FINAL REPORT:
Vehicle Infrastructure Integration Proof-of-Concept Results and Findings – Infrastructure

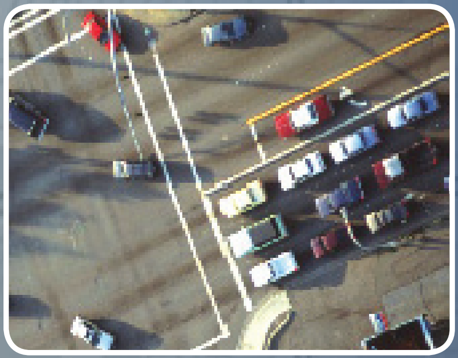

VOLUME 3B

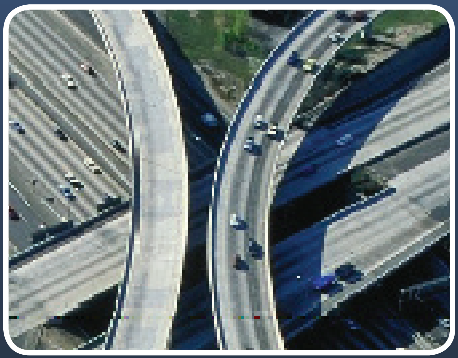

U.S. Department of Transportation
Research and Innovative Technology Administration

Technical Report Documentation Page

1. Report No. FHWA-JPO-09-057	2. Government Accession No.	3. Recipient's Catalog No. EDL# 14488

4. Title and Subtitle Final Report: Vehicle Infrastructure Integration Proof-of-Concept Results and Findings – Infrastructure	5. Report Date May 2009
	6. Performing Organization Code
7. Author(s) Ram Kandarpa, Mujib Chenzaie, Justin Anderson, Jim Marousek, Tim Weil, Frank Perry, Ian Schworer, Joe Beal, Chris Anderson	8. Performing Organization Report No.
9. Performing Organization Name and Address Booz Allen Hamilton 8283 Greensboro Drive McLean, VA 22102	10. Work Unit No. (TRAIS)
	11. Contract or Grant No.
12. Sponsoring Agency Name and Address Research and Innovative Technology Administration (RITA) U.S. Department of Transportation 1200 New Jersey Avenue, SE Washington, DC 20590	13. Type of Report and Period Covered
	14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

In 2005, the US Department of Transportation (DOT) initiated a program to develop and test a 5.9GHz-based Vehicle Infrastructure Integration (VII) proof of concept (POC). The POC was implemented in the northwest suburbs of Detroit, Michigan. This report outlines the purpose and procedures for various tests, identifies the test articles, and documents the results of that testing. It also discusses the implications of those test results relative to the overall viability of the VII concept and identifies recommendations for future work, including whether and how designs and standards need modification. This volume is intended for engineering managers and practicing engineers interested in the design and development of VII systems and applications.

17. Key Word Vehicle Infrastructure Integration (VII), Connected Vehicle, Cooperative Vehicle Infrastructure, IntelliDrive^{SM}	18. Distribution Statement

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

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Table of Contents

Table of Figures

Table of Tables

Front Cover Photos: *Booz Allen stock photos*
Front and Back Cover Maps: *Booz Allen-generated plot points on Google map*

1. INTRODUCTION

The Vehicle Infrastructure Integration (VII) program[1] is a joint government-industry research effort focused on developing standardized wireless vehicular communications for two primary purposes:

- **Among vehicles (vehicle-to-vehicle)** – Mostly as a means of enabling advanced crash avoidance applications
- **Between vehicles and various entities that notionally reside within the broader infrastructure (vehicle-to-infrastructure)** – Mostly as a means of collecting enhanced roadway condition information and broadcasting various alerts and related traveler information back to vehicles.

The proposed wireless media is based on dedicated short-range communications (DSRC) technology operating at 5.9GHz, and the communication standards are based on Wi-Fi (802.11xx) family of standards. However, the new vehicular communications standards have been heavily modified to allow for fast connection times between entities, optimize the available bandwidth, ensure security and privacy of communications, and support a wide variety of public and private sector applications. The standards and protocols that embody the VII wireless communications concept are captured in the IEEE 802.11p, IEEE P1609 suite of standards, and Society of Automotive Engineers (SAE) J2735, which describe physical through application layer standards needed to enable VII messaging.

Leveraging the specialized DSRC communications and VII standards, vehicles would continuously broadcast their exact position, speed, and heading (at about 10 times per second). Surrounding vehicles could then receive these “heartbeat” messages so that any particular vehicle would have an accurate, real-time “awareness” of all of the other vehicles in its vicinity. This dynamic information would then be analyzed and used to determine whether an unsafe driving condition was developing—and subsequently warn the driver to take action, or even automatically take control of the vehicle’s steering or braking in order to mitigate or avoid the accident.

As part of the overall VII concept, in addition to the new vehicular wireless communications media, a complementary communications network is also proposed. This network supports unique communication services that help manage messaging between the vehicle and the infrastructure including message priorities; ensure security and privacy for both “publishers” (vehicles that are sources of vehicle behavior information) and “subscribers” (those entities that use the vehicle probe data to derive information and communicate with vehicles); and augment the multi-channel and other specialized features of the VII communication standards. The VII network provides an appropriate interface between a VII-equipped vehicle and broader public communication networks, so that messaging with the vehicle can be completed in a safe, secure, and efficient manner.

Under the VII concept, vehicles will be equipped with a DSRC radio, a highly accurate on-board positioning system, and an appropriately configured on-board computer to facilitate communications, support various applications, and provide an interface to the driver (collectively, this equipment is called the OBE, or on-board equipment). Vehicles would communicate with each other and with roadside

---

[1] The program name, Vehicle Infrastructure Integration (VII), was in official use at the inception of and during the execution of the work described in this report. The US Department of Transportation has initiated a new program titled “IntelliDrive^{sm},” which now encompasses all activities that were previously part of VII. The IntelliDrive^{SM} logo is a service mark of the U.S. Department of Transportation.

transponders (or roadside equipment, RSEs), which would be linked to the specialized VII network. RSEs would be positioned at major signalized intersections and along interstates and major arterials.

The VII-equipped vehicles periodically record their precise location, speed, heading, and other information (such as temperature, wiper, and headlight status) at frequent intervals (based on time and speed factors), and then transmit this data to the infrastructure where it could be “subscribed” (requested) by public and private entities. The information packets (also called vehicle probe data) could then be collected, analyzed, and used to report (and predict) travel times along road segments, roadway incidents, and other localized road surface and weather information. Such information could then be broadcast back to vehicles to allow drivers (and an on-board navigation system if so equipped) to make better trip decisions. Because such probe data are accurate, detailed, near real time, and (eventually) ubiquitous, they could represent a huge change in the reporting and prediction of roadway conditions—and could even be used to control and coordinate traffic signals and other traffic management systems, thereby allowing for increased throughput on the nation’s roadways.

Based on early government-industry discussions, the long-range deployment plan called for VII communications capability to be built into new vehicles as individual models are periodically updated (model changes), and for governments (local, state, and federal) to sponsor deployment of the communications infrastructure. While business and operating models have not yet been fully developed, early concepts called for the VII network to be operated by some type of joint (and non-profit) government-industry partnership. Since the VII program was initiated, its concept has expanded to include additional communications means between the vehicle and infrastructure. For example, the evolving VII concept allows for VII-enabling communication standards (potentially including, but not mandating, DSRC) to also be integrated into aftermarket devices such as personal navigation devices or even cell phones—thus accelerating deployment. Further, the VII communications concept has expanded to allow for multiple wireless media (such as cellular, satellite, and/or WiMAX) to augment DSRC-based messaging and for various applications and services to be delivered to vehicles and drivers using the most efficient delivery channel (or media) that meets each application’s particular performance, security, and reliability requirements. Finally, business and operating models are being revisited to encourage more private sector involvement, sponsorship, and operating authority.

The US Department of Transportation (DOT) initiated a program to develop and test a VII proof of concept (POC) in support of a nationwide deployment decision. The POC test was executed by the U.S. DOT and its private sector partner, the VII Consortium (VIIC)—an organization comprising of several light-duty original equipment manufacturers (OEMs) and industry participants. The US DOT has engaged Booz Allen to design the national VII network architecture, act as system integrator, and implement a POC system as part of a VII development test environment (DTE) near Detroit, Michigan.

2. ORGANIZATION OF THE FINAL REPORT

The *Final Report: Vehicle Infrastructure Integration Proof-of-Concept* is organized into three volumes, as described below:

- ***Volume 1B – Final Report: Vehicle Infrastructure Integration Proof-of-Concept Executive Summary – Infrastructure***, provides an overview of the key findings and recommendations from the POC testing. This volume is intended for executives and managers of organizations interested in the deployment of VII.
- ***Volume 2B – Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Infrastructure***, describes the overall approach undertaken to prove the VII concept through a structured testing program. It describes the overall experimental design used in proving the VII concept by providing an overview of the system architecture and the design of systems, subsystems, and components, as well as the public sector applications developed to prove some of the system concepts. This volume is intended for engineering managers and practicing engineers interested in the design and development of VII systems and applications.
- ***Volume 3B – Final Report: Vehicle Infrastructure Integration Proof-of-Concept Results and Findings – Infrastructure*** (this document), outlines the purpose and procedures for various tests, identifies the test articles, and documents the results of that testing. It also discusses the implications of those test results relative to the overall viability of the VII concept and identifies recommendations for future work, including whether and how designs and standards need modification. This volume is intended for engineering managers and practicing engineers interested in the design and development of VII systems and applications. It assumes the reader has read and is familiar with the technical description in Volume 2.

3. VII POC PROGRAM OVERVIEW

In 2005, the US DOT initiated a program to develop and test a 5.9GHz based VII POC, in support of a nationwide deployment decision. The POC test was executed by the US DOT and its private sector partner, the VIIC—an organization comprising of several light-duty OEMs and industry participants. The US DOT has engaged Booz Allen to design the national VII network architecture, act as system integrator, and implement the POC system.

The POC DTE was implemented in the northwest suburbs of Detroit, Michigan. Fifty-two RSEs were installed within 45 square miles, 27 vehicles were configured with OBEs, and a communications network was established. Further, a limited number of public and private sector applications were developed, primarily as a means of testing the end-to-end functionality and performance of the VII system.

3.1. POC Goals and Objectives

The primary goals of the POC included demonstrating the technical performance and functionality of the VII architecture and associated concept and proving that safety, mobility, and commercial (private) applications can be effectively implemented. Key among the technical and functional requirements was the need for vehicular users to maintain their anonymity for applications that track a vehicle’s location from point to point, while the overall system ensures a high level of security.

The design of the POC test environment and associated system software and applications was completed based on objectives established to meet the aforementioned goals. The following are the key objectives identified for the POC:

- **Validate Standards:** The IEEE 1609 and 802.11p suite of standards were established for the vehicle-vehicle and vehicle-infrastructure communication paths used in VII. In addition, the SAE J2735 standard was established to provide a common message set for these communications. A POC objective was to validate that these standards properly function in the 5.9GHz band allocated for use with VII.
- **Provide Core Services:** A set of core services are part of the VII architecture. These include publish/subscription services, network management, mapping and positioning services, and a Certificate Authority. An objective of the POC was to test these services to ensure they properly function.
- **Support Applications:** Successful implementation of VII at the national level depends on the architecture being capable of supporting simultaneous operation of the safety, mobility, and commercial (private) applications expected to run on the system. An objective of the POC was to demonstrate that a basic set of applications, representative of those expected to be part of the initial set used on VII, can concurrently run on VII.
- **Demonstrate Security and Privacy:** The ability to protect against malicious intrusion, while maintaining privacy, is an essential part of the VII concept. An objective of the POC was to test the susceptibility of the system to intrusion (hackers) while ensuring the anonymity of users.

3.2. POC Scope

The scope of the POC test was developed based on the goals and objectives described above. While the VII communication standards and technology (as originally envisioned) are based on existing Wi-Fi platforms, there are significant modifications made to enable high-speed, low-latency communications with moving vehicles. In addition, the envisioned public and private sector applications that would use VII technology require special provisions related to security, privacy, and reliability to ensure safety and anonymity of motorists. The realization of VII wireless standards and implementation of the roadside network that enables communications to and from vehicles required addressing special challenges that have been unique to the VII program—and therefore required significant development and innovation as well as thorough testing of the concepts and technology solutions.

The development of the VII wireless standards (IEEE 802.11p and the IEEE P1609 suite of standards) began in earnest in early 2000. The design of the VII infrastructure network began in 2005. While both of these efforts continue as of this writing, the design concepts, prototype hardware, and prototype software had progressed sufficiently by late 2006 that it was deemed appropriate to develop and test a small-scale, integrated VII system that embodied all of the key VII design elements, thereby testing the overall VII concept. The resulting POC evaluation plan focused on first testing VII subsystems and components under controlled laboratory conditions, then testing the integrated VII concept under controlled garage and test track environments, and finally implementing a model of a full-scale VII system, including VII-equipped vehicles and a prototype communications infrastructure network in a real-world, uncontrolled environment. The host site for the real-world testing is a suburb in the northwest area of Detroit, Michigan, where 52 RSEs were installed within 45 square miles encompassing both highways and arterials, 27 vehicles (25 core test vehicles and 2 Toyotas used as spares) were configured with VII technology (OBEs), and a communications network was established. Further, a limited number of public and private sector applications were developed, primarily as a means of testing the end-to-end functionality and performance of the VII system.

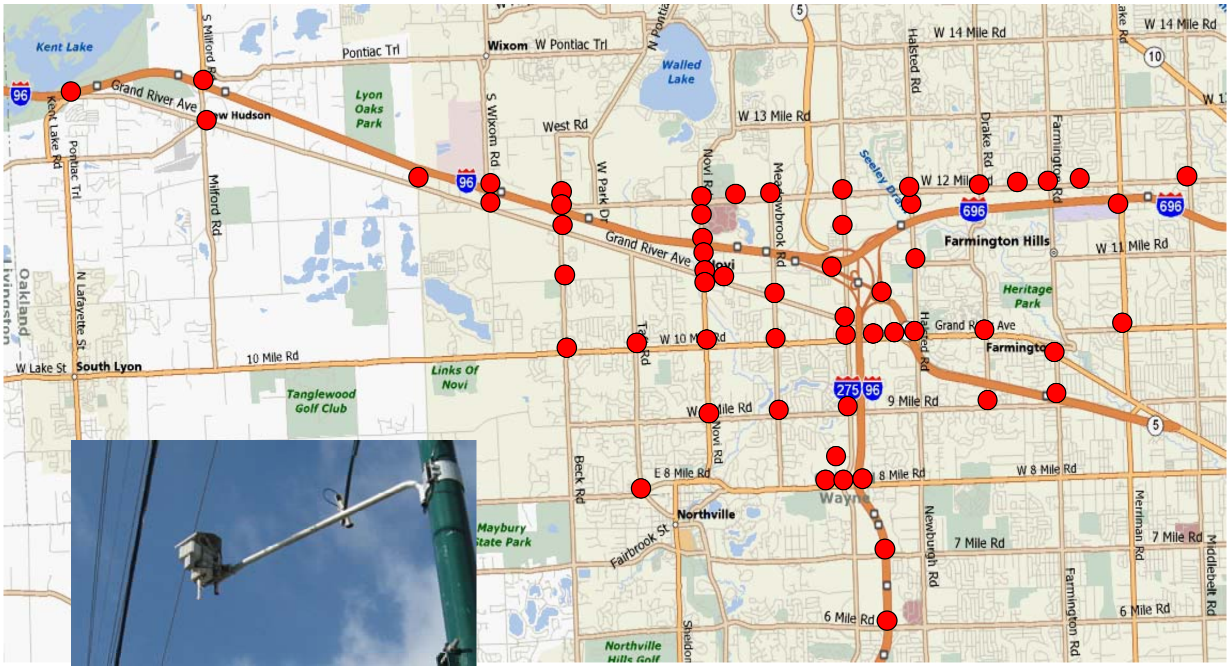

Booz Allen staff photo;
Booz Allen-generated plot points on Google map

Figure 1: VII POC DTE

By testing the VII concept and systems in a real-world uncontrolled environment in Michigan (see Figure 1; red dots represent RSEs), designers and developers were able to: 1) determine whether the new communication standards, technologies, hardware and software solutions do in fact support the intended applications (i.e., that the VII system "works"); 2) obtain valuable information for refining the concept and solutions; and 3) demonstrate the potential value of this new wireless vehicular communications media to both government and private stakeholders.

The VII testing program was structured into three main phases—subsystem test, system integration and test, and public and private applications testing. Subsystem testing was divided into three parts to coincide with the three software build deliveries for the infrastructure and RSEs. The testing of each build corresponded to the functionality of the software delivered in each build. During Build 1 subsystem testing, only the most basic functionality was tested, such as the ability to receive and forward a probe data message. Build 2 software added additional functionality, like the ability to parse probe data messages into the different probe data elements. Build 2 subsystem testing focused on verifying the new functionality and regression testing the old functionality. These testing efforts led to the final subsystem test of the Build 3 software, which included the integration among all of the infrastructure elements—Advisory Message Delivery Service (AMDS), Probe Data Service (PDS), Information Lookup Service (ILS), and Identity and Access Management (IdAM)—and the RSE proxies. Build 3 subsystem testing also included a regression test of all previous functionality. After completion of subsystem testing, the VII POC moved into the system integration and test phase. The goal of this phase was to validate the VII national system requirements (NSR) and demonstrate that the end-to-end VII system was functional and ready to support user application tests. The goal of the applications testing phase was to demonstrate the ability of the VII architecture to support multiple public and private sector applications. Public and private applications testing were conducted from May 2008 to September 2008.

3.3. Participants

The US DOT and its private sector partner the VIIC executed the POC. The US DOT team included the ITS Joint Program Office (JPO), Noblis, and Booz Allen Hamilton. Booz Allen acted as the system integrator for the infrastructure and public applications, while the VIIC led the development and testing of the OBE and private and safety applications. The Booz Allen team consisted of Raytheon, Iteris, Sirit, Telcordia, and TechnoCom (now Kapsch). The Michigan Department of Transportation (MDOT), Road Commission for Oakland County (RCOC), and Wayne County – Division of Roads (WCDR) provided local support in Michigan. The VIIC organization during the POC consisted of nine of the primary vehicle OEMs and a diverse team of suppliers from across the industry. This group of government agencies, contractors, suppliers, and OEMs guided and executed the VII POC.

3.4. Roles and Responsibilities

Tables 1 and 2 summarize the roles and responsibilities of the Booz Allen team and the VIIC organization respectively.

Table 1: Roles and Responsibilities of Booz Allen Team

	Pgm Mgmt	System Eng	SDN	ENOC & CA	RSE	DTE	SIT	Public Apps	POC Testing	O&M	Stds
Booz Allen Hamilton	X	X	X	X	X	X	X	X	X	X	X
Raytheon		X			X	X	X				X
Iteris		X						X			X
Telcordia		X	X	X			X				
Sirit					X		X				
TechnoCom / Kapsch					X		X				
MDOT / RCOC / WCDR						X			X	X	

Table 2: Roles and Responsibilities of VIIC Organization

Participant	Work Order									
	Pgm Mgmt	System Eng	Radio	Policy Support	OBE	Apps	Positioning	Security	Test	Private Service Enablers
ABSS Inc.	X									
ARINC Inc.							X			
Battelle Institute		X							X	
BMW of North America Inc			X		X					
Cogenia Partners LLC	X	X						X	X	
Delphi Automotive Systems Inc					X	X				
Denso International America			X		X					
Dykama Gossett				X						
Intel Americas Inc.										X
Honda R&D Americas Inc.									X	
TechnoCom / Kapsch			X					X		X
Mark IV IVHS Inc						X				
Mercedes Benz			X		X					
Moser Racing	X									

Participant	Work Order									
	Pgm Mgmt	System Eng	Radio	Policy Support	OBE	Apps	Positioning	Security	Test	Private Service Enablers
LLC										
MTS LLC						X				
Navteq North America LLC						X				
Nissan Technical Center NA				X			X			
Ntru Cryptosystems Inc								X		
Parvus Corporation					X					
Prosyst Software Gmbh					X					
Raytheon Company						X			X	
Roush Industries									X	
Sirit Technology Inc.				X						
Telcordia Technologies								X	X	
Toyota Motor Engineering and Manufacturing NA Inc									X	
Transcore LP.			X							
VIIC Membership	X									
WFET Group		X							X	
Wind River Systems Inc.					X					

3.5. Key Concepts

The POC system was scoped and designed to assess and prove the following core concepts, with the intent of providing critical inputs for a viability decision:

- The viability of the IEEE 1609 emerging standards and **DSRC performance characteristics** of individual links and collectively within a LAN. These characteristics include link performance, channel switching, service advertisements, service join/un-join.

- The collection of **anonymous J2735 probe data** generated by vehicles through roadside wireless access points (RSEs) for near real-time distribution to remote data users via the VII infrastructure. These data consumers subscribe to topics of interest, and VII infrastructure will provide them with a copy of any probe data that matches their criteria.
- The remote input and distribution of SAE J2735-defined **transportation operations center (TOC) advisory messages** through the VII infrastructure for subsequent DSRC broadcast to vehicles in targeted geographic locations for a specified time interval, using a prioritization scheme. The reception of those messages by vehicles and prioritizing for display to driver.
- The viability of various classes of **end-to-end communications** between applications running on vehicles, remote network servers (TOCs), and roadside infrastructure (signal controllers).
- **Authentication** of vehicle-to-infrastructure, infrastructure-to-vehicle, and vehicle-to-vehicle communications.
- The security and **privacy** of end user information as it traverses through the VII system.
- The utility and performance impacts of vehicle-to-vehicle beacon (“heartbeat”) messages to provide vehicles with situational tactical awareness in support of critical safety applications.
- Generation, distribution, and operational validation (anomaly detection) of **localized micro maps** containing detailed roadway geometries for intersections and roadway segments.
- The accurate, real-time **positioning** of operational vehicles and the viability of positioning augmentation through positioning correction data provided by the VII infrastructure.
- The **remote provisioning and management** from a central platform of (prototype) RSE units, deployed in various configurations in typical roadside locations.

3.6. POC Standards

The radio portion of the POC RSE and OBE were based on IEEE standards. Additionally, selected data exchanged between the RSE and OBE and between OBEs are formatted using SAE standards. The following list of standards was implemented (with some variations and/or extensions) for the POC:

- **SAE J2735-DSRC Message Set:** The POC application/data message set layer is based on a variation of J2735 Version 1. Variations are captured in *VII POC DSRC Msg Set.pdf*.
- **IEEE 1609.3-Networking:** The POC network layer is based on a variation of 1609.3/D18 and includes aspects from 1609.3/D21. Variations are captured in *SYS120-04 OBE-RSE IRS.pdf*.
- **IEEE 1609.2-Security:** The POC security layer is based on 1609.2 Version 1. Variations are captured in *SYS120-04 OBE-RSE IRS.pdf*.
- **IEEE 1609.4-Upper MAC:** The POC upper MAC layer is based on a variation of 1609.4/D07. Variations are captured in *SYS120-04 OBE-RSE IRS.pdf*.
- **IEEE 802.11p-Lower MAC and PHY:** The POC lower MAC and PHY layers are based on a variation of 802.11p/D1.0. Variations are captured in *SYS120-04 OBE-RSE IRS.pdf*.

In the RSE, the radio handler is the interface between the application (POC services, SAE J2735 message set) layer and the IEEE 1609.3 layer of the radio handler. The radio handler manages the registration of all services with the radio.

In the OBE, the communication manager (Comm Manager) is the interface between the application (SAE J2735 message set) layer and the IEEE 1609.3 layer of the radio. The Comm Manager manages the registration of all applications with the radio. It notifies applications when the services for which they have registered are available and passes data from the radio to the applications and vice versa. Figure 2 depicts the communications stacks for the POC RSE and OBE.

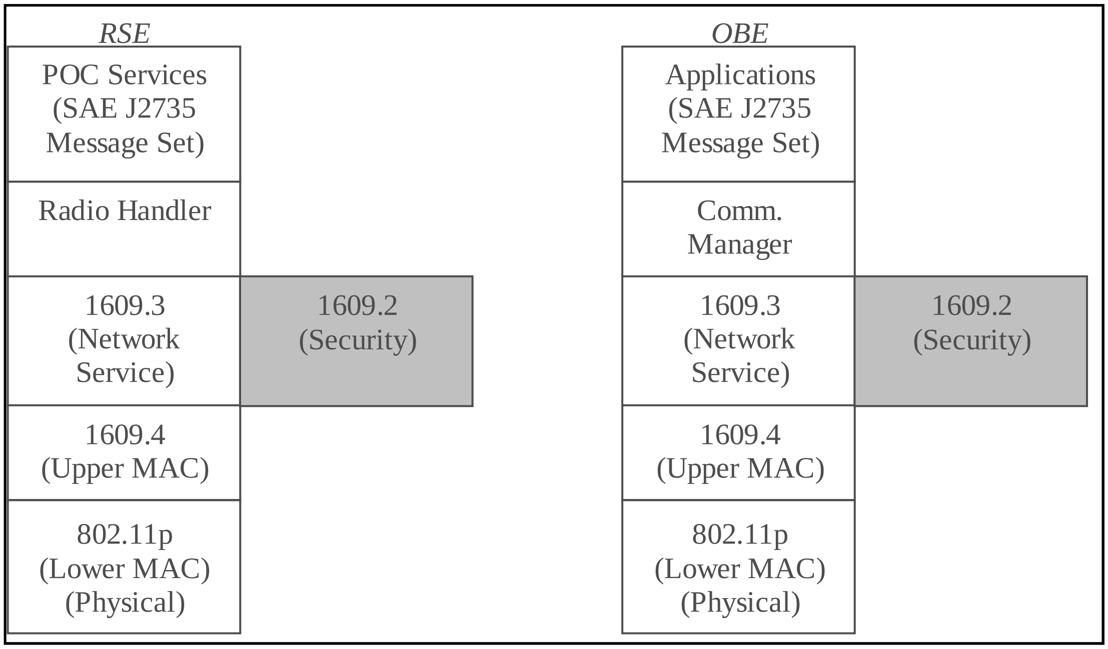

Figure 2: RSE and OBE Communication Stack

3.7. Testing Phases

The VII testing program was structured into three main phases—subsystem test, system integration and test, and public and private applications testing. These three phases followed a logical progression starting with testing the subsystems individually, with a gradual integration of each subsystem culminating in the testing of the entire infrastructure subsystem. Once subsystem testing was completed, the test program moved to integrating all of the VII subsystems and conducting an end-to-end system test. The testing phase ended with a series of tests exercising the public and private applications throughout the DTE using multiple vehicles.

Development of the RSE, Service Delivery Node (SDN), and Enterprise Network Operations Center (ENOC) subsystems was divided into three software builds. Each build provided additional functionality and increased integration and included PDS, AMDS, Positioning Service (POS), Map Element Distribution Service (MEDS), Network Management System (NMS), IdAM, Certificate Authority, and Network and Security testing. The goal of subsystem testing was to validate that the implemented subsystems satisfied the requirements in the SDN, ENOC, Certificate Authority, and RSE Subsystem Specifications (SSS), as well as the associated Interface Requirement Specifications (IRS). Verification of the SSS and IRS provided a base level of confidence that each subsystem was functional and operating correctly before attempting integration between the subsystems. SDN and ENOC subsystem testing was conducted at Booz Allen's Center for Network and Systems Innovation (CNSI) Laboratory in Herndon, VA. RSE subsystem testing was conducted at Raytheon Laboratory facilities in Orange County, CA, with network connectivity to the CNSI Lab.

After completion of subsystem testing, the VII POC moved into the system integration and test phase. Integration activities started with the deployment of production versions of two SDNs located in Herndon, VA, and Waterford, MI, and the production ENOC located in Herndon, VA. Following these activities, integration of RSEs began with RSE installations throughout the DTE as well as RSEs at Battelle's facilities to support OBE testing activities. The goal of the system integration and test phase was to verify the NSR and demonstrate that the end-to-end VII system was functional and ready to support user application tests. The majority of system integration and test activities occurred in the Michigan DTE, with early integration tests occurring at the CNSI Lab and Battelle and VIIC test track facilities. Early integration activities began in October 2007 and continued through November 2008.

The final phase of the VII test program was the public and private applications testing. These tests utilized applications developed by both the VIIC and Booz Allen, and included probe data collection, advisory message delivery, off-board navigation, electronic payments, and Signal Phase and Timing (SPAT)/Geometric Intersection Descriptions (GIDs) delivery. Tests were conducted using up to 25 vehicles equipped with OBEs and utilized each RSE installed throughout the DTE.

3.8. Analysis Phase

3.8.1. POC Scope Changes

The initial POC plan included extensive testing of the following public sector applications:

- ***Signal Timing Optimization –*** The goal of this application is to determine whether the VII system can collect data useful for optimizing traffic signals. The application uses collected probe data to analyze traffic flows in a particular area and then, based on that analysis, adjusts the signal timing. This process can be carried out on an ongoing basis to maintain optimal flow over time.
- ***Ramp Metering –*** The goal of this application is to determine whether the VII system can collect data useful for controlling freeway on-ramp metering lights to maximize the freeway throughput. This application is similar to signal timing optimization except the traffic situation is significantly different. One element to be measured is whether the PDS can provide data to accurately model the merging behavior of vehicles at the ramp region.
- ***Traveler Information*** – The goal of this application is to determine whether the VII system can deliver data that can then be used to generate useful advisory messages. This application uses the PDS to collect a variety of information and then, based on that information, generate advisory messages that are then distributed back out to the roadway environment using the AMDS subsystem.
- ***Weather Information –*** The goal of this application is to determine whether the VII system can deliver data that can then be used to generate a useful assessment of weather conditions in a given region. The results of this assessment can then be converted into weather advisory messages as appropriate, which can then be distributed to the appropriate roadway regions using the AMDS.
- ***Corridor Management Planning Assistance –*** The goal of this application is to determine whether the VII system can deliver travel pattern data that can then be used to support and improve road-planning activities. By understanding the general origin and destination information of travelers in a

region, it is expected that road planners can better determine how to manage the roads, as well as where and how to invest in road improvements.

- ***Corridor Management Load Balancing –*** The goal of this application is to determine whether the VII system can deliver road data that can then be used to support and improve transportation management activities. Corridor management load balancing is principally concerned with making the best use of the roadway infrastructure through the redistribution of traffic through advisory messages distributed using the AMDS.

The original objectives of the POC public sector applications were to: a) validate the technical viability of the VII system and b) provide limited demonstration of the potential utility of VII to determine various transportation management measures. All six public applications described above were developed based on detailed requirements specifications. However, the full-scale testing of these applications was de-scoped from the VII program. As a result, the analytical software (i.e., algorithms) necessary to assess the data and produce the transportation management measures was developed but not tested. For the POC, only a subset of the initially planned developments and subsequent tests were performed. These tests emphasized safety and generally used existing developed services and applications. The application and system service functions examined in the POC public applications testing are briefly described below. Section 4.0 includes a detailed discussion of the tests and the results.

- ***Probe Data Management –*** Examine the potential for enhancing intersection safety by directing vehicles to provide specific data at specific intervals for a given intersection.
- ***Advisory Broadcast Strategies –*** Provide in-vehicle warnings and advisories under a variety of traffic and other conditions to see how quickly the system can provide localized information directly to vehicles.
- ***Probe Data Reception –*** Determine whether probe data can be collected in near real time under a variety of traffic and other conditions, which could potentially help monitor and detect traffic and weather conditions that affect driver safety.
- ***Probe Data Vehicle Trajectories – Turning Movements –*** Determine whether probe data can be collected in near real time at intersections, which could potentially help monitor and detect traffic conditions (such as queue overflows and suboptimal signal phase and timing) at the intersection that adversely affect safety at that intersection.
- ***Probe Data Vehicle Stop and Start –*** Determine whether probe data can be collected in near real time in stop-and-go traffic, which could potentially help monitor and detect roadway congestion and, in turn, help motorists avoid the congested roadway segments and potential accident locations.
- ***Probe Data Vehicle Trajectories: Lane Change –*** Determine whether probe data can be collected in near real time from vehicles making unusual lane changes in the roadway, which could potentially help monitor and detect traffic conditions such as roadside hazard that adversely affect safety along the roadways.
- ***Probe Data Event-Generated Snapshots/Icy Road –*** Determine whether probe data, based on vehicle kinematics data (such as anti-lock braking system, traction control, etc.) can be collected in near real time, which could potentially help monitor and detect adverse weather conditions that affect driver safety.
- ***Trip Path –*** Collect trip path (O-D) data for future analysis.

As a result of the scope reduction, the following test cases, described in the test cases document *VII POC Applications Test Cases v0 5.doc,* were not carried out:

- TC-PDS-01
- TC-PDS-03
- TC-PDS-08
- TC-AMDS-01
- TC-TP-01
- TC-ILS-01.

3.8.2. General Deviations from the Test Cases

Any specific deviation from the test cases is noted in the appropriate analysis section. The following deviations apply to all test cases that were carried out.

At the start of testing, it was believed that a significant amount of probe data could not be transmitted from the OBE to the RSE if the vehicles traveled closely together, due to the limitations imposed by the Vehicular Datagram Transport Layer Security (V-DTLS) protocol. This belief was based on the Probe Data Gathering report presented by the VIIC in early August. Thus, the first 4 of the 6 days of PDS testing were conducted using vehicles dispersed (by 3 minutes). The fifth day was conducted with vehicles very closely dispersed (by 1 minutes), then batched in groups of five, and the last day with vehicles traveling as a platoon. This is in contrast to the original specification in “Test Cases and Conditions Matrices,” which called for 2 days each of dispersed, batched, and platoon testing.

Due to time and vehicle dispersion constraints, five new test routes were developed. The original test route was very similar to the new test route number four.

The weather for all days of PDS testing was clear and sunny, so none of the weather conditions could be recorded.

3.8.3. Reason for Exclusion of Some Data from Analysis

During the analysis phase, it became apparent that at times, certain OBEs, RSEs, and/or applications did not perform according to specifications. After re-checking the data and establishing that in fact there was a problem with the hardware unit or application, a note of the issue was made and the compromised data was not included in the analysis.

3.8.4. Data Sources Used For Analysis

Data analysis was conducted immediately following the conclusion of POC testing. The following data sources were analyzed:

- **OBE OSGi Logs** – Roush, the company that provided the drivers for testing, also provided the technicians that downloaded the OBE logs at the end of each testing day and uploaded those logs to the VIIC remote servers for further analysis.
- **RSE Proxy Manager (PM) Logs** – The RSE PM logs were turned on prior to the start of testing each morning and logs were downloaded off of the RSE at the end of each testing day. They were

uploaded to VIIC remote servers for further analysis. Due to storage limitations on the RSEs, the logs could not run continuously.

- **SDN Traffic Logs** – The SDN server had an instance of Wireshark, a network protocol analyzer, running throughout the POC testing that collected all incoming and outgoing network traffic.
- **AMDS Broker Database** – Database records were used to verify exact times of AMDS message transmissions for calculation of end-to-end latencies.
- **Public Apps Database** – Database records were used to verify reception of probe data snapshots.
- **Test Coordinator Notes** – The test coordinator kept daily logs, that, among other things, described:
    - Vehicle Issues
    - Driver Issues
    - Actual Test Times
    - Vehicle/Driver/Observer Batch Assignments
- **Test Observer Forms** – Seven test observers were assigned to test vehicles, and they kept a log that included:
    - Weather conditions
        - Sky cover
        - Pavement conditions
        - Air temperature
        - Precipitation type
        - Degree of precipitation
    - Traffic conditions
    - Unusual traffic events
        - Traffic backups
        - Emergency vehicles
        - Accidents
    - During AMDS testing
        - Message reception times and locations.

3.8.5. Data Analysis Tool Description

A Battelle analysis tool was used for the analysis of some of the test results. The tool logged data from the test vehicles, as well as infrastructure, and end-user applications, which were uploaded into a centralized SQL database for analysis. Within the database, customized C# programs as well as SQL stored procedures were used to parse and link information between sources. Once processed, the captured data was included in a Master Analysis SQL database table and made available for statistical analysis. Statistical analysis, including summaries, was conducted through the use of custom programs written in SAS© code. The analysis code included procedures for directly accessing the data in the SQL database and manipulating it, including merging it with other linked information so that summaries and custom tabulations could be developed.

The tool was tested prior to use. First, the code used for the public applications analysis was essentially the same as that used for the private applications testing. For the private applications testing, testers hand validated a sample of the results against the calculated numbers. Two different programmers (a primary and a reviewer) reviewed the code for accuracy and consistency.

There was an issue with heartbeat analysis because the OBEs were not in the correct mode during testing. The heartbeats sent and received could not be matched and therefore the developers constructed their own “pseudo matching ID.” However, it is possible that some of the heartbeats were incorrectly matched (or not matched). There is no way of validating this.

4. VII POC PROGRAM KEY FINDINGS

4.1. Dedicated Short-Range Communications

DSRC testing was performed during the system integration and test phase. The scope of this testing included demonstrating DSRC interoperability between an RSE radio built by one vendor and an OBE radio built by another vendor. This testing focused on fundamental DSRC standards and included Wireless Access in Vehicular Environments (WAVE) Service Advertisement (WSA) broadcasts, WAVE Short Message (WSM) Broadcasts, WAVE Basic Service Set (WBSS) generation, and Internet Protocol version 6 (IPv6) with both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) communications. Sirit performed the testing, using custom-made DSRC test software at its test facility in Carrollton, Texas.

Initial DSRC tests showed that the communications range between OBE and RSE was limited to between 3 and 15 meters. After several weeks of focused analysis and re-testing by TechnoCom, Denso, Transcore, and Sirit, several patches to both RSE and OBE radio software were developed, and the communications range limitation issue was resolved. Final range testing results showed solid radio communications from RSE to OBE up to 1,100 meters, with multipath effects degrading communications at 660 meters, 850 meters, 900 meters, and 1,000 meters. This communications range was much larger than initial projections. These results also showed a link imbalance with OBE to RSE communications only available up to 400 meters. Figure 3 provides the detailed results from this testing including the WSA packet error rate (PER) percentage.

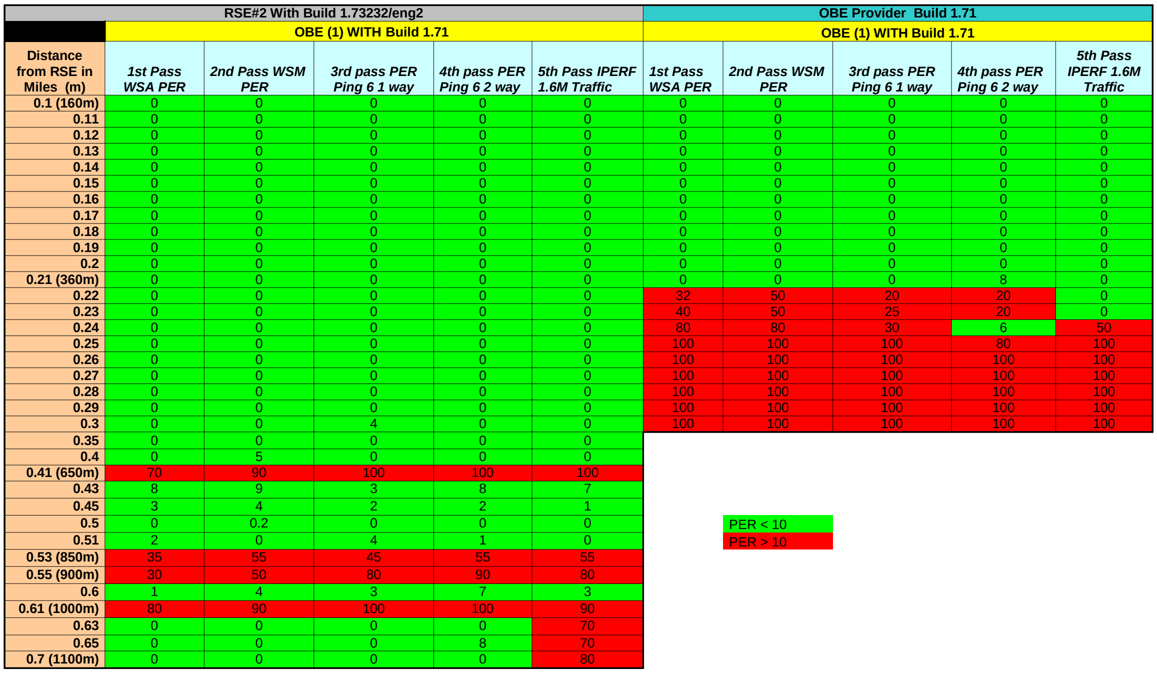

RSE#2 With Build 1.73232/eng2						OBE Provider Build 1.71				
	OBE (1) WITH Build 1.71					OBE (1) WITH Build 1.71				
Distance from RSE in Miles (m)	*1st Pass WSA PER*	*2nd Pass WSM PER*	*3rd pass PER Ping 6 1 way*	*4th pass PER Ping 6 2 way*	*5th Pass IPERF 1.6M Traffic*	*1st Pass WSA PER*	*2nd Pass WSM PER*	*3rd pass PER Ping 6 1 way*	*4th pass PER Ping 6 2 way*	*5th Pass IPERF 1.6M Traffic*
0.1 (160m)	0	0	0	0	0	0	0	0	0	0
0.11	0	0	0	0	0	0	0	0	0	0
0.12	0	0	0	0	0	0	0	0	0	0
0.13	0	0	0	0	0	0	0	0	0	0
0.14	0	0	0	0	0	0	0	0	0	0
0.15	0	0	0	0	0	0	0	0	0	0
0.16	0	0	0	0	0	0	0	0	0	0
0.17	0	0	0	0	0	0	0	0	0	0
0.18	0	0	0	0	0	0	0	0	0	0
0.19	0	0	0	0	0	0	0	0	0	0
0.2	0	0	0	0	0	0	0	0	0	0
0.21 (360m)	0	0	0	0	0	0	0	0	8	0
0.22	0	0	0	0	0	32	50	20	20	0
0.23	0	0	0	0	0	40	50	25	20	0
0.24	0	0	0	0	0	80	80	30	6	50
0.25	0	0	0	0	0	100	100	100	80	100
0.26	0	0	0	0	0	100	100	100	100	100
0.27	0	0	0	0	0	100	100	100	100	100
0.28	0	0	0	0	0	100	100	100	100	100
0.29	0	0	0	0	0	100	100	100	100	100
0.3	0	0	4	0	0	100	100	100	100	100
0.35	0	0	0	0	0					
0.4	0	5	0	0	0					
0.41 (650m)	70	90	100	100	100					
0.43	8	9	3	8	7					
0.45	3	4	2	2	1					
0.5	0	0.2	0	0	0					
0.51	2	0	4	1	0					
0.53 (850m)	35	55	45	55	55					
0.55 (900m)	30	50	80	90	80					
0.6	1	4	3	7	3					
0.61 (1000m)	80	90	100	100	90					
0.63	0	0	0	0	70					
0.65	0	0	0	8	70					
0.7 (1100m)	0	0	0	0	80					

🟩 PER < 10
🟥 PER > 10

Figure 3: DSRC Range Test Results

Additionally, both DSRC testing activities and testing of the overall POC system uncovered an issue with situations where multiple RSEs had overlapping coverage. This was an area not addressed in the DSRC standard, and resulted in issues with receiving AMDS messages at the OBE. Section 4.3 of this document includes the details related to AMDS testing. Currently, when in range of multiple RSEs, an OBE will create a WBSS with the first RSE it hears. It will only switch to another RSE if that RSE is advertising a service with a higher priority. If all other RSEs have lower-priority services, the OBE will not create a WBSS with any other RSEs, and will miss any service channel messages or services offered on those RSEs. The issue has been briefed to the DSRC community, and they are currently exploring solutions.

The POC RSE deployment locations were based on Noblis (Mitretek) studies commissioned by the US DOT to study RSE deployment scenarios. These studies indicated that coverage overlap will most likely be common in a nationwide deployment and was intentionally factored into the POC RSE location strategy. It was noted that with RSE coverage areas much larger than initial projections, there are many more instances of multiple RSE coverage overlap than anticipated.

Testing throughout the VII program has shown that both OBE and RSE are interoperable and fully capable of transmitting and receiving traffic in accordance with the DSRC standard. VII services and public and private applications have consistently utilized IPv6, TCP/UDP, and WSM message traffic throughout the DTE.

Table 3 summarizes the key findings in this area.

Table 3: Key Findings for DSRC

Finding Number	Finding
F-DSRC-1	Final range testing results showed solid radio communications from RSE to OBE up to 1100 meters, with multipath effects degrading communications at 660 meters, 850 meters, 900 meters, and 1,000 meters. These results also showed a link imbalance with OBE to RSE communications only available up to 400 meters.
F-DSRC-2	DSRC radio findings indicate that adequate range can be achieved in most conditions, although typical roadside link quality is significantly different from that measured in open field testing due to roadside furniture.
F-DSRC-3	Testing in situations where an OBE came within the overlapping range of multiple RSEs highlighted an issue involving prioritization and service selection. While the OBE successfully created a WBSS based on the highest priority service currently being advertised from the RSEs within range, it will not create a WBSS with any other RSE, and will miss any service channel messages or services offered on those RSEs.
F-DSRC-4	Testing throughout the VII program has shown that both OBE and RSE are interoperable and fully capable of transmitting and receiving traffic in accordance with the POC baselined DSRC standards. VII services and public and private applications have consistently utilized IPv6, TCP, UDP, and WSM message traffic throughout the DTE. With the notable exception of the OBEs "confirm before join" discrepancy, the OBE and RSE conform to the POC baselined DSRC standards. However, testing showed that the DSRC standards do not adequately address functionality for multiple overlapping RSE coverage areas.
F-DSRC-5	WAVE/DSRC-based security mechanisms generally work as intended and can effectively provide the authentication and authorization features required to meet stated privacy and security objectives.
F-DSRC-6	Communication quality was reduced by an "unbalanced link" situation whereby the OBE would commence transmission of data after coming in range of an RSE's broadcast, but at a distance too far for the RSE to receive the OBE's data.

4.2. Probe Data Service

PDS testing was conducted in all three phases of the VII test program. The scope of PDS testing was to show that OBEs could generate and send probe data messages to RSEs, and that Network Users could subscribe to all probe data elements or subsets of probe data elements and receive the data that accurately matched the subscriptions. Testing was accomplished at the CNSI Lab, the Battelle and VIIC test tracks, and the DTE. Both the VIIC and Booz Allen teams conducted the testing.

Initial PDS tests involving vehicles equipped with OBEs showed probe data loss rates greater than 60 percent. Subsequent analysis showed that the imbalance between the RSE and OBE radios resulted in the majority of the probe data loss. OBE-equipped vehicles would receive a PDS advertisement from an RSE at approximately 1,100 meters away and would immediately reply by sending the full buffer of probe data snapshots to the RSE. Testing showed that the RSE could not reliably receive those messages until the OBE was within 400 meters, resulting in the loss of most probe data messages sent between 400 and 1,100 meters away. Per the VII specifications, once the OBE sent the messages, the buffer was deleted and all probe data snapshots not received by the RSE were lost. Implementation of V-DTLS greatly improved the probe data loss rates by requiring the OBE and RSE to set up a V-DTLS connection before transmitting probe data.

Testing PDS with V-DTLS uncovered further key findings. V-DTLS as currently designed is a single thread process, allowing only one V-DTLS session per RSE at a time. This issue affects large numbers of vehicles trying to send probe data to the same RSE, forcing some vehicles to wait until the next RSE interaction to transmit their probe data. Probe data may be lost if a vehicle’s buffer reaches capacity prior to the next RSE interaction; however, that loss would still be less than the probe data loss without V-DTLS. Additionally, one to four probe data snapshots are lost per failed V-DTLS transaction; however, V-DTLS transaction failures were not prevalent in final DTE testing. The final finding was a small amount of data loss on the SDN to Network User interface. After troubleshooting this issue, the cause was found to be the need to send an Address Resolution Protocol (ARP) message when sending the first probe data snapshot to a Network User and the use of UDP as the protocol for sending probe data snapshots. After 30 seconds, the SDNs ARP cache clears and must resend an ARP, which would be sent at the same time as the first outgoing probe data snapshot. Since probe data is sent using UDP and then deleted from the SDN buffers, the snapshot is lost and no re-transmissions are possible. Subsequent messages sent within 30 seconds keep the ARP cache from clearing and prevent the loss of additional probe data. In test cases utilizing large numbers of vehicles driving through the DTE, ARP-related loss was minimal.

PDS testing showed that OBE-equipped vehicles can successfully generate and send probe data to RSEs and that Network Users can create, update, and delete PDS subscriptions and receive probe data snapshots in accordance with those subscriptions. Table 4 summarizes the key findings in this area.

Table 4: Key Findings for PDS

Finding Number	Finding
F-PDS-1	Initial PDS tests involving vehicles equipped with OBEs showed probe data loss rates greater than 60 percent. Implementation of V-DTLS greatly improved the probe data loss rates by requiring the OBE and RSE to set up a V-DTLS connection before transmitting probe data.
F-PDS-2	PDS testing showed that OBE-equipped vehicles can successfully generate and send probe data to RSEs and that Network Users can create, update, and delete PDS subscriptions and receive probe data snapshots in accordance with those subscriptions.
F-PDS-3	Testing of PDS during public applications showed that typically 80-to-90 percent of probe data sent during successful V-DTLS connections was received by Network Users. Additionally, the average end-to-end latency of probe data snapshots was found to be ranging from just over 0.5 second to almost 1.5 seconds.
F-PDS-4	A small amount of probe data is lost between the SDN and Network User due to the need for sending an ARP message when sending the first probe data snapshot to a Network User and the use of UDP as the protocol for sending probe data snapshots.

The following sections provide a detailed analysis of the results from final testing of PDS during the public applications integration testing conducted in Detroit, Michigan, between August 18, 2008, and September 12, 2008.

4.2.1. Probe Data Reception (Test Case PDS-02)

Test Case Objectives

- Test the application's capability to receive probe data under a wide variety of test conditions
- Test the accuracy and completeness of received probe data.

Original Data Outputs

- Comparison of probe data received by the application and probe data sent by the OBE
- Screenshots documenting verification of non-algorithmic requirements
- Comparison of accuracy of stop/start locations, stop delay, time and position trajectories, arterial travel time, speed, location, heading computed by applications with the same values, computed by OBE reference GPS
- Application-computed same probe segment number (PSN) vehicle traces for all vehicles
- Probe Data logs at one RSE.

Revised Data Outputs

- Comparison of probe data received by the application and probe data sent by the OBE
- Probe data logs at one RSE.

Original Exit Criteria

The test is passed if the probe data received by the application includes an appropriate percentage of the expected probe data based upon the OBE log. The percentage will be determined from OBE-to-RSE integration testing and dry run testing, and represents the percentage of snapshots that RSEs collect from vehicles under realistic traffic conditions, accounting for packet loss. If the percentage of loss data is greater than the percentage experienced in integration testing, then a map showing the locations of lost packets will be created.

The test will fail if the application receives data from RSEs to which it is not subscribed, or if it does not receive data from all RSEs to which it is subscribed, when vehicles are communicating with the RSE. The application operator will check several probe data subscription requirements. If any of these requirements fail, the test will fail.

The test also will fail if probe data position, speed, and heading for each vehicle are not recorded with an appropriate percentage of accuracy, as determined in laboratory position, speed, and heading testing. The test will also fail if the application fails to construct same-PSN vehicle trajectories correctly. The accuracy of trip paths in the merge area onto I-275 Northbound is of particular interest. Accuracy will be determined by comparing positioning logs to probe data.

Revised Exit Criteria

The test is passed if the probe data received by the application includes an appropriate percentage of the expected probe data based upon the OBE log. The percentage will be determined from OBE-to-RSE integration testing and dry run testing, and represents the percentage of snapshots that RSEs collect from vehicles under realistic traffic conditions, accounting for packet loss. If the percentage of loss data is greater than the percentage experienced in integration testing, then a map showing the locations of lost packets will be created.

The test will fail if the application receives data from RSEs to which it is not subscribed, or if it does not receive data from all RSEs to which it is subscribed, when vehicles are communicating with the RSE. The application operator will check several probe data subscription requirements. If any of these requirements fail, the test will fail.

Deviations from Test Case

Heavy traffic and precipitation were not encountered during the POC. Platoon formation was cancelled due to limitations of RSE-to-OBE communication caused by V-DTLS.

Hypotheses

1. A significant percentage of probe data snapshots that are securely transmitted by OBEs reach subscriptions.
2. Probe data snapshots take less than 1 minute to propagate from an OBE to a Network User under all conditions. During system integration and test, all probe data snapshots took less than 1 minute to reach the public applications. Based on that observation, it was hypothesized that it would take less than 1 minute for probe data under all conditions to propagate from an OBE to a Network User.

Data Analysis Strategy

The following original data outputs were removed:

1. Screenshots documenting verification of non-algorithmic requirements. *Reason: Non-algorithmic requirement verification was de-scoped.*

2. Comparison of accuracy of stop/start locations, stop delay, time and position trajectories, arterial travel time, speed, location, heading computed by applications with the same values, computed by OBE reference GPS. *Reason: No reference GPS information or other information for computing probe data accuracy was available.*
3. Application-computed same-PSN vehicle traces for all vehicles. *Reason: Algorithm testing was de-scoped.*

However, to verify that the test objectives were met, it was necessary to test the Public Apps' capability to receive probe data under a wide variety of test conditions and test the accuracy and completeness of received probe data. In addition, it was crucial to verify that probe data was delivered in a timely manner, so that the data could be used to detect real-time conditions.

Measure 1: Probe Data Loss

Measure Objective: Determine the average number of probe data snapshots lost when propagating from an OBE to a Network User and also the average loss between the OBE and RSE.

Compute separate figures for the following conditions:

1. Daily averages
2. Overall averages by hour 6 AM to midnight, for 15-minute intervals
3. Summary by day for weather conditions
4. Summary by vehicle configuration (dispersed, batched, platoon)
5. Summary when PDM active
6. By snapshot type (periodic, start/stop, event triggered) – summary value for all days
7. By RSE—to see if any RSEs had worse loss or latency
8. Freeway versus arterial (road type).

Measure 2: Probe Data Latency

Measure Objective: Determine the average time it takes for a probe data snapshot to propagate from an OBE to a Network User and also the average latency between the OBE and RSE.

Compute separate figures for the following conditions:

1. Daily averages
2. Overall averages by hour 6 AM to midnight, for 15-minute intervals
3. Summary by day for weather conditions
4. Summary by vehicle configuration (dispersed, batched, platoon)
5. Summary when PDM active
6. By snapshot type (periodic, start/stop, event triggered) – summary value for all days
7. By RSE—to see if any RSEs had worse loss or latency
8. Freeway versus arterial (road type).

Perform analyses for a sample of snapshots on a sample of vehicles under each condition.

Methods

During 6 days of probe data testing, have vehicles drive in various vehicle configurations on five separate PDS routes. The vehicles generated probe data, and the Public Apps subscribed to all probe data being generated.

The Google Earth screenshot in Figure 4 shows the five routes driven:

- Route 1 – Purple
- Route 2 – Blue
- Route 3 – Green
- Route 4 – Red
- Route 5 – Yellow.

Booz Allen generated routes on Google map

Figure 4: PDS-02 Routes

Table 5 shows the routes, vehicle configurations, and timing of each test run during the 6 days of PDS-02.

Table 5: PDS-02 Run Details

Date	Run	Route	Vehicle Configuration	Start	End	Loops
8/20	1	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	10:30am	12:00pm	2
8/20	2	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	12:52pm	2:22pm	2
8/20	3	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	2:46 pm	4:16pm	2
8/20	4	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	4:31pm	5:16pm	1
8/21	1	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	10:20am	11:50am	2
8/21	2	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	12:49pm	2:19pm	2
8/21	3	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	2:34pm	4:04pm	2

Date	Run	Route	Vehicle Configuration	Start	End	Loops
8/22	1	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	7:07am	8:28am	2
8/22	2	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	8:50am	10:05am	2
8/22	3	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	10:30am	11:49am	1
8/22	4	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	12:30am	1:53pm	2
8/25	1	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	1:50pm	3:15pm	2
8/25	2	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	3:30pm	4:40pm	2
8/25	3	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	5:10pm	6:40pm	1
8/25	4	5 vehicles on each of 5 routes	Dispersed: 3-minute spacing	7:30pm	9:00pm	2
8/26	1	5 vehicles on each of 5 routes	Dispersed: 1-minute spacing	9:55am	11:15am	2
8/26	2	5 vehicles on each of 5 routes	Dispersed: 1-minute spacing	11:15am	12:30pm	2
8/26	3	5 vehicles on each of 5 routes	Batched: 5 vehicles together	1:15pm	2:45pm	2
8/26	4	5 vehicles on each of 5 routes	Batched: 5 vehicles together	2:52pm	4:05pm	2
8/26	5	5 vehicles on each of 5 routes	Batched: 5 vehicles together	4:20pm	5:05pm	1
8/27	1	25 vehicles on route 5	Platoon—all together	10:00am	11:20am	2
8/27	2	25 vehicles on route 4	Platoon—all together	11:35pm	12:48pm	2
8/27	3	25 vehicles on route 3	Platoon—all together	1:59pm	2:47pm	1
8/27	4	25 vehicles on route 2	Platoon—all together	3:00pm	3:41pm	1
8/27	5	25 vehicles on route 1	Platoon—all together	4:00pm	4:52pm	1

Results and Analysis

The analysis was performed on all of the PDS-02 days detailed above, but also on the days of probe data management (AMDS-02), turning movements (PDS-04), start-and-stop traffic (PDS-05), lane change (PDS-06), and slippery conditions (PDS-07) testing. For details on those days of testing, see the report sections for each of those test cases.

Daily Averages – Loss

The average OBE-to-RSE reception rates ranged from nearly 60 percent to nearly 100 percent, with typical values in the 90-to-100-percent range. Figure 5 illustrates values for each day of testing.

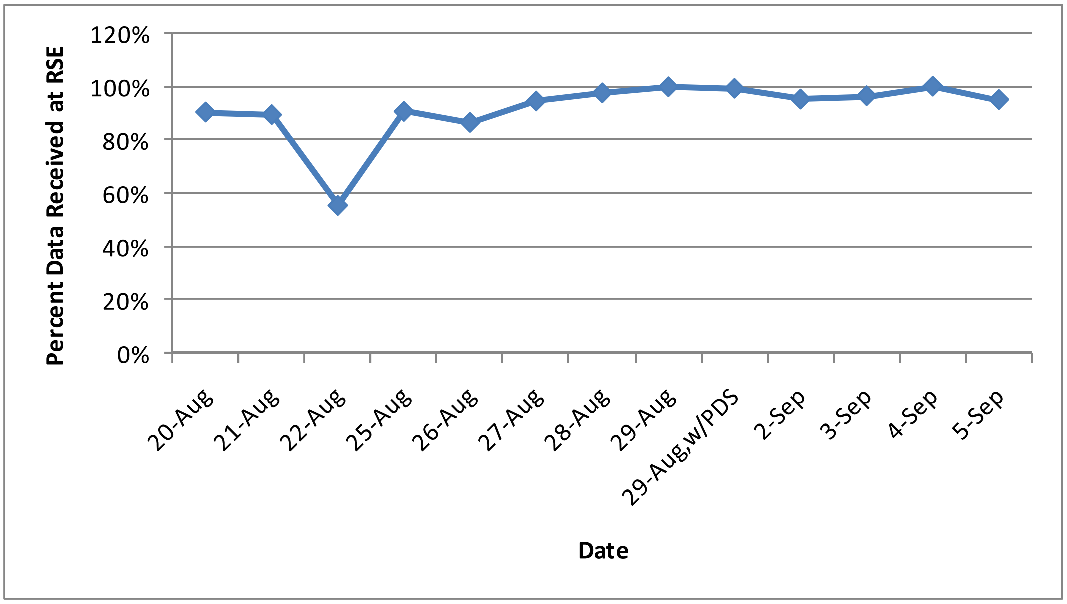

Figure 5: OBE-to-RSE Probe Loss by Day of Testing

On August 29, the PDS was not operational until 1 pm EDT. The data above shows the percentage of data received at RSEs before and after the PDS was restarted. It is apparent that the PDS has no significant effect on probe data loss between the OBE and RSE, as expected. It is not clear what caused the unusually low probe data success rate on August 22.

The average end-to-end probe data success rates—the percentage of probe data successfully transmitted from the OBE to the applications database—ranged from 60 percent to over 90 percent, with typical values in the 80-to-90-percent range. Days with low probe data success rates were due to problems with the PDS (e.g., PDS not operational, RSE not connected to PDS).

Figure 6 illustrates the average data transmission percentages by day.

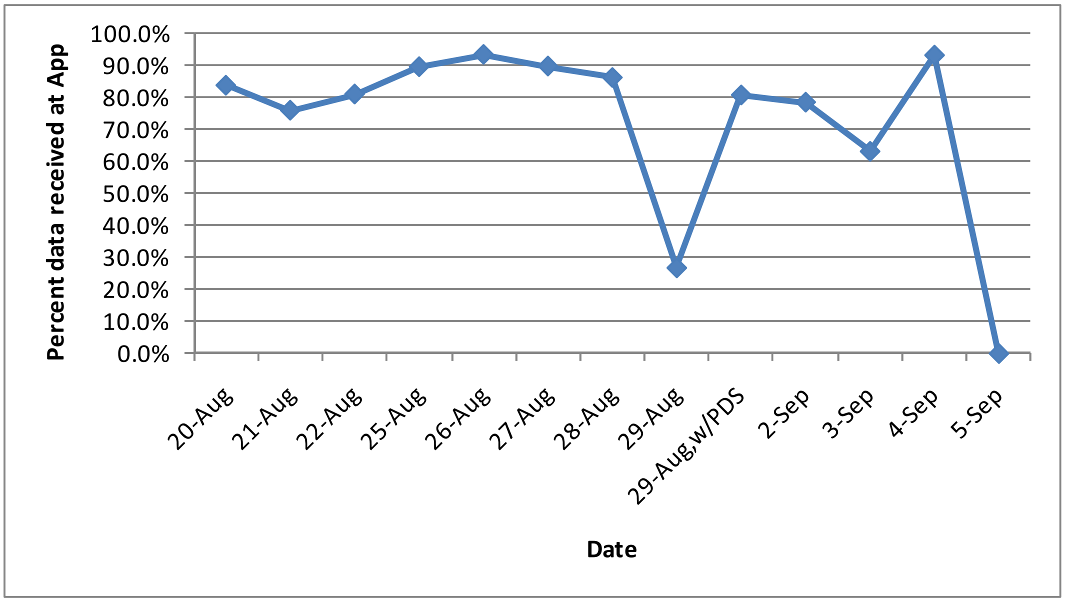

Figure 6: Average OBE – Applications Database Daily Transmission Percentages

As confirmed by the data from Aug. 29, the PDS has a significant effect on probe data loss between the OBE and Public Apps. The RSEs used for the slippery conditions test on September 5 was not correctly connected to the PDS, so the probe data success rate for that day shows as 0 percent because none of the probe data was captured in the Public Apps data. However, the probe data has been captured in a text format. Based on these observations, many of the results presented later in this section will not include these 2 days (August 29 and Sept. 5) in order to keep the PDS malfunctions from skewing the results.

An unusually low day of just 60-percent probe data success rate occurred on Sept. 3rd, the day of turning movements testing. This could be because of RSE issues which are explained in the **Analysis by RSE** section.

Values from just below 80 percent to just above 90 percent transmission were seen on all other days. The major sources of data loss were airlink loss, ARP loss, and the loss of broker connections at some RSEs.

It should be noted that on August 22, 25, and 26, the Public Apps database received more probe data than were recorded in the RSE logs. After looking at the details of this discrepancy it was concluded that missing RSE logs caused this to happen.

Daily Averages – Latency

The average end-to-end daily latency for probe data snapshots (OBE transmission to Public Apps database reception) ranged from just over 0.5 second to almost 1.5 seconds. It is not always clear what caused some days to have greater latency than others. However, all days fell within this range.

On September 5, routing of RSE data to the Public Apps database was down. The data was captured in text files. However that data was not uploaded or analyzed and therefore the loss and latency values for September 5 between the OBE and Public Apps could not be calculated.

