

Establishing Monitoring Programs for Mobility and Travel Time Reliability (L02)
Strategic Highway Research Program 2
Transportation Research Board
National Academies

Atlanta Case Study Validation: Travel Time Reliability Monitoring

Authors:

Tiffany Barkley, Berkeley Transportation Systems, Inc.
Rob Hranac, Berkeley Transportation Systems, Inc.
Eric Mai, Berkeley Transportation Systems, Inc.
Berkeley, CA; February 2012

George List, North Carolina State University
Nagui Rouphail, North Carolina State University
Billy Williams, North Carolina State University
Raleigh, NC; February 2012

CONFIDENTIAL

Transportation Research Board of the National Academies Privileged Document

This report, not released for publication, is furnished only for review to members of or participants in the work of SHRP 2. This report is to be regarded as fully privileged, and SHRP 2 must approve dissemination of the information included herein.

Acknowledgements:

The Federal Highway Administration in cooperation with the American Association of State Highway and Transportation Officials sponsored this work. It was conducted in the second Strategic Highway Research Program (SHRP 2), which is administered by the Transportation Research Board of the National Academies. Bill Hyman and Mename Hedhli managed this project.

TABLE OF CONTENTS

1. Monitoring System.....	5
Study Description	5
Site Overview	6
Sensors	8
Data Management.....	11
Systems Integration.....	12
Other Data Sources.....	13
Summary.....	13
References.....	13
2. Methodological Advancement.....	15
Overview	15
Site Description	15
Method	16
Results	18
3. Use Case Analysis	19
Use Case 1: Integrating ATMS Data Into A Travel Time Reliability Monitoring System....	19
Use Case 2: Determining Travel Time Regimes And The Impact Of The Seven Sources Of Congestion	24
Use Case 3: Quantifying And Explaining The Statistical Difference Between Multiple Sources Of Vehicle Speed Data.	33
4. Lessons Learned.....	43
Overview	43
Systems Integration.....	43
Methodological Advancement.....	43
Probe Data Comparison	43

LIST OF FIGURES

Figure 1-1: Travel Time Reliability System Description	5
Figure 1-2: Map Of Atlanta Freeways	7
Figure 1-3: GDOT Traffic Detector Network	10
Figure 1-4: GDOT Managed Lane Detector Network	10
Figure 2-1: Downtown Connector Study Route	15
Figure 3-1: AA Multi-State Normal Pdfs	25
Figure 3-2: Midday Multi-State Pdfs	25
Figure 3-3: PM Multi-State Normal Pdfs	26
Figure 3-5: Midday Travel Times By Source	29
Figure 3-4: AM Peak Travel Times By Source	29
Figure 3-6: PM Peak Travel Times By Source	30
Figure 3-8: Locations Of Navteq TMC Paths (Longitudinal Black Lines) And Traficon Video Detectors (Perpendicular Black Lines)	36
Figure 3-7: Study Area On I-285	36
Figure 3-9: Common Temporal Aggregation Of Comparison Data	37
Figure 3-10: Comparison Of Speeds From Video (Black) And Probe (Gray) Sources	38
Figure 3-11: Comparison Of Speeds From Video (Black) And Probe (Gray) Sources	39
Figure 3-12: Cross-Correlation Of Data From Figure 3-5	39
Figure 3-13: Data From Figure 3-5 After Shifting Probe Data	39
Figure 3-14: Scatter Plots Comparing Correlation Of Speed Measurements With Distance Between Detectors	41
Figure 3-15: Difference In Speed Measurements (Video – Probe)	41

LIST OF TABLES

Table 1-1: GDOT Sensor Network Summary	8
Table 3-1: Mapping Of Lane Types From Navigator To Pems	20
Table 3-2: Data Reporting Frequencies By Device Type	21
Table 3-3: Event Data Format	22
Table 3-4: Event Data Set By Event Type	23
Table 3-5: Regime Parameters By Time Period	26
Table 3-6: Five-Minute Travel Time Samples By Time Period And Source	27
Table 3-7: Source Contributions To AM Peak Regimes	30
Table 3-8: Source Contributions To Midday Regimes	31
Table 3-9: Source Contributions To PM Peak Regimes	31

1. MONITORING SYSTEM

STUDY DESCRIPTION

This case study is the fourth of five performed by the project team in order to validate the approaches to travel time reliability monitoring described in the Travel Time Reliability Monitoring Guidebook. The goal of each case study is to illustrate how agencies apply best practices for: monitoring system deployment; travel time reliability calculations; and agency use and analysis of the system. To accomplish this goal, the team is implementing prototype travel time reliability monitoring systems at each of the five sites. These systems take in sensor data 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. This case study 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 Figure 1-1.

