UC Berkeley
Research Reports

Title

Cooperative Intersection Collision Avoidance System (CICAS): Signalized Left Turn Assist and Traffic Signal Adaptation

Permalink

https://escholarship.org/uc/item/30c8j1kc

Authors

Misener, Jim
Barnes, M.
Chan, Ching-Yao
et al.

Publication Date

2010-04-01

eScholarship.org Powered by the California Digital Library
University of California

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Cooperative Intersection Collision Avoidance System (CICAS): Signalized Left Turn Assist and Traffic Signal Adaptation

Jim Misener, et al.

California PATH Research Report
UCB-ITS-PRR-2010-20

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation, and the United States Department of Transportation, Federal Highway Administration.

The contents of this report reflect the views of the authors who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the State of California. This report does not constitute a standard, specification, or regulation.

Final Report for Task Order 6608

April 2010

ISSN 1055-1425

Cooperative Intersection Collision Avoidance System (CICAS): Signalized Left Turn Assist and Traffic Signal Adaptation

Task Order 6608

FINAL REPORT

Jim Misener
M. Barnes
Ching-Yao Chan
Delphine Cody
Susan Dickey
R. Goodsell
Tim Gordon
Zu Whan Kim
Tom Kuhn
Thang Lian
David Nelson
Christopher Nowakowski
K. Nubukawa
Ashkan Sharafsaleh
Steven Shladover
John Spring
Joel VanderWerf
Wei-Bin Zhang
Liping Zhang
Kun Zhou

University of California, Berkeley
1357 S. 46th Street, Bldg. 190; Richmond, CA 94804-4648

March 2010

Table of Contents

Table of Figures

Table of Tables

EXECUTIVE SUMMARY:

Overview of the CICAS Project

The Cooperative Intersection Collision Avoidance (CICAS) program is a multi-year, cooperative research program including federal, state, academic, and industry partners. The goal of the research program is to use ITS technologies to address the problem of intersection crashes. The program is funded through an 80/20 cost share, typically split between the U.S. Department of Transportation (D.O.T.) and a local state D.O.T. The program began in 2003, and has been divided into three functional segments based on crash type. The largest programmatic segment is CICAS-V (Violation) which is led by CAMP, with support from researchers at Virginia Tech, and aims to address the problem of straight crossing path collisions which tend to be the result of stop sign or stop light violators. A second programmatic segment is CICAS-SSA (Stop Sign Assist) which is led by the Minnesota D.O.T., with support from researchers at the University of Minnesota, and aims to address the problems associated with crossing or entering a high-speed, rural road from a stop sign at a minor road. The final programmatic segment is CICAS-SLTA (Signalized Left Turn Assist) which is lead by Caltrans, with support from the California PATH Program (Partners for Advanced Transit and Highways) at the University of California, Berkeley. This CICAS-SLTA segment aims to address crashes caused by vehicles making left turns at signalized intersections where there is no protected left-turn signal.

This report focuses on the activities related to the CICAS-SLTA program segment. It builds upon work that was performed during earlier segments (Task Orders 5600, 5601 and 6607) which included the previous Intersection Decision Support (IDS) project and CICAS. The CICAS-SLTA project is currently continuing research in the areas of sensor testing, field data collection, human factors data collection and testing, and warning algorithm design and simulation. The project aimes scheduled to lead to the design of a Field Operational Test (FOT) ready system and detailed FOT work plan, depending on the success of this and a subsequent effort.

In this final report, we describe Cooperative Intersection Collision Avoidance System-Signalized Left Turn Assistance (CICAS-SLTA) activities to address so-called Left Turn Across Path-Opposite Direction (LTAP-OD) problem, where drivers have cognitive, time pressure or obstruction challenges to make safe left turn decisions with a permitted signalized left turn. We also summarize salient portions of the signal violation case where CICAS would assist, suggesting an all-red interval which would be triggered by a CICAS-equipped vehicle which is in turn providing a violation warning to its driver; in this manner, all intersection users – equipped and unequipped – could benefit. We call this CICAS-Traffic Signal Adaptation (TSA).

Specifically, we describe results from our now-close alignment of this research to other US DOT Vehicle-Infrastructure Integration (VII) activities, where the intersection is equipped with low latency, highly reliable wireless Dedicated Short Range Communication with vehicles. We detail our Concept of Operations which recognizes that because the LTAP-OD involves at least two interacting vehicles, to make CICAS-SLTA effective and to address the very important question of how to use some roadside surveillance to account for other roadway users which may affect the movement of the left-turning vehicle. We also describe the CICAS-TSA Concept of Operations. However, we note that the preponderance of TSA work was conducted under the

related Task Orders 5210 and 6210. More detailed TSA information are available in those task order reports.

The ultimate goal of any intersection collision avoidance system is to reduce intersection crashes by providing drivers with information to support turning decisions. The work we describe directly supports this and is the product of several years' work in naturalistic driving data collection, driver modeling, and *in toto* developing an integrated intersection safety approach while investigating the applicability of a large set of already- or nearly- available "commercial off the shelf" systems towards designing a CICAS. Our particular CICAS research has focused and continues to focus on the Left-Turn Across Path with Opposite Direction Traffic (LTAP-OD) crash type (see Misener and Shladover (2006)), although there is significant other CICAS work in North America that investigates other crashes, according to Chang, et al (2007). To that extent, we also investigate a version of signal violation safety that countermeasures this type of crossing path crash by dynamically extending the all-red interval, TSA. Both elements of the work are described herein.

Why our focus on LTAP/OD? This is a common crash scenario in which a driver makes a left turn across the path of a vehicle approaching from the opposite direction (i.e., "Left Turn Across Path/ Opposing Direction" or LTAP/OD crash scenario). LTAP/OD crashes account for 27.3 % of all US intersection-related crashes and two-thirds of all LTAP/OD crashes occur at signalized intersections (see Misener, et al (2007)). Before designing an infrastructure countermeasure system, the reasons for such crashes were considered including:

- driver failure to judge safe time gaps correctly,
- driver failure to judge the speeds of closing vehicles,
- driver failure to see the oncoming vehicle (i.e., "looked but did not see"), and
- obstruction of the driver's view by an opposing vehicle.

The CICAS-SLTA Concept of Operations (ConOps) is summarized in Figure 1 below. At its core is the concept of a staged or graduated alert, changing in intensity (warning-to-decision support), depending on the subject vehicle (SV) maneuver stage. In the figure, the degree of alert is connoted by a progression of color, yellow (low intensity), to orange and to red (high intensity). The ConOps illustrates the need: the research challenge and direction of the project is that to determine a graduated alert, drivers, their movements and the potential means to interact with the driver should be determined as a function of road type, traffic volume and presence of other road users. Clearly, to fulfill this model, a data-driver research plan is needed; this is the plan.

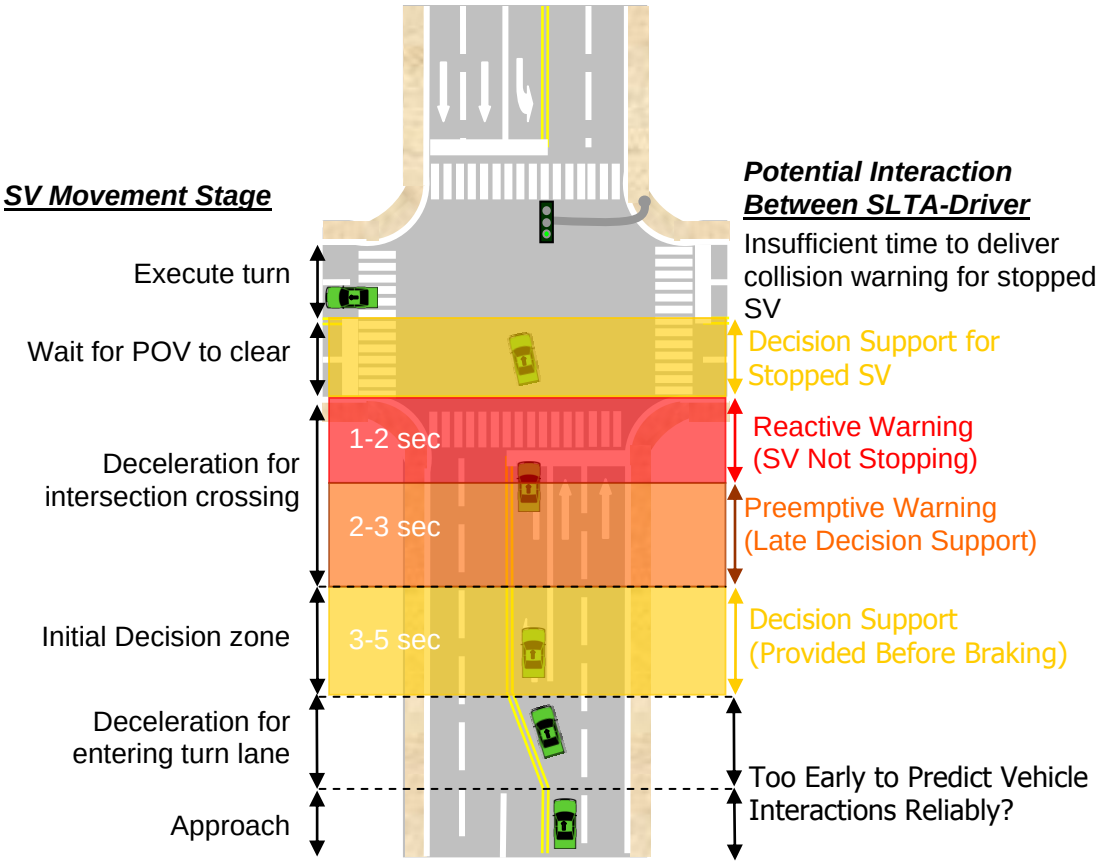

Figure 1. Illustration of CICAS-SLTA ConOps

It is useful to consider that research activities within the CICAS-SLTA project may be clustered in three areas:

1. Increased emphasis on accounting for other road users – pedestrians and bicyclists – and their interaction with left-turning vehicles. Investigations of this are described by I. Banerjee et al (2005) and Chan, et al (2006). Methods of reliable real-time detection of pedestrians and bicyclists are important elements of CICAS-SLTA. In the end, in collaboration with CICAS researchers focused on signal violation, we plan on developing a practical and inclusive urban intersection state map.

2. Further outreach and refinement/fusion of commercially available roadside sensing system to address CICAS-SLTA requirements. The resultant system should be able to detect and provide to the intersection state map precise positions and other state information of road users, with particular focus on these other road users and the "threat" or principal other vehicle (POV).

3. An extensive experimental and measurement program spanning a range of observation types, from roadside observations of large quantities of vehicles at different intersection settings, to controlled settings with naïve drivers at our Richmond Field Station

Intelligent Intersection, and finally to examining behaviors of naïve drivers with an instrumented car on public roads.

Recently, we have rethought and are detailing our ConOps to more closely align to the US DOT's Vehicle-Infrastructure Integration (VII) program. Because the LTAP-OD involves at least two interacting vehicles, to make CICAS-SLTA effectiveness to approach the effectiveness of individual vehicles approaching the intersection and communicating via Dedicated Short Range Communication (DSRC) transceiver. Because of issues of market penetration and to address the very important question of how to account for other roadway users which may affect the movement of the left-turning vehicle, the CICAS-SLTA requires some sensor system on the roadside that acquires and tracks range and range rate of approaching POVs and other roadway users. However, with VII, we need not sense the VII SV, which are equipped with DSRC radios. In this manner, CICAS-SLTA leverages off the VII program and actually extends that program by integrating pedestrian and bicyclist state tracking.

As a case in point, in recent years the research reported by Chang, (2007) has been expanded with the US DOT and consortium of automobile manufacturers into an effort called Cooperative Intersection Collision Avoidance Systems – Violation (CICAS-V) (see Neal et al (2006)). These researchers focused on enhancing the conspicuity of the red signal by providing an additional roadside warning display and an in-vehicle audible or visual alert. Additionally, research is underway as part of the CICAS program in a concept called Traffic Signal Adaptation (TSA), where the TSA works with the CICAS-V system to additionally provide a dynamic all-red extension with the objective to enact a safe clearance interval (and provide a photo citation to the signal violator).

Certainly, as outlined in Chang (2007) there are other intersection crossing path crash types that could be addressed by VII. In the past several months, the left-turn across path CICAS countermeasure effort called Signalized Left Turn Advisory (SLTA) has also been transformed into a VII-enabled effort. Very much like CICAS-V, the intersection phase and timing would be given to the "subject vehicles". Unlike CICAS-V, the subject vehicles would be left-turning vehicles that are not necessarily violating traffic signals; rather, the CICAS-SLTA vehicles would be given left turn advisory, with additional information given to the SLTA-based or roadside computers that either sense or are communicated via DSRC the trajectories of straight through vehicles which may present a crossing path conflict. A complicating and important factor for CICAS-SLTA would be the presence of other road users, most specifically pedestrians and bicyclists.

Signal Phase and Timing Acquisition

Both CICAS-SLTA and CICAS-V rely on acquisition of traffic signal timing. Prior to the introduction of novel concepts as CICAS, this was regarded as a mature, though still-evolving field. Many existing traffic controllers already run software that supports communication protocols, such as the National Transportation Communications for ITS Protocol (NTCIP) [17] and California's Assembly Bill 3418 (AB3418) protocol, designed to communicate signal phase and loop detector information to traffic management centers.

These protocols were designed for monitoring and control of actuated and coordinated intersection timings, from a central location over a modem, at a time granularity of seconds and were not designed to notify vehicles of real-time signal phase count-down, needed for CICAS and over a high bandwidth DSRC connection. Traffic signal controller protocol messages also do not contain geographical/mapping information needed by vehicles in order to recognize what part of the signal phase information applies to the vehicle's approach to the intersection. Importantly, many intersections contain legacy equipment that do not support external communication.

To address this and in the vein of conducting research that will lead to on-the-road safety deployment, we have at different installations with different legacy systems we have used the NTCIP protocol, the AB3418 protocol and a special "sniffer" circuit that non-invasively detects signal phase without any interaction with the traffic signal controller.

Interface to VII and DSRC

All three methods of acquiring signal phase timing information are integrated with the same phase count-down and broadcast software. It is desirable to have well-specified interfaces between the signal phase acquisition, the phase countdown and the broadcast in order to isolate changes required in different installations. The modular system design and logical connections shown in Figure 2 have been used at all three of our roadside testbed intersections, two of which use the sniffer and one of which uses AB3418, as well as at the experimental intersection, where we have tested all three methods. In each case the complexity of the software lies in the phase countdown, as we are enacting a parallel processing of the internal logic of the traffic signal controller in order to duplicate the countdown that the traffic signal controller conducts to set signal phases.

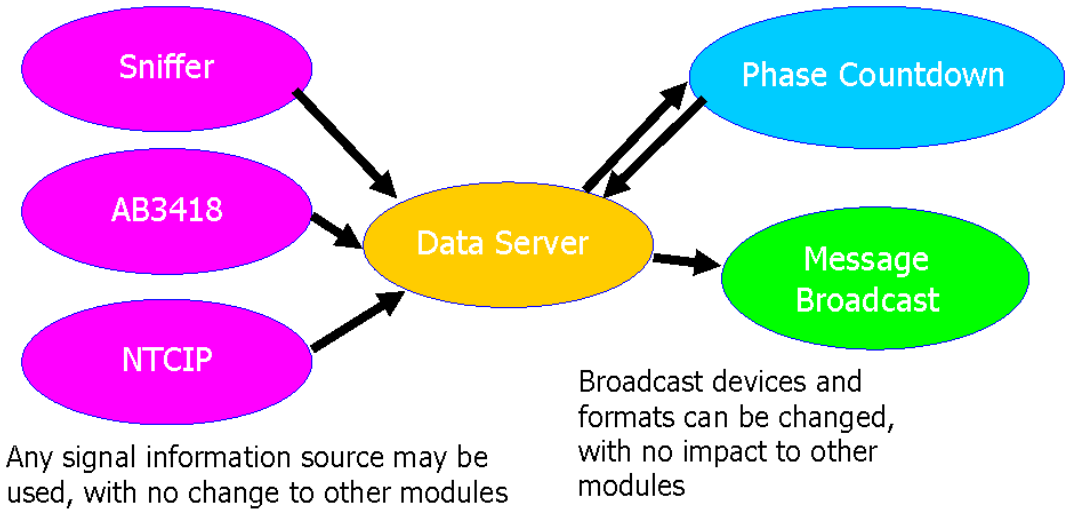

Figure 2. Modular Design of Signal Phase Acquisition and Broadcast System

All-Red Clearance

As an additional component to left-turn warning, the project also addresses providing a dynamic all-red clearance interval by DSRC communication from the vehicle to the traffic signal controller such that the "victim" vehicle is not involved in a side collision. This is called Traffic Signal Adaptation (TSA) and is enabled by a CICAS-V violating vehicle-to-traffic signal controller message. Our TSA ConOps is illustrated in brief in Figure 3. Because VII privacy concerns may preclude automated or wireless ticketing, an important adjunct of TSA is the inclusion of a red light photo enforcement system to dissuade drivers of DSRC-equipped vehicles who may wish to take advantage to the potential safety from broadside collision that an all-red interval may provide.

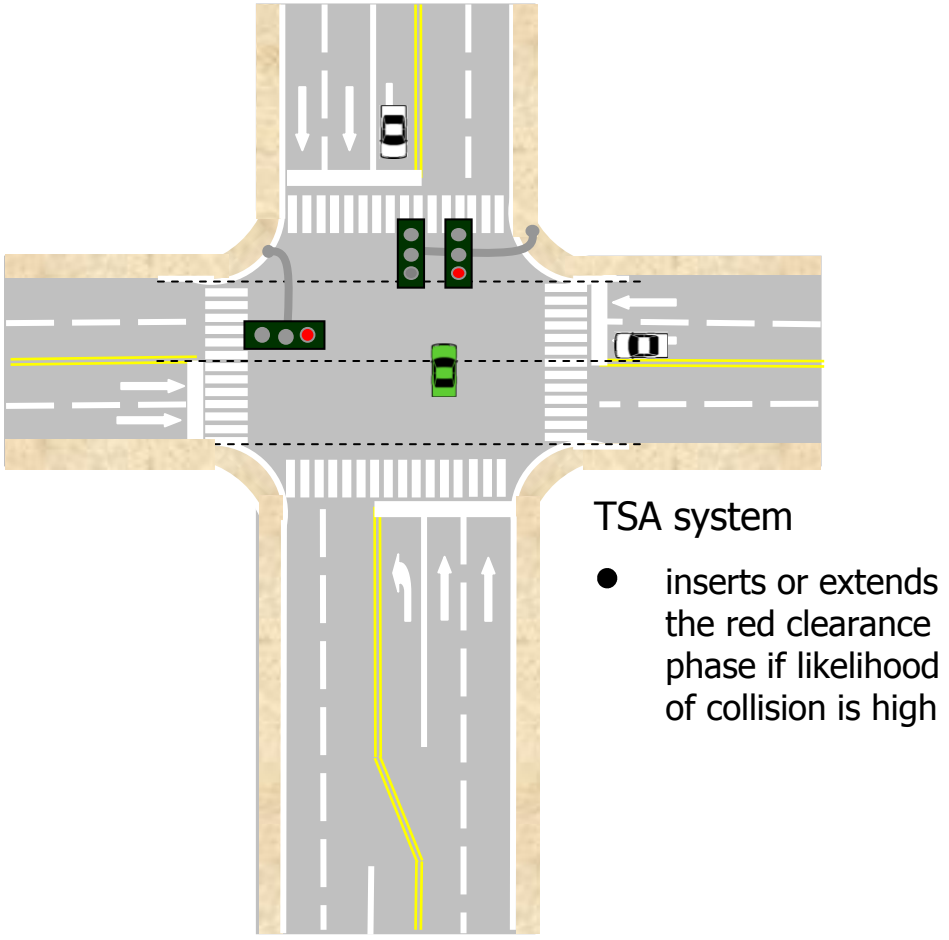

Figure 3. Illustration of CICAS-TSA ConOps

Outside of developing and revamping a VII-focused CICAS-SLTA (and TSA) ConOps, recent progress with CICAS-SLTA includes:

- Establishment of a "sniffer" working with a 170-type controller (and conceivably with any controller), combined with a message set, that provides DSRC signal state information to approaching, equipped cars (Page Mill Rd and SR 82, El Camino Real).

- Installation of an extensible 2070-type controller interface (adaptable to NTCIP, AB 3418 or other protocols) to provide signal state information via DSRC link directly from controller-to-computer-to-radio roadside equipment, benefiting CICAS-V as well as future aspects of CICAS-SLTA (5^{th} Av and SR 82).
- Close work with the standards communities to establish CICAS ConOps and message set definitions that work for the ongoing research and the potential larger research and deployment communities. In the application domain, this includes SAE J2735 and ISO TCP WG 14 Working Group. In the radio domain, this includes contributions to the IEEE 802.11p and 1604.2-3 standards. With SAE J2735, this include discussion, testing and facilitation of the CICAS-V-developed Geometric Intersection Description (GID) and Signal Phase and Timing (SPAT) messages.
- Decision and downselection of criteria, siting and instrumentation to begin more detailed observational data. Two intersections with different characteristics have been chosen for initial instrumentation, and an experimental protocol and accompanying intersection instrumentation suite procurement decision has been completed. We expect to use long-range semi-portable roadside radar installation for approach leg trajectory data collection and very portable inside-the-intersection box laser radar – both purchased from European vendors – for our near-term data collection.

The work reported here extends prior research work to provide the additional breadth and depth of work needed in preparation for a potential follow-on Field Operational Test (FOT) effort. One of the enabling technologies behind CICAS-SLTA is understood to be DSRC, which will be delivered as part of VII. Abiding by our ConOps, the research is structured with the assumption that DSRC transceivers will be available at intersections (as "Roadside Equipment" or RSE) and onboard (as "Onboard Equipment" or OBE), and that signal phase and timing, other dynamic data and basic intersection maps will be available to the OBE. Therefore, one basic premise is that the subject vehicle will have DSRC communication with the aforementioned information, supporting a driver-vehicle interface. Another is that a variety of sources, to include VII and also roadside sensing, will provide information from additional road users whose state will affect CICAS.

In the end, we have conducted research to transform US DOT's VII concept of DSRC transceivers located on the roadside and in vehicles to enable high reliability, low latency communication of intersection safety messages, in an important set of intersection safety countermeasures.

SECTION 1. SIGNALIZED LEFT TURN ASSIST

We structure our report of the Signalized Left Turn Assist by introducing our Concept of Operations, which provides an overview and conceptual framework. This is divided into sections of ConOps Scope, System Users and Stakeholders and Operational Scenarios.

We follow with a detailed discussion of the main thrust of our effort, which was to provide data for and models for the crucial aspect of a SLTA alert or advice algorithm: observations and subsequent gap assessment models. Our work in this area is essentially divided into two portions. There is a "PATH portion" of four sections: Description of Data Sets, Interfaces, Analysis Models, then Results (*Characteristics of the Predicted Post-Encroachment Metric*). There is a section which we subcontracted to the University of Michigan Transportation Research Institute (UMTRI) or the "UMTRI portion" (*Analysis of Left-Turn Gap Acceptance at Signalized Intersections Using Naturalistic Driving Data*). The "PATH Portion" was based upon mining and interpreting data collected at the PATH Richmond Field Station Intelligent Intersection through the course of several years, and the "UMTRI Portion" was similarly based upon mining and interpreting data, theirs however from antecedent Road Departure Collision Warning (RDCW) Field Operational Test data.

We next cover the CICAS-SLTA field test data collection, then we focus on detecting the primary other road user: the pedestrian. This section is followed by a description of the SAE J2735 interaction (*Messages and Maps*). We end the SLTA report by describing the upgrades undertaken with the Richmond Field Station Intelligent Intersection.

1.1 CONCEPT OF OPERATIONS

Scope

This concept of operations (ConOps) addresses the Cooperative Intersection Collision Avoidance System - Signalized Left Turn Assistance (CICAS-SLTA). The ConOps serves as a foundation document, describing CICAS-SLTA for the system to be uniformly understood by researchers, project participants and stakeholders alike. It serves three functions: (i) as an outreach tool, (ii) as a 'living document' to identify and capture researchers and stakeholders (sponsors, deployers and users) interests and concerns, and (iii) from a project systems engineering perspective, as input to the sequence of steps necessary to conduct the project. The position of the ConOps is shown in the context of the entire process in the following sequence (with associated project tasks given parenthetically):

- A detailed ConOps (Task 3 *System Design*)
- System Requirements (Task 2 *Research*)
- System Design (Task 2 *Research* & Task 3 *System Design*)

- Implementation and Testing (Task 2 *Research*, Task 3 *System Design* and Task 4 *System Development and Prototyping*)
- System Verification (Task 4 *System Development and Prototyping*)
- Field Operational Test (FOT) (Task 5 *FOT Planning* and Task 6 *FOT*)

To be clear, the ConOps provides a detailed narrative description of the CICAS-SLTA System, not a narrative of the program or research plan. Moreover, the ConOps does not describe requirements, or the testing or verification plan. The ConOps, however, develops the foundation for these. Following the definition of scope, purpose and goals and objectives, the bulk of the ConOps does so with the following organization:

- System Definition
- Assumptions and Constraints
- Operational Description (per stakeholder category)
- Operational Scenarios (for normal as well as difficult or failure scenarios)

Purpose of CICAS-SLTA

The primary purpose of the CICAS-SLTA countermeasure system is to provide support to the driver to avoid collisions during permissive left turn movements with oncoming traffic. The particular crossing crash type addressed by the system is classified as Left Turn Across Path/Opposite Direction (LTAP/OD), and the conflicting movement is depicted in Figure 4. The LTAP/OD crash type accounts for 27% of intersection crashes (Ragland & Zabyshny, 2003). The CICAS-SLTA countermeasure addresses vehicle-vehicle head-on, sideswipe crashes and potential crashes from conflicts of left-turning vehicles with other roadway users such as pedestrians and bicyclists.

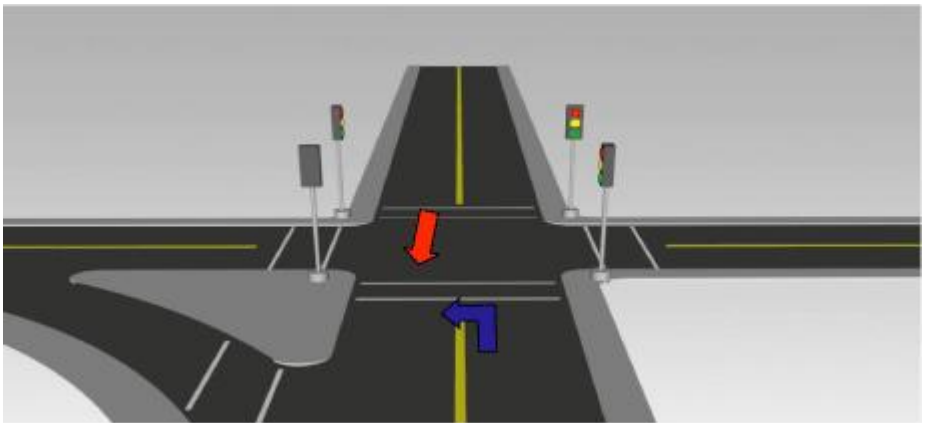

Figure 4. The LTAP/OD Scenario. The blue Arrow Represents Subject Vehicle, and red Arrow Represents Principal Other Vehicle.

CICAS-SLTA provides a way of addressing left turn safety problems at intersections without having to alter the signal phasing, such as replacing a permissive left by a protected left turn. Replacing a permissive left turn with a protected left turn typically leads to traffic flow reductions and might not be doable in some circumstances. If the CICAS alert criteria are properly tuned to help drivers avoid unsafe encounters, they may even increase intersection capacity by helping drivers better distinguish acceptable turning gaps, reducing the number of adequate gaps that are rejected.

From this perspective, it offers one more option for the traffic engineer to consider when an intersection has left turn safety problems.

Goals and Objectives

Goals:

- Reduce frequency and severity of LTAP/OD crashes at intersections.
- Reduce stress for drivers trying to make unprotected left turns.
- Harmonize gap acceptance behaviors across drivers, so that more drivers choose gaps that promote both safety and efficiency.
- Facilitate traffic flow by providing most of the safety advantages of a protected left turn phase, without incurring the interruption of opposing traffic.

Objectives:

- Address the sources of LTAP-OD conflicts by:
    - improving drivers' assessments of the traffic situation by providing them with better information about the presence and approach speed of opposing traffic,
    - discouraging drivers from making left turns with inadequate gaps in the opposing traffic.
- Alert drivers to the presence of other road users such as those in the crosswalks on the destination leg, so that they are less likely to endanger pedestrians or bicyclists, and less likely to get stranded in the path of oncoming traffic while waiting for these users.
- Provide the left-turning drivers with better information on whether or not an oncoming vehicle is going to stop during a signal phase transition, and continuing to alert the left-turning driver if the oncoming vehicle may violate the red shortly after that phase transition.
- Reduce the number of safe gaps in opposing traffic that are rejected by left turning drivers.

System Definition

Description

CICAS-SLTA is a system that helps drivers make unprotected left turns at signalized intersections in the face of approaching traffic by enhancing their awareness of the presence, location and speed of the approaching vehicles, as well as other road users who may come in conflict with the left-turning drivers. The system is intended to help avoid LTAP/OD crashes that might otherwise occur for a variety of reasons that can be summarized in the following causal behaviors:

- The Subject Vehicle (SV), or left-turning driver may not have perceived the approaching Principal Other Vehicle (POV) because of inattentiveness or intersection geometric features, occlusions blocking the field of view, or other perceptual or conspicuity issues;

- If the SV driver perceived the approaching POV, the SV driver may have misjudged its distance or approach speed, thereby underestimating the closeness of the turning conflict; or

- The SV driver may have been uncertain about how to judge the closeness of the encounter, so CICAS-SLTA could potentially improve his comfort and confidence in making a turning decision.

- Additionally, the SV driver may be unaware of the presence of pedestrians until he/she has committed to making a left turn.

The CICAS-SLTA system includes the following elements:

- Detection subsystems to identify the locations and speeds of the turning SVs and approaching POVs.

- Detection subsystems to identify the presence of pedestrians that could influence the turning movement of the SV.

- A computer hosting software to process the detection information, estimate the future trajectories of the vehicles, predict the closeness of their encounters, and determine if and when to display information to the SV drivers.

- A Driver-Vehicle Interface (DVI), which is a dynamic display to convey information to the SV drivers. The specific form and appearance of the display are TBD, based on the results of the research.

- An optional Driver-Infrastructure Interface (DII) that provides similar functionality to the DVI, but serves all vehicles approaching the intersection to make a left turn,

These elements are represented in Figure 5 below.

If the companion CICAS-Violation (CICAS-V) crossing path countermeasure is deployed, its DSRC-based communication and positioning could contribute to a single CICAS system that combines CICAS-V and CICAS-SLTA functions into one device with a common DVI. The information displayed by the DVI may include indicators of the distance and/or time to an impending conflict with a POV, indicators of the presence of pedestrians in the destination crosswalk, or a warning to deter particularly hazardous turning maneuvers, in addition to the CICAS-V imminent violation warning.

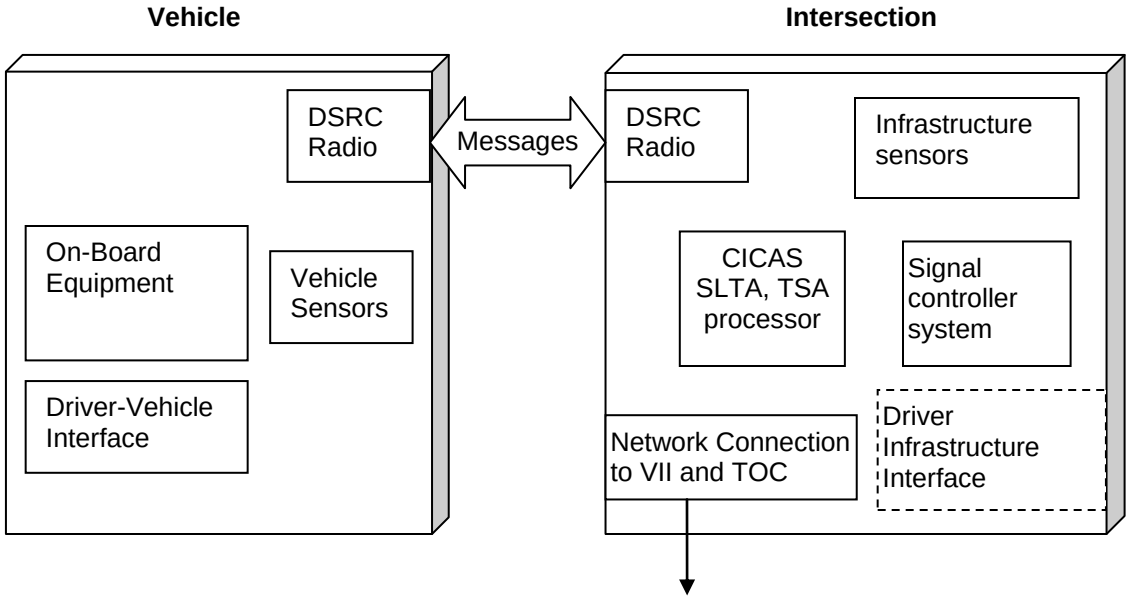

Figure 5. CICAS SLTA System

Modalities for Displaying Information to the Driver

The decision on display modalities and the possibility to switch between modalities will be based on the results of data collection, analysis and simulation work in the project. Based on the results of this research, we will develop and test messages to display to the driver. We expect that the display will provide different information based on the threat level and the threat source.

Architectural Alternatives

CICAS provides an opportunity to make best use of the strengths of infrastructure and vehicle-based approaches to detecting and avoiding intersection conflicts. The keys to defining an effective CICAS-SLTA ConOps are maximizing flexibility and incorporating the ability to address the full range of intersection conflicts, including traffic control device violations and turning drivers' gap estimation errors.

Flexibility is important because of the wide diversity of intersection characteristics that are relevant to the behavior of the system, such as:

- Urban, suburban or rural driving environments
    - Pedestrian and bicyclist density
    - Traffic speed and density
- Traffic volumes, ranging from very high to very low
- Wide variations in legacy traffic control infrastructure
- Geometric design
    - Number of lanes
    - Visibility of approaching traffic

It is also important to accommodate changes over time as intersection traffic conditions and the underlying intersection infrastructure change, providing smooth paths for growth in capabilities. This means that the CICAS-SLTA system should accommodate the needs of that intersection throughout a logical growth path from stand-alone fixed-time signal control, through actuated and semi-actuated signal control and corridor- or area-wide coordinated signal controls. CICAS-SLTA could also benefit from the implementation of additional sensing capabilities for identifying specific intersection conflicts, one example being real-time adaptive signal control, with the signal phase changes based on detected motions of platoons of vehicles on the various approaches. The equipped intersections could also operate with varying degrees of cooperation with and from the vehicles, based on the market penetration of VII-capable vehicles.

Infrastructure-based systems gain the following advantages for providing intersection collision avoidance support:

- Infrastructure-based displays can provide information to all drivers approaching the intersection, not just the drivers with the most capable vehicles;

- The center of the intersection has line-of-sight contact with all approach legs for purposes of sensing and wireless communications with approaching vehicles;

- By combining sensor or detector data with information from the traffic signal controller and information communicated from vehicles, the intersection infrastructure can know the complete state map of the intersection; and

- Investments in infrastructure installations can be prioritized based on the safety record of each intersection and its existing infrastructure, so that the first deployments can be targeted where they are likely to produce the maximum benefits.

Vehicle-based systems benefit from a different set of natural advantages:

- A vehicle-based system can know the state of the vehicle (location, speed, acceleration and fault conditions) more accurately and earlier than an infrastructure-based sensor would be able to detect it;

- Vehicle-based alerts to drivers can be more salient than infrastructure-based alerts (using audible and haptic information channels as well as visual);

- Vehicle-based alerts can be integrated with other in-vehicle safety systems to save costs and to optimize driver workload demands under stressful conditions; and

- If systems are introduced on a significant fraction of the new vehicles each year, it is possible to reach broad market penetration faster than by relying entirely on infrastructure systems.

Furthermore, future vehicle-based systems may take advantage of the following potential features:

- Driver patterns of past driving behavior and intentions (e.g., turn signals) can be detected and used to enhance the accuracy of conflict predictions; and

- Vehicle-based alerts can be tailored to individual driver behavior and preferences, so that a single alert criterion does not need to satisfy the broad diversity of the driving population as a whole.

These strengths are largely complementary, so that combinations of both vehicle and infrastructure elements should make it possible to expand the performance envelope of CICAS beyond what would be possible if vehicles and infrastructure were not coordinated with each other.

It is therefore possible to define an inclusive CICAS concept of operations that can address both violation and turning conflicts with varying allocations of responsibility between the vehicle and infrastructure elements. No reasonable alternatives are precluded, and flexibility is provided to enable full interoperability among vehicles and intersections with varying levels of capability.

Assumptions and Constraints

- CICAS-SLTA is assumed to rely on infrastructure-based sensors to detect trajectories of approaching vehicles. It is likely that no single sensor will be able to provide all the needed measurements, but the specific selection of sensors is TBD, based on results of the research. Over time, as VII market penetration increases, it should become possible to augment this information with information from in-vehicle sensors, communicated to the intersection infrastructure using DSRC communications.

- Information is provided to drivers through a driver-vehicle interface (DVI), which is a dynamic display in the vehicle, relying on threat information communicated from the intersection using DSRC. The exact combination of display modalities for interacting with the driver will be determined as part of the research.

- The display will alert SV drivers to the presence of approaching POVs that represent potential threats to the turning SV and to the approximate severity of the threat they represent. It may also warn SV drivers against making a turn in particularly hazardous conditions (e.g. imminent crash).

- The display will alert SV drivers to the presence of vulnerable road users, such as pedestrians and bicyclists in their path.

- The display will not indicate to drivers that it is safe to turn or that they should proceed into their turn, but will leave the turning initiative to the driver.

- The DVI and the logic that governs its contents must be designed to be suitable for use by a large fraction of the driving population and will therefore be a compromise. The logic (and/or parameters) will be adjustable based on the class of the subject vehicle and its performance, and it may also be tunable to the preferences and capabilities of individual drivers.

- CICAS-SLTA should be able to operate under normal driving conditions (large range of weather and illumination).

- CICAS-SLTA is intended for use at intersections with permissive left turns, but not for protected left turns, which would obviate the need for the system.

- CICAS-SLTA should be applicable for the full range of traffic speeds and densities that would be expected at intersections with permissive left turns.

- CICAS-SLTA parameter values (e.g., alert timing, alert threshold values) will not be uniform across all intersections, but will be adjusted to match the needs of individual intersections, based on their geometry, signal phasing, and traffic volumes and speeds.

System Users and Stakeholders

The CICAS-SLTA system operation and its impacts can be viewed from microscopic through macroscopic perspectives, progressing from the direct encounter with the system by the road user, through the operation of the system by the infrastructure operator and up to the State or Federal governments, who may see a public safety impact. We will distinguish the system end users, or users who will interact directly with the system and the general road users who will benefit from its deployment, from the operators who will be involved in the system selection and operation. The system is described from each identified user group's perspective.

Identification and Description

Intersection Users

The primary system user will be the driver intending to make a left turn at the intersection. The driver will view the system through the progression of alerts described in Section 3. He or she will be primarily concerned with usability, in the sense of providing drivers with a tool that they will deem effective, efficient and satisfactory to employ; accuracy and consistency, which relates directly to the degree of false alarms, which must be low.

It is important to note that 'vehicle drivers' describes a heterogeneous population, including individuals of any age and skill. Older drivers are somewhat over-represented in crossing path crashes at intersections. In fact, drivers age 65 and older represented 11% of crossing path crashes compared to 6.4% of non-crossing path crashes[1]. It is important that all classes of users accept the system and that the needs of older drivers be considered during the design of the system.

Secondary users of the system are the drivers of the POV and vulnerable road users who will benefit from the SV drivers' use of the system. In terms of POVs, motorcyclists are over represented in the LTAP/OD crash statistics, presumably because they are less conspicuous because of their smaller visual cross section. Pedestrians and bicyclists constitute a special class of other vulnerable road users who are often overlooked when they attempt to cross the SV's destination leg of the intersection. Public acceptance and adoption of CICAS-SLTA may be heavily influenced by this group's perceptions of the system. If the implicit "go ahead" given to the SV by the system does not reliably account for vulnerable road users, then it may not be accepted by these other road users.

Intersection Operators

State DOTs and local jurisdictions will use the CICAS-SLTA primarily for increasing the safety of their intersections. Secondarily, they will support the burden of installation, operation and maintenance, to include training of personnel for this equipment. They will also want to preserve and eventually improve capacity and will want the CICAS-SLTA system to provide demonstrable safety benefits (with little downside or tradeoff in, for example, increased rear-end crashes in exchange for reduced crossing path crashes), plus no reduction in intersection throughput.

The operator will more likely accept some safety benefit if other operational benefits such as increased capacity could also be realized. It will also likely accept the CICAS-SLTA system if it does not substantially change or revamp legacy systems and operations. In the end, standards and associated warranties, such as the one provided by the MUTCD or as adopted in California by the California Traffic Control Devices Committee (CTCDC), remain the reference for the practitioner.

Other Stakeholders

Intersection Equipment Suppliers

These suppliers will have to integrate the CICAS-SLTA equipment suite in a form that will be well understood by the traffic control device industry and interface easily with their current products so that it can be offered to public agency customers at modest cost and risk.

Their roles are expected to include:

- Participation in the development of standards and certification procedures
- Training of their sales and maintenance staffs and customers' staffs about the system
- Marketing to their public agency customers
- Development of reliable, cost-effective hardware and software implementations, compatible with existing products.

Vehicle Manufacturers and Suppliers

Any acceptable, deployable system that has demonstrated public safety benefits, particularly if CICAS-SLTA functions can be implemented using common components with other in-vehicle systems, will be more likely to succeed, e.g., by combination with CICAS-V or VII use cases (such as dynamic in-vehicle signage).

From this perspective, their roles include:

- Participation in the development of standards and certification procedures.
- Training of their sales and maintenance staffs about the system.
- Development of reliable, cost-effective hardware and software implementations, for installation in new vehicles
- Marketing the system to consumers (vehicle purchasers).

US DOT

Any acceptable, deployable system that has demonstrated public safety benefits, particularly those that show dramatic benefits compared to costs and are accepted by the aforementioned stakeholders, will be of interest, which is why US DOT is currently keenly interested in CICAS-SLTA. In support of CICAS-SLTA implementation, US DOT will:

- Sponsor development, field testing and evaluation of the effectiveness of the system

- Participate in the development of standards and certification procedures.

- Assist state and local agencies in the installation, operation and maintenance of CICAS-SLTA system by developing guidelines.

- Develop training materials and courses related to CICAS-SLTA system installation, operation and maintenance

- Develop tools supporting the identification of intersections benefiting from CICAS-SLTA technology, design of CICAS-SLTA system at actual intersections and monitoring performance of the system.

MUTCD and State Standard Committees (e.g., CTCDC)

The MUTCD will provide guidance and guidelines on the appearance of the optional DII.

Interactions
The interactions between users and other stakeholders are illustrated below in Figure 6. These interactions are from the viewpoint of a deployed system, and as such does not represent the implicit role of the research underlying the development of the system.

At this time, we expect that the equipment developers and suppliers (traffic control infrastructure and vehicle industries) will be offering their products to state DOTs and local jurisdictions and that they will both follow relevant recommendations and guidelines. In turn, the product will be deployed by State DOTs and local jurisdictions in order to improve the safety of the end users. We also expect that the end users will demand the in-vehicle systems from the automotive OEMs.

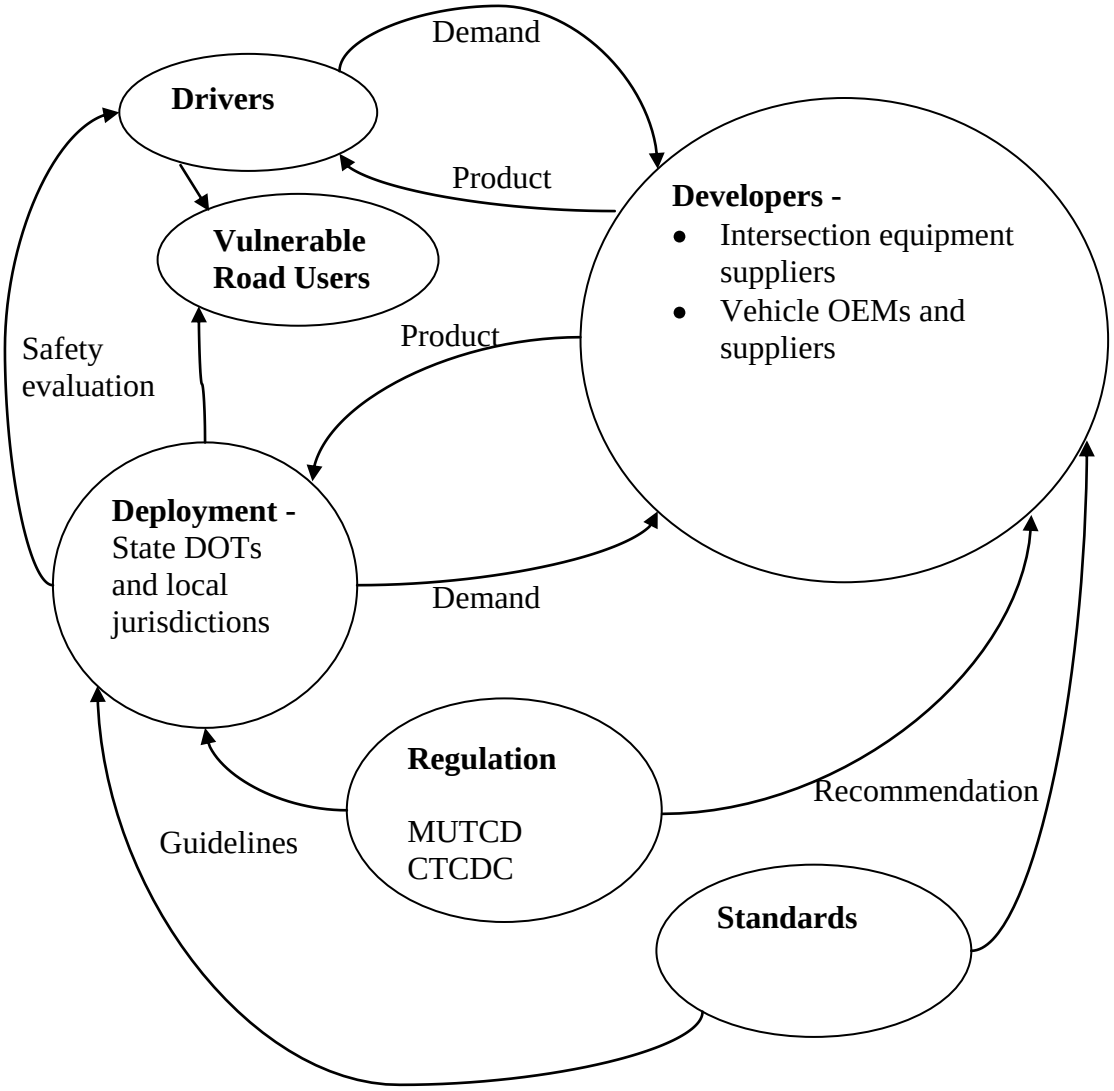

Figure 6. Interactions among Stakeholders

Issues

Cost/Benefit, Effectiveness, Marketability: CICAS/SLTA needs to be perceived favorably by the intersection users and operators and other stakeholders on both the infrastructure and vehicle sides in order to be deployed. This means that it must be able to satisfy a variety of concerns:

- Perceived by local traffic engineers to be a cost-effective strategy compared to other alternatives for addressing intersection safety and traffic problems;

- Perceived to be safe and effective by the general public, including drivers, vulnerable road users and community activists;

- Perceived by the traffic supply industry to be marketable to public agencies based on its costs and performance;

- Considered to be a desirable feature by vehicle purchasers;

- Perceived by the automotive industry to be sufficiently marketable that they want to offer it as a product.

Reliability and Litigation: This issue will concern all of the users and other stakeholders at different levels:

- The end user, or SV driver, could be negatively affected by responding to the system, e.g. the vehicle following the SV does not expect the SV driver to stop and rear-ends the SV.

- The intersection equipment supplier, intersection operator and vehicle supplier (manufacturer and/or vendor) will have liability exposure if end users do not understand well the range of capabilities of the SLTA system, in which case they could not integrate the information provided and therefore end up in a crash or overestimate the capability of the system.

- The intersection equipment supplier, intersection operator and vehicle supplier could be involved in case of failures.

Installation and Maintenance: CICAS/SLTA should be no more difficult to install or maintain than any other systems that are already accepted in the infrastructure and vehicle industries. Since the vehicle-side equipment is likely to differ little from other existing in-vehicle components (radio, display and processor), primary attention needs to be focused on the infrastructure, to address issues such as ensuring that:

- The needed detectors can be mounted within existing infrastructure constraints, rather than requiring significantly new mounting fixtures;

- Installation can be accomplished quickly, without requiring extensive labor, and with a minimum of lane closure time;

- The detectors can operate for a long time without requiring manual recalibration, adjustment or realignment;

- The computational requirements can be met by existing traffic signal controllers or new devices that can be readily installed within existing traffic signal control cabinets;

- All components are able to meet the existing environmental requirements for field devices.

CICAS Operational Scenarios

This section is composed of three subsections. In the first one, the baseline, variation to the baseline scenarios and extension scenarios are described. The interactions with CICAS -Traffic Signal Adaptation (TSA) are also discussed in this subsection. The second subsection introduces the more complex scenarios, and finally the third section presents the system failure scenarios.

Operational Scenarios for Normal Conditions

Baseline Scenario

The CICAS SLTA system senses two vehicles approaching an intersection from opposite directions during a green phase. The driver of one of the vehicles intends to turn left, and that vehicle is considered the SV. The vehicle on the opposite leg continues to approach in one of the through lanes and is considered a POV. If the SV driver were to turn without stopping at the intersection, these two vehicles may be on a collision path.

The CICAS SLTA system will provide the SV driver information about the approach of the POV. At the same time, the system will evaluate the likelihood of a collision between the two vehicles' projected trajectories. Two cases might then result:

1. The SV driver integrates the information in the decision process stage and either slows down to let the POV clear or stops in the middle of the intersection if necessary to let the POV(s) clear the path.

2. The SV driver decides he/she can beat the POV. The SLTA system determines that there is a high likelihood of a collision by a time pre-defined as the warning onset, and then triggers a warning. The SV driver then either slows down or stops to let the POV clear the intersection.

For cases when the SV driver stops in the intersection, the SLTA system will continue to deliver information to the driver about all categories of potential hazards to turning movements (POVs, bicyclists or pedestrians), regardless of SV driver actions.

Figure 7 illustrates the processing steps that the system will conduct as the SV approaches the intersection and after it has stopped in the intersection. The display of CICAS SLTA information will be provided by a DVI, and may also be provided by an optional DII at some intersections.

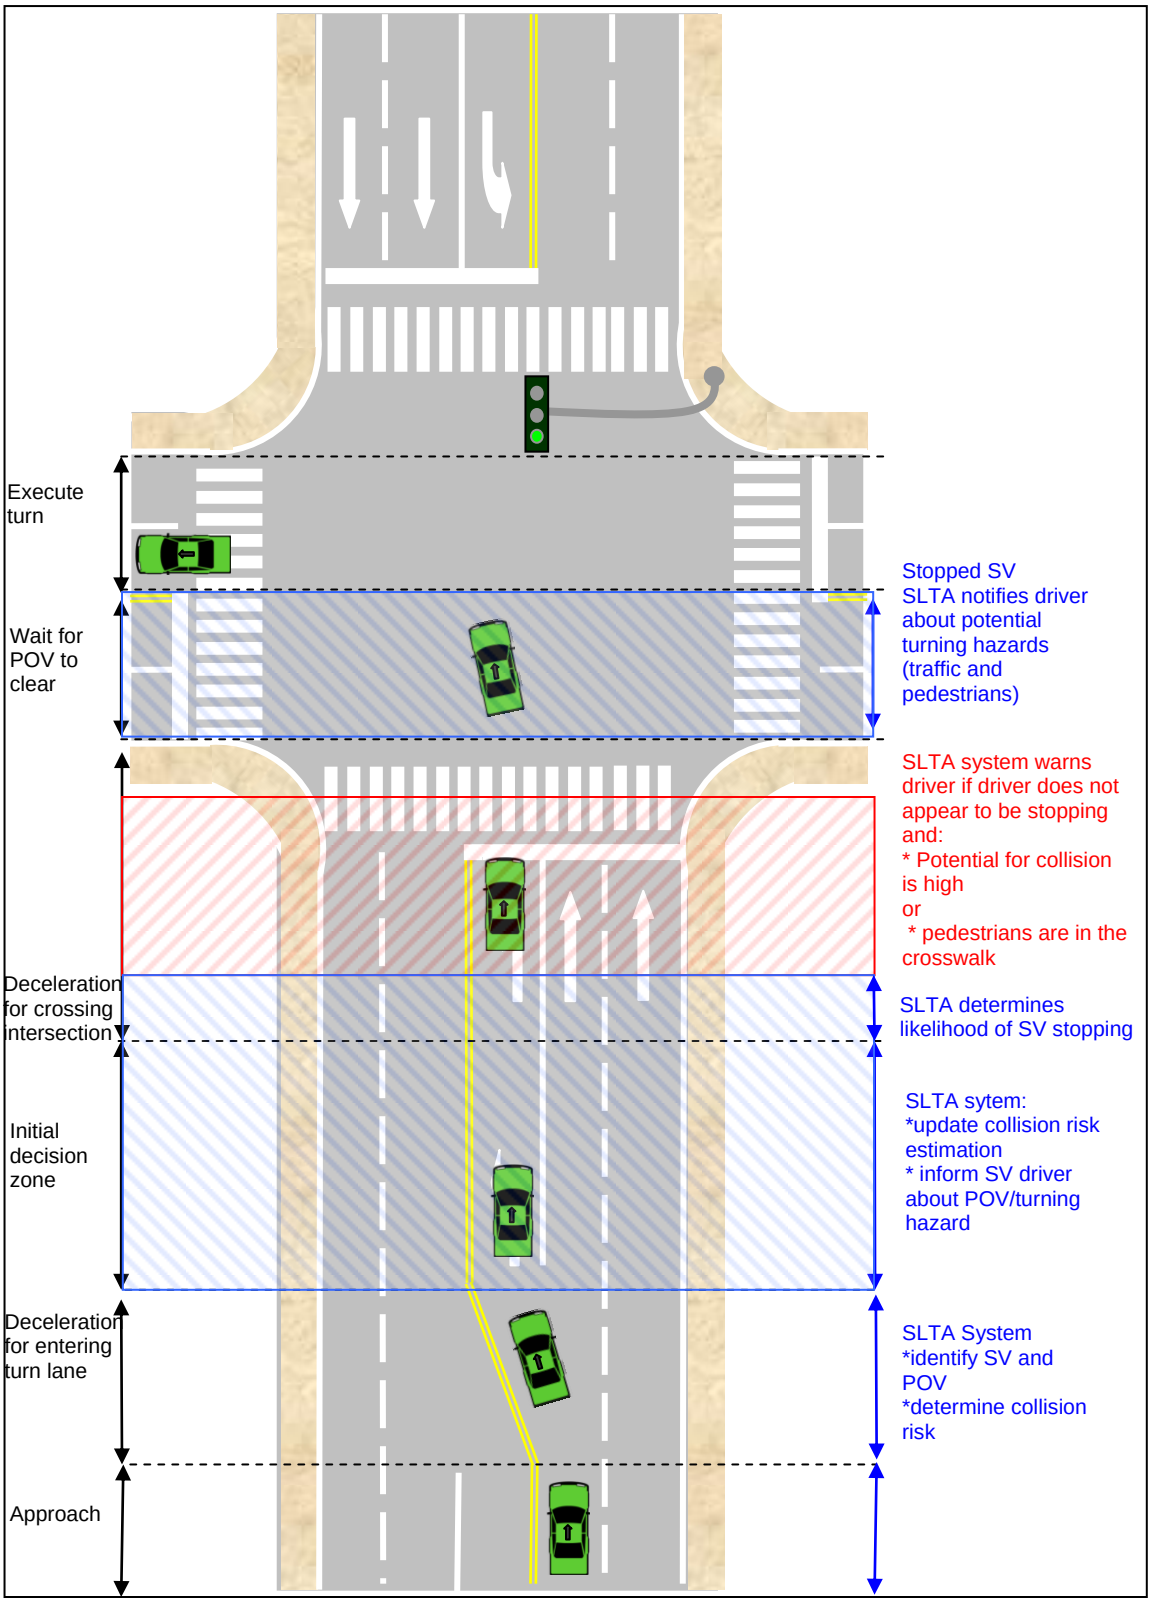

Figure 7. The SLTA Baseline Scenario (not drawn to scale)

Variations of the Baseline Scenario

Intersection without a left turn pocket:

The behavior of the system will be similar to its behavior in the baseline scenario described above. The main difference is that the system will have to resolve the uncertainty about whether the SV driver is planning to make a left turn in order to determine whether to display CICAS-SLTA information, based on identification of directional signal status and steering and braking actions.

Intersection with multiple turning lanes:

In some locations, a permissive left turn may be implemented at an intersection with multiple turning lanes. One of the lanes may be a shared through and turn lane, in which case its vehicles would be handled the same as in the case without a left turn pocket. The system behavior for multiple turning lanes will be similar to the baseline case, except that the zone that the system will have to consider for the determination of conflicts will be different, as it will include the multiple lanes of the SV approach.

Uncontrolled Intersections:

The absence of control at the intersection provides another variation. In this situation, the system will proceed with a similar identification of the SV and POV potential conflict as in the baseline. The main difference is that there is no traffic control that would constrain behavior.

Arrival on traffic signal phase transition:

For situations when the SV reaches the intersection and the traffic signal is transitioning from green to amber or from amber to red, the logic of the system will rely first on the trajectories of the vehicles, and then on the traffic signal status in order to determine what message to display to the drivers.

Stopped on red:

For situations when the SV is waiting at a red light, there will be no possibility to predict a potential collision, as the SV and POV are stationary when the signal transitions to green and the SV driver behavior will be difficult to predict from the measurements that are available (e.g., is the driver creeping to the middle of the intersection or going to turn in front of a POV?). Hence, the CICAS system will display potential hazards to the driver after the signal turns green (see Figure 8 below).

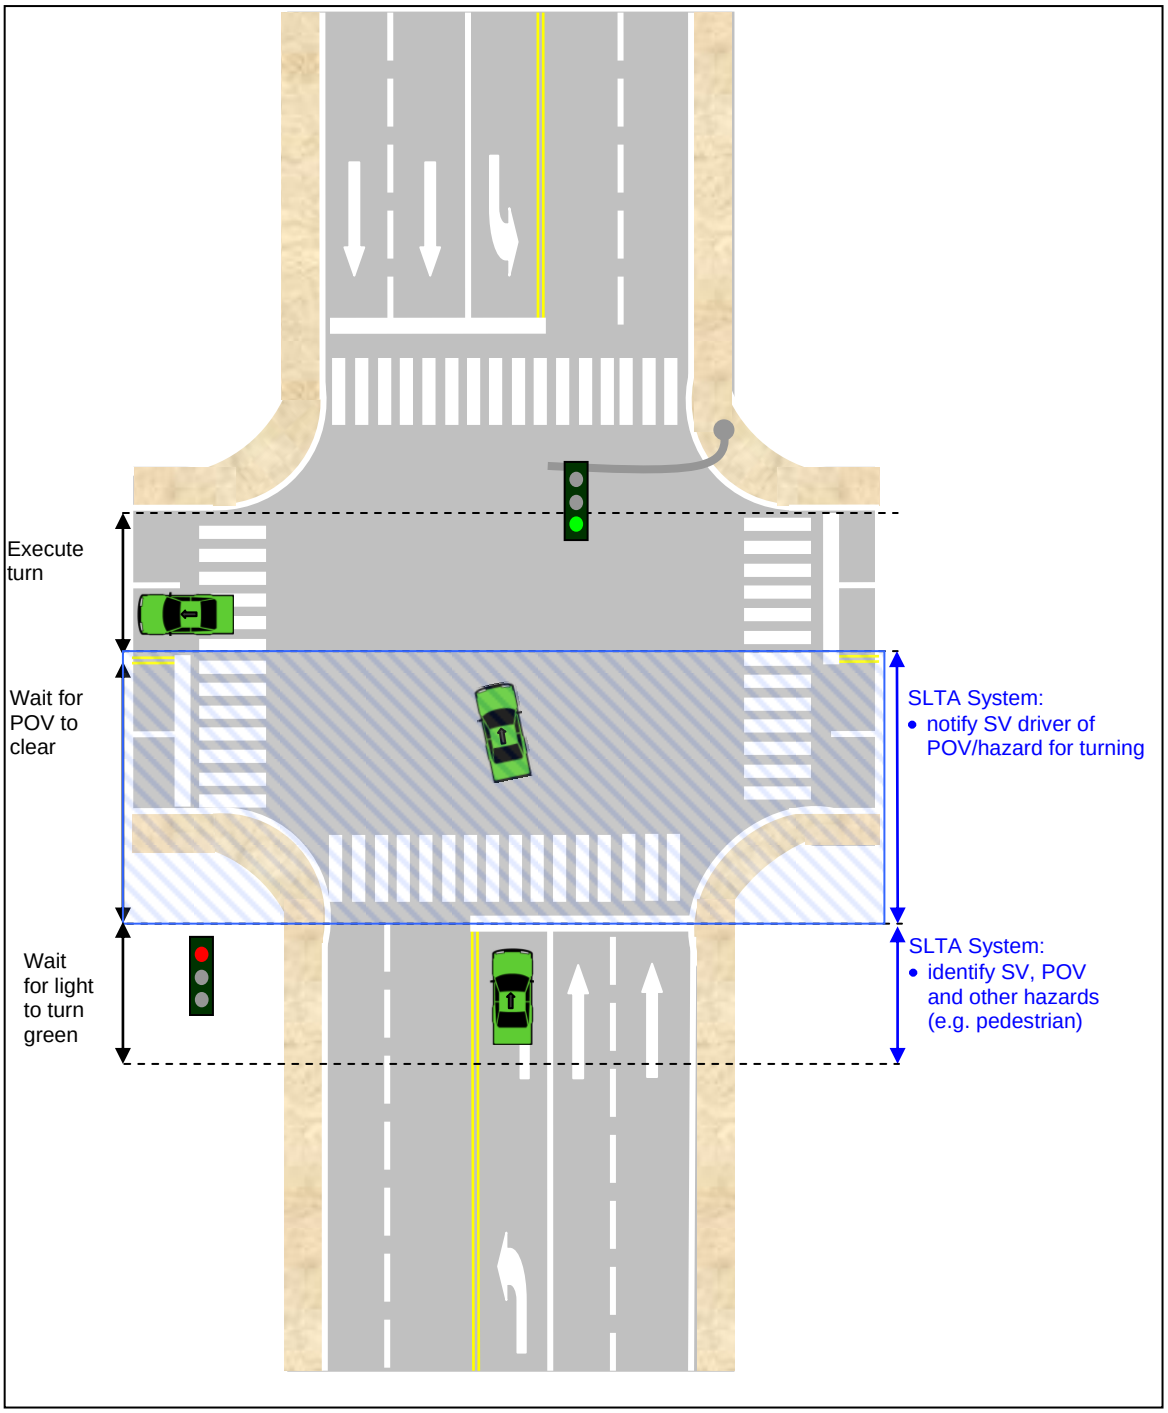

Figure 8. SLTA Waiting on Red

Other road users:

The other road users that the system will consider are pedestrians and bicyclists. If either of these two hazards is detected within a range of concern, then the SV driver will be informed of the presence of this hazard. Alternative methods of informing the driver will be evaluated in the CICAS-SLTA research work to determine the most effective approach (for example, deciding

whether it is better to explicitly inform the driver of the presence of vulnerable road users or to incorporate this presence into the general advice about turning hazards).

Combination of POV and the vulnerable road user:

A typical scenario is a case in which the presence of either a pedestrian or bicyclist triggers an alert by the SLTA system, leading the SV driver to postpone the turning maneuver. As a result of the SV driver's decision to stop or slow down, an incoming POV becomes a threat to the left turning maneuver. In that situation, the SV driver will be informed about the additional threat and that the turning maneuver is still not recommended.

Extension Scenarios Enabled by Future Applications of New Technology

1. Vehicles communicate their classes, speeds and locations to the intersection so that the intersection can combine this information with the infrastructure-based detector data make better estimates of potential conflicts and of the needs of the approaching vehicles' drivers.
2. Vehicles display customized information to drivers, adapted to vehicle and/or driver characteristics, so that information can be customized beyond the "compromise" alerts decided upon by the infrastructure system.
3. CICAS-SLTA functions are implemented at uncontrolled intersections with left-turn conflict and/or visibility problems, as a stand-alone system. In this case, there would be no pre-existing intersection signal infrastructure, so the entire infrastructure expense would have to be allocated to the CICAS-SLTA system.
4. The SV communicates its turn signal status to the intersection, so that the driver's intention to turn may be more clearly revealed at intersections where there is no left turn pocket.
5. Direct Vehicle to Vehicle (V2V) communication are enabled, allowing the SV and POV to directly communicate and potentially advise both drivers how to avoid the impending conflict.

CICAS – SLTA, CICAS-V and TSA Interactions

There are three cases in which the multiple CICAS systems will interact:

- The left turning vehicle (SV for the CICAS-SLTA system) is stranded in the intersection when the signal is about to turn red and another vehicle is going to run the red light from the opposite direction (SV for CICAS-V and for TSA). In this case, the driver of the turning vehicle would be informed of the threat by the SLTA system. The threat assessment would be determined by the CICAS processor at the intersection, which would supply this information to the SLTA system using the VII wireless link, but would simultaneously give the TSA advisory to the signal controller to extend the all-red interval. This would reduce the danger to the turning vehicle, which would have more time to clear the intersection before the cross traffic receives its green signal. The signal phase and timing (SPAT) message broadcast by the intersection could be augmented to include the increased risk of running of the signal due to the presence of the left turning vehicle, and the CICAS-V driver could potentially be notified of a higher threat than a signal violation (the threat of colliding with the turning vehicle).

- When the left turning vehicle is predicted to run a red light, it is the subject vehicle for receiving the in-vehicle CICAS-SLTA and CICAS-V alerts, as well as for the TSA implementation of all-red extension.
- The CICAS-TSA detectors and software are focused on identifying whether an approaching vehicle is stopping or proceeding through the intersection on a red signal. This information is very useful for the CICAS-SLTA alert algorithms to know about the approaching vehicles that could threaten a left turning vehicle. When CICAS-TSA determines whether the approaching vehicle is stopping or running the red light, that information can be provided to CICAS-SLTA to modify its threat assessment accordingly.

The CICAS-TSA, CICAS-V and CICAS-SLTA systems will be designed to work collaboratively; therefore, no contradictory information will be given to drivers. We will carefully examine how the systems should interact and how the information flow regarding the signal phase change conditions will be integrated with the SLTA system logic.

Operational Scenarios for More Complex Situations

The identification of operational scenarios for more complex situations will be based on the data analysis of odd cases and simulation of these cases to evaluate how they impact the behavior of the system. Based on the results of this research, we expect to be able to i) describe the characteristics of the complex scenarios and ii) suggest mediation methods. We expect that some of the elements that will contribute to complex situations are:

- vehicles changing speed at a higher rate than normal (stronger accelerations)
- high density of SV and POVs, with gap suitable for only one SV at a time

System Fault Scenarios

Faults cannot be entirely avoided with any system designed and implemented by humans. The designers and implementers need to anticipate the full range of faults that are possible and determine how to address them before the system is implemented. Well-established methods can be applied to detect the occurrence of faults, identify their specific nature, and respond automatically to mitigate their consequences. Application of these methods has to be specific to the design of the individual system, and cannot be specified meaningfully before the system is designed. The faults of greatest concern are those that are not (or cannot be) identified, because it is not possible to implement corrective actions to minimize their consequences.

CICAS - SLTA

1. Problems detecting approaching vehicles
    - inaccurate speed or position measurements
    - partial loss of detection (limited locations)
    - missed detection

- false detection
- temporary roadwork affecting the traffic flow and possibly occluding sensing systems at the intersection
- vandalism

2. Problems detecting vulnerable road users (pedestrians, bicyclists)
    - partial loss of detection (incomplete coverage)
    - missed detection
    - false detection
    - problems may be identified or not identified
3. Problems with infrastructure-based threat assessment
    - Corrupted sensor data
    - Incorrect calibrations
    - Computer malfunctions (may be consistent or intermittent)
    - Interface or connection problems with sensor inputs, traffic signal controller or wireless system
4. In-vehicle system malfunctions
    - Positioning system errors
    - no signal displayed to driver
    - false signal displayed to driver
    - problem may be identified or not identified
    - may or may not be possible to alert drivers to malfunction (default display?)
5. Wireless (VII/DSRC) communication system malfunctions
    - Loss of downlink capability from intersection to vehicles
    - Loss of uplink capability from vehicles to intersection

- Intermittent availability

6. Complete loss of power or system function
    - backup system to notify drivers of inoperability?
7. Inappropriate parameter settings at installation
    - may be difficult to identify

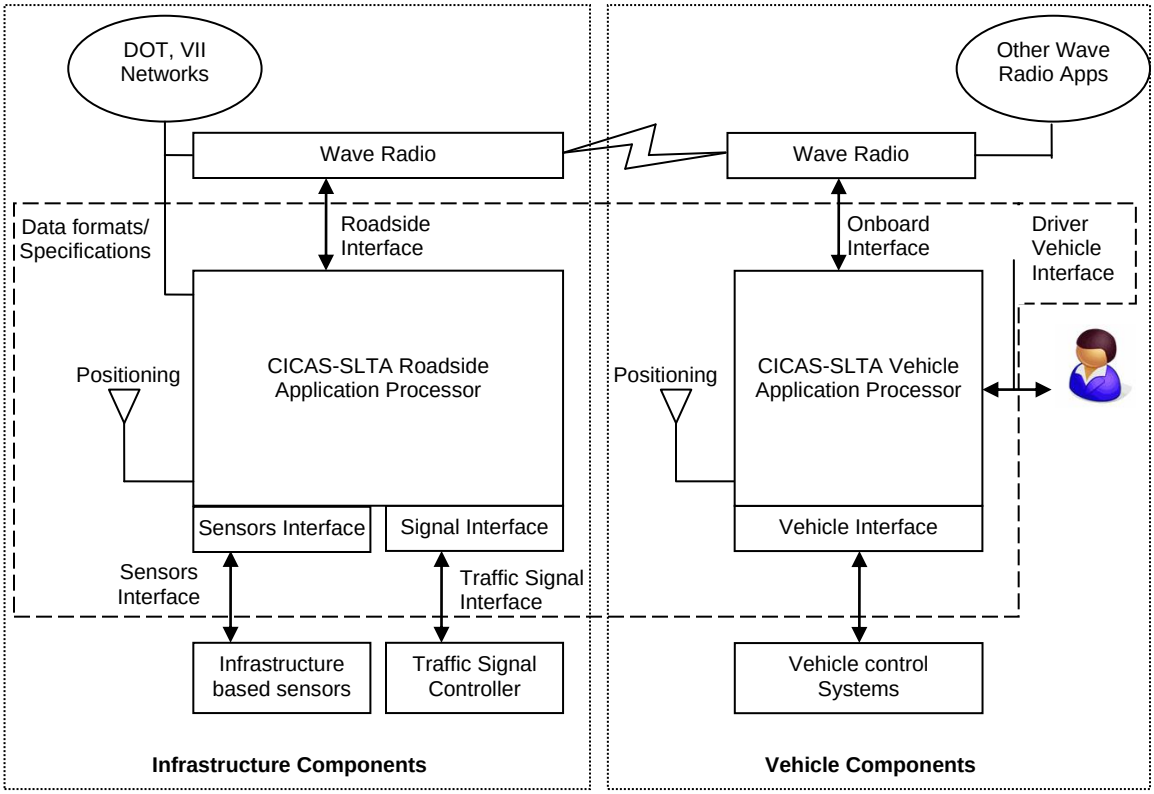
DOT, VII Networks
Other Wave Radio Apps
Wave Radio
Wave Radio
Data formats/ Specifications
Roadside Interface
Onboard Interface
Driver Vehicle Interface
Positioning
CICAS-SLTA Roadside Application Processor
Positioning
CICAS-SLTA Vehicle Application Processor
Sensors Interface
Signal Interface
Vehicle Interface
Sensors Interface
Traffic Signal Interface
Infrastructure based sensors
Traffic Signal Controller
Vehicle control Systems
Infrastructure Components
Vehicle Components

1.2 BASELINE OBSERVATION STUDIES AND GAP ASSESSMENT MODELS

1.2.1 PATH Portion

This part of the CICAS final report focuses on the data mining activities related to the CICAS-SLTA (Signalized Left Turn Assist) program segment. Two sets collected during prior projects were reanalyzed to address questions relating the driver interface and alert algorithm design.

The first research question addressed in was whether the decision support should be provided to drivers through an infrastructure or an in-vehicle display. While both types of interfaces have both advantages and disadvantages, no evidence was found to rule out the potential for either design.

The second research question addressed in this report was whether or not the driver assistance provided could take the form of a last-second warning. While the analyses in this report do not entirely rule out the possibility of providing last-second warning to a driver who intends to turn left in a less than ideal situation, the results were not promising. Additional metrics would need to be developed to detect the drivers intent to turn both quicker and more reliably.

The final research question addressed in this report was whether or not to continue to pursue the predicted post-encroachment time metric as the primary means for determining whether or not it is safe for the driver to turn in front of the oncoming vehicles. Although there is much fine tuning of the metric that is still possible, overall, the research showed that the predicted post-encroachment time metric provided greater predictive power than other, simpler models.

A series of three documents were written on the results of the California PATH IDS research program with a primarily focus on the creation of a suburban left-turn assistant. The first report (Chen, et al., 2005) was an interim report describing the initial research progress aimed at understanding the extent and causes of left-turn crashes. In that report the following topics were discussed:

- An analysis of frequency of the various intersections crash types in the state of California as compared to the frequency of those crash types in U.S. overall.

- The results of a literature review examining the human factors issues, or more specifically, driver errors, which have been proposed as potential root causes for intersection crashes.

- The results of the work to conceptualize, design, build, and demonstrate a prototype intersection collision avoidance countermeasure to address the Left Turn Across Path with Opposite Direction traffic (LTAP-OD) conflict.

The second report (Misener, et al., 2007) constituted the final IDS project report which described the overall systems engineering research performed during the project. In that report the following topics were discussed:

- A methodology and the results of several radar-based field observations

- A methodology and the results of several video-based field observations

- The results of a series of laboratory experiments to fine tune various parameters related to the DII looming effect

- The results of the effort to locate and test a number Commercial Off-The-Shelf (COTS) technologies and devices which could potentially be used to fulfill the sensing requirements needed by an IDS LTAP-OD countermeasure system

- The system architecture conceptualization, functional requirements, performance specifications, and algorithm design considerations for an IDS LTAP-OD countermeasure system

The third and final report (Bougler, Cody, and Nowakowski, 2008) was a supplement to the IDS final project report that described, in detail, the two human factors, instrumented-vehicle and instrument intersection, data collections conducted during the IDS project. These two data collections form the basis for the data mining activities detailed in this current report. In that report the following topics were discussed:

- First, the supplemental report covered the results of a field test experiment in which drivers drove an instrumented vehicle, out in real traffic, around a two-block radius, ten times, making a total of 40 left turns. This study was meant to help characterize normal, left-turn behavior while driving.

- Second, the supplemental report covered the results of a more controlled experiment conducted on a closed test track, where drivers were asked to make turning decisions based on the speed and distance of an oncoming vehicle, both with and without the aid of an infrastructure-based warning.

Based on these preliminary analyses of the field test data and the test track data, three peer-reviewed reports have also been written on these results at the time of this report. First, a report describing the preliminary analysis of typical left-turn maneuvers was published at the Transportation Research Board annual meeting (Cody, 2004), and follow-up, more focused on gap acceptance, appeared at the annual Driving Assessment conference (Cody, Nowakowski, and Bougler, 2007). Second, based on the initial analysis of the test track experiment, a report describing the potential benefits of an infrastructure-based, left-turn warning was published at the Human Factors and Ergonomics Society annual meeting (Nowakowski, 2006).

Current Research Goals

This report section documents a number of seemingly disparate data mining efforts which have progressed under the CICAS-SLTA project since the completion of the IDS project. Two data sets were combined and examined in further detail during these data mining efforts. This report describes the two data sets that were used and the results of the analyses that were conducted. The overall goal of the data mining effort was to learn as much as possible with the existing data before collecting new sets.