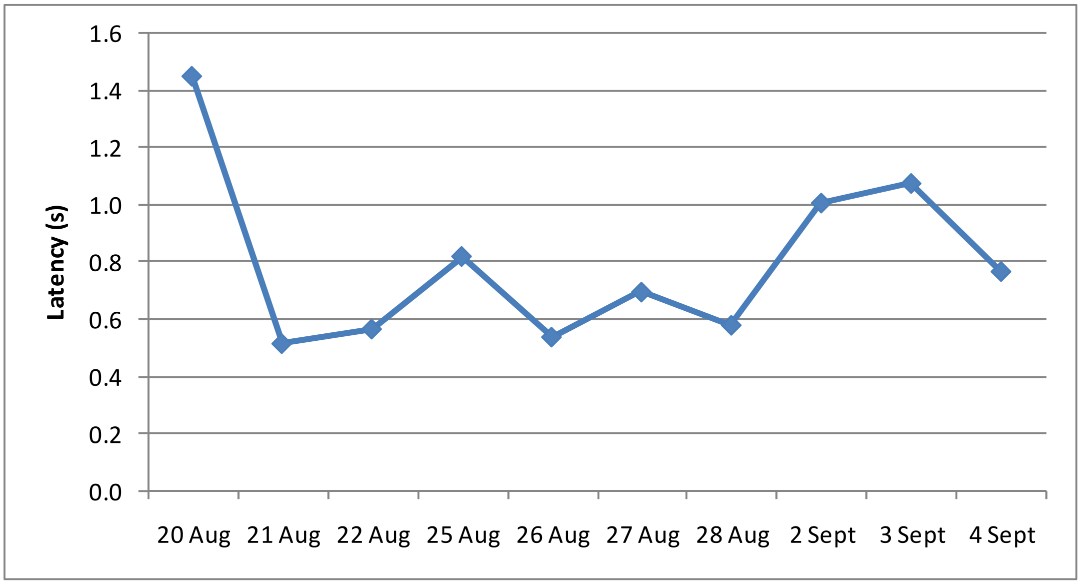

Figure 7: End-to-End Latency Excluding August 29

Overall Averages – Loss

Table 6 includes the percentage of data successfully received by the RSE and by the Public Apps using the total number of messages across all days. The aggregate probe data success rate to the RSE was about 87 percent. The aggregate probe data success rate to the Public Apps database was about 83 percent.

Table 6: Percentage of Probe Data Received by RSE and Public Apps

	% Received by RSE from OBE	% Received by Apps from OBE	% Received by Apps from RSE
Average all Received	87.8%	83.2%	94.8%

Table 7 displays the overall probe data loss values that can be computed.

Table 7: Overall Probe Data Loss

Avg OBE-RSE loss:	12.2%
Avg RSE-App loss:	12.4%
Avg non ARP RSE-App loss:	5.5%

The first value is 100 percent minus the value of RSE reception, 87.8 percent, indicating an average loss of 12.2 percent between the OBE and the RSE. The second value reflects the loss between the probe data received at the RSE and data received at the Public Apps (100 percent minus 87.6 percent). The final value is the loss rate adjusted for the 6.9-percent ARP loss. This rate was computed by counting one snapshot loss for every pause in data reception of more than 30 seconds in the Public Apps. It may be a slight underestimate of the ARP loss, since for pauses of over 1 minute, only one snapshot loss is counted, so a singleton snapshot every 30 seconds would not be counted in the estimate.

Loss between the OBE and the Public Apps database is due to some RSEs not being included in probe data subscriptions, and some RSEs not transmitting probe data, despite being subscribed to and having PDS connections. The tests were not re-run because September 12th was set as the absolute deadline by the US DOT for Public Apps testing. The following section will analyze this.

Overall Averages – Latency

Table 8: Average Latency for Probe Data

Average of all received:	0.81s

When performing normally, the average latency is about 0.8 second.

15-Minute Averages

Illustrating the probe data success rates and latency by 15-minute time segment will allow for the analysis of the effects of time of day. Note that the times presented in the charts in this section are the actual test times in Eastern Standard Time. Because only probe data snapshots that were successfully transmitted by the RSE were considered, it is not expected that there will be a clear relationship between the time of day and the probe loss or latency.

15-Minute Averages – Loss

Probe data success rate values for most 15-minute periods ranged from 40 percent to nearly 100 percent. A majority of the 15-minute periods had values in the 75-to-95-percent range. Early morning and late evenings had greater variability in data points, due to the fact that there were fewer days of testing, and hence fewer data points in these sections.

A few 15-minute periods had low probe data success rates. This can be explained by the fact that at these times of day, a disproportionate amount of testing was done on RSEs from which there was no or very little data transmission.

No clear pattern is visible in the data, suggesting that time of day and traffic level do not influence latency or loss. This result confirms the expectation explained above.

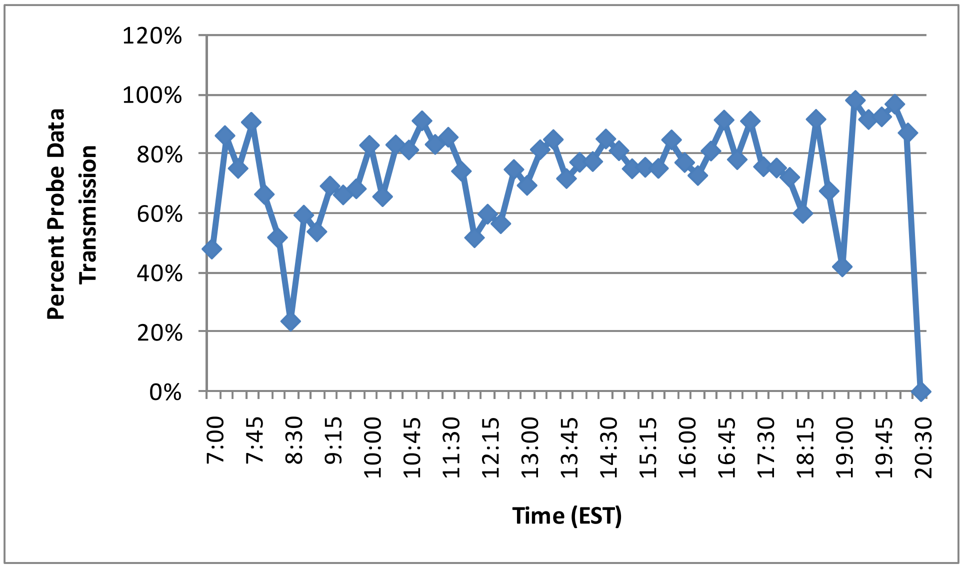

Figure 8: Probe Data Transmission

15-Minute Averages – Latency

Probe data latency values ranged from 0.4 second to just over 1 second, but a few outlier measurements occurred after 4 PM on a few days, during a few 15-minute periods, enough so that some 15-minute periods spiked to latency values of 3 and even 7 seconds.

There were fewer data points after 4 PM than before 4 PM, allowing outlier values to skew the results more heavily than at other times. Further analysis is needed to determine what caused these outlier values to occur. A few OBEs seemed to record the wrong timestamps on some occasions, so some of these outlier values may be caused by those OBEs.

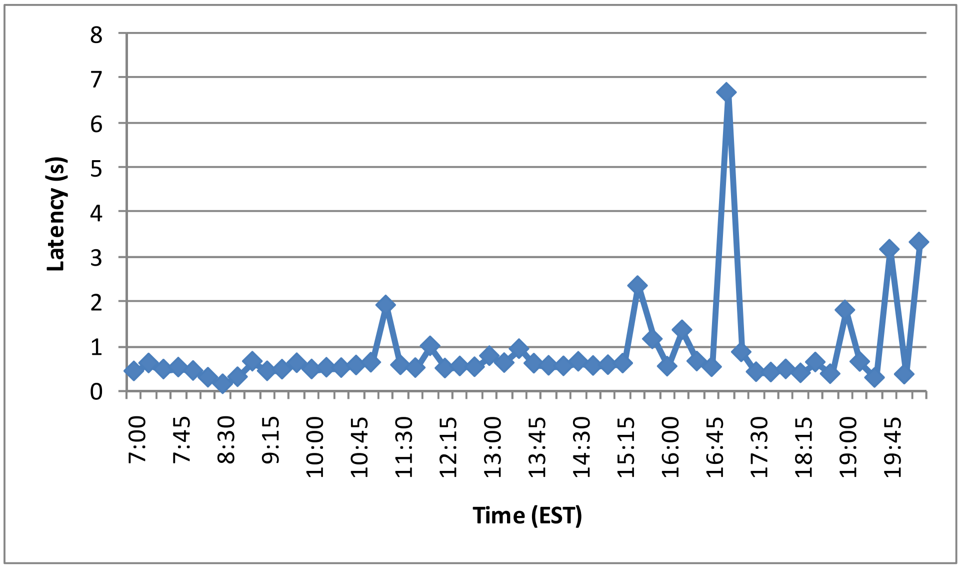

Figure 9: Probe Data Latency by 15-minute Period

RSE Averages – Loss

Analyzing probe data by RSE allows the determination of outlier or bad RSEs, which had unusually low probe data success rates. It also gives an indication of what typical airlink loss is.

The results show that many RSEs had no loss at all (100-percent probe data success rate), some had a 0-percent probe data success rate (mostly due to faulty broker connections), and some had probe data success rates ranging from 59 percent to 99 percent, due to a mixture of airlink loss and faulty broker connections or missing probe data subscriptions.

Figure 10 shows the RSEs with zero snapshots transmitted to them as 0 percent, and the remaining RSEs are shown with the actual transmission percentage. The remaining RSEs were fairly evenly distributed in the 5-percent bins between 75 percent and 99 percent.

There was significant variability in the probe data success rates by RSE. Almost half of the RSEs (22) had a 100-percent probe data success rate—no loss between OBE and RSE. Of the 52 RSEs, 10 had no snapshots transmitted to them, although of these 10, two were not yet active (I-96 and Kent Lake and M5&Farmington), and one was not on any route (9-Mile and Drake).

The remaining seven RSEs must be investigated to determine why they did not receive any snapshots. In some cases, the PDS broker connection was down during testing (Grand River and Novi, I-96 and Beck, perhaps others). In other cases, further investigation will be necessary.

For the remaining RSEs, which had probe data success rates between 65 percent and 99 percent, specific analysis will be needed to determine how much of the loss was due to pure airlink loss (loss in transmission between the OBE and the RSE), missing broker connections, or missing subscriptions.

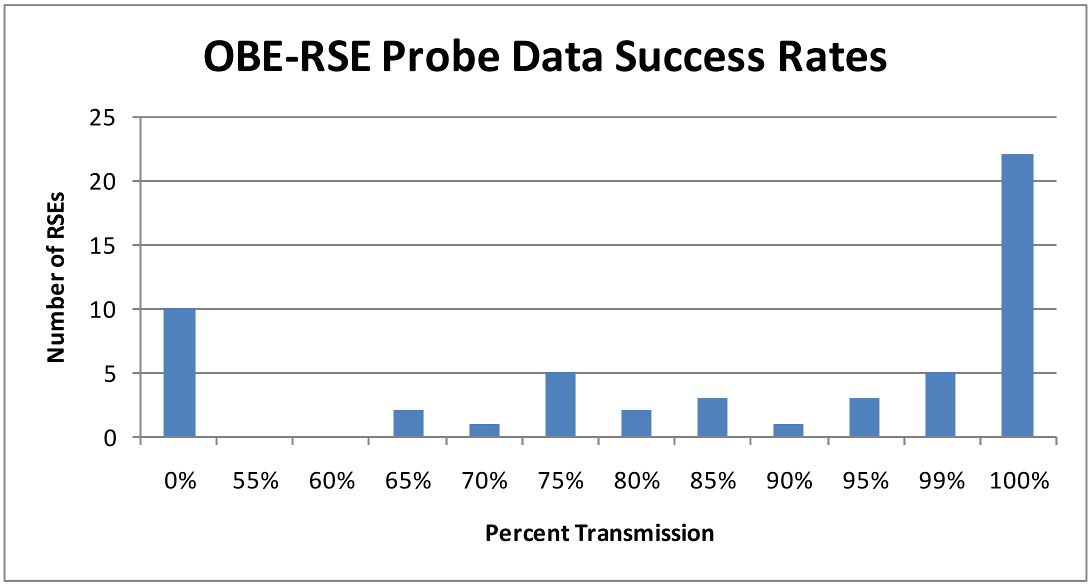

Figure 10: OBE-to-RSE Probe Data Success Rates

RSE Averages – Latency

Analyzing latency among RSEs will help to determine outlier RSEs and to characterize the effect an RSEs backhaul connection has on the latency of probe data transmitted to the RSE. Some distribution is expected because of the different latencies of the different backhaul networks, with hardwire connections expected to have the lowest latency, 3G cell phone connections expected to have the next lowest latency, and WiMAX connections expected to have the highest, although 3G and WiMAX latencies will depend on the geometry of the antenna for the given RSE.

Figure 11 shows probe data latency from the RSE to the Public Apps by RSE backhaul type.

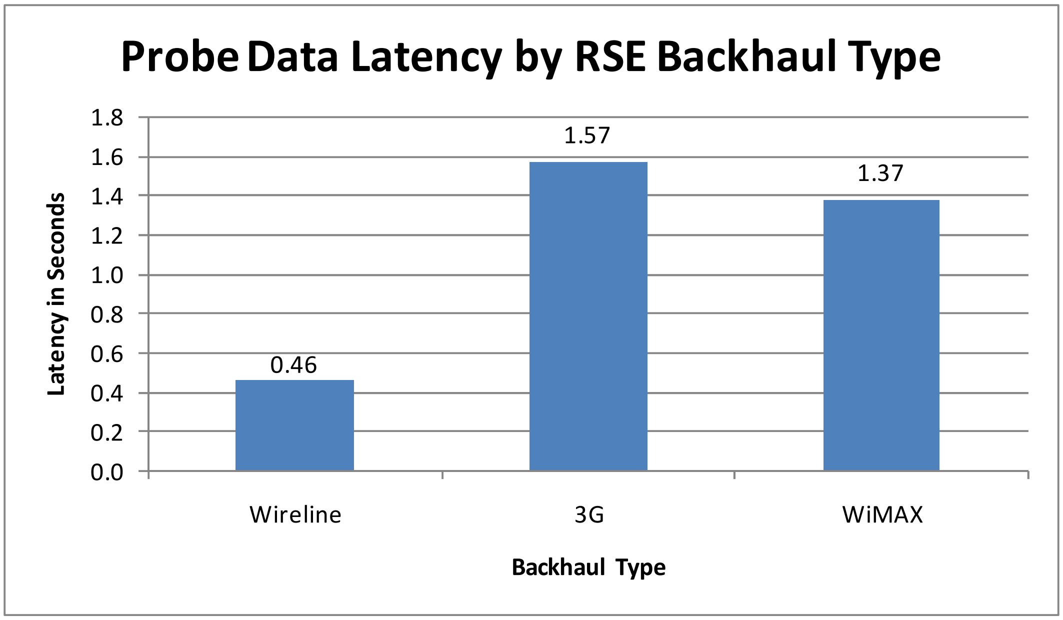

Figure 11: Probe Data Latency by RSE Backhaul Type

Figure 12 shows the distribution of OBE to Public Apps database probe data latency, by the RSE from which it is transmitted.

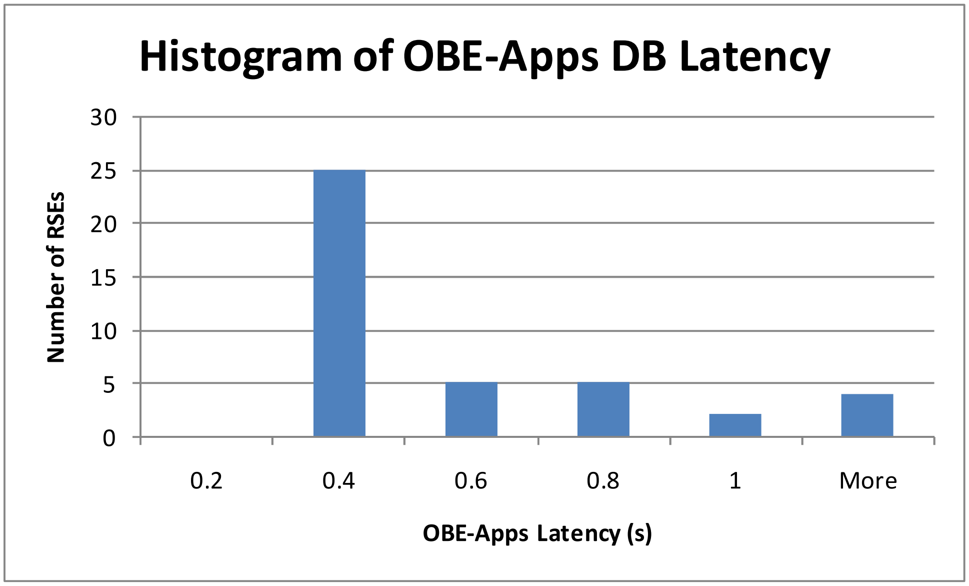

Figure 12: OBE-Application Database Latency

The results show that typical end-to-end latencies vary from 0.4 second to 1 second depending on the RSE. Unit testing indicated typical latency values of 1 second or less for most RSEs, and the Public Apps testing confirmed this value for larger number of vehicles. A majority of the RSEs had between 0.4 second and 0.6 second latency. Five RSEs had between 0.6 second and 0.8 second, and five RSEs had between 0.8 second and 1 second. Six RSEs had latencies higher than 1 second, and all six had either WiMAX or 3G backhaul connections. Such variation is explained by the difference in backhaul connections, above.

Freeway versus Arterial

This measure characterizes the difference between probe data transmitted to RSEs located on freeways and RSEs located on arterials.

Since snapshot transmission only takes place once a V-DTLS connection is established, no difference in performance was expected between snapshots sent to freeway RSEs and arterial RSEs. The loss rates were nearly identical. However, the latency rates were somewhat different. Latency rates for freeway RSEs were, on average, 0.1 second longer than the arterial RSEs. This difference is almost certainly due to the fact that freeway RSEs were exclusively WiMAX backhaul, while arterial RSEs were a mixture of T1, 3G, and WiMAX backhauls.

Freeway versus Arterial – Loss

Table 9: Probe Data Loss by Road Type

	Number of Snapshots	Number at RSE	% Received at RSE	Number at Apps	% Received at Apps
Freeway	10,428	9,048	86.7%	9,638	92.4%
Arterial	62,424	54,101	86.7%	54,920	88.0%

Freeway versus Arterial – Latency

Table 10: Probe Data Latency by Road Type

Latency	Average Latency
Freeway	0.81
Arterial	0.69

Snapshot Type

This measure characterizes how the three different types of snapshots performed with respect to probe loss and latency. The three snapshot types are periodic (created while vehicle is driving), start/stop (created when a vehicle first stops, regularly during stops, and then again when a vehicle starts), and event-triggered (created when there is an antilock brake system (ABS), traction control, stability control, or brake boost event).

Since all snapshots are transmitted the same way, over the same protocol, and the type is purely internal to the snapshot, no difference in performance between the snapshot types was expected. This is in fact what was observed, with both periodic and start/stop snapshots having about an 89-percent reception rate at the Public Apps database, and an average latency of about 0.78 second.

The values for event-triggered snapshots were somewhat different, but because the sample size was so much smaller—64 snapshots compared to over 22,000 start/stop and over 50,000 periodic—this difference is not statistically significant.

Snapshot Type – Loss

Table 11: Probe Data Loss by Snapshot Type

Snapshot type	Number of Snapshots	Number at RSE	% Received at RSE	Number at Apps	% Received at Apps
Periodic	50,355	43,765	86.9%	44,418	88.2%
Start/Stop	22,433	19,330	86.2%	20,079	89.5%
Event-triggered	64	57	89.1%	61	95.3%

Snapshot Type – Latency

Table 12: Probe Data Latency by Snapshot Type

Snapshot type	Latency
Periodic	0.84
Start/Stop	0.72
Event-triggered	0.46

Vehicle Configuration

This condition characterizes how vehicles were spaced apart, with respect to each other during testing.

Vehicle Configuration – Loss

While it was expected that vehicle configuration would have a significant effect on the probe data success rates due to V-DTLS limiting one vehicle at a time to connect to the RSE, it was not expected that vehicle configuration would have any significant impact on data loss, since a successful V-DTLS connection was a prerequisite for the analysis done here.

The testing proved that this was the case. All vehicle configurations resulted in a 92 percent +/-6 percent probe data success rate between the OBE and Public Apps.

Table 13: Probe Data Loss by Vehicle Configuration

Vehicle Configuration	Number of Snapshots	Number at RSE	% Received at RSE	Number at Apps	% Received at Apps
Dispersed – 3-minute spacing	28,179	21,974	78.0%	24,212	85.9%
Dispersed – 1-minute spacing	4,616	4,150	89.9%	4,529	98.1%
Batched	10,344	9,164	88.6%	9,950	96.2%
Platooned	13,328	12,763	95.8%	12,180	91.4%

Vehicle Configuration – Latency

It was expected that latency would not be affected by vehicle spacing, since snapshots are only counted towards latency if they are successfully transmitted to the RSE. The average latency of dispersed and batched vehicle configurations was 0.5 second +/-0.06 second, while the Platooned configuration had a somewhat greater latency. This is probably due to the fact that about half of the data from Platooned testing came from the day of start/stop testing, on which vehicles spent a disproportionate amount of time on freeways, and hence, transmitting to RSEs with WiMAX connections, which were known to have longer latencies.

Table 14: Probe Data Latency by Vehicle Configuration

Snapshot type	Average Latency(s)
Dispersed – 3-minute spacing	0.56
Dispersed – 1-minute spacing	0.47
Batched	0.55
Platooned	0.75

Weather

It was assumed that precipitation could negatively impact probe data success rates and probe data latency, since it sometimes interferes with communication channels, and that cloud cover could also impact probe data success rates. However, this data analysis could not support or deny these expectations, due to a lack of data (only 3 days were recorded with adverse weather conditions) and the inability to disambiguate the effects of adverse weather from other issues.

Most data was collected on sunny clear days. The test observer forms indicated that there were 3 days that were not sunny and clear. Table 15 presents data from these days.

Table 15: Probe Data Loss by Weather Type

Day	Weather Condition	OBE-Apps DB Loss	OBE-Apps DB Latency
22-Aug	Mostly Cloudy	82.1%	0.57
25-Aug	Partly Cloudy	90.7%	0.82
29-Aug	Rain	86.8%	23.6

The end-to-end probe data success rates ranged from 82 percent to 91 percent, which were typical values for probe data loss. August 22 experienced an average latency of .57 second, very typical values, and August 25 had a latency of 0.82 second, slightly longer than typical, but well within the range of standard values.

The value for latency on August 29 was 23.6 seconds. This latency is much longer than for all other days, and as discussed above, is likely related to the PDS issue earlier in the day. There is not enough evidence to attribute the unusual latency solely to rain.

Thus, this data analysis is insufficient to conclude whether or not adverse weather has a significant impact on probe data loss or latency.

Conclusions

The VII system implemented for POC works well as a probe data collection tool. Some individual cases need more detailed analysis, but some general conclusion can be reached.

OBE to RSE communications often have 100-percent success, but sometimes have values in the 90-to-100-percent range values due to airlink loss. Cases where lower probe data success rates were observed during testing require further analysis, but the lower success rates are likely due to factors such as lack of a subscription or PDS broker connection issues.

OBE to Public Apps probe data success rates were around 80 percent. The reasons for OBE to Public Apps loss include airlink loss, missing subscription loss, and PDS broker connection loss. In addition, the OBE to Public Apps database experienced loss due to the ARP issue. Because the average reception rate to the RSE and to the Public Apps database were very similar, the analysis did not indicate any other sources of loss.

The goal of this test was to compare the results, in this case 90-percent OBE to Public Apps success rate, to the results of the OBE-to-RSE integration testing, dry-run testing, and test track testing. It has been established that the current data is insufficient for a real comparison.

Battelle did some testing at the track and has some data; however, all of that data was collected without V-DTLS installed in the OBE and RSE. Once an OBE became available for integration testing in the DTE, it was plagued by an instability issue. The OBE would not stay on long enough for data collection. At the same time, V-DTLS was just released and its functionality was being tested at the RSEs. A revised OBE version became available in June of 2008, which caused simultaneous integration testing and OBE software build testing. However, testing was stopped by the US DOT after 1 week. During that week, for successful V-DTLS connections, the probe data success rate was about 95 percent. However, the success rate was extremely hard to determine, due to problems with the positioning system and the probe data application itself, which were later fixed before start of applications testing.

The average end-to-end latencies were 0.8 second, with most data latencies in the range of about 0.4 second to 1 second. A small number of longer latencies occurred. Further analysis is needed to determine what caused outliers, but their small number indicates that these are anomalous circumstances, possibly caused by incorrect OBE timestamps.

Both of the hypotheses that “A significant percentage of probe data snapshots that are securely transmitted by OBEs reach subscriptions” and “Probe data snapshots take less than 1 minute to propagate from an OBE to a network subscriber under all conditions” were validated by the results.

4.2.2. Probe Data Management (Test Case AMDS-02)

Probe Data Management (PDM) Messages Explained

PDM messages are used to enable network users to modify probe data snapshot generation characteristics of passing vehicles. This is useful for determining vehicle trajectories through an intersection, updating roadway maps, etc.

Sample Size

This allows the PDM message to apply its settings to a random sample of vehicles. This uses the last two digits of the current PSN to determine whether probe management is to be used. If the current PSN falls between these two values, then the PDM policy should be applied. The numbers are inclusive (e.g., using 41 and 43 would provide a 3-percent sample and 00.49 would be a 50-percent sample).

Start/Stop Strategy

This allows the user to modify the start and stop thresholds that determine when to generate a “stop” and/or “start” probe data snapshot. Vehicle stops and starts are treated as probe data trigger events since they could imply a traffic backup and indicate the travel time of a particular roadway link, among other things. “Start” is defined as when the vehicles speed exceeds a specified “start” threshold (default is 10 miles per hour, or mph), and a “stop” is defined as when the vehicle does not move forward for a specified “stop” time (default is 5 seconds) and no other stops have occurred within the “last stop” specified time (default is 15 seconds). The “last stop,” or stop lag, threshold is intended to prevent multiple snapshots when vehicles are creeping forward. No probe data snapshots are taken between a “stop” and “start” event.

Termination Method

Distance-based termination provides a distance-to-live type of time-out. It allows users to provide the distance driven until the probe management process ceases and the default condition is applied. The time-based termination method provides a time-to-live type of time-out. It allows users to provide the number of seconds at which time the probe management process ceases and the default condition is applied.

Collection Method

Distance-based collection allows Network Users to change the snapshot collection policy based on speed and distance. Two distances and two speeds are included in this data frame (D1, S1 and D2, S2) to be used by the OBE as follows:

- If speed is ≤ S1, then distance to next snapshot is D1 – default 20 mph (8.9 m/s) and 100m, used when distance values are in error.
- If speed is ≥ S2, then distance to next snapshot is D2 – default 60 mph (26.8 m/s) and 500m, used when distance values are in error.
- If speed is > S1 and < S2, then distance to snapshot is linearly interpolated between D1 and D2.
- If S1 is set to zero, then the distance to the next snapshot is always D1.

Time-based collection allows Network Users to change the snapshot collection policy based on elapsed time. Two times and two speeds are included in the message (T1, S1 and T2, S2) to be used by the OBE as follows:

- If speed is ≤ S1, then time to next snapshot is T1 – default 20 mph (8.9 m/s) and 6 seconds.
- If speed is ≥ S2, then time to next snapshot is T2 – default 60 mph (26.8 m/s) and 20 seconds.
- If speed is > S1 and < S2, then time to snapshot is linearly interpolated between T1 and T2.
- If S1 is set to zero, then the time to the next snapshot is always T1.

Test Case Objective

Demonstrate that VII can support the probe generation management capabilities specified in SAE J2735 for POC.

Original Data Outputs

- Comparison of probe data received by the application and probe data requested by the application.
- Screenshots documenting verification of non-algorithmic requirements.

Revised Data Outputs

- Comparison of probe data received by the application and probe data requested by the network user.

Exit Criteria

The test passes if the probe data received by the application corresponds to the probe data requested based upon the log in the application, and the OBE responds correctly to the PDM request parameters. The test will not fail due to packet losses within the range of those that can be expected to occur during real traffic conditions.

Data Analysis Strategy

Analyze OBE OSGi logs to see whether OBEs responded correctly to the PDM requests. Data from OBE logs should be analyzed because not all OBE log data is transmitted to the Public Apps due to V-DTLS and other connectivity issues. Probe generation variables altered in PDM requests can be grouped into five broad categories—periodic snapshot generation frequency, start/stop snapshot generation frequency (optional), PDM snapshot transmission frequency (optional), PDM snapshot sampling (optional), and PDM termination. A measure of the success of OBEs at implementing each of these categories was created.

Deviations from Test Case

Heavy traffic was not encountered during scheduled days of testing. RSE density in DTE did not change during testing.

Hypotheses

1. OBEs under PDM did generate periodic snapshots at the frequency specified in the PDM message.
2. PDM strategies expired on time.
3. Only vehicles with the correct PSN at the time of PDM reception responded to the message.
4. Stop snapshots were only generated when the vehicle speed was below the stop threshold.
5. Stop snapshots were generated past the specified lag time.
6. Start snapshots were only generated when the vehicle speed was above the start threshold.

Data Analysis Process Description

OBE logs were analyzed to determine the performance of PDM messages. The analysis included the following for each OBE log and each PDM request:

- Timestamp of PDM request reception, activation, and expiration
- Determination of whether probe data generation rates corresponded with rates specified in the PDM request
- *If start/stop generation rates were modified in the PDM request, determination of start/stop generation rates*: Did it correspond with request specified in the PDM request? (Note: This could only be tested if a stop occurred.)

- *If PDM PSN sampling was modified in the PDM request, determination of whether the appropriate vehicles corresponded to the PDM request*
- *If transmission rates were modified in the PDM request, determination of whether the transmission rate corresponded to the one specified in the PDM request*
- Total number of messages generated per vehicle per PDM request
- Total number of messages transmitted per vehicle per PDM request.

PDM Analysis Tool

A PDM analysis tool was written in Microsoft Excel's "Visual Basic for Applications" programming language. The purpose of the tool was to parse through an OBE log and extract relevant information for analysis. The following information was extracted for each OBE log for one particular PDM message:

- Total number of times the message became active
- Total snapshots in the OBE log
- Total number of snapshots created while the vehicle was under PDM that expired based on time or distance
- Listing of every instance of the PDM message's activation:
    - Activation and expiration times
    - Duration (calculated based on OBE log timestamps)
    - PSN at activation
    - Reason for message expiration (this is a line item in the OBE log)
    - Activation group ID (every instance with an activation and expiration time was assigned a sequential ID)
- Listing of every snapshot that was generated while the OBE was under PDM:
    - Entire content of snapshot
    - The following columns were added to aid in analysis
        - Snapshot time
        - Snapshot type
        - Latitude
        - Longitude
        - Current speed (mph) – this refers to the speed of the vehicle at the time of the snapshot
        - Time or distance from previous snapshot
        - Next snapshot in seconds or meters – this is calculated based on current speed and formula put forth in SAE J2735
        - Average speed between current snapshot and previous snapshot
        - Activation group ID
        - Real snapshot type – the snapshot type specified in the snapshot compared to the OBE log entry that says what type of snapshot was generated.

Strategies and Messages

For the 2 days of PDM testing, 12 distinct strategies were created (6 for each day of testing) and sent out using 33 unique messages. Table 16 shows the details of the 12 strategies; highlighted in yellow are the strategies that were analyzed.

Table 16: All Broadcast Strategies for PDM

		Time-based collection						Start/Stop is optional, but all parameters required if Start/Stop included		
Strategy	**Day**	**Speed 1 (mph)**	**Time 1 (s)**	**Speed 2 (mph)**	**Time 2 (s)**	**Term Time (s)**	**Sample Start / Stop**	**Start / Stop threshold (mph)**	**Stop / Lag Time (s)**	**TX Time Interval (optional)**
1	Th	10	7	40	7	75	0/99			
2	Th	20	20	40	20	75	0/99	5/3	1/10	
3	Th	15	4	30	10	30	15/99	5/5	2/10	
4	Th	10	6	27	20	60	0/99	2/2	1/5	
5	Th	1	4	20	20	45	0/99			5
6	Th	5	10	10	4	200 (m)	0/99			15
		Distance-based collection						**Start/Stop is optional, but all parameters required if Start/Stop included**		
		Speed 1 (mph)	**Dist 1 (m)**	**Speed 2 (mph)**	**Dist 2 (m)**	**Term Dist (m)**	**Sample Start / Stop**	**Start / Stop threshold (mph)**	**Stop / Lag Time (s)**	**TX Time Interval (optional)**
1	Fri	5	20	10	200	200	0/99			
2	Fri	1	4	20	80	1000	0/99	1/1	1/1	
3	Fri	5	40	20	120	500	50/99	5/5	5/1	
4	Fri	5	20	5	20	200	0/99	5/2	1/10	10
5	Fri	10	40	20	80	200	0/99			5
6	Fri	1	4	5	20	30(s)	0/99			

Routes and RSEs

Table 17 shows the analyzed messages that were sent to the RSEs.

Table 17: RSEs Used for Analysis of PDM

RSE Location	RSE Number	Backhaul
I-275_8 Mile Rd	us.mi.csnv.rse.0040	WiMAX
9 Mile Rd_Meadowbrook Rd	us.mi.csnv.rse.0058	3G
12 Mile Rd_Kendallwood Dr	us.mi.csnv.rse.0060	3G
Hills Tech Dr_Haggerty Rd	us.mi.csnv.rse.0080	Wireline

It is worth noting that RSE 0080 was down until 2:44 pm on August 28. The logs for that RSE were not available for analysis.

RSE 0060's internal clock was off during testing. On Thursday, August 28, 2008, it was set to 01-05-07. This should not have had any effect on its ability to pass along PDM messages to the radio handler.

Vehicles on all routes except route 5 should have received the messages that were analyzed. Figure 13 provides a screen shot from Google Maps, which shows all four routes and four RSEs.

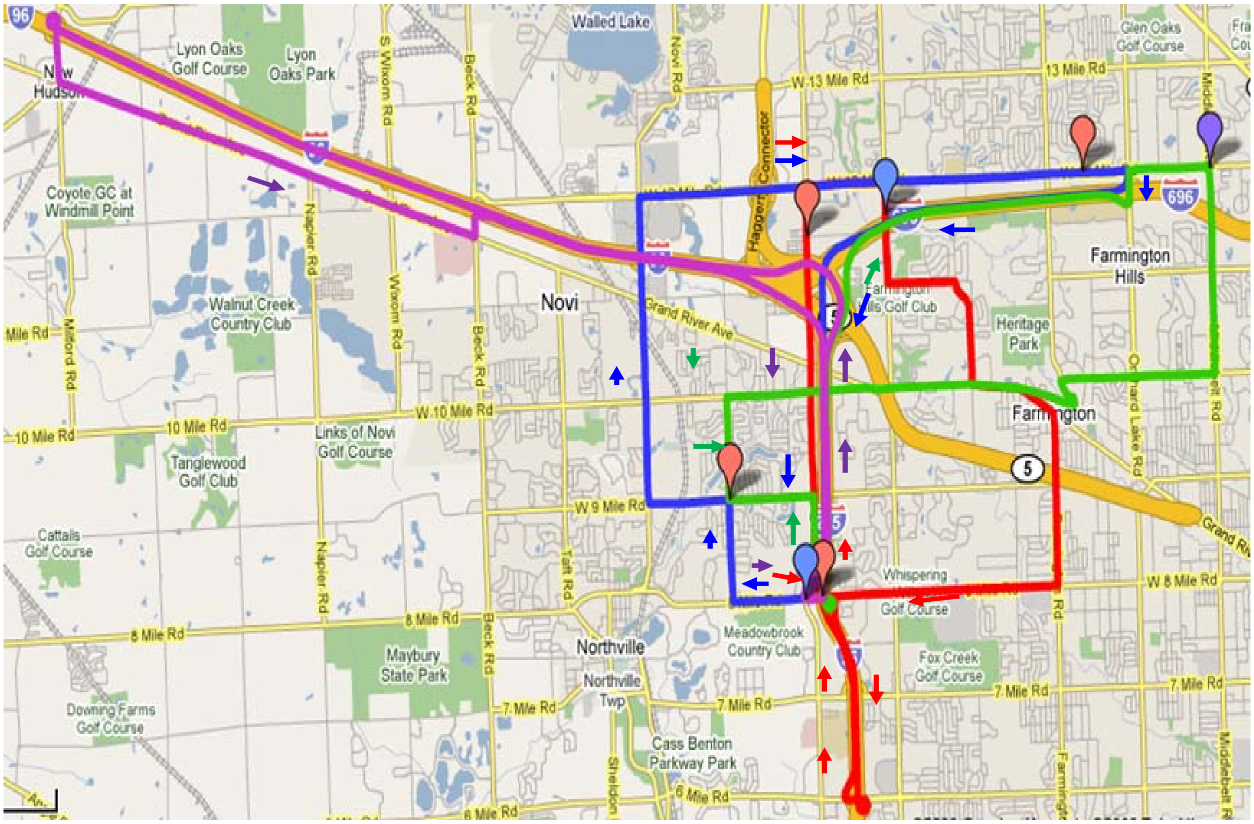

Booz Allen generated routes and plot points on Google map

Figure 13: Routes and RSEs Used for Analysis of PDM

Analysis for Thursday, August 28, 2008

PDM Message Selected

The PDM message analyzed for August 28 had the parameters shown in Table 18.

Table 18: Contents of PDM Message on August 28

Sample Start/End		Termination Method	
15	99	Time	30 s
Collection Method			
Speed 1	Speed 2	Time 1	Time 2
15 mph	30 mph	4 s	10 s
Start/Stop Strategy			
Stop Threshold	Stop Time	Stop Lag	Start Threshold
2 mph	10 s	10 s	5 mph
Message ID: 3973328C-4DB3-4704-925D-FC3D43DF78BD			

This message was selected for detailed analysis because it had the highest occurrence in the OBE logs. It appeared a total of 18,140 times across all OBEs. The message with the second highest occurrence in the OBE logs came in 8,003 times.

Methods

On Thursday, August 28, PDM messages with time-based collection methods were sent to the vehicles. The vehicles were sent out for a total of four runs, and all runs except run 4 had PDM messages broadcasting from the RSEs. Run 4 did not include PDM messages because very little data was received in the morning. Turning off the PDM message was used to determine its effect on probe data generation and/or transmission. The vehicles were divided into groups of five and sent on five routes. On runs 1 and 2, the vehicles were spaced 3 minutes apart, and there was no spacing during runs 3 and 4.

Table 19: Test Start and End Times for PDM on August 28

Run	Start	End	Total Loops
Run 1	10:00 am	11:20 am	2
Run 2	12:15 pm	1:32 pm	2
Run 3	2:00 pm	3:19 pm	2
Run 4	3:40 pm	5:05 pm	2, except vehicles on route 4 did only 1 loop

Note about OBE Logs

For Thursday, August 28, data was received from 21 vehicles that had received the PDM message. Only 20 vehicles (4 routes with 5 vehicles per route) should have responded, but as seen on the Google Maps screenshot in Figure 14, vehicles on route 5 came very close to the I-275 and 8 Mile RSE.

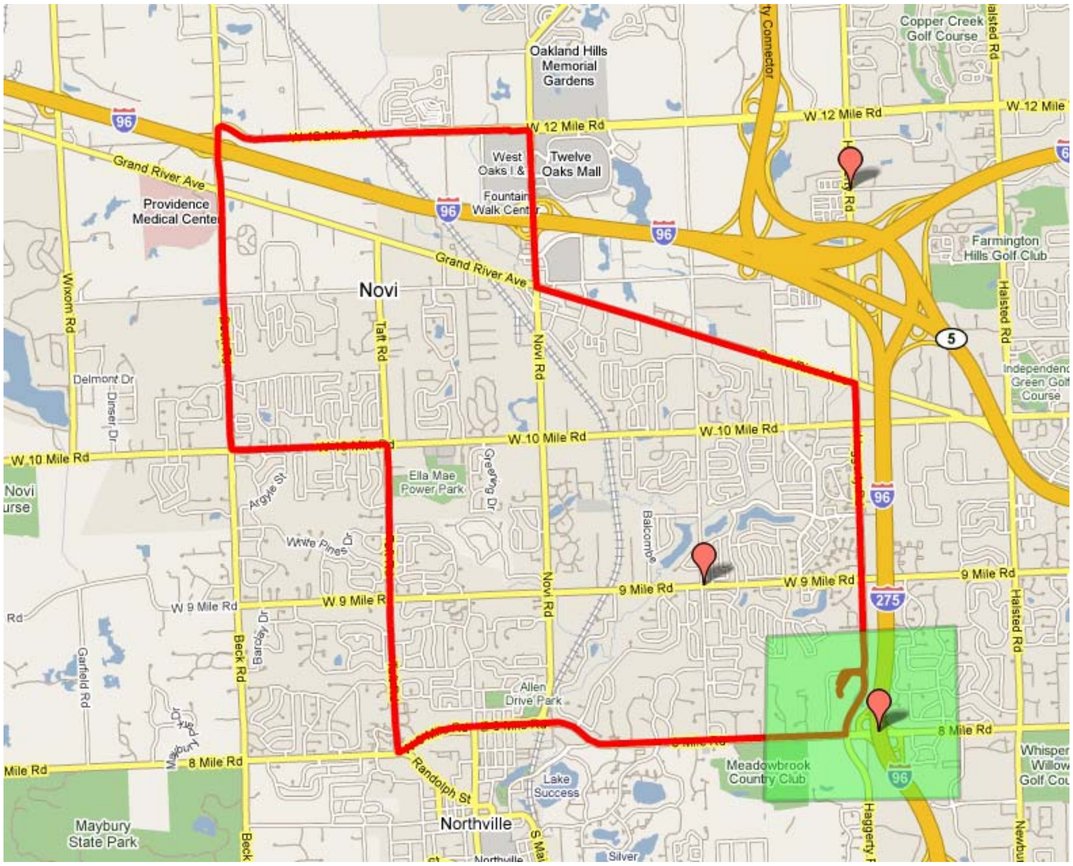

Booz Allen generated routes and plot points on Google map

Figure 14: Route 5 to RSE Proximity

OBE B856 was disqualified from inclusion in this analysis due to apparent positioning system errors.

Table 20 presents the issues that were experienced with other vehicles during testing.

Table 20: Vehicle Issues on August 28

C1 (C453)	Lost GPS comm's 1 pulse per second (pp) active before session 2
E4 (B325)	No Positioning, 1pp active after session 1
P1 (C482)	Lost GPS comm's 1pp active after session 1
P10 (C590)	Probe Generation Error after Session 1
P12 (C366)	Probe Generation Error after session 3
P12 (C366)	Screen locked up, OBE dead after session 2
P8 (C548)	Screen locked up, OBE dead after session 2
P9 (C082)	No vehicle speed (CAN) data, after session 2

Measure 1: Percentage of Periodic Snapshots within 4 Seconds of Intended Time

Measure Objective: Determine whether OBEs under PDM generated periodic probe data at the frequency specified in the PDM message.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For the RSEs broadcasting PDM messages with a fixed snapshot generation frequency (e.g., 4 seconds), verify in the OBE logs that snapshots were generated at that frequency.

Results: For the message in question, the OBEs were to generate a periodic snapshot every 4 seconds if the speed was less than 15 mph. They were also to generate a periodic snapshot every 10 seconds if the speed was more than 30 mph. And for speeds between 15 and 30 mph, the snapshot time was to be linearly interpolated.

Figure 15 shows the number of seconds after a snapshot's intended generation time that it was actually generated.

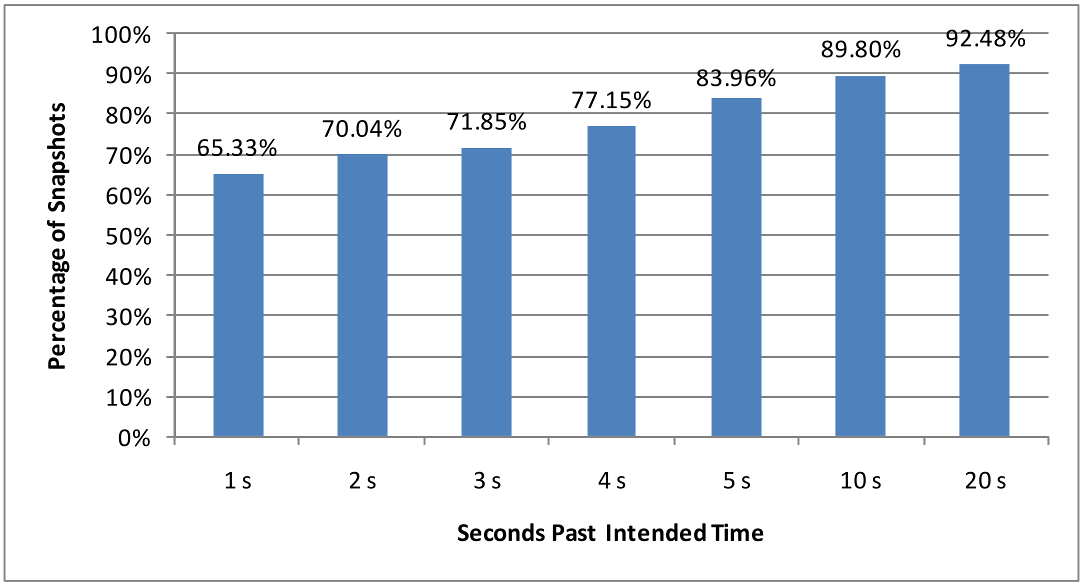

Figure 15: Cumulative Distribution of Snapshots – PDM August 28

No snapshots were generated before the intended time, over 77 percent of snapshots were generated within 4 seconds of the intended time, and almost all snapshots were generated within 20 seconds.

Analysis: The results show that almost 77 percent of snapshots are generated within 4 seconds of their intended time. This means that if an operator requests snapshots to be generated every 10 seconds, they can expect the majority of them to be generated between 10 and 14 seconds. Having the ability to alter periodic snapshot generation will indeed be useful. However, it would be most useful if the operators were not limited by OBE shortcomings. It is recommended that, in the future, OBE implementations support faster snapshot generation and by default store more snapshots (recommendation of 300) in memory.

Measure 2: Standard Deviation of Differences between Intended Time and Actual Time

Measure Objective: Determine whether OBEs under PDM generated periodic probe data within a reasonable standard deviation of the anticipated time.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed generation frequency (e.g. 4 seconds), verify in OBE logs that messages were generated within a reasonable standard deviation of the anticipated time.

Results: Figure 16 shows the standard deviation between anticipated snapshot generation times and actual snapshot generation times.

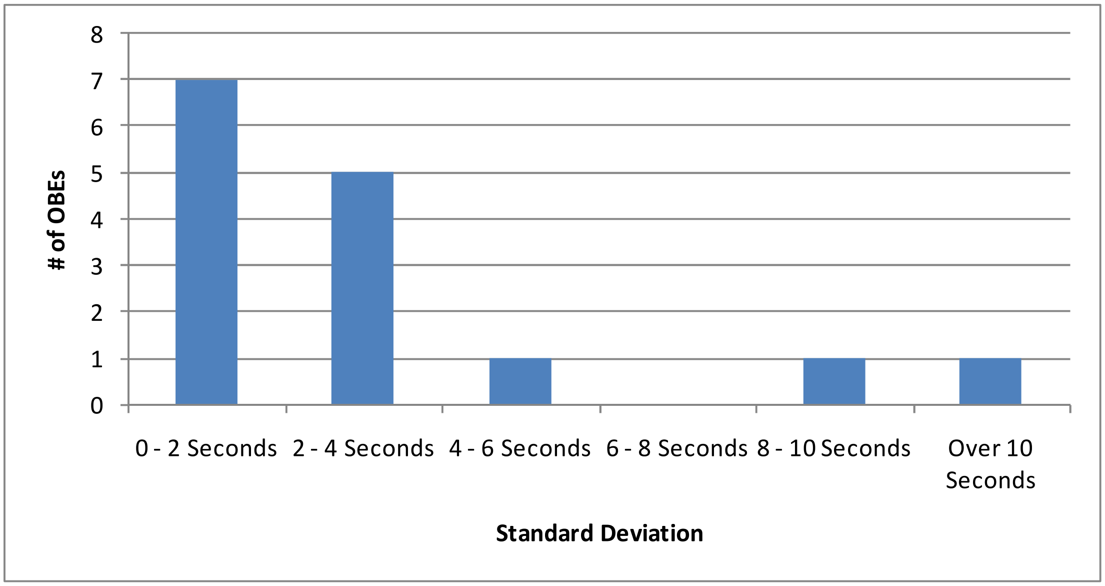

Figure 16: Standard Deviation from Anticipated Time – PDM August 28

A majority of OBEs generated periodic snapshots within a standard deviation of 4 seconds from their intended time.

Measure 3: Percentage of Time-Based Expirations – Percent of all PDM Activations for a Given Serial Number that Expired Based on Time

Measure Objective: Determine the percentage of PDM activations that expired based on time.

Concept: Every PDM message has either an expiration time or expiration distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: In the OBE OSGi logs, divide the total number of time-based expirations of a particular serial number by the total times that serial number became active.

Results: Time-based expiration is defined as the OBE enabling a PDM message until its intended expiration time.

Time-based expirations are affected by overlapping RSEs and close-by RSEs. If the OBE receives another unique message, the newly received message becomes active and the first PDM message expires.

The following screen shot from Google Maps (Figure 17) shows the routes and RSEs (in red) used for the PDM message in question as well as close-by RSEs (in blue) that were also broadcasting PDM messages.

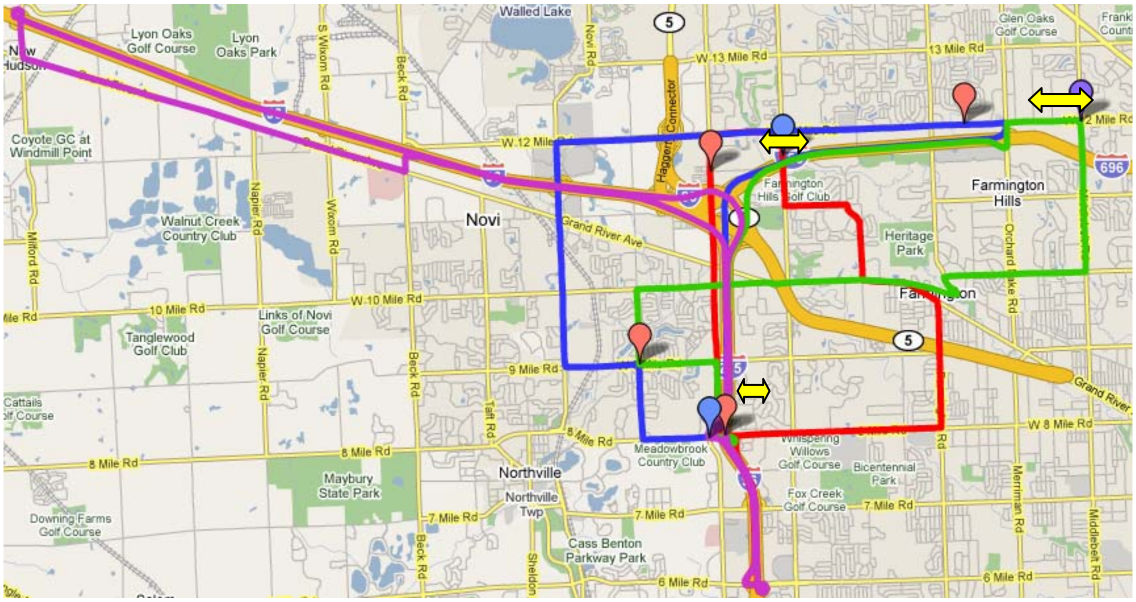

Booz Allen generated routes and plot points on Google map

Figure 17: Nearby RSE Affecting Time-Based Expiration – PDM August 28

The three yellow arrows in decreasing length are 1.6 miles, 1.06 miles, and 0.23 miles. In at least the 0.23 mile case, it is very likely that the RSE footprints were overlapping and therefore decreasing the percentage of time-based expirations.

Figure 18 shows the time-based expiration of PDM messages by OBE.

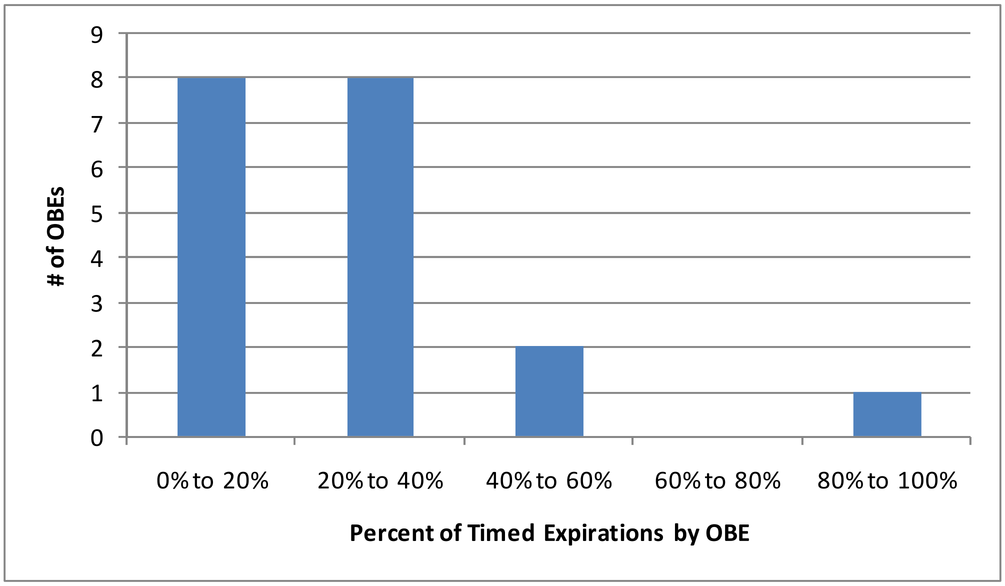

Figure 18: Percentage of Timed Expiration by OBE – PDM August 28

Analysis: On average, only less than 40 percent of PDM messages expired based on time, and this can be attributed to overlapping RSEs that were also broadcasting PDM messages.

Measure 4: Percentage of Time-Based PDM Expirations that Expired Correctly

Measure Objective: For PDM messages that expired based on time, determine whether their expiration was at exactly 30 seconds.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in OBE logs that each time-based expiration met the criteria defined in the PDM message.

Results: The results showed that if a message expired based on time, it did so correctly 100 percent of the time.

Expiration time allows the TMC operator to capture more or less probe data for a certain period. The results show that the time specified in the PDM message will be executed correctly.

Analysis: 100 percent of messages that were not overridden by new PDM messages expired at the correct time.

Measure 5: Percentage of Correctly Sampled PDM Messages (i.e., percentage of PDM that had a PSN within the allowed range when they became active)

Measure Objective: Determine whether only the vehicles with the appropriate PSN at the time of PDM reception responded to the message.

Concept: Sampling is used to limit the number of vehicles that respond to a particular PDM request and thereby controls the total number of snapshots that the TMC operator receives. Sampling makes senses when a majority of vehicles on the roadway are VII-enabled, because if every vehicle responded to a PDM request, the ensuing onslaught of snapshots from thousands of vehicles could overwhelm the infrastructure.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For the sample specified in the PDM message (15-99), verify that only OBEs with PSNs that ended in that range of numbers (15-99) responded to the message.

Results: Data from 20 OBEs for the PDM message in question were analyzed. The results showed that the sample specified in the PDM message was correctly met 100 percent of the time.

Analysis: Sampling is a very useful tool for the TMC operator to limit the number of snapshots passing through the system. When combined with alteration of the probe data generation scheme, it can achieve its goals of managing probe data traversing the system.

Measure 6: Percentage of Stop Snapshots with Current Speed below Stop Threshold

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at speeds below the stop threshold specified in the PDM message.

Concept: Stop threshold is used to indicate the minimum speed that indicates vehicle movement. For example, on highways with a speed limit of 65 mph, a vehicle traveling at less than 10 mph may be considered “stopped.” The vehicle’s instantaneous speed is used to determine snapshot type. The resulting stop snapshots can be used to detect traffic congestion, accidents, and other unplanned events.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each stop snapshot generated, verify that the snapshot speed was less than the stop threshold speed.

Results: All OBEs met the goal of generating a stop snapshot when the vehicle speed was below the stop threshold speed 99.99 percent of the time.

Figure 19 shows all stop snapshots and whether or not they met the stop threshold.

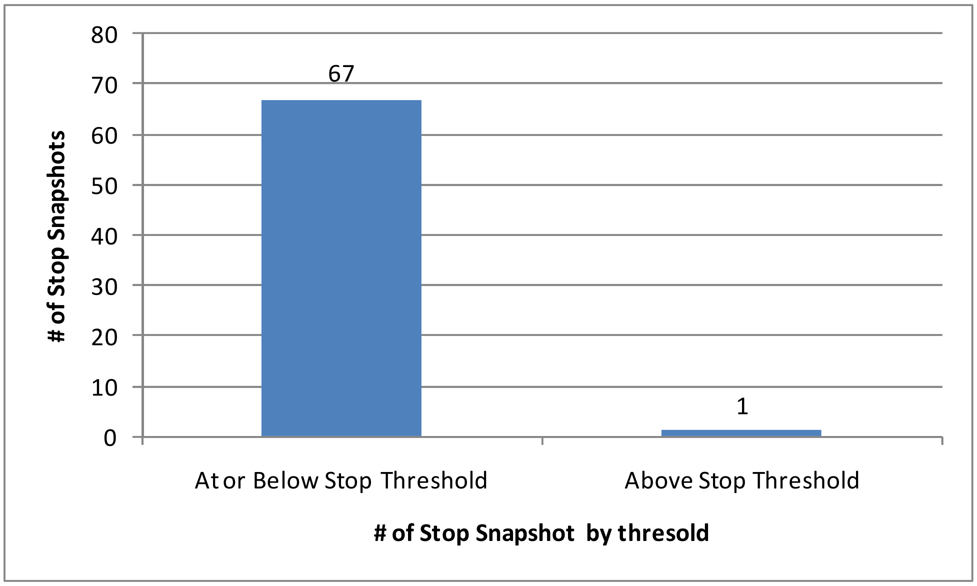

Figure 19: Number of Stop Snapshots by Threshold – PDM August 28

Most stop snapshot speeds were at 0 mph. Since stop time was 10 seconds, it can be concluded that most snapshots were generated at red lights, traffic delays, and/or highway exit ramps.

Figure 20 is a Google Earth representation of some of the stop snapshots that were generated. As the vehicles approach the traffic light at 12 Mile and Kendallwood Drive or move away from it, the speeds are at or slightly above zero.

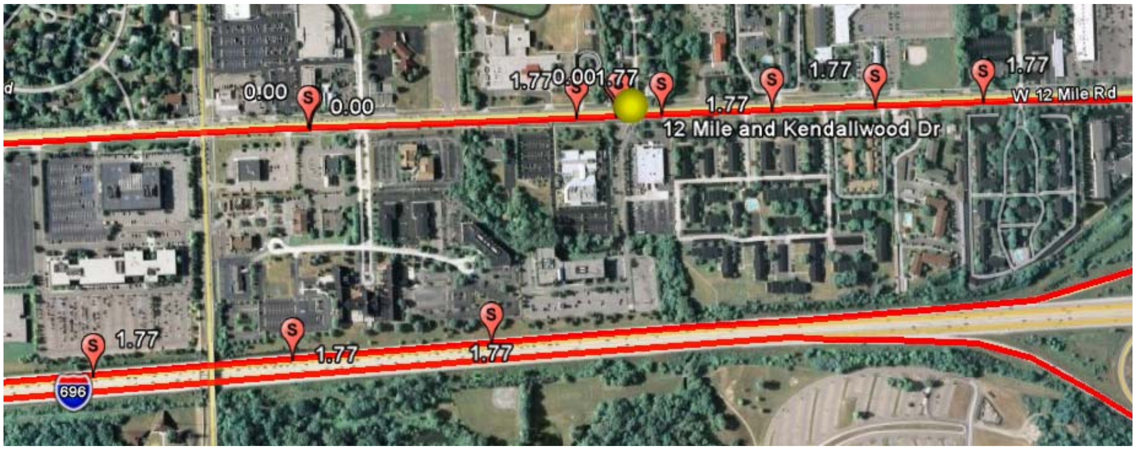

Booz Allen generated routes and plot points on Google map

Figure 20: Stop Snapshots on Google Earth (1) – PDM August 28

The majority of the rest of the stop snapshots were generated by vehicles starting out the run around I-275 and 8 Mile, an area prone to traffic delays.

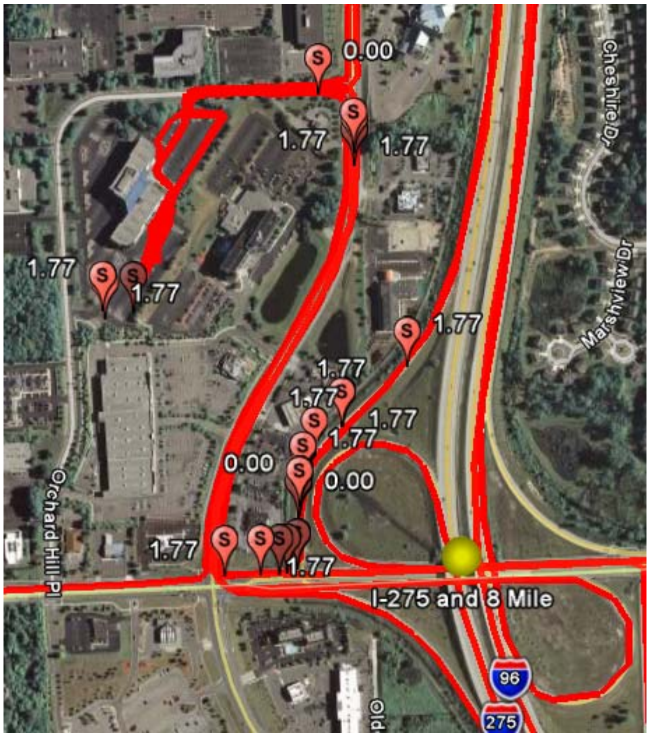

Booz Allen generated routes and plot points on Google map

Figure 21: Stop Snapshots on Google Earth (2) – PDM August 28

Analysis: The results show that stop thresholds were implemented correctly and OBEs responded with stop snapshots. However, stop snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear that such units are commercially available at a reasonable cost.

Measure 7: Percentage of Stop Snapshots Generated after Stop Lag

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at the frequency (stop lag) specified in the PDM message.

Concept: Stop lag is used to prevent the OBE from generating too many stop snapshots when the vehicle is creeping along at speeds below the stop threshold. Stop lag prevents the OBE buffer from filling up with nearly identical stop snapshots and overwriting periodic and/or event-generated snapshots.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate the difference between the current and previous stop snapshots and verify that it is after the specified lag time.

Results: 100 percent of the stop snapshots exceeded the 10-second stop lag time. Figure 22 shows how many seconds after the stop lag time the next stop snapshot was generated.

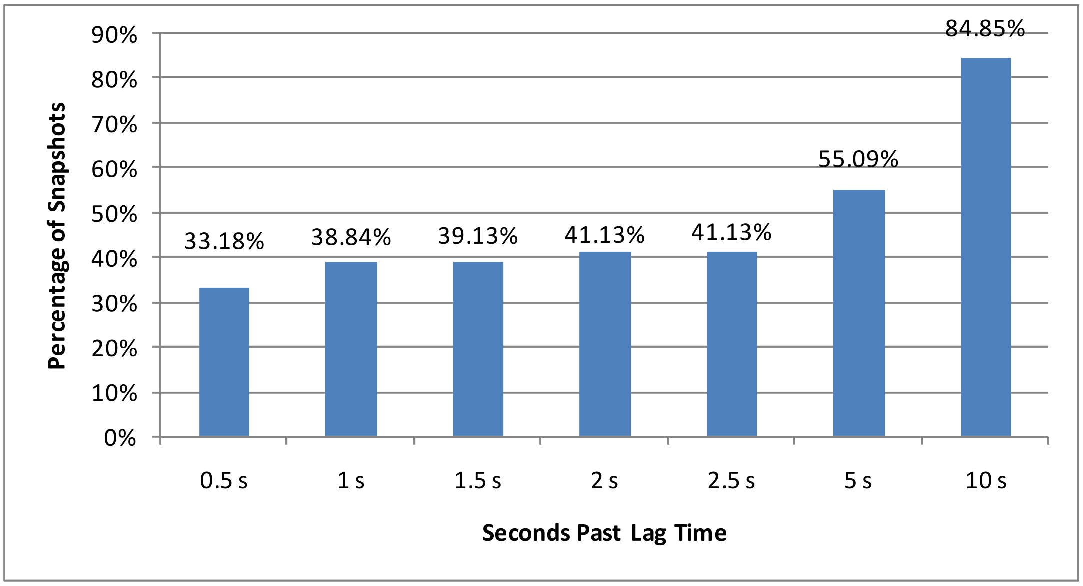

Figure 22: Cumulative Percentage of Snapshots past Lag Time – PDM August 28

A majority of snapshots were generated within 5 seconds past the lag time.

Analysis: Stop lag is useful in preventing the fill up of OBE memory with almost identical snapshots; however, the default setting of 30-snapshot buffer size is still very limiting. Increasing the default OBE buffer size setting to 300 snapshots is recommended.