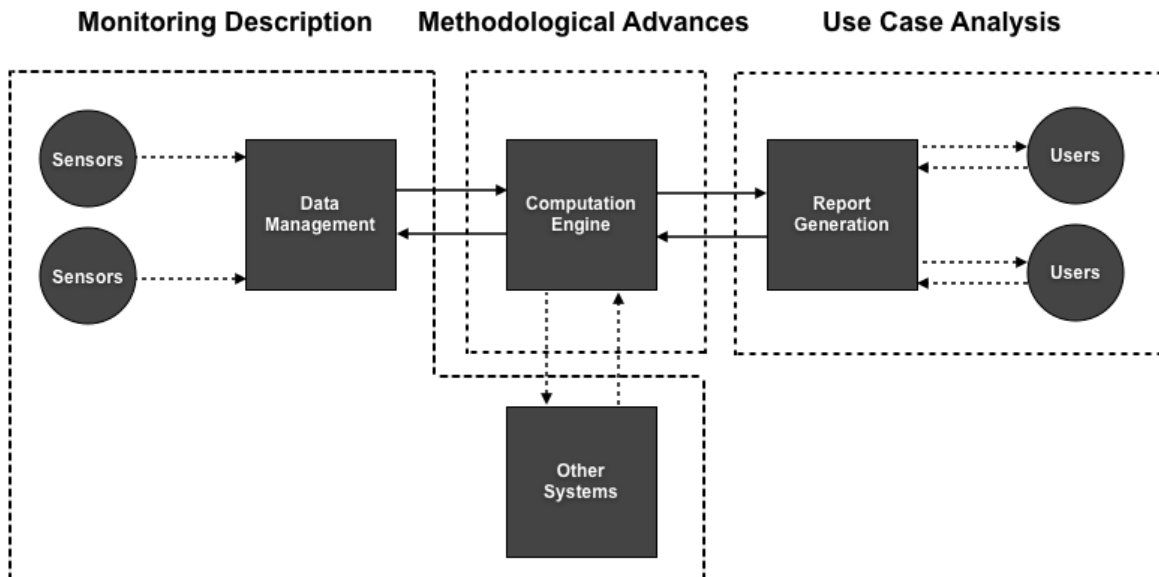


Figure 1-1: Travel Time Reliability System Description

This *monitoring system description* 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.

SITE OVERVIEW

The team selected the Atlanta 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. 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 Figure 1-2. 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

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.

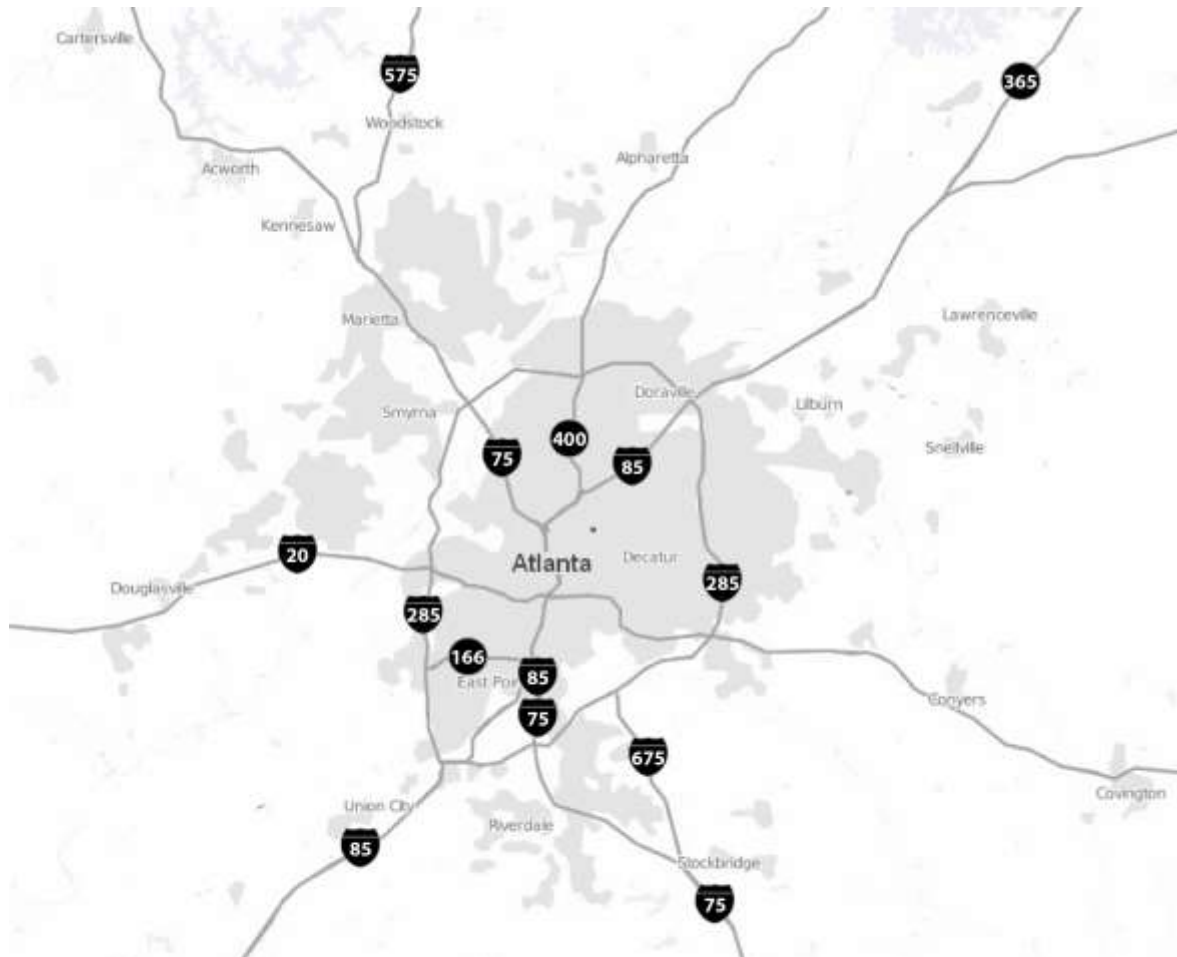


Figure 1-2: Map of Atlanta Freeways

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 1-1.