The data mining analyses described in this report cover a myriad of topics including the following:

- Initial insights comparing the effectiveness of a Driver-Infrastructure Interface (DII) versus a Driver-Vehicle Interface (DVI)

- An initial analysis into the accuracy and predictive power of the various components of the trailing buffer or predicted post-encroachment time metric

- An analysis of the characteristics of the trailing buffer or predicted-post encroachment time metric

- An analysis of the feasibility of a last-second collision warning in the context of a CICAS-SLTA system

Although the individual data mining efforts described in this report may seem to be less than strongly related, the overall themes or common threads among the analyses lies in the prediction of the Subject Vehicle's (SV) trajectory and the design of the interface to the SV driver. After all, the entire CICAS-SLTA problem boils down to two main tasks. First, the system must detect and predict the potential for an unsafe turn, and second, the system must effectively communicate that information to the driver.

This report describes the data sets that were used in the subsequent analyses, and it discusses the limitations of those data sets. It should be noted that there were also a number of additional data mining analyses were conceived and planned for this report, but they were unable to be completed. These planned analyses centered around replicating the types of analyses done during the field tests using the data collected during the more controlled experiments conducted at the Richmond Field Station. However, a key prerequisite for these planned analyses involved merging data sources. Unfortunately, in some cases, this was impossible because the synchronization across data sources was simply not feasible within a reasonable amount of effort.

The first analysis discussed in this report made use of a combination of two data sets that had been merged. The first data set was funded under the IDS project and consisted of 20 drivers who drove through the instrumented, test-track intersection at the Richmond Field Station during an experiment that tested the drivers reactions to left-turn warnings presented in the infrastructure (a DII). The second data set was collected under a project funded by the Toyota Motor Company, and a complete report of the project can be found in Bougler, Cody, et al. (2005). In this project, the same subjects that were tested in the DII experiment, were tested under similar conditions, only using a left-turn warning presented in the vehicle (a DVI). By combining the data sets, each driver made approximately 48 intersection crossings with a single oncoming vehicle either with no assistance, with the assistance of a DII, or with the assistance of a DVI, allowing for a direct comparison of the potential effects of using infrastructure versus in-vehicle based warnings.

The second analysis discussed in this report also used the combined data set from the two test-track intersection experiments conducted at the Richmond Field Station. The main goal of this third analysis was to examine the impacts of vehicle speeds, distances, and times to intersection

by using a series of regression analyses, to determine the predictive power of each component used in the predicted trailing buffer model.

The third analysis discussed in this report examined the trailing buffer or predicted post-encroachment time metric under a variety of both real and simulated cases. The goal of this section was to understand and describe how the metric evolves during the intersection approach. The data used in this section came from the combined data set from the two test-track experiments conducted at the Richmond Field Station. Simulations were then based on the real data taken during the experiments to show what would happen to the metric in cases when the Principal Other Vehicle (POV) driver decides to speed up or slow down.

The fourth and final analysis discussed in this report examined the feasibility of providing drivers with a last-second warning, potentially in response to the driver ignoring some initial decision support also provided by the system. Again, the combined data set from the two test-track experiments conducted at the Richmond Field Station served as the basis from which to provide supporting evidence for this analysis.

Data Set Descriptions

Two sets were used in the analyses described in this report. The data collection or experiments that are described in this section were not conducted as part of this CICAS data mining project and report. The experiments described in this section were all conducted under separate projects. The descriptions provided in this section simply review the key points of each experiment to better understand where each data set comes from, and they describe how the different data sets were combined, merged, or otherwise transformed to allow for the new data analyses that are described later in this report. Where appropriate, references to the original project reports are given which provide far more detail on how each experiment was conducted.

Overview of Test Equipment Common To All Data Sets

In both of the data sets, the test participants drove the California PATH Ford Taurus instrumented research vehicle. While a more complete description on this vehicle can be found in Bougler, Cody, and Nowakowski (2008), a summary of the key vehicle components is provided below in Figure 7.

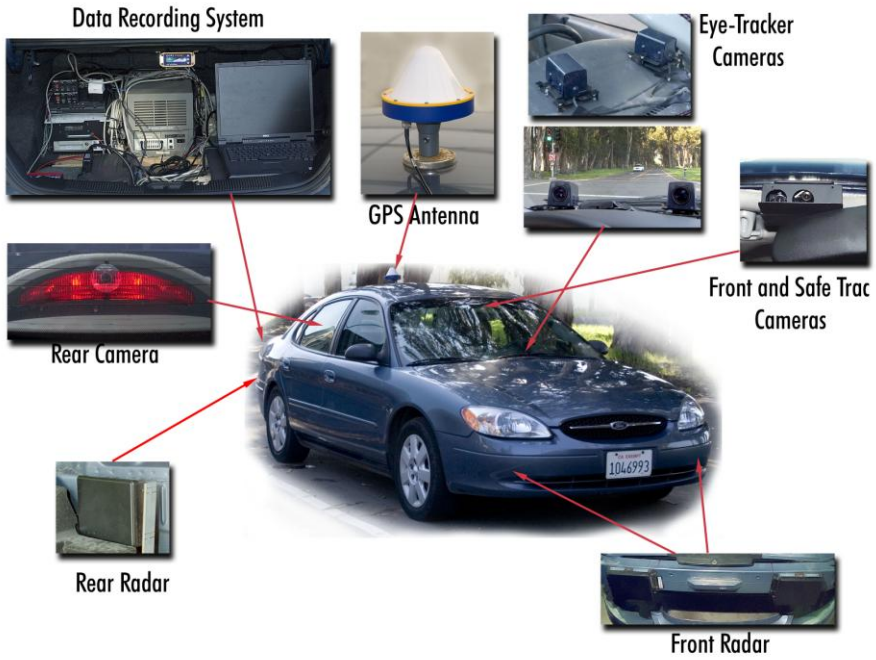

Figure 9. California PATH instrumented Ford Taurus (model year 2000).

Deleted:

To the driver, most of the instrumentation was not obvious. The only visible instrumentation consisted of two cameras on the dashboard which were used by the eye tracking system and an array of antennas on the outside of the vehicle. The instrumented vehicle's standard configuration recorded the following information at a rate of about 13 Hz:

- Speed
- Steering angle
- Brake pedal (on/off)
- Turn signals
- Throttle percent
- 3-axis acceleration
- Differential GPS
- Two forward looking EVT-300 radars (12 degree field of view & 150 meter range)
- 5 camera channels of video (3 forward, 1 rear, and 1 on the driver)
- FaceLab eye tracking system (version 3)

IDS Richmond Field Station DII & DVI Experiment Data Sets

Overview of Data Sets

Two separate experiments were conducted at the Richmond Field Station intelligent intersection. The first experiment was funded under the IDS project and focused on testing drivers reactions to left-turn warnings presented in the infrastructure (a DII). A second experiment, of identical design and recruiting the same subjects, was funded by the Toyota Motor Company and focused on testing drivers reactions to left-turn warning presented in the vehicle (a DVI). For the

purposes of the data mining analyses described in this report, the data from these two experiments were merged into a single data set, where each of 20 drivers made approximately 48 intersection crossing, both with and without the aid of warnings.

Test Participants

Although the details of the test participants have been covered in a previous IDS report (Bougler, Cody, and Nowakowski, 2008), in summary, twenty licensed drivers from two age groups, 10 younger (20 to 38 years old, mean of 28.3) and 10 older (65 to 84 years old, mean of 75.2), participated in the experiments. The same drivers were recruited to participate in both experiments. Within each age group, there were 5 men and 5 women drivers. Thirty percent of the drivers reported driving less than 5,000 miles per year, 45 percent reported driving between 5,000 and 10,000 miles per year, and 25 percent reported driving over 10,000 miles per year. While the reported mileage may seem low, it should be noted that the older drivers were typically retired. Additionally, in an urban area, a 30 mile per day commute (7800 miles per year) can easily equate to 1 to 2 hours per day with traffic. Participants were recruited through email and word of mouth advertisements. All participants were paid $15 in each study for their participation.

Test Equipment and Procedures

Both experiments utilized the California PATH instrumented intersection located at the Richmond Field Station, in Richmond, CA. The intersection is a typical four-leg intersection with one 12-foot wide lane in each direction. Using a suite of in-pavement magnetic loops, 3M microloops, and EVT-300 radars, a roadside computer recorded the approaching vehicles' speeds and distances. The roadside computer then calculated warnings and relayed them to either a sign in the infrastructure (DII) or wirelessly over an 802.11b network to the test participant's vehicle to be displayed on an in-vehicle display (DVI). The traffic signal was always green during the experiments. A picture of the Richmond Field Station instrumented intersection as seen from the test participant's perspective is shown in Figure 8.

Figure 10. Test equipment used during the Richmond Field Station experiments.

Deleted:

The DII was a LED-based, looming, no-left-turn sign, that was mounted on the traffic signal pole (slightly below the pedestrian signals) on the far-left corner of the intersection from the Subject Vehicle (SV) driver's point of view. The design and operation of the DII has been described in detail in previous IDS reports, specifically, Misener, et al (2007).

The DVI used a 7" LCD display (Xenarc Model 700YV), mounted in a high central position (as shown in Figure 1) to approximate a typical navigation system display. The visual aspect of the warning also used a looming no-left-turn graphic identical in nature and function to the DII (although the DVI looming states are not depicted in the figure). Additionally, the DVI warning screen was titled "Vehicle Approaching," and a countdown bar representing the Principal Other Vehicle's (POV) current distance to the intersection was provided to the left of the looming no-left-turn graphic. An auditory cue consisting of a pair of 2000 Hz tones at a 200 ms cadence (with the volume adjusted to a comfortable level for each driver) was also provided once to indicate that the DVI warning screen had activated. Further details are provided in Bougler, et al. (2005).

In summary, for both experiments, a trial started with the SV and POV parked approximately 80 and 260 m from the intersection, respectively. The start of the approach was coordinated by computer. To the test participant, the start of a trial was signified by a change in the DVI from a pretrial screen to a simulated navigation system screen accompanied by the voice instruction "Left Turn Ahead." The SV driver was then instructed to accelerate to 20 mph (8.9 m/s) while the confederate POV driver accelerated to 25 mph (11.2 m/s). The normal variations in the initial driver reaction time (the starting of the vehicles), the acceleration profiles, and the actual approach speeds for a given trial all affected the outcome of how the vehicles reached the intersection and the exposure to warnings. The test participants were both free and encouraged to violate either the DII or DVI if they felt it was wrong, and there was plenty of time to turn in front of the POV.

Summary of the Data Collected During the Experiments

During each of the experiments conducted at the Richmond Field Station, data were collected from two primary sources. First, there was data collected by the sensors mounted in the intersection (primarily EVT-300 radars), and second, there was data collected by the SV's (Taurus's) onboard data recording system. Table 1 contrasts the different measures available from each recording source.

Table 1. Data collected by the intersection compared to onboard the vehicle.

Data Collected by the Intersection	Data Collected by the SV (Taurus)
• SV speed • SV distance to intersection • POV speed • POV distance to intersection • DII or DVI warning activations • Approach outcome (whether or not the SV turned in front of the POV) • Acceleration derived from speed	• Speed (from vehicle wheel pulses) • Steering angle • Brake pedal (on/off) • Turn signals • Throttle pedal (percent depressed) • 3-axis acceleration (from an accelerometer) • Differential GPS (SV position) • 5 camera channels of video (3 forward, 1 rear, and 1 on the driver) • FaceLab v.3 eye tracking system

As can be seen from comparing the columns in the above table, there is some overlap in data collected to track the SV, but measurements of the oncoming traffic (POV) were only available from the infrastructure sensors. Additionally, there were differences in the accuracy and reliability of the information provided by each sensor. As an example, the DGPS tended to lose signal during sections of the intersection approach depending on the orientation of the satellites at the time of the test, while the intersection-based radars tended to lose track of the SV as it got too close to the intersection, as it slowed down, and once it started to make the turn. Furthermore, actual driver inputs such as steering, throttle, and brake were directly sensed onboard the vehicle, but could only be inferred by looking at the data collected by the intersection-based sensors.

During the original analysis of the IDS experiment which has been previously cited, only the intersection-based measures were analyzed since the method that had been set up to merge and synchronize the data sources failed to work correctly. As part of this data mining effort, one of the questions of interest was whether or not it would be possible, with a reasonable amount of effort, to merge the data sources using other means. The merging of data sources would have allowed a direct comparison between the IDS Berkeley Field Test data set and the Richmond Field Station experiments data sets. Some of the questions posed that were of interest and would have required a synchronization of the two data sources included the following:

1 Given the slower approach speeds and more contrived nature of the Richmond Field Station experiments, it is possible that the driving behavior seen could be different from what was observed in the more naturalistic Berkeley Field Test. Being able to rerun some of the analyses performed during the IDS Berkeley Field Test with the Richmond Field Station experimental data would have helped to provide some external validity to the experiment results.

2 A comparison of initial braking measures was desired to determine what kind of delay would be seen from the time that a driver put his foot on the brake pedal until the time that the radar detected deceleration.
3 An examination of the steering wheel movements with respect to the distance to intersection was desired. The EVT-300 radar used to track the SV distance to intersection provided very little accuracy when it came to lateral movements. One hypothesis to explain why turns were made during trials with small (or negative) predicted trailing buffers was that the drivers began their turning movement early. Thus, by cutting the corner and shortening the travel distance, the SV was able to clear the zone of conflict faster than might be predicted.
4 Although anecdotal evidence was observed on where drivers were looking during the approach, verifying these observations with the eye tracker data was desired.

Data Source Synchronization

During the Richmond Field Station experiments, the data acquisition systems were located in two physically separated places, the instrumented Ford Taurus and the instrumented intersection. Furthermore, the data files were recorded at different intervals with different start and stop times. The Taurus engineering data files were recorded and stored in 15-minute files starting from the moment that the car ignition was turned on. The video on the Taurus was recorded and stored in one-minute files. At the intersection, data were stored in 15-minute files starting from the moment that the experimental control software was switched on.

The attempted synchronization process was broken into three stages. The first stage was to synchronize the 15-minute engineering files from the Taurus with the 1-minute video files. Although the processes to do this phase were already in place and well tested, there was a catastrophic error in this software which ended up accidentally deleting most of the IDS data. This caused a delay from July to almost October of 2007 for data recovery. However, in the end, all of the data were recovered and synchronized in this first phase.

The second stage of the synchronization process involved a rough synchronization of the now 1-minute Taurus engineering and video files with variable length data files recorded by the intersection which corresponded to a single trial (intersection approach by the Taurus). This rough synchronization was done manually using a combination of rough timestamps and by watching the videos to determine when each Taurus intersection approach occurred. This method, however, only synchronized the data files to within one-minute.

The third and final stage of the synchronization process was to use the two overlapping measures of vehicle position or speed to get a sub-second synchronization of the data sources. Vehicle position was measured simultaneously by the intersection-based radar and by the vehicle-based DGPS. However, synchronization based on position alone proved to be impossible as there were a significant number of trials in which the DGPS signal had completely dropped out due to not seeing any satellites.

Alternatively, vehicle speed was simultaneously measured by both the intersection-based radar and the vehicle-based wheel-pulse counter. By knowing the approximate starts of a trial or intersection approach, the speed profiles as measured by both sources could be synchronized. However, just looking for the start of vehicle movement as detected by both the radars and wheel pulse counters was not very accurate from a synchronization standpoint. The radar, being

Doppler based, had quite a bit of lag in detecting the start of vehicle movement, since the vehicle needed to be traveling at a speed of over 0.4 m/s (1 mph) for a period of time before the vehicle was even registered by the radar. Conversely, the wheel-pulse counter on the vehicle was somewhat noisy at speeds below 1 m/s (2 mph).

As shown in Figure 9, what could be reliably matched was the overall acceleration and deceleration profiles of the vehicle during each trial. The graph below shows a simple visual best fit for overlaying the wheel speed profile to the radar speed profile as the vehicle approached the intersection from about 85 meters down to about 5 meters past the stop bar (designated as 0 m of distance to intersection on the graph).

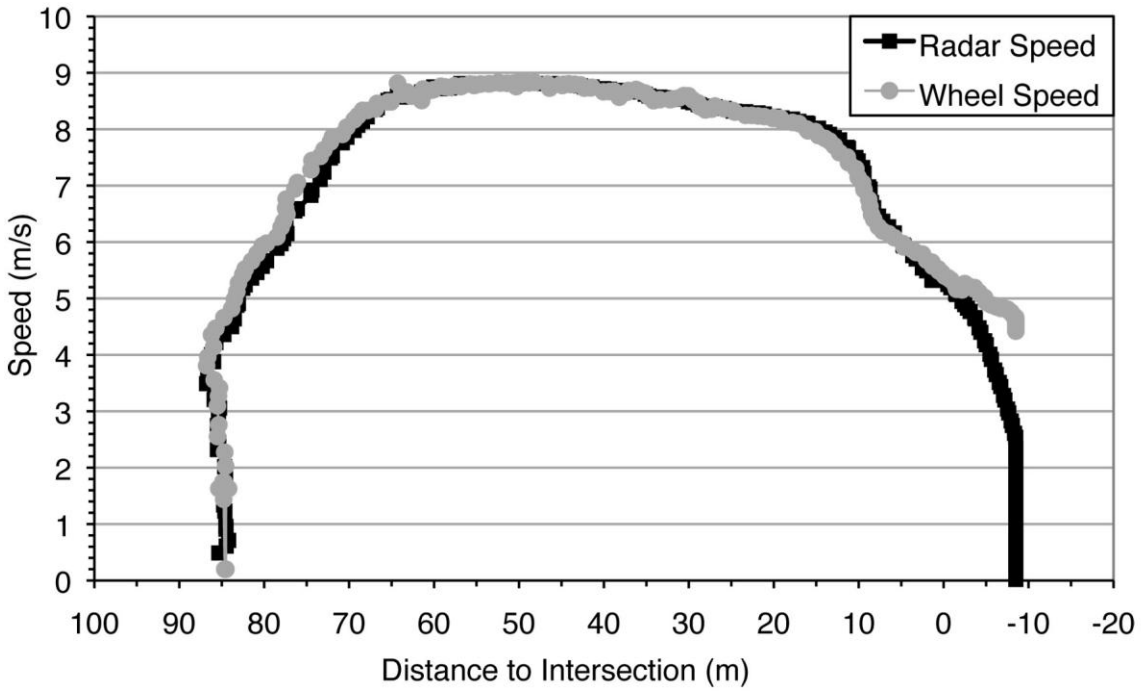

Figure 11. Typical visual synchronization of intersection- and vehicle-based speed measures.

Deleted:

Given that a visual synchronization of the speed curves could be made, then it would most assuredly follow that an automated mathematical process, such as using least squares regression, could be written to provide for the synchronization of these two data sources. However, at this point, it was determined that although merging of the intersection-based and vehicle-based data sources was technically feasible, it was not feasible within a reasonable amount of effort for the potential gains.

This was especially true given the advent of newer and more promising intersection-based sensor technologies, such as the Smart Micro Sensors radar (http://smartmicro.de/) and IBEO 270-degree laser scanner (http://ibeo-as.de/). Both of these sensor solutions promise to do a much better job at tracking vehicles both up to the intersection and through making a turn at the intersection. In light of these advancements and the continued difficulty in getting the intersection and vehicle data for the previous experiments synchronized, it was decided to abandon the parts of the data mining project that required synchronization in order to focus on installing the newer, more advanced sensors and collecting new sets of data to answer the questions of interest.

Data Set Limitations

Since the synchronization of the intersection-based and vehicle-based data sources was abandoned, it is worth pointing out the limitations of the intersection-only-based measures which have been used for the analyses to follow in this report. There are two major limitations in trying to extend the data collected during the Richmond Field Station DII and DVI experiments. The first and most obvious limitation of the data sets comes from the fact that DII and DVI warnings most likely influenced drivers' turning decisions, and subsequently, their choices of speed and braking. Furthermore, the warnings came at three different timings (distances to the intersection). In summary, when examining data when the SV is 36 m or closer to the intersection, any potential effects of a DII or DVI warning need to be considered.

The second limitation of the data set comes from the characterization of the EVT-300 radar performance as the SV reaches the intersection stop bar. As shown below in Figure 10, tracking the radar speed as the SV reaches the stop bar (a distance to intersection of 0 m), it appears that the vehicle came to stop several meters into the center of the intersection. However, in this case, the vehicle actually turned left in front of the POV without ever stopping. This difference is highlighted by watching the divergence of the onboard wheel speed measure as the vehicle passes the stop bar. According to the wheel speed sensor, the vehicle never slowed below 4 m/s (9 mph). This discrepancy can easily occur as the radar switches reflection targets on the car and eventually loses track of the car as it crosses and leaves the path of the radar beam. To summarize, the intersection-based radar measurements of vehicle speed and distance which were used in the Richmond Field Station experiments can only be trusted to within about 5 m from a vehicle actually reaching the stop bar.

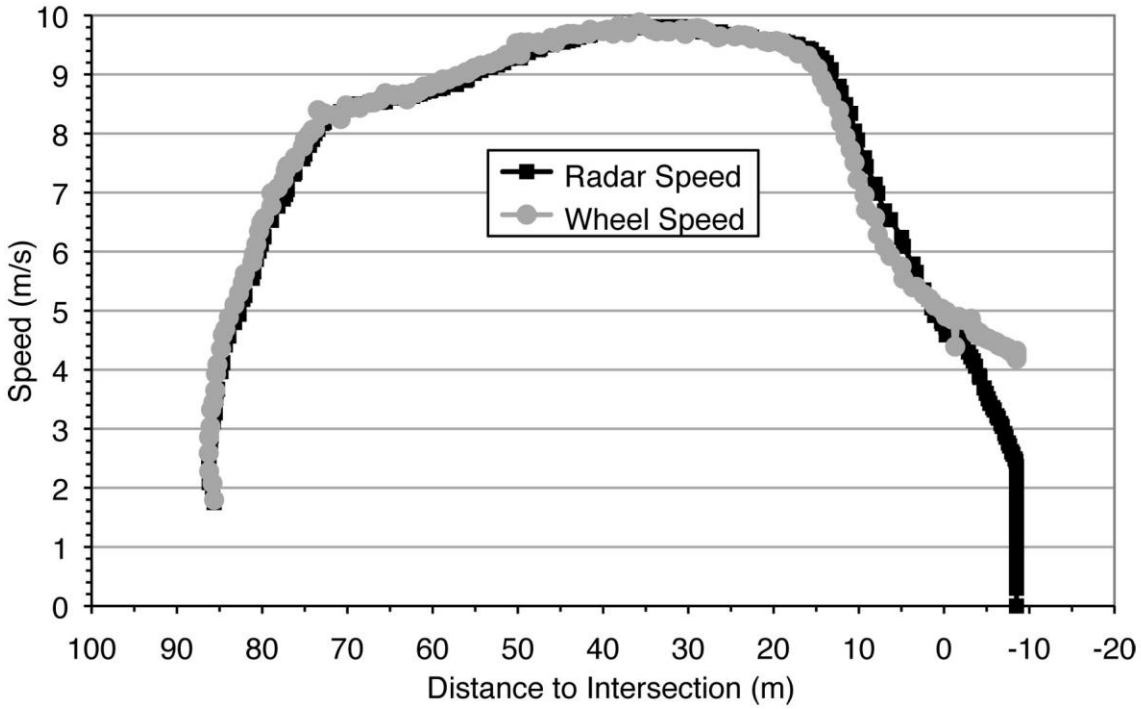

Figure 12. Typical trial when radar diverges from wheel speed after reaching the stop bar.

Deleted:

Driver Infrastructure vs. In-vehicle Interfaces

Overview

The goal of the analysis described in this section was to compare the effects of an infrastructure-based left-turn collision warning to a similar in-vehicle variant of the same warning. More simply, this analysis tried to tackle the question of whether it was better to give a left-turn advisory or warning through a roadside sign or using an in-vehicle display. At the time of this report, at least one report based on this analysis had been published at the 87th Annual Transportation Research Board Meeting (Nowakowski, Cody, and O'Connell, 2008).

The infrastructure based warning, also referred to as a Driver Infrastructure Interface (DII), was a looming no-left-turn sign. The in-vehicle variant, also referred to as a Driver-Vehicle Interface (DVI), showed a graphic of the same looming sign on an in-vehicle display, but added an auditory component to alert the driver to a change in the display state. This analysis used the combined data set from two experiments conducted at the Richmond Field Station. The details of those original experiments, as well as images of the DII and DVI, were reviewed earlier in Section 2.3.

The differences between the DII and DVI warning locations were examined along three lines:

1. How effective was each location (DII vs. DVI) at influencing the decision to turn?
2. How did the warning location affect the requirements for the timing of the information delivery (warning onset)?
3. How did the warning location and timing affect braking (driver response)?

Experimental Design

Three factors were manipulated when these two experiments are combined: vehicle arrival as measured by predicted trailing buffer (predicted post-encroachment time), warning presence (None, DII, or DVI), and warning onset (based on the SV's distance to the intersection). The experimental design (see Table 2) was composed of eight unique test conditions, with three repetitions per condition, for a total of 24 trials presented in a fixed randomized order (not including six additional practice trials). Multiplying this design by two experiments, there were a total of 48 trials. By combining the two experiments, there were simply more repetitions for the conditions when the warning was not shown as the DII and DVI Off columns could be combined and coded a single condition in the analysis.

Table 2. Experimental design (number of trials by condition).

Experiment 1: DII					Experiment 2: DVI				
Trailing Buffer (s)	DII On (with onset of)			DII Off	Trailing Buffer (s)	DVI On (with onset of)			DVI Off
	20 m	28 m	36 m			20 m	28 m	36 m	
-1.0	3	3	3	-	-1.0	3	3	3	-
0.0	3	3	3	3	0.0	3	3	3	3
1.0	-	-	-	3	1.0	-	-	-	3

The first factor manipulated was the arrival of the vehicles which was measured by the average predicted trailing buffer (predicted post-encroachment time), a prediction of how much spare time would remain if the SV made a typical turn in front of the POV. For any given SV position,

the predicted trailing buffer was calculated by subtracting the SV's predicted time to clear the zone of conflict from the POV's time to reach the zone of conflict. This metric possesses the unique quality of being relatively invariant, i.e., it remains relatively constant throughout the approach until the SV driver decides to stop and wait for the POV to pass.

For the RFS intersection, the zone of conflict was defined as a box at the quadrant of the intersection box where the two vehicles would cross. So, logically, a predicted trailing buffer of 0.0 s would mean that the SV's rear bumper would just clear the intersection as the POV's front bumper entered the intersection. However, the specific algorithm used to calculate predicted trailing buffer was based on a mean turning speed of 13.2 mph (5.9 m/s) and a mean deceleration rate of 0.16 g (1.61 m/s2). As originally discussed in Bougler, Cody, and Nowakowski (2008Error! Bookmark not defined.), the use of means biased the prediction, allowing successful turns to be made in front of the POV with predicted trailing buffers of less than 0 s when the SV driver was more aggressive than the mean.

Deleted:

While the predicted trailing buffer could be calculated for any point during the SV's approach, an average value was used for this analysis. The predicted trailing buffer, for any particular trial, was computed as the average instantaneous trailing buffer during the SV's approach (between 36 and 20 m prior to entering the intersection). This corresponded to the time period when a warning was likely to go off. However, averaging over this period also likely introduced some bias since that is also the time period when the decision to stop was made and executed. Once the driver begins to execute a stop, the invariant properties of the predicted trailing buffer metric cease.

Although the planned experiment design called for 3 levels of predicted trailing buffer to be tested, the reality of the experiment was that vehicle arrivals could not be tightly controlled, resulting in trials with a much greater dispersion of predicted trailing buffer (ranging from -6.0 to 4.5 s). For predicted trailing buffers below -1.7 s, all trials results in the SV slowing and stopping to let the POV pass. For predicted trailing buffers above 1.0 s, almost all trials resulted in the SV turning in front of the POV. For predicted trailing buffers in between, drivers sometimes turned and sometimes stopped.

The second factor manipulated was the warning presence or location (None vs. DII vs. DVI). However, to accommodate the variability in trailing buffer, when the predicted trailing buffer was below 1.0 second, the warning was always given, and when it was above 1.0 second, the warning was always suppressed. For trailing buffers between -1.0 and 1.0 seconds, most drivers experienced some trials both with and without a warning. Effectively, although there was a experimental plan in which test conditions were laid out in a factorial design, in reality, the actual conditions for each trial were somewhat randomly distributed and did not perfectly match the design for each participant.

The third factor manipulated was the warning onset or the point during the SV's approach when the warning would first illuminate. Three warning onset points were tested, 20, 28, and 36 m prior to entering the intersection box, corresponding to approximately 2.5, 3.5, and 4.5 seconds before reaching the intersection at the nominal speed of 20 mph (8.9 m/s). It should be noted that the stop bar (or in this case, crosswalk), was roughly at 4 m. However, during the left-turn maneuver, the crosswalk was not a factor as drivers always choose to stop in the middle of the intersection (at approximately -2 m) when waiting for a POV to pass.

The Effect of Warning Location on Warning Effectiveness

Overview

The first goal of this experiment was to examine the effect of warning location on warning effectiveness as measured by the influence exerted by the warning over the driver's decision to stop. For each trial there were two possible outcomes: either the driver could turn in front of the oncoming vehicle or stop and wait for it to pass. If the driver chose to stop, they were asked the following question: "Did you think there was enough time to turn in front of that car?" Responses fell into the following three categories:

1 The SV stopped but the driver felt that he could have turned
2 The SV stopped but might have turned if he was in a hurry
3 The SV stopped and the driver felt that there was not enough time to turn.

Thus, besides simply looking at whether or not the driver turned, the trial response could be further differentiated by the opinion that was solicited on whether or not the driver thought that there was enough time to turn. It should be noted that drivers never expressed regret after choosing to make a turn in front of the POV in either the DII or DVI experiment.

Analysis Methodology

One of the opportunities in this data mining project was to explore different analysis methodologies which could be applied to the data sets. The previous IDS DII experiment attempted to apply an analysis of variance (ANOVA) methodology; however, there were a number of caveats discussed in the report as to why the data did not fit perfectly into the bounds of this methodology. ANOVAs require that the measurement be both continuous and normally distributed, however, in this case, the measurement is really binary or categorical, either the driver turned or stopped or the driver selected one of three finite choices. Second, too many missing cells in a factorial design become problematic.

In this data mining project, the use of generalized linear models was selected to be more flexible. The general linear model is based on least squares regression and allows for the selection of a non-normal distribution (such as a binomial or multinomial distribution) which can be related through a link function. This, in effect, relaxes the normality assumption and handles the problem with missing cells by allowing predicted trailing buffer (a factor that could not be tightly controlled) to be modeled as continuous covariate.

Furthermore, the analysis of warning effectiveness by location was also broken into two separate analyses. The first analysis looked at just the decision to turn or stop, while the second analysis looked at the driver rating given if and only if the decision was to stop. In the original analysis of the IDS DII experiment, the responses were viewed as simple continuum from turned (and confident about the decision) to stopped (and confident about the decision). However, there was criticism on this notion of a single response continuum. Given the temporal separation between the decision to turn or stop and the decision to answer the rating question, it is possible (and likely) that these are separate decisions which may be based on entirely different underlying phenomena. Thus, the decision was made to separate the analyses.

Typically, when dealing with a binary response, such as the decision to turn or stop, it is modeled using a binomial distribution with either a logit or a probit link function. Logit link functions are

used to model the proportion of cases in each category. It assumes equal probability to be in either category, and the typical examples use a logit function to model political party affiliation. Probit models have assumptions that are consistent with having a binary dependent variable which is assumed to be a proxy for some underlying continuous normal distribution. It is typically used in medicine with the dose-response paradigm. The effect of the dose (medicine) may be continuous, but the output (response) may only seem categorical, such as feeling better or feeling worse. Additionally, probit models have also been used in signal detection theory and business success vs. failure.

In modeling decision to turn, the analysis described in this report used a probit link function. The reasoning follows the logic that the decision to turn or stop is based on evidence, such as a driver's estimation of his or her time required to clear the intersection. This "evidence" of the situation, that is gathered by the driver, could be thought of an underlying continuous variable as suggested by the typical uses of the probit model. While there are some reports that have modeled certain driving and turning decisions using logit models (Yan and Radwan, 2007), the examples of when to use probit models seem more consistent, at least *prima facie*, with turning decision scenario being analyzed in this situation. However, in practice, there is usually very little difference in the analysis results due to selecting probit versus logit link functions.

Results

In the first analysis, the turning rate, or the binary decision made by the driver to either turn in front of the POV or to stop and let it pass, was analyzed using a repeated measures, generalized linear model assuming a binomial distribution with a probit link function. The model tested a factorial design between age, warning presence and location (no warning vs. DII vs. DVI), and average predicted trailing buffer (which was modeled as a covariate). Simple plots of response by trial number showed that practice effects were negligible, and therefore not need to be included in the model.

Gender effects were removed from the model as there were not enough evenly distributed cases to reliably model both age and gender. Any gender trends would likely be an artifact of the small driver sample size and of the three most extreme drivers. One younger female and one older male were extremely aggressive in their turning behaviors during both experiments, choosing to make many of their turns with very little trailing buffer. One older female was extremely conservative, choosing to reject very large trailing buffers.

While driver age was not significant, both the main effects for warning presence and location, Wald χ^2_2=35.7, p<.001, and predicted trailing buffer, Wald χ^2_1=40.4, p<.001, were significant. As expected, across all conditions, as the predicted trailing buffer increased, so did the turning rate. (See Figure 3-1.) However, the implication of the main effect for warning presence and location was a little more difficult to interpret. A pairwise comparison of estimated marginal means (using the Bonferroni correction) was used to determine which conditions were significantly different from each other. This analysis indicated that the presence of either the DII, p<.001, or of the DVI, p<.01, resulted in lower turning rates when compared to the conditions when no warning was given. However, there was not a significant difference in turning rate between the conditions when the DII was present as compared to the conditions when the DVI was present. Both warning locations were similarly effective at influencing the driver to slow or stop and let the POV pass.

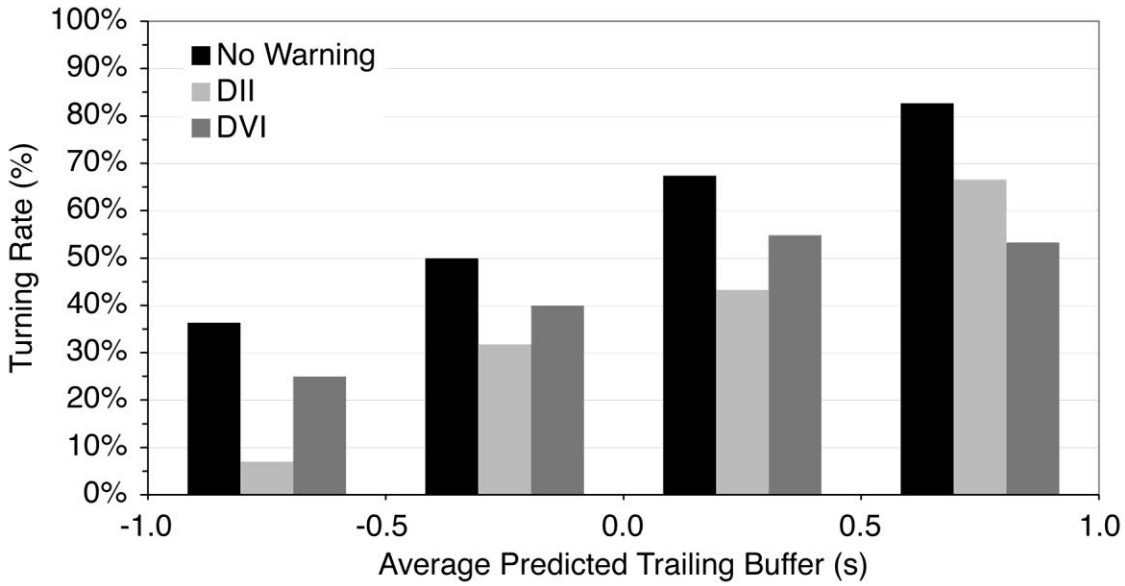

Figure 13. Driver turning rate by warning type and average predicted trailing buffer.

Deleted:

The second analysis examined the drivers' post-trial ratings on the appropriateness of their decision to stop and let the POV pass. If the driver decided to stop, a rating (from 1 to 3) corresponding to "enough time to turn", "might turn in a hurry", or "not enough time to turn" was solicited. The rating response was modeled using a repeated measures, generalized linear model assuming a multinomial distribution with a probit link function (using the same model effects and covariates as described in the first analysis). Again, as expected, trailing buffer was significant, Wald χ^2_1=12.2, p<.001, which simply indicates that the larger the trailing buffer, the more likely the drivers were to feel that they might have turned in a hurry or had enough time to turn, even though they actually stopped.

As shown in Figure 12, when the predicted trailing buffer was greater than 0.0 s, nearly 80 percent of the time when drivers stopped (and no warning was given), the drivers said that they either could have turned or might have turned if they were in a hurry. Although Figure 3-2 also shows that there was typically small decrease in the percent of responses indicating that there was enough time to turn when either the DII or DVI was present, warning presence or location did not have any significant effects or interactions on the post-trial driver rating of their decision to stop. This may be due to the low sample sizes in these experiments, especially in the very low number of trials that fell into the category of trailing buffer less than -0.5 s and no warning given.

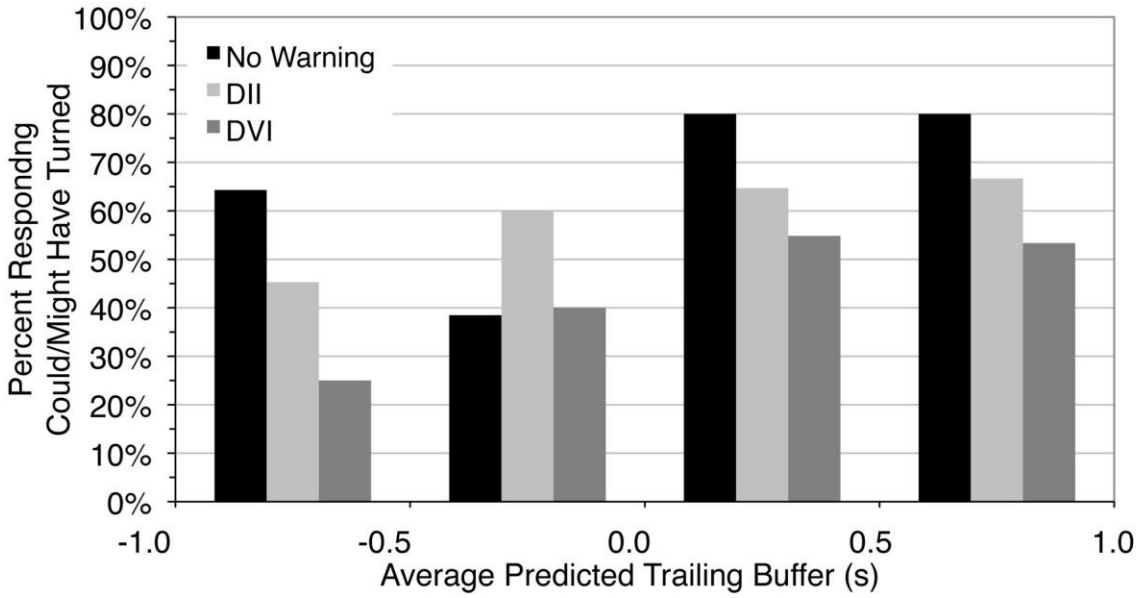

Figure 14. Post-trial "could/might have turned" response percentage by warning type and average predicted trailing buffer.

Deleted:

The Effects of Warning Location and Onset on Driver Ratings

Overview

The second goal of this study was to examine the effect of warning onset to determine when the optimal time would be to give a driver an indication that it was unsafe to proceed with the left turn. In both the DII and DVI experiments, an opinion that was solicited by the driver after each trial where a warning was presented. Responses were coded on a five-point rating scale ranging from "too early" to "too late" with "Just Right" being in the middle. Although three different warning onsets were used, the driver was not made aware of which timing was used for a particular trial. Ratings were given regardless of whether or not the driver choose to obey or ignore the advice.

Analysis Methodology

In this analysis the driver ratings were again modeled using a repeated measures, generalized linear model assuming a multinomial distribution with a probit link function. The model contained a full factorial of age, gender, warning location (DII vs. DVI), and warning onset (20, 28, or 36 m from the intersection). The main effect of trailing buffer was included in the model as a covariate.

Results

The main effects for age, gender, and warning location were not significant, but the main effect of warning onset, Wald χ^2_2=68.2, p<.001, was significant. Overall, none of the warning onset points were appreciably rated as too early, so the main effect of warning onset can be summarized by saying that the later the warning came, the more frequently it was rated as a little late or too late. As shown in Figure 13, the 20 m onset (latest warning) was more frequently rated on the late side across both age and warning location as compared to the 28 and 36 m onsets, which tended to be rated as "just right" more frequently.

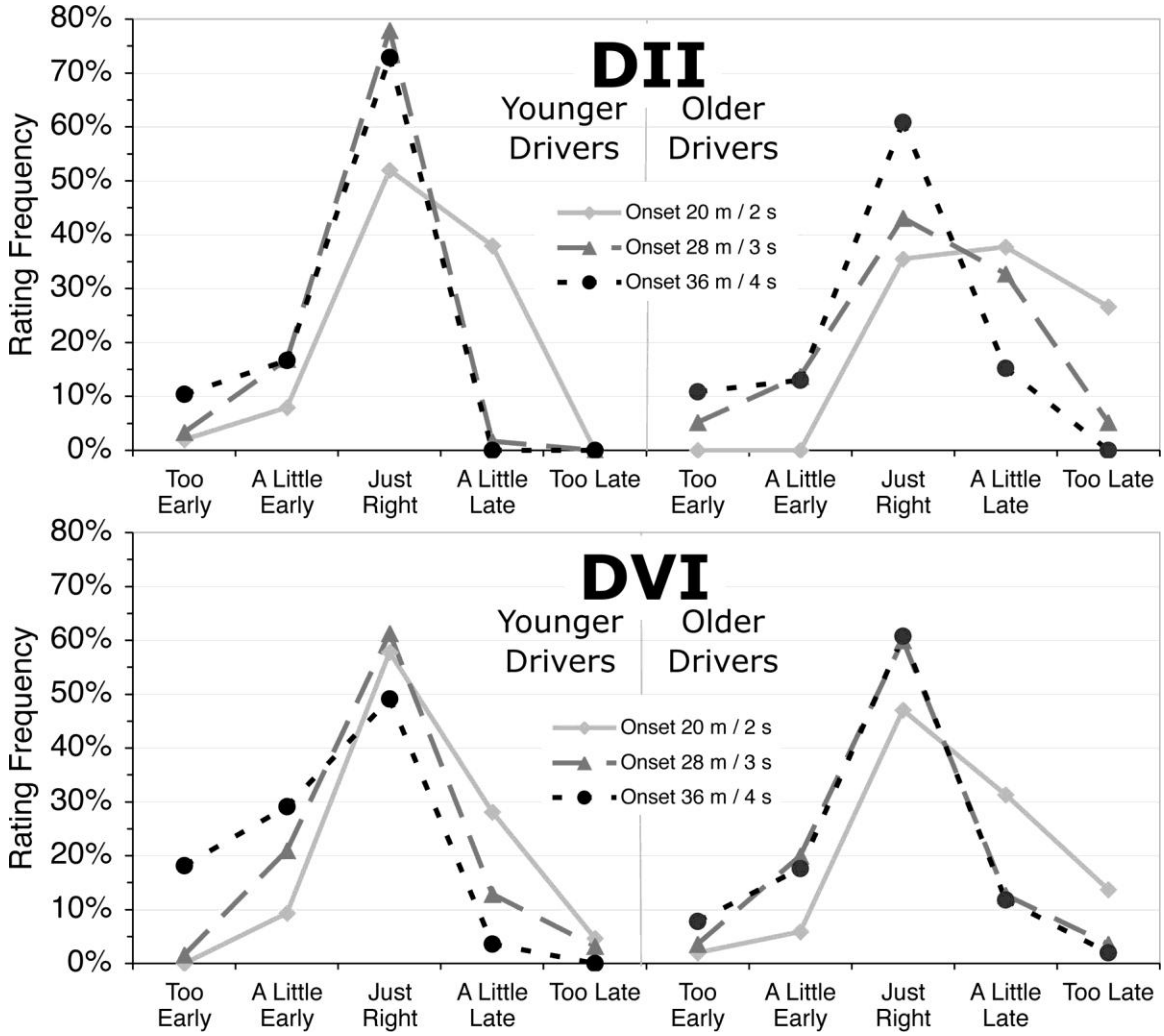

Figure 15. Frequency of driver ratings as a function of warning onset.

Deleted:

Although the main effects of age and warning location were not significant, the interaction between age, warning location (DII vs. DVI), and warning onset, Wald χ^2_2=10.5, p=.005, was significant. While this may seem like a very complex interaction, it can be fairly simply explained. Older drivers showed a clear preference for the earliest warning onset, 36 m, when the warning was given in the infrastructure (DII). However, when the warning was given in-vehicle (DVI), the older drivers rated both the 28 m and 36 m onsets as similarly acceptable. In contrast, younger drivers rated both the 28 m and 36 m onsets as similarly acceptable when the warning was located in the infrastructure (DII). However, when the warning was given in-vehicle (DVI), the 36 m onset was rated as being on the early side almost 50 percent of the time.

To summarize, older drivers showed a clear preference for earlier warnings, but when the warning was given in-vehicle (DVI), they were more willing to accept slightly later warnings. Younger drivers were more willing to accept later warning onsets for both warning locations; however, they were less willing to tolerate earlier warnings when using the DVI.

Although the test participants had been asked to rate the timing of the warning onset irrespective of whether or not they agreed with the appropriateness of the warning, the main effect of trailing

buffer, Wald χ^2_1=5.1, p=.02, was also significant. As shown in Figure 14, when the trailing buffer was large, between 0.5 and 1.0 s (indicating a larger lag), warnings were less likely to be rated as "just right," and the distribution of responses was weighted more towards the early side. More simply put, when the drivers were more likely to disagree with the need for a warning, they were also more likely to disagree with the timing of the warning causing them to feel that the warning came too early.

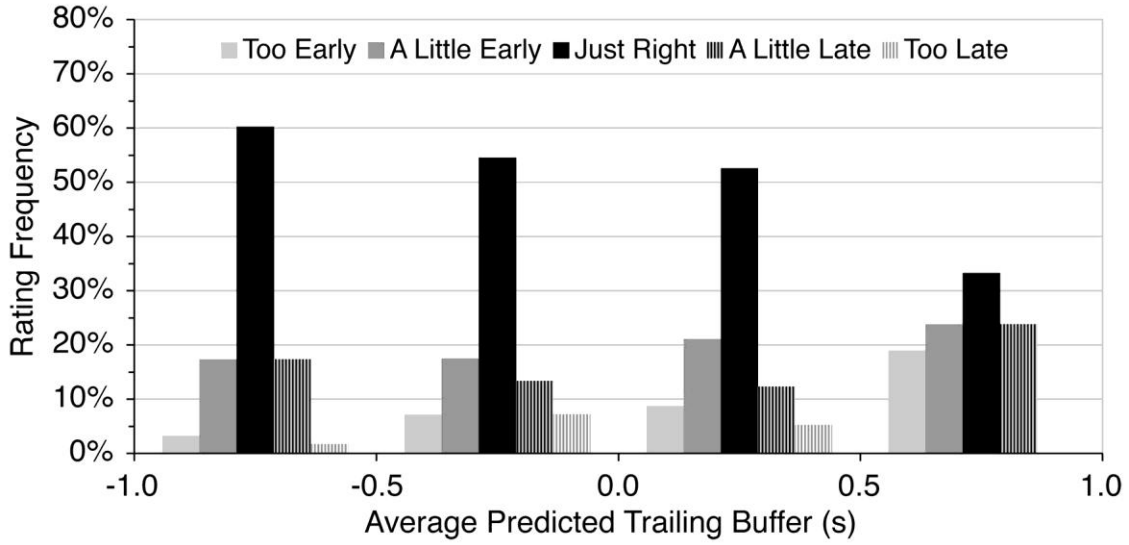

Figure 16. Frequency of driver ratings as a function of trailing buffer.

Deleted:

The Effect of Warning Location and Onset on Braking Profiles

Overview

While the preceding section described the effects of warning location and onset on subjective driver ratings, a more objective performance measure might be found in examining the SV driver's braking profile. By looking at the location of the driver's onset of initial braking, some insights might be provided as to when the decision to turn or stop is being made. If, for example, drivers typically started braking before a warning was even given, then it might indicate that the warning was occurring too late. However, care needs to be taken to not over extrapolate the results of the analysis beyond this experiment. In reality, the onset of initial braking is highly dependent on intersection geometry, approach speeds, and traffic conditions. None of these were variables in this experiment.

Analysis Methodology

Two analyses were run using the onset of initial braking as the dependent measure. The onset of initial braking was defined as the distance from the intersection at which the SV's deceleration exceeded 0.03 g as measured by the infrastructure-based radar. Although this threshold may seem low, it was chosen based on the relative consistency of the SV approach and the braking characteristics of the test vehicle. The test vehicle had a relatively high idle due to the increased electrical load on the alternator from the data recording systems. At speeds of 20 to 25 mph, the vehicle did not seem to slow much by just letting off the accelerator.

The first analysis used a repeated measures ANOVA with age and gender as between-subjects factors and a single treatment category, test condition, as the within-subject factor. In this treatment category, four test conditions were compared: 1) No warning and the SV driver turned in front of the POV, 2) No warning and the SV driver decided to stop and let the POV pass, 3) The DII warning was given and the SV driver decided to stop, and 4) The DVI warning was given and the SV driver decided to stop.

A second analysis of the initial braking onset focused only on the difference between the DII and DVI combined with the differences between warning onsets for the conditions for cases when the SV stopped. A repeated measures ANOVA was used with age and gender as between-subjects factors and warning location and warning onset as within-subjects factors.

Unlike the previous analyses, these two analyses of onset of initial braking could be modeled using simple ANOVA techniques because the assumptions of a continuous response and of normality were not violated. (See Appendix A for further discussion and analysis.) There were also no issues with missing cells since trailing buffer was not really a factor that was considered.

Results

In the first analysis of onset of initial braking, the only effect that was significant was the main effect of test condition, $F(3,48)=6.5$, $p<.001$. Test condition, as described earlier, was the comparison of four different treatment, or in this case, trial outcomes (see Figure 15). Although the main effect was significant, to find out which treatments or test conditions were different from each other, a post-hoc Tukey-Kramer was used. This test revealed that the only test condition that was significantly different from the others was the case when no warning was given and the driver decided to turn in front of the POV.

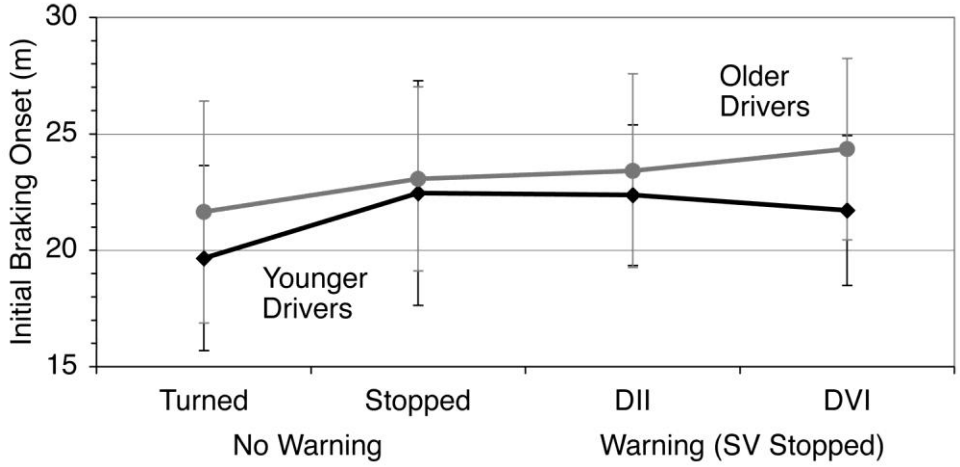

Figure 17. Initial braking onset by test condition.

Deleted:

As shown in Figure 15, drivers who decided to turn in front of the POV initiated braking slightly later than those who decided to stop. Although this difference was statistically significant, the mean difference in braking onset between drivers deciding to turn and drivers deciding to stop was only on the order of 3 m or less than 0.3 s at the nominal approach speed of 9 m/s (20 mph). Although this difference is not practically significant, because there is a slight and significant difference between the turned and stopped conditions, we can conclude from this analysis that the initial decision to either turn in front of the POV or stop and wait for it to pass, is made, on

average, between 20 and 25 m from entering the intersection (at least during this experiment). This would also lead us to conclude that the latest warning onset tested (20 m before the stop bar) should have resulted in warnings that arrived on the order of 0.3 to 0.5 seconds after the initial decision to turn or stop had been made.

The second analysis focused only on comparing the last two conditions shown in Figure 15, the DII vs. the DVI for cases when the driver decided to stop. By concentrating on only these two conditions, the effect of warning onset could also be tested. The main question here was whether or not having an earlier warning influenced drivers to stop earlier. However, consistent with the results of the first analysis, there appeared to be no significant influence on the initial braking onset for either the warning location or the timing of the warning onset. In effect, earlier warnings did not necessarily lead to earlier braking.

Discussion and Conclusions of Analysis

The first goal of this analysis was to examine the potential effectiveness of two warning locations (DII vs. DVI) and the primary measure used was the drivers' responses as measured by observed turning rate and a post trial rating scale (for trials when the driver decided to stop). The test participants were free to violate any warnings given during the experiment if they disagreed, so extrapolating directly from these results to estimates of warning compliance rate would not be appropriate. Overall, the turning rate increased as the predicted trailing buffer or post-encroachment time (a measure of the available lag or "gap") increased. The presence of a warning in either the infrastructure (DII) or in the vehicle (DVI) resulted in a reduction in mean turning rate of 35 percent as compared to similar trials when no warning was given.

The reduction in the turning rate was also greater when the trailing buffer was shorter, 55 percent as opposed to around 28 percent. However, when the driver did decide to stop, the presence or absence of a warning did not seem to influence the driver's *ex postfacto* rating on whether or not there had been enough time to turn. This was disappointing as it would have been beneficial for the decision support and warning system to provide for both a reduction in the number of turns and an increase in the driver's confidence about the correctness of their decision to stop.

While the results of this study lead to the conclusion that both infrastructure and in-vehicle warnings have potential to reduce unsafe left turns, there were number of design lessons learned. First, the drivers commented that the DII was intuitive and felt both trustworthy and compulsory since people are used to obeying signs placed in the infrastructure, although the compulsory aspect may have been due to the fact that the DII was modeled after an existing, enforceable, regulatory sign. The DVI concept was less familiar, and there was no inherent trust in any decision support provided by the vehicle. Second, drivers commented that the DII had a distinct advantage in visual location as it had been placed very near the driver's focus of attention (tracking the POV). In contrast, the DVI implemented in this experiment required the drivers to glance away from the road, which most were not comfortable doing during the turning maneuver. Finally, the auditory alert provided by the DVI was often viewed as annoying, with some drivers commenting that they would prefer this type of system to be silent. This preference is not surprising given the advisory nature of the information, combined with the potential frequency of alerts.

The second goal of this analysis was to examine how the warning location might influence the timing requirements for the information delivery. The warning was provided at onsets of 20, 28, and 36 m from the intersection. Driver ratings were the primary performance measure analyzed.

Older drivers showed a strong preference for the earliest warning (36 m), but were willing to tolerate a slightly later warning (28 m) when using the DVI. Younger drivers were more willing to tolerate later warnings for both the DII and DVI, but the latest warning (20 m) was still considered a little too late when using the DII.

To summarize, these results indicated that older drivers would generally prefer the warning up to a second earlier than younger drivers, which makes sense given what we know about aging. These results also suggest that infrastructure-based warnings should come earlier than in-vehicle warnings by about a second; however, this may have been due to the fact that the DVI had a distinct salience advantage due to its auditory component. If both the DII and DVI had been designed as visual only, then the timing preferences might have been similar.

The third and final goal of this analysis was to examine how the location and timing of the warning affected the onset of initial braking. It was expected that this measure might provide insight into both warning effectiveness and warning onset requirements. However, the analysis of initial braking onset provided little insight to warning effectiveness since there was no mean difference between warning or no warning conditions. There may still be an effect such as a reduction in the number of outliers or extremely late braking onsets related to indecision, but this study did not provide enough samples to perform this type of analysis. The analysis of the mean onset of initial braking did, however, show that a decision to stop, both with and without a warning, was initiated approximately 23 m from the intersection, clearly showing that the warnings that were initiated at 20 m (2.5 s) from the intersection came too late to be integrated into the driver's initial decision. Information that is to be provided as decision support (as opposed to a warning) will likely need to be provided on the order of 3.5 to 4.5 seconds before the SV reaches the intersection, and this value may need to be adjusted for higher approach speeds which were not tested.

In conclusion, there was no compelling evidence to suggest that either a DII or a DVI would be either more or less effective for an LTAP-OD warning. This is contrary to the results published by Virginia Tech (Neale, Perez, Lee, and Doerzaph, 2007) regarding a clear advantage by the DVI for a straight-crossing-path, red-light-violation warning. This is due to the fact that there are differences in the assumptions of the causal factors between straight crossing path crashes and the crashes resulting from left-turn decisions. Straight-crossing-path crashes are mostly thought to be related to either distraction or intentional violation. Since drivers are already distracted from looking at the infrastructure, the only effective solution is warning them using an in-vehicle device, specifically an auditory warning. In the LTAP-OD scenario, distraction is not the primary cause of the crash. Drivers are generally already attentive, and therefore an effective visual warning can be given either in the infrastructure or in the vehicle.

Based on these results, there are a number of challenges to designing an effective DVI for an LTAP-OD warning. First, there is the question of whether or not the system should be providing proactive decision support, which would be active during any left-turn maneuver at an equipped intersection, or whether the system is providing a reactive warning, which is only active when the driver is sensed as making a turn that should not be made. While these two system are not mutually exclusive, the DVI design differs depending on the system's intent. Systems that provide proactive information need to be designed to be less intrusive, as drivers do not really want their vehicle to be beeping and buzzing at them every time they try to make a left turn.

Additional challenges for an effective DVI design stem from the problem of location and visual attention. The DVI design cannot compete for the driver's attention while the driver is focused on judging gaps in the oncoming vehicles. However, the DII had a distinct advantage in attracting the driver's attention at the Richmond Field Station intersection because the intersection is fairly small and relatively uncluttered. While in this experiment, drivers preferred the DII, that is not to say that similarly acceptable DVI designs cannot be found. A centrally located, high-mounted display may be an option, and if not, Head-Up Display (HUD) technology has made much progress in recent years and could eventually also be a viable option.

Analysis of Models to Predict SV Turning Rate

This section describes a number of different analyses which were run using the combined data set from the two experiments conducted at the Richmond Field Station. Prior analyses in both this report and in other CICAS-SLTA reports have focused on describing the timing of an intersection approach through the predicted trailing buffer, a metric that was developed during the IDS project. This metric is an attempt to describe the turning scenario and turning decision by accounting for the movements of both the SV and the POV in a single common reference, time. This metric was born from the notion that the real decision being made by drivers in this turning situation is based on some estimate of time, i.e., drivers ask themselves something along the lines of, "Is there enough time to turn in front of this oncoming car?"

One traditional intersection conflict metric discussed in the traffic engineering community is the post-encroachment time. The classic definition of post-encroachment time (Allen, Shin, and Cooper, 1978) is the time from the point where the SV clears the zone of conflict to the time that the POV arrives to the zone of conflict. Thus, the post-encroachment time can only be calculated once the POV reaches the zone of conflict. The trailing buffer metric, as previously defined, is, in essence, a prediction of the post-encroachment time, made at the point where the SV clears the zone of conflict. As long as the POV maintains it's current speed, the trailing buffer will equal the post-encroachment time. However, if the POV decides to brake, perhaps in response to the turning SV, then the trailing buffer will not accurately predict the actual post-encroachment time. It is a subtle distinction, but a distinction none the less.

Using a prediction of the Post-Encroachment Time (predicted PET) as a warning decision criterion has a number of nontrivial issues. First, as suggested by its name, the metric is based entirely on a prediction of driver behavior at any point during the approach. As the situation evolves, the prediction is only as valid as the driver's choice to behave in the predicted manner. Since the predicted post-encroachment time is based on the assumption that the driver is going to attempt to turn in front of the POV, the prediction begins to diverge once the driver decides to stop and let the POV pass (because the driver is no longer following the predicted behavior).

Additionally, although drivers may frame the turning decision in terms of time, humans are not perfect sensors and not perfect information processors. As discussed in Chan, Cody, et al. (2005)Error! Bookmark not defined. there are a number of biases that have been reported in the literature on a driver's ability to make turning decisions, specifically, there is a bias towards insensitivity to the speed of the oncoming POV. The analyses described in this section attempt to take a step back and compare the predicted post-encroachment time to less sophisticated metrics, such as simply looking at POV distance to intersection. In effect, this section examines the influences of the various data components that go into creating the trailing buffer or post-

Deleted:

encroachment prediction to determine which components hold the most predictive power during each stage of the intersection approach.

Method

Many of the analyses in this section are purely descriptive in an attempt to describe how the predicted post-encroachment time, a fairly complex metric, compares to more simple metrics which could be used to predict turning behavior. Since a turning scenario in which both the SV and POV are moving is difficult to describe, there were three snapshots in time taken during the SV's approach to the intersection at the following distances: 36 m, 28 m, and 20 m. The distances were measured, not from the stop bar, but from the point where the SV crossed into the intersection (0 m). Technically, the stop bar would be located at +4 m, and the typical SV stopping point in the intersection would be located at -2 m.

For each snapshot location, the descriptive summaries of the key potential metrics (such as POV distance and speed) are provided and discussed in an order of increasing complexity. From this cursory analysis, metrics that might appear to offer some predictive power in explaining why the SV driver might decide to turn in one situation versus stop in another are then modeled using a general estimating equation regression technique. Similar to the analysis described in Section 4.3 of this report, the probability distribution was assumed to be binomial with a probit link function. Since the metrics are not independent, each metric became its own unique model. The models were then compared on their goodness of fit using the Quasi Likelihood under Independence Model Criterion (QIC) and Corrected Quasi Likelihood under Independence Model Criterion (QICC).

Snapshot: SV is 36 m from the Intersection

Overview

The first snapshot of the LTAP-OD turning scenario was taken when the SV was 36 m from the intersection. Based on the initial braking point distribution (Appendix A), this snapshot was taken before significant braking was ever applied by the driver of the SV. It was also taken a moment before the earliest DII or DVI warning was ever given. It can, therefore, be assumed that this snapshot corresponds to what the situation looked like to the driver perhaps a second before an initial turning decision is typically made.

At this snapshot in time, there were only three independent objective measures that were recorded about the turning scenario: the POV D2I (Distance to Intersection), the POV speed, and SV speed. (The SV D2I was 36 m.) From these three objective measures, two computed or derived estimates could also be created. First, POV D2I and POV speed can be combined to estimate POV T2I (Time to Intersection). Second, POV T2I and SV speed can combined to form the predicted trailing buffer metric which has been discussed in prior work. This results in the following five metrics which may or may not have the ability to help predict whether or not the SV driver will decide to turn in front of the POV:

1 POV D2I
2 POV Speed
3 POV T2I
4 SV Speed
5 Predicted Trailing Buffer (Predicted Post-Encroachment Time)

Descriptive Statistics

Figure 16 shows the distribution of POV D2Is (distances to intersection) for the cases when the SV driver decided to turn and the cases when the SV driver decided to stop. Although there is some overlap between the two distributions, there is a clear separation between the means of the two distributions. Mean POV D2I when the SV driver decided to stop was 65.9 m, and the mean POV D2I when the SV driver decided to turn was 82.6 m for SV turned. The overlap of the distributions is such that the 50th percentile of the decided to stop distribution corresponds to the 5th percentile of the decided to turn distribution, and the 50th percentile of the decided to turn distribution corresponds to the 95th percentile of the decided to stop distribution. This more or less defines the range of indecision about whether to turn or stop as being between the means of the two distributions or between 65 and 83 m from the intersection.

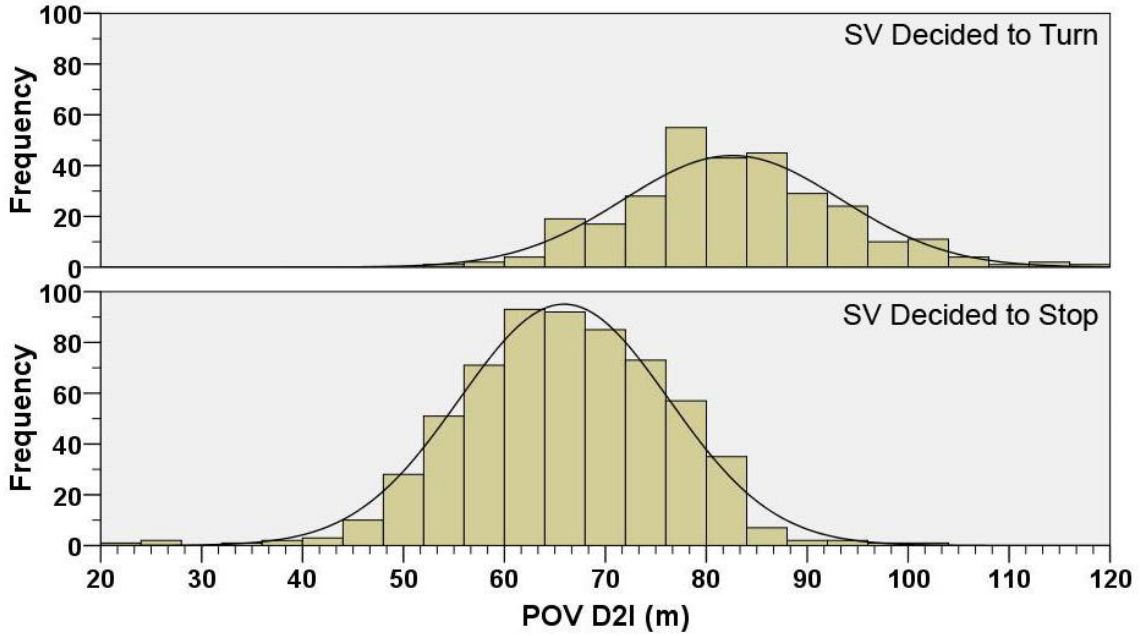

Figure 18. Distribution of POV D2Is by SV turning decision when the SV is 36 m away.

Deleted:

As shown in Figure 17, the distributions of POV speeds did not vary when comparing cases when the SV driver decided to turn against cases when the SV driver decided to stop. Both distributions had a mean POV speed of 10.9 m/s (24.4 mph), which is not surprising since the confederate POV driver was asked to approach at a constant speed of 25 mph. In effect, the range of POV speed variation was only from 10 to 12 m/s (22.4 to 26.8 mph), which would likely have been barely perceptible to the SV driver.

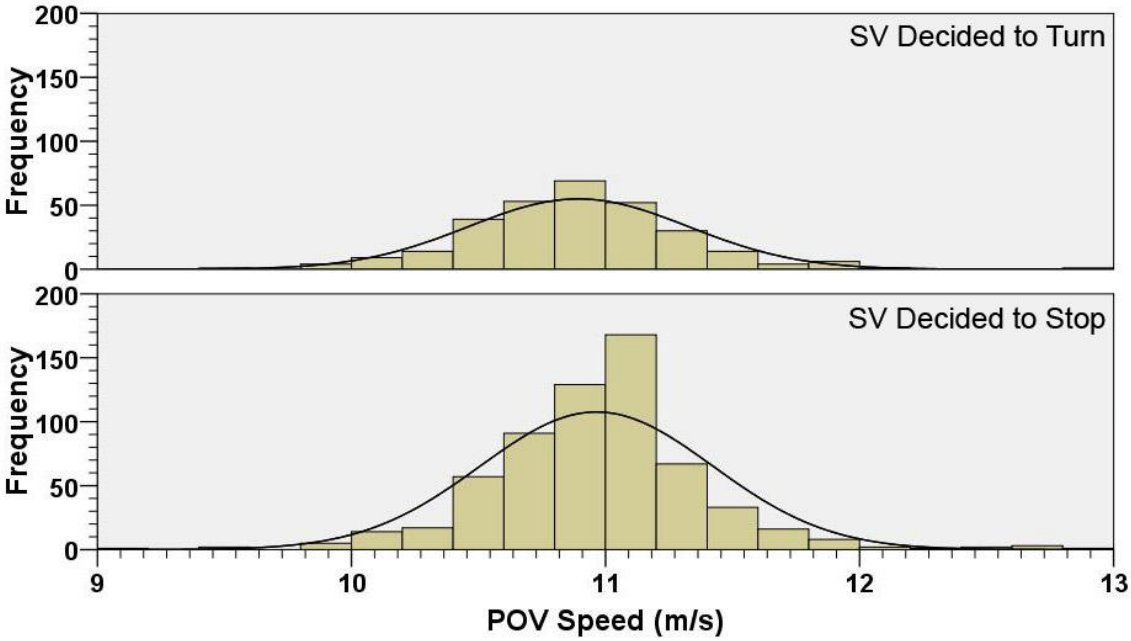

Figure 19. Distribution of POV speeds by SV turning decision when the SV was 36 m away.

Deleted:

Although there was only a 2 m/s variance in the POV speeds, travelling 65 m (the mean POV distance when the SV drivers decided to stop) at 10 m/s yields a T2I (Time to Intersection) of 6.5 seconds. Conversely, travelling 83 m (the mean POV distance when the SV drivers decided to turn) at 12 m/s yields a T2I of 6.9 seconds, or only 0.4 seconds more time. Thus, even though the speed variance from trial to trial was small, it is possible that combining the distance and speed to get T2I would explain more of the variability in the turning response.

Figure 18 shows the distribution of POV T2Is for the cases when the SV driver decided to turn and the cases when the SV driver decided to stop. The mean POV T2I when the SV driver decided to stop was 6 s, while the mean T2I when the SV driver decided to turn was 7.6 s. However, the overlap between the distributions remained relatively consistent with the overlap seen in the D2I distributions. Again, the overlap of the distributions is such that the 50^{th} percentile of the decided to stop distribution corresponds to the 5^{th} percentile of the decided to turn distribution, and the 50^{th} percentile of the decided to turn distribution corresponds to the 95^{th} percentile of the decided to stop distribution. This would tend to suggest that for these cases, adding POV speed to POV distance to form POV T2I will not help much in predicting whether or not drivers will decide to turn.

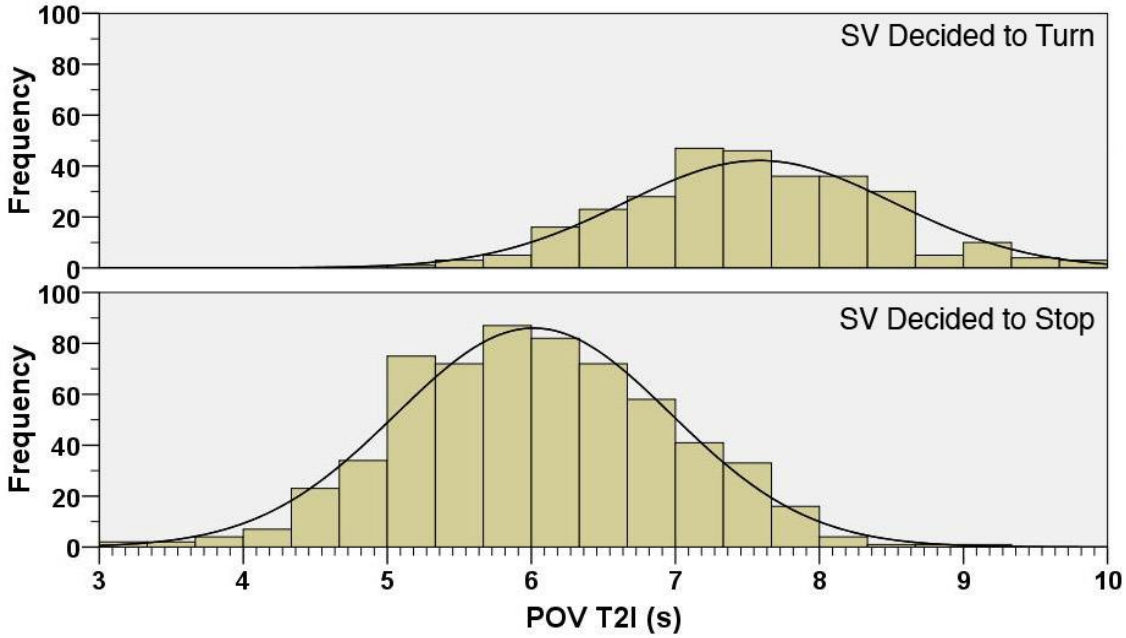

Figure 20. Distribution of POV T2Is by SV turning decision when the SV was 36 m away.

Deleted:

Figure 19 shows the distribution of SV speeds for the cases when the SV driver decided to turn and the cases when the SV driver decided to stop. This is a small difference in the means of these two distributions. When the SV driver eventually decided to turn, the mean SV speed was 9.4 m/s (21 mph), and when the SV driver eventually decided to stop, the mean SV speed was only 9.0 m/s (20.1 mph). Additionally, the distributions overlapped in range nearly perfectly, with the 0.4 m/s offset. Since the SV was at 36 m from the intersection in this snapshot, a mean difference of 0.4 m/s only equates to a quarter second difference in the SV's predicted T2I. Given these two distributions, it is unlikely that the SV's initial speed when the SV is still 36 m from the intersection will play much of a role alone in predicting whether or not the SV driver will decide to turn in front of the POV or not.

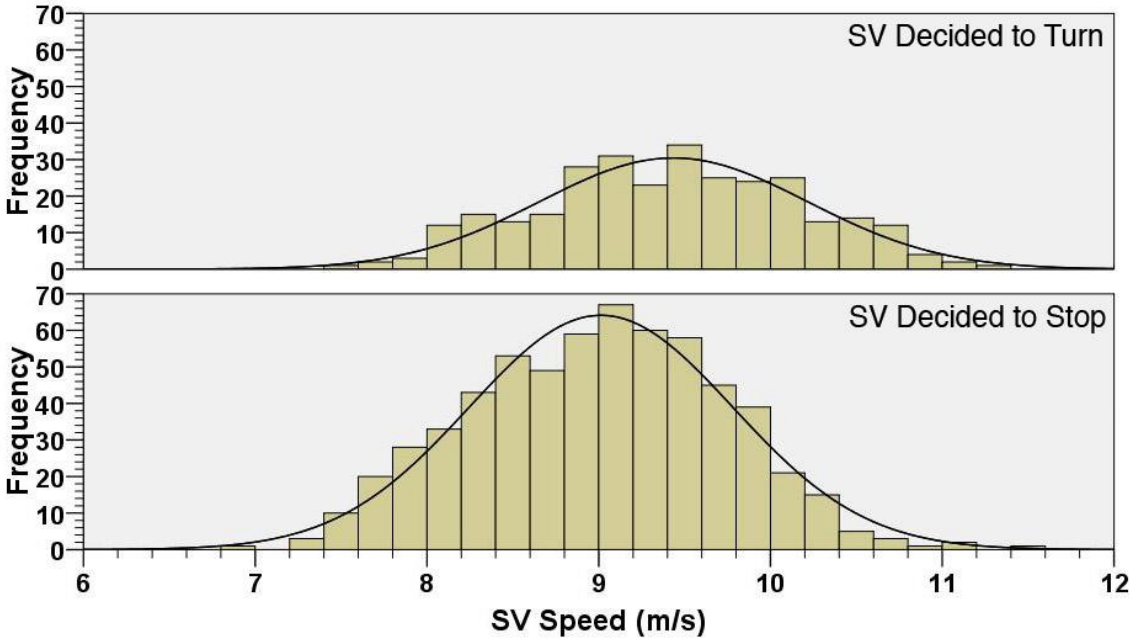

Figure 21. Distribution of SV Speeds by SV turning decision when the SV was 36 m away.

Deleted:

Although the distribution of trailing buffers has been shown before in previous reports, it is repeated below in Figure 20 for easier comparison with the distributions provided in the previous figures in this section.