Measure 8: Percentage of Start Snapshots with Current Speed above Start Threshold Speed

Measure Objective: Determine whether OBEs under PDM generated start event probe data at speeds above the start threshold speed specified in the PDM message.

Concept: Start threshold speed is used to indicate the minimum speed at which the vehicle is no longer considered “stopped.” The resulting start snapshots can be used to detect the end of traffic congestion and clearing of accidents among other things.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each start snapshot generated, verify that the snapshot speed was more than the start threshold speed.

Results: The OBEs generated start snapshots with speeds higher than the start threshold speed 100 percent of the time.

The majority of start snapshots had speeds more than the start threshold speed. However, only 23 start snapshots from 7 OBEs were received for the entire day. The snapshots with speeds less than the start threshold speed were generated by faulty OBEs.

Of the non-faulty OBEs that reported any start snapshots, all snapshots speeds were above the start threshold speed.

Analysis: The results show that start threshold speeds were implemented correctly and OBEs responded with start snapshots. However, start snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear that such units are commercially available at a reasonable cost.

Analysis for Friday, August 29, 2008

PDM Message Selected

The PDM message analyzed for August 29 had the following parameters shown in Table 21.

Table 21: Contents of PDM Message on August 29

Sample Start/End		Termination Method	
50	99	Distance	500 m
Collection Method			
Speed 1	Speed 2	Distance 1	Distance 2
5 mph	20 mph	40 m	120 m
Start Stop Strategy			
Stop Threshold	Stop Time	Stop Lag	Start Threshold
5 mph	1 s	1 s	5 mph
Message ID: 1EDAC8FB-0F01-4EBB-B324-4D5C4FAAA145			

This message was selected for detailed analysis because it had the highest number of occurrences in the OBE logs. It appeared a total of 36,391 times across all OBEs. The message that had the second highest occurrence in the OBE logs came in at 26,514 times.

Methods

On Friday, August 29, PDM messages with distance-based collection methods were sent to the vehicles. The vehicles were scheduled to go out for four runs; however, due to a PDS server issue, run 1 was cancelled. The vehicles were divided into groups of five and sent on five routes. On runs 2 and 3, the vehicles were spaced 3 minutes apart, and there was no spacing during run 4.

Table 22: Test Start and End Times – PDM August 29

Run	Start	End	Total Loops
Run 1	N/A	N/A	Cancelled due to PDS server issues
Run 2	8:20 am	9:26 am	2
Run 3	9:36 am	11:15 am	2
Run 4	12:00 pm	2:00 pm	3, except vehicles on route 1 did only 2.75 loops

Note about OBE Logs

For Friday, August 29, 13 vehicles sent uncompromised data for analysis.

OBE B856 was disqualified from inclusion in the analysis due to apparent positioning system errors. The OBE reported a constant speed of 40.39 mph for all snapshots generated.

Table 23 shows the issues experienced with other vehicles during testing.

Table 23: OBE issues on August 29

P2 (C862)	Probe generation error after session 2; resolved via reboot
C3 (B422)	No vehicle speed (CAN) after session 2; resolved via reboot
A2 (C594)	Probe application did not come back after vehicle shut down for lunch; resolved via reboot
E1 (B450)	Probe application did not come back after vehicle shut down for lunch; resolved via reboot
P10 (C590)	"Invalid location object validation" error after session 2; resolved via reboot
P12 (C366)	Heading, speed, throttle, POS, time confidence, and elevation, not equipped, HMI (screen) not responsive after initial boot up @ 4:30 am; resolved via reboot

Measure 1: Percentage of Periodic Snapshots within 20 Meters of Intended Distance

Measure Objective: Determine whether OBEs under PDM generated periodic probe data at the frequency specified in the PDM message.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This in turn allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed generation frequency (e.g., 30 meters), verify in OBE logs that messages were generated at that frequency.

Results: For the message in question, the OBEs were to generate a periodic snapshot every 40 meters if the speed was less than 5 mph. They were also to generate a periodic snapshot every 120 meters if the speed was more than 20 mph. And for speeds between 5 and 20 mph, the snapshot distance was supposed to be linearly interpolated.

Figure 23 shows how many meters after a snapshot's intended generation distance it was actually generated.

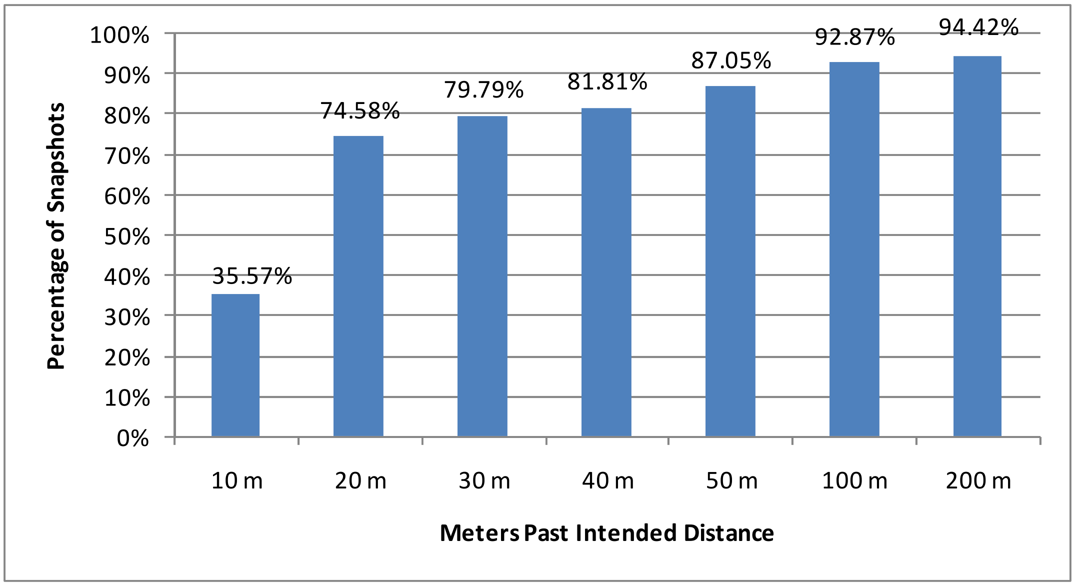

Figure 23: Cumulative Distribution of Snapshots – PDM August 29

No snapshots were generated before the intended distance, over 74 percent of snapshots were generated within 20 meters of the intended distance, and almost all snapshots were within 100 meters.

Analysis: The results show that over 74 percent of snapshots are generated within 20 meters of their intended distance. If an operator requests snapshots to be generated every 100 meters, they can expect the majority of them to be generated between 100 and 120 meters. Having the ability to alter period snapshot generation will indeed be useful. However, it would be most useful if the operators were not limited by OBE shortcomings. It is recommended that, in the future, OBEs can generate snapshots faster and by default store more snapshots in memory.

Measure 2: Standard Deviation of Differences between Anticipated Distance and Actual Distance

Measure Objective: Determine whether OBEs under PDM generated periodic probe data within a reasonable standard deviation of the anticipated distance.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This in turn allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with an intended generation frequency (e.g., 30 meters), verify in OBE logs that messages were generated within a reasonable standard deviation of the anticipated distance.

Results: Figure 24 shows the standard deviation between anticipated snapshot generation distances and actual snapshot generation distances.

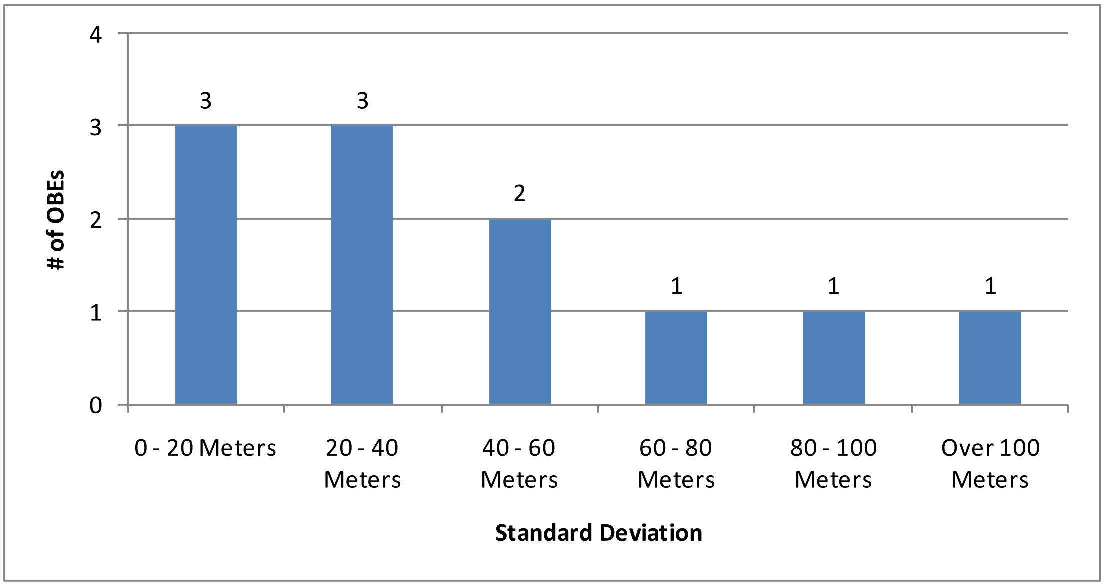

Figure 24: Standard Deviation from Intended Distance – PDM August 29

A majority of OBEs generated periodic snapshots within a standard deviation of 40 meters from their anticipated distance.

Measure 3: Percentage of Distance-Based Expirations – Percent of All PDM Activations for a Given Serial Number That Expired Based on Distance

Measure Objective: Determine the percentage of PDM activations that expired based on distance.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: In OBE OSGi logs, divide the total number of distance-based expirations of a particular serial number by the total times that serial number became active.

Results: Distance-based expiration is defined as the OBE responding to a PDM message until its intended expiration distance.

Distance-based expirations are affected by overlapping RSEs and close-by RSEs. If the OBE receives another unique message, then the newly received message becomes active, overriding the first PDM message.

The following screen shot from Google Maps (Figure 25) shows the routes and RSEs (in red) used for the PDM message in question, as well as close-by RSEs (in blue) that were also broadcasting PDM messages.

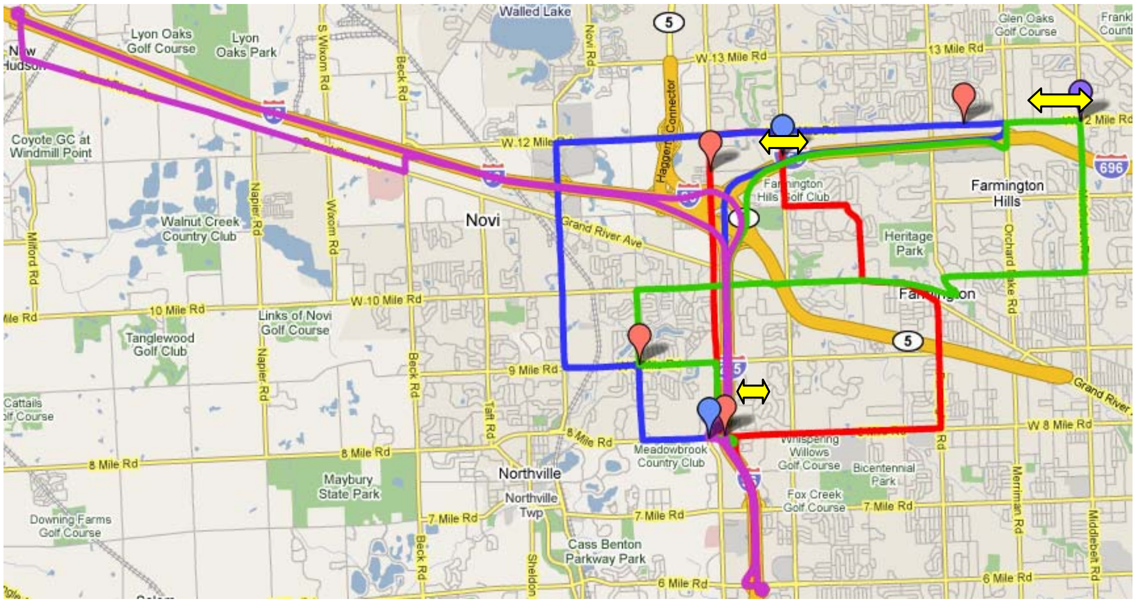

Booz Allen generated routes and plot points on Google map

Figure 25: Nearby RSE Affecting Distance-Based Expiration – PDM August 29

The three yellow arrows in decreasing length are 1.6 miles, 1.06 miles, and 0.23 miles. In at least the 0.23 mile case, it is very likely that the RSE footprints were overlapping and therefore decreasing the percentage of distance-based expirations.

Figure 26 shows the distance-based expiration of PDM messages by OBE.

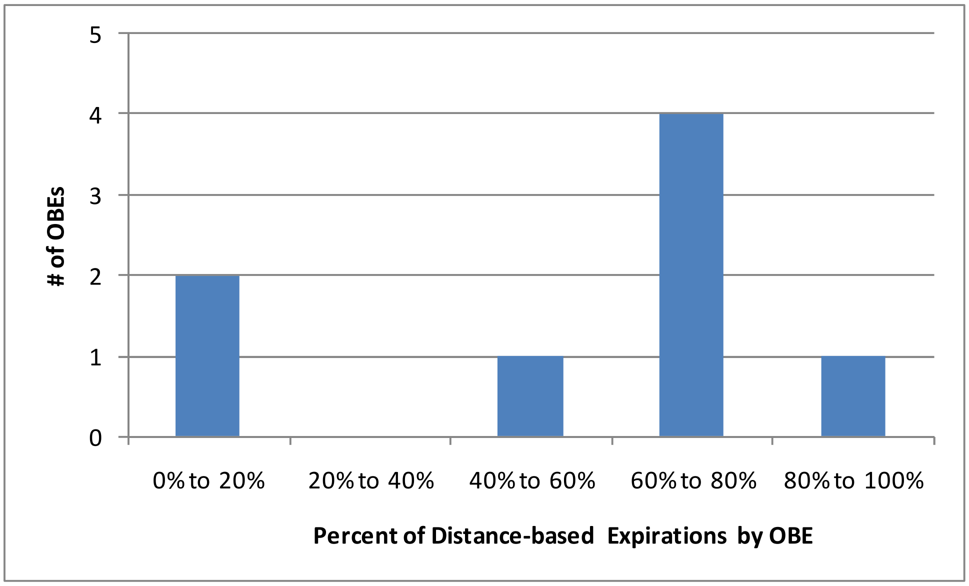

Figure 26: Percentage of Distance-Based Expirations by OBE

<u>Analysis:</u> A majority of PDM messages expired based on distance. The reason message expiration by distance is not 100 percent could be attributed to overlapping RSEs that were also broadcasting PDM messages.

Measure 4: Percent of Distance-Based PDM Expirations That Expired Correctly

Measure Objective: For PDM messages that expired by distance, determine whether their expiration was at exactly 500 meters.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in OBE logs that each distance-based expiration met the criteria defined in the PDM message.

Results: The results showed that if a message expired based on distance, it did so correctly 100 percent of the time.

The results also showed that all PDM messages expired at their anticipated expiration distance.

Analysis: The PDM message expiration times were carried out successfully 100 percent of the time.

Measure 5: Percentage of Correctly Sampled PDM Messages (i.e., percent of PDM that had a PSN within the allowed range when they became active)

Measure Objective: Determine whether only the vehicles with the correct PSN at the time of PDM reception responded to the message.

Concept: Sampling is used to limit the number of vehicles that respond to a particular PDM request and thereby controls the total number of snapshots that the TMC operator receives. Sampling makes senses when a majority of vehicles on the roadway are VII-enabled, because if every vehicle responded to a PDM request, the ensuing onslaught of snapshots from thousands of vehicles could overwhelm the infrastructure.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For the sample specified in the PDM message (50-99), verify that only OBEs with PSNs that ended in that range of numbers (50-99) responded to the message.

Results: Data from 12 OBEs for the PDM message in question was analyzed. The results showed that the sample specified in the PDM message was met 100 percent of the time.

Analysis: Sampling is a very useful tool in limiting the number of snapshots going through the system. When combined with alteration of the probe data generation scheme, it can achieve its goals of managing probe data traversing the system.

Measure 6: Percentage of Stop Snapshots with Current Speed below Stop Threshold

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at speeds below the stop threshold specified in the PDM message.

Concept: Stop threshold is used to indicate the minimum speed that indicates vehicle movement. For example, on highways with a speed limit of 65 mph, a vehicle traveling less than 10 mph may be considered "stopped." The vehicle's instantaneous speed is used to determine snapshot type. The resulting stop snapshots can be used to detect traffic congestion, accidents, and other unplanned events.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each stop snapshot generated, verify that the snapshot speed was less than the stop threshold speed.

Results: All OBEs met the goal of generating a stop snapshot 100 percent of the time, when the vehicle speed was below the stop threshold speed.

Stop lag and stop time were also decreased to 1 second each (See Table 21) . This change caused stop snapshots to be generated every second after the initial snapshot, which resulted in 2,054 stop snapshots on Friday versus only 68 stop snapshots on Thursday where the stop lag was 10 seconds.

Analysis: The results show that stop thresholds were implemented correctly and OBEs responded with stop snapshots. However, stop snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear whether such units are commercially available at a reasonable cost.

Measure 7: Percentage of Stop Snapshots Generated after Stop Lag

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at the frequency (stop lag) specified in the PDM message.

Concept: Stop lag is used to prevent the OBE from generating too many stop snapshots when the vehicle is creeping along at speeds below the stop threshold. Stop lag prevents the OBE buffer from filling up with nearly identical stop snapshots and overwriting periodic and/or event-generated snapshots.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate the difference between the current and previous stop snapshots and verify that it is after the specified lag time.

Results: 100 percent of stop snapshots met the 1-second stop lag time. Figure 27 shows how many seconds after the stop lag time the next stop snapshot was generated.

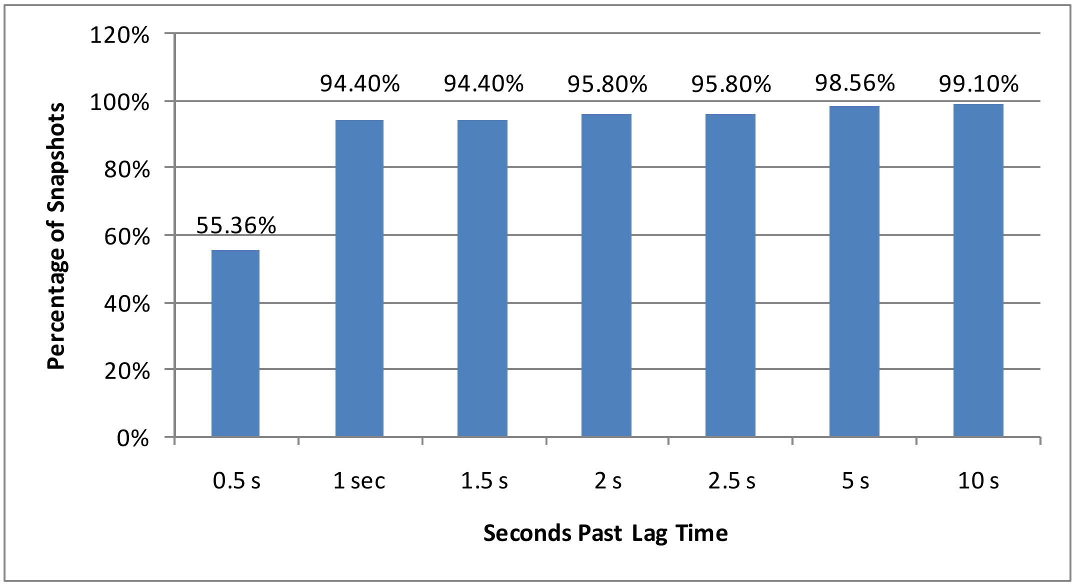

Figure 27: Cumulative Distribution of Snapshots Past Lag Time – PDM August 29

Over 94 percent of snapshots were generated within 1 second past the lag time.

Analysis: Stop lag is useful in preventing the fill up of OBE memory with almost identical snapshots; however, the default setting of 30-snapshot buffer size is still very limiting. Increasing the default OBE buffer size setting to 300 is recommended.

Measure 8: Percentage of Start Snapshots with Current Speed above Start Threshold

Measure Objective: Determine whether OBEs under PDM generated start event probe data at speeds above the start threshold specified in the PDM message.

Concept: Start threshold is used to indicate the minimum speed at which the vehicle is no longer considered “stopped.” The resulting start snapshots can be used to detect the end of traffic congestion and clearing of accidents among other things.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each start snapshot generated, verify that the snapshot speed was more than the start threshold speed.

Results: The OBEs met the start threshold goal 50 percent of the time. In contrast, on August 28, a 100-percent success rate was observed. The difference could be due to the fact that both the start and stop thresholds were the same for August 29 and both stop lag and stop time were decreased to 1 second.

Analysis: The results from August 28 showed that start thresholds were implemented correctly and OBEs responded with start snapshots. However, having the same start and stop threshold value as well as a low stop lag and stop time can lead to a steep decline in OBE responsiveness. This may be due to either OBE processing power or how the OBE software was implemented.

Start snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear whether such units are commercially available at a reasonable cost.

Conclusions

The hypothesis was that PDM messages were received and implemented by the OBEs correctly. To analyze the data, PDM messages were broken down into eight distinct measures and looked at separately. The measures were:

- Frequency of probe data generation
- Standard deviation of generated probe data from intended time or distance
- Percent of time-based expirations
- Implementation of PDM message expiration by time
- Sampling of OBEs
- Stop snapshot threshold implementation
- Stop snapshot lag time implementation
- Start snapshot threshold implementation.

Based on the results, it is safe to conclude that PDM was properly implemented and OBEs responded correctly.

4.2.3 Vehicle Stop and Start Probe Data Messages (Test Case PDS-05)

Original Test Case Objectives

- Verify that the Public Apps can trace vehicle movements as they traverse an area with heavy traffic that causes stops and starts.
- Exercise the OBE’s capability to generate start and stop probe data messages.

Revised Test Case Objectives

- Exercise the OBE’s capability to generate start and stop probe data messages.

Original Data Outputs

- Application-computed same-PSN vehicle traces for all vehicles
- Comparison of vehicle start and stop locations from application traces and from reference GPS measurements.

Revised Data Outputs

- Output of vehicle’s stops that can be read by a graphical map interface, such as Google Earth.

Original Exit Criteria

The test passes if the application correctly constructs the same-PSN vehicle trajectories for all vehicles and the start and stop locations exhibit the same level of accuracy observed for vehicle location in laboratory testing. The test will not fail due to communication failures occasioned by packet loss, as long as the packet loss does not exceed the value observed in integration and testing, with adjustments made for "real world" traffic situations.

Revised Exit Criteria

The test passes if the start and stop locations exhibit the same level of accuracy observed for vehicle location in laboratory testing. The test will not fail due to communication failures occasioned by packet loss, so long as the packet loss does not exceed the value observed in integration and testing, with adjustments made for "real world" traffic situations.

Deviations from the Test Cases

The test case and conditions matrices imply that there will be one run on an arterial route and one run on a freeway route. Instead, three runs were executed, each encompassing both freeway and arterial roads. This was done in order to use the same routes that had been used on the PDS-02 test case, which were familiar to drivers, and because it was very simple to distinguish freeway data from arterial data, simply based on location of the data points.

The original data outputs were removed for the following reasons:

- Application-computed same-PSN vehicle traces for all vehicles – *Algorithm testing was de-scoped*.
- Comparison of vehicle start and stop locations from application traces and from reference GPS measurements – *No reference GPS information was available*.

Hypothesis

- In heavy traffic conditions, at locations where traffic backs up, the percentage of start and stop snapshots will increase significantly (double or more) over when there is no traffic backup.

Data Analysis Strategy

Algorithm testing was a prerequisite for verification of the original test objective, that a vehicle's movements in stop-and-go traffic be traced. However, because algorithm testing was de-scoped, it was only possible to produce an output of the vehicle stops that can be read by a graphical map interface, such as Google Earth. This was the output that was created.

In addition, the original objective implied ability to locate stopped traffic based on received vehicle data. Thus, for selected locations along the testing routes, the percentage of stop snapshots among all snapshots was computed in 5-minute intervals.

Measure 1: Number of Start/Stop Snapshots in Given Roadway Segments as Percentage of All Snapshots Received in a Given Location for 5-Minute Windows throughout the Testing Period

Methods

For 1 day of testing (September 2), during late afternoon and evening hours, vehicles traveled in platoons (all 25 vehicles together) along the three probe data routes that included freeway and arterial locations. The three routes used were PDS Route 1, PDS Route 3, and PDS Route 4.

Results and Analysis

Probe data was received along each of the routes traversed. Stop and start snapshots from probe data received at the application were extracted.

Figure 28 from Google Earth shows the location of all stop snapshots for the day of start/stop testing.

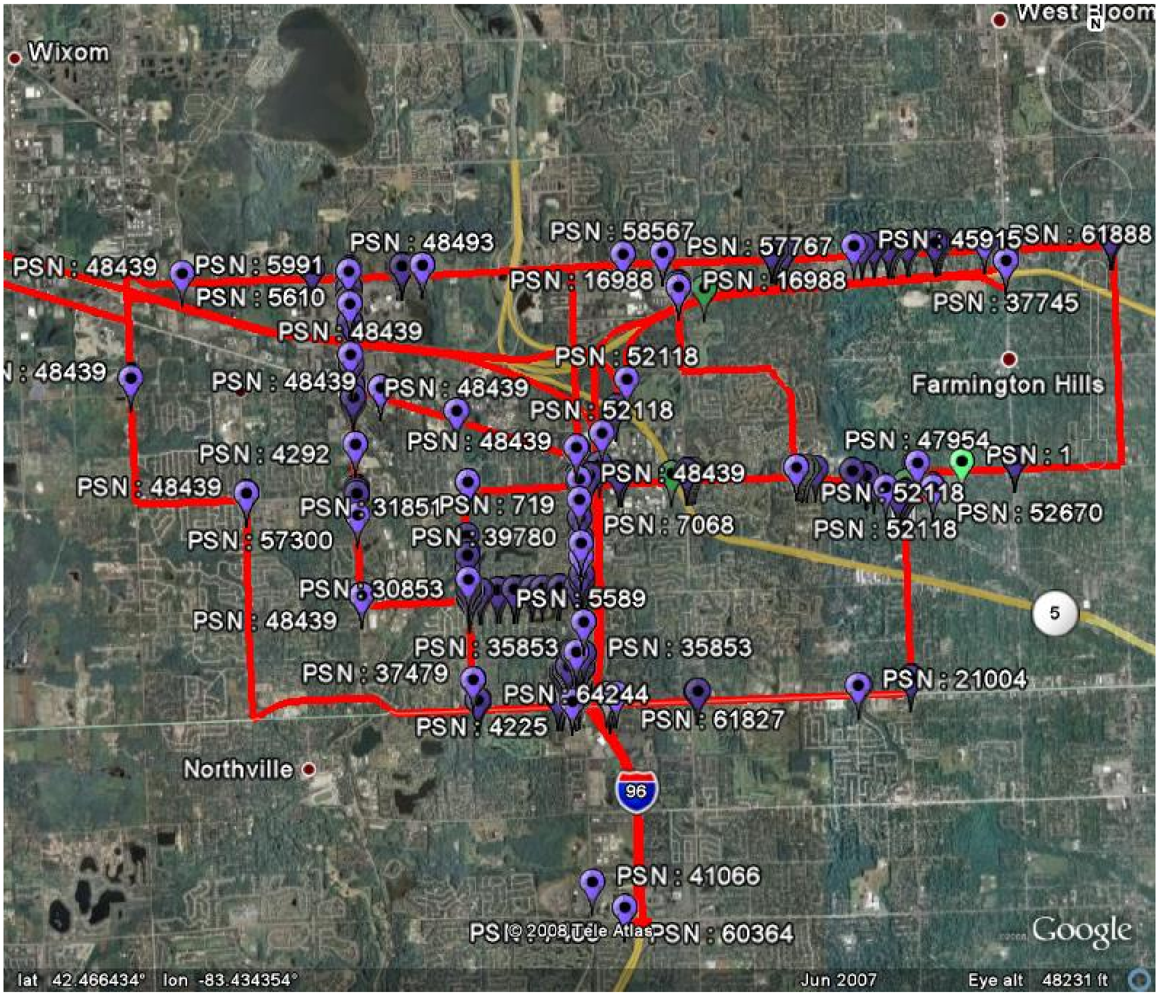

Booz Allen generated routes and plot points on Google map

Figure 28: Location of All Stop Snapshots

Using Google Earth, locations containing high densities of start/stop snapshots were located. One location that contained a high density of start/stop snapshots was Grand River and Novi Road. This is an intersection with a stoplight, and the vehicles transmitted many stop snapshots at this light.

Figure 29 provides a zoomed-in view of the area.

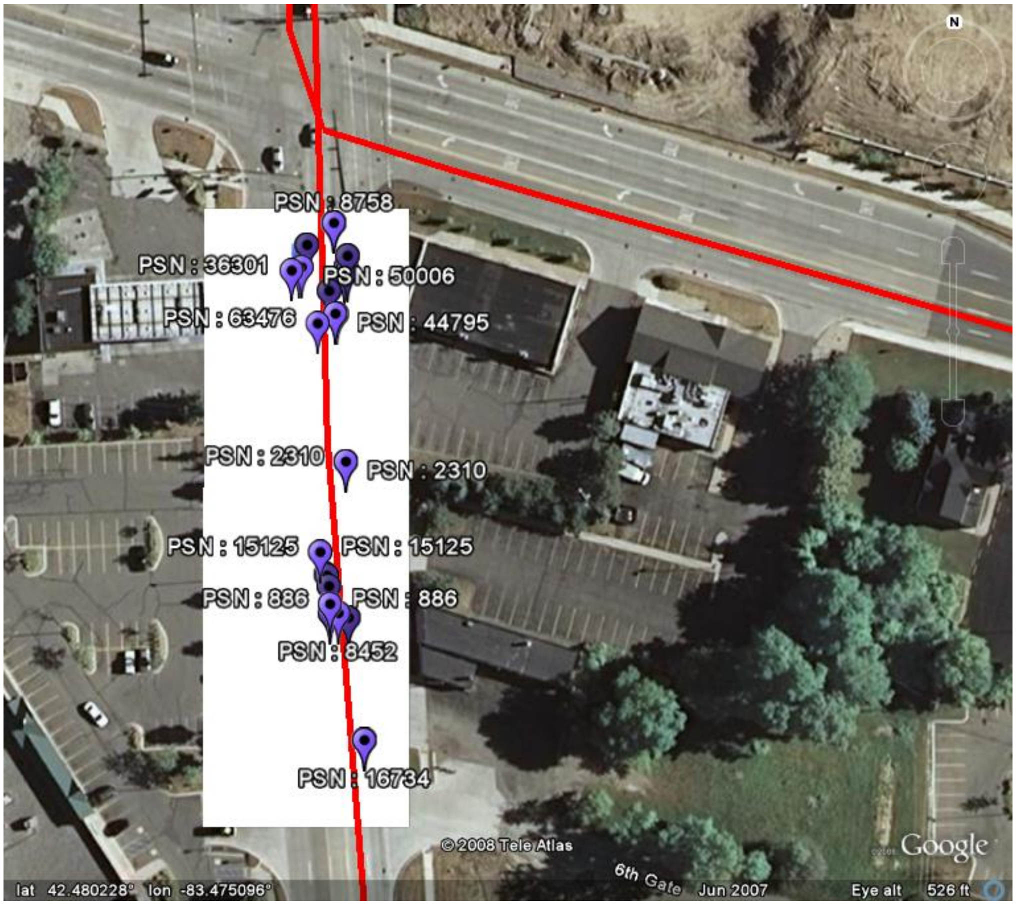

Booz Allen generated routes and plot points on Google map

Figure 29: Stop Snapshots at Grand River and Novi

Figure 30 shows stop snapshots as a percentage of all snapshots for this location.

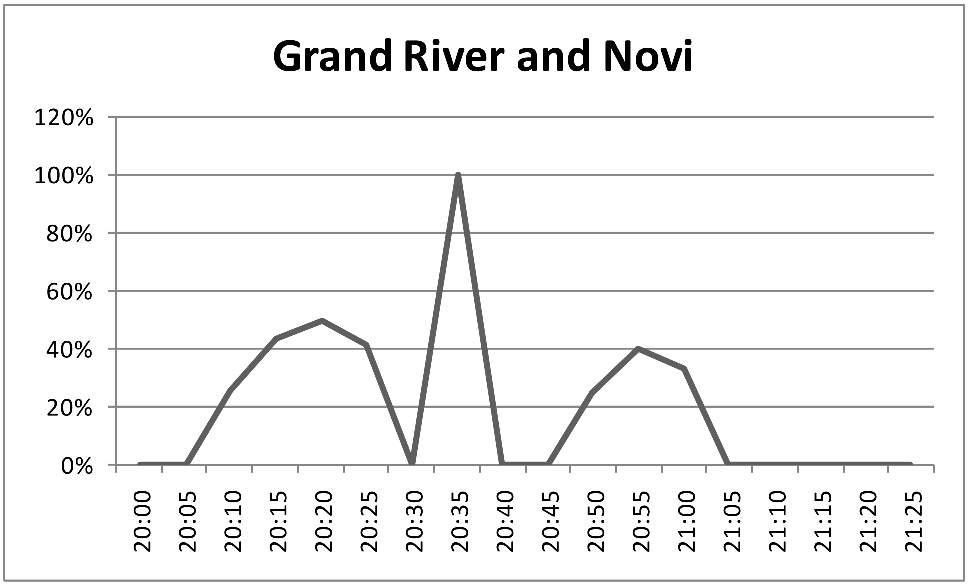

Figure 30: Stop Snapshots as a Percentage of All Snapshots (1)

The spike to 100 percent at the 20:35 to 20:40 time slot was caused by a single vehicle, a vehicle well behind the other vehicles in the first loop.

For comparison, here is the same computation for an area the same size as the above area, not adjoining the intersection, and thus not experiencing stops.

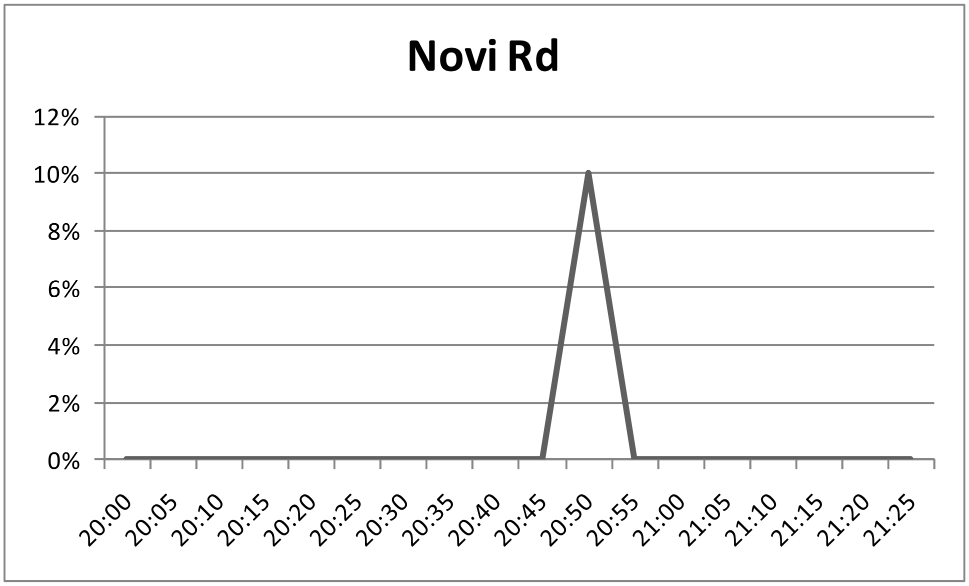

Figure 31: Stop Snapshots as a Percentage of All Snapshots (2)

In this data, only one stop snapshot is recorded, and the percentage of stop snapshots never exceeds 10 percent.

The above analysis was done for an arterial. There were few small locations of stopping on freeways on the day in which explicit start/stop tests were executed. However, during one of the days of probe data reception (PDS-02) testing, there was significant stopping on the freeway. That was on August 25, 2008.

The I-96 connection, where I-96 West diverges from I-696 and I-275, proved to be a very fertile area for stops. Figure 32 presents an image of all stop snapshots collected on one segment for the entire day of August 25.

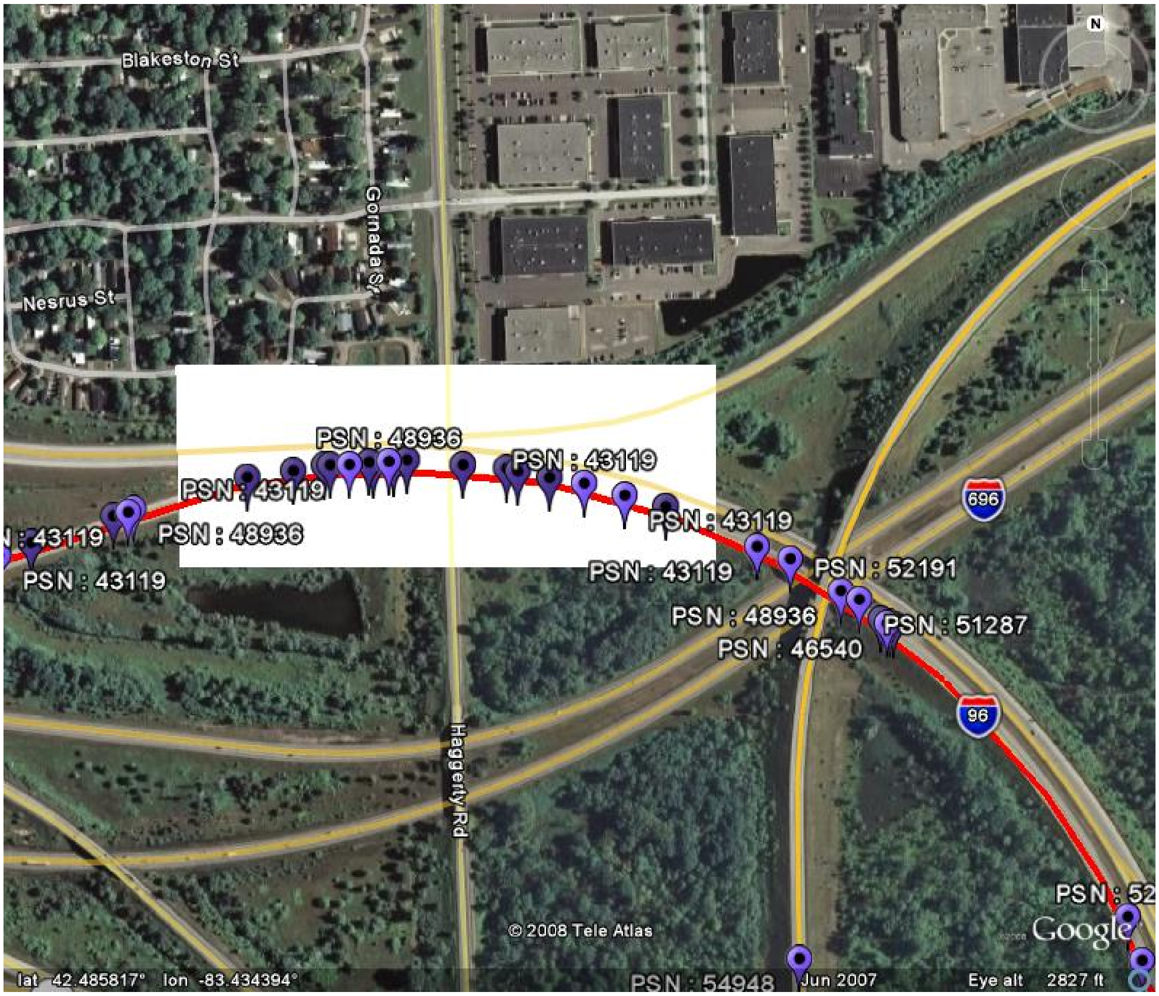

Booz Allen generated routes and plot points on Google map

Figure 32: Stop Snapshots on I-275 N to I-696 W

Figure 33 illustrates the stop percentages for this area of I-96.

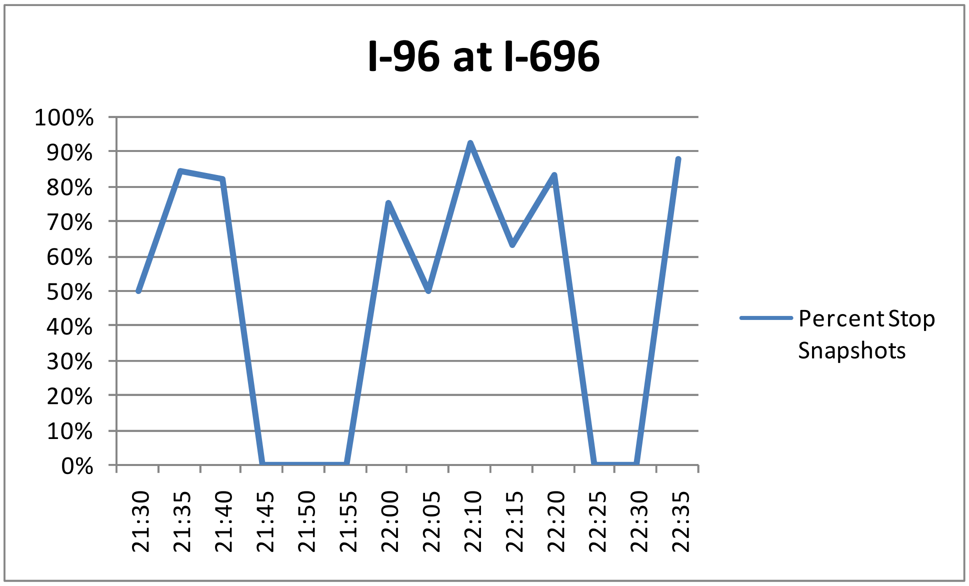

Figure 33: Stop Snapshots as a Percentage of all Snapshots on I-275 N to I-696 W

Clearly, during stopping times, percentages of stop snapshots exceed 50 percent of all snapshots.

In comparison, Figure 34 presents data for another freeway segment after the stopping location.

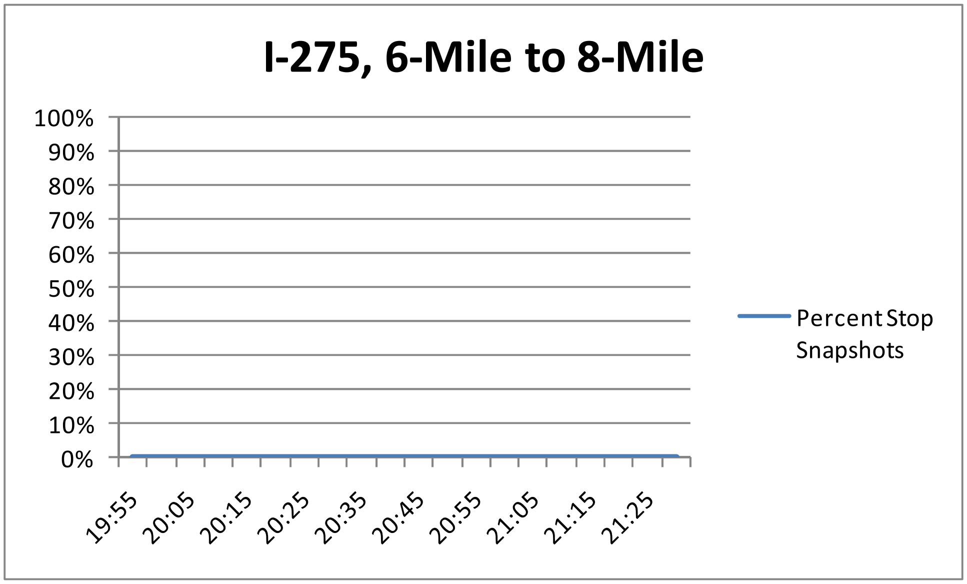

Figure 34: Stop Snapshots as a Percentage of All Snapshots on I-275 between 6 and 8 Mile

A comparison of the two charts shows the bottlenecks along roadways.

Conclusions

The probe data gathering OBE software, coupled with the VII PDS, provide a platform that allows Network Users to detect vehicle stops. Vehicle stops can be seen visually.

Further system improvements will be needed before algorithms can be developed. Specifically, the lack of probe data (due to the two issues -- V-DTLS and infrequent generation of probe data -- discussed elsewhere) prevents any scientific rigorous computation of stop snapshot percentage threshold values where stopping occurs.

However, because stop snapshots are handled no differently than other snapshots, it may be possible to develop algorithms that detect stopping areas using this data.

Tweaking the default stop snapshot generation values in the OBE software (stop speed threshold, stop time, stop lag, start speed threshold) may also be desirable so that vehicles create data that is more useful for stop snapshot generation.

4.2.4 Event-Generated Snapshots/Slippery Road (Test Case PDS-07)

Due to difficulties in scheduling test cases to coincide with inclement weather, this test case was revised considerably. Table 24 presents the revised test case about event-generated snapshots and slippery roads.

Table 24: Revised PDS-07 Test Case

PDS Test Case 7: Event-Generated Snapshots / Icy Road			
Test Plan	Applications Integration and Test	**Test Originator**	Booz Allen Hamilton
Test Case Number	TC-PDS-07	**Test Location**	Detroit DTE
Test Objective	• Test the Public Apps ability to distinguish event-generated snapshots. • Exercise the OBE's capability to generate event-triggered messages.	**Approximate Test Duration**	Data Collection: 6.75 hours Data Analysis: 8 hours
Resources/ Equipment Required	Vehicles: 9 Vehicle Operators: 5 Test Coordinator: 2 Data Analysts/SME: 1 RSE Sites: 2 Field Data Collection Personnel: 0	**Data Collection Needs**	• OBE Logs (including vehicle input signals) • SDN Network Traffic Logs (filtered for PDS) • Application Database
Data Outputs	• Graphical output of vehicle's event-triggered snapshots	**Test Risks/ Contingencies**	• None
Requirements	See Requirements spreadsheet.		
Capabilities Tested	8		
Brief Description	The purpose of this test case was to verify that the Public Apps correctly received event-generated probe data snapshots and to exercise the OBE's ability to create event-generated snapshots. ABS and traction control events, by default, cause event-generated snapshots to be created by the OBE. These events typically occur in slippery conditions, such as on icy roads or bridges. However, because the POC period occurred in August and September, no days of slippery road conditions were experienced during POC. Thus, it was necessary to simulate slippery conditions. Slippery conditions could only be simulated in a controlled testing environment. The application collected probe data from the controlled test environment.		
Entrance Criteria	See Entrance Criteria section above.		
Preconditions	1. The PDS can collect and provide probe data to Network Users throughout the geographic extent. This data includes vehicle location, heading, and speed for each OBE-equipped vehicle in the roadway network. 2. Vehicles are transmitting probe data to the PDS when they make contact with RSEs. 3. Probe data subscription profiles have been developed and saved in a Network User application.		
Precedent Test Cases	TC-PDS-02		
Test Conditions	Time and budget constraints did not permit test execution for all combinations of test conditions. Test conditions were simulated as described in the Test Method/Configuration field below.		

PDS Test Case 7: Event-Generated Snapshots / Icy Road	
Test Method/ Configuration	For 1 day of testing, slippery conditions were simulated in a controlled area and drivers caused traction control and ABS events. The controlled area used was an access-controlled parking lot at the Roush facility. Vehicles lined up and drove into a "Traction Control System (TCS) Area" that was strewn with sand, and the drivers revved their engines and took off quickly, spinning their tires, and typically causing the TCS to engage—in turn, creating an event-generated snapshot. After driving about 50m to the "ABS Area," drivers then braked hard on an area of pavement that had been covered with water, causing an ABS event. Drivers then left the access-controlled parking lot, drove around the Roush building, going by the two portable RSEs that had been set up at Roush, and returned to the access-controlled area.
Route	See Test Method.
Specific Driving Instructions	See Test Method.
Minimum Time or Distance for Data Collection	Vehicles completed 17 laps around the lot.
Actions of Field Data Collection Personnel	N/A
Data Recording by Test Observers	N/A
Exit Criteria	The test passes if the application receives all event-generated snapshots. The test will not fail due to communication failures occasioned by packet loss, as long as the packet loss does not exceed the value observed in integration and testing.

Hypothesis

- In simulated conditions, at locations where vehicles line up awaiting their next run, the percentage of start and stop snapshots will increase significantly (double or more) over when there is no traffic backup.
- Event-generated snapshots will have their highest concentrations around the ABS and traction control system (TCS) areas.

Data Analysis Strategy

To verify that vehicle movements causing event-triggered snapshots can be traced, it is only necessary to produce an output of the vehicle movements that can be read by a graphical map interface, such as Google Earth. Thus, this is the output that was created.

Measure 1: Number of Event-Generated Snapshots in Given Roadway Segments as Percentage of All Snapshots Received in a Given Location

Methods

During run-for-record testing, five professional technicians from Roush drove a total of nine vehicles—the five Nissan Altimas and the four Ford Edges. This was done because Roush required that this type of driving be done by professional technicians. Also, during dry runs of the test, it was very difficult to make the Jeep Cherokees engage ABS or TCS; thus, they were not used during the run-for-record testing.

The test coordinator waited until the TCS area and ABS area were clear of vehicles, and then allowed vehicles to perform the driving maneuvers as described above. The deputy test coordinator waited at the entrance to the access-controlled parking lot to be sure there were no vehicular incidents as vehicles entered and exited the parking lot.

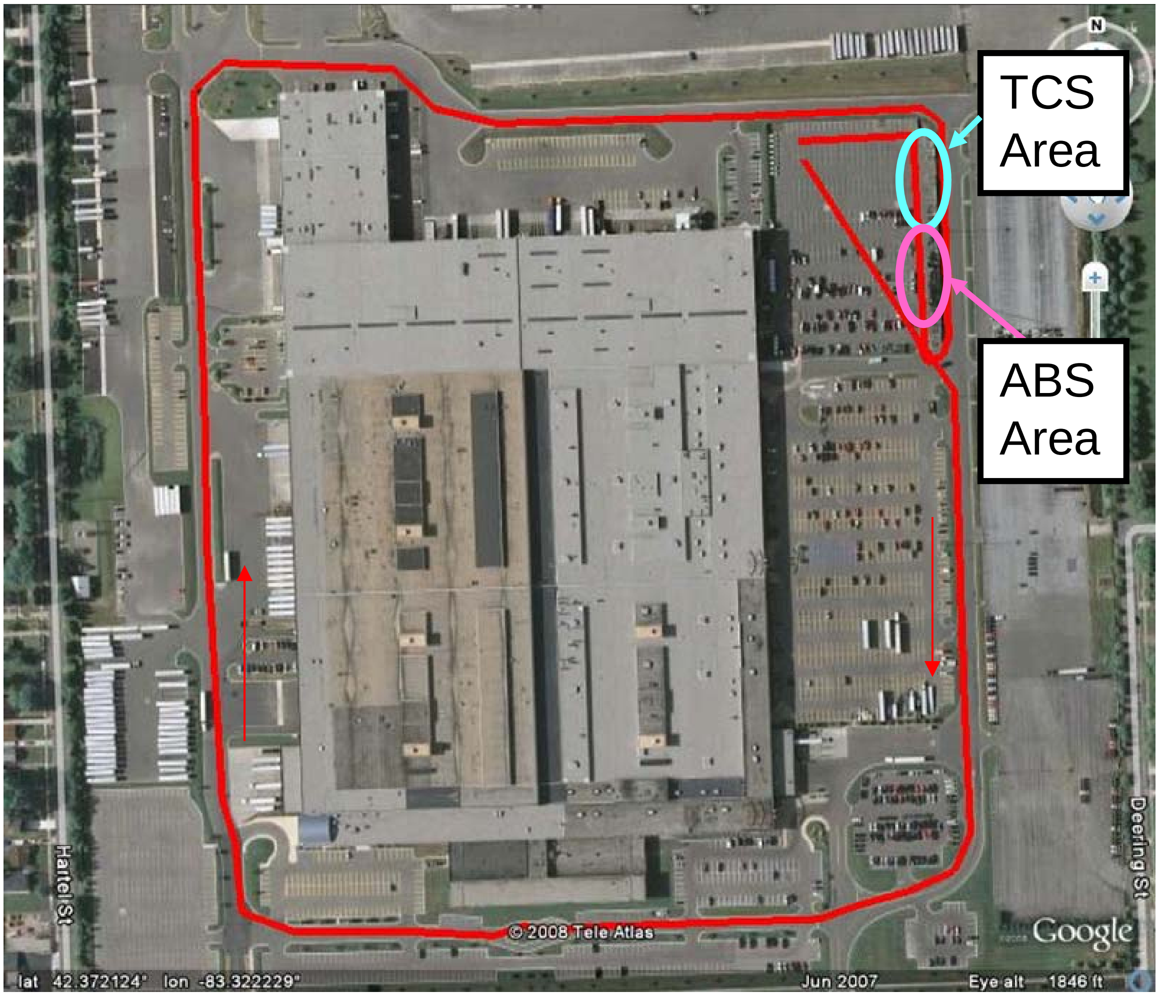

Booz Allen generated routes and plot points on Google map

Figure 35: Slippery Condition Test Area

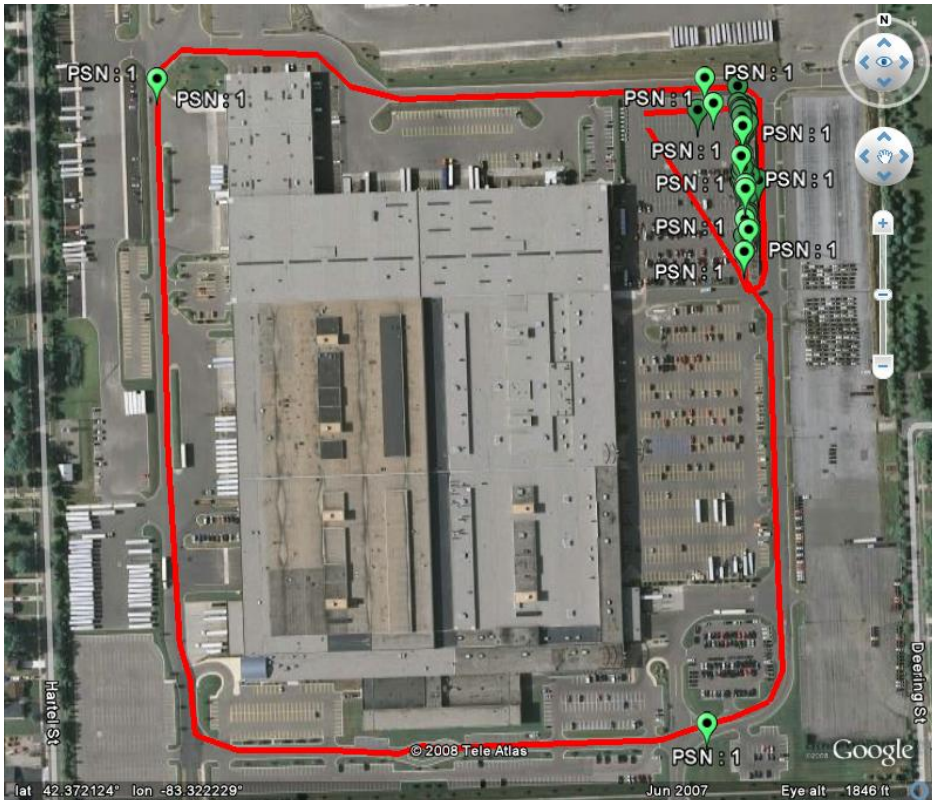

Booz Allen generated routes and plot points on Google map

Figure 36: Slippery Conditions Route with All Data Points Highlighted

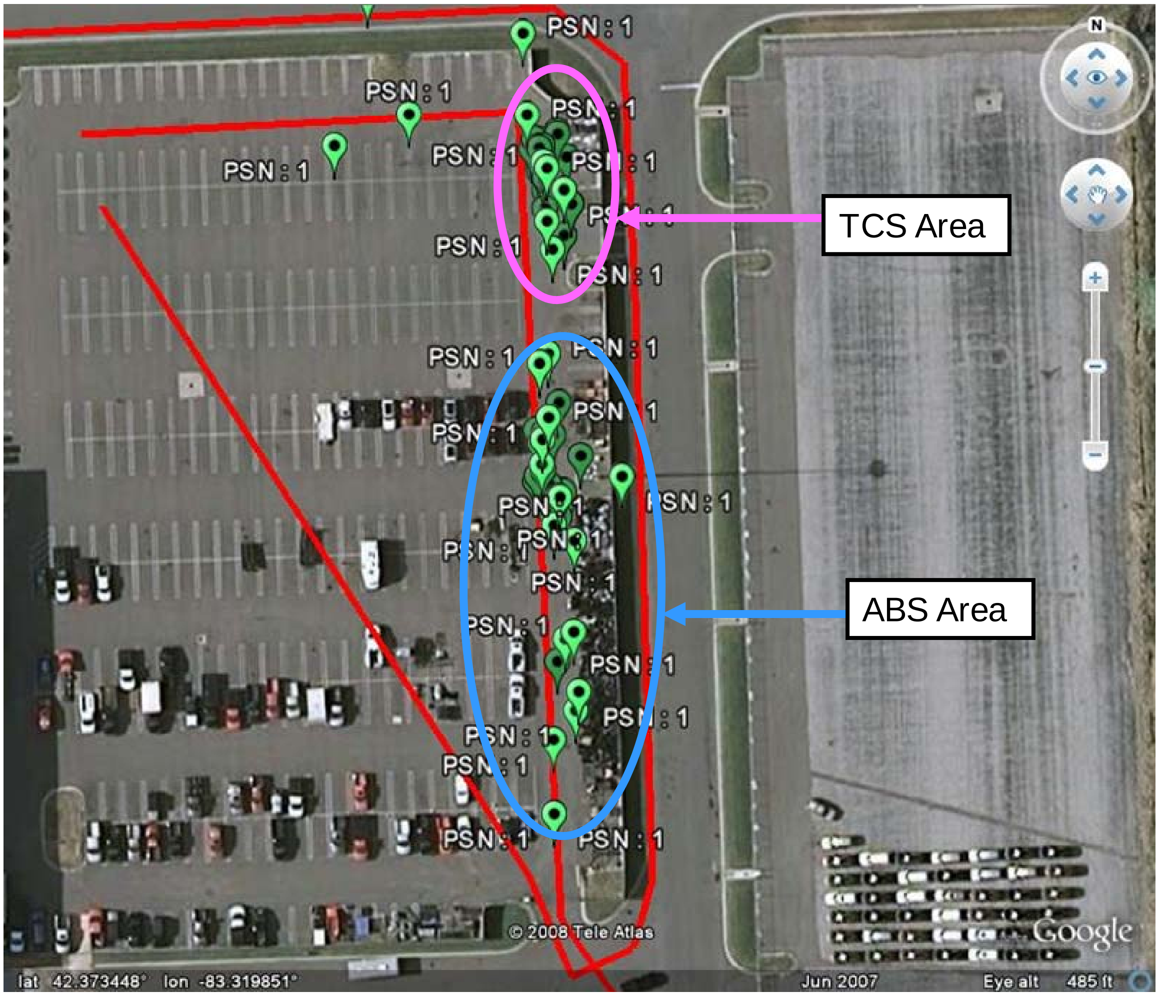

Booz Allen generated routes and plot points on Google map

Figure 37: Slippery Data Detail

Measure 1 was computed for the area of the access-controlled parking lot where the slippery conditions were simulated.

It should be noted that the majority of the snapshots were stop snapshots generated by the vehicles while sitting and waiting to be taken on the next loop through the TCS and ABS areas and the route as pictured above. Vehicles were left running during all test runs, even when not in use. Since there were nine vehicles, but only five drivers, the vehicles were not in use for about 44 percent of the time.

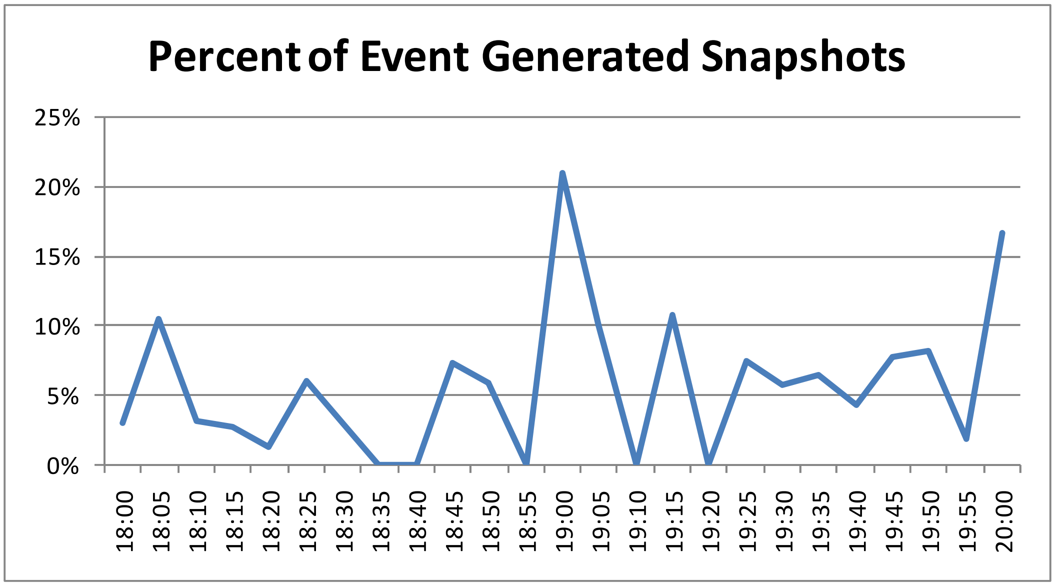

Figure 38: Percentage of Event-Generated Snapshots

While events were being simulated, the percentage of event-generated snapshots varied from 0 percent to 20 percent of all snapshots, as seen in Figure 38. Since this was a simulation, these values are by no means indicative of values that would occur in real slippery pavement conditions. However, it shows that the VII system correctly transmits event-generated snapshots.

During other testing days, when no slippery condition simulation took place, roughly eight event-generated snapshots per day were experienced. These were typically caused by drivers braking hard on dry pavement during the course of normal driving (e.g., to make stoplights, avoid accidents). For example, during the 6 days of PDS-02, 47 event-triggered snapshots were recorded. These snapshots almost always occurred in isolation from each other, leading to low percentages of event-generated snapshots. For example, if there were a 5-minute window with 50 snapshots received, including 1 event-generated snapshot, only 2 percent of snapshots were event-generated.

However, in 5-minute time segments when few snapshots were received, this could lead to relatively high percentages (e.g., when there were eight snapshots received, one event-generated snapshot could mean 12.5 percent).

Conclusions

This test shows that the VII system correctly transmits event-generated snapshots in groups. It may be possible to construct an algorithm that identifies slippery conditions when event-generated snapshots are received in concentrations.

To detect slippery pavement, it may be useful to have the count of event-generated snapshots and the count of total snapshots received (or different-PSN snapshots received) as well as the percentage of event-generated snapshots. For example, one may be able to say that slippery conditions are occurring only when there are at least 4 different event-generated snapshots received with at least 4 different PSNs, that the percentage of event-generated snapshots received is above 10 percent, and at least 20 total snapshots have been received for a given event-generated snapshot.

4.2.5 Same-PSN Vehicle Trajectories: Turning Movements (Test Case PDS-04)

Test Case Objective

- Verify that the application can trace vehicle movements as they traverse an intersection.

Original Data Output

- Application-computed same-PSN vehicle traces for all vehicles.

Revised Data Output

- Manually computed same-PSN vehicle traces for all vehicles.

Original Exit Criteria

The test passes if the Public Apps correctly constructs the same-PSN vehicle trajectories for all vehicles. Vehicle location data received by the application should exhibit the same level of accuracy observed for vehicle location in laboratory testing for turning, non-static vehicles. The test will not fail due to communication failures occasioned by packet loss, as long as the packet loss does not exceed the value observed in integration and testing, with adjustments made for “real world” traffic situations.

Revised Exit Criteria

The test passes if the operator can correctly construct the same-PSN vehicle trajectories for all vehicles. Vehicle location data received by the Public Apps should exhibit the same level of accuracy observed for vehicle location in laboratory testing for turning, non-static vehicles. The test will not fail due to communication failures occasioned by packet loss, as long as the packet loss does not exceed the value observed in integration and testing, with adjustments made for “real world” traffic situations.

Deviations from the Test Cases

The test case and conditions matrices implied that drivers would choose their direction of turn each time they approached the focal intersection, and that each driver would thus make all three movements. However, the test coordinator, in consultation with the Roush program manager, strongly discouraged this procedure, since it had large potential for driver uncertainty and error.

Instead, vehicles were assigned a turning direction and asked to complete as many turns of the route in that direction as possible during the time allotted.