Table 1-1: 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. Figure 1-3 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. Figure 1-4 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.

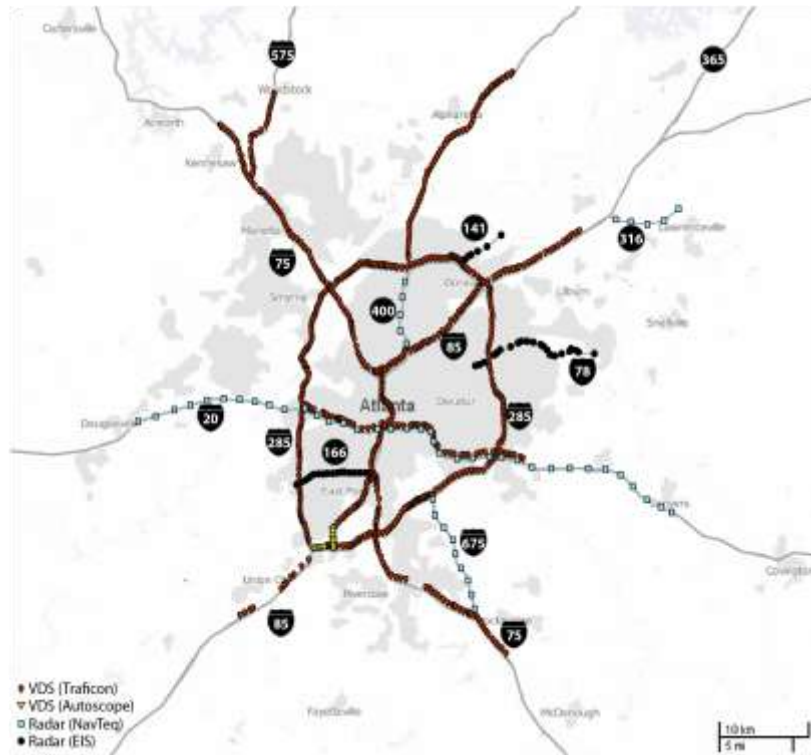


Figure 1-3: GDOT Traffic Detector Network

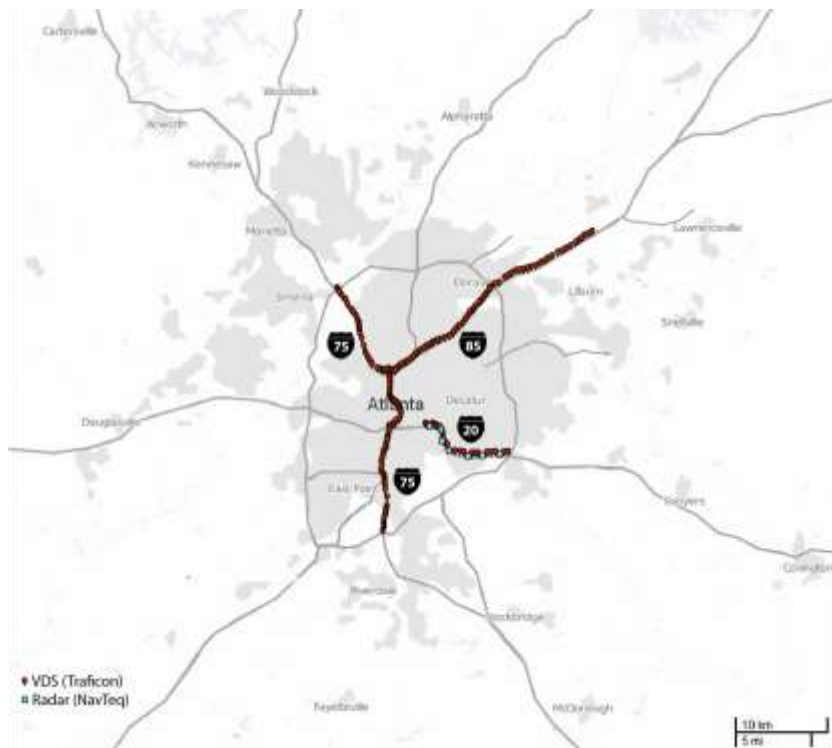


Figure 1-4: GDOT Managed Lane Detector Network

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 length of the TMC segments vary, but they 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.

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] Brunswick MPO 2030 Long Range Transportation Plan. Section 12. Glynn County, Georgia. 2005. <http://www.glynncounty.org/index.aspx?NID=1024>

2. Methodological Advancement

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. 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

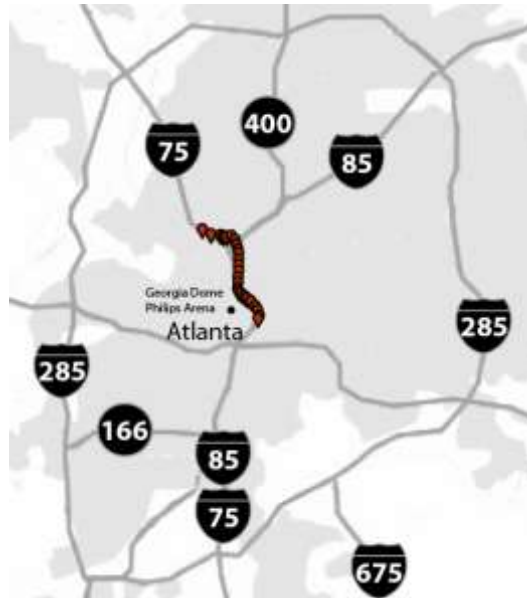


Figure 2-1: Downtown Connector Study Route

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,

$$(1) f(T|\lambda, \theta) = \sum_{K=1}^K \lambda_K f_K(T|\theta_K)$$

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 K th state, λ_K is the probability of the K th state occurring, and θ_K is the distribution parameters for the K th 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 1/10th 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; and
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.