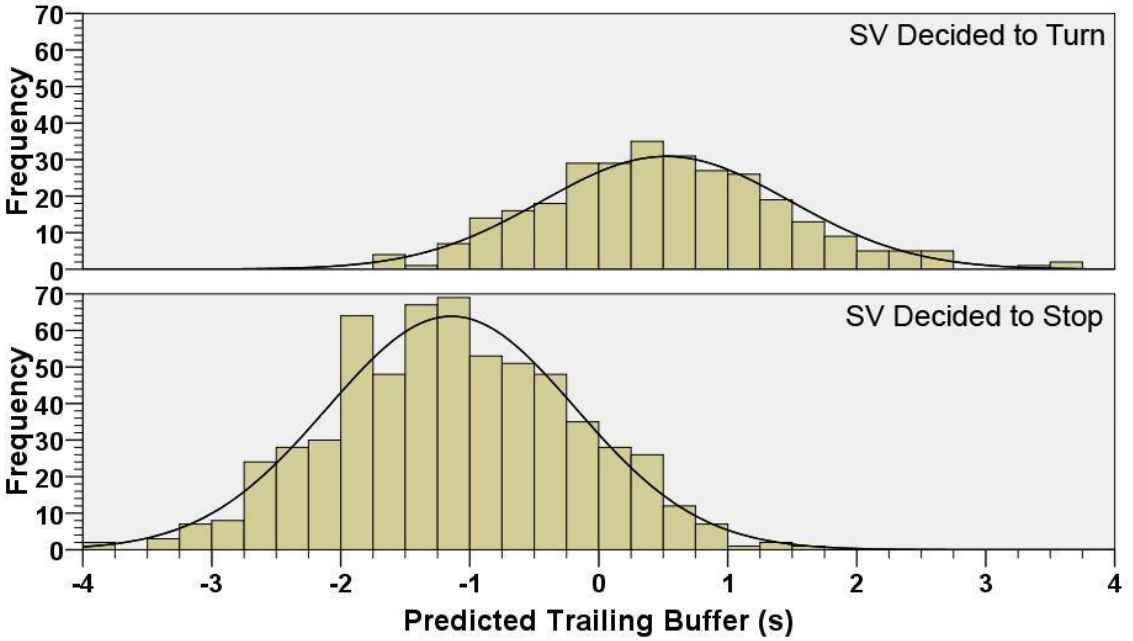

Figure 22. Distribution of predicted trailing buffers by SV turning decision.

Deleted:

There was a noted difference in the means of the distributions. When the SV driver decided to turn in front of the POV, the mean predicted trailing buffer was 0.5 s, while when the SV driver decided to stop, the mean predicted trailing buffer was -1.1 s. However, the spread of the distributions again almost identically mirrored the spread of the distributions for POV D2I. The

50th percentile of the decided to stop distribution corresponds to the 5th percentile of the decided to turn distribution, and the 50th percentile of the decided to turn distribution corresponds to the 95th percentile of the decided to stop distribution, suggesting that the trailing buffer metric may be no better than simply looking at the POV distance.

Modeling Results

Based on the analysis of the five distributions discussed in the previous section, three metrics appeared to have some potential for predicting the driver's decision to turn when the SV was 36 m away from the stop bar: 1) the POV D2I, 2) the POV T2I, and 3) the predicted trailing buffer. To complete this analysis, three repeated measures general estimating equation models were created by combining each of the three metrics (treated as a within subjects factor) with the additional factor of driver age (a between subjects factor). As a reminder, warnings were not an issue here because at this distance, no warnings had yet appeared.

The results of these models are listed below in Table 3. Age was not significant for any of the models; however, in all three cases, the primary metric was significant. Interestingly, the goodness of fit statistic (QIC) improved as the metric complexity increased. (Smaller QIC values are better.) Thus, using the predicted trailing buffer metric was slightly more predictive of the SV turning behavior than simply looking at POV T2I or POV D2I.

Table 3.

Model	Significance Test		Goodness of Fit	
	Wald X^2_1	p	QIC	QICC
POV D2I	69.6	< .001	782.1	741.7
POV T2I	94.1	< .001	765.3	722.1
Predicted Trailing Buffer	118.2	< .001	735.6	690.0

Discussion

The goodness of fit statistic is not necessarily a very strong test for comparing models. Similar to looking at an r^2 statistic, there is no set guideline for saying that one model is a good model while another model is a bad model. The statistic simply shows that one model appears to explain more variance than another. While there did appear to be a slight advantage in the power to predict whether or not drivers would choose to turn or stop when using the predicted trailing buffer metric (over the simpler metrics), it should be noted that several of the factors used in computing the predicted trailing buffer metric were held relatively constant due to the nature of experiments from which the data were gathered. The POV speed was fixed at 24±2 mph and the SV speed was fixed at 20±2 mph. Were these speeds allowed to vary like they do in the real world, by as much as 10 mph, then the predicted trailing buffer or other time based measures would probably be much better than simply looking at the POV distance.

Snapshot: SV is 28 m from the Intersection

Overview

The second snapshot of the LTAP-OD turning scenario was taken when the SV was 28 m from the intersection. However, one of the problems with taking this snapshot is that there are a number of confounding effects that start to enter into play. First, at a distance of 28 m, some of

the trials, approximately 23 percent, will already have been influenced by either the DII or the DVI when it activated at 36 m from the intersection. Although it was debated as to whether or not to remove these trials from the analysis, it was decided to leave them in because drivers were free to ignore the warnings (and did so about 17 percent of the time).

Second, the mean initial braking point was 25.6 m. By the time the SV reached the 28 m snapshot, the initial braking had already begun on 25 percent of the trials. Although the initial braking point does not necessarily correspond to a turning decision having been made, the trend was for the initial braking point to come later when the decision was to turn in front of the POV. In effect, the snapshot taken at this location is more than likely confounded with behavior taken after the initial turning decision had been made, at least on up to 25 percent of the trials. However, for 75 percent of the trials, this snapshot can be thought to show what the situation looked like just before or right as a typical initial turning decision was being made.

Similar to the last section, five metrics were available which may or may not have the ability to help predict whether or not the SV driver will decide to turn in front of the POV. Of those five metrics, only the following four are discussed because POV speed, as in the last analysis, remained relatively constant.

1. POV D2I
2. POV T2I
3. SV Speed
4. Predicted Trailing Buffer (Predicted Post-Encroachment Time)

Descriptive Statistics

Figure 21 shows the distribution of POV D2Is (distances to intersection) for the cases when the SV driver decided to turn and the cases when the SV driver decided to stop. The mean POV D2I when the SV driver decided to stop was 56.3 m, while the mean POV D2I when the SV driver decided to turn was 73.7 m. These means are relatively similar to what was seen in the snapshot of the SV at 36 m from the intersection, except now both of the vehicles are roughly 8-10 m closer to the intersection. The spread of the distribution also matched what was seen in the previous snapshot. The range of indecision is still between the two distribution means.

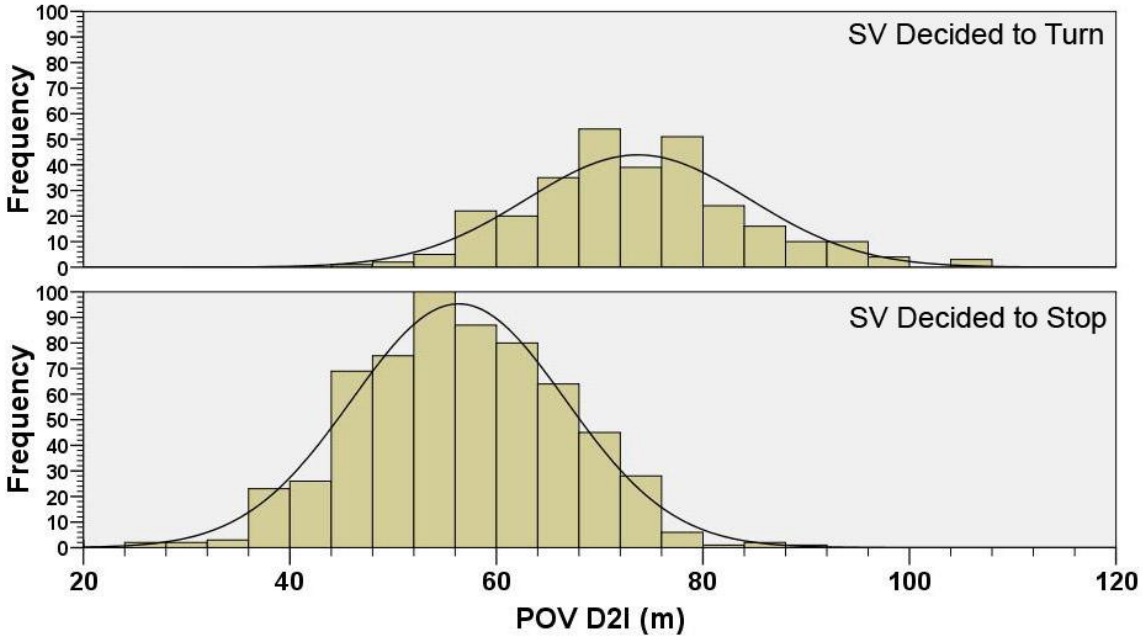

Figure 23. Distribution of POV D2Is by SV turning decision when the SV is 28 m away.

Deleted:

The POV speed distribution means and spreads were also identical between decisions to turn and decisions to stop, so the detailed analysis is not shown. Ninety percent of the POVs were travelling between 10.2 and 11.6 m/s (22.8 and 25.9 mph). Figure 22 shows the resulting distributions when POV D2I and POV speed were combined to yield POV T2I. Compared to the previous snapshot, the means have decreased by just under a second. The mean POV T2I when the SV driver decided to stop was 5.1 s, while the mean T2I when the SV driver decided to turn was 6.8 s. The overlap of the distributions is still roughly the same as the overlap seen during the previous snapshot.

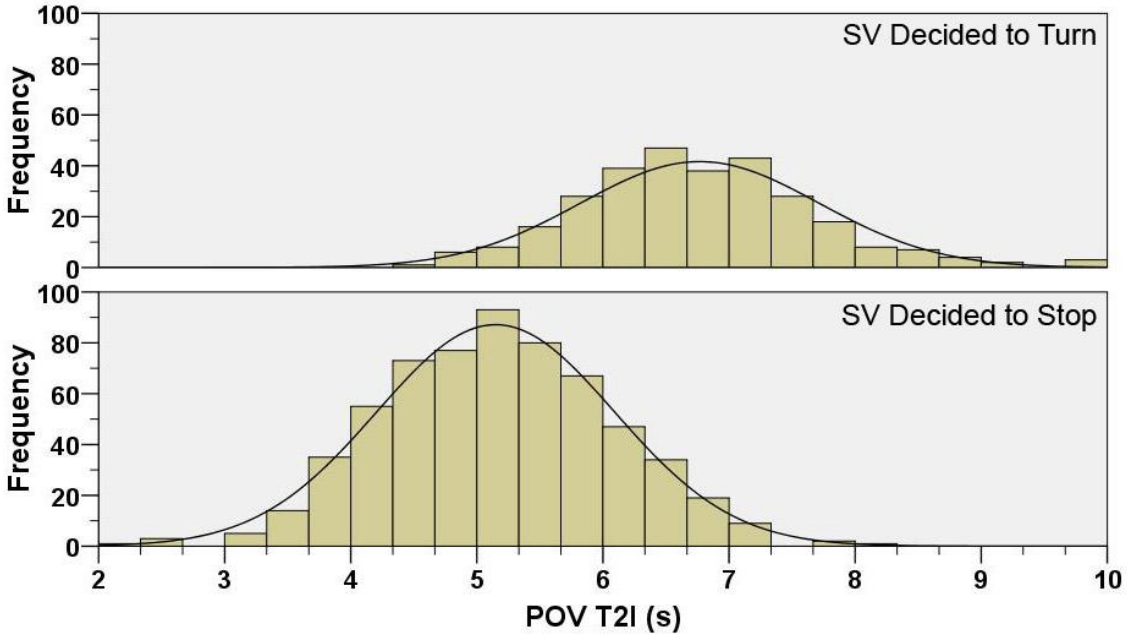

Figure 24. Distribution of POV T2Is by SV turning decision when the SV was 28 m away.

Deleted:

Figure 23 shows the distribution of SV speeds for cases when the SV driver decided to turn compared to cases when the driver decided to stop. Unlike the previous snapshot when the SV was at 36 m from the intersection, this snapshot clearly shows a difference in SV speeds developing between the two cases. In the cases when the driver decided to turn, the SV was travelling, on average, 0.5 m/s (1.1 mph) faster. The mean for the stopped distribution was 8.9 m/s (19.9 mph), while the mean for the turned distribution was 9.4 m/s (21 mph).

Figure 24 shows the distribution of predicted trailing buffers for cases when the SV driver decided to turn and cases when the driver decided to stop. Since the predicted trailing buffer is supposed to be a metric that is relatively invariant throughout the approach, it would be expected that the distributions would look nearly identical to the previous snapshot. In fact, when the SV driver decided to turn, the mean predicted trailing buffer was 0.5 s which was identical to the previous snapshot. When the SV driver decided to stop, the mean predicted trailing buffer was also identical to the previous snapshot at -1.1 s.

In effect, although the distribution of SV speeds taken during this snapshot showed that a difference in behavior between drivers who would eventually decide to turn and drivers who would eventually decide to stop, these subtle differences did not yet play a large role in influencing the outcome of the trailing buffer metric. In fact, the predicted post-encroachment time that was calculated by the trailing buffer metric remained relatively constant between snapshots, as evident by Figure 25 which shows the distribution of differences between the two snapshot. The mean change in the predicted trailing buffer from the 36 m snapshot to the 28 m snapshot was on the order of 0.02 s and 90 percent of the differences were under 0.2 s.

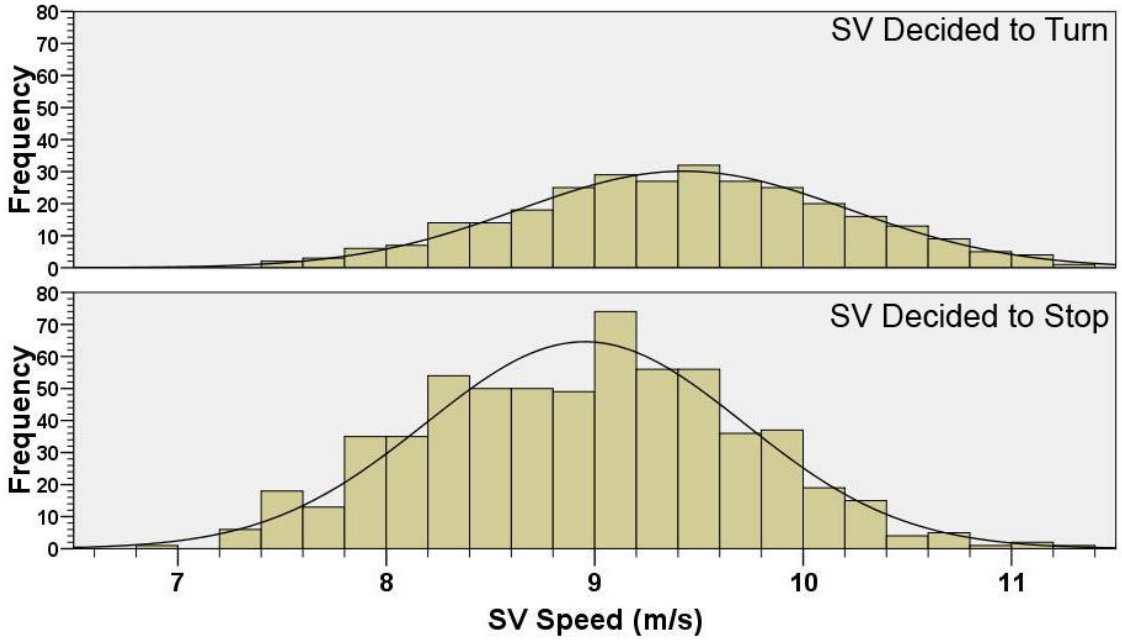

Figure 25. Distribution of SV Speeds by SV turning decision when the SV was 28 m away.

Deleted:

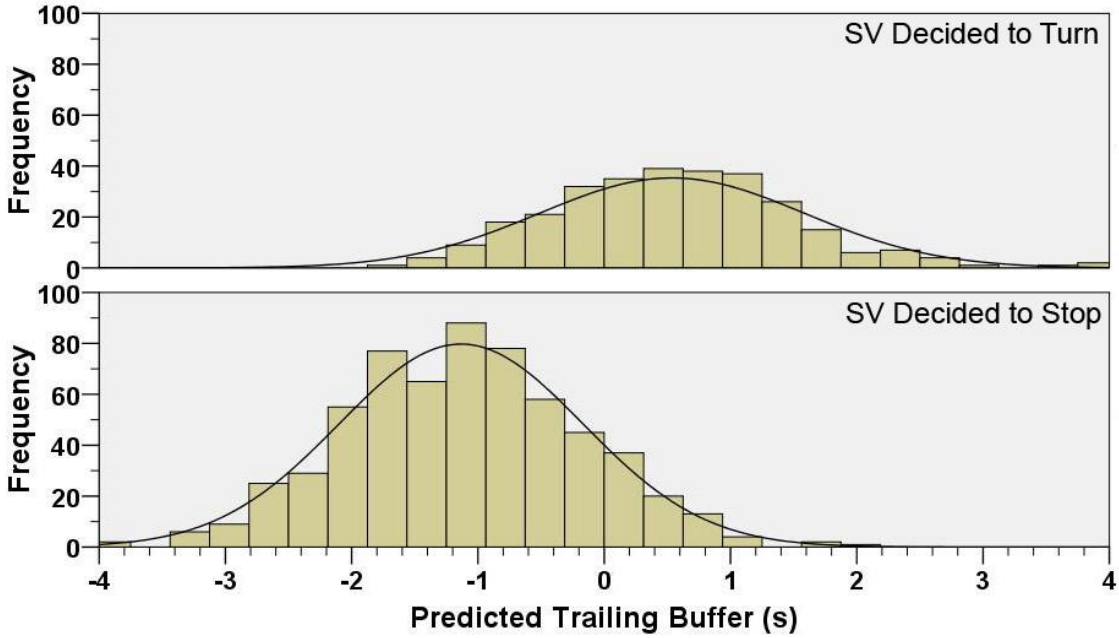

Figure 26. Distribution of predicted trailing buffers by SV turning decision.

Deleted:

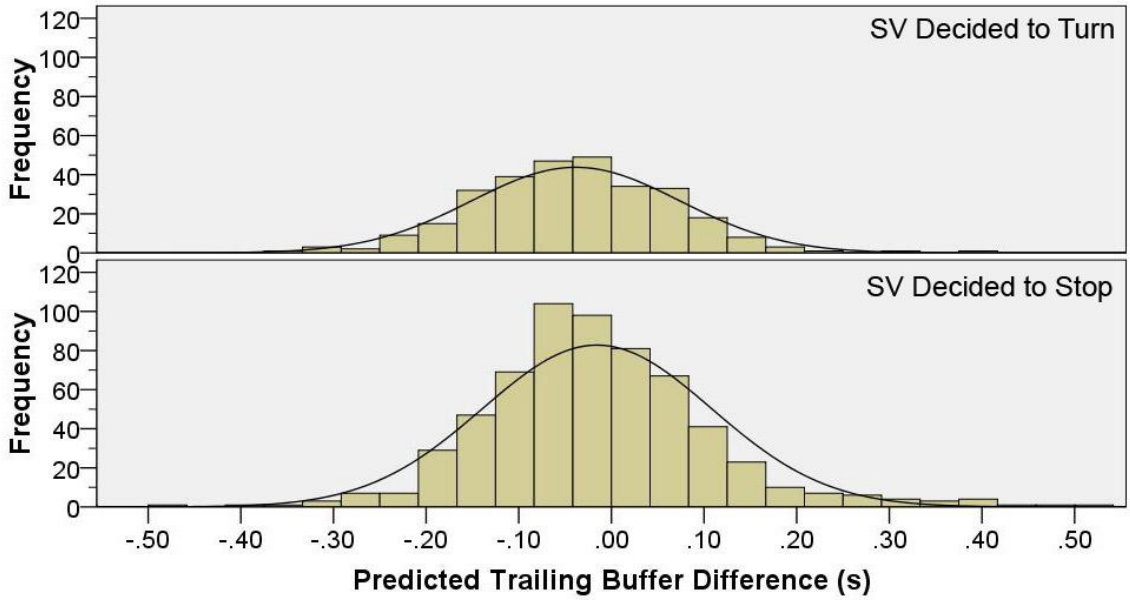

Figure 27. Distributions of differences between predicted trailing buffers at 36 and 28 m.

Deleted:

Modeling Results

Based on the analysis of the four distributions discussed, again, all three metrics (POV D2I, POV T2I, and predicted trailing buffer) appeared to have some potential for predicting the driver's decision to turn when the SV was 28 m away from the stop bar. Similar to the analysis in Section 4.2, three repeated measures general estimating equation models were created by combining each of the three metrics (treated as a within subjects factor) with the additional factor of driver age (a between subjects factor).

The results of these models are listed below in Table 4. Age was not significant for any of the models; however, in all three cases, the primary metric was significant. The goodness of fit

statistic (QIC) still improved as the metric complexity increased. Thus, using the predicted trailing buffer metric at this distance from the intersection was slightly more predictive of the SV turning behavior than simply looking at POV T2I or POV D2I.

Table 4.

Model	Significance Test		Goodness of Fit	
	Wald X^2_1	p	QIC	QICC
POV D2I	76.4	< .001	756.3	715.6
POV T2I	93.9	< .001	741.0	695.6
Predicted Trailing Buffer	42.1	< .001	724.6	732.6

Discussion

The snapshot discussed in this section was taken when the SV was 28 m from the intersection. On all counts, the results were fairly similar to the snapshot taken at 36 m from the intersection. The only major difference was that at 28 m from the intersection, there started to be a noticeable half-second difference in mean SV speed when looking at the speed distributions for cases when the driver eventually decided to turn and cases when the driver eventually decided to stop. Similar to what was seen in the previous snapshot, the predicted trailing buffer was still the best metric for predicting whether the driver would turn or stop. Additionally, the predicted trailing buffer metric did not appear to change much between the previous snapshot and the current snapshot.

Snapshot: SV is 20 m from the Intersection

The third snapshot of the LTAP-OD turning scenario was taken when the SV was 20 m from the intersection. Similar to the previous snapshot some of the trials will already have been influenced by either the DII or the DVI when it activated at 36 or 28 m from the intersection. In this snapshot, either a DII or DVI had already been shown to the driver 50 percent of the time (with a 17 percent rate of non-compliance).

However, more importantly, when the SV was at 20 m from the intersection, it had already passed the mean initial braking point (which ranged from 20 to 22 m from the intersection). In fact, by the time the SV reached the 20 m snapshot, the initial braking had already begun on almost 75 percent of the trials. It is highly likely that this snapshot can be interpreted as what the situation looked like up to a half-second after the driver made his or her initial decision to turn or stop.

Similar to the last section, POV speed was ignored as it remained relatively constant throughout the approach. The analyses focused on POV D2I, POV T2I, SV speed, and predicted trailing buffer.

Descriptive Statistics

Figure 26 shows the distribution of POV D2Is (distances to intersection) for the cases when the SV driver decided to turn and the cases when the SV driver decided to stop. The mean POV D2I when the SV driver decided to stop was 47.1 m, while the POV D2I when the SV driver decided to turn was 65.0 m. Again, these means are relatively similar to what was seen in the snapshot of the SV at 28 m from the intersection, except now both of the vehicles are roughly 8.5-9.5 m closer to the intersection, which makes sense since both vehicles are travelling roughly 8 to 11

m/s for about 1 s in between the two successive snapshots. The spread of the distribution also matched what was seen in the previous two snapshots, with the range of indecision still being between the two distribution means.

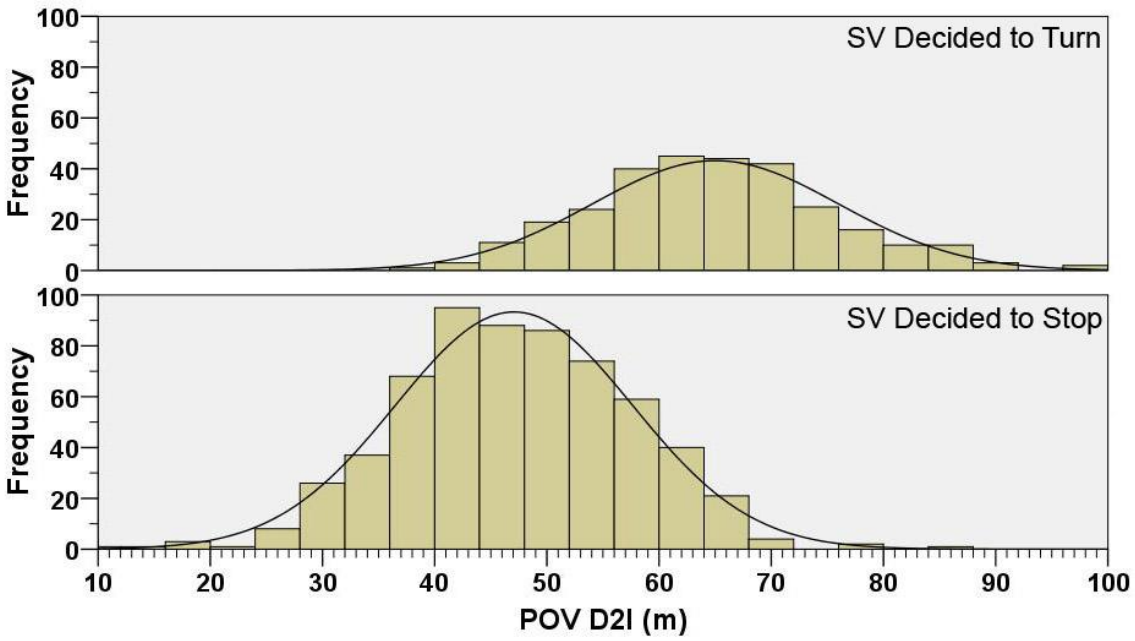

Figure 28. Distribution of POV D2Is by SV turning decision when the SV is 20 m away.

Deleted:

Figure 27 shows distributions of POV T2Is. Compared to the previous snapshot, the means have decreased by just under a second. The mean POV T2I when the SV driver decided to stop was 4.3 s, while the mean T2I when the SV driver decided to turn was 5.9 s. The overlap of the distributions is still roughly the same as the overlap seen during the previous snapshot.

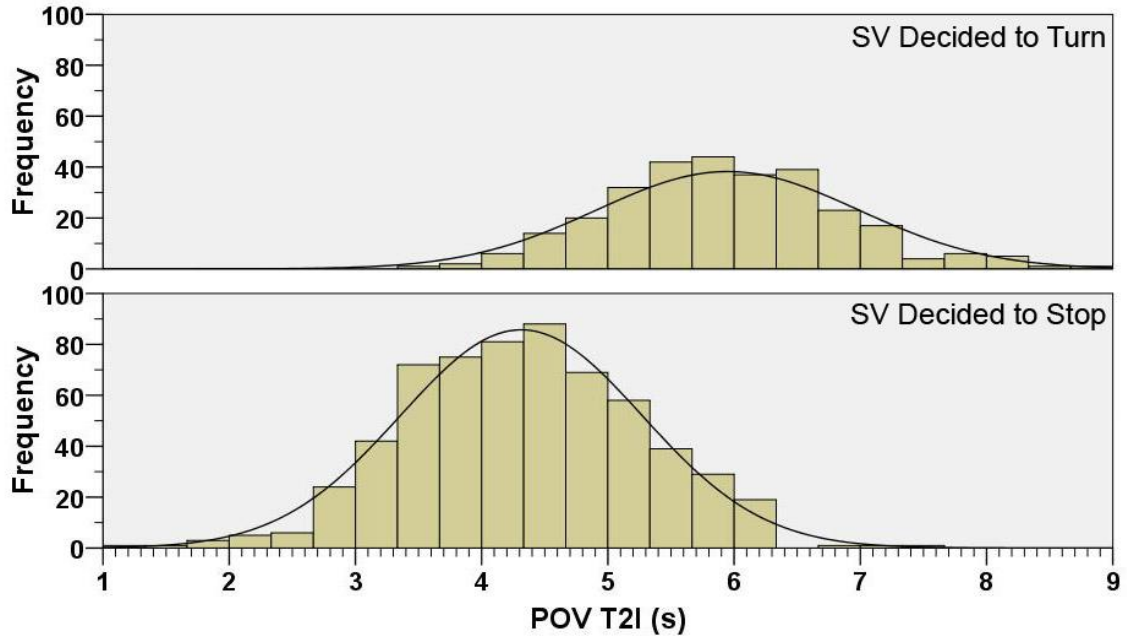

Figure 29. Distribution of POV T2Is by SV turning decision when the SV was 20 m away.

Deleted:

Figure 28 shows the distribution of SV speeds for cases when the SV driver decided to turn compared to cases there the driver decided to stop. Similar to the previous snapshot taken when the SV was at 28 m, there is a growing difference between the means of cases when the SV driver decided to turn and cases when the SV driver decided to stop. In the previous snapshot, the difference was 0.5 m/s, but in the current snapshot the difference has grown to 0.6 m/s. The mean SV speeds for the stopped and turned distributions were 8.7 m/s (19.4 mph) and 9.3 m/s (20.8 mph), relatively unchanged from the previous snapshot. This would suggest that although the drivers were on the brakes by this point, significant slowing had not yet occurred.

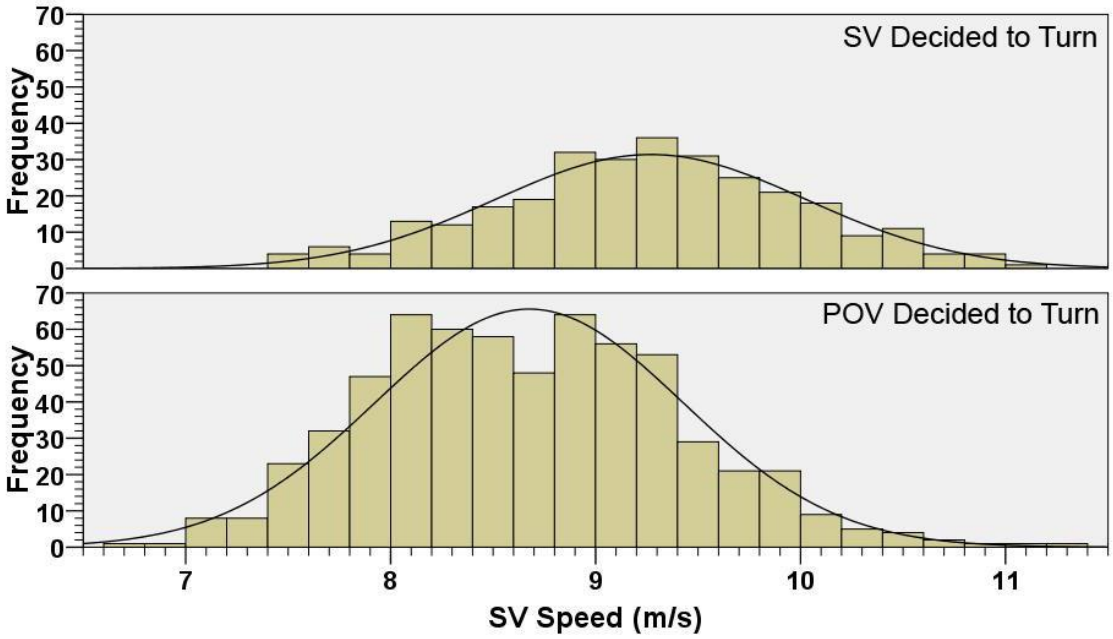

Figure 30. Distribution of SV Speeds by SV turning decision when the SV was 20 m away.

Deleted:

Consistent with the assumption that the predicted trailing buffer will remain invariant throughout the approach, the predicted trailing buffer distributions shown in Figure 29 appear nearly identical to the previous two snapshots. When the SV driver decided to turn, the mean predicted trailing buffer was 0.48 s (which was identical to the previous snapshot's 0.5 s), and when the SV driver decided to stop, the mean predicted trailing buffer was -1.2 s (again, nearly identical to the -1.1 s in the previous snapshots). Similarly, as shown in Figure 30, the mean difference in the predicted trailing buffer between the snapshots taken at 35 and 20 m was on the order of 0.05 when the SV decided to stop and 0.01 s when the SV decided to turn.

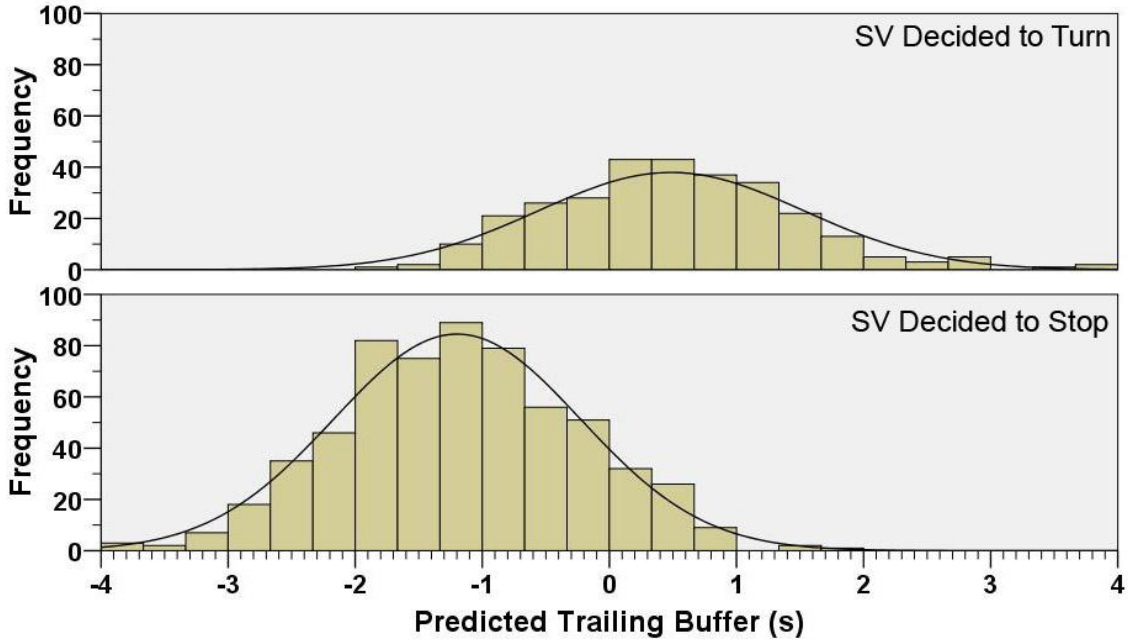

Figure 31. Distribution of predicted trailing buffers by SV turning decision.

Deleted:

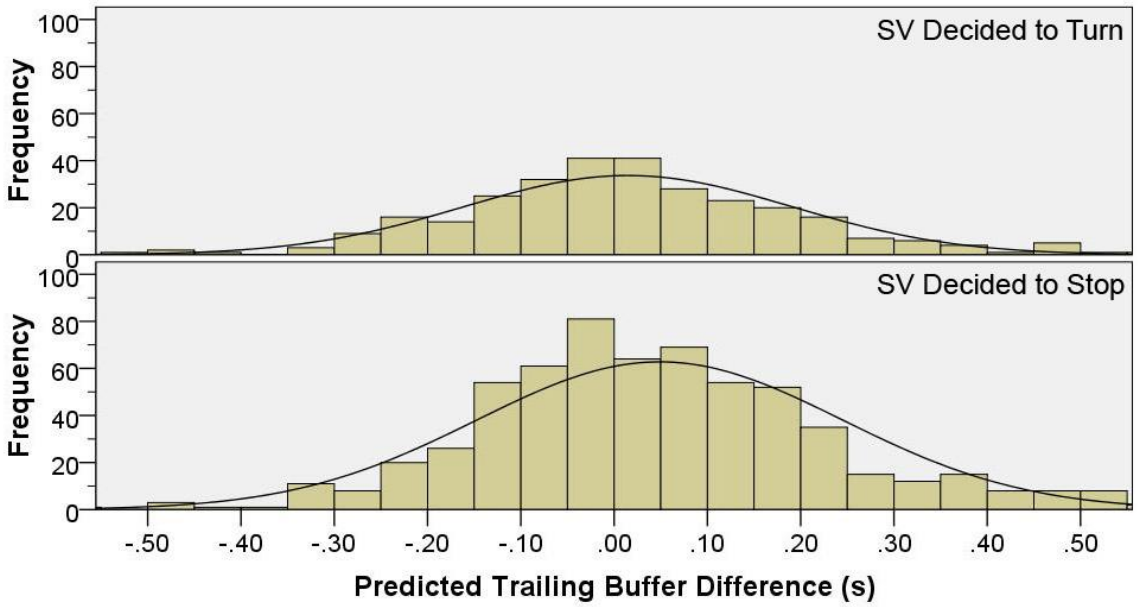

Figure 32. Distributions of differences between predicted trailing buffers at 36 and 20 m.

Deleted:

Since the predicted trailing buffer is based on the assumption that the SV driver will try to turn, it would be expected that once a driver decides to stop and executes that decision, the metric will no longer remain an invariant. As the SV slows to a stop, the predicted trailing buffer should decrease resulting in a larger difference when comparing the prediction made at 36 m with the prediction made at 20 m. This is exactly what the data appear to start showing in this snapshot. The distribution of differences in the prediction between snapshots for cases when the SV decided to stop is starting to shift to the right, while the distribution for the cases when the SV decided to turn (and is thus following the predicted path) remains centered around zero.

Modeling Results

Similar to the analyses in Sections 4.2 and 4.3, three repeated measures general estimating equation models were created by combining each of the three metrics, POV D2I, POV T2I, and predicted trailing buffer, (each being treated as a within subjects factor) with the additional factor of driver age (a between subjects factor).

The results of these models are listed below in Table 5. Age was, again, not significant for any of the models; however, in all three cases, the primary metric was significant. The goodness of fit statistic (QIC) still improved as the metric complexity increased. Thus, using the predicted trailing buffer metric at this distance from the intersection was slightly more predictive of the SV turning behavior than simply looking at POV T2I or POV D2I.

Table 5.

Model	Significance Test		Goodness of Fit	
	Wald X^2_1	p	QIC	QICC
POV D2I	83.7	< .001	744.0	709.8
POV T2I	42.6	< .001	728.3	736.3
Predicted Trailing Buffer	47.6	< .001	714.9	722.9

Discussion

The snapshot discussed in this section was taken when the SV was 20 m from the intersection. On all counts, the results were fairly similar to the snapshots taken at both 36 and 28 m from the intersection. Like the snapshot that was taken at 28 m from the intersection, there remained a noticeable half-second difference in mean SV speed when looking at the speed distributions for cases when the driver eventually decided to turn and cases when the driver eventually decided to stop. Similar to what was seen in the previous snapshots, the predicted trailing buffer was still the best metric for predicting whether the driver would turn or stop. However, when the SV has reached this distance to the intersection, the estimate of the predicted trailing buffer begins to diverge from previous estimates for the cases when the driver eventually decides to stop. The predicted trailing buffer metric is only invariant up to the point when the SV driver decides to stop. After that point, the driver is no longer following the predicted path, and therefore, the metric diverges.

Interim Conclusions

One of the most difficult challenges in regards to the LTAP-OD crash problem comes from the fact that there are no standard metrics which can be used to describe a turning maneuver where both approaching vehicles are moving. The gap or lag between vehicles has traditionally only been measured when the vehicle attempting to make the turn was already stopped and waiting. Prior work in the CICAS-SLTA program has suggested a metric referred to as the predicted trailing buffer or predicted post-encroachment time, which is basically an estimate of the spare time that would remain before a POV would reach the intersection once the SV cleared the intersection. The goal of the analyses described in this section was to examine the components of the predicted trailing buffer or post-encroachment metric to determine the predictive power of each component.

The predicted trailing buffer or post-encroachment time combines the SV and POV speeds and distances along with an assumption of their intended movements to predict the likely outcome,

should the SV decide to turn. The main concern with the methodology behind this metric lies in the necessity of prediction. In essence, the question that has been asked is whether or not all of this prediction really results in a metric that has more predictive power over any of the component or raw data measurements used to describe the situation. To analyze this question, three static snapshots in time were taken when the SV was 36, 28, and 20 m from the intersection. By freezing the location of the SV, the relative predictive power of four different models could be compared.

The results showed that the largest contribution from a single component came from the POV distance to intersection. The POV speed alone played very little role in predicting whether or not the SV would turn; however, it should be noted that the POV speed was relatively fixed during the data collection. It only varied by about ±2 mph within one standard deviation. Nonetheless, for all SV snapshots, combining POV speed and POV distance to produce POV time to intersection resulted in a slightly more predictive model. Furthermore, combining SV speed and distance with POV speed and distance to produce the predicted post-encroachment time metric was the most predictive model in explaining SV drivers' decisions to turn. These results held true for all of the different snapshots taken at the various SV distances to the intersection. In all cases, the predicted trailing buffer or post-encroachment time metric turned out to be the model with the most predictive power.

Characteristics of The Predicted Post-Encroachment Metric

The analysis described in the prior saection concluded that the Predicted Post-Encroachment Time (PET) metric appeared, statistically, to be the best model to predict SV turning behavior throughout the SV approach. This section describes and illustrates the characteristics of the predicted PET metric using data gathered during the two experiments conducted at the Richmond Field Station during the IDS project.

The analyses in the prior sections of this report used a first version of the predicted PET metric based on a calculation of a prediction of the trailing buffer. This metric assumed that the SV would keep a constant speed until it needed to decelerate to a mean turning speed. While this methodology works well when the SV is far from the intersection, it did not address what happens as the SV reaches the intersection and begins its deceleration. In order to complete the analyses described in this section, a second generation of a predicted PET metric needed to be created that would be able to handle the SV's movement all the way to the intersection.

Additionally, the first version of the predicted PET metric assumed that the SV cleared the intersection when it cleared the pedestrian crosswalk. In the version of the predicted PET used in this section, the SV was said to have cleared the Zone of Conflict (ZoC) with the POV as soon as the rear bumper of the SV cleared the POV's lane. This shortened the time that the SV needed to travel to clear the ZoC by approximately 0.5 s. Thus, the predicted PET values used in this section are on the order of 0.5 s greater than the predicted PET values that were used in Sections 4 and 5. The actual calculation method used for the second generation predicted PET that is used the analyses in this section is discussed in later in this report.

The figures in this section typically only plot a subset of the trials collected during experiments conducted at the Richmond Field Station in order to maintain an appropriate level of clarity for interpretation. The plots typically start when the SV was between 36 or 20 m from the intersection (corresponding to the range where DII or DVI warnings were given to the driver) and end when the SV reached the intersection (the 0-point on the graphs). This is because the

radar tracking the SV during the experiments typically lost the SV as a target when it was near the 0 point on the graph or when the SV speed dropped below approximately 2 m/s (5 mph).

For reference, the 0-point shown on the graphs is approximately 3 m before entering, and 13 m before clearing, the ZoC with the POV. Thus, the point at which the SV's rear bumper cleared the ZoC would be -13 m on the graph. The stop bar in the SV's lane would be located at 4 m on the graph.

Analysis of the Predicted PET Using Data from the RFS Intersection

Overview – Cases When it is Clear that the SV Can Turn or Should Stop

Figure 31 shows how the predicted Post-Encroachment Time (PET) changes as the SV approaches the intersection. When the PET was above 1.5 s (shown in blue), the SV driver almost always chose to turn in front of the POV. When the PET was below -1.0 s (shown in red), the SV driver always chose to stop. When the PET was between -1.0 and 1.5 s (data not shown on this graph), the SV driver sometimes to chose to turn and sometimes chose to stop. These cases will be discussed in subsequent sections.

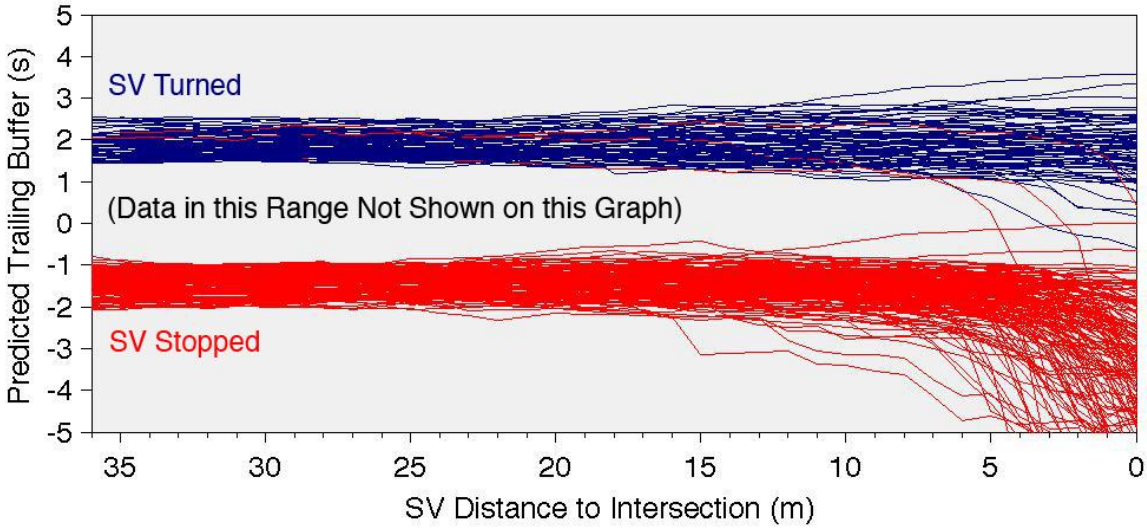

Figure 33. Example of tracking the predicted PET during the SV approach.

Deleted:

As seen in the Figure 31, the predicted PET stays relatively constant throughout the SV approach (from 36 m down to about 5 m) because both the SV and POV are, more or less, following the trajectories prescribed by the predicted PET algorithm. When the SV driver decided to turn (cases in blue), the predicted PET continues to remain invariant throughout the turn. However, when the SV driver decided to stop and let the POV pass, the predicted PET decreases asymptotically since the metric is now predicting what would happen if the SV driver decided to reaccelerate and attempt a turn in front of the POV. This can be seen more clearly in Figure 32, which shows plots of three cases when the SV driver started with plenty of time to turn in front of the POV (predicted PETs between 1.5 and 2 s), but eventually decided to stop.

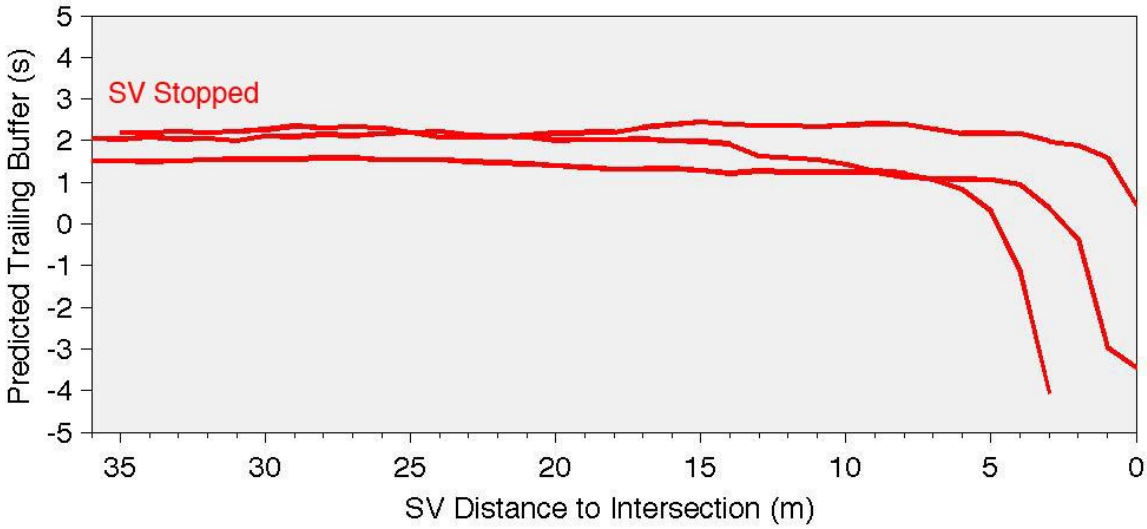

Figure 34. Evolution of the predicted PET for cases when the SV decided to stop.

Deleted:

Cases When the Predicted PET is Between -1.0 and -0.75 s

Figure 33 shows the evolution of the predicted PET for cases when the predicted PET was first calculated to be between -1.0 and -0.75 seconds when the SV was 36 m from the intersection (around the time that decision support information might be provided to the driver). As the scenario evolves, although the predicted PET remains relatively constant throughout the approach, it does change. By the time the SV reaches 10 to 15 m from the intersection, the predicted PET now varies from -2.0 to -0.5 s. As shown in this graph, only three turns were actually made when the predicted PET started in this range. In each of those cases, the predicted PET increased as the scenario evolved, ending with a predicted PET at 0 of nearly -0.5 s. Although this still indicates a likely collision, further analysis can be found in Figure 34, which shows the actual speed profile of the approaching SV.

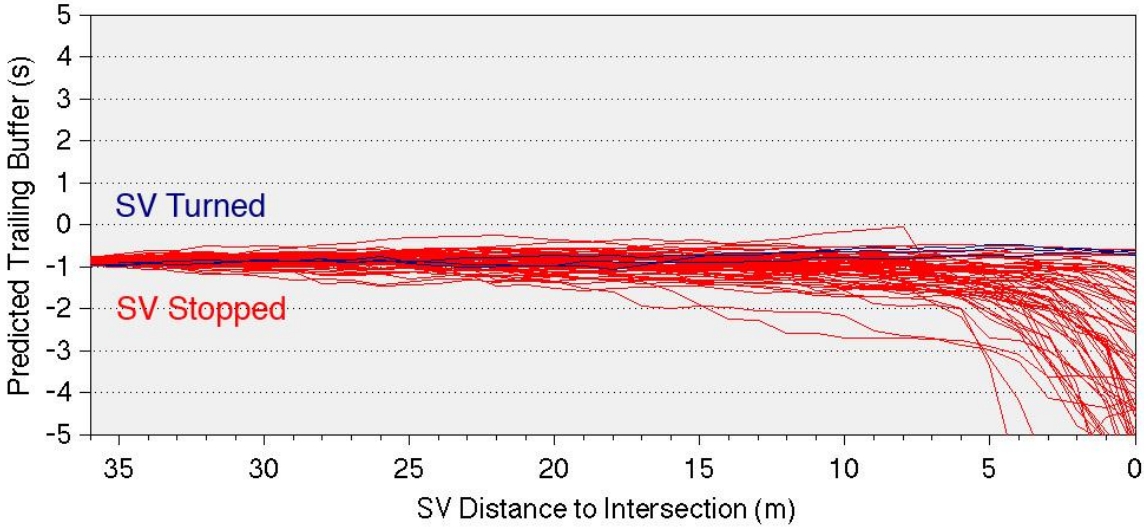

Figure 35. Evolution of the predicted PET when it started between -1.0 and -0.75 s.

Deleted:

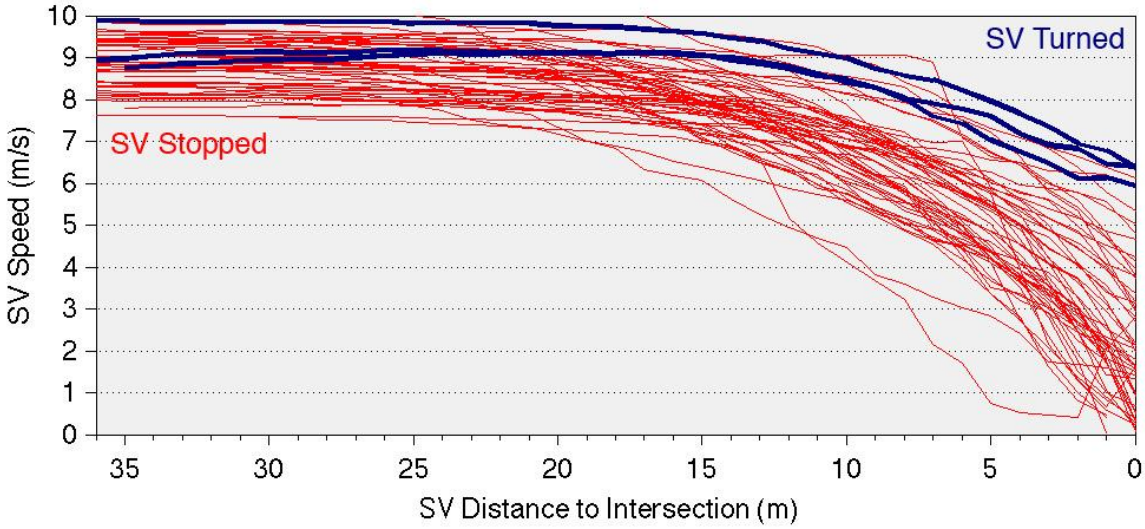

Figure 36. SV speed profile for cases when the predicted PET started between -1.0 and -0.75 s.

Deleted:

As shown in Figure 34, for the three cases when the SV actually decided to turn, even though the predicted PET was tracking between -1.0 and -0.5 s (a likely collision), the SV was travelling on the faster end of the normal speed range, and likely turned with a speed in the range of 6.0 to 6.5 m/s, rather than the predicted mean of 5.89 m/s. In these cases, the faster turning speed could account for up to a 0.25 seconds of prediction error. Other time saving strategies, such as cutting the corner, might also account for some of the prediction error, and there is always the possibility of measurement errors. However, the sensing technology used to generate the data set during the experiments failed to produce a reliable measure of the actual PET for validation. Thus, any conjecture drawn on either the appropriateness of these turn or potential errors in the predicted PET metric, are still inconclusive.

Cases When the Predicted PET is Between -0.75 and 0.0 s

Figures 5-5, 5-6, and 5-7 show the evolution of the predicted PET metric for cases when it started between -0.75 and 0.0 s at increasing quarter-second intervals. As the starting predicted PET increases from graph to graph, the number of SV decisions to turn (shown in blue) increases. These three graphs were grouped together in this section because, given the interpretation of the predicted PET metric, a predicted PET of 0.0 s would be an obvious cut-off point for determining which turning scenarios would be safe and which turning scenarios would be unsafe. Given this warning criterion, all of the cases shown in this section would receive some indication that the turn was unsafe. Turns in front of the POV were made in only 17 percent of the cases. However, as discussed previously, without validation of the predicted PET metric, a true determination of whether the turns made by the SV drivers in these situations were actually OK (and would be considered false alarms if a warning was provided) or if these were actually turns that should have been prevented, is still unknown.

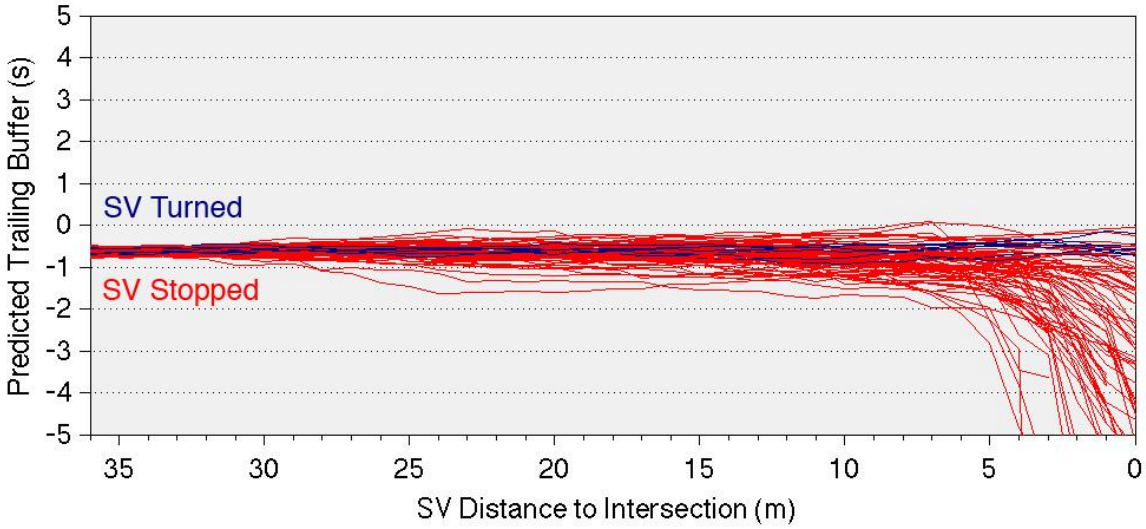

Figure 37. Evolution of the predicted PET when it started between -0.75 and -0.5 s.

Deleted:

There are number of cases when the predicted PET increases dramatically as the SV approaches the intersection (particularly in Figure 36 and 37). Since the POV was not braking during the experiment, these predicted PET increases were due the SV accelerating and reaching speeds on the order of 12 m/s (25 mph) instead of the nominal 9 m/s (20 mph) that was more typical during the experiment. The drivers that accelerated, ended up braking later and harder than prescribed by the predicted PET metric, and consequently, entered the turn at speeds between 6 and 7 m/s, also higher than the 5.89 m/s prescribed by predicted PET metric.

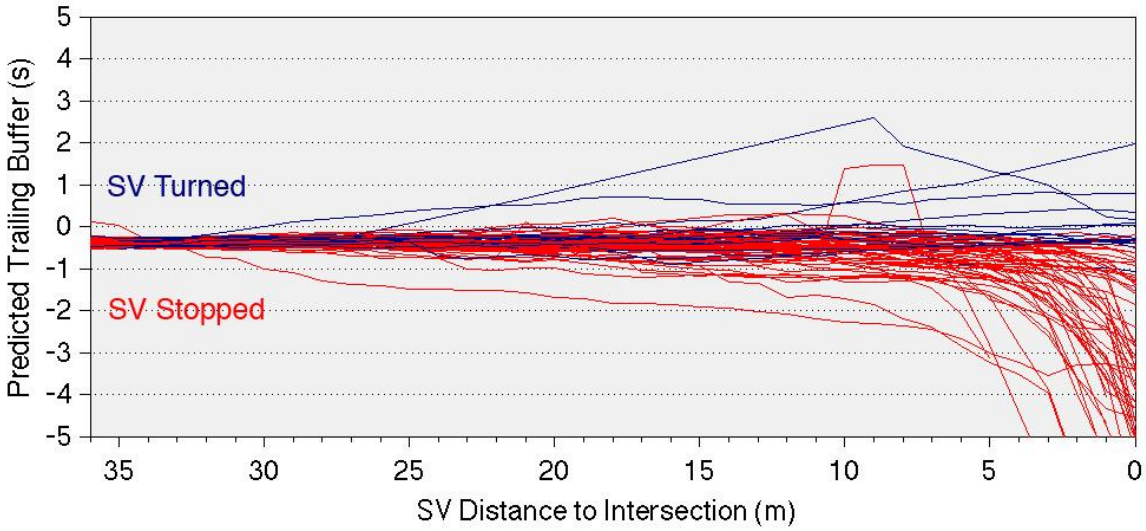

Figure 38. Evolution of the predicted PET when it started between -0.5 and -0.25 s.

Deleted:

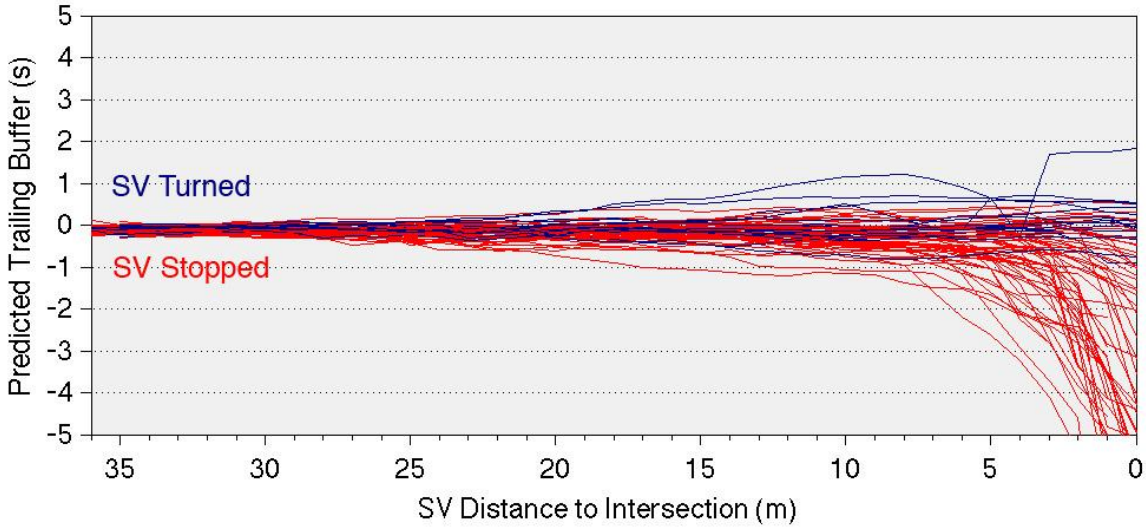

Figure 39. Evolution of the predicted PET when it started between -0.25 and 0.0 s.

Deleted:

A final note in examining the above graphs is that the predicted PET metric is not necessarily the most sensitive metric for predicting when the SV has decided to stop. The predicted PET attempts to describe the likely scenario outcome if the driver decides to turn. In a sense, the predicted PET metric masks stopping behavior by predicting what would happen if the driver decides to reaccelerate. As shown in Figures 35 to 37, the predicted PET begins to decrease rapidly (indicating that the SV is stopping) only after the SV has reached a point 5 m or less from entering the intersection. As shown in Figure 38, simply looking at the speed profile of the SV during the approach might allow for discrimination between decisions to turn or stop somewhat earlier, perhaps as early as 8 or 9 m depending one's tolerance for false alarms and missed detections. This would correspond to 10 to 11 m before entering the ZoC.

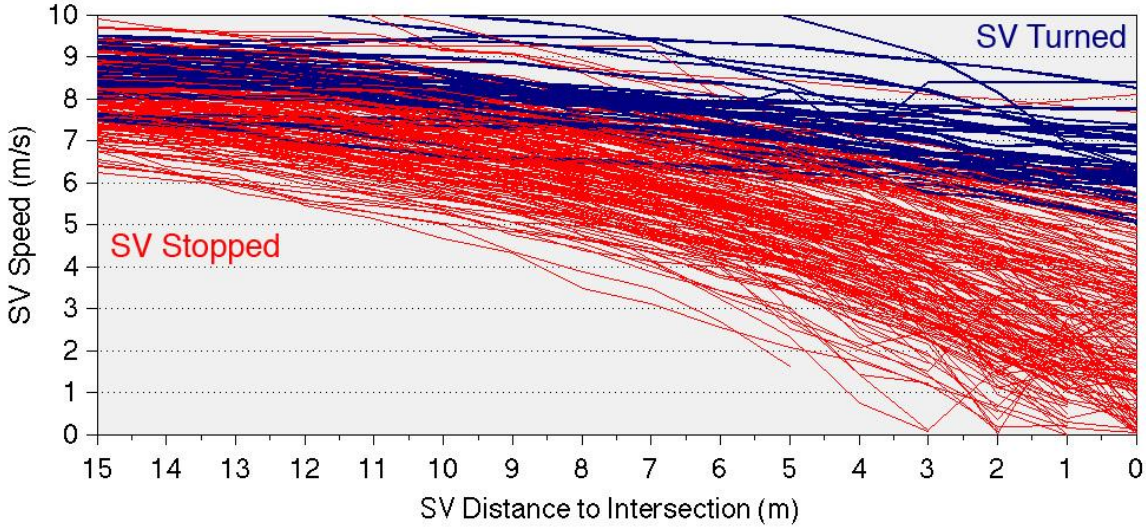

Figure 40. SV speed profile for cases when the predicted PET started between -0.75 and 0.0 s.

Deleted:

Cases When the Predicted PET is Between 0.0 and 1.5 s

Figures 39, 40, and 41 show the evolution of the predicted PET metric for cases when it started between 0.0 and 1.5 s at increasing half-second intervals. As the starting predicted PET increases from graph to graph, the number of SV decisions to turn (shown in blue) increases. If a predicted PET of 0.0 s was used as the warning cut-off point n. Given this warning criterion, all of the cases shown in this section would receive some indication that the turn was safe.

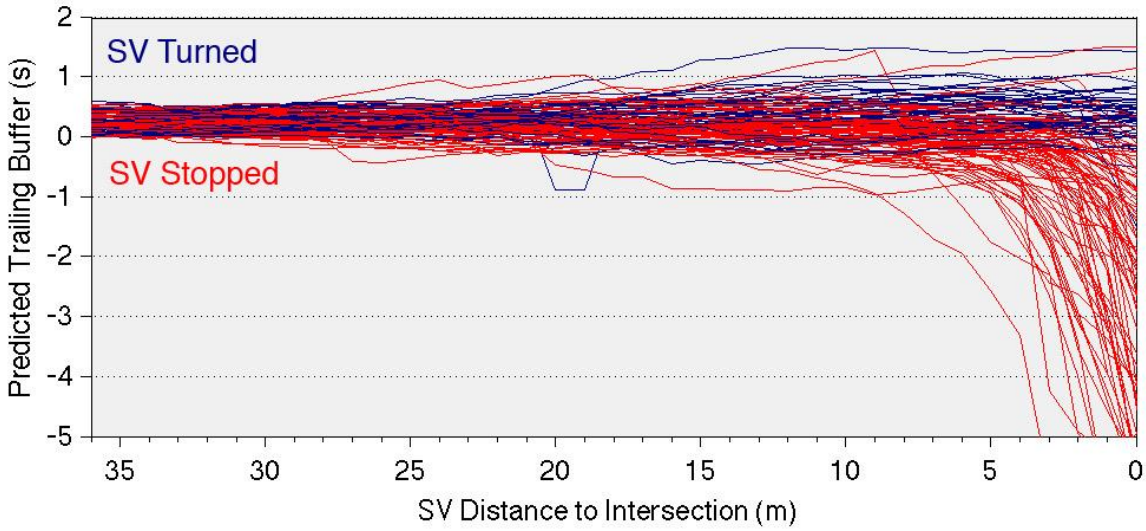

Figure 41. Evolution of the predicted PET when it started between 0.0 and 0.5 s.

Deleted:

Figure 39 shows the cases when the predicted PET, when the SV was 36 m from the intersection, was between 0.0 and 0.5 s, indicating that the SV driver could complete the turn in front of the POV with little spare time. In only 34 percent of these cases, the SV driver decided to turn in front of the POV, and in 66 percent of the cases, the SV driver decided to stop. Of the cases when the driver decided to stop, half of the time the trial was rated as "definitely not enough time to turn" while the other half of the time, the trial was rated along the lines of "I probably could have turned."

In Figure 40, showing cases when the predicted PET was between 0.5 and 1.0 s at the desired time to give the driver decision support, the number of cases when a turn was made by the SV driver increased to 53 percent. Of the cases when the SV driver decided to stop, only 25 percent were rated as "definitely not enough time to turn." In Figure 41 showing cases when the predicted PET was between 1.0 and 1.5 s at the desired time to give the driver decision support, the number of cases when a turn was made by the SV driver increased to 72 percent, with nearly 80 percent of the drivers who decided to stop having rated the trial along the lines of "probably could have turned."

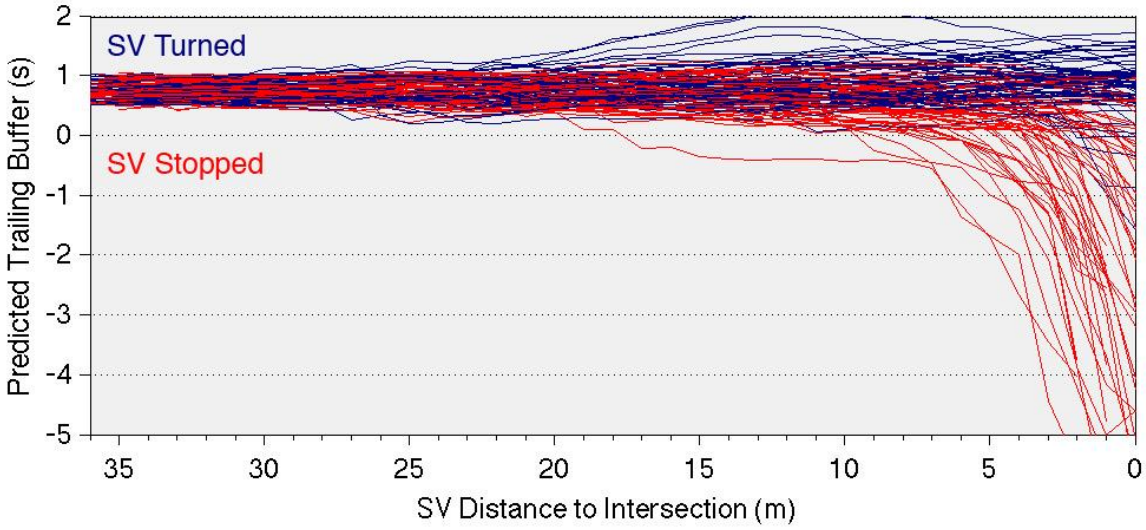

Figure 42. Evolution of the predicted PET when it started between 0.5 and 1.0 s.

Deleted:

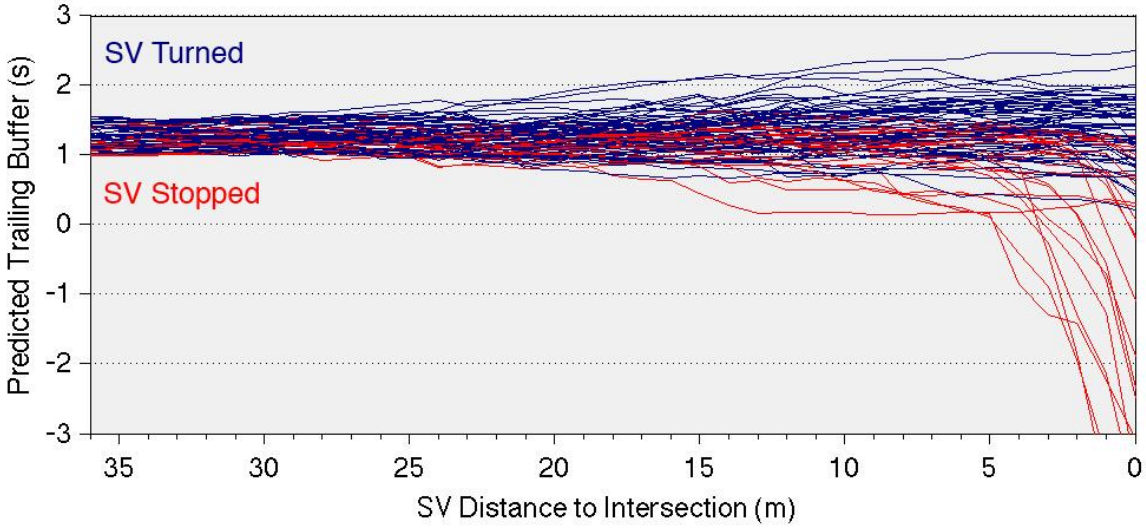

Figure 43. Evolution of the predicted PET when it started between 1.0 and 1.5 s.

Deleted:

One remarkable observation in examining the three graphs in this section is for the cases when the SV driver decided to turn, the predicted PET does not usually change dramatically throughout the approach. The predicted PET for an SV that eventually turns, does not usually drop, though it may increase slightly. When looking at a half-second range of predicted PET when the SV is at 36 m from the intersection, the predicted PET by the time the SV reaches 0 m from the intersection is typically still within the originally predicted half-second range or slightly greater. Thus, the predicted PET tends to be somewhat conservative, which is not surprising since it was based on mean desired turning speeds and mean desired deceleration rates.

Analysis of the Predicted PET Using Simulated Cases

Overview

In this section, the effects of changes in the POV speed on the predicted PET metric are explored using simulations. For the POV, there are three possibilities. First, the POV could maintain its current speed and heading, which has been the assumption of the predicted PET metric as

discussed to this point. Second, the POV could decide to accelerate, perhaps in an attempt to clear the intersection before an impending signal phase change. And finally, the POV could decide to decelerate in order to make a turn at the intersection. The focus of this section is to understand how the predicted PET metric reacts to these changes.

The simulations used in this section are based on the data collected at the RFS intersection during the DII and DVI experiments. The initial conditions for each simulation were based on the actual SV and POV measurements taken during the experiments, starting when the SV was 36 m from the intersection. The POV trajectory was then modified as per the conditions of the case to be simulated, and the predicted PET was recalculated for the new SV and POV trajectories.

The Case of the Accelerating POV

As described earlier in this report, the intersection approach and driver's decision to turn in front of the POV or to stop and let the POV pass is not necessarily a single decision made at a particular point in time. More likely, the decision can evolve over time, and based on the speed profiles involved and typical vehicle deceleration characteristics, the driver likely has ample time to change his or her initial decision all the way up until the vehicle is only a few meters from the zone of conflict. One situation in which the driver might need to change his or her initial decision is when the POV decides to accelerate.

The simulation shown in Figure 42 started with trials in which the initial predicted PET was between 0 and 0.5 s when the SV was 36 m from the intersection. This corresponds to a scenario when the SV would just barely clear the zone of conflict before the POV entered it. For argument's sake, if we were to use 0 s of predicted PET as our warning threshold, and we were to provide the SV driver with some sort of feedback at the preferred location of 36 m from the intersection, all of the drivers in the cases plotted would have received an initial indication that the turn was OK to make. (Note: only cases when the SV decided to turn in front of the POV are plotted). However, as per the simulation, once the SV reached 30 m, the POV decides to accelerate at 0.9 m/s^2 until reaching a maximum speed of 15.6 m/s (35 mph).

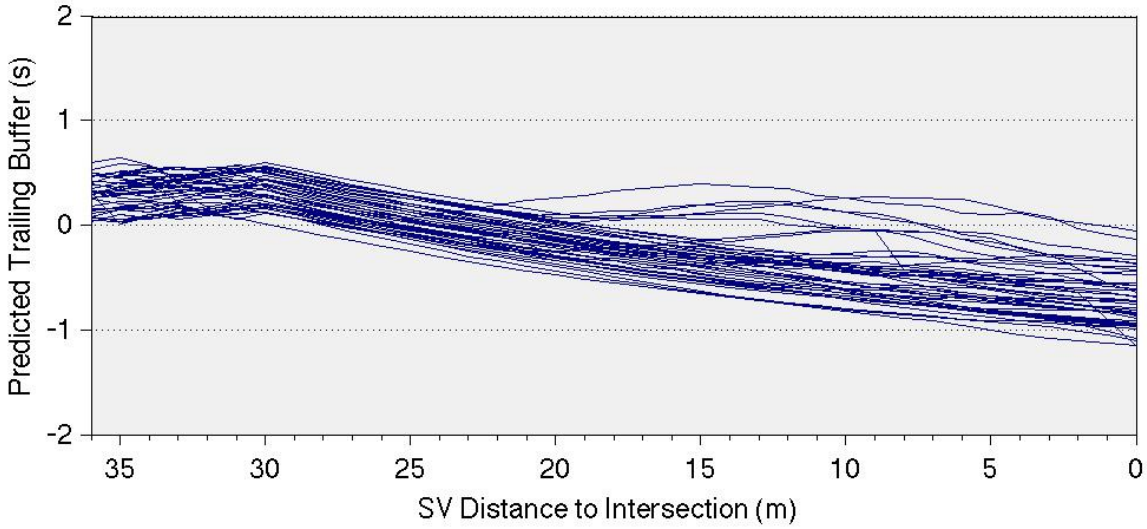

Figure 44. The effect of the POV accelerating at 0.9 m/s^2 to 15.6 m/s on the predicted PET.

Deleted:

As shown in Figure 42, the effect of a constant acceleration in the POV translates to a relatively linear decrease in the predicted trailing buffer. With an acceleration of 0.9 m/s^2 (a fairly fast acceleration), the predicted PET drops by about 1.0 seconds between the time that the SV is at 30 m and the time that the SV has reached the intersection. If a warning to the driver was based on the predicted PET with a threshold of 0 s, an indication that the turn was no longer safe would be given, at the latest, when the SV was somewhere around 20 m from the intersection, which would be plenty of time for the SV driver to change his or her mind and come to a stop.

However, it should be noted that the predicted PET as shown in Figure 42 does not take into account any POV acceleration. The calculation method for the predicted PET, which has been described thus far, only looks at the current POV speed and distance to the intersection, and then assumes that the POV will maintain its current speed. The assumption of a POV behavior, which, in this case turns out to be false, adds an element of delay between the point when the POV starts accelerating (when the SV is at 30 m) and the point when predicted PET recognizes the situation as no longer safe (when the SV is at 20 m).

In contrast, as shown in Figure 43, if there could be perfect prediction of the POV's intended acceleration, the lag in predicting the change to an unsafe condition would be effectively reduced to a single time step, resulting in the prominent step function shown in the graph. In Figure 43, the SV and POV speed and distance to intersection profiles are identical to the ones used in Figure 42. The only difference between the graphs is how the predicted PET metric was calculated. In Figure 43, an assumption was made based on the POV's detected acceleration. If the POV was accelerating, it was assumed that the POV would continue its acceleration until it cleared the intersection.

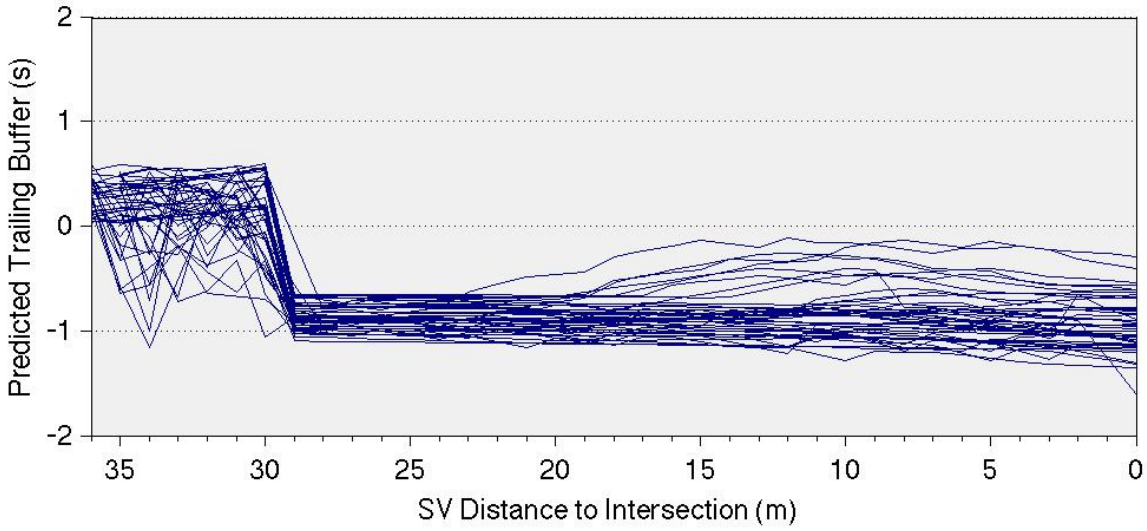

Figure 45. The effect of the POV accelerating at 0.9 m/s^2 to 15.6 m/s on the predicted PET (using POV acceleration prediction).

