U.S. Department
of Transportation
**National Highway
Traffic Safety
Administration**

People Saving People
www.nhtsa.dot.gov

DOT HS 809 414

August 2000

Development of Performance Specifications for Collisions Avoidance Systems for Lane Change Crashes

Technical Report Documentation Page

1. Report No.
DOT HS 809 414

2. Government Accession No.

3. Recepient's Catalog No.

4. Title and Subtitle
Development of Performance Specifications for Collisions Avoidance Systems for Lane Change Crashes.

5. Report Date
August 2000

6. Performing Organization Code

7. Authors
S. Talmadge, R. Chu, C. Eberhard, K. Jordan, P. Moffa

8. Performing Organization Report No.

9. Publishing Organization Name and Address
TRW Space and Defense
One Space Park
Redondo Beach, CA 90278

10. Work Unit No. (TRAIS)

11. Contract or Grant No.
DTNH22-93-X097922

12. Sponsoring Agency Name and Address
National Highway Traffic Safety Admin. (NHTSA)
Office of Vehicle Safety Research
400 Seventh Street SW, Room 6220
Washington, D.C. 20590

13. Type of Report and Period Covered
Final Report - August 1993 through August 2000

14. Sponsoring Agency Code
NRD-12

15. Supplementary Notes

16. Abstract

This report presents preliminary guidelines to develop an effective lane change Collision Avoidance System (CAS). These guidelines were developed, to a large part, based on the experiences from a small number of naïve drivers tested on public roads with a testbed vehicle. A functional lane change CAS was implemented on the testbed utilizing today's technology. This was done to demonstrate utility and to perform realistic testing. However, the implementation described in this report is not the only acceptable one and it is certainly not the least expensive. Nevertheless, it demonstrates that an effective lane change CAS can be built today.

In order to avoid lane change collisions, the lane change CAS must monitor the areas on either side of the vehicle to determine the presence of another vehicle that could interfere with a planned lane change. It must also determine if a vehicle is approaching those areas with enough relative speed to potentially be in conflict with the instrumented vehicle. Those two tasks must be accomplished in any and all driving environments.

The effectiveness of the lane change CAS was estimated using drivers' errors as a surrogate for collisions. This method led to an estimate of 43% effectiveness, which is consistent with previous estimates found in the literature. Although the cost saving per vehicle is estimated to be relatively small for a lane change CAS, eliminating lane change crashes is still deemed to be a worthwhile endeavor, one that focus groups say they want and will pay for.

17. Key Words
Collision Avoidance System, Lane Change, Guidelines, Benefits

18. Distribution Statement
This document is available to the public through the National Technical Information Service, Springfield, VA 22161

19. Security Classif. (of this report)
Unclassified

20. Security Classif. (of this page)
Unclassified

21. No. of Pages
178

22. Price

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized

This page intentionally left blank

This page intentionally left blank

1. EXECUTIVE SUMMARY 1

2. INTRODUCTION 2
  2.1 Scope 3
  2.2 Objective 3
  2.3 Background 3
  2.4 Approach 5
  2.5 Outline 6

3. DEFINITIONS 7
  3.1 System Definitions 7
  3.2 Driver Definitions 9
  3.3 Applicable Documents 9
    3.3.1 Project Documents 9
    3.3.2 Other References 10

4 OVERALL SYSTEM GUIDELINES 13
  4.1 Simplicity for Driver Comprehension 13
  4.2 Bias Towards Safety 14
  4.3 Minimize False Positives 14
  4.4 Practicality and Utility 15
  4.5 Some Measured Results 15

5 PERFORMANCE GUIDELINES 17
  5.1 Modes of Operation 17
  5.2 Detection of Conflicts 17
  5.3 Mitigation of False Positives 26
  5.4 Clarity of Warning 29
  5.5 Public Safety and Avoidance of Interference 30
  5.6 Availability and Reliability 31
  5.7 Failsafe Operation 31
  5.8 Other Considerations 32
  5.9 Lane Change CAS Test Procedures 33
    5.9.1 Environmental Conditions 33
    5.9.2 Required Equipment 33
    5.9.3 CAS Configuration 34
    5.9.4 Test Procedures 34

6. BENEFITS ESTIMATES 37
  6.1 Review of Statistics Pertaining to Proximity Crashes 39
  6.2 Review of Statistics Pertaining to Fast Approach Crashes 39
  6.3 Review of Benefits Models 40
    6.3.1 Taxonometric Approach 40
    6.3.2 Kinematic Approach 40
    6.3.3 Human Usage Approach 41
    6.3.4 Collision Surrogates 41

6.4 Benefits Analysis 42
  6.4.1 Benefits Estimation Methodology 42
  6.4.2 Test Run Results 48
  6.4.3 Model Limitations 51
6.5 Costs 52

7. SUMMARY AND CONCLUSIONS 54

APPENDIX A: TESTBED TEST RESULTS A-1
A-1 Introduction and Executive Summary A-1
A-2 Testbed Description A-5
  A-2.1 Scanning Laser Rangefinder A-8
  A-2.2 Eye-tracker A-15
  A-2.3 Global Positioning System A-19
  A-2.4 Data Acquisition System A-22
  A-2.5 DSP and Host System Software A-24
  A-2.6 Detection, Tracking and Warning A-27
  A-2.7 Eye-tracker Data Correction Algorithm A-38
A-3 Static and Semi-Dynamic Tests A-49
  A-3.1 Detection Area and Range Verification A-49
  A-3.2 Detection Latency A-52
  A-3.3 Range and Speed Verification A-56
A-4 Track Tests A-66
  A-4.1 Dynamic Warning Test A-66
  A-4.2 Dynamic Tests with Interference A-71
  A-4.3 Three Lane Test A-74
  A-4.4 Approach and Pass Test A-75
  A-4.5 Modifications to SV A-77
  A-4.6 Test Track Preferences A-80
A-5 Open Road Testing A-85
  A-5.1 Experimental Design A-85
  A-5.2 Testing Protocols A-88
  A-5.3 Lane Change Dynamics A-90
  A-5.4 Eye Tracker Results A-102
  A-5.5 System Performance and Driver Reactions A-108
A-6 Summary and Conclusions A-113
A-7 References A-114

APPENDIX B: DISPLAY MODALITY QUESTIONNAIRE B-1

APPENDIX C: DRIVER INFORMATION LETTER C-1

Acknowledgement

TRW would like to acknowledge the many fruitful discussions about collision avoidance with Drs. August Burgett and David Smith of the National Highway Traffic Safety Administration. Special mention must go to Louis Tijerina of Ford Motor Company, who, when he was at the Transportation Research Center in East Liberty Ohio, did the experimental design and contributed in many other ways to the successful completion of this work.

This page intentionally left blank

1. EXECUTIVE SUMMARY

When lane change collisions occur, the driver usually has no awareness of the danger. This is because he or she "looked, but did not see." However, a lane change collision avoidance system (CAS) could also be "looking" and, if used properly, would require that both the driver and the CAS miss the threat for a collision to occur. The lane change CAS would thus represent another set of "eyes" watching the area around the subject vehicle and alerting the driver to potential conflicts. Seen in this way, a lane change CAS would be a very attractive addition to future vehicles. This report presents guidelines that when followed will lead to an effective lane change CAS.

These CAS performance specifications guidelines were developed, to a large part, from test track and on-road test data taken using engineering staff and 12 naïve subjects driving a TRW lane change CAS testbed vehicle. This vehicle implemented the functions required in a lane change CAS utilizing today's technology. This was done to demonstrate utility and to perform realistic testing. Of course, the TRW implementation described in this report is not the only acceptable one and it is certainly not the least expensive. Nevertheless, it does demonstrate that an effective lane change CAS can be built today.

In order to avoid lane change collisions, the lane change CAS must monitor the areas on either side of the vehicle to determine the presence of another vehicle that could interfere with a planned lane change. It must also determine if a vehicle is approaching those areas with enough relative speed to potentially be in conflict with the instrumented vehicle. Those two tasks must be accomplished in any and all driving environments.

By using drivers' errors as a surrogate for collisions, the effectiveness of a lane change CAS was estimated at 43%. This number is consistent with previous estimates found in the literature, but should, nevertheless, be regarded as a preliminary estimate due to the small number of data points used in this study to quantify it. Although the cost saving per vehicle achieved is relatively small for a lane change CAS, eliminating this type of crash remains a worthwhile endeavor, and is a service that the public says it wants and will pay for.

2. INTRODUCTION

As part of the study of potential Intelligent Transportation Systems (ITS) technology countermeasures designed to alleviate various crash types, the Space and Electronics Group of TRW was funded by The Office of Vehicle Safety Research of the National Highway Traffic Safety Administration (NHTSA) to study collision avoidance systems (CAS) for dealing with lane change collisions. Previous efforts on this contract are summarized below and a list of reference documents can be found in sections 3.3.1 and 3.3.2:

- As part of Task 1, extensive analyses of lane change crash data from the 1992 NHTSA General Estimates System (GES) and case-by-case analyses of lane change crashes from the 1992 NHTSA Crashworthiness Data System (CDS) were performed. A taxonomy of lane change collision classifications and crash-related events was developed. Based on these analyses of the crash data, lane change conflict scenarios and opportunities for collision avoidance were identified. The taxonomy, the size assessment data and the statistical descriptions of lane change crashes are published in the Task 1 report [1].

- The Task 2 report [2] also presented the results of the analyses of the crash scenarios and developed a set of CAS functional goals for countermeasures designed to prevent or mitigate these crashes.

- Under the Task 3 effort [3], we evaluated a significant number of existing CAS devices. The types of devices tested here fell into two general categories: proximity sensors looking to the side of the vehicle, and longer range sensors to detect higher speed vehicles approaching from behind the instrumented vehicle in an adjacent lane. These devices were attempting to address significant crash types utilizing technology that can be fielded today.

- The objective of Task 4 [4] was to develop and test preliminary performance specifications for a lane change CAS to achieve each of the identified functional goals. Using Monte Carlo simulation techniques, the parameters of various CAS configurations were tested against potential crash scenarios to determine those parameters producing the optimal effectiveness.

- For Task 5, a technology assessment and evaluation of the currently available technologies applicable to a lane change CAS were generated [5].

- Under Tasks 6 and 7, TRW designed and built a fully automated data acquisition system for assessing the effectiveness of a lane change collision avoidance system [6]. The system, which has been installed in a government furnished vehicle, features two subsystems that allow for automated acquisition and analysis of data:
    - a scanning laser rangefinder that serves both as the collision warning sensor and collector of ground truth information, and
    - an eye-tracker which outputs, at 30 Hz, the driver's gaze direction.

- This report describes, in Appendix A, the work performed in Task 8, wherein a plan was developed for the calibration and use of the testbed briefly described above with a representative sample of naïve drivers. The results of the driving tests and data reduction from the automated data acquisition system were used to evaluate the various display modes and the drivers' acceptance and behavior adjustments with a prototype CAS [7,8].

This work culminates with the performance specifications presented in this report.

2.1 Scope

This document addresses those CAS that deal with lane change crashes involving other vehicles (called the Principal Other Vehicle or POV) in the lane adjacent to that of the instrumented subject vehicle (SV) and those in which the POV is initially behind the SV in the adjacent lane (into which the lane change is to be made) and is traveling faster than the SV in the same direction. Only systems that provide a warning to the driver when a vehicle is in potential conflict to a lane change are covered. No systems that take active control of the vehicle (for example, by inhibiting steering) are considered.

2.2 Objective

The objective of a lane change CAS is to reduce the number and severity of lane change crashes in order to save lives, lessen injuries, and lower the associated costs of these collisions.

2.3 Background

In 1992, there was a total of 306,372 lane change crashes (standard error 23,645) as determined from the GES analysis [1]. This represents approximately 5.1 % of the 5,982,606 police reported crashes represented by the GES. Drifting crashes and crashes when leaving a parked position are not usually included as lane change maneuvers. Without these, the total is 235,950 crashes (standard error 18,854) which is 3.9 % of the total. This is consistent with previous studies [9,10].

The lane change CAS could impact the crash categories shown in Table 2-1.

Table 2-1: Summary of lane change crash categories which could be addressed by a lane change CAS (GES, 1992).

SV striking (angle)	87,264	
SV striking (sideswipe)	80,676	
SV struck (angle)	26,410	
SV struck (sideswipe)	9,803	
SV rear end struck	10,656	
Both vehicles changing lanes	4,790	(will detect presence inside proximity zone but further identification of a threat requires determination of lateral acceleration)
Drifting (angle)	26,003	(requires determination of lateral acceleration)
Drifting (sideswipe)	22,614	(requires determination of lateral acceleration)
Subtotal	268,216	or 4.5% of 5,982,606 police reported crashes
SV rear end striking	16,351 *	(may be indirectly affected; since the SV driver has better situational awareness of what is going on behind him, he can direct more attention to what is going on in front of him.)
SV leaving parking	21,805 *	(Large angle between SV axis and POV bearing)
Total	306,372	

* Lane change CAS may not be appropriate

The case by case analyses of crash data and police crash reports from the 1992 CDS and the 1992 GES suggest that most lane change crashes are due to the driver's being unaware of the potential crash threat posed by the other vehicle. For this reason, any system that detects the threat and warns the driver with sufficient lead time to perform corrective action will mitigate the effects of these collisions. Although neither database contains the information on when the driver first perceived the threat, the fact that in less than 8 % of the traditional lane change crashes did the driver even attempt any corrective action indicates that the driver was caught by surprise or did not have sufficient time to respond. It is entirely possible that the threat was never perceived. Many CDS case analyses indicate this.

In all these cases, the function of the countermeasure system is to warn the driver of imminent danger in a timely manner. Our goal is to analyze the system parameters in order to determine the requirements for a CAS that warns the driver with sufficient time to avoid the crash while minimizing false positives.

In this report, we will quantify the performance specifications for a comprehensive lane change CAS consisting of two detection and warning subsystems:

- the *proximity* warning subsystem detects the presence of vehicles in the lane adjacent to the instrumented vehicle at short distances and warns the driver, if appropriate, prior to initiation of a lane change maneuver.
- the *fast approach* warning subsystem detects and issues a warning on higher closing speed vehicles approaching from behind the instrumented vehicle in the adjacent lane.

Of course, for either subsystem both sides of the instrumented vehicle must be covered.

Some lane change categories listed in Table 2-1 above will not be directly addressed by this system. The "lane change, rear end striking" collisions could be avoided or mitigated through the use of a forward looking or longitudinal CAS. The drifting crashes would be better avoided by a lane keeping system that would alert the driver whenever a lane boundary was being crossed while no turn signal was activated. Crashes when two vehicles are simultaneously changing lanes are very rare, and we believe that many of them could be avoided by the suite of systems we've considered. Active systems which take control of the vehicle are not included in this study explicitly. Most likely, they would provide the same warning as the systems discussed here. The only difference is that the active systems would take some level of control of the vehicle if the driver has not responded appropriately to the warning within some predetermined time.

2.4 Approach

Specifications are generated based on information gathered from a variety of sources. These include the analyses of the crash databases that allowed us to develop a taxonomy of lane change crash types and an understanding of the causal factors and kinematics associated with them. We also obtained valuable information from the testing of available lane change CAS and from Monte Carlo analyses performed during the earlier portions of this contract. A survey of the existing technology gave us insight into what capabilities existed for the sensing and computing parts of the CAS. Finally, the test drives conducted with our testbed vehicle led to a clearer understanding of the features that a CAS must have in order for it to be useful.

The testbed system features two subsystems that allow for automated acquisition and analysis of data. The first is a scanning laser rangefinder that serves both as the collision warning sensor in the CAS and as a collector of ground truth information. The second subsystem is an eye-tracker to monitor the driver's head and eye motion. The eye-tracker outputs the driver's gaze direction at 30 Hz. The system is tied together by a fully automated data acquisition system for assessing the effectiveness of the lane change CAS portion of the system and collecting ground truth. The prototype CAS has been integrated into a passenger car.

The hardware was integrated into a 1993 Chevrolet Caprice Classic equipped with in-mirror displays [16] and mirror-edge mounted LED's as collision warning displays to the driver. Both the displays and the car itself were furnished by the U.S. DOT. The core of the system is a Pentium computer and a digital signal processor (DSP) board. The DSP is sized to accept the continuous stream of laser data at 12 kHz. Real-time data processing and the collision warning algorithm are resident in the DSP. Thus all the sensor outputs required for the warning algorithm are brought into the DSP, including vehicle speed, steering angle, and turn signal usage. Based upon the algorithm's decision to warn the driver, the DSP controls the displays via digital output lines. An independent video system multiplexes and records the images from four cameras monitoring the area behind and to the side of the car. The video record is secondary to the computerized system and is used for scene confirmation.

The data collected falls into roughly three categories. The first concerns driver eye glance behavior, such as mirror glance time and glance frequency per mirror, eyes on the road time, and head turn incidence. The second category is in-vehicle driver behavior measures, such as turn signal use and turn signal onset time with respect to lane change start. The last category is driver-vehicle measures, such as lane change completion time, forward and rear gap distances, brake applications, peak decelerations per application, and travel speed.

2.5 Outline

The remainder of this document is divided as follows. Section 3 contains definitions and references, including those to the previous reports generated on this program. The overall system guidelines approach is presented next in Section 4. Section 5 includes the performance guidelines for the comprehensive CAS. Benefits estimates are addressed in Section 6, followed by a summary and conclusions.

3. DEFINITIONS

The following are brief definitions of terms encountered in this report.

3.1 System Definitions

Availability . The fraction of time the CAS is functioning when it should be operational.

Benefits Analysis. A technique to estimate the reduction in the number and/or severity of a kind of crash when a CAS or other safety system is employed.

Closing velocity. The difference between the velocity of the POV and the velocity of the SV, i.e., ($V_{POV} - V_{SV}$). Closing velocity is positive if the gap distance is decreasing and negative if the gap distance is increasing.

Comprehensive CAS. A CAS that performs the functions of the proximity CAS subsystem and the long-range CAS subsystem.

Crash avoidance system (CAS). Any system which improves the driver's situational awareness in order to prevent or mitigate the damage done by an automotive crash.

Crashworthiness Data System (CDS). A database designed to provide extensive crashworthiness information, such as deformations and penetrations of the passenger compartment, glass breakage, etc. and occupant data. Estimates based on the CDS when weighted are used to quantify the losses due to motor vehicle crashes during a given year.

Detection Probability (P_D). The fractional number of times the system correctly issues a warning when a valid threat is present. $P_D = \frac{TP}{TP + FN}$ (see section 3.2)

Detection zone. The specific area in which a sensor can reliably detect another vehicle or object.

Fast approach crash. A lane change crash in which the closing velocity, ($V_{POV} - V_{SV}$), is greater than 15 mph [2].

Fast approach sensor. A lateral CAS sensor whose detection zone lies behind the SV. Typically this sensor detects both location and closing velocity.

Gap distance. The distance from the rear bumper of the SV to the front bumper of the POV.

General Estimates System (GES). A database of vehicular crashes based on Police Accident Reports (PARs) which can be used to study the conditions surrounding any kind of motor vehicle crash. The data is sampled and each sample is weighted to accurately represent all police reported crashes occurring on public roads during a given year.

Latency time. The time required by the CAS to process the data and present a warning to the driver.

Lateral CAS. A CAS whose detection zone lies in the areas adjacent to the SV, to the left hand side and the right hand side. It monitors the lanes adjacent to the SV's lane for avoiding lane change and roadway departure collisions.

Longitudinal CAS. A CAS whose detection zone lies in the same lane either in front of or behind the SV. It monitors the SV's lane for maintaining safe vehicle headways and avoiding rear end and backing collisions.

Long-range CAS Subsystem. That part of the CAS that monitors the adjacent lane and warns if a vehicle with the appropriate closing speed is present.

Monitor Mode. A CAS approach where the lane change warning is presented even when the turn signal is not activated. When the turn signal is on, the presentation character of the warning changes to a more immediate format.

Principal other vehicle (POV). The vehicle detected by the CAS sensor which may pose a threat to the SV. If there is more than one sensor, there may be more than one POV.

Proximity CAS Subsystem. That part of the CAS that monitors the proximity zone and warns if another vehicle is present.

Proximity crash. A lane change crash in which the relative speed of the SV and POV, $|(V_{POV} - V_{SV})|$, is less than or equal to 15 mph [2].

Proximity sensor. A lateral CAS sensor whose detection zone (proximity zone) lies between the front of the SV up to a given distance from the rear of the SV.

Rangefinder. A device to measure distance of an object from the user. It usually utilizes either a radar or a laser radar (ladar).

Rejection Ratio. The fraction of non-valid threats that the CAS does not warn on.

Repetition rate and/or update time. The rate at which the CAS checks for obstacles in the detection zone. Update time is the interval between checks.

Subject vehicle (SV). The vehicle equipped with the CAS.

Turn-signal-activation Mode. A CAS approach where the lane change warning is only presented when the turn signal is activated.

Vehicle reaction time. The time interval between the driver's action (braking, steering, etc.) and the beginning of the vehicle response.

Warning Display. The informational display produced by the CAS that a vehicle is located in the zone adjacent to the SV or that a vehicle is rapidly approaching the zone adjacent to the SV from the rear. The warning display may have several modes indicating various levels of urgency or indicating the source of the warning (proximity or fast approach subsystem) or both.

3.2 Driver Definitions

Driver reaction time (informed or surprised). The time interval required for the driver to acquire and process new information and begin to react to it. Studies have derived several distributions of drivers' reaction time [11]. Frequently, the surprise reaction times have been used. However, it may be argued that a driver of a vehicle equipped with a CAS is not surprised when the CAS delivers a warning, and therefore, the reaction times characterized by an alert reaction times are more appropriate. Also, when a driver is performing a more stressful maneuver such as a lane change, the driver is more alert and hence the alert driver reaction times distribution may be more appropriate. Simulation runs have been conducted for both sets of reaction times.

True Positive (TP). A warning that is issued when a valid crash threat is actually present.

False Positive (FP). A warning that is issued when no valid crash threat is present.

True Negative (TN). A n instance when no valid threat is present and no warning is given.

False Negative (FN). An instance when a valid crash threat is present and no warning is given.

Nuisance Alarm. An inappropriate warning that is issued when the system determines that a crash threat is present but the driver does not perceive it as a threat, such as a parked vehicle.

Glance. A shift in visual regard which entails a movement of the head and/or eyes.

Glance Duration. The amount of time the driver's attention is focused in one direction.

3.3 Applicable Documents and References

3.3.1 Project Documents

1. C.D. Eberhard, K.M. Luebkemann, P.J. Moffa, S.K. Young, R.W. Allen, E.A. Harwin, J.P. Keating, and R. Mason, "Development of Performance Specifications for Collision Avoidance Systems for Lane Change, Merging, and Backing; Task 1

Interim Report: Crash Problem Analysis", DOT HS 808 431, NHTSA, Washington, D.C., 1994.

2. S. K. Young, C. D. Eberhard, and P.J. Moffa, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 2 Interim Report: Functional Goals Establishment.", Contract No. F04606-90-D-0001/0052 and DTNH22-93-X-07022, NHTSA, Washington, D.C., 1995.

3. S. Talmadge, K.E. Yokoyama, G.A. Shreve and S. Johnston, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 3 Interim Report: Test of Existing Hardware Systems", DOT HS 808 434, NHTSA, Washington, D.C., 1995.

4. C.D. Eberhard, P.J. Moffa, S.K. Young, and R.W. Allen, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 4 Interim Report: Development of Preliminary Performance Specifications", DOT HS 808 430, NHTSA, Washington, D.C., 1995.

5. P.J. Moffa, J. Austin, G.S. Dow, I. Ikizyan, and M. Hibben, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 5 Interim Report: Crash Countermeasure Technology Investigation", DOT HS 808 506, NHTSA, Washington, D.C., 1996.

6. S. Talmadge, D. Dixon, and B. Quon, "Development of Performance Specifications for Collision Avoidance Systems for Lane Change Crashes; Task 6 Interim Report: Testbed Systems Design and Associated Facilities", DOT HS 808 369, NHTSA, Washington, D.C., November, 2001.

7. S. Talmadge, R.Y. Chu, and K. Jordan, "Development of Performance Specifications for Collision Avoidance Systems for Lane Change, Merging and Backing; Task 8 Interim Report: Testbed Test Results", included as Appendix A of this report.

8. S. Talmadge, R. Y. Chu, and R. S. Riney, "Description and Preliminary Data from TRW's Lane Change Collision Avoidance Testbed", Proceedings 10th Annual Meeting ITS America, Boston, May, 2000.

3.3.2 Other References

9. J-S Wang and R. R. Knipling, "Lane Change/Merge: Problem Size Assessment and Statistical Description", DOT HS 808 075, National Highway Traffic Safety Administration, Washington, D.C., January 1994.

10. J.D. Chovan, L. Tijerina, and G. Alexander, D.L. Hendricks, "Analysis of Lane Change Crashes", Contract No. DTRS-57-89-D-00086, U.S. Department of Transportation, October, 1993.

11. L.M. Forbes, "Discussion: The Driver's Role in Collision Avoidance Systems", Collision Avoidance Systems, Issues & Opportunities, Workshop Proceedings, ITS America – Safety and Human Factors Committee and the National Highway Traffic Safety Administration, March 21-22, 1994. Examples of several "Alert" and "Surprise" reaction time distributions are given in this document, which also includes the individual references to the studies in which each of the reaction time distributions were derived.

12. L. Tijerina and W.R. Garrott, "A Reliability Approach to Estimate the Potential Effectiveness of a Crash Avoidance System to Support Lane Change Decisions", SAE Paper 970454, presented at the SAE International Congress and Exposition, Detroit, February, 1997.

13. L. Tijerina, "Sensitivity Analysis of a Reliability Model of the Potential Effectiveness of a Crash Avoidance System to Support Lane Change Decisions", presented at the NHTSA and ITS Expert Panel on Modeling Safety Benefits, ITS America Conference, Washington, June, 1997.

14. L. Tijerina, W.R. Garrott, D. Stolzfus, S. Johnston, G. Forkenbrock, E. Parmer, and F. Brickman, "Driver Acceptance and Preferences for SCAS Visual Alerts in the Case of Normal Lanekeeping, No Desire to Change Lanes, and Awareness of Principal Other Vehiclein an Adjacent Lane", Internal working paper, East Liberty, OH, December, 1996.

15. N.D. Lerner, B.M. Cotwal, R.D. Lyons, D.J. Gardner-Bopnneau "Interim Report: Preliminary Human Factors Guidelines for Crash Avoidance Warning Devices", DOT HS 808 342, National Highway Traffic Safety Administration, Washington, D.C., January 1996.

16. The display icon chosen was a red triangle, an international symbol of caution. This icon was installed into a Muth mirror.

17. R. Huey, J. Harpster, and N. Lerner, "In-Vehicle Crash Avoidance Warning Systems; Human Factors Considerations", Summary Report DOT HS 808 531, Washington D.C., February 1997.

18. L. Tijerina and S. Hetrick, "Analytical Evaluation of Warning Onset Rules for Lane Change Crash Avoidance Systems", presented at the Human Factors and Ergonomics Society Meeting, Albuquerque, September, 1997.

19. H. Jula, E. Kosmatopoulos, and P. Ioannou, "Collision Avoidance Analysis for Lane Change and Merging", Technical Report, University of Southern California, April, 1999.

20. S. Blackman, Multiple Target Tracking with Radar Applications, Artech House, pp. 21–25, 1986.

21. I. Nair, "Assessment of Potential Biological Effects from Electromagnetic Radiation in the Microwave and Infrared Regions", Contract No. DTNH22-93-D-07007, National Highway Traffic Safety Administration, Washington, D.C., May, 1995.

22. A. Burgett, "Methodologies for Evaluating the Impact on Safety of Intelligent Vehicle Safety Systems", Proceedings of the 14th International Conference on Enhanced Safety of Vehicles, Munich, May, 1994.

23. A. Burgett, "DOT's Approach to ITS Safety Evaluations", Proceedings of ITS America Workshop on ITS Safety Evaluations, May, 1995.

24. T.A. Dingus, "Empirical data Needs in Support of Model Building and Benefits Estimation", presented at the NHTSA and ITS Expert Panel on Modeling Safety Benefits, ITS America Conference, Washington, June, 1997.

25. J. Koziol, V. Inman, M. Carter, J. Hitz, W. Najm, S. Chen, A. Lam, M. Penic, M. Jensen, M. Baker, M. Robinson, and C. Goodspeed, "Evaluation of Intelligent Cruise Control, Volume 1 – Study Results", DOT HS 808 969, NHTSA, Washington D.C., October, 1999.

26. J-S Wang, R. Sherrer, and L. Blincoe, "Monetary Benefits of Collision Avoidance Systems (CAS)", Preliminary Assessment of Crash Avoidance Systems Benefits, NHTSA Benefits Working group, October, 1996.

27. "Consumer Acceptance of Automotive Crash Avoidance Devices", U.S. Department of Transportation, Charles River Associates Project No. 852-05, April, 1998.

28. L. Barr, W. G. Najm and J. D. Smith, "Crash Problem Definition for the Intelligent Vehicle Initiative", DOT-VNTSA-HW00Q-PM-99-01, FHWA, Washington, D.C., December 1999

4. OVERALL SYSTEM GUIDELINES

Our approach to developing a lane change CAS is driven by simplicity, safety, practicality, and utility. In the following sections we will elaborate on our philosophy.

4.1 Simplicity for Driver Comprehension and Acceptance

A CAS is useful only if the driver easily understands the displayed information and accepts the displayed information as reliable. The display must be easily detected and interpreted by the driver during the routine activities associated with driving, and the interpretation of the display should not interrupt the driver's train of thought or distract him from routine monitoring of the driving environment. Indeed, it should assist him in maintaining a high level of awareness of the traffic around him. The driver must also understand the limitations of the CAS, e.g., the detection zones are finite and well defined, the system will not warn of a fast approach POV in the adjacent lane until it is within a certain distance (adjustable) of the SV, the system will not warn of a driver changing lanes into the SV driver's target lane until the separation is less than a lane width, and so on. The SV driver should appreciate these limitations. The SV driver should not depend only on the CAS to determine if the lane change is safe, and in all cases he should verify for himself that the lane change is safe before continuing. This mode of use which is the safest is termed parallel usage [12,13].