3. Use Case Analysis

USE CASE 1: INTEGRATING ATMS DATA INTO A 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 associated 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 appropriately categorize Navigator detectors in PeMS, as shown in Table 3-1. 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 3-1: 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 3-2 shows the observed reporting frequencies by vendor type.

Table 3-2: 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 to handle this issue was to change the process for extracting data from the GDOT raw database table. Instead of extracting data from all 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 3-3. It contained 21,540 event records summarizing Navigator events within the Atlanta metropolitan region.

Table 3-3: 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 3-4 (grouped and summed into event types in similar categories).

Table 3-4: 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 AM, midday, and PM weekday 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. Figure 3-1, Figure 3-2, and Figure 3-3 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 3-5 summarizes the regime parameters (probability of occurrence and mean travel time) by time period.

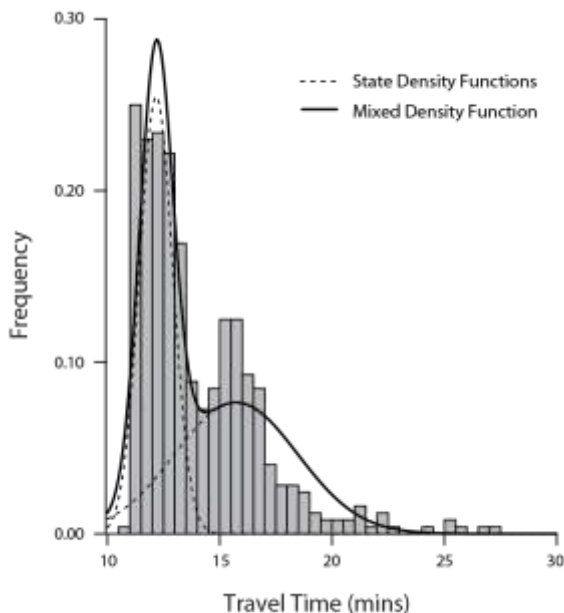


Figure 3-1: AM Multi-state Normal PDFs

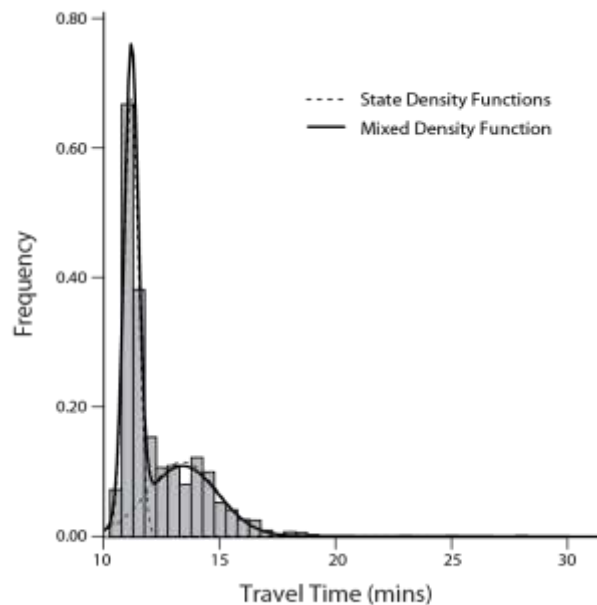


Figure 3-2: Midday Multi-state PDFs

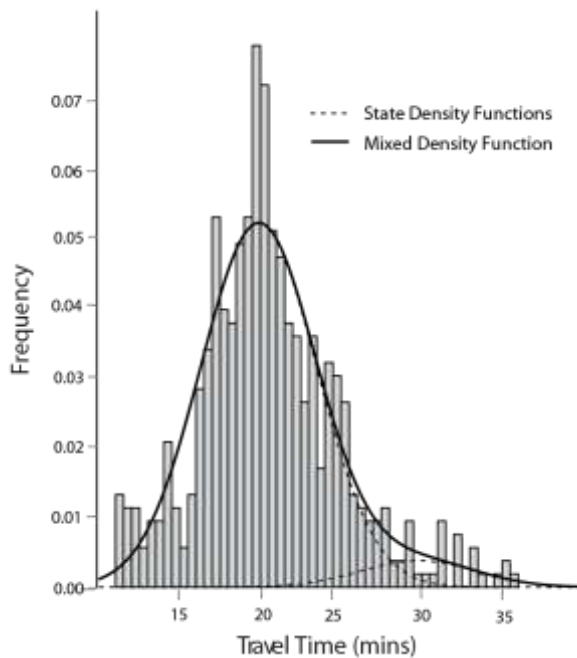


Figure 3-3: PM Multi-state Normal PDFs

Table 3-5: 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 Figure 3-2.). 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 3-6 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 3-6: 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. Figure 3-5, Figure 3-4, and Figure 3-6 illustrate the breakdown of travel times by source within each state. Table 3-7, Table 3-8, and Table 3-9 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 Figure 3-6 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.