Hypothesis

- Turning movements can be computed for a significant percentage of vehicles for which probe data is received before and after an intersection.

Data Analysis Strategy

Same-PSN vehicle chains are computed by the algorithms, whose testing has been de-scoped.

However, to verify the original test objective (that turning movements can be traced using the data), it is only necessary to know the percentage of vehicle trips through the intersection that have the same PSN before and after going through the intersection. Thus, a measure should be constructed to determine how many vehicles send probe data snapshots with the same PSN before and after going through the intersection.

Measure 1: Percentage of Vehicles with Same PSN before and After Traversing a Focal Intersection

Methods

Divide vehicles into groups turning right, going straight, and turning left (7 left-turn vehicles, 10 straight-through vehicles, and 8 right-turn vehicles). Instruct drivers to approach the focal intersection (12-Mile and Farmington, from 12-Mile Eastbound), make their assigned turn, and return to the focal intersection on the specified route. Instruct drivers to complete as many loops as possible of the route during the run time.

During the first of the three runs, allow default probe data settings to be used. During the remaining two runs, use different probe management settings to determine the effects of PDM on the ability to compute turning movements.

Since more probe data will increase the likelihood of having data with the same PSN at the start and finish of an intersection, choose PDM requests that increase the frequency of probe data generation from the default settings.

Table 25 presents the run times and corresponding PDM requests.

Table 25: Run Time and Corresponding PDM Requests – Turning Movements

Run	Start	End	PDM Request
Run 1	10:10 AM	12:09 PM	None
Run 2	12:53 PM	2:42 PM	Time-based periodic snapshot generation: Snapshots constantly generated at 4s Termination time: 60s Start-Stop settings left at default Omni-directional (no directional limitation) No sampling (100% of vehicles should respond to PDM request)
Run 3	3:04 PM	4:35 PM	Distance-based periodic snapshot generation: Distance between snapshots rises from 4m at 1m/s to 20m at 5 m/s Termination Distance: 500 m Start stop settings: Stop threshold 1 m/s Stop Time 5s Stop Lag 5s Start threshold 1 m/s Omni-directional (no directional limitation) No sampling (100% of vehicles should respond to PDM request)

Data Analysis Steps

OBE probe gathering OSGi logs were analyzed for a given start and end time on a given day. Based on the OBE logs, the number of “loops” per vehicle and number of “same-PSN turning movements” per vehicle were counted. Same-PSN turning movements are defined as the number of times that each vehicle records a snapshot within the four borders of the start box and another snapshot within the four borders of the left, right, or straight-through boxes with the same PSN. Loops are defined as the number of times each vehicle has two snapshots in the start box, separated by more than 5 minutes.

Table 26 defines the four boxes, which are illustrated in Figure 39. The start box is white, the left-turn box is green, the straight-through box is yellow, and the right-turn box is blue.

Table 26: Definition of Turning Movement Areas

	West Border	East Border	North Border	South Border
Start Box	83.398118° W	83.378593° W	42.498810° N	42.497752° N
Left-Turn Box	83.379287° W	83.378098° W	42.506000° N	42.498864° N
Straight-Through Box	83.377975° W	83.369000° W	42.499373° N	42.498559° N
Right-Turn Box	83.378404° W	83.377197° W	42.498463° N	42.490000° N

Figure 39 illustrates the routes of the vehicles traveling in the three directions specified here.

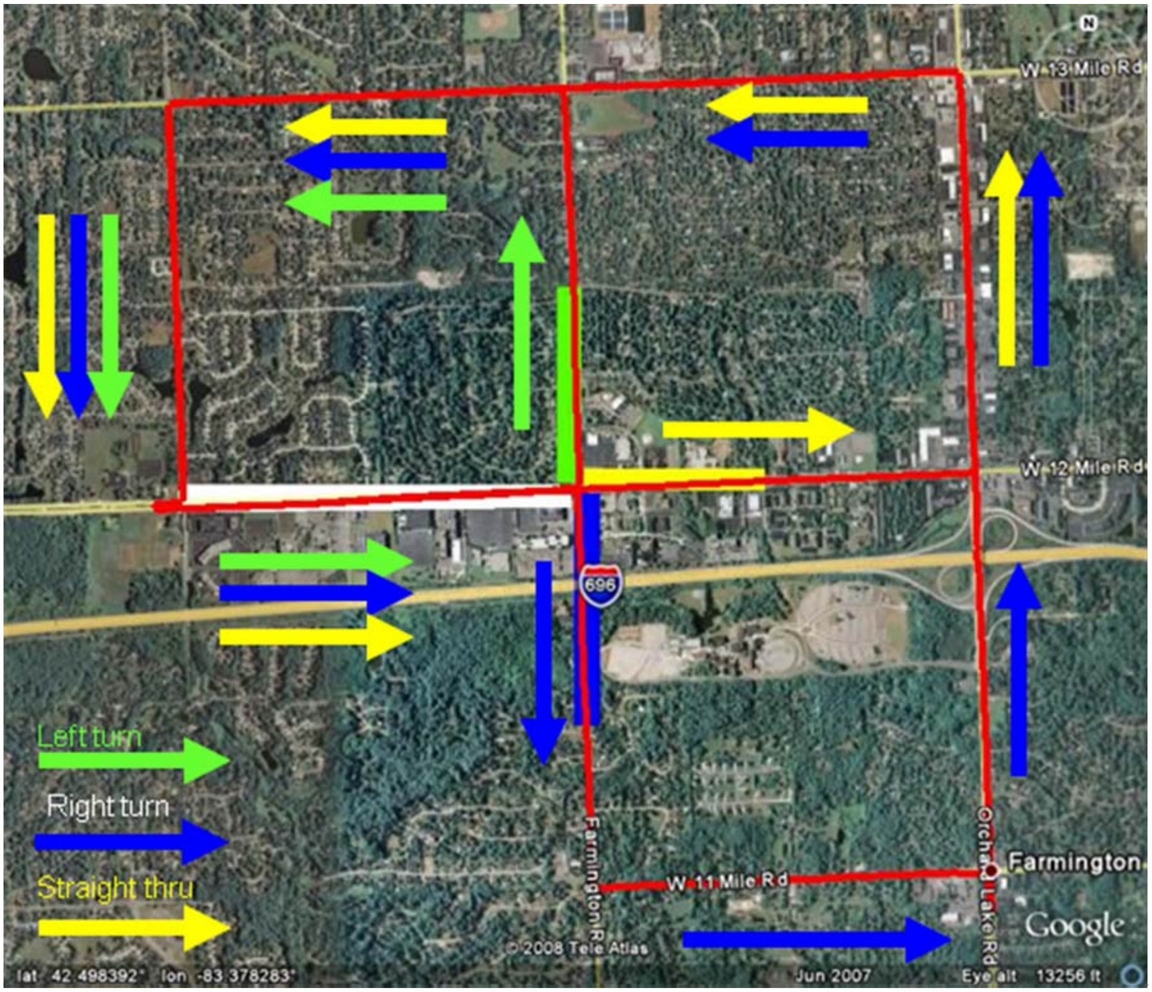

Booz Allen generated routes and plot points on Google map

Figure 39: PDS-04 Routes

For test run 1, vehicles turning right took another route—they turned around in the Oakland County Community College entrance, halfway between 12-Mile and 11-Mile on Farmington road, drove back to 12-Mile and Farmington, and then proceeded on the route for left-turning vehicles.

Because of the scarcity of real-time probe data from right-turning vehicles, the test coordinators decided to reroute the vehicles according to the routes displayed above, so that the right-turning vehicles would be certain to be out of the range of the RSE when they turned around, and so that they would contact an unencumbered RSE (12-Mile and Orchard hill) where they could dump their probe data relatively soon after they crossed the focal intersection).

Figure 40 illustrates the focal intersection in close detail. If any snapshots are found in both the start box (white) and any of the turning boxes (left/green, straight/yellow, right/blue) having the same PSNs, a turn is detected.

Booz Allen generated routes and plot points on Google map

Figure 40: PDS 04 Detection Zones

Results and Analysis

Several remarks on the data should be made before the discussion of the results.

Several OBEs behaved unusually during testing. These OBEs' logs were not used in the analysis, as follow: Two OBEs (C590, C545) generated no snapshots during the entire day, and their results are thus not included here. Another OBE's log (C453) contained unreliable data (snapshots located hundreds of meters apart time stamped at the same millisecond). Data from this OBE's log was not used in the analysis, since the unreliable data could have been misinterpreted by the data analysis tool.

This analysis was performed on OBE logs, so that the effects of the V-DTLS probe data reception limitations would not impact this analysis. However, the "infrequent generation of probe data"—the fact that OBEs generated probe data less often than they should have per the default settings could have affected the outcome of this analysis, since this analysis was based on the probe data generated by the vehicles. If vehicles failed to generate probe data before their PSN changes and after they left the start box, their turning movements would not be recognized, so any factor that slows down the generation of probe data harms the likelihood of creating turning movements with the same PSN on both sides of the turn.

With the relatively small number of detected turning movements in this test, randomness can play a significant role in the results observed. It should also be noted that even under ideal circumstances, this test is subject to randomness. Vehicles change PSNs at random times, per the probe data generating specification, and traffic flows cause vehicles to spend differing amounts of time in intersections. The scarcity of the probe data exaggerates this randomness by creating a layer of sampling on top of the already random data.

It should also be noted that vehicles generated snapshots less frequently when stopped than when moving. By default, vehicles generated snapshots every 10 seconds when stopped. This number compares with every 4 seconds when moving for runs 1 and 2 or even more frequently (every 4 meters to 20 meters) for run 3. (Run 1 was entirely default probe generation, so as speed increased, the time between snapshots increased from 4 seconds.)

Table 27 presents the results of the analysis. The first block of results shows the number of “same-PSN turning movements” (movements with the same PSN in the start box as in the indicated direction) for each of the directions, while the second block of results includes all turning movements that would have been detectable from the probe data generated in each of the OBEs (called loops in the above analysis tool description).

The third block of results is the percentage of all turning movements that would have been detectable (Measure 1 as described above).

Table 27: PDS-04 Results

	Same PSN turning movements				All turning movements (loops)				Percentage of all turning movements with same PSN			
	Straight Through	Right Turn	Left Turn	Totals	Straight Through	Right Turn	Left Turn	Totals	Straight Through	Right Turn	Left Turn	Totals
Run 1	4	2	0	6	20	5	21	46	20.0%	40.0%	0.0%	13.0%
Run 2	7	6	2	15	16	23	31	70	43.8%	26.1%	6.5%	21.4%
Run 3	8	2	4	14	15	8	21	44	53.3%	25.0%	19.0%	31.8%
Totals	19	10	6	35	51	36	73	160	37.3%	27.8%	8.2%	21.9%

From the results, two main trends can be noted: (1) the percentage of detectable turning movements increases with the runs; (2) and the straight through is the most detectable direction, followed by right turns.

The first trend is explained simply by the fact that having more probe data generated creates a greater likelihood of detectible turning movements. The PDM request for run 2 lowered the periodic generation time to 4 seconds regardless of vehicle speed. The PDM request for run 3 lowered it even further, requiring a snapshot every 4 meters at 1 mph, or every 20 meters at 5 mph. In addition, the PDM for run 3 lowered the time between stop snapshots to 5 seconds. Clearly, the more data that is generated, the more likely it is that there will be one snapshot before the intersection with the same PSN as after the intersection. It should be noted that this difference due to the PDM request is strong confirmation that using PDM is an effective way to change data generation to create more useful probe data.

The second trend is explained by the fact that vehicles traveling straight through the intersection were the least likely to be required to stop and wait—and spend time in the "no-man's land" of the intersection where their PSN might change and disqualify them from creating a traceable turning movement. Left-turning vehicles, having to wait for oncoming traffic, where the most likely to have to stop in the "no-man's land," and thus were the least likely to be detectable. Right-turning vehicles could often turn right without waiting, but still had to slow down, giving them results in between those of straight-through vehicles and left-turning vehicles.

Conclusions

More frequent generation of probe data using PDM allows more turns to be detected. This test demonstrated that PDM is an effective way to alter probe generating behaviors of OBEs to create more useful data at a given intersection.

The results of this test suffer from the "infrequent generation of probe data" issue. Resolving this issue will improve the results of the test.

The frequency of probe data generation during stops has a real influence on the outcome of turning movements measurements. Data should be generated as frequently as practical during stops to make the turning movements most detectible.

The area of the start and turning boxes should include a part of the area in the intersection to improve the likelihood of detecting turning movements. This should be possible without increasing the likelihood of incorrectly detecting turning movements.

More turning movements should be observed to be able to rule out the effects of randomness on the outcome of this test.

Even with all of the issues described here, a detection percentage over 50 percent was recorded for straight-through vehicles using the PDM request on run 3. Thus, the data indicate that with all of these improvements, consistent results well above 50 percent can be achieved.

4.2.6 Same-PSN Vehicle Trajectories: Lane Change (Test Case PDS-06)

Test Case Objectives

- Verify that the application can trace vehicle movements as they change lanes.

Original Data Outputs

- Application-computed same-PSN vehicle traces for all vehicles.

Revised Data Outputs

- Manually computed same-PSN vehicle traces for all vehicles.

Exit Criteria

The test passes if the application correctly constructs the same-PSN vehicle trajectories for all vehicles. The test will not fail due to communication failures occasioned by packet loss, as long as the packet loss does not exceed the value observed in integration and testing, with adjustments made for "real world" traffic situations.

Deviations from Test Cases

The tests were conducted with simulated lane changes. MDOT and RCOC could not guarantee lane closures times and locations.

All vehicles were instructed to perform the same lane changes along the route, as indicated in figure 46 below. This step conflicts with the vehicle configuration described in the test cases, which specifies that four vehicles travel in four separate lanes in parallel. This deviation was made for two reasons—it became clear to the testing staff that it would be difficult for drivers to follow different paths in parallel, and by minimizing the number of different lane change maneuvers, the testers hoped to get more probe data per maneuver, which would allow the detection of lane changes.

Test observers were not asked to record all lane closure time data mentioned in the test cases. Testers felt that it would be very difficult or impossible for observers to record this data in a consistent way, and creating the data would only lead to confusion during data analysis.

Hypotheses

- Lane changes can be detected if enough snapshots with the same PSN are generated at times when vehicles perform a lane change.

Data Analysis Strategy

The data can be analyzed by snapping snapshots of specific roadway lanes. If enough snapshots with the same PSN are attached to separate lanes in locations where drivers are instructed to change lanes, then the lane change can be detected.

Snapping data to travel lanes can be accomplished using the following procedure, developed by Walt Fehr of the VIIC. The following example is computed from position data collected while running the tolling application on June 13, 2008, from the vehicle with OBE C482 in the Michigan DTE. The vehicle makes a lane change as it approaches the Novi Road overpass on I-96 Westbound. In the images, the wavy red lines are vehicle trajectories, the blue lines with circles are the lane center lines, the red lines with circles are the shoulder center lines, and the large black Xs are the locations of the probe data snapshots.

1. For roadway segments where lane-level analysis will be performed, find the lane center line segments on a high-resolution map for each lane. Also find lane center line segments for the road shoulder, which will be used to determine when a vehicle is off the road.

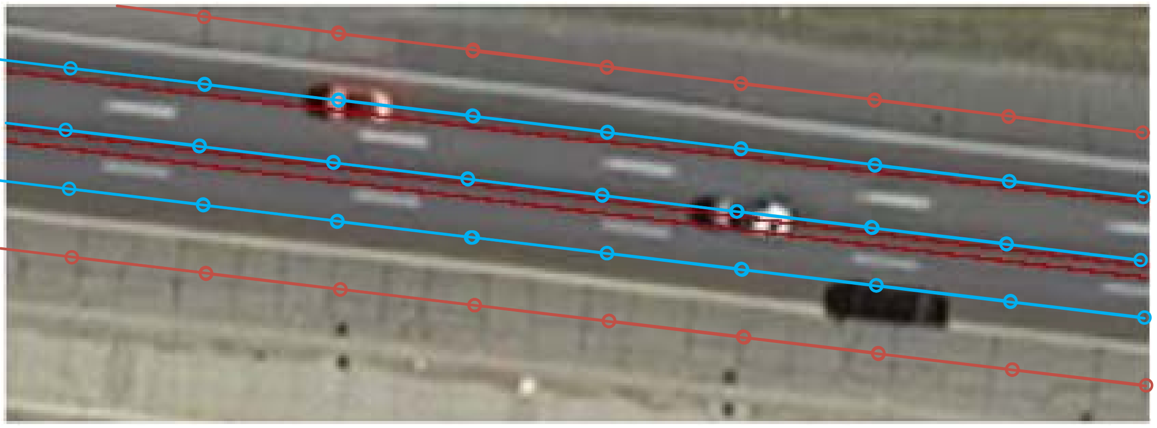

Booz Allen generated routes and plot points on Google map

Figure 41: Lane Change – Determining Center Lines

2. To determine which lane a probe data snapshot occurs in, draw a line through its vehicle location in the snapshot perpendicular to the vehicle heading in the snapshot. Compute the distance to the lane center line segment of the two neighboring lanes (or shoulder). The lane (or shoulder) with the shortest distance on the perpendicular line is the lane to which the snapshot is snapped.

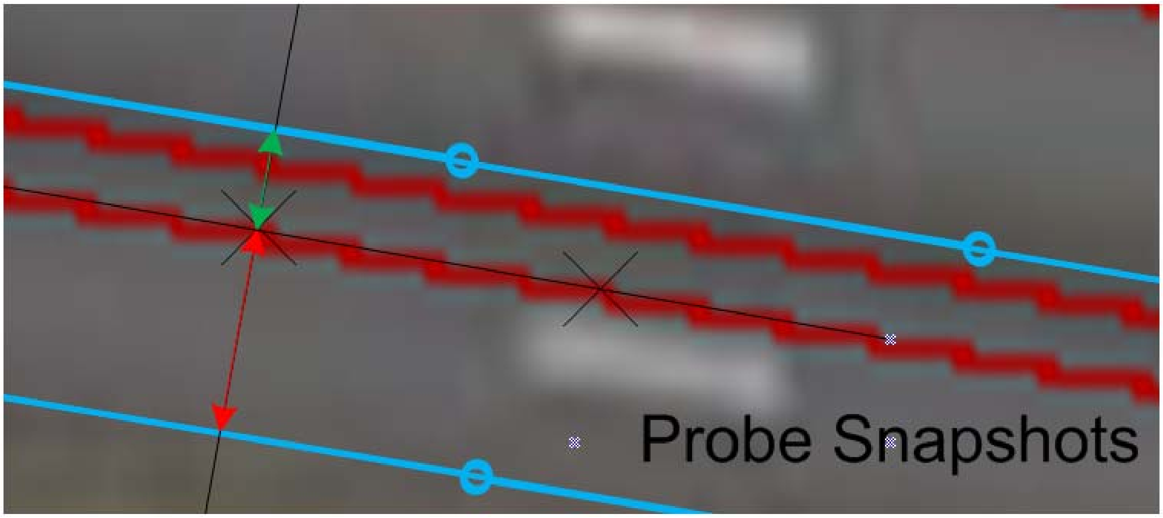

Booz Allen generated routes and plot points on Google map

Figure 42: Lane Change – Snapping Snapshots

3. For this snapshot, the vehicle's recorded position is near the right side of the lane but still closest to the lane center line segment for the lane it is in.

Once snapshots are affixed to a given lane, groups of snapshots with the same PSN will be examined to determine whether a lane change has occurred. Only groups of four or more snapshots will be considered to ensure that the observed movement was truly a lane change, and not a brief swerve into the other lane. All groups of four or more same-PSN snapshots will be considered valid lane changes, but other snapshots will not.

Methods

Drivers were instructed to drive together as a platoon of 25 vehicles and to change lanes in specific areas. The following sheet (shown in Figure 43) was given to drivers to explain where they should change lanes.

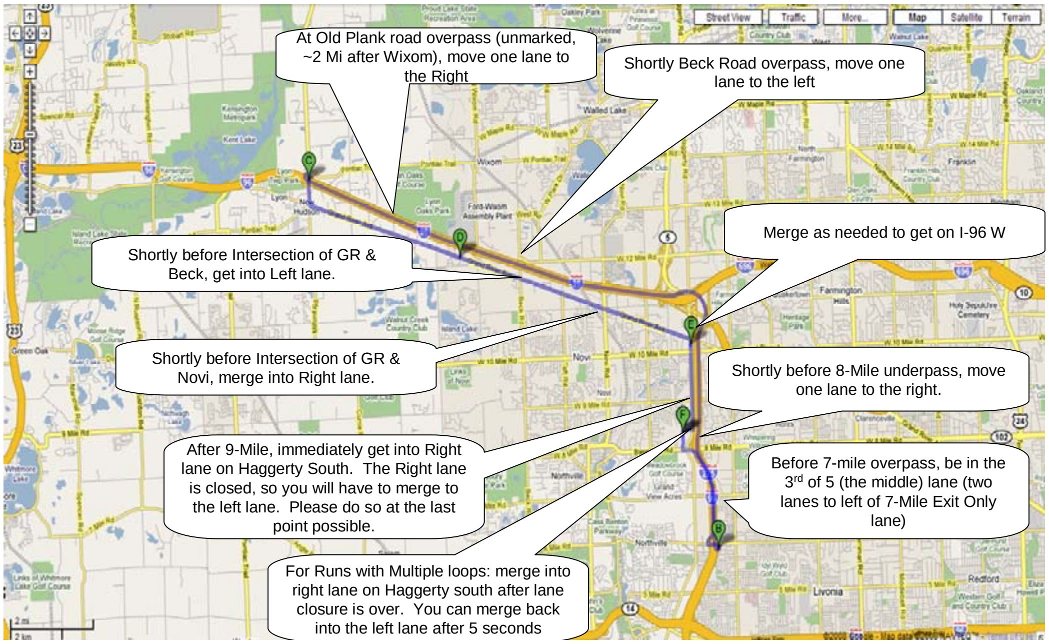

Booz Allen generated routes and plot points on Google map

At Old Plank road overpass (unmarked, ~2 Mi after Wixom), move one lane to the Right

Shortly Beck Road overpass, move one lane to the left

Merge as needed to get on I-96 W

Shortly before Intersection of GR & Beck, get into Left lane.

Shortly before Intersection of GR & Novi, merge into Right lane.

Shortly before 8-Mile underpass, move one lane to the right.

After 9-Mile, immediately get into Right lane on Haggerty South. The Right lane is closed, so you will have to merge to the left lane. Please do so at the last point possible.

Before 7-mile overpass, be in the 3rd of 5 (the middle) lane (two lanes to left of 7-Mile Exit Only lane)

For Runs with Multiple loops: merge into right lane on Haggerty south after lane closure is over. You can merge back into the left lane after 5 seconds

Figure 43: Lane Change Areas

Three runs of the test case were executed over the course of the testing day.

Table 28: Lane Change Test Times

Run	Start	End	Total Loops
Run 1	10:06 AM	11:46 AM	2
Run 2	12:38 PM	2:18 PM	2
Run 3	2:35 PM	4:14 PM	2

Results and Analysis

To determine the areas on the test route most likely to have visible lane changes, a visual inspection of the data points on Google Earth was performed. The segment of Grand River Avenue between Wixom Road and Market Street was the part of the route with the biggest clusters of data points around areas where lane changes were specified to occur. Using the technique described above, all snapshots were snapped to a lane in this roadway segment.

Note that the above driver instruction sheet tells drivers to drive in certain lanes of Grand River between Wixom Road and Novi Road: “Shortly before Intersection of GR & Beck, get into Left lane,” and “Shortly before Intersection of GR & Novi, merge into Right lane.”

Table 29: Number of Snapshots in Lanes or Off-Road

Location	#Snapshots Snapped to Right Lane	#Snapshots Snapped to Left Lane	#Snapshots Snapped to Off Road	Notes
1. Wixom to second entrance before Beck Rd	7	2	0	Vehicles stay mostly in Right lane as instructed
2. Second entrance before Beck Rd to Beck Rd	5	3	3	Vehicles begin shift to left lane, which occurs at intersection with Beck
3. Beck Rd to 100m before Novi Rd	9	86	8	Vehicles remain mostly in left lane as instructed
4. 100m before Novi Rd to Novi Rd	25	10	20	Vehicles move to right lane as instructed
5. Novi Rd. to Market St.	86	84	24	Vehicles begin return to left lane because right lane ends after Market St.

The lanes that the snapshots were snapped to fit the expected pattern on Grand River between Wixom Road and West 11 Mile Road.

1. Between Wixom Road and the second entrance into the hospital prior to Beck Road, all of the snapshots but two were in the right lane.
2. Between the second entrance to the hospital prior to Beck Road and Beck Road, there were three data points off of the roadway. There were now three data points in the left lane, with the rest still in the right lane.
3. Between Beck Road and a point 100 meters prior to Novi Road, the vast majority of data points were in the left lane. There was a small number in the right lane and a similar small number off of the roadway.
4. Between that 100-meter point and Novi Road, there was a noticeable shift in the data points from the left lane to the right lane. There were also clusters of data points off of the roadway.
5. Between Novi Road and West 11 Mile Road, the vast majority of data points were in the right lane. There were several in the left lane or off of the roadway. After West 11 Mile Road, there was a noticeable shift to the left lane. This lane change was not specified in the driver instructions, but was necessary because the right line of Grand River Eastbound ends 150 meters after the Market Street intersection. As the vehicles reached Market Street, approximately 80 percent of the data points were in the left lane. There was also an unusual group of data points well off of the roadway.

There were 13 probe segment numbers with 4 or more probe data snapshots in the group for the entire day in the roadway segment of Grand River between Wixom and Market Street. Only four of these groups showed lane changes:

- Two groups showed the left lane to right lane change that is expected prior to Novi Road (point 4 from the list above).
- Two groups showed the lane change from the right lane to the left lane after Novi Road (point 5 from the list above).

Of the nine groups of four or more, two of the PSN groups contained data points that were off of the roadway, one group had all of the points in the left turn lane, and one group had all of the points off to the right of the road. Seven groups had all of the data points in one lane, so no lane change could be detected.

Conclusions

Given that the positional information reported in probe data snapshots lacked lane-level accuracy, this technique used to snap snapshots to lanes works well as an alternative for purposes of this analysis.

Individual vehicle lane changes cannot be detected because of the sparseness of the data points. There were no observable “tracks” in the data with same-PSN snapshots in both lanes, illustrating a lane change. Only four lane changes could be observed from the entire day of testing.

Gross behaviors of vehicles could be detected. If the roadway is divided into segments and the number of data points collected during a given time period were compared to historical data for similar time periods, an unusual event might be detected. In the data, the shift to the left lane immediately after Novi Road was not explained by the instructions given to the drivers. One would not have expected it so soon since the narrowing of Grand River to two lanes does not occur until well after Market Street.

4.3. Advisory Message Delivery Service

AMDS testing was accomplished in all three phases of the VII test program. The scope of AMDS testing was to show that Network Users could generate AMDS messages with varying broadcast strategies to send to RSEs, and that OBEs could receive and correctly display those messages according to the message instructions. Testing was conducted by both the VIIC and Booz Allen teams at the CNSI Lab, Battelle and VIIC test tracks, and DTE.

Initial testing with OBEs and a single RSE successfully demonstrated AMDS functionality; however, as testing moved to multiple RSEs with overlapping coverage areas, issues arose. Initial tests found multiple cases of OBE instability and crashing when in range of multiple RSEs. Troubleshooting of the issue resulted in a new OBE build that fixed this issue, but led to the discovery that the DSRC standard does not adequately address instances where an OBE is in range of multiple RSEs at the same time. This led to the key finding that if an OBE enters an area with multiple RSEs broadcasting AMDS messages on the service channel, it may never receive the AMDS messages from one of the RSEs. The issue results from the OBE joining the WBSS of one of the RSEs and not being able to exit that WBSS and join the WBSS of another RSE unless a message of higher priority is being advertised by the second RSE. No mechanism currently exists within DSRC to mitigate this issue, although the DSRC community is aware and working on a solution. Another finding during AMDS testing was that control channel messages that were advertised in the WSA caused instability in the OBE resulting in a software crash. Since WSAs should advertise only service channel messages, a fix was developed for the RSE to that effect.

During AMDS testing, multiple broadcast strategies were used, which varied the broadcast region, presentation region, message activation time, and priority. All messages that were received by the RSEs were correctly broadcast by them and displayed by the OBEs according to the message instructions. Additionally, currently active messages were correctly updated or deleted, and those changes were accurately reflected by both the RSE and OBE. Table 30 summarizes the key findings in this area.

Table 30: Key Findings for AMDS

Finding Number	Finding
F-AMDS-1	If an OBE enters an area with multiple RSEs broadcasting AMDS messages on the service channel, it may never receive the AMDS messages from one of the RSEs.
F-AMDS-2	Control channel messages that were advertised in the WSA caused instability in the OBE resulting in a software crash. Since WSAs should only contain service channel advertisements, a fix was developed for the RSE to that effect.
F-AMDS-3	During AMDS testing, multiple broadcast strategies were used which varied the broadcast region, presentation region, message activation time, and priority. All messages that were received by the RSEs were correctly broadcast by them and displayed by the OBEs according to the message instructions. Additionally, already active messages were updated and deleted, and those changes were accurately reflected by both the RSE and OBE.

The following sections provide a detailed analysis of the results from final testing of AMDS during the public applications integration testing carried out in Detroit, Michigan, between August 18, 2008, and September 12, 2008.

4.3.1. Broadcast Strategies (Test Case AMDS-03)

Public signage testing was conducted on August 18 and August 19. A total of 41 messages were sent with varying parameters. Of those, 35 made it to the designated RSEs. Of the remaining six messages, four were sent to FW5 and FW3 RSEs, which had problems and did not receive any messages. The remaining two messages were sent to RSE 84 (Grand River and Novi), which was down.

Test Case Objectives

- Test the application’s capability to send advisory messages based on multiple combinations of broadcast strategies (e.g., measure type, location, RSEs, and valid region).
- Test the application’s ability to prioritize between public-sector advisory messages and civil emergency notifications.

Data Output

- Comparison of received message logs with transmitted message logs, accounting for broadcast strategies.

Exit Criteria

The test passes: (1) if the traveler advisory messages received by vehicles match the traveler advisory messages that have been transmitted, and vehicles do not receive messages that they should not have received, but do receive all messages that they should have received, appropriate to the broadcast strategy; (2) if the simulated civil emergency message is given priority over the traveler information messages; and (3) if the message is successfully cancelled. The test will not fail due to incomplete data reception if the Application Database comparison with OBE logs indicates incomplete transmission reception from the OBE, such as a communication cutoff due to a truck driving in front of an OBE antenna, as long as packet losses reflect the level of packet loss experienced in system integration testing, with adjustments for real-world conditions.

Deviations from Test Cases

Heavy traffic and precipitation were not encountered during testing.

Hypothesis

1. All messages sent from Public Apps were received by the OBEs.
2. Roundtrip latency from Public Apps to RSEs should be a matter of a few seconds.
3. Message durations were in line with operator-specified values.
4. All messages were only displayed to the driver when the vehicle was in a valid active region.
5. Only RSEs defined in the broadcast strategy played the messages.
6. The correct broadcast priority was received by the RSEs.
7. Priority order was implemented correctly and highest priority messages were displayed on top.
8. Messages were received, decoded, and displayed correctly.

Data Analysis Strategy

To verify the original test objective, it is necessary to know how each aspect of the broadcast strategy behaved. The following eight measures were created to ascertain whether broadcast strategies were implemented correctly:

1. **Loss** – What percentage of messages made it from the Public Apps to the OBE?
2. **Latency** – How fast were messages received by the RSE from the Public Apps?
3. **Activation/Expiration of Messages** – Did messages become active and expire on time, as defined by the operator?
4. **Active Region and Direction** – Were messages displayed in the correct active region and only to the vehicles traveling in the direction specified by the operator?
5. **RSE Selection within Broadcast Strategy** – Did only RSEs selected in the broadcast strategy receive the messages?
6. **Broadcast Priority within Broadcast Strategy** – Did the RSEs receive the correct broadcast priority?
7. **OBE Prioritization of Messages** – Did the OBEs prioritize messages correctly?

8. **Encoding and Decoding of Messages** – Was the content of the message correctly encoded by the application and decoded by the OBE?

Routes and RSEs

Table 31 lists the RSEs that were used for testing.

Table 31: List of RSEs

RSE Site ID and Name	Backhaul	RSE Number
82 (Grand River Ave @ Milford Rd)*	3G	0024
84 (Grand River Ave @ Novi Rd)*	WIMAX	0056
FW3 (I-96 @ Beck Rd)	WIMAX	0053
544 (9 Mile Rd @ Haggerty Rd)*	T1	0062
FW5 (I-96 @ Novi Rd)	WIMAX	0069
FW6 (I-96 @ Rt-5)	WIMAX	0041
FW7 (I-96 @ I-275)	WIMAX	0045
FW9 (I-275 @ 8 Mile Rd)	WIMAX	0040
355 (10 Mile Rd @ Haggerty Rd)	T1	0026
FW10 (I-275 @ 7 Mile Rd)**	WIMAX	0010
FW11 (I-275 @ 6 Mile Rd)**	WIMAX	0023
444 (Grand River Ave @ Haggerty Rd)	T1	0005
991 (8 Mile Rd @ Haggerty Rd)	T1	0034
213 (Grand River Ave @ Beck Rd)*	3G	0078
857 (Grand River @ Napier Rd (Clover)*	3G	0054
180 (Grand River Ave @ Market St)	3G	0066
481 (Grand River @ Meadowbrook Rd)	T1	0082
992 (8 Mile Rd @ Orchard Hill Pl)	T1	0055

The following Google Earth screenshot (Figure 44) shows all of the above RSEs on the route that was used.

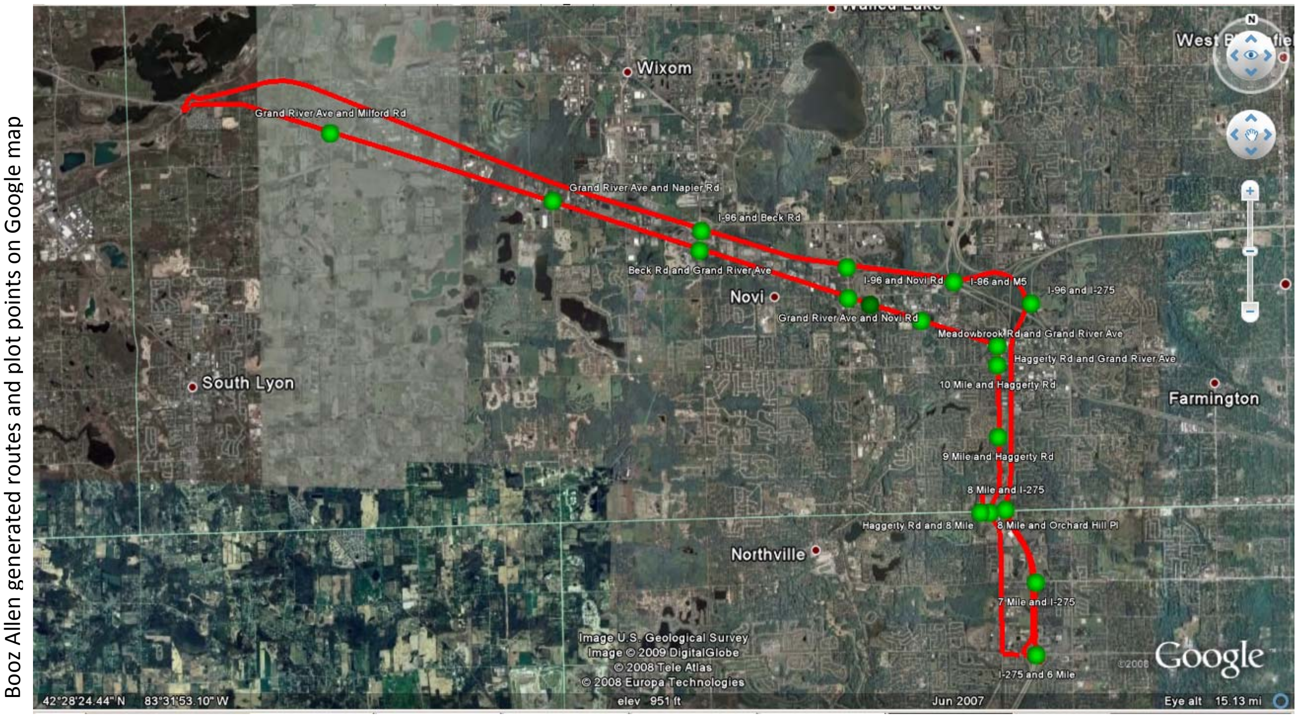

Figure 44: All RSEs on Test Route

On both days, the same two routes were driven as shown in the following figures.

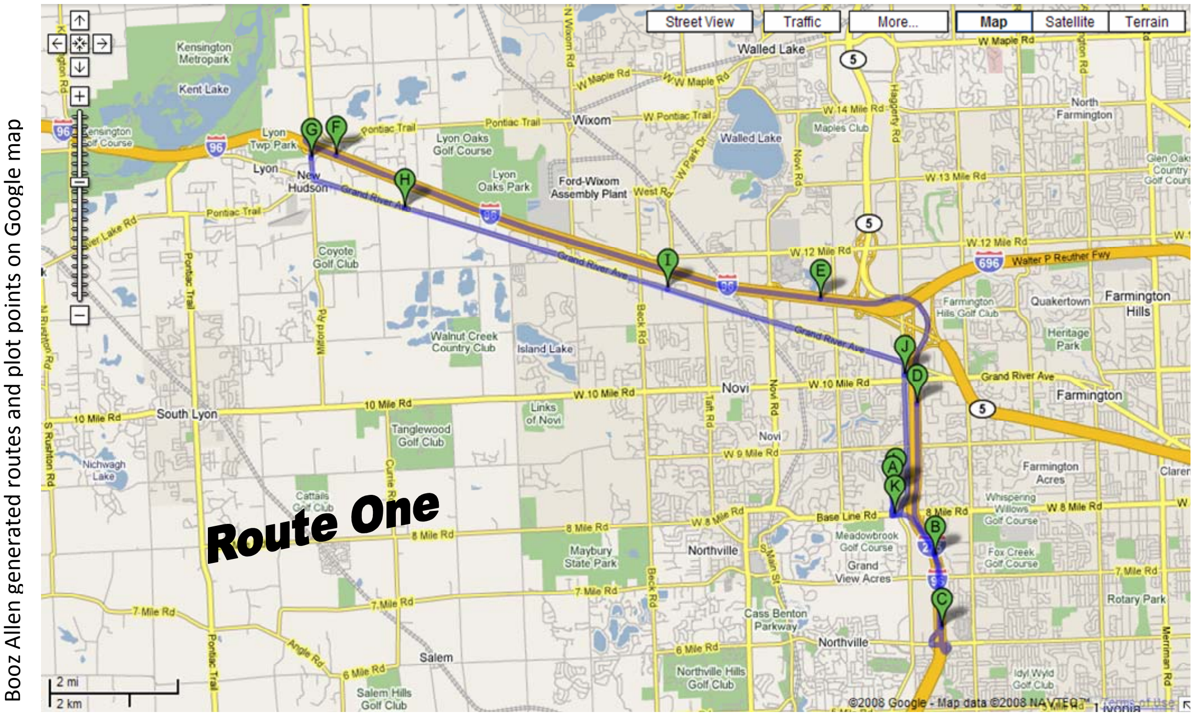

Figure 45: Route 1

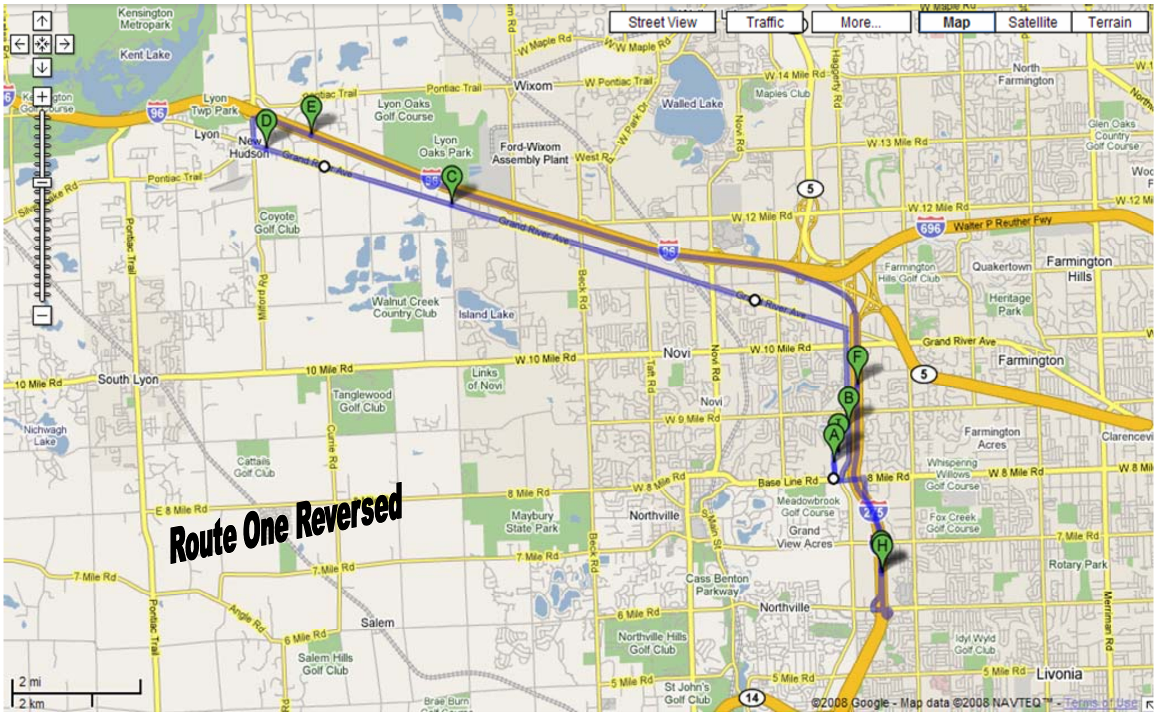

Booz Allen generated routes and plot points on Google map

Figure 46: Route 2

The routes look similar because the same route was used but the directions of some vehicles were reversed. The bubbles on the routes only indicate the direction of travel starting at “A” through “H” or “J” (depending on the route).

Note about OBE Logs

On both days of testing, data was received from all 25 vehicles. However, on August 19 at 11:03 am, 10 vehicles (listed in Table 32) were taken away from testing and used for the working group demonstrations.

Table 32: Vehicles for Demo

OBE ID	Vehicle
B450	E1
B042	E2
C453	C1
C482	P1
C594	A2
B193	E3
B856	C2
C862	P2
C779	A6
C545	P11
C694	P3

Note about RSE Logs

Due to an oversight by the test coordinator, logging on the RSEs were not turned on for August 18; therefore, only the RSE logs for August 19 were available for analysis.

The following messages were active on August 19.

Table 33: Active AMDS Messages on August 19

MSG ID	DESCRIPTION	RSE_LIST	ACT.	EXP.	PRIORITY	PSID
1334	traffic flowing freely I-275 & 6-Mile traffic flowing freely 0023-1	us.mi.csnv.rse.0029 us.mi.csnv.rse.0023	8/17/2008 0:00	9/14/2008 0:00	1	34603009
1335	heavy traffic I-275 & 6-Mile heavy traffic 0023-2	us.mi.csnv.rse.0029 us.mi.csnv.rse.0023	8/18/2008 17:00	8/19/2008 17:00	6	34603009
1333	incident cleared I-275 & 7-Mile incident cleared 0010-2	us.mi.csnv.rse.0010	8/17/2008 0:00	9/13/2008 0:00	6	33619969
1329	delays I-275 & 8-Mile delays 0040-1	us.mi.csnv.rse.0040	8/17/2008 0:00	9/14/2008 0:00	1	34603009
1330	delays cleared I-275 & 8-Mile delays cleared 0040-2	us.mi.csnv.rse.0040	8/18/2008 17:00	8/19/2008 17:00	1	34603009
1337	closed 9 & Haggerty closed 0062-1	us.mi.csnv.rse.0062 us.mi.csnv.rse.0046	8/17/2008 0:00	9/14/2008 0:00	8	33685505
1306	delays I-96 & Rt5 delays 0041-1	us.mi.csnv.rse.0041	8/16/2008 0:00	8/23/2008 0:00	3	34603009
1307	delays cleared I-96 & Rt5 delays cleared 0041-2	us.mi.csnv.rse.0041	8/16/2008 0:00	8/23/2008 0:00	8	34603009
1308	incident I-96 & Novi incident 0077-1	us.mi.csnv.rse.0077	8/16/2008 0:00	8/23/2008 0:00	4	33619969
1309	incident cleared I-96 & Novi incident cleared 0077-2	us.mi.csnv.rse.0077	8/16/2008 0:00	8/23/2008 0:00	7	33619969
1310	traffic flowing freely I-96 & Beck traffic flowing freely 0075-1	us.mi.csnv.rse.0075	8/16/2008 0:00	8/23/2008 0:00	5	34603009
1311	traffic building I-96 & Beck traffic building 0075-2	us.mi.csnv.rse.0075	8/16/2008 0:00	8/23/2008 0:00	6	34603009

MSG ID	DESCRIPTION	RSE_LIST	ACT.	EXP.	PRIORITY	PSID
1312	heavy traffic GR & Milford heavy traffic 0024-1	us.mi.csnv.rse.0024	8/16/2008 0:00	8/23/2008 0:00	6	34603009
1313	traffic congestion GR & Milford traffic congestion 0024-2	us.mi.csnv.rse.0024	8/16/2008 0:00	8/23/2008 0:00	5	34603009
1314	delays cleared GR & Napier delays cleared 0054-1	us.mi.csnv.rse.0054	8/16/2008 0:00	8/23/2008 0:00	7	34603009
1315	delays GR & Napier delays 0054-2	us.mi.csnv.rse.0054	8/16/2008 0:00	8/23/2008 0:00	4	34603009
1318	incident cleared GR & Beck incident cleared 0078-1	us.mi.csnv.rse.0078	8/16/2008 0:00	8/23/2008 0:00	8	33619969
1319	construction work I-96 at Rt5, Novi & Beck construction work 417775-3	us.mi.csnv.rse.0075 us.mi.csnv.rse.0077 us.mi.csnv.rse.0041	8/16/2008 0:00	8/19/2008 14:00	6	34603009
1320	closed GR at Milford, Napier & Beck closed 245478-3	us.mi.csnv.rse.0078 us.mi.csnv.rse.0024 us.mi.csnv.rse.0054	8/18/2008 18:00	8/23/2008 0:00	6	34603009
1344	traffic congestion GR & Market traffic congestion 0066-2	us.mi.csnv.rse.0082	8/17/2008 0:00	9/14/2008 0:00	1	34603009
1322	delays GR & Meadow delays 0082-2	us.mi.csnv.rse.0082	8/16/2008 0:00	8/23/2008 0:00	2	34603009
1321	delays GR & Meadow delays 0082-1	us.mi.csnv.rse.0082	8/16/2008 0:00	8/23/2008 0:00	3	34603009
1323	construction work GR & Meadowbrook construction work 000082-3	us.mi.csnv.rse.0082	8/16/2008 0:00	8/23/2008 0:00	6	34603009
1325	traffic building 8Mile & Orchard traffic building 0055-2	us.mi.csnv.rse.0055	8/16/2008 0:00	9/13/2008 0:00	1	34603009
1350	travel time 8 & Haggerty travel time 0034-2	us.mi.csnv.rse.0034	8/18/2008 0:00	9/14/2008 0:00	1	34603009
1326	8 Mile & Haggerty 8 Mile & Haggerty travel time 0034-2	us.mi.csnv.rse.0034	8/16/2008 0:00	9/20/2008 0:00	1	34603009

MSG ID	DESCRIPTION	RSE_LIST	ACT.	EXP.	PRIORITY	PSID
1328	closed 8Mile at Hagg, Orch closed 5534-3	us.mi.csnv.rse.0055 us.mi.csnv.rse.0034	8/17/2008 0:00	9/14/2008 0:00	8	33685505
1336	icy patches I-275 6 to 8 Mile icy patches 401023-3	us.mi.csnv.rse.0029 us.mi.csnv.rse.0010 us.mi.csnv.rse.0040 us.mi.csnv.rse.0023	8/17/2008 0:00	9/14/2008 0:00	3	34603009
1340	travel time 10 & Haggerty travel time 0026-2	us.mi.csnv.rse.0026	8/17/2008 0:00	9/14/2008 0:00	1	34603009
1342	incident cleared GR & Haggerty incident cleared 0005-2	us.mi.csnv.rse.0005	8/17/2008 0:00	9/14/2008 0:00	6	33619969
1341	incident GR & Haggerty incident 0005-1	us.mi.csnv.rse.0005	8/17/2008 0:00	9/21/2008 0:00	6	33619969
1348	icy patches Hagg GR to 9Mile icy patches 052662-3	us.mi.csnv.rse.0062 us.mi.csnv.rse.0046 us.mi.csnv.rse.0026 us.mi.csnv.rse.0005	8/17/2008 0:00	9/14/2008 0:00	3	33685505
1345	delays GR & Novi delays 0056-1	us.mi.csnv.rse.0079	8/17/2008 0:00	9/14/2008 0:00	1	34603009
1346	delays cleared GR & Novi delays cleared 0056-2	us.mi.csnv.rse.0079	8/17/2008 0:00	9/14/2008 0:00	1	34603009
1343	traffic building GR & Market traffic building 0066-1	us.mi.csnv.rse.0066	8/17/2008 0:00	9/14/2008 0:00	6	34603009
1349	traffic congestion GR & Market traffic congestion 0066-2	us.mi.csnv.rse.0066	8/18/2008 0:00	9/14/2008 0:00	6	34603009
1347	construction work GR & Novi-Mkt construction work 5666-3	us.mi.csnv.rse.0066 us.mi.csnv.rse.0079	8/18/2008 5:00	8/19/2008 5:00	8	33685505

MSG ID	DESCRIPTION	RSE_LIST	ACT.	EXP.	PRIORITY	PSID
1362	civil emergency civil emergency	us.mi.csnv.rse.0029 us.mi.csnv.rse.0075 us.mi.csnv.rse.0066 us.mi.csnv.rse.0062 us.mi.csnv.rse.0078 us.mi.csnv.rse.0024 us.mi.csnv.rse.0046 us.mi.csnv.rse.0079 us.mi.csnv.rse.0054 us.mi.csnv.rse.0055 us.mi.csnv.rse.0082 us.mi.csnv.rse.0026 us.mi.csnv.rse.0061 us.mi.csnv.rse.0034 us.mi.csnv.rse.0005 us.mi.csnv.rse.0010 us.mi.csnv.rse.0040 us.mi.csnv.rse.0039 us.mi.csnv.rse.0023 us.mi.csnv.rse.0041	8/19/2008 12:45	8/19/2008 13:00	1	33751041
1361	civil emergency civil emergency	us.mi.csnv.rse.0029 us.mi.csnv.rse.0075 us.mi.csnv.rse.0066 us.mi.csnv.rse.0062 us.mi.csnv.rse.0078 us.mi.csnv.rse.0024 us.mi.csnv.rse.0046 us.mi.csnv.rse.0079 us.mi.csnv.rse.0054 us.mi.csnv.rse.0055 us.mi.csnv.rse.0082 us.mi.csnv.rse.0026 us.mi.csnv.rse.0061 us.mi.csnv.rse.0034 us.mi.csnv.rse.0005 us.mi.csnv.rse.0010 us.mi.csnv.rse.0040 us.mi.csnv.rse.0039 us.mi.csnv.rse.0023 us.mi.csnv.rse.0041	8/19/2008 12:00	8/19/2008 12:15	1	33751041
1359	travel time I-96 West to Kent Lake Rd travel time 12 minutes	us.mi.csnv.rse.0052	8/18/2008 0:00	8/20/2008 0:00	1	34603009

MSG ID	DESCRIPTION	RSE_LIST	ACT.	EXP.	PRIORITY	PSID
1358	civil emergency Novi Michigan Amber Alert civil emergency	us.mi.csnv.rse.0052	8/18/2008 0:00	8/20/2008 0:00	1	33751041
1357	civil emergency Novi Michigan Amber Alert civil emergency	us.mi.csnv.rse.0049	8/18/2008 0:00	8/20/2008 0:00	1	33751041
1360	accident Grand River Ave at Novi Rd. accident Proceed with Caution	us.mi.csnv.rse.0052	8/18/2008 0:00	8/20/2008 0:00	6	33619969
1356	accident Grand River at Novi Rd accident Proceed with caution	us.mi.csnv.rse.0049	8/18/2008 0:00	8/20/2008 0:00	6	33619969

Measure 1: Loss of Public Messages

Measure Objective: Determine whether all messages sent by Public Apps were received by OBEs.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Using the OBE logs, find out if every message appeared at least one time in one of the OBE logs.

Results: 100 percent of all messages that were received by the RSEs were also received by the OBEs. Almost 15 percent of messages did not make it to the OBEs. They were due to RSE being down and/or incorrectly configured RSEs.

Based on the above results, as long as the RSEs are online and correctly configured, it is reasonable to assume that all messages, including high priority civil emergency messages, will make it to the RSE and ultimately to the OBE.

Measure 2: Roundtrip Latency of Messages from Public Apps to RSE

Measure Objective: Determine the latency between the times the Public Apps sent out the signage message, it was received by the RSE, and the Public Apps received a reply from the RSE.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Only the ones that received a response from the RSE should be counted. Also, since the Public Apps database does not record exact sent time and RSE acknowledgement receipt time, the data in the AMDS broker database, which is only one step removed from the Public Apps, was used.

Results: Submissions of advisory message requests to the AMDS were immediately acknowledged. The average roundtrip latency was 3 minutes and 17 seconds for submitted messages to be propagated to target RSEs for subsequent broadcasts, which appears low enough to handle all types of messages, including civil emergencies. However, some responses from RSEs did not come back for several hours. This was expected for RSEs that were down when messages were sent and only came online several hours later. As soon as they register with the broker and receive the message, they send back an acknowledgment message, which will have a latency of several hours. Further detailed analysis is required to assess why some RSEs did not respond with an acknowledgment message, even though they did receive the messages.

Measure 3: Messages Became Active and Expired at Times Specified by the Operator

Measure Objective: Determine whether the activation and expiration times of messages on the OBEs were the same as specified by the operator.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in the OBE logs that the messages display time falls within the active time frame.

Results: 99.95 percent of messages were displayed and removed at times specified in the AMDS message. Out of 2,146 instances when a message was displayed to the driver, only once was it after the expiration time, and it was only off by 12 seconds. This could be a possible OBE glitch. It is important for an operator to be able to control exactly when a message becomes active and when it expires. Given the results, they can be highly confident that messages will be displayed and removed from the OBEs on time.

Measure 4: Messages Were Only Displayed within Correct Active Regions and Direction

Measure Objective: Determine whether the messages were only displayed within the active region and removed once the vehicle exited that active region.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: The percentage of HMI displays is calculated as follows: ((Valid Display + Valid Reject)/ (Valid Display + Valid Reject + Invalid Display + Invalid Reject)) * 100.

Results: 100 percent of messages were displayed in the correct activation regions and 100 percent of messages that were received but should not have been displayed, due to directionality and/or activation time, were not displayed.

Out of the 44 messages received, 27 were correctly displayed and 17 correctly rejected.

It is very important that drivers only receive messages that apply to them, in order not to overload them with unnecessary information. To do this, the operator needs to be able to specify exact presentation location and direction of travel for vehicles. The results show that both active regions and directionality are implemented correctly within the OBEs.

Measure 5: Only RSEs Defined in the Broadcast Strategy Played the Messages

Measure Objective: Determine whether the messages were only played on RSEs that were defined in the broadcast strategy.

Data Source: RSE PM logs.

High-Level Data Analysis Plan: Count the messages that were sent out on all RSE logs, and then verify that only RSEs that should have received a certain message only show those messages and not others.

Results: In a fully rolled out VII implementation, there will be a lot of traffic going through the network, and it will be imperative that the system is not overloaded with unnecessary information.

The tests results showed that 100 percent of broadcast strategies were implemented correctly. The messages were only played on RSEs defined by the broadcast strategy, and messages that expired on August 18 were not played on RSEs on August 19.

Measure 6: Did the RSE Receive the Correct Broadcast Priority (0-7) as Sent by the Operator

Measure Objective: Determine whether the correct broadcast priority was received by the RSEs.

Data Source: RSE PM logs.

High-Level Data Analysis Plan: In the RSE PM logs, “Message Priority” refers to broadcast priority. Using the message priority and message call back ID, verify that correct broadcast priorities were received by the RSEs.

Results: Broadcast priorities define how often a message is played on the RSE playlist. A priority of 0 means that a message will be played every 100 millisecond, and a priority of 7 means that a message will be played every 100,000 millisecond. The broadcast priority is also used to order the playlists; lower numbers have higher priorities—so, for example, a civil emergency message with a priority of 0 would be played before a travel time message is played.

One hundred percent of all broadcast strategy priorities were received, decoded, and successfully implemented by the RSEs.

Measure 7: Did the OBE Display the Highest Priority Messages First

Measure Objective: Determine whether the messages with the highest priority were displayed on top of other messages.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify for at least 1 day that all lower-number (higher priority) messages were displayed before the higher-number messages (lower priority). Verify that the prioritization performed by the OBE, and logged by the OBE, resulted in the display of messages in an ascending order of priority.

Results: Once the OBE receives a list of messages, it has to prioritize which one will be displayed on top of the other messages. For example, an Amber Alert will take precedence over speed limit signs.

The results showed that 100 percent of messages were prioritized successfully.

Measure 8: Message Content – Did the OBE Display the Message Content (e.g., International Traveler Information Systems (ITIS) codes and free text) as Sent by the Operator

Measure Objective: Determine whether the message content stayed intact during transmission to the OBE and subsequent decoding.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Using the OBE logs, determine whether the OBE received the correct ITIS code, category, and free text components of the messages.

Results: Just as important as message priority is message content. Drivers have to see exactly what the operator intended for them to see. A public signage consists of a series of ITIS codes and free text. The results showed that 100 percent of the message content was successfully received and decoded by the OBE and ultimately displayed to the drivers.

Conclusion

The results of the tests showed an overall successful implementation of AMDS. The following is a list of the hypotheses and their conclusions:

1. *All messages sent from Public Apps were received by the OBEs* –Verified by the results; 100-percent success.
2. *Roundtrip latency from Public Apps to RSEs should be a matter of a few seconds* – Contradicted by the results; over 3-minute roundtrip latency.
3. *Message durations were in line with operator specified values* – Verified by the results; 99.9-percent success.
4. *All messages were only displayed to the driver when the vehicle was in a valid active region* –Verified by the results; 100-percent success.
5. *Only RSEs defined in the broadcast strategy played the messages* –Verified by the results; 100-percent success.
6. *The correct broadcast priority was received by the RSEs* –Verified by the results; 100-percent success.
7. *Priority order was implemented correctly and highest priority messages were displayed on top* – Verified by the results; 100-percent success.
8. *Messages were received, decoded and displayed correctly* –Verified by the results; 100-percent success.

4.4. Communications Service

Throughout the system life cycle of the VII POC, various backhaul technologies were employed to provide VII RSE and OBE devices as well as VII partners and stakeholders with access into the VII POC environment. Infrastructure-to-infrastructure and vehicle-to-infrastructure communications benefited from the use of multiple communication mechanisms such as T-1 data-circuits, a closed WiMAX network, and various 3G connections. Each technology was leveraged to support infrastructure where the other means of communication were impossible or cost prohibitive.

Backhaul connectivity to the various VII subsystem locations (Virginia DTE, Michigan DTE, and California SDN) utilized DS-3 interconnections to create a true VII Community of Interest (COI) Network. Furthermore, these communication technologies also supported the dual stack IPv4/IPv6 nature of VII architecture. It should also be noted, that for VII stakeholders without the ability to support an in-house IPv6 capability, utilizing more than 40 individual site-to-site IPSEC virtual private networks (VPNs) and IPv6 to IPv4 tunnels, the VII COI became an ad hoc IPv6 service provider.

However, climate and environmental conditions at various RSE deployment sites introduced new challenges with each implementation. Line-of-sight issues and interference (usually caused by the weather) between the WiMAX subscriber and base stations caused bandwidth fluctuations and RSE degradation during POC testing.

Furthermore, VII operations staff identified ongoing issues with the 3G backhaul that plagued the infrastructure as a whole. Bandwidth congestion and contention during times of high vehicle traffic (rush hour) left 3G RSEs practically inoperable. The 3G service provider routers had major limitations in their support of the IPv6 protocol stack and basic security options such as VPN interoperability, SSL certificate capabilities, and the inability to allow the integration of basic non-proprietary routing protocols, which effectively mitigated the ability to communicate routing information in a secure manner.

Finally, as systems interfaces stabilized and basic security issues and workarounds were created, the ability to use the WiMAX and 3G technologies improved to that of the traditional T-1 land lines.

Figure 47 depicts the backhaul technologies leveraged to support the VII POC and vehicle-to-infrastructure communication.

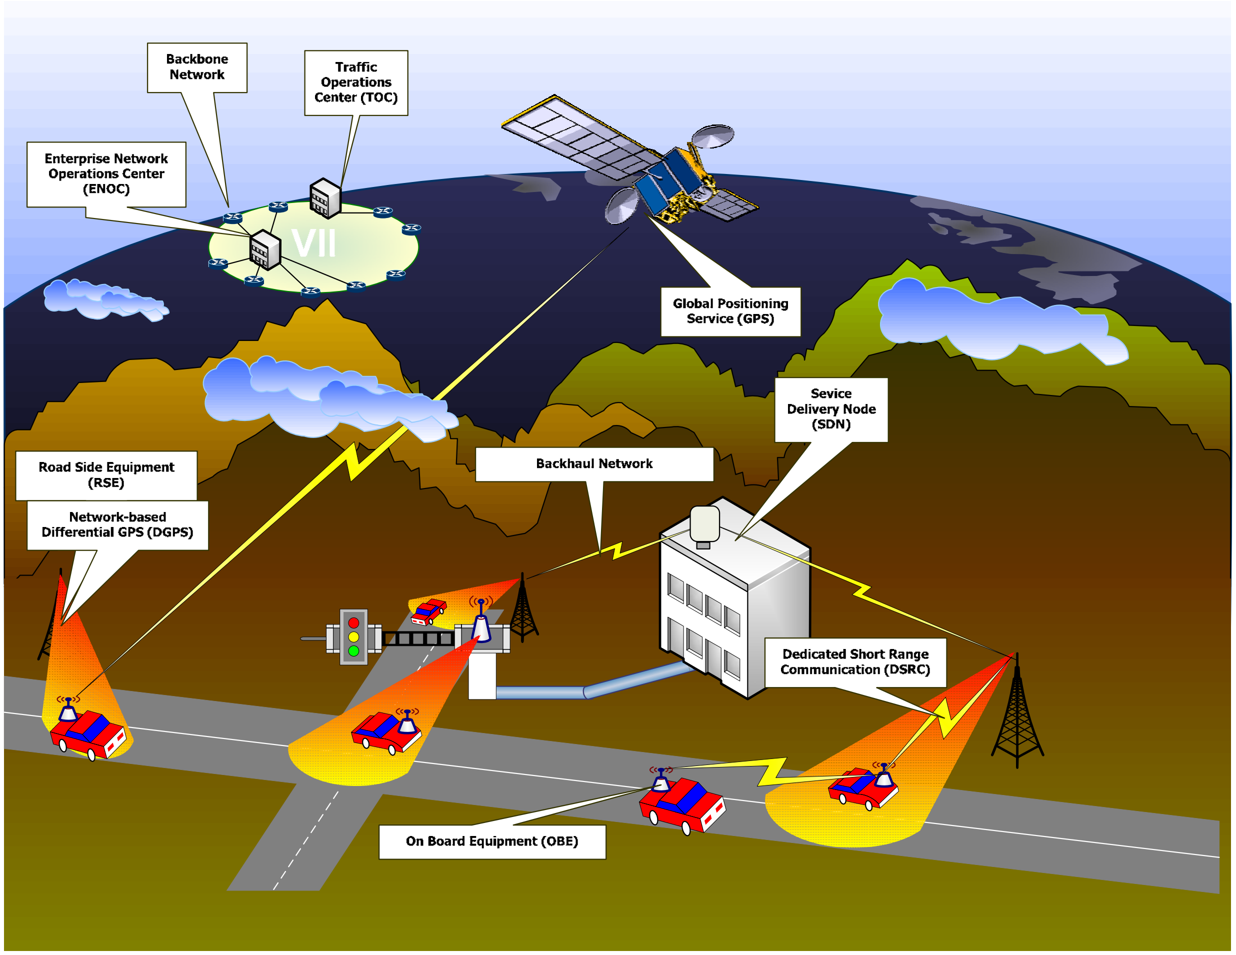

Figure 47: POC Backhaul Technologies

Table 34 summarizes the key findings in this area.

