the 10pm to 12am data from method 1 is even more pronounced in travel times synthesized with this method. Both of these methods leverage individual trip travel times across multiple links. The low quantity of data at night exaggerates

the influence of individual trips on the data, creating these spikes. This method produces very narrow travel time distributions that are offset slightly from those generated by the other two methods. Here, we see travel times during the AM and midday time periods are faster by 5 minutes compared to the other methods, with dramatically fewer long travel times. The 5pm to 7pm travel time distribution is again the most widely distributed, but travel times are shifted to the right (slower) by 10 minutes compared with the results from methods 1 and 2.

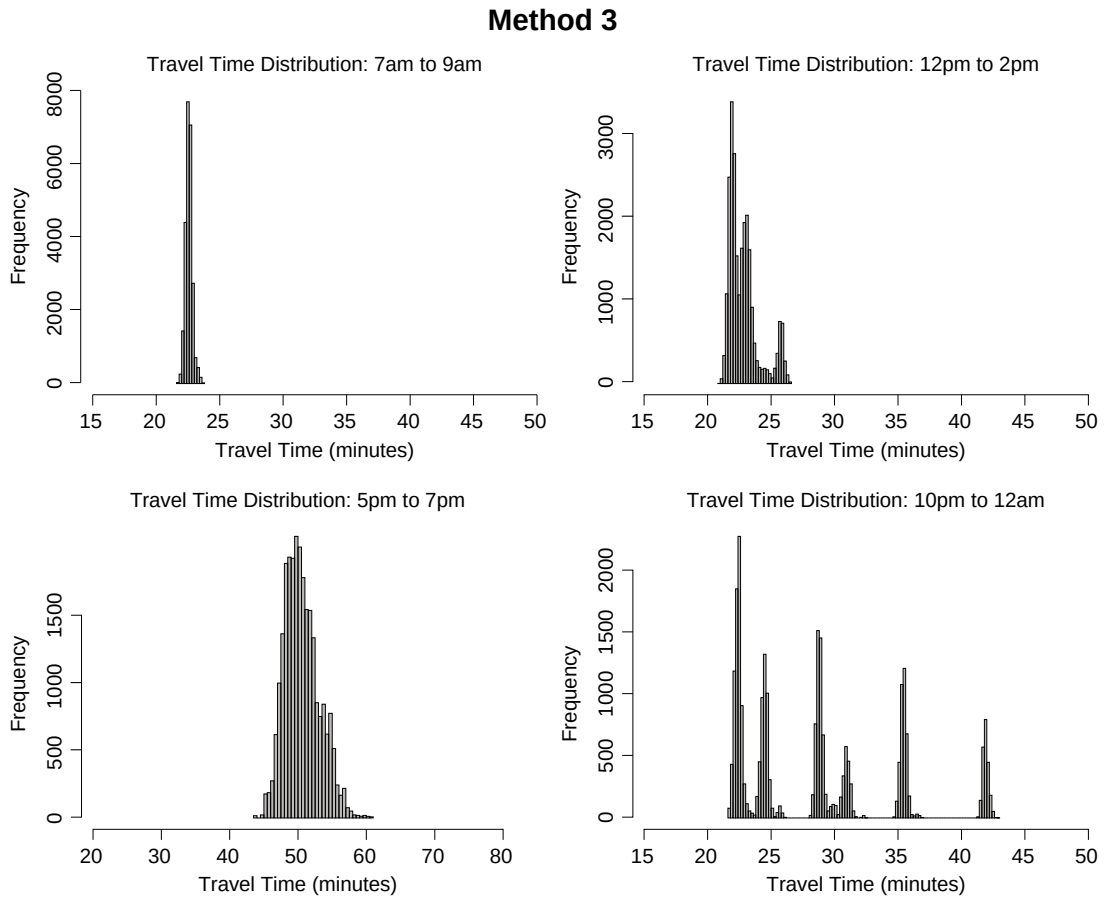

Exhibit C6-10: Route PDF generation method 3

Conclusions

Each of the three methods presented for assembling route travel time probability density functions from probe vehicle data is enabled by the techniques introduced in the methodology section. Constructing these PDFs requires identification of the data points corresponding to a particular route, separation of data by time of day when possible, and an understanding of the relationships between link speed distributions and route speed distributions. These tools, combined with the research team's findings related to speed correlations between consecutive links within a trip, led to the development of these three PDF-generation methods.