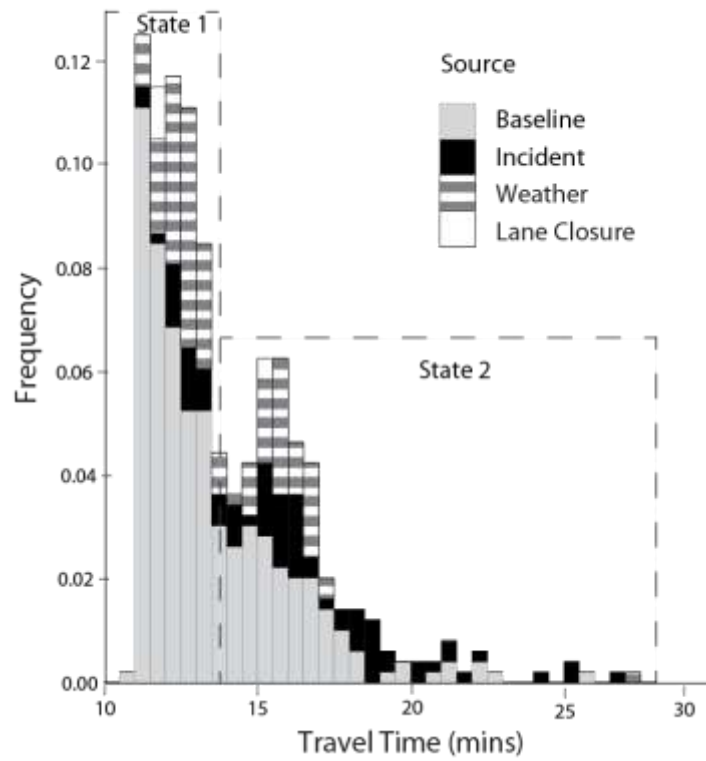


Figure 3-5: AM Peak Travel Times by Source

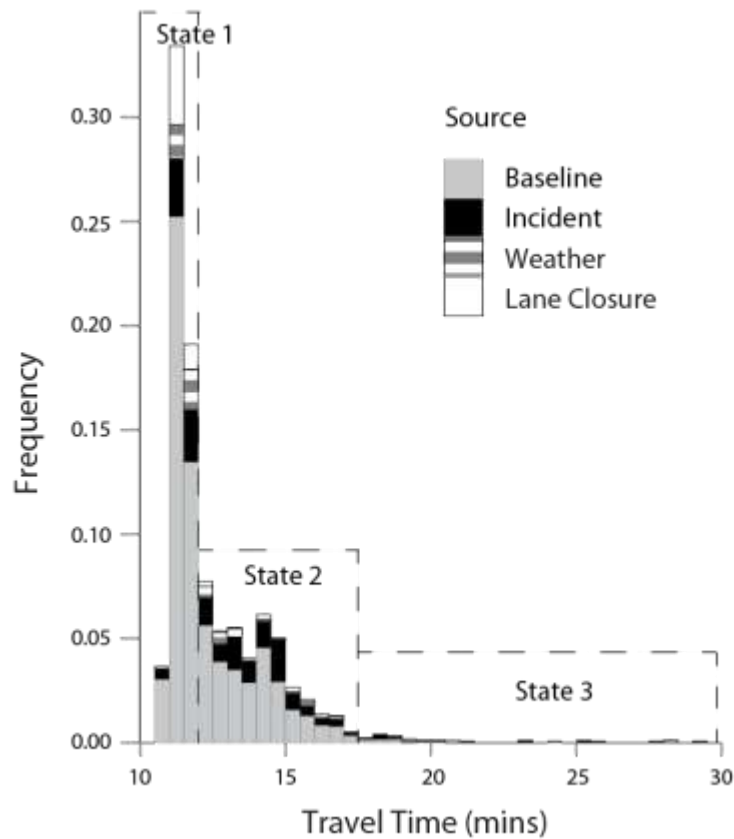


Figure 3-4: Midday Travel Times by Source

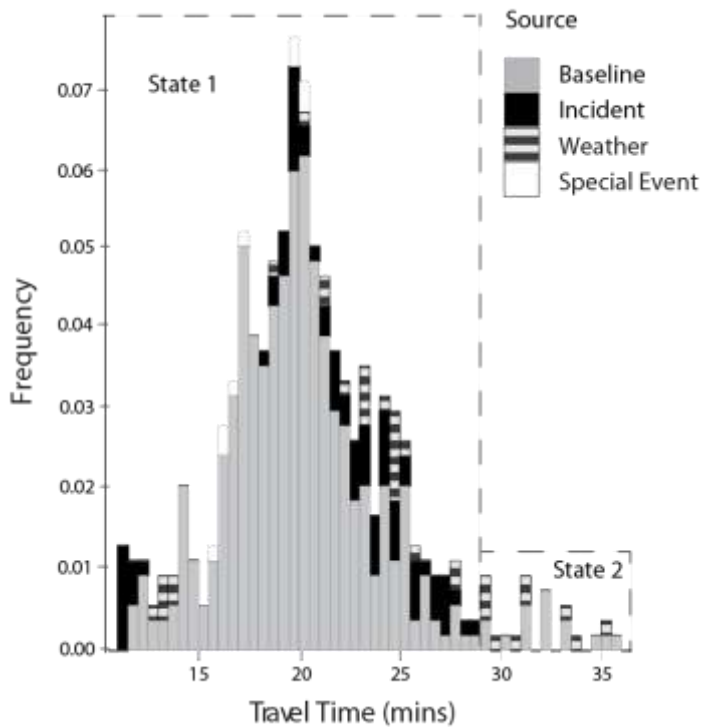


Figure 3-6: PM Peak Travel Times by Source

Table 3-7: 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 3-8: 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 3-9: 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 postmile. 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 postmile 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 in Chapter 1, 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 postmiles 25 to 35 in the northbound direction, and 45 to 55 in the southbound direction. Although these postmile ranges differ, they represent the same stretch of roadway (see Figure 3-8). 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 between 80,000 and 90,000 vehicles/day in the northbound direction and approximately 100,000 vehicles/day in the southbound direction.

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, 7 of which lie within a TMC path. In the Southbound direction, 5.3 of the 10 miles in the study area are covered by 8 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 Figure 3-7).