Table 34: Key Findings for Communications Service

Finding Number	Finding
F-COMM-1	Communication technologies supported the dual stack IPv4/IPv6 nature of VII architecture.
F-COMM-2	Management of network communications resources for multiple simultaneous applications is more complex than expected.

4.5. Map Element Distribution Service

MEDS testing was primarily accomplished in the subsystem testing phase. The Booz Allen team conducted this testing at its CNSI Lab. VIIC and Nissan conducted tests using the MEDS GID distribution functionality in the DTE.

A key component of MEDS testing was not accomplished due to the lack of lane positioning accuracy in probe data. Initial plans were to test the ability of the MEDS to modify a stop bar at an intersection based on probe data received from the system. This functionality required lane-level accurate probe data to determine where vehicles stopped when waiting for a red light and to then analyze multiple instances of stopped vehicles to determine whether the stop bar location needed to be updated.

MEDS testing successfully showed that GIDs could be imported into MEDS and sent to a corresponding RSE, via the AMDS service. Additionally, MEDS integration with PDS was shown in the CNSI Lab.

Table 35 summarizes the key findings in this area.

Table 35: Key Findings for MEDS Service

Finding Number	Finding
F-MEDS-1	GIDs could be imported into MEDS and sent to RSE, via the AMDS service.
F-MEDS-2	Lane-level positional accuracy is needed to test automated GID updates made from probe data (e.g., Stop Bar Update).
F-MEDS-3	Inaccurate or null vehicle dimensions, both size and shape, within probe data, mitigate the effectiveness of geo-statistical processing of probe data for developing or validating the accuracy of GIDs.
F-MEDS-4	Knowledge of the location of the GPS antenna on the vehicle is critical to determine the vehicle’s exact location on the roadway and in relation to other vehicles.

4.6. Positioning Service

POS testing was accomplished in the subsystem and system integration and test phases. The VIIC and Booz Allen teams conducted the testing at the CNSI Lab and the Battelle and VIIC test track facilities. The VIIC conducted extensive testing of the OBE and vehicle positioning systems, while Booz Allen testing focused on receiving a HA-NDGPS feed at the SDN and re-broadcasting that feed to RSEs in the SAE J2735 defined format.

The key finding from POS testing was that commercially feasible positioning technologies in vehicles currently do not support lane-level accuracy. The VIIC performed extensive testing, with close support from both positioning system manufacturers, Ublox and Sirf. The final conclusion from the VIIC positioning team, Ublox, and Sirf is that the current hardware and software available is not sufficient to support lane-level accurate positioning data. Additionally, VIIC test results showed that having VII vehicles receive and process HA-NDGPS data degraded the accuracy of positioning data.

Positioning testing did successfully show that the SDN is able to receive HA-NDGPS data over a network and re-broadcast that data to RSEs. Testing done at the VIIC Battelle test track showed that RSEs were able to receive this data and broadcast the HA-NDGPS corrections in the SAE J2735 format, and that the OBE would receive those corrections and go into differential correction mode.

Table 36 summarizes the key findings in this area.

Table 36: Key Findings for POS

Finding Number	Finding
F-POS-1	Currently, commercially feasible positioning technologies in vehicles do not support lane-level accuracy, which is required for some applications. Commercial products from two different vendors were used in the testing, and neither one met the lane-level accuracy requirement.
F-POS-2	The SDN was able to receive HA-NDGPS data over an Internet Protocol (IP) network, and forward the relevant Real-Time Correction Messages (RTCM) positioning correction data to RSEs. Testing done at the VIIC Battelle test track showed that RSEs were able to receive this RTCM data and broadcast the HA-NDGPS corrections in the SAE J2735 format, and that the OBE would receive those corrections and enter into differential correction mode.
F-POS-3	While, the end-to-end delivery of the RTCM positioning corrections was successfully tested, limitations of the OBE's GPS receiver resulted in HA-NDGPS positioning corrections actually degrading the accuracy of OBE positioning data. Consequently, positioning accuracies resulting from application of the corrections was not fully tested.

4.7. Security

The results of testing conducted for the VII security services are detailed in the following documents, which are located on the VII Portal:

- *Vehicle Infrastructure Integration (VII) Infrastructure Build 1 Test Report*, Sections 2.4 to 2.6 and 2.9.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 2 Test Report*, Sections 2.4 to 2.6.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 3 Test Report*, Sections 2.4 to 2.6.

VII security and privacy services are designed to provide trusted data to the vehicle operators and to prevent the misuse of vehicle information by unauthorized users of the system. RSE and onboard units utilize strong identity credentials to reject bogus messages and threats from would-be attackers of the network. VII POC security service features provide secure management of VII assets, network services, and the security of data transmitted across the network. This includes support for two communication protocols (V-DTLS and Vehicle-Host Identity Protocol, or V-HIP) developed for the VII POC system, which provide for secure DSRC communications. The VII security service provides for defense-in-depth protection of hardware, software, and information components against attackers, unauthorized users, and physical or electronic sabotage.

Included within the VII POC security services operations is an IdAM service providing VII network access account lifecycle management for RSE devices, Network Users (data subscribers, advisory providers, and transaction service providers), and administrative users. The VII security services, using a public key infrastructure (PKI), mediate access to the VII infrastructure by enabling the following activities to occur:

- Integrated identity lifecycle with the VII certificate authority to support X.509 certificates for managing provisioning, secure messaging, and applications security services
- Integrated identity lifecycle with the VII certificate authority to support 1609.2 WAVE certificates for managing provisioning, secure messaging, and applications security services. Using PKI-enabled services, the 1609.2 WAVE security capabilities include digital signature, encryption, certificate validation, and certificate management

- Automated and controlled provisioning of newly installed RSE devices and corresponding certificates, location directory lookup services for RSE devices, and centralized deprovisioning for RSE devices as they eventually get decommissioned
- Integrated account and access provisioning to SDN devices and services, allowing for simplified or single sign-on (SSO) capabilities for system administrators, remote system users, or applications.

By deploying a complete WAVE 1609 PKI, VII POC security has met (and exceeded) earlier US DOT Vehicle Safety Communications (VSC) study (2006) criteria, which states that "One of the anticipated major uses of this technology is for safety information such as emergency braking warnings, traffic signal violation warnings, and curve speed warnings. Security is a primary concern in any vehicle-safety application and it is especially important when the critical data is being transmitted over the air." The VII POC has demonstrated both the infrastructure-to-vehicle and vehicle-to-vehicle secure communications requirements described in the earlier 2006 VSC study.

Table 37 summarizes the key findings in this area.

Table 37: Key Findings for Security Service

Finding Number	Finding
F-SECU-1	The VII POC security IdAM service successfully provided VII network access account lifecycle management for RSE devices, Network Users (data subscribers, advisory providers, and transaction service providers), administrative users, and service provider management systems.
F-SECU-2	VII service provisioning (AMDS) using the 1609.3 Provider Service ID (PSID) demonstrated the successful interoperability of the 1609.3 (network) and 1609.2 (security) protocols for providing POC security services using WAVE.
F-SECU-3	1609 Security Configuration Management of multiple devices (certificate authority, Lightweight Directory Access Protocol, RSE, OBE) requires the training of multiple stakeholders on the use of 1609.2 cryptographic libraries.
F-SECU-4	By provisioning "anonymous certificates," the VII POC certificate authority enabled anonymity and privacy for vehicle-to-vehicle communication. Anonymous certificates remain outside the current definitions provided by the 1609 Security Protocol standards.

4.8. Enterprise Network Operations

ENOC testing was accomplished in the subsystem and system integration and test phases. The Booz Allen team conducted this testing at the CNSI Lab, and all SDN, ENOC, and RSE equipment was monitored by the ENOC at the DTE. The scope of ENOC testing was to demonstrate that all VII equipment, with the exception of OBEs, can be remotely monitored, updated, and maintained at a separate operations center.

The results of testing conducted for the VII POC ENOC are detailed in the following documents, which are located on the VII Portal:

- *Vehicle Infrastructure Integration (VII) Infrastructure Build 1 Test Report*, Sections 2.8.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 2 Test Report*, Sections 2.7 and 2.10.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 3 Test Report*, Sections 2.7 and 2.10.

The key finding from ENOC testing was the complexity and difficulty of integrating the RSE, which relied on multiple custom-developed applications to provide the VII services, with a commercial-off-the-shelf (COTS) network management product. While functionality like fault and security monitoring were straight forward to develop, the more advanced abilities like the configuration of specific applications and mass updates to applications via ENOC software were extremely complex and time consuming. After extensive testing, and day-to-day operations and maintenance, it was found that the most efficient way to manage RSE configurations and software updates was to make changes to each RSE individually. ENOC procedures were developed to simplify and expedite this process, allowing public and private applications testing to continue with minimal disruption; however, on a larger scale, this approach may not be feasible, and support from COTS manufacturers may be necessary to help develop tools to manage large deployments of RSEs.

ENOC testing successfully showed that a central operations center can monitor, maintain, and update all equipment in the DTE. Through the course of the VII test program, numerous patches and full software builds were deployed from the ENOC to RSEs already installed at POC intersections. Additionally, multiple configuration changes were made to those RSEs throughout the testing program to support specific test objectives and to troubleshoot issues seen during system integration and test.

Table 38 summarizes the key findings in this area.

Table 38: Key Findings for ENOC

Finding Number	Finding
F-ENOC-1	Installing, configuring, and maintaining the RSE was a more complex and difficult task than originally envisioned. It relied on multiple custom-developed applications to provide the VII services, and a COTS network management product extensively custom configured to interact with the RSE. It was found that the most efficient way to manage RSE configurations and software updates was to make changes to each RSE individually. However, on a larger scale, this approach may not be feasible, and support from COTS manufacturers may be necessary to help develop tools to manage large deployments of RSEs.
F-ENOC-2	A central operations center can monitor, update, and re-configure RSEs in the Detroit DTE remotely using an ENOC. Additionally, an ENOC can monitor and maintain the SDN remotely.

4.9. Certificate Authority

Both the Booz Allen team and the VIIC support team conducted the certificate authority testing. The Booz Allen team tested the X.509 certificate authority, which issues certificates to Network Users for access to the VII services. The VIIC support team tested the 1609.2 certificate authority, which provides the certificates to both OBE and RSE for DSRC communications. The X.509 certificate authority testing was accomplished at the CNSI Lab, and 1609.2 certificate authority testing was accomplished at Telcordia facilities, the CNSI Lab, and VIIC test track facilities.

The results of testing conducted for the VII POC certificate authority are detailed in the following documents, which are located on the VII portal:

- *Vehicle Infrastructure Integration (VII) Infrastructure Build 1 Test Report*, Section 2.10.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 2 Test Report*, Section 2.8.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 3 Test Report*, Section 2.8.

The X.509 certificate authority testing successfully showed that the SDN and ENOC can issue and revoke certificates to network users. Additionally, testing showed that without a valid certificate, Network Users were unable to access and use VII services.

Table 39 summarizes the key findings in this area.

Table 39: Key Findings for Certificate Authority

Finding Number	Finding
F-CA-1	The X.509 certificate authority testing successfully showed that the SDN and ENOC can issue and revoke certificates to network users. Additionally, testing showed that without a valid certificate, Network Users were unable to access and use VII services.
F-CA-2	The 1609 certificate authority testing demonstrated that the certificate authority can issue 1609.2 certificates (anonymous, identified, CSR signing, and WSA) and distribute to RSEs and OBEs. The certificate authority can also revoke any of these certificates and distribute via certificate revocation lists.

4.10. Test Bed Deployment

The VII network was primarily tested during subsystem testing and system integration and test, but utilized by all services during all phases of testing. The scope of the VII network testing was to demonstrate the ability of the network to support IPv6 and demonstrate a range of backhaul technologies. VII network testing was conducted in the CNSI Lab and at the DTE.

The results of the testing conducted for the VII POC network are detailed in the following documents and are located on the VII portal:

- *Vehicle Infrastructure Integration (VII) Infrastructure Build 1 Test Report*, Section 2.7.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 2 Test Report*, Section 2.11.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 3 Test Report*, Section 2.11.

Three key findings were discovered during VII network testing and use. The first is that WiMAX backhaul technology is susceptible to interference when using publicly available, unlicensed spectrum. Although WiMAX interference was not common, occurring twice over a 9-month period, these instances resulted in the interruption of RSE to SDN connectivity for the majority of WiMAX RSEs in the DTE. The second key finding was that 3G backhaul performance was dependant on the proximity to the closest Sprint cell phone tower and the number of users on the 3G service at any given time. Although rarely affecting the performance of VII services, large data transfers could be affected on 3G RSEs depending on location and time of day. The final key finding was the network's inability to support HIP natively. After troubleshooting issues with V-HIP sessions, it was discovered that, as implemented for the POC OBE, V-HIP used an experimental value in a protocol field (ethertype flag), which was not supported by COTS network equipment in the VII POC infrastructure. Individual VPNs between RSEs offering HIP and the

SDN were set up to mitigate this issue; however, on a large scale deployment of RSEs, this approach would not be practical.

The VII network successfully demonstrated that T1, 3G, and WiMAX backhauls can be fully supported. Of the three backhaul technologies, T1 was the most reliable and consistently provided 1.544 Mbps of bandwidth, although it required long deployment lead times and was the most expensive, month over month.

WiMAX was the second most reliable connection, providing an average bandwidth of 6 Mbps. A commercial WiMAX network was not available in the DTE, so a private WiMAX network was deployed. The initial set up of the network was costly; however, there were no month-to-month costs. Since this was a private network, POC network staff had to maintain all WiMAX equipment. Line of sight between the base station (network) and the subscriber unit (RSE) was critical for the 5.8 GHz frequency band and/or the particular hardware deployed.

Although 3G was the least reliable, connections were still active a majority of the time. Maximum bandwidth of 3G was 800 Kbps; however, bandwidth varied throughout the day. Overall, 3G was the cheapest of the backhaul technologies, provided adequate performance for most of the VII testing activities, and was the simplest to deploy.

The RSE for POC was mounted as an entire unit (i.e., compute platform, power supply, GPS antenna, and DSRC antenna) approximately 25 feet above the ground. While this made the RSE simpler to install, it made servicing the RSE hardware difficult. When the hardware required service, an RCOC/WCDR bucket truck was required. If the compute platform and the power supply were mounted approximately 10 feet above the ground, this would allow network personnel to service the devices without the need of a bucket truck.

Table 40 summarizes the key findings in this area.

Table 40: Key Findings for Test Bed Deployment (TBD)

Finding Number	Finding
F-TBD-1	The VII WiMAX backhaul network experienced periodic interference using the public unlicensed 5.8 GHz spectrum. Sources of interference were not always able to be determined; however, in one instance, a local radio station broadcasting at 100MHz did interfere with the WiMAX intermediate frequency band of one base station.
F-TBD-2	3G backhaul performance was dependant on the proximity to the closest cell phone tower and the number of users on the 3G service at any given time.
F-TBD-3	There is a requirement to suspend an OBE to TSP transaction session when an OBE leaves the range of an RSE, and resume the transaction session when the OBE subsequently comes within range of the same or a different RSE. An experimental V-HIP protocol was successfully tested; however, it was a qualified success, because it could not be supported as implemented by the COTS network equipment used in the VII backhaul and backbone network. Workarounds, in the form of IPv4 “SIT” tunnels, were required to successfully route V-HIP traffic through the POC infrastructure.
F-TBD-4	The VII network successfully demonstrated that T1, 3G, and WiMAX backhauls can be fully supported. Of the three backhaul technologies, T1 was the most reliable and consistently

Finding Number	Finding
	provided 1.544 Mbps of bandwidth, although it required long lead times to set up and was the most expensive. WiMAX was the second most reliable and provided an average bandwidth of 6 Mbps, but was not available in the DTE and required a private WiMAX network to be installed. Initial set up of the WiMAX network was costly; however, there were no month-to-month costs. Although 3G was the least reliable, those connections were still active the majority of the time. Maximum bandwidth of 3G was 800 Kbps; however, bandwidth varied throughout the day. Overall, 3G was the cheapest of the backhaul technologies, provided adequate performance for most of the VII testing activities, and was the simplest to deploy.
F-TBD-5	The current mounting location of the RSE compute platform and power supply made servicing the hardware difficult and required an RCOC/WCDR bucket truck.

4.11. Network User

The results of testing conducted for the VII POC Network User are detailed in the following documents, which are located on the VII portal:

- *Vehicle Infrastructure Integration (VII) Infrastructure Build 1 Test Report*, Sections 2.2 and 2.3.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 2 Test Report*, Sections 2.2, 2.3, and 2.9.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 3 Test Report*, Sections 2.2, 2.3, and 2.9.

POC testing showed that the AMDS service works well for network users. Testing showed that network users could send traveler advisory messages based on a mixture of ITIS codes and free text as specified in SAE J2735 to vehicles, and that vehicles would receive the messages quickly. Network users could cancel advisory messages; however, if a vehicle had already received the message, it would still display it. Network users could send advisory message sync requests to the RSE. Network users could query the delivery status of their advisory messages. Network users received error messages regarding message delivery failures.

POC testing showed that the PDS worked well for network users. Testing proved that network users could subscribe to probe data by choosing the RSE where the probe data was received, the time of the probe data reception, and the desired probe data elements included in each snapshot. Network users could also query the status and content of their subscriptions and modify their subscriptions before their expiration times. The modified subscriptions would take effect immediately. Network users could receive data in real time, and could cancel their subscriptions. In addition, network users could query supported probe data elements by the PDS, prior to or after placing their subscriptions. Network users received error messages containing error IDs and descriptions.

POC testing showed that the network user trip path application worked well. Network users received all trip path data from vehicles transmitting data (through an opt-in service). The anonymization of trip path data in the OBE worked. Only two vehicle tests of the trip path service were conducted.

Table 41 summarizes the key findings in this area.

Table 41: Key Findings for Network User

Finding Number	Finding
F-NU-1	AMDS service works well for network users.
F-NU-2	Network users could send traveler advisory messages (based on a mixture of ITIS phrases and free text, as specified in SAE J2735) to vehicles, and vehicles could receive the messages quickly.
F-NU-3	Network users could cancel advisory messages; however, if a vehicle had already received the message, it would still display it.
F-NU-4	Network users could send advisory message sync requests to the RSE.
F-NU-5	Network users could query the delivery status of their advisory messages.
F-NU-6	Network users received error messages regarding message delivery failures.
F-NU-7	The PDS worked well for network users.
F-NU-8	Network users could subscribe to probe data by choosing the RSE where the probe data was received, the time of the probe data reception, and the desired probe data elements included in each snapshot.
F-NU-9	Network users could query the status and content of their subscriptions.
F-NU-10	Network users could modify their subscription before its expiration time and the modified subscription would take effect immediately.
F-NU-11	Network users could receive data in real time and could cancel their subscriptions.
F-NU-12	Network users could query supported probe data elements by the PDS, prior to or after placing their subscriptions.
F-NU-13	Network users received error messages containing error IDs and descriptions.
F-NU-14	The trip path application worked well for network users.
F-NU-15	Network users received all trip path data from vehicles transmitting data (through an opt-in service).

4.12. Privacy

The scope of the privacy testing was to demonstrate the ability to protect against malicious intrusion, while maintaining privacy, which is an essential part of the VII concept. It also included testing the susceptibility of the system to intrusion (hackers), while ensuring the anonymity of users.

The results of testing conducted for the VII security and privacy services are detailed in the following documents which are located on the VII portal:

- *Vehicle Infrastructure Integration (VII) Infrastructure Build 1 Test Report*, Sections 2.4 to 2.6 and 2.9.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 2 Test Report*, Sections 2.4 to 2.6.
- *Vehicle Infrastructure Integration (VII) Infrastructure Build 3 Test Report*, Sections 2.4 to 2.6.

The following table summarizes the key findings in this area.

Table 42: Key Findings for Privacy

Finding Number	Finding
F-PRIV-1	POC was designed not to collect personally identifiable information (PII). As documented in the VII System Infrastructure Privacy Impact Assessment (PIA), the POC test environment demonstrates that it does not collect PII.
F-PRIV-2	The VII PIA provides adequate guidance for privacy protection in a scaling national deployment.
F-PRIV-3	Inspection of probe data during testing provided a means of validating the absence of PII content in POC.

4.13. Standards

Standards testing was accomplished in all three phases of the VII test program. The scope of this testing included validating and studying how the IEEE 1609 and 802.11p suite of standards operate in the 5.9 GHz spectrum allocated by the FCC for VII use, as well as studying and verifying how SAE J2735 standards (POC DSRC message sets) function in support of VII applications.

The following table summarizes the key findings in this area.

Table 43: Key Findings for Standards

Finding Number	Finding
F-STD-1	V-DTLS only allows one vehicle to send probe data at a time.
F-STD-2	The probe data message, as specified in the POC DSRC message set, can successfully accommodate sending vehicle sensor data to network users.
F-STD-3	The PSN, as specified in the POC DSRC message set, can successfully be generated randomly and changed at predetermined times and conditions.
F-STD-4	The PDM message, as specified in the POC DSRC message set, can successfully accommodate directing vehicles to change probe data collection policies.
F-STD-5	The start/stop strategy for probe data collection, as specified in the POC DSRC message set, can successfully accommodate suspension of collection and sending of probe data while the vehicle is stopped.
F-STD-6	The start gap for probe data collection, as specified in the POC DSRC message set, can successfully accommodate prolonging the collection of probe data until a distance traveled threshold is met.
F-STD-7	The traveler information message, as specified in the POC DSRC message set, can successfully accommodate sending relevant information/signage to vehicles.
F-STD-8	Java, as implemented in POC, is not suitable for sending/receiving low-latency messages such as SPAT/GID and heartbeat, as it has too much overhead.
F-STD-9	The trip path message, as implemented in POC, can successfully provide the route a vehicle has traveled.

4.14. Multiple Simultaneous Applications Impacts

The goal of multiple simultaneous applications (multi-apps) impacts testing was to observe the impact of running multiple applications simultaneously on the OBEs, RSEs, and the Network User side and comparing the results of each individual application when executed independently and while in the presence of other applications. Specifically, the goal was to characterize the impact:

- Of simultaneous UDP, TCP/IP, and WSM communications on the OBE, RSE, infrastructure, and network users, in a multi-vehicle environment
- On reception of probe data by Public Apps
- Of heartbeat generation density versus heartbeat messages reception, under real traffic conditions
- Of high-priority infrastructure messages on the higher-priority heartbeat messages
- Of WSAs on high-priority infrastructure messages and higher-priority heartbeat messages
- Of PDM messages on high-priority infrastructure messages and higher-priority heartbeat messages, while in the presence of WSAs
- Of multi-apps on the reception of probe data by Public Apps
- Of several control channel messages on high-priority infrastructure messages and higher-priority heartbeat messages, while in the presence of WSAs.

Tests were conducted while vehicles were both parked and driving. The parked tests provided results that served as a baseline for comparison with the results of driving tests. Parking formations were dictated by available space and testing needs.

The following sections provide a detailed analysis of the findings from the multiple multi-apps impacts testing carried out in Detroit, Michigan, on September 11 and 12, 2008.

4.14.1. Heartbeat, Probe, and Off-Board Navigation (OBNA) Tests

Test Case Objectives

- Characterize the impact of simultaneous UDP, TCP/IP, and WSM communications on the OBE, RSE, infrastructure, and network users, in a multi-vehicle environment.
- Characterize the impact on reception of probe data by Public Apps.

Data Outputs

- Comparison of heartbeat (HB) transmission rate and latency by HB frequency and vehicle concentrations
- Comparison of probe data loss and latency versus single application (single-app) testing data
- Comparison of probe data loss and latency and roundtrip latency of OBNA
- Comparison of HB transmission rate and latency, probe data loss and latency, and roundtrip latency of OBNA.

Exit Criteria

The tests pass if HB, probe, and OBNA performances remain in line with their baseline results. Data from HB-only testing will be used as baseline for HB. Data from probe-only testing (both during multi-apps and single-app testing) will be used as baseline for probe. There is no baseline data for OBNA.

Hypotheses

1. Roundtrip latency for OBNA will increase as more applications become active on the OBE.
2. A significant percentage of probe data snapshots that are securely transmitted by OBEs reach subscriptions.
3. Probe data snapshots take less than 1 minute to propagate from an OBE to a network subscriber under all conditions.
4. For HB, as the number of vehicles in range increases, the actual transmission rate drops.
5. For HB, as the number of vehicles in range increases, the percentage of heartbeats transmitted drops.
6. For HB, as the transmission rate increases, the percentage of heartbeats transmitted drops.
7. For HB, as the number of vehicles in range, increases the percentage of heartbeats received drops.
8. For HB, as the transmission rate increases, the percentage of heartbeats received drops.
9. For HB, as the number of vehicles in range increases, average latency increases.
10. For HB, as the number of transmission rate increases, average latency increases.

Data Analysis Strategy

Probe data loss and latency will be calculated using the same techniques that were used during single-app testing. HB transmission rate and latency data will be calculated using the results of the analysis tool. OBNA's roundtrip latency will be calculated using the results of the analysis tool.

Applications

On September 11, the testing included probe, OBNA, and HB applications. For OBNA, the drivers were told to request the same route multiple times.

Routes and RSEs

Table 44 shows the RSE that was used for testing.

Table 44: List of RSE(s)

RSE Location	RSE Number	Backhaul
12 Mile Rd & Meadowbrook Rd	us.mi.csnv.rse.0067	T1

The following screen shots from Google Maps show the routes and RSE used for multi-apps testing on September 11.

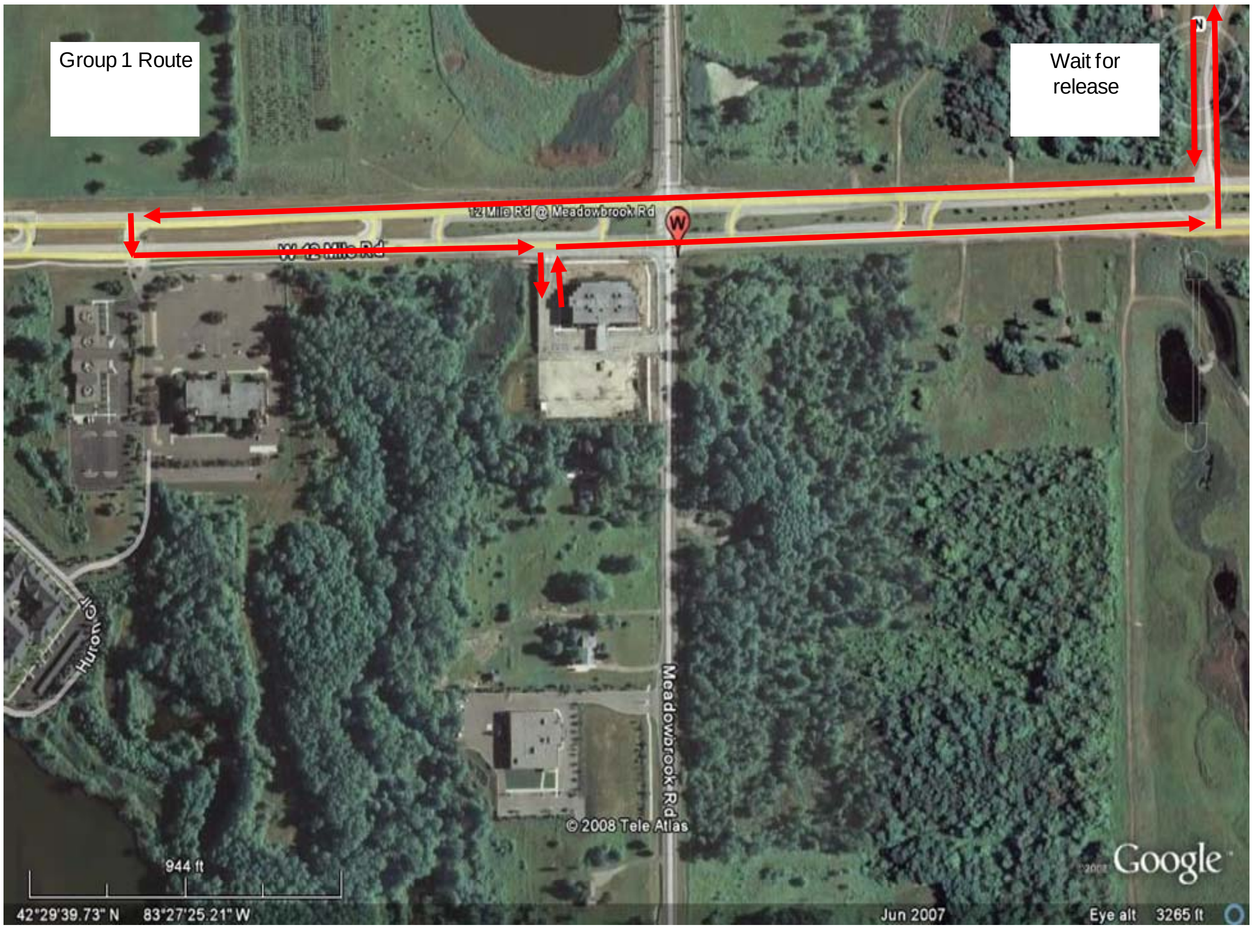

Booz Allen generated routes and plot points on Google map

Figure 48: Group 1 Route and RSE

Booz Allen generated routes and plot points on Google map

Figure 49: Group 2 Route and RSE

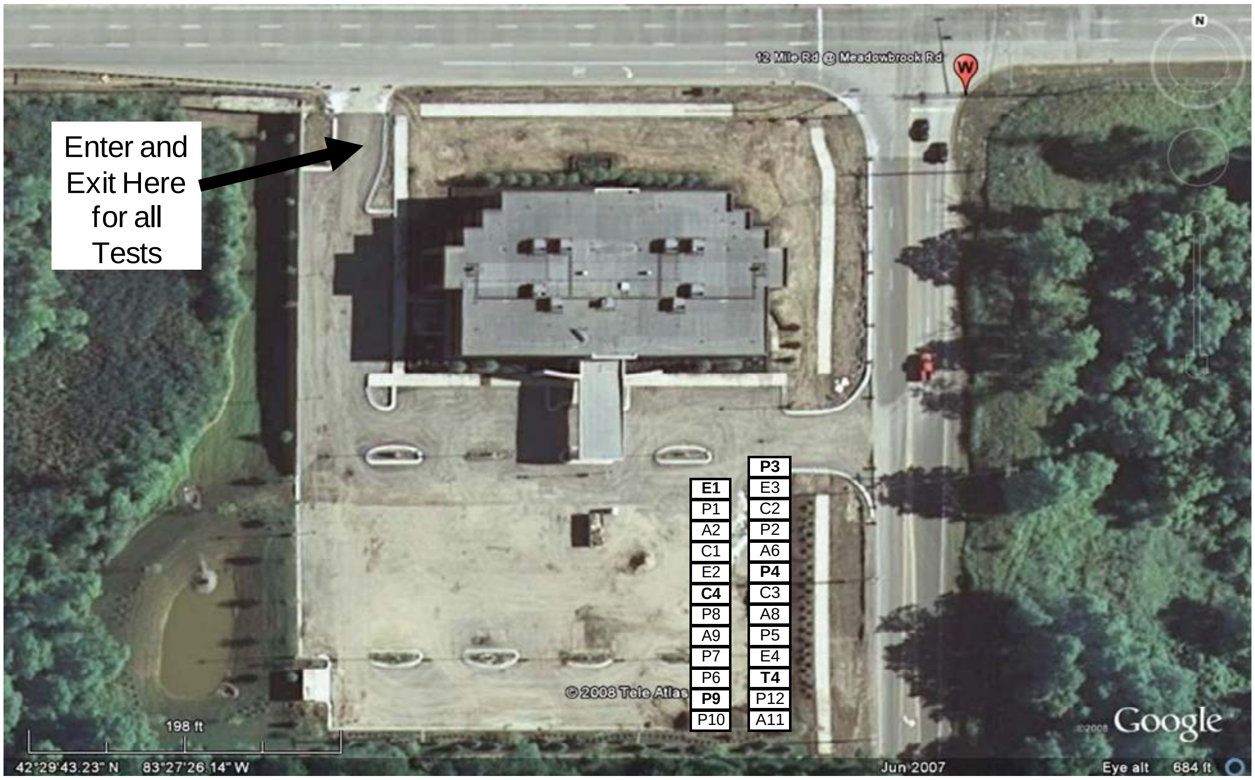

Booz Allen generated plot points on Google map

Figure 50: Parked Configuration on September 11

Vehicles drove the above route. Applications were turned on in the sequence shown in Table 45.

Table 45: Sequence of Application Start Ups

Parked HB
Traffic HB
Parked Probe
Traffic Probe
Parked Probe and OBNA
Traffic Probe and OBNA
Parked Probe and OBNA and HB*
Traffic Probe and OBNA and HB*
*Tests carried out with 6 vehicles

Note about the RSE Log

The RSE timestamp was off by around 3 hours. Because of this, the latency from the OBE to the RSE and from the RSE to the Public Apps could not be calculated for probe data.

Note about OBE Logs

During probe and OBNA testing, 22 of 25 vehicles were able to download navigation routes. However, when HB was added to the list of applications, none of the vehicles were able to download navigation routes. After some troubleshooting, it was decided to finish the testing with only six vehicles. According to an observer, during one six-vehicle run on the roadway, only one vehicle received the navigation message, and then only after 5 minutes and 40 seconds.

Probe data snapshots were received from seven vehicles that had received our PDM message. However, 25 vehicles should have responded. Further, detailed analysis of OBE logs is necessary to determine the reason. For HB, test mode was not turned on, so the analysis tool did not have the sequence number to work with but had to try and match heartbeats using information from the heartbeat itself. This was not a perfect process and it could have led to some mismatches.

The following table lists issues with vehicles during testing.

Table 46: OBE Issues – September 11

Nav Only, T5: HIP not working, which caused Nav not to work
P9: turned off during lunch, when it should not have been
Parked Nav, P12: did not receive map
Parked Nav, P5: did not receive map
Parked Nav, E1: did not receive map
Parked Nav, C2: OSGI crashed
Parked Nav, A6: OSGI crashed

Off-Board Navigation

Roundtrip Latency

Measure Objective: Determine the roundtrip latency between the time when the navigation request was sent and maneuvers were received.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Using the OBE logs, find out the roundtrip latency using "Route Request Timestamp" and the "Maneuvers Received Timestamp."

Results: To understand the results, first the vehicle formations and their effect on average latency will be analyzed.

Probe and OBNA in parked formation (21 vehicles): In this formation, the average latency was about 11 seconds.

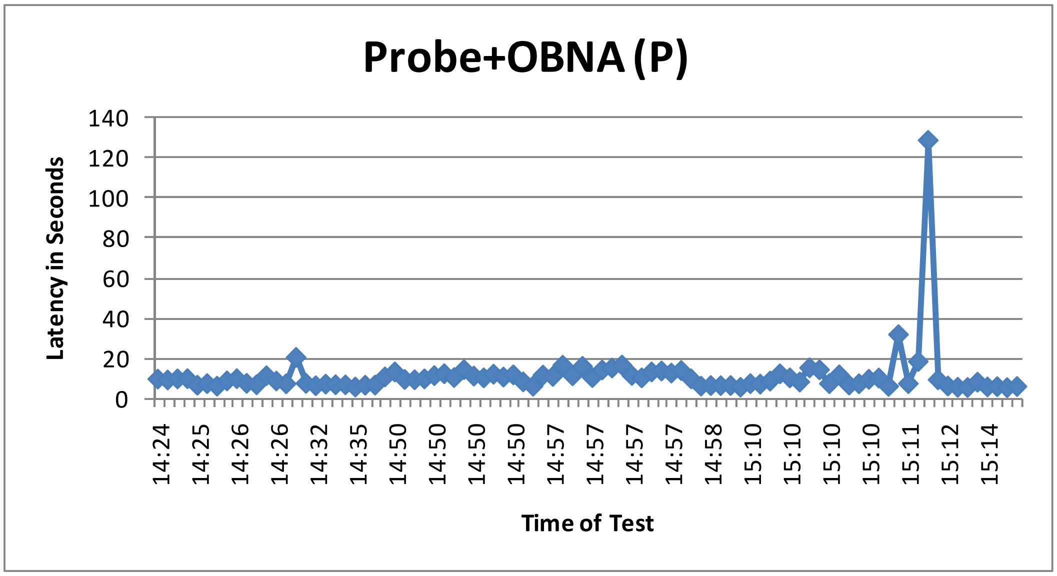

Figure 51: Average Latency PDS and OBNA Parked – excl. Faulty Data

Probe and OBNA in traffic formation (16 vehicles): In this formation, the average latency was about 13 seconds. This could be due to the fact that the RSE on the route was not overloaded with too many simultaneous requests.

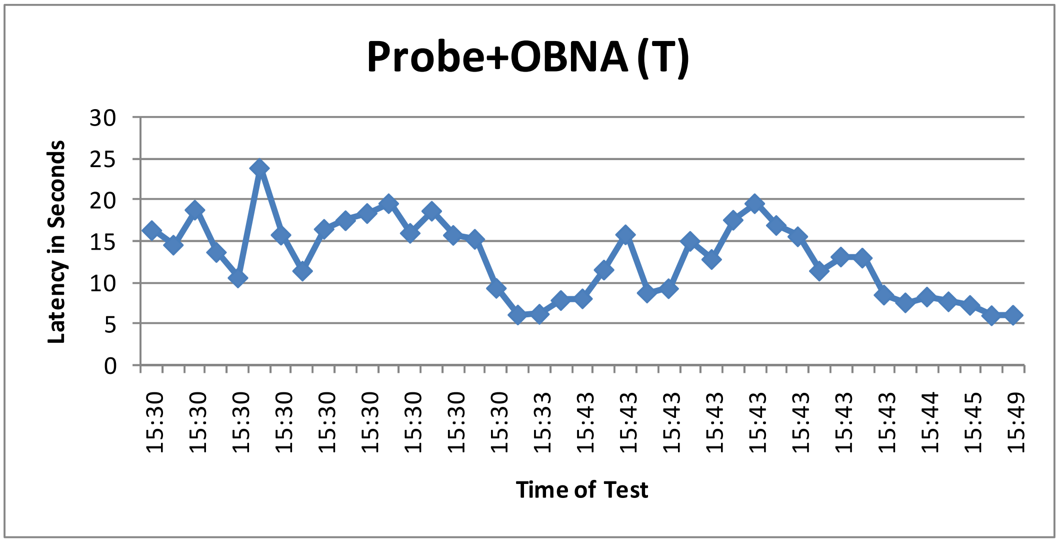

Figure 52: Average Latency PDS and OBNA Traffic

Probe, OBNA, and HB in parked formation (6 vehicles): In this formation, the average latency was about 57 seconds as shown in Figure 53.

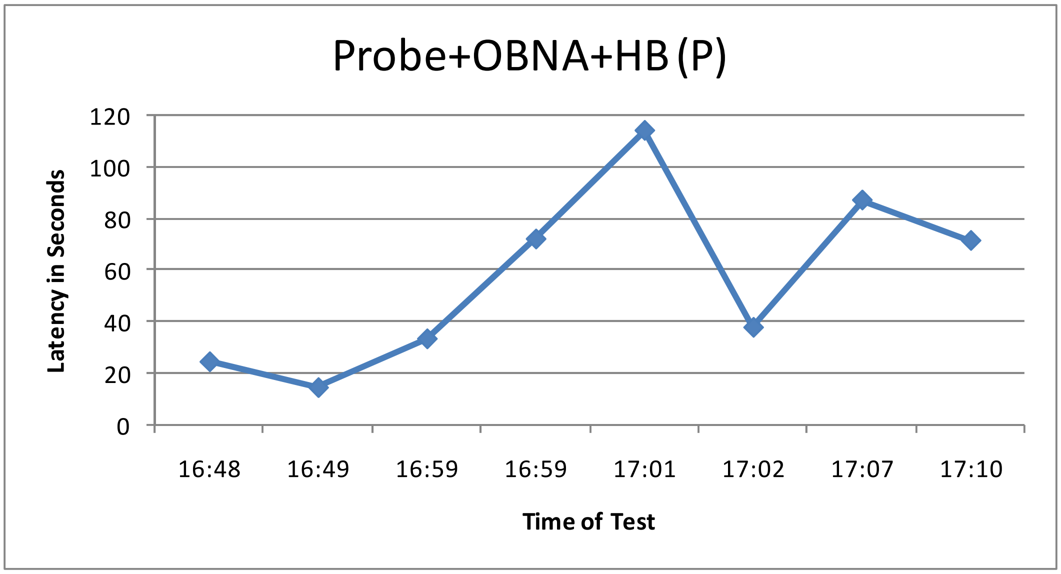

Figure 53: Average Latency PDS, OBNA, and HB Parked – excl. Faulty Data

Probe, OBNA, and HB in traffic formation (5 vehicles): In this formation, the average latency was 8.4 seconds as shown in Figure 54.

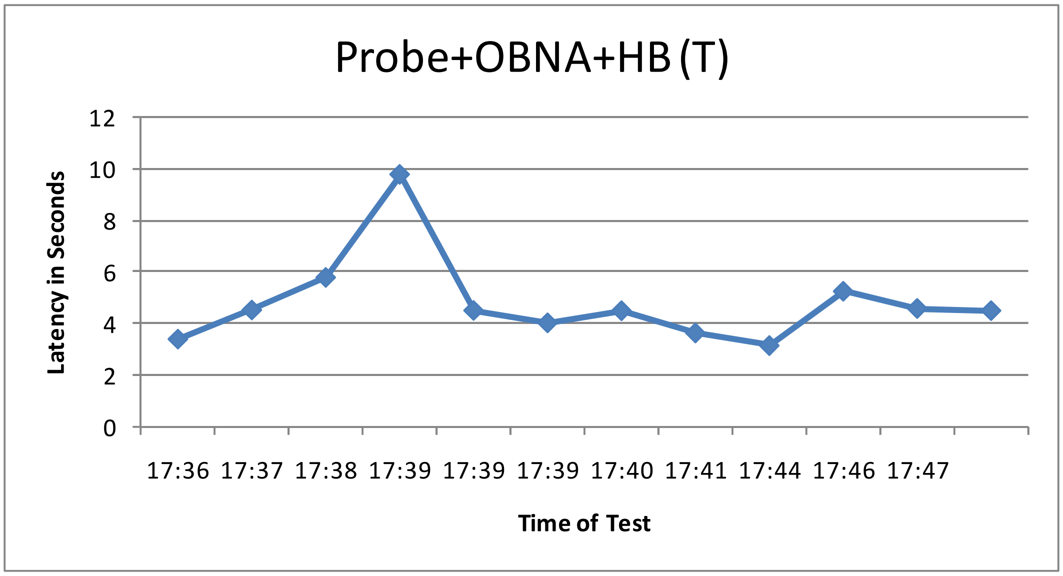

Figure 54: Average Latency PDS, OBNA, and HB Traffic

Conclusion

The results show that when HB, probe, and OBNA are running simultaneously on a given vehicle, the OBE-to-RSE interaction becomes unreliable. The underlying issue may be in the OBE's implementation of the various applications in parallel and/or the RSE's capability to serve 25 vehicles with OBNA data while receiving probe data from those vehicles. This issue warrants further testing and analysis.

Figure 55 shows how the overall average roundtrip latency of OBNA requests was affected by the different testing scenarios.

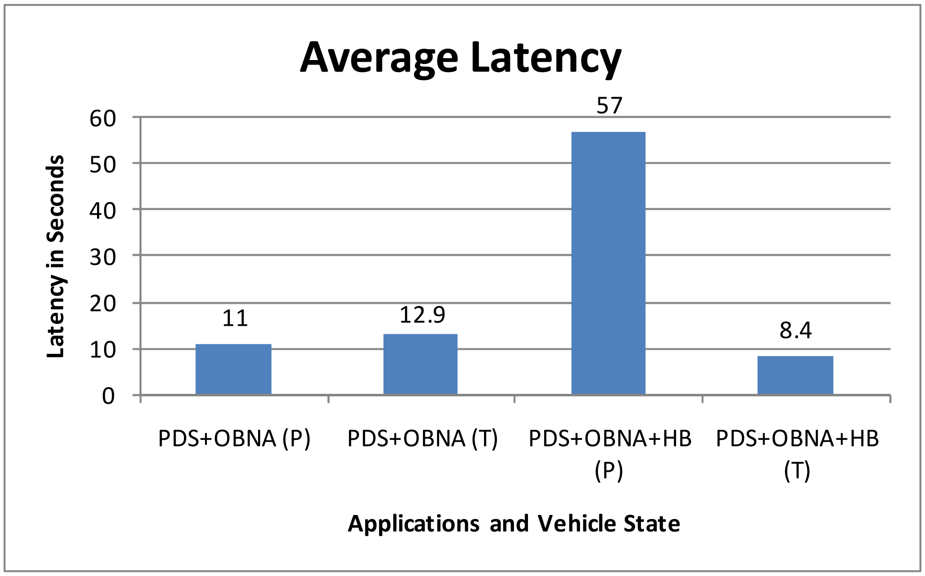

Figure 55: Average Latency, All Data

On average, roundtrip latency was around 22 seconds; however, parked HB and OBNA add a significant amount of lag to the results. Given these limitations, it becomes even more important that OBEs are able to download routes from multiple RSEs.

Probe

Measure 1: Probe Data Loss

Measure Objective: Determine the average number of probe data snapshots lost when propagating from an OBE to a network user.

Data Source: OBE OSGi logs, Public Apps database.

High-Level Data Analysis Plan: Using summary loss data, calculate probe data loss for each test scenario.

Results and Analysis: Figure 56 show the percentage of probe data that did not make it to the RSE (blue bars) and to the Public Apps (red bars).

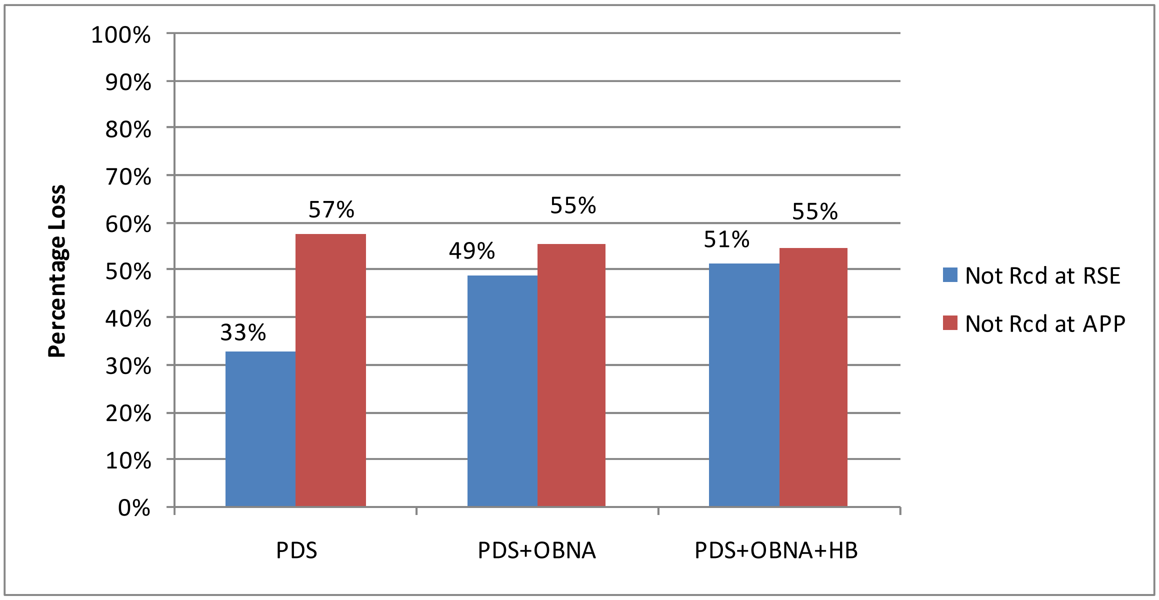

Figure 56: Probe Data Loss Day 2 Multi-Apps

The percentage of probe data loss from the OBE to the RSE increases as more applications are turned on. One conclusion that can be drawn is that as the "chatter" between the OBE and RSE increases, the probe data loss percentage increases. The loss percentages are also generally higher than what was observed during PDS-only testing days. This could be due to the fact that all of the probe data was being transmitted to only one RSE and during parked tests; 25 vehicles were attempting to communicate with that RSE at the same time.

A consistent loss of over 50 percent was observed between the OBE and Public Apps. A close look at the results reveals that the loss of probe data, after it was received by the RSE, was greatest when only the probe gathering application was turned active on the OBE.

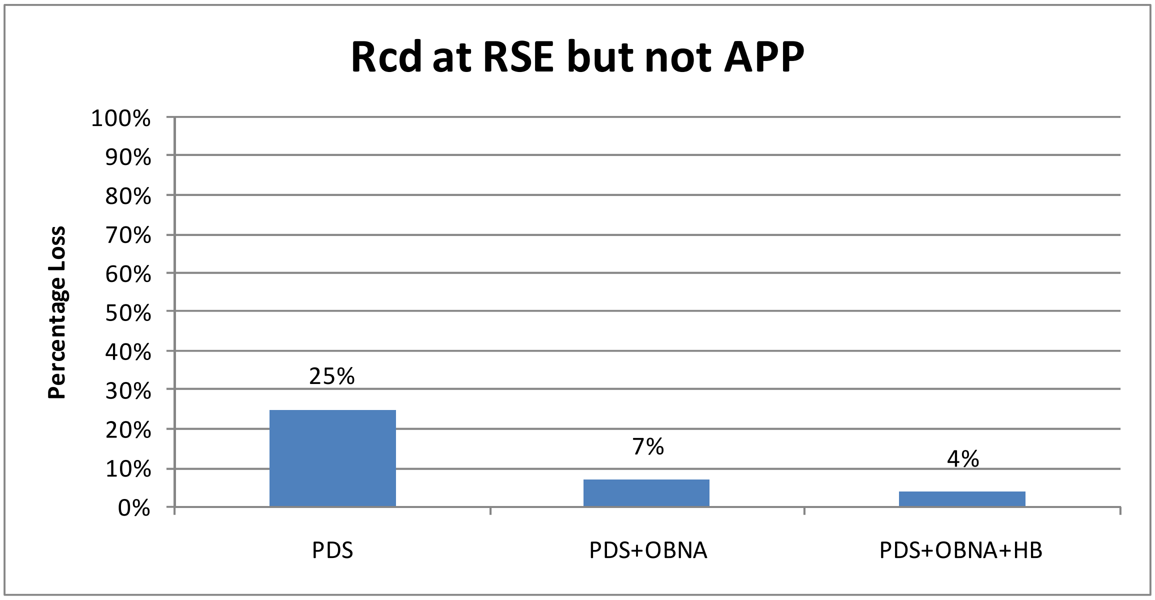

Figure 57: Probe Data Loss to the Public Apps Day 2 Multi-Apps

To explain this anomaly, further analysis of the network is required. It is possible that the WiMAX backhaul connection was experiencing periods of downtime.

Breaking down the tests further into parked and traffic scenarios shows a similar result.

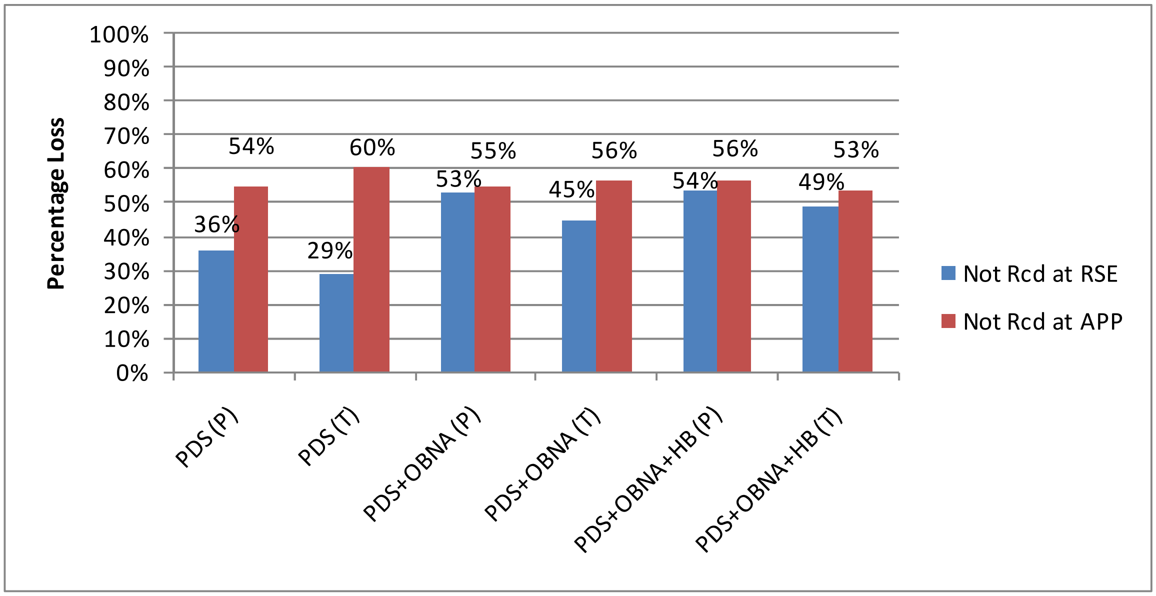

Figure 58: Probe Data Loss by Driving Condition Day 2 Multi-Apps

Conclusion: Generally, probe data loss during driving tests is lower than parked tests, most likely due to fewer vehicles communicating with the RSE at the same time.

To get a better understanding of the results, the test would have to be repeated with different backhaul RSEs. Generally speaking, WiMAX is more prone to loss than 3G or T1 backhaul connections.

Measure 2: Probe Data Latency

Measure Objective: Determine the average time it takes for a probe data snapshot to propagate from an OBE to a network user.

Data Source: OBE OSGi logs, Public Apps database.

High-Level Data Analysis Plan: Using summary latency data, calculate probe data latency for each test scenario.

Results and Analysis: Due to an RSE clock synchronization error, the latency between OBE and RSE could not be calculated.

The latency between OBE and Public Apps ranged from 0.5 to 1.2 seconds.

The following presents the breakdown by test and driving condition.

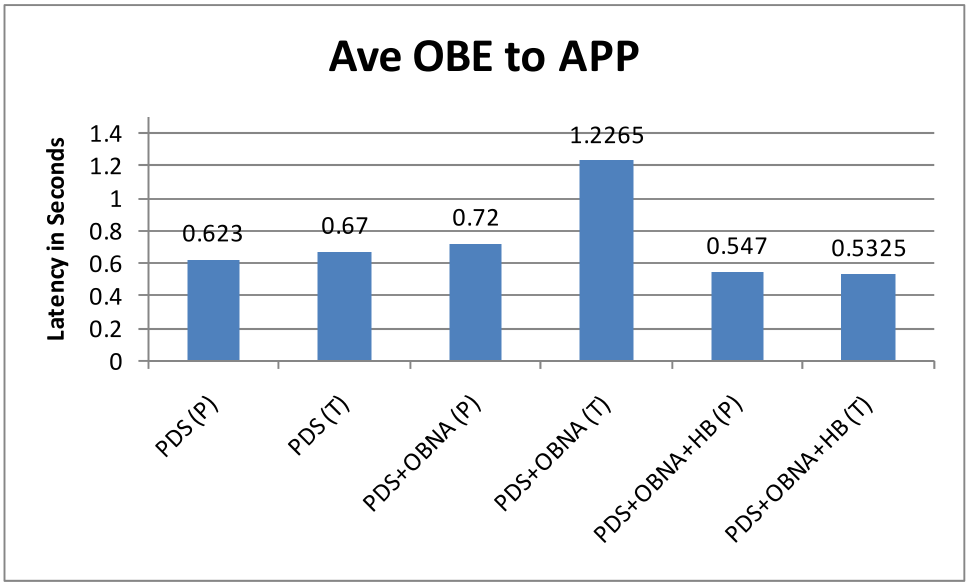

Figure 59: Probe Data Latency by Driving Condition Day 2 Multi-Apps

Conclusion: These numbers are in line with the observation of 0.7-second average latency during PDS-only testing. It appears that multiple applications on the OBEs do not adversely affect probe data latency.

Heartbeat

Methods

Vehicle staging and ramp-up tests took place in Roush parking lot on September 11 between 6:30 am and 9 am. Vehicles were arranged in an arc due to space limitations.

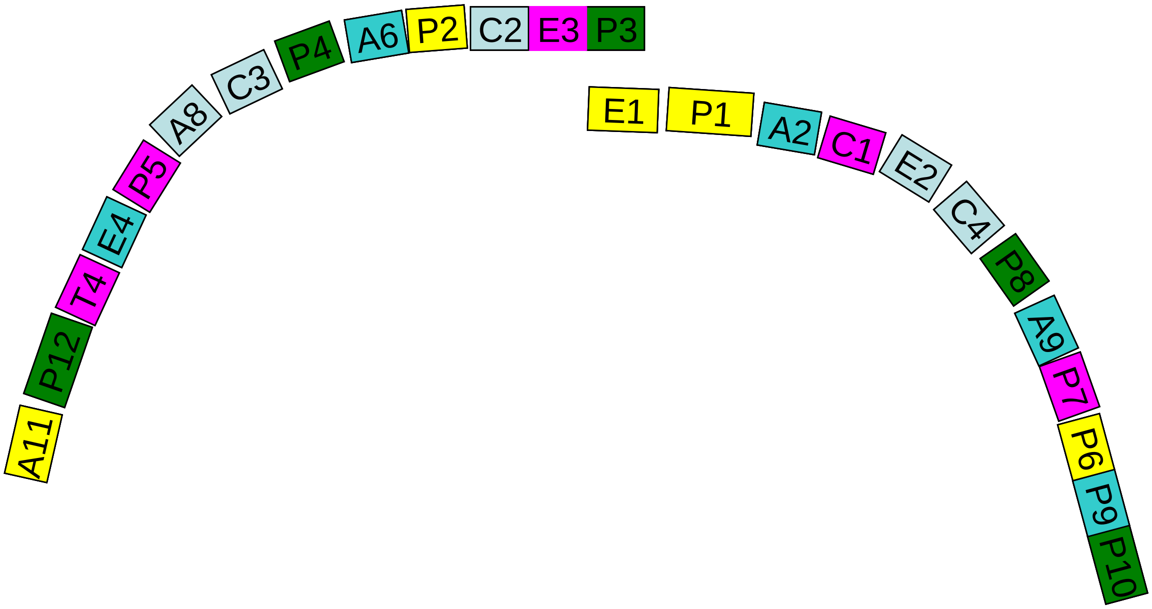

Figure 60: Ramp-Up Tests Parked Formation

The vehicles were grouped together in the batch formations shown in Table 47.

Table 47: Vehicle Batch Assignments September 11

OBE ID	Vehicles	Batch Assignment
C590	P10	BATCH A
C366	P12	BATCH A *started w/ Batch C
C694	P3	BATCH A
C256	P4	BATCH A
C548	P8	BATCH A
C594	A2	BATCH B
C779	A6	BATCH B
C141	A9	BATCH B *started w/ Batch D
B325	E4	BATCH B
C082	P9	BATCH B
C453	C1	BATCH C
B193	E3	BATCH C
C832	P5	BATCH C
C992	P7	BATCH C
C291	T4	BATCH C
CC94	A11	BATCH D
B450	E1	BATCH D
C482	P1	BATCH D * Didn't start
C862	P2	BATCH D
C194	P6	BATCH D
C985	A8	BATCH E
B856	C2	BATCH E
B422	C3	BATCH E
B420	C4	BATCH E
B042	E2	BATCH E

Measure 1: Actual Transmission Rate, as a Function of the Number of Vehicles within Range at 10Hz

Measure Objective: Verify whether the transmission rate was affected by the number of vehicles within range.

Concept: Transmission rate is set within the application, and the default transmission rate is 10Hz or roughly 10 heartbeats per second. However, it is possible that fewer heartbeats are generated as more vehicles come online due to OBE overload.

Hypothesis: As the number of vehicles in range increases, the actual transmission rate drops.

Data Source: OBE OSGi logs for ramp up tests.

High-Level Data Analysis Plan: Compare average actual transmission rate for the ramp-up tests between groups of vehicles in range.

Results and Analysis: Figure 61 shows the actual transmission rate as a function of the number of vehicles in range.

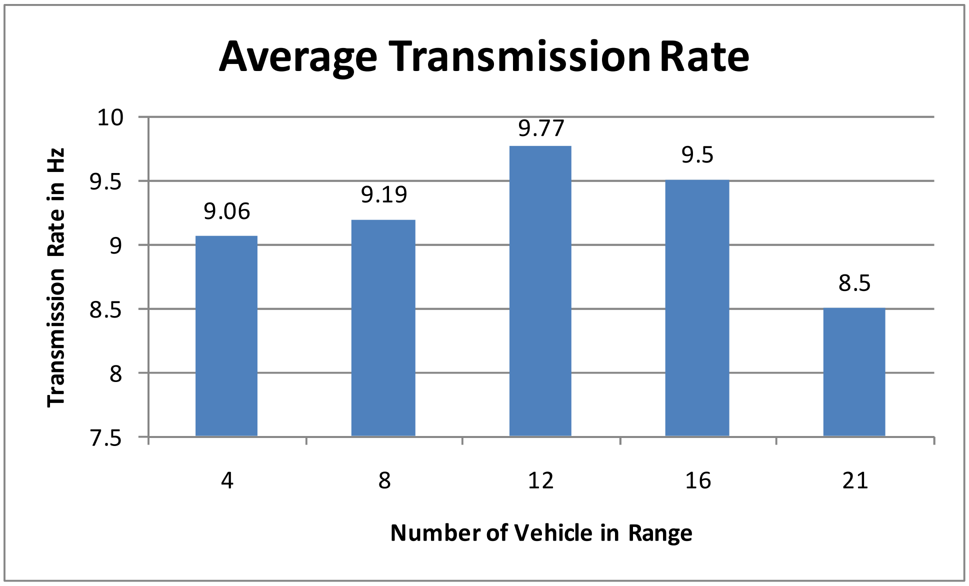

Figure 61: Actual Transmission Rate, as a Function of # of Vehicles within Range @10Hz

Heartbeat generation and transmission does not seem to be directly affected by the number of vehicles in range.

Measure 2: Percentage (Compared to Theoretical) of Heartbeats Transmitted, as a Function of the Number of Vehicles within Range at 10Hz

Measure Objective: Determine whether the percentage of heartbeats transmitted was affected by the number of vehicles within range.

Concept: Similar to Measure 1, this measure looks at the percentage of heartbeats transmitted.

Hypothesis: As the number of vehicles in range increases, the percentage of heartbeats transmitted drops.

Data Source: OBE OSGi logs for ramp up tests.

High-Level Data Analysis Plan: Compare average percentage of heartbeats transmitted for the ramp-up tests between groups of vehicles in range.

Results and Analysis: Figure 62 shows the percentage of heartbeats transmitted versus theoretical limit as a function of the number of vehicles in range.

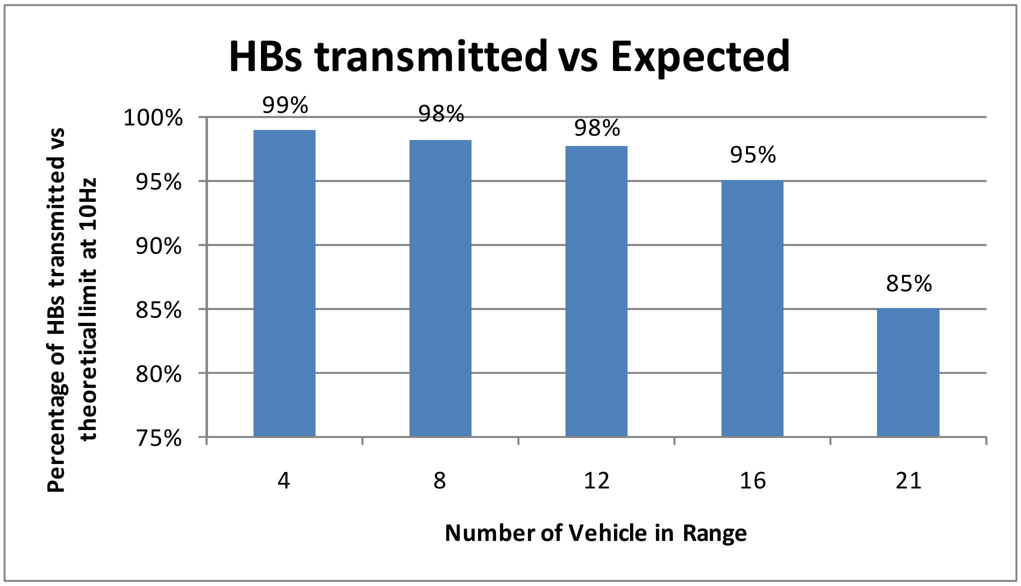

Figure 62: Percentage of Heartbeats Transmitted, as a Function of # of Vehicles within Range @10Hz

As the number of vehicles increases, the percentage of heartbeats generated and transmitted goes down. This is in line with the hypothesis and means that, as more and more vehicles send out heartbeats in a given area, the total number of heartbeats by any one vehicle goes down. However, even at a 21-vehicle concentration, an 85-percent response rate can be observed, which may be sufficient for algorithms running on vehicles.