We have investigated positioning the warning display in either the rearview mirror or sideview mirrors or both [14-17]. These are among the natural places where the driver will look prior to a lane change maneuver. For some modes of operation, the warning display is present, if appropriate, even when there is no indication that a lane change is imminent. This is the monitor mode. The immediacy of the warning changes when the turn signal indicator is activated. The display is transformed from a steady red light to a flashing one. For the other mode, termed the turn-signal-activation mode, the warning is only presented when the turn signal is on, and it is a flashing red light. Flashing red lights are generally associated with an immediately dangerous situation [15-17], such as a railroad crossing when the train is imminent.

Presenting these displays in the mirrors should minimize confusion with other CAS displays that may also be present in the vehicle since, with the exception of a backing CAS, their displays would naturally be located differently. For a backing CAS, the warning could be presented in the same mirror(s) since it would only be activated when the SV was in reverse, a time when the lane change CAS would not be operational.

Finally, a small indicator light could be located in the warning display area(s) to remind the driver of the existence of the CAS and to assure the driver that the CAS is operational and has satisfactorily passed all of the internal self tests. Its nature should be such that it could not be confused with the actual warning display. Alternatively, an indicator light (located with the other indicator lights) could be illuminated for a short time when the vehicle is started and then go out if the CAS is operational. If a self-test failed, then the indicator light would stay illuminated at startup or light up when the fault was first detected.

4.2 Bias Towards Safety

Any system that is designed to detect an object must also confront the chance of generating a false positive. A false positive is a warning display initiated when no object or threat is present. One way to generate a false positive is when the detection threshold is set so low that sensor noise can trigger a detection. The simplest way to avoid these false positives is to raise the detection threshold. However, raising the threshold could cause legitimate targets with small signal returns to be missed. Any CAS must be biased to not missing true detections for obvious safety reasons. A missed detection could be fatal, whereas a false positive can only be annoying.

Our approach has always been to keep missed detections down to an absolute minimum by using discrimination algorithms to differentiate appropriate detections from false ones. In the testbed system, our key methodology to accomplish this is to track targets and subject them to a persistence criterion before they are declared as obstacles. Random noise detections will tend not to persist through this filter, and thus our approach which utilizes a very low detection threshold in order to avoid misses can still have an acceptably low false positives rate. Additional approaches to reducing false positives are discussed in the next section.

4.3 Minimize False Positives

False positives can be caused by the system detecting a real object and issuing a warning even though the object presents no collision threat. A simple example to demonstrate this concept is the case where a tree on the side of a curving road is detected by a forward-looking sensor. The tree is a valid collision threat type that appears to be in the direct path of the vehicle. However, because the road is curving, it poses no real danger. Ways to avoid issuing a warning in this case have been investigated. They include using a video camera or other sensor to determine the roadway geometry, tracking a number of vehicles in front of the SV to determine when a curve is present, and monitoring the steering of the SV.

For a lane change CAS, false positives can be generated by parked cars, hedges, guardrails, and other objects along the side of the roadway, oncoming traffic, vehicles more than one lane over or directly behind the SV (especially on curves), and overhead signs and the roadway itself.

We have used the following approaches to avoid these false positives. To avoid warning on fixed objects, we measure the speed of objects detected ahead of the SV in the adjacent lane. This is enabled in our approach by utilizing a scanning laser rangefinder and by tracking the objects detected. The scanning system allows us to look ahead of the SV. Thus we can detect objects before they enter the detection zones and can establish tracks on them. Tracking is necessary because we have chosen a sensor that only measures distance, and we need a relative speed determination to implement the long-range CAS. In order to determine if an object is fixed, we must determine its relative speed and compare it to the SV's ground speed. Thus, we also need to monitor the SV ground speed. When an object is detected ahead in an adjacent lane and it is determined to be fixed, a warning is suppressed even when it enters the proximity warning zone. In addition, when the SV is travelling at less than 10 mph, no warning is issued. This avoids false positives when the SV is stopped at a traffic light and other vehicles are present.

When executing a left turn on a two-way street, oncoming vehicles in the adjacent lane could trigger a warning when they entered the proximity detection zone. Those warnings are not necessary and would be classified as false positives. In the same way that fixed objects are suppressed, a vehicle entering the proximity zone from in front of the SV with a velocity opposite to that of the SV can be ignored because they have been previously detected and tracked. Also, below a specified SV speed (10 mph as discussed above), no warnings are given.

Other false positives are avoided by carefully tailoring the proximity and long-range detection zones. In our implementation, this is accomplished by using a very narrow-beam sensor that is scanned and accurately measures range. Knowledge of the range and pointing angle of a narrow beam allows us to "paint" the detection zones very accurately. This lets us discriminate between vehicles in our own lane or two lanes over even at long ranges and on curves. Vertical control of the detection zone also provided by narrow beams avoids detection of the roadway itself or overhead signs, even at the longest detection ranges.

4.4 Practicality and Utility

The operation of the CAS must be automatic and reliable. For this reason, the CAS must be initiated whenever the vehicle is started and operational whenever the transmission is engaged in any of the forward gears. The CAS must be capable of self test when it is initially turned on and should repeat the self tests at regular intervals during continuous operation. The results of the self tests should be relayed to the driver as discussed in Section 4.1.

The CAS sensor(s) must operate in all possible lighting conditions and weather conditions. Although some degradation in sensor performance may be tolerated during inclement weather, the sensor(s) should be packaged so that there is no permanent damage caused by exposure to normal amounts of rain, snow, sleet, dust, mud or blowing sand.

4.5 Some Measurement Results

In order to evaluate our implementation of a lane change CAS, a fraction of our data was analyzed in detail [7]. One day of driving for each of three of the naïve drivers was selected. Those days were chosen since they were the only ones where a missed detection (false negative) was noted. As such, the estimates based on missed detections are an upper bound. This selected data covered 12,477 seconds of driving. The results of the analysis are given below.

True Positives	413
False Positives	146
False Negatives	3

Here true positives are defined as the cases where a valid threat was present and warned against. False positives arose when objects on the side of the road like parked cars and poles were incorrectly warned against. This typically happened on curves where our algorithm has trouble computing the relative speed well enough to eliminate them. Also, false positives occurred when the algorithm incorrectly assumed that a target was present when in fact there was none. See [7] for details on the actual algorithms that were implemented. Finally, the few false negatives (misses) arose when a valid threat was not warned against. Another point derived from the data was that 995 stationary targets were detected and correctly not warned against. This represents only a tiny fraction of the total number of true negatives. A true negative is the situation wherein the system determines that there is no threat present and does not issue a warning. In principle a CAS has an opportunity to issue warnings once every cycle. One can readily see that over the course of an extended drive, the duration of time when there is no object, let alone a valid target, in the detection zone is apt to be considerable, i.e. an extremely large number of cycle periods.

This data reflects on the performance of the lane change CAS we implemented. As mentioned above, we were very concerned with missed detections. We only missed three targets out of 416. This yields a detection probability of 0.9928 (413/416).

What is most interesting is that the false positive rate is just over 42 per hour. In our questioning of all naïve drivers (i.e., those not informed about the intent or design of the experiment), none complained about these false warnings [7]. In the literature [17] it has been stated that more than a few incorrect warnings per hour would be unacceptable. That analysis was specifically for an audible warning that served no real purpose. Our naïve drivers see value in the lane change CAS and our warning is described as, and recognized by drivers as, a situation awareness warning and is intentionally designed to be unobtrusive. As such, our false warnings are much more easily tolerated.

5. PERFORMANCE GUIDELINES

Performance guidelines for a lane change CAS will be presented in this section. As for other CAS, there are three critical and interacting subsystems that must be addressed:

- The sensor subsystem,
- The algorithm/processor subsystem, and
- The display subsystem.

Guidelines will be presented for the overall lane change CAS system and then flowed down into the major subsystems as appropriate.

The primary requirements for the lane change CAS are the reliable detection of potential conflicts prior to a lane change decision, the minimization of false and nuisance alarms to facilitate driver acceptance, clarity of warning, ease of use, public safety and avoidance of interference, availability and reliability, and failsafe operation. Before guidelines that address these areas are presented, two possible modes of operation will be discussed.

5.1 Modes of Operation

During our driving tests two modes of operation were investigated. They differ in when warnings are presented. For the monitor mode, warnings were given to the driver whenever a vehicle with the appropriate characteristics (as discussed in Section 5.2) was detected in any of the zones defined below. This mode provides situational awareness to the driver even when he or she has no intention of making a lane change. As discussed in Section 5.4, the character of the warning was changed when the turn signal corresponding to the same side of the SV as the detection was activated. A standard technique to convey increased severity was utilized. For the turn-signal-activated mode, the warning was only presented when the appropriate turn signal was activated, and only the increased severity warning was used. In either case, turn-signal-activation should be determined by the sensor subsystem including knowledge of the intended direction of lane change.

1. For monitor mode, the lane change CAS should issue a warning whenever a valid target is detected, and the warning should change if the corresponding turn signal is activated.

2. For turn-signal-activation mode, the lane change CAS should issue a warning whenever a valid target is detected, and the corresponding turn signal is activated.

3. The sensing subsystem should determine if the turn signal is activated and the intended direction of lane change.

****

5.2 Detection of Conflicts

The guidelines for the proximity and fast-approach zone sensing (see Figure 5-1) will be presented separately but they should all be considered as part of the guidelines for the fully functional lane change CAS. The issue here is the detection of another vehicle in the adjacent lane that may interfere with the driver's intended lane change and the presentation of a clear warning.

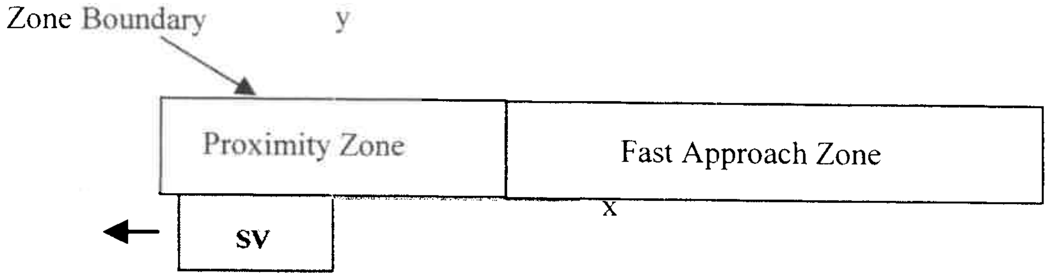

Figure 5-1: Schematic of the proximity and fast approach zones for the right side of the subject vehicle.

The proximity zone is defined as the adjacent lanes on either side of the vehicle and extending from in front of the front bumper to 30 feet behind the SV's rear bumper. It includes the areas where another vehicle may not be seen in any of the SV's mirrors, the so-called blind spots. The choice of 30 feet was based on the results of driving tests with engineers who designed the CAS system [7]. These preferences are shown in Figure 5-2. A number of POVs were presented to the engineers at various ranges and they decided on the range at which they no longer felt comfortable making a lane change. These tests were performed at two different test tracks and on a freeway as noted in the legend. Ignoring the two extreme values, their choices fall between 20 and 30 feet. We picked 30 feet as a conservative value. Taking into account the variations on this parameter, we believe that it should be adjustable. This will be discussed below as part of guideline 20. We found that extending the proximity zone in front of the SV was useful in some instances. The exact distance is TBR but we believe that four feet is reasonable. In our implementation, the lateral extent of the proximity zone was chosen as 11 feet in order to cover most of the adjacent lanes and yet not spill over into the lanes two over from the SV. Its lateral extent starts at the edge of the SV. Any vehicle the size of a bicycle or larger should be detected in that region.

Noting that the coverage zone is defined in terms of covering the adjacent lane, it is worth commenting that the length of the fast approach zone (162 ft. – see guideline #6 below) may be affected by curves in the road. A fixed coverage zone will spill outside the lane markings in a curve and could give rise to false positive alarms. It seems desirable that the coverage zones be made conformal to the road topology. Realizing that this might be a difficult goal in the short term, it could be implemented using either information from the steering angle, highly accurate GPS or accelerometer data.

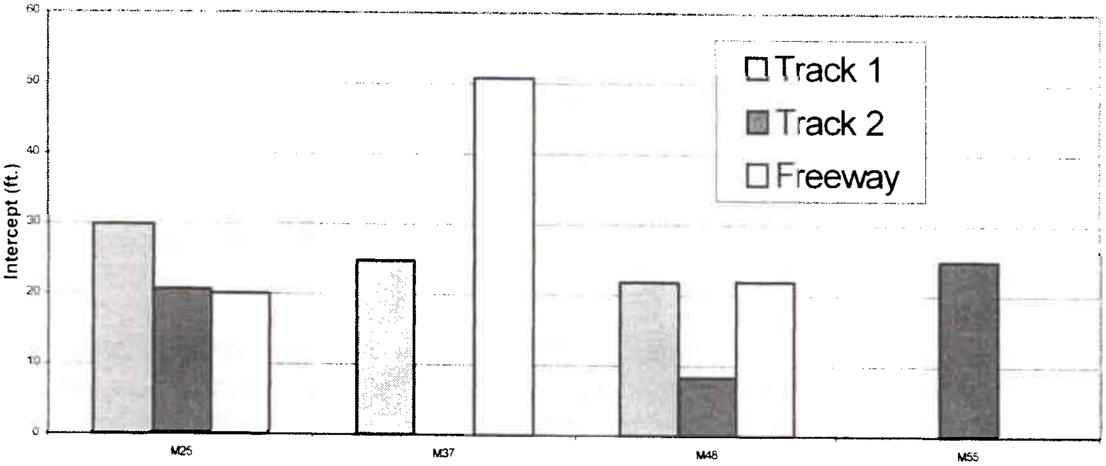

Figure 5-2: Preferences for maximum distance of proximity zone from drives with test engineers.

4. The lane change CAS should be able to detect any vehicle in the proximity zones which are the adjacent lanes on either side of the SV running longitudinally from four feet (TBR) in front of the front bumper of the SV to a distance 30 feet behind the rear bumper of the SV.

****

Any vehicle in the proximity zone is considered a threat regardless of its speed. This is clearly the case for one that is next to the SV. The proximity zones extend behind the SV in order to protect against another vehicle that may intersect the path of the SV during the lane change due to its higher speed. The same is true for slower speed vehicles in front of the SV, but they should easily be seen by the driver. For reasons to be discussed in the next section, stationary objects are not warned against.

5. The lane change CAS should issue a warning whenever a moving vehicle is detected in a proximity zone.

****

In addition to the proximity zone detection, there are guidelines related specifically to the function of detecting fast-approach POVs. The fast-approach zones extend on both sides of the SV. They are defined (see Figure 5-1) as having the same lateral extent as the proximity zones, that is, covering the adjacent lanes. The fast approach zones start at the far end of the proximity zones (nominally 30 feet but adjustable down to 20 feet) and extend to 162 feet beyond the SV's rear bumper. Any vehicle the size of a bicycle or larger should be detected in that region. In our implementation, we have chosen to make the fast approach zones to have less lateral extent than the proximity zones. This is done to avoid detecting vehicles in the same lane as the SV, especially on curves. Also, it avoids detecting vehicles in the same lane as the SV that are driving near the edge of the lane.

The choice of 162 feet is based on several factors we determined during test drives. Figure 5-3 displays the distribution of closing velocities we measured during all of the test drives. Whenever the SV was in motion, the passing speeds of all vehicles in the monitored lane were determined. The vast majority (over 99%) of measurements shows closing velocities below 30 mph (44 feet/second). One could use a lower closing speed and still cover most closing velocities encountered but then the warning on the fastest approaching vehicles would come late. This would give rise to a potentially dangerous situation. This danger is exacerbated by the fact that statistically the worst injuries occur for the collisions with the highest closing speeds.

In the same way as was done to determine the maximum extent of the proximity zone, test engineers were asked to state their preference for fast approach warning. Here the parameter of interest is the time until the fast approach vehicle will enter the proximity zone [18]. The results are tabulated in Figure 5-4.

An average value of 3 seconds was chosen for our algorithm. This means that the criteria for warning on a fast approach vehicle is the following. If the velocity is greater than a value given by the distance from the end of the proximity zone divided by 3 seconds, then a warning is given. This warning boundary is just a straight line in distance-velocity space. Most naïve drivers found this setting very reasonable with a few suggesting it should be shorter. Adjusting it will be discussed under guideline 20. Vehicles in the fast approach zone whose range and velocity are such that they would enter the proximity zone within 3 seconds are warned against. This definition of the fast approach warning space is shown graphically in Figure 5-5 along with the proximity zone [19]. The proximity zone is represented as region I, and the fast approach zone is marked as II. No warnings are issued for region III. The top edge of the proximity zone occurs at 30 feet and the slope of the fast approach border is 3 seconds. As discussed in guideline 20, both of those parameters will be adjustable. As the maximum range of the proximity zone is pulled in, so will the longitudinal extent of the fast approach zone. Also, the slope of the warning boundary changes when the 3 second threshold is lowered, and the maximum extent of the fast approach zone decreases. Noting the markings in Figure 5-5, this can be expressed mathematically as:

$$X_f^{max} = X_p^{max} + 44 \text{ ft/sec} * t_{warn},$$

where t_{warn} is nominally 3 seconds.

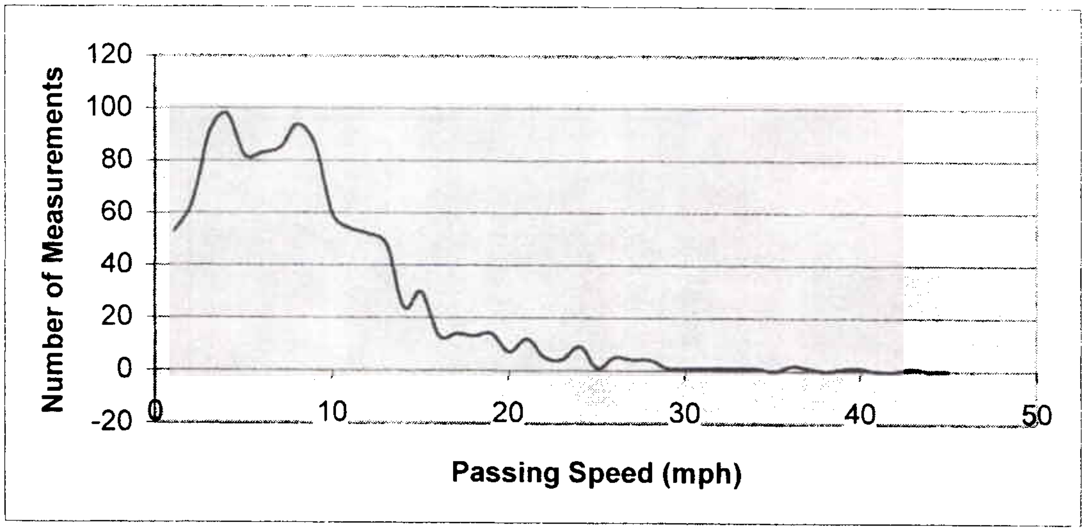

Figure 5-3: Distribution of passing speeds measured from all the test drives.

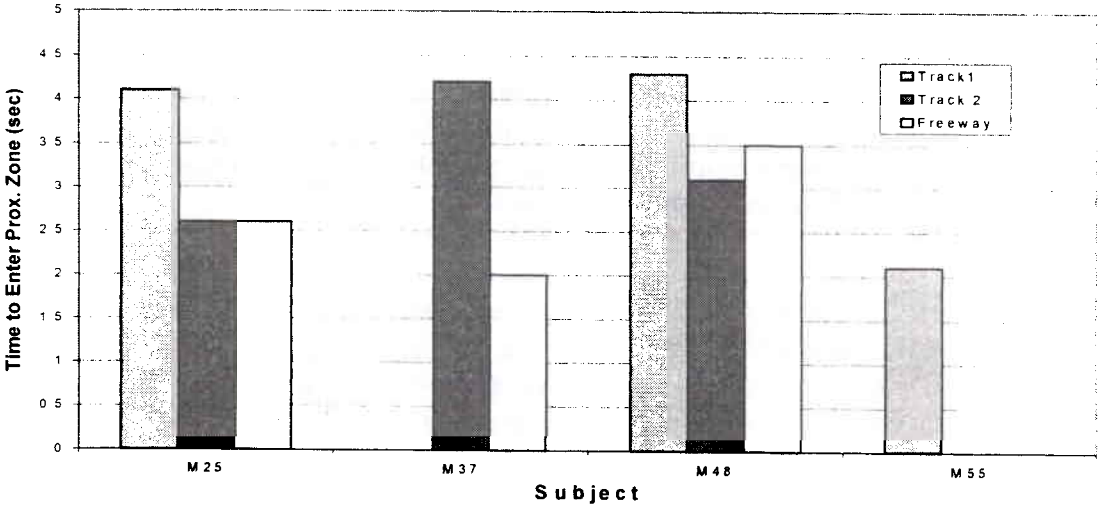

Figure 5-4: Preferences for time to enter proximity zone from drives with test engineers.

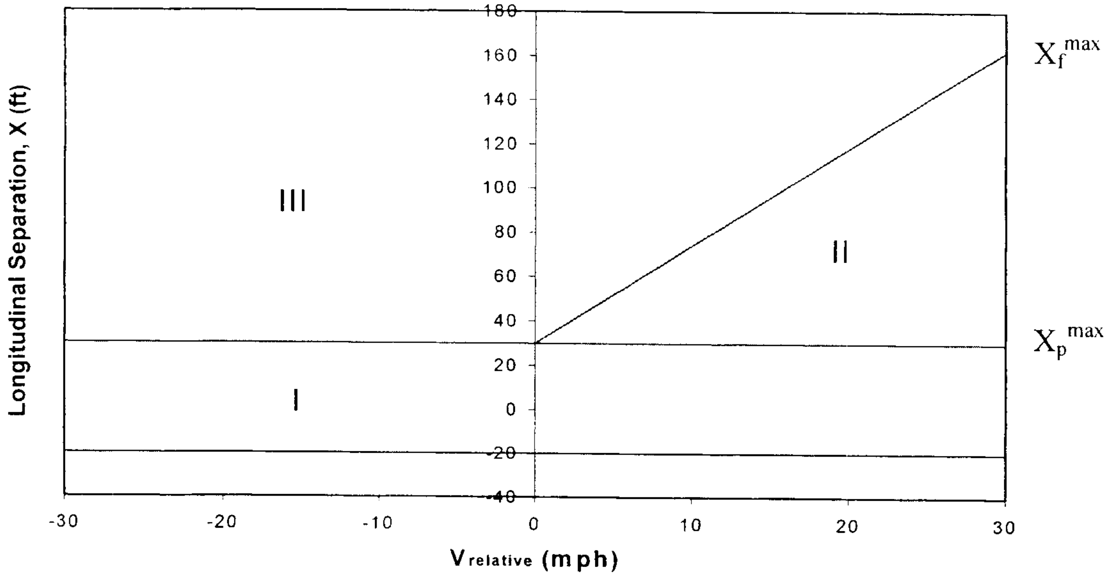

Figure 5-5: Warning zones for the lane change CAS.

6. The lane change CAS should be able to determine the presence and velocity of any vehicle in the fast approach zones which are defined to be the adjacent lanes on either side of the SV running longitudinally from the edge of the proximity zone at 30 feet to 162 feet beyond the rear bumper of the SV.

****

If all vehicles detected in the fast-approach zone were warned against, then a large number of nuisance alarms would result. Only those vehicles that are closing on the SV and are predicted to enter the proximity zone within the specified time are considered threats. An algorithm must take into account both the longitudinal distance of the detected vehicle from the SV and the closing speed. The closer the distance, the lower the closing velocity that has to be warned against. We have chosen to implement the simple expression below, where the nominal values are set to 30 feet for X_p and 3 seconds for t_{warn}. A warning is given if:

$$X \le X_p + v_{relative} * t_{warn}.$$

7. The lane change CAS should issue a warning for fast-approach vehicles based on their distance and closing speed.

****

We want to keep the maximum error in the computation of the time to enter the proximity zone to 0.5 seconds, based on the variations seen in the driver's preferences in Figure 5-4.

When utilizing a system that measures relative distance, relative velocity can also be determined. This can be done directly from a Doppler measurement or from an analysis of successive range measurements. The accuracy of the relative velocity measurement is driven by the fast approach warning algorithm. Errors in the estimates of range or velocity will lead to variations in the computed time to enter the proximity zone. Since we derive the velocity from the range in our implementation, there is a relationship between their errors. As reported in [20], it is given by:

$$\delta v = 1.3 * \delta R.$$

The variation in the time to enter the proximity zone is given by:

$$\delta t = -(R/v)*(\delta v/v) + \delta R/v.$$

For a desired δt of 0.5 seconds, these equations yield:

$$\delta v/v = 0.133$$

and

$$\delta R/R = 0.034.$$

Based on the warning algorithm, the accuracy for the range and velocity should be as given above. Since the relative velocity accuracy is only driven by the algorithm it will not be discussed further. We will defer the range accuracy discussion until the accuracy of the detection zones is considered.

8. The lane change CAS should be able to determine the time to enter the proximity zone to an accuracy of 0.5 seconds.

****

It is difficult to set an exact probability of detection. It is clear from any reasonable analysis of crash statistics that humans are excellent at detecting the presence of other vehicles around their own vehicle [13]. It would be extremely difficult to specify a probability of detection that is comparable to the driver's. However, some minimal performance needs to be set based on further testing and analysis.

If the driver's utilize the lane change CAS in a parallel mode wherein they check both the CAS display and the area into which they plan to make a lane change, then any detection probability will result in positive benefits. It is when the lane change CAS is used as a substitute for driver vigilance that additional risk is incurred. It would take a very high probability of detection to induce drivers to solely rely on the CAS.

The fact that the CAS cannot match the average driver's capability has extremely important safety ramifications. It cannot be stressed too strongly that the CAS should always be utilized as an adjunct to the driver's own vigilance and never as a substitute. If this is done, then a failure of both the driver and the CAS are necessary for a crash to occur.

In our implementation, we enhance the detection probability by maintaining the warning for 0.8 seconds after the detected object disappears. Hence, if a target is missed on one or two scans, it is not eliminated from the trajectory analysis. By maintaining the track, it can be reacquired at a later time without loss of warning. In this way, our probability of detection is enhanced since it requires several continuous misses before a false negative occurs. Since false negatives are so potentially dangerous, we recommend that any lane change CAS be implemented with some type of persistence criteria to eliminate this "flickering" effect in which a target disappears for short periods. This can happen for any ranging technology due to destructive interference from several scatterings off the same target, multipath, etc.

9. The probability of detection of a vehicle located in either of the lane change CAS detection zones should be greater than TBD percent.

****

It is difficult to precisely locate an object in the detection zone when using a typical ranging device. In order to avoid missed detections in the proximity zone and nuisance alarms from vehicles two lanes over, the edges of the proximity zone must be determined to an accuracy of one foot. Much less than this is very difficult and unnecessary, and much more will cause missed detections and raise the nuisance alarm level.

A ranging device usually determines distance from the sensor by emitting some form of radiation and then sensing reflected energy. The simplest way to visualize the process is to think of a very short burst of energy sent out and then reflected back. The round trip time is then directly related to the distance of the object given that the speed of propagation of the radiated energy is known. (For radars and ladars this is just the speed of light.)

Knowing the distance from the sensor does not pin down location. This is accomplished by focusing the emitted energy into a beam. Then knowledge of the pointing direction and the range allows one to locate the detected object in two dimensions. In order to assure that the object detected is in one of the proximity zones, required accuracies must be determined for the sensor measurement. The accuracy of one foot in determining the edge of the proximity zone will lead to range and pointing accuracies for the sensor.

For the proximity zone, it would be possible to utilize a presence detector. It could be something like a video camera. Then the processing of the imagery would lead to a determination of the presence of a threatening vehicle. However, it would be extremely difficult, if not impossible, to discriminate against fixed objects, and such a system would be plagued with false positives. Such a system would not utilize ranging as described above, but would have to define the proximity zone by pointing knowledge and field of view considerations. This points up the fact that we can not set range or pointing accuracy guidelines since they are strongly implementation dependent. Only the overall accuracy of defining the detection zones is required.

For long ranges it is difficult to determine which lane a detected vehicle is in. For a laser system, the beamwidth is typically very narrow. This allows for precise determination of the angular position of the detected vehicle but requires scanning or multiple beams to cover the entire fast approach zone. On the other hand, radar beams are much wider and are shaped like the contours shown in Figure 5-6. They can readily cover the entire fast approach zone, but their angular accuracy is worse. Again, there are implications for the pointing and ranging accuracies associated with determining the edges of the fast approach zone. In our ipmlementation, we have chosen a narrower zone to avoid some of these problems and this relaxes the accuracy required to a reasonable 2 feet. Again, this has implications on the pointing and ranging accuracies which must be met along with the restrictions due to the warning algorithm accuracy.

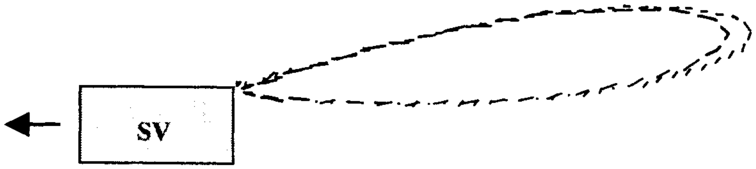

Figure 4-6: Typical coverage zone for a radar beam.

10. The lane change CAS should determine the edge of the proximity zone to an accuracy of 1 foot.

11. The lane change CAS should determine the edge of the fast approach zone to an accuracy of 2 feet.

****

Another important factor to be considered is latency. Even if the sensor has the ability to detect every object, the CAS cannot prevent crashes unless the warning is given to the driver in a timely and clear manner. Clearness will be discussed in Section 5.4 and timeliness will be presented here. The timeliness of the warning is related to the latency of the CAS, that is, the time it takes to report a warning after a target first appears. The latency is made up of the time it takes for the sensor to get the initial return from the target, the time it takes the processor to determine that a warning should be issued, and the time it takes for that warning to be presented to the driver. In our case, we use a scanning system and so the latency is driven by the scan speed of 10 Hz. As discussed in Section 4.3, we keep false alarms down by looking for a consistent set of detection before a valid detection is declared. The complexity of the algorithm used and speed of the processor chosen both contribute to the processing time. Finally, the display must be conspicuous enough and the warning clear enough so that the driver perceives and understands it very quickly. This has implications for the human factors aspects of the warning display. From our experience with testing other CAS systems as part of Task 3, we have determined that the latency of the CAS should be less than or equal to 0.2 seconds.