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.



Figure 3-8: Study Area on I-285

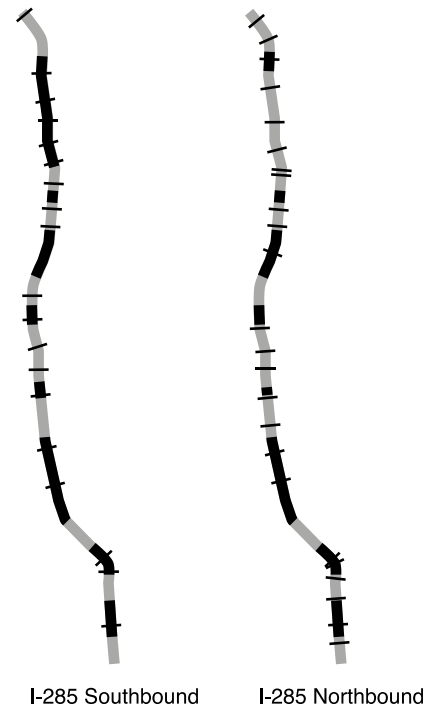


Figure 3-7: Locations of Navteq TMC paths (longitudinal black lines) and Traficon video detectors (perpendicular black lines)

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 Figure 3-9). 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.

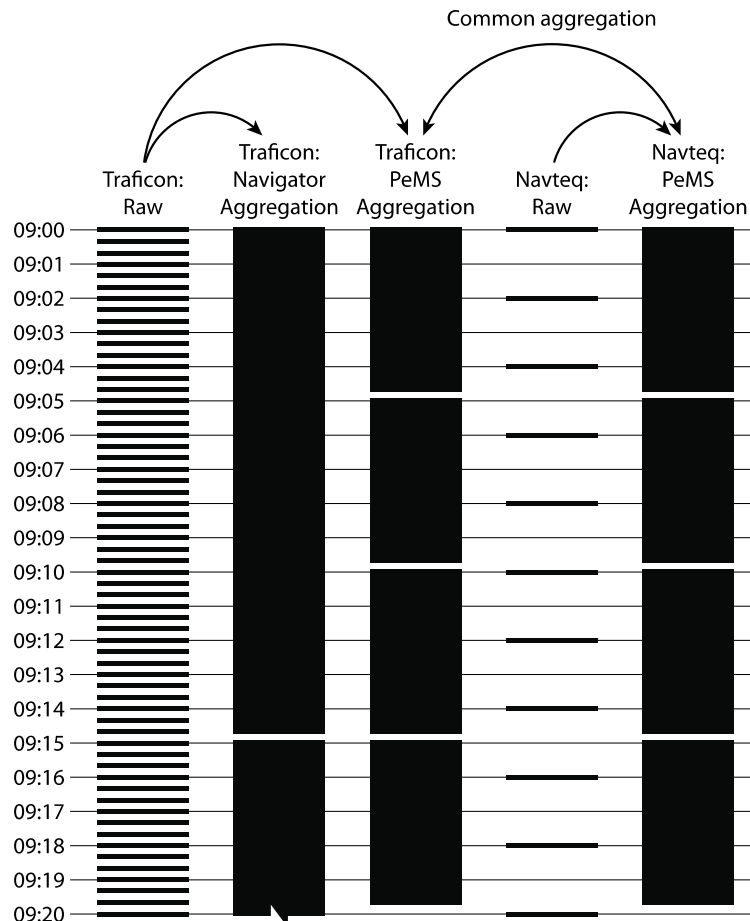


Figure 3-9: Common temporal aggregation of comparison data

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 Figure 3-8). There were 7 pairs of video detectors and TMC paths in the northbound direction and 12 in the southbound direction.

With video and probe detectors paired by location, their speed measurements can be plotted and compared visually. Figure 4 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. Figure 3-11 plots the results of this filtering on the time range and data from Figure 3-10, showing all of the time points from Figure 3-10 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.

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.