Deleted:

Unfortunately, while this assumption of the POV maintaining its current acceleration may seem reasonable, acceleration measurements are inherently noisy, and the assumption can have a very large impact on the predicted PET, especially when the POV is far from the intersection. This can be seen in comparing the predictions when the SV was between 36 and 30 m between Figures 42 and 43. In Figure 42, the predicted PET is relatively stable ranging from 0 to 0.5 s,

while in Figure 42, the predicted PET varies wildly from -1.0 to 0.5 s due to the addition of POV acceleration predictions. Although filtering of the acceleration might help smooth out a predicted PET which uses POV acceleration, the metric will still be highly sensitive to the accuracy of any assumptions made about the POV accelerations. As an example, contrast the graphs shown in Figures 44 and 45.

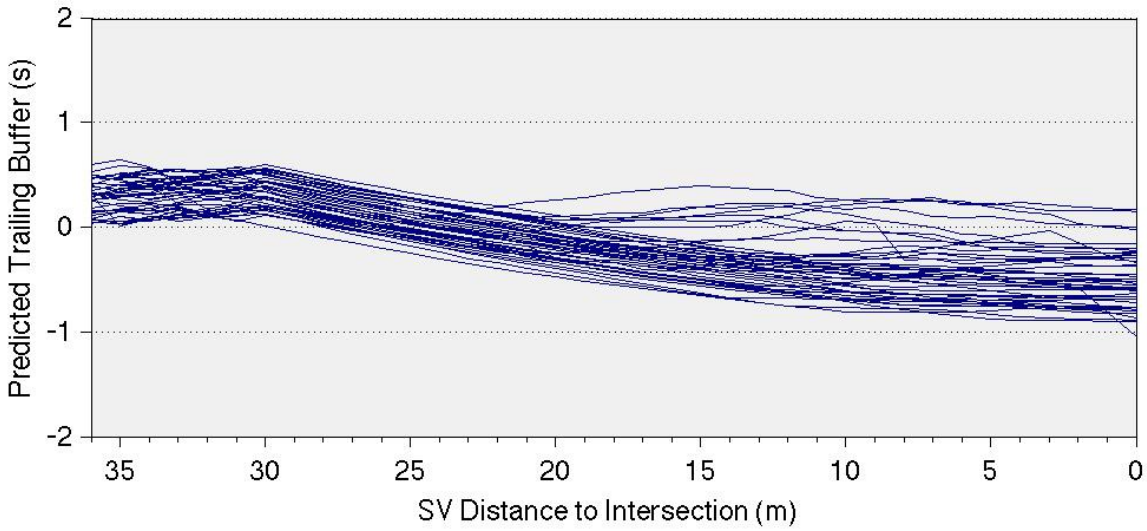

Figure 46. The effect of the POV accelerating at 0.9 m/s^2 to 13 m/s on the predicted PET.

Deleted:

In Figure 44, the same set of data was used as in the previous graphs. The simulation again began accelerating the POV when the SV had reached 30 m from the intersection, but this time, the simulation stopped the acceleration of the POV when the POV reached 13 m/s (29 mph), which typically occurred before the POV reached the intersection. Similar to what was seen in Figure 42, the POV acceleration appears as a linear decrease in the predicted PET, but instead of continuing throughout the SV approach, the predicted PET stabilizes once the POV stops accelerating (which occurs when the SV reaches about 10 m from the intersection).

The only difference between Figure 45 and 44 is whether or not the predicted PET metric included POV acceleration prediction. The net effect in Figure 45 of including the POV acceleration prediction is an instantaneous underestimation of the predicted PET once the POV starts to accelerate (when the SV was at 30 m). Then, once the POV stopped accelerating, the underestimation (caused by the faulty assumption that the POV would continue to accelerate at its current rate through the intersection) was corrected. Basically, when the POV started to accelerate, the predicted PET dropped by about 1.0 s, which would have triggered a warning as the predicted PET dropped below 0.0 s. Then when the POV stopped accelerating, the predicted PET increased by 0.5 s. The immediate underestimation of the predicted PET would have resulted in a number of cases in which a warning would have been given during the evolution of the scenario, but that warning would have turned out to have been a false alarm.

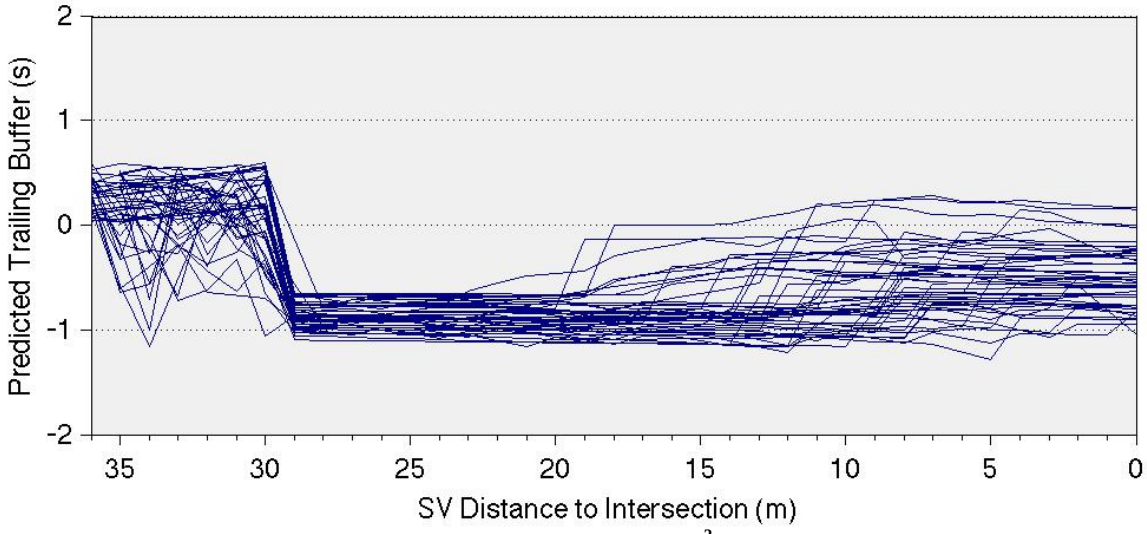

Figure 47. The effect of the POV accelerating at 0.9 m/s^2 to 13 m/s on the predicted PET (using POV acceleration prediction).

Deleted:

The Case of the Decelerating POV

The case of a decelerating POV is less dangerous than the case of an accelerating POV, since the deceleration of the POV roughly translates to an increase in the time available for the SV to turn in front of the POV. The most common case of a decelerating POV would be the case in which the POV slows to make a turn at the intersection. The relevant question posed in this section is how well and how quickly the predicted PET metric can detect a decelerating POV?

In the previous section regarding the case of the accelerating POV, the simulation looked at cases when the SV driver decided to turn when the predicted PET was just above the warning threshold (arbitrarily set at 0.0 s of predicted PET). This was because these cases were the ones to be most likely impacted by the acceleration of the POV. Conversely, when looking at the case of a decelerating POV, the opposite would seem true. The cases that would seem most relevant would be when the SV driver decides to stop and the predicted PET is just below the warning threshold. In this case, the slowing of the POV would now make the situation safe for the driver to turn in front of the POV.

The simulation shown in Figure 46 started with trials when the initial predicted PET (when the SV was 36 m from the intersection) was between -0.5 and 0.0 s. This corresponds to a scenario when the SV driver would have received a warning, since for the purposes of our discussion, the warning threshold has been arbitrarily set at 0.0 s (meaning that the SV would just finish clearing the zone of conflict before the POV enters it). The simulation of the POV assumed that the POV would maintain its original speed profile until it reached the point where it needed to start deceleration in order to reach a desired turning speed based on a desired deceleration rate. For the purposes of this simulation, the desired turning speed was set at 5.89 m/s (13 mph), and the desired deceleration rate was set at 1.61 m/s^2. Both of these values are the same as the ones that were used in the prediction of the SV's turning profile. Thus, these values would be appropriate if the POV was attempting to turn left, but the desired turning speed may be a little high if the POV was attempting to turn right.

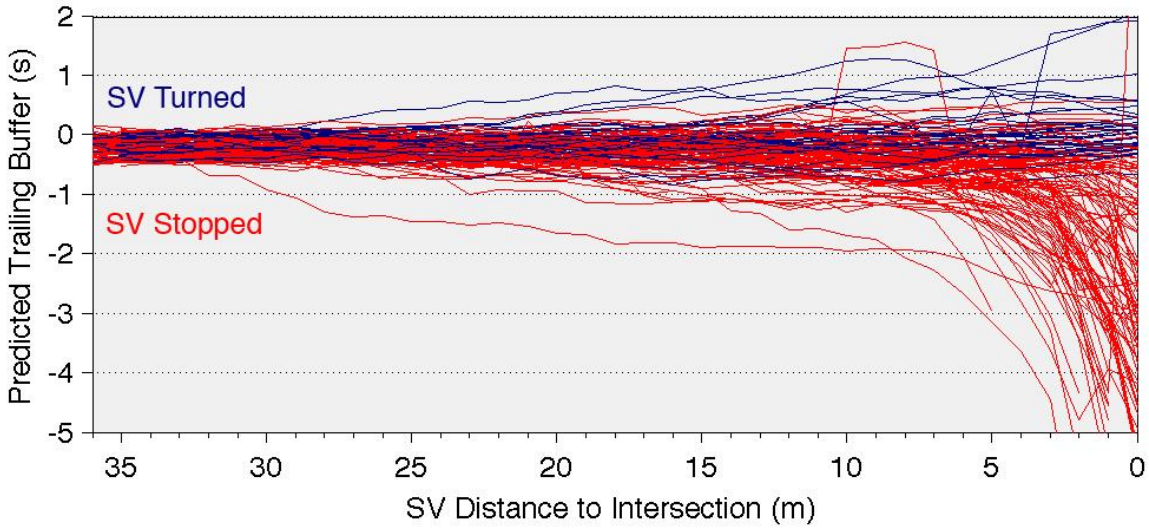

Figure 48. The effect of the POV decelerating as if to make a turn at the intersection.

Deleted:

As shown in Figure 46, there was actually very little impact on the predicted PET due to the POV deciding to slow as if it were to make a turn at the intersection. For the cases when the SV decided to turn (shown in blue), there was a slight increase, on the order of 0.5 s, in the predicted PET starting when the SV was around 10 m from the intersection. However, for the cases when the SV decided to stop, the predicted PET provided very little evidence that the POV was slowing.

While at first this result may seem a little puzzling, it is simply a reflection on the fact that the POV does not need to start slowing down until it reaches 20 to 30 m from the intersection. In the scenario that set up, the POV is only just starting to decelerate as the SV arrives at the intersection. This is really not surprising considering how many times one might recall encountering this same situation while driving. An initial decision is made by the SV driver that there is, just barely, not enough time to turn. However, after coming to a near stop in the intersection, the SV driver observes that the POV looks like it is starting to slow, and this results the SV driver thinking that he could have turned if only he'd known that the POV was going to slow down. What this simulation shows is that it would be unlikely that an LTAP-OD system could detect the POV's intent to slow in this situation because the POV does not even start to decelerate until well after the system would have advised the SV driver to stop and let the POV pass.

Discussion and Interim Conclusions

The predicted trailing buffer or predicted post-encroachment time (PET) is the primary metric that has been proposed to model and, ultimately, predict whether or not a driver will, or should, turn in front of a POV. The PET metric basically predicts the expected difference between how long it will take the POV to reach the intersection versus how long it will take the SV to clear the path of the POV. Thus, if the metric predicted that the PET would be 0 s, then the SV would just barely complete a turn in front of the POV. Negative values would indicate a less safe turning scenario. Positive values would indicate that the SV would clear the intersection with some amount of spare time before the POV arrived.

Unlike other CICAS applications, in the LTAP-OD scenario the interactions of two moving vehicles must be described and predicted while they are both still approaching the intersection. In CICAS-V, the main decision is whether or not an SV will stop for a traffic signal. In CICAS-SSA, the SV is already assumed to be at a stop in the intersection. While there are cases in CICAS-SLTA where the SV will be stopped, the most difficult case to handle is the one where the SV driver decides to turn without stopping, and decision support must be given to the driver before the SV reaches the intersection. In this case, the LTAP-OD could be said to be constantly evolving. Vehicle speeds may be changing rapidly as drivers speed up or slow down.

The analyses that were performed in this section examined the evolution of the predicted PET as the SV approaches the intersection under a number of likely scenarios. If the SV driver eventually decides to turn, the predicted PET metric showed the emergent feature of remaining relatively invariant throughout the SV's approach. If the SV driver eventually decides to stop, the predicted PET metric also remained relatively invariant until the SV reached 5 m or so from the intersection, at which point the predicted PET dropped rapidly as the SV came to a stop.

While the predicted PET may be a good metric for predicting whether or not the conditions are right for the SV to attempt a turn in front of the POV, the predicted PET was not very good at detecting when the SV driver decided to stop. Based on the predicted PET alone, distinctions between SV drivers who have decided to stop and those who have decided to turn can be made once the SV has reached 5 m or so from the intersection. Other metrics, such as even looking at the SV speed alone, might make more accurate predictions earlier, possibly as early as 8 or 9 m from the intersection. Thus, the predicted PET is a promising metric determining early on whether or not the vehicles will conflict with each other as they both reach the intersection, but if the driver interface called for something like a reactive, last-second warning, then the predicted PET would not be an ideal metric for this purpose.

The other analyses performed in this section examined how the predicted PET reacted to events such as POV accelerations or decelerations using simulations. Typically, POV acceleration or deceleration resulted in a respective increase or decrease in the predicted PET metric that was more or less linear in nature lasting for the duration of the acceleration/deceleration event. In effect, this means that there will be some lag between the time when the POV starts to accelerate and the time that the predicted PET drops below any sort of warning threshold.

If the predicted PET metric were to take into account POV acceleration data, then POV acceleration or deceleration events would result in a step function, thereby eliminating any lag in the detection of the evolving unsafe turn scenario. However, there were two problems associated with the use of POV acceleration data which will require more research. First, acceleration data is very noisy and will need to be filtered. Second, knowing what the current acceleration of the POV is, does not necessarily contain any information that will help decide how much or for how long the POV will continue the acceleration. Since using the POV acceleration data (which is inherently noise to begin with) results in a step function output to the predicted PET metric, the use of this information makes the predicted PET metric highly unstable. A small quick POV acceleration might result in a large change in the predicted PET, only to be cancelled later as the POV stops accelerating. Clearly, there will need to be more research and analysis into the trade-offs between how quickly the predicted PET metric can be made to detect changes in the situation versus the potential for the metric to become too sensitive and lead to increased false alarms.

POV decelerations also prove to be somewhat problematic to the predicted PET metric. If the POV has decided to slow down and turn at the intersection, and the conflict between the SV and POV is close (with a predicted PET near or just under 0 s), then the SV has more than likely already started slowing down before the POV would typically start to slow down. In essence, the predicted PET cannot be used to detect a POV that is slowing to turn at the intersection because the POV does not need to start slowing until well after the SV driver has practically already come to a stop at the intersection.

Feasibility of a Last-Second LTAP-OD Warning

There are a number of ways that a safety system can interact with a driver ranging on a continuum from providing information and warnings to taking control of the vehicle and automatically maneuvering it out of harm's way. However, our language is imprecise in describing this continuum. The word "warning" is often used rather generically to encompass a broad range of interactions with the driver. As an example, illuminating an icon in the driver's side mirror to indicate the presence of a vehicle in the blind spot might be referred to as a warning. But there is a large difference in the intent, requirements, and design of a side obstacle detection warning interface and a forward collision warning interface. In some circles, a side obstacle detection system warning might be classified as a situational awareness display (as opposed to a true warning). The side obstacle display attempts to provide the driver with information at a time that a decision is being made, while the forward collision warning system attempts to intervene at the last second to prod the driver into taking some corrective action to avoid a collision.

Rather than attempt to split hairs on the definition of the word "warning," it will suffice in this section to make a distinction between a situational awareness or decision support type of display (such as the blind spot awareness example) and a last-second warning (such as a forward collision warning system). To date, the work on the CICAS-SLTA project has been pursuing the concept of providing the SV driver with some sort of decision support. A system design following the decision support interaction model would provide the driver with information or advice on a recommended course of action during the time that the driver is trying to make a decision. Alternatively, another model for interacting with the driver might be to employ a last-second warning. Using this model, the system would attempt to detect that the driver has made a poor decision or is doing something wrong, and then attempts to intervene before the situation becomes unrecoverable. This is the type of interaction employed in the CICAS-V red-light running application. A last-second warning will need to be more intense to ensure that the driver perceives it quickly and reacts appropriately. However, it should be noted that there is also a trade-off between the intensity of the interaction and the acceptability of warnings that might be perceived as false alarms. With a more intense interaction, the drivers' tolerance for false alarms will likely drop.

The main question posed in this section is whether or not a CICAS-SLTA system could potentially employ a last-second warning to stop the SV driver from attempting to turn in front of an approaching POV. The last-second warning may be conceived as stand-alone or in conjunction with providing decision support as part of a multi-staged warning system. However, the first step is to determine whether or not a last-second warning would even be feasible. The answer to this question depends on two factors. First, what is the last possible moment that a

warning can be given while still allowing enough time for the driver to perceive and react to that warning by taking the appropriate corrective action? Second, at this last possible moment, can the SV driver's intent to stop or turn be detected with an acceptable trade-off between missed detections and false alarms?

Potential answers to these questions were analyzed in this section using the data collected during the Richmond Field Station DII and DVI experiments. First, there was an examination of normal stopping behavior based on cases when the driver of the SV appeared to change his or her mind about turning in front of the SV. Second, the issue was approached from the viewpoint of finding out how late a last-second warning could be given. And finally, the issue was approached from the standpoint of examining the trade-offs between missed detections and false alarms.

Normal Stopping Behavior

During the RFS DII and DVI experiments, the mean SV deceleration rates ranged from 0.1 to 0.28 m/s^2 (0.1 to 0.28 g); however, peak decelerations may have been higher, especially at the slow speeds typically seen during the left-turn maneuver. Typically the drivers appeared to make an initial decision on whether or not to turn in front of the POV and then stuck with that decision. However, there were a few instances when this was not the case. As was reported in Bougler, Cody, and Nowakowski (2008)[3], there were a number of trials noted when it appeared to the experimenter that the SV driver initially decided to turn in front of the POV, but then changed his or her mind and stopped to let the POV pass. Further investigation into those trials revealed that in these situations, the SV driver tended to change the deceleration rate, braking harder, as late as 8 to 12 m from the intersection (10 to 14 m from the zone of conflict).

To establish a range of what might be considered normal stopping behavior, the data from the RFS experiments were first sorted to find cases when the SV decided to turn in front of the POV without stopping. The SV speed distribution profile as a function of distance to zone of conflict was then calculated and used to compute a required deceleration profile (if the SV driver decided to stop) as a function of distance to zone of conflict. The required deceleration at any given point was computed using the following kinematic formula:

$$a_{req} = (V_f^2 - V_0^2) / (2*(X_f - X_0))$$

Where:
a_{req} = SV required deceleration
V_f = SV final speed (0 m/s^2)
V_0 = SV initial speed
X_f = SV final distance to intersection
X_0 = SV initial distance to intersection

As shown in Figure 47, two required deceleration profiles were plotted based on either the 10^{th} or 90^{th} percentile of the SV speed distribution profile. As observed during the experiments, if the driver changed his or her mind when the SV was between 14 and 10 m from the zone of conflict, the deceleration required to stop could be on the order of 2.5 m/s^2 (0.26 g) if the vehicle was travelling on the low end of the speed distribution. On the high end of the speed distribution, the required deceleration could be as large as 3.5 m/s^2 (0.36 g). For comparison, both of these decelerations are well within what might be considered normal deceleration. As an example, under conditions when a driver might decelerate for a green to amber traffic signal change, mean deceleration values typically ranged from 0.1 to 0.4 g according to data collected by White and

Ferlis (2004). (As a side note, the bump in the curves around the 8 m mark were due to random noise in the distribution of SV speed profiles.)

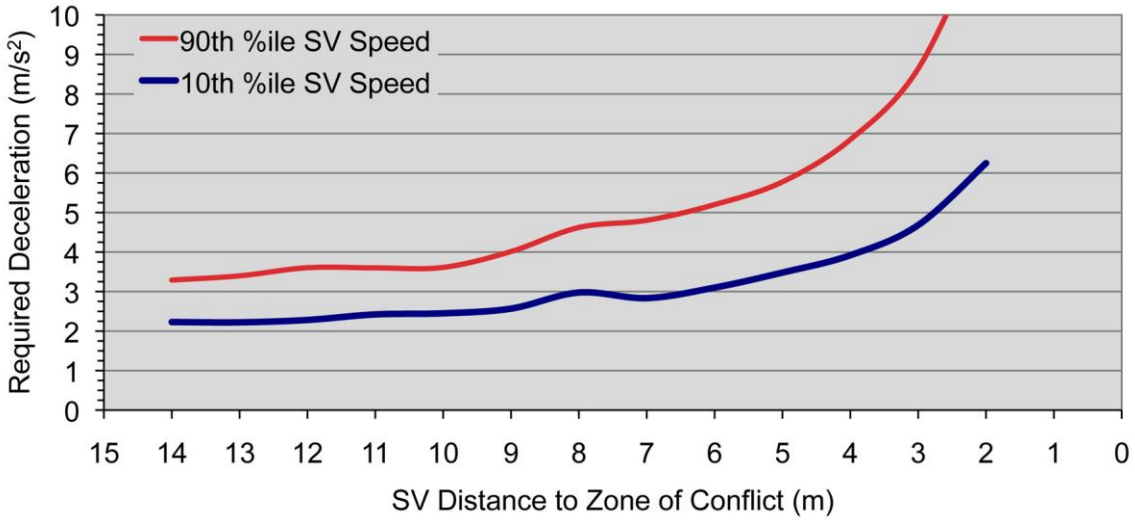

Figure 49. Required deceleration profile as a function of distance to zone of conflict.

Deleted:

As the SV gets closer to the intersection, the required deceleration, if the driver were to decide to stop, exceeds 0.4 g when the SV reaches 4 m from the zone of conflict (based on the 10th percentile of the SV speed distribution). This is because at 4 m from the zone of conflict, the low end of the SV speed distribution is only 5.6 m/s (12.5 mph). However, based on the 90th percentile of the SV speed distribution, 0.4 g of required deceleration is exceeded when the SV reaches 8 m from the zone of conflict. At 8 m from the zone of conflict, an SV who intends to turn may still be travelling as fast as 8.5 m/s (19 mph). Unfortunately, when considering the design of a warning, the fast travelling SV is probably the most likely scenario when a last-second warning would be needed, as this might be the case of a driver attempting to beat a POV to the intersection.

Required Deceleration in Response to a Last-Second Warning

Although Figure 47 shows that reasonable deceleration rates would be required of the SV if the driver initiated stopping as late as 8 m from the zone of conflict, the calculations used to generate that figure were based on the instant initiation of braking. If a last-second warning were to be provided to the driver, it would be provided because the system detected that the driver did not appear to be planning to stop, even though a dangerous situation was predicted. In this case, the driver would need some time to perceive and react to the warning before the initiation of braking could occur. Based on the detection of, and reaction to, green to amber traffic signal changes, the shortest perception-response that could be expected would be on the order of 0.5 s (based on unpublished data collected during the IDS Berkeley Field Test). Although this may seem fast when compared to something like a brake reaction time to a forward collision warning (generally accepted to be somewhere in the range of 0.75 to 1.0 s), it should be remembered that in this case, the driver is already decelerating. The required response is simply a change in brake pressure, since the driver already has his or her foot on the brake when the warning activated.

In general, to account for driver response time in the required deceleration calculation, the general assumption is that the driver will continue on his or her previous or default course of

action for the duration of the assumed perception-response time. In this case, at the time of the warning, the SV should already be decelerating to a desired turning speed, and this constitutes the default behavior that should continue for the duration of the perception-response time. Thus, the equation used to calculate the required deceleration in Section 6.2 needed to be modified as described below:

$$a_{req} = (V_f^2 - (V_0 + a_0*t_{PRT})^2) / (2*(X_f - X_0 - V_0*t_{PRT} - .5*a_0*t_{PRT}^2))$$

Where:
a_{req} = SV required deceleration
V_f = SV final speed (0 m/s^2)
V_0 = SV initial speed
X_f = SV final distance to intersection
X_0 = SV initial distance to intersection
t_{PRT} = Assumed driver perception-response time (0.5 s)
a_0 = SV's mean deceleration under the assumption that it will turn

Based on the above equation and the SV approach speed distributions calculated earlier, the required deceleration profiles, assuming a 0.5 s perception-response time, are shown Figure 48. In order to allow a driver travelling at the 90th percentile of the approach speed distribution enough time to stop at a reasonable 0.4 g deceleration rate after issuing a last-second warning, the warning would need to be given when the SV was at least 14 m from the zone of conflict. If the SV was travelling at the low end of the speed distribution, the warning could be given as late as 8 m from the zone of conflict.

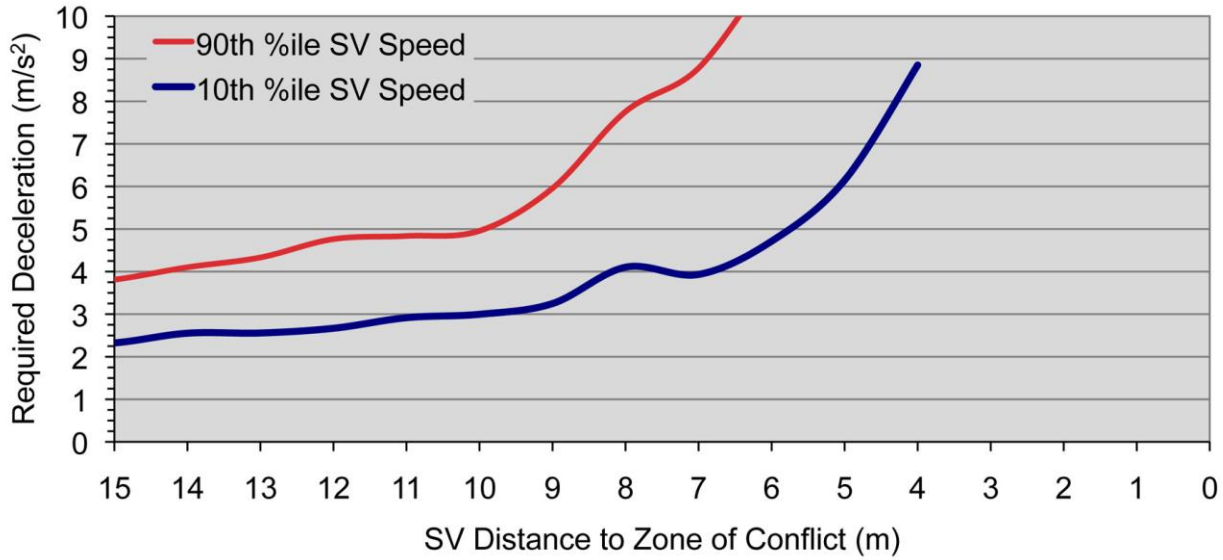

Figure 50. Required deceleration to stop with a 0.5 s RT for 10th and 90th %ile SV speeds.

Deleted:

Detection of Driver Intent at the Required Last-Second Warning Point

Section 6.3 presented a method to establish the moment when a last-second warning could be given to a driver, while still allowing that driver enough time to bring the vehicle to a stop before entering the zone of conflict with a reasonable rate of deceleration. Based on those findings, it was determined that a last-second warning should occur somewhere around 14 m from the zone of conflict. Figure 49 shows the approach speed distributions when the SV was 14 m from the

zone of conflict, for drivers who turned in front of the POV (blue) and for those who decided to stop (red).

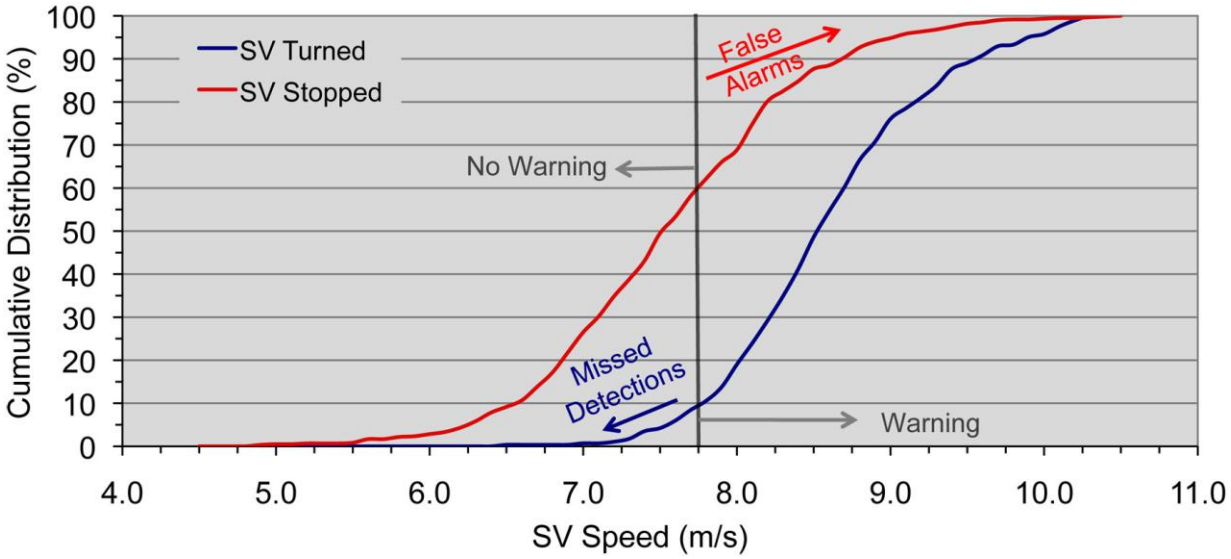

Figure 51. Approach speed distribution when the SV is 14 m from the zone of conflict.

Deleted:

As discussed in Section 3.0 of this report, the mean SV deceleration rates for drivers who decide to turn and those who decide to stop are mostly indistinguishable from each other. Thus, if the SV speed alone is used to distinguish intent to turn or stop, an arbitrary cut-off point could be set, and any SV traveling slower than the cut-off point would be assumed to be stopping while any SV traveling faster would receive a warning. Anything to the left of the cut-off point on the blue line would indicate a missed detection, and anything to the right of the cut-off point on the red line would indicate a false alarm. As shown in Figure 49, one arbitrary cut-off point could be 7.75 m/s which translates to the acceptance of a missed detection rate on the order of 10 percent. However, as shown in this example, this translates to a false alarm rate of nearly 40 percent. The SV speed profile distributions simply do not have a large distinction between drivers intending to stop and those intending to turn when the SV is still 14 m from the zone of conflict.

Discussion and Interim Conclusions

The main question posed in this section was whether or not a CICAS-SLTA system could potentially employ a last-second warning to stop the SV driver from attempting to turn in front of an approaching POV. The answer to this question hinged on a two factors. First, what was the last possible moment that a warning be given, and second, at this warning point, could the SV's intent to stop or turn be detected?

In response to the first question, the analysis showed that a last-second warning could be given as late as 14 m from the zone of conflict to keep the deceleration required by the driver at a reasonable level (0.4 g). For comparison, the warnings that were given during the experiments which conformed more to the situational awareness or decision support model of interaction were tested at 22, 30, and 38 m from the zone of conflict (with 30 m being the preferred location). Thus, one possibility for a multi-staged CICAS-SLTA system would be to give decision support when the SV was at 30 m from the zone of conflict (about 4 seconds away),

followed by a last-second warning when the SV was 14 m from the zone of conflict (about 2 seconds away) if the driver did not appear to be stopping.

Based on these calculations, the last-second warning would clearly need to follow the initial decision support fairly closely. It is difficult to predict how drivers might perceive such a system. One possibility is that drivers might perceive the system as "too chatty," especially if there is a high instance of false alarms. In the case of a false alarm, the driver would have received the initial decision support, but then felt that the second warning was too soon (and thus, overly annoying) since it would likely occur right as the driver began decelerating to a stop.

Unfortunately, the analysis also showed that when the SV was 14 m from the zone of conflict, it is very difficult to predict whether or not the SV is going to stop based on speed alone. Although there is always a trade-off between missed detections and false alarms, the overlap between the speed distributions for drivers who stopped and drivers who turned, was such that a 10 percent missed detection rate would be accompanied by a 40 percent false alarm rate.

To be fair, these calculations were all based on the worst case scenario: a SV driver who was travelling at the upper 90^{th} percentile of the approach speed distribution. The false alarm rate over the entire range of SV approach speeds could be considerably less. Additionally, if the SV was not travelling at the upper end of the speed distribution, the last-second warning point could be delayed by almost a second, activating when the SV reached only 8 m from the zone of conflict. Discrimination between drivers who intend to stop and those who intend to turn at this point is better, with a false alarm rate of 20 percent corresponding to a missed detection rate of 10 percent.

From these analyses, one thing is clear. The most challenging part of implementing a last-second warning for the LTAP-OD scenario is finding a way to reliably predict the driver's intent to attempt a turn in front of the POV or to come to a stop and let the POV pass.

Conclusions and Next Research Steps

This "PATH Portion" of the data mining efforts covered a number of topics such as providing further descriptions of the left-turn maneuver, a comparison of the potential effectiveness of a driver-vehicle interface (DVI), and detailed analyses aimed at the characterization of the predicted trailing buffer or predicted post-encroachment time metric. These efforts utilized driving behavior data collected during two experiments that had been previously completed. These data collections served as a basis for the current efforts and were conducted either during the CICAS-SLTA's predecessor program or from private funding provided by Toyota Motor Company. The overall goal of the data mining effort was to learn as much as possible from the existing data before collecting new sets.

There was a significant amount of effort put into processing and combining the different data sets used for the data mining analyses. Not all of the data sources could be combined with a reasonable amount of effort, and thus, not all of the analyses that were planned could be completed. Effectively, the data collected by the vehicles (including detailed measures of speed, acceleration, and throttle/brake use) could not be synchronized with the data collected by the infrastructure which described vehicle movements relative to the intersection. This prevented a set of planned analyses that would have replicated the types of analyses done during the Field Test conducted in the city of Berkeley using the data collected during the more controlled test track experiments.

However, as a byproduct of the data synchronization efforts, the limitations of the data sets collected during the test track experiments were well characterized. Specifically, the infrastructure based radars collected speed and distance measurements, but the accuracy of these measures became less than reliable once the vehicles reached 5 or 6 m from the intersection. When the vehicles were very close to the intersection, the senor accuracy appeared to drop significantly, and often the infrastructure reported speed profiles departed significantly from what the vehicle was actually doing. This was likely due to the fact that the infrastructure radars lost the targets near the intersection due to the vehicle's turning movement or the vehicle's speed dropping below some threshold of the device to get an accurate reading from the Doppler shift.

On the design and effectiveness of driver-infrastructure vs. in-vehicle interfaces

The first analysis discussed in this part of the final report compared the effects of an infrastructure-based left-turn collision warning (DII) to a similar in-vehicle variant of the same warning (DVI). The analysis used the combination of two data sets from experiments conducted using the Richmond Field Station instrumented intersection. The analysis had three goals. The first goal of this analysis was to examine the potential effectiveness of two warning locations (DII vs. DVI). Overall, the turning rate increased as the predicted trailing buffer or post-encroachment time (a measure of the available lag or "gap") increased. The presence of a warning in either the infrastructure (DII) or in the vehicle (DVI) resulted in a reduction in mean turning rate of 35 percent as compared to similar trials when no warning was given. However, there was not a significant difference in the observed effectiveness between the DII and the DVI; both have the potential to reduce unsafe left turns.

The second goal of this analysis was to understand the lessons learned regarding the designs of the DII and DVI. Specifically, the experiments determined that the optimal time to provide the driver with decision support information was between 3.5 and 4.5 s before the vehicle reached the intersection. Older drivers and the use of a primarily visual DII favored earlier information, while younger drivers and the use of a more salient DVI allowed for the information to be delivered later. The DII was more intuitive, trustworthy, and compulsory in the opinions of the drivers, whereas the DVI concept was inherently less familiar as drivers are not accustomed to receiving advice from their vehicles. Additionally, the DII had a distinct visual location advantage as it was very near the driver's focus of attention, whereas glances inside the vehicle to the DVI were less desirable, and the auditory alert was often viewed as annoying. These results do not conclude that a DVI would be inappropriate for this application, but simply conclude that the preference for this type of decision support interface would be primarily visual with a very high mounted, display that is central to the driver's field of view.

The final goal of this analysis was to examine any effects the DII, DVI, or various warning points might have had on the driver's initiation of braking. However, there was little evidence to support that any of these factors had much influence, which is not surprising given the conclusions from the Berkeley Field Test deceleration profile analysis also described in this report.

In conclusion, there was no compelling evidence to suggest that either a DII or a DVI would be either more or less effective as a means to present LTAP-OD decision support information or warnings. Either location presents a number of advantages and disadvantages. The DII can be used to address any vehicle, not just DSRC equipped vehicles; however, it can only address a single vehicle at a time. Additionally, there are more challenges in creating a sign that conveys a

salient and clear meaning to the driver, but remains consistent with all of the MUTCD conventions and guidelines. The DVI, on the other hand, allows individual drivers to be targeted, and provides a number of design options, such as the use of audio, which might be unavailable to a DII, but there will be significant challenges in finding a balance of design where the system is useful and not perceived as annoying.

What is the best metric to predict SV turning rate?

One of the most difficult challenges in regards to the LTAP-OD crash problem comes from the fact that there are no standard metrics which can be used to describe a turning maneuver when both approaching vehicles are moving. Prior work in the CICAS-SLTA program has attempted to develop a metric referred to as the predicted trailing buffer which is a prediction of the post-encroachment time. This metric is basically an estimate of the spare time that would remain before the Principal Other Vehicle (POV) reaches the intersection once the SV cleared the intersection. The second analysis discussed in this report examined the components used to calculate the predicted trailing buffer to determine the predictive power of each component.

The test track studies conducted during the IDS project found that there was significant overlap in the trailing buffer ranges when comparing the distributions of trials in which drivers stopped, as opposed to trials when the drivers turned in front of the POV. The implication was that there is no single trailing buffer based warning threshold that could be set to eliminate perceived false alarms. One of the concerns regarding the use of the predicted trailing buffer or post-encroachment time metric was that the metric essentially included assumptions about future vehicle movements to generate a predicted outcome. The fear was that the combination of the raw vehicle measurements was somehow obfuscating the situation, and a simpler model, relying on fewer assumptions, might actually have an increased predictive power.

However, the results showed that combining SV speed and distance with POV speed and distance to produce the predicted trailing buffer metric produced the most predictive model to explain the SV drivers' decisions to turn. POV distance alone accounted for the most variance in the turning decision, but the more complicated models, such as time to intersection and predicted post-encroachment time, always turned out to have a better goodness of fit, and thus, more predictive power than more simple vehicle measures taken alone.

On the characterization of the predicted post-encroachment metric

The fourth analysis discussed in this report examined the evolution of the predicted post-encroachment time (PET) metric during the intersection approach under a variety of both real and simulated cases. The predicted PET is the primary metric that has been proposed to determine whether or not the situation is safe for the SV driver to turn in front of the POV. Fortunately, the metric proved to be fairly stable and relatively invariant throughout the SV's approach, at least until the SV driver decided to stop and let the POV pass. However, the metric is not very good at quickly or accurately detecting the SV's stopping behavior. In a sense, the predicted PET metric predicts what would happen if the driver decided to turn. As the SV slows to let the POV pass, the metric transitions to predicting what would happen if the SV driver decided to attempt the turn by reaccelerating from a stop. Thus, if a proposed driver interface design for a CICAS-SLTA system requires the detection of whether or not the SV is stopping, then new metrics will need to be developed to detect that behavior more directly.

Another analysis performed in this section examined how the predicted PET reacted to events such as POV accelerations or decelerations using simulations. Typically, POV acceleration or deceleration resulted in a respective increase or decrease in the predicted PET metric that was more or less linear, lasting for the duration of the acceleration/deceleration event. In effect, this indicated that there would be some lag between the time when the POV starts to accelerate and the time that the predicted PET drops below any sort of warning threshold. Various strategies were discussed on how to incorporate measured POV acceleration into the predicted PET model; however, there will need to be more research into the trade-offs between how quickly the predicted PET metric can be made to detect changes in the situation versus the potential for the metric to become too sensitive such that small changes in the POV speed lead to increased false alarms.

Finally, the case when a POV decides to slow down and turn at the intersection was analyzed. In this case, POV deceleration proved to be somewhat problematic for the predicted PET metric. In essence, the predicted PET metric does not detect a POV that is slowing to turn at the intersection because the POV does not need to start slowing until well after the SV driver has decided to come to a stop at the intersection. The influence of the slowing SV on the predicted PET metric overshadows any influence that come from the slowing of the POV. Thus, if the CICAS-SLTA system requires the detection of a POV that intends to slow and turn at the intersection, then new metrics will need to be developed to detect and predict that behavior.

How feasible is a last-second LTAP-OD warning?

The fifth and final analysis discussed in this report examined the feasibility of providing drivers with a last-second warning in addition to some form of initial decision support. The idea behind this type of warning is that it would be provided only if the system detected that the driver was still intending to turn in front of the POV in a dangerous condition. This type of warning could then be designed to be more aggressive, possibly utilizing the auditory modality, which was not favored by drivers when presented with decision support information. The feasibility of this type of warning depends on two factors. First, it depends on how late a warning can be given to a driver while still allowing the driver enough time to take corrective action. Second, it depends on how well the driver's intent to stop or turn can be detected at the last-second warning point.

To answer the first question, an analysis was performed to find out what rates of deceleration might be required by the SV driver to change from intending to turn to intending to stop. The analysis was based on setting a maximum required deceleration at 0.4 g (which is typical of a brisk stop for a change in traffic signal phase), assuming a driver response time of 0.5 s (which is short but also typical of the response time to a change in traffic signal while attempting a left turn), and using the SV speed distribution by distance to intersection curves collected during the Richmond Field Station experiments. Based on these assumptions, the analysis found that a last-second warning could be given to the SV driver as late as 14 m (or about 2 s) from the zone of conflict with the POV. In contrast, the ideal point to give the driver decision support was on the order of 30 to 38 m from the zone of conflict.

Unfortunately, the answer to the second question was less promising. Based on a last-second warning point at 14 m from the zone of conflict, it proved very difficult to predict whether or not the SV was going to stop based on speed alone. Even allowing for a missed detection rate as high as 10 percent, the false alarm rate was still on the order of 40 percent. From these analyses, one thing is clear. The most challenging part of implementing a last-second warning for the

LTAP-OD scenario is finding a way to reliably predict the driver's intent to attempt to turn in front of the POV or to come to a stop and let the POV pass. In the absence of such a capability, it does not appear to be possible to simultaneously achieve a low enough missed detection rate and a low enough false alarm rate for the system to be perceived as both safe and acceptable to drivers.

Recommendations

Although this data mining report covered a number of topics in detail, significant conclusions and recommendations can be drawn, which may influence the direction of the CICAS-SLTA research and, ultimately, the design of a prototype left-turn assistance system. The first question that was answered in this report was whether to pursue the design of a DII or DVI as the means for communicating the information in this application to the left-turning drivers. While both types of interfaces have advantages and disadvantages, no evidence was found to rule out the potential for either design. Continuation of work on both types of driver interfaces is therefore recommended.

Keeping with the topic of driver interfaces, a second question that was answered in this report was whether or not to pursue a last-second warning, either instead of, or in addition to, providing the driver with early decision support information. While the analyses in this report do not entirely rule out the possibility of providing last-second warning to a driver who intends to turn left in a less than ideal situation, the results were not promising. Further research must be invested into finding metrics that can more quickly and accurately detect whether a driver intends to attempt a turn or to stop and wait. If such metrics can be developed, then the last-second warning might become a viable DVI option.

The final question that was answered in this report was whether or not to pursue the predicted post-encroachment time metric as the primary means for determining whether it is safe for the driver to turn in front of the oncoming vehicle. Although much fine tuning of the metric is still possible, the research showed that the predicted post-encroachment time metric provided greater predictive power than other, simpler, models. The CICAS-SLTA alert algorithms should therefore continue to pursue the development of the predicted post-encroachment time metric.

However, the research also confirmed that the predicted post-encroachment time metric does not allow for the detection of a number of critical scenarios. As an example, the metric is insensitive to conditions when the POV might be slowing to make a turn, and the metric is slow to detect that the SV driver has decided to stop. Future CICAS-SLTA programmatic research needs to eventually address the development of these additional metrics.

Normality Check on the SV's Initial Braking Point Distribution

One of the questions posed in the report was whether or not the initial braking point could be analyzed using an ANOVA due to the assumption that responses were normally distributed. As shown in Figure 50, the frequency distribution of initial braking points did appear to follow a normal distribution (as shown by the line on the graph). Although there was a difference in means between trials when the Subject Vehicle (SV) decided to turn in front of the Principal Other Vehicle (POV) (20.6 m) and trials when the SV decided to stop (22.9 m), there was little difference in the overall shape of the distribution due to this factor. Based on these graphs, it was determined that the assumption of normality held and that the ANOVA method could be used for the analysis.

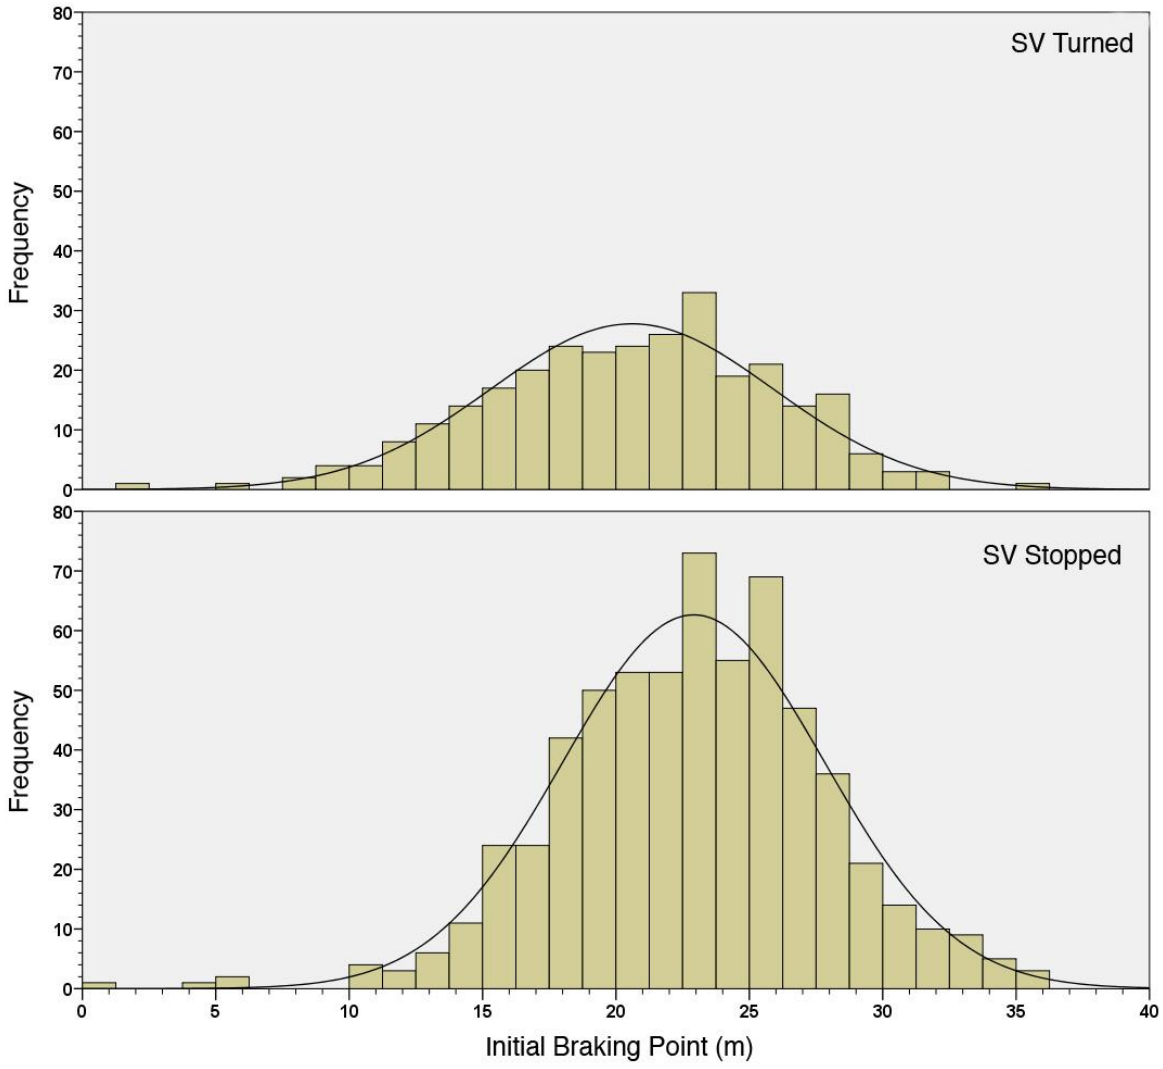

Figure 52. Distribution of initial braking points split by the SV's decision to turn or stop.

Deleted: 50

Second Generation Predicted PET Calculation Method

The concept of the Predicted Post-Encroachment Time (PET) metric is to predict the outcome, at any point during the Subject Vehicle's (SV) approach, if the SV decided to turn in front of the Principal Other Vehicle (POV). The metric roughly equates to a prediction of how much spare time would remain between the time it takes for the SV to clear the Zone of Conflict (ZoC) before the POV enters the ZoC. Thus, a Predicted PET of 0 s would indicate that the SV's rear bumper cleared the ZoC just before the POV's front bumper enters the ZoC.

In order to do this the arrivals of both the SV and POV to the ZoC must be predicted. This prediction is composed of two components which must be calculated at each point during the SV's approach: 1) an estimate of the SV's required time to clear the ZoC and 2) an estimate of the POV's time to reach the ZoC. The second component was simply computed by assuming that the POV will maintain its current speed through the intersection. Thus, the POV's time to reach the ZoC can be calculated based on its speed and distance to intersection (D2I). For the experiments at the Richmond Field Station, the POV D2I was equivalent to the POV distance to ZoC. For larger intersections, this may not be the case.

The SV ZoC clearance time was modeled using the sum of the three phases (t_1, t_2, and t_3). The first phase corresponded to the SV driver maintaining his current speed until needing to decelerate to a desired turning speed given a desired deceleration rate. The second phase corresponded to the acceleration/deceleration to reach the desired turning speed, and the third phase corresponded to coasting through the turn at the desired turning speed.

The desired turning speed used in the predicted PET calculation was 5.89 m/s which was the average turning speed measured during the Richmond Field Station experiments and is unchanged from the first version of the predicted PET metric. As discussed in Section 3 of this report, this compared fairly well to the typical turning speeds measuring during the Berkeley Field Test of approximately 6 m/s. The desired deceleration used in the predicted PET calculation was 1.61 m/s^2, also based on the mean deceleration rate measured during the Richmond Field Station experiments and unchanged from the previous version of the metric. As discussed in Section 3 of this report, although this value fell within the typical range of deceleration rates observed during the Berkeley Field Test, a more typical value observed for those intersections was on the order of 1.2 m/s^2.

One new parameter needed for the predicted PET calculation is a desired re-acceleration rate. As the SV approaches the intersection and starts to decelerate, there is the potential, especially if the SV driver intends to stop, for the SV's speed to drop below the desired turning speed. In this case, the predicted PET calculation assumes that the driver changes his or her mind and reaccelerates up to the turning speed, and then holds the turning speed constant throughout the remainder of the turn. Based on the analysis of the Berkeley Field Test data in Section 3, the desired re-acceleration rate was set at 0.9 m/s^2.

The formulas needed to compute the predicted SV time to clear the ZoC were based entirely on simple kinematic equations. The computation method for calculating this

prediction at any given time step based on a measured SV location and speed has been described below:

$$\text{SV ZoC Clearance Time} = t_1 + t_2 + t_3$$
$$t_2 = (v_i - v_{turn}) / a_{desired}$$
$$d_2 = ½ * (v_i + v_{turn}) * t_2$$
$$t_1 = (d_i - d_2) / v_i$$
$$t_3 = (¼ * 2\pi * r_{intersection} + l_{vehicle}) / v_{turn}$$

Where:

v_i = The initial SV speed as measured (m/s)
v_{turn} = The desired turning speed for the SV (5.89 m/s)
$a_{desired}$ = The desired SV deceleration/reacceleration rate (1.61 or -0.9 m/s^2)
d_2 = The distance covered by the SV while decelerating during t_2
d_i = The initial SV distance as measured to turn turning point (m)
$r_{intersection}$ = The radius joining the center of the SV turn lane to the center of the destination lane (5.45 m for the RFS Intersection)
$l_{vehicle}$ = SV vehicle length (5 m for the Ford Taurus sedan used as the SV)

The first step of the prediction is to determine the value of t_2 or the time required for the SV to slow down or speed up to the desired turning speed (the deceleration/reacceleration phase). Once this value has been calculated it can be used to determine how much distance (d_2) the SV has covered during this phase.

The second step of the prediction is to determine the value of t_1 which can come from one of three equations depending the SV's speed and distance to intersection.

First, if the SV is traveling faster than the turning speed and d_2 is less than d_i, then it is assumed that the SV will maintain its current speed for the duration of the t_1 phase, and t_1 can be calculated using the following equation.

$$t_1 = (d_i - d_2) / v_i$$

Second, if the SV is travelling slower than the during speed and d_2 is less than d_i, then it is assumed that the SV will immediately speed up, and maintain the desired turning speed for the duration of the t_1 phase.

$$t_1 = (d_i - d_2) / v_{turn}$$

Finally, if d_2 is greater than d_i, then there is not enough time to complete the deceleration/reacceleration phase before entering the turn, and t_1 goes to 0.

The third step of the prediction is to determine the value of t_3 which can come from one of two equations depending the SV's speed and distance to intersection. First, if the SV completes the deceleration/reacceleration phase before entering the turn, then it is assumed that the turn will be completed at the desired turning speed, and t3 is calculated using the following equation:

$$t_3 = (¼ * 2\pi * r_{intersection} + l_{vehicle}) / v_{turn}$$

Second, if the SV cannot complete the deceleration/reacceleration phase before entering the turn, then the distance travelled during t_3 is appropriately reduced, and t_3 is calculated using the following equation:

$$t_3 = (d_i - d_2 + ¼ * 2\pi * r_{intersection} + l_{vehicle}) / v_{turn}$$

The final step in calculating the predicted PET is to sum the three phases that make up the SV's predicted time to clear the ZoC and subtract the POV's predicted time

1.2.2 UMTRI Portion: Analysis of Left-Turn Gap Acceptance at Signalized Intersections using Naturalistic Driving Data

Executive Summary

In this work, naturalistic driving data was analyzed to establish benchmarks for turning behavior and conflict metrics for permissive left turns at signalized intersections. The following specific outcomes were achieved:

1) Automatic detection methods were developed to locate relevant left turn events within the FOT database
2) Detailed reconstruction of vehicle trajectories was carried out for the turning subject vehicle (SV) and principal other vehicles (POV's) approaching the intersection from the opposite direction
3) Conventional conflict metrics such as Gap Time (GT) and Post Encroachment Time (PET) were evaluated
4) A brief investigation was made on the influence of age and gender on gap acceptance
5) Leading Buffer (LB) and Trailing Buffer (TB) metrics were calculated as a candidate *predictive* measure of left turn timing.

Several technical challenges were overcome in the course of this work:

- Locating suitable events in the naturalistic driving database required specific information on intersection types, turning priorities and locations in SE Michigan; this was not generally available, so data tables were constructed for one particular county
- Fusion of data from various vehicle sensors required detailed estimation of timing offsets, as well as application of multiple Kalman filters to merge data and reduce the effects of noise (radar data is particularly noisy, and GPS data is only accurate to within a few meters)
- Because of occlusions and target dropping, direct use of radar data was only able to provide broken trajectories; hence algorithms were developed to provide the necessary associations and extrapolations to create continuous trajectories
- Speed changes for the subject vehicle are critical to making predictions of gap timing via the trailing and leading buffers; hence a numerical model was developed based on the event data analysis

Once vehicle trajectories had been constructed, a number of analyses were conducted. Summary results are:

- PET result showed that Michigan drivers exhibited a typical 2 second minimum gap acceptance when turning ahead of the approaching vehicle, similar to that reported previously for Californian drivers
- Traditional measures such as GT are not appropriate to making useful predictions of gap timing for CICAS-SLTA, and estimation of reference speed during the turn was necessary in the LB/TB calculations
- LB is mainly relevant in cases where there are multiple POV's
- Male drivers accepted shorter buffer times than female drivers in the case where the SV crossed first.

1. Introduction

In the SE Michigan region alone some 96,000 intersection crashes occurred over the past nine years. Developing technologies to mitigate intersection related crashes would substantially benefit the driving public.

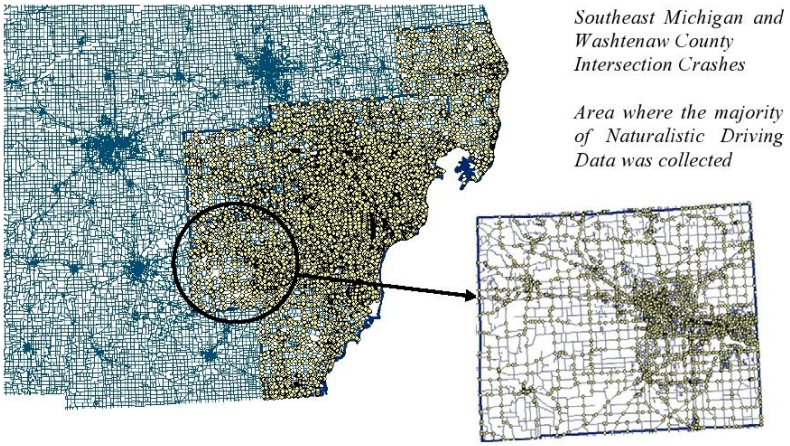

Figure 1.1. Locations of intersection crashes in SE Michigan and Washtenaw County in the last nine years.

The aim of this study has been to capture and analyze left turns made at signalized intersections using naturalistic driving data. The purpose is to characterize the range of such behaviors to support the development of a signalized left turn assistant (SLTA) system enabled via wireless communication between the subject vehicle (SV) and a site-based sub-system for vehicle monitoring and motion estimation. Because the field studies that generated the driving data were for unaccompanied drivers (driving without the presence or intervention of experimenters) the range of driver performance for gap acceptance and vehicle control is believed to be typical of the general population of SE Michigan.

The data used in this analysis was originally collected for the purpose of evaluating driver assistance systems, but measurements include a great deal of naturalistic driving that is largely or wholly unaffected by the presence of these systems. In particular there are many examples of left turns across the path of oncoming traffic, from which it is possible to estimate gap timing for cases where the turning vehicle either waits or does not wait for the oncoming vehicle to pass. Extracting these events and analyzing the kinematics of the turning and opposing vehicles is the purpose of this study.

Left turns at urban and sub-urban intersections can involve relatively complex estimation and decision making when the turn is on an unprotected (permissive) green light. Three critical aspects for the driver to monitor are: signal state, presence of oncoming traffic, and pedestrians in the cross walk.

The goal here is to characterize the range of normal behavior and expand on what has been gathered in California. In particular,

- the driver in the turning (subject) vehicle is largely unaware of being an experimental subject (no accompanying researcher present, has long-term use of the test vehicle, and no intrusive instrumentation present)
- a broad spectrum of high quality data is obtained from the subject vehicle – radar, vehicle motion, controls – plus video of the forward scene and the driver's face to provide contextual support for sampled cases
- baseline data is available for the individual drivers (e.g. age and gender)
- the data is derived from Michigan intersection characteristics and driving styles

Section 2 summarizes the data sources available, and explains why one particular naturalistic driving data source was used for the analysis; additional data on intersections (location, type, priorities) were also vital to the process of isolating suitable unprotected left turn events. Section 3 is at the core of the analysis, describing how somewhat noisy vehicle motion and radar data were used to generate trajectories for the conflicting vehicles. Once achieved at an acceptable level of accuracy, it is then quite straightforward to determine relevant conflict metrics. These metrics are reviewed in Section 4, and specific output is presented in Section 5. Finally some specific examples are analyzed in greater depth in Section 6.

2. Data Sources

Given the aim of this research, two instrumented vehicle studies at UMTRI are relevant: Automotive Collision Avoidance System (ACAS) [5] and Road Departure Crash Warning (RDCW) [7]. These field operational tests (FOT's) were designed to test the performance and acceptability of driver assistance systems in the normal driving environment, by subject drivers who use a test vehicle in a naturalistic setting, i.e. as a direct replacement for their own private car. ACAS tested forward collision warning coupled to adaptive cruise control, while RDCW tested warning for lane and road departure prevention. Each study captured a wide range of vehicle variables from CAN, inertial sensors and DGPS. Forward video provided lane tracking, and (together with driver face video) contextual information for manual review. Further, radar was used to detect relative position and velocity information for other vehicles and also fixed roadside objects.

Figure 2.1 shows a screenshot view of a turning event obtained from the RDCW data. This shows some window controls, radar "dots", forward video and further radar target data. Many other variables (more than 200 data channels) were obtained from the subject vehicle. The SV is being proposed as a moving data acquisition platform used to capture its own absolute motion and the relative motions of other vehicles, both with sufficient accuracy to determine useful conflict metrics related to the gap timing, gap acceptance and associated speed and acceleration information for the Subject Vehicle (SV) (and potentially the POV).

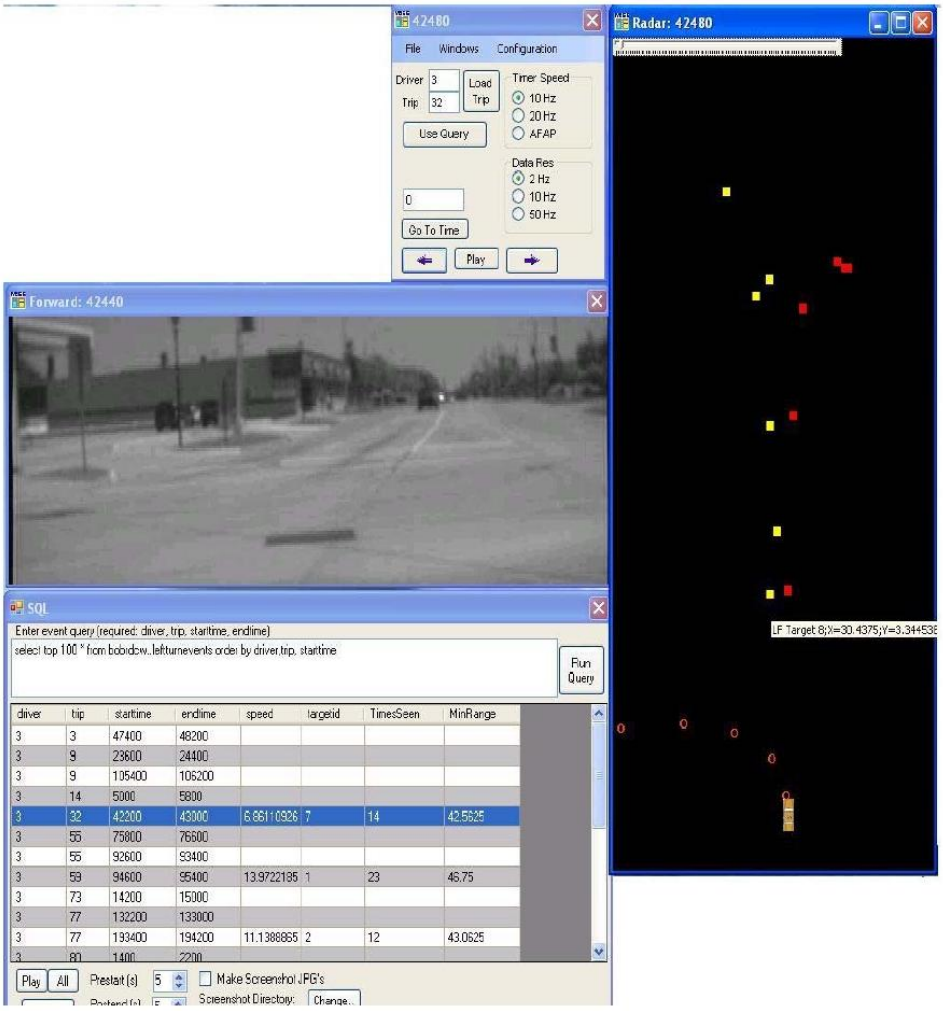

Figure 2.1. Video and radar traces obtained from RDCW test data.

There are just two options for capturing gap metrics and other naturalistic information for left turns at intersections: SV turning or POV turning. In the case where the SV is turning, it allows decision making and control behavior for the turning event to be related to driver characteristics, such as age and gender. By contrast, in the POV turning case, the decision and control is in the remote vehicle therefore contextual data is not available. In this case a dedicated site-based data capture is probably more appropriate (such data collection was carried out within this project by California PATH using stationary radar located at trial intersections) and since this is not part of the proposed study in Michigan, it will not be pursued further here.

When this analysis was first considered, the initial question was that of viability, and this had to be reviewed in some depth from the perspective of data quality and availability. Of course UMTRI has access to the ACAS and RDCW databases, but the details of the data types, sampling frequencies etc. may impact the viability of conducting worthwhile analysis. The SV motion can be characterized by on-board instrumentation which determines speed, acceleration (x,y components), yaw rate, location (via DGPS). Therefore, a portion of the analysis will be to fuse and refine these data so that the motion of the SV provides a suitable platform for defining the relative motion of the POV's using radar. Yaw angle measurement accuracy is particular important because of the strong sensitivity of azimuth angle to POV location in space; this issue is well understood, even before trajectories are analyzed. Also, relative timing of all sensors is vitally important for fusion, and again the yaw angle sensitivity means that even a minor delay in DGPS (relative to the radar data) will induce large variations in the global position estimates of remote objects. Fortunately, radar reflections from fixed objects can be used to help calibrate such errors, and Section 3 will discuss this in some detail.

The issues mentioned above were expected at the start of the project, before actual data were analyzed. However, during the review of data sources, a fundamental problem was discovered with the data captured from the ACAS vehicle systems. When the SV is at rest the forward radar was automatically switched off, rendering it completely useless for tracking the POV in cases where the SV is waiting to turn. Some events would be captured, but the radar switch-off would be a strong source of bias in establishing normal turning behavior. For this reason no detailed analysis of ACAS data was undertaken in the project.

Fortunately, RDCW radar instrumentation did not switch off in this way, and therefore the focus of this analysis has been based on the RDCW FOT dataset. Summary information of the RDCW instrumentation, including radar and cameras, are provided the Appendix. Reference [7] provides further details.

In addition to the RDCW FOT driver data, the project made use of geo-located intersection data and aerial photos; data tables have been compiled to provide intersection lat/long locations, type of signalization, traffic volumes and in some cases whether the intersection was equipped with permissive and protected left turns. The aerial photos allowed for measuring various points at the intersection and to overlap driver trajectory data to test validity.

3. Algorithms

3.1 Outline of Procedures for Conflict Metric Analysis

The analysis of the driver behavior on left turns is comprised of four major tasks:

- Data retrieval from the RDCW database,
- Reconstruction of SV trajectories,
- Reconstruction of POV trajectories, and
- Computation of the conflict metrics.

The overall analysis process is represented in Figure 3.1. After retrieving the data from the RDCW database, the trajectory, velocity and acceleration of the SV were computed by the Kalman filters (KF) in the blocks labeled (1), (2) and (4). Latency time (time offset between GPS and on-board instrumentation) is estimated in (3) and validated in (5). The calculation of the POV motion required another latency time calculation (6) in order to locate the target points in conjunction with the SV motion already found in (7). These target point radar reflections are noisy, including reflections from multiple objects and may include broken sub-sections; hence in (8) the sub-tracks are associated or grouped, and then refined via the same type of KF used in (1). The subcomponent(10) is where the conflict metrics are calculated, essentially a post-processing of the trajectory data. Each process is explained in detail in later sections.

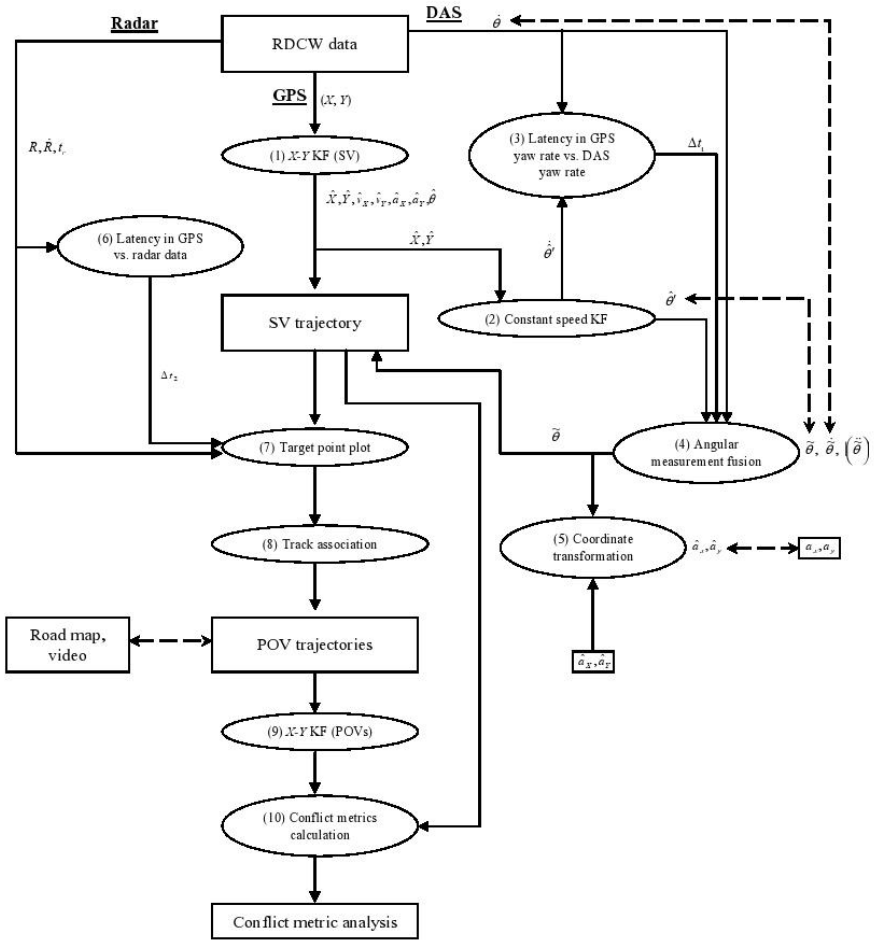

Figure 3.1. Flow chart of the conflict metric analysis.

3.2 Initial Screening of RDCW database

3.1.1 Data retrieval

To retrieve the left turn event data from the RDCW database, MATLAB® was used to send SQL queries with parameters to identify particular events. Those parameters were Driver #, Trip #, Start time and End time (recorded in centi-seconds) chosen from a list of all relevant left turn events detected in the RDCW database. Figure 3.2 shows a sample from the list of such events in Washtenaw County, MI. The total number of events was 1576. In the figure, SEMCOG ID is a label code that uniquely identifies the individual

intersection, and Type represents the type of the traffic control of the intersection indicated by either TS (traffic signal) or FL (flashing light). Additional information about whether an intersection of the TS type was protected or unprotected was given in a separate intersection database table.

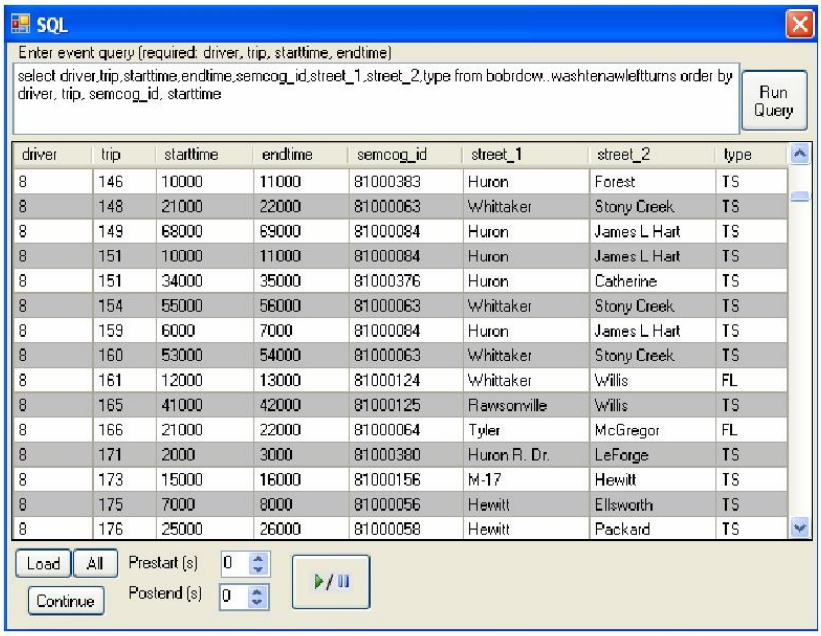

driver	trip	starttime	endtime	semcog_id	street_1	street_2	type
8	146	10000	11000	81000383	Huron	Forest	TS
8	148	21000	22000	81000063	Whittaker	Stony Creek	TS
8	149	68000	69000	81000084	Huron	James L Hart	TS
8	151	10000	11000	81000084	Huron	James L Hart	TS
8	151	34000	35000	81000376	Huron	Catherine	TS
8	154	55000	56000	81000063	Whittaker	Stony Creek	TS
8	159	6000	7000	81000084	Huron	James L Hart	TS
8	160	53000	54000	81000063	Whittaker	Stony Creek	TS
8	161	12000	13000	81000124	Whittaker	Willis	FL
8	165	41000	42000	81000125	Rawsonville	Willis	TS
8	166	21000	22000	81000064	Tyler	McGregor	FL
8	171	2000	3000	81000380	Huron R. Dr.	LeForge	TS
8	173	15000	16000	81000156	M-17	Hewitt	TS
8	175	7000	8000	81000056	Hewitt	Ellsworth	TS
8	176	25000	26000	81000058	Hewitt	Packard	TS

Figure 3.2. List of left turn events in RDCW database.

Retrieved data contained necessary variables to reconstruct the SV and POV trajectories sensed by the GPS, motion sensors and forward radars. Although very time-consuming, each left turn event was manually examined to test if it had an appropriate time range for the conflict metric analyses. The Start and End times were previously found automatically via SQL queries of turning motions. For analysis here, the retrieved data was extended by ten seconds before the Start time and after the End time, mainly to improve the latency time estimations between variables from different onboard modules.

3.2.2 Variables used in trajectory reconstruction

In order to reconstruct the vehicle trajectories, kinematic variables must be obtained by either directly accessing desired variables or deriving ones from retrieved variables. In order to obtain as good results as possible, the most suitable set of variables were to be determined before beginning reconstructing the trajectories. Variables required for absolute location of the vehicle were position and heading angle, and for the SV the only direct measurement was via GPS, which included differential correction. The accuracy of GPS is typically within a few meters, but varies depending on external conditions such as weather, the number of visible satellites, and the relative location of the receiver.

Random errors contained in GPS data could be removed using a filter, and for this purpose a standard Kalman Filter was adopted; this has the advantage that velocity and acceleration estimates are determined directly from the filter. In the SV trajectory reconstruction process, the accurate heading angle of the SV was considered crucial to maintaining the relative orientations of the SV and POV; otherwise, the error (increasing in proportion to the distance between the vehicles) would cause a severely distorted POV trajectory, since the POV motion was available only relative to the SV. For SV heading angle, there were three choices: 1) GPS heading angle directly available in the database, 2) obtained from integrating the on-board yaw-rate sensor $\dot{\theta}$ and 3) from the fusion of various angular measurement sources, including the KF generated trajectory. The first choice had a significant defect due to a truncation in the speed range of less than 5 m/s – see Figure 3.3 for an example.