Meeting this latency for vehicle that cut into the fast approach zone from adjacent lanes was at first difficult for our implementation. This is because our tracking algorithm takes some time to settle down before an accurate estimate of the relative velocity can be obtained. In our implementation, we have chosen to "pre-track" vehicles in the regions adjacent to the fast approach zones so that if they enter a fast approach zone, a reliable velocity estimate will be available and the warning can be issued within the latency time.

12. The lane change CAS should have an overall latency of less than or equal to 0.2 seconds.

****

5.3 Mitigation of False Positives

False postitves are inevitable but they must be carefully controlled. Too many of them will lead to the CAS being turned off or ignored. Remember that false positives can occur when spurious signals from the sensor are interpreted as true detections by the detection software. One way to limit them is to choose a sensing subsystem with enough energy such that returned energy from the weakest targets is much larger than the noise level. In the sensor world, this is termed as having a good signal to noise ratio. Another helpful approach is to utilize detection algorithms that look for a consistency to the detections. In other words, a warning is not declared by the algorithm until several detections are measured with a consistent time history. Noise detection is a random process and this consistency should be missing from noise-related energy returns.

False positives can also be limited by careful control of the detection zones so that objects outside of them are not warned against. As discussed below, we further reduce false positives by controlling the times when warnings are given and by eliminating fixed objects and opposing traffic.

It is difficult to determine an exact limit for the number of false positives since it is based on the subjective value of the CAS as determined by each driver and is related to the nuisance alarm rate and the mode of warning presented. As pointed out in Section 4.5, even seemingly large rates can be tolerated by drivers.

13. The lane change CAS should have a false positive rate of less than TBD per hour.

****

For a lane change CAS, false positives could arise from the detection of fixed objects in either adjacent lane, approaching traffic one lane over and to the left of the CAS, vehicles in the fast-approach zone that are not threatening, and detections of objects outside of specified detection zones.

The CAS must determine if a detected object is stationary. This same guideline will allow the CAS to determine if a detected object is traveling in the opposite direction and thus suppress any warning. Since it is unlikely that many detected objects will be moving very slowly, the accuracy of the relative speed of other vehicle is not stressing. To differentiate fixed objects from moving vehicles, an accuracy of 5 mph is adequate. Knowledge of the SV speed is also required to decide if the detected object is stationary. This measurement requires an easy to achieve accuracy of 2 mph.

14. The sensing subsystem should be able to determine the SV speed to an accuracy of 2 mph.

15. The lane change CAS should not issue a warning on fixed objects.

16. The lane change CAS should not issue a warning on opposing traffic.

****

In our implementation, we determine if an object is fixed by detecting it ahead of the SV prior to its entering the proximity zone and then determining its relative speed. When the SV is turning, it is not possible to detect the object early enough to determine its speed before it enters the proximity zone. In order to eliminate false positives associated with fixed objects when the SV is turning, the sensing subsystem must monitor the steering and detect that a turn is being made. A turn is defined as a maneuver that occurs when the steering wheel is turned more than TBD degrees, and all warnings are suppressed during a turn.

17. The sensing subsystem should monitor the steering system in order to determine if a turn is being made.

18. The lane change CAS should not issue a warning during a turn.

****

It was found that a large number of potential nuisance alarms could result when the SV is stopped or moving very slowly. We determined that below an SV speed of 10 mph, it was better to suppress all warnings. This is another reason for measuring the SV speed.

19. The lane change CAS should only issue warnings when the SV is traveling greater than 10 mph.

****

The lane change CAS should come programmed to a very conservative definition of proximity and fast approach threat. If the driver finds this inappropriate, then in order to reduce these perceived nuisance alarms, some adjustment should be allowed. Clearly, limits should be set on how aggressive that setting can be made for the sake of safety. As shown in Figure 5-2, the preference for the extent of the proximity zone mostly ranged from 20 to 30 feet. We would recommend allowing adjusting the maximum extent of the proximity zone between 20 and 30 feet beyond the SV's rear bumper with the factory setting at 30 feet.

Similarly, as shown in Figure 5-4, some drivers preferred a "time to arrival" parameter as low as 2 seconds. Again, to limit perceived nuisance alarms, adjustment of this parameter should also be allowed. We would recommend providing for the adjustment of this parameter between 2.5 and 3 seconds, with the setting from the factory being at the most conservative value. The 2.5 second lower limit was chosen so that even with the maximum error in estimating the time to enter the proximity zone (0.5 seconds), the minimum allowable time will be 2 seconds.

20. The parameters utilized in the warning decision algorithm should be adjustable within limits, with their factory setting being at the most conservative (3 seconds for the time to enter the proximity zone, and 30 feet for the maximum longitudinal extent of the proximity zone).

****

The algorithm adjustment should be easily achieved by the driver through a clearly marked dial or other reasonable means. The labeling should utilize simple to understand terminology and the vehicle manual should describe the operation.

21. The lane change CAS should have a driver interface to adjust the warning decision algorithm within limits.

22. The lane change CAS should come with documentation that explains its operation and the adjustment.

****

In addition to detecting and warning on objects in the proximity and fast-approach zones, it is important to not warn on objects outside of those zones. This will help to limit false positives. In our implementation, we use a scanning laser radar with good range resolution and very narrow beamwidth. This allows us to clearly define the detection zones of interest. With a radar system, the beamwidth will always be much wider and great care will have to be exercised to achieve the necessary level of knowledge of the detected object's position. Of course, control of the coverage in both azimuth and elevation is required. The azimuthal control will help meet guidelines 10 and 11. The elevation control will help avoid false positives from the detection of overhead signs and the roadway itself.

23. Objects outside the proximity or fast-approach detection zones must not cause a warning to be issued.

****

5.4 Clarity of Warning

As mentioned in Section 3, a CAS is only useful if its warning can be easily detected and interpreted by the driver. There is a body of research on human factors that addresses how information can best be presented to a person. For a CAS, we are concerned with the tradeoff between getting the driver's attention without distracting him or her from the overall driving task. The display icon used in our tests [16] are good examples. Any CAS display should build on this body of information. Good overall sources of this data can be found in the literature [15,17].

24. The lane change CAS should follow well established human factors guidelines.
****

While a driver is preparing for a lane change, he or she will be looking in the rearview and sideview mirrors to assess the environment. A natural place to position the warning from the lane change CAS would be in one or both of those locations. Half the naive drivers preferred using only the sideview mirrors, but in order to accommodate those drivers who don't utilize those mirrors (for example, some prefer to look over their shoulder and some only use the driver side mirror) we propose to utilize both locations. The placement of the warning in the mirrors also encourages parallel usage since the driver's eyes can see the warning and perform his or her own surveillance at the same time.

25. The warning display for the lane change CAS should be placed in both the rearview mirror and the sideview mirrors.
****

The display should indicate on which side the warning is coming from. This can be readily accomplished by placing the warnings in the respective sides of the rearview mirror and in the relevant sideview mirror.

26. The display for the lane change CAS should clearly indicate which side of the vehicle the warning is coming from.
****

For the monitor mode, when the turn signal is activated, the warning should become more immediate and attention getting. That same level of warning should always be used for the turn-signal-activation mode. For our choice of a visual display, this urgency was conveyed by changing the warning from a steady burning icon to a flashing one.

27. Whenever the turn signal is activated, the lane change CAS warning should have a higher degree of urgency.
****

However the displays are presented in the mirrors, they should not interfere in the mirror's primary function. That is, they should block as little of the mirror surface as possible and be located off to the side of the reflecting surface. Mirrors must conform to 49 CFR Paragraph 571.111.

28. The warning display for the lane change CAS should not interfere with the primary function of the SV's mirrors.

****

5.5 Public Safety and Avoidance of Interference

As discussed in the Task 5 Interim Report, the most natural way to measure range and relative velocity of an object is by irradiating it with energy and measuring the scattered energy coming back to the sensor. For a comprehensive lane change CAS, we believe that a radiating system will be utilized. This leads to two more guidelines. One is that the system must meet stringent restrictions on how much radiation humans can be exposed to. These are given in [21] for both RF (radar) and electromagnetic (ladar) radiators in terms of acceptable energy density. These restrictions are a function of the duty cycle of the radiator (continuous or pulsed) and the frequency (wavelength).

29. The sensing subsystem should meet all safety restrictions on radiated energy.

****

Also, the radiated energy should not interfere with any other system on the SV or other vehicles on the roadway, even those with the same kind of CAS. Careful coordination among manufacturers and/or government regulations may be necessary to delineate the form of the energy radiated by various CAS and other on-board vehicle systems. There are FCC guidelines for which RF bands are available to the CAS designer. A region around 77 GHz has been set aside for use in automotive systems. However, even within that band, interference can occur. For example, a forward looking radar on a vehicle that is behind the lane change CAS-equipped one could easily overwhelm the energy it is trying to receive from its own reflected signal. The direct energy from the other system could easily be orders of magnitude larger because of the well known phenomenon that energy density diminishes like one over the distance traveled squared. The energy radiated by the other CAS directly at the lane change CAS may travel a much shorter distance than the energy the lane change CAS radiates which must reflect off a target and then return. Also, the reflected energy off of the target is usually only a small fraction of the energy radiated.

Similar arguments hold for a laser radar. The probability of being interfered with is a complex function of the number and density of other radiating systems on the roadway, the type of waveforms utilized, the duty cycle, the bandwidth of the emitters and receivers, and the frequency agility of the various radiators. These terms are defined in the Task 5 Interim Report [5].

30. The sensing subsystem should not interfere with any other system on the SV or other vehicles on the roadway.

****

5.6 Availability and Reliability

The lane change CAS should be reliable. It should be on whenever the SV is in a forward gear and traveling more than 10 mph.

31. The sensing subsystem should be available greater than TBD percent of the time when the SV is in forward gear and traveling greater than 10 mph.

32. The sensing subsystem should determine if the SV is in a forward gear.

****

The detection of a vehicle in the proximity or fast-approach zones should be reliable day and night under normal operating conditions including inclement weather, such as rain, hail, snow, dust, and fog. It shall be able to function in the presence of road spray.

33. The detection of a vehicle in the proximity and fast-approach zones should be reliable day and night under normal operating conditions including inclement weather, such as rain, hail, snow, dust, and fog; low-horizon sun; and in the presence of road spray.

****

5.7 Failsafe Operation

If the sensing system is not functional, then the entire lane change CAS cannot perform. The sensing subsystem should have the ability to know when it is not functional. This could be accomplished by monitoring the detections for consistency and/or by providing a way to sample the emitted energy periodically.

This could be accomplished for a scanning system like ours by putting an internal reflector in a part of the scan that is not needed for CAS performance and making sure energy is returned from that part of the scan. Also, if very short- range detections are being constantly received, that would indicate a problem. Perhaps mud is blocking part of the emitted energy aperture and reflecting some of it. Regardless of how it is accomplished, the lane change CAS indicator should alert the driver when a problem is determined to exist.

34. The lane change CAS should be able to self diagnose a failure.

35. The lane change CAS should have a means to alert the driver when a system failure condition is detected.

****

5.8 Other Considerations

There are several other guidelines that we have considered proposing. We do not have any direct evidence for their necessity, but we feel they may prove reasonable. At this point, we include them for completeness, but have no relevant data.

At this time, there is no clear preference for which mode to implement. Our naïve drivers were approximately split evenly in terms of their preference for either mode. Some experts fear that monitor mode will discourage use of the turn signal. On the other hand, many drivers don't always signal, and the turn-signal-activation mode would be useless in those cases. The turn-signal-activation mode is seen as a means to reduce the annoyance of false positives since the warning display is only operational for the short period when a turn signal is activated. On the other hand, for those instances when the naïve drivers felt that the warning was inadequate, 80% were for the turn-signal-activation mode. Until further data is acquired, we believe that a system with both modes would be best. A conservative approach to this issue would be to have the system start up in monitor mode, and then allow the driver to switch to turn-signal-activation mode. This may prove annoying to the driver that continually drives such an equipped vehicle and who prefers the turn signal mode. On the other hand, it ensures that different drivers of the same vehicle will always be presented with the option that provides maximum information.

36. The lane change CAS should have a switch to allow the driver to choose which mode is being employed. If further research deems it desirable, the start-up configuration should be monitor mode.

****

Some discussion has evolved around differentiating the warnings coming form the proximity zones and those from the fast approach zones. On the one hand, it may help the driver determine what vehicle is causing the warning and allow him or her to make an informed decision on whether to execute the lane change maneuver. This alleviates the concern that a warning could be caused by vehicles in both zones simultaneously, and a perfunctory evaluation by the driver might dismiss one of the warnings and miss the other. On the other hand, we may be presenting too much information to the driver and reducing the simplicity of the system. The following guideline is open to further research and is only preliminary.

37. Further research should determine if the lane change CAS should present a different warning to the driver for each proximity zone and fast approach zone.

****

There are certain lane change crashes like drifting ones where the driver may not be looking in the mirrors where the warning display is presented. We have chosen to use lights to display the warnings since that modality is less intrusive and hence the driver should tolerate more nuisance and false alarms. Also, the lights in the mirrors are conspicuous enough to be easily seen during the typical lane change maneuver.

For those cases where the visual warning is not adequate, there may be a need for an audible or haptic warning. This would be a third level of warning to supplement the solid and flashing light alerts already prescribed. This most urgent modality would be utilized when a crash was imminent. It would require the implementation of an algorithm based on the relative distance and speed of the SV and the POV. No details of that processing have been determined, and no testing with an audible alarm has been done. Nonetheless, we speculate that that type of warning may be appropriate under certain conditions.

38. Further research should determine if the lane change CAS should have an audible or haptic warning when a crash is imminent.

****

5.9 Lane Change CAS Test Procedures

The purpose of these test procedures is to advance a methodology whereby any lane change CAS can be readily evaluated to determine if it meets a minimum set of standards for functionality. They are designed to test for adherence to the specification guidelines and are based upon common lane change scenarios. These procedures are the product of experience in testing commercially available prototype systems as well as the refinement of those procedures necessitated by the development of a more advanced lane change system for human subject testing.

5.9.1 Environmental Conditions

For the following basic performance tests, the environmental conditions can be clear and dry on either concrete or asphalt road surface. However, a subset of the performance tests should be performed under adverse weather conditions with test results compared to clear weather conditions to note any degradation. These tests should, in principle, be technology independent. Electro-optic systems should be evaluated in strong sunlight, or shortly after sunrise or shortly before sundown to test for interference with natural light. This may be difficult to control but the tester should be cognizant of possible limitations and try to stress the technology.

5.9.2 Required Equipment

A minimum of one other vehicle (the POV) in addition to the vehicle with the lane change CAS installed (the SV) is necessary to perform these tests. This second vehicle will be used as a probe to measure the limits of the system under test.

When taking the measure of any system an appropriate yardstick must be used. By this, it is meant that the accuracy of the measurement made on the system under test must be at least as accurate as the system being measured. For this purpose it has been found that a differential GPS is most useful. Since any lane change driver warning algorithm will be based upon measurements made relative to the vehicle with the CAS installed, this lends itself nicely to differential measurements. It is not necessary that the differential solutions be available in real time; post processing is perfectly acceptable. Therefore two complete GPS systems must be procured – one installed on each of the two vehicles involved. The receive antennas must be installed at those points on the two cars with the shortest distance between them. For example, if one is testing lane changes on the right side of the subject vehicle, then one antenna is to be installed on the right rear corner of the subject vehicle, while the other should be situated on the left front corner of the POV. The GPS receivers must be capable of outputting pseudo range, phase and phase rate (Doppler shift). It is not necessary that the receivers be dual frequency, although the details should depend on the software used to process the raw data and arrive at the differential solution for the temporal evolution of range between the two receivers.

When properly used, systems of this type can achieve accuracies on the order of several centimeters with both vehicles in motion. Perhaps the most important requirement for accurate measurements with GPS is the necessity of acquiring static data at the beginning of each test segment. The more static data that the processing software has to use, the more accurate the solution upon which it converges. Acquisition times of 20 minutes yield the accuracy just mentioned.

The data acquisition system to be used should be self-contained with the input from the CAS in the form of a step voltage when the warning is given. This allows the warnings to be accurately timed. Timing is crucial since one is measuring distance and speed to sub-meter and meter/sec resolutions. Therefore it would be ideal if the timing clock of the data acquisition system could be periodically referenced to GPS time since ultimately it is the GPS system that will be calibrating those speeds and distances.

5.9.3 CAS Configuration

Configuration and calibration (if required) should be performed prior to the tests according to the manufacturer's specifications. If the CAS has any adjustable parameters, it is important to note these settings so that the proper evaluation can be performed. In addition, it is also important to note whether the manufacturer claims to reject nuisance warnings. One of the tests will be to determine the ability of the CAS to reject those objects not deemed a threat by the system.

For the purposes of these tests it is assumed that the only output that the CAS will provide will be a warning. No output concerning speed or distance is assumed.

5.9.4 Test Procedures

These tests are designed to verify the major specifications for a lane change CAS. These are latency, field of regard, and the accuracy of its warning algorithm.

5.9.4.1 Latency

Detection latency will be measured by recording the response to a POV moving at successive velocities on a path perpendicular to the longitudinal axis of the subject vehicle at rest, as shown in Figure 4.9-1. Since there is no speed determination necessary for warning on an object entering the proximity zone, the delay between the moment the POV crosses the boundary and the appearance of a warning is a measure of the system latency. The principle behind this test is that as the POV enters the testbed field of view at increasing speeds, the point at which the system gives a positive indication will also move gradually across the field of view. The detection latency is obtained by measuring the slope of the line of the positions at which the target is recognized, as a function of the velocity of the POV. This should be done at speeds ranging from 5 to 35 mph in increments of 10 mph, with at least five repetitions at each speed.

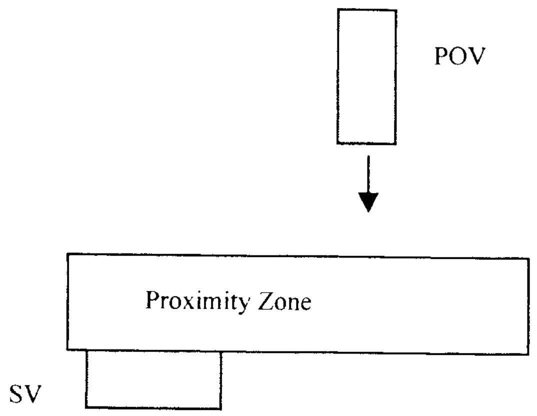

Figure 5.9-1 Test arrangement for determination of detection latency.

5.9.4.2 Static Determination of Algorithm Accuracy

This test shall take place on a straight section of the test facility, with the subject vehicle at rest. The POV is driven toward the subject vehicle as if it were in the adjacent lane. The POV is driven at constant speed toward the testbed vehicle and the time at which a warning is issued is recorded. POV speed shall be varied from 5 mph to 45 mph in increments of 10 mph, and the test at each speed shall be repeated five times. Plotting the range and speed as determined by the GPS at the moment of warning will yield a curve that can be compared to the one specified by the manufacturer.

5.9.4.3 Dynamic Determination of Algorithm Accuracy

Basically this is simply a repetition of the previous test, with the important distinction that both vehicles are now in motion. The test has the POV passing the SV at relative speeds of between 5 and 45 mph. The SV itself will maintain speeds at 20, 35 and 50 mph. Some combinations of speeds may be excessive for the facility used. Under no circumstances must safety be compromised. Each combination of SV and POV passing speed must be repeated five times in order to accumulate enough statistics to determine exactly how the CAS under test is performing.

5.9.4.4 Dynamic Determination of Proximity Zone

In contrast to the previous test where the POV passed the subject vehicle, in this test the situation will be reversed. As the subject vehicle passes the POV, the rear of the POV enters the leading edge of the proximity zone first. As it leaves, the front of the car exits the trailing edge of the proximity zone. Once it leaves the proximity zone, there should be no warning since the gap between the two vehicles is increasing. The basic functionality of this scenario will be tested as well as quantitatively verifying the accuracy of the boundaries of the proximity zone. With the POV traveling at 30 mph, the subject vehicle should pass at closing velocities of 10 and 20 mph. Again, five repetitions should be performed for each case.

5.9.4.5 Resistance to False Positive Alarms

Objects that give rise to false positive alarms include parked cars, poles, trees, guardrails and bushes. In short, anything that is stationary is almost by definition not a threat. If the CAS under test claims to be able to reject these objects, then it is important to measure this rejection ratio. For this test a number of obstacles can be arranged so that the subject vehicle can drive by to see if they elicit a warning response. A number of cars, trucks and/or vans of varying sizes and heights should be parked in a line with varying gaps between them. Metal and wood poles should be set up as well. Extended objects such as a row of bushes, guardrails and walls may be utilized where they exist. The subject vehicle should be driven past these objects at 30 mph while varying the lateral distance. Two values of the lateral distance, 4ft and 8 ft, should suffice. The tester must keep track of the cumulative number of objects driven past and the number of warnings triggered. Some objects may have a greater probability of triggering a warning, and these should be noted.

6. BENEFITS ESTIMATES

The thorough investigation of the 1992 General Estimate System (GES) and Crash Worthiness Data System databases was performed for lane change crashes for light vehicles only[1]. Although our discussion is based on this analysis, a recent revisiting of the crash statistics database [28] might modify somewhat our conclusions, primarily in the area of cost analysis. A taxonomy of collision subset classifications and crash-related events was developed. Based on these analyses, the lane change crash problem may be divided among the categories shown in Table 6-1:

Table 6-1: Lane change taxonomy.

SV* striking (angle)	87,264	
SV* striking (sideswipe)	80,676	
SV* struck (angle)	26,410	
SV* struck (sideswipe)	9,803	
SV* rearend struck	10,656	
Total	214,809	

* SV is the vehicle changing lanes.

The total of 214,809 crashes represents 3.6 % of 5,982,606 police reported crashes in 1992. In addition, there were 16,351 crashes in which the SV struck the rearend of the vehicle in front (SV rearend striking) after changing lanes. These crashes also may be indirectly affected. If the SV driver has better situational awareness of what is going on behind him, he may direct more attention to what is going on in front of him. Collisions due to lane change to the right were as frequent as those due to lane change to the left. In a majority (80%) of all lane change crashes, the lane-changing SV struck the POV.

Comprehensive damage and loss directly due to these crashes are assessed in terms of fatal crash equivalents (FCEs) and dollars as shown in Table 6-2 [9]. However, these estimates do not include other losses incurred by uninvolved motorists, such as time lost to roadway tie-ups, added pollution, etc.

Table 6-2: Fatal crash equivalent severity scale (originated in 1988).

Crash Severity (Most Severely-Injured Occupant)	Comprehensive $ Value Per Crash (1988 Dollars)	Fatal Crash Equivalent (FCE)
Fatality (K)	$ 2,722,548	1.0000
Incapacitating (A)	$ 228,568	0.0840
Non-incapacitating (B)	$ 48,333	0.0178
Possible Injury (C)	$ 25,228	0.0093
No Injury Reported (O)	$ 4,489	0.0016
Injury Unreported	$ 4,144	0.0015

As shown in the table above, there is damage or loss associated even with crashes in which there is "no injury" reported or the injury is unreported. This is not surprising since even minor crashes frequently produce anxiety, soreness and inconvenience to the drivers and passengers involved. No collision, even the most minor, is without some damage.

If the manners of collision (angle and sideswipe) are summed together for the "SV striking" and the manners of collision (angle, sideswipe and rearend) are summed for "SV struck" categories, the totals and per crash damage assessments are shown in Table 6-3. The separate category "SV rearend striking" is also shown.

Table 6-3: Damage Assessments (FCEs and 1999 Dollars)

Crash Type **	Crashes	FCEs	FCE / crash	No Injuries ***	Total Dollar Amount ****
SV striking cr	170,014	711	4.3 E-3	64 %	$ 2,725.5 M
SV struck crashes	48,633	310	6.4E-3	40 %	$ 1,188.4 M
SV rearend st crashes	16,351	329	2.0E-2	16 %	$ 1,261.1 M

** Not including drifting or leaving parking
*** No confirmed injuries associated with the crash
**** Calculated with "KABCO" injury severity FCE equivalent, then converted to 1988 dollars which are then converted to 1999 dollars (factor of 1.408) through information found on
< http://woodrow.mpls.frb.fed.us/economy/calc/cpihome.html >

The dollar amounts are stated in terms of millions of dollars. Also one should note that column 5 indicates that most lane change crashes do not cause serious injury or death.

From these analyses of the crash databases, a description of the "average" lane change crash emerges. For the crash categories described above,

- 76 % of lane change crashes occur on level roadways.
- 94 % of lane change crashes occur on straight roadways.
- 85 % of lane change crashes occur in clear weather.
- 77 % of lane change crashes occur on dry road surfaces.
- 75 % of lane change crashes occur in daylight lighting conditions.
- 69 % of lane change crashes occur in non-junction areas of the roadway.

In less than 8% of the crashes did the driver even attempt any corrective action indicating that the driver was caught by surprise or did not have sufficient time to respond.

Most lane change crashes occurred on surface streets (84%) as opposed to interstate highways or within interchange areas. About 70% of lane-change crashes occurred in non-junction zones. In about 80% of all cases, there were no traffic controls present. This applied equally to junction-related lane change crashes as well as the more frequent non-junction related crashes. As most collisions occurred on regular streets, the posted speed limit in 69% of the cases was found to be ≤ 45 mph.

These findings taken together suggest that it is entirely possible that the threat was never perceived and that most lane change crashes are due to the SV driver's being unaware of the potential crash threat posed by the POV.

The SV driver needs better situational awareness. A CAS that detects the potential threat and warns the SV driver with sufficient lead time to perform corrective action will prevent or mitigate the effects of these collisions. Driver reaction times range from a few tenths of a second to a few seconds. The action (in many cases, aborting the lane change, but sometimes braking, steering, etc.) also requires another increment of time-before-crash to accomplish the avoidance.

6.1 Review of Statistics Pertaining to Proximity Crashes

In the majority of lane-change crashes, the driver seems unaware of the impending collision and takes no evasive action. The small relative speed between the colliding vehicles recorded in the majority of the crashes implies that the vehicles are in close proximity to one another before the lane change begins, and that there is little or no longitudinal gap between the SV and the POV. Proximity crashes in which the closing speed is less than 15 mph, constitute 78 % of the 1992 GES lane-change collisions. For a closing speed of 15 mph, the relative distance closed in a typical human reaction time of 1.0 s is 22 ft, about a car length. This is consistent with our proximity zone that drivers felt safe with in our test of about 30 feet.

6.2 Review of Statistics Pertaining to Fast Approach Crashes

For collisions in which closing speeds are in excess of 15 mph, the term "fast approach crashes" is used. Of all lane-change crashes, 16 % have closing speeds between 15 and 30 mph while 6% involve closing speeds greater than 30 mph. This may indicate that a longer fast approach zone would be appropriate. This contrasts with our measurements of passing vehicles (Figure 5-3) where only 1% of the closing vehicles had relative speeds greater than 30 mph. Our decision to base the maximum extent of the fast approach zone on 30 mph represents a tradeoff between trying to limit the maximum range of the CAS sensor and catching most of the higher approaching speed vehicles.

When one takes into account the severity of the crashes, two speed-related factors stand out as potential contributors to the more severe crashes:

- closing velocities in excess of 35 mph,
- the speed of either vehicle is in excess of 55 mph.

Although the numbers of "fast approach" crashes is less than those of "proximity" crashes, the prevention or mitigation of "fast closing" collisions must be considered because of the severity of the crashes involved.

The manner of collision for most "fast approach crashes" during lane changes is rearend (striking and struck). Of all rear-end collisions that result after a lane-change, 78% have closing speeds in excess of 15 mph. For a nominal vehicle speed of 55 mph and a typical lane width of 12 ft, the time required to execute a lane change ranges from 0.85 s for a vehicle pointing angle of 10 degrees to 8.5 s for a vehicle pointing angle of 1 degree. For closing longitudinal velocities of 15 mph and 30 mph, the times required for the faster vehicle to gain a longitudinal distance equal to one car length (20 ft) are 0.9 s and 0.45 s respectively. From these considerations, the significant time scale ranges from a fraction of a second to several seconds.

6.3 Review of Benefits Models

It is useful to estimate the benefits that would accrue if a lane change CAS was provided on all vehicles on US roadways. From an estimate of the number of crashes avoided or mitigated, one can estimate: savings in lives, reduction in insurance and other societal costs, improvement in traffic flow and the concomitant decrease in pollution, etc. The benefit estimate is usually quoted as the fraction of crashes avoided when the CAS is employed [22,23]. However, difficulties in making this estimate arise from many factors, including uncertainties in how effective the warnings will be, how the driver will react to the inevitable false and nuisance alarms, the performance of the CAS in real-world situations, and mode of usage by the driver.

There are several standard approaches to estimating the CAS benefits. They will be described in the following sections.

6.31 Taxonometric Approach

This approach addresses the types of collisions in a specific class that could be avoided if a CAS were employed. Each subset of the type of collision to be avoided is considered as a class and a determination of the effectiveness of a CAS against that subclass is made. These subsets may differ from each other due to the cause of the crash, the geometry of the crash, or the state of the driver.

For lane change, we might argue that the CAS would be totally ineffective against collisions arising when two vehicles are simultaneously making lane changes (geometrical classification), for drifting crashes when the driver has no intention of changing lanes (causal classification), or for the case where the driver is intoxicated (driver state classification).

6.3.2 Kinematic Approach

For this method the dynamics of various collision types are considered and a calculation of which could be avoided is performed. Here, the geometry of the collision, the dynamics of the vehicles, and the reaction times of the driver and vehicle are all taken into account. Then a determination is made as to what fraction of the collisions can be avoided based on values of the parameters mentioned above. A natural approach to this computation is Monte Carlo (or the related Latin hypercube) where distributions of a number of variables (e.g., driver reaction time, braking level, gap between vehicles) are randomly sampled and a large number of calculations leads to statistical estimates of effectiveness.