Measure 3: Percentage (compared to total) of Heartbeats Received, as a Function of the Number of Vehicles within Range

Measure Objective: Determine whether the number of heartbeats received goes down as more vehicles come online.

Concept: Each OBE has a limited capacity for handling incoming data. This measure checks for an upper bound for incoming heartbeat data from varying numbers of vehicles.

Hypothesis: As the number of vehicles in range increases, the percentage of heartbeats received drops.

Data Source: OBE OSGi logs for ramp up tests.

High-Level Data Analysis Plan: Compare the percentage of heartbeats received for the ramp-up tests between groups of vehicles in range.

Results and Analysis: Figure 63 shows the average percentage of heartbeats received as a function of the total number of heartbeats sent. OBE C291 was not included in this measure because it did not receive any snapshots during test times.

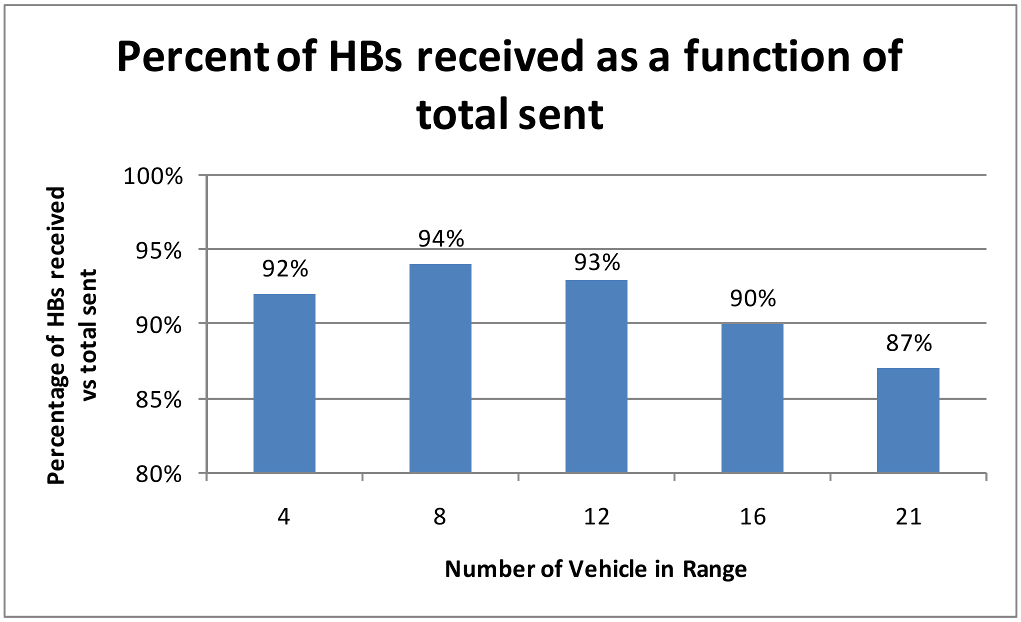

Figure 63: Percentage of Heartbeats Received, as a Function of # of Vehicles within Range

The data does indicate a slight downward trend as the number of vehicles increase. Although this is in line with the hypothesis, it is recommended that this test be repeated with more vehicles for a longer period of time to determine if the trend down is permanent or temporary.

Measure 4: Average Latency, as a Function of the Number of Vehicles within Range

Measure Objective: Determine whether heartbeat latency was affected by the number of vehicles in range.

Concept: It is imperative that the latency between a heartbeat sent and the same heartbeat received by another vehicle is miniscule. This measure checks for average latency as the number of vehicles in range increases.

Hypothesis: As the number of vehicles in range increases, average latency increases.

Data Source: OBE OSGi logs for ramp-up tests.

High-Level Data Analysis Plan: Compare average latency for the ramp-up tests between groups of vehicles in range.

Results and Analysis: Several OBEs were excluded from the analysis because the analysis tool identified negative latencies in their logs (i.e., the heartbeat was received by receiving vehicles before it was generated). Negative latencies could be the result of one of the following:

- The timestamps between the two vehicles were not synchronized. (This was also observed in VIIC testing.)
- The heartbeat application did not capture time correctly.
- The wrong heartbeats were matched to each other.

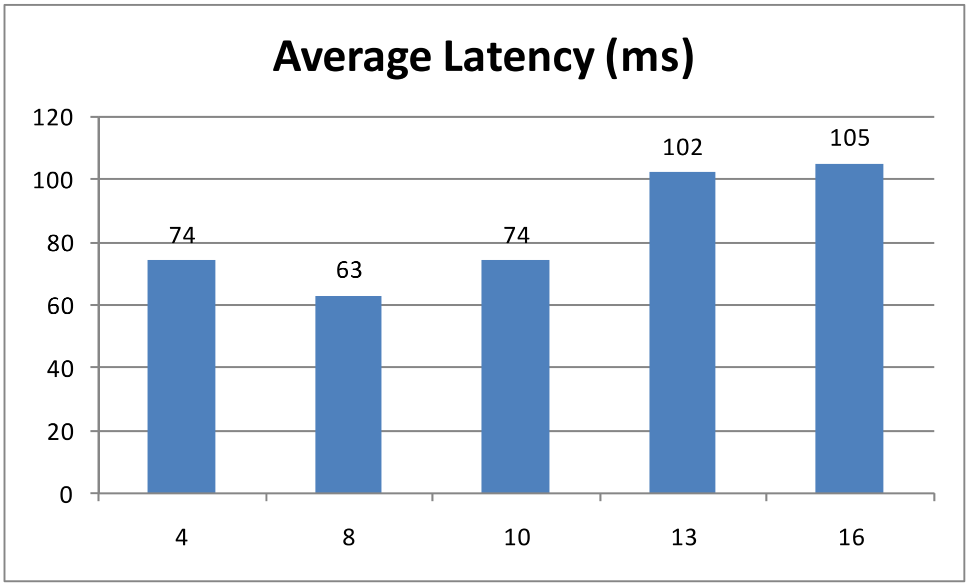

Figure 64: Average Latency, as a Function of # of Vehicles within Range

From Figure 64, it can be concluded that, in general, the latency between when a heartbeat is sent and received is less than 1 second. The number of vehicles in range does not seem to have a discernible effect on latency.

4.14.2. Heartbeat, SPAT, WSA, PDM, and AMDS Tests

Test Case Objectives

- Characterize the impact of heartbeat generation density versus heartbeat message reception, under real traffic conditions.
- Characterize the impact of high-priority infrastructure messages on the higher-priority heartbeat messages.
- Characterize the impact of WSAs on high-priority infrastructure messages and higher-priority heartbeat messages.
- Characterize the impact of PDM messages on high-priority infrastructure messages and higher-priority heartbeat messages, while in the presence of WSAs.
- Characterize the impact of multi-apps on reception of probe data by Public Apps.
- Characterize the impact of several control channel messages on high-priority infrastructure messages and higher-priority heartbeat messages, while in the presence of WSAs.

Data Outputs

- Comparison of heartbeat transmission rate and latency by heartbeat frequency and vehicle concentrations
- Comparison of SPAT loss and latency and heartbeat transmission rate and latency
- Comparison of SPAT loss and latency and heartbeat transmission rate and latency with active WSA

- Comparison of SPAT loss and latency and heartbeat transmission rate and latency with active WSA and PDM messages
- Comparison of SPAT loss and latency and heartbeat transmission rate and latency with active WSA, PDM, and public AMDS messages.

Exit Criteria

The tests pass if HB, SPAT, PDM, and AMDS performances remain in line with their baseline results. Data from HB-only testing will be used as a baseline for HB. (Note: HB results for September 12 were not yet available for inclusion in this report.) Data from PDM-only testing will be used as a baseline for PDM. Data from AMDS-only testing will be used as a baseline for AMDS. Data from VII communications safety link testing on June 19, 2008, will be used as a baseline for SPAT.

Hypotheses

1. The frequency at which SPAT messages are received by the OBEs decreases as the number of active applications on the OBE increases.
2. The latency of SPAT messages increases as the number of active applications on the OBE increases.
3. OBEs under PDM did generate periodic probe data at the frequency specified in the PDM message.
4. OBEs under PDM did generate periodic probe data within a reasonable standard deviation of the anticipated time.
5. The majority of PDM messages expired naturally, unless the vehicle came in contact with another RSE or left the RSE footprint early.
6. PDM messages that expired naturally did so correctly.
7. Stop snapshots were only generated when the vehicle speed was below the stop threshold.
8. Stop snapshots were generated past the specified lag time.
9. Start snapshots were only generated when the vehicle speed was above the start threshold.
10. All messages sent from Public Apps were received by the OBEs.
11. Roundtrip latency from Public Apps to RSEs should be a matter of a few seconds.
12. Message durations were in line with operator-specified values.
13. All messages were only displayed to the driver when the vehicle was in a valid active region.
14. Only RSEs defined in the broadcast strategy played the messages.
15. The correct broadcast priority was received by the RSEs.
16. Priority order was implemented correctly and highest-priority messages were displayed on top.
17. Messages were received, decoded, and displayed correctly.

Data Analysis Strategy

PDM implementation data was calculated using the same techniques that were used during single-app testing.

SPAT frequency and latency data was calculated manually using spreadsheets.

AMDS implementation and execution data was calculated using the same techniques that were used during single-app testing.

HB transmission rate and latency data was calculated using the results of the analysis tool. (Results not yet available.)

Applications

On September 12, the testing included HB, SPAT, private signage (WSA), PDM, and public signage applications.

Routes and RSEs

Table 48 lists the RSE that was used for testing.

Table 48: List of RSEs Day 2 Multi-Apps

RSE Location	RSE Number	Backhaul
9 Mile Rd & Meadowbrook Rd	us.mi.csnv.rse.0058	3G

The following screen shot from Google Maps (Figure 65) shows the route and RSE used for multi-apps testing on September 12.

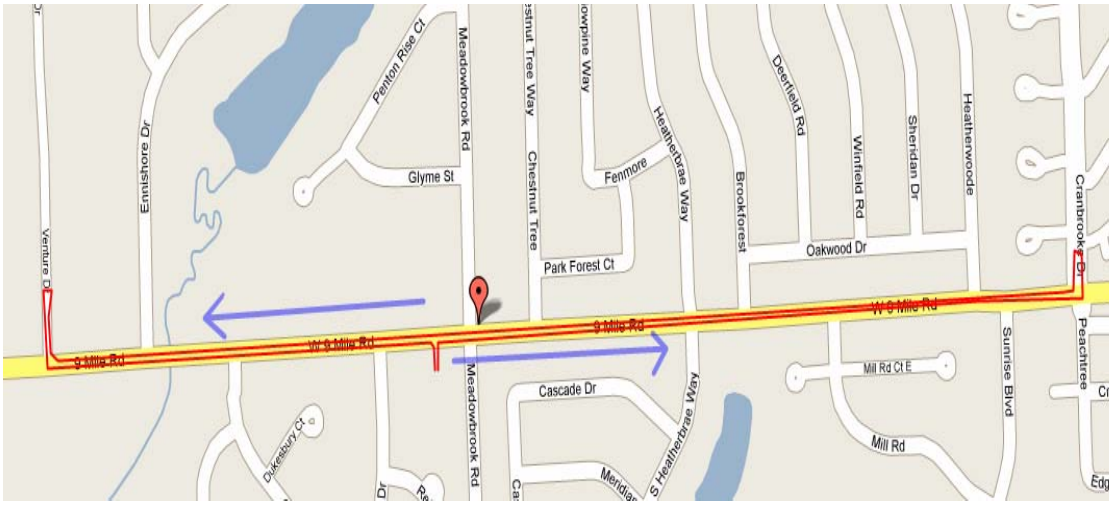

Booz Allen generated routes and plot points on Google map

Figure 65: Route and RSE Day 2 Multi-Apps

The following table lists issues with vehicles during testing.

Table 49: Vehicle Issues during Testing

HB, SPAT & WSA	
	P1: Did Not receive messages
HB, SPAT, WSA, & PDM	
	P9: HB not Transmitting
	P2: not logging
	P1: Parked, not receiving messages
	C4: Received Messages, but did not display them
	C1: Received Messages, but did not display them. Wrong Logging parameters
	P6: Received Messages, but did not display them. Wrong Logging parameters
	P8: Received Messages, but did not display them
	P7: Received Messages, but did not display them
	P4: Received Messages, but did not display them
HB & SPAT & WSA & PDM & Signage	
	P9: HB not Transmitting
	P2: not logging
	P1: Parked, not receiving messages
	C4: Received Messages, but did not display them
	C1: Received Messages, but did not display them.
	P6: Received Messages, but did not display them.
	P8: Received Messages, but did not display them
	P7: Received Messages, but did not display them
	P4: Received Messages, but did not display them

SPAT

SPAT messages were sent at 10Hz or about 10 times per second by the only RSE on the route. SPAT data collected from the VII communications safety link testing on June 19 at the same RSE was used as the baseline.

For this analysis, the time period covered should have been 11:23 am on September 12, when SPAT was turned on until 3:30 pm. However, only one vehicle reported SPAT data for the 11:23 am to 12:39 pm period. Thus, it was not included in this analysis. Therefore, only data from 12:39 pm until 3:30 pm was used.

Note about OBE Logs

SPAT messages were only logged by 7 OBEs; however, 25 vehicles drove the route. It is possible that the SPAT application did not start on the remaining OBEs. OBE C453 (C1) was not included in the analysis because it only reported partial data. OBE CC94 (A11) was excluded because its numbers were outside of the norm; at one point, it reported receiving 211 SPAT messages in a 4-second period, which is about five times faster than the rate at which the RSE was broadcasting. OBE C779 (A6) was excluded for a similar reason. It reported 828 SPAT messages in a 50-second period, which is about 1.6 times the sent rate, and the average latency for this period was 50 seconds.

Although these logs were excluded for now, they may prove to be useful in the context of application interplay on the OBEs if further, more detailed, analysis is undertaken.

It should be noted that to calculate latencies in milliseconds, one needs extremely accurate clocks on the RSEs and OBEs. The times must also be synchronized before the start of the test. Figure 66 shows the percentage of SPAT messages that were received "after" they were sent. Although the numbers are better than the baseline, they still never reached the 50-percent mark. The clocks were either not synchronized, not accurate enough, or both.

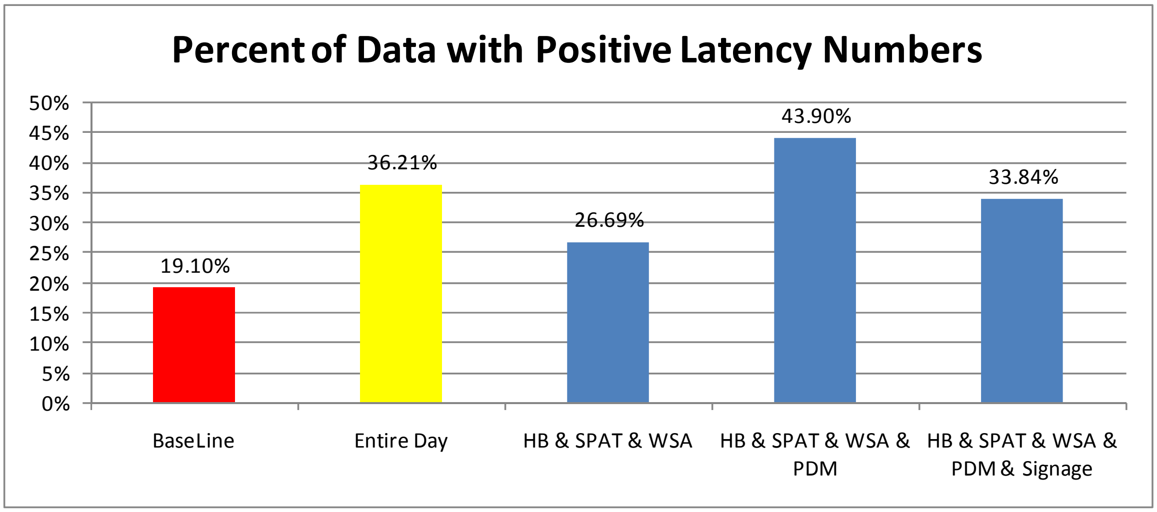

Figure 66: Percentage of Data with Positive Latency Numbers

Note about RSE Logs

SPAT listener logs were not collected from the RSE because it was not turned on. Therefore, only the OBE logs were used for analysis. This may or may not be a problem, but to verify the findings, the tests would need to be repeated and re-analyzed with both OBE and RSE logs.

Measure 1: Frequency at which SPAT Messages Were Received

Measure Objective: Determine the frequency at which OBEs received SPAT messages from the RSE.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate, using the OBE logs, the frequency at which SPAT messages were received.

Results: Figure 67 shows the frequency at which SPAT messages were received by OBEs.

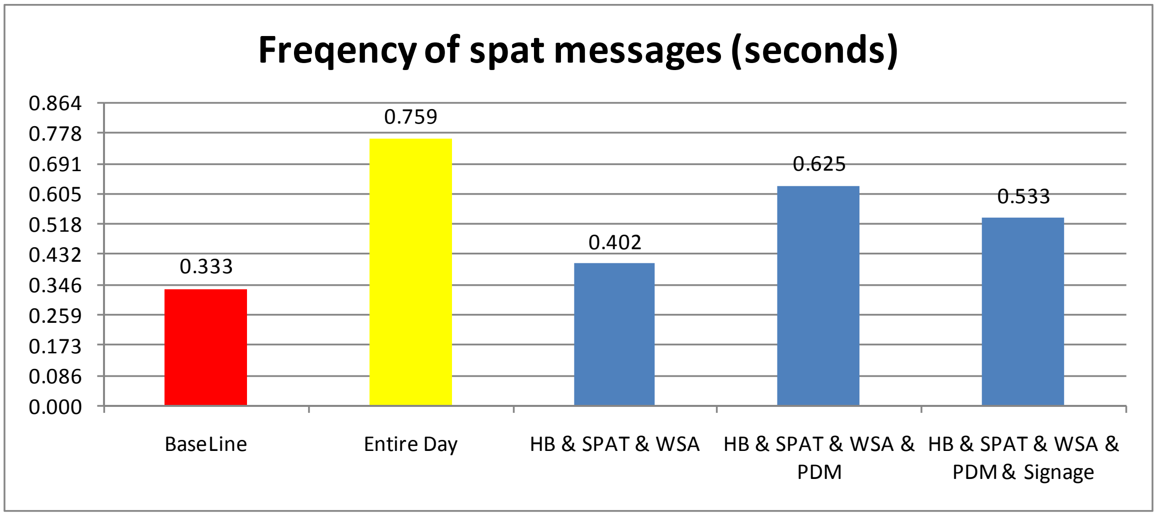

Figure 67: Frequency of SPAT Messages

Overall, frequency rates were better than the baseline data, but the frequency rates never got close to the 10Hz theoretical limit. Frequency rates did not decrease as the number of active applications on the OBE increased; instead, they went up and down. The frequency increased when PDM was added and decreased as public messages were added.

Figure 68 breaks down the data further into parked and traffic testing.

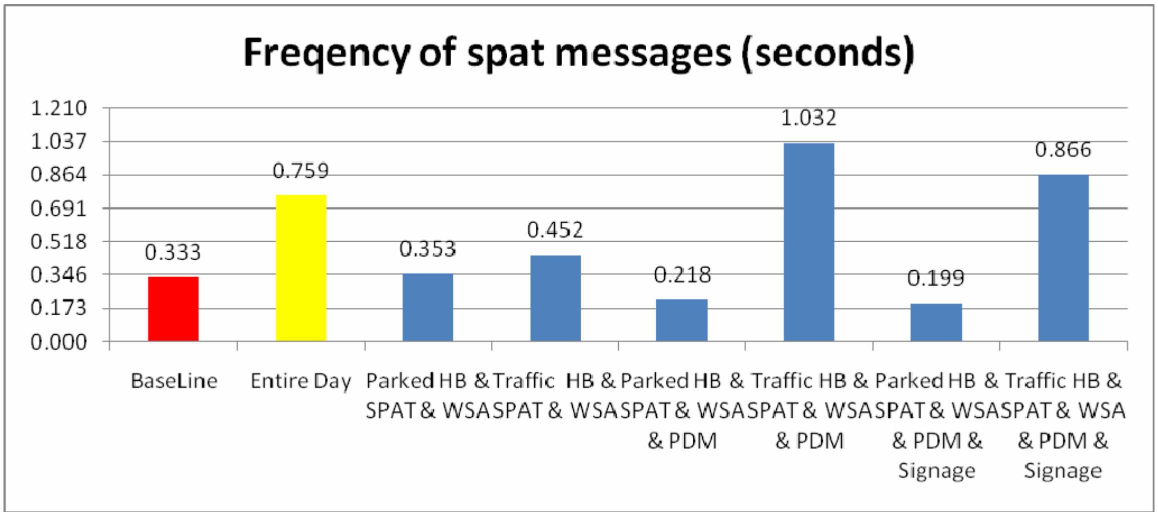

Figure 68: Detailed Frequency of SPAT Messages

The frequency rate increases during driving tests and decreases during parked tests. This finding is not in line with the hypothesis.

To completely understand the results, more analysis needs to be done on the interplay of applications on the OBE and each application's priorities as it pertains to OBE's logging capability.

Measure 2: Latency of SPAT Messages

Measure Objective: Determine the latency between the time the RSE sent out the SPAT message (as seen in SPAT message payload) and the time it was received by the OBE.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate, using the OBE logs, the latency of SPAT messages sent by the RSE and received by the OBE.

Results: Figure 69 shows average latency in seconds.

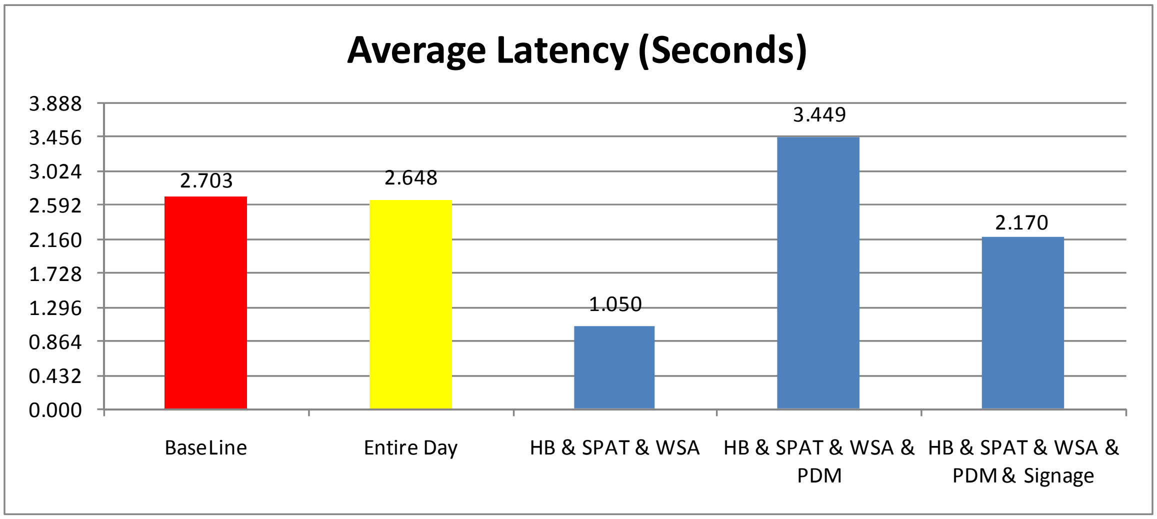

Figure 69: Average Latency

The average latency achieved for the entire testing period (11:23 am to 3:30 pm) is almost identical to the baseline data.

Broken down by groups of applications, the results show that HB, SPAT and WSA actually achieved a lower latency than the baseline. When PDM was added to the mix, a spike of nearly 2.5 seconds could be observed. However, what is interesting is the fact that when public messages are sent to the RSE, the SPAT latency drops to what is considered normal levels.

To get better insight into the rise and fall of SPAT latency, the numbers were broken down further into parked and traffic testing.

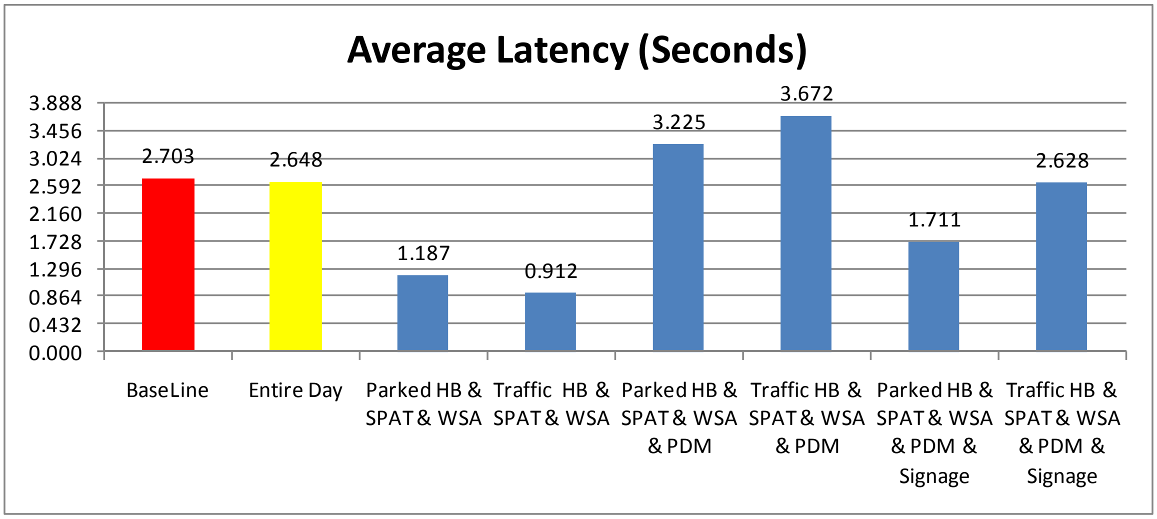

Figure 70: Detailed Average Latency

The results paint a similar picture, with the same up and down behavior after adding PDM and public messages.

Analysis: The hypothesis was that as more applications are turned on, the latency would continue to rise linearly. However, that is not the case. To better understand this phenomenon, more detailed analysis of all application logs needs to be carried out.

Conclusion: SPAT frequency and latency displayed some unexpected behavior. The results did not support the hypothesis that the frequency would fall and latency would rise as more applications were turned on. Further testing and analysis is necessary to explain this phenomenon.

PDM

On Friday, September 12, a PDM message with time-based collection method was sent to the vehicles. The testing continued for about 2 hours after the PDM message was sent.

The PDM message included the information shown in Table 50.

Table 50: Contents of PDM Message on August 28

Sample Start/End		Termination Method	
0	99	Time	180 s
Collection Method			
Speed 1	Speed 2	Time 1	Time 2
4 mph	15 mph	4 s	10 s
Start Stop Strategy			
Stop Threshold	Stop Time	Stop Lag	Start Threshold
2 mph	5 s	15 s	10 mph

Note about OBE Logs

For Friday, September 12, data was received from seven vehicles that had received the PDM message. However, 25 vehicles should have responded to the message.

Table 51 lists the issues experienced with vehicles during testing.

Table 51: Vehicle Issues during Testing

P9 (C082): HB not transmitting
P2 (C862): not logging
P1 (C482): Parked, not receiving messages
C4 (B420): Received messages, but did not display them
C1 (C453): Received messages, but did not display them. Wrong logging parameters
P6 (C194): Received messages, but did not display them. Wrong logging parameters
P8 (C548): Received messages, but did not display them
P7 (C992): Received messages, but did not display them
P4 (C256): Received messages, but did not display them

Measure 1: Percentage of Periodic Snapshots within 4 Seconds of Anticipated Time

Measure Objective: Determine whether OBEs under PDM generated periodic probe data at the frequency specified in the PDM message.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This in turn allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed-generation frequency (e.g., 4 seconds), verify in OBE logs that messages were generated at that frequency.

Results: For the message in question, the OBEs were to generate a periodic snapshot every 4 seconds if the speed was less than 4 mph. They were also to generate a periodic snapshot every 10 seconds if the speed was more than 15 mph. For speeds between 4 mph and 15 mph, the snapshot time was to be linearly interpolated.

Figure 71shows in how many seconds a snapshot was actually generated, after its anticipated generation time.

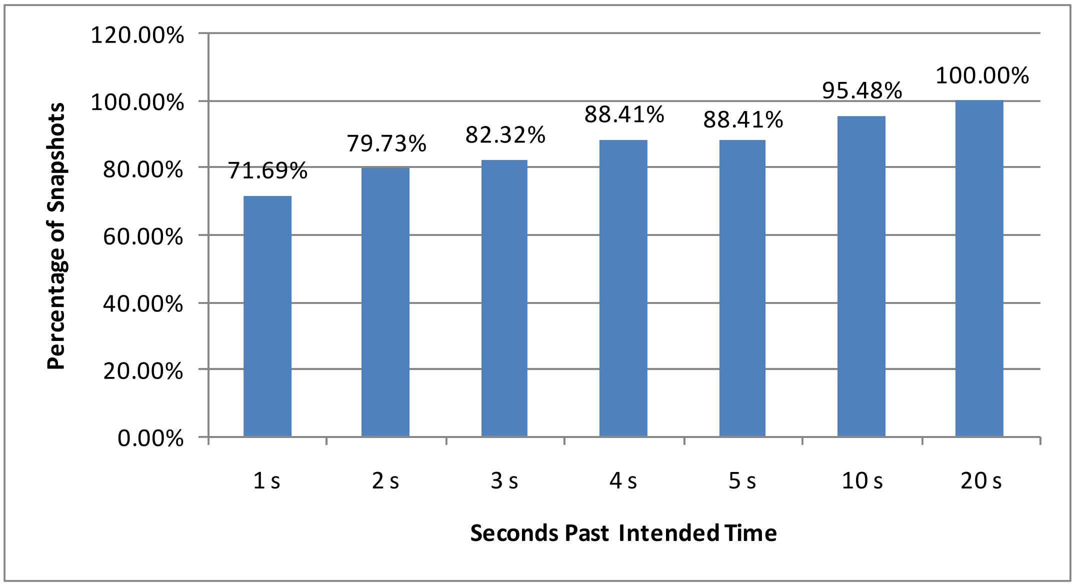

Figure 71: Cumulative Distribution of Snapshots

No snapshots were generated before the intended time, almost 90 percent of snapshots were generated within 4 seconds of the anticipated time, and all snapshots were generated within 20 seconds.

Analysis: The results show that most snapshots are generated within 4 seconds of their anticipated time. This means that if an operator requests snapshots to be generated every 10 seconds, they can expect the majority of them to be generated between 10 and 14 seconds.

Measure 2: Standard Deviation of Differences between Intended Time and Actual Time

Measure Objective: Determine whether OBEs under PDM generated periodic probe data within a reasonable standard deviation of the anticipated time.

Concept: See Measure 1.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed-generation frequency (e.g., 4 seconds), verify in OBE logs that messages were generated within a reasonable standard deviation of the anticipated time.

Results: Only four OBEs reported enough data to calculate this measure. However, the computed standard deviation of 2.8 seconds is in line with what was observed during PDM-only testing.

Measure 3: Percentage of Time-Based Expirations -- Percent of All PDM Activations for a Given Serial Number That Expired Based on Time

Measure Objective: Determine the percentage of PDM activations that expired based on time.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message. PDM messages also expire when the OBE leaves the RSE footprint, which is usually about 700 meters wide.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: In OBE OSGi logs, divide the total number of time-based expirations of a particular serial number by the total times that serial number became active.

Results: Time-based expiration is defined as the OBE responding to a PDM message until its anticipated expiration time.

Time-based expirations are affected by overlapping RSEs and close-by RSEs. If the OBE receives another unique message, then the latest message becomes active and the first PDM message expires. The same is true for situations where the OBE leaves the RSE footprint prematurely.

All OBEs had a time-based expiration rate of 100 percent.

Analysis: A success rate of 100 percent was achieved, compared to only a 40-percent success rate during PDM-only testing. This most likely is due to the fact that there were no overlapping RSEs on the route, and only one PDM message was playing on the only RSE.

Measure 4: Percentage of Time-Based PDM Expirations That Expired Correctly

Measure Objective: For PDM messages that expired based on time, determine whether their expiration was at exactly 30 seconds.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message. PDM messages also expire when the OBE leaves the RSE footprint, which is usually about 700 meters wide.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in OBE logs that each time-based expiration met the criteria defined in the PDM message.

Results: The results showed that if a message expired based on time, it did so correctly 100 percent of the time. However, if an OBE leaves the RSE footprint before the anticipated expiration time, then that message expires immediately. This behavior is according to specifications, and the OBE data shows that it was implemented correctly.

Expiration time allows the TMC operator to capture more or less data for a certain period. The results show that their specified time will be executed correctly.

Analysis: The PDM message expiration times were carried out successfully 100 percent of the time.

Measure 5: Percentage of Stop Snapshots with Current Speed below Stop Threshold

Measure Objective: Determine whether OBEs under PDM generated stop probe data at speeds below the stop threshold specified in the PDM message.

Concept: Stop threshold is used to indicate the minimum speed that indicates vehicle movement. For example, on highways with a speed limit of 65 mph, a vehicle traveling less than 10 mph may be considered "stopped." The resulting stop snapshots can be used to detect traffic congestion, accidents, and other unforeseen events.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each stop snapshot generated, verify that the snapshot speed was less than the stop threshold speed.

Results: The stop threshold goal was met 98.8 percent of the time. Only three snapshots had speeds higher than the stop threshold. They were generated by OBEs C985 (A8) and C594 (A2). Neither of these vehicles had reported any problems during testing.

Analysis: The results show that stop thresholds were implemented correctly, and OBEs responded with stop snapshots. However, stop snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear that such units are commercially available at a reasonable cost.

Measure 6: Percentage of Stop Snapshots Generated after Stop Lag

Measure Objective: Determine whether OBEs under PDM generated stop probe data after the stop lag specified in the PDM message.

Concept: Stop lag is used to prevent the OBE from generating too many stop snapshots when the vehicle is creeping along at speeds below the stop threshold. Stop lag prevents the OBE buffer from filling up with nearly identical stop snapshots and overwriting periodic and/or event-generated snapshots.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate the difference between the current and previous stop snapshots and verify that it is after the specified lag time.

Results: 100 percent of stop snapshots met the 10-second stop lag time successfully. Figure 72 shows how many seconds after the stop lag time the next stop snapshot was generated.

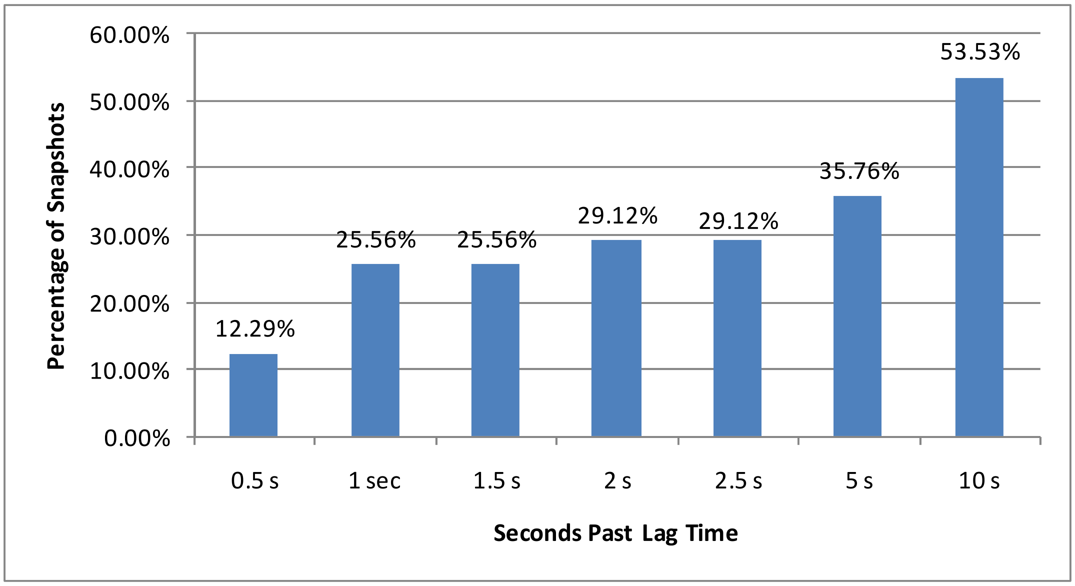

Figure 72: Percentage of Snapshots past Lag Time

Only 35 percent of snapshots were generated within 5 seconds past the lag time.

Analysis: Stop lag is useful in preventing the fill up of OBE memory with almost identical snapshots. However, it is recommended that, in the future, OBE implementations support faster snapshot generation and by default store more snapshots (recommendation of 300) in memory.

Measure 7: Percentage of Start Snapshots with Current Speed above Start Threshold

Measure Objective: Determine whether OBEs under PDM generated start probe data at speeds above the start threshold specified in the PDM message.

Concept: Start threshold is used to indicate the minimum speed at which the vehicle is no longer considered “stopped.” The resulting start snapshots can be used to detect the end of traffic congestion and clearing of accidents among other things.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each start snapshot generated, verify that the snapshot speed was more than the start threshold speed.

Results: Only six start snapshots were received for the duration of the test. All six snapshots had speeds higher than the start threshold.

Analysis: Start snapshot generation on vehicles may not be working properly. Twenty-three start snapshots were received during one of the PDM-only testing days, and only six start snapshots were received during multi-apps PDM testing. The OBE software needs to be retested to make sure it corresponds to specifications.

Conclusion

The numbers received for multi-apps testing compared to pure PDM testing were better in terms of on-time percentages. The difference between the PDM message sent for PDM testing and the one sent for multi-apps testing was the start stop strategy. It was not altered for multi-apps testing. To obtain a better comparison, other PDM messages from the PDM-only testing days should be analyzed.

Advisory Messages (AMDS)

Public signage for multi-apps was conducted on September 12. A total of 20 messages with varying parameters were sent to the RSE. One hundred percent of them made it to the designated RSE successfully.

Note about OBE Logs

Data from only 11 vehicles was received, as listed in Table 52.

Table 52: OBEs with Data

OBE ID	Group ID
C366	P12
C453	C1
C548	P8
C590	P10
C594	A2
C694	P3
C832	P5
C862	P2
C985	A8
C992	P7
CC94	A11

Measure 1: Loss of Public Messages

Measure Objective: Determine whether all messages sent by Public Apps were received by OBEs.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Using the OBE logs, find out if every message appeared at least one time in one of the OBE logs.

Results: 100 percent of all messages sent were received successfully by the OBEs. The TMC operator can be confident that all messages, including civil emergency messages, will make it to the RSE and ultimately to the OBE, as long as the RSE is online.

Measure 2: Roundtrip Latency of Messages from Public Apps to RSE

Measure Objective: Determine the latency among the times that the Public Apps sent out the signage message, it was received by the RSE, and Public Apps received a reply from the RSE.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Only the messages that received a response from the RSE should be counted. Also, since the Public Apps database does not record exact sent time and RSE acknowledgement receipt time, the data in the AMDS broker database, which is only one step removed from the Public Apps, was used.

Results: The average latency could not be calculated because the “Sent” and “Received” timestamps in the AMDS broker database reported incorrect data.

The “Sent” timestamp for all 20 messages was 13-SEP-08 10:15:45.711. The “Received” timestamp was either 14/SEP/08 12:47:13.350 or 13/SEP/08 10:15:36.350, resulting in a latency of either 1 day 2:31:27.639 or a negative latency of 00:00:09.361. Further, detailed analysis of the data is necessary to determine the cause of this anomaly.

Measure 3: Messages Became Active and Expired at Times Specified by the Operator

Measure Objective: Determine whether the activation and expiration times of messages on the OBEs were the same as specified by the operator.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in the OBE logs that the messages display time falls within the active time frame.

Results: 100 percent of messages were displayed and removed at times specified by the operator. It is important for an operator to be able to control exactly when a message becomes active and when it expires. Given the results, operators can be highly confident that messages will be displayed and removed from the OBEs on time.

Measure 4: Messages Were Only Displayed within Correct Active Regions and Direction

Measure Objective: Determine whether the messages were only displayed within the active region and removed once the vehicle exited that active region.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: The percentage of HMI displays is calculated as follows: ((Valid Display + Valid Reject)/ (Valid Display + Valid Reject + Invalid Display + Invalid Reject)) * 100.

Results: 100 percent of messages were displayed in the correct activation regions, and 100 percent of messages that were received but should not have been displayed, due to directionality and/or activation time, were not displayed.

It is very important that drivers only receive messages that apply to them, in order not to overload them with unnecessary information. To do this, the operator needs to be able to specify exact presentation location and direction of travel for vehicles. The results show that both active regions and directionality are implemented correctly within the OBEs.

Measure 5: Only RSEs Defined in the Broadcast Strategy Played the Messages

Measure Objective: Determine whether the messages were only played on RSEs that were defined in the broadcast strategy.

Data Source: RSE proxy manager logs.

High-Level Data Analysis Plan: Count the messages that were sent out on all RSE logs, then verify that only RSEs that should have received a certain message only shows those messages and not others.

Results: In a fully rolled out VII implementation, there will be a lot of traffic going through the network, and it will be imperative that the system is not overloaded with unnecessary information.

The test results showed that 100 percent of broadcast strategies were implemented correctly. The messages were only played on RSEs defined by the broadcast strategy.

Measure 6: Did the RSE Receive the Correct Broadcast Priority (0-7) as Sent by the Operator?

Measure Objective: Determine whether the correct broadcast priority was received by the RSEs.

Data Source: RSE proxy manager logs.

High-Level Data Analysis Plan: In the RSE proxy manager logs, "Message Priority" refers to broadcast priority. Using the message priority and message callback ID, verify that correct broadcast priorities were received by the RSEs.

Results: Broadcast priorities define how often a message is played on the RSE playlist. A priority of 0 means that a message will be played every 100 millisecond, and a priority of 7 means that a message will be played every 100,000 millisecond. The broadcast priority is also used to order the playlists; lower numbers have higher priorities—so, for example, a civil emergency message with a priority of 0 would be played before a travel time message is played.

One hundred percent of all broadcast strategy priorities were received and decoded by the RSE successfully.

Measure 7: Did the OBE Display the Highest Priority Messages First?

Measure Objective: Determine whether the messages with the highest priority were displayed on top of other messages.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify that all lower-number (higher priority) messages were displayed before the higher priorities.

Results: Once the OBE receives a list of messages, it has to prioritize which one will be displayed on top of the other messages. For example, an Amber Alert will take precedence over speed limit signs.

The results showed that 100 percent of messages were prioritized successfully.

Measure 8: Message Content – Did the OBE Display the Message Content (e.g., ITIS codes and free text) as Sent by the Operator?

Measure Objective: Determine whether the message content stayed intact during transmission to the OBE and subsequent decoding.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Using the OBE logs, determine whether the OBE receive the correct ITIS code, category, and free text components of the messages.

Results: message content is just as important as message priority. Drivers have to see exactly what the operator intended for them to see. A public signage consists of a series of ITIS codes and free text. The results showed that 100 percent of the content of the messages were successfully received, decoded, and ultimately displayed to the drivers.

Conclusion

The results of the tests were perfectly in line with the results from the single-app testing days. It seems that multiple applications on the OBE do not affect the correct processing of AMDS messages.

5. RECOMMENDATIONS FOR FUTURE WORK

The following sections provide recommendations for future work, based on the POC testing led by Booz Allen. Recommendations based on the POC testing led by the VIIC are documented in the VIIC's final report. The reader is encouraged to also read the VIIC's report to obtain a complete picture of the POC recommendations.

5.1. Dedicated Short-Range Communications

DSRC radio findings indicate that adequate range can be achieved in most conditions, although typical roadside link quality is significantly different from that measured in open field testing due to roadside furniture. The effects of multipath can clearly be discerned, and further study is required to determine the full impact and to assess the potential mitigation approaches.

It is also recommended that further analysis should be conducted to determine the appropriate protocol mechanisms needed within the WAVE/DSRC protocol suite to effectively measure and leverage link quality to improve data transfer under suboptimal conditions. For example, this might include the incorporation of an RSSI (Received Signal Strength Indicator) feature to better control when selected communications links should be established.

WAVE/DSRC-based security mechanisms generally work as intended and can effectively provide the authentication and authorization features required to meet stated privacy and security objectives. It is recommended that the two security-related IP protocols, V-DTLS and V-HIP, developed as part of the VII POC system, be further developed with the end goal of submission to the appropriate standards bodies.

In summary, the VII POC system successfully demonstrated the core DSRC-based vehicle-to-vehicle and vehicle-to-infrastructure communications functions. However, POC testing identified issues with the prototype WAVE/DSRC radio implementations, mostly due to limitations and shortcomings in the 1609 WAVE protocol standards suite. The majority of these shortcomings resulted from the dynamic nature of the mobile radio relative to the stationary radio and to other mobile radios. The WAVE/DSRC standards and the next-generation radio implementations need to be refined to include signal quality measures, enhanced service processing logic (unbalanced links), and multi-RSE service arbitration logic. The two secure communications protocols (V-HIP and V-DTLS) developed for the POC, and optimized for use with the WAVE/DSRC 1609 protocols, need to be further developed with the end goal of submission to the appropriate standards bodies.

5.2. Probe Data Service

The PDS showcased the inherent advantages of collecting anonymous vehicle sensor information as automobiles traverse the nation's roadways. While the POC effectively demonstrated core PDS functionality, there were some areas with room for improvement. For the VII network to operate more effectively, future VII efforts should investigate the following:

- **Scalability Testing of Probe Data Messages –** Further studies should be conducted to load and stress test the PDS in order to determine the optimal number of probe data messages that a single subsystem can physically process. These performance metrics are important to determine realistic software and hardware requirements for a production system.

- **Probe Date Message Set Optimization –** The probe data message set was defined in an ASN.1[2] format by the Society of Automotive Engineers (SAE). The current message structure contains multiple nested data elements resulting in significant data overhead. Additionally, the current nested structure creates unnecessarily complicated parsing and data extraction conditions. If the message set was defined in a more flat structure, the PDS would see performance improvements in encoding, decoding, and parsing.
- **Creation of a Data Abstraction Layer for Probe Data Messages –** The probe data message is a dynamic data structure with the potential for hundreds of different types of sensor data elements. The probe data message should be defined in a flexible way such that sensor data elements can be easily added or removed based on industry or public demand. To create a more flexible probe data message structure, an abstraction layer should be developed, thus creating a more flexible software architecture and allowing data to be easily transformed from one type to another.
- **PDS Subsystem Integration with Web Service (WS) Standards** – Instead of using custom-developed Web services, the PDS could benefit from leveraging existing, industry-recognized Web service protocols. The PDS subsystem could potentially use WS-Notification and WS-Resource Management for probe data delivery and subscription services. Further investigation is needed, but utilizing maturing industry specifications could lead to greater interoperability and faster time to market for VII applications that use PDS.
- **Probe Data Loss between the SDN and the Network User** – Modifications to the communication session flow should be investigated to mitigate the data loss issue due to the ARPs.
- **Probe Data Service Outages** – There should be some way in the VII architecture for people subscribing to probe data at any SDN to know about PDS system failures at other SDNs where they are subscribing to probe data.
- **V-DTLS –** The POC V-DTLS implementation should be updated to support multiple threads and allow multiple vehicles to establish V-DTLS sessions with a POC RSE concurrently.

In summary, the PDS adequately met the requirements of the VII system after the implementation of V-DTLS, throughout VII POC testing. Once the RSE received the probe data, losses throughout the system were minimal. Those losses that did occur were found between the SDN and Network User interface, and due to the implemented use of UDP as the transport protocol. A review of alternative protocols to use for this interface and analysis of increased bandwidth required for a new protocol should be conducted to determine whether there are viable options other than UDP. Additionally, changes to the structure of ASN.1 probe data messages should be reviewed due to the complexities found in designing the parsing and data extraction mechanisms resident in the PDS. Although the parsing functionality proved adequate for the POC, increased message throughput could adversely impact future performance with larger-scale probe data tests.

---

[2] The IntelliDrive^{SM} POC Message Sets were developed and interpreted using an OSS Nokalva's ASN.1 compiler and runtime tools.

5.3. Advisory Message Delivery Service

The AMDS enables the distribution of advisory messages from VII network applications to a targeted set of RSEs. The RSEs then transmit messages to any vehicles in their wireless broadcast range. The POC effectively demonstrated core AMDS functionality; however, there were some design and performance limitations. To optimize the AMDS functionality, future VII efforts should investigate the following:

- **Improvements in AMDS Synchronization –** As a fundamental requirement, the AMDS must maintain the status of all advisory messages that are currently being broadcast on the RSEs. Currently, Network Users may submit a single request containing multiple advisory messages to multiple target RSEs. When coupled, these requirements introduce extraordinary complexity on the system, especially under error conditions. It may be more advantageous to limit the Network User to only placing one advisory message per request. A more thorough investigation of this approach and its resulting ramifications should be conducted.
- **Message Recall Function** – The version of the J2735 message set used for the POC did not include a “Message Recall” message. This meant that once a vehicle’s OBE received a message, the OBE would display that message until it expired (or the vehicle was turned off), and there was no way for the Network User who sent that message to cancel it from the vehicles. This problem was identified early in the POC, and new versions of the J2735 message set will include a message recall message that serves this purpose.
- **Advisory Message Status Delay –** Due to the distributed nature of the VII network, there is an anticipated time lag in determining the status of “active” advisory messages. As the AMDS does not respond back to the Network User until determining the status of all RSEs associated with the advisory message, this delay can be somewhat lengthy. The delay is significantly extended when one or more RSEs experience backhaul connectivity issues. The AMDS should be enhanced to intermittently provide “partial” status updates, effectively providing status for those RSEs that respond earlier.
- **Advisory Message Subscriptions** − The AMDS could be enhanced to support Network User subscription functionality, this would allow authorized Network Users to determine what advisory messages are currently active at specific RSEs. A more thorough investigation of this approach and its resulting ramifications should be conducted.
- **Directionality Based on Instantaneous Vehicle Heading –** In the current version of the J2735 message set, if messages are distributed based on vehicle heading, the vehicle heading used is always the heading at the time the vehicle enters the valid (or active) region. However, it is sometimes desirable to send a message instantaneously to all vehicles traveling a given direction in a valid region, regardless of which heading they were traveling when they entered the active region. Consideration should be given to adding a parameter to the J2735 traveler advisory message that specifies whether the heading in the message should be applied only to vehicle heading at entry to the active region, or whether it should be applied continuously while the vehicle is in the active region.

In summary, the AMDS met the basic requirements of the VII system; however, due to deficiencies in the DSRC standard regarding overlapping RSEs, there were instances of message loss. As a result, the DSRC standards should be revisited and a study to determine what mechanisms can be put in place to allow vehicles to receive messages from multiple RSEs with overlapping coverage. Additionally, POC testing found a gap in functionality that should be further explored. In the current architecture, once an

AMDS message is received by a vehicle, there is no way to tell that vehicle that the message has been canceled or revoked by the user who created it. Further study to determine a mechanism to allow RSEs to broadcast a message canceling or revoking AMDS messages should be conducted.

5.4. Communications Service

In general, the VII POC communications service met the basic requirements of the dual stack IPv4/IPv6 nature of the VII architecture. However, prior to national rollout, the conventional Internet mobility and security protocols need to be refined and modified to support dynamic and anonymous mobile users.

5.5. Map Element Distribution Service

The POC MEDS design goal was development of fundamental geo-statistical algorithms proving the viability of updating intersection geometries based on probe data. Given the inherent complexity of the geospatial processing functions, development of the MEDS capability required several primary design assumptions. As MEDS is further developed, these assumptions should be addressed to refine the original geo-statistical algorithms.

- **Vehicle Dimensions** – Vehicle dimensions, both size and shape, have significant impact on the geo-statistical processing. During POC, MEDS assumed each vehicle was the same shape and infinitesimally small, introducing significant error into MEDS results. These dimensions need to be more adequately addressed during MEDS processing to increase the accuracy of updated intersection geometries.
- **GPS Antenna Location** – Given the positional accuracies needed to support MEDS functionality, knowing the location of the antenna on the vehicle is critical to determine its exact location on the roadway and in relation to other vehicles. However, as antenna location has not been standardized across vehicle manufactures, MEDS assumed antenna location in the center of a simplified vehicle, introducing error into MEDS results.

In addition to addressing major design assumptions, continued work on MEDS should include closer coordination with other organizations across the ITS community. Many industry organizations, academic institutions, and commercial vendors are actively performing research into the use of probe data for map updates. Leveraging these efforts in combination with continued MEDS development will facilitate development of the MEDS capability supporting multiple VII use cases.

In summary, the vehicle dimensions, both size and shape, have significant impact on the MEDS geo-statistical processing. These dimensions need to be more adequately addressed to increase the accuracy of updated intersection geometries.

5.6. Positioning Service

The POC POS design was based upon delivery of standard RTCM corrections through the VII infrastructure using the Networked Transport of RTCM via Internet Protocol (NTRIP) standard. While the end-to-end delivery of the RTCM corrections was successfully tested during the POC, the expected accuracies resulting from application of the corrections was not fully realized due to issues with the vehicle OBE used to process the correction data. To eliminate this problem during future POS development and testing phases, the software and hardware in the OBE need to be better aligned with the NTRIP standard to optimize use of the correction information by the vehicle. In addition, several other considerations should be addressed during future POS activities.

- Evaluate security issues related to “spoofing” of GPS signals and potential impact on vehicle position information included in the probe message.
- Evaluate alternate positioning solutions and emerging technologies enabling enhanced delivery of RTCM corrections through the VII infrastructure or other delivery mechanisms. This includes use of alternative correction sources besides the HA-NDGPS.

In summary, the POC POS design was based upon delivery of standard RTCM corrections through the VII infrastructure using the NTRIP standard. Future POS development and testing phases should evaluate alternative positioning solutions and emerging technologies, enabling enhanced delivery of RTCM corrections through the VII infrastructure or other delivery mechanisms, including use of alternative correction sources besides the HA-NDGPS.

5.7. Security

For the VII POC, the lessons learned in deploying WAVE security suggest several areas for continuing research and investigation. The VII 1609 PKI model should be assessed relative to the capability of the 1609.2 protocol to address the VII performance requirement for 250 signature verifications per second.

Methods to detect intrusion and eliminate attackers from the vehicle segment were only briefly studied during the VII POC, but key issues about the robustness of the 1609 certificate system in the face of intrusion were identified and need to be addressed prior to any public deployment. Other topical areas include standardizing security interfaces, performing a critical analysis on anonymity and privacy for vehicle-to-vehicle communications, and implementing Operations System Support (OSS) for automated provisioning, monitoring, and maintenance. Finally, the scalability of the VII 1609 PKI for regional deployments and the ability to provide centralized management of the deployed RSE security services need to be addressed.

A primary objective of the VII POC was the deployment of wireless security using WAVE 1609.2 Security and related protocol services (V-DTLS and V-HIP). The development of these protocols should be continued with the goal of becoming standards. Additionally, the effectiveness of this communications security architecture should be assessed in terms of mitigation of the potential risks to VII services including identity theft, tracking, disruption, and tampering of VII information. Performance criteria regarding latency need to be established regarding the use of 1609 PKI for digital signature and encryption. Enhancements to the VII POC IPv6 protocol architecture should include an assessment of available IPv6 security functions and controls, as applicable, leading to the deployment of a mature IPv6 stack for end-to-end communications. Examples of this might include hardening DNS services and secure routing protocols.

In summary, the VII POC security services have demonstrated the ability to provide trusted data to the vehicle operators and to prevent the misuse of vehicle information by unauthorized users of the system. RSEs and OBEs thus equipped with strong identity credentials have rejected bad messages and avoided threats from would-be attackers. The “Defense in Depth” model deployed for the VII POC security architecture should be used as a baseline and assessed against the National ITS Architectures for Information Security, Operational Security, and Threat Mitigation criteria. The analysis and reporting of quantifiable metrics regarding VII POC security services should be developed to further demonstrate the secure management of VII assets and network services, as well as the security of user data transmitted across the VII POC system from misuse by attackers, unauthorized users, and physical or electronic sabotage. In particular, the use of IPv6 security services should be evaluated for further enhancements to the VII POC "Defense in Depth" model.

5.8. Enterprise Network Operations

The VII ENOC subsystem was the epicenter of the VII POC. Its role was to provide network management and security information services to the underlying VII subsystem architecture. As the primary entry point for access into the VII enclave, the ENOC was responsible for tracking the network health and system security while providing a centralized location for investigation and remediation of VII-related component issues. In the future, large-scale build outs should consider all aspects of a network operations implementation. Identification of the correct people, creation of tested processes, adherence to established security policies, and incorporation of the right technologies are paramount to the deployment of a successful ENOC capability.

5.9. Certificate Authority

The certificate authority requirements of the VII POC system were met though independent linked certificate authority servers, providing a small-scale PKI service, deployed to securely manage the different classes of messages, applications, and services traversing the VII system. These certificate authority capabilities included lifecycle management for public and private keys, digital signatures, and a variety of encryption mechanisms. This certificate authority model (which included both X.509 and 1609 certificates) was proven to be a viable security design given the limited number of RSEs, OBEs, and services deployed in the various VII test beds utilized during the POC testing. As a follow-on to the VII POC program, this security architecture should be assessed in terms of providing a scalable solution for enterprise-wide deployment. Performance modeling and quantitative analysis of the certificate authority are suggested as a means for looking at network throughput, hardware encryption devices, certificate revocation lists (CRLs), and digital signature verification operations.

5.10. Test Bed Deployment

All deployed backhaul technologies (i.e., 3G, WiMAX, and T1) proved capable of supporting the VII network and services requirements. However, the following items should be considered in future deployments:

- T1
    - Initial deployment requires considerable lead time. Be sure to plan 6 to 8 months for the T1 service provider to install required cabling to the RSE site.
    - Specify data circuit as opposed to voice circuit. Be sure the service provider understands the equipment requires a data circuit.
    - Be sure the network and all related hardware supports IPv6 connectivity.
    - There can be high recurring charges of approximately $290 per month per location. This is relatively inexpensive for a T1; typical T1 costs are between $800 and $1,000 per month. Budget accordingly.
- WiMAX

  - Use a licensed frequency band to increase the reliability and stability of network. This will, however, result in higher equipment costs and possibly monthly charges. Budget accordingly.
  - Alternative backhaul solutions should be available for cases when line of sight between the base station (network) and subscriber unit (RSE) cannot be achieved.
  - Alternative WiMAX hardware manufacturers should be researched to ensure the most reliable equipment is deployed.
  - Be sure the network and all related hardware supports IPv6 connectivity.
- 3G
  - Be sure the service provider offers static IP addresses for its 3G devices; this is usually a nominal additional cost per month (e.g., between $3 and $5 on top of the standard 3G data subscription).
  - Sending IPv6 traffic over an IPv4 connection capabilities must be in place, since today's 3G networks do not support IPv6 natively.
- General
  - Many backhaul products do not have the ability to support an implementation of the IPv6 protocol stack. During the implementation of the VII POC, IPv6/IPv4 hybrid solutions took the place of true native IPv6 solutions. Vendors should be encouraged to customize/tailor their products to utilize an IPv6 stack. Moving forward, a mechanism should be in place to track vendor offerings and technical roadmaps and formulate a process to overcome these types of issues.
  - RSE compute platform and power supply should be mounted 10 feet above the ground. This will keep the equipment out of the reach of most vandals while allowing network personnel to service the equipment with a ladder.

In summary, the VII backhaul infrastructure satisfied the basic requirements of the VII POC. Through the use of traditional T1s, 3G technologies, and the private WiMAX network, the VII architecture maintained the various interfaces, devices, and communication paths required for the VII POC network services. The implementation of these technologies, however, was a challenge and, at times, was cost prohibitive. The T1 and WiMAX networks provided the most availability and ease of use due to VII engineers having administrative control. Both networks offered the stability that the RSE and OBE devices required; however, they were also the most cost prohibitive. Additional issues revolved around the lack of IPv6 vendor support. Moving forward, the ITS steward should work with its partners to push IPv6 maturity in its technical roadmaps.

5.11. Network User

Possible improvements to the AMDS specification discovered during the POC testing are the inclusion of an instantaneous directionality field in advisory messages (in addition to the "entering active region directionality" field already present); a clearer statement in the specification that messages should be displayed to vehicles that receive them when the vehicles are already within the active region (in addition to upon entering the active region); a mechanism that would allow the network users to send advisory messages and guarantee themselves that the messages will be displayed to the driver correctly, regardless of how the message displays are implemented in different vehicles; a mechanism to cancel messages that are on the OBE when the cancellation request is sent, not when the OBE is turned off; and a map-based display of all active advisory messages.

Possible improvements of the PDS specification discovered during the POC testing include adding the ability for the network user to subscribe to probe data based on the time and location of snapshot generation (instead of snapshot transmission to RSE); the availability of lane-level maps in the future, which will enable the development of network user applications that rely on lane-level road network data; and a map-based display of all RSEs included in the subscription(s).

In summary, most of the network user needs were met by the Public Apps. However, POC tests revealed some areas in need of further development. Most of these relate to additional features that would greatly enhance the network users' ability to make better-informed decisions. For both the AMDS and PDS, the network users need more fine-tuned control to better manage advisory messages and probe data subscriptions. Specifically, for PDS, further customization of subscriptions would need to be added, and for AMDS, the network user should have the ability to further limit the affected vehicles for a given message.

5.12. Privacy

The following recommendations are made to maintain compliance with privacy expectations and the stated principles and limitations for VII:

- **Manage Privacy Protections Throughout the Life Cycle** – Future enhancements to the system will create many changes both to the data collected and the method of transmission; the information sharing partners and new technologies will be introduced. At such times, the Privacy Impact Assessment will need to be revisited to ensure that new privacy risks have not been introduced.
- **Maintain the Openness Principle** – The VII infrastructure does not collect personal information, nor can it be used to track individuals—therefore it has no privacy risk. The privacy-related risks for VII are perceived risks, not actual risks. The VII Privacy Policy Framework hints at a communication plan. It is recommended that the VII infrastructure team participate and ensure that communication around the VII infrastructure is reviewed for accuracy, as the communications plan gets written.
- **Draft Agreements** – Memorandum of agreements defining data retention practices and data usage need to be drafted for network data subscribers. This is to ensure that measures are in place to protect the probe data collected by the subscribers.
- **Further Evaluate Privacy Risk** – Privacy risk should be further evaluated in the management of anonymity and vehicle-to-vehicle communications. Full assessment of the anonymous certificate authority within applicable Governance, Risk, and Compliance (GRC) policies and legislation should be considered.

In summary, the privacy requirements have been assessed in both the vehicle-to-infrastructure and vehicle-to-vehicle segments of the network. The term "anonymity" describes the POC design goal that, as far as possible, broadcast transmissions from a vehicle operated by a private citizen should not leak information that can be used to identify that vehicle to unauthorized recipients. While vehicle-to-infrastructure privacy requirements were met by examining traceability and identity in probe data transmission, the vehicle-to-vehicle privacy aspects of the VII POC suggest broader operational requirements beyond existing standards (IEEE 1609 Security). As such, technical innovations to protect the anonymity of the driver and vehicles require further examination of VII POC privacy protection technology (use of protocol innovations and anonymity protection mechanisms). To meet the privacy requirement of the VII POC, further assessment of available trust models (privacy by design, privacy by trust) are needed to evaluate certificate management, "bad actor," and scalability factors for the network. Privacy constraints in the vehicle-to-vehicle segment need further research to consider various issues including how the system can sustain more compromised vehicles with reasonable overall system performance, how misbehaving vehicles can be eliminated without compromising their privacy, and how consistent or predictable levels of privacy can be maintained.

5.13. Standards

Based on lessons learned during POC development, deployment, and testing, modifications will be recommended to the standards.

5.13.1. SAE

Lessons learned from the POC resulted in recommendations for modifying the following message sets:

- Probe data
- PDM
- Traveler advisories
- (In-vehicle) Signage.

Recommendations and lessons learned were fed back to the SAE throughout the duration of the POC. The majority of the feedback will be incorporated into Version 2 of the standard.

A "Developer Notes" section was added to applicable appendices in the standard to advise developers of some pitfalls to avoid when creating applications based on J2735.

1609

Recommendations from lessons learned during POC include, but are not limited to, the following:

- Accommodations should be made for overlapping RSEs.
- The "User" should have a mechanism for stabilizing which WBSS is joined when the same Provider Service ID (PSID) is advertised on overlapping RSEs or RSEs in close proximity.
- Use and number of priorities should be better defined.
- PSID and PSC needs should be better defined, possibly reworked entirely.

802.11p

Recommendations from lessons learned during the POC include, but are not limited to, the following:

- The size of the contention window should be increased to reduce the likelihood of collisions (i.e., two devices transmitting at the same time).