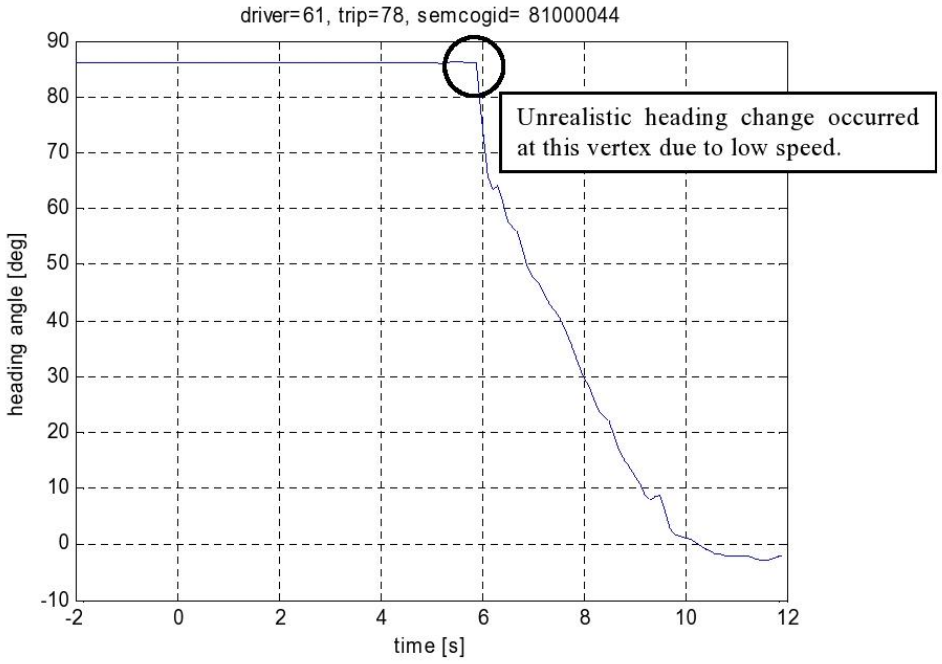

Figure 3.3. Inaccurate GPS heading angle.

For the second option, simple integration of the yaw rate creates a smooth curve but requires additional information for the initial heading, and is likely to cause an absolute drift error which accumulates as integration time increases. These motivate the third choice, angular measurement fusion, which actually performed best among the three. A procedure of the application of the angular measurement fusion is described in section 3.3.4.

The kinematic variables available from radar, used to generate the POV trajectory, were range R, range-rate $\dot{R}$ and lateral offset t_r. In addition, longitudinal acceleration a_x and lateral acceleration a_y collected from the onboard accelerometer were retrieved from the database in order to validate the heading angle of the SV resulting from the angular measurement fusion. Table 3.1 summarizes the sets of variables used in the process of reconstructing the SV and POV trajectories.

Table 3.1. Variables retrieved for the trajectory reconstruction process.

Vehicle	Variables for trajectory reconstruction	Variables for validation
SV	X, Y, $\dot{\theta}$	a_x, a_y, $\dot{\theta}$
POV	R, $\dot{R}$, t_r	

The sampling rates of these variables were 10Hz for (X,Y) and $\dot{\theta}$, and 20Hz for $(R,\dot{R},t_r)$ and (a_x,a_y). However, the sampling rate for the GPS was slower - the GPS data were updated at 4Hz and in a way that was not fully synchronized with the other variables, so that that samples, while recorded at 10Hz, this value typically stays constant for 2 or 3 time-steps. This fact (due to GPS clock update limitations) motivated the use of a time-varying KF for those variables with varying time steps.

3.2.3 Selection and fine-tuning of left turn events

Since unprotected left turns at signalized intersections are the focus of this analysis, the intersection type (protected or unprotected) is checked using available intersection database information. This data was not available for the entire region of Southeast Michigan where the FOT took place, only for Washtenaw County; for this reason, analysis of left turn events was restricted to this particular region. (In any case a certain amount of manual review was required in this analysis, so a complete analysis of the FOT data was not in fact practicable). MATLAB® was used to select candidate left turn events as shown in the screenshot below – figure 3.4. In the RDCW database, the radar data were sometimes unavailable (e.g. because of device limitations or malfunction in a specific test vehicle during a specific test period) and it also happened that a constructed SV trajectory was corrupted due to excessively noisy GPS data; such turning events were not considered amenable to analysis and were also removed from the candidate left turn event set. Moreover, the established list of turning events in the RDCW database sometimes do not contain real left turn events – e.g. a sharp curve in the road was being negotiated. Moreover, on occasion the RDCW system gave a warning to the driver during a left turn, and the driving behavior would not be representative, so such cases were excluded. Figure 3.5 shows examples of the three problematic cases: (a) corrupted trajectory due to large errors in the GPS data, (b) improperly assigned time range and (c) non left turn case. The stars indicate the trajectories for the period of time of an event and the solid lines show the trajectories in a time range with the extra 10s for both the

Start time and End time. The detailed method of trajectory reconstruction is explained in the later sections of this chapter.

As mentioned earlier, the time range between the Start time and End time could be modified from the original one. These times were checked visually for each left turn event and, if necessary, adjusted manually so that the trajectory was fully included within the time period. Criteria for proper time range were: 1) the trajectory was to start and end at small values of yaw rate, 2) the lowest speed before the turn was to be included and 3) any following or previous turn was to be excluded where necessary. Note that even if an original pair of the Start time and End time of a left turn event was changed to different values, it was specified by the original times along with the other numbers, Driver # and Trip #, to fully identify the event in the initial event table.

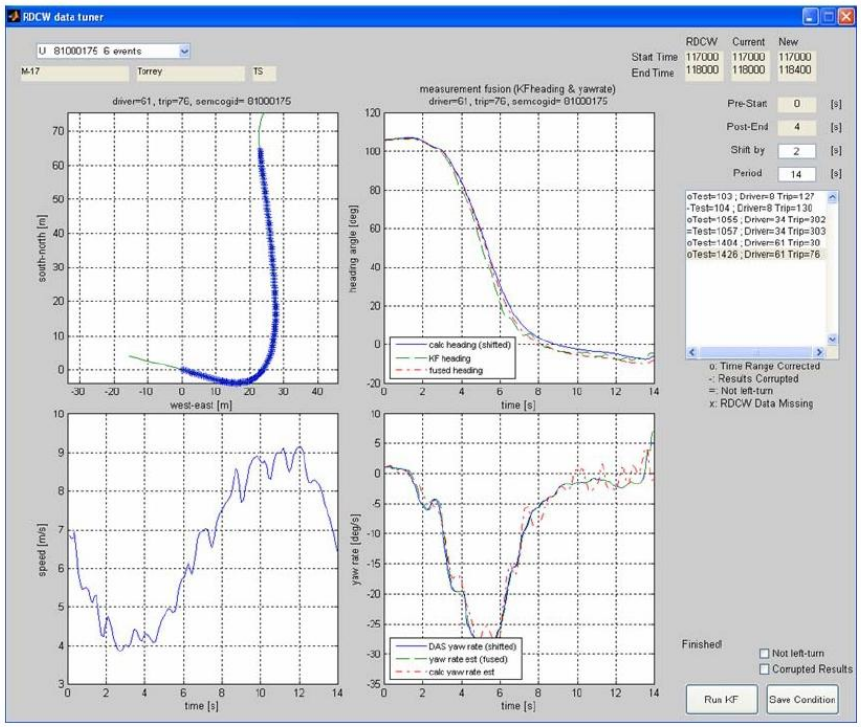

Figure 3.4. Screenshot of event screening process.

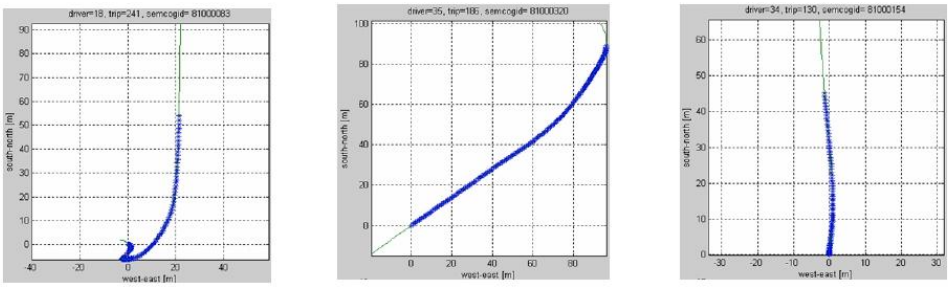

(a) Corrupted trajectory (discarded)

(b) Improper time range (corrected)

(c) Not a left turn (discarded)

Figure 3.5. Examples of inappropriate cases to be corrected or removed.

3.2.4 Summary of left turn events in RDCW

The tables below show summary information about all detected left turns in the FOT database after the initial data checks and screening were performed as mentioned in the last section. The initial set of left turns were detected via a simple algorithm based on vehicle kinematics: first each quadrant with a horizontal axis in the west-east direction and a vertical axis in the south-north direction was quantized. Then, if the heading angle of the SV changed from one quadrant to another in counter clockwise direction, the motion was considered as a candidate left turn. These events were then matched to locations in the database of Washtenaw intersections to generate the initial candidate left-turn set. Approximately 76% of all these left turns were usable for the analyses (Table 3.2), more than half of the intersections were provided with unprotected signals (Table 3.3) and approximately 51% of all the left turns detected were permissive left turns (Table 3.4).

Table 3.2. Summary of left turns in Washtenaw County in RDCW database.

	# of turns	Percentage
Valid left turns	1191	75.57 [%]
Corrupted results	116	7.36 [%]
Non-left turn events	149	9.45 [%]
Left turns but data missing	115	7.30 [%]
RDCW warning on	5	0.32 [%]
Total	1576	100 [%]

Table 3.3. Types of intersection in Washtenaw County in the constructed intersections database.

Intersection type	# of intersections	Percentage
Unprotected	199	58.19 [%]
Protected	108	31.58 [%]
Flashing light	35	10.23 [%]
Total	342	100 [%]

Table 3.4. Candidate left turns for analysis.

Intersection type	# of left turns detected	Percentage
Unprotected	594	51.12 [%]
Protected	476	40.96 [%]
Flashing light	92	7.92 [%]
Total	1162	100 [%]

In the following sections, methodologies of reconstructing trajectories of the SV and POVs are explained by using an example left turn (Driver=8, Trip=165, Semcog ID=81000125, Start time=41000 and End time=42000). In this example, there were the SV, an oncoming POV and a large articulated truck in front of the SV. The truck made a left turn before the SV, and the SV followed the truck without waiting for the oncoming POV. After these two vehicles made left turns, the POV appeared to enter the intersection (there was no absolute guarantee that the POV actually did, since it disappeared from the forward radars and video camera; on the other hand the continuous motion of the SV – as determined from the trajectory estimation – renders it highly likely).

SEMCOGID#81000125
Rawsonville Rd@ Willis Rd
Washtenaw County, Michigan
Driver=8, Trip=165
Start time=41000
End time=42000

Figure 3.6. Aerial view (above) and video image (bottom) of the example left turn event.

3.3 Subject Vehicle Trajectory Estimation

3.3.1 X-Y Kalman Filter (KF)

This first step of the overall analysis estimated the state variables of the SV by using a time-varying KF. Those states were position, velocity and acceleration variables in global intersection coordinates, $\hat{X}, \hat{Y}, \hat{v}_X, \hat{v}_Y, \hat{a}_X$ and $\hat{a}_Y$. The filtering process was achieved by using the KF twice by sweeping backwards in time to provide settled initial conditions, and then forward in time to create the estimated trajectory. Since the initial conditions were unknown, the backward filter gave these initial estimates for the forward KF to evolve from (see Appendix for more detailed information).

From the resulted KF estimates, the SV trajectory, speed and acceleration magnitude were plotted in Figure 3.7 along with the original path without filtering from the raw GPS data (see Appendix for coordinate transformation from the geographical coordinate system to Cartesian coordinate system), DAS speed and acceleration magnitude derived from the longitudinal and lateral acceleration components from the DAS. On the scale used, the position estimate is hardly changed, the speed estimate closely follows that measured on-board the vehicle, as does the acceleration estimate. The filter certainly smoothes these latter estimates, and there appears to be a slight time delay in the acceleration estimate. However, overall the results indicate a high level of correlation between direct measurement of speed and acceleration compared to those inferred from GPS measurements.

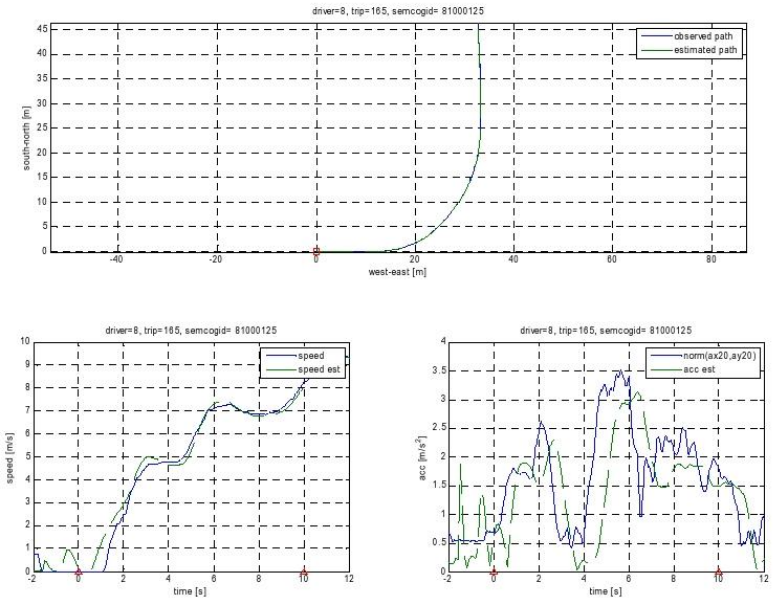

Figure 3.7. Results of the X-Y KF for the SV motion. The SV trajectory (top), speed (bottom-left) and acceleration magnitude (bottom-right).

3.3.2 Constant speed KF

Estimating yaw angle $\hat{\theta}$ from the velocity estimate $(\hat{v}_X, \hat{v}_Y)$ from the above X-Y KF causes a large heading error whenever vehicle speed is low - at low speed the following equation, used for the estimate, becomes ill-conditioned:

$$\hat{\theta} = \tan^{-1} \frac{v_Y}{v_X} \tag{3.1}$$

This can also be understood based on the fact that GPS location of the SV is obtained at a constant sampling frequency, even when the SV was stopped, and in this case the path direction is completely lost. In the example left turn event, the SV was stationary for approximately 3.5s in the early part of the event.

When the vehicle speed is small (so changes of the SV GPS position are small) the raw vehicle trajectory tends to possess a dense cluster of points in a small area, with values dominated by small errors. By magnifying the corresponding region in the SV path, it can be recognized that it was severely corrupted in the order of one centimeter (Figure 3.8). As a result, $\hat{\theta}$ obtained directly from equation (3.1) contains large errors in the low speed range (Figure 3.9).

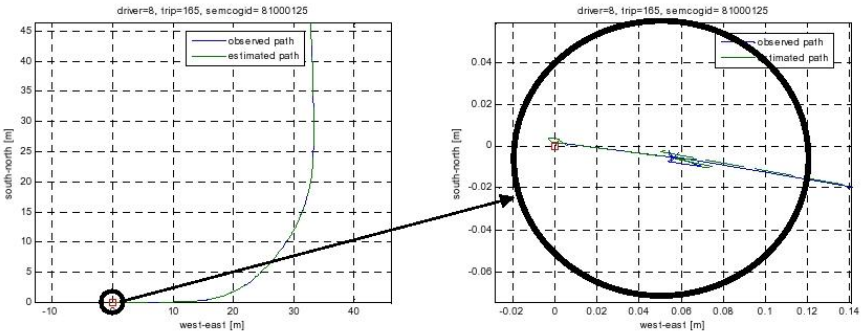

Figure 3.8. Corrupted trajectory due to the slow SV speed.

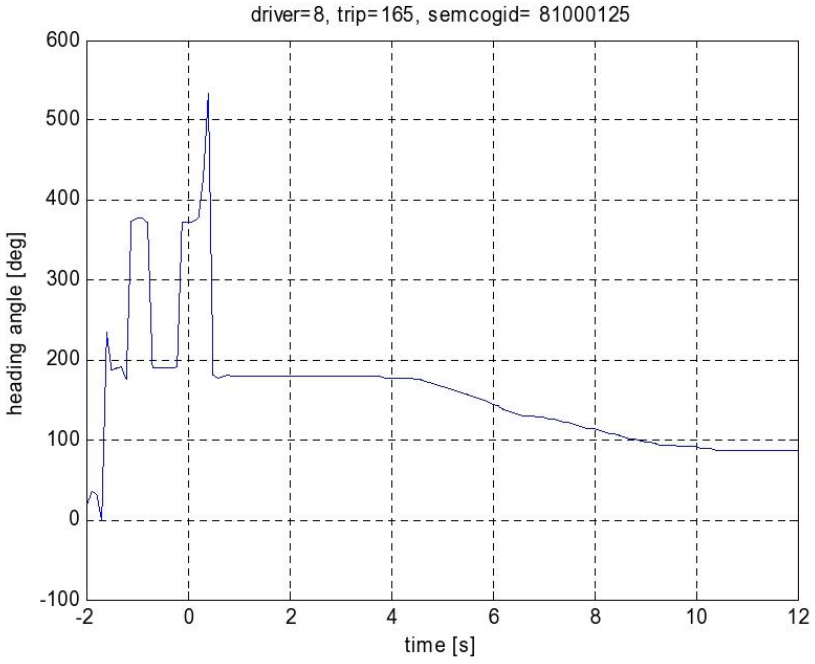

Figure 3.9. Heading error due to the low SV speed.

In order to eliminate such errors, a "constant speed KF" was used in which the trajectory for the measurement in the KF was modified to have equally spaced points with a sufficiently large spatial interval (1m) along the SV path (Figure 3.10(a)) with a fictitious time domain which was also equally spaced with an arbitrary value for the intervals (e.g. 0.1s). Consequently, after recovering the real time domain, the errors were removed in the resulted smoothed heading angle (Figure 3.10(b)). It can be seen that the new speed estimates in the fictitious time domain were roughly constant (Figure 3.11), but with the result that heading estimates become stable.

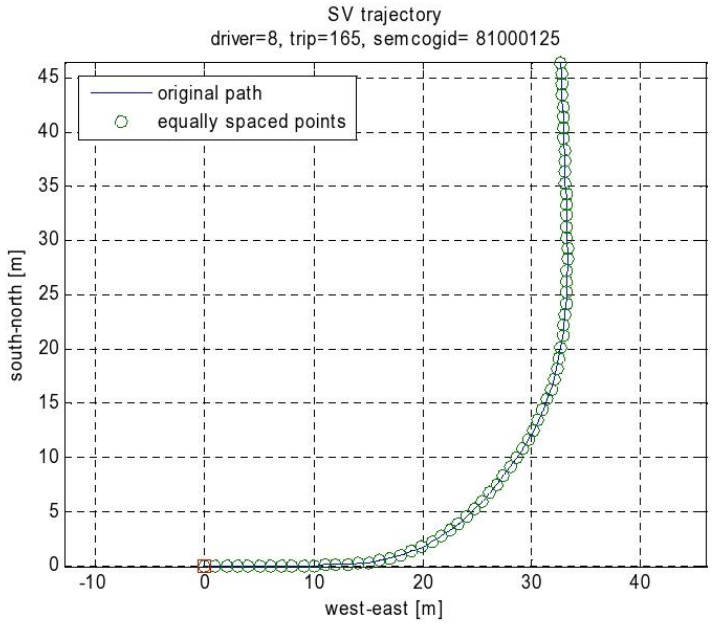

(a)

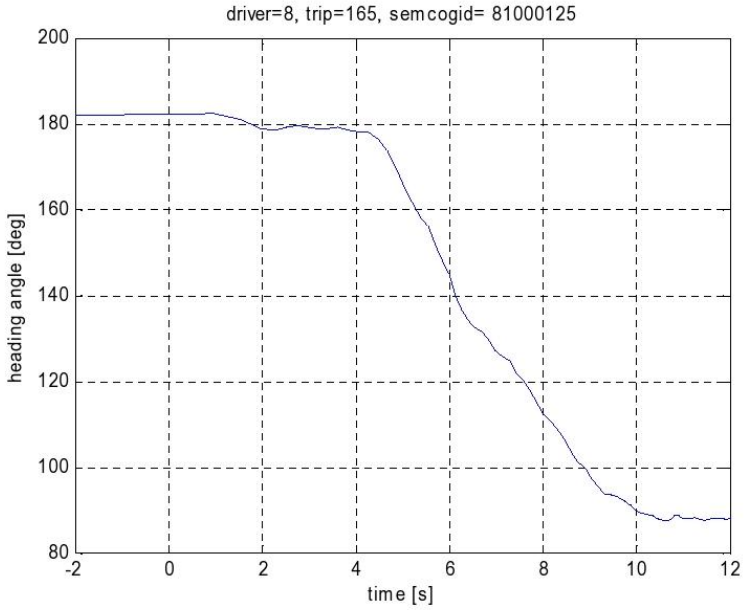

(b)

Figure 3.10. (a) Equally spaced points to avoid the heading error and (b) resulted improved heading angle.

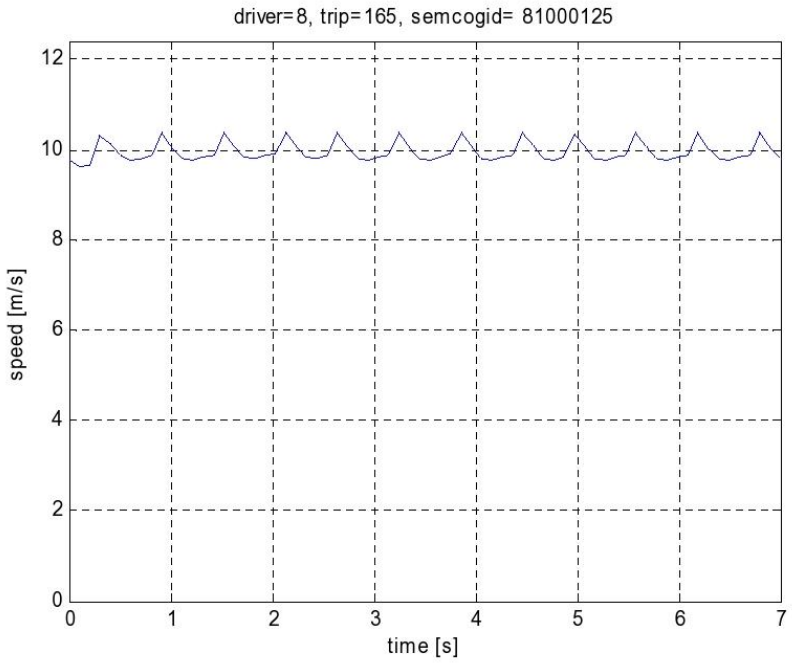

Figure 3.11. Speed estimate obtained from filtering the equally spaced trajectory in the fictitious time domain.

3.3.3 Latency time difference between the GPS and DAS

A further improvement was applied for the heading angle to be more reliably estimated, namely fusion of the above GPS derived estimation with on-board yaw rate sensor. The above estimated heading angle $\hat{\theta}'$ (GPS-based) and $\dot{\theta}$ (on-board yaw rate) were the measurements selected to be fused. Details are described in the next section. In this section, latency time estimation between these variables is explained. This is a preliminary process for the angular measurement fusion process. Due to different modules used in the data processing for the GPS and DAS, they were assumed to have distinct latency times. Therefore, the difference between those latency times, Δt_1, needed to be estimated. This latency time was then used to shift the DAS yaw rate along its time axis; once the variables have been aligned, they can be combined into a single estimate.

The yaw rate is first derived from the data in the constant speed KF, based on the notion that the curvatures of the SV path were the same for both of the time scales, the real time scale for the X-Y KF and the fictitious time scale for the constant speed KF. The curvature equation that connects the two cases is as follows

$$\kappa = \frac{\dot{\hat{\theta}}'}{\hat{U}} = \frac{\dot{\hat{\theta}}_{eq.spaced}}{\hat{U}_{eq.spaced}} \qquad \Rightarrow \qquad \dot{\hat{\theta}}' = \frac{\hat{U} \cdot \hat{a}_{y,eq.spaced}}{\hat{U}^2_{eq.spaced}}. \tag{3.2}$$

Then, Δt_1 is calculated by minimizing the cost function defined as,

$$\min_{\Delta t_1} f = \sum_{k=1}^{N} \left\{ \dot{\hat{\theta}}'(t_k) - \dot{\theta}(t_k - \Delta t_1) \right\}^2, \tag{3.3}$$

where N is the number of sampling points in the left turn event.

In the example event, the time shift between the GPS yaw rate, $\dot{\hat{\theta}}'$, and the DAS yaw rate, $\dot{\theta}$, was 0.45s by which $\dot{\hat{\theta}}'$ was delayed. By choosing the GPS time as the reference time, $\dot{\theta}$ was shifted by that amount to the right as shown in Figure 3.12.

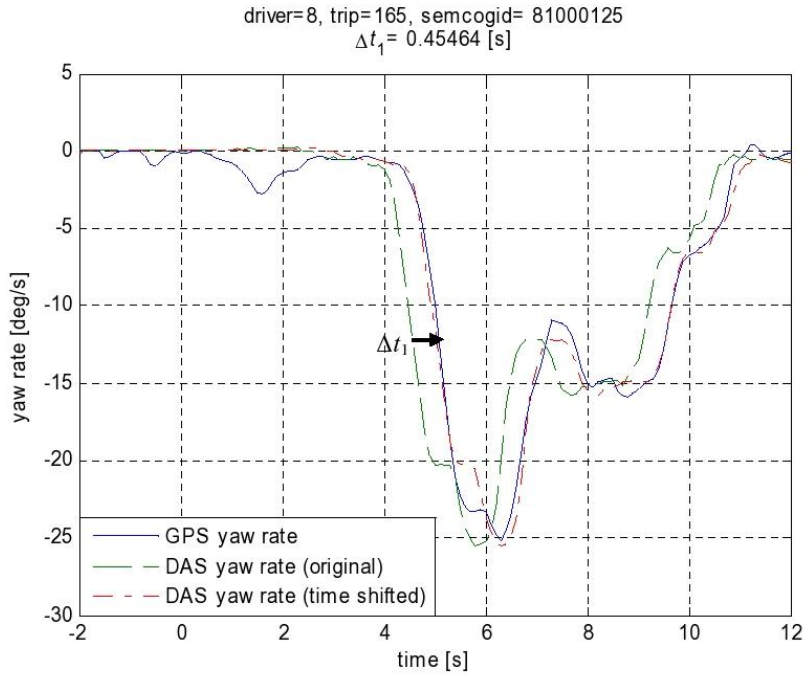

Figure 3.12. Result of the latency time calculation.

3.3.4 Angular measurement fusion

The motivation behind the fusion process were that the GPS always gave an absolute position over the entire time range but contained significant noise, while the yaw rate transducer could be used to derive an accurate angular change as long as the integration time was. The fusion was to provide a pooled estimate that improves either sensor used in isolation. The measurement fusion was adopted here because of its superior performance to the state variable fusion [8].

Fusion starts with measurements $\hat{\theta}'$ and $\dot{\theta}$. A Wiener process acceleration model was used for the angular motion, similar to the analysis of planar motion in the previous applications of the KF (see Appendix). The form of the state equation used here is that of 1-D motion in discrete-time:

$$\boldsymbol{\varphi}(k+1) = \mathbf{A}(k)\boldsymbol{\varphi}(k) + \mathbf{G}v \tag{3.4}$$

where $\boldsymbol{\varphi}$ is a 3 component state vector:

$$\boldsymbol{\varphi}(k) = \begin{Bmatrix} \phi(k) \\ \dot{\phi}(k) \\ \ddot{\phi}(k) \end{Bmatrix},\ \mathbf{A}(k) = \begin{bmatrix} 1 & T(k) & \frac{T(k)^2}{2} \\ 0 & 1 & T(k) \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{G}(k) = \begin{bmatrix} \frac{T(k)^2}{2} \\ T(k) \\ 1 \end{bmatrix}. \tag{3.5}$$

Here $\phi, \dot{\phi}$ and $\ddot{\phi}$ are the state variables of the system corresponding to the angle, angular velocity and angular acceleration, respectively. The fused measurement equation is described as

$$\mathbf{y}(k) = \mathbf{C}\boldsymbol{\varphi}(k) \tag{3.6}$$

where

$$\mathbf{y}(k) = \begin{Bmatrix} \hat{\theta}'(k) \\ \dot{\theta}(k) \end{Bmatrix},\ \mathbf{C} = \begin{bmatrix} \mathbf{C}_{KF} \\ \mathbf{C}_{DAS} \end{bmatrix} \text{ and } \mathbf{w} = \begin{Bmatrix} w_{KF} \\ w_{DAS} \end{Bmatrix}. \tag{3.7}$$

The components of the measurement matrix $\mathbf{C}$ are $\mathbf{C}_{KF} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}$ and $\mathbf{C}_{DAS} = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$. The fused measurement noise covariance is $\mathbf{R}$=diag(R_{KF}, R_{DAS}) where R_{KF} and R_{DAS} are the individual measurement noise covariance. Two example results are shown below including the original example used in the previous sections (Figure 3.13). In the second example, it can be seen clearly that the heading angle was smoothed by the measurement fusion particularly in the beginning part of the event. More importantly, the time rate of change of these signals are consistent with the yaw rates measured on the vehicle.

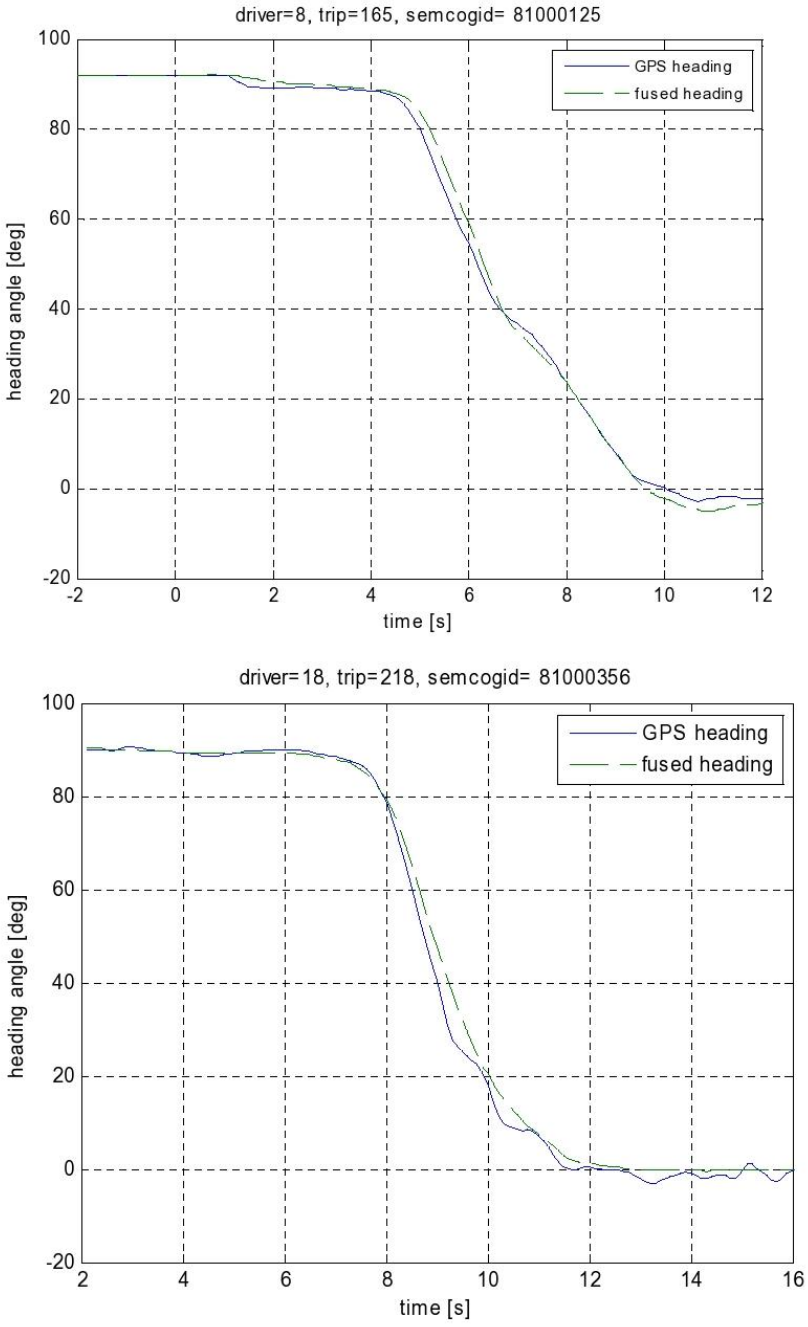

Figure 3.13. Results of the smoothed headings after the angular measurement fusion.

3.3.5 Validation of the fused heading

As the first part of the validation process, the fused heading angle, $\tilde{\theta}$, was examined by comparing it to the GPS-based heading, $\hat{\theta}'$, along with a comparison of the fused yaw rate, $\dot{\tilde{\theta}}$, with the DAS yaw rate, $\dot{\theta}$. As the second part, an independent validation was based on the resolved acceleration components $(\hat{a}_x, \hat{a}_y)$ being compared with those obtained directly from the vehicle-based sensor (a_x, a_y). In this second case, $\tilde{\theta}$ was required to transform $(\hat{a}_X, \hat{a}_Y)$ in the global coordinate system into $(\hat{a}_x, \hat{a}_y)$ in the local one.

Figure 3.14(a) shows that the mean error in $\tilde{\theta}$ vs. $\hat{\theta}'$ was 0.54 degrees and the standard deviation was approximately 2 degrees. The largest deviation occurred while the SV was turning. In the worst case, the error might be fully attributed to the estimation of $\tilde{\theta}$ in the measurement fusion process, e.g. because of noisy data in the GPS. It is highly unlikely that the errors are all due to $\tilde{\theta}$ but this does give a reasonable upper bound on the errors in estimated heading angle. In any case, remaining errors in the heading angle estimate will have some effect on the POV trajectory. As mentioned later, some further refinement of the POV trajectory will also be possible: the speed of the POV can be improved by using $\dot{R}$ measurements and a wrongly directed heading angle for the POV can be corrected by using an appropriate extrapolation method in order for the POV trajectory to reach the intersection for reasonable results of the conflict metrics. In terms of the lateral position error, the deviation would be about 3.5m for 2 degrees of heading error for the POV located at 100m away from the SV. At this stage we believe heading angle estimates are as accurate as can reasonably be constructed from the available data. The following figures show examples of heading angle and yaw rate comparisons.

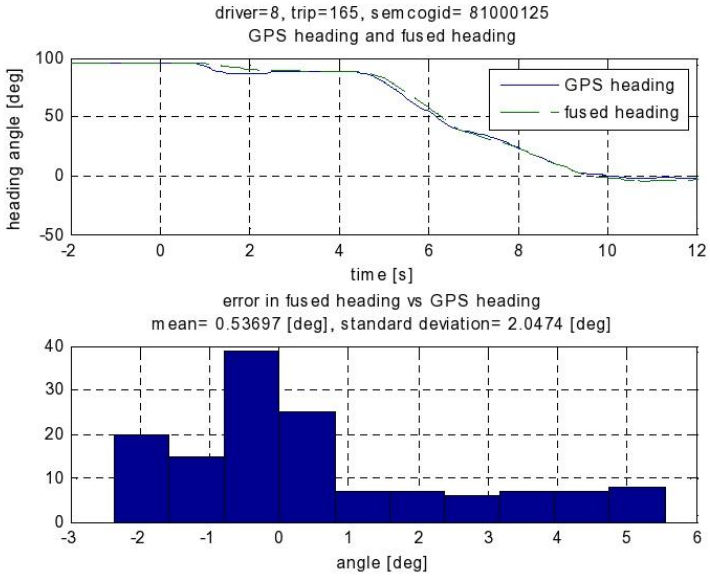

(a)

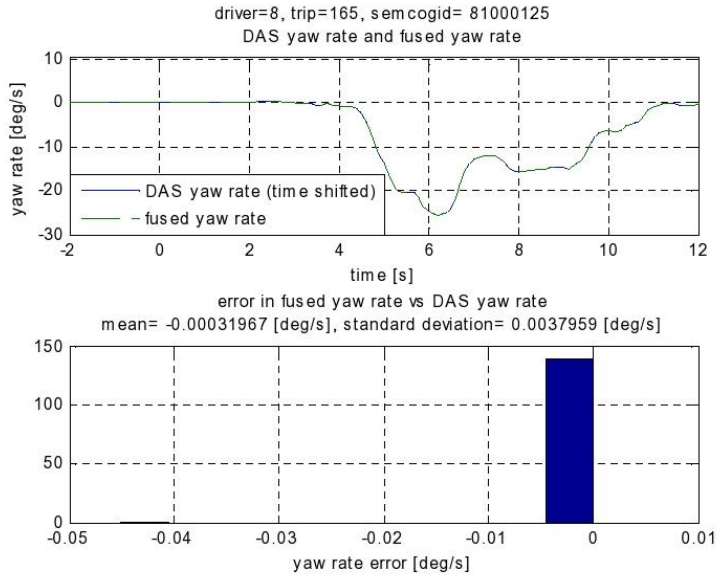

(b)

Figure 3.14. Comparisons of (a) fused heading vs. GPS heading (b) fused yaw rate vs. DAS yaw rate.

For further validation, the coordinate transformation from the global to local coordinate system was achieved by,

$$\begin{Bmatrix} \hat{a}_x \\ \hat{a}_y \end{Bmatrix} = \begin{bmatrix} \cos\tilde{\theta} & \sin\tilde{\theta} \\ -\sin\tilde{\theta} & \cos\tilde{\theta} \end{bmatrix} \begin{Bmatrix} \hat{a}_X \\ \hat{a}_Y \end{Bmatrix}. \tag{3.8}$$

Figure 3.15 shows that the coordinate transformation results reasonably represent each component of the SV acceleration. For the lateral components, it can be recognized that a_y was biased by about 0.5 m/s^2 which might be a calibration error or the result of cross-slope (since the SV was at rest in the beginning period).

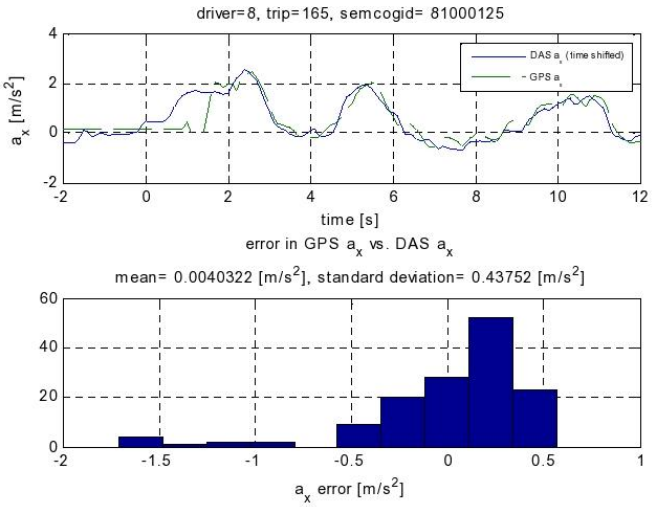

Figure 3.15a. Comparisons of the longitudinal components of the KF based acceleration vs. those of on-board (DAS) measurement

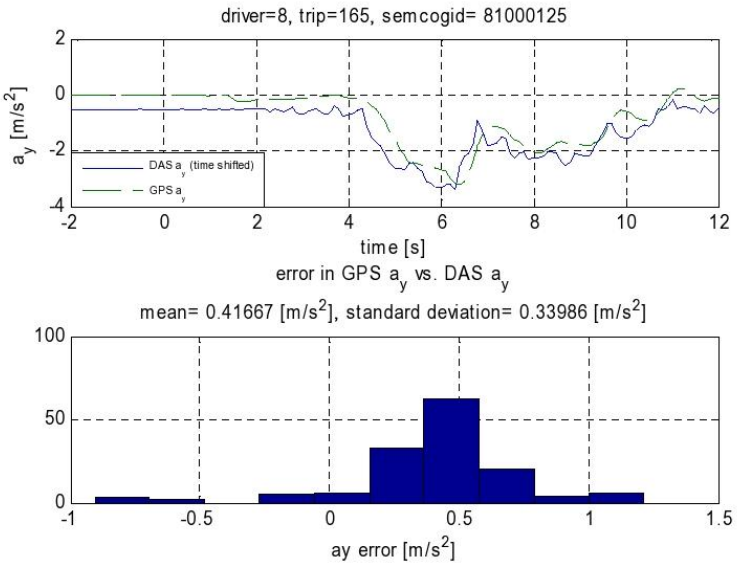

Figure 3.15b. Comparisons of the lateral components of the KF based acceleration vs. those of on-board (DAS) measurement

3.4 Trajectory Estimation for Principal Other Vehicles

3.4.1 Locating the target points

Before proceeding to the latency time difference calculation between the GPS and radars, it is worth explaining how the target points were located on the two dimensional surface. By using the SV trajectory reconstructed in Section 3.2, the target points detected by the forward radars had been located on the two dimensional plane. A diagram of relative locations of the SV and a target point is shown in Figure 3.16. Each radar gave information of the target motion by R, $\dot{R}$ and t_r (range, range-rate and lateral deviation – a surrogate for azimuth: see Figure 3.16. The radar was installed behind the fog light on each side located at 0.65m away from the mid-point of the frontal edge, with the center line of the field of view (FOV) canted 3deg outward. Therefore azimuth angle ϕ was calculated by the following equations:

$$\alpha = \sin^{-1}\left(t_r / R\right)$$
$$\beta = \alpha + 3 \times \pi / 180$$
$$L = R\cos\beta \quad . \qquad (3.9)$$
$$t_c = R\sin\beta + 0.65$$
$$\phi = \tan^{-1}\left(t_c / L\right)$$

Note that the radar system was established based on the ISO coordinate system while the SAE coordinate system was applied to any other parts in the analysis.

The resulted target point distribution is shown in Figure 3.17. In the figure, LF and RF indicate the radars on the left and right respectively and stationary objects (triangles) detected by the radars were also plotted for the use in the next section to remove the latency time difference between the GPS and radars. Previously, Δt_1 was used as an estimate of the latency time difference between the SV motion based on the GPS position and DAS, but it might not be the best choice – see the distorted POV trajectory in Figure 3.17. This result brought a need for a second adjustment of time difference between GPS and radar timing. This issue will be discussed in the next section.

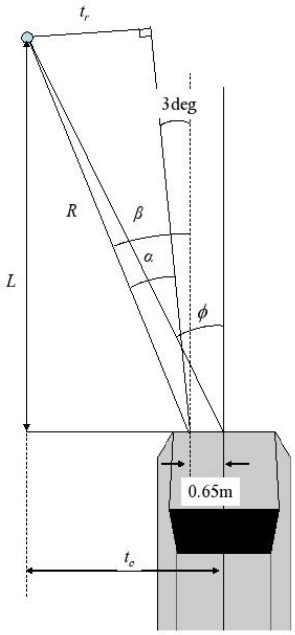

Figure 3.16. Geometry used in the azimuth calculation.

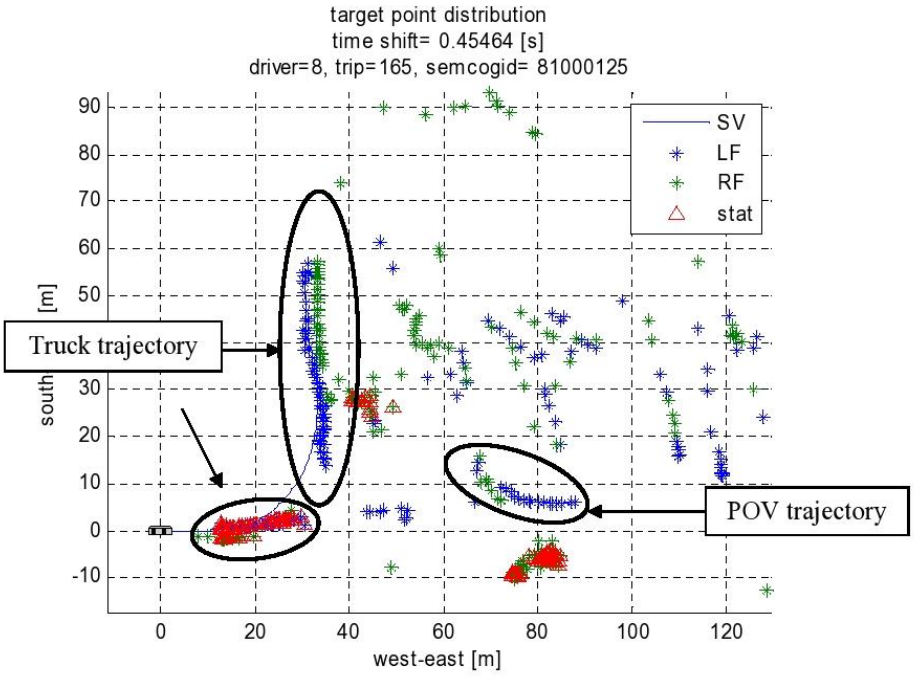

Figure 3.17. Target point distribution and detected static (slow) objects by using Δt_1 as the initial estimation of the latency time between the GPS and radars.

3.4.2 Latency time difference in GPS and radars

The above figure suggests that Δt_1 might not provide the correct latency adjustment between the SV and radars, so a second optimization problem was to be formulated in order to estimate an improved time shift Δt_2. This issue was addressed by applying a similar methodology as for calculation of Δt_1. The variable used here was, however, the scattering of the target points of which stationary objects consisted and the degree of the scattering of the target points was defined as the sum of the distances to the points of a stationary object from its geometric centroid. Therefore, for the j-th stationary object the cost function was defined as,

$$f_j = \sum_{k=1}^{n} \left\{ \left(x_{j,k}(\Delta t_2) - x_{j,c}(\Delta t_2) \right)^2 + \left(y_{j,k}(\Delta t_2) - y_{j,c}(\Delta t_2) \right)^2 \right\} \tag{3.10}$$

where $(x_{j,k}, y_{j,k})$ represented the k-th point in the j-th stationary object and $(x_{j,c}, y_{j,c})$ was the coordinates of the geometric centroid of the same object. Consequently, a multi-objective optimization was formulated for multiple objects, using the weighted sum method,

$$\min_{\Delta t_2} f = \sum_{j=1}^{m} w_j f_j \,, \tag{3.11}$$

where w_j was the weighting factor for the j-th stationary object and its value was defined as the reciprocal of the product of the average range, $\overline{R}_j$, and number of points, N, in the object, that is, $w_j = 1/(\overline{R}_j N)$. Figure 3.18 shows the corrected POV trajectory with Δt_2=0.09 s while Δt_1=0.45 s.

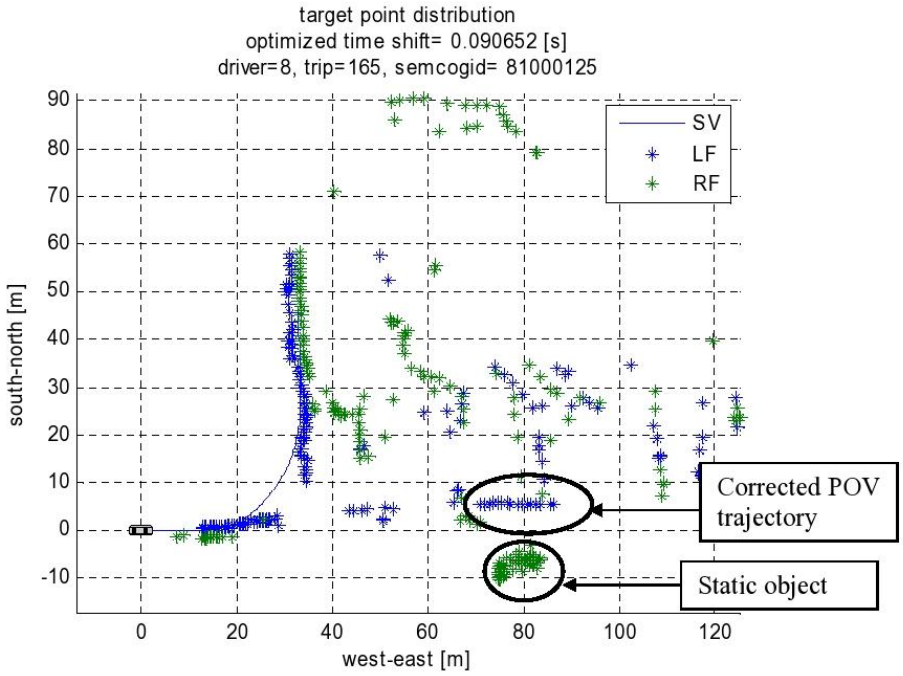

Figure 3.18. Result of the latency estimation by the multi-objective optimization to minimize the distributions of the static objects.

3.4.3 Target point association

For each of the radars, the following 4-step procedure was performed to construct tracks.

(1) Tracklet construction : target points are associated point-wise based on constructed connectivity matrices that contained connectivity information for each target point at the current time to other target points from one time step before and after.
(2) Ambiguity resolution : split or merged segments are resolved. (The connectivity matrices were constructed by using relatively loose threshold values for the connectivity criteria shown below in (1) and its result tends to contain such segments.
(3) Tracklet connection : fragmented segments (or remaining isolated points) due to focusing only on one time step for the point association in (1) are connected by extending the considered time range to a larger one by including a few more time steps, since an object might not be scanned at every time step.
(4) Connection by kinematic variables : speed and acceleration filters are applied, since there might be wrong connections due to the lack of information about the azimuth rate from the radars.

Details of each step are now explained

(1) Tracklet construction
Tracklets were constructed by the following criteria - neighboring points are connected provided:

$$\begin{aligned} &\left|\Delta t_r\right| \le \varepsilon_1 \\ &\left|\Delta R\right| \le \varepsilon_2 \\ &\left|\Delta \dot{R}\right| \le \varepsilon_3 \\ &\left|\Delta R - \dot{R}'\Delta t\right| \le \varepsilon_4 \end{aligned} \tag{3.12}$$

where for each time step between $t=t_1$ and $t=t_2$ ($t_2 > t_1$)

$$\begin{aligned} &\Delta t_r = t_r(t_2) - t_r(t_1) \\ &\Delta R = R(t_2) - R(t_1) \\ &\Delta \dot{R} = \dot{R}(t_2) - \dot{R}(t_1) \\ &\dot{R}' = \left(\dot{R}(t_2) + \dot{R}(t_1)\right)/2 \\ &\Delta t = t_2 - t_1 \end{aligned} \tag{3.13}$$

The target point IDs of the predecessor and successor of a point in the current time step are then stored in a *connectivity matrix*. The structure of the connectivity matrix is described in the following simplified example. Let us assume, in the time steps t_0, t_1, t_2, t_3, and t_4, there are sets of target points, (1, 2), (3, 4, 5), (6, 7, 8), (9, 10) and (11, 12), respectively, and the connectivity of each point, as defined in equation (3.12), is given by

arrows as shown in the tree diagram in Figure 3.19. Each arrow then defines a row in the corresponding connectivity matrix. For a given root the matrix expands at each time step until all branches terminate, e.g. at time t_4.

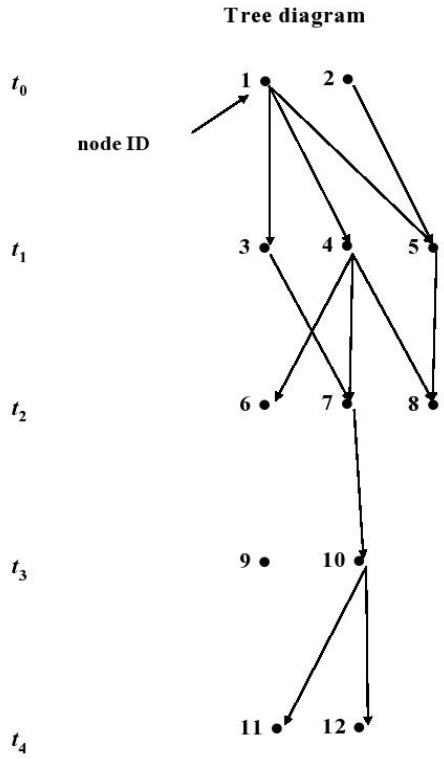

For the first root

At $t=t_1$

1 3
1 4
1 5

At $t=t_2$

1 3 7
1 4 6
1 5 8
1 4 7
1 4 8

At $t=t_3$

1 3 7 10
1 4 6 10
1 5 8
1 4 7 10
1 4 8

At $t=t_4$

1 3 7 10 11
1 4 6 10 11
1 5 8
1 4 7 10 11
1 4 8
1 3 7 10 12
1 4 6 10 12
1 4 7 10 12

For the second root

At $t=t_1$

2 5

At $t=t_2$

2 5 8

Figure 3.19. Structure of the connectivity matrix and example result of the breadth first search method.

To fully reconstruct the trajectories, a tree search method was needed and the "breadth first" search was adopted. This algorithm searches all connecting points in one level and proceeds to the deeper level, starting with a root node. Since there are two root nodes in this example, two sets of search must be performed. The result of this tree search method for this example consists of 8 paths for the first root node and 1 path for the second root node. At this stage, all candidate paths are retained.

Figure 3.20 shows target trajectories in the example left turn event, including multiple connections (which are hard to see in this figure). The target points (represented by stars) detected by the left radar were connected by lines. The circles and squares indicate the start and end points of the segments respectively. The trajectories of the truck in front of the SV and vehicle on the opposite traffic observed in the video were created automatically by using the methodology explained above. There were two separate trajectories of the truck, since the radar once lost it when it disappeared outside of the radar range. Also, there were many fragments that were unnecessary for the conflict metric calculation in the outside of the roads and these were to be removed.

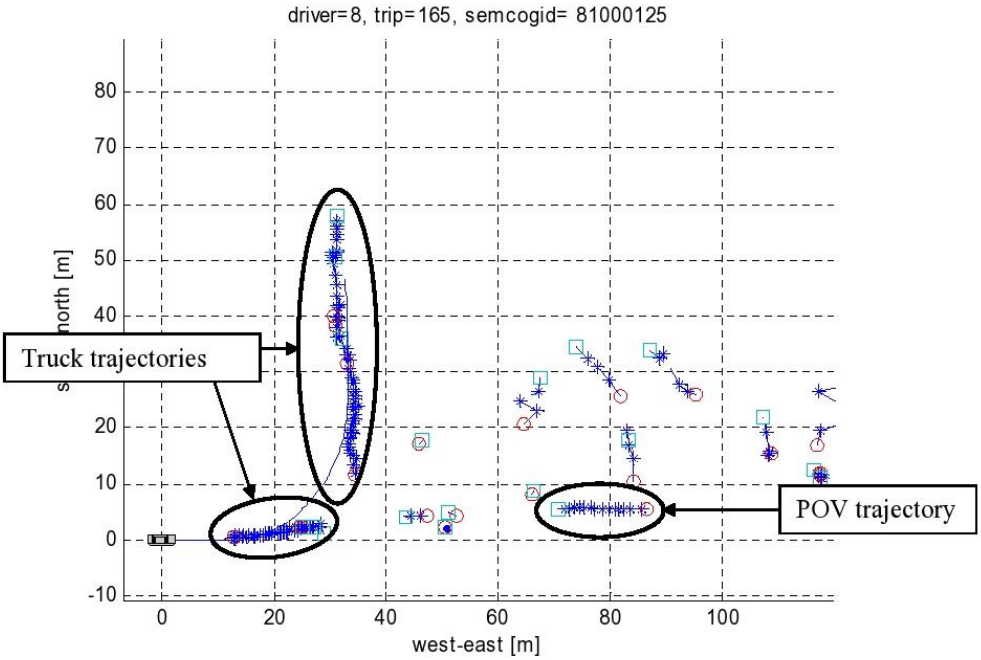

Figure 3.20. Connected target points by the tracklet construction.

(2) Ambiguity resolution

Split or merged segments were eliminated by using the normalized distance squared (NDS) to the ambiguous node [2]. The ambiguous node was located either at the next split point or at the merging point. The NDS was defined as,

$$d_{i,j} = d\mathbf{y}_{i,j}{}^{T}\mathbf{S}_i^{-1}d\mathbf{y}_{i,j} \tag{3.14}$$

where

$$d\mathbf{y}_{i,j} = \mathbf{y}_i(j) - \mathbf{y}_i(j-1)\,. \tag{3.15}$$

The last equation for $d\mathbf{y}_{i,j}$ was the measurement residual and $\mathbf{S}_i$ was the innovation covariance at the j-th point (ambiguous node) of the i-th segment. Figure 3.21 shows a flowchart of the ambiguity resolution process.

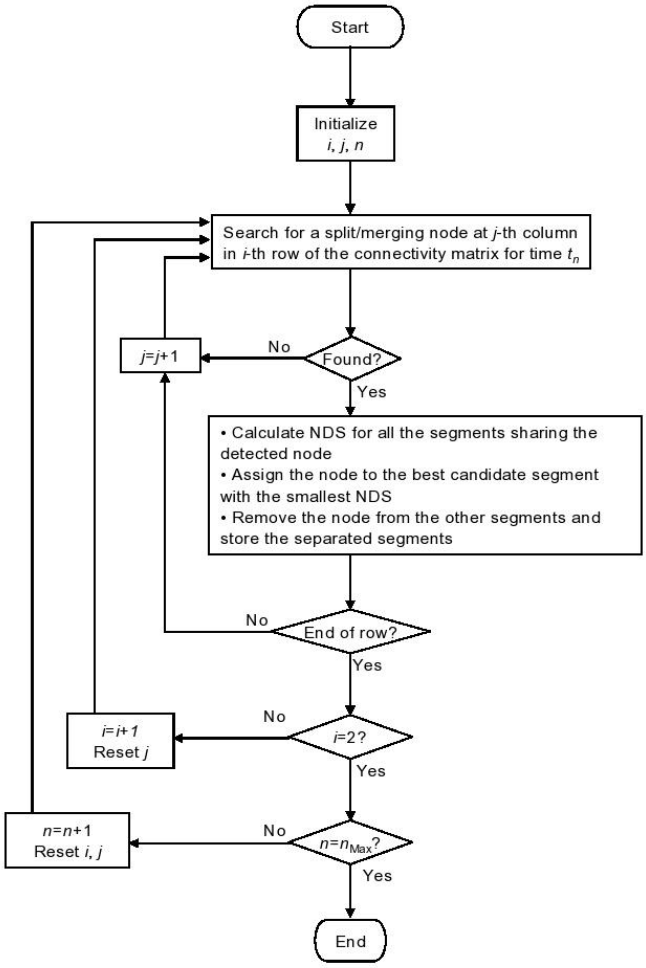

Figure 3.21. Flowchart for ambiguity resolution.

A typical result of the ambiguity resolution is shown in Figure 3.22. After the tracklet construction step in (1), two ambiguous nodes were found indicated by diamond marks in the figure. Several isolated points can be recognized at points of the overlapped marks of a circle (initial point of a segment) and square (end) as remainders after the separation of segments to which the ambiguous node was not assigned. In this particular case, the resulted segment after the resolution process started from the ambiguous node 2. The reason for this is that this ambiguous node was both merging and splitting node and was assigned to different segments in the merging resolution and splitting resolution. Note that here there is no intention to remove noise from the signal, just to define a single track, one point per time-step.

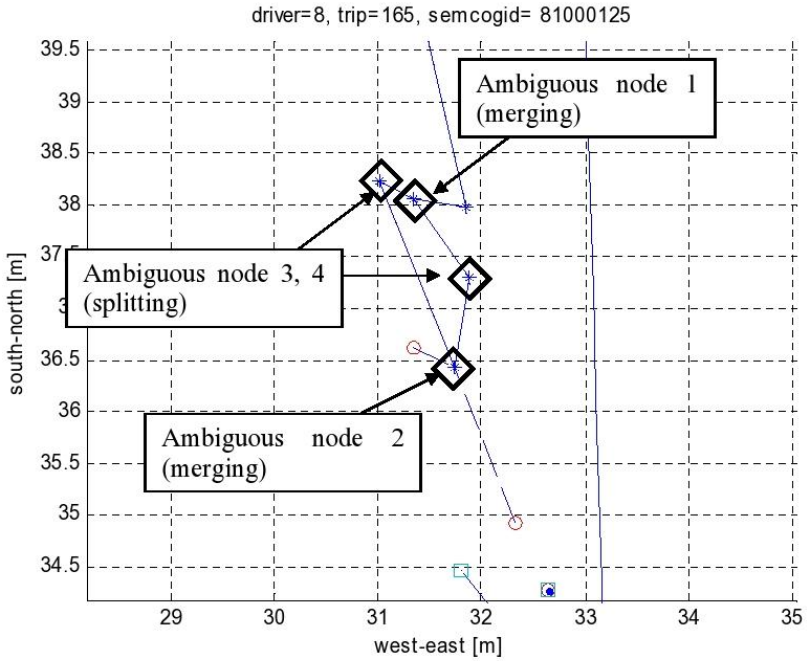

Figure 3.22 (a) Ambiguity resolution - ambiguous nodes

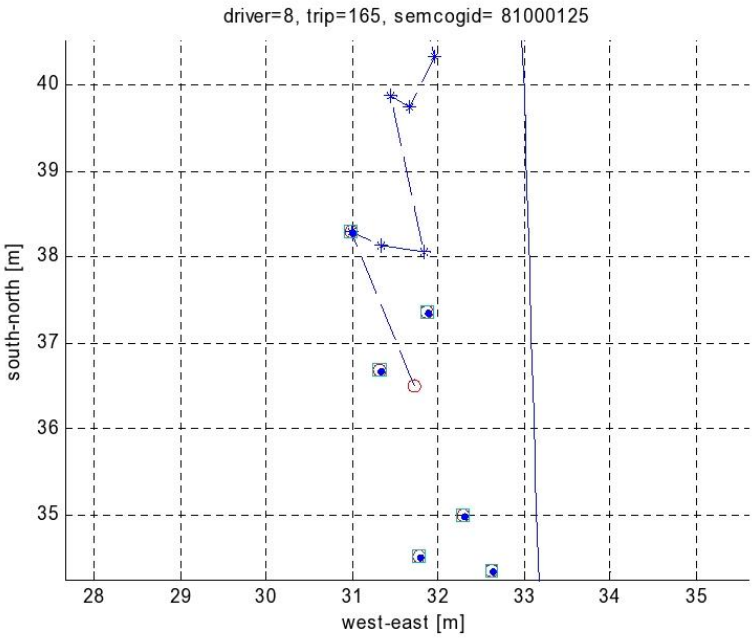

Figure 3.22 (b) Ambiguity resolution - results

(3) Tracklet connection

Here the problem is to connect tracklets when they are split due to missing data. The same criteria used in the first step were adopted, but with an extended time range. When multiple choices were possible, the ambiguity resolution described in (2) was employed. Unlike the tracklet registration in (1), partial tracks were available in this stage. Therefore, the ambiguity resolution could be applied without preparing all possible associated segments.

(4) Connection by kinematic variables

Speed between every two contingent points in a segment was calculated geometrically using the target-point positions and, if there were any segments whose speed was faster than a threshold value, the segment was separated at the points. This was continued until no separation was found. The issue which motivated the use of this filtering was that it was possible for target points from different objects during the SV being turning to be connected wrongly due to the lack of the angular speed of the azimuth. In other words, if there are points in distant locations of which R, $\dot{R}$ and ϕ of a point are close to those of the other, they may be wrongly connected based on the methodology in the previous stage. There was one separation in the example case as shown in Figure 3.23.

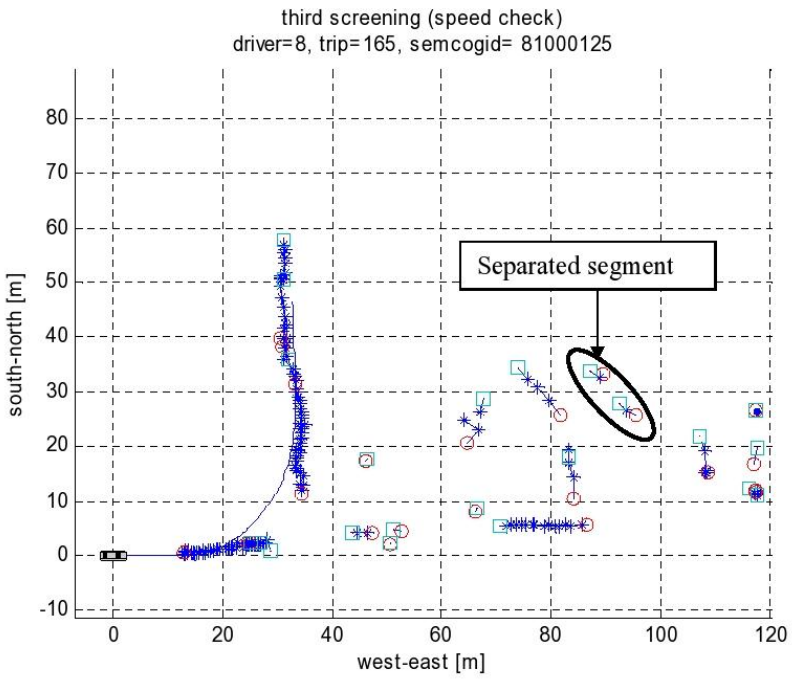

Figure 3.23. Result of the segment separation by the speed filter.

Similarly, the acceleration was utilized to search possible additional connections between the segments. It was also found geometrically in the same manner as in the speed filter. In addition, since a length of one time-step (0.1s) was considered on searching the target-point connectivity, it was expanded to a longer time period (e.g. 0.5s).

For a pair of segments to be evaluated for their association, two acceleration values were calculated by choosing the last point from one segment and the first two points from the other for the first acceleration and the last two points from the former segment and the first point from the latter for the second acceleration. Then, the difference of these values was compared with a threshold value. In the example case, there was no segment associations found, so the following result was taken from another left turn event (Figure 3.24).

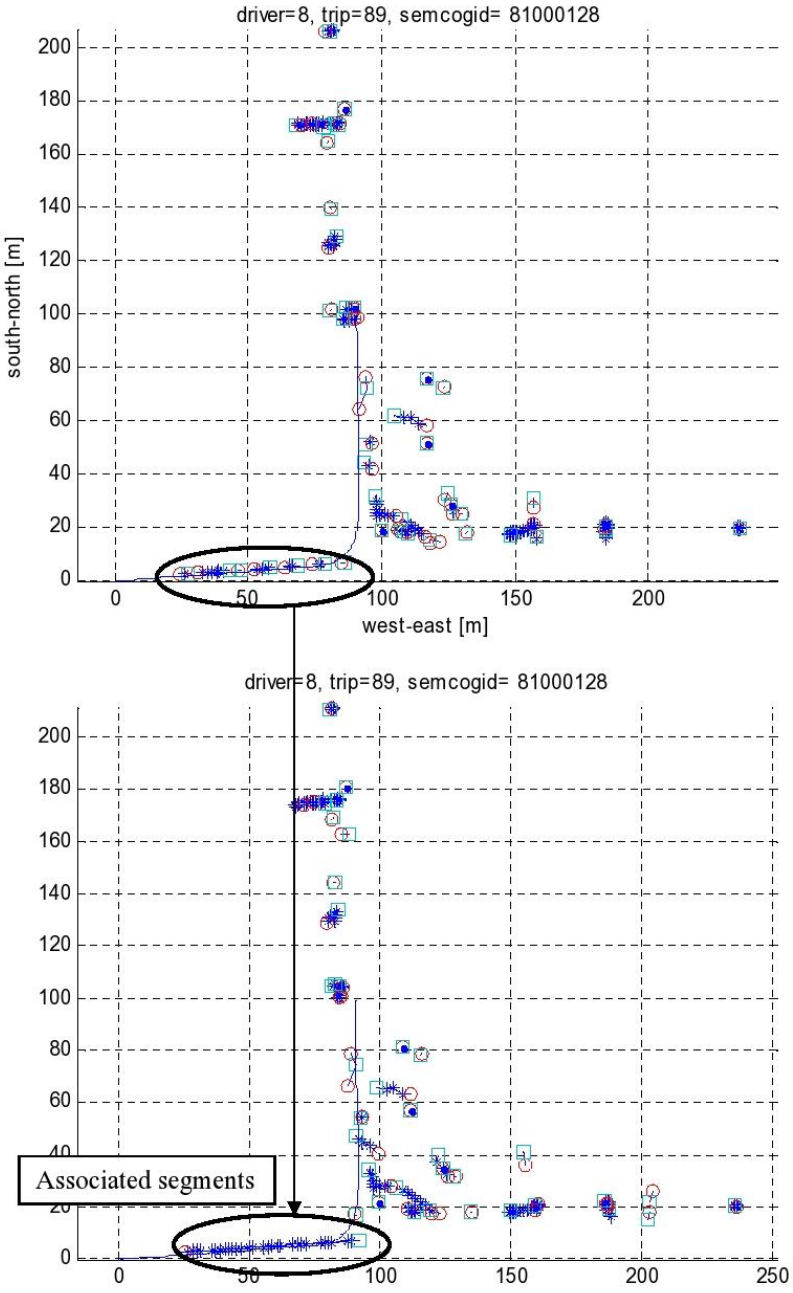

Figure 3.24. Segment association by the acceleration filter.

3.4.4 X-Y Kalman Filter for POVs

As in the reconstruction of the SV motion, the X-Y KF was used for the POVs to estimate their kinematic variables (position, velocity and acceleration in 2D global coordinates). The relative speed and distance of the POV must be reproduced properly to make it possible to calculate an accurate conflict metric. Figure 3.25 shows comparisons between the KF results and radar data for the range and range rate for the POV (track=1) and the large truck (track=6). Although, in both cases, the ranges were accurately found from the KF position estimate, the range rates contained large errors in both of the POVs. This is seen to be due to a lack of data in the POV trajectory – there are insufficient data points available in the POV tracks, and they did not converge with accurate initial estimates and state covariance obtained from the backward KF.

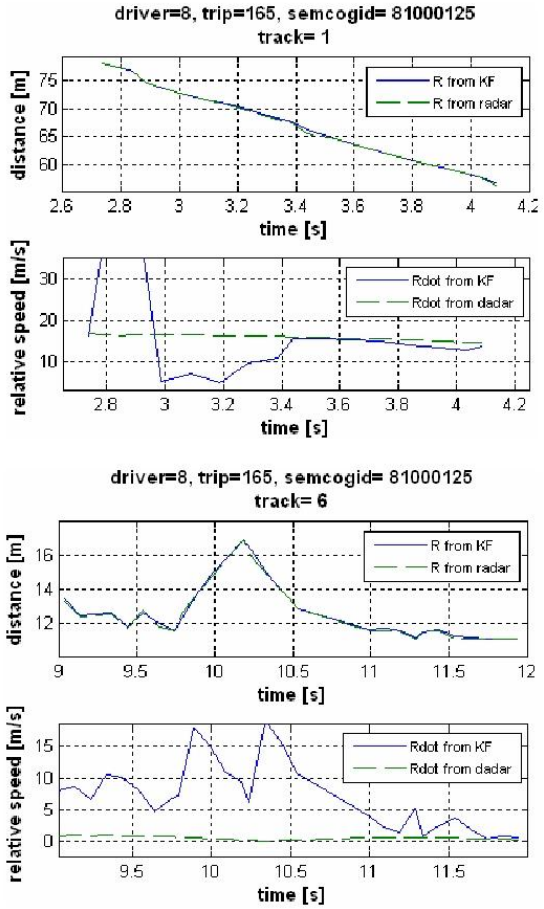

Figure 3.25. Comparisons of speed estimates of POV trajectories of the thru-vehicle (track=1) and the truck (track=6) with the relative speeds from the KF velocity estimates.

Two ways might be considered to resolve this problem in a simple manner. One way would be a use of any suitable curve fitting technique followed by an interpolation to increase the number of points for the KF to have enough number of data points. The other way simply replaces the velocity estimate with another velocity made independently of the SV velocity estimate and it can be calculated from the POV heading angle and $\dot{R}$ from the radars. In this work, the second approach was adopted because of its simplicity and anticipated improvement in accuracy. The new POV velocity estimate was computed by the following equations,

$$\begin{aligned} \mathbf{v}_{POV} &= \mathbf{v}_{SV} + \mathbf{v}_{POV/SV} \\ \mathbf{v}_{POV/SV} &= U_{POV/SV}\cos\theta_{POV}\mathbf{I} + U_{POV/SV}\sin\theta_{POV}\mathbf{J} \\ U_{POV/SV} &= \frac{\dot{R}}{\cos(\theta_{POV} - (\theta_{SV} + \phi))} \end{aligned} \tag{3.16}$$

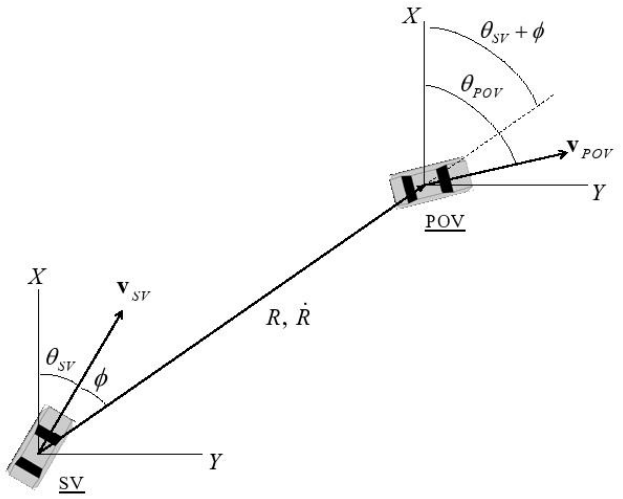

Figure 3.26. Orientations of the SV and POV.

We included an assumption that the POV heading angle was $\theta_{POV} \approx \theta_{SV} + \pi$. Then, since the FOV of each forward radar was 8deg, and it was mounted with an offset outward by 3deg (to give 7deg of coverage each side from the SV heading direction) the speed error would be approximately 0.75% if the SV orientation was parallel to the POV's direction of travel. However, this error would become larger as the SV heading angle changed during a left turn due to the lack of information of the time rate of change of azimuth. For example, if the SV heading angle is 10deg and if the POV is located within the FOV of the forward radar, the error becomes about 4.4%.

The result for the through vehicle after replacing the KF velocity estimate is shown in Figure 3.27. The inversely calculated range rate from the velocity-replaced POV (non-relative motion) very closely matched the radar $\dot{R}$. This result indicates the usefulness of the method to replace the noisy KF velocity estimates at least if involved vehicles are moving satisfying the assumption explained above. Note that, because of the assumption, the result for the truck, which was also turning left and followed by the SV in the same lane, was omitted.

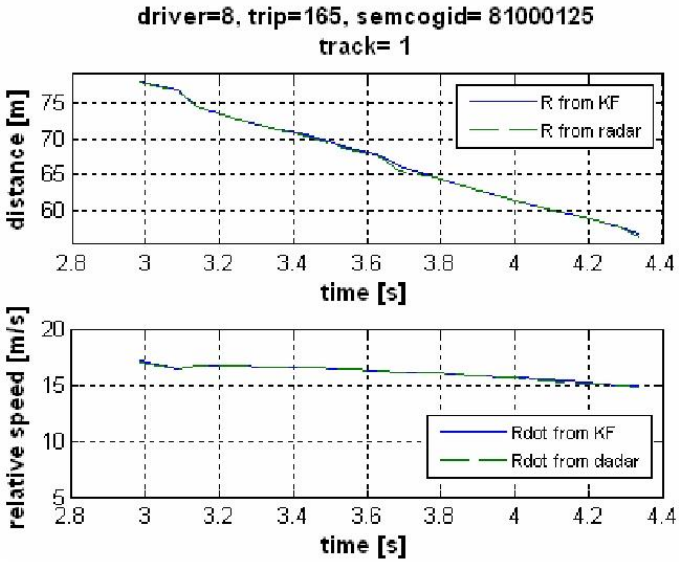

Figure 3.27. Modified results of the X-Y KF for the through vehicle with the alternative velocity.

3.4.5 Extrapolation of POV trajectory

With the FOV of the forward radars (8deg), the POV cannot be detected all the time during the SV being turning, but its future trajectory must be predicted for the conflict metrics. Extrapolation of the trajectory was achieved by assuming that the POV kept the velocity of the last point of the trajectory after it disappeared from the radar, with some adjustment made to keep the POV in its current lane. Such assumptions appear inevitable, since the POV motion is not available once the SV has turned.

As mentioned, a primary issue in the extrapolation process using the POV velocity is direction rather than speed- the final part of such POV trajectory may direct it off the road and cause a significant error in an extrapolated trajectory. Hence, for extrapolation, the magnitude of the velocity vector was used, and, for the direction, the POV was assumed to travel on an imaginary straight path which was created by connecting two points: the last point of the POV trajectory and a point defined at 3.6 m (typical lane width) away from the left turn lane where the SV was located to the left of the SV in its lateral direction. This is illustrated in Figure 3.28.