It should be mentioned that this approach is very useful to determine the *relative* effectiveness of various approaches to implementing a CAS as we did in Task 4 [4].

6.3.3 Human Usage Approach

Here one looks at how people might use the CAS and then determines what change to the number of collisions would result. A number of papers have been published by Tijerina and his coworkers [12-14] that address the effectiveness of a lane change CAS. In their approach, the effectiveness of the CAS hinges on the way the drivers utilize it. For example, it may have an overall negative effectiveness (meaning that the number of collisions *increases* when it is used) if most drivers blindly rely on it and stop visually checking before they make a lane change.

In this approach, the effectiveness of the driver is estimated from the number of lane change crashes versus the number of lane change attempts, and then that data is used with estimates on the probability that the driver will employ the CAS in one of four modes. Estimates of about 20 to 80% effectiveness result from this approach and there are many caveats about the estimation of a number of parameters.

6.3.4 Collision Surrogates

Extrapolation from changes in the driver error rate during roadway tests to number of crashes avoided can be used to arrive at a benefits estimate. In this approach, road test experience is employed to estimate the benefits of the CAS being tested. This technique is also referred to as the Traffic Conflict Technique.

Since it would be irresponsible to simulate crashes or even near-crashes, field testing of a CAS must rely on indirect evidence of its potential for avoiding crashes. Fortunately, crashes and near-crashes are rare events. This indirect evidence can come from an evaluation of the number of "driver errors" or "conflict situations" encountered with and without the CAS operational.

"Driver errors" or "conflict situations" are defined as lane changes made when another vehicle is present in the proximity zone close enough in the fast-approach zone to trigger a warning as discussed above. Thus, driver errors are used as safe surrogates for near-misses or collisions [24]. Commonly used is Heinrich's Triangle [24] (as shown in Figure 6-1) in which errors and near-crashes are related to collisions.

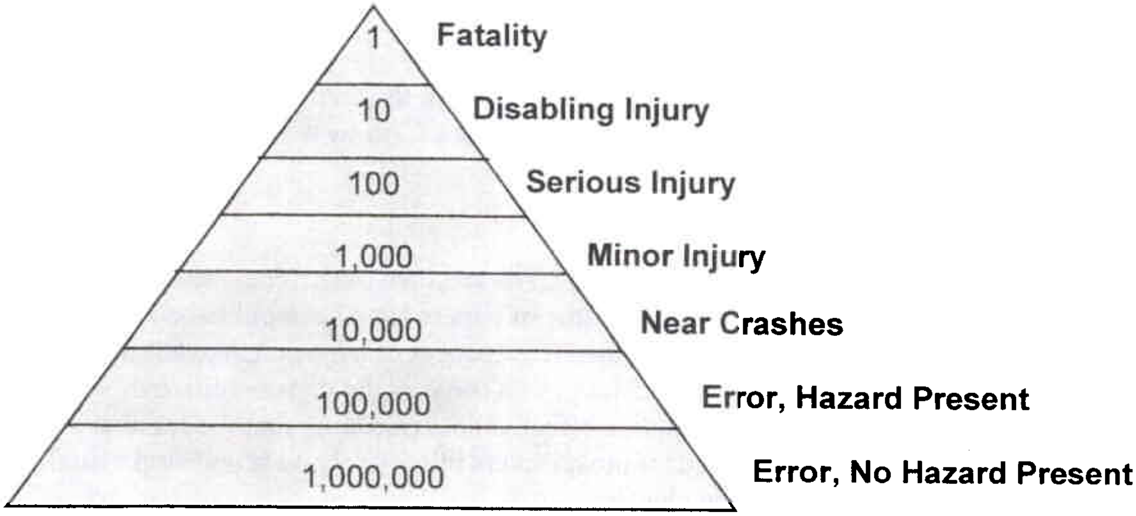

Figure 6-1. Heinrich's Triangle.

6.4 Benefits Analysis

A benefits estimation methodology developed by us is described in Section 6.4.1. The method is then applied to our road test data in Section 6.4.2. Limitations to the methodology are discussed in Section 6.4.3.

6.4.1 Benefits Estimation Methodology

The definition of the effectiveness of a CAS, *E*, may be given by

$$E = \frac{\text{Number of target crashes prevented by the CAS}}{\text{Number of crashes if no CAS existed}}$$

or

$$E = \frac{\text{Total number of crashes without CAS - Total number of crashes with CAS}}{\text{Total number of crashes without CAS}}$$

If the CAS is effective in preventing crashes, then the upper limit of E would be given when using the "total number of crashes with CAS" assuming usage by 100% of the vehicles on the road. We have assumed here 100% usage. The expression above may be rewritten as

$$E = 1 - \frac{(\#\text{ lane changes with CAS}) * \text{probability(lane change with CAS resulting in crash)}}{(\#\text{ lane changes w/o CAS}) * \text{probability(lane change w/o CAS resulting in crash)}} \tag{1}$$

By assuming an equal number of lane changes with and without a CAS, we can focus on the quality of the lane changes accepted. Thus, the effectiveness may be reduced to

$$E = 1 - \frac{\text{probability(lane change with CAS resulting in crash)}}{\text{probability(lane change w/o CAS resulting in crash)}} \tag{2}$$

To examine the structure of the probability terms, interactions between the SV and the POV are modeled, on the most basic level, as two-body problems. Each lane change is characterized by two parameters x_i and v_i:

- x_i = the vehicle longitudinal separation or gap distance between the SV and POV at the start of the lane change, t_i, (x_i is positive if the rear bumper of the SV is in front of the front bumper of the POV) and
- v_i = the closing velocity at the start of the lane change, t_i. (v_i is positive if $v_{POV} > v_{SV}$, i.e., the gap is diminishing or negative if $v_{POV} < v_{SV}$, i.e., the gap is increasing).

The case with no POV present is represented by $x_i = \infty$.

The start of a lane change, t_i, is defined to be the transition point between the decision phase and the execution phase of a lane change as illustrated in Figure 6-2 [12]. The decision phase is the period of time beginning when the driver desires to perform a lane change. It continues until the driver actually starts to move the steering wheel to move the SV laterally into the new lane. During this phase, one of the major activities of the driver is to detect either present or upcoming traffic or obstacles in the planned destination lane. Based upon this assessment, the driver either proceeds to the execution phase or temporarily postpones execution of the lane change. The execution phase begins when the driver starts to make the move into the new lane. It continues until the SV has been laterally stabilized in a lane at the conclusion of the maneuver.

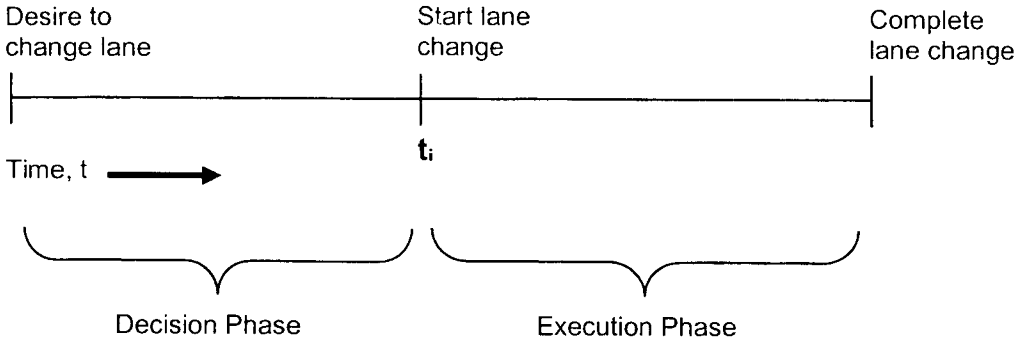

Figure 6-2: Time-line showing the two phases of a lane change maneuver.

The probability of a lane change started by the SV driver at the phase point (v_i, x_i) ending in a crash without a CAS is the result of SV and POV actions:

$$P^{WO}(v_i, x_i) = P_{SV}^{WO}(v_i, x_i) P_{POV}(v_i, x_i)$$

where

$P_{SV}^{WO}(v_i, x_i)$ = probability of SV driver (without a CAS) failure to take adequate evasive actions when required during the execution phase, and

$P_{POV}(v_i, x_i)$ = probability of POV driver failure to take adequate evasive actions when required during the execution phase of the lane change.

Similarly, The probability of a lane change started by the SV driver at the phase space point (v_i, x_i) ending in a crash with a CAS is also the result of both driver's actions:

$$P^{W}(v_i, x_i) = P_{SV}^{W}(v_i, x_i) P_{POV}(v_i, x_i)$$

where

$P_{SV}^{W}(v_i, x_i)$ = probability of SV driver (with a CAS) failure to take adequate evasive actions when required during the execution phase.

Since the POV driver's response depends only on v_i and x_i and not directly on whether a CAS is deployed in the SV, the same expression for the POV driver failure is used in both cases with and without a CAS, and a crash occurs when both the SV and POV fail to take adequate actions. Thus, $P^{W}(v_i, x_i)$ and $P^{WO}(v_i, x_i)$ are the probabilities of a lane change starting at phase space point (v_i,x_i) ending in a crash with and without a CAS, respectively. To account for lane changes at various phase space points made by an individual driver as well as different drivers in the population, the averaged crash probabilities are:

$$< P^{WO} >= \iint L_{WO}(v_i, x_i) P_{SV}^{WO}(v_i, x_i) P_{POV}(v_i, x_i) dv_i dx_i$$

$$< P^{W} >= \iint L_{W}(v_i, x_i) P_{SV}^{W}(v_i, x_i) P_{POV}(v_i, x_i) dv_i dx_i$$

where

$L_{WO}(v_i,x_i)$ = normalized distribution of lane changes for all drivers in the population without a CAS, and

$L_W(v_i,x_i)$ = normalized distribution of lane changes for all drivers in the population with a CAS.

The distributions are normalized such that:

$$\iint L_{WO}(v_i,x_i)dv_i dx_i = 1$$

$$\iint L_W(v_i,x_i)dv_i dx_i = 1$$

Effectiveness of a CAS as expressed in Equation (2) can now be written as:

$$E = 1 - \frac{\iint L_W(v_i,x_i)P_{SV}^{W}(v_i,x_i)P_{POV}(v_i,x_i)dv_i dx_i}{\iint L_{WO}(v_i,x_i)P_{SV}^{WO}(v_i,x_i)P_{POV}(v_i,x_i)dv_i dx_i} \tag{3}$$

In principle the distributions $L_{WO}(v_i,x_i)$ and $L_W(v_i,x_i)$ can be measured directly if extensive road tests are performed. However, our limited test runs did not yield enough lane changes to accurately define the distributions in detail. $P_{SV}^{WO}(v_i,x_i)$, $P_{SV}^{W}(v_i,x_i)$ and $P_{POV}(v_i,x_i)$ may be estimated by Monte Carlo simulations [4] which involve modeling the dynamics of the vehicles and the drivers' reaction. Difficulties with such simulations arise from the high level of uncertainty associated with the model inputs and assumptions. To circumvent these problems, simplifications are made in evaluating the expression in Equation (3).

The first model simplification is the postulation that the main benefit of a CAS is to assist the driver during the decision phase to postpone the execution of potentially dangerous lane changes. In other words, the main effect is the modification of the lane change distribution from $L_{WO}(v_i,x_i)$ to $L_W(v_i,x_i)$. Positive benefits come from moving some of the lane changes that may have started in the potentially dangerous region of the (v_i, x_i) phase plane to somewhere safer. Figure 6-3 shows some examples of postponed lane changes with the aid of a CAS. The top and bottom lines are the boundaries of the warning zone of the CAS. (Remember that the origin is defined at the SV's rear bumper, and therefore the proximity zone extends into regions of negative x_i.) In the lower right region of the figure, the case with a fast approaching POV is represented. Without the CAS, the SV driver may not be aware of the POV and may change lanes right in front of the approaching vehicle. With the aid of the CAS, the lane change may now be postponed after the POV has safely overtaken the SV. The case illustrated in the upper right region of the figure represents the situation when the SV driver postpones the lane change until he has accelerated or the POV driver has braked to reduce the relative speed between the two vehicles. The case illustrated in the left side of the figure represents the situation when the SV driver postpones the lane change until there is a larger separation between the two vehicles.

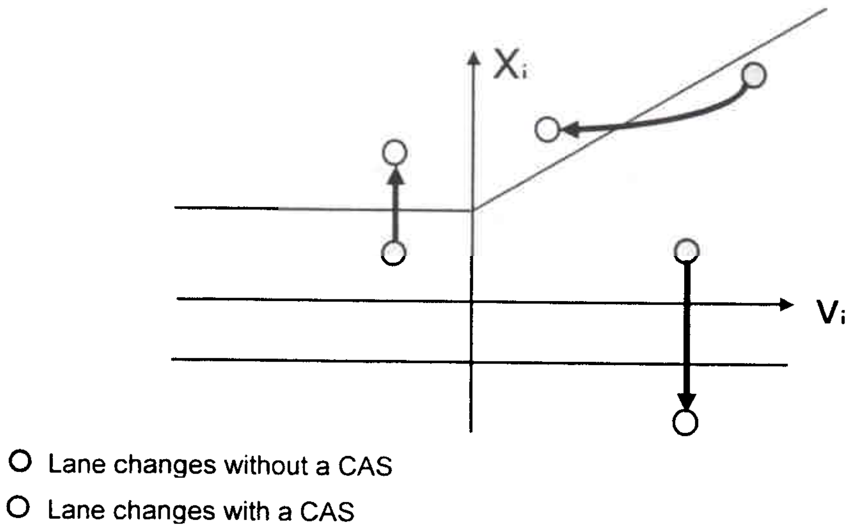

Figure 6-3: Examples of postponed lane changes with the aid of a CAS.

A CAS may also help prevent lane change crashes during the execution phase by warning the SV driver to abort a lane change after he has started maneuvering the vehicle. The benefits of the CAS prior to the lane change decision are contained in the distribution terms, L. However, the benefit of the CAS during the lane change execution is postulated to be small and is ignored in our simplified model (leading to a lower bound on the CAS benefit) yielding:

$$P_{SV}^{W}(v_i,x_i) = P_{SV}^{WO}(v_i,x_i) = P_{SV}(v_i,x_i)$$

The effectiveness in Equation (3) now reduces to:

$$E = 1 - \frac{\iint L_W(v_i,x_i)P_{SV}(v_i,x_i)P_{POV}(v_i,x_i)dv_i dx_i}{\iint L_{WO}(v_i,x_i)P_{SV}(v_i,x_i)P_{POV}(v_i,x_i)dv_i dx_i} \tag{4}$$

Another simplification of the model is to divide the (v_i,x_i) phase plane into two zones, namely the conflict zone and the no-conflict zone as illustrated in Figure 6-4. It is assumed that lane changes started in the no-conflict zone will almost certainly not end in crashes. In other words, either the SV driver or the POV driver would take adequate evasive actions when required if the lane change starts in the no-conflict zone:

$$P_{SV}(v_i,x_i)P_{POV}(v_i,x_i) \cong 0 \qquad \text{in the no-conflict zone}$$

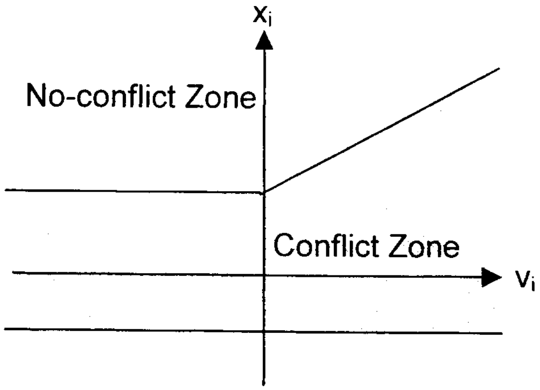

Figure 6-4: Diagram of (v_i,x_i) phase plane showing the conflict and no-conflict zones.

With this assumption:

$$E = 1 - \frac{\iint_{conflictzone} L_W(v_i,x_i)P_{SV}(v_i,x_i)P_{POV}(v_i,x_i)dv_i dx_i}{\iint_{conflictzone} L_{WO}(v_i,x_i)P_{SV}(v_i,x_i)P_{POV}(v_i,x_i)dv_i dx_i} \tag{5}$$

The last model simplification assumes that the lane change distribution with a CAS, $L_W(v_i,x_i)$, is proportional to the distribution without a CAS, $L_{WO}(v_i,s_i)$:

$$L_W(v_i,x_i) = BL_{WO}(v_i,x_i) \quad \text{in the conflict zone.}$$

B<1 represents positive benefits while B>1 represents negative benefits.

With this assumption, the effectiveness in Equation (4) simply reduces to:

$$E = 1 - B \tag{6}$$

Integrating both sides of the equation $L_W(v_i,x_i) = BL_{WO}(v_i,x_i)$ over the conflict zone gives:

$$\iint_{conflictzone} L_W(v_i, x_i)dv_i dx_i = B \iint_{conflictzone} L_{WO}(v_i, x_i)dv_i dx_i$$

$$\Rightarrow B = \frac{\text{fraction of lane changes started in the conflict zone with a CAS}}{\text{fraction of lane changes started in the conflict zone without a CAS}}$$

Substituting the expression for B into Equation (6) gives the effectiveness as:

$$E = 1 - \frac{\text{fraction of lane changes started in the conflict zone with a CAS}}{\text{fraction of lane changes started in the conflict zone without a CAS}} \tag{7}$$

These fractions of lane changes started in the conflict zone with and without a CAS represent the "Error, Hazard Present" level of the Heinrich triangle and can be observed for each driver in the test group, resulting in an effectiveness estimate for each driver, or weighted and summed together to provide an effectiveness estimate for the whole group.

6.4.2 Test Run Results

As derived above, we will use the change in driver behavior as manifested in the fraction of lane changes made under conflict with and without a lane change CAS to estimate the benefit derived from the use of such a CAS. We will do this separately for the comprehensive lane change CAS and the proximity subsystem alone. We expect that the comprehensive lane change CAS will be more effective since it covers more dangerous situations.

For each lane change for all drivers in the driving tests, the (v_i,x_i) point is taken 1 second prior to the wheels crossing the lane definition line. Looking at that time, we have analyzed all the lane changes that were made during our test runs when a POV was present. The data was divided into three segments. The first one includes all the runs made without any functional CAS. In Figure 6-5, we have plotted the separation and relative velocity of the POV. The conflict zone is taken to be identical to the warning zone of the comprehensive CAS. Any point below the heavy line is considered under conflict. Similar plots for all the runs that were performed with either the comprehensive or proximity lane change CAS are shown in Figures 6-6 and 6-7. Not shown on these graphs are the cases for which there was no POV within sight of the sensor system.

We see numerous under conflict lane changes "below the line", where the driver changed lanes in spite of the warning. If drivers can be induced to make safer lane changes with the lane change CAS, then one might infer that when deployed in the field, the number of lane change crashes would decrease.

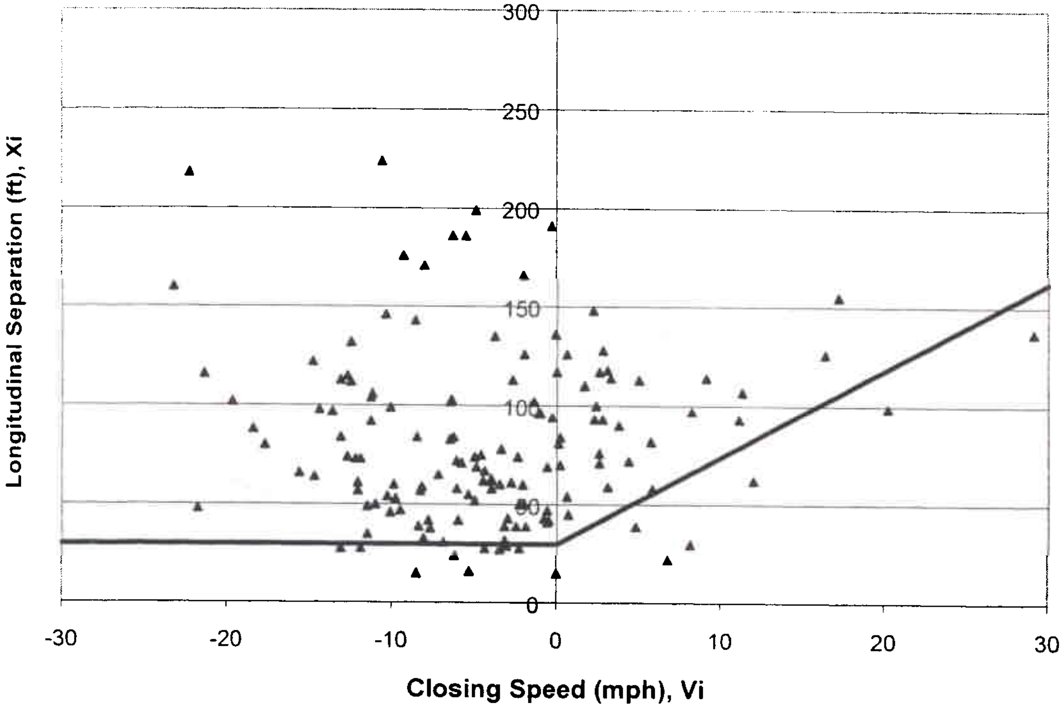

Figure 6-5: Phase space coordinates at 1 sec before lane crossing for all test runs without a CAS.

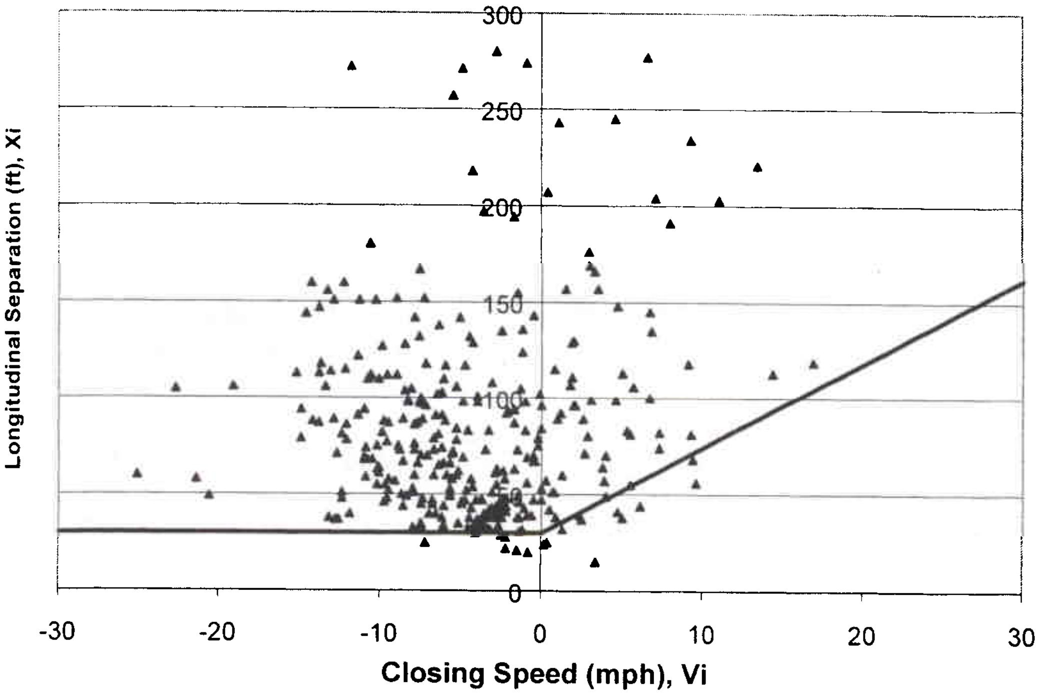

Figure 6-6: Phase space coordinates at 1 sec before lane crossing for all test runs with a comprehensive CAS.

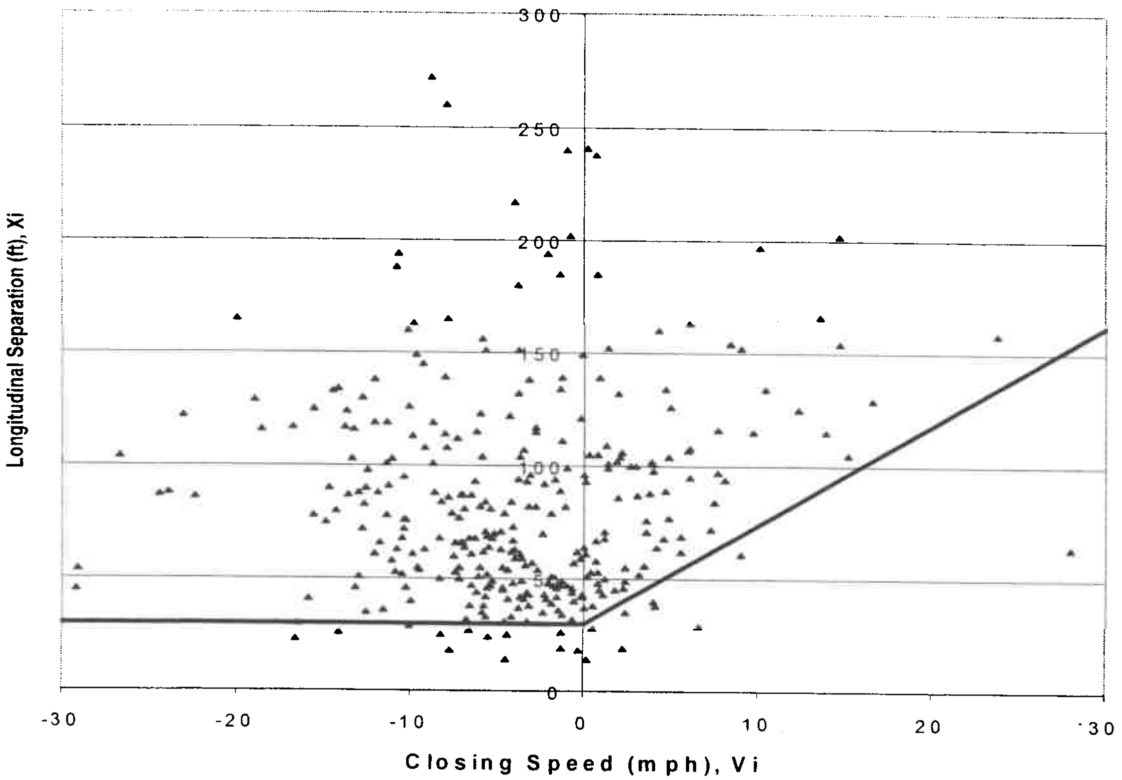

Figure 6-7: Phase space coordinates at 1 sec before lane crossing for all test runs with a proximity CAS.

The numbers of under conflict lane changes are presented in Table 6-1. Note that 5 of the 22 lane changes under conflict with the proximity lane change CAS involved fast approaching vehicles outside of the proximity warning zone which extends 30 ft behind the SV. In those cases, the drivers were not warned by the proximity lane change CAS.

	No CAS Runs	CAS Runs	
		Comprehensive	Proximity
Lane changes under conflict	16	18	22
Total number of lane changes	273	541	552
Fraction of lane changes under conflict	0.059	0.033	0.040

Table 6-1: Tabulation of conflict/non-conflict totals for all right lane changes.

The expression derived in Section 6.4.1 can now be used to estimate the effectiveness. Substituting into that expression yields:

$$E_{\text{comprehensive}} = 0.43$$
$$E_{\text{proximity}} = 0.32$$

Since the numbers of under conflict lane changes are relatively small, the uncertainty in these calculations could be large. It is, however, reassuring to see that the CAS makes a positive effect on the drivers' behavior and that the comprehensive CAS is more effective than the proximity one alone. These promising results warrant further testing.

The statistical error, σ, associated with the measured number of lane changes under conflict, N, is given by:

$$\sigma = \sqrt{N}$$

The percentage error is:

$$\frac{\sqrt{N}}{N} x100\% = \frac{1}{\sqrt{N}} x100\%$$

A small N would therefore have a large percentage error. For example, the statistical error associated with the 16 lane changes under conflict for the no CAS runs is 25%. An error analysis gives:

$$E_{\text{comprehensive}} = 0.43 \pm 0.20$$
$$E_{\text{proximity}} = 0.32 \pm 0.22$$

6.4.3 Model Limitations

Limitations of the benefits estimation methodology derived in Section 6.4.1 are discussed in this section with reference to our test run data.

The first assumption we make in our model is that the number of lane changes is the same whether a CAS is in place or not. We did not check this assumption against our test data since one of the two no CAS runs for each test subject had a shorter than standard route. The change in route made direct comparison difficult. Despite the lack of direct experimental evidence, we feel that this assumption is a valid one. However, if future tests prove otherwise, the effect of the change in the number of lane changes when a CAS is used can be accounted for in the model by keeping the two number of lane change terms in Equation (1) and carrying them through the analysis.

It is also assumed in our model that a CAS has little effect during the execution phase of a lane change. However, it is probable that the crash avoidance ability of a SV driver during the execution phase may improve if a CAS alerts him of the POV that he may otherwise be unaware of. On the other hand, it is also possible that the crash avoidance ability during the execution phase may suffer if the SV driver does not bother to monitor the situation during the execution phase himself but depends solely on a CAS that is not

100% reliable. The effect therefore depends on how a CAS is used and also on its reliability.

In our model, it is assumed that $L_W(v_i,s_i) = BL_{WO}(v_i,s_i)$ in the conflict zone and B is determined from the test run data as the fraction of lane changes started in the conflict zone with a CAS divided by the fraction of lane changes started in the conflict zone without a CAS. As can be seen in Figures 6-5, 6-6 and 6-7, most of the lane changes started in the conflict zone in our test runs were close to the zone boundary and there were few lane changes made near the region shown schematically in Figure 6-8, with or without a CAS. However, the shaded region in Figure 6-8 is probably the place where a significant number of lane change crashes start. Without significant number of data points near this region, we can only extrapolate the effect of a CAS on the lane change distribution from elsewhere in the conflict zone to this more critical region of the phase plane. More extensive road tests should be able to provide data in this region which corresponds to the "Near Crashes" level of the Heinrich triangle.