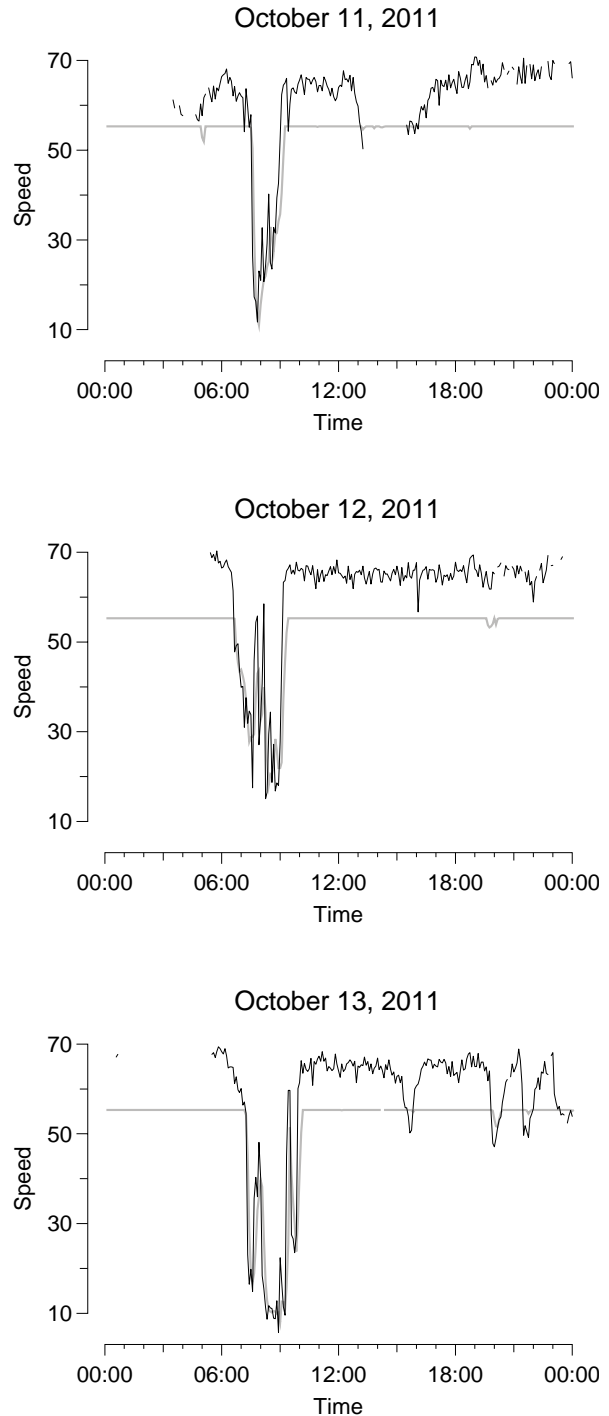


Figure 3-10: Comparison of speeds from video (black) and probe (gray) sources

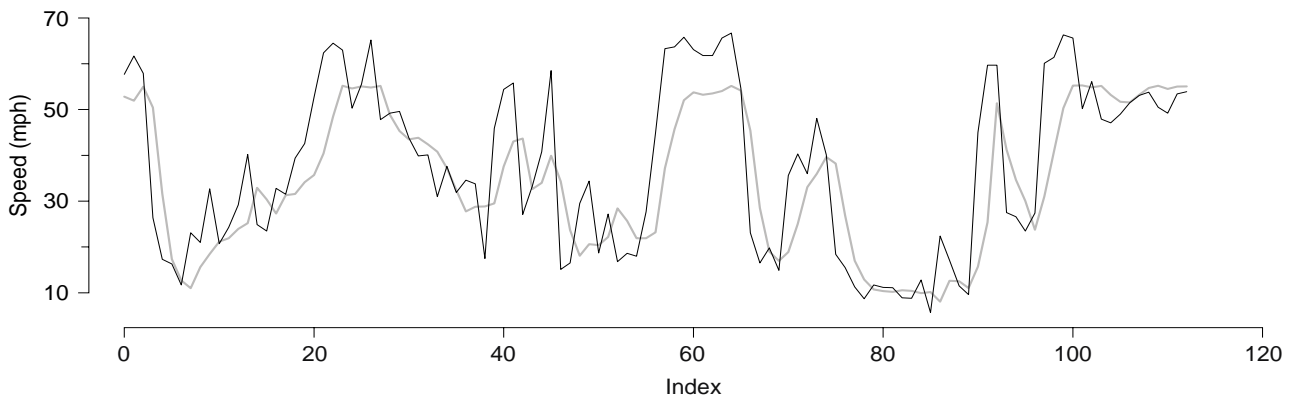


Figure 3-11: Comparison of speeds from video (black) and probe (gray) sources

Upon inspection of Figure 3-5, 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 Figure 3-5 was computed. It can be seen in Figure 3-6 that the peak correlation occurs at a lag of -1. The unshifted data, as shown in Figure 3-5 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 Figure 3-6). This technique can be used to calibrate sensor measurements.