In summary, some refinements are required to the DSRC standards implemented for the POC (i.e., SAE J2735, IEEE 1609.2, 1609.3, 1609.4, and IEEE 802.11p) based on lessons learned during the POC. The usage descriptions and content of some J2735 messages need to be updated to avoid uncertainty during development and implementation. A constant line of communications was open between POC developers and the SAE during development, resulting in the majority of the issues being addressed in J2735 version 2.0. Refinements to the IEEE standards are required to better accommodate the nature of the vehicular environment. Specifically, standards refinements should be made in the area of vehicle-to-vehicle communications, where both transmitter and receiver are moving simultaneously and in the area of vehicle-to-infrastructure communications regarding overlapping RSE coverage areas where multiple RSEs are broadcasting the same PSIDs.

APPENDIX A: ACRONYMS

AMDS	Advisory Message Delivery Service
ARP	Address Resolution Protocol
BSP	Bootstrap
CA	Certificate Authority
CICAS	Cooperative Intersection Collision Avoidance System
CNSI	Center for Networks and System Innovation
COMM	Communications Service
COTS	Commercial Off the Shelf
CRL	Certificate Revocation List
CSIRC	Computer Security Incident Response Capability
CSP	Credentials Service Provider
DSRC	Dedicated Short-Range Communications Link
DOT	Department of Transportation
DSRC	Dedicated Short Range Communications Link
ENOC	Enterprise Network Operations Center
EPS	Electronic Payment Service
FHWA	Federal Highway Administration
GID	Geographic Intersection Description
GPS	Global Positioning System
HIF	Hardware Interface
HIP	Host Identity Protocol
HMM	Health Management Manager
IdAM	Identity and Access Management
IDD	Interface Design Description
IEEE	Institute of Electrical and Electronics Engineers
ILS	Information Lookup Service
IP	Internet Protocol
IPSec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6

IRS	Interface Requirements Specification
ITIS	International Traveler Information Systems
LAN	Local Area Network
LCS	Life Cycle Service
LDAP	Lightweight Directory Access Protocol
LSS	Local Safety System
MIB	Management Information Base
MEDS	Map Element Distribution Service
MTBF	Mean Time Between Failure
MTBM	Mean Time Between Maintenance
NAP	Network Access Point
NMS	Network Management System
NTRIP	Networked Transport of RTCM via Internet Protocol
OBE	Onboard Equipment
OBNA	Off-Board Navigation
OEM	Original Equipment Manufacturer
PDC	Probe Data Collector
PDM	Probe Data Management
PDS	Probe Data Service
PER	Packet Error Rate
PFM	Platform
PKI	Public Key Infrastructure
POC	Proof of Concept
POS	Positioning Service
PRXY	Proxy
PSOBE	Public Service Onboard Equipment
PSN	Probe Segment Number
RFC	Request for Comment
RIS	Roadside Infrastructure Service
RSE	Roadside Equipment
RSE PM	RSE Proxy Manager
RSR	RSE Subsystem Requirements
RSU	Road Side Unit

RTCM	Real Time Correction Messages
RTG	Routing Services
SAE	Society of Automotive Engineers
SDN	Service Delivery Node
SEC	Security Service
SIT	System Integration and Testing
SNMP	Simple Network Management Protocol
SPAT	Signal Phase and Timing
TAM	Transportation Area Maps
TCP	Transmission Control Protocol
TCS	Traction Control System
TOC	Transportation Operations Center
TSR	Transaction Service Router
UDP	User Datagram Protocol
US DOT	U.S. Department of Transportation
V-DTLS	Vehicular Datagram Transport Layer Security
V-HIP	Vehicle-Host Identity Protocol
VII	Vehicle Infrastructure Integration
VIIC	Vehicle Infrastructure Integration Consortium
VPN	Virtual Private Network
WAN	Wide Area Network
WAVE	Wireless Access in Vehicular Environments
WBSS	WAVE Basic Service Set
WSA	Wave Service Advertisement
WSM	Wave Short Message

APPENDIX B: COMPLETE DISCUSSION OF APPLICATION TESTING RESULTS

Probe Data Management (Test Case AMDS-02)

Analysis for Thursday, August 28, 2008

PDM Message Selected

The PDM message analyzed for August 28 had the parameters shown in Table 53: Contents of PDM Message on August 28.

Table 53: Contents of PDM Message on August 28

Sample Start/End		Termination Method	
15	99	Time	30 s
Collection Method			
Speed 1	Speed 2	Time 1	Time 2
15 mph	30 mph	4 s	10 s
Start/Stop Strategy			
Stop Threshold	Stop Time	Stop Lag	Start Threshold
2 mph	10 s	10 s	5 mph
Message ID: 3973328C-4DB3-4704-925D-FC3D43DF78BD			

This message was selected for detailed analysis because it had the highest occurrence in the OBE logs. It appeared a total of 18,140 times across all OBEs. The message with the second highest occurrence in the OBE logs came in 8,003 times.

Methods

On Thursday, August 28, PDM messages with time-based collection methods were sent to the vehicles. The vehicles were sent out for a total of four runs, and all runs except run 4 had PDM messages broadcasting from the RSEs. Run 4 did not include PDM messages because very little data was received in the morning. Turning off the PDM message was used to determine its effect on probe data generation and/or transmission. The vehicles were divided into groups of five and sent on five routes. On runs 1 and 2, the vehicles were spaced 3 minutes apart, and there was no spacing during runs 3 and 4.

Table 54: Test Start and End Times for PDM on August 28

Run	Start	End	Total Loops
Run 1	10:00 am	11:20 am	2
Run 2	12:15 pm	1:32 pm	2
Run 3	2:00 pm	3:19 pm	2
Run 4	3:40 pm	5:05 pm	2, except vehicles on route 4 did only 1 loop

Note about OBE Logs

For Thursday, August 28, data was received from 21 vehicles that had received the PDM message. Only 20 vehicles (4 routes with 5 vehicles per route) should have responded, but as seen on the Google Maps screenshot in Figure 73, vehicles on route 5 came very close to the I-275 and 8 Mile RSE.

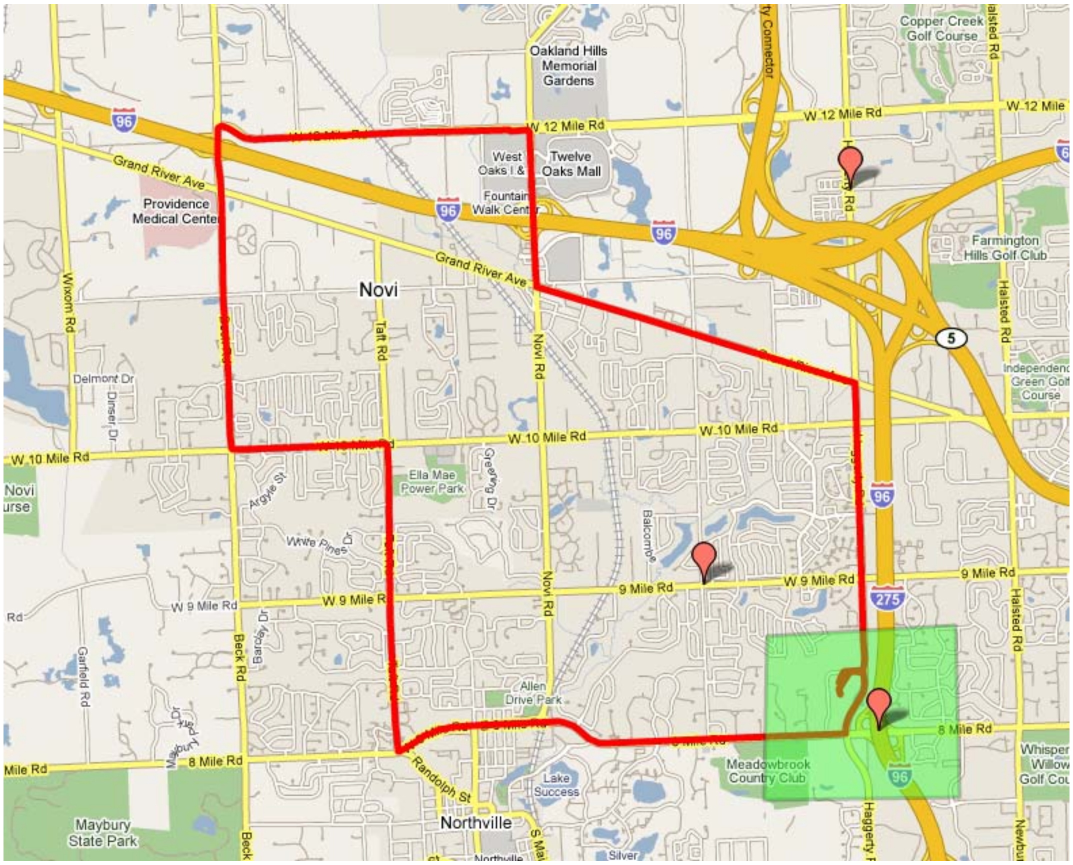

Booz Allen generated routes and plot points on Google map

Figure 73: Route 5 to RSE proximity

OBE B856 was disqualified from inclusion in this analysis due to apparent positioning system errors. One of the PDM messages that it received, stayed active for over 6 minutes, but the OBE reported it as expiring within 30 seconds.

Table 55 presents the issues that were experienced with other vehicles during testing.

Table 55: Vehicle Issues on August 28

C1 (C453)	Lost GPS comm's 1 pulse per second (pp) active before session 2
E4 (B325)	No Positioning, 1pp active after session 1
P1 (C482)	Lost GPS comm's 1pp active after session 1
P10 (C590)	Probe Generation Error after Session 1
P12 (C366)	Probe Generation Error after session 3
P12 (C366)	Screen locked up, OBE dead after session 2
P8 (C548)	Screen locked up, OBE dead after session 2
P9 (C082)	No vehicle speed (CAN) data, after session 2

Measure 1: Percentage of Periodic Snapshots within 4 Seconds of Intended Time

Measure Objective: Determine whether OBEs under PDM generated periodic probe data at the frequency specified in the PDM message.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For the RSEs broadcasting PDM messages with a fixed snapshot generation frequency (e.g., 4 seconds), verify in the OBE logs that snapshots were generated at that frequency.

Results: For the message in question, the OBEs were to generate a periodic snapshot every 4 seconds if the speed was less than 15 mph. They were also to generate a periodic snapshot every 10 seconds if the speed was more than 30 mph. And for speeds between 15 and 30 mph, the snapshot time was to be linearly interpolated.

Figure 74 shows the number of seconds after a snapshot's intended generation time that it was actually generated.

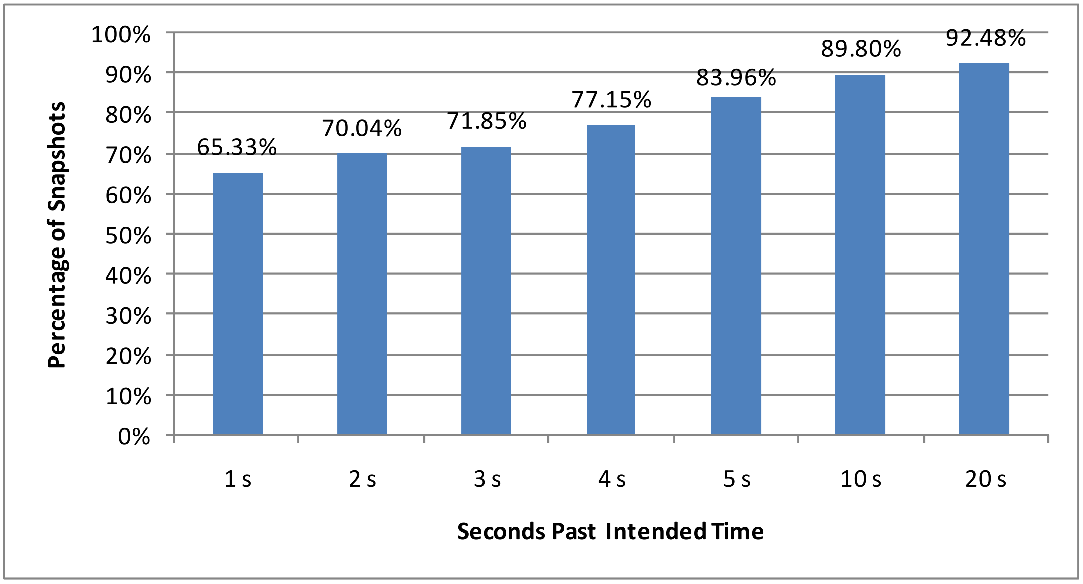

Figure 74: Cumulative Distribution of Snapshots – PDM August 28

No snapshots were generated before the intended time, almost 77 percent of snapshots were generated within 4 seconds of the intended time, and almost all snapshots were generated within 20 seconds.

Analysis: The results show that almost 77 percent of snapshots are generated within 4 seconds of their intended time. This means that if an operator requests snapshots to be generated every 10 seconds, they can expect the majority of them to be generated between 10 and 14 seconds. Having the ability to alter periodic snapshot generation will indeed be useful. However, it would be most useful if the operators were not limited by OBE shortcomings. It is recommended that, in the future, OBE implementations support faster snapshot generation and by default store more snapshots (recommendation of 300) in memory.

Measure 2: Standard Deviation of Differences between Intended Time and Actual Time

Measure Objective: Determine whether OBEs under PDM generated periodic probe data within a reasonable standard deviation of the anticipated time.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed-generation frequency (e.g., 4 seconds), verify in OBE logs that messages were generated within a reasonable standard deviation of the anticipated time.

Results: Figure 75 shows the standard deviation between anticipated snapshot generation times and actual snapshot generation times.

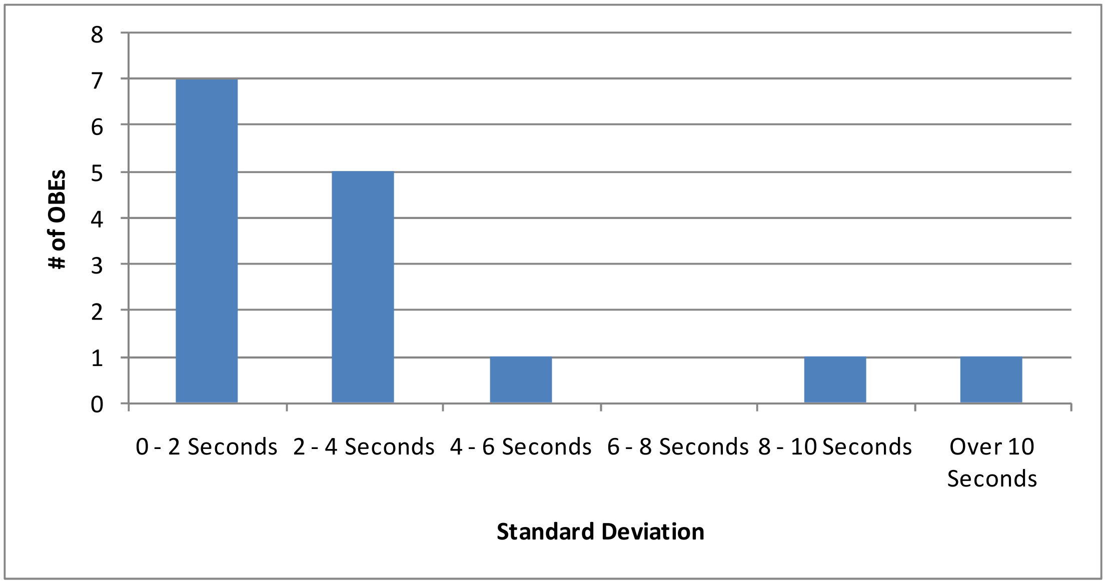

Figure 75: Standard Deviation from Anticipated Time – PDM August 28

A majority of OBEs generated periodic snapshots within a standard deviation of 4 seconds from their intended time.

OBEs C992, C548, and 482A skewed the data because they reported erroneous numbers. The data was considered erroneous because all three OBEs reported instances of PDM message expirations that were longer than the 30 seconds specified in the message. In the case of OBE C992, one message did not expire for over 10 minutes.

Measure 3: Percentage of Time-Based Expirations – Percentage of all PDM Activations for a Given Serial Number that Expired Based on Time

<u>Measure Objective:</u> Determine the percentage of PDM activations that expired based on time.

<u>Concept:</u> Every PDM message has either an expiration time or expiration distance. This allows the operator to control the region that is affected by the PDM message.

<u>Data Source:</u> OBE OSGi logs.

<u>High-Level Data Analysis Plan:</u> In the OBE OSGi logs, divide the total number of time-based expirations of a particular serial number by the total times that serial number became active.

<u>Results:</u> Time-based expiration is defined as the OBE enabling a PDM message until its intended expiration time.

Time-based expirations are affected by overlapping RSEs and close-by RSEs. If the OBE receives another unique message, the newly received message becomes active and the first PDM message expires.

The following screenshot from Google Maps (Figure 76) shows the routes and RSEs (in red) used for the PDM message in question as well as close-by RSEs (in blue) that were also broadcasting PDM messages.

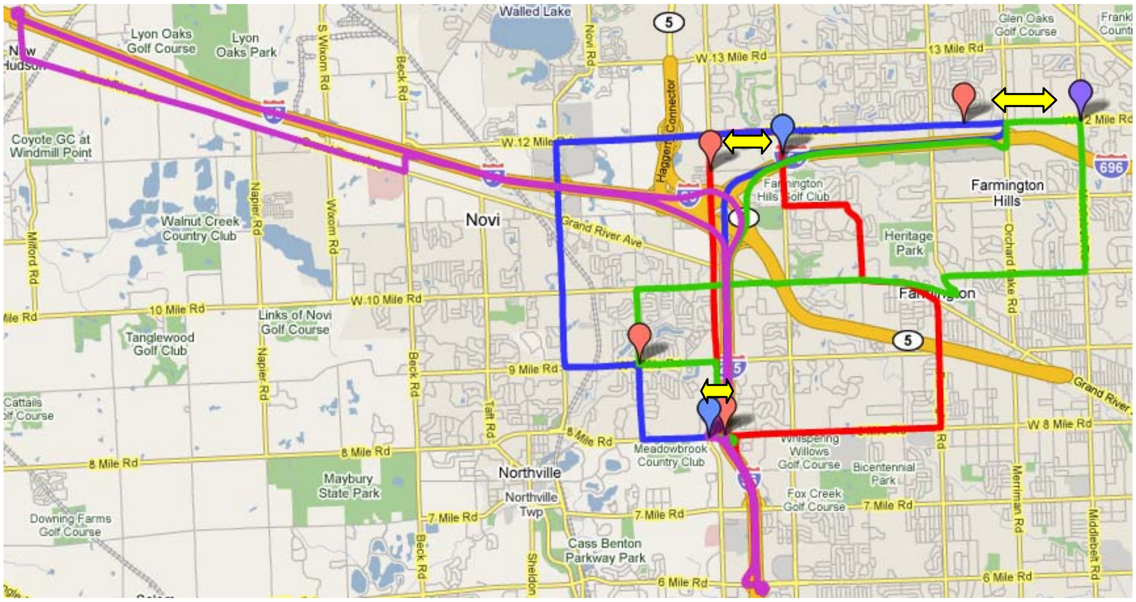

Booz Allen generated routes and plot points on Google map

Figure 76: Nearby RSE Affecting Natural Expiration – PDM August 28

The three yellow arrows in decreasing length are 1.6 miles, 1.06 miles, and 0.23 miles. In at least the 0.23-mile case, it is very likely that the RSE footprints were overlapping and therefore decreasing the percentage of time-based expirations.

Figure 77 shows the time-based expiration of PDM messages by OBE.

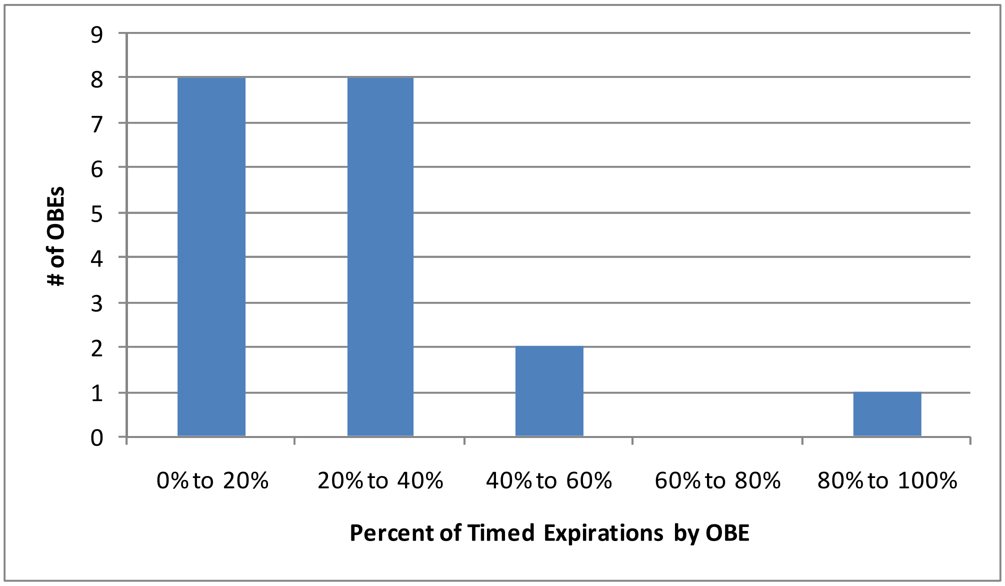

Figure 77: Percentage of Timed Expiration by OBE – PDM August 28

Most OBEs had a time-based expiration rate of less than 60 percent; however, OBE 482 (Group ID P1), which drove route 1, achieved a time-based expiration rate of almost 90 percent. This vehicle encountered a GPS comm. error and did not make the first run. In addition, this vehicle only generated 4 snapshots while under this PDM directive for the entire day, and therefore its numbers are unreliable.

Analysis: On average, only less than 40 percent of PDM messages expired based on time, and this can be attributed to overlapping RSEs that were also broadcasting PDM messages.

Measure 4: Percentage of Time-Based PDM Expirations that Expired Correctly

Measure Objective: For PDM messages that expired based on time, determine whether their expiration was at exactly 30 seconds.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in OBE logs that each time-based expiration met the criteria defined in the PDM message.

Results: The results showed that if a message expired based on time, it did so correctly 100 percent of the time.

Expiration time allows the TMC operator to capture more or less probe data for a certain period. The results show that the time specified in the PDM message will be executed correctly.

Figure 78 shows the percentage of messages that expired by time, not by overlapping RSEs.

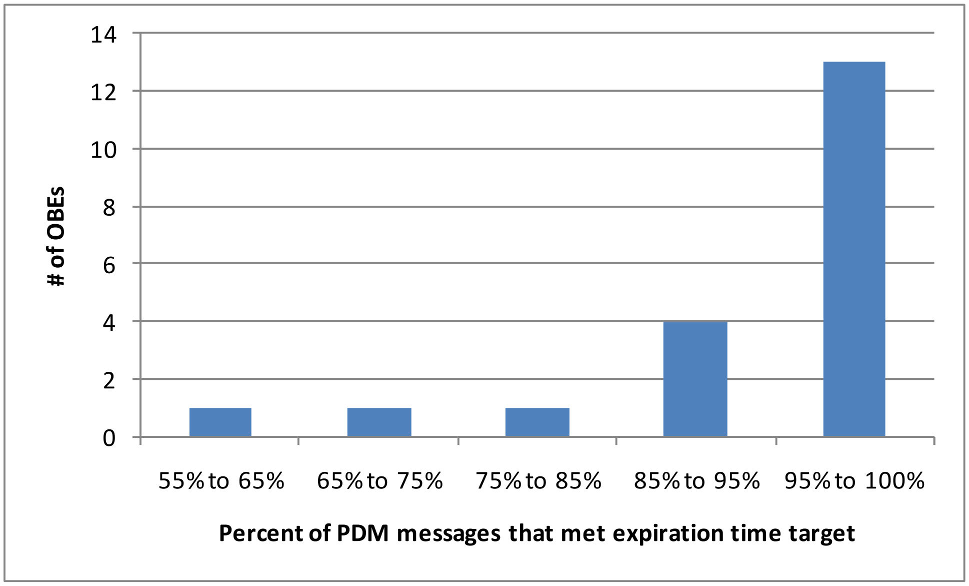

Figure 78: Percentage of PDM Messages That Met Expiration Time Target – PDM August 28

PDM messages on OBE C548 (Group ID P8) met their intended expiration time only 60 percent of the time. This OBE stopped working after the second run and only had five time-based expirations. The numbers from this OBE seem unreliable, but cannot be excluded without further detailed analysis of the OBE logs.

Analysis: 100 percent of messages that were not overridden by new PDM messages expired at the correct time.

Measure 5: Percentage of Correctly Sampled PDM Messages (i.e., percentage of PDM that had a PSN within the allowed range when they became active)

Measure Objective: Determine whether only the vehicles with the appropriate PSN at the time of PDM reception responded to the message.

Concept: Sampling is used to limit the number of vehicles that respond to a particular PDM request and thereby controls the total number of snapshots that the TMC operator receives. Sampling makes senses when a majority of vehicles on the roadway are VII-enabled, because if every vehicle responded to a PDM request, the ensuing onslaught of snapshots from thousands of vehicles could overwhelm the infrastructure.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For the sample specified in the PDM message (15-99), verify that only OBEs with PSNs that ended in that range of numbers (15-99) responded to the message.

Results: Data from 20 OBEs for the PDM message in question were analyzed. The results showed that the sample specified in the PDM message was correctly met 100 percent of the time.

Analysis: Sampling is a very useful tool for the TMC operator to limit the number of snapshots passing through the system. When combined with alteration of the probe data generation scheme, it can achieve its goals of managing probe data traversing the system.

Measure 6: Percentage of Stop Snapshots with Current Speed below Stop Threshold

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at speeds below the stop threshold specified in the PDM message.

Concept: Stop threshold is used to indicate the minimum speed that indicates vehicle movement. For example, on highways with a speed limit of 65 mph, a vehicle traveling at less than 10 mph may be considered "stopped." The vehicle's instantaneous speed is used to determine snapshot type. The resulting stop snapshots can be used to detect traffic congestion, accidents, and other unplanned events.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each stop snapshot generated, verify that the snapshot speed was less than the stop threshold speed.

Results: All OBEs met the goal of generating a stop snapshot when the vehicle speed was below the stop threshold speed 99.99 percent of the time.

The only stop snapshot above the stop threshold was generated by OBE C992 (Group ID P7). It is not clear why this OBE generated a faulty snapshot, especially because it wasn't reported to be defective during testing.

OBE 548 was not included in the results because it reported a speed of zero for most of its snapshots, even though it was moving, evidenced by the varying lat/longs.

OBE C453 was not included in the results because it reported a constant speed of 1.77 mph for all snapshots.

Figure 79 shows all stop snapshots and whether or not they met the stop threshold.

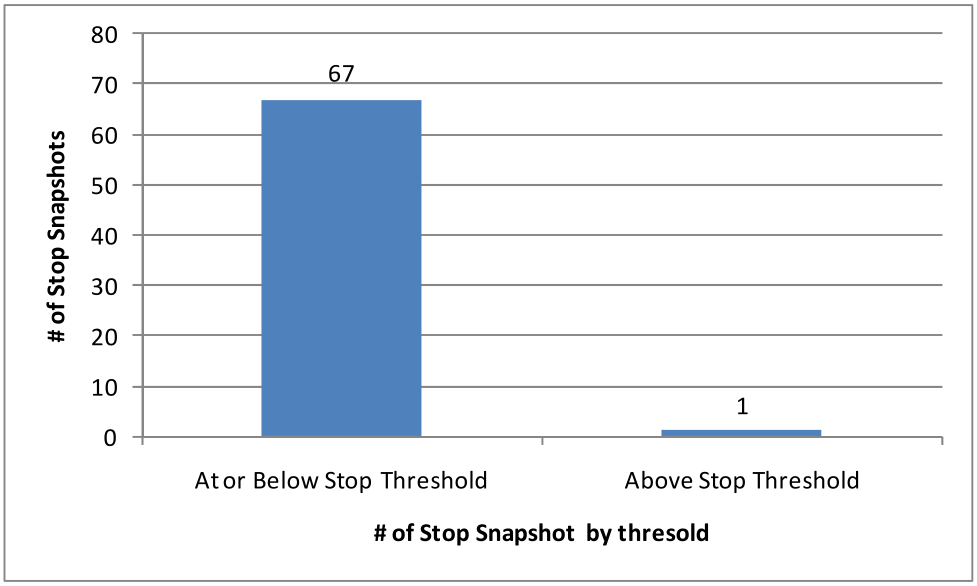

Figure 79: Number of Stop Snapshots by Threshold – PDM August 28

Most stop snapshot speeds were at 0 mph. Since stop time was 10 seconds, it can be concluded that most snapshots were generated at red lights, traffic delays, and/or highway exit ramps.

Figure 80 is a Google Earth representation of some of the stop snapshots that were generated. As the vehicles approach the traffic light at 12 Mile and Kendallwood Drive or move away from it, the speeds are at or slightly above zero.

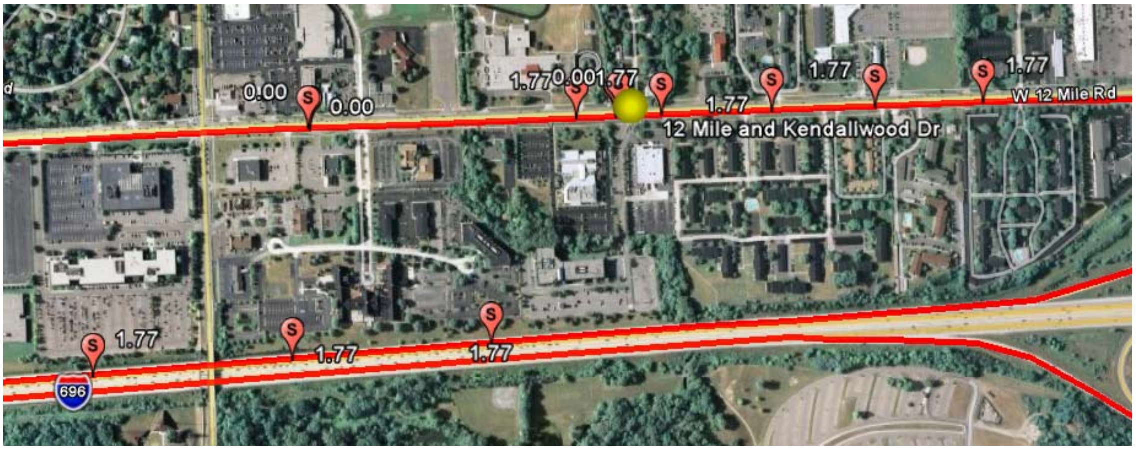

Booz Allen generated routes and plot points on Google map

Figure 80: Stop Snapshots on Google Earth (1) – PDM August 28

The majority of the rest of the stop snapshots were generated by vehicles starting out the run around I-275 and 8 Mile, an area prone to traffic delays.

Booz Allen generated routes and plot points on Google map

Figure 81: Stop Snapshots on Google Earth (2) – PDM August 28

Analysis: The results show that stop thresholds were implemented correctly and OBEs responded with stop snapshots. However, stop snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear that such units are commercially available at a reasonable cost.

Measure 7: Percentage of Stop Snapshots Generated after Stop Lag

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at the frequency (stop lag) specified in the PDM message.

Concept: Stop lag is used to prevent the OBE from generating too many stop snapshots when the vehicle is creeping along at speeds below the stop threshold. Stop lag prevents the OBE buffer from filling up with nearly identical stop snapshots and overwriting periodic and/or event-generated snapshots.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate the difference between the current and previous stop snapshots and verify that it is after the specified lag time.

Results: 100 percent of stop snapshots met the 10-second stop lag time. Figure 82 shows how many seconds after the stop lag time the next stop snapshot was generated.

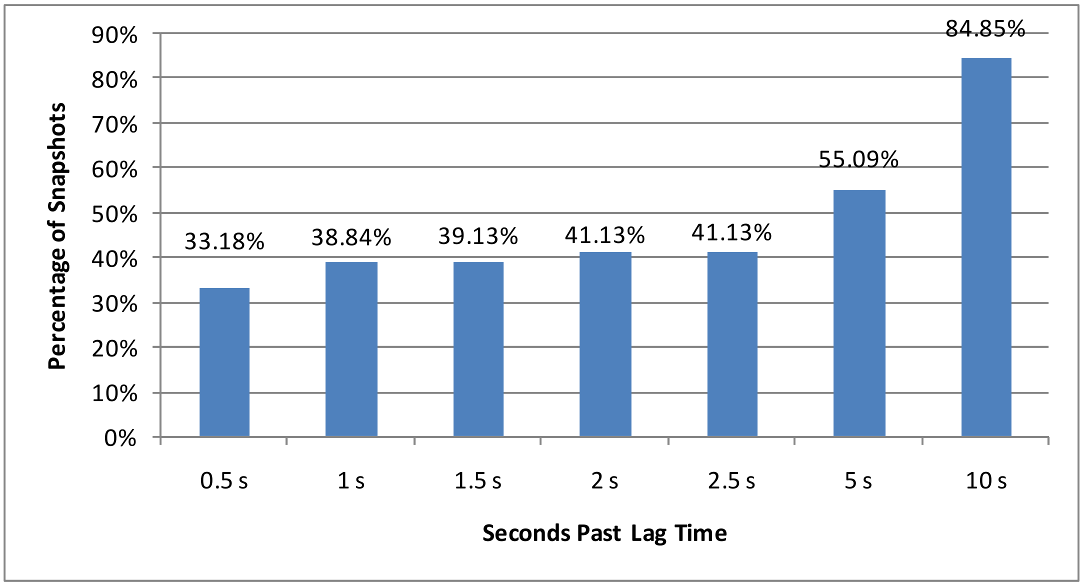

Figure 82: Percentage of Snapshots Past Lag Time – PDM August 28

A majority of snapshots were generated within 5 seconds past the lag time.

Analysis: Stop lag is useful in preventing the fill up of OBE memory with almost identical snapshots; however, the default setting of 30-snapshot buffer size is still very limiting. Increasing the default OBE buffer size setting to 300 snapshots is recommended.

Measure 8: Percentage of Start Snapshots with Current Speed above Start Threshold

Measure Objective: Determine whether OBEs under PDM generated start event probe data at speeds above the start threshold specified in the PDM message.

Concept: Start threshold is used to indicate the minimum speed at which the vehicle is no longer considered “stopped.” The resulting start snapshots can be used to detect the end of traffic congestion and clearing of accidents among other things.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each start snapshot generated, verify that the snapshot speed was more than the start threshold speed.

Results: The OBEs generated start snapshots with speeds higher than the start threshold speed 100 percent of the time.

The majority of start snapshots had speeds more than the start threshold. However, only 23 start snapshots from 7 OBEs were received for the entire day. The snapshots with speeds less than the start threshold were generated by faulty OBEs.

Of the non-faulty OBEs that reported any start snapshots, all snapshots speeds were above the start threshold speed.

OBEs B450 and B325 exhibited unexpected behavior and were not included in the analysis. For example, OBE B450 generated multiple snapshots within the same second and OBE B325 generated multiple start snapshots in a row and generated start snapshots following periodic snapshots.

Analysis: The results show that start thresholds speeds were implemented correctly and OBEs responded with start snapshots. However, start snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear that such units are commercially available at a reasonable cost.

Analysis for Friday, August 29, 2008

PDM Message Selected

The PDM message analyzed for August 29 had the following parameters shown in Table 56.

Table 56: Contents of PDM Message on August 29

Sample Start/End		Termination Method	
50	99	Distance	500 m
Collection Method			
Speed 1	Speed 2	Distance 1	Distance 2
5 mph	20 mph	40 m	120 m
Start Stop Strategy			
Stop Threshold	Stop Time	Stop Lag	Start Threshold
5 mph	1 s	1 s	5 mph
Message ID: 1EDAC8FB-0F01-4EBB-B324-4D5C4FAAA145			

This message was selected for detailed analysis because it had the highest number of occurrences in the OBE logs. It appeared a total of 36,391 times across all OBEs. The message that had the second highest occurrence in the OBE logs came in at 26,514 times.

Methods

On Friday, August 29, PDM messages with distance-based collection methods were sent to the vehicles. The vehicles were scheduled to go out for four runs; however, due to a PDS server issue, run 1 was cancelled. The vehicles were divided into groups of five and sent on five routes. On runs 2 and 3, the vehicles were spaced 3 minutes apart, and there was no spacing during run 4.

Table 57: Test Start and End Times – PDM August 29

Run	Start	End	Total Loops
Run 1	N/A	N/A	Cancelled due to PDS server issues
Run 2	8:20 am	9:26 am	2
Run 3	9:36 am	11:15 am	2
Run 4	12:00 pm	2:00 pm	3, except vehicles on route 1 did only 2.75 loops

Note about OBE Logs

For Friday, August 29, 13 vehicles sent uncompromised data for analysis.

OBE B856 was disqualified from inclusion in the analysis due to apparent positioning system errors. The OBE reported a constant speed of 40.39 mph for all snapshots generated.

Table 58 shows the issues experienced with other vehicles during testing.

Table 58: OBE issues on August 29

P2 (C862)	Probe generation error after session 2; resolved via reboot
C3 (B422)	No vehicle speed (CAN) after session 2; resolved via reboot
A2 (C594)	Probe application did not come back after vehicle shut down for lunch; resolved via reboot
E1 (B450)	Probe application did not come back after vehicle shut down for lunch; resolved via reboot
P10 (C590)	"Invalid location object validation" error after session 2; resolved via reboot
P12 (C366)	Heading, speed, throttle, POS, time confidence, and elevation, not equipped, HMI (screen) not responsive after initial boot up @ 4:30 am; resolved via reboot

Measure 1: Percentage of Periodic Snapshots within 20 Meters of Intended Distance

Measure Objective: Determine whether OBEs under PDM generated periodic probe data at the frequency specified in the PDM message.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This in turn allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed generation frequency (e.g., 30 meters), verify in OBE logs that messages were generated at that frequency.

Results: For the message in question, the OBEs were to generate a periodic snapshot every 40 meters if the speed was less than 5 mph. They were also to generate a periodic snapshot every 120 meters if the speed was more than 20 mph. And for speeds between 5 and 20 mph, the snapshot distance was supposed to be linearly interpolated.

Figure 83 shows how many meters after a snapshot's intended generation distance it was actually generated.

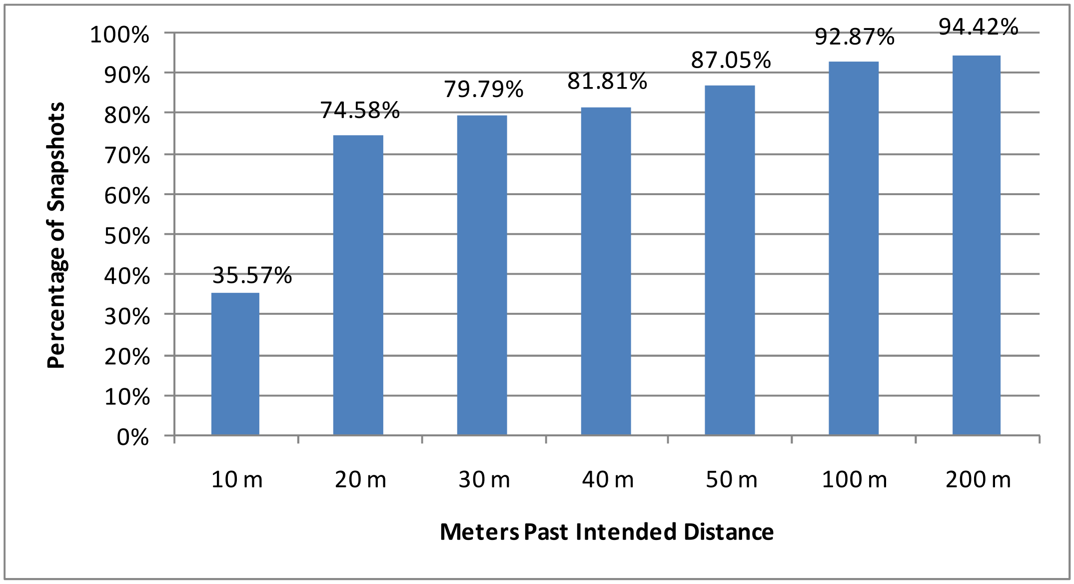

Figure 83: Percentage of Snapshot within 20 m – PDM August 29

No snapshots were generated before the intended distance, over 74 percent of snapshots were generated within 20 meters of the intended distance, and almost all snapshots were within 100 meters.

Analysis: The results show that over 74 percent of snapshots are generated within 20 meters of their intended distance. If an operator requests snapshots to be generated every 100 meters, they can expect the majority of them to be generated between 100 and 120 meters. Having the ability to alter period snapshot generation will indeed be useful. However, it would be most useful if the operators were not limited by OBE shortcomings. It is recommended that, in the future, OBEs can generate snapshots faster and by default store more snapshots in memory.

Measure 2: Standard Deviation of Differences between Anticipated Distance and Actual Distance

Measure Objective: Determine whether OBEs under PDM generated periodic probe data within a reasonable standard deviation of the anticipated distance.

Concept: The ability to alter probe data generation rates is designed to give the traffic management operators control over the amount of probe data that they receive from a specific location. This in turn allows them to manage the balance between network efficiency and network effectiveness.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For RSEs broadcasting PDM messages with fixed generation frequency (e.g., 30 meters), verify in OBE logs that messages were generated within a reasonable standard deviation of the anticipated distance.

Results: Figure 84 shows the standard deviation between anticipated snapshot generation distances and actual snapshot generation distances.

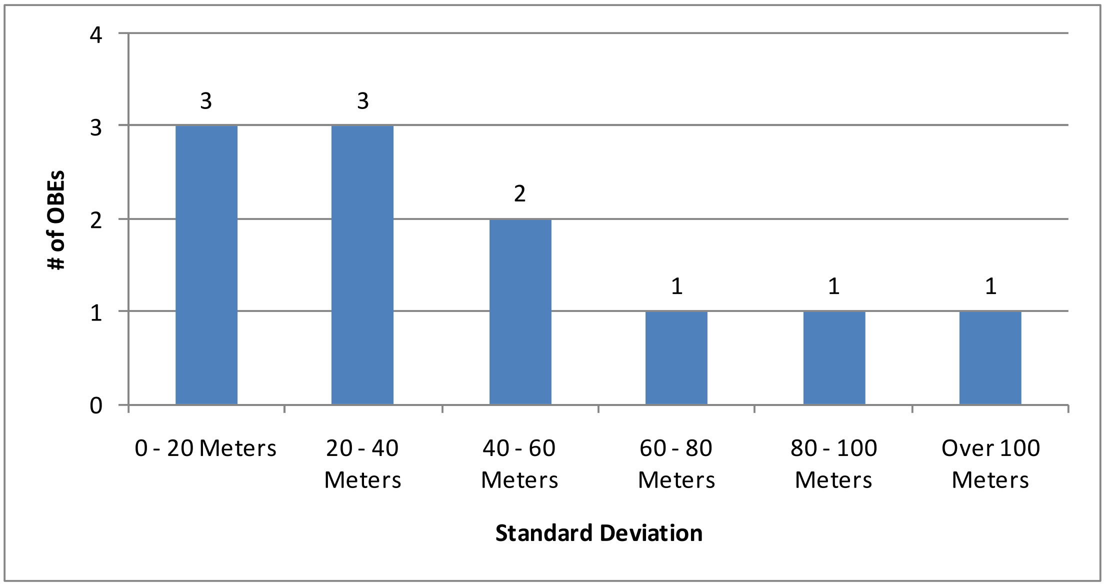

Figure 84: Standard Deviation from Intended Distance – PDM August 29

A majority of OBEs generated periodic snapshots within a standard deviation of 40 meters from their anticipated distance.

OBEs C985, B450, and C082 exhibited unexpected behavior and were not included in the analysis. For example, OBEs C985 and B450 generated snapshots that would only be possible if the vehicles were traveling at nearly 300 mph and 212 mph, respectively, and OBE C082 at one point reported traveling only 501 meters; whereas its positioning service reported that it had traveled over 8 miles.

The issues above lead to doubts about the accuracy of the remaining snapshots from those three OBEs.

Measure 3: Percentage of Distance-Based Expirations – Percentage of All PDM Activations for a Given Serial Number That Expired Based on Distance

Measure Objective: Determine the percentage of PDM activations that expired based on distance.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: In OBE OSGi logs, divide the total number of distance-based expirations of a particular serial number by the total times that serial number became active.

Results: Distance-based expiration is defined as the OBE responding to a PDM message until its intended expiration distance.

Distance-based expirations are affected by overlapping RSEs and close-by RSEs. If the OBE receives another unique message, then the newly received message becomes active and the first PDM message expires.

The following screenshot from Google Maps (Figure 85) shows the routes and RSEs (in red) used for the PDM message in question, as well as close-by RSEs (in blue) that were also broadcasting PDM messages.

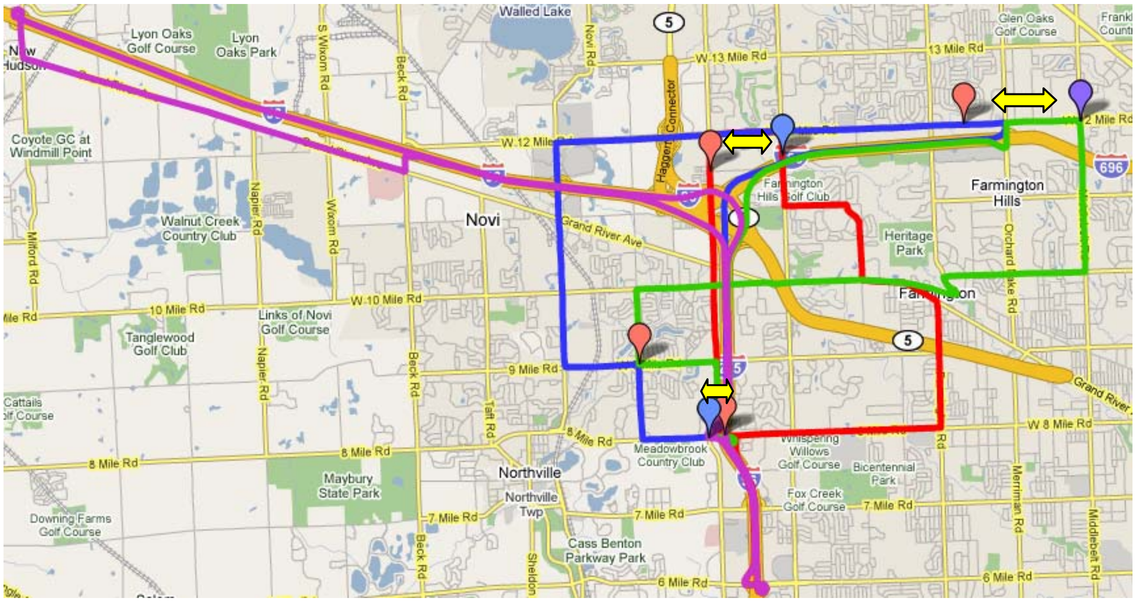

Booz Allen generated routes and plot points on Google map

Figure 85: Nearby RSE Affecting Natural Expiration – PDM August 29

The three yellow arrows in decreasing length are 1.6 miles, 1.06 miles, and 0.23 miles. In at least the 0.23-mile case, it is very likely that the RSE footprints were overlapping and therefore decreasing the percentage of distance-based expirations.

Figure 86 shows the distance-based expiration of PDM messages by OBE.

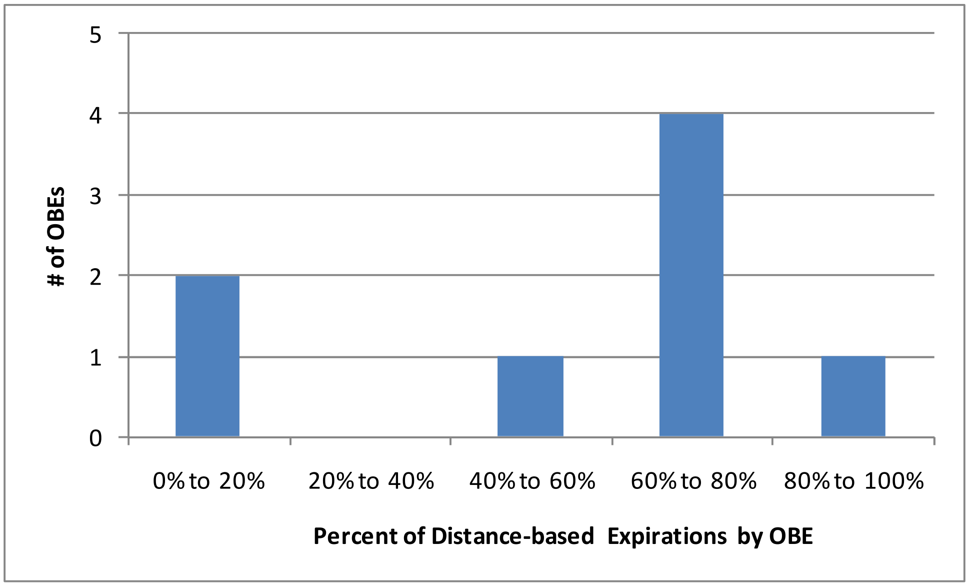

Figure 86: Percentage of Distance-based Expirations by OBE

Analysis: A majority of PDM messages expired based on distance. The reason message expiration by distance is not 100 percent could be attributed to overlapping RSEs that were also broadcasting PDM messages.

Measure 4: Percentage of Distance-Based PDM Expirations That Expired Correctly

Measure Objective: For PDM messages that expired by distance, determine whether their expiration was at exactly 500 meters.

Concept: Every PDM message has either an expiration time or distance. This allows the operator to control the region that is affected by the PDM message.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Verify in OBE logs that each distance-based expiration met the criteria defined in the PDM message.

Results: The results showed that if a message expired based on distance, it did so correctly 100 percent of the time. The results also showed that all PDM messages expired at their anticipated expiration distance.

Analysis: The PDM message expiration times were carried out successfully 100 percent of the time.

Measure 5: Percentage of Correctly Sampled PDM Messages (i.e., percentage of PDM that had a PSN within the allowed range when they became active)

Measure Objective: Determine whether only the vehicles with the correct PSN at the time of PDM reception responded to the message.

Concept: Sampling is used to limit the number of vehicles that respond to a particular PDM request and thereby controls the total number of snapshots that the TMC operator receives. Sampling makes senses when a majority of vehicles on the roadway are VII-enabled, because if every vehicle responded to a PDM request, the ensuing onslaught of snapshots from thousands of vehicles could overwhelm the infrastructure.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For the sample specified in the PDM message (50-99), verify that only OBEs with PSNs that ended in that range of numbers (50-99) responded to the message.

Results: Data from 12 OBEs for the PDM message in question were analyzed. The results showed that the sample specified in the PDM message was met 100 percent of the time.

Analysis: Sampling is a very useful tool in limiting the number of snapshots going through the system. When combined with alteration of the probe data generation scheme, it can achieve its goals of managing probe data traversing the system.

Measure 6: Percentage of Stop Snapshots with Current Speed below Stop Threshold

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at speeds below the stop threshold specified in the PDM message.

Concept: Stop threshold is used to indicate the minimum speed that indicates vehicle movement. For example, on highways with a speed limit of 65 mph, a vehicle traveling less than 10 mph may be considered "stopped." The vehicle's instantaneous speed is used to determine snapshot type. The resulting stop snapshots can be used to detect traffic congestion, accidents, and other unplanned events.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each stop snapshot generated, verify that the snapshot speed was less than the stop threshold speed.

Results: All OBEs met the goal of generating a stop snapshot when the vehicle speed was below the stop threshold speed 100 percent of the time.

Stop lag and stop time were also decreased to 1 second each (See Table 21). This change caused stop snapshots to be generated every second after the initial snapshot, which resulted in 2,054 stop snapshots on Friday versus only 68 stop snapshots on Thursday where the stop lag was 10 seconds.

OBEs B042, C779, and B193 exhibited behavior not in line with J2735. At one time, all three generated more than one stop snapshot within one second and at different lat/longs. Because of this apparent error, the remaining snapshots generated by these OBEs are deemed unreliable and are not included in the analysis.

The following Google Earth screenshot (Figure 87) depicts all stop snapshots for August 29.

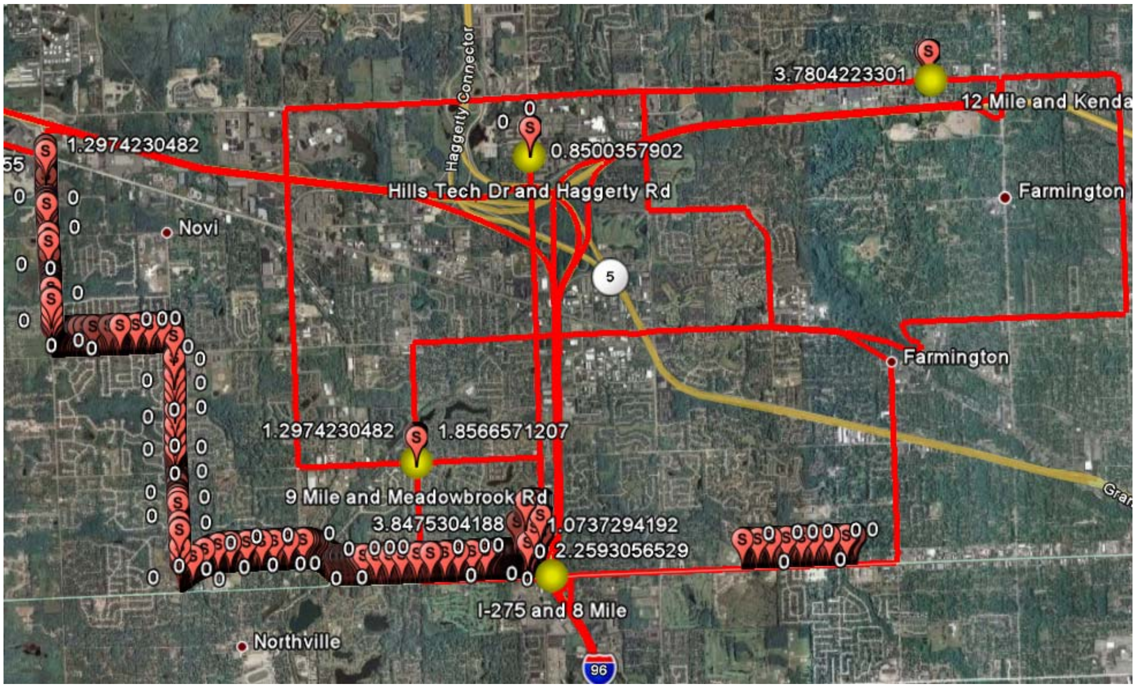

Booz Allen generated routes and plot points on Google map

Figure 87: All Stop Snapshots on Google Earth – PDM August 29

It is noteworthy that a trip path going from Grand River Ave (top left of image) back to the staging area (center of image by yellow RSE) using arterials can be observed.

Of the 2,054 stop snapshots generated, almost 47 percent were by OBE B450 and a little more than 49 percent were by OBE C082.

OBE C082 reported a speed of zero between 1:39 pm and 1:56 pm. However, the information in its snapshots indicate that it traveled over 8 miles. Even though the positioning service malfunctioned during this time, the snapshot generation service did not. Given that the stop time and stop lag were both lowered to 1 second and the test time lasted for 17 minutes, one would expect around 1,020 total snapshots (1,020 seconds in 17 minutes). The vehicle generated 1,013 snapshots, exactly as expected.

The Google Earth screenshot (Figure 88) depicts only the stop snapshots received from OBE C082.

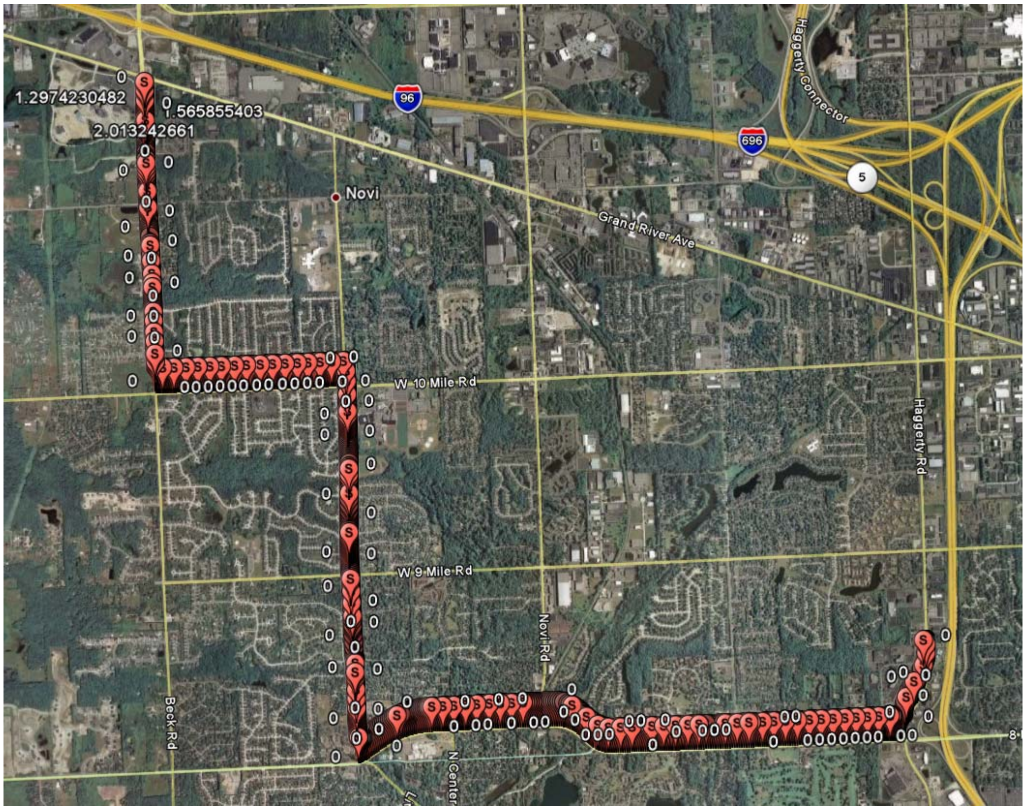

Booz Allen generated routes and plot points on Google map

Figure 88: OBE C082 Stop Snapshots – PDM 8/29

A closer look at OBE B450 reveals that most of its 959 stop snapshots were generated between 5:17 pm and 5:52 pm, in the middle of evening rush hour commute. This OBE performed as expected, and the number of stop snapshots received is on par with expectations, given the 1-second stop time and stop lag.

The Google Earth screenshot (Figure 89) shows all 959 stop snapshots. Most were generated while the vehicle was parked at the staging area or waiting to get onto Haggerty Rd from Orchard Hill Pl.

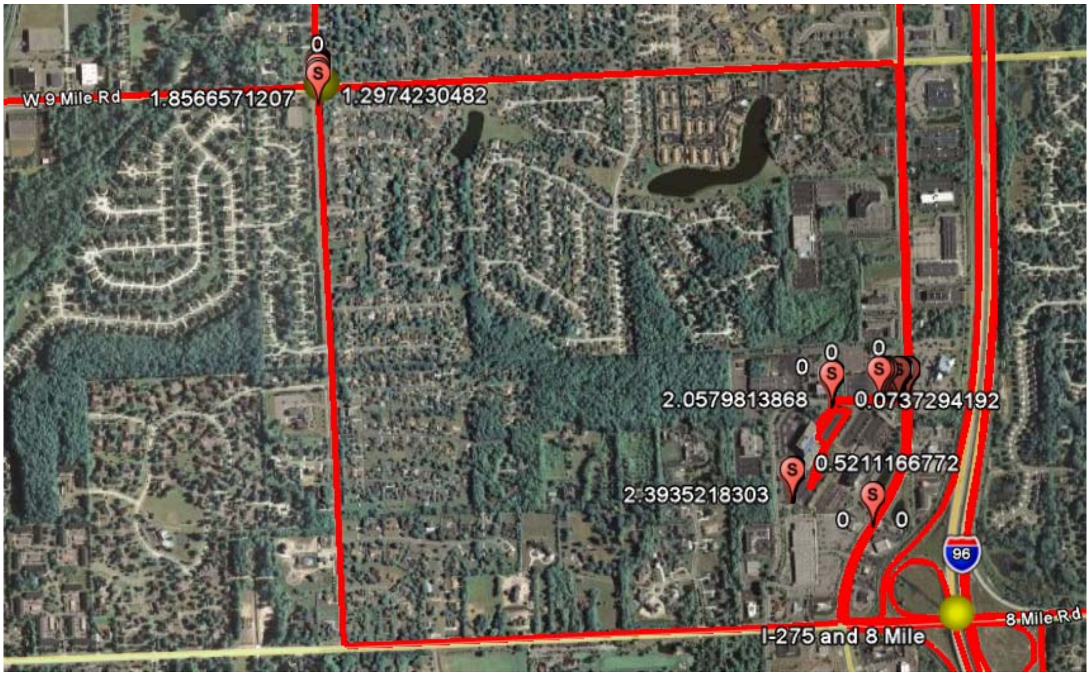

Booz Allen generated routes and plot points on Google map

Figure 89: OBE B450 Stop Snapshots – PDM 8/29

The screenshot above does not represent the total number of snapshots generated by B450 adequately. Clicking on "one" snapshot reveals how many snapshots were actually generated at that particular location (Figure 90).

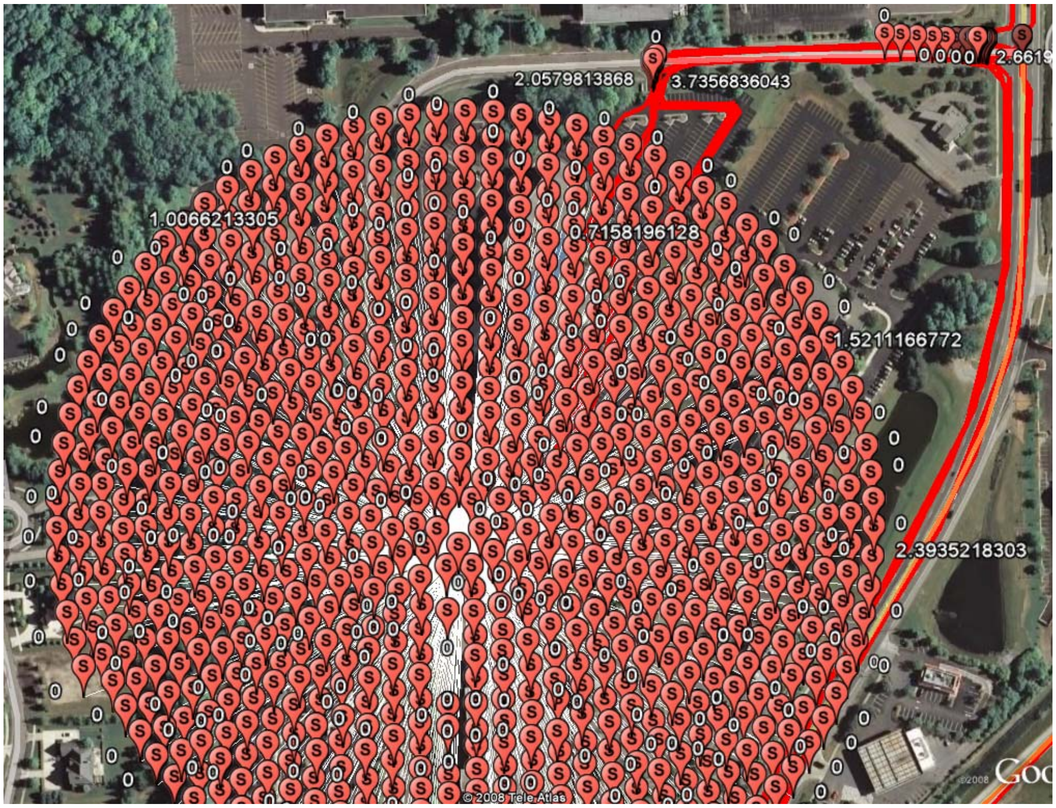

Booz Allen generated routes and plot points on Google map

Figure 90: OBE B450 Stop Snapshots – Detailed – PDM 8/29

Analysis: The results show that stop thresholds were implemented correctly and OBEs responded with stop snapshots. However, stop snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear whether such units are commercially available at a reasonable cost.

Measure 7: Percentage of Stop Snapshots Generated after Stop Lag

Measure Objective: Determine whether OBEs under PDM generated stop event probe data at the frequency (stop lag) specified in the PDM message.

Concept: Stop lag is used to prevent the OBE from generating too many stop snapshots when the vehicle is creeping along at speeds below the stop threshold. Stop lag prevents the OBE buffer from filling up with nearly identical stop snapshots and overwriting periodic and/or event-generated snapshots.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Calculate the difference between the current and previous stop snapshots and verify that it is after the specified lag time.