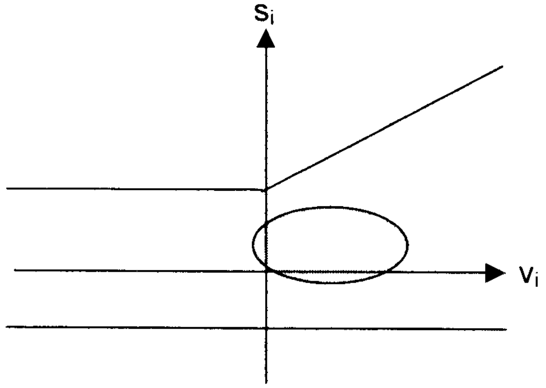

Figure 6-8: A critical region inside the conflict zone.

None of the discussions above consider benefits derived from mitigating the collision. Even with the improved situational awareness, the lane change driver may not be able to prevent the crash. Nevertheless, corrective or evasive actions can lessen the injuries and damage which occur in the crash. Modeling this kind of effect is possible using techniques, such as Monte Carlo calculations, which draw on experimentally derived distributions describing the traffic states.

6.5 Costs

It is too early to predict the cost of building a lane change CAS. A report on intelligent cruise control [25] utilized costs of existing devices in limited production and learning curve arguments to predict costs when large market penetration is achieved. There are no comprehensive lane change CAS products available to the best of our knowledge that would allow us to perform the same analyses. We will instead look at potential cost benefits per vehicle associated with an effective lane change CAS as a means to estimate reasonable cost.

Following [26], we can estimate the total cost to society of all lane change crashes in a given year. Based on Table 6-3 the total cost to society (in 1999 dollars) of lane change crashes that can be addressed by our system is $5.2B. Using our estimated effectiveness of 43% (not far from our estimate) leads to a cost benefit of $2.2B per year if all vehicles were equipped with a lane change CAS.

To compute the value per vehicle of a lane change CAS also requires the number of vehicles per year produced and an estimate of the discount factor [26]. A discount factor is utilized to account for the fact that the cost of the lane change CAS is paid for when the vehicle is purchased, but the benefit is attained when a crash is avoided. The discount factor covers the potential gains the purchaser could have accrued if the CAS price had been invested and then the crash was paid for when it occurred. The authors estimated a discount factor of 0.8218.

In [26], they use 14.5 million for the number of new vehicles sold in the U.S. per year. Discounting the potential cost savings of $2.2B and dividing by the number of vehicles sold per year, leads to a value of $126. This is the amount of money saved per vehicle if every vehicle sold was equipped with a lane change CAS. We cannot envision a comprehensive lane change CAS where the cost to the driver would be that small within the next decade.

This analysis is discouraging if taken at face value. Since covered lane change crashes account for less than 4% of all crashes and since lane change crashes are usually less severe in terms of damage and injury, they only account for a small fraction of the total cost of all crashes. Based on cost savings only, a lane change CAS would never be cost-effective. However, cost savings are only a part of the equation. Many drivers typically pay around $200 for conventional cruise control for the added convenience even if they use it only occasionally. The overall benefit of a lane change CAS to the driver must be evaluated. This includes cost savings, convenience, perceived safety, and peace of mind.

In a recent study [27] in which eight focus groups were questioned about the desirability of crash avoidance devices, backing and lane change CAS were preferred quite strongly over rear end and running off the road CAS. The cost benefit of the preferred systems is only a small fraction of the other two. Clearly, the people surveyed were using other criteria when stating their preferences. One can speculate that drivers find making lane changes and backing up more stressing than regular driving because of the limited visibility. Avoiding rear end and running off the road collisions involves paying attention to the roadway in front of the driver and may seem less stressing. Regardless of the reason, the perceived benefits of a lane change CAS were strong in these groups and that must be taking into account when determining “what the market will bear” in terms of lane change CAS cost to the consumer. It is too early to predict a price for a comprehensive lane change CAS based on existing technology and too difficult to do it based on real and perceived benefits.

This page intentionally left blank

7. SUMMARY AND CONCLUSIONS

When lane change collisions occur, the driver usually has no awareness of the danger. This is because he or she "looked but did not see". A lane change CAS will also be "looking" and, if used properly, will require that both the driver and the CAS miss the threat for a collision to occur. The lane change CAS will represent another set of "eyes" watching the area around the SV and alerting the driver to potential conflicts. Seen in this way, a lane change CAS would be a very attractive addition to future vehicles. As mentioned in Section 6.5, a previous study showed that a number of consumers already see it as highly desirable. This report has presented guidelines that when followed will lead to an effective lane change CAS.

These guidelines were developed, to a large part, based on the experiences we had during a number of naïve driver test runs with our testbed. We implemented a functional lane change CAS utilizing today's technology. This was done to demonstrate utility and to perform realistic testing. Our implementation is not the only acceptable one and it is certainly not the least expensive. Nevertheless, it demonstrates that an effective lane change CAS can be built today.

In order to avoid lane change collisions, the lane change CAS must monitor the areas on either side of the vehicle to determine the presence of another vehicle that could interfere with a planned lane change. It must also determine if a vehicle is approaching those areas with enough relative speed to potentially be in conflict with the instrumented vehicle. Those two tasks must be accomplished in any and all driving environments. This requires the lane change CAS to reject some targets in order to limit nuisance alarms.

By using driver error as a surrogate for collisions, we have been able to estimate the effectiveness of a lane change CAS. We arrived at 43%, a number consistent with previous estimates found in the literature. Although the potential cost saving per vehicle of $126 is relatively small for a lane change CAS, eliminating crashes is still a worthwhile endeavor, and one we believe that the public will want and will pay for.

We have acquired a large set of useful data with our testbed with relatively few test drives. Supplementing that data with many more test runs would be helpful in further defining the guidelines for a lane change CAS and also in obtaining valuable information that could support a number of other investigations and eventually lead to the enhancement of overall driver safety.

As a result of our work we have identified a number of areas for future research. Current data did not support the choice of situational awareness (monitor mode) versus turn-signal-activated mode. There is some indication that turn signal mode may lead to insufficient warning of a hazard.

In order to make the benefits estimation tractable, a number of simplifying assumptions had to be made. Primary among these is that these short term tests are meaningfully

This page intentionally left blank

correlated with long term driver behavior. Performing these types of tests over long periods (at least one week) would allow one to study the learning curve as well as examining long term adaptability.

The precision of the benefits estimation was seen to be dependent on some very small number of lane changes in conflict. A longer term study would allow for better statistics. Conflict situations are still considerably removed from crashes, but longer term studies might uncover pre-crash events from a number of near-crashes.

This page intentionally left blank

Appendix A – Testbed Test Results

A-1.0 Introduction and Executive Summary

This report represents the documentation of the testing of a lane change collision avoidance system and the human subject testing carried out in support of the development of performance specifications for a lane change (CAS). Although a design report had been previously issued [1] this report goes into extensive detail concerning the performance of each subsystem as well as the total system performance.

A government furnished 1993 Chevrolet Caprice was transformed into an instrumented testbed designed to acquire dynamic data in the area surrounding the vehicle as it traveled, eye glance data from the driver and vehicle data concerning speed, acceleration, steering etc. The subsystem collecting the dynamic data was centered around a scanning laser rangefinder. A subset of this data together with a subset of the vehicle data formed the basis by which a collision warning algorithm, based in an on-board digital signal processor (DSP), made decisions as to whether or not to warn the driver of a potential hazard. Special warning displays were fitted on to the car by the Vehicle Research and Test Center in East Liberty, OH. They consist of red triangles in the mirror that are only visible when lit. Otherwise, the mirror looks identical to an ordinary mirror.

The testbed was first run through a series of system validation tests. The accuracy of the range and speed determination was measured by reference to differential GPS. It was found that under most conditions the range accuracy to a vehicle closing from the rear while the testbed was also in motion was between 1 and 2 feet. Speed determination was to within 1 mph. In practice the greatest factor limiting range determination had to do with pointing. For targets at large distances (> 150 ft) the motion of the testbed would effectively scan the laser in the vertical direction. Under normal driving this might result in a vertical scan of about 1-2 feet. On a small car a variation of this amount could shift the point where the laser strikes the car from the front grill to somewhere on the front hood. This would effectively vary the range by as much as 5 feet. Under conditions where the road would vary in elevation, or the car hits a bump, the laser could wind up pointing momentarily at the ground or at the sky.

Pilot testing was done with members of the project to assess the parameters of the driver warning algorithm. In particular we were looking to determine how far back in the adjacent lane from the rear of the testbed, should the proximity warning zone extend. Also we wanted to ascertain how much warning time we should give the driver before a fast approaching car in the adjacent lane enters this proximity zone. A series of tests were performed at two different test tracks and on Los Angeles freeways. Drivers were told to indicate the "last moment that I would be willing to change lanes" in front of a fast approaching confederate vehicle by pressing a button that marked the time in the data file. The tests showed a number of interesting results. Data from the controlled track tests showed no dependence on subject vehicle speed in determining the amount of warning time desired. It depends solely on the relative velocity in a linear fashion. The slope of that line corresponds to the warning time while the intercept corresponds to the backward extent of the proximity zone, discussed above. Ultimately what was presented to the naïve drivers was a 3 second warning and a 30 foot proximity zone.

A series of tests were run with 12 naïve drivers. It was decided to have three independent variables. The first was the location of the warning display. The display could be located in the center mirror, right side mirror or both mirrors. The second variable was the mode of presentation of the warning. One mode was monitor mode, wherein a warning was always given if warranted. However, when the driver put on the turn signal indicator, the steady illuminated red triangle would change to a blinking triangle. The other mode was turn signal activated mode wherein there would be no warning at all unless the driver put on the turn signal indicator. In this case the triangle would blink. Finally the last variable was the system variable. We could warn only on the presence of another vehicle in the proximity zone or we could add to that the ability to warn on fast approaching vehicles. The latter configuration is called the comprehensive system.

We found that in studying the dynamics of lane changes that all the lane changes could be classified into one of five categories. The first is making a lane change where there is no one behind you, or at least so far back as to be non-threatening. The second is ostensibly the same as the first with the proviso that the subject vehicle driver has just allowed another vehicle to pass and then pulled in behind that vehicle. The third is where the subject vehicle passes another vehicle and then pulls in front of that vehicle. Number four is where the subject vehicle pulls in front of a fast approaching vehicle forcing some rapid deceleration of that vehicle, possibly coupled with acceleration of the subject vehicle. Finally the last category is where the speeds of the subject and principal other vehicle are roughly comparable and the subject vehicle changes lanes in front of that vehicle. All five of these categories have distinct trajectories in longitudinal separation-closing velocity phase space.

Clearly the desired result of any collision avoidance system is that it reduce crashes. It is rather difficult to project benefits to the population at large from limited studies but we have developed a methodology that allows us to project benefits from the measurements that we took. In particular since a lane change system is designed to influence the decision making of the driver, the measure of those decisions is the movement of conflict lane changes to non-conflict lane changes. We define conflict in this context as a lane change plotted in phase space one second before crossing the lane marking that lies in the warning region of the driver warning algorithm. Clearly if we can influence the driver to make safer lane changes, this should reduce crashes. By this measure our data shows an effectiveness for the comprehensive system of 43%, and for the proximity system an effectiveness of 32%. The exact values are subject to considerable (~50%) error because of the small number of lane changes made in conflict. Studies of much longer duration will be required to refine these numbers, but the main importance here is the methodology.

In trying to asses the anticipated closing speeds of other vehicles in order to set the maximum range for the sensor system, we tabulated the relative speeds of every vehicle that passed the testbed on either side. It was seen that accommodating a 30 mph passing speed would account for 99% of all passing vehicles. For a three second warning and a 30 ft proximity zone, this sets the range to about 160 ft. One implication here is that a proximity system alone, accommodating perhaps only 5 mph passing speeds only accounts for about 30% of all the passing vehicles. This is clearly an area that needs further definition as to the implications of not warning on the balance of those vehicles. On the questionnaires that the drivers completed, almost half the drivers said that they preferred a proximity only system. Anecdotally it was mentioned by a couple of the drivers that "you just tell me what's in my blind spot, and I'll handle the rest". Whether driver behavior will compensate for the lack of a fast approach warning remains to be seen.

If the drivers did show their ability to be influenced by the system then it should also show up in their eye glance behavior. We did, in fact see that the driver's attention during lane changes to the mirrors when they were lit up is increased over the situation where there was no warning system engaged (baseline).

By selecting a few runs for detailed analysis, we were able to derive the following performance statistics for the collision warning system.

Detection Probability	99.3%
Rejection Ratio	90.7%
False Positive Alarm Rate	42 / hr

The figure for detection probability is almost certainly a lower bound since the runs selected for analysis were ones in which the test conductor noted possible false negatives. The rejection ratio is the fraction of stationary objects that were correctly rejected for warning by the algorithm. These include objects such as parked cars, poles, walls, etc. The average false positive alarm rate of 42/hour obscures the fact that during normal driving these typically come in bursts such as when driving next to a row of bushes. During freeway driving, there are almost no such alarms.

Given such an alarm rate, one would think that drivers would show significant annoyance. On only 3% of the runs did drivers indicate any annoyance in their response to the questionnaire. In fact on only 1 out of 3 runs did the drivers even notice any false positive alarms at all. This is almost certainly due to the unobtrusiveness of the warning, the fact that people can ignore it if they do not want the information and that the drivers felt that the system held great utility for them. If we multiply the false positive alarm rate by the fraction of runs wherein the drivers indicated some annoyance (0.03) we get the true nuisance alarm rate of 1.3/hour.

Analysis of the drivers' preferences shows an almost even split between having the display in both the center and side mirrors versus the side only. Again there was an almost even split (60/40) between monitor/turn signal mode and between comprehensive/proximity system (60/40). Although the drivers felt that in 9 out of 10 runs, the system gave them adequate warning of potential hazards, in those cases where they felt otherwise, almost 80% of the time the system was in turn signal mode. This point bears further watching since if the use of turn signal mode has safety implications, then that is reason enough to exclude that mode from being implemented in a lane change collision warning system.

There are a number of intriguing preliminary conclusions as a result of this work. However the statistics upon which these conclusions are based is marginal. This alone suggests the need for further, more extensive work in this area. Perhaps more importantly, this study makes no statement about long term adaptability. Only by studying the convergence of an individual's habits when using a lane change collision warning system, can one make reliable judgements as to an optimal system and the likely effectiveness of that system in preventing crashes.

A-2.0 Testbed Description

This chapter will describe the testbed in detail. It is broken up into broad categories of hardware and software, with introductory remarks covering the vehicle itself.

All the hardware was mounted in a USDOT furnished 1993 Chevrolet Caprice. A photograph of the vehicle is shown in Figure A-2.0-1. The scanning laser rangefinder, which serves as the sensor for the CAS system, was mounted on the right rear corner of the car. This affords the laser an almost 270° field of view. As can be seen in the photograph, the rear doors and the rear window are coated with a dark film. The purpose of this is twofold. The first is to limit the sunlight that enters the car and the second is to hinder onlookers in other cars from seeing the equipment and operator in the rear seat and thus distract them from their task of driving. There were two externally mounted, in addition to the two internally mounted, cameras to record the scene. One was near the front wheel well and pointed backward, while the other was mounted underneath the laser scan head and was pointed forward. These two cameras were both fitted with telephoto lenses. The interior cameras, fitted with wide angle lenses, monitored the right side of the testbed and the area behind the car.

Figure A-2.0-1: Side View of Testbed

An operator, whose function it was to set up and monitor the operation of the data acquisition system, rode along in the back seat. The operator station is shown in a photograph in Figure A-2.0-2. The TV monitor to the left of the operator served two functions. At the beginning of each run, it was used to set up the eye-tracker. During the run it was used to monitor the combined image of the four scene cameras. The operator switched images by front panel controls on two quad combiners located underneath the TV monitor. The data acquisition computer monitor was mounted directly in front of the operator. The display shows the scene dynamics as sensed by the laser rangefinder. In addition, the status of all the vehicle mounted sensors plus the output of the eye-tracking system were displayed.

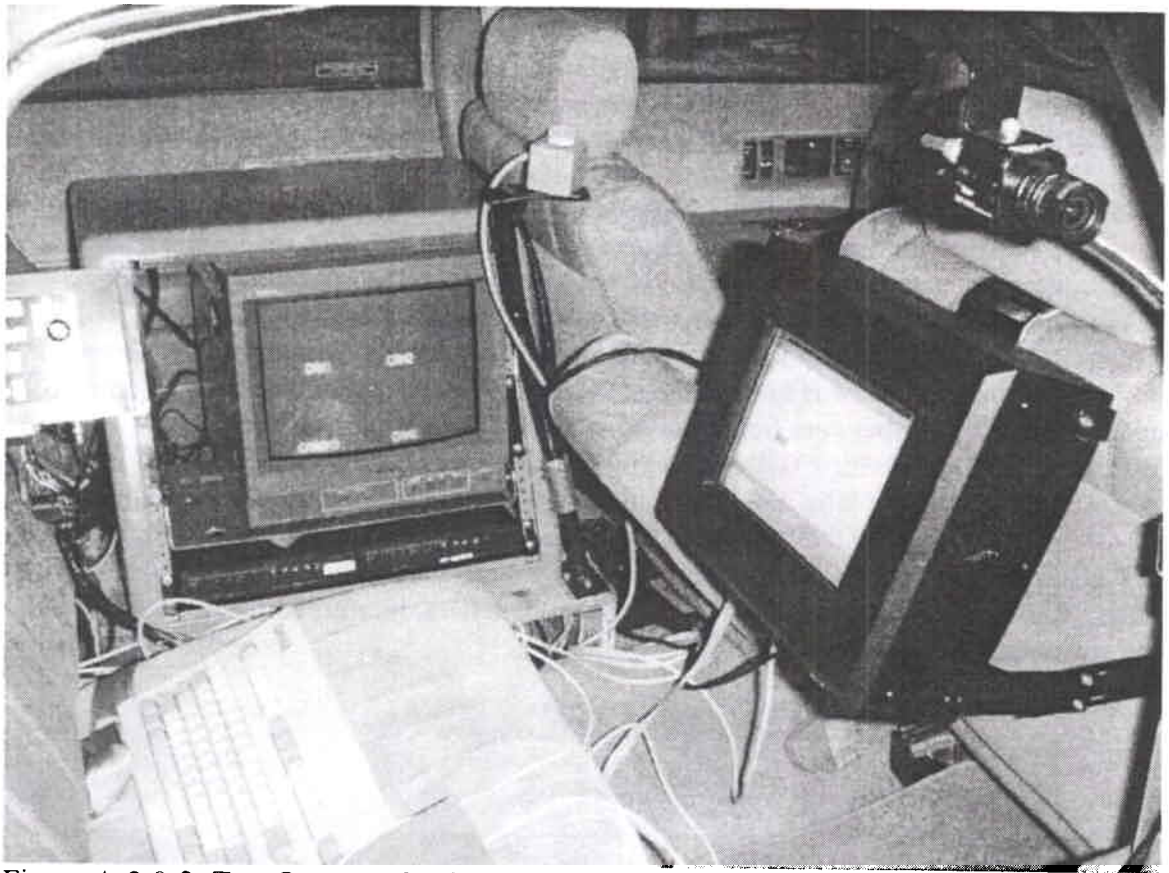

Figure A-2.0-2: Test Operator Station

The testbed was fitted with two types of warning displays by the Vehicle Research and Test Center of East Liberty, OH. Figure A-2.0-3 shows the area between the center mirror and the right side mirror. The mirrors are in fact Muth mirrors which allow for an icon, which in this case is a red triangle, to be displayed on command. When the display icon was not illuminated, the mirror was indistinguishable from any ordinary mirror. The brightness of the display was adjustable. The center mirror had two triangles; one for right lane changes and one for left lane changes. On the inboard side of the right side mirror was an alternative display of a row of LEDs. The alternative display for the center mirror was simply a triangle symbol mounted above the mirror and independently illuminated. For these tests, only the Muth mirrors were used. Since we monitored only the right side of the testbed, only lane changes made to the right were warned against. The left side icons were never illuminated.

Figure A-2.0-3: View from Driver's Seat of Warning Icons in and Next to Mirrors

A-2.1 Scanning Laser Rangefinder

The scanning laser rangefinder used was a Riegl LD-90-3100-GF/HP. A picture of the laser mounted on the testbed is shown in Figure A-2.1-1. The laser system was in fact a Riegl laser rangefinder, with fiber optic output, mated to a scanning mechanism, manufactured by K^2T, Inc. of Pittsburgh, PA. This company is an outgrowth of the Carnegie-Mellon Robotics Institute. The fact that the scanning mechanism was remote from the laser and its electronics allows for the smallest package possible to be mounted on the exterior of the testbed. The housing was painted the same color as the car to minimize its conspicuity. The rest of the system was fixed inside the trunk of the testbed, as shown in Figure A-2.1-1. The block diagram of the laser system is shown in Figure A-2.1-2. The host computer system controls the laser function via an RS-232 port while the data was sent to the DSP via an enhanced parallel (ECP) port. The specifications for the laser are summarized in Table A-2.1-1.

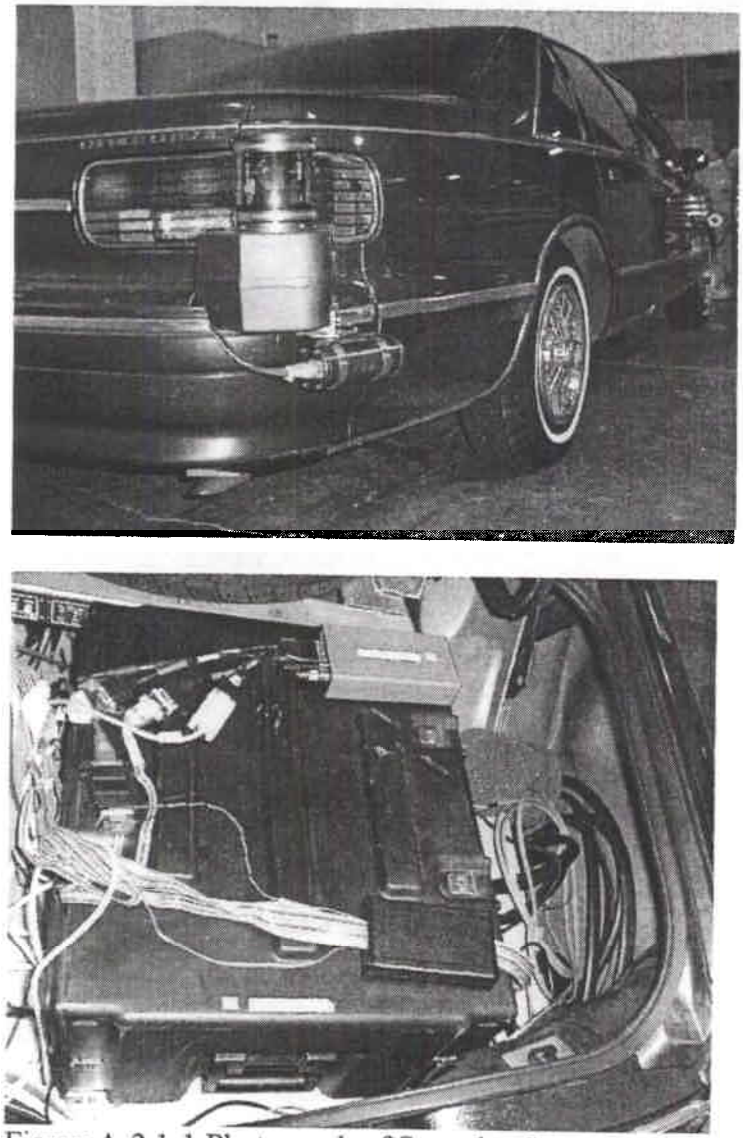

Figure A-2.1-1 Photograph of Scanning Laser Rangefinder and Associated Electronics

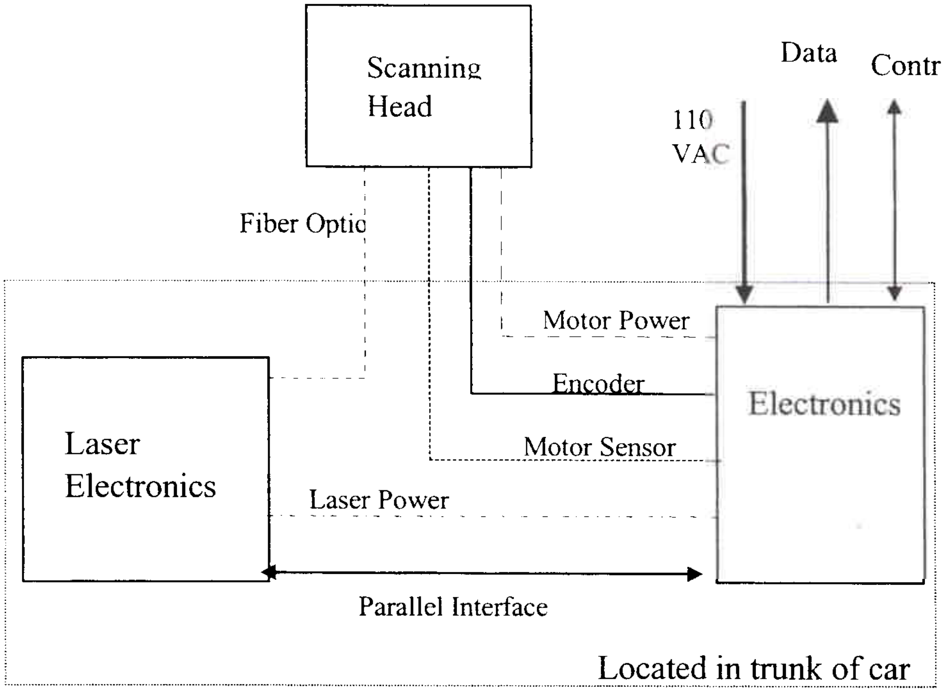

Figure A-2.1-2. Laser system block diagram

Parameter	Specification
Range	1m to > 80m
Accuracy	4 cm typical, 10 cm worst case
Sample Rate	12 kHz max.
Horizontal FOV, scan rate	360°, ≤ 40 Hz (software selectable)
Vertical FOV, scan rate	± 15° determined by hardware changeout, 6 Hz max.
Wavelength	900 nm
Peak Power	approximately 28 W
Pulse format/width	Pulse train / 12 nsec
Eye Safety	Class 3B, eyesafe when scanned

Table A-2.1-1 Laser Specifications

The laser beam was delivered to the scanning mechanism via fiber optic (the orange colored cable in Figure A-2.1-2). A second fiber optic carried the laser return to the detecting optics. In the scanning mechanism were upcollimating optics for both the transmit and receive channels. The exit diameter of the beam was about 1.25 inches. The beam hit a fixed 45° mirror that redirected it vertically, and then another 45° mirror that was rotating so as to scan 360° in azimuth. By changing a cam in the scanning mechanism, one could vary the extent of the vertical scan. Some experimentation was done with a cam that would scan from 0 to 10 degrees. This was soon eliminated because at distances in excess of 25 feet, the laser was pointing above the roofs of most cars. As a result, a simple line scan was used.

The laser, as delivered, did not meet the range specification. In particular, there appeared to be an exclusion zone around the laser, such that no object could be detected below 7 feet, and objects with low reflectivity, such as a car struck by the laser at grazing incidence, would disappear below 10 to 12 feet. The source of the difficulty was the fact that all the optics located in the scanning mechanism reflect some radiation back toward the receiving fiber optic/detector. Insofar as the laser rangefinder is concerned, it appeared as if there was an object at a range of a few centimeters. This reflected light masked the returns from close-in targets. Only after the range was increased to that of one-half the distance covered by light during a pulse width of approximately 13 nsec, would the object appear. In addition it was noticed that the number of returns, both at close distances and far distances, was much less with the provided window than without. Because the scanning head was mounted outside the car, it was clearly necessary that some window be used. Therefore tests were performed that focused on the effect of the window material.

A test was done to determine the cause of the discrepancy. The experimental configuration was as shown in Figure A-2.1-3. The laser was positioned next to another car such that the relative orientation was as if the other vehicle was about to pass in the adjacent right lane. The laser was oriented downward in the vertical direction so that it would hit the ground about 80 feet beyond the other car.

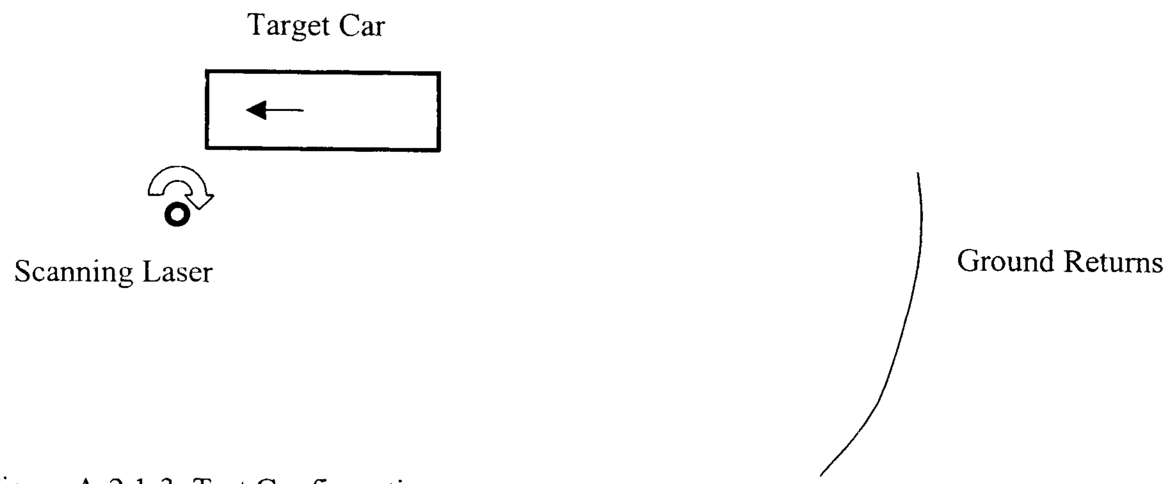

Figure A-2.1-3: Test Configuration

The target car subtended 60.75° in the azimuthal direction, while the ground sample was 54° in extent. Sampling at a rate of 12 kHz and rotating at 10 Hz meant that there can only be a maximum of 202 returns from the car and 180 returns from the ground. Three different window materials were used and compared to the case with no window. Approximately 180 scans were averaged to obtain the data listed in Table A-2.1-2.