Methods 1 and 2 compared well with each other, while the results of method 3 differed in terms of travel time magnitude and variability. The differences in the shapes of the distributions across methods, particularly in the night-time period when data was sparse, demonstrates the

strong influence of the correlations of speeds along consecutive links within a route. With most of the nighttime coverage made by full trips composed of two or more points, the timestamp-based travel times dominated the night-time data set. The modes of these unusually shaped distributions reveal individual trips in the data.

Although results were not validated with a different data source, the probability density functions generated using methods 1 and 2 appear to match expectations. An online trip planner estimates the travel time on this route to be 28 minutes, which generally agrees with the distributions seen here. They resemble typical route travel time distributions, even though no trips were observed traveling along the entire route.

It is possible to extract quantitative travel time reliability metrics from the time of day travel time distributions compiled and presented in this section. Knowing the distribution of travel times on a route enables the data user to compute any reliability metric, such as planning time or buffer time.

LESSONS LEARNED

Overview

This case study demonstrates that it is possible to obtain trip reliability measures based on probe data, even when that probe data is sparse. The travel time distribution for the route is constructed from vehicles that only travel on a portion of the route, and takes into account the linear dependence of speeds on consecutive links. This case study also contributes techniques for creating time-space contour plots based on probe speeds. These contour plots can be made to represent any measured speed percentile, so that contours for the worst observed conditions can be compared with typical conditions.

Probe Data Characteristics

Much of this case study effort focused on understanding the aggregation steps used to convert data from GPS receivers into link-based speeds. Understanding the way raw GPS data is processed and aggregated is vital for proper interpretation of data elements. It also enables all components of the data set to be utilized to increase the richness of travel time PDFs.

As probe data finds wider adoption for travel time monitoring, it is important for users to understand that data from these sources is still sparse. This sparseness necessitates complex processes for determining travel time distributions on routes of interest. When GPS and other technologies reach a certain penetration rate in the population and more vehicles traverse entire routes, the assemblage of route travel time distributions will be simplified. Currently, however, the construction of well-formed PDFs requires that every element in the data set (from speeds on single links to complex travel times across multiple links) should be used to generate the distribution.

Probe data sparseness also increases the minimum level of temporal aggregation that can be supported by the data set. For example, in this case study, the quantity of data was not sufficient to measure route travel time reliability at a granularity of five-minutes, which was the common reporting unit for case studies that relied on loop detector data. Instead, aggregation had to be done at the peak period, multi-hour level. Additionally, in this case study, weekend trips could not be removed from the data set, as there were not sufficient weekday data points to generate full PDFs. Finally, in order to generate the presented results, all data points collected

over the 10 year span of the data set had to be used. In practice, this long time-frame does not allow for trend analysis. Transportation planners and operators often require an understanding of how route travel times vary on a day-by-day, week-by-week, and month-by-month basis.

One probe data characteristic that counteracts the sparseness problem is that data coverage is highest during the time periods, and at the locations, where the most vehicles are traveling on the roadway. These are also the time periods and locations at which reliability monitoring is the most critical. As probe technologies become more common in vehicles, the availability of data points and route-level trip data will naturally increase, resulting in richer data sets that can be analyzed at a finer-grained interval than were possible in this case study.

REFERENCES

1) U. S. Census Bureau. (2000). Population, Housing Units, Area, and Density: 2000 - State -- County / County Equivalent Census 2000 Summary File 1 (SF 1) 100-Percent Data. Retrieved February 22, 2012, from http://factfinder2.census.gov/faces/tableservices/jsf/pages/productview.xhtml?pid=DEC_00_SF1_GCTPH1.ST05&prodType=table.
2) American Community Survey. (2010). ACS S0802: Means of Transportation to Work by Selected Characteristics, 1-Year Estimates. Retrieved February 22, 2012, from http://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=5&ved=0CDwQFjAE&url=http%3A%2F%2Fwww.adcogov.org%2FDocumentView.aspx%3FDID%3D665&ei=3C9NT5nsOuOwiQKws4iZDw&usg=AFQjCNEOt2eTvtMIetEmYC9_631M00Cz1w.