Results: 100 percent of stop snapshots met the 1-second stop lag time. Figure 91 shows how many seconds after the stop lag time the next stop snapshot was generated.

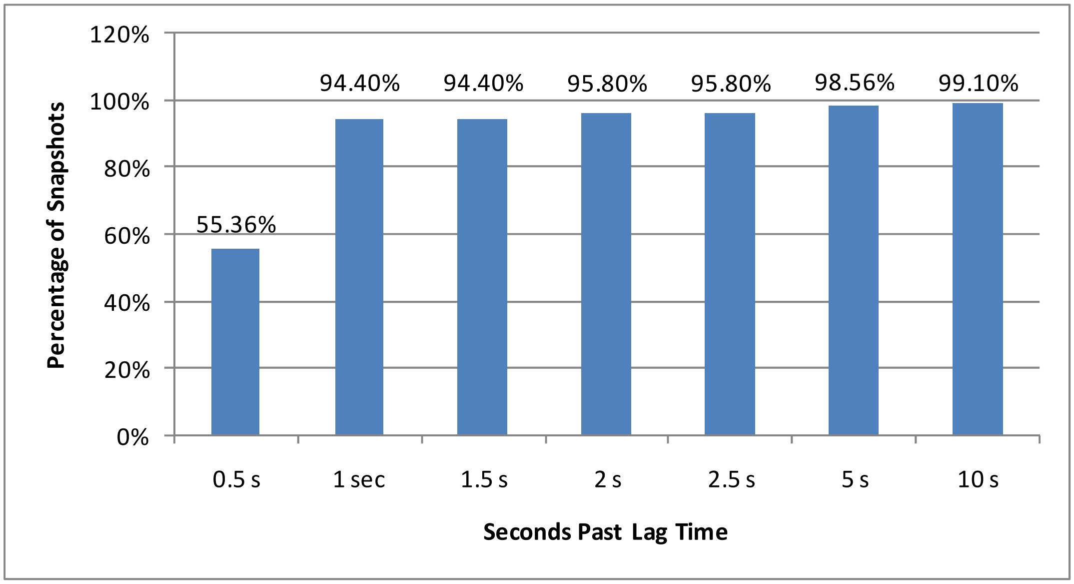

Figure 91: Percentage of Snapshots past Lag Time – PDM August 29

Over 94 percent of snapshots were generated within 1 second past the lag time.

Analysis: Stop lag is useful in preventing the fill up of OBE memory with almost identical snapshots; however, the default setting of 30-snapshot buffer size is still very limiting. Increasing the default OBE buffer size setting to 300 is recommended.

Measure 8: Percentage of Start Snapshots with Current Speed above Start Threshold

Measure Objective: Determine whether OBEs under PDM generated start event probe data at speeds above the start threshold specified in the PDM message.

Concept: Start threshold is used to indicate the minimum speed at which the vehicle is no longer considered "stopped." The resulting start snapshots can be used to detect the end of traffic congestion and clearing of accidents among other things.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: For each start snapshot generated, verify that the snapshot speed was more than the start threshold speed.

Results: The OBEs met the start threshold goal 50 percent of the time. In contrast, on August 28, a 100-percent success rate was observed. The difference could be due to the fact that both the start and stop thresholds were the same for August 29 and both stop lag and stop time were decreased to 1 second.

Analysis: The results from August 28 showed that start thresholds were implemented correctly and OBEs responded with start snapshots. However, having the same start and stop threshold value as well as a low stop lag and stop time can lead to a steep decline in OBE responsiveness. This may be due to either OBE processing power or how the OBE software was implemented.

Start snapshot concentrations only become useful when they are combined with accurate lane-level data. For future OBEs, it is recommended that the POC GPS receivers be replaced with units that can process differential signals to achieve better accuracy. However, it is not clear whether such units are commercially available at a reasonable cost.

Conclusions

The hypothesis was that PDM messages were received and implemented by the OBEs correctly. To analyze the data, PDM messages were broken down into eight distinct measures and looked at separately. The measures were:

- Frequency of probe data generation
- Standard deviation of generated probe data from intended time or distance
- Percentage of time-based expirations
- Implementation of PDM message expiration by time
- Sampling of OBEs
- Stop snapshot threshold implementation
- Stop snapshot lag time implementation
- Start snapshot threshold implementation.

Based on the results, it is safe to conclude that PDM was properly implemented and OBEs responded correctly.

Off-Board Navigation – Complete

Roundtrip Latency

Measure Objective: Determine the roundtrip latency between the time when the navigation request was sent and maneuvers were received.

Data Source: OBE OSGi logs.

High-Level Data Analysis Plan: Using the OBE logs, find out the roundtrip latency using “Route Request Timestamp” and the “Maneuvers Received Timestamp.”

Results: The average roundtrip latency between the time that the request was sent and the time when maneuvers were received was 91.31 seconds.

Figure 92 shows average latency as number of application and vehicles states are changed.

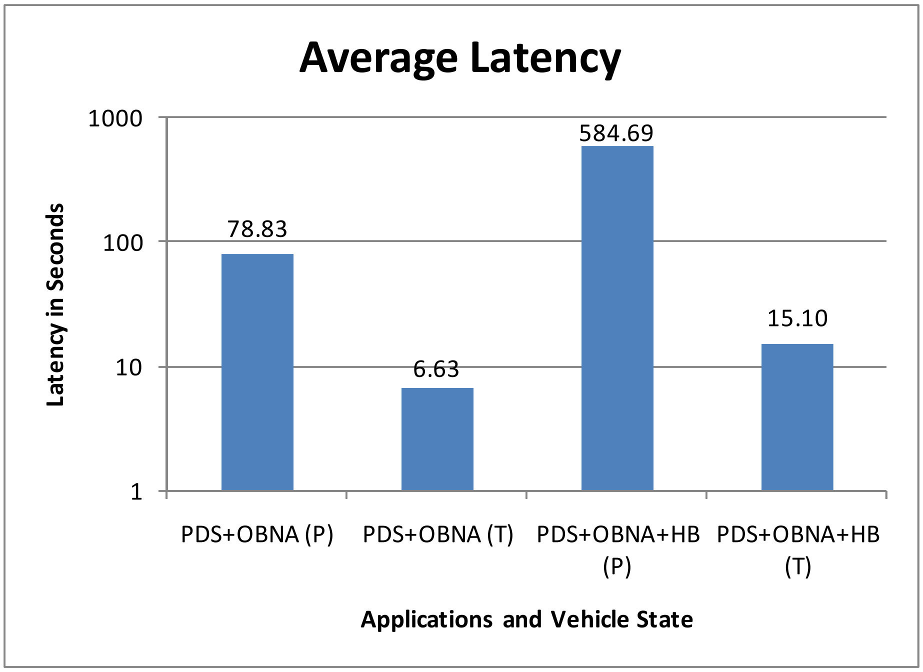

Figure 92: Average OBNA Latency for All Data

Between 13:10 and 14:30 the vehicles ran both probe and OBNA while parked. Data was received from 21 vehicles, and the average latency was 78.83 seconds. Running the same test while vehicles were driving caused a significant drop in latency, probably because the RSE was not overburdened with requests.

HB was then added to the mix of applications. The latency spiked to almost 10 minutes while probe, OBNA, and HB were running simultaneously and the vehicles were parked. This number dropped to a more reasonable 15 seconds once the vehicles started driving.

A closer look at each one of the different formations is needed to better understand the results:

Probe and OBNA in parked formation (21 Vehicles):

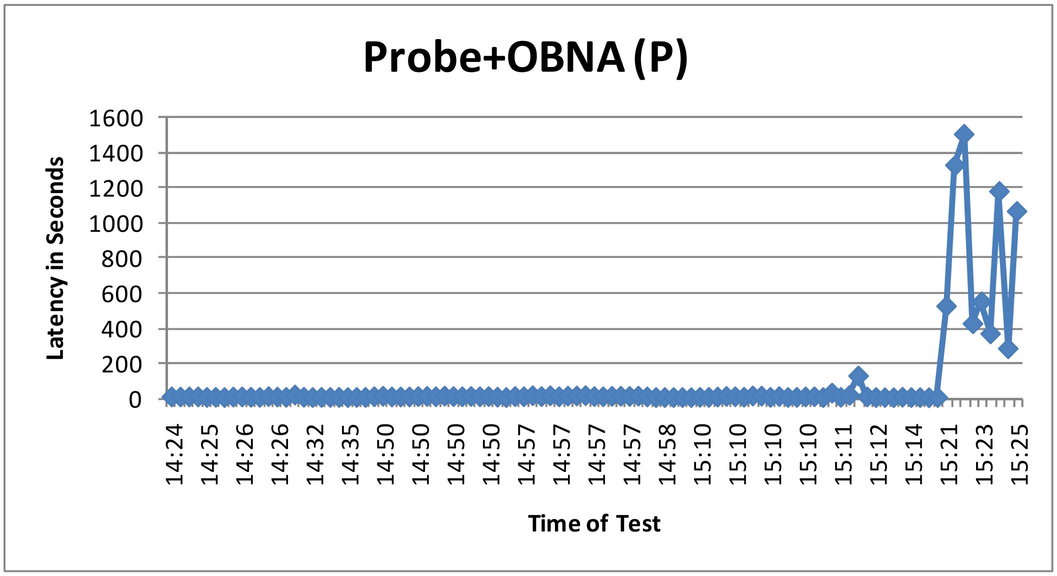

Figure 93: Average Latency PDS+OBNA Parked

The average latency for this formation (Figure 93) was 78.83 seconds. The numbers spiked in the last 4 minutes of testing, which skewed the results considerably. In those 4 minutes, multiple near-simultaneous requests were made from the same vehicles, which could have negatively affected roundtrip latency. The tests relied on drivers pressing the route request button at certain times, and the multiple requests were most likely caused by drivers pressing the button multiple times. Based on this theory, nothing prevented this situation from happening prior to the last 4 minutes.

Figure 94 shows the data excluding the last 4 minutes of testing. The average latency dropped from over 78 seconds to just about 11 seconds.

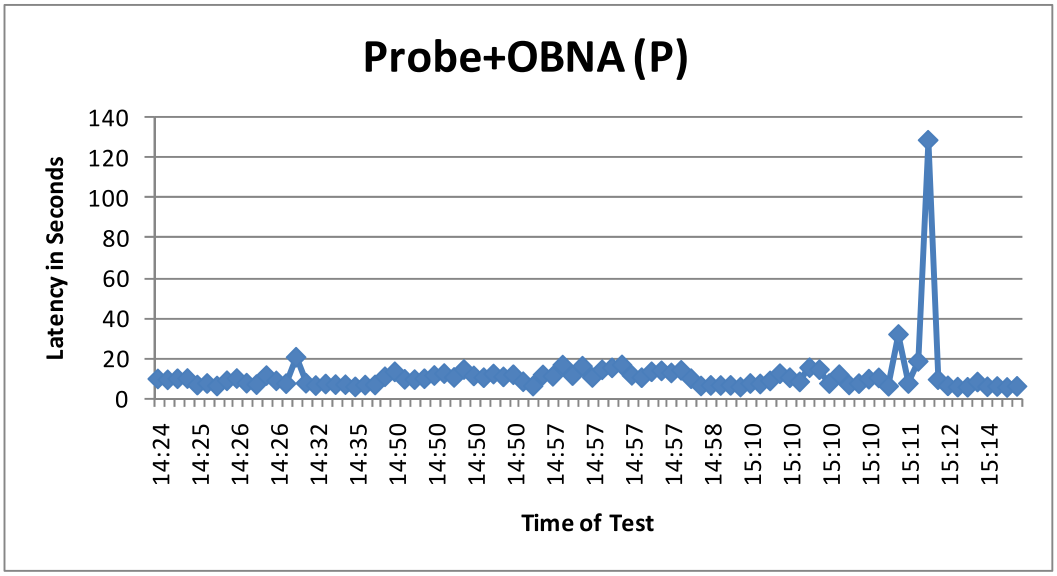

Figure 94: Average Latency PDS+OBNA Parked excluding Faulty Data

Probe and OBNA in traffic formation (16 vehicles):

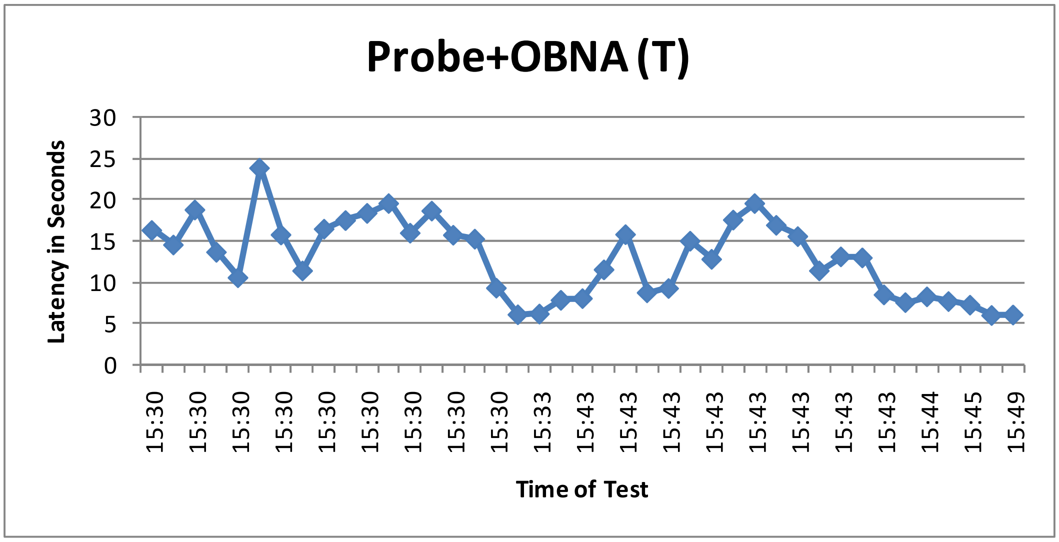

Figure 95: Average Latency PDS+OBNA Traffic

During traffic testing of probe and OBNA, the average latency was about 13 seconds (Figure 95). This could be due to the fact that the RSE on the route was not overloaded with too many simultaneous requests.

Probe, OBNA and HB in parked formation (6 Vehicles):

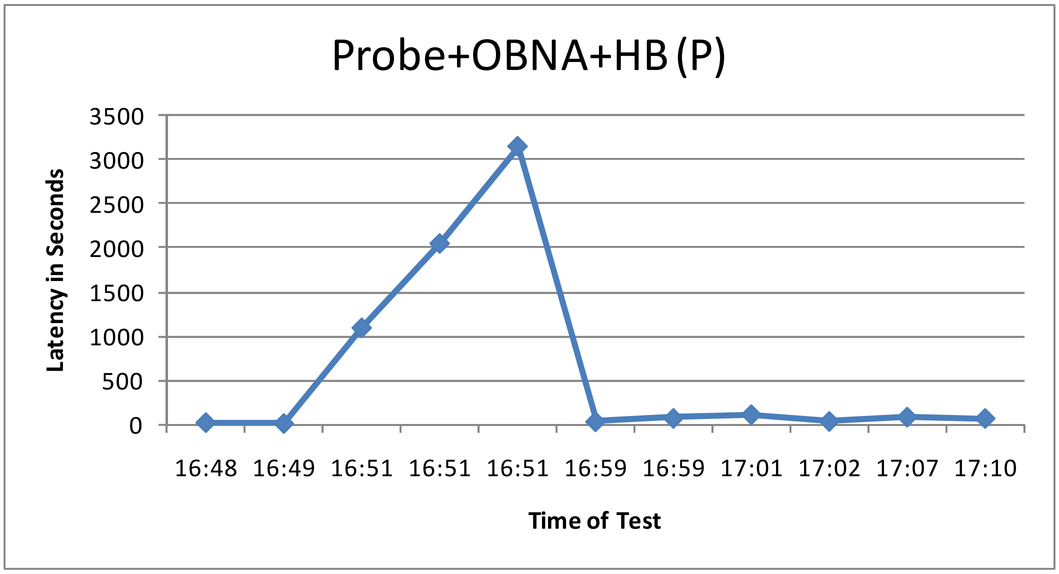

Figure 96: Average Latency PDS+OBNA+HB Parked

The average latency for this formation was 584 seconds (Figure 96). However, the data was skewed by OBE B856, which sent out three route requests in a matter of 6 seconds. It seems that the current OBNA implementation is not well suited for multiple route requests from a single vehicle in a short span of time. Driver error is believed to have caused the unusually long latencies seen above. Assuming that this theory is correct and excluding OBE B856, the average latency goes to 57 seconds, as seen in Figure 97.

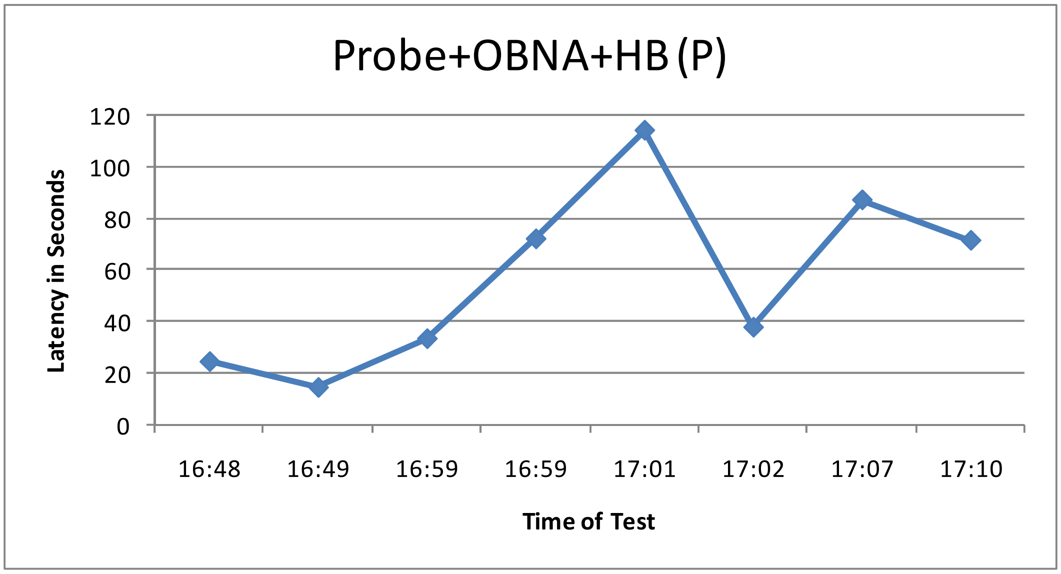

Figure 97: Average Latency PDS+OBNA+HB Parked – excluding Faulty Data

Probe, OBNA and HB in traffic formation (5 Vehicles):

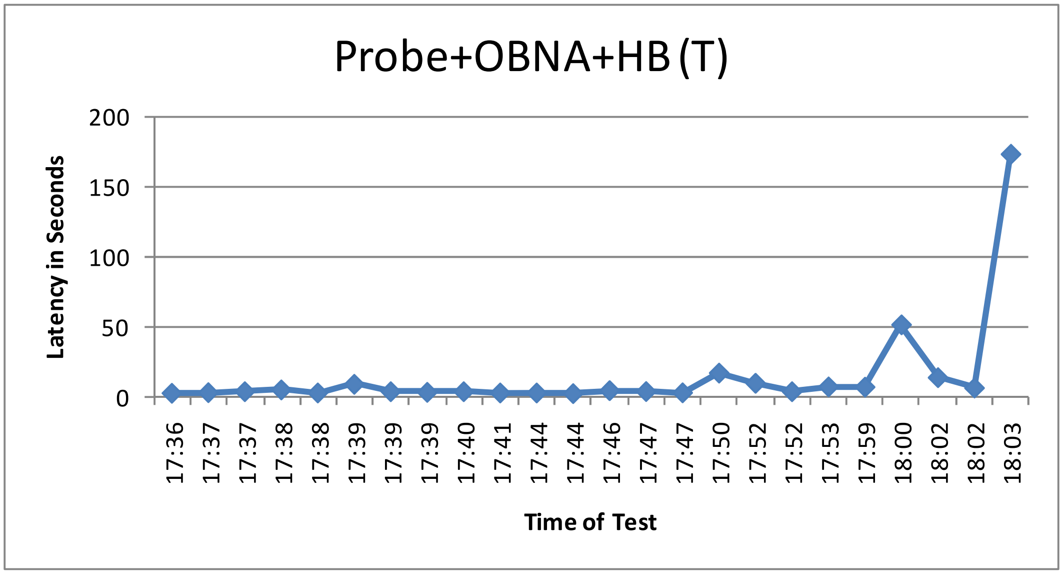

Figure 98: Average Latency PDS+OBNA+HB Traffic

During traffic testing and with HB on, the average latency dropped to around 15 seconds. However, even this low number was adversely affected by OBE C941s unusually high number of requests. In a period of about 26 minutes, there were 24 route requests by 5 vehicles. However, 50 percent of those requests were by OBE C941, and the last two were especially close (22 seconds apart), which possibly led to the spike at the end.

Excluding OBE C941, the average latency falls to 8.4 seconds, as seen in Figure 99.

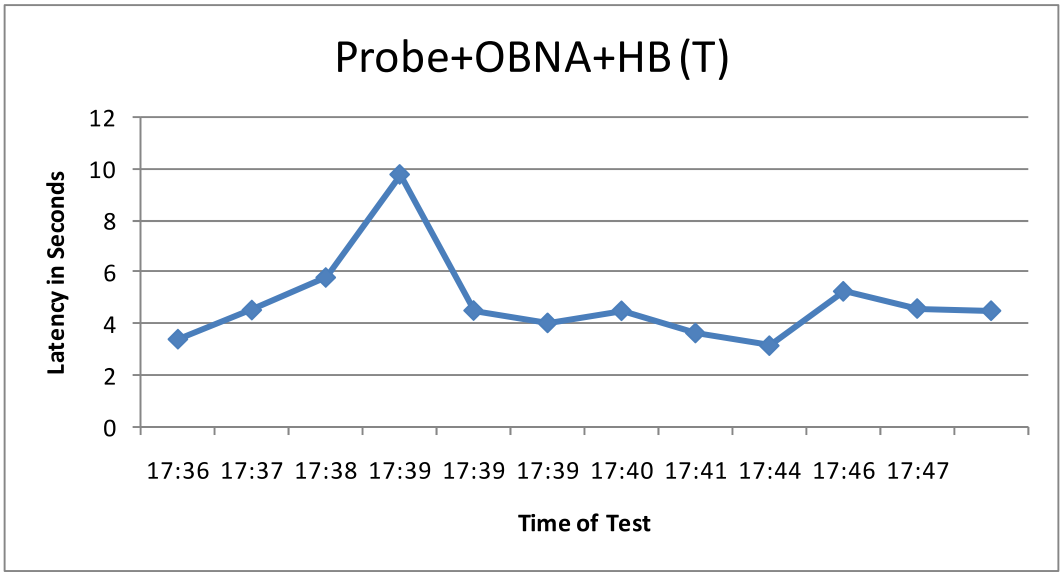

Figure 99: Average Latency PDS+OBNA+HB Traffic – excluding Faulty Data Day 2 Multi-apps

Conclusion

The results show that when HB, probe, and OBNA are running simultaneously on a given vehicle, the OBE to RSE interaction becomes unreliable. The underlying issue may be in the OBE's implementation of the various applications in parallel and/or the RSE's capability to serve 25 vehicles with OBNA data while receiving probe data from that many vehicles. This issue warrants further testing and analysis.

Figure 100 shows how the average roundtrip latency of OBNA requests was affected by the different testing scenarios, excluding faulty data.

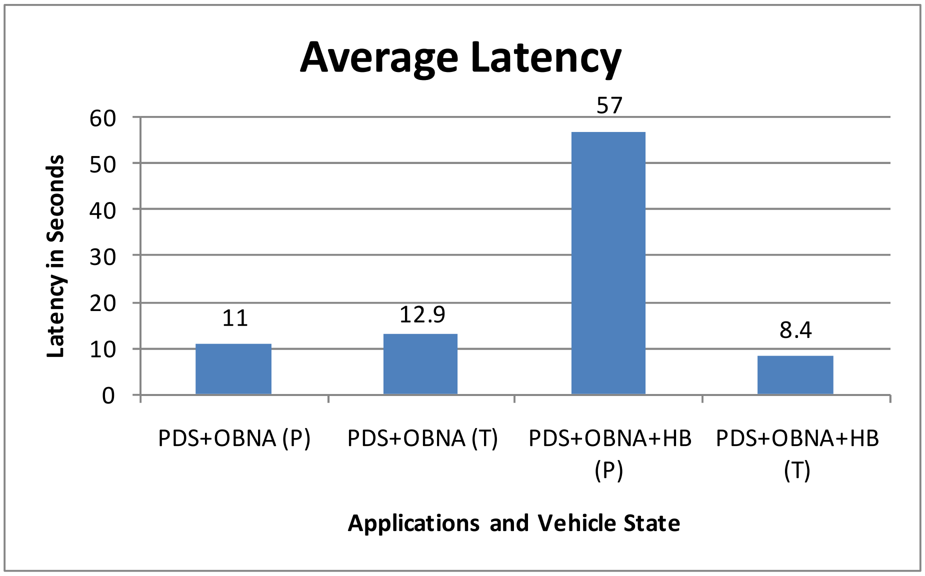

Figure 100: Average Latency all Data – excluding Faulty Data Day 2 Multi-apps

On average, roundtrip latency is around 22 seconds; however, parked HB and OBNA add a significant amount of lag to the results. Given these limitations, it becomes even more important that OBEs are able to download routes from multiple RSEs.

Probe Data Reception (Test Case PDS-02)

Daily Averages – Loss

The average OBE-to-RSE reception rates ranged from nearly 60 percent to nearly 100 percent, with typical values in the 90-to-100-percent range. Figure 101 illustrates values for each day of testing.

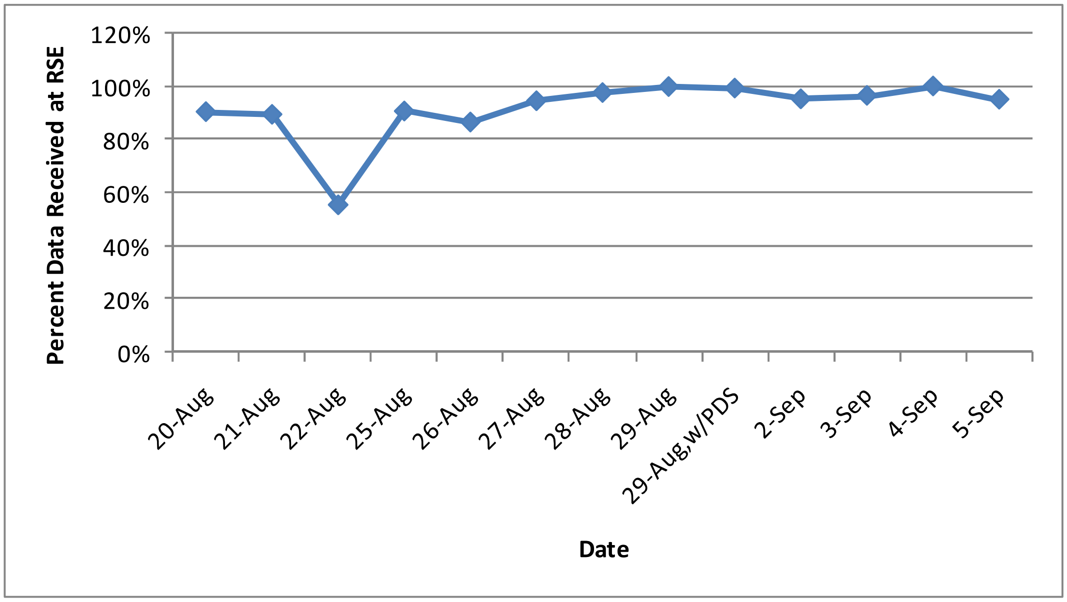

Figure 101: OBE to RSE Probe Loss by Day of Testing

On August 29, the PDS was not operational until 1 pm EDT. The data above shows the percentage of data received at RSEs before and after the PDS was restarted. It is apparent that the PDS has no significant effect on probe data loss between the OBE and RSE, as expected. It is not clear what caused the unusually low probe data success rate on August 22.

The average end-to-end probe data success rates—the percentage of probe data successfully transmitted from the OBE to the applications database—ranged from 60 percent to over 90 percent, with typical values in the 80-to-90-percent range. Days with low probe data success rates were due to problems with the PDS (e.g., PDS not operational, RSE not connected to PDS).

Figure 102 illustrates the average data transmission percentages by day.

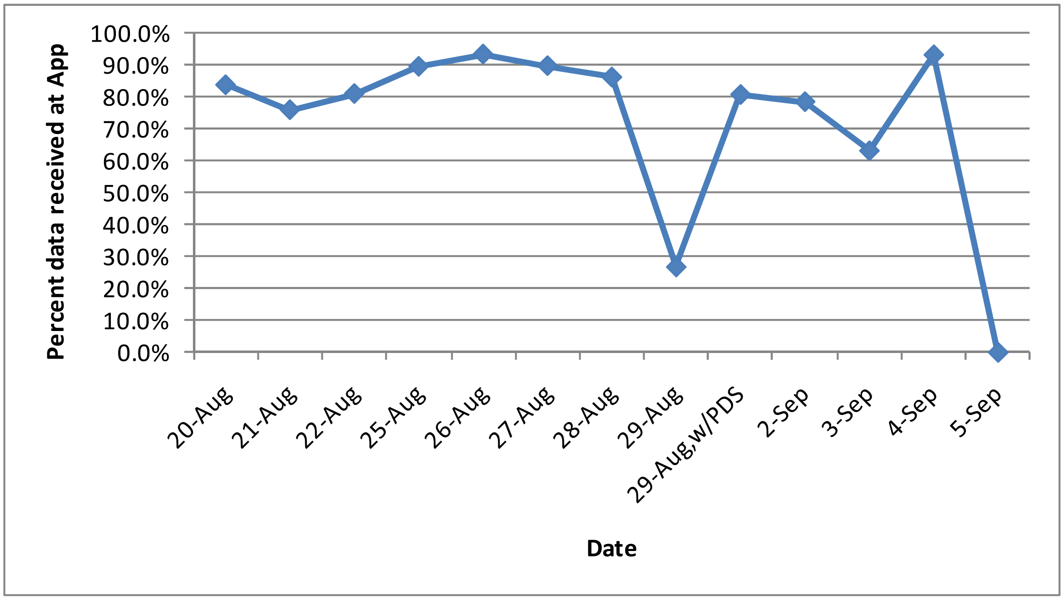

Figure 102: Average OBE-Applications Database Daily Transmission Percentages

As confirmed by the data from August 29, the PDS has a significant effect on probe data loss between the OBE and Public Apps. The RSEs used for the slippery conditions test on September 5 was not correctly connected to the PDS, so the probe data success rate for that day shows as 0 percent because none of the probe data was captured in the Public Apps data. However, the probe data have been captured in a text format. Based on these observations, many of the results presented later in this section will not include these 2 days (August 29 and September 5) in order to keep the PDS malfunctions from skewing the results.

An unusually low day of just 60-percent probe data success rate occurred on September 3, the day of turning movements testing. This could be because of RSE issues, which are explained in the **Analysis by RSE** section.

Values from just below 80 percent to just above 90 percent transmission were seen on all other days. The major sources of data loss were airlink loss, ARP loss, and the loss of broker connections at some RSEs.

It should be noted that on August 22, 25, and 26, the Public Apps database received more probe data than were recorded in the RSE logs. After looking at the details of this discrepancy, it was concluded that missing RSE logs caused this to happen.

Daily Averages – Latency

The average end-to-end daily latency for probe data snapshots (OBE transmission to Public Apps database reception) ranged from just over 0.5 second to almost 1.5 seconds. It is not always clear what caused some days to have greater latency than others. However, all days fell within this range.

On August 29, issues were experienced with the probe data service that prevented transmission of data until late in the afternoon. Once the probe data service was restarted, a large number of snapshots were delivered with a latency of several hundred to several thousand seconds. This caused the average latency of data on August 29 to be around 200 seconds. This was clearly due to the restarting of the Probe data service. On September 5, routing of RSE data to the Public Apps database was down. The data was captured in text files. However, that data was not uploaded or analyzed, and thus the loss and latency values for September 5 between the OBE and Public Apps could not be calculated.

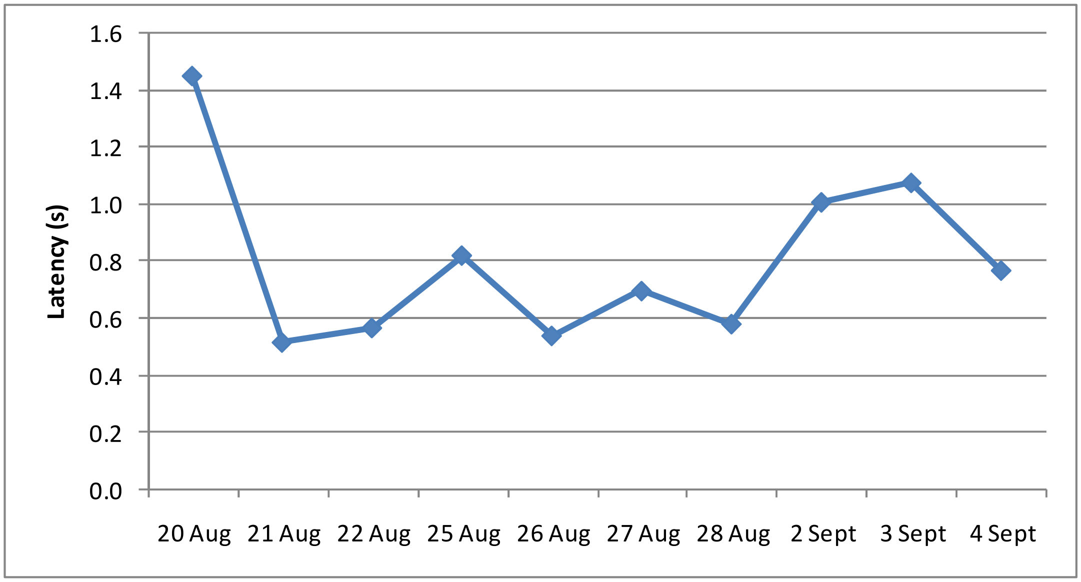

Figure 103: End-to-End Latency excluding August 29

Overall Averages – Loss

Table 59 includes the percentage of data successfully received by the RSE and by the Public Apps using the total number of messages across all days. The aggregate probe data success rate to the RSE was about 87 percent. The aggregate probe data success rate to the Public Apps database was about 75 percent when the days of PDS malfunction were included, but about 83 percent when they are not included.

Table 59: Percentage of Probe Data Received by RSE and APP

	% Received by RSE from OBE	% Received by Apps from OBE	% Received by Apps from RSE
Average all Received	86.9%	76.2%	87.6%
Omitting bad days (Aug 29, Sept. 5)	87.8%	83.2%	94.8%

Table 60 displays the overall probe data loss values that can be computed, omitting August 29 and September 5.

Table 60: Overall Probe Data Loss excluding August 29 and September 5

Avg OBE-RSE loss:	12.2%
Avg RSE-App loss:	12.4%
Avg non ARP RSE-App loss:	5.5%

The first value is 100 percent minus the value of RSE reception, 87.8 percent, indicating an average loss of 12.2 percent between the OBE and the RSE. The second value reflects the loss between the probe data received at the RSE and data received at the Public Apps (100 percent minus 87.6 percent). The final value is the loss rate adjusted for the 6.9-percent ARP loss. This rate was computed by counting one snapshot loss for every pause in data reception of more than 30 seconds in the Public Apps. It may be a slight underestimate of the ARP loss, since for pauses of over 1 minute, only one snapshot loss is counted. Therefore, a singleton snapshot every 30 seconds would not be counted in the estimate.

Loss between the OBE and the Public Apps database is due to some RSEs not being included in probe data subscriptions, and some RSEs not transmitting probe data, despite being subscribed to and having PDS connections. The tests were not re-run because September 12 was set as the absolute deadline by the US DOT for Public Apps testing. The following section will analyze this.

Overall Averages Omitting Bad RSEs

It is also insightful to look at aggregate averages when the six RSEs with unusually low probe data success rates are omitted (out of a total of 50 RSEs that were used during probe data testing). These six RSE are discussed in the section on RSE data transmission.

Table 61: Aggregate Averages excluding 6 RSEs with Unusually Low Probe Data Success Rates

	% Received by RSE	% Received by Apps from OBE
Average all Received	93.6%	80.2%
Omitting bad days (Aug 29, Sept. 5)	87.0%	88.6%

When these RSEs are omitted, the Apps database reception level rises to 80 and 89 percent (with and without the bad days of August 29 and September 5, respectively). The RSE rates are higher when those bad days are included, but actually lower when they are not included.

Table 62: Overall Probe Data Loss excluding 6 RSEs with Unusually Low Probe Data Success Rates

Avg OBE-RSE loss:	13.5%
Avg RSE-App loss:	-1.6%
Avg non ARP RSE-App loss:	-9.8%

The average OBE-RSE loss is slightly higher, at 13.5 percent, but surprisingly, the average OBE-App loss is -1.6 percent, and subtracting the ARP loss makes it an even more negative, -9.8 percent.

These two negative numbers are the result of the unusual results on August 22, 25, and 26, when the Public Apps database received more data than the total snapshots in the captured RSE logs. This fact is due to missing RSE logs. For example, if OBEs send out 10 snapshots, five each to two RSEs, and only one RSE log is captured, it would cause the average RSE-App loss to be -100 percent because the app would have 10 snapshots, and the RSE log shows a total of five received.

Overall Averages – Latency

Table 63: Average Latency for Probe Data

Average of all received:	6.19 seconds
Average of all received excluding August 29:	0.81 second

When the "bad day" of August 29 is included in the analysis, the average latency is over 6 seconds per snapshot, but as discussed above, the PDS service was not performing normally on August 29. When performing normally, the average latency is about 0.8 second.

15-Minute Averages – Loss

Probe data success rate values for most 15-minute periods ranged from 40 percent to nearly 100 percent by 15-minute period. A majority of 15-minute periods had values in the 75-to-95-percent range. Early morning and late evenings had greater variability in data points, due to the fact that there were fewer days of testing, and hence fewer data points in these sections.

A few 15-minute periods had low probe data success rates. This can be explained by the fact that at these times of day, a disproportionate amount of testing was done on RSEs from which there was no or very little data transmission.

No clear pattern is visible in the data, suggesting that time of day and traffic level do not influence latency or loss. This result confirms the expectation explained above.

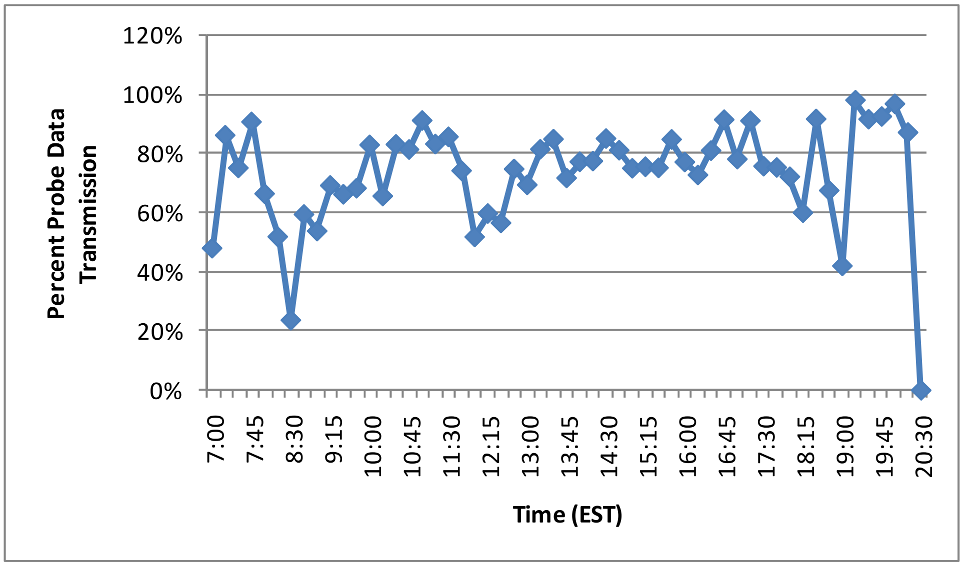

Figure 104: Probe Data Transmission excluding "Bad Days"

15-Minute Averages – Latency

Probe data latency values ranged from 0.4 second to just over 1 second, but a few outlier measurements occurred after 4 PM on a few days, during a few 15-minute periods, enough so that some 15-minute periods spiked to latency values of 3 and even 7 seconds.

There were fewer data points after 4 PM than before 4 PM, allowing outlier values to skew the results more heavily than at other times. Further analysis is needed to determine what caused these outlier values to occur. A few OBEs seemed to record the wrong timestamps on some occasions, so some of these outlier values may be caused by those OBEs.

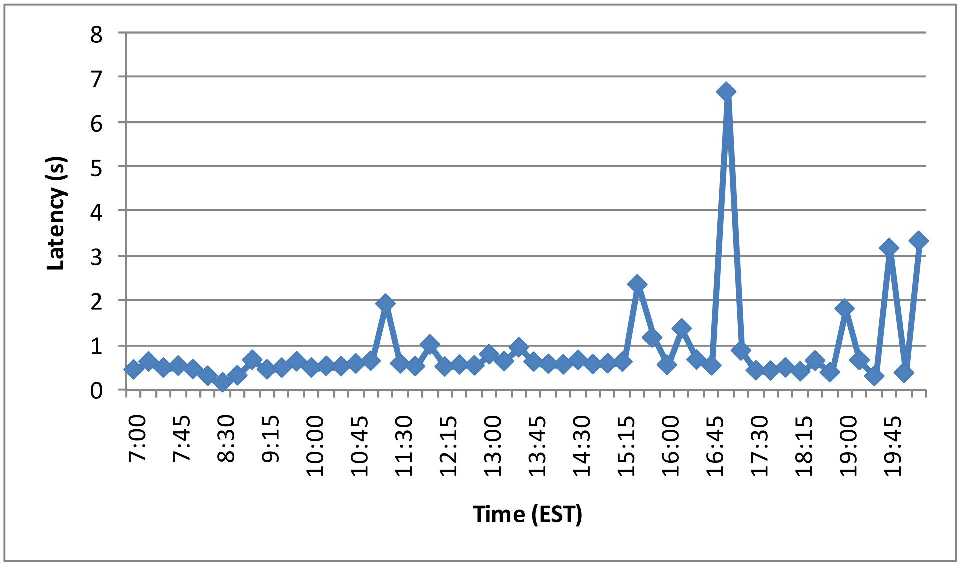

Figure 105: Probe Data Latency by 15-Minute Period excluding September 29

RSE Averages – Loss

Analyzing probe data by RSE allows the determination of outlier or bad RSEs, which had unusually low probe data success rates. It also gives an indication of what typical airlink loss is.

The results show that many RSEs had no loss at all (100-percent probe data success rate); some had a 0-percent probe data success rate (mostly due to faulty broker connections); and some had probe data success rates ranging from 59 percent to 99 percent, due to a mixture of airlink loss and faulty broker connections or missing probe data subscriptions.

Figure 106 shows the RSEs with zero snapshots transmitted to them as 0 percent, and the remaining RSEs are shown with the actual transmission percentage. The remaining RSEs were fairly evenly distributed in the 5-percent bins between 75 percent and 99 percent.

There was significant variability in the probe data success rates by RSE. Almost half of the RSEs (22) had a 100-percent probe data success rate—no loss between OBE and RSE. Of the 52 RSEs, 10 had no snapshots transmitted to them, although of these 10, two were not yet active (I-96 and Kent Lake and M5&Farmington), and one was not on any route (9-Mile and Drake).

The remaining seven RSEs must be investigated to determine why they did not receive any snapshots. In some cases, the PDS broker connection was down during testing (Grand River and Novi, I-96 and Beck, perhaps others). In other cases, further investigation will be necessary.

For the remaining RSEs, which had probe data success rates between 65 percent and 99 percent, specific analysis will be needed to determine how much of the loss was due to pure airlink loss (loss in transmission between the OBE and the RSE), missing broker connections, or missing subscriptions.

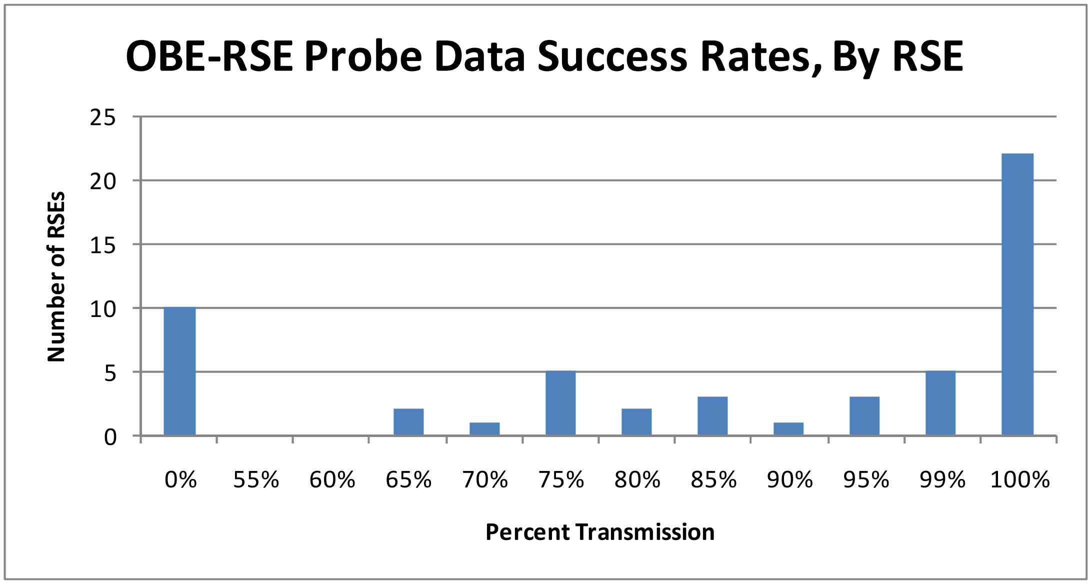

Figure 106: OBE-RSE Probe Data Success Rates by RSE

RSE Averages – Latency

Analyzing latency among RSEs will help to determine outlier RSEs and to characterize the effect an RSE's backhaul connection has on the latency of probe data transmitted to the RSE. Some distribution is expected because of the different latencies of the different backhaul networks, with hardwire connections expected to have the lowest latency, 3G cell phone connections expected to have the next lowest latency, and WiMAX connections expected to have the highest, although 3G and WiMAX latencies will depend on the geometry of the antenna for the given RSE.

Figure 107 shows probe data latency from the RSE to the Public Apps by RSE backhaul type.

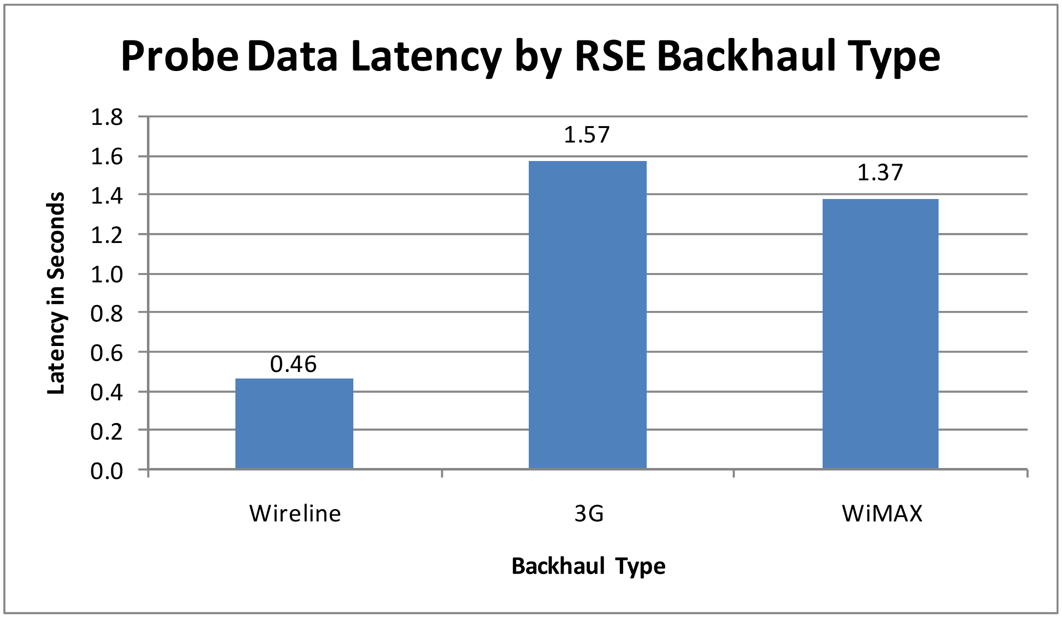

Figure 107: Probe Data Latency by RSE Backhaul Type

Error! Reference source not found. Figure 108 shows the distribution of OBE to Public Apps database probe data latency, by the RSE from which it is transmitted. For this diagram, data from August 29 has been omitted to prevent distortion as described above.

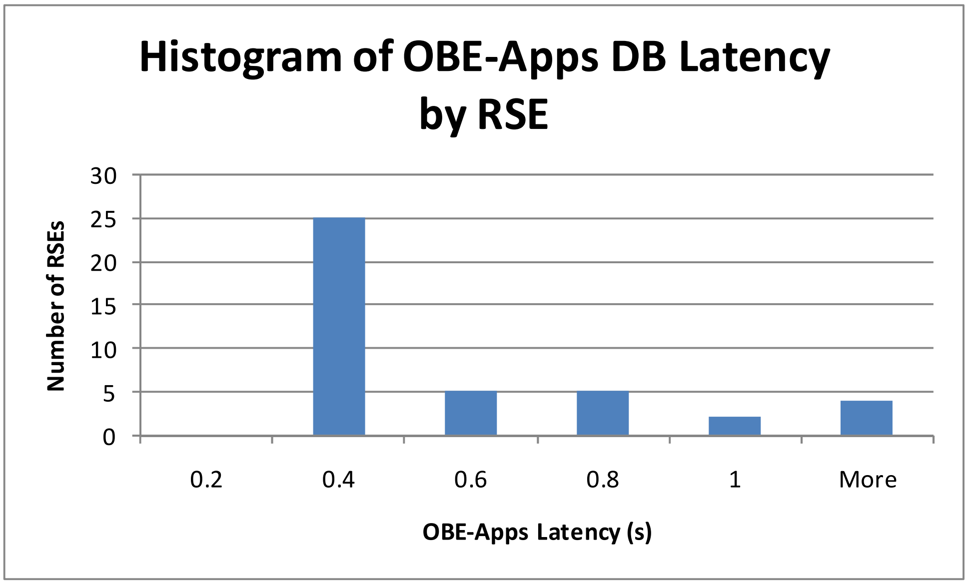

Figure 108: OBE-Application Database Latency

The results show that typical end-to-end latencies vary from 0.4 second to 1 second depending on the RSE. Unit testing indicated typical latency values of 1 second or less for most RSEs, and the Public Apps testing confirmed this value for larger number of vehicles. A majority of the RSEs had between 0.4 second and 0.6 second latency. Five RSEs had between 0.6 second and 0.8 second, and five RSEs had between 0.8 second and 1 second. Six RSEs had latencies higher than 1 second, and all six had either

WiMAX or 3G backhaul connections. Such variation is explained by the difference in backhaul connections, above.

Freeway versus Arterial

This measure characterizes the difference between probe data transmitted RSEs located on Freeways and RSEs located on arterial routes. For these measures, the "bad days" (Aug 29, Sept 5), and "bad RSEs" were removed from the computation, since they are not indicative of the performance of the system in normal operation.

Since snapshot transmission only takes place once a V-DTLS connection is established, no difference in performance was expected between snapshots sent to freeway RSEs and arterial RSEs. The loss rates were nearly identical. However, the latency rates were somewhat different. Latency rates for freeway RSEs were, on average, 0.1 second longer than the arterial RSEs. This difference is almost certainly due to the fact that freeway RSEs were exclusively WiMAX backhaul, while arterial RSEs were a mixture of T1, 3G, and WiMAX backhauls.

Freeway versus Arterial – Loss

Table 64: Probe Data Loss by Road Type

	Number of Snapshots	Number at RSE	% Received at RSE	Number at Apps	% Received at Apps
Freeway	10,428	9,048	86.7%	9,638	92.4%
Arterial	62,424	54,101	86.7%	54,920	88.0%

Freeway versus Arterial – Latency

Table 65: Probe Data Latency by Road Type

Latency	Average Latency
Freeway	0.81
Arterial	0.69

Snapshot Type

This measure characterizes how the three different types of snapshots performed with respect to probe loss and latency. The three snapshot types are periodic (created while vehicle is driving), start/stop (created when a vehicle first stops, regularly during stops, and then again when a vehicle starts), and event-triggered (created when there is an ABS, traction control, stability control, or brake boost event). For these measures, the "bad days" (August 29, September 5), and "bad RSEs" were removed from the computation, since they are not indicative of the performance of the system in normal operation.

Since all snapshots are transmitted the same way, over the same protocol, and the type is purely internal to the snapshot, no difference in performance between the snapshot types was expected. This is in fact what was observed, with both periodic and start/stop snapshots having about an 89-percent reception rate at the Public Apps database, and an average latency of about 0.78 second.

The values for event-triggered snapshots were somewhat different, but because the sample size was so much smaller—64 snapshots compared to over 22,000 start/stop and over 50,000 periodic—this difference is not statistically significant.

Snapshot Type – Loss

Table 66: Probe Data Loss by Snapshot Type

Snapshot type	Number of Snapshots	Number at RSE	% Received at RSE	Number at Apps	% Received at Apps
Periodic	50,355	43,765	86.9%	44,418	88.2%
Start/Stop	22,433	19,330	86.2%	20,079	89.5%
Event-triggered	64	57	89.1%	61	95.3%

Snapshot Type – Latency

Table 67: Probe Data Latency by Snapshot Type

Snapshot type	Latency
Periodic	0.84
Start/Stop	0.72
Event-triggered	0.46

Vehicle Configuration

This condition characterizes how vehicles were spaced apart, with respect to each other during testing.

For these measures, the “bad days” (August 29, September 5), and “bad RSEs” were removed from the computation, since they are not indicative of the performance of the system in normal operation.

Vehicle Configuration – Loss

While it was expected that vehicle configuration would have a significant effect on the probe data success rates due to V-DTLS limiting one vehicle at a time to connect to the RSE, it was not expected that vehicle configuration would have any significant impact on data loss, since a successful V-DTLS connection was a prerequisite for the analysis done here.

The testing proved that this was the case. All vehicle configurations resulted in a 92 percent +/-6 percent probe data success rate between the OBE and Public Apps.

Table 68: Probe Data Loss by Vehicle Configuration

Vehicle Configuration	Number of Snapshots	Number at RSE	% Received at RSE	Number at Apps	% Received at Apps
Dispersed – 3-minute spacing	28,179	21,974	78.0%	24,212	85.9%
Dispersed – 1-minute spacing	4,616	4,150	89.9%	4,529	98.1%
Batched	10,344	9,164	88.6%	9,950	96.2%
Platooned	13,328	12,763	95.8%	12,180	91.4%

Vehicle Configuration – Latency

It was expected that latency would not be affected by vehicle spacing, since snapshots are only counted towards latency if they are successfully transmitted to the RSE. The average latency of dispersed and batched vehicle configurations was 0.5 second +/-0.06 second, while the Platooned configuration had a somewhat greater latency. This is probably due to the fact that about half of the data from Platooned testing came from the day of start/stop testing, on which vehicles spent a disproportionate amount of time on freeways, and hence, transmitting to RSEs with WiMAX connections, which were known to have longer latencies.

Table 69: Probe Data Latency by Vehicle Configuration

Snapshot type	Average Latency(s)
Dispersed – 3-minute spacing	0.56
Dispersed – 1-minute spacing	0.47
Batched	0.55
Platooned	0.75

Weather

It was assumed that precipitation could negatively impact probe data success rates and probe data latency, since it sometimes interferes with communication channels, and that cloud cover could also impact probe data success rates. However, this data analysis could not support or deny these expectations, due to a lack of data (only 3 days were recorded with adverse weather conditions) and the inability to disambiguate the effects of adverse weather from other issues.

Most data was collected on sunny clear days. The test observer forms indicated that there were 3 days that were not sunny and clear. Table 70 presents data from these days.

Table 70: Probe Data Loss by Weather Type

Day	Weather Condition	OBE-Apps DB Loss	OBE-Apps DB Latency
22-Aug	Mostly Cloudy	82.1%	0.57
25-Aug	Partly Cloudy	90.7%	0.82
29-Aug	Rain	86.8%	23.6

Note that data for August 29 was for the tests run after 3PM only, since PDS problems were experienced before 3PM.

The end-to-end probe data success rates ranged from 82 percent to 91 percent, which were typical values for probe data loss. August 22 experienced an average latency of .57 second, very typical values, and August 25 had a latency of 0.82 second, slightly longer than typical, but well within the range of standard values.

The value for latency on August 29 was 23.6 seconds, even when only the data from after 3 PM, when PDS service issues were resolved, were considered. This latency is much longer than for all other days, and as discussed above, is likely related to the recently resolved PDS. There is not enough evidence to attribute the unusual latency solely to rain.

Thus, this data analysis is insufficient to conclude whether or not adverse weather has a significant impact on probe data loss or latency.

Conclusions

The VII system implemented for POC works well as a probe data collection tool. Some individual cases need more detailed analysis, but some general conclusion can be reached.

OBE to RSE communications often have 100-percent success, but sometimes have values in the 90-to-100-percent range values due to airlink loss. Cases where lower probe data success rates were observed during testing require further analysis, but the lower success rates are likely due to factors such as lack of a subscription or PDS broker connection issues.

OBE to Public Apps probe data success rates were around 90 percent. The reasons for OBE to Public Apps loss include airlink loss, missing subscription loss, and PDS broker connection loss. In addition, the OBE to Public Apps database experienced loss due to the ARP issue. Because the average reception rate to the RSE and to the Public Apps database were very similar, the analysis did not indicate any other sources of loss.

The goal of this test was to compare the results, in this case 90-percent OBE to Public Apps success rate, to the results of the OBE-to-RSE integration testing, dry-run testing, and test track testing. It has been established that the current data is insufficient for a real comparison.

Battelle did some testing at the track and has some data; however, all of that data was collected without V-DTLS installed in the OBE and RSE. Once an OBE became available for integration testing in the DTE, it was plagued by an instability issue. The OBE would not stay on long enough for data collection. At the same time, V-DTLS was just released and its functionality was being tested at the RSEs. A revised OBE version became available in June of 2008, which caused simultaneous integration testing and OBE software build testing. However, testing was stopped by the US DOT after 1 week. During that week, for successful V-DTLS connections, the probe data success rate was about 95 percent. However, the success rate was extremely hard to determine, due to problems with the positioning system and the probe data application itself, which were later fixed before start of applications testing.

The average end-to-end latencies were 0.8 second, with most data latencies in the range of about 0.4 second to 1 second. A small number of longer latencies occurred. Further analysis is needed to determine what caused outliers, but their small number indicates that these are anomalous circumstances, possibly caused by incorrect OBE timestamps.

Both of the hypotheses that “A significant percentage of probe data snapshots that are securely transmitted by OBEs reach subscriptions” and “Probe data snapshots take less than 1 minute to propagate from an OBE to a network subscriber under all conditions” were validated by the results.

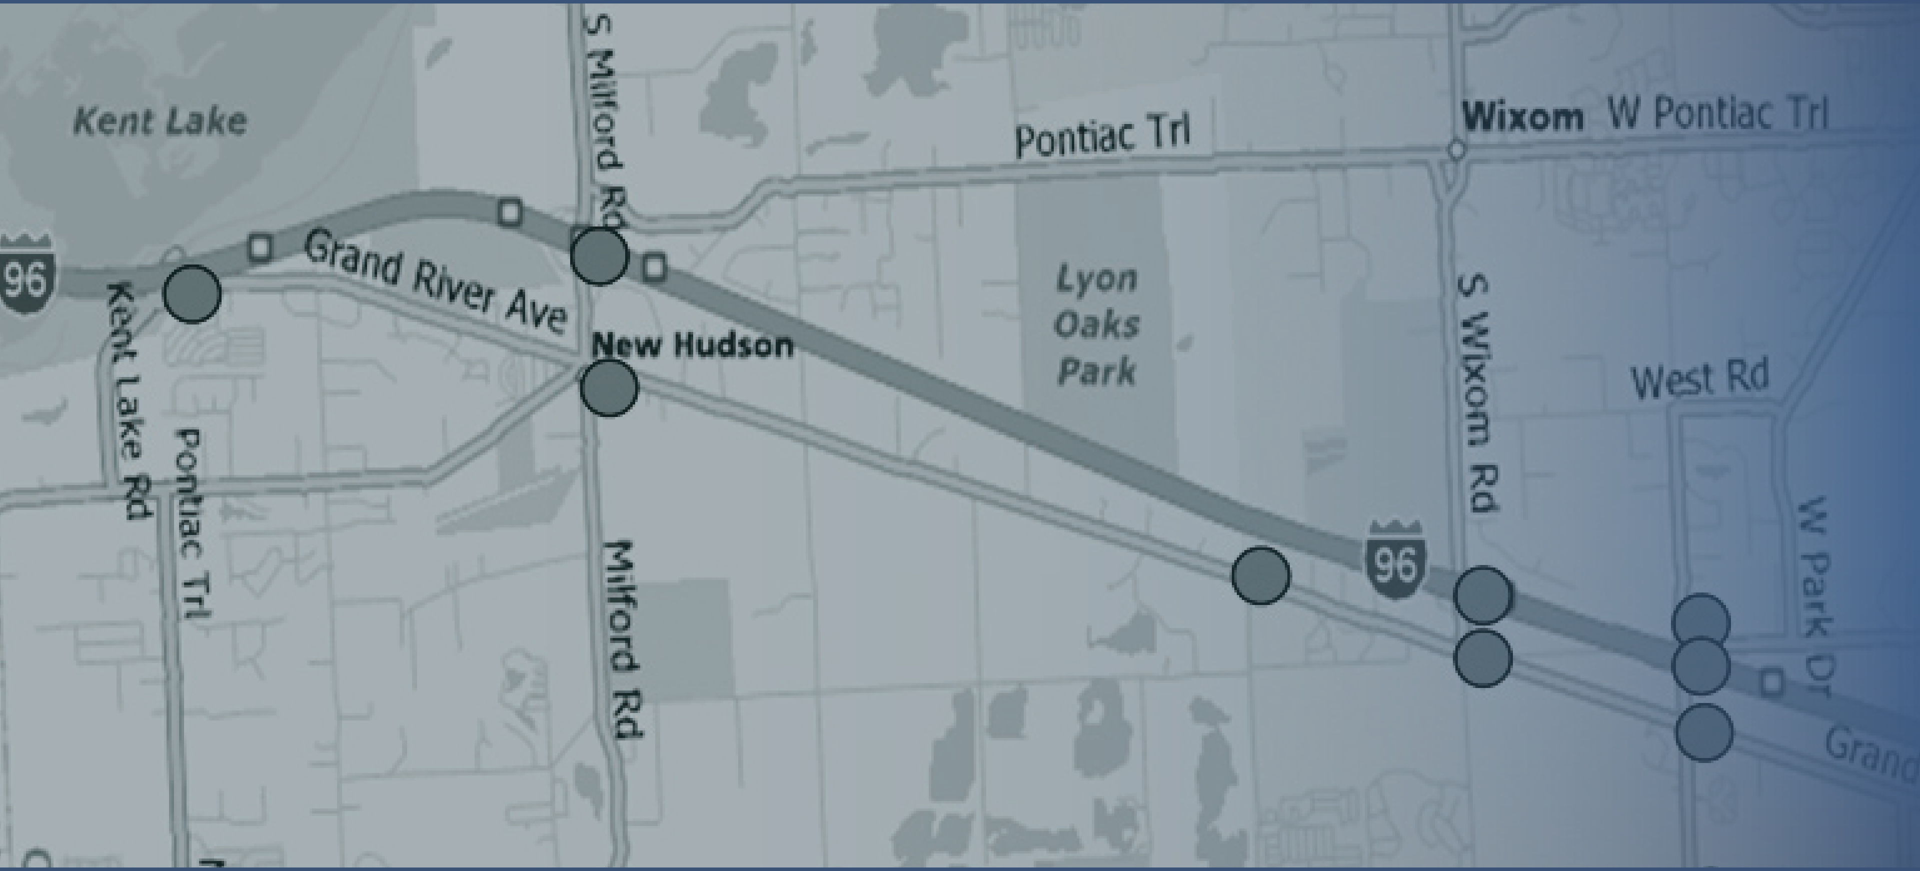

U.S. Department of Transportation
Research and Innovative Technology Administration

U.S. Department of Transportation | 1200 New Jersey Avenue, SE | Washington, DC 20590

Publication No. FHWA-JPO-09-057, EDL No. 14488