Material	Car Returns / scan	Ground Returns / scan
No window	114.8	125.1
1/8" plexiglass (supplied)	0.4	33.4
1/16" acrylic	2.1	56.4
10 mil red filter (mylar)	24.7	104.8

Table A-2.1-2: Number of laser returns with different window materials.

Neither the returns from the car, which extended partly into the exclusion zone, nor from the ground approach the theoretical maximum. This is not surprising since small variations in the surface can significantly affect the strength of the return. However, it was very clear that all the windows had a significant effect on the number of returns. Each of the window materials was tested for transmission efficiency using a spectrophotometer. Transmission as a function of wavelength from 750nm to 1000nm was measured. Surprisingly, all of the materials exhibited a flat response over the interval and had the same transmission efficiency of approximately 90%. Focussing on the number of ground (distant) returns for the red filter, we see that the fraction of returns compared to no window material is about 84%. Realizing that the path of a return signal involves two passes through the window, this is almost exactly the fraction expected from two passes through a 90% transmissive material. The fact that the other materials have significantly less returns compared to the null case (i.e. no window) suggests strongly that for distant returns at least, the thickness of the material is significant. The reason for this is that a thick material that is curved into a cylindrical shell acts as a diverging lens, with a focal length f, given by the relation

$$f = \frac{R(R+t)}{t(n-1)}$$

where R is the radius of curvature, t is the thickness and n is the index of refraction. For an index of refraction of 1.58 for plexiglass and a radius of 2.5 inches, the 1/8 inch window cover that was supplied has a focal length of approximately 86 inches. This is sufficient to mismatch the incoming returns with the receiver collecting optics and therefore reduce the strength of the returns and thereby the number. Measurements were made of the beam size with the scanner held immobile. Using a special infrared viewer, the beam was made visible. At a distance of 20 feet the beam expands in the horizontal direction from 1.5 inches to 4.0 inches, confirming the fact that the plexiglass cover acts like a diverging cylindrical lens.

The situation at close range is in fact closely related. We see from the table above that the number of returns per scan for the closer target is even more strongly correlated with window thickness than for the one at the farther range. While in the case of the ground returns, the optical mismatch tends to reduce the return signal strength in comparison to the detector noise, for close in returns the comparison is not to the noise but to the returns from the optics in the laser scan head. If the return from a real target is strong enough then the laser software, which looks at the rising edge of the laser pulse, will interpret that return as a new object. If however that return is weak then the software will think that it is part of the slow decay of the laser trailing edge and ignore it. It has been observed that diffuse scatterers, such as an individual's clothes, placed normally incident to the beam will be identified as a real target at 6.5 ft. A car, with its highly reflective surface will actually direct most of the radiation away from the laser, unless it hits exactly at normal incidence and is essentially retro-reflected.

The result of these tests was to construct a new window "frame" and substitute the red filter for the plexiglass. An unforeseen advantage of using this red filter is that it obscures the rotating mirror that can attract the attention of passing drivers.

While doing these tests, it was noticed that the line scanning mechanism had in fact a small vertical excursion. Pointing the laser at the ground at a distance of about 70 feet away, and plotting the closest point within a narrow azimuthal range results in Figure A-2.1-4. The laser itself was 37.5 inches above the ground at the scanning mirror.

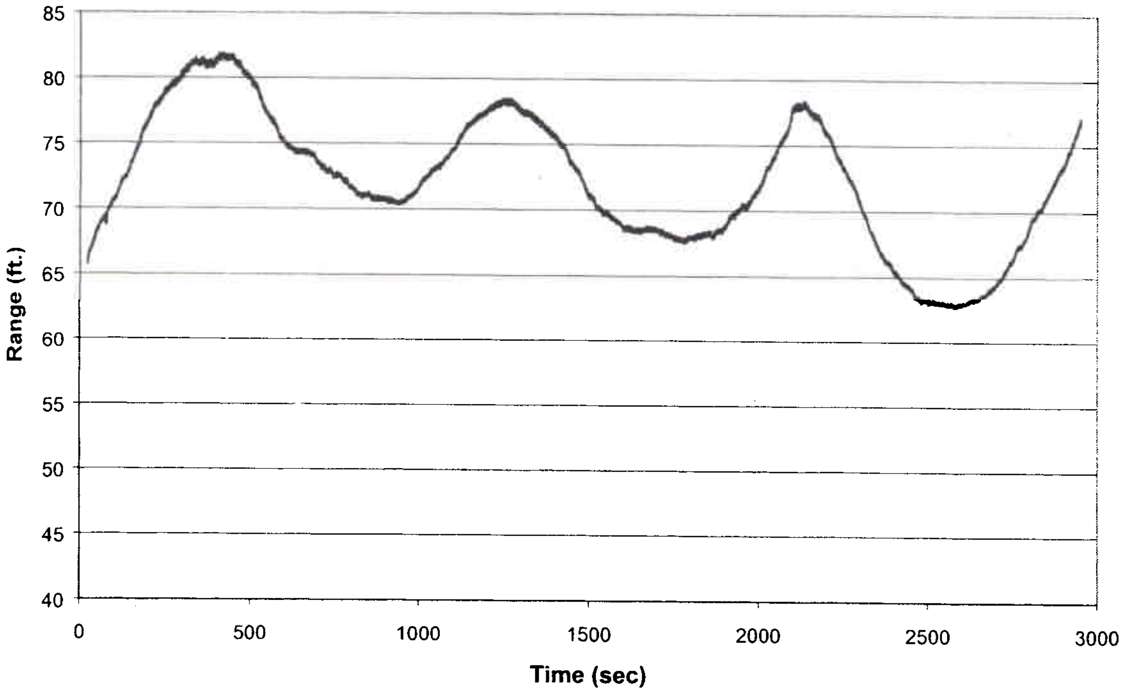

Figure A-2.1-4: Period and extent of vertical scan of the laser

The data indicates that the laser scans vertically with a period of about 15 minutes and a maximum extent of 0.6°. This needed to be factored into consideration when deciding how to align the laser. The criteria used was that when the laser was at the bottom of its scan, it should be aligned with the bumper of a target car positioned about 200 feet away. In principle the laser will scan up from here and maximize the amount of time spent "painting" the car. In practice the motion of the testbed (in the vertical direction) and the geometry of arterials and freeways meant that the vertical excursions were probably considerably larger than that determined statically.

The next test prior to actual use of the laser system was range calibration. A flat target, oriented perpendicular to the laser beam was moved along a tape measure laid along the floor. The returned values of the laser (as determined by the firmware supplied by the vendor) was compared to the tape measure. The result is shown in Figure A-2.1-5. There is a distinct non-linear trend at close range which can be seen more clearly when the difference between the laser measurement and the tape is plotted, as in Figure A-2.1-6. Also in the latter figure we show a piecewise linear fit to the calibration data that we use in interpreting the data from the laser, both in real time in the testbed and in post-processing.

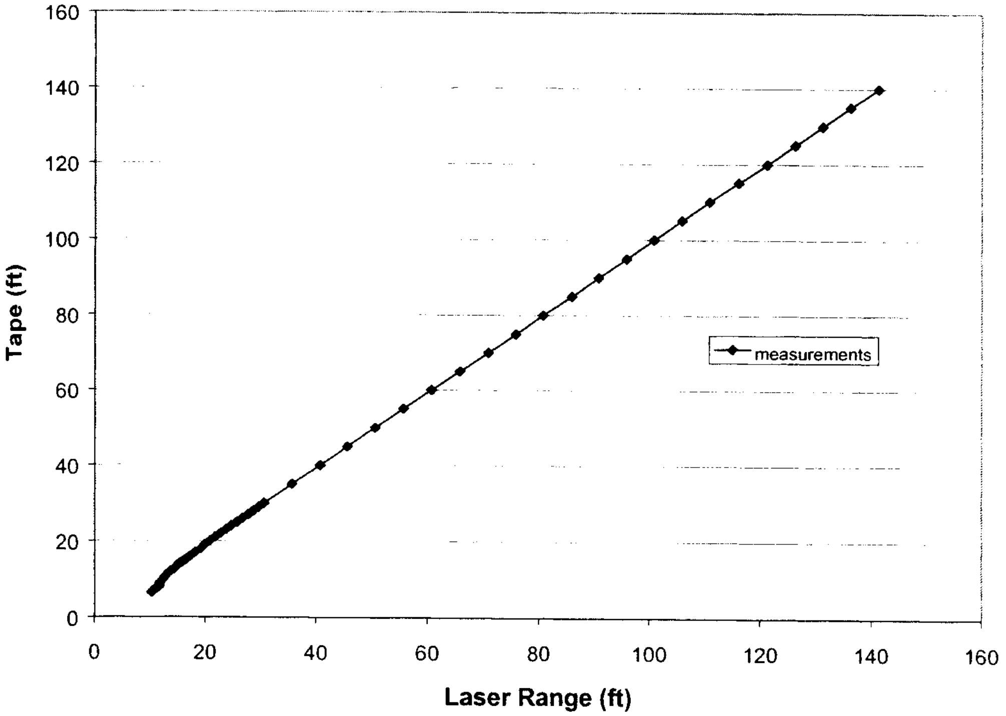

Figure A-2.1-5: Range Calibration of the Laser

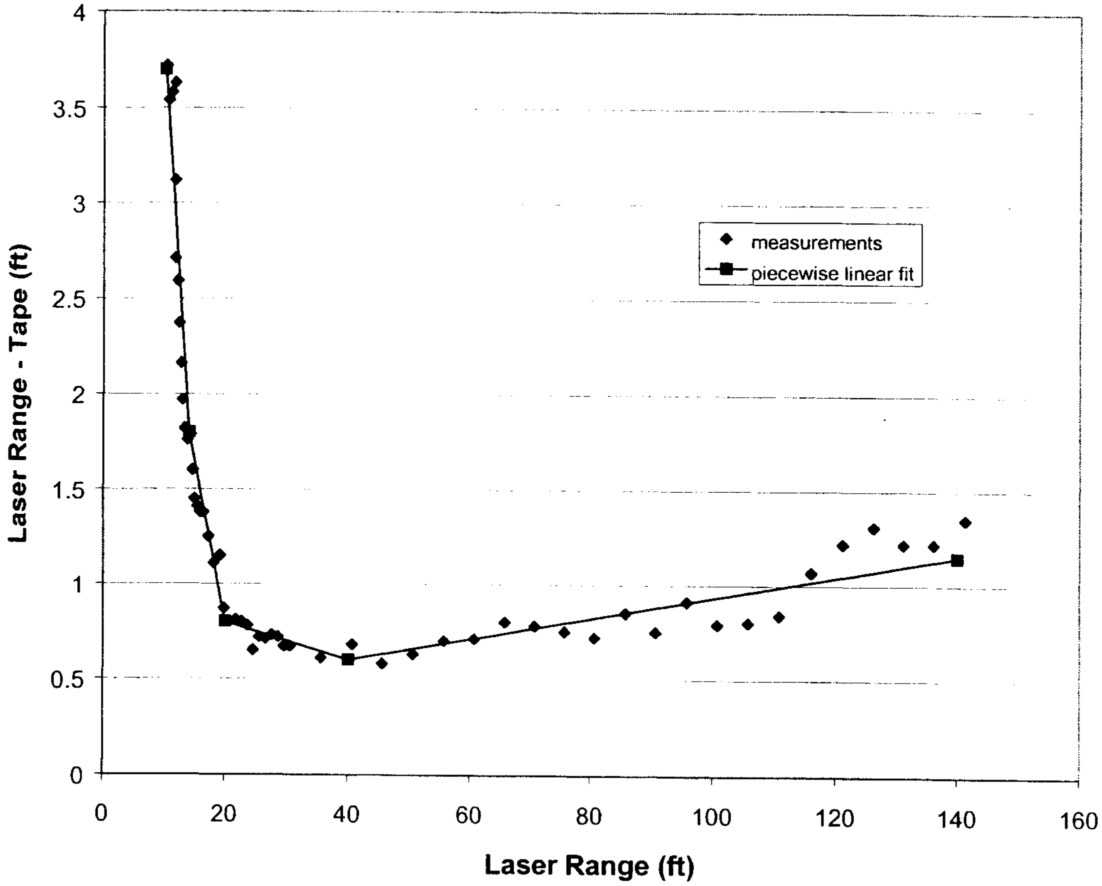

Figure A-2.1-6: Laser measurement error

A-2.2 Eye-tracker

The eye-tracker system was purchased from ISCAN, Inc. of Burlington, MA. The eye-tracker does two measurements. It first calculates the head orientation by measuring the magnetic field vector of a source fixed in the vehicle with a sensor that is mounted on the cap that the driver must wear. Also mounted on the brim of the cap is an infrared (IR) illuminator (less than 1mw/cm2 irradiance) and camera. The system is illustrated schematically and photographically in Figure A-2.2-1. The infrared source illuminating the eye is not coaxial with the eye imaging camera, therefore the pupil acts as a sink to the IR and the surrounding areas reflect the IR back to the camera yielding a "dark pupil" image. A dichroic mirror is placed in front of the eye to reflect the IR radiation, so that the eye can be viewed by the camera on the brim of the hat. The mirror is transparent to visible radiation so that the subject sees very little of the mirror because he is focused at a much greater distance. In Figure A-2.2-1 the pupil image is represented by the larger grey circle, while the corneal reflection is represented by the smaller white circle. Near the center of the eye, the cornea is nearly spherical and remains fixed with respect to the subject's head. The eye-tracker automatically computes the position of the pupil and corneal reflection over the two dimensional matrix of the eye imaging camera. This position corresponds to the direction of the eye with respect to the head, which is vectorially added to the head orientation vector yielding a final glance direction. This direction is computed at 60 Hz and is stored in the eye-tracker data acquisition computer.

We have pre-determined the important planes of regard around the car, as illustrated in Figure A-2.2-2. For example, the driver's windshield area (1), rear view mirror (5), right side view mirror (8), left (12) and right (9) front side windows are readily identified. The system computes the current plane of regard and outputs that information serially to the testbed data acquisition computer.

The section to follow describes the usage of and calibration procedure for the eye tracker system and provides commentary concerning limitations of the system. The magnetic head tracking subsystem is a unit from Polhemus called Fastrak. This subsystem outputs x,y,z and azimuth, elevation and roll information corresponding to the absolute position of a magnetic sensor with respect to a fixed reference source cube. The reference cube is mounted to a plate attached to the driver's headrest support rods. Because the sensor must always be positioned to the front and left of the source cube, and remain within 30 inches of the cube, we have built flexibility into the mounting locations of both the source and sensor in order to accommodate all sizes of drivers. The source cube can be moved laterally and the sensor can be positioned with or without extensions on the cap. The downside of using such a magnetic based position subsystem is that the steel in the body of the car tends to distort the field. This distortion is dependent both on the distance between source and sensor and on the position of the driver's seat, and is non-linear. There are two saving factors. The first is that the overall effect is not large and the second is that our requirements for pointing accuracy are not restrictive. Essentially we are interested in the plane number of the point of regard, rather than the exact position within that plane that the subject is viewing. The solution then is to define the planes loosely (i.e. with significant borders around the smaller ones) so that the defined planes can accommodate all drivers. After this was done, one needed to only take into account the position of the seat (and hence the source cube, which is the origin of the coordinate system). In effect the eye-tracker system had to be customized in software for each driver.

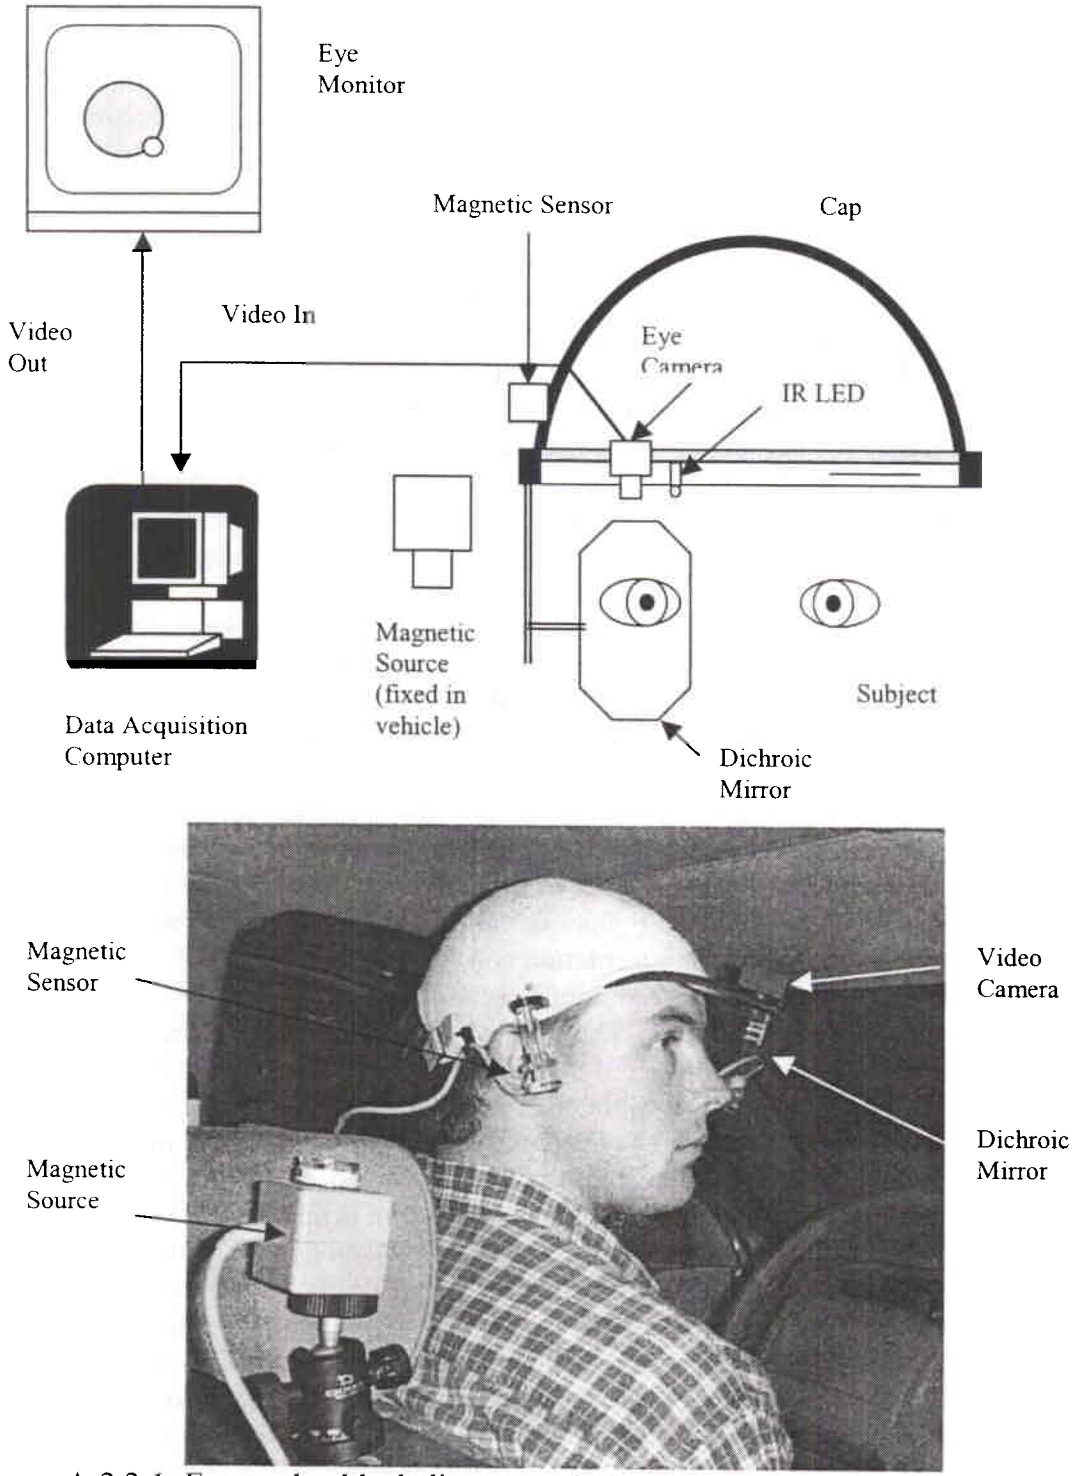

Figure A-2.2-1: Eye-tracker block diagram and Photograph

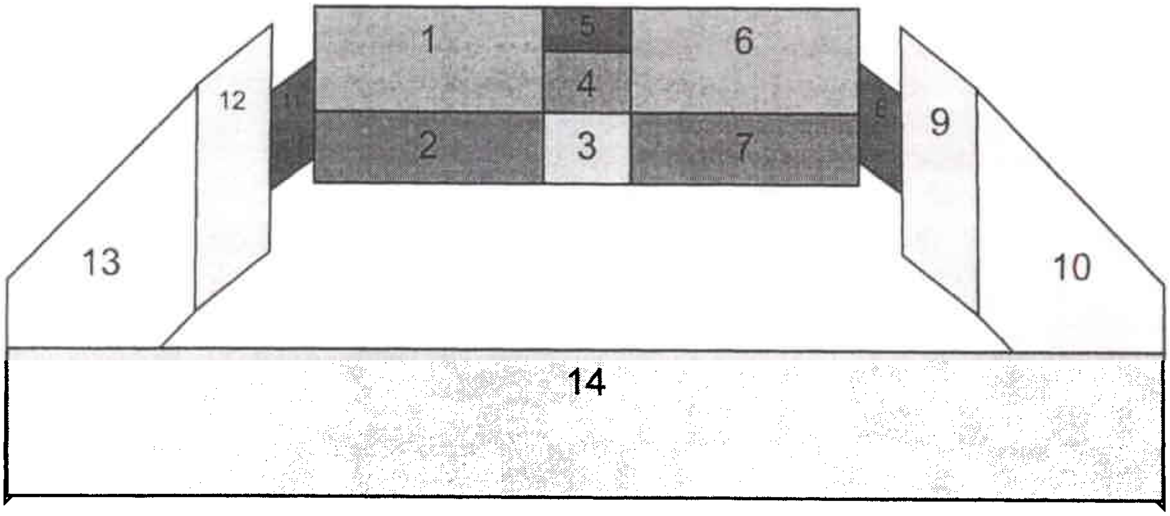

Figure A-2.2-2: Eye-tracker planes of regard

The calibration of this device must take place outdoors, so that the subject's pupil size will be stable during the test. This calibration was performed at the beginning of every run. The calibration procedure required the subject to adjust the headgear and the angle of the dichroic mirror so that the eye was well centered in the field of view of the camera. Subjects wearing glasses must first adjust the vertical position of the dichroic mirror so that the reflection from the eye was separated from that of the glasses. In general this meant lowering the mirror position, and of course adjusting the tilt of the mirror to center the image of the eye. The first step in the calibration procedure was to adjust the intensity levels and thresholds for the pupil and corneal reflection to account for the amount of ambient sunlight. The subject was unaware of any changes; this is solely for setting the instrument levels. The effect of sunlight is to reduce the contrast of the image. As a result, there are areas adjacent to the eye that compete in brightness with either the pupil or corneal reflection, and will under certain conditions confuse the eye tracker. Next, the subject was asked to look straight ahead to establish the baseline head position. A target with five dots arranged as indicated below in Figure A-2.2-3 was placed directly in front of the subject. The dots subtend a full angle of between 20 and 30 degrees, depending on the seat position of the driver. He was asked to stare at each of the five dots, in sequence, without moving his head. This calibrated the eye movement. At this point the subject was asked to direct his gaze at a number of fixed points around the car in order to assess the accuracy of the calibration. If necessary, the calibration was repeated. It was good practice to repeat this calibration check at the end of the run to verify that it had not changed.

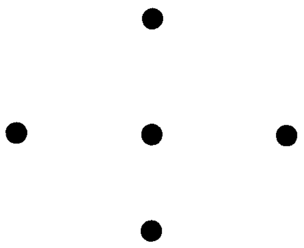

Figure A-2.2-3 Eye Calibration Pattern

Data output is of two forms. The first is a plane number that is continually transferred, via RS-232 output to the central data acquisition computer. This number is then folded in to each line of data which is tied to the laser scan rate (see section A-2.4). The second form is a complete set of data that is stored on the eye-tracker computer itself. This data includes the azimuth and elevation angle of the line of sight, and separately the head azimuth, elevation and roll angles, and the eye azimuth and elevation angles. In order to make use of this file it is time tagged with the GPS time by the host computer.

Software algorithms that we developed to process the eye-tracker data will be described in section A-2.7.

A-2.3 Global Positioning System (GPS)

The primary use for the GPS was to validate the workings of the laser and associated target acquisition software. In order to verify the accuracy of a given system it is necessary to use a yardstick that is at least as accurate. In differential mode, which is how we used it, the accuracy of the GPS under circumstances pertinent to our tests can range from the centimeter level to about 1 meter. In general we employed two independent GPS receivers and antennas. One was permanently mounted in the testbed, with the antenna placed directly over the laser scan head. The other was fixed on a confederate vehicle, with the antenna usually placed close to the front left corner (which is the point on the POV closest to the testbed when passing on the right).

In this section we will describe the hardware and software used, as well as give an elementary overview of GPS, especially as it applies to the manner in which we have used it in these tests. This description is taken liberally from the Waypoint Consulting, Inc. manual. (Waypoint 1999).

The general principle behind the use of satellites for determining position is, at the top level, fairly simple. The GPS system may be viewed, simply, as a continuous series of radio signals broadcast from a constellation of orbiting satellites to a radio receiver on the ground. These signals contain ephemeris information on the known position of each satellite, as well as measurement data indicating the range to each satellite and information describing the relative velocity of the satellites with respect to the receiver. Ephemeris data is a set of parameters used by a GPS receiver to predict the location of a GPS satellite and its clock behavior. Each GPS satellite contains and transmits ephemeris data, its own orbit and clock. Ephemeris data is more accurate than the almanac data but is applicable over a short time frame (four to six hours). Ephemeris data is transmitted by the satellite every 30 seconds. The GPS solution then reduces to a familiar problem in trigonometry, wherein if one measures the distance to three known points from an unknown point, the x, y, and z coordinates of the unknown point may be computed. The GPS problem is slightly more complicated by the fact that the radio signal travel time is unknown, necessitating the addition of one more known point to solve for the four unknowns x, y, z, and Δt.

Some GPS receivers make only pseudo-range or code measurements. These types of receivers are designed basically for navigation and have an accuracy of from 1m to 100m. Other receivers output phase and phase rate (Doppler) measurements as well as pseudo range and are designed for surveying as well as navigation. Depending on how they are used, they have an accuracy of from 0.01m to 100m.

We used the Ashtech G12 receiver, which follows a single frequency (L1 – 1575.42 MHz), can track up to 12 satellites, updates at 10 Hz, and is capable of recording the precise time of an input pulse. We used this last feature to periodically synchronize the internal timing of our data acquisition system. The receiver is differential capable and outputs standard code plus carrier phase and Doppler information. In operation we download and store the pseudo range and phase information for post processing.

In simplest form the expressions for pseudo range R and phase Φ, which form the basis of GPS measurements are

$$R = c * \Delta t + \varepsilon_1$$

$$\Phi = \Theta + \Sigma\Theta + N + \varepsilon_2$$

In the equation for pseudo range c is the speed of light, Δt is the transmission time, and ε_1 is the error term. In the equation for phase all the terms are given in terms of wavelength, which for the L1 frequency is 0.19m. It is this equation that is the basis for all high accuracy measurements. The quantity Θ is the instantaneous phase at the moment of lock-on and is measured in the receiver by mixing an internal reference frequency with the incoming GPS signal. The term $\Sigma\Theta$ is the sum of the incoming carrier waves from the moment the satellite has been locked by the receiver. This measured portion of the phase data is generally accurate to the sub-millimeter level. The term ε_2 is the error term that enters into the phase measurement.

The N term, which is called the phase ambiguity, is not a measurable quantity but rather must be computed from many observations. N represents the number of integer wavelengths from the receiver to each satellite at the moment of lock-on to that satellite. This means that N represents the integer portion of the first range or distance to the satellite at the first measurement epoch. The fractional portion of the initial distance is measured internally as the value Θ. Subsequent changes in the measured ranges, as the satellite and/or the receiver moves, are recorded in the term $\Sigma\Theta$.

Errors in GPS measurements come from a number of sources including errors in the broadcast ephemeris, propagation errors in the ionosphere and troposphere, receiver clock biases and receiver noise. The other major error is associated with selective availability, which is the process whereby the Department of Defense "dithers" the satellite clock and/or broadcasts erroneous orbital ephemeris data to create a pseudo range error. This was originally designed to prevent non-US military personnel from receiving accurate position information. Recently (May 3, 2000) the Secretary of Transportation announced a Presidential directive to end the practice of selective availability.

For a single receiver operating by itself, it can be shown that the effect of all these error sources can be dramatic. For instance, a timing error of only 0.1 microsecond in the receiver or satellite clock results in a coordinate difference on the order of 30m. Therefore a single receiver typically can be out of position by as much as 100m. The simplest way to eliminate most of the above mentioned errors is by using two receivers, and solving only for the difference in coordinates from a base station to the remote receiver. By knowing the absolute position of the base receiver, one can survey an area very accurately by using this method. If the satellite measurements from two receivers are combined, either post mission or in real time, errors from ephemerus, receiver/satellite clock, atmosphere, ionosphere and selective availability are virtually eliminated.

If the baseline length between base and remote receivers is less than 15 km, accuracies of 1-2 parts per million are possible with single frequency receivers. As the baseline increases the inaccuracy will tend to grow because the errors will no longer be the same magnitude at both receivers.

In our case, we were only interested in differential measurements. We were also fortunate in that the baseline was never more than 1 km between the SV and the POV. Furthermore it was not necessary to have real time information, so that eliminated the need for a radio link between the two receivers in order for the information to be combined and processed in one processor. The software for doing this job is GrafNav/GrafMov, developed by Waypoint Consulting Inc of Calgary, Alberta in Canada. GrafMov is a generalization of GrafNav wherein both receivers are moving. If only one receiver was moving then GrafNav was used.