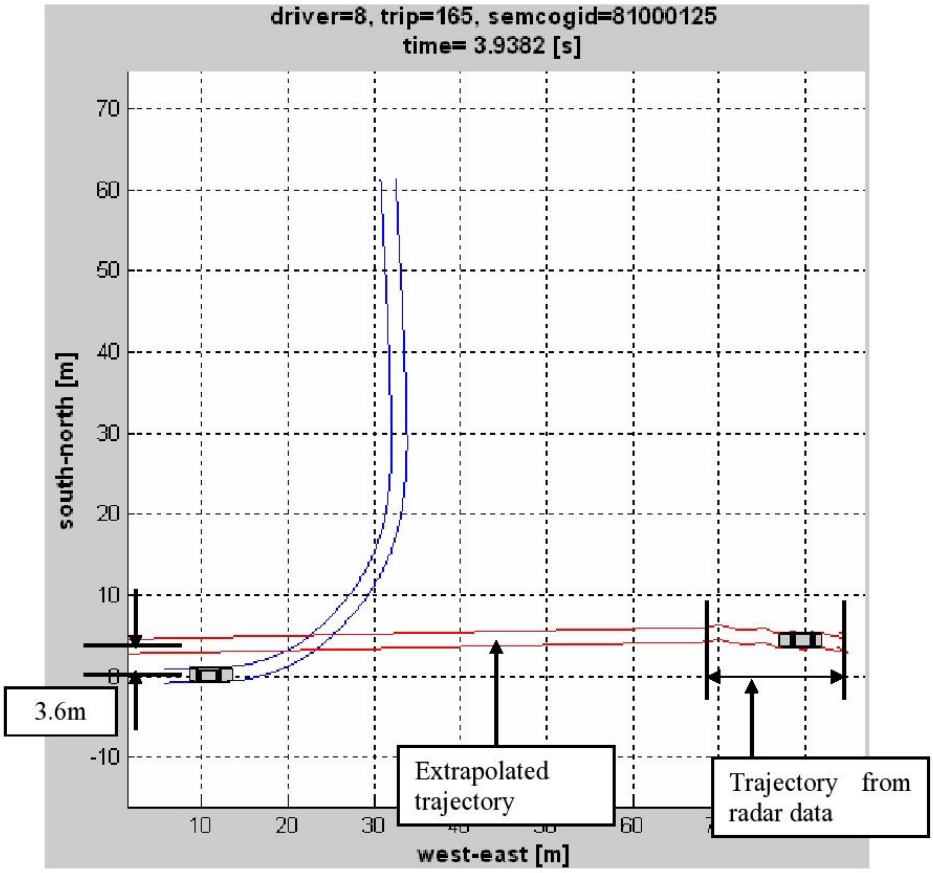

Figure 3.28. Extrapolated POV trajectory for conflict metrics.

Finally, the obtained trajectories were drawn on an aerial map of the intersection as well as of two other left turn examples. It can be noticed that vehicle trajectories have not fit in the traffic lane perfectly due to errors in the GPS and map data. Nevertheless, it seems fairly acceptable for conflict metric analysis. Besides, in terms of the relative positions between the SV and POV, each of the results shows a consistent relationship with the linear extrapolation to compensate the error of the lateral position of the POV. Therefore, if the heading angle of the SV is accurate, the error of the heading angle of the POV is small at least at locations near the SV.

SEMCOGID#81000125
Rawsonville Rd@ Willis Rd
Washtenaw County, Michigan

Driver=8, Trip=165
Start time=41000,
End time=42000

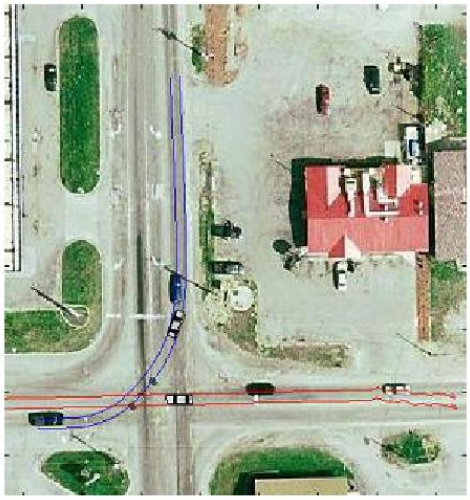

SEMCOGID#81000083
Huron River Dr@Superior Rd
Washtenaw County, Michigan

Driver=39, Trip=25
Start time=37000,
End time=38000

SEMCOGID#81000148
M153@Prospect Rd
Washtenaw County, Michigan

Driver=39, Trip=42
Start time=79000,
End time=80000

Figure 3.29. Aerial pictures of intersections with reconstructed vehicle trajectories.

4. Conflict Metrics

4.1 Review of Conflict Metrics

4.1.1 Conflict point

In this work we consider path crossing conflicts, where the geometric paths of a subject vehicle (SV) and principal other vehicle (POV) are considered fixed. Conflict metrics are then based on actual or predicted speeds along these paths from some initial point, and hence depend on the timing of arrival and departure of both vehicles from the encroachment zone - the overlapped area made by trajectories. A collision point exists if at least one point in the encroachment zone is simultaneously occupied by both vehicles – the first such occurrence defines a (predicted) collision point. Then the simplest measure of conflict is the time to collision (TTC): the time required for both vehicles to reach the collision point, normally based on the assumption of constant speed.

In car-following or merging conflicts, the encroachment zone is extended in space and TTC exists in most cases where a traffic conflict exists – if there is no collision point then the conflict is generally absent. By contrast, in path crossing conflicts, a "near miss" condition may be predicted without any collision point existing. Hence TTC is a relatively poor measure of conflict in such cases. Rather, the more natural measure of conflict relates to time gap – either with the SV (the turning vehicle) exiting the encroachment zone ahead of the POV (the thru vehicle), or with the SV turning after the POV. For the left turn across path – opposite direction (LTAP/OD) cases considered in this study, such metrics may include gap time (GT), post-encroachment time (PET) and leading and trailing buffers (LB and TB).

All forms of gap metric (defined more precisely below) are determined from (i) current vehicle positions and speeds (ii) assumptions or knowledge about time evolution of speed vs. time (iii) information about vehicle shape. The importance of the first two factors is obvious, while vehicle shape has only a subtle effect on the timing and location of collision or conflict points – see Figure 4.1. Compared to the uncertainties in vehicle trajectory measurements, the effect of shape on the collision point is very small. Since we do not necessarily predict or detect a collision, we need to define conflict points based on the location of a fictitious collision point when the motion of the SV is either delayed (Figure 4.2(a)) or advanced (Figure 4.2(b)). In these cases it is clear that – depending on the boundary shape of the vehicles – the collision would take place close to the relevant vertex of the encroachment zone. Hence we assume that the conflict point is precisely at one such vertex, and hence all forms of gap timing can be determined by comparing arrival and departure times at the relevant vertex.

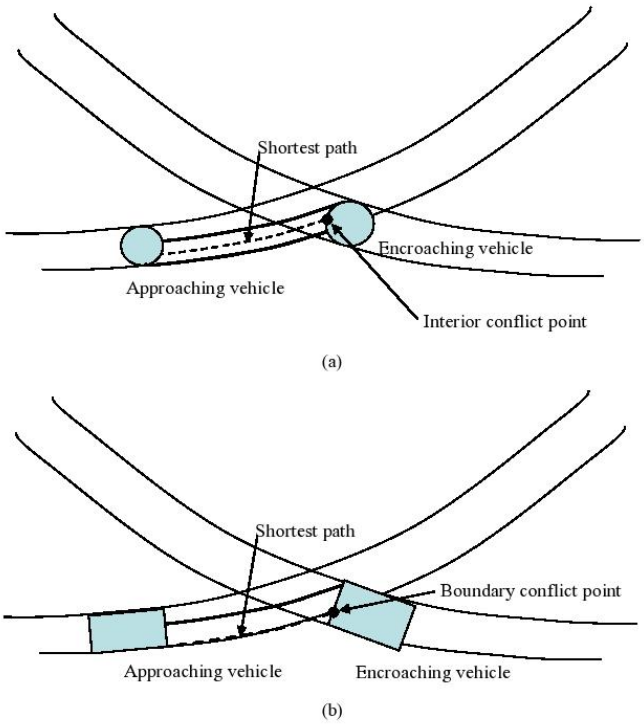

Figure 4.1. Varying locations of conflict points depending on the vehicle shapes.

4.1.2 Gap Time

The gap time (GT) is defined as

"Time lapse between completion of encroachment by turning vehicle and the arrival time of crossing vehicle if they continue with same speed and path." [6]

Therefore, the GT changes over time during a conflict event. In this definition, the GT is considered only when the SV completes (or is expected to complete) its turn before the POV enters the encroachment zone, but it can be used for the alternative case in which the order of encroachments of the vehicles is switched. The former case gives a positive value and the latter negative, assuming no collision occurs. Provided no collisions occur, and the GT is not calculated too far ahead of the conflict, the distributions of positive and negative GT's do not overlap and the same variable name can be used without ambiguity. Clearly the appropriate GT should be based on the corresponding conflict point shown in Figure 4.2. In this study, no collisions were detected, so the actual sequence (SV crosses the encroachment zone first or second) is uniquely know, and this determines which

"flavor" of GT is calculated. The important point here is that the GT calculation assumes constant speed for both vehicles, and this produces potentially misleading results, especially when the SV is stationary of moving slowly before making a turn – the constant speed "prediction" ignores the reality of the turning vehicle accelerating as part of the turning maneuver.

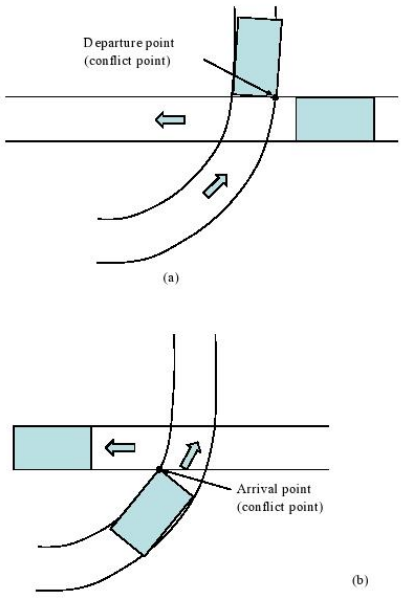

Figure 4.2. Left turn conflict points. (a) SV crosses first (conflict point is departure point of the SV) (b) SV crosses second (conflict point is arrival point of the SV)

4.1.3 Post Encroachment time

The post-encroachment time (PET) is defined as [1] "…the time from the end of encroachment to the time that the through vehicle actually arrives at the potential point of collision."

Unlike GT, PET does not include a predictive element; instead it provides a single value - the actual spare time from measured trajectory data. In normal usage, PET applies to the case where the SV turns first; in this case PET coincides with GT computed at the conflict point shown in Figure 4.2(a), assuming POV speed is in fact constant. Otherwise it is the time between the event shown in Figure 4.2(a) and the arrival of the SV at the conflict point

When the SV turns after the POV, a second post-encroachment time exists (we can also designate this PET) – the GT at the conflict point shown in Figure 4.2(b); the roles of the two vehicles are simply reversed, and in this case the PET will be treated as a negative quantity. Although this case may not be as significant as the case where the SV turns

first – the SV has a much greater degree of control in managing the conflict – the measure is needed to complete the overall gap acceptance picture determined from measured data.

4.1.4 Reference motion profile

As mentioned above, a specific limitation for the GT metric is that it is calculated using a constant speed assumption. If the SV is stopped at the stop bar or moves slowly before the start of the turn, the GT computed at this time tends to become small or negative, indicating a large but inappropriate conflict level - once the SV starts turning, it can completes a left turn in a few seconds, with the outcome being a positive PET – clearly the human driver does not predict based on constant speed! Moreover, even when the SV approach speed to the intersection is higher, the prediction of PET is only useful when a more realistic prediction of the vehicle speed profile is available – for faster approach we expect to see braking to a lower speed before turning and then acceleration away from the intersection. Therefore, it is essential to have a more realistic method of estimating speed and hence predicting the PET, rather than using the constant speed assumption. We approach this via a pre-determined *reference profile* of the SV motion to resolve the problem and implement this via a predicted time to conflict point (T2C) parameter – to predict the time for the SV to locations (a) or (b) in Figure 4.2. T2C for the POV is also needed, but there the constant speed assumption will be kept.

A reference motion profile for the SV may be based on time, speed or acceleration variables associated with an expected SV path. A speed profile has been used in this report, since it turns out to be easier to characterize the typical vehicle. The approach is to characterize typical free-flow speed during a left turn maneuver from measurements of multiple such events – these may of course be influenced by intersection geometry and driver preference, but common elements exist and a mathematical function with adjustable parameters can be applied to represent the relationship between the path position and speed. Detailed explanation of the reference speed profile will be given in the Section 5.

4.1.5 Leading and Trailing Buffers

Once a realistic speed profile has been found for predicting turning movements, the corresponding analogs of GT can be found. These are termed the *leading buffer* (LB) and *trailing buffer* (TB) – again predicting the spare time between the SV and a POV at a conflict point [9 (p.147)]. The LB is defined for cases that the SV enters an encroachment zone after the POV and the TB is for the case where the SV turns first. Therefore, the LB and TB respectively predict negative and positive PET parameters as defined above. Although the aims of TB/LB and GT are the same - to predict the PET, the difference is that the LB/TB calculations use more realistic information about likely future SV motion obtained from observations. Again, detailed analysis is presented in the next chapter.

4.2 GT and PET results

4.2.1 Single Left Turn Event

Results of the GT and PET (for the same left turn example described in Chapter 3) based on the trajectory reconstruction are shown in Figure 4.3. Both axes are in units of seconds; the horizontal axis represents the independent variable: elapse time during the conflict (starting at the somewhat arbitrary event start time of the event data). The vertical axis represents three dependent variables: the solid line is the predicted time to the SV exiting the encroachment zone, the sloping dashed line is the predicted time for the POV to enter the encroachment zone, and gap time is the difference between the two, positive if the turning SV is anticipated to exit first (collision avoided). These predictors presume a simple fixed trajectory path and a constant speed from the instant in question, so as SV speed varies, the SV line deviates from a unit negative slope. The POV is modeled with constant speed when it is outside of the detection zone of the forward radar, so the unit slope here is as expected (predicted time to arrival decays uniformly). These "predictors" remain well defined after the two vehicles reach these reference locations, so we naturally include the negative values as part of the conflict reconstruction data. The actual time between SV departure and POV arrival is the PET and is shown as the time gap on the horizontal (elapse time) axis, corresponding to where the two curves go through zero.

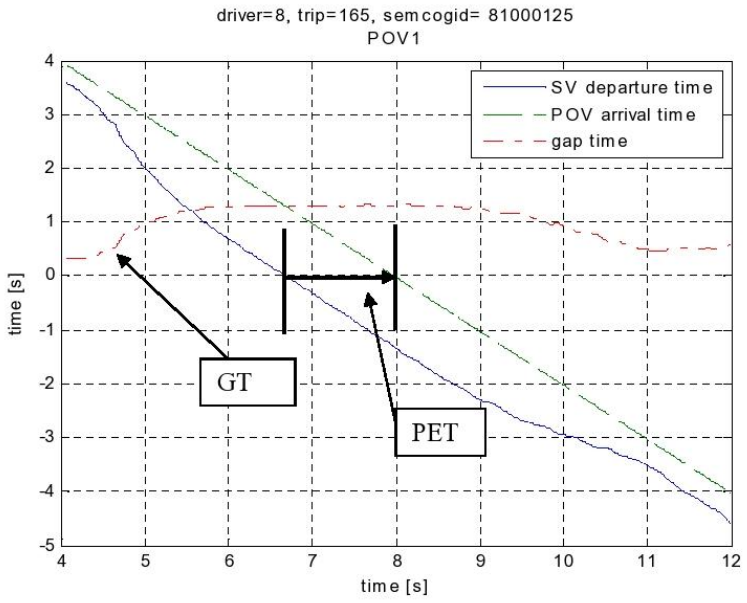

Figure 4.3. Examples of GT and PET.

At t=4s, the (predicted) gap time is quite small – around 0.3s. Then, as the SV accelerates to make the turn, the GT increases by about 1.25s and this value is maintained until the POV reached the conflict point. This demonstrates the limited actual predictive performance of GT during an actual turning event.

4.2.2 Distributions of GT and PET for All Detected POVs

In all the valid left turn events at unprotected signalized intersections in the database, there were 162 POVs detected in the automated way. Figure 4.4 shows the resulting distributions of GT_{min} (minimum gap time predicted during the 2 seconds before the conflict point) GT_{max} (maximum gap time over the same interval) and PET of those events with the values within [-10, 10]s (158 POVs).

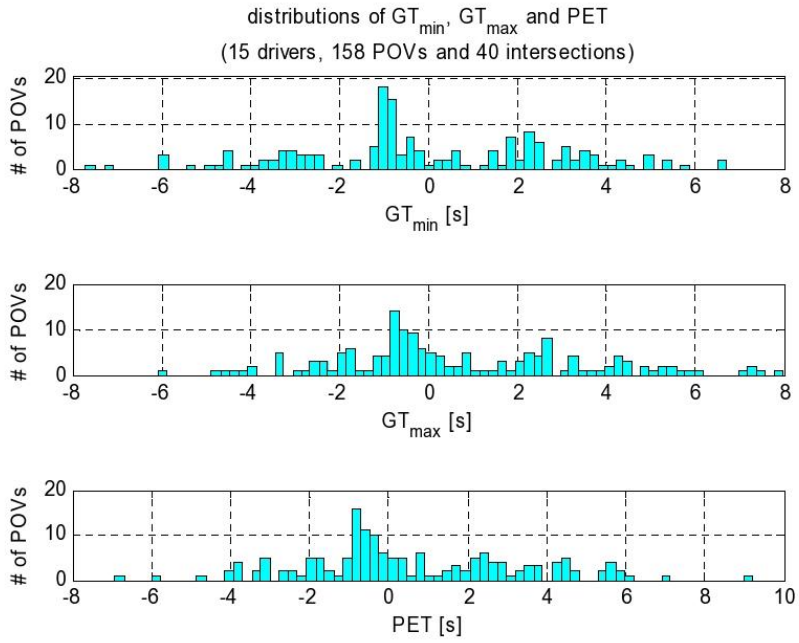

Figure 4.4. Distributions of GT_{min}, GT_{max} and PET for all detected POV.

It can be noticed that several POVs were involved within a range close to 0s, but it turns out that many of them are erroneous. Therefore, every event was checked by video observation and removed if an obvious error was found. Situations which caused erroneous detections of POVs were:

1) Turning POV onto the same road as the SV or onto a driveway near the intersection and
2) Erroneous radar reflection - no such POV was confirmed in the video.

Figure 4.5 shows the distributions that are a modified version of Figure 4.4 after removing erroneous POVs irrelevant to the LTAP/OD scenario. Cumulative percentages for the right hand side (positive) and left hand side (negative) of the PET distribution are shown in Figure 4.6. It is noticeable that there are only a few cases in which SVs reached conflict points with the PET of 2s before POVs did. This observation agrees with the result by Chan [4]. Note that in Figure 4.6(b) there are a few cases where a positive PET is predicted, and for the SV turning behind the POV this would indicate a collision. This highlights a sensitivity problem with the data analysis, particularly for the second case. The timing of the POV departing the encroachment zone is predicted from its radar track assuming constant speed - if the POV accelerates at all the actual timing will drift away from the prediction during this "blind" period – sufficient in some cases to turn a small negative PET into a positive estimate. In the same case, we are also required to estimate the length of the POV, and this is another potential source of error.

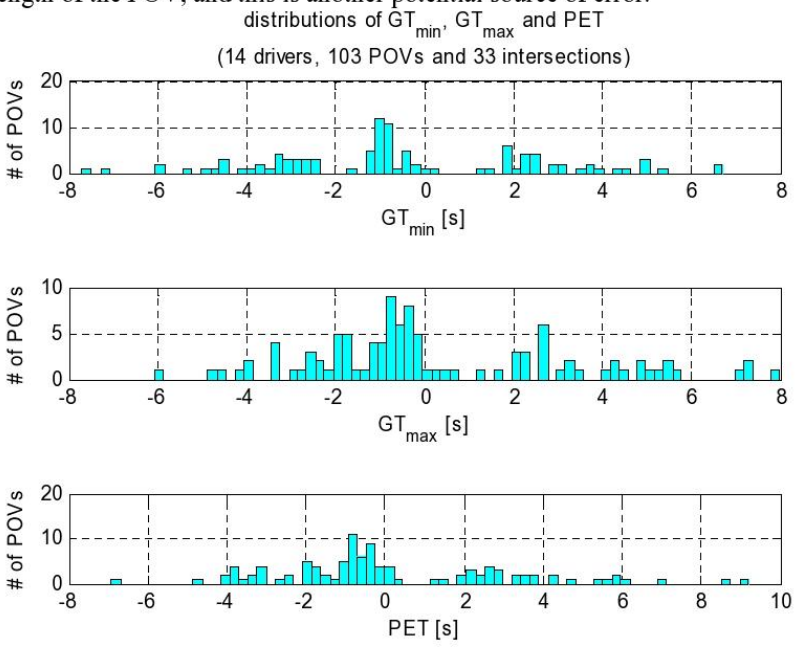

Figure 4.5. Distributions of GT_{min}, GT_{max} and PET with thru-POVs confirmed by video observation.

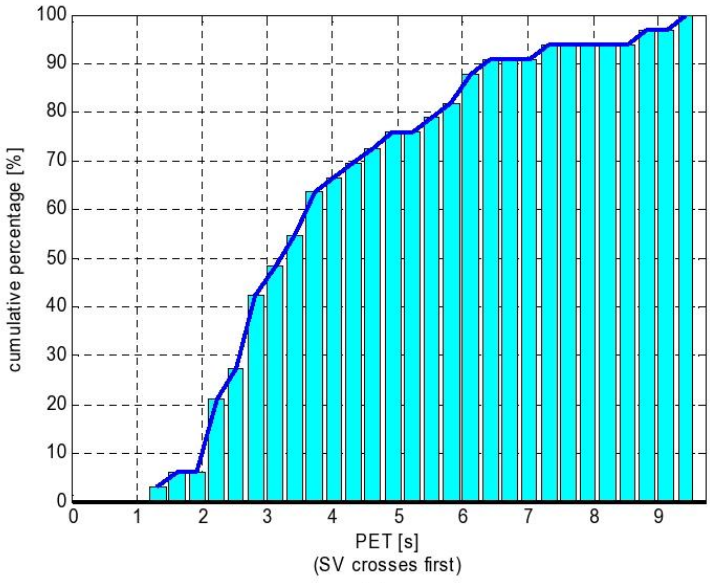

(a)

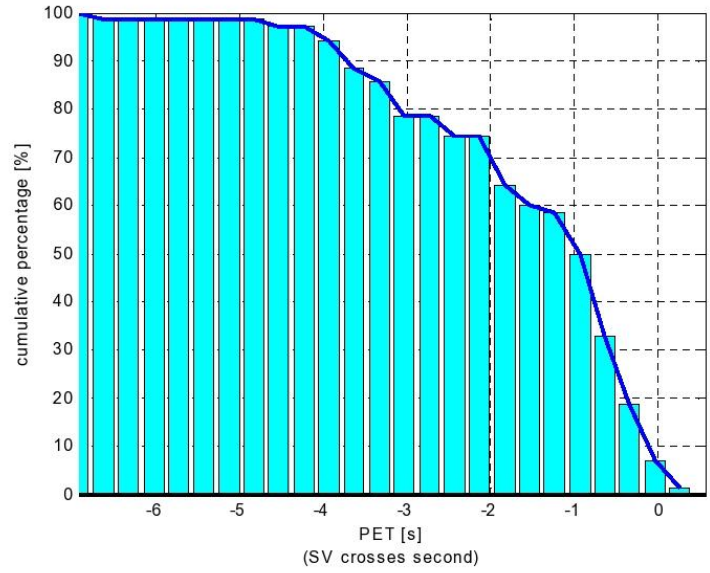

(b)

Figure 4.6. Cumulative percentage. (a) SV crosses first and (b) SV crosses second.

Table 4.1 shows a summary of the POVs detected in all left turns found in the RDCW database. According to video observations, not all POVs were genuine cases of a thru vehicle; all 162 cases were reviewed and the just over 30% were rejected for the reasons stated in the table. For the case where the POV stopped, it appears that the signal phase for the opposite traffic was no longer green when the SV turned, but, this cannot be verified, since the signal phase status is unavailable in the RDCW data.

Table 4.1. Contents of the detected POVs.

	# of POVs	Percentage
Successful detections	106	65.43 [%]
No actual POV	18	11.11 [%]
POV turning	36	22.22 [%]
POV stopped	2	1.23 [%]
Total number of detected POVs	162	100 [%]

4.2.3 PET by Age and Gender

Age and gender are considered to affect risk-taking behavior in vehicle driving. In RDCW, ages of test drivers were classified into three groups: younger (20-30yr), middle aged (40-50) and older (60-70).

The following figures show cumulative histograms of PET.

In Figure 4.7, cumulative histograms were made based on subsets for gender and a turning order (i.e. the SV crossed first or second) but including all age groups. In this figure, male drivers appear to be more "aggressive" than female drivers as observed in their minimum PETs for the case where the SV crossed first: about 1s and 2s, respectively. On the other hand, for the case where the SV crossed second, no significant difference was found between male and female drivers. Note that the latter case has little impact in terms of crash risk, apart from the case where there are multiple POV's present and the driver must find a gap related to LB for the leading vehicle and TB for the subsequent vehicle.

In terms of the age factor, although the younger group (Figure 4.8) shows a similar tendency to that of the combined age groups (Figure 4.7), the middle-aged and older groups had only small numbers of detected POVs in LTAP/OD events. While the data presented is clearly limited, the methods exist to extract such data in future studies.

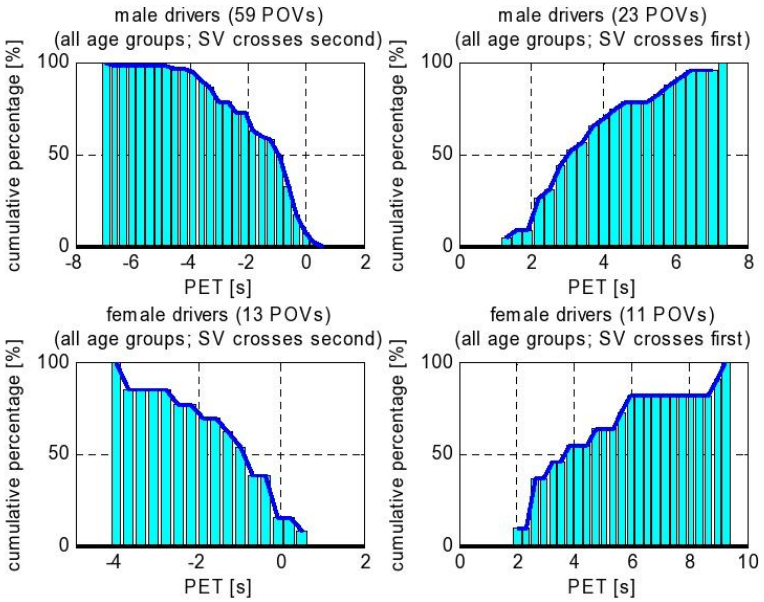

Figure 4.7. All age groups.

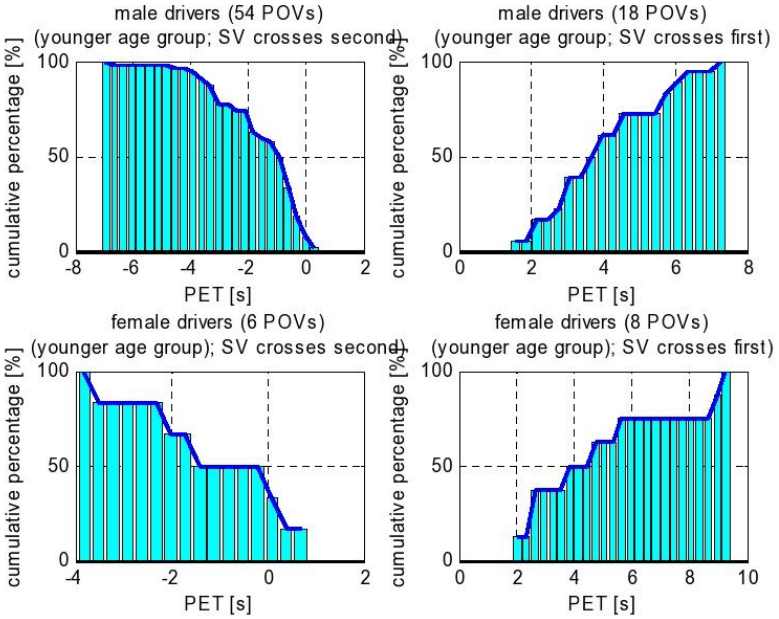

Figure 4.8. Younger drivers.

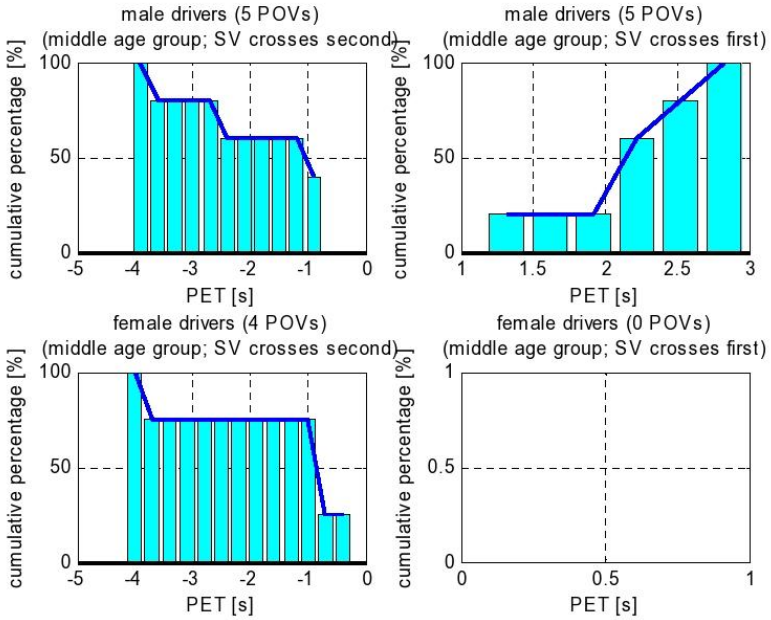

Figure 4.9. Middle-aged drivers.

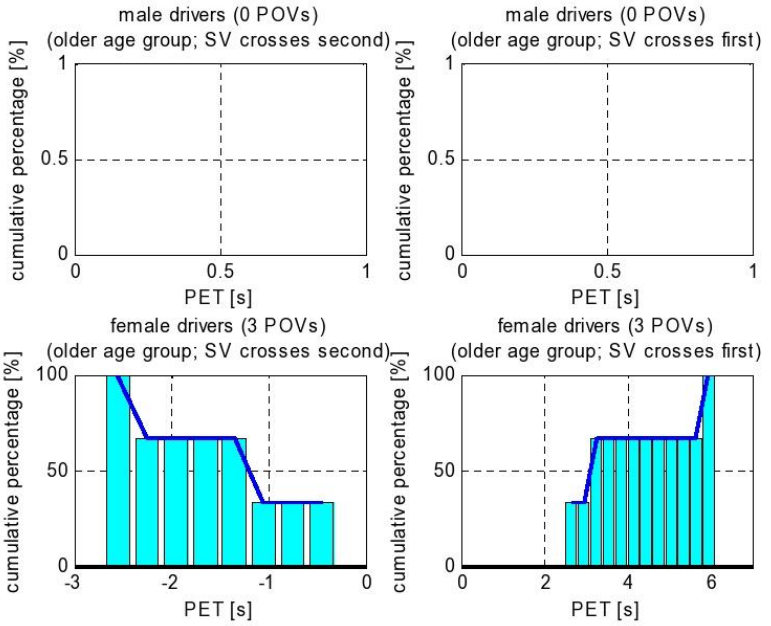

Figure 4.10. Older drivers.

4.3 Conflict Metrics for Left Turn Assistance

The conflict metrics might be used as key information to determine whether the CICAS/SLTA system should give a turning vehicle a warning or not. In this case, the conflict metrics must be able to predict a spare time at the conflict situation sufficiently accurately. Factors required would include kinematic variables of turning and thru vehicles, namely, position (distance from the intersection), speed and acceleration before the SV enters the intersection as well as other crucial information (e.g. signal state, other turning vehicles, intersection geometry, presence of pedestrians). Driver intention will not be known, but a warning is specific to a *potential* decision to turn, so it is plausible to assume that the warning system would continuously update its advice condition on the assumption that the driver is "just about to turn". The next section includes a deeper analysis of typical speed control through the turn, and uses this as a basis for computing real-world distributions of leading and trailing buffers.

5. Analysis of Leading Buffer (LB) and Trailing Buffer (TB)

5.1 Outline of LB/TB Computation

The computation of the LB/TB takes the following procedure:

1) Breakdown of left turns by their turning directions for each intersection,
2) Construction of a reference speed profile for each left turn direction,
3) Prediction of speed profile using the reference speed profile and
4) Computations of the LB/TB.

The left turn direction is defined by roads on which the SV was traveling before and after making a left turn. It is used to provide proper geometric information for the construction of a *reference speed profile* – a generic predictor of speed for the forward motion of the subject vehicle, used in place of the simplistic *constant speed* assumption traditionally considered, and applied in Section 4 above. Once a reference speed profile is defined, then this is to be used along with the SV speed observed at any instant to make a prediction of future speed. This in turn determines the predicted time for the SV to reach the relevant conflict point (departure point for case (a) when the SV turns first, arrival point for (b) when the SV turns second) as shown in Figure 4.2 This "time to conflict point" for the subject vehicle is denoted $T2C_{SV}$. Similarly, the time for the POV to reach the relevant conflict point is denoted $T2C_{POV}$ and corresponds to the predicted time until the POV enters the encroachment zone in (a), or exits in case (b) in Figure 4.2. As with the gap time (GT), the differences between these conflict times predicts the leading and trailing buffer magnitude:

(a) $|TB| = T2C_{POV} - T2C_{SV}$

(b) $|LB| = T2C_{SV} - T2C_{POV}$

As with gap time, TB and LB are continuous functions of time during the conflict, based on future prediction of the conflict event; the key difference here is that a less simplistic speed predictor is used for the SV in the "time gap" calculation. See Appendix A for further notation.

5.2 Construction of Reference Speed Profile

5.2.1 Breakdown of Left Turns by Left Turn Directions

Before constructing the reference speed profile, left turns are broken down by their directions of turns. For example, for a 4-leg intersection, there are 4 possible turn directions (Figure 5.1) and a unique reference speed profile is provided to each of them. A left turn direction is identified by finding the closest stop bar node to the initial SV position. The stop bar nodes are available in geographical coordinates in the intersection database. The example result (Figure 5.2**)** shows that there are two left turn directions

existing in the database (Direction 1=solid and Direction 2=dashed) at this intersection (Semcog ID=81000083). Therefore, two distinct reference speed profiles are to be created for this intersection.

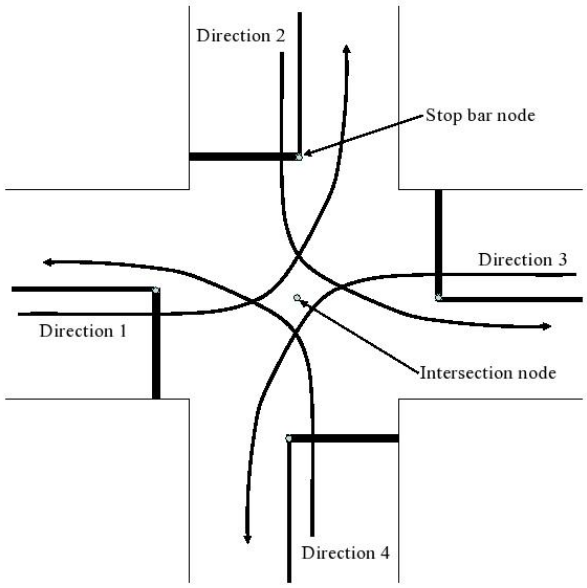

Figure 5.1. Schematic of left turn directions

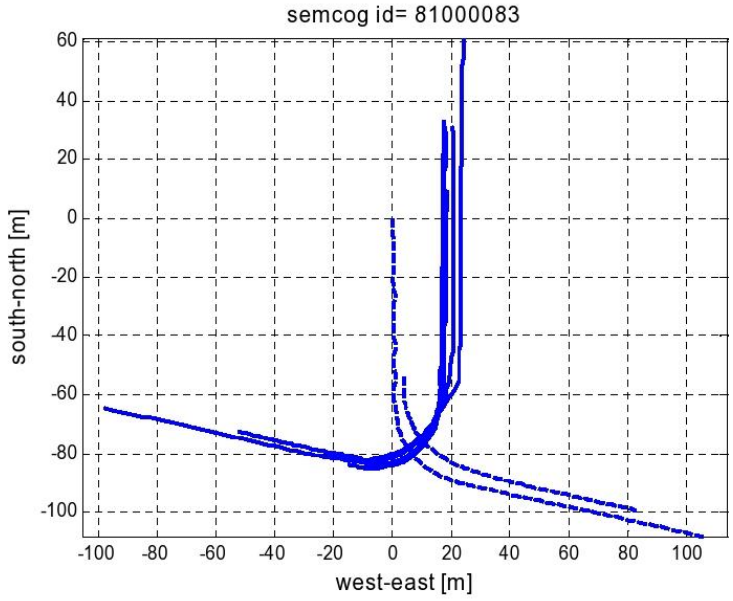

Figure 5.2. Two left turn directions in Semcog ID=81000083.

5.2.2 Function for Reference Speed Profile

In order to prepare a reference speed profile, three assumptions are made; 1) there is no significant interaction with oncoming vehicles, 2) traffic signal is in permissive (green) phase and 3) there are no pedestrians present. This is expected to create a free turn behavior in which the reference speed profile consist of three parts: 1) deceleration, 2) constant speed and 3) acceleration. In this case, a general form of the reference speed profile (speed vs. distance through the turn) will be based on three phases: entry into the turn, steady-state turn, exit from the turn. The following structure is therefore chosen

$$U_{ref} = f_1(s)\Phi(s_{c1} - s) + f_2(s)\{\Phi(s - s_{c1}) - \Phi(s - s_{c2})\} + f_3(s)\Phi(s - s_{c2}) \qquad (5.1)$$

where $f_i(s)$ is a sub-function for the i-th part of the free turn, s is the path coordinate, $\Phi(x)$ is the Heaviside step function defined as

$$\Phi(x) = \begin{cases} 0, & x < 0 \\ 1, & x \geq 0 \end{cases}, \qquad (5.2)$$

and s_{c1} and s_{c2} are path coordinates at the intersections between $f_1(s)$ and $f_3(s)$, and $f_2(s)$ and $f_3(s)$, respectively.

It should be clear that a position (or distance) based model is greatly preferred to a time-based model; in the latter case for example, when the driver approaches slower than standard, the model would predict acceleration to take place even though the turn had not been completed. Also, as we described below in Section 5.3.1, the reference speed is not the same as the predicted speed – this depends on current speed as well, and a transition from current speed to reference speed is to form part of the predictive model.

The function $f_1(s)$ represents a deceleration stage on approaching the turn, and can be described by

$$f_1(s) = \mathrm{Re}\left(\sqrt{2p_1(s-p_2)}\right), (p_1 < 0) \tag{5.3}$$

with an assumption of constant deceleration. Likewise, assuming a constant acceleration, the exit phase is described as

$$f_3(s) = \mathrm{Re}\left(\sqrt{2q_1(s-q_2)}\right), (q_1 > 0). \tag{5.4}$$

Note that p_1 and q_1 represent acceleration magnitudes, while p_2 and q_2 indicate the distance coordinates at which minimum values of the functions are given. The second part is simply written by

$$f_2(s) = U_{th} \tag{5.5}$$

where U_{th} is a constant speed. The above equations simply represent a set of idealized functions designed to represent typical speed control for real intersection turning movements, so that predictions can be made when the actual data is not available (as in computing TB estimates in real-time).

U_{ref} was found by fitting the above function to a representative curve of actual left turns. The representative curve was selected by averaging all left turns with respect to s. Then, in order to remove sharp corners at the connecting points of the functions, a time-based moving average was computed.

Figure 5.3 shows an example reference speed profile. In the bottom figure, the solid curve is the representative curve and the dashed curve is the result of fitting the aforementioned function.

Note that in designing a warning system, more or less conservative estimates can be made for TB and LB prediction, or even adapting these predictions based on trends specific to the driver. Here we target the average conditions as a reference that can be perturbed later at the design stage of the CICAS SLTA system.

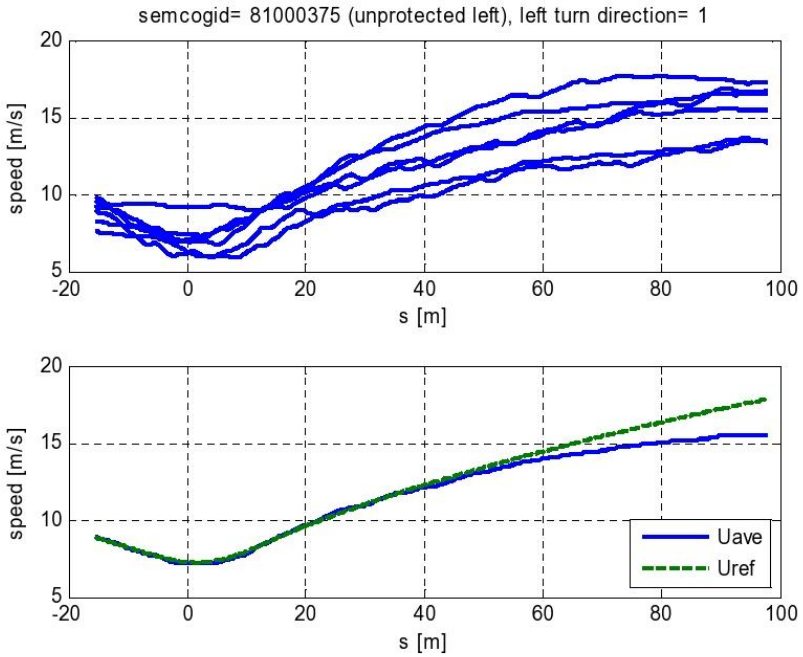

Figure 5.3. Free turns in SEMCOG ID=81000375 (top) and resulted U_{ref} fit to average free turn curve (bottom).

5.2.3 Evaluation of Reference Speed Profile

It may be reasonable to anticipate that the turning time varies depending on the intersection geometry affecting on the length and curvature of a turning path. In this section, possible relationships between parameters for the intersection geometry and those for the reference speed profile is investigated.

The size of an intersection is specified by the lengths of segments connecting the intersection node and stop bar node and the shape is defined as the angle made by those segments. In Figure 5.4, the point A is a stop bar node on a road on which the SV is initially located, the point B is that of the other road onto which the SV is turning, and the point C is the intersection node. The length d_1+d_2, a sum of the lengths of the segments AC and BC, is selected to indicate the intersection size and the angle δ is the angle between those segments to represents the shape of the intersection.

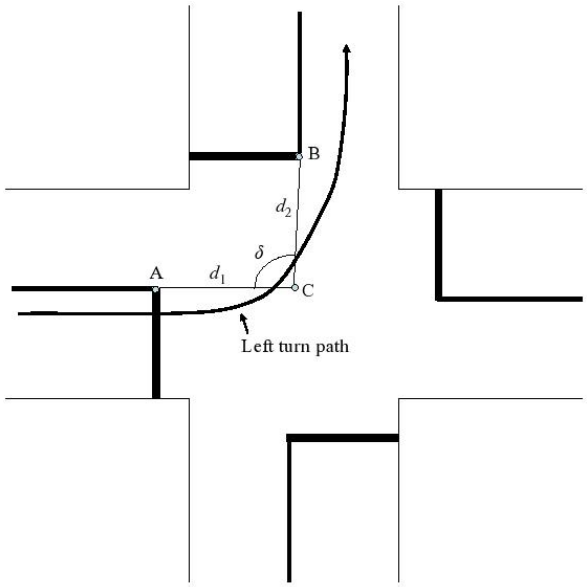

Figure 5.4. Intersection geometry.

Figure 5.5 shows relationships between the parameters for the reference speed profile, p_1, q_1, and U_{th} and those for the intersection, d_1+d_2 and δ. Note that p_1 and q_1 are acceleration values in the deceleration and acceleration phases during a left turn.

Contrary to the original expectation, neither the size nor shape of the intersection show a clear and direct relationship to the parameters in the reference speed profile. While it might be concluded that, within the range of intersections considered, the turning vehicle follows a uniform reference speed profile, it is also clear that the variations in the parameters(Figure 5.5) are large compared to the mean. These variations could be arising from sampling errors (since in many cases the samples are small), or from other factors such as road class, traffic density, presence of pedestrians, road roughness, lane marking quality etc. Most likely, both sources of variation are present, so the approach taken is to treat each intersection (and indeed each turning movement within that intersection) as a distinct case that requires the construction of a unique reference speed function. For further comments, see Section 5.5.

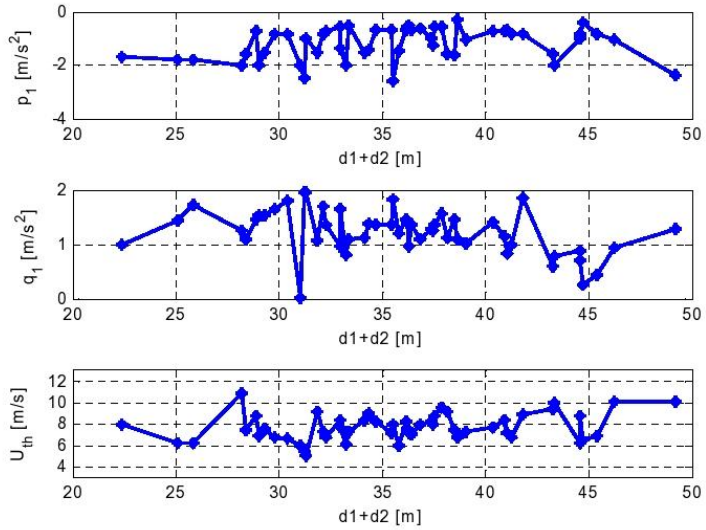

(a)

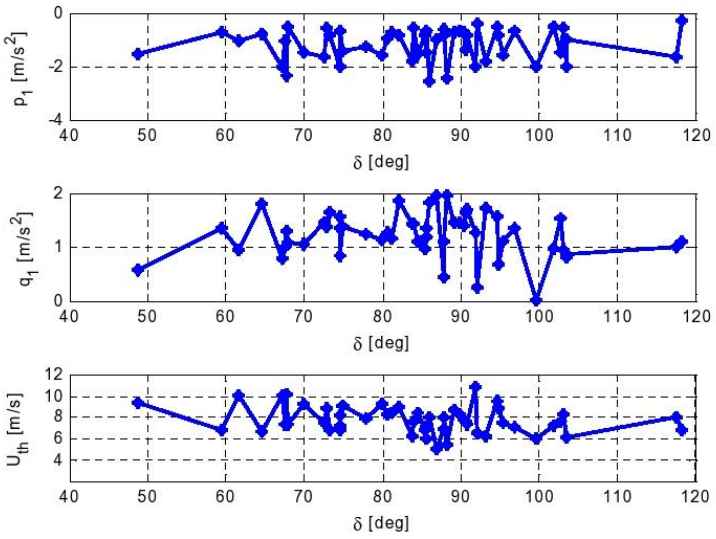

(b)

Figure 5.5. Relationships between the reference profile parameters and the intersection size (a) and shape (b).

5.3 LB/TB Calculation

5.3.1 Computation of Predicted Speed Profile from Reference Speed

In order to evaluate the LB/TB, a predicted speed profile is to be made at every SV location. Since speed prediction depends on current speed as well as typical speed, a blending function is applied to transition between the two. The following exponential function implements this model

$$U_{pred}(s;s_0) = \left(k_1 - U_{ref}(s_0)\right)e^{-k_2(s-s_0)} + U_{ref} \tag{5.6}$$

where s is the path coordinate of the SV, s_0 is the current SV position at which the LB/TB is computed and k_1 and k_2 are parameters. At position $s=s_0$ the model "predicts" current speed, provided the first parameter k_1 is set to the current observed speed (i.e. $k_1=U(s_0)$) ; for s sufficiently large, the exponential function becomes small and the prediction tends towards U_{ref}. The second parameter k_2 determines how quickly this transition takes place; for example. k_2=0.3 corresponds to a "relaxation length" $(0.3)^{-1} = 3.33m$ (as a rule of thumb, the vehicle must travel through 4 times this length to achieve 98% of the transition from $U(s_0)$ to U_{ref}) . A smaller value of k_2 corresponds to a larger relaxation length, and hence a greater distance for the SV to reach U_{ref}. Note that the prediction is to be updated at each location on the SV path.

5.3.2 LB/TB Calculation

By using the U_{pred} explained above, a predicted $T2C_{SV}$ (to arrive at the conflict point for an LB case or to leave that point for a TB case) can be computed by integrating the inverse of U_{pred} along s;

$$T2C_{SV}(t) = \int_{s_0|_{s=t}}^{s_C|_{s=t_C}} \frac{ds}{U_{pred}(s;s_0)} \tag{5.6}$$

where s_0 is the longitudinal track coordinate at t (time of LB/TB computation) and s_C and t_C are those for the SV at the conflict point. $T2C_{POV}$ can be calculated in the similar way, but the result becomes a straight line because of the POV trajectory extrapolated using a constant speed as described in Chapter 3. Therefore,

$$T2C_{POV}(t) = \frac{D(t)}{U_{POV}} \tag{5.7}$$

where $D(t)$ is the length of path between the POV and conflict point at the time of LB/TB computation and U_{POV} is the POV speed, that is constant.

The LB/TB can be calculated now by the equation (see Section 5.1)

$$LB/TB = T2C_{POV}(t) - T2C_{SV}(t). \quad (5.8)$$

where now the value is defined to become negative in an LB case. This convention is the same as that in the GT (Section 4) and is retained so we may compare the LB/TB with the GT in a consistent way.

5.4 LB/TB Results

5.4.1 Individual Left Turns

In the following examples, the LB/TB and GT are compared for their capabilities of predicting the PET. The first example shows a LB case (Figure 5.6) and the second one shows a TB case (Figure 5.7). In the figures, $T2C_{SV,GT}$ represents the time to conflict point for the SV (see Appendix for notations) obtained from the GT calculation – i.e. using a constant speed assumption for the SV, while $T2C_{SV,LB/TB}$ uses the prediction of variable speed. For the POV, $T2C_{POV}$ assumes constant speed for its prediction in all cases. In each figure, a triangle was placed to indicate the time at which the SV reaches its conflict point. Clearly, the GT curve and LB/TB curve meet each other at this triangular mark..

In the first example, the SV was initially moving slowly and then started accelerating at about t=1s (Figure 5.6(a)). At the beginning, since the LB was positive (Figure 5.6(b)), it was anticipated that the actual conflict was a TB case. It seems that there was not enough spare time for the SV (less than 1.2s) to make a left turn in front of the POV and thus it waited until the POV had passed the encroachment zone first. As a result, there appeared a transition point in the LB curve at around t=1.4s from positive to negative. This transition can also be noticed in the crossing point between the $T2C_{SV,LB}$ and $T2C_{POV}$ in Figure 5.6(c). This assessment of the turning decision could be made because of the employment of U_{pred} (dotted curve in Figure 5.6(a)). In contrast, the GT did not give early information about the PET, since no realistic prediction about the future speed was used.

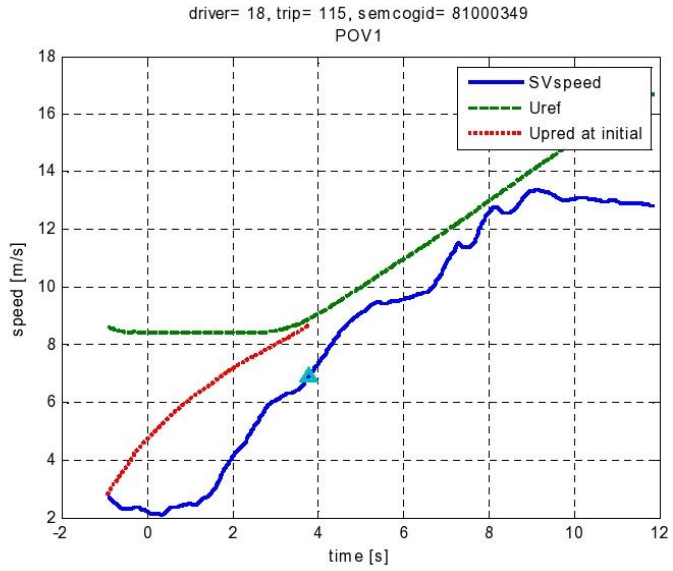

Figure 5.6 (a). Example 1 (LB case)

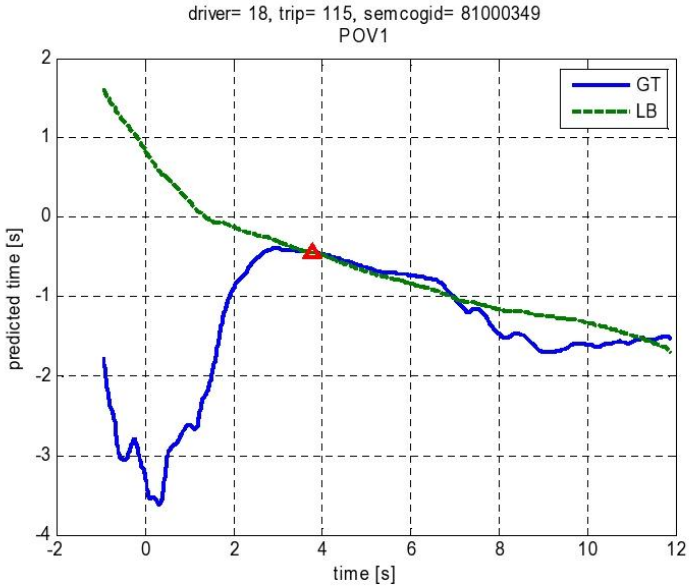

Figure 5.6 (b). Example 1 (LB case)

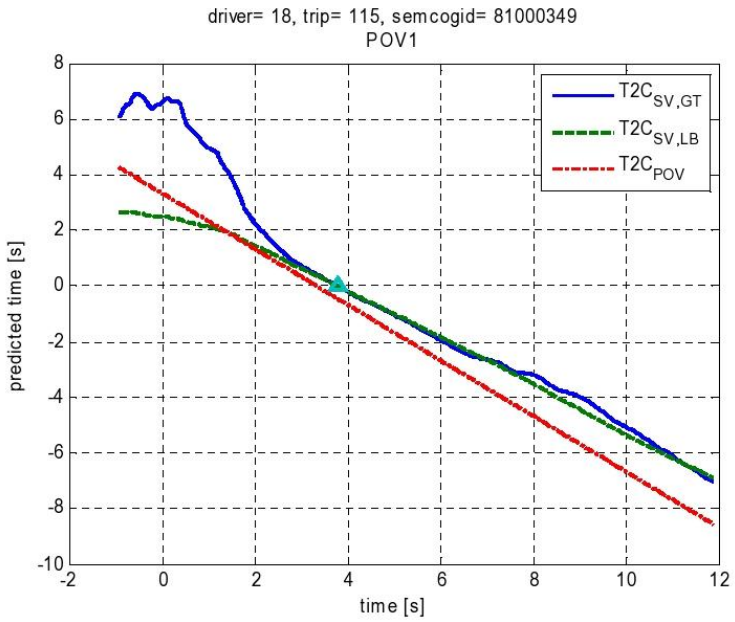

Figure 5.6 (c). Example 1 (LB case).

In the second example, a TB case is shown (Figure 5.7). Again, the GT in the early period deviated significantly from the PET value at a triangle but, the TB provided a fairly accurate prediction. According to the TB curve, we could expect that the SV would turn before the POV, since the TB value was larger than 2s which might be the threshold value for a decision making of a left turn in TB cases as confirmed in section 4.2.2 and in [4].

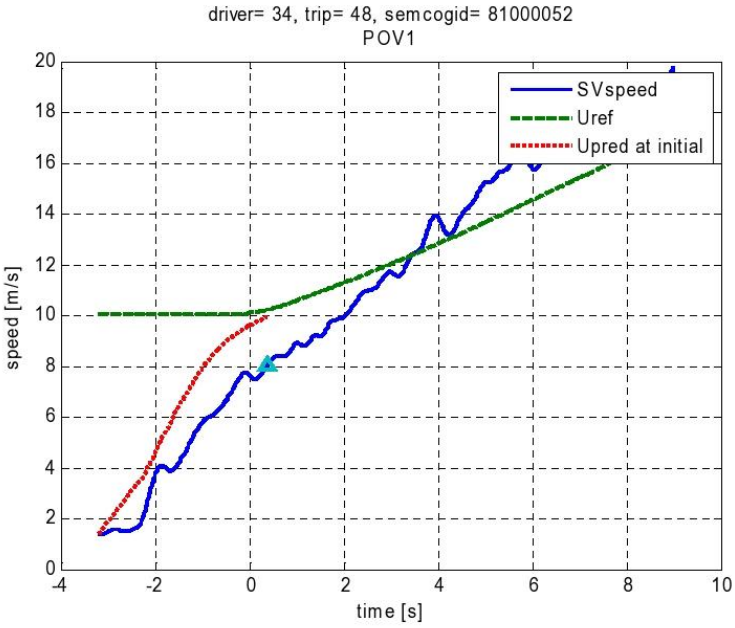

Figure 5.7 (a). Example 2 (TB case)

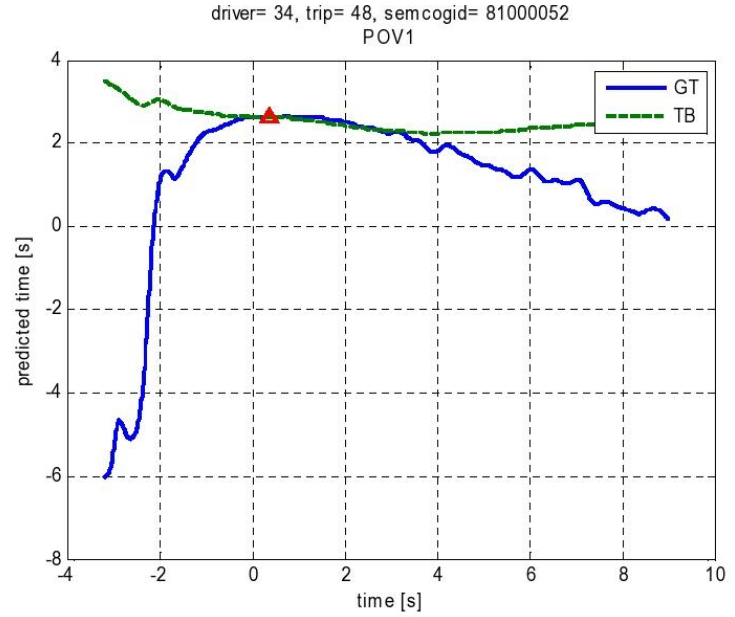

Figure 5.7 (b). Example 2 (TB case)

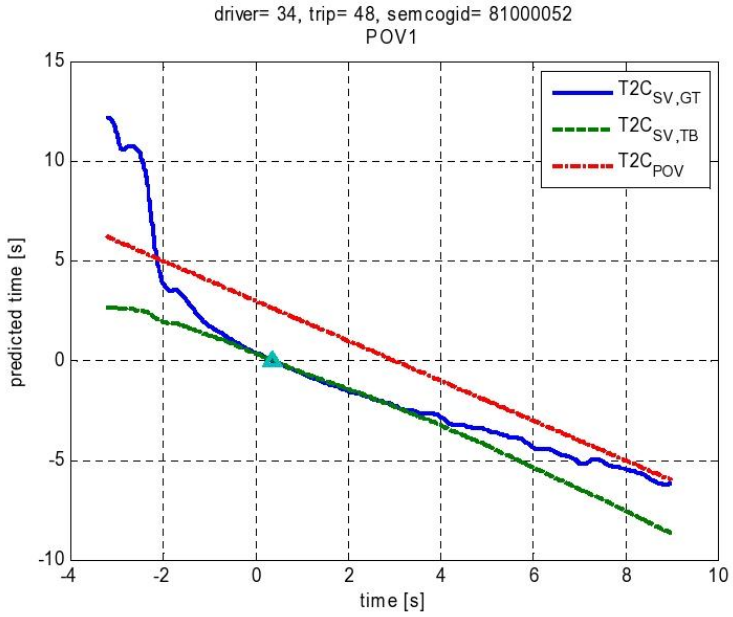

Figure 5.7 (c). Example 2 (TB case)

5.4.2 Distribution of LB/TB for Multiple Left Turns

Figure 5.8 shows distributions of errors in the LB/TB predictors at the start of a number of left turn events. The error is defined as LB/TB minus PET (upper plots); for comparison, the GT error (GT minus PET) are shown in the lower plots. The events are all from one particular driver (driver #18, who presented the highest number of relevant turning events). The LB cases provided 19 events (Figure 5.8(a)) and the TB cases 7 (Figure 5.8(b)). Events where the vehicle was initially stationary (and hence GT predictions are meaningless) were excluded.

According to the definitions of leading buffer in equation (5.8), a positive error in the LB indicates that the time separation predicted at the start of the maneuver is shorter than the actual one. In Figure 5.8(a), the whole distribution is located on the positive side in the horizontal axis. In other words, the LB tended to underestimate the spare time before the SV actually started turning. This is completely natural, since the prediction at any instant assumes the driver will act as though the turn was free of oncoming traffic. An SLTA using this measure would advise against turning, just as in Figure 5.6(a) the driver actually "advises himself" against turning – in this figure the driver holds the SV at a very low velocity to allow the POV to pass in front before initiating the turn. Thus we need to be careful in the interpretation of "errors" here – in all 19 cases the LB calculation would indicate that the driver should delay the normal turn profile, and in all cases he did. GT produces both positive and negative errors, indicating that constant

speed is not only a poor predictor of actual turning behavior, it also offers no clear basis for activating an advisory system

For the trailing buffer case – Figure 5.8(b) – a positive error indicates that the predicted time separation for the conflict is greater than actually occurred. In these cases it is the GT that is one-sided – a constant speed assumption almost always predicts a worse time separation than actually occurs, and again no useful predictions emerge. The TB gives a much smaller variation and seems to offer a reasonable metric for advising on driver action before the turn is made.

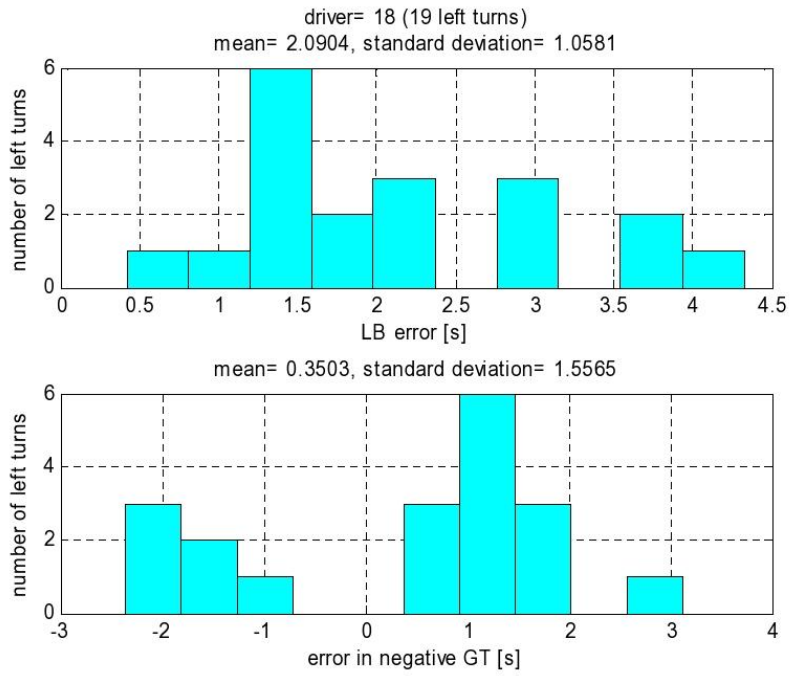

Figure 5.8. (a) Distributions of LB and GT errors for Driver=18.

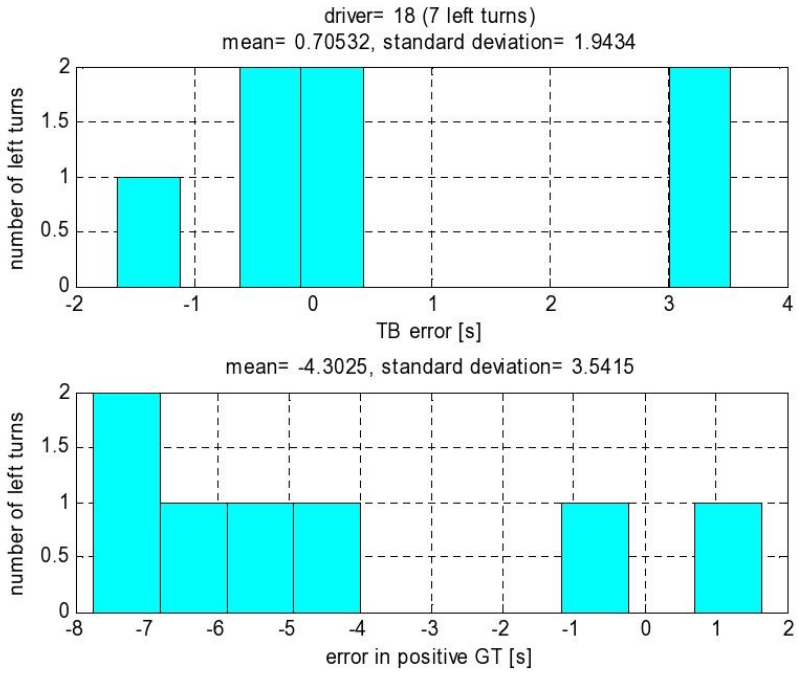

Figure 5.8. (b) Distributions of TB and GT errors for Driver=18.

5.5 Variability in Estimated Reference Speed Profiles

The accuracy of LB/TB results is directly related to the performance and relevance of U_{ref}. The results in Section 5.2.3 show no clear relationships between the obtained U_{ref} and intersection geometry, even if they would intuitively have some relationship. Of course, other factors, e.g. visibility of the intersection influenced by surrounding buildings, degree of flatness of the road surface and surface roughness due to wheel ruts, might change U_{ref} parameters. Besides all the related geometric factors, a more significant reason would be the relatively large variance in multiple turns. The following figure shows actual speed profiles of 52 left turns made by a single driver at a single intersection. At s=10m, the spread in the ordinate axis is about 10m/s. Thus, the reference speed cannot be used as a faithful predictor of actual speed, rather a reference that guides conditional warning states ("if the driver accelerates normally, the PET should provide a safe margin …"). It certainly provides significant improvement over simple GT estimates, based on a constant speed assumption.

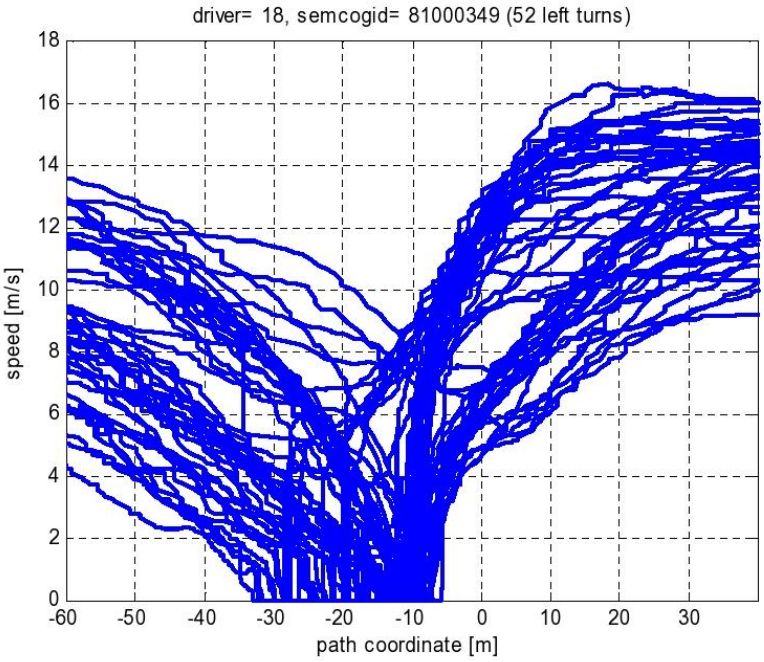

Figure 5.9. Actual speed profiles of multiple left turns.

6. Specific Turning Events for Gap Acceptance

6.1 Buffer Bands and Band Gaps for Multiple POV Case

If there is only one relevant POV, the significant conflict metrics will be the TB. This is because it is a relatively simple task for the SV to cross the encroachment zone after the POV without causing a collision. In this case, the LB might be used only to determine when a warning by the left turn assistance should be turned off. On the other hand, if there are multiple POVs, it needs to be evaluated as well as the TB because it may be crucial to establishing suitable gaps for a safe left turn between POVs of a particular pair. In this section, a concept of the buffer band and corresponding band gap are introduced for this purpose. They would be helpful to visualize the gaps between POVs in order to find an acceptable gap for a left turn.

An example buffer band is shown as an envelope made by the LB curve and TB curve for a single POV case (Figure 6.1). In the figure, the LB was 1.00s and the TB was -1.44s at t=2s. Since the LB must be negative if the SV crosses second and the TB must be positive if it crosses first, the SV could not initiate a turn at this instant. In equation (5.8) the signs of the LB and TB are determined in the same way by the difference

$$T2C_{POV}(t) - T2C_{SV}(t)$$

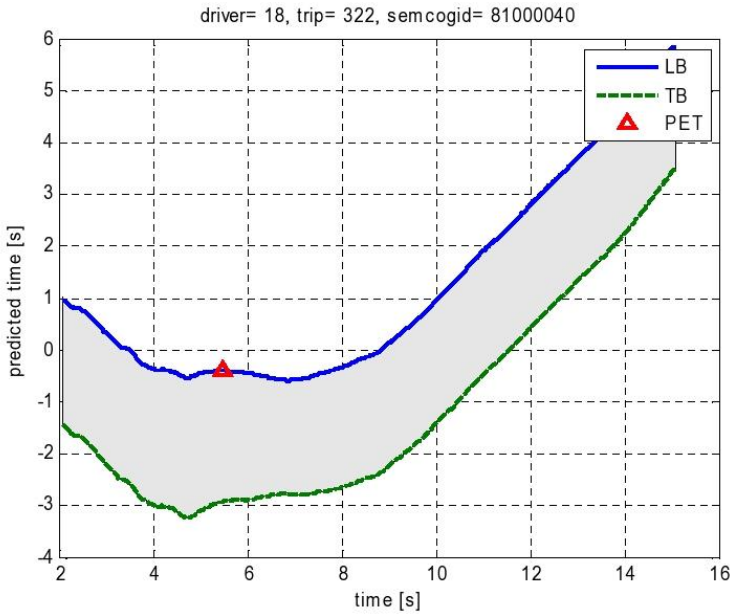

Figure 6.1. Buffer band in LB/TB.

It is clear that the TB is always smaller than the LB, i.e. 0<TB<LB for the case that the SV crosses first and TB<LB<0 for the case that the SV crosses second. A left turn is achievable only when both have the same signs. In other words, when the SV turns left, the buffer band must exist only in the positive range (SV crosses first) or negative range (SV crosses second). In the example case, the SV waited until the LB curve became sufficiently large and negative to provide an enough time separation. At completion of the maneuver, the whole band was directed to the negative side with PET= -0.40s at t=5.50s.

For a multiple POV case, if there is an enough gap between two POVs with which a safe left turn can be expected, the conditions for this situation are,

$$\begin{cases} \mathrm{LB}_i \le \mathrm{Th}_{\mathrm{LB}} \\ \mathrm{TB}_{i+1} \ge \mathrm{Th}_{\mathrm{TB}} \end{cases}, \ (1 \le i \le n-1). \tag{6.1}$$

where n(>2) is the number of approaching POVs, LB_i is the LB of the i-th POV, TB_{i+1} is the TB of the (i+1)-th POV which follows i-th POV, $\mathrm{Th}_{\mathrm{LB}}$(<0) is that for maximum acceptable LB and $\mathrm{Th}_{\mathrm{TB}}$(>0) is a threshold for minimum acceptable TB. Since the LB and TB for those POVs holds a relationship, $\mathrm{TB}_i<\mathrm{LB}_i<0<\mathrm{TB}_{i+1}<\mathrm{LB}_{i+1}$, the band gap between the buffer bands of the two POVs lies across the horizontal line at the ordinate of 0s on the predicted-time vs. time plot. The BG value is now defined as

$$\mathrm{BG}_i = \mathrm{TB}_{i+1} - \mathrm{LB}_i, \ (1 \le i \le n-1) \tag{6.2}$$

where BG_i is the band gap between the buffer bands of the i-th POV and the (i+1)-th POV. The BG_i must satisfy the following inequality,

$$\mathrm{BG}_i \ge \mathrm{Th}_{\mathrm{TB}} - \mathrm{Th}_{\mathrm{LB}} \quad (= |\mathrm{Th}_{\mathrm{TB}}| + |\mathrm{Th}_{\mathrm{LB}}| = \Delta\mathrm{Th}). \tag{6.3}$$

The terms in parentheses makes clear that the band gap is the sum of two positive terms, one for an adequate trailing buffer, and one for an adequate leading buffer. The band gap depends on the time spacing between POV's and the anticipated speed of the SV as it makes the turn. Condition (6.3) is a necessary requirement but it is not sufficient – the SV driver must also time the turn correctly to make use of the available band gap. The pair of conditions (6.1) are both necessary and sufficient, provided the SV driver follows the predicted speed profile within some tolerance.

6.2 Example Events

In the following, two examples with multiple POVs for each are considered. In the examples we assume thresholds $\mathrm{Th}_{\mathrm{LB}}$=-1s, $\mathrm{Th}_{\mathrm{TB}}$=2s and hence ΔTh=3s.

In the first example, there were two POVs and the SV turned between these POVs (Figure 6.2); the PET of the POV1 was -3.9s and that of the POV2 was 1.32s. At $t = -1.17s$, at which the initial detection of the POV2 was achieved, LB_1=-3.67s and

TB_2=1.34s and BG_1=5.01s. Thus, BG_1 was larger than ΔTh and LB_1>Th_{LB}, but TB_2 did not satisfy the second condition in expression (6.1). Consequently, although the SV actually turned between these POVs, it would have been advisable to have made the turn earlier, if possible; given the "late start" of the executed turn, and the parameters assumed above, the SLTA would naturally have advised against making this particular left turn.

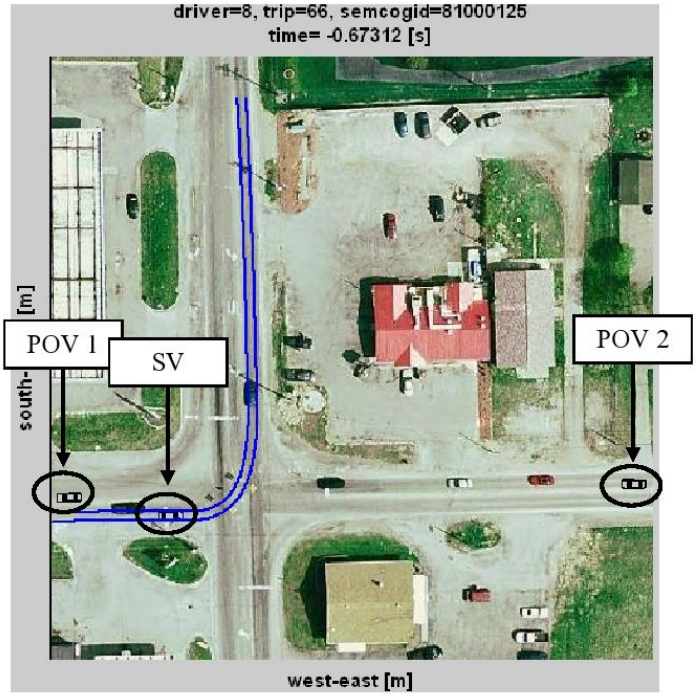

Figure 6.2. Vehicle configuration (Example 1).

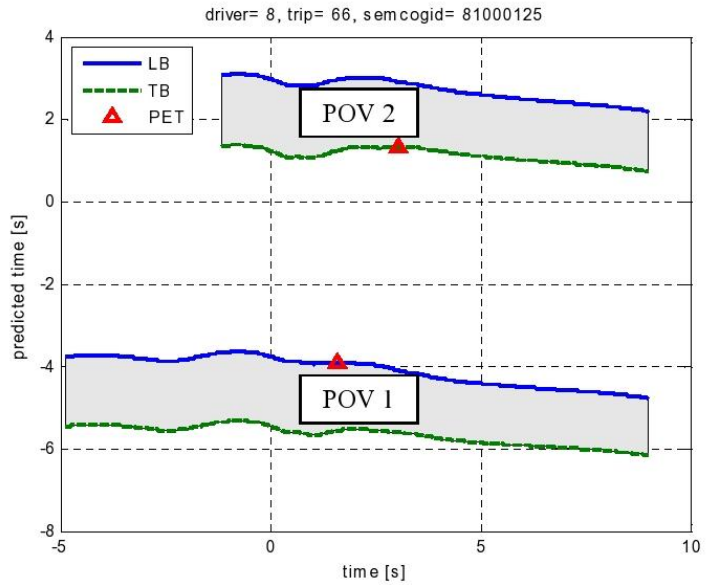

Figure 6.3. Buffer bands for two POVs (Example 1).