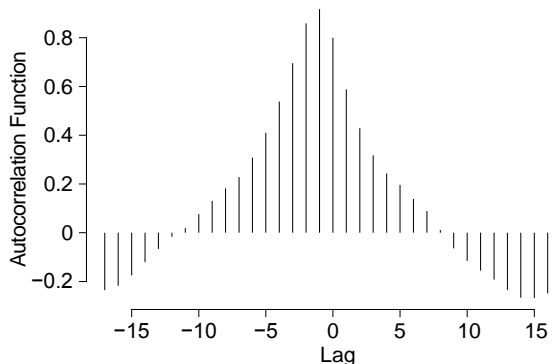


Figure 3-12: Cross-correlation of data from Figure 3-11

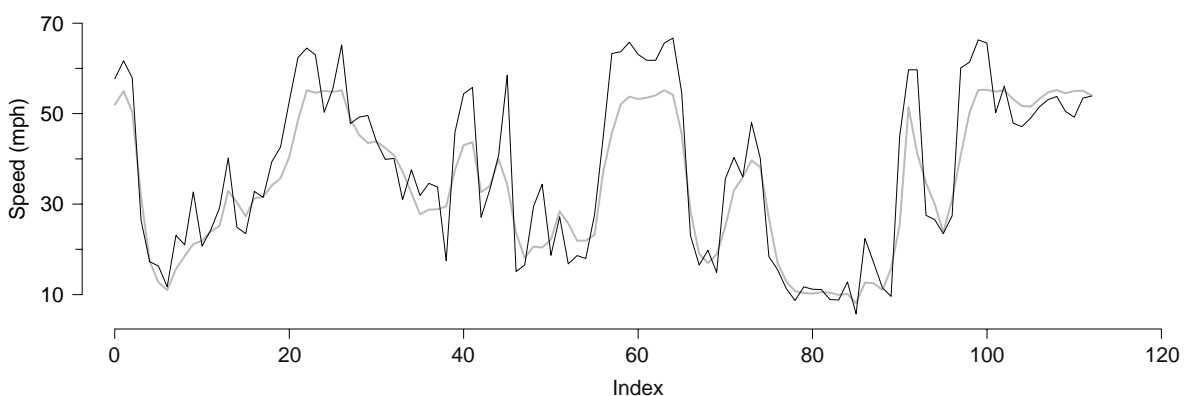


Figure 3-13: Data from Figure 3-11 after shifting probe data

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 Figure 3-11). 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. Figure 3-13 shows the difference in speed measurements for the same pair of detectors and time range as in Figure 3-10 and Figure 3-11. Speed measurements from this pair of detectors matched well, with a correlation coefficient of 0.85. Figure 3-11 shows both speed profiles in general agreement. However, when the difference in speed measurements is plotted in Figure 3-13, 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.

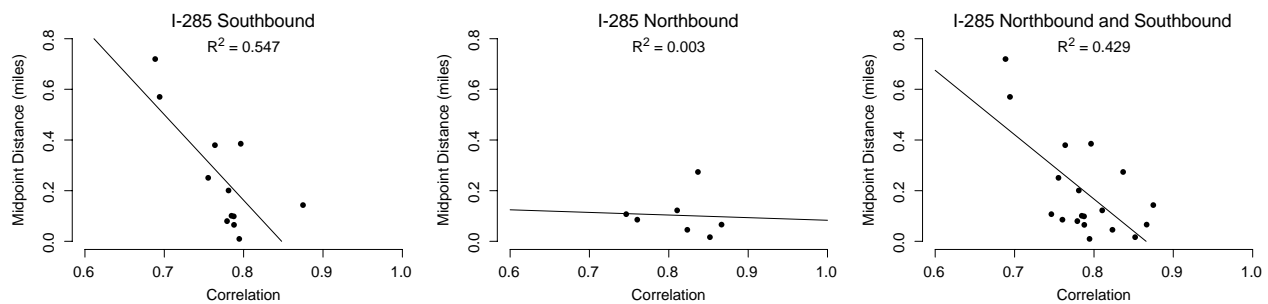


Figure 3-14: Scatter plots comparing correlation of speed measurements with distance between detectors

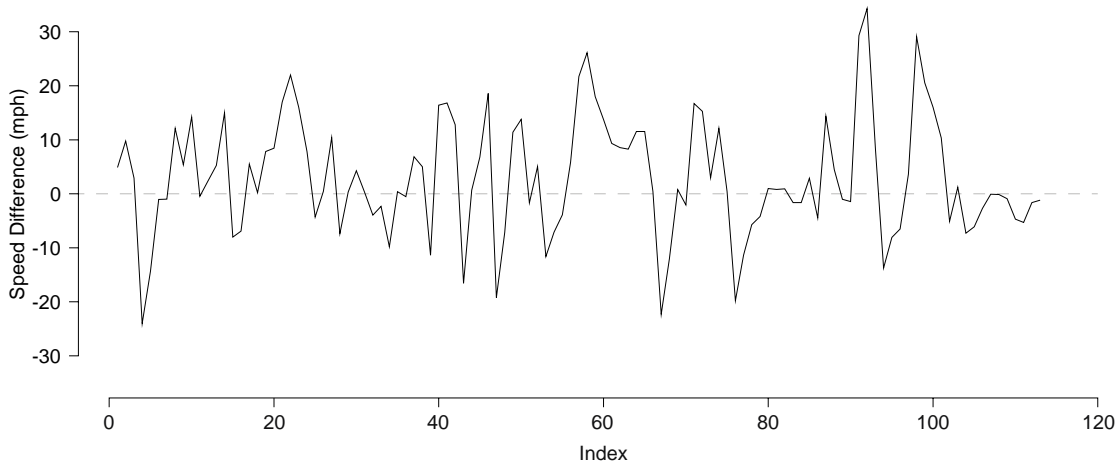


Figure 3-15: Difference in Speed Measurements (video – probe)

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, the hazards of comparing data from individual 5-minute periods was seen by plotting the difference between two data sets.

4. 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 dealt 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.