If, as we have seen, the function of the post-processing software is to converge to a solution of the phase ambiguity, the best way to insure that one's solution is correct is to collect sufficient data wherein the two receivers are stationary. By requiring the phase ambiguity to be an integer, accuracies of .5 – 2 cm can be obtained by collecting 20 minutes of static data.

There are certain precautions that need to be observed when taking GPS data. First and foremost is to avoid loss of lock during data collection. Loss of lock occurs when a satellite is momentarily blocked from direct line of sight. It is possible to recover from this but the processed solution could suddenly jump to a new value. The second precaution to be observed is the avoidance of multipath signals. This occurs when the signal from the satellite bounces off an object such as a nearby building and arrives slightly delayed from the same signal which has taken the direct path. In such a case the receiver could become confused as to the nature of the true signal and will introduce time delays (and hence distance errors). Satellites that are near the horizon are the most prone to this problem, and they can be masked out in the post processing software. We typically rejected all satellites below an elevation angle of 10°.

Our procedure for taking GPS data was to set up the two receivers for static data taking, and measure the distance between them with a tape measure. In this way we have independent knowledge of the distance between them so that we should expect that the solution for the static portion of the data should converge to that value. We would then collect static data for 20 minutes. The tests under consideration would then be carried out. At the end of the testing, we would again collect static data for 20 minutes. Collecting static data before and after the tests allows us to compare the forward and reverse solution. The difference at any point in time between the forward and reverse solutions is a measure of the accuracy of the solution. Of course since we measured the distance between the two receivers at the beginning and end we know which solution to believe if there is a discrepancy.

A-2.4 Data Acquisition System

Figure A-2.4-1 is a block diagram of the data acquisition and control system. The core of the system is a Pentium computer with a digital signal processor (DSP) board. The DSP is sized to accept the continuous stream of laser data at 12 kHz. Real time data processing and the collision warning algorithm are resident in the DSP. As a consequence, all the sensor outputs that are required for the warning algorithm are also brought into the DSP. These include vehicle speed, steering angle, and turn signal usage. Based upon the algorithm's decision to warn the driver, the DSP will control the displays via digital output lines. All inputs not essential to the operation of the real time algorithm are communicated directly to the host computer. These are the eyeplane and GPS data, which in fact are collected in separate files and combined only during post processing. The host computer is a 300 MHz Pentium II, with 128 MB of RAM. It is configured with a 1GB removable Jaz drive that serves as a convenient way in which to transport data.

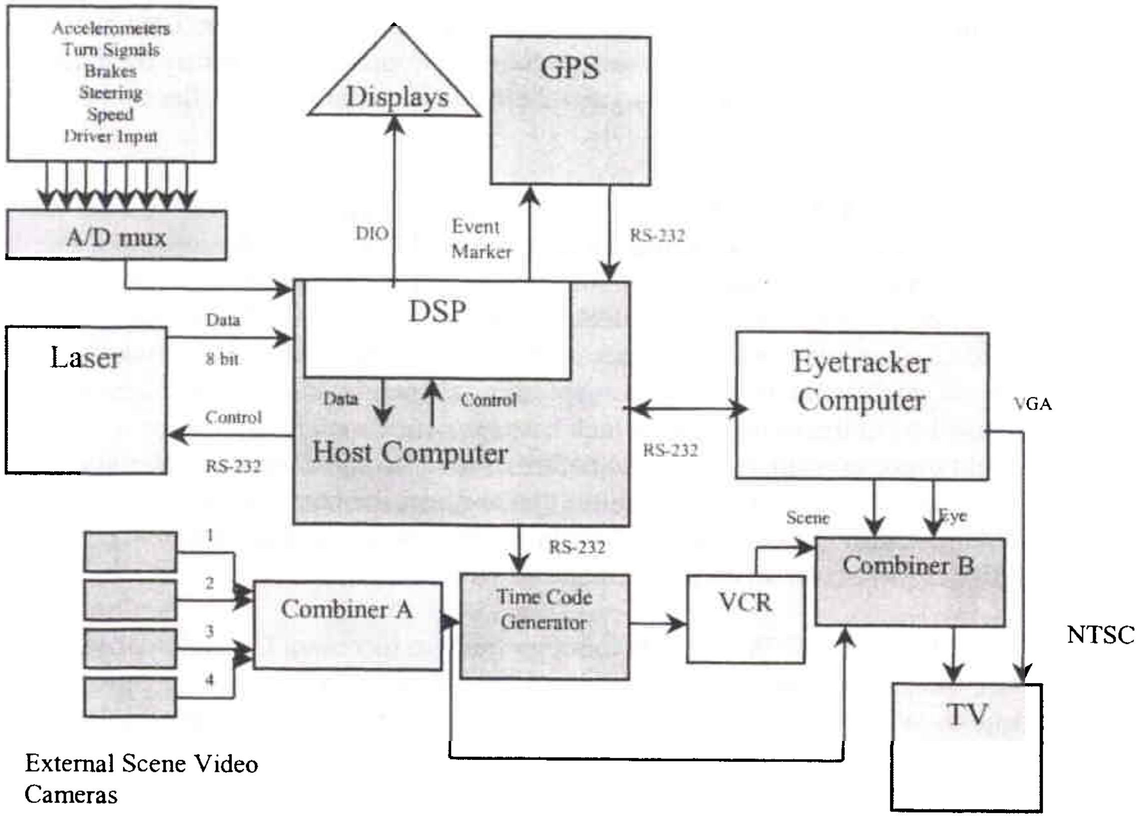

Figure A-2.4-1: Block Diagram of Control and Data Acquisition System

The DSP is an M44 model from Innovative Integration in Westlake Village, CA. It includes an on-board Texas Instruments TMS320C44 floating point processor which provides up to 60 MFLOPS of computational power, six DMA channels and four interprocessor communication ports. The DMA channels are used for data transfers from the laser embedded computer. The laser data rate consists of 8 bytes at a 12 kHz sample rate. The DSP board also has 32 bits of digital I/O available. We use these channels to control the displays and to send an event marker to the GPS, which synchronizes the DSP clock to the GPS. Added on to the DSP are two daughter boards providing 32 A/D single ended channels of 12 bit resolution.

An independent video system multiplexes and records the images from four cameras that monitor the area behind and to the side of the car. The CCD cameras are WAT-202B from Eatec Corp. The video record is secondary to the computerized system and is only used for scene confirmation. There are two externally mounted, in addition to the two internally mounted, cameras to record the scene. One is near the front wheel well and pointed backward, while the other is mounted underneath the laser scan head and is pointed forward. These two cameras are both fitted with telephoto lenses. The interior cameras, fitted with wide angle lenses, monitor the right side of the testbed and the area behind the car. The four cameras are then combined using a quad video combiner (MV-85) from RobotVision. The resulting scene is then time tagged by a Horita VLT-50PC. The host computer obtains the time from the GPS unit which hands it off to the Horita. The Horita then increments the time from that point on.

Time synchronization is crucial, particularly during the time of validation of the laser system as an accurate collision warning system. The synchronization process is started by operator command from a drop down menu in the acquisition and control program. A 2 kHz counter on the DSP is started at the same instant that a digital pulse is sent to the GPS. The GPS receiver records the exact time at which it receives the pulse and puts this into a message that it sends to the host computer via an RS-232 line. The host receives and records the GPS time, sets the time code generator (Horita) to that time and also sends a serial marker byte to the eye-tracker computer to mark the time in that data file. Once the DSP counter starts, the DSP sends a serial marker byte to the eye-tracker computer once per second. The marker bytes are embedded in the eye data for subsequent synchronization with the laser data. It was found that the 2 kHz counter on the DSP runs slightly fast when compared to the GPS. Hence it is necessary to re-synchronize every 5 minutes so that errors do not accumulate.

There are a group of relatively slow signals that are recorded by the A/D on the DSP board. They are accelerometer, turn signal, brake usage, steering angle, vehicle speed and auxiliary input. The accelerometers are from Omega (ACC102) with 100mv/g output and a ± 75g range. We have two accelerometers that monitor lateral and longitudinal acceleration.

A-2.5 DSP and Host System Software

As described in Section A-2.4, the DSP is responsible for (1) acquiring the laser data and other sensor outputs, (2) calibrating and transforming the laser data, (3) executing the collision warning algorithm, (4) controlling the driver warning display, and (5) sending the data to the host computer for archival. All the DSP software is written in C.

The logic flow of the system is illustrated in Figure A-2.5-1. Laser data is transferred from the laser embedded computer to the DSP via the DMA channels. As shown in Figure A-2.5-2, about a 90° span of the laser scan is blocked by the testbed itself and therefore contains no useful information. For this reason, the DSP subdivides a scan into two parts. It starts the data collection process near the rear edge of the testbed. Exactly 1000 samples (about 300° in span) are collected. The DSP then calculates how many samples it has to skip before it should start collecting data again. By adjusting the number of samples skipped, data for every scan always starts at the same azimuth. At the end of each scan, the laser data is calibrated and transformed into rectangular coordinates. The calibrated and transformed data is then fed into the collision warning algorithm along with other required sensor data that include the testbed speed, the steering angle and the turn signal usage. If a warning is warranted, warning displays will be energized. The raw laser data for each scan along with the processed outputs of the collision warning algorithm are stored in a storage buffer for transfer to the host computer. Two storage buffers are allocated so that data for the current scan can be stored while the data for the previous scan are being transferred to the host computer. Once the data are received by the host computer, they are unpacked and stored in the hard drive for archival. The data are also displayed in graphical form on a flat panel mounted at the back of the passenger seat.

Besides data archival, data display and interface with other system components as described in Section A-2.4, the host computer also provides the interface with the test operator. Through the use of a graphical interface, the operator can vary the configuration of the collision warning system. The changeable parameters include the display location (side mirror, center mirror, or both), the mode of operation (turn-signal-activated mode or monitor mode) and the system type (proximity system or comprehensive system). The modes of operation and the system types will be defined in section A-5.1. The blink rate of the warning display can also be adjusted. Once set, these parameters are passed to the DSP. All host software is written in Visual C++. Windows 95 is the operating system on the host computer.

A-2.5.1 Data File Structure

The raw data and the processed outputs from the collision warning system are archived in data files. Each file contains 5 minutes worth of data. The first line of each data file lists the parameters defining the size of the detection zones. The second line lists the 2 kHz DSP counter reading when the file is created. At the same instant, a digital pulse is sent to the GPS. The GPS receiver records the exact time at which it receives the pulse.

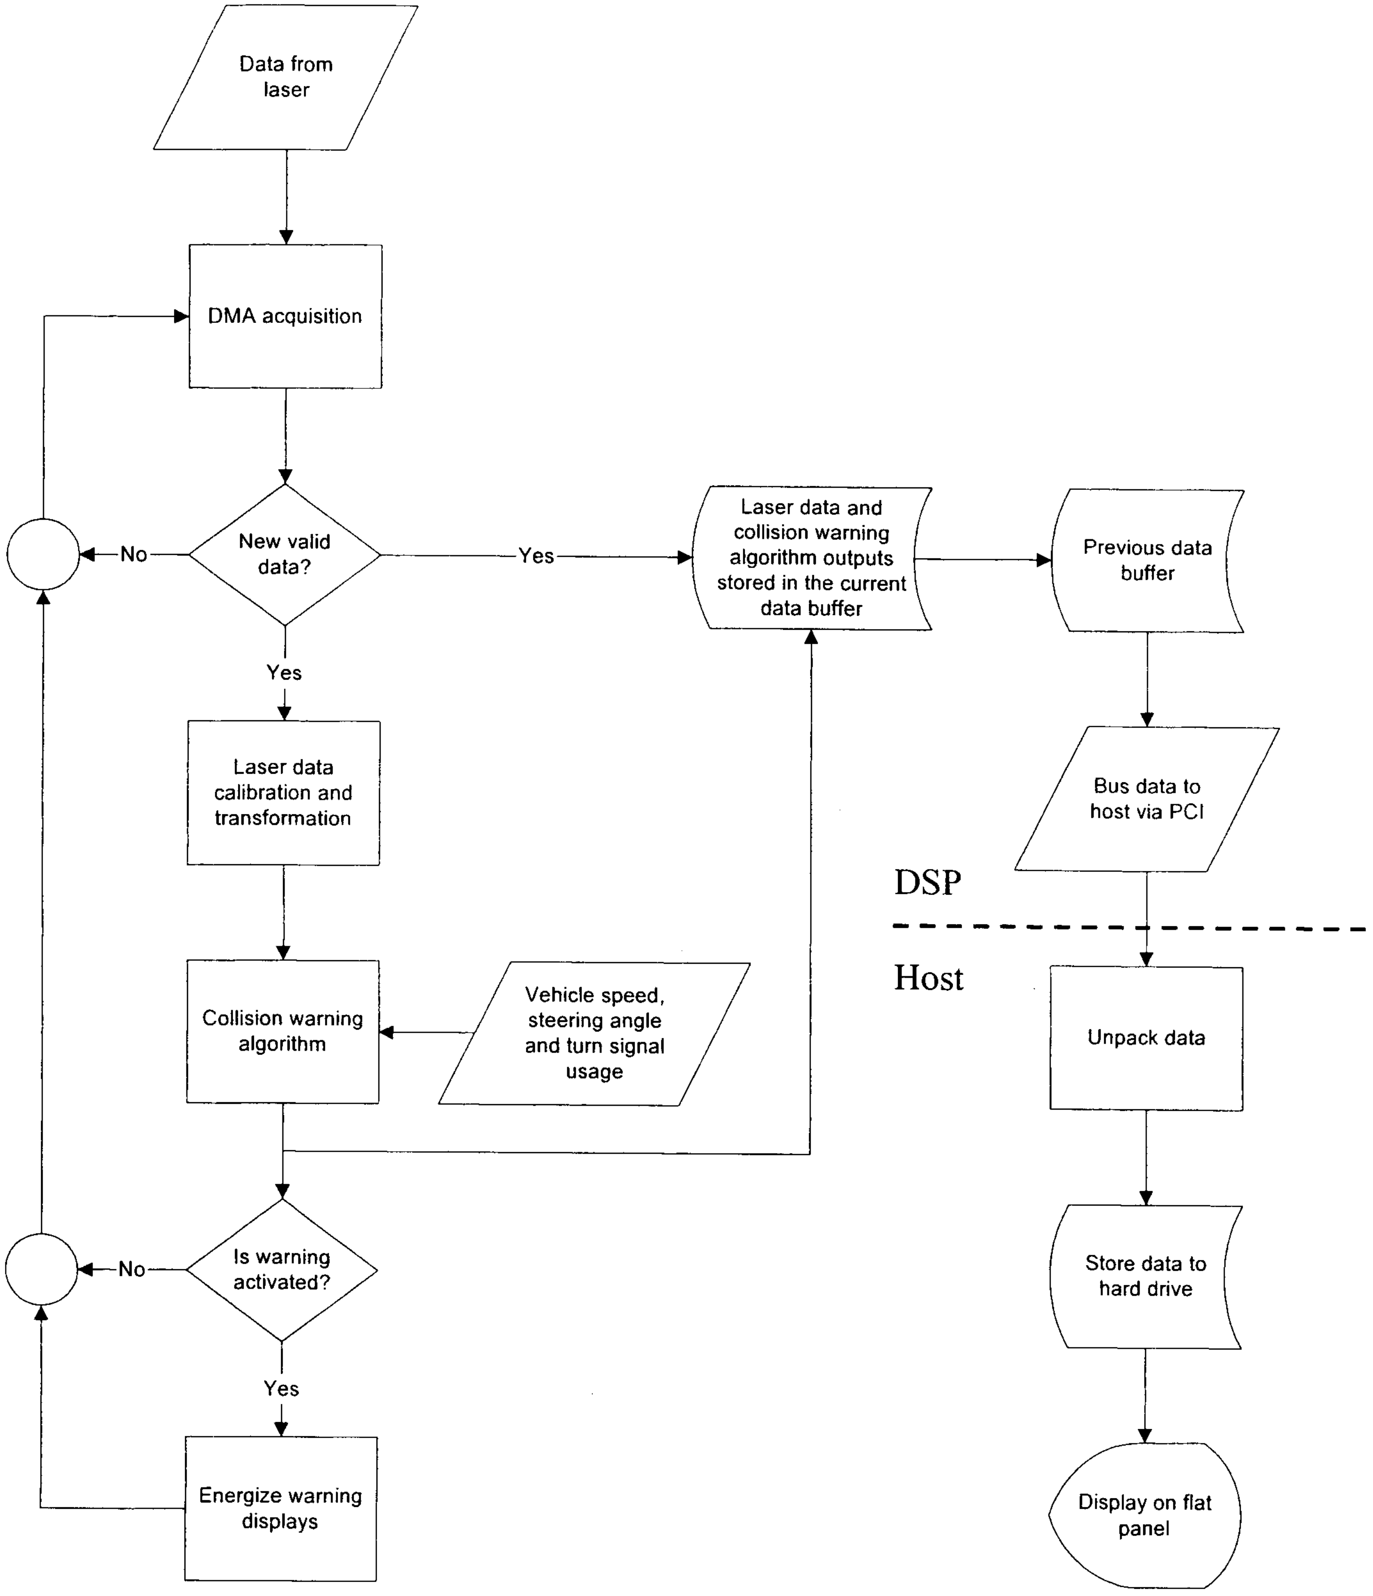

Figure A-2.5-1 DSP and Host System Logic Flow

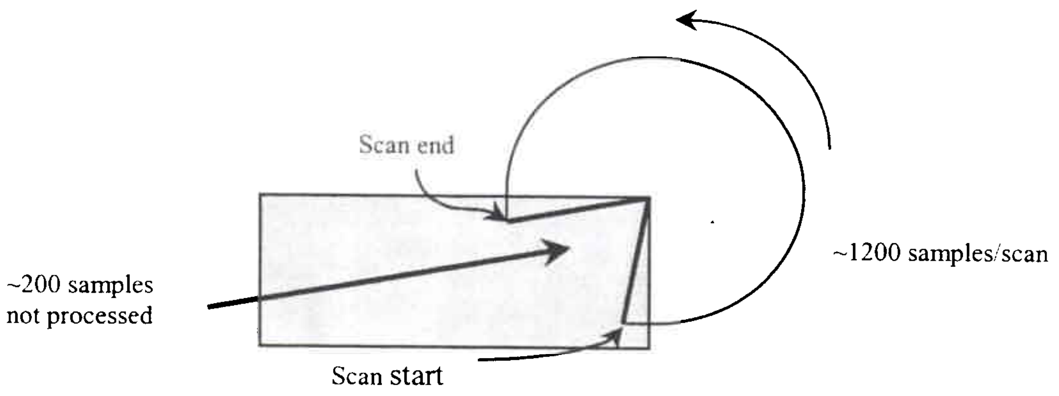

Figure A-2.5-2 Laser Scan Geometry

By comparing the DSP counter reading with the recorded GPS time, precise synchronization of the laser data and the GPS data can be achieved. The two header lines are then followed by data records. Each record contains data for one scan. The first 10 lines of the record contains the x and y coordinates of the 10 tracked points. Details of the tracking algorithm will be discussed in the next section. The 11th line lists the counter reading taken at the end of the scan, the left turn signal usage, the right turn signal usage, the brake usage and the approaching speed of the tracked vehicle in the right adjacent lane. The 12th line lists the steering angle, the testbed speed, the proximity warning output by the warning algorithm, the fast approach warning, the plane-of-regard output by the eye-tracker, an auxiliary input activated by a push button and another auxiliary input activated by the host keyboard. The 13th line lists the longitudinal and lateral accelerations as measured by the accelerometers. The last field in the line is not used. The ASCII data are then followed by the binary laser data. Each of the 1000 samples is 10 bytes in length. Two bytes are used for each of the 5 output fields. They are the azimuth, the range, the elevation, the amplitude and the status. Only the azimuth and the range data are used in the warning algorithm.

A-2.6 Detection, Tracking and Warning

The area around the subject vehicle is divided into zones, as illustrated in Figure A-2.6-1. Since the laser scanner is mounted on the right side of the vehicle, these zones are defined to facilitate warning against potentially dangerous right hand lane change only. The proximity zone (1) is 11 feet wide extending 30 feet back behind the car, and is directly adjacent to the car. The front edge of the zone protrudes a few feet in front of the car. The zone is wide enough to include any car that may be present in the adjacent lane, but not too wide as to include any car two lanes over. Cars detected in the proximity zone will trigger warnings. Details of the collision warning algorithm will be discussed in Section A-2.6.2. There is a forward zone (5), wherein objects are tracked in order to determine whether they are stationary with respect to the ground. Since the driver of the subject vehicle should be well aware of all stationary objects that enter the proximity zone from the front, warnings against these objects are considered nuisance alarms. Objects deemed to be stationary will, therefore, not trigger any warning when they pass from the forward zone into the proximity zone. Objects in the fast approach zone (2) are tracked in order to alert the driver of fast approaching vehicles in the adjacent lane. This zone has a length of 275 feet, long enough to cover the maximum range of the laser scanner. Both the forward zone and the fast approach zone are offset 4 feet to the right of the subject vehicle so that no car travelling in the same lane as the subject vehicle will be included. Finally zones 3 and 4 are extensions of zone 2. Since it takes some time for the tracking algorithm to settle after it starts tracking an object, tracking cars in zones 3 and 4 offer the advantage of an early start if these vehicles eventually change lane into the fast approach zone. Parameters for all the zones are listed in Table A-2.6-1. X and Y are the coordinates of the left front corner of each zone, with the position of the laser scanner being the origin of the coordinate system as shown in Figure A-2.6-1. The width and length of each zone are also listed in the table.

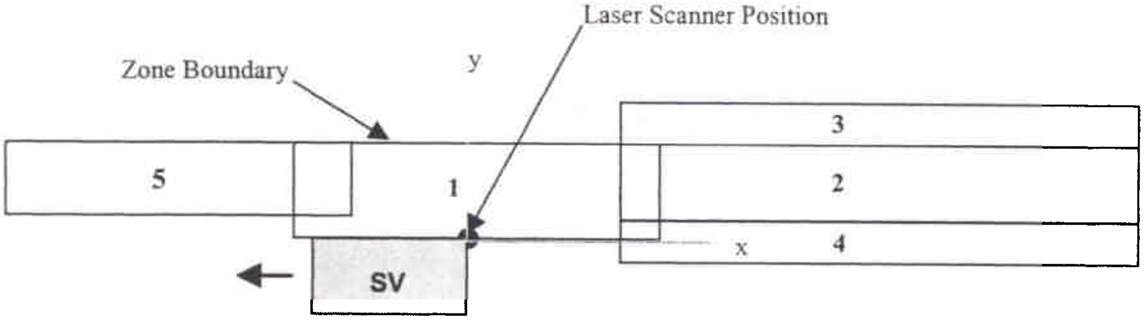

Figure A-2.6-1: Detection and Tracking Zones

	X (ft)	Y (ft)	Width (ft)	Length (ft)
Proximity Zone (1)	-20	0	11	50
Fast Approach Zone (2)	25	4	7	275
Fast Approach Zone Extension (3)	25	11	10	275
Fast Approach Zone Extension (4)	25	-6	10	275
Forward Zone (5)	-200	4	7	195

Table A-2.6-1 Zone Parameters

A-2.6.1 Tracking System

As mentioned above, objects in zones 2, 3, 4 and 5 are tracked. Figure A-2.6-2 is a block diagram showing the elements of the recursive tracking system. In general, the tracking process takes place independently in each zone. Recursive processing assumes that tracks have been formed on the previous scan. Observations are derived from data received from the laser scanner and then the processing loop shown in Figure A-2.6-2 is executed. Observations are first considered for the update of existing tracks. Gating tests and correlation algorithm determine the observation-to-track pairings. Observations not assigned to existing tracks can initiate new tentative tracks. A tentative track becomes confirmed when the number of observations included in the track satisfies confirmation criteria. Similarly, there is a deletion criteria that eliminate tracks not paired with observations for a set number of scans. Finally, after inclusion of the new set of observations, tracks are predicted ahead to the arrival time for the next set of observations. Gates are placed around these predicted positions and the processing cycle repeats. These elements are next discussed in more detail.

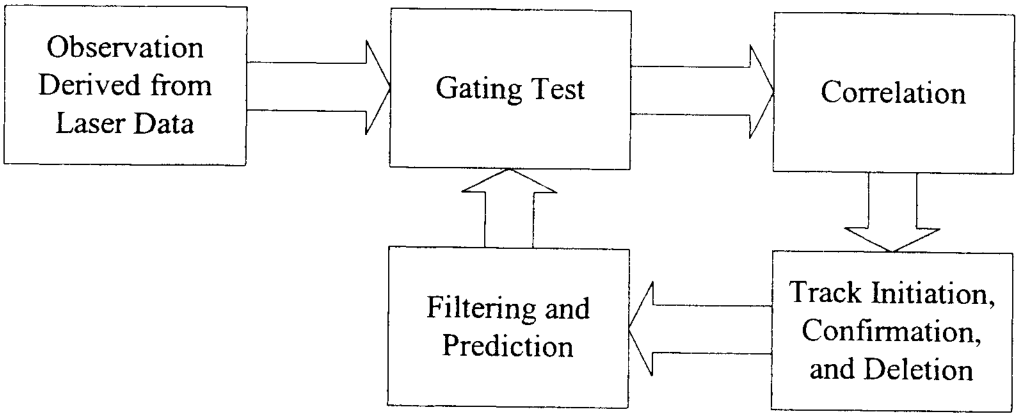

Figure A-2.6-2 Elements of the Recursive Tracking System

The objective of the system is to track the closest object in each zone. Hence, the relevant observation for each zone is simply the closest return (i.e., the point with the minimum range from the laser scanner) in that zone. The rest of the returns are not used by the tracking system. Figure A-2.6-3A shows a car in the fast approach zone. Geometrically, the target consists of 2 lines nearly perpendicular to each other, one line formed by returns from the front of the car and the other line formed by returns from the left side. The closest return is, therefore, around the front left corner of the vehicle. As the vehicle approaches the subject vehicle, the range of the closest return decreases accordingly. However, once the front of the vehicle exits the fast approach zone, the closest return inside the zone will shift to the intersection formed by the left side of the vehicle and the forward zone boundary as shown in Figure A-2.6-3B. The range to this point will no longer decrease but will stay more or less constant until the car completely emerges from the zone. Simply processing the range to the closest return inside the zone will in this case yield the erroneous result that the car comes to a stop when, in fact, it is exiting the zone. To counter this problem, when the closest return is within 2 feet of the forward boundary, the system will ignore it and will treat the case as if there is no return inside the zone at all. The track of the vehicle will then be predicted based on data collected in previous scans. Details of the track prediction process will be discussed later. Similarly, closest returns within 2 feet of the forward boundaries of zones 3 and 4 and within 2 feet of the rear boundary of zone 5 are ignored.

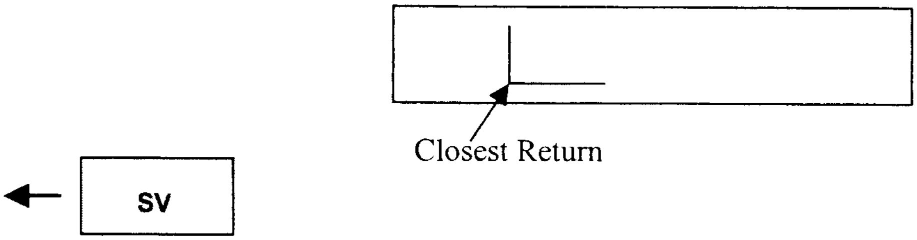

Figure. 2.6-3A A Car Inside the Fast Approach Zone

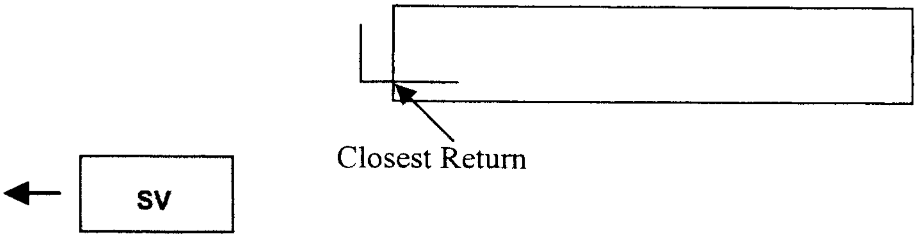

Figure A-2.6-3B A Car Exiting the Fast Approach Zone

Gating is the first part of the correlation algorithm used to decide if the closest return belongs to a previously established target track in that zone or to a new target. Up to a maximum of 2 tracks are maintained in each zone. To account for the possibility that the laser scanner may lose track of a vehicle momentarily due to poor reflection from the target, a confirmed track is extended for a short period of time after a missed detection occurs. Establishing a second track for the next car in the zone in the meantime covers the case that the first car may have exited the zone.

The closest return in each zone may satisfy the gates of one or both existing tracks. If one gate is satisfied the observation becomes a candidate for association with that track. If both gates are satisfied, then it becomes a candidate for the track it is closest to. If the observation does not satisfy the gate of any track, it becomes a candidate for the initiation of a new target track. The x-position of an existing track is estimated for the current scan using the filtered velocity and position for the last scan established by the tracking filter. This process of prediction will be discussed in detail later. The observation is deemed to satisfy the gate of an existing track if its x-position falls within the gate width around the predicted track position and its y-position is within 6 feet of the last observed value. The gate width is normally set to 4 feet around the predicted position. However, it is widened if the system has lost track of the target for a number of scans. The widened gate width accounts for the increased uncertainty of the predicted track position. The gate width in this case is set to $\sqrt{N+1}$ x 4 feet where N is the number of missed observations. For a tentative track that has only one prior observation and therefore no established velocity estimate, the gate width is set to 15 feet around the first observation to accommodate a range of velocity up to ~100 mph.

The correlation function takes the output of the gating function and makes final observation-to-track assignment. In the case where the observation is within the gate of a single track, the assignment is made immediately. If the observation satisfies the gates of both tracks, a confirmed track will take precedence over a tentative track. Classification of tracks will be discussed later. If both are confirmed tracks, the track that has its predicted position closest to the observation will take precedence.

Normally, the observation in each zone is tested only with tracks in the same zone. However, observation in the fast approach zone (2) is also tested against the tracks in zones 3 and 4 if no correlation is found within the zone 2. If correlation is established with an existing track in either zone 3 or 4, it is assumed that the target vehicle is in the process of changing lanes into the fast approach zone. The existing track parameters are then copied from the existing track into a new track in zone 2. This provision avoids the time delay and initial errors associated with track initialization. In fact, zones 3 and 4 are defined solely to serve this purpose.

An observation not assigned to existing tracks is used to form a new tentative track. Once a tentative track is initialized, two correlating observations from the following two scans are required to turn it into a confirmed track. Hence, the track confirmation requires a total of three observations. As described earlier, the second observation has to fall within 15 feet of the first observation and the third observation has to fall within 4 feet of the position predicted based on the first two observations. If the confirmation criteria are not met, the tentative track will be deleted. Once a track is confirmed, it will only be deleted under one of two conditions. A confirmed track will be deleted if it has not received any correlating observation for 9 scans. In other words, a confirmed track is maintained up to 0.8 second after its last observation. The second condition happens when a new track needs to be established but there are already two existing tracks in the zone. One of the existing tracks has to be deleted to make room for the new one. A tentative track, if present, will be deleted. But if both tracks are confirmed tracks, the one that is further away will be deleted.

The filtering step incorporates the correlated observations into the updated track parameter estimates. The α-β tracker is chosen for its simplicity. This filter is defined by the following equations [2]

$$\begin{aligned} x_s(k) &= x_p(k) + \alpha[x_o(k) - x_p(k)] \\ v_s(k) &= v_s(k-1) + \frac{\beta}{T}[x_o(k) - x_p(k)] \\ x_p(k+1) &= x_s(k) + Tv_s(k) \end{aligned}$$

where x_s is the smoothed position of the target for scan k based on its predicted position, x_p, and its measured position, x_o. The smoothed velocity, v_s, for scan k is a function of its smoothed velocity for the previous scan, the sampling interval, T = 0.1 second, the measured position and the predicted position. The last equation in the group generates the predicted position to the next scan. The initialization process is defined by

$$\begin{aligned} x_s(1) &= x_p(2) = x_o(1) \\ v_s(1) &= 0 \\ v_s(2) &= \frac{x_o(2) - x_o(1)}{T} \\ x_s(2) &= x_o(2) \end{aligned}$$

The filtering equations are used when an observation is received on scan k. If no observation is correlated to the track, the smoothed position is set equal to the prediction, $x_s(k) = x_p(k)$, and $v_s(k)$ is unaltered. This effectively amounts to assuming $x_o(k) = x_p(k)$. The prediction, $x_p(k+1)$, to the next scan is computed as before.

The values of α and β are normally set to values between 0 and 1. The choice of values represents a compromise between noise reduction and responsiveness of the filter. The latter quality includes the filter's capability to respond to accelerating or decelerating targets as well as its capability to damp out initial errors. A filter with low α and β values has high performance with respect to measurement noise but is slow in damping out initial errors and is also sluggish in response to accelerating or decelerating targets. On the other hand, a filter with high α and β values is responsive but has poor noise performance. With these considerations in mind, α is set to 0.4 and β is set to 0.1. Assuming that a target vehicle is maintaining a constant speed and the noise in position measurement is 1 foot, such a filter can reduce the velocity estimate error from an unfiltered value of 10 mph to 1 mph. With these filter coefficients, initial errors are typically damped out in 0.5 second.

The y-position of the target is kept as one of the track parameters, but is not filtered. As mentioned earlier, the y-position is compared to the last observed value as part of the gating test. An observation whose y-position lies more than 6 feet from the value last observed will fail the gating test.

A-2.6.2 Collision Warning Algorithm

The software on-board the vehicle functions to process the laser and other data in real time in order to evaluate whether a warning shall be presented to the driver. A warning is given when there is an object in the proximity area of the subject vehicle or when there is a fast approaching vehicle in the adjacent lane. The guiding principles behind the warning algorithm are 1) to have no false negative, 2) to eliminate as many nuisance alarms and false positives as possible without giving rise to false negatives and 3) to construct a system that is easy for the driver to comprehend.

The overall logic of the warning algorithm is illustrated in Figure A-2.6-4. The DSP executes this instruction set at the end of every scan (0.1 sec). The algorithm first initializes the Proximity Flag and the Fast Approach Flag to OFF. The Proximity Flag will subsequently be turned on if the algorithm detects any object in the proximity zone, while the Fast Approach Flag will be turned on if there is a fast approaching object in the fast approach zone. Other flags are used to indicate the status of the system. These status flags retain their states from scan to scan unless the criteria for resetting them are met. They include the Proximity Warning Disable Flag, the All Warning Disable Flag and the Persistent Warning Flag. The uses of these flags as well as their setting and resetting criteria will be discussed below.

As described in section A-2.1 the optics in the laser scanning mechanism reflects part of the outgoing radiation back toward the detector. For this reason, the laser scanner will return a close range under any one of the following conditions: (1) there is no target in range; (2) the target in range does not reflect sufficient amount of radiation back either due to surface properties or geometry; or (3) the target is inside the exclusion zone extending about 10 feet from the scanner. Hence, the first step in processing the laser data is to discard the superficial returns with close ranges. Examination of the data indicates that 3.5 feet is the appropriate cutoff range. The algorithm then finds clusters of returns that it can identify as a single object. Contiguous points within 5 feet in the x-direction and 3 feet in the y-direction are grouped together as part of an object. It then screens out objects that are that are too long or too wide to be vehicles. This eliminates objects such as walls and the ground returns when the subject vehicle is going up a relatively steep incline which forces the laser beam to hit the ground behind the car.

After screening out close returns and large objects, the remaining returns are passed to the tracking algorithm described above. Two system types are selectable. For the proximity system, only the presence of a target in the proximity zone 1 (see Figure A-2.6-1) will trigger a warning. For the comprehensive system, a fast approaching object in zone 2 or an object in zone 1 will trigger a warning. The decision to warn is based solely upon presence in the proximity zone 1 and upon time to enter the proximity zone from zone 2. If a tracked object in zone 2 is going to enter zone 1 within 3 seconds, the Fast Approach Flag will be set ON indicating a fast approach vehicle in the right lane.

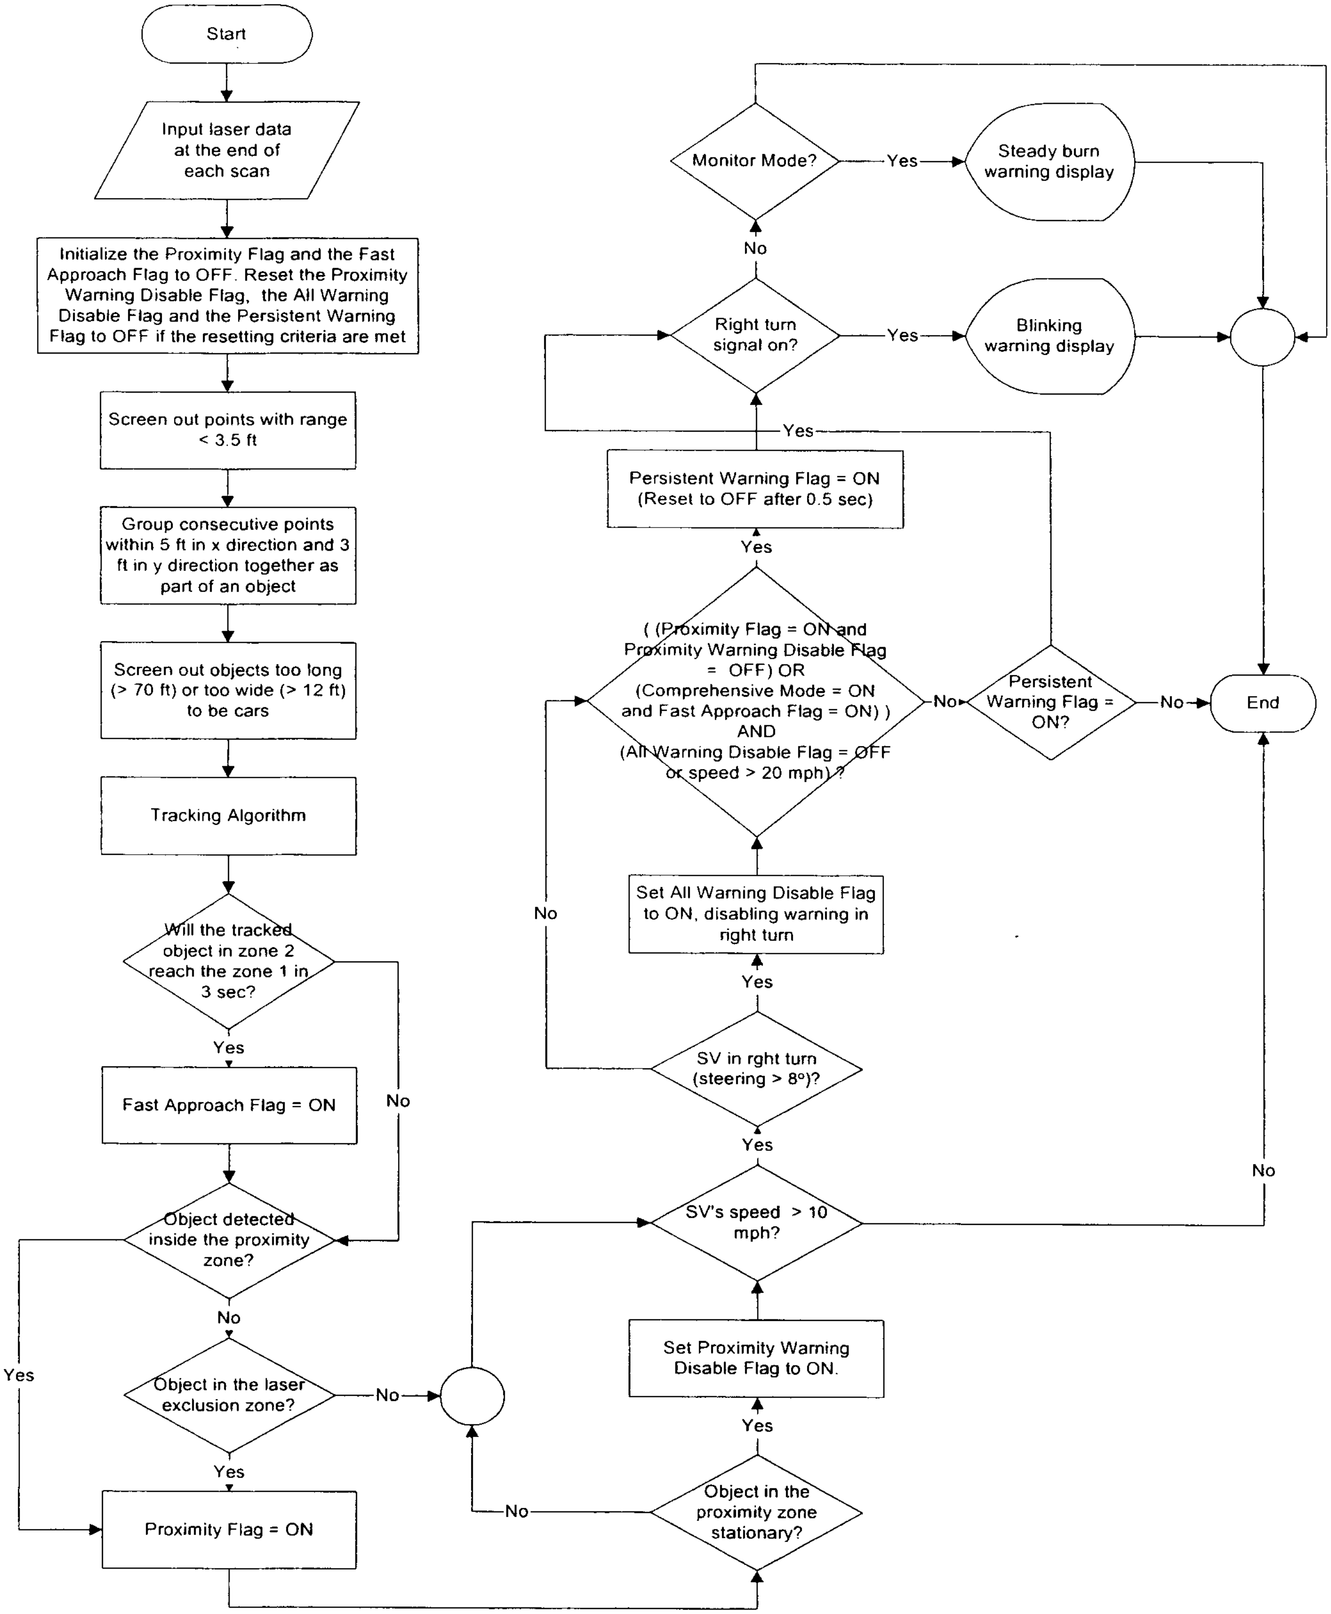

Figure A-2.6-4: Warning Algorithm Logic Flow

The proximity zone 1 is monitored such that any return in the zone will set the Proximity Flag ON indicting the presence of an object. The two systems are summarized in Figure A-2.6-5, which shows the X-V phase space diagram. X is the longitudinal distance measured rearward from the laser scanner to the front bumper of the target vehicle. V is the closing speed of the target vehicle. The horizontal line at X = -20ft approximately corresponds to the front bumper of the SV. A vehicle in region I will trigger a warning when the proximity system is selected. When the comprehensive system is selected, a vehicle in regions I or II will trigger a warning. A vehicle in region III will not trigger a warning under any circumstances.

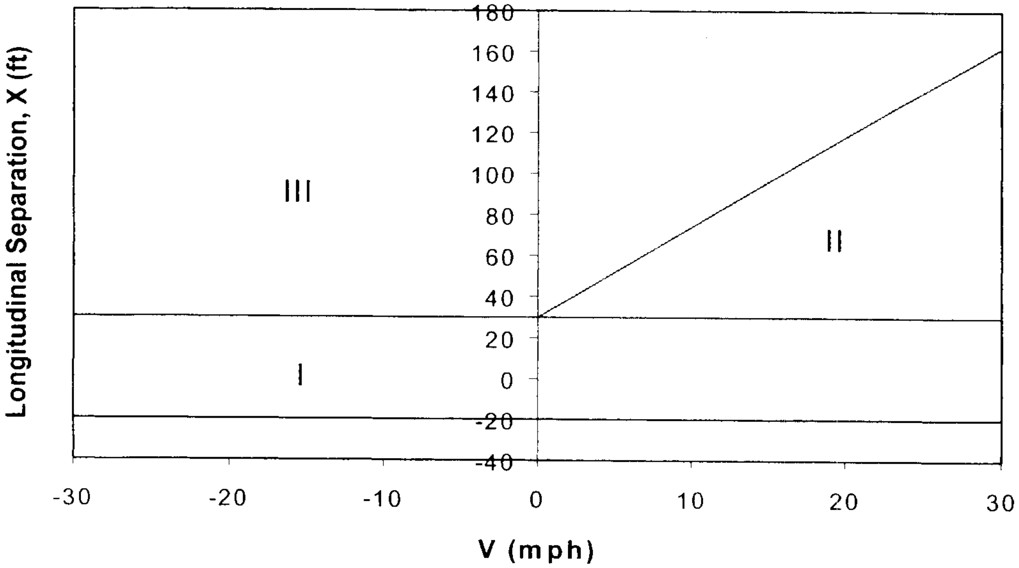

Figure A-2.6-5 Warning Regions

To solve the problem that a car may hide completely inside the exclusion zone of the laser, hence, avoiding detection, an algorithm has been developed to infer the presence of such a vehicle. As discussed previously, the scanner will return a close range (less than 3.5 feet) when the laser beam is reflected off an object inside the exclusion zone due to its failure to isolate the object returns from those reflected back by its own optics. Based on this effect, the algorithm treats any contiguous section on the passenger side that is made up of close returns only as a possible indication of a hidden vehicle. If the section is less than 10 degree in angular span, it is deemed too small to be even a motorcycle and therefore is ignored. It should be noted that a wide enough sector of close returns is not a definite indication of a hidden vehicle since the absence of a target within the laser range produces the same signature. To reduce the chance of issuing a false warning caused by the empty background, the algorithm checks whether there was a target in the proximity zone in the previous scan. If the Proximity Flag wasn't on, the algorithm will assume that the sector of close returns in this scan is caused by empty space. If the Proximity Flag was on, the algorithm will assume that the vehicle previously detected has disappeared into the exclusion zone.

The algorithm will next determine whether the object in the proximity zone is stationary with respect to the ground. Warnings against stationary objects such as parked cars or poles close to the edge of the road are considered nuisance alarms and are thereby suppressed. An object is considered to be stationary if its approaching speed is within 5 mph of the subject vehicle's speed. The Proximity Warning Disable Flag is then set to ON indicating that the proximity warning should be disabled. This flag is kept on over subsequent scans until the proximity zone is clear of all objects for 3 consecutive scans. The Proximity Warning Disable Flag will also be reset if the speed of the subject vehicle drops below 10 mph.

No warning is given when the subject vehicle is traveling below 10 mph. This criteria is designed to stop the warnings while the car is not in motion. Warnings are also disabled when the car is making a right turn. This criteria is designed to eliminate warnings from objects on a corner that enter the proximity zone on an angle during a turn and hence avoid being tracked in zone 5 and identified as stationary objects. The algorithm treats the situation as a right turn situation when the steering angle goes above 8 degrees. The All Warning Disable Flag is then set on to disable all warnings. The threshold of 8 degrees is set low enough to include all right turns but high enough to exclude road curves. The All Warning Disable Flag is kept on over subsequent scans until the subject vehicle has traveled 50 feet after the steering drops back below 8 degrees. The extra 50 feet allows the algorithm some time to start tracking objects on the roadside after the subject vehicle comes out of the turn.

If the object passes all the tests and a decision is made to warn the driver, the algorithm will have the warning persist 0.5 sec even after the warning criteria are no longer met.

Two modes of operation are software selectable. When the monitor mode is selected, a steady warning is displayed when required. The severity of the warning is elevated to a blinking display when the right turn signal is energized. When the turn signal mode is selected, a warning is displayed only when the turn signal is energized, and it is a blinking one.

The display location is also selectable. It can be in either the rearview mirror, or the right side mirror, or both.

As stated earlier, the guiding principle of the warning algorithm is to eliminate all false negatives while minimizing the occurrence of false positives and nuisance alarms. False positives are warnings triggered by non-existent threats, the bulk of which come from stationary objects. There are a few types of false positives not properly eliminated by the algorithm. The first type is a proximity warning triggered by empty background. Since an empty background and a hidden vehicle in the exclusion zone have the same pattern of returns, the algorithm will mistakenly conclude that there is a vehicle inside the exclusion zone if the empty background spans more than 10 degrees and, by coincidence, there was a vehicle detected inside the proximity zone in the last scan. After the vehicle exits the proximity zone, the warning will persist until the empty background passes. Setting a time limit (e.g. 1 sec) on how long the warning would persist after a vehicle is assumed to have disappeared completely into the exclusion zone would reduce the duration of such a false warning, but only at the expense of missing some real targets that may stay hidden in the exclusion zone longer than the set time limit.

Another type of false positive is a proximity warning triggered by a vehicle following in the same lane, but close enough and slightly offset in such a way that it cuts the corner of the proximity zone. The situation is illustrated in Figure A-2.6-6. This problem may be avoided by clipping the corner of the proximity zone as shown in Figure A-2.6-7. However, such a system may miss a motorcycle riding in the adjacent lane close to the lane divider.

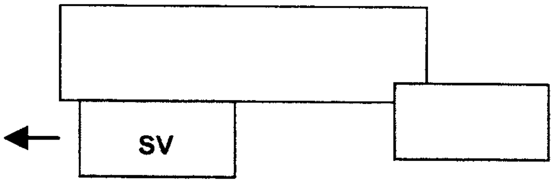

Figure A-2.6-6 False Proximity Warning Triggered by Vehicle Following in Same Lane

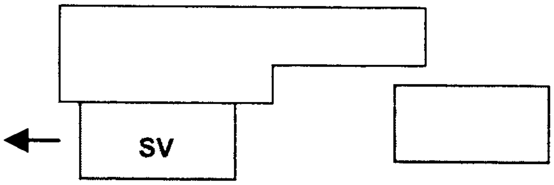

Figure A-2.6-7 Proximity Warning Zone with a Corner Clipped

The third type of false positive is a fast approach warning that may be caused by a long object such as a guard rail intersecting the boundary of zone 2 as shown in Figure A-2.6-8. The two snapshots are taken 1 scan (0.1 sec) apart. The subject vehicle is traveling along a curve. As a result, the guide rail along the roadside cuts into zone 2. The closest return in that zone comes from the intersection between the rail and the zone boundary. In the upper snapshot, the tracked point has a range of about 85 feet. In the next scan, the point tracked has shifted to a different location and now has a range of about 80 feet. The algorithm therefore mistakenly concludes that there is a fast approaching object in that lane. The track initiation process which requires 3 consecutive returns satisfying the gating criteria before a confirmed track is established eliminates a large portion of these cases, but not all. Discrimination against long objects did not screen out the rail in this case since the returns are broken up into pieces. This type of false warning may also be triggered by a car following in the same lane in a curve. The situation is illustrated in Figure A-2.6-9. This type of problem can be alleviated by requiring a greater number of consistent returns before a track is confirmed at the expense of system latency. Another way to alleviate this problem is to delete an approaching track once an observation is missed if the track is close in range (e.g. within 75 ft) where real targets, in general, have reliable returns. Doing so will prevent false tracks from getting too close to the subject vehicle. Of course, one will run the risk of missing a real but poorly reflecting target.

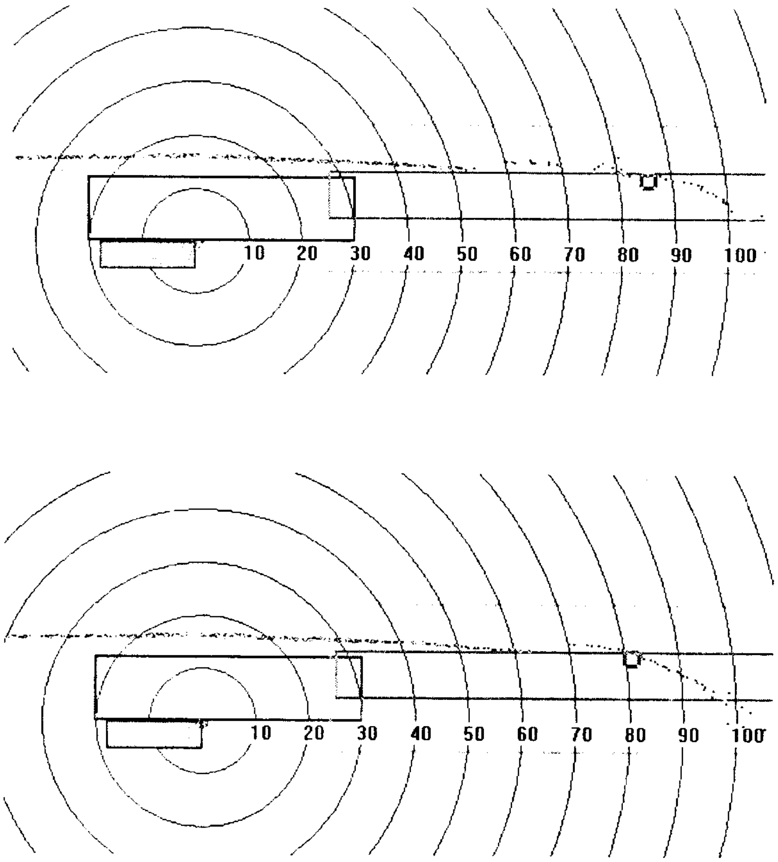

Figure A-2.6-8: False Fast Approach Warning Triggered by a Long Wall

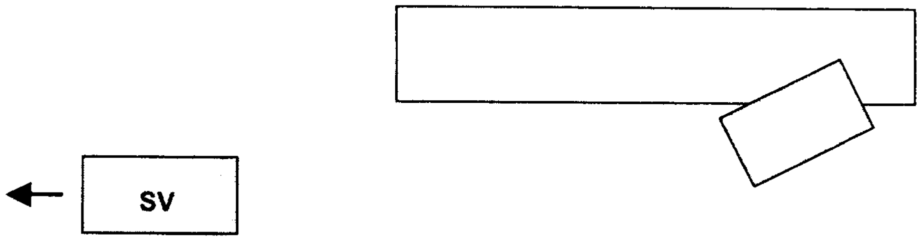

Figure A-2.6-9: False Fast Approach Warning Triggered by a Vehicle Following in the Same Lane in a Curve

Most of the false positives are triggered by stationary objects and are properly suppressed. However, there are a few situations when these warnings are not eliminated. This happens when the algorithm cannot properly track a stationary object in zone 5 to determine its approaching speed before it enters the proximity zone. An extended object such as a guard rail intersecting zone 5 will pose a tracking problem. A stationary object entering the proximity zone from the side due to the sideways movement of the subject vehicle will also trigger a false warning.

A-2.7 Eye-tracker Data Correction Algorithm

The principles of operation and the eye-tracker hardware are described in Section A-2.2. The purpose of the eye-tracker system is to record and output the plane of regard that the subject is looking at. However, field experience with the system has uncovered limitations peculiar to the automobile environment that require part of the plane data to be corrected during the data analysis process. Both the limitations and the correction algorithm that recovers the data to the specified level of accuracy will be discussed in this section.

As discussed in Section A-2.2, the pupil acts as a sink to the IR illumination while the cornea acts as a reflecting surface. The eye-tracker searches for the pupil location by first identifying all regions in the eye image that are darker than the operator set pupil threshold. If more than one dark region exists, the system assumes the largest region to be the pupil and identifies the center of that region as the pupil position. Only one region in the eye image shown in Figure A-2.7-1A is darker than the pupil threshold. Figure A-2.7-1B shows that the system correctly marks the center of the dark region with a cross as the pupil position. The system searches for the corneal reflection in a similar fashion by identifying regions brighter than the operator set corneal threshold. Corneal reflection is also marked with a cross in Figures 2.7-1A and 2.7-1B.

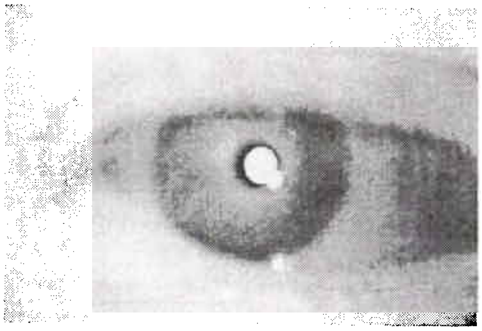

Figure A-2.7-1A Corneal Reflection and Area Darker the Pupil Threshold

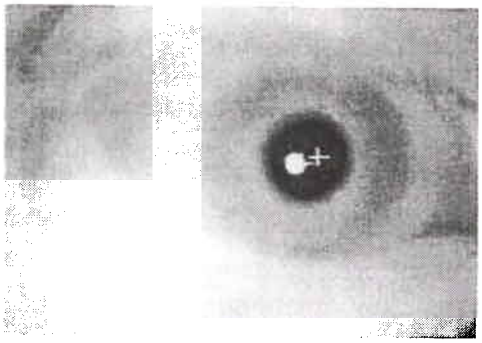

Figure A-2.7-1B Corneal Reflection and Pupil Position

There are a number of conditions under which the eye-tracker may fail to identify the pupil position and the corneal reflection correctly. Figure A-2.7-2 illustrates one such condition when the system identifies the dark area at the corner of the eye as the pupil. In general this problem occurs when the subject vehicle enters a shaded area. A region that is normally brighter than the pupil threshold may become dark enough to be mistaken as the pupil when the ambient illumination drops.

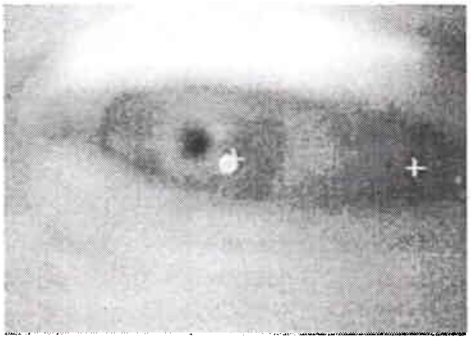

Figure A-2.7-2 Dark Area at the Corner of the Eye Mistaken as the Pupil

Near the center of the eye, the cornea is nearly spherical and therefore only reflects the infrared source at a single location. However, the surface becomes bumpy near the edge of the cornea. When the subject looks to the side, the IR source may be reflected at multiple locations as illustrated in Figure A-2.7-3. When that happens, the eye-tracker may lose track of the original reflection and picks up another reflection instead. With more extreme eye movement, the IR reflection may roll off the cornea onto the non-reflective sclera. In that case, the corneal reflection will disappear entirely.

Figure A-2.7-3 Multiple Corneal Reflections

The eye-tracker may sometimes mistakenly treat other bright areas in the eye image as corneal reflections. An example is shown in Figure A-2.7-4. In this case, the IR reflection off the lower eye lid is taken as the corneal reflection. This may happen as the subject squints when the car turns into the sun. Subjects with narrow eyes are particularly problematic. The eye-tracker also has problems with subjects who wear glasses. The fundamental problem here is that the glass provides another interface (and a strong one) from which light from the LED can be reflected. Figure A-2.7-5 shows an example in which the reflection off the eyeglass at the lower edge of the eye image is mistaken as the corneal reflection.

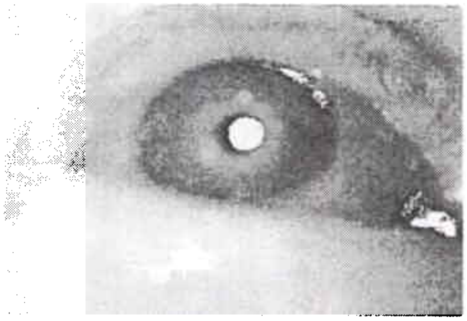

Figure A-2.7-4 Reflection from the Lower Eye Lid Mistaken as the Corneal Reflection