In the second example, two POVs were again detected and both crossed the encroachment zone earlier than the SV (Figure6.4). PET values of POV1 and POV2 were -1.91s and -0.32s, respectively. In Figure 6.5, a noticeable difference from the last example can be seen in the overlapped region. This indicates the SV could not turn between these POVs and BG_1 was negative (-0.55s). As a result, the SV crossed the encroachment zone after both the POVs did. At the initial time (t=1.01s), the union set of the two buffer bands spread from -2.57s to 1.41s - crossing the zero time and hence making the turn infeasible (Figure 6.5). As the SV waited, the whole band shifted downward until it became completely on the negative side at t=4.06s. The SV started turning about t=4.8s as indicated by the change of the slope of the LB_2 curve and reached the conflict point at t=6.54s with PET = -0.32s as described above.

These two examples indicate a plausible analysis method for setting true gap acceptance – where the gap is to be between a pair of POV's not just a gap time relative to a single POV. In congested traffic conditions this is an important consideration, especially as the existence of a steady stream of oncoming traffic naturally leads drivers to look for gaps and potentially take risks to avoid what are perceived as excessive delays. Objective information on gap feasibility would undoubtedly help the driver, particularly if this information were available early to allow smooth execution of the turn.

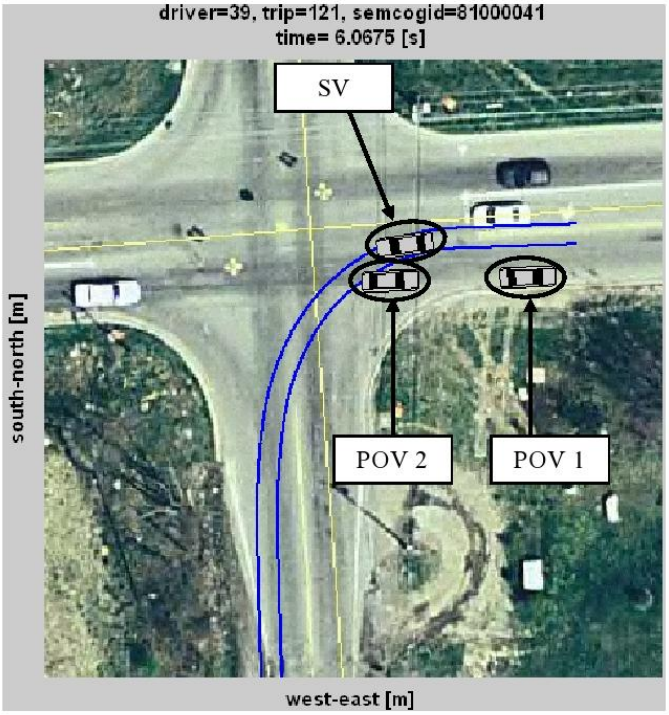

Figure 6.4. Vehicle configuration (Example 2).

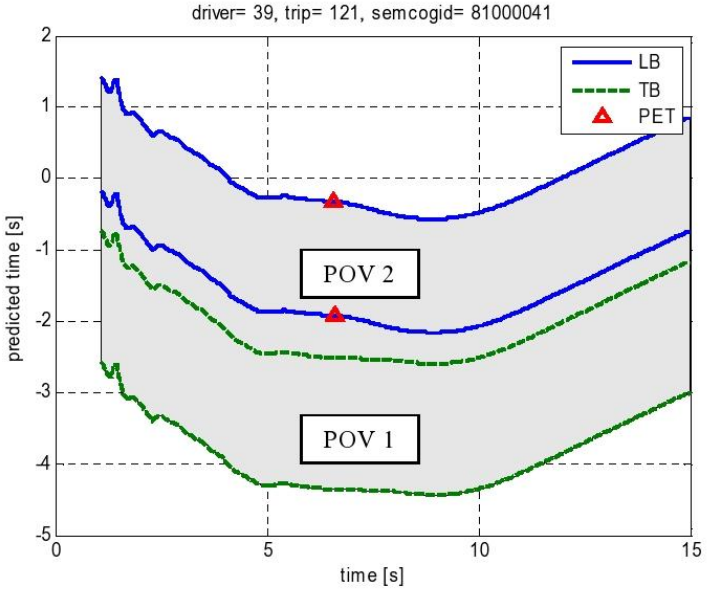

Figure 6.5. Buffer bands for two POVs (Example 2).

7. Summary and Conclusions

In this report, the following have been achieved:

1) Detection of left turns in the naturalistic driving data,
2) Reconstruction of SV and POV trajectories,
3) Computation of conventional conflict metrics such as GT and PET,
4) Brief investigation of influence of age and gender on gap acceptance, and
5) Computation of LB/TB for prediction of left turn timing.

By removing inherent noise and timing uncertainty in GPS positioning data via Kalman filtering, the reconstructed SV trajectories from the smoothed position data showed plausibly accurate results, as layered onto aerial maps. Since the POVs were detected only relative to the SV motion, and including relatively small numbers of data points, the accuracy of the resulting POV trajectories is not expected to be as high as that of the SV trajectories. Nevertheless, by correcting the vehicle speed and using an appropriate directional extrapolation, reasonable conflict metrics results were obtained with what appear to be acceptable errors.

The conflict metrics were computed using the concept of conflict point at which the separation time between the SV and POV became shortest for non-collision cases. The PET result implied that Michigan drivers had a similar 2 s gap to California drivers as described in [4] in the case of the SV crossing first.

With respect to the effect of gender, male drivers showed more "aggression" – accepted shorter buffer times – than female drivers in the case where the SV crossed first, a case which suggests increased crash risk compared with the alternate case. On the other hand, no clear conclusion was drawn on the age influence in LTAP/OD, since the number of examples from middle-aged and older age groups was very low.

The leading and trailing buffers (LB/TB) were computed by introducing a *reference speed profile*, obtained by averaging the speed curves of free turning movements in the measured data; the speed profile is related to location (distance) in the turn, rather than elapsed time. It enabled improved predictions of available time (PET) even if the SV was initially moving at a slow speed, a case when the GT would generate a meaningless large negative value in its prediction. Moreover, both the LB and TB could be computed for each vehicle at each time for gap acceptance application, in the case where there were multiple approaching POVs. Equation (6.1) appears to be a satisfactory criterion for defining acceptable gap, though at this stage only limited conflict situations have been analyzed.

The above analysis provides new reference data for gap acceptance in permissive left turns at signalized intersections. In the future, additional data would be useful and could be collected using more dedicated on-vehicle instrumentation (e.g. with wider field of

view radar on the SV), using a driving simulator, or even from site-based trajectory estimation.

7.1 Issues and Next Steps

While substantial progress has been made on this research and development topic, there are a number of immediate extensions and developments motivated by the results presented above. In terms of the original proposed Statement of Work, an accompanied driver study with enhanced vehicle instrumentation could provide information on additional factors, such as signal phase, during the turn. However, there are several other options available, and Figure 7.1 is put forward as a high level reference for these options, mapping potential research needs to the available methods. The research needs emerge as current limitations of the study presented here:

- Driver factors were only partially addressed due to sparseness of the data once it is broken down by driver types; and driver types could be expanded to include evidence of aggressive or assertive driving. This would help to resolve questions such as: do algorithms need to be tuned to individual driver performance, and in that case which indicator variables should be used to adapt algorithm parameters?
- Signal phase was not available in the source data, and clearly the effect of a signal change from permissive green to red greatly affects the decision process for the driver who is already partially committed to making the turn. Other transient factors such as lighting and presence of pedestrians in the cross-walk were also unavailable, so expanding data in this area would be desirable
- The current study has suggested an objective basis for providing go/no-go advisories to the driver, but it is not clear how appropriate or useful these would be in practice. Assuming that a full-blown field operational test is not feasible at this time, a number of opportunities exist to conduct small-scale analysis and evaluation.
- Significant variability was found when estimating reference speeds from measured data on left turns. At least part of these variations appears due to the nature of the intersection, though not simply on its size or the angle of turn. A more thorough analysis of this issue appears feasible and relevant.

Against this background, the research team will choose the best course of action to leverage these developments into an SLTA system.

- The current study can be expanded to include more intersections – this requires additional effort to gather reliable intersection data from multiple counties in SE Michigan, but essentially this is a low-risk venture building on the progress reported above. It has a particular advantage of leveraging the naturalistic data already available, though of course it inherits a number of disadvantages of the present analysis, e.g. signal phase is unknown, no interaction existed between the (real) driver and the (planned) advisory system
- The UMTRI driving simulator could be used to analyze interactions between naïve drivers and candidate systems. Issues exist with using a fixed-base simulator, particularly to do with event realism and motion sickness.

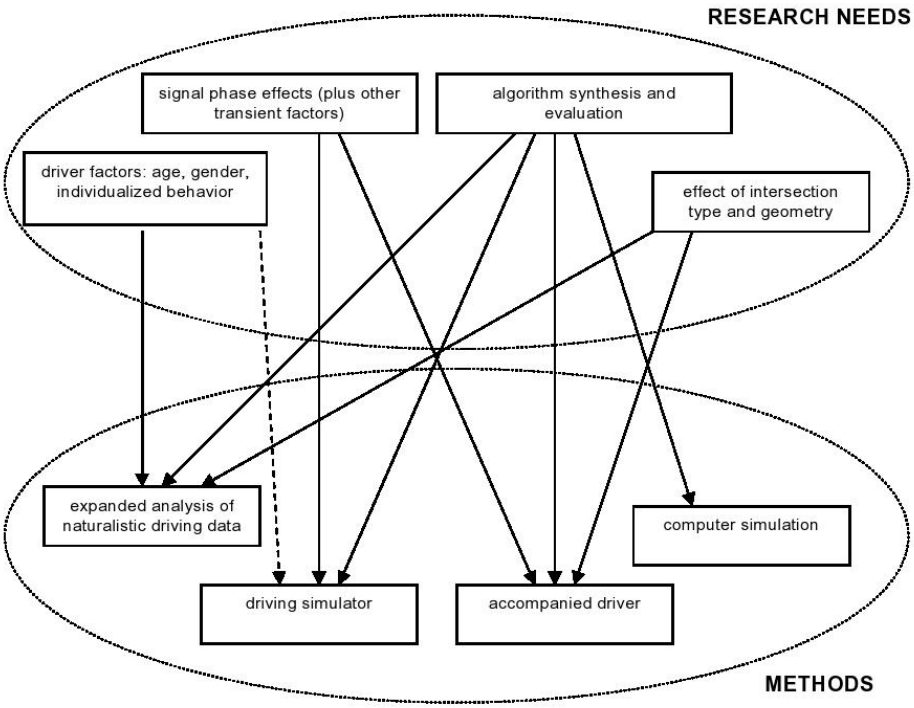

Figure 7.1. Research Problems and Methods for Cooperative Intersection Safety

- Accompanied driver studies on public roads would enable improved data capture via more optimally located radar (side radar to the right of the SV) and video (to capture signal phase). However, the data set would not be large, conflicts would typically be few, and it would be questionable how naturalistic the collected data would really be.
- Computer simulation, particularly if used in combination with naturalistic data analysis, allows the possibility of expanding real datasets to virtual ones. A suitable driver model could be developed directly from the data captured already in this study. It is likely that simulator data would also be needed to model the effects of more severe conflicts not seen in the naturalistic driving data.
- Not included in the figure, UMTRI (together with CA PATH and others) is currently developing an accurate trajectory data capture system using site-based synchronized digital video cameras. This system will not be available for another 12 months or so, but it does offer the future possibility of collecting large datasets for naturalistic left turn behavior at urban or suburban intersections.

References

[1] Allen, B.L., Shin, B.T. and Cooper, P.J., "Analysis of Traffic Conflicts and Collisions", *Transportation Research Record*, **667** (1978) 67-74.

[2] Bar-Shalom, Y. and Li, X.R., "Multitarget-Multisensor Tracking: Principles and Techniques", YBS Publishing, (1995).

[3] Bar-Shalom, Y., Li, X.R. and Kirubarajan, T., "Estimation with Applications to Tracking and Navigation", Wiley, (2001).

[4] Chan, C.Y., "Characterization of Driving Behaviors Based on Field Observation of Intersection Left-Turn Across-path Scenarios", *IEEE Transactions on Intelligent Transportation Systems*, **7(3)** (2006) 322-331.

[5] Ervin, R., Sayer, J., LeBlanc, D., Bogard, S., Mefford, M., Hagan, M., Bareket, Z. and Winkler, C., "Automotive Collision Avoidance System Field Operational Test Methodology and Results", *National Highway Traffic Safety Administration*, UMTRI-2005-7-1 (2005).

[6] Gettman, D. and Head, L., "Surrogate Safety Measures from Traffic Simulation Models, Final Report", *Federal Highway Administration*, FHWA-RD-03-050 (2003).

[7] LeBlanc, D., Sayer, J., Winkler, C., Ervin, R., Bogard, S., Devonshire, J., Mefford, M., Hagan, M., Bareket, Z., Goodsell, R. and Gordon, T., "Road Departure Crash Warning System Field Operational Test: Methodology and Results, Volume 1: Technical Report", *National Highway Traffic Safety Administration*, UMTRI-2006-9-1 (2006).

[8] Li, X.R., "Comparison of Two Measurement Fusion Methods for Kalman-Filter-Based Multisensor Data Fusion", *IEEE Transactions of Aerospace and Electronic Systems*, **37(1)** (2001) 273-280.

[9] Misener, J.A., Chan,C.Y., Cody, D., Dickey, S., Nowakowski, C., Greenhouse, D., Ragland, D., Shladover, S.E. and Vanderwerf, J., "California Intersection Decision Support: A Systems Approach to Achieve Nationally Interoperable Solutions II", *California PATH*, (2007).

Appendix

A.1 Notations

a	DAS acceleration magnitude.
a_x, a_y	DAS acceleration components in local coordinate system.
$\hat{a}$	Acceleration magnitude from X-Y KF estimate.
$\hat{a}_x, \hat{a}_y$	Acceleration estimates by X-Y KF in local coordinate system.
$\hat{a}_X, \hat{a}_Y$	Acceleration estimates by X-Y KF in global coordinate system.
ACAS	Automotive Collision Avoidance System.
BG	Band gap between buffer bands.
BG_i	Band gap made by i-th and (i+1)-th POVs.
CAN	Controller area network.
d_1, d_2	Intersection parameters (distance).
DGPS	Differential GPS.
FOT	Field operational test.
FOV	Field of view.
GPS	Global positioning system.
GT	Gap time.
k_1, k_2	Parameters for U_{pred}.
LB	Leading buffer.
LB_i	Leading buffer for i-th POV.
LTAP/OD	Left turn across path – opposite direction
p_1, p_2	Parameters for U_{ref} for deceleration stage.
PET	Post-encroachment time.
POV	Principal other vehicle.
q_1, q_2	Parameters for U_{ref} for acceleration stage.
$R, \dot{R}$	Range and range rate of a target from the radar.
RDCW	Road Departure Crash Warning System.
s	Path coordinate.
SV	Subject vehicle.

TB	Trailing buffer.
TB_i	Trailing buffer for i-th POV.
Th_{TB}, Th_{LB}	Threshold for acceptable LB/TB.
t_r	Transversal distance of a target from the center line of FOV of the radar.
TTC	Time to collision.
$T2C_{SV}$	Time for SV to arrive at encroachment zone (if SV turns second), or leave encroachment zone (if SV turns first).
$T2C_{POV}$	Time for POV to leave encroachment zone (if SV turns second), or arrive at encroachment zone (if SV turns first).
U	DAS speed.
$\hat{U}$	Speed from the velocity estimates by X-Y KF.
U_{pred}	Predicted speed profile.
U_{ref}	Reference speed profile.
U_{th}	Parameters for U_{ref}.
UMTRI	University of Michigan Transportation Research Institute.
$\hat{v}_x, \hat{v}_y$	Velocity estimates by X-Y KF in local coordinate system.
$\hat{v}_X, \hat{v}_Y$	Velocity estimates by X-Y KF in global coordinate system.
X, Y	X (west-east) and Y (south-north) coordinates transformed from geographical coordinates from GPS.
$\hat{X}, \hat{Y}$	Position estimate by X-Y KF.
δ	Intersection parameter (angle).
Δt_1	Latency difference estimate between GPS and DAS.
Δt_2	Latency difference estimate between GPS and radar.
ΔTh	Difference, $Th_{TB} - Th_{LB}$.
ϕ	Azimuth to a target obtained from R and t_r.
$\dot{\theta}$	DAS yaw rate.
$\hat{\theta}$	Heading angle estimate by X-Y KF.
$\hat{\theta}'$	Heading angle estimate by constant speed KF.
$\dot{\hat{\theta}}'$	Yaw rate estimate by constant speed KF.
$\tilde{\theta}, \dot{\tilde{\theta}}, \ddot{\tilde{\theta}}$	Modified angular estimate by angular measurement fusion.

A.2 RDCW Instrumentations

A.2.1 Overall system

The RDCW was intended to develop two main functions: lateral drift warning (LDW) and curve speed warning (CSW). Test instruments were prepared with 11 Nissan Altima 3.5SE sedans (model year 2003) (Figure A.1). Figure A.2 shows the RDCW system with architecture installed in the above vehicles.

Figure A.1. Vehicle fleet. (Reproduced from [7])

Table A.1 shows the primary and supporting sensors in the LDW and CSW. Among these, the GPS, yaw rate sensor and accelerometer in the DAS were used to reconstruct SV trajectories. For POVs, necessary information about their motions was given by the forward radars. In addition, video clips by the CCD camera were accessed to confirm manually if the reconstructed trajectories were reasonable. Note that there were two GPSs installed in the vehicles. One with differential corrections was installed in the DAS and the other was in the CSW. For the SV trajectory reconstruction, the former GPS was adopted.

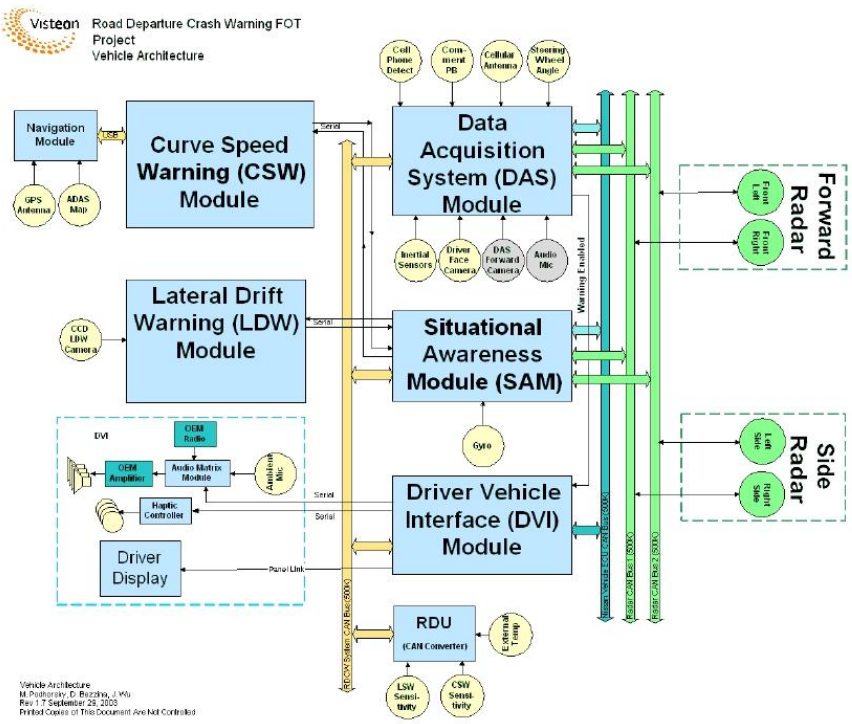

Figure A.2. Schematic diagram of the RDCW system architecture (Reproduced from [7])

Table A.1. Key sensors used by the LDW and CSW systems (Reproduced from [7])

Sensor	LDW system sensors		CSW system	
	Primary	Supporting	Primary	Supporting
Forward CCD camera	X			X
GPS		X	X	
Digital map		X	X	
Digital map look-aside database		X		
Vehicle speed	X	X	X	
Yaw rate gyro		X	X	
Driver brake switch		X		
Driver turn signal switch		X		X
Forward-looking radars		X		
Side-looking radars		X		

The CCD camera was installed inside the cabin on the front windshield. The forward radars and the side radars were mounted inside the front bumper as shown in Figure A.3.

Figure A.3. Physical installation of the forward and side radars. (Reproduced from [7])

It is important to convince that the RDCW data contains naturalistic driver behaviors on left turns at intersections without warnings by the LDW or CSW. The LDW suppresses warning requests in particular situations;

1) The vehicle speed is 11.1m/s or has been $\leq$ 11.1 m/s within the last 5s,
2) The vehicle has completed a lane change within the last 2s,
3) The lateral speed has exceeded 1m/s within the last 2s,
4) A turn signal has been on within the last 5s,
5) The brake pedal has been pressed within the last 5s and
6) The curve radius is $\leq$ 250 m.

Also, the CSW suppresses warning requests in the following cases;

1) The vehicle speed is below 8m/s and
2) The vehicle deceleration is less than the requested deceleration.

The CSW determines the vehicle position and a predicted path using digital map information. Given a predicted path, it provides a warning to slow down the vehicle speed if an excessive vehicle speed is anticipated by comparing with a threshold for the lateral speed.

With those conditions, there were no events detected with warnings from the LDW. On the other hand, there were 5 events with warnings from the CSW by 2 drivers at 2 intersections during left turns. Excluding these 5 events, it can be said that the driver behavior on left turns in the RDCW data is assumed to be naturalistic.

A.2.2 Specifications of instruments

Specifications of the instruments relevant to the left turn analyses are shown in the tables below.

Table A.2. Accelerometer

Input range	±2g
Zero g drift	±30mV
Sensitivity	1V/g
Noise RMS	1mg
Bandwidth	DC -50Hz
Data sampling rate	20Hz

Table A.3. CCD camera

Color	Monochrome
Horizontal FOV	40±4deg
Vertical FOV	30±3deg
Data sampling rate	10Hz

Table A.4. Forward radars

Range (moving objects)	120±1m
Range (stationary objects)	60±1m
Range rate	±48.88±0.17m/s
Horizontal FOV	8±0.6deg from 40 to 100m
Vertical FOV	3deg
Frequency	77GHz
Data sampling rate	20Hz

Table A.5. GPS

Position error (horizontal direction)	Significantly less than 1m at 2σ (95%) and close to 1m at 3 σ (99%)
Position error (vertical direction)	2 to 2.5 times greater than horizontal error

Table A.6. Yaw rate transducer

Sensitivity (-30 to 80 ºC)	25.00±2.5mV/deg/sec
Sensitivity drift (-30 to 80 ºC)	±5%
Dynamic range	±60deg/s
Linearity (at -50 [deg/s])	±0.5%
Linearity (at -60 [deg/s])	±5%
Frequency response at 7 [Hz]	-7.0dB (minimum) and -2.0dB (maximum)
Cross axis sensitivity	±5.0%
Data sampling rate	10Hz

A.3 Transformation of Geographical Coordinates into Cartesian Coordinates

The GPS coordinates given in the geographical coordinates (θ, γ) can be transformed into the Cartesian coordinates as follows.

A point on the earth surface, *P*, is denoted as,

$$\{\mathbf{r}_P\}_{spherical} = R_e \mathbf{e}_r \tag{A.1}$$

by using the spherical coordinate system. R_e is the radius of the earth and $\mathbf{e}_r$ is a unit vector which can be rewritten by a set of unit vectors of the *XYZ* global (fixed) coordinate system as,

$$\begin{aligned} \mathbf{e}_r &= \cos\phi\mathbf{K} + \sin\phi\mathbf{n} \\ &= \sin\phi\cos\theta\mathbf{I} + \sin\phi\sin\theta\mathbf{J} + \cos\phi\mathbf{K}. \end{aligned} \tag{A.2}$$

Thus, the point, *P*, can be given in the *XYZ* global coordinate system by,

$$\begin{aligned} \{\mathbf{r}_P\}_{XYZ} &= R_e\left(\sin\phi\cos\theta\mathbf{I} + \sin\phi\sin\theta\mathbf{J} + \cos\phi\mathbf{K}\right) \\ &= R_e\left(\cos\gamma\cos\theta\mathbf{I} + \cos\gamma\sin\theta\mathbf{J} + \sin\gamma\mathbf{K}\right). \end{aligned} \tag{A.3}$$

Then, a rotation matrix to rotate the global axes about the *Z* axis by $-(\pi/2 - \theta)$ followed by another rotation about the *X* axis by $(\pi - \phi)$ is given by,

$$\mathbf{R} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -\cos\phi & \sin\phi \\ 0 & -\sin\phi & -\cos\phi \end{bmatrix} \begin{bmatrix} \sin\theta & -\cos\theta & 0 \\ \cos\theta & \sin\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{A.4}$$

Therefore,

$$\{\mathbf{r}_P\}_{X'Y'Z'} = \mathbf{R}\{\mathbf{r}_P\}_{XYZ}. \tag{A.5}$$

The resulted orientation becomes the one shown in Figure A.4(b).

For the reconstruction of the SV trajectory, the rotation matrix needs to be defined only once for each left turn by a datum position specified by, (θ_0, γ_0), such as the intersection node or initial SV coordinates in a left turn event and this matrix can be used for the other positions too, since the vehicle motion is considered approximately planar. Therefore, the transformation of SV coordinates from the *XYZ* global coordinate system to the *X'Y'Z'* local coordinate system is described by,

$$\{\mathbf{r}_{SV'}\}_{X'Y'Z'} = \mathbf{R}_0 \{\mathbf{r}_{SV'}\}_{XYZ} \tag{A.6}$$

where

$$\mathbf{R}_0 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -\cos\phi_0 & \sin\phi_0 \\ 0 & -\sin\phi_0 & -\cos\phi_0 \end{bmatrix} \begin{bmatrix} \sin\theta_0 & -\cos\theta_0 & 0 \\ \cos\theta_0 & \sin\theta_0 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{A.7}$$

The radius of the earth used here is defined as 6371000 [m]. Although the value is approximate because of the non-spherical shape of the earth, the distance between a final position of the SV in a typical left turn event obtained by using the semimajor axis and another obtained by using the semiminor axis is about 15 [cm]. Therefore, this deviation is negligible in the conflict metric analysis.

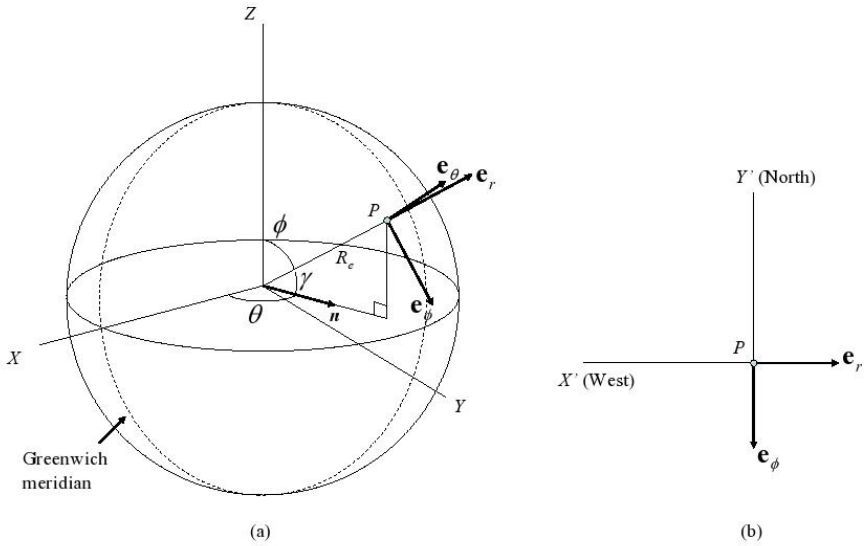

Figure A.4. Coordinate systems used in the coordinate transformation. (a) three dimensional view and (b) top view of the designated coordinate orientation for the planar motion.

A.4 Equations for X-Y Kalman Filter

The first KF was applied to the SV coordinates given by the GPS in order to obtain the kinematic estimates, namely, position, velocity and acceleration estimates. The Wiener process acceleration model [3] was used for the plant equation where the time derivative of the acceleration (jerk) was assumed to be a zero-mean white noise. Thus, the system equations, for one time step, $T(k)=t_{k+1}-t_k$, are

$$\mathbf{x}(k+1) = \mathbf{A}(k)\mathbf{x}(k) + \mathbf{G}v \tag{A.8}$$

where

$$\mathbf{x} = \begin{Bmatrix} X(k) \\ v_X(k) \\ a_X(k) \\ Y(k) \\ v_Y(k) \\ a_Y(k) \end{Bmatrix}, \quad \mathbf{A}(k) = \begin{bmatrix} 1 & T(k) & 0.5T(k)^2 & 0 & 0 & 0 \\ 0 & 1 & T & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & T(k) & 0.5T(k)^2 \\ 0 & 0 & 0 & 0 & 1 & T(k) \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{G}(k) = \begin{bmatrix} 0.5T(k)^2 \\ T(k) \\ 1 \\ 0.5T(k)^2 \\ T(k) \\ 1 \end{bmatrix} \tag{A.9}$$

with the measurement equation,

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) \tag{A.10}$$

where

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \tag{A.11}$$

The covariance matrices **Q** and **R** for the process noise, v, and measurement noise are given by

$$\mathbf{Q}(k) = \sigma_v \mathbf{G}(k)\mathbf{G}(k)^T \tag{A.12}$$

and

$$\mathbf{R} = \sigma_w \mathbf{C}\mathbf{C}^T, \tag{A.13}$$

respectively. The state prediction and measurement prediction are represented by

$$\begin{aligned}\hat{\mathbf{x}}(k+1\,|\,k) &= \mathbf{A}(k)\mathbf{x}(k\,|\,k)+\mathbf{G}(k)v \\ \hat{\mathbf{y}}(k+1\,|\,k) &= \mathbf{C}\hat{\mathbf{x}}(k+1\,|\,k)\end{aligned} \tag{A.14}$$

and the updated state estimate is given by

$$\hat{\mathbf{x}}(k+1\,|\,k+1) = \hat{\mathbf{x}}(k+1\,|\,k)+\mathbf{W}(k+1)\{\mathbf{y}(k+1)-\hat{\mathbf{y}}(k+1\,|\,k)\} \tag{A.15}$$

where **W** is the filter gain that is found by

$$\mathbf{W}(k+1) = \{\mathbf{A}(k)\mathbf{P}(k\,|\,k)\mathbf{A}^{T}(k)+\mathbf{Q}(k)\}\mathbf{C}^{T}\{\mathbf{R}+\mathbf{C}\mathbf{P}(k+1\,|\,k)\mathbf{C}^{T}\}^{-1}. \tag{A.16}$$

The state covariance matrix **P** is described as

$$\begin{aligned}\mathbf{P}(k+1\,|\,k) &= \mathbf{A}(k)\mathbf{P}(k\,|\,k)\mathbf{A}^{T}(k)+\mathbf{Q}(k) \\ \mathbf{P}(k+1\,|\,k+1) &= \mathbf{P}(k+1\,|\,k)-\mathbf{W}(k+1)\{\mathbf{R}(k+1)+\mathbf{C}\mathbf{P}(k+1\,|\,k)\mathbf{C}^{T}\}\mathbf{W}^{T}(k+1).\end{aligned} \tag{A.17}$$

Since appropriate initial conditions must be given for the KF (forward KF), another KF is swept from the last point of the SV path (backward KF) to find suitable initial conditions for the forward KF. The combination of these KFs can increase the overall accuracy of the estimates of the SV motion. The system equations for the backward KF are described similarly to the forward KF as follows,

$$\mathbf{x}(k) = \mathbf{A}^{-1}(k+1)\mathbf{x}(k+1)+\mathbf{G}v \tag{A.18}$$

with the same measurement equation as for the forward KF,

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k). \tag{A.19}$$

A.5 Trajectories of Vehicle Corners

In order to calculate the conflict metrics, the size of the vehicles must be considered. Thus, the trajectories of the corner points of the vehicles were constructed by using $\left(\hat{X},\hat{Y}\right)$, $\tilde{\theta}$ and the geometry of the vehicle body. According to Figure, the position vectors of the corner points are given by

$$\begin{aligned}
\mathbf{r}_1 &= \left(\hat{X} + L_a' \cos\tilde{\theta} + L_c \sin\tilde{\theta}\right)\mathbf{I} + \left(\hat{Y} + L_a \sin\tilde{\theta} - L_c \cos\tilde{\theta}\right)\mathbf{J} \\
\mathbf{r}_2 &= \left(\hat{X} + L_a' \cos\tilde{\theta} - L_c \sin\tilde{\theta}\right)\mathbf{I} + \left(\hat{Y} + L_a' \sin\tilde{\theta} + L_c \cos\tilde{\theta}\right)\mathbf{J} \\
\mathbf{r}_3 &= \left(\hat{X} - L_b' \cos\tilde{\theta} + L_c \sin\tilde{\theta}\right)\mathbf{I} + \left(\hat{Y} - L_b' \sin\tilde{\theta} - L_c \cos\tilde{\theta}\right)\mathbf{J} \\
\mathbf{r}_4 &= \left(\hat{X} - L_b' \cos\tilde{\theta} - L_c \sin\tilde{\theta}\right)\mathbf{I} + \left(\hat{Y} - L_b' \sin\tilde{\theta} + L_c \cos\tilde{\theta}\right)\mathbf{J}
\end{aligned} \tag{A.20}$$

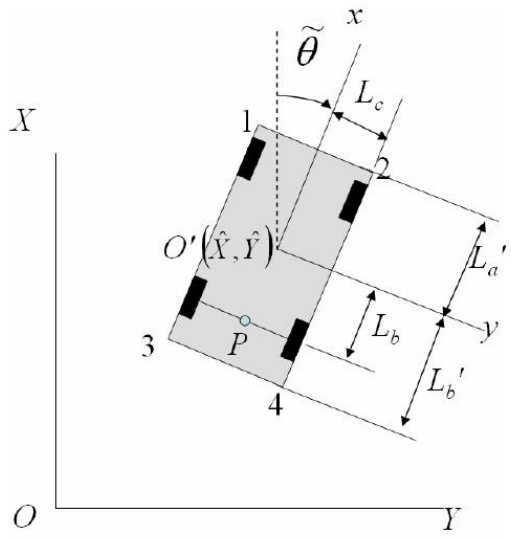

Figure A.5. Geometry of the SV and coordinate system for the corner trajectories (top view).

Instead of applying the KF to each corner trajectory, the first and second time derivatives are computed to find the velocity and acceleration by using the trajectory of O'. The velocity vectors can be represented as

$$
\begin{aligned}
\dot{\mathbf{r}}_1 &= \left(\dot{\hat{X}} - L'_a\dot{\tilde{\theta}}\sin\tilde{\theta} + L_c\dot{\tilde{\theta}}\cos\tilde{\theta}\right)\mathbf{I} + \left(\dot{\hat{Y}} + L_a\dot{\tilde{\theta}}\cos\tilde{\theta} + L_c\dot{\tilde{\theta}}\sin\tilde{\theta}\right)\mathbf{J} \\
\dot{\mathbf{r}}_2 &= \left(\dot{\hat{X}} - L'_a\dot{\tilde{\theta}}\sin\tilde{\theta} - L_c\dot{\tilde{\theta}}\cos\tilde{\theta}\right)\mathbf{I} + \left(\dot{\hat{Y}} + L'_a\dot{\tilde{\theta}}\cos\tilde{\theta} - L_c\dot{\tilde{\theta}}\sin\tilde{\theta}\right)\mathbf{J} \\
\dot{\mathbf{r}}_3 &= \left(\dot{\hat{X}} + L'_b\dot{\tilde{\theta}}\sin\tilde{\theta} + L_c\dot{\tilde{\theta}}\cos\tilde{\theta}\right)\mathbf{I} + \left(\dot{\hat{Y}} - L'_b\dot{\tilde{\theta}}\cos\tilde{\theta} + L_c\dot{\tilde{\theta}}\sin\tilde{\theta}\right)\mathbf{J} \\
\dot{\mathbf{r}}_4 &= \left(\dot{\hat{X}} + L'_b\dot{\tilde{\theta}}\sin\tilde{\theta} - L_c\dot{\tilde{\theta}}\cos\tilde{\theta}\right)\mathbf{I} + \left(\dot{\hat{Y}} - L'_b\dot{\tilde{\theta}}\cos\tilde{\theta} - L_c\dot{\tilde{\theta}}\sin\tilde{\theta}\right)\mathbf{J}.
\end{aligned}
\tag{A.21}
$$

The acceleration vectors are

$$
\begin{aligned}
\ddot{\mathbf{r}}_1 &= \left(\ddot{\hat{X}} - L'_a\ddot{\tilde{\theta}}\sin\tilde{\theta} - L'_a\dot{\tilde{\theta}}^2\cos\tilde{\theta} + L_c\ddot{\tilde{\theta}}\cos\tilde{\theta} - L'_c\dot{\tilde{\theta}}^2\sin\tilde{\theta}\right)\mathbf{I} \\
&\quad + \left(\ddot{\hat{Y}} - L'_a\ddot{\tilde{\theta}}\sin\tilde{\theta} - L'_a\dot{\tilde{\theta}}^2\cos\tilde{\theta} + L_c\ddot{\tilde{\theta}}\cos\tilde{\theta} - L'_c\dot{\tilde{\theta}}^2\sin\tilde{\theta}\right)\mathbf{J} \\
\ddot{\mathbf{r}}_2 &= \left(\ddot{\hat{X}} - L'_a\ddot{\tilde{\theta}}\sin\tilde{\theta} - L'_a\dot{\tilde{\theta}}^2\cos\tilde{\theta} - L_c\ddot{\tilde{\theta}}\cos\tilde{\theta} + L'_c\dot{\tilde{\theta}}^2\sin\tilde{\theta}\right)\mathbf{I} \\
&\quad + \left(\ddot{\hat{Y}} - L'_a\ddot{\tilde{\theta}}\sin\tilde{\theta} - L'_a\dot{\tilde{\theta}}^2\cos\tilde{\theta} - L_c\ddot{\tilde{\theta}}\cos\tilde{\theta} + L'_c\dot{\tilde{\theta}}^2\sin\tilde{\theta}\right)\mathbf{J} \\
\ddot{\mathbf{r}}_3 &= \left(\ddot{\hat{X}} + L'_b\ddot{\tilde{\theta}}\sin\tilde{\theta} + L'_b\dot{\tilde{\theta}}^2\cos\tilde{\theta} + L_c\ddot{\tilde{\theta}}\cos\tilde{\theta} - L'_c\dot{\tilde{\theta}}^2\sin\tilde{\theta}\right)\mathbf{I} \\
&\quad + \left(\ddot{\hat{Y}} - L'_b\ddot{\tilde{\theta}}\cos\tilde{\theta} + L'_b\dot{\tilde{\theta}}^2\sin\tilde{\theta} + L_c\ddot{\tilde{\theta}}\sin\tilde{\theta} + L'_c\dot{\tilde{\theta}}^2\cos\tilde{\theta}\right)\mathbf{J} \\
\ddot{\mathbf{r}}_4 &= \left(\ddot{\hat{X}} + L'_b\ddot{\tilde{\theta}}\sin\tilde{\theta} + L'_b\dot{\tilde{\theta}}^2\cos\tilde{\theta} - L_c\ddot{\tilde{\theta}}\cos\tilde{\theta} + L'_c\dot{\tilde{\theta}}^2\sin\tilde{\theta}\right)\mathbf{I} \\
&\quad + \left(\ddot{\hat{Y}} - L'_b\ddot{\tilde{\theta}}\cos\tilde{\theta} + L'_b\dot{\tilde{\theta}}^2\sin\tilde{\theta} - L_c\ddot{\tilde{\theta}}\sin\tilde{\theta} - L'_c\dot{\tilde{\theta}}^2\cos\tilde{\theta}\right)\mathbf{J}.
\end{aligned}
\tag{A.22}
$$

It is needed to find analytical expressions for $\dot{\tilde{\theta}}$ and $\ddot{\tilde{\theta}}$ in the above equations. This can be achieved by using the constraint on the center of the rear wheels by assuming that they are connected by a rigid shaft. Since they are velocity-constrained (non-holonomic), the constraint equation can be found from the differential of the position vector. The position vector of the point *P* is

$$
\mathbf{r}_P = \left(\hat{X} - L_b\cos\tilde{\theta}\right)\mathbf{I} + \left(\hat{Y} - L_b\sin\tilde{\theta}\right)\mathbf{J} \tag{A.23}
$$

and thus, after the coordinate transformation from the global coordinate system to the local coordinate system, the differential of this vector is written by

$$
\begin{aligned}
d\mathbf{r}_P &= \left(d\hat{X} + (L_b\sin\tilde{\theta})d\tilde{\theta}\right)\mathbf{I} + \left(d\hat{Y} - (L_b\cos\tilde{\theta})d\tilde{\theta}\right)\mathbf{J} \\
&= \left(\cos\tilde{\theta}d\hat{X} + \sin\tilde{\theta}d\hat{Y}\right)\mathbf{i} + \left(-\sin\tilde{\theta}d\hat{X} + \cos\tilde{\theta}d\hat{Y} - L_b d\tilde{\theta}\right)\mathbf{j}
\end{aligned}
\tag{A.24}
$$

The rear wheels are constrained in the lateral direction by assuming the tire slip is negligible, so the second term of the last equation becomes zero. Consequently, $\dot{\tilde{\theta}}$ and $\ddot{\tilde{\theta}}$ are given by

$$\dot{\tilde{\theta}} = \frac{-\dot{\tilde{X}}\sin\tilde{\theta} + \dot{\tilde{Y}}\cos\tilde{\theta}}{L_b}, \text{ and} \tag{A.25}$$

$$\ddot{\tilde{\theta}} = \frac{-\ddot{\tilde{X}}\sin\tilde{\theta} - \dot{\tilde{X}}\dot{\tilde{\theta}}\cos\tilde{\theta} + \ddot{\tilde{Y}}\cos\tilde{\theta} - \dot{\tilde{Y}}\dot{\tilde{\theta}}\sin\tilde{\theta}}{L_b}. \tag{A.26}$$

1.3 Field Testing

This section documents an update on the research activities especially on the selection of sites for conducting field data collection, where a combination of sensing technologies are deployed to detect and measure traffic conflict scenarios that are candidate situations for offering warnings to the drivers. The first part of the report describes the procedures taken to carry out this task and the selection criteria for candidate intersection locations for field testing. This task is a continuing effort from earlier investigations when a methodology was developed to extract scenario-specific parameters of interests. An enhanced set of sensing devices are evaluated and combined to acquire more accurate and informative data for the developments of warning algorithms. The outcome of this work will provide valuable inputs for the implementation of intersection safety countermeasures

Selection of CICAS-SLTA Candidate Intersection for Data Collection and Future Field Operational Tests

The role of CICAS-SLTA will be to supplement standard protection measures and also provide additional safety, cost, and delay benefits. By having reviewed the currently available left-turn treatments and the literature examining their benefits and drawbacks, clearer guidelines can be developed for the CICAS-SLTA testing process. To effectively test the different aspects, the first goal is to choose intersections that will accommodate the range of characteristics that are vital for proper functioning of CICAS-SLTA. The preliminary steps completed for this selection process are described in this section.

Prototype testing of CICAS-SLTA has been conducted at the Richmond Field Station complex in Richmond, California by PATH. The complex includes an experimental intersection that is fully equipped with sensing, computing and communication equipment to allow the evaluation of commercially-off-the-shelf devices as well as testing the response and acceptance of drivers in CICAS-SLTA. To prepare for field-operational tests (FOT) for potential deployment of CICAS-SLTA, intersections in the San Francisco Bay Area were screened to identify the potential candidates most suitable for field experimentations.

Screening Process

The intersection screening process consisted of two parts:

1) Identifying the best candidate intersections based on historical collision data and project requirements.
2) Conducting field visits to the sites to investigate the attributes of such candidates.

The data review portion required analyzing traffic collision databases known as the Statewide Integrated Traffic Reporting System (SWITRS) that are published by the California Highway Patrol. SWITRS contains data elements from all traffic collisions in the state which are reported using a standard system. Data for the years 2002-2004 from counties comprising the Bay Area was used to identify LTAP-OD collisions at intersections by determining the collision location, direction of vehicles involved and the

movement preceding the collision. The intersections with high numbers of such LTAP-OD collisions were first identified. Furthermore, greater considerations are given to locations that provide synergistic linkage to other projects. For example, intersections that are included in the development of VII California testbed or already in the paths of Alameda County Congestion Management Corridor will be given higher priority.

Following the initial screening, approximately 30 intersections were visited to explore the intersection attributes, such as:

- Left-turn pockets
- Left-turn lanes
- Intersection geometry
- Line of sight
- Pedestrian levels
- POV volume, speed
- POV # of lanes
- % of left-turning vehicles
- % of heavy-duty vehicles
- Clustered/dispersed traffic flow

Features of Candidate Intersections

After collision data analysis and field visits, the list of sites was reduced to sixteen candidates. These sites were determined to have the best potential for future Field Operational Tests (FOTs) with appropriate locations for equipment placement and geometries suited for CICAS-SLTA intervention.

To better present and visualize the candidates, diagrams were created to show the distribution of left-turn collisions at the intersection. Figure 51 shows a sample diagram for one of the intersections in the study, which highlights left-turn collisions over the three year period from 2002-2004. The main State Route is in the North/South direction and the Cross Street directs traffic East and West. The blue arrows signify that the left-turns in all directions are permitted left-turns. CICAS-SLTA would focus on the left turns from the State Route to the Cross Street and the considerable number of LTAP-OD collisions shown by an ‘X’ for each collision. There is also an ‘o’ for the left-turn in the South-bound direction that symbolizes a rear-end collision occurred while the vehicle was trying to turn. In total, this intersection had 23 left-turn related collisions over the three year period, with seven other collisions not shown on the diagram that are usually the result of vehicles running a red light. This intersection represented one of the highest number of LTAP-OD collisions and is a candidate for a CICAS-SLTA system or other standard protection safety countermeasures to improve the level of safety.

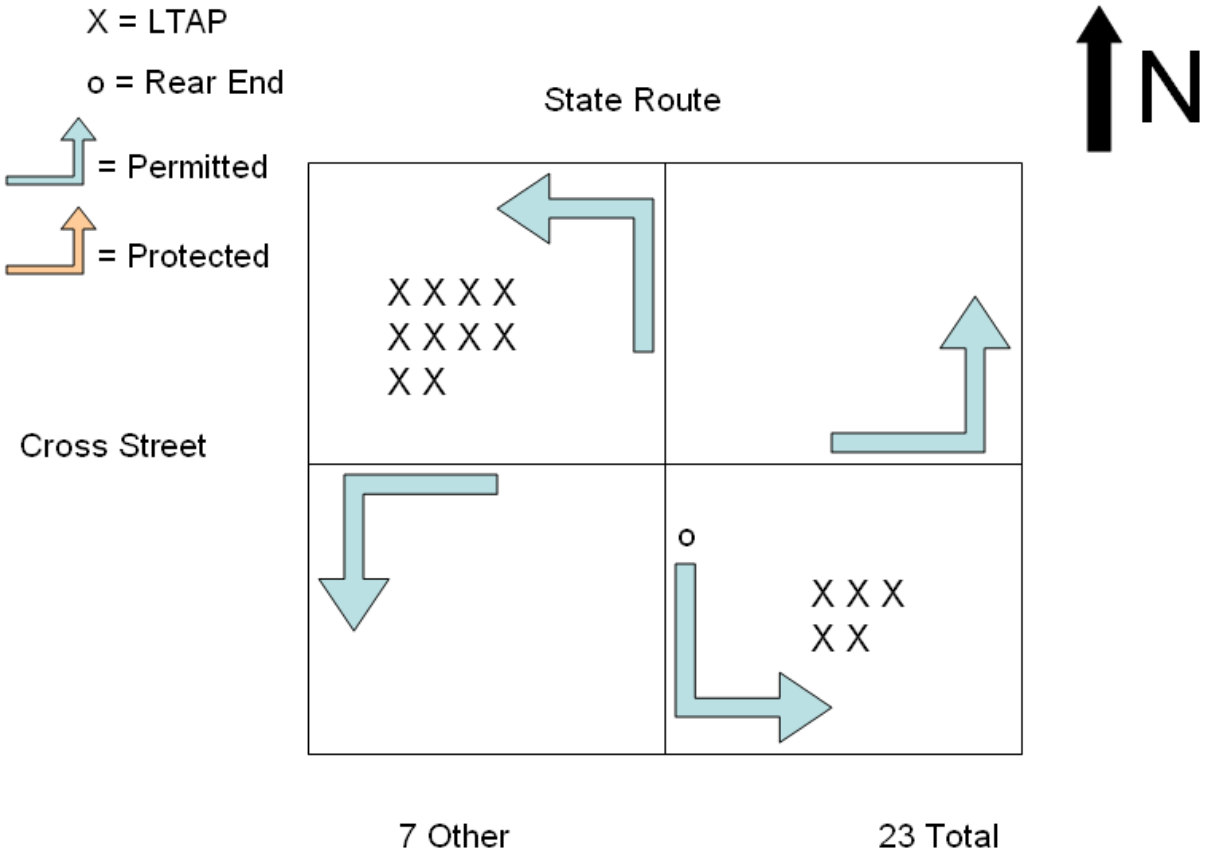

Figure 53, Diagram of Left-Turn Collisions.

Deleted: 51

Several general characteristics of the candidate intersections were evident after completing the investigations and reviewing the diagrams:

1) Most of these sites had very high traffic volume in the primary traffic direction, likely contributing the high number of collisions. High concentrations of collisions were also found to occur when vehicles are turning left from the direction of the primary corridor onto cross streets. The CICAS-SLTA designed for such locations will need the capability to monitor a relatively large number of vehicles at relatively high speeds.
2) Most of the sites have two lanes of opposite traffic, but only half of them had a left-turn pocket. This observation implies that locations exceeding two lanes with safety problems are most likely already being considered for or safeguarded by protective left turn setups. This is in line with many published protected left-turn warrants that state three or more lanes would require a protected left-turn [14]. It is an important aspect that will also imply a typical CICAS-SLTA design can focus on intersections of similar sizes without resulting in any selection bias. The

turning time needed for completing left turns, which depends on the intersection size or more specifically the number of crossing lanes, is an essential parameter in determining the warning threshold.

3) Few intersections have a high pedestrian presence. This means that only selective sites for CICAS-SLTA consideration need to include pedestrian sensing capabilities.

The candidate locations were found to be located almost exclusively along three major corridors. Table 6 classifies the intersections by corridor A, B, and C (SR-82, SR-123, and SR-13 respectively) and shows the distribution of collisions in each direction. Intersections 13 and 16 consist of major routes in both the North/South and East/West directions and subsequently have collisions listed for each direction. Table 7 provides a qualitative description of various parameters associated with the intersections. Only qualitative assessment is provided because accurate quantitative values are not available.

Table 6. LTAP-OD Collisions at Selective Locations

Site Number	Intersection	Number of LTAP-OD Collisions			
		North	South	East	West
1	SR 82 / Rosedale	2	11		
2	SR 82 / Broadway	1	3		
3	SR 82 / Howard	2	7		
4	SR 82 / Tilton	2	3		
5	SR 82 / Chapin	1	4		
6	SR 82 / Floribunda	10	5		
7	SR 82 / 12th	0	0		
8	SR 123 / W Grand	0	3		
9	SR 123 / Brighton	1	6		
10	SR 123 / Cedar	3	1		
11	SR 123 / Potrero			5	2
12	SR 13 / SR 123			5	3
13	SR 13 / Telegraph	3	0	3	1
14	SR 13 / Sacramento			1	0
15	SR 13 / Claremont			2	3
16	Eastshore Blvd / Potrero	3	5	2	4

Table 7. Characteristics of Selective Locations

Intersection	Left turn pocket	Ped Vol	Veh Vol	POV Speed	POV lanes	Left turn %	% heavy veh	Geometry	Line of sight
SR 82 / Rosedale (Ray)	No	Low	High	High	2	Low	Low	Standard	Good
SR 82 / Broadway	No	Low	High	High	2	Low	Low	Standard	Good

SR 82 / Howard	No	Low	High	High	2	Low	Low	Standard	Good
SR 82 / Tilton	No	Low	High	High	2	Low	Low	Standard	Poor
SR 82 / Chapin	No	Low	High	High	2	Low	Low	Standard	Good
SR 82 / Floribunda	No	Low	High	High	2	Low	Low	Standard	Poor
SR 82 / 12th	No	Low	High	High	2	Low	Low	Standard	Good
SR 123 / W Grand	Yes	High	High	Low	2	Low	Low	Offset	Poor
SR 123 / Brighton	Yes	Low	Low	High	2	Low	Low	Offset	Good
SR 123 / Cedar	Yes	Low	High	High	2	High	High	Offset	Good
SR 123 / Potrero	Yes	Low	Low	Low	2	High	High	Offset	Poor
SR 13 / SR 123	Yes	Low	High	Low	2	Low	Low	Offset	Good
SR 13 / Telegraph	Yes	High	High	Low	2	High	Low	Offset	Good
SR 13 / Sacramento	Yes	Low	High	High	2	High	Low	Standard	Good
SR 13 / Claremont	No	Low	High	Low	2	High	Low	Standard	Good
Eastshore Bl / Potrero	Yes	Low	Low	Low	1	Low	Low	Offset	Good

Distinct features are present among the corridors, irrespective of the aforementioned common attributes. To summarize:

1) Corridor A (SR-82) has very narrow lanes, no left-turn pockets, high traffic volume, and relatively low percentage of left-turn traffic when compared to the straight-through traffic. On this corridor, there is considerable pressure to make a left turn from vehicles following or waiting behind.
2) Corridor B (SR-123) is relatively wide and provides left-turn pockets. Some intersections on this corridor have a relatively large percentage of left-turn traffic. The offset geometry and quality of lines of sight vary.
3) Corridor C (SR-13) has intersection features that vary significantly from one to another. The speed also varies from certain blocks to others. A few intersections have some pedestrian presence.

Technical Approach for Field Observation

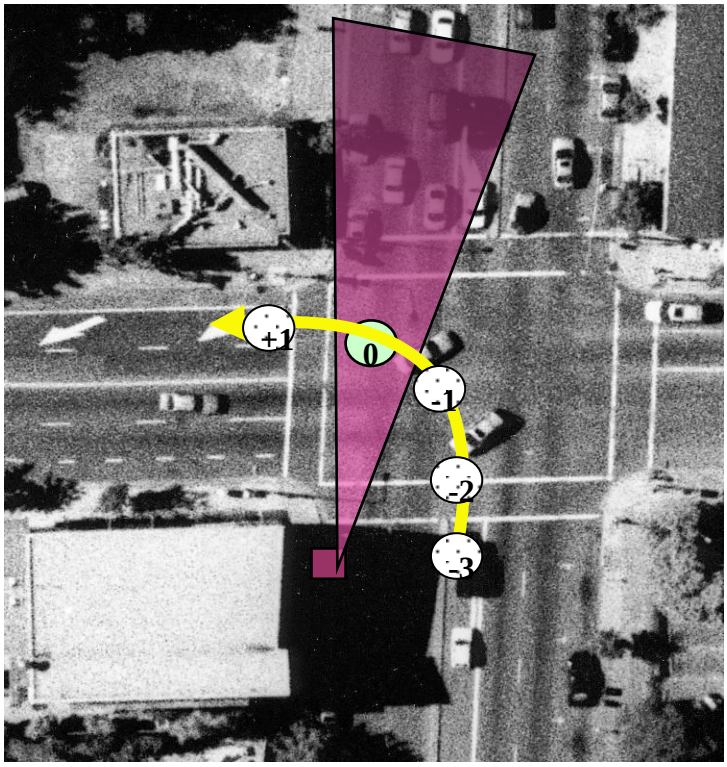

As an important part of the studies in the intersection safety systems, considerable efforts were dedicated to the observation and testing of driver decisions and response under various traffic conditions. For example, a methodology was developed at PATH to facilitate the understanding of driving behaviors in intersection turning scenarios. Figure 52 depicts an exemplar setup at an intersection for filed data collection. A mobile platform, consisting primarily of a radar sensor and a data acquisition computer, was stationed at selected intersections. As a target vehicle makes a left turn at the intersection, the movements of the target vehicle relative to the oncoming traffic can be measured and analyzed.

Figure 54. Schematic Diagram of Data Collection Setup

Deleted: 52

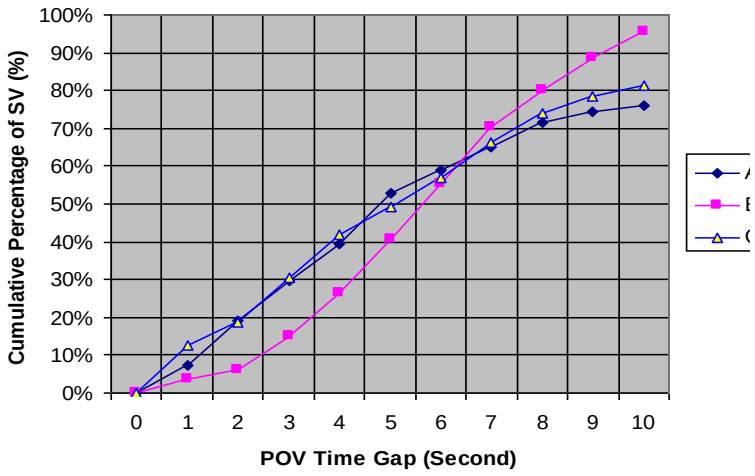

Figure 55. Accepted Gaps of Oncoming Traffic by Left-Turn Drivers at Sites A, B, and C

Deleted: 53

Subsequently, the interaction of the turning vehicle and the oncoming traffic can be quantified by calculating the time gap or time lag observed in the oncoming traffic when the subject vehicle (SV) moves into the zone of potential conflicts, which in this case is located at the point denoted by "0" in Figure 51. The time gap or lag is defined by the projected arrival time of the nearest oncoming vehicle, which is called the principal other vehicle (POV). Figure 52 shows the distribution of time gaps accepted by all drivers making a left turn at three different sites. The accepted gap curve may be influenced by traffic conditions, intersection layout, and the behaviors of individual drivers.

The methodology described above was proven to be very revealing and illuminating, and considerable insights have been gained by deploying the same methodology at several observation sites in previous efforts. However, due toe the limited availability of commercially-off-the-shelf products, the primary sensor that has been used was a Doppler radar (EVT-300 ^{TM} by Eaton-Vorad). The EVT-300 utilizes mono-pulse radar technology in conjunction with digital signal processing, resulting in a forward-looking collision warning system. It is mainly developed and used by freight vehicles to alert drivers of imminent collisions. In this study, we adopted the radar component with in-house data acquisition codes to acquire target information, which exhibited several limitations on its own:

(1) This Doppler radar only detects moving targets, and stationary targets are dropped out. The disappearance and reappearance of targets as they stop and restart can easily cause confusion in defining the exact interaction among various targets in the range.
(2) The radar only outputs the measurements of seven targets at one time. Thus, certain targets will not be shown when traffic is heavy.
(3) Like many other sensing devices, the radar can not detect targets that are out of line-of-sights. Even though multiple sensors can be placed at different locations in an integrated setup, the data processing for sensor fusion is no easy task.
(4) Also like many other radar devices, the radar sensor is subject to target splitting or surface scattering, therefore occasional misidentification is unavoidable.
(5) Radar can detect targets with an identifiable radar cross section, yet it has no capabilities to identify vehicle classification or size.

In consideration of the above constraints and limitations, it became apparent that an enhanced set of sensor suite will be strongly desirable even though it becomes increasingly expensive and cumbersome for integration when a more sophisticated sensor suite is required. It should be noted here that the current philosophy of the CICAS project is to develop and utilize an operational concept that is based on vehicle-infrastructure cooperation. In this case, wireless communication links will allow vehicles and infrastructure to exchange data so that the burden of sensing vehicle motions does not fall completely on the infrastructure side. Nevertheless, for the interests of understanding of driver behaviors and developing CICAS, the utilization of enhanced sensor configurations will be beneficial.

Enhanced Sensor Configuration for Field Observation

Figure 54 depicts an exemplar setup of sensor suite to capture vehicle movements near the intersection. The sensor suite includes:

1 A long-range radar to capture the oncoming traffic stream
2 Two short-range radar located at different corners of the intersection to capture the left-turn movements of the left-turning vehicle
3 A laser scanner to provide a 240-degree profiling of all targets around the intersection
4 Video cameras to capture images of opposite views from different directions

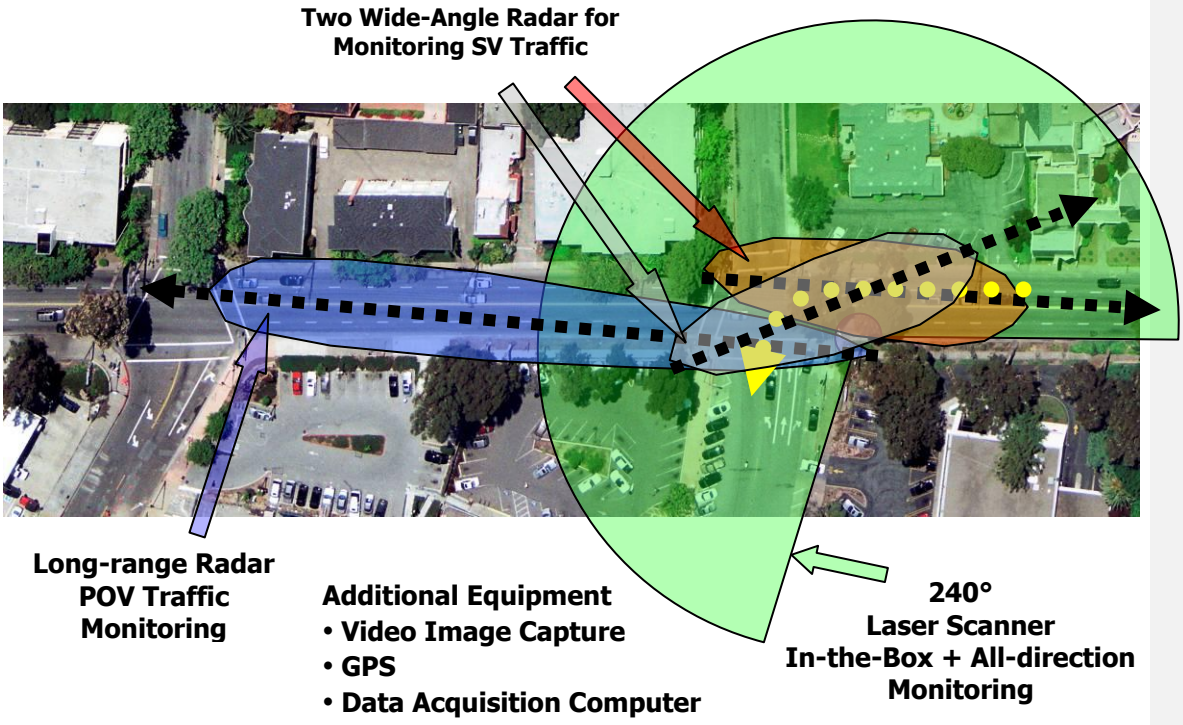

Figure 56, Schematic Diagram of Enhanced Sensor Configuration for Field Observation

Deleted: 54

Sensing Suites

During the data collection period, an industrial computer system will be interfaced withseveral types of sensing devices to record synchronized data with a Global Positioning System incorporated to provide standard time stamps, if necessary.

- SmartMicro Universal Medium Range Radar (UMRR SMS ™ radar) http://www.smartmicro.de/html/umrr_platform.html
  - This UMRR sensor is able to measure range, speed and angle of multiple objects simultaneously. It provides a hardware (antenna) and software selectable field of view, true object separation capability in range and Doppler Stand-alone or network operation is possible.

    - General features:
        - Measurement cycle time selectable between 20ms and 200ms
        - Frequency Band - 24GHz ISM Band
        - UMRR family has multi mode capability and may be switched between pulse- and FMCW based operational modes within one measurement cycle.
        - Operational modes are:
            - Pulsed
            - Pulse-Doppler
            - CW
            - FMCW
- IBEO ALASCA ^{TM} Laser Scanner http://www.ibeo-as.com/english/products_alascaxtsingle_educational.asp
    - A laser pulse with a defined duration is sent and reflected by an object. The reflection from the object is captured by a photodiode and transformed into signals in an optoelectronic circuit. The time interval between the pulse of light being sent and its reflection being received, making due allowance for the speed of light, indicates the distance to the object that reflected the light. By way of the rotating mirror, the laser range finder operates as a scanner, because the mirror deflects each outgoing beam. The mirror's continuous rotation, in conjunction with the pulsing laser, generates a complete environmental profile of the vehicle within the laser scanner's visible range.
    - The IBEO Automobile Sensor split the laser beam into four vertical layers. Distance measurements are taken independently for each of these layers with an aperture angle of 3.2°. This allows any pitching of the vehicle, caused by an uneven surface or driving maneuvers such as braking and accelerating, to be fully compensated.
- Visible light cameras and infrared camera
- EVT-300 Doppler radar, when applicable

Recent Field Work and Instrumentation Suite

A search for candidate intersections was conducted to look for patterns in collisions history where the target scenarios are meaningfully represented. The first intersection, where field data collection was conducted, is at Brighton and San Pablo Ave located in the city of Albany, California. This location is a T-shaped intersection, where the target traffic scenario is the left-turning traffic from southbound San Pablo Ave onto eastbound Brighton.

For each of the candidate location, the work plan is to install a set of traffic monitoring devices for a period of 3-4 weeks, and run the data acquisition system continuously for that duration. After the data collection is completed at one location, the equipment set will be moved to the next location.

Over the course of CICAS and its predecessor, IDS (intersection Decision Support), California PATH have evaluated a number of commercial sensing devices. For the data collection at the candidate intersections, the CICAS team will deploy several devices to capture the traffic movements. The potential list of traffic monitoring devices includes the following:

(1) One industrial computer PC-104 + all cables and connectors
(2) 2-3 CCD cameras + cables
(3) One Wavetronix radar
(4) 1-2 IBEO laser scanner
(5) 3 SmartMicro radar
(6) 4 Numetrics NC-200 ground traffic sensors
(7) One NEC Infrared camera + laptop + cables (depending on availability)
(8) One EVT-300 radar

Figure 57. Placement of Equipment at Candidate Intersection

Deleted: 55

In Figure 55, the positions of various devices are indicated by the color icons shown.

(1) The PC-104 will be ideally placed inside (if space available) in the signal controller cabinet.
(2) Two-three cameras will be placed: One from high-up on one light pole near the intersection to provide overhead view of the intersection box; a second and a third looking at two directions on San Pablo Avenue.
(3) One Wavetronix traffic sensor will be placed at the northwestern corner to monitor the approaching traffic.
(4) One IBEO laser scanner also placed at the northwestern corner to cover the movements of left-turning vehicles and the approaching opposite traffic.
(5) Three SmartMicro radar sensors will be placed at the northeastern, southwestern corner, and northwestern corners to provide coverage of traffic movements.
(6) Four Numetrics surface traffic sensors will be installed on the pavement surface of the two approaching traffic lanes, with one position near (50 feet or 15 m from) the intersection and the other position farther (100 feet or 30 m) from the intersection.
(7) One NEC infrared camera may also be used, which will be placed any appropriate position.
(8) One EVT-300 sensor may also be used, which will be placed any appropriate position.

Note: The use of items (6-8) are optional and only for limited time periods to capture data for comparative evaluation.

Field Data Illustration

This section provides an overview of traffic data collected in recent field observation.

EVT-300

- EVT-300 is a radar sensor, originally developed and marketed for truck on-board collision warning systems. It has been adopted for the CICAS project to be used a target detection and tracking device.
- The sensor is placed at Location A, which is near the northeastern corner of the intersection, facing a southerly direction. The shaded area in Figure 55 roughly represents the field of view of the radar device.
- For roadside installation, it can be mounted on a pole.

Sensor Data – Range and Range Rate

Figure 57 shows the range and range rate detected by EVT-300 in a time span of approximately 5 minutes. On the left chart, the targets are detected with a maximum range of approximately 130 meters. On the right chart, the speed in generally in the range of -15 to +15 m/sec (or 33 mph). Each color icon in the chart represents a data point detected by the radar. Durign the hours when this set of data was obtained, the traffic signal cycle was 108 seconds at this intersection. The time span included in these

charts cotnain 2-plus cycles. Figure 58 is an illuistration of vehcile movements based on the same data set from Fgirue 57.

Figure 58, Sensor Placement near the intersection of Brighton and San Pablo Ave
(Radar Antenna is placed at Location A, which is near the northeastern corner of the intersection, facing a southerly direction)

Deleted: 56

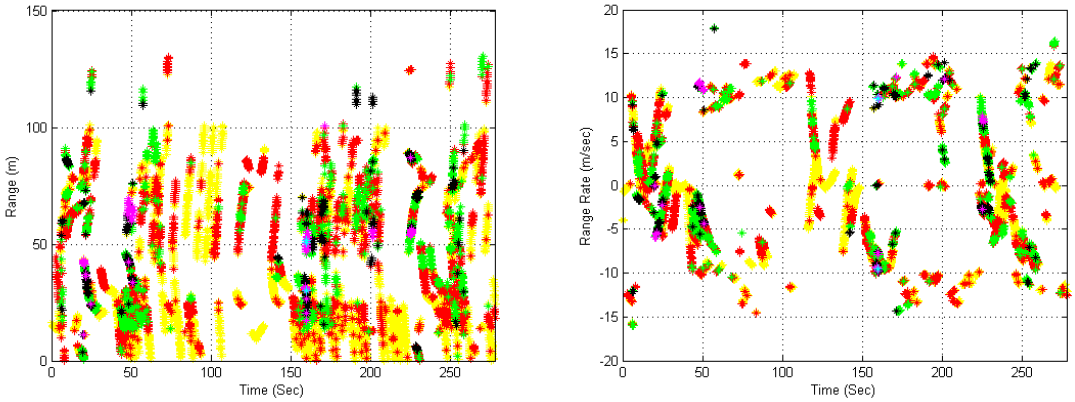

Figure 59, Range and Range Rate of Detected Targets Near the Intersection

Deleted: 57

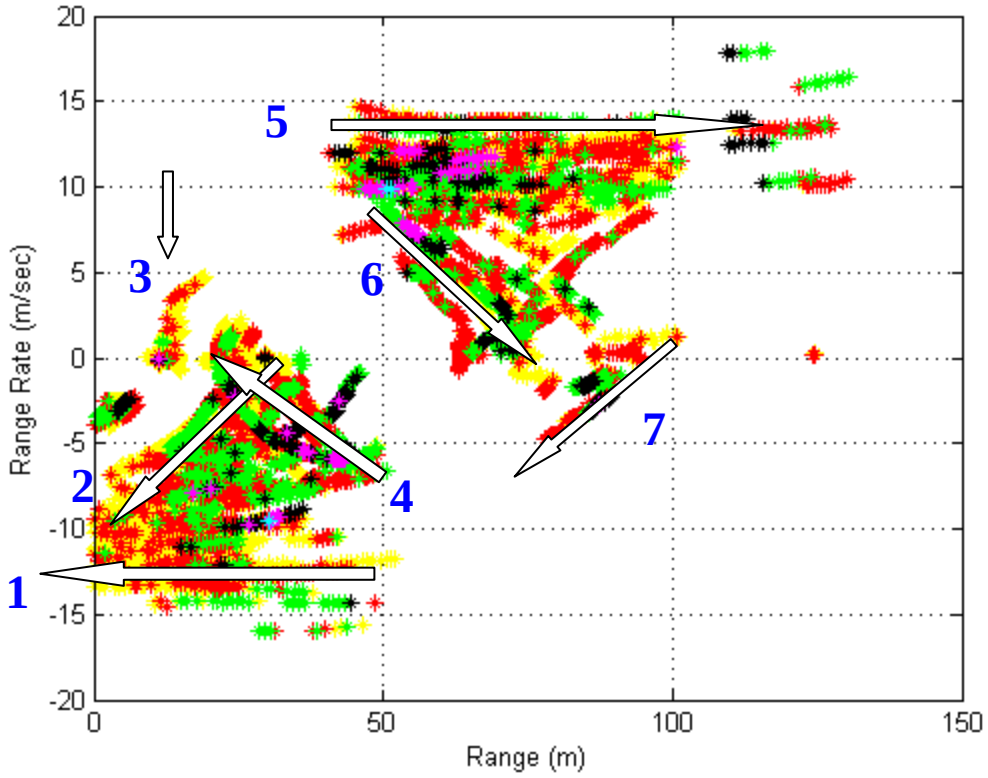

Figure 60, Range Rate versus Range of Detected Targets near the Intersection

Deleted: 58

Numetrics NC-200

- Nu-Metrics NC-200 ™ (*Quixote Technology*) Traffic Sensors is a portable traffic analyzer designed to be placed directly in the traffic lanes to provide traffic data. The NC-100/200 utilizes Vehicle Magnetic Imaging to detect vehicle count, speed and classification.
- Approximate size: 8" X 5 " X 1", 12" X 10" including rubber cover.
- Sensors are placement in two lanes, with two at 50 feet from the near crosswalk and two at 100 feet from the crosswalk. (A and B at 50 feet from crosswalk; C and D at 100 feet from crosswalk, with A and C on the inside lane while B and D on the outside lane)

Speed Data

Figure 59 shows the comparison of speed range at three sensor lcoations, A, B, and C. The data from A and B provides a comparative distribution of vehicle speed on the inside and outside lanes at a distance relatively close (50 feet or 15 meters) to the intersection. Compared to A, Location B has relatively lower speed vehicles and fewer higher speed vehicles, probably due to some vehicles slowing down and preparing for right turns. The speed range at Location C is quite similar to Location A, except for a few vehicles traveling at speeds greater than 40 mph (18 m/sec). A portion of traffic traveling at

higher speeds is expected because of the greater distance from the intersection. The difference is not as significant as anticipated, partially because the sensor locations are within a short block between two signalized intersections. The traffic flow is regulated by the signals and not showing much gain in speed for those vehicles stopped previously at the first intersection.

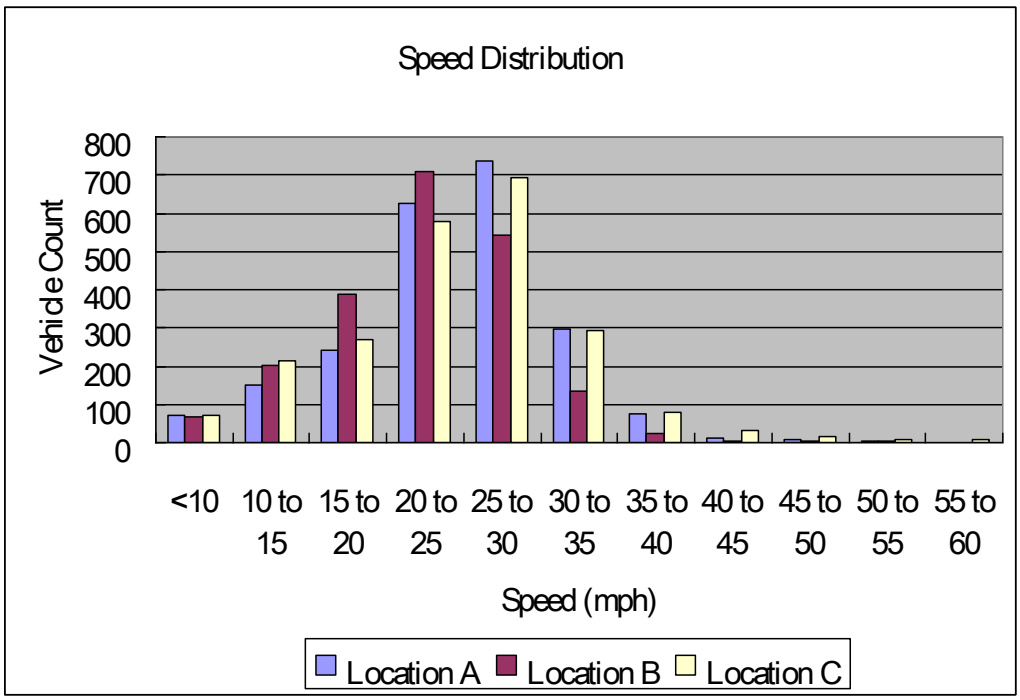

Figure 61, Traffic Speed Distribution and Vehicle Counts

Deleted: 59

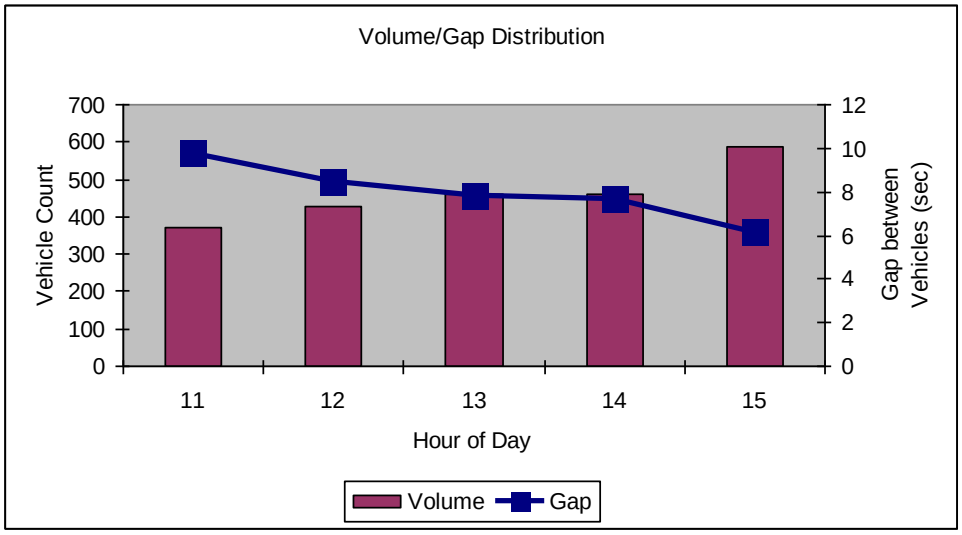

Figure 62, Traffic Gap and Vehicle Counts

Deleted: 60

Traffic Volume and Gap between Vehicles

Figure 60 depicts the traffic volume at different horus during the period of observation. The count is lowests before noon and increased to the highest level in the mid afternoon hour. The average gap in seconds dropped inversely.

Vehicle Length Distribution

Figure 61 below shows the counts of vehicles and the corresponding vehicle lengths according to the preset length parameters for different classifications. Roughly speaking, the category with lengths less than 10 feet (3 meters) long is either motorcycle or vehicles too short to be identified as passenger cars. The great majority (1852 count or 83%) are passenger cars.

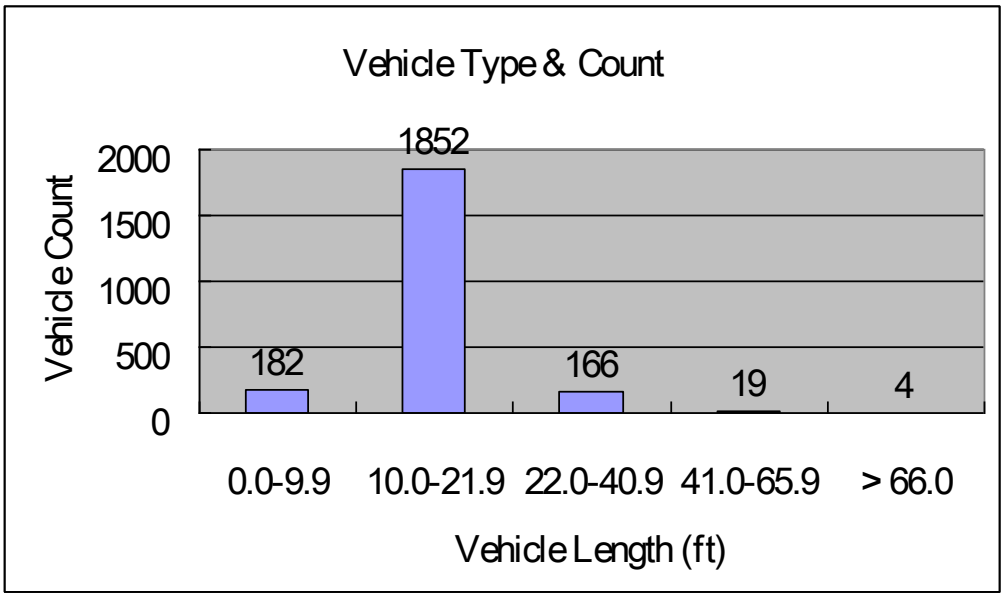

Figure 63, Vehicle Length Distributions by Vehicle Counts and Percentage

Deleted: 61

NEC Infrared Camera

The specifications of NEC Infrared Camera are given in the table below.

Temperature Measurement Range	-20℃ to 100℃
Resolution	0.06℃(at 30℃, 60Hz)
Detector	Uncooled focal plan array (microbolometer)
Spectral Range	8 to 14 μ m
Instantaneous Field of View	1.2mrad
Field of View	21.7° (H) × 16.4° (V)
Thermal Image Pixels	320(H) × 240(V) pixels

Some of images captured in the field are shown in Figure 62. Figure 62 shows vehicles under sunny weather on the left and rainy weather on the right. Figure 63 shows the images of pedestrians in good weather on the left and rainy weather on the right, respectively.

Figure 64. Infrared Images of Vehicles in Sunny and Rainy Weather Conditions

Deleted: 62

Figure 65. Infrared Images of Pedestrians in Sunny and Rainy Weather Conditions

Deleted: 63

Interim Conclusions

Selection criteria based on the potential for intersection safety countermeasures were established. A screening process was conducted to narrow down the choices of candidate sites for data collection and future field operational tests. An enhanced set of sensor suite is being developed and utilized for conducting field observation and for capturing vehicle trajectory data near intersections. The task of enhanced and extended field observation for CICAS is still ongoing. More detailed results of associated data analysis will be reported in the updated version of future reports when they are available. The outcome of this study, in conjunction with the experience and knowledge gathered in the previous phase of related projects will provide valuable inputs for future safety projects.

1.4 Accommodating Other Road Users: Pedestrians

Object detection and tracking is a major research area in computer vision. One of its application areas is traffic scene analysis, which we used for CICAS because cameras are less costly and easier to install than most other sensors, so many are already installed on the roadside, particularly at intersections. Resultant video images are used to estimate traffic flows, to detect vehicles and pedestrians for signal timing, and to track vehicles and pedestrians for safety applications.

For decades various vehicle and pedestrian detection and tracking algorithms have been introduced, (Antonini, Beymer, Ervin, Kamijo, Kim, Koller, Tuzel, Weber, Worrall,) and there are many commercial systems to detect vehicles and pedestrians. Most of the above systems (and also many of other object tracking applications) are based on the *background subtraction* algorithm. It first extracts a static background hypothesis from a sequence of images, then calculates a difference between the background hypothesis and the current image to find foreground objects.

The background subtraction algorithm requires a relatively small computation time and shows robust detection in good illumination conditions. However, it suffers from the problems such as occlusions, the presence of shadows, and sudden illumination changes. Many efforts have been made to solve the occlusion problem, for example by applying a Markov Random Field model (Koller), but the result has not been satisfactory with significant occlusions, as it is an extremely difficult problem to segment occluded objects only by considering only the background subtraction result.

Furthermore, it is difficult to deal with problems such as sudden illumination changes and stopped vehicles. For example, vehicles moving slowly in traffic congestions or stopped vehicles at an intersection will eventually be recognized as background objects. This is because the usually, tracking is determined by linking the center of the object blobs; oftentimes resulting in zigzaggy trajectories.

Another approach is based on feature tracking and grouping (Beymer, Antonini). This is done by extracting and tracking individual *corner features* (Forstner), and grouping them based on their trajectories. The grouping algorithms are applied to full trajectories of the corner features for post-processing applications. Since this method uses a set of long trajectories, segmentation among occluded objects is easier to peform than with background subtraction.

A challenge in applying this algorithm is that it is not always easy to robustly track a corner feature over a long period of time, especially when vehicles turn at intersections or are occluded by other vehicles or pedestrians. In addition, keeping a set of long corner trajectories can be burdensome. For example, when a vehicle is waiting at an intersection for over a minute the corner feature trajectories must be kept for more than 1000 frames. Another challenge is that the grouping algorithms often group nearby vehicles moving together or separate a large vehicle into two because corner features are not evenly distributed over the vehicles in many cases. Most importantly, many applications require

real time detection and tracking, and the computational and storage needs for these algorithms makes direct application not possible.

Finally, there is an approach based on appearance-based vehicle detection (Kim, Koller, Worrall). Kim and Malik introduced a model-based vehicle detection algorithm eventually adopted by the NGSIM (Next Generation SIMulation) traffic modeling effort to generate a large number of long trajectories (Kim). However, such an approach only works on a certain classes of vehicles (e.g., passenger cars and trucks for Kim) from a limited range of view points (for example, a bird-eye view for Kim). In addition, the tracking, based on template matching does not work for intersection videos where vehicles turn, causing their sizes on the image change significantly.

There is also active research on appearance-based pedestrian detection, for example in Tuzel's *Proc. IEEE Conf. Computer vision and Pattern Recognition.* However, many of the reliable detection algorithms require heavy computation and the robust tracking is still a problem. The algorithm we will introduce can effectively combine non-motion-based object detection algorithms or interactive detection results to get higher quality of results and to reduce computation.

We ised an approach to combine the background subtraction and the feature tracking and grouping algorithms. Our algorithm is distinguished from the previous ones that:

- it gives more robust background subtraction result by using a feature tracking result as an additional cue;
- at the same time, it gives a better feature detection and grouping result by using a background subtraction cue;
- it provides a multi-level feature grouping algorithm to deal with various sizes of objects, such as for detecting passenger vehicles and bicycles at the same time; and
- it introduces a dynamic feature grouping which can be applied to realtime applications and produce a high quality of object trajectories from fragmented feature tracks (as opposed to previous work's using nice long trajectories).

An approach of combining background subtraction and feature tracking is also found in Kanhere's *ICPR Volume 1*, where Kanhere et al. used the background subtraction result to estimate the 3D heights of corner features for a low-angle off-axis camera. To estimate the heights it was assumed that the bottom of the background subtraction region is also the bottom of the object. Although such an assumption often fails due to occlusions, promising results were reported from a number of challenging highway video clips because of redundant estimation over multiple frames. The trajectory generation is accomplished by matching frame-by-frame detections, but it does not deliver high quality trajectories. However, such a height estimation algorithm may be incorporated in our framework as an additional cue to further improve the result.

Background Subtraction

A typical background subtraction algorithm applies a Kalman filter (or -blending) to the pixel intensities to find the background (Kilger)Such an approach works well when foreground objects appear infrequently, but when the background is occluded by an object for a significant time, the algorithm begins to fail. Another problem is that Mt is usually generated from Dt by thresholding and applying morphological operators. Such self-feedback can make the filtering unstable. For example, a single detection failure or a sudden illumination change can result in a permanent failure (or a *ghost*), which may even grow in size until it covers up the entire image. Sudden illumination changes commonly occur in many field video images because most video cameras have an auto-iris feature.

Various augmentations have been applied to the background subtraction, for example, to use temporal median. More recently, in *IEEE Intelligent Transportation Sstems Conference,* Batista et al. introduced various augmentations including the use of multi-layer background models and dynamic thresholding. Such augmentations significantly improve the robustness of the system, but the intrinsic problem of self-feedback is still there.

Therefore, we incorporate an external cue (the corner features) to generate more robust Mt. In addition, we made simple modifications to improve the robustness:

- the temporal median approach is combined with the alpha-blending;
- an illumination correction procedure is added to deal with sudden/temporary illumination changes.

After the object blobs are found, we apply an additional validation step to remove the ghosts. We assume that within the entire non-ghost foreground region there exists at least one *valid corner*, i.e., a corner feature which is not found from the background image. For more details on the valid corner, see Feature Tracking and Grouping.

An example result is shown in Figure 1 where the illumination challenge caused by an auto-iris camera. The two white vehicles in the bottom changes the entire scene darker, and it causes significant false alarms. However, the error is minimized by applying the illumination correction.

The resulting object blobs are not the final result but they are used as supplementary cues in the feature tracking and grouping which will be discussed in the next section.

Previous feature grouping work, Antonini and Beymer groups corner features directly into objects using proximity and motion history. Such a single-level grouping is difficult and/or computationally expensive, especially when we deal with objects of different scales (for example, bicycles, passenger cars, and trucks). For instance, the distance between two corner features that belong to the same vehicle can be much larger than the two corner features that belong to two nearby vehicles, which can confuse the grouping algorithm. However, when we apply a sophisticated grouping algorithm to handle this it

will bring in computational burden, particularly when comparing long trajectories of corner features.

To efficiently deal with the problem, we present a multilevel grouping where individual corner features are first grouped into small clusters ("cluster features") then the cluster features are grouped into object-level features.

The grouping is performed on each frame (*dynamic grouping*), as opposed to the previous efforts by Antonini and Beymer, which applied the grouping algorithms on the final tracking results. Therefore, the proposed algorithm can be applied in real time. Note that one of our main goals is to generate a trajectory of an object and directly applying conventional grouping algorithms, such as K-means and the Normalized-Cut (Shi). Frame-by-frame basis will not provide solid trajectories.

The lowest-level corner features are detected by finding the eigenvalues of the local sums of derivatives (Forstner). The corner features are only detected in the *foreground region*, which is estimated by the background subtraction algorithm. The detected corners are tracked by applying cross-correlation template matching on a small image patch (9×9 in our implementation) about the corner. The search window sizes for the match vary from the applications but we first apply a large window (for example, 15 × 15) when the direction/speed of the corner is not known, then a small window (for example, 7 × 7) near the expected position estimated by the previous movement.

The tracked corners are validated with comparison of the background image. Another template matching search is performed on the background image with a small search window size (3 × 3 in our implementation). When a corner has a match in the background image, it is considered invalid and removed. Such invalid corners are often generated by errors by tracking failures such as drifting or errors in estimating the foreground region.

We consider a corner feature 'valid' when it is tracked over a number of (three in our implementation) consecutive frames, does not have a match in the background image, and is neither moving nor picked-up by an existing cluster (see Cluster Tracking). When a corner feature matching fails over a number of (five in our implementation) consecutive frames it is no longer used. Corner features are detected in each and every frame, and those not overlapping with existing ones are subject to tracking.

Note that a feature trajectory can comprise several thousand frames long in traffic video images due to long waiting times at signalized intersections. However, we do not need to keep the whole thousand frames of corner trajectories but just for, say, 20 frames in our implementation because of our dynamic multi-level grouping approach.

The next step is to group the corner tracks into small clusters, with each represented by a circle. An ellipse is also used for an application with single-level grouping. We apply a variation of an Expectation- Maximization (EM) algorithm (Dempster) for the grouping. For each cluster, its expected position and size in the new frame are determined from its current "member" corner features. Then, for each corner feature, its membership in the

current frame is re-determined by comparing its position, history of motion, past trajectory, previous cluster membership, and background-subtraction-blob membership history to their "parent cluster". The position and the size of the cluster are re-determined by applying the updated membership. Then after a small number of iterations (3 in our implementation), the final position and the size of the cluster is determined.

The individual probability distributions were estimated by using a semi-supervised learning procedure. First the algorithm was run with manually-assigned parameters to obtain a reasonable result. In addition, a user interface was developed to correct observably erroneous membership assignments. Thus we obtain a reasonably good grouping dataset with a few manual interventions. Finally, parametric probability distributions were fit to the collected dataset. We used half-Gaussian, gamma, beta, and step (for filtering) distributions to approximate the individual distributions. For example, the positive and negative background-subtraction region history was fit to a pair of beta distributions. Note that the cluster grouping is a hidden-level with no ground truth. Therefore, the parameter, proximity, was manually modified that each cluster maintains a reasonable size.

Finally, the clusters are grouped into objects. The same EM-style algorithm for the corner feature grouping is used for grouping the clusters but with a few differences:

- the shape of an object is fit to an ellipse instead of a circle;
- the position and the shape of the final object is determined not by the positions of the member clusters but by the positions of the member corner features of the member clusters; and
- 3D criteria is used in determining the minimum and maximum sizes of the ellipse. In addition, the trajectories are compared in 3D space (assuming a fixed height, say 0.5 meter in our implementation) to determine the membership.

It is necessary to group clusters based on their 3D trajectories because the disparity among the trajectories can be quite large in the image space (in pixels) due to camera perspective.

We present a result on an intersection video clip of about 80 *s* with vehicles, bicycles, and pedestrians. This video clip is particularly challenging because of occlusions among all road users in additional to having a cluttered background. It also contains some challenging illumination cases as presented. The video clips contain nine vehicles and seven bicycles on the road and many other passing-by pedestrians. Even bicycles on the crossings and backgrounds are included. Our algorithm correctly detected the positions, sizes, and trajectories of all these road users/objects. Most of the pedestrians and bicycles in the cluttered background were also detected by the algorithm with a small number of under-segmentation errors and missed detections. The original and processed video clips can be downloaded at `http://path.berkeley.edu/~zuwhan/ztracker`.

The average processing time was 45ms on a 1.8GHz Pentium Core processor. The algorithm was applied at 10 frames per second. The processing time rather depends on

the scene complexity than the size of the image as the corner tracking takes a significant portion of the total computation.

Our algorithm can be applied to various object detection and tracking applications; therefore, we next present its applications on pedestrian detection and tracking, bicycle detection and classification, and interactive trajectory extraction. Detecting, counting and/or tracking pedestrians can be used for various safety and traffic engineering applications. Since the shape and the size of pedestrians are roughly uniform, we use a single-level grouping with a fixed cluster shape (a vertically ellipse). The video clip from the website referenced earlier contained total 17 pedestrians crossing the road, and there was only one false alarm and one grouping error (shown in the left) where three people walking together were initially detected as two (which was subsequently corrected). There was also one tracking failure ,but the person was also re-detected later.

There are many safety and traffic engineering applications that require a complete set of high-level trajectories, but most of them do not require on-line processing. Therefore, a vehicle tracking system that allows a minimal-level of user-interaction to validate and correct the detection/tracking result can be very useful (Kim). We developed such an interactive vehicle detection system based on our algorithm that allows a user to locate an undetected vehicle, remove a false detection, reposition the center of the vehicle, modify trajectories, and even modify the corner-feature membership for a high-level user.

We used a single-level grouping for the interactive system with a fixed vehicle size. When a user manually detects an object by specifying its position and the direction, a search is performed to find the member features – nearby corner features of the same motion.

1.5 Messages and Maps

Potential Vehicle and Map Add-ons to J2735

To implement CICAS-SLTA may require changes to the MSG_MapFragment message and subordinate structures. One approach may be to define an intersection in terms of approaches and egresses closer to that described in "CICAS-V DSRC Message Descriptions and Examples (January 23, 2008)" than the J2735 Draft 24 specification, which introduces intermediate DF_ApproachesObjects in-between the intersection and paired approaches and egresses. This has in fact been accommodated in Draft 27.

We propose extending the definition of a DF_Approach structure so that it can specify the permissible egresses under different signal conditions. From any intersection entrance and exit geometry, we believe we can compute a plausible trajectory across an intersection, and from the mathematical intersection of two such trajectories we can calculate a zone of conflict where collisions may occur. When two vehicles on different approaches are predicted to be in the same zone of conflict at the same time, an alert can be issued to the driver or drivers.

Addtionally, MSG_VehicleDigest is a new message that describes how an intelligent intersection can inform nearby DSRC-enabled vehicles of the basic position and speed of nearby vehicles. The DF_BasicReport structure describes one vehicle in seven bytes, excluding surrounding tags. It describes vehicles with a granularity of 0.1m relative to the intersection's primary reference point, and 0.2m/sec in speed. The MSG_VehicleDigest message immediately improves the situational awareness of DSRC-enabled vehicles as a received message fills gaps in knowledge originating from non-DSRC enabled vehicles, which have no means to broadcast their dynamics, and from the failed receipt of vehicle-to-vehicle basic safety messages from DSRC-enabled vehicles. The MSG_VehicleDigest message is a pure addition to the proposed J2735 Draft 24 specification; it does not directly change any structure defined in that document.

MSG_MapFragment *(Modifications)*

Data Frame: DF_Intersection *(modifications in italics)*

Use: A complete description of an intersection roadway geometry and its allowed navigational paths (independent of any additional regulatory restrictions that may apply over time or from user classification).

ASN.1 Representation:

```
Intersection ::= SEQUENCE {
    name         DescriptionName OPTIONAL,
    id           IntersectionID,
    refPoint     ReferencePoint,
    config       NodeConfig OPTIONAL,
    approaches   SEQUENCE (SIZE(1..32)) OF Approach,  -- approach
geometries
    egresses     SEQUENCE (SIZE(1..32)) OF Approach,  -- egress
geometries
    …
}
```

Used By: This entry is used directly by one other data structure in this standard, a MSG called MSG_MapFragment <ASN> <XML>. In addition, this item may be used by data structures in other ITS standards.

Data Frame: DF_Approach (modifications in italics)

Use: The Approach data structure is used to bundle related motor vehicle lanes (both reference lanes and computed lanes are described) within the intersection for an Approach or Egress description which is part of an intersection. It also allows expressing information about any barriers found between lanes (medians), other types of lanes (such as train crossings), and information about pedestrian and bicycle lanes or walkways, all of which may cross the described motor vehicle lanes. *When a DF_Approach describes*

the geometry of an entry into an intersection, the fields ending with "egressIDs" explicitly define the egress routes from the intersection which are permitted when the MSG_SPAT (Signal Phase and Timing) message says the corresponding turning maneuver is allowed. From knowledge of the approach geometry and the egress geometry, it is possible to infer a plausible vehicle trajectory through an intersection. The area where two independent trajectories intersect defines a zone of conflict where vehicles on those trajectories are at risk of colliding.

MSG_VehicleDigest *(New)*

Message: MSG_VehicleDigest

Use: An intersection RSU may transmit a MSG_VehicleDigest message to communicate position and speed of as many nearby vehicles as possible to nearby DSRC-enabled vehicles. The message also establishes the primary mechanism for an intelligent intersection with local sensors, such as radar, to communicate the most important information from those sensors to DSRC-enabled vehicles. At the receiving vehicle, the MSG_VehicleDigest can compensate for a missing MSG_BasicSafetyMessage that failed to arrive because of communication range or line-of-sight issues. When local sensors are available, an intersection can report the position and speed of non-DSRC enabled vehicles that would otherwise be invisible to nearby DSRC-enabled vehicles.

ASN.1 Representation:

```
NearbyVehicles ::= SEQUENCE {
        DE    intersectionID IntersectionID,
        DF    vehicleReports      SEQUENCE (SIZE (0..128)) OF
BasicReport,
        DF    intersectionReports SEQUENCE (SIZE (0..128)) OF
BasicReport,
        DF    fusedReports        SEQUENCE (SIZE (0..128)) OF
BasicReport
        …
}
```

Remarks: Reports on individual vehicles are organized into three groups depending on the origin of the data. The intersection constructs vehicleReports entirely from MSG_BasicSafetyMessages broadcast by nearby vehicles, intersectionReports from stationary sensors local to the intersection, and fusedReports from a combination of MSG_BasicSafetyMessages and stationary sensor data. An individual vehicle may be represented simultaneously in any of the groups, or none, at the discretion of the intersection.

Data Frame: DF_BasicReport

Use: A BasicReport structure compactly encapsulates the position and speed of one vehicle at the intersection described in the enclosing MSG_VehicleDigest message.

Position is expressed in a Cartesian co-ordinate system oriented to true north centered on the intersection's primary reference point. Heading is not provided because it can be inferred from position if the position accuracy is sufficient to place the vehicle on a specific intersection approach or egress.

ASN.1 Representation

```
BasicReport ::= SEQUENCE {
    hashID       Integer (0..255),       -- 1-byte non-unique
identifier
    secMark      DSecond                 -- 2-byte ms after current
minute
    east         Integer (-2048, 2047),  -- 12-bits, LSB in units of
0.1m
    north        Integer (-2048, 2047),  -- 12-bits, LSB in units of
0.1m
    speed        Integer (0..255),       -- 1-byte, LSB in units of
0.2m/sec
    ...
}
```

Used By: This entry is used directly by the following 1 other data structure in this standard (record type, descriptive name, ASN.1, and XML name (if present) of each):

```
        DF    MSG_VehicleDigest    <ASN>    <XML>, and
```

Remarks: When the DF_BasicReport structure's position and speed originate in whole or part from a MSG_BasicSafetyMessage, the intersection assigns the low-order byte of the MSG_BasicSafetyMessage structure's temporaryID to the hashID of the DF_BasicReport structure for that vehicle. When the reported vehicle dynamics originate entirely from stationary sensors at the intersection, the intersection assigns its own one-byte tracking number to the hashID of the DF_BasicReport structure for that vehicle. The hashID is NOT guaranteed to be unique within a single MSG_VehicleDigest section. Nonetheless, a receiving vehicle may use it along with the secMark to help associate DF_BasicReports originating from the same vehicle at different times, and to identify when a DF_BasicReport structure is redundant with the original and more complete MSG_BasicSafetyMessage available from a direct vehicle-to-vehicle communication.

The secMark is the number of milliseconds after the start of the current minute. When the intersection constructs a DF_BasicReport structure in whole or part from the MSG_BasicSafetyMessage, it sets the secMark to the sectMark of the original DF_BasicReport. In the case of a DF_BasicReport structure constructed only from data using sensors local to the intersection, the intersection assigns a secMark corresponding to the reported position and speed estimates.

ASN.1 Representation:

```
Approach ::= SEQUENCE {
    name               DescriptiveName OPTIONAL,
    id                 ApproachNumber,
    drivingLanes       SEQUENCE (SIZE(0..32)) OF VehicleReferenceLane,
    computedLanes      SEQUENCE (SIZE(0..32)) OF VehicleComputedLane
OPTIONAL,
    trainsAndBuses     SEQUENCE (SIZE(0..32)) OF SpecialLane OPTIONAL,
    barriers           SEQUENCE (SIZE(0..32)) OF BarrierLane OPTIONAL,
    crosswalks         SEQUENCE (SIZE(0..32)) OF CrosswalkLane OPTIONAL,
    leftTurnEgressIDs      SEQUENCE (SIZE(0..3)) OF  -- IDs of
permissible
                           ApproachNumber OPTIONAL,  --   left turn
egresses
    softLeftEgressIDs      SEQUENCE (SIZE(0..3)) OF  -- IDs of
permissible
                           ApproachNumber OPTIONAL,  --   soft-left
egresses
    straightAheadEgressIDs SEQUENCE (SIZE(0..3)) OF  -- IDs of
permissible
                           ApproachNumber OPTIONAL,  --   straight
egresses
    softRightEgressIDs     SEQUENCE (SIZE(0..3)) OF  -- IDs of
permissible
                           ApproachNumber OPTIONAL,  --   soft-right
egresses
    rightTurnEgressIDs     SEQUENCE (SIZE(0..3)) OF  -- IDs of
permissible
                           ApproachNumber OPTIONAL,  --   right-turn
egresses
    …
}
```

Used By: This entry is used directly by one other data structure in this standard, a DF called *DF_Intersection* <ASN> <XML>. In addition, this item may be used by data structures in other ITS standards.

Remark: Note that the integer value give to each described item (lane, barrier, crosswalk, etc.) is used in other messages and data frames to refer to that object within the context of this globally unique intersection that this data frame is used in.

Data Frame: DF_ApproachesObject *(removed)*

1.6 CICAS RFS Intelligent Intersection Build and Status

One of the several tasks included under CICAS-SLTA project includes controlled experiments of driver gap acceptance when making left turns against opposing traffic. A great deal of work has been done to instrument and test both the experimental intersection and primary subject vehicle at UC Berkeley's Richmond Field Station such that precise

data can be collected and carefully evaluated. This effort can be grouped into three areas: 1) Development of the 1998 Ford Taurus instrumented subject vehicle, 2) Development of the experimental intersection, and 3) Development of the vehicle to roadside communication link for experimental control and data gathering.

The effort to develop the Taurus includes integration and testing of a large number of independent and interdependent sensor systems. There has been additional software development to provide data "fusion" of some of these sensors to give a refined estimate of the vehicle's position. This was a very demanding task, and required unexpected effort to get consistent data recorded in the vehicle and communicated to the roadside.

The effort to develop the experimental intersection required integration of a new radar system from SMS, light beam sensors, and vehicle-reported position information. An additional software process to fuse these inputs (the "roadside fusion" process) is used to provide refined estimates of the SV's and POV(s)' positions and velocities. The experiments will also benefit from independently developed software that reports traffic signal states and allows requests for signal state changes. Additional work was done to develop a more precise description of the intersection and its conflict zones, including a software framework to provide this information to the experimental control framework or to experimental warning algorithms. Unexpected effort was required to keep the SMS radar system functioning, and development of the roadside fusion algorithm was delayed by slow progress with the SV and with vehicle-roadside communication.

The development of the vehicle to roadside communication uses DSRC band wireless signals for sending and receiving both experimental protocol and vehicle sensor data. This effort integrated DSRC hardware from different manufacturers to maintain active communication with all experimental vehicles (both SV and POVs) and the roadside for the duration of the experiment. The communication framework is reliable and responsive enough to allow close coordination of vehicle start times for the experiments.

These three efforts converged for data collection trials in mid-June 2009 that proved the communication model and gathered data from both the roadside and vehicle (SV and POV) sensor systems. The data from these trials and smaller subsequent ones was used to complete initial development of the roadside fusion algorithm. The June trials also provided a first test of the communication links that underlie experimental control and the software framework for getting experimental parameters to all vehicles in a timely manner and displaying them to the driver(s).

Work after the June data collection included completion of the initial version of the roadside fusion algorithm as well as initial development of a data management and analysis framework in MATLAB so experimental data can be effectively analyzed. Analysis and debugging of this data also lead to refinement of some of the raw sensor inputs in the Taurus SV.

Instrumentation of the Taurus Subject Vehicle

The Taurus test subject vehicle required significant development to support data collection experiments. The goal was to have a vehicle with many different forms of sensors to detect both the driver's actions and precise information about the position, velocity and status of various vehicle systems. To this end, the car has been outfitted with two general purpose computers, multiple cameras, and a wide variety of sensor systems. These can be broadly broken down into video, network, serial and A/D inputs to the primary data collection computer.

The primary data recording computer is a PC104 industrial CPU running the Linux operating system. It has special purpose cards to allow particular kinds of inputs: a card for 12 A/D inputs, a card for 8 serial inputs, a card for 4 ethernet inputs, a card for 4 MPEG video inputs, and a timer/counter card. It also has video/audio output for a driver interface unit. A description of the sensors these cards enable is below. A separate laptop computer is used to run a digital video-based eyetracker system using two dash-mounted cameras. A separate communication CPU enables WiFi, DSRC, and wired networking and serves as the communication hub between the roadside and the vehicle.

A/D and Timer/Counter Inputs

The A/D card reads voltages from hardware taps of the vehicle's brake light, left and right turn signal lamps and reverse lamp, a potentiometer to measure steering angle, and taps to the ignition system (on/off), cruise control (on/off), throttle valve position, and 4 radar-based side detectors. An additional timer/counter card reads raw inputs from the ABS system to determine accurate vehicle wheel speed. Each of these raw voltage inputs is conditioned with additional circuitry to insure proper values could be read. Of these 13 inputs, all but the four radar-based side detectors are fully enabled and collect useful data.

Serial Inputs

The serial input card receives separate inputs from four EVT/300 radar systems mounted at each corner of the vehicle, as well as a roof-mounted GPS unit, accelerometer, gyro, and a SafeTrac lane-tracking system. Each of these inputs requires additional interface software to manage the serial connection and record the inputs to a central inter-process communication system. Of these inputs, all but the four radar units are fully enabled and collect useful data.

Video Inputs

The video card collects up to four video inputs from among four cameras and the output to the driver video display. These inputs may include forward facing cameras in the (top) center, left and right of the windshield, a rear-facing camera from the bottom center of the rear window. All software to record and save the video data is fully enabled and collecting useful data. An additional driver-facing camera from by the rear-view mirror may be installed in the future.

Network I/O

The network I/O includes connections to the eyetracker computer, which provides synchronization data to the primary data recording CPU, the communications device (a Savari On Board Unit (SOBU)) that handles DSRC communication and manages the experimental control logic, as well as additional network connections for experimenter's laptops and a WiFi interface. There is additional work to be done to integrate and calibrate the eyetracker computer into the overall system.

Position-Calculation Software

In addition to reading and recording the data from all the sensors and network devices, the primary computer also runs an additional process to give a higher-accuracy sense of the vehicle's position and heading. This process reads the current values of the GPS system, gyro, and speed sensor and computes a more accurate vehicle position than is possible with just GPS. This revised position and heading are amongst the values reported by the vehicle to the roadside.

Instrumentation for the POVs

The experimental protocols call for at least two additional cars to act as POVs. For the sake of these experiments, these cars will be only lightly instrumented with portable equipment. This will consist of a Savari OBU (SOBU) similar to that in the SV, incorporating a 1 Hz GPS input and DSRC communication link to the roadside. The SOBU's WiFi interface will be used connect to a portable wireless display which will display experimental control information to the POV driver. Both the SOBU and wireless display are portable and can be installed in any available vehicle. While any vehicle might be used as a POV, for the sake of these experiments, the POVs will be a pair of Nissan Altimas for all experiments to reduce variability in POV performance (acceleration and braking).

Instrumentation of the Intersection

The Richmond Field Station experimental intersection has many possible equipment configurations and a large number of different sensor systems. For the purposes of the CICAS SLTA experiments, only a couple of sensor systems will be used in addition to network links to the traffic signal controller and DSRC communication systems. The primary data recording computer on the roadside is a PC104 industrial computer running Linux.

Inputs to the Roadside CPU

For this project, the roadside computer reads four A/D inputs from light beam sensors – one on each edge of the intersection box. These provide ground-truth signals for when vehicles cross into or out of the intersection box. The rest of the inputs to the roadside CPU come over its network interface. It receives ethernet input from an SMS radar system that has radar heads covering all the approaches to the intersection and the intersection itself. The network interface is also used to interact with the traffic signal controller, receiving countdown timers for the signal phases and sending requests for specific signal phases and timing. The traffic signal controller enabling these capabilities is a 2070 controller running custom-developed software.

The roadside computer also receives updates of position and heading from the vehicles in the experiment. These updates are limited to GPS data from the POVs, but for the SV include GPS, revised position information, speed, and status information for the brake and turn signals. The roadside computer gets the updates over a DSRC radio link provided by a networked roadside unit. There are at least two roadside DSRC units that can be used for the CICAS experimental work – both a Kapsch/Technocom MCNU and a Savari RSU.

All inputs to the roadside CPU, including the RSUs, are fully enabled and are collecting useful data.

Roadside Fusion for Tracking Vehicles

One of the primary jobs of the roadside computer is to record all these inputs for later analysis. It also uses this data to compute a more precise estimate of the vehicles positions, velocities, and associated data. This more precise estimate will be used in both experimental analysis and for experimental control. This software is now in its initial version, and will be refined and further integrated as more data is collected.

Communication Framework and Experimental Control

While each part of the communication framework has been described above, some more explanation of how it all functions and how it enables experimental control is helpful. Note that all the communication processes have been through initial development and have been tested in several rounds of data collection. As additional data is collected, the experimental control applications will be refined, but the necessary infrastructure is all built and running.

POV Communication

The POVs are set up with a portable SOBU and wireless display device (such as an iPod Touch or netbook). The SOBU communicates over DSRC to the roadside, sending and receiving messages from/to the roadside DSRC computer (either a Kapsch/Technocom MCNU or Savari RSU). The SOBU runs several additional software processes to cache the messages and experimental data as well as a simple webserver which enables a web application n the wireless display device. All messages sent and received by the SOBU use UDP broadcast protocol over the DSRC band, with incoming messages on a separate port from the outgoing messages. Separate processes send and receive messages.

The SOBU has a GPS input and it gathers the GPS data and sends it out over DSRC using one of two message types generated by the POV. The second message it generates is the system's current idea of the experimental state. This experimental state is derived from information messages received by the SOBU from the roadside as well as data entered with interaction from the POV driver. The driver typically only indicates "ready" or "abort" before the trial begins, though the web interface may be extended for the driver to indicate more about the POVs starting position or other characteristics.

The POV receives primarily the roadside CPU's experimental state data, which may include trial number, condition variables, etc. This message is generated on the roadside data collection computer, forwarded over UDP unicast to the roadside DSRC computer (MCNU or Savari RSU), and then UDP broadcast on the DSRC band to the POVs (and SV).

SV Communication

In contrast to the POV, communication and data messages are more complicated in the SV. In this case, the display is controlled by the in-vehicle data recording computer, which also is responsible for gathering vehicle data to report back to the roadside. The SOBU is responsible for maintaining experimental control information (what trial, what conditions, etc) and running the webserver that is used for the experimental control application on the display. The test subject does not interact with the display – that role is the responsibility of the experimenter, who sits in the rear of the car.

As with the POV, the SV generates and broadcasts two messages. The first puts together revised position and heading, raw GPS, speed, turn signal and brake status from values

on the in-vehicle CPU. This message is sent over UDP unicast on a wired link from the in-vehicle CPU to the SOBU, which forwards it over UDP broadcast on the DSRC band to the roadside DSRC computer (MCNU or Savari RSU). The second SV message is the web application's information about the experimental control state of the SV, which is generated on the SOBU and sent over UDP broadcast to the roadside. The experimental control state is derived from information received from the roadside as well as the experimenter's input to the web application.

The SV receives two messages from the roadside, both UDP broadcast from the roadside DSRC computer (MCNU or Savari RSU) after being generated on the roadside data recording computer and forwarded to the DSRC computer. These messages are the experimental control and the current traffic signal status.

By inference then, both the SOBU and the in-vehicle data recording computer send and receive messages, so there are bi-directional links between the data recording computer and the SOBU(wired), and the SOBU and the roadside (wireless DSRC). These all maintain separate UDP ports to avoid confusion between messages meant for the SOBU but not the data recording computer and vice-versa.

Roadside Communication

The roadside data collection CPU acts as a primary clearinghouse for all the vehicle's incoming messages. It maintains the final status of the experimental control data, and generates the final "start" messages for each vehicle after determining that all vehicles are in a consistent state (they are ready, on the right trial with the right conditions, etc). The roadside CPU is also responsible for generating timing messages based on data from the traffic signal controller.

The messages generated by the roadside CPU are sent on a wired link by UDP unicast to the roadside DSRC computer (MCNU or Savari RSU). The DSRC computer the forwards them on UDP broadcast over the DSRC band. The DSRC computer is also receiving the messages broadcast by the vehicles and sending them by UDP unicast on the wired link to the roadside data recording computer.

SECTION 2. TRAFFIC SIGNAL ADAPTATION CONOPS

The concept of operations (ConOps) described here addresses the Cooperative Intersection Collision Avoidance System – Traffic Signal Adaptation (CICAS-TSA). The ConOps serves as a foundation document, describing CICAS-TSA for the system to be uniformly understood by researchers, project participants and stakeholders alike. It serves three functions: (i) as an outreach tool, (ii) as a 'living document' to identify and capture researchers and stakeholders (sponsors, deployers and users) interests and concerns, and (iii) from a project systems engineering perspective, as input to the sequence of steps necessary to conduct the project. As generally illustrated in Figure 72, the position of the ConOps is shown in the context of the entire process or Vee-Diagram in the following sequence (with associated project tasks given parenthetically):

- A detailed ConOps (Task 3 *System Design*)
- System Requirements (Task 2 *Research*)
- System Design (Task 2 *Research* & Task 3 *System Design*)
- Implementation and Testing (Task 2 *Research*, Task 3 *System Design* and Task 4 *System Development and Prototyping*)
- System Verification (Task 4 *System Development and Prototyping*)
- Field Operational Test (FOT) (Task 5 *FOT Planning* and Task 6 *FOT*)

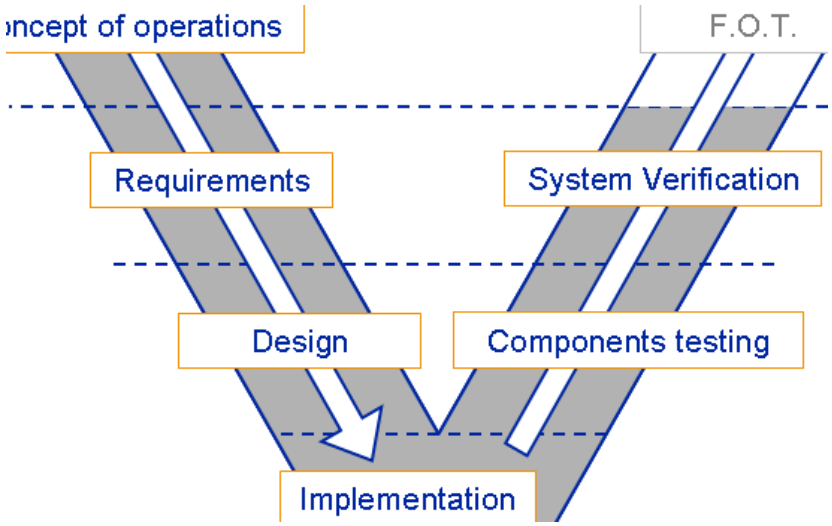

Figure 66. Context of ConOps: The Systems Engineering Vee Diagram

Deleted: 72

The ConOps provides a detailed narrative description of the CICAS-TSA System, not a narrative of the program or research plan. Moreover, the ConOps does not describe requirements, or the testing or verification plan. The ConOps, however, develops the

foundation for these. Following the definition of scope, purpose and goals and objectives, the bulk of the ConOps does so with the following organization:

- System Definition
- Assumptions and Constraints
- Operational Description (per stakeholder category)
- Operational Scenarios (for normal as well as difficult or failure scenarios)

Purpose of CICAS-TSA

Crossing Path (CP) collisions at traffic signal-controlled intersections account for about 9% of all police-reported collisions in the United States. Among all signalized CP collisions, Straight Crossing Path (SCP) and Left Turn Across Path (LTAP) of the Subject Vehicle (SV) account for 61% and 25% of CP crashes, respectively. These movements are the two most prevalent among those maneuvers leading to crashes. Red-light-running (RLR) is a major cause of these. Notable, of those signalized CP collisions that included charged violations, 44% of collisions with SV performing straight pre-crash maneuvers and 4% of collisions with SV performing LTAP pre-crash maneuvers were caused by the SV running a red light. An additional 57% of collisions with SV performing LTAP pre-crash movements were caused by SV's failure to yield right-of-way (ROW). (Najm, Smith and Smith, 2001).

Under the CICAS program, while the CICAS-SLTA (Signalized Left Turn Assistance) system intends to reduce Left Turn Across Path/Opposite Direction (LTAP/OD) collisions due to failure to yield ROW by providing the driver with better information for making safer turning judgment, the CICAS-TSA is to reduce intersection collisions due to red-light-running (RLR). With the CICAS-TSA countermeasure, we focus on the dynamic extension of the red clearance interval between conflicting traffic flows; this is also called the all-red interval. Importantly, the CICAS-TSA countermeasure addresses the two most dominated pre-crash maneuvers: signal controlled SCP and LTAP. This approach prevents the release of any conflicting vehicles to the intersection when a red-light-running incident occurs with high likelihood of collision, and therefore reduces the possibility of such a collision that might otherwise occur.

Goal and Objective

The goal of CICAS-TSA is to reduce frequency and severity of intersection crashes due to red-light-running. The objective is to enable the signal systems to detect a red-light-running incident, to estimate the likelihood of crash, and to dynamically extend the red clearance phase to delay the entering time of signal controlled conflict vehicles and therefore to avoid the potential crash.

System Definition

CICAS-TSA is a system that extends the red clearance interval when the probability for a collision is high. It is designed to prevent RLR related collisions, where the SV runs a red light and cuts across the path of a Principle Other Vehicle (POV) entering the intersection from a conflicting approach at the same time.

The CICAS-TSA countermeasure is built upon the existing traffic signal control system, and it includes the following functional elements:

- Detection to identify the locations and speeds of vehicles approaching an intersection;
- Traffic signal control with additional functions that adjust signal timing for TSA purposes;
- A traffic signal interface allowing the TSA processor to obtain real-time signal status from the signal controller and to communicate the message of adjusted timing to the signal controller for execution;
- A TSA processor processing the detection and signal status information, estimating the projected vehicle trajectories and predicting the probability of red-light-running, assessing the likelihood of collision, then determining when to trigger the signal controller and how to adjust its timing; and
- Communication links that connect all TSA functional elements.

The detection functional element of CICAS-TSA will be designed to best use the infrastructures that either exist already or should be available in the near future. The opportunities for cooperation between CICAS-TSA and the other CICAS projects will also be carefully considered.

Assumptions and Constraints

- CICAS-TSA is assumed to primarily use both vehicle-based and infrastructure-based sensors to detect trajectories of approaching vehicles. It is likely that no single sensor will be able to provide all the needed measurements. The specific selection of sensors is TBD. Over time, it should become possible to augment the infrastructure-based sensor information with information from in-vehicle sensors, communicated to the intersections' infrastructure using DSRC.
- CICAS-TSA should be able to operate under normal driving conditions (full range of weather and illumination, except for extreme environmental or physical conditions, in which driving itself becomes hazardous or TSA can no longer properly function).

- CICAS-TSA should be applicable for the full range of traffic speeds and densities that would be expected at signal controlled intersections.

- CICAS-TSA parameter values will not be uniform across all intersections. They will be adjusted to match the needs of individual intersections, based on their geometry, signal phasing, and traffic volumes and speeds. These values are TBD, based on the results of the research.

Operational Descriptions, by Stakeholder Interest

The CICAS-TSA system operation and its impacts can be viewed from microscopic through macroscopic perspectives, progressing from the direct encounter with the system by the road user, through the operation of the system by the infrastructure operator and up to the State or Federal governments, who may see a public safety impact. We will distinguish the system end users, or users who will interact directly with the system and the general road users who will benefit from its deployment, from the operators who will be involved in the system selection and operation. The system is described from each identified user group's perspective.

Intersection Users

End users will be vehicle drivers. The operations of the CICAS-TSA system will change the situation of driver approaching and within the intersection and the vehicle driver will experience the safety effects of the change, but should not be conscious of any perceptible change in driving conditions. It is important that intersection users accept the system and not change their driving behavior to try to take advantage of its features.

System Operators

State DOTs and Local Jurisdictions

State DOTs and Local Jurisdictions have a two-fold function and use of CICAS-TSA, as traffic intersection operators and being responsible for public safety. In the operations role, the State and local DOTs are concerned about safety as well as capacity and will want the CICAS-TSA system to provide demonstrable safety benefits and little or no reduction in intersection throughput.

The operator will also be concerned about installation and operations and maintenance burdens, to include training of personnel for this equipment. It will be more likely to accept some safety benefit if other operational benefits such as increased capacity could also be realized. It will also likely accept the CICAS-TSA system if it can be perceived as an overlay and does not substantially change or revamp legacy systems and operations.

From the more macroscopic perspective of enhanced public safety, systematic safety benefits will motivate the agency to accept the system.

Intersection Equipment Suppliers

To the extent that CICAS-TSA provides safety benefits and enhances the already offered products, it will be accepted. The TSA equipment suite should be well understood by the traffic control device industry and interface easily with the current products so that it can be accepted and offered to public agency customers at modest cost and risk. Note that the CICAS-TSA countermeasure of dynamic extension of the red clearance interval does not change the stop-go decision of a red-light runner but rather prevents a potential crash due to red-light-running. Therefore, the system in nature can co-exist with red-light cameras (RLC).

US DOT

Any acceptable, deployable system that has demonstrated public safety benefits, particularly those that show dramatic benefits compared to costs and are accepted by the aforementioned stakeholders will be of interest, which is why US DOT is currently keenly interested in CICAS-TSA.

CICAS-TSA Operational Scenarios

Normal Operational Scenarios

Baseline Scenario: SCP red-light-running

If a SV approaches an intersection on a through lane during the green and yellow phase and then enters the intersection on the red phase, the SV may pose a hazardous condition under which the SV may collide with an POV entering the intersection either from a lateral leg or making a left turn from the opposite direction (with protected left turn phase). The CICAS-TSA system tracks the movement of the SV and detects whether the SV is likely to run the red light. The system also identifies the POV, estimates the likelihood of collision with the SV, and predicts the time-to-collision. The CICAS-TSA system then may trigger the traffic signal controller to extend (or insert) the red clearance phase if the likelihood of a collision is high.

Figure 73 illustrates the CICAS-TSA operation under the baseline scenario. (The green vehicle represents the SV, while the white and gray vehicles represent POVs respectively from the lateral direction and opposing protected left-turn lane).

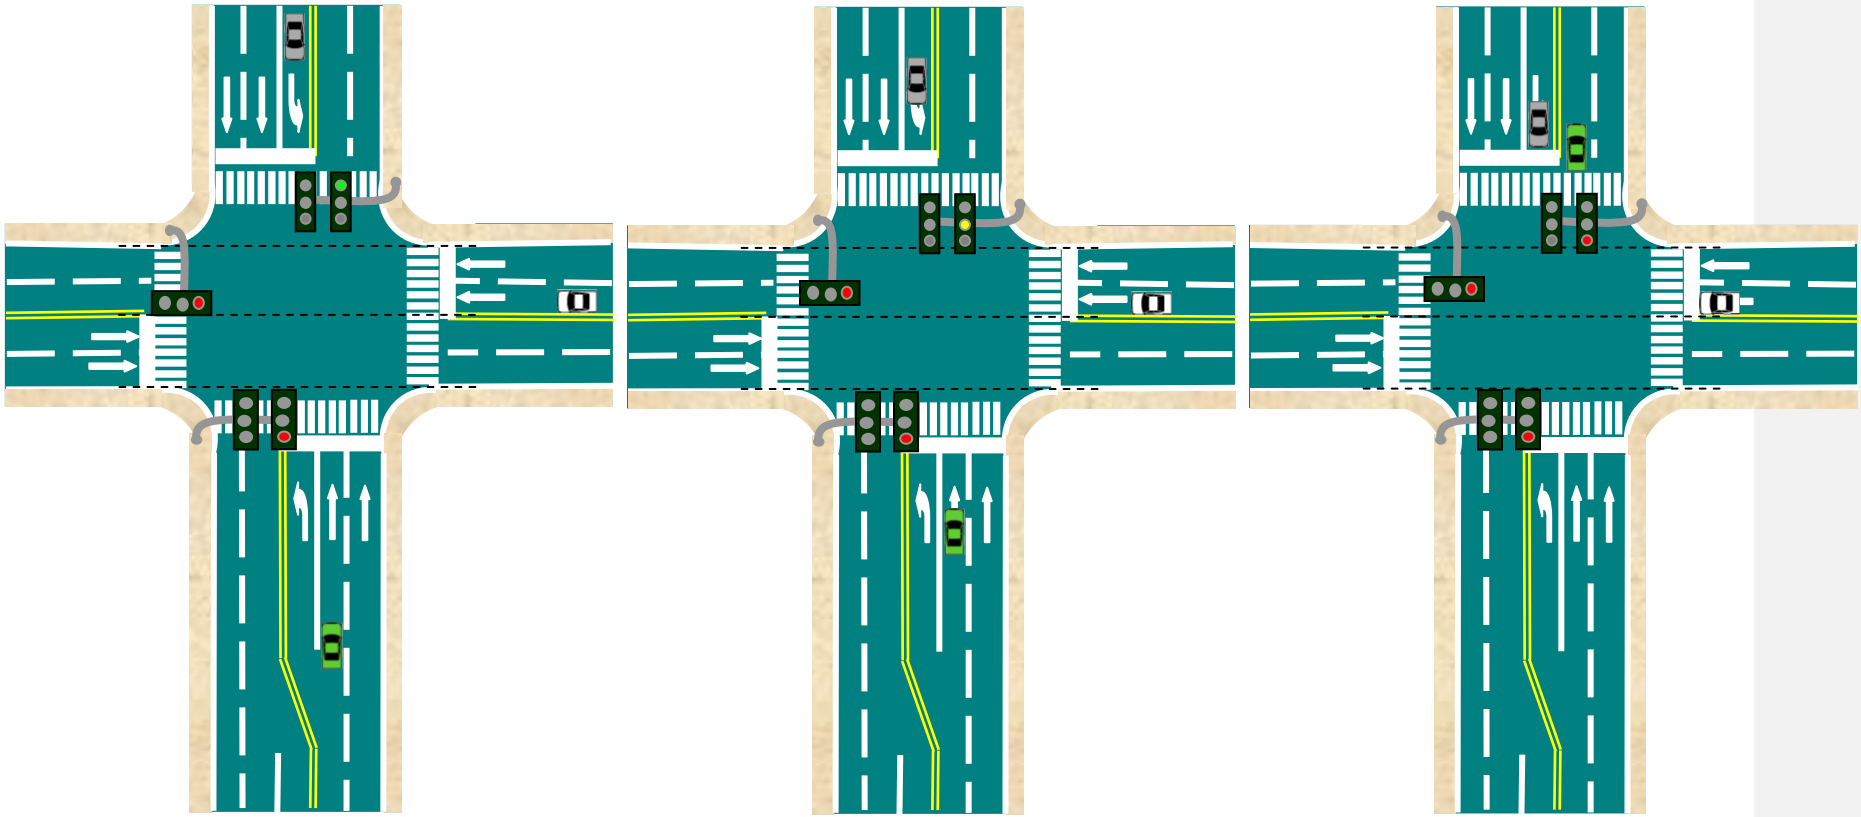

a) TSA system identifies SV (green vehicle) and POV(white or gray vehicle)

b) TSA system

- estimates the probability of SV running red-light
- estimates the likelihood of collisions with POV
- estimates desired red clearance interval

c) TSA system inserts or extends red clearance phase if likelihood of collision is high, and the SV clears conflict zone during red clearance phase

Figure 67. The TSA Baseline Scenario (not drawn to scale)

Deleted:

Variations of the Baseline Scenario

SCP red-light-running in conflict with a left-turning POV which is stranded inside an intersection

If the driver of the POV does not find an acceptable gap during the green or yellow phase, and starts the left-turn movement at the onset of the red phase, the POV might be on a collision path with an SCP red-light-running vehicle from the opposite direction. The CICAS-TSA system identifies the presence of the stranded POV that is stopped inside the intersection and is about to make a left turn. It tracks the movement of the SV and determines the likelihood of its running the red light. The CICAS-TSA system then may trigger the traffic signal controller to extend (or insert) the red clearance phase if the likelihood of the SV running the red light is high, in order to allow both the red-light runner and the stranded POV to clear the intersection safely.

Figure 74 below illustrates the CICAS-TSA operation under this scenario. (The green vehicle represents the SV and the gray vehicle represents the POV).

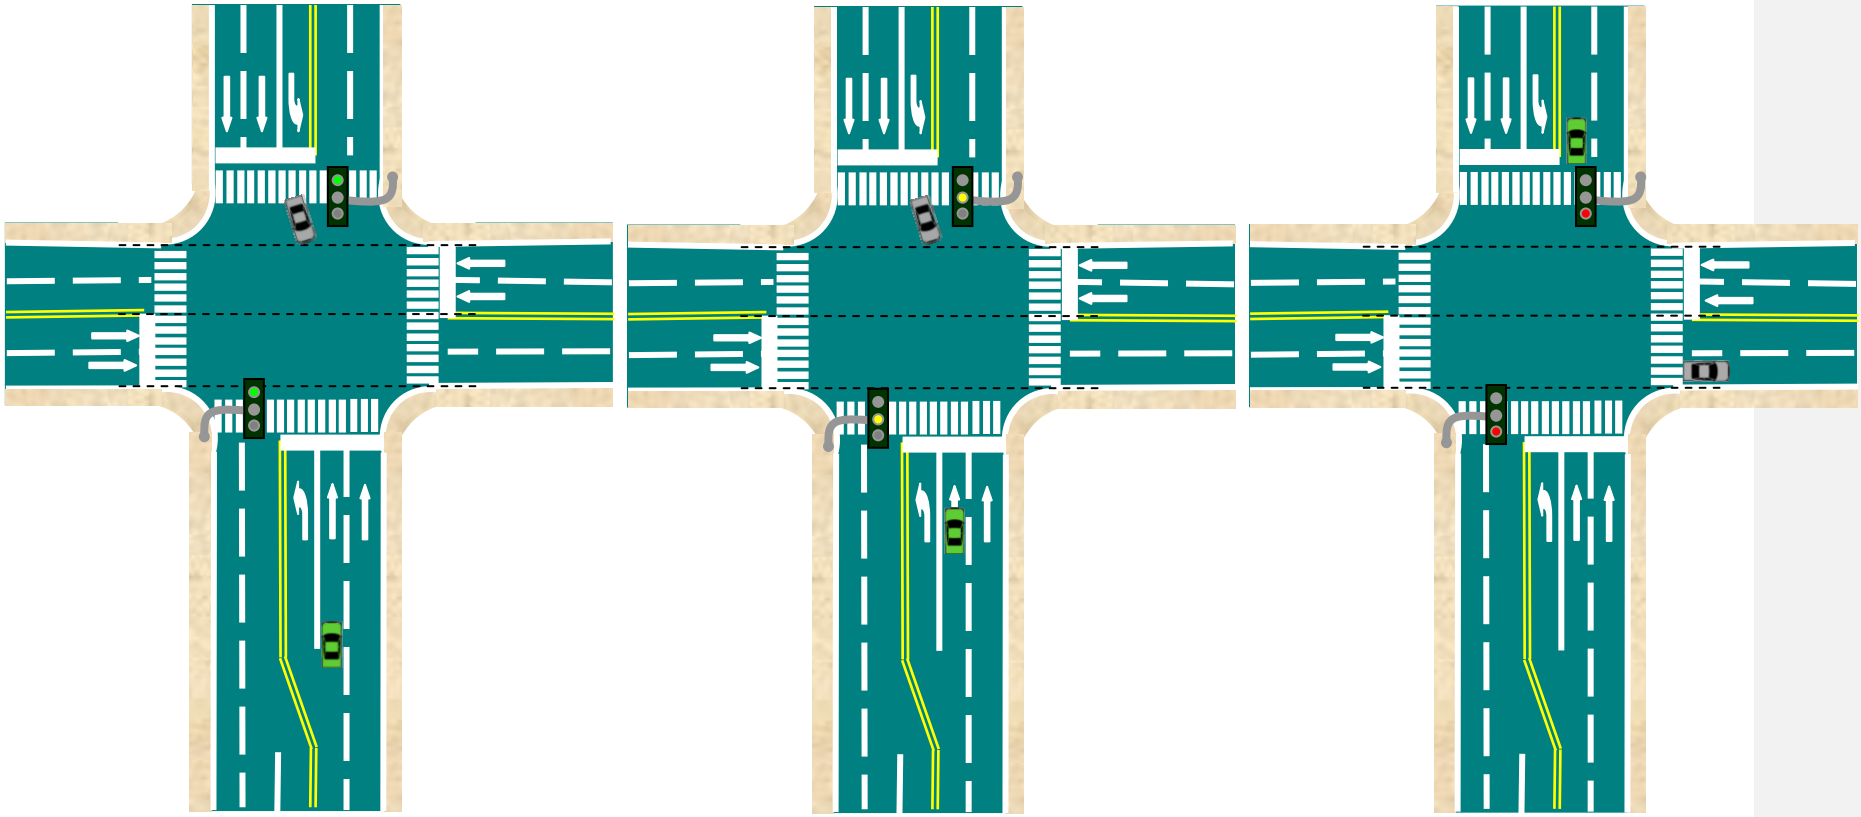

a) TSA system identifies POV (the gray vehicle) and SV (the green vehicle)

b) TSA system

- estimates the probability of SV running red-light
- estimates desired red clearance interval

c) TSA system inserts or extends red clearance phase if likelihood of SV running red-light is high. Both SV and POV clear the intersection during red clearance phase

Figure 68, SCP Red-Light-Running in conflict with a Stranded Left-Turning POV (not drawn to scale)

Deleted:

LTAP red-light-running from dedicated left-turn lane

The left-turn movement is either protected or permissive. Note that the case in which the left-turn vehicle is at the intersection when the signal changes to red is not considered as red-light-running and is excluded from this scenario.

Under this scenario, the operation of the CICAS-TSA system will be similar to that for the baseline scenario, as the system will be designed to identify the lane on which the SV (left-turn vehicle) is traveling so that it can associate SV movement with the correct signal phase and potential conflicts.

Extended Scenarios Enabled by Future Applications of Technologies

- Each vehicle communicates their class, speed and location to the intersection to enable the intersection to better use timing strategies and reduce potential conflicts to provide safer service.

- The intersection communicates via DSRC its state map data and/or driver advisories to the approaching vehicles to augment information available in the vehicles, such that in-vehicle advisories augment the TSA system.

CICAS-SLTA and CICAS-TSA Interactions

There are two cases under which the CICAS-SLTA and CICAS-TSA system will interact:

- The SLTA SV is stranded in the intersection when the signal is about to turn red and the TSA SV is going to run a red light from the opposite direction.

- The SLTA SV and TSA SV is the same vehicle. In other wards, the left turn vehicle is predicted to run a red light from the dedicated left turn lane.

The CICAS-TSA and CICAS-SLTA system will designed to work collaboratively therefore no contradictory information to drivers will be given. We will carefully examine how the two systems should interact and how the information flow regarding the signal phase change conditions will be integrated with the SLTA system logic.

System Failure Scenarios

The TSA system will be designed to handle typical failures. The design of the TSA system should consider the following failure conditions:

Typical System Failures for CICAS-TSA

Sensors) failure - Problems with detecting approaching vehicles (may be difficult to identify)

- Inaccurate speed or position measurements
- Partial loss of detection (limited locations)
- Missed detection

Communication failure

- TSA processor does not receive sensing data
- TSA processor does not receive signal status (phase and interval) data
- Traffic signal controller does not execute TSA timing

Traffic signal in special operation modes

- Flashing mode
- Preemption mode

SECTION 3. CONCLUSIONS AND NEXT STEPS:

This work conducted to date towards assessing the technical feasibility of developing a left turn assist system for signalized intersections with permissive left turns and to enact traffic signal adaptation has progress significantly – but is incomplete. The end goal of this project is to show the feasibility of providing left-turn decision support and warnings to the driver through an in-vehicle interface in order to make an objective decision on whether to proceed to a Field Operational Test.

In short, the CICAS-SLTA project consists of three major phases:

- Phase 1: This phase is called "Conduct Detailed Investigation" and its overall objective is to translate research, field observations, and experimental results into a testable and deployable system. In this phase, research is aimed at an operational understanding of gap acceptance from the human factors perspective of driver attitudes and behavior; as well as the vehicle perspective by evaluating vehicle positions, speeds, and zones of conflict for left turns. Additionally, the components needed for vehicle detection, pedestrian and bicyclist sensor suites are studied. A key output of Phase 1 is an initial SLTA alert algorithm that will be tested in Phase 2.

At the end of Task Order 6608, we have nearly – but not fully – completed Phase 1. Remaining work includes integrating the components, then establishing a set of tests at the Richmond Field Station Intelligent Intersection facility to determine whether emerging SLTA algorithms are sufficient to proceed to Phase 2.

We have reported on work to date by starting with a Concept of Operations, which provides an overview and conceptual framework. We followed with a detailed discussion of the main thrust of our effort, which was to provide data for and models for the crucial aspect of a SLTA alert or advice algorithm: observations and subsequent gap

assessment models. Our work in this area was essentially divided into two portions, (i) a "PATH portion" of four sections: Description of Data Sets, Interfaces, Analysis Models, then Results (Characteristics of the Predicted Post-Encroachment Metric) and (ii) a section which we subcontracted to the University of Michigan Transportation Research Institute (UMTRI) or the "UMTRI portion" (Analysis of Left-Turn Gap Acceptance at Signalized Intersections Using Naturalistic Driving Data). The "PATH Portion" was based upon mining and interpreting data collected at the PATH Richmond Field Station Intelligent Intersection through the course of several years, and the "UMTRI Portion" was similarly based upon mining and interpreting data, theirs however from antecedent Road Departure Collision Warning (RDCW) Field Operational Test data.

We next covered the CICAS-SLTA field test data collection, then we focus on detecting the primary other road user: the pedestrian. This section is followed by a description of the SAE J2735 interaction (Messages and Maps). We ended the SLTA report by describing the upgrades undertaken but not yet completed with the Richmond Field Station Intelligent Intersection.

If this next, experimental step is successful, the following sequence of events may occur -

- Phase 2: This phase is to "Assess the Technical Feasibility" of implementing a CICAS-SLTA system. In this phase all the pieces of an SLTA system that was researched in Phase 1 will be assembled and a system of vehicle and infrastructure components will be tested and evaluated using a test-bed in California. The primary output of this Phase will be a Feasibility Report with recommendations on the next steps for CICAS-SLTA.
- Phase 3: This Phase is optional since it is dependent on the results of the feasibility assessment in Phase 2; and mutual agreement between the USDOT and Caltrans that it is beneficial to further CICAS-SLTA research. The recommendations of Phase 2 will determine specifics that should be done to "Assess Field Deployability" and prepare for a full scale Field Operational Test (FOT) that will use naïve drivers. If Phase 3 is exercised, the primary tasks will be: further engineering field tests of the CICAS SLTA sensing, prediction, and warning algorithms at intersections which were instrumented in Phase 2; a pilot FOT; and plans for a full scale FOT.

REFERENCES

Alameda County Congestion Management Agency. http://www.accma.ca.gov/pages/AboutIntro.aspx. Accessed July 15, 2007.

Allen, B., Shin, B., and Cooper, P. (1978). Analysis of Traffic Conflicts and Collisions, *Transportation Research Record*, 667: pp. 67-74.

Agent, K. R. (1987). "Guidelines for the use of protected/permissive left-turn phasing." *ITE Journal* 57, no. 7 (1987): 37-43.

G. Antonini and J. P. Thiran. Counting pedestrians in video sequences using trajectory clustering. *IEEE Trans. Circuits and Systems for Video Technology*, 16(8):1008–1020, 2006.

I. Banerjee, S.E. Shladover, J.A. Misener, C.-Y. Chan, D.R. Ragland, "Impact of Pedestrian Presence on Movement of Left Turning Vehicles: Method, Preliminary Results and Possible Use in Intersection Decision Support", Transportation Research Board Annual Meeting, Washington DC, January 2005, Report No. 05-2199.

J. Batista, P. Peixoto, C. Fernandes, and M. Ribeiro. A dual-stage robust vehicle detection and tracking for real-time trafficmonitoring. In *IEEE Intelligent Transportation Systems Conference*, pages 528–535, 2006.

D. Beymer, P. McLauchlan, B. Coifman, and J. Malik. A real time computer vision system for measuring traffic parameters. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 495–501, 1997.

Bonneson J., Zimmerman, K. and Brewer M., (Aug. 2002) Engineering Countermeasures to Reduce Red-Light-Running, Federal Highway Administration, Washington, DC. FHWA/TX-03/4027-2.

Bougler, B., Cody, D., Geyer, J., Horne, J., Misener, J., Nowakowski, C., Rodier, C., Ragland, D., and Shaheen, S. (2005). *Investigation of Elderly Driver Safety and Comfort: In-Vehicle Intersection "Gap Acceptance Advisor" and Identifying Older Driver Needs* (UCB-ITS-PRR-2005-23). Berkeley, CA: California PATH Program, Institute of Transportation Studies, University of California, Berkeley.

Bougler, B., Cody, D., and Nowakowski, C. (2008). *California Intersection Decision Support: A Driver-Centered Approach to Left-Turn Collision Avoidance System Design* (Technical Report UCB-ITS-PRR-2008-1). Berkeley, CA: California PATH, Institute of Transportation Studies, University of California, Berkeley.

California SWITRS database. http://www.chp.ca.gov/switrs/. Accessed July 15, 2007.

Chan, C-Y, et al., "Observations of Driver Time-Gap Acceptance at Intersections in Left-Turn Across-Path Opposite-Direction Scenarios," TRB Transportation Records, No. 1910, pp. 10-19.

Chan, C.-Y. (2005). "Intersection Decision Support: A Systems Approach to Achieve Nationally Interoperable Solutions." *UCB-ITS-PRR-2005-11*, California Partners for Advanced Transit and Highways (PATH).

Chan, C-Y, "An Investigation of Traffic Characteristics and Their Effects on Driver Behaviors in Intersection Crossing-Path Maneuvers," IEEE Intelligent Vehicles Symposium, to be held in Istanbul, Turkey, June 2007.

Chan, C-Y, "Defining Safety Performance Measures of Driver-Assistance Systems for Intersection Left-Turn Conflicts," IEEE Intelligent Vehicles Symposium 2006, Tokyo, Japan, June 2006.

Chan, C-Y, Bougler, B., "Evaluation of Cooperative Roadside and Vehicle-Based Data Collection for Assessing Intersection Conflicts," IEEE Intelligent Vehicles Symposium 2005, Las Vegas, June 2005.

Chan, C., Cody, D., Cohn, T., Dickey, S., Greenhouse, D., Mak, T., Marco, D., Nguyen, K., Misener, J., Nowakowski, C., Ossenbrugen, P., Ragland, D., Shladover, S., tan, S., Vanderwerf, J., Wang, X, Zabyshny, A., and Zennaro, M. (2005). *California Intersection Decision Support: A Systems Approach to Achieve Nationally Interoperable Solutions* (Technical Report UCB-ITS-PRR-2005-11). Berkeley, CA: California PATH, Institute of Transportation Studies, University of California, Berkeley.

C.-Y. Chan, F. Bu and S. Shladover, "Experimental Vehicle Platform for Pedestrian Detection", California PATH Research Report UCB-ITS-PRR-2006-16, August 2006.

Chang, J., Cohen, D. Blincoe, L., Subramanian, R., and L. Lombardo (2007). "CICAS-V Research on Comprehensive Costs of Intersection Crashes" (Report No. 07-0016), National Highway Traffic Safety Administration, Washington D.C. Available from http://www-nrd.nhtsa.dot.gov/pdf/nrd-01/esv/esv20/07-0016-O.pdf.

Cody, D. (2004). Analysis of Exploratory Field Test Data from an Instrumented Vehicle: Description of Left Turn Maneuvers for Intersection Safety Countermeasure Design. *Proceedings of the 83rd Annual TRB Meeting. Washington, D.C.: Transportation Research Board, National Research Council.*

Cody, D. Nowakowski, C., and Bougler, B. (2007). Observation of Gap Acceptance During Intersection Approach. *Proceedings of The 4th International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle design.* Iowa City, IA: The University of Iowa, Public Policy Center. pp 321-327. [Online: http://ppc.uiowa.edu/driving-assessment/2007/]

R. Cucchiara, C. Grana,M. Piccardi, and A. Prati. Detecting moving objects, ghosts, and shadows in video streams. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(10):1337–1342, 2003.

A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39 (Series B):1–38, 1977.

R. Ervin, C. MacAdam, J. Walker, S. Bogard, M. Hagan, A. Vayda, and E. Anderson. System for assessment of the vehicle motion environment (SAVME): volume i, ii, 2000.

W. Forstner and E. Gulch. A fast operator for detection and precise location of distinct points, corners, and centers of circular features. In *Proc. Intercommision Conf. on Fast Processing of Photogrammetric Data*, pages 281–305, 1987.

Fuerstenberg, K.C., "Intersection Driver Assistance System - The EC project INTERSAFE," 13th World Congress & Exhibition on Intelligent Transport Systems, London, United Kingdom, October 2006.

E. R. Green, K. R. Agent, and J. G. Pigman. Evaluation of auto incident recording system (AIRS). Research Report KRC-05-09, Kentucky Transportation Center, University of Kentucky, Sept. 2005.

Intelligent Transportation Systems - Cooperative Intersection Collision Avoidance Systems. U.S. Department of Transportation. http://www.its.dot.gov/cicas/index.htm. Accessed July 15, 2007.

S. Kamijo, Y. Matsushita, K. Ikeuchi, and M. Sakauchi. Occlusion robust tracking utilizing spatio-temporal markov random field model. In *icpr*, volume 01, page 1140, 2000.

N. K. Kanhere, S. J. Pundlik, and S. T. Birchfield. Vehicle segmentation and tracking from a low-angle offaxis camera. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005.

M. Kilger. A shadow handler in a video-based realtime traffic monitoring system. In *IEEE Workshop on Applications of Computer Vision*, 1992.

Z. Kim, G. Gomes, R. Hranac, and A. Skabardonis. A machine vision system for generating vehicle trajectories over extended freeway segments. In *12th World Congress on Intelligent Transportation Systems*, 2005.

Z. Kim and J.Malik. Fast vehicle detection with probabilistic feature grouping and its application to vehicle tracking. In *Proc. IEEE Intl. Conf. Computer Vision*, volume 1, pages 524–531, 2003.

Z. Kim. Realtime obstacle detection and tracking based on constrained delaunay triangulation. In *IEEE Intelligent Transportation Systems Conference*, pages 548–553, 2006.

D. Koller, J. Weber, and J. Malik. Robust multiple car tracking with occlusion reasoning. In *Proc. European Conf. on Computer Vision*, pages A:189–196, 1994.

D. Koller, K. Daniilidis, and H.-H. Nagel. Model– based object tracking in monocular image sequences of road traffic scenes. *International Journal of Computer Vision*, 10(3):257–281, 1993.

Lee S.E., Knipling R.R., DeHart M.C., Perez M.A., Holbrook G.T., Brown S.B., Stone S.R., and Olson R.L., (March 2004) Vehicle-Based Countermeasures for Signal and Stop Sign Violations: Task 1 and Task 2, U.S. Department of Transportation, National Highway Traffic Safety Administration, Washington DC. DOT HS 809 716.

Maze, T. H., Henderson, J.L., Sankar, S. (1994). "Impacts on Safety of Left-Turn Treatment at High-Speed Signalized Intersections." Iowa State University.

Misener, et al, "Toward a Smarter, Safer Intersection: Application of Intelligent Countermeasures to the Intersection Crash Problem," *ITS World Congress*, Madrid, Spain, Nov. 2003.

James A. Misener, Steven E. Shladover, "PATH Investigations in Vehicle-Roadside Cooperation and Safety: A Foundation for Safety and Vehicle-Infrastructure Integration Research," *Proceedings of the IEEE Intelligent Transportation Systems Conference*, Toronto, Canada, September 2006.

Misener, J., Chan, C., Cody, D., Dickey, S., Nowakowski, C., Greenhouse, D., Ragland, D., Shladover, S., and Vanderwerf, J. (2007). *California Intersection Decision Support: A Systems Approach to Achieve Nationally Interoperable Solutions II* (Technical Report UCB-ITS-PRR-2007-1). Berkeley, CA: California PATH, Institute of Transportation Studies, University of California, Berkeley.

Najm W.G., Smith J.D. and Smith D.L., (July 2001) Analysis of Crossing Path Crashes, U.S. Department of Transportation, National Highway Traffic Safety Administration, Washington DC. DOT HS 809 423.

Neale, V., Perez, M., Lee, S. and Doerzaph, Z. (2007). Investigation of Driver-Infrastructure and Driver-Vehicle Interfaces for an Intersection Violation Warning System. *Journal of Intelligent Transportation Systems*, 11(*3*): pp. 133 - 142.

Nowakowski, C. (2006). A Preliminary Evaluation of a Driver-Infrastructure Interface for a Left-Turn Warning. *Proceedings of the 50th Annual Meeting of the Human Factors and Ergonomics Society.* Santa Monica, CA: The Human Factors and Ergonomics Society.

Nowakowski, C., Cody, D., and O'Connell, J. (2008). A Comparison of Infrastructure Versus In-Vehicle Driver Interfaces for Left-Turn Warnings. *Proceedings of the 87th Annual TRB Meeting.* Washington, D.C.: Transportation Research Board, National Research Council.

Oda, T., Tajima, T., "Driving Safety Support System for Senior Society, " *ITS World Congress*, Toronto, Canada, October 1999.

David R. Ragland and Aleksandr A. Zabyshny, "Intersection Decision Support Project: Taxonomy of Crossing-Path Crashes at Intersections Using GES 2000 Data" (August 1, 2003). *UC Berkeley Traffic Safety Center*. Report UCB-TSC-RR-2003-08.

Richard, C.M., Michaels, E.F. and Campbell, J.L., (Nov. 2005) Driver Attitudes and Behaviors at Intersections and Potential Effectiveness of Engineering Countermeasures, Federal Highway Administration, Washington, DC. FHWA-HRT-05-078.

Schagrin, M., "Cooperative Intersection Collision Avoidance Systems Initiative, " Presentation at Intelligent Transportation Systems America Annual Meeting, May 2005.

J. Shi and J. Malik. Motion segmentation and tracking using normalized cuts. In *International Conference on Computer Vision*, pages 1154–1160, 1998.

Schulze, M., "PReVENT: Prevention and Active Safety Application to Improve Traffic Safety," IEEE Intelligent Vehicles Symposium, Las Vegas, June 2005.

O. Tuzel, F. Porikli, and P. Meer. Human detection via classification on riemannian manifolds. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2007.

VII California. http://www.viicalifornia.org/. Accessed July 15, 2007.

White, B. and Ferlis, R. (2004). Algorithm for determining inattentive signal violators to be used in infrastructure-based intelligent system for signal violation prevention. *Proceedings of the 83th Annual TRB Meeting.* Washington, D.C.: Transportation Research Board, National Research Council.

A.Worrall, G. Sullivan, and K. Baker. Pose refinement of active models using forces in 3d. In *Proc. 3rd European Conf. Computer Vision*, pages 341–352, 1994.

Yamamoto, S., Mizutaru, K., Seto, M. "AICHI DSSS (Driving Safety Support System) Field Verification Test," 13th World Congress & Exhibition on Intelligent Transport Systems, London, United Kingdom, October 2006.

Yan, X., and Radwan, E. (2007). Effect of Restricted Sight Distances on Driver Behaviors During Unprotected Left-Turn Phase at Signalized Intersections. *Transportation Research Part F*, 10: pp. 330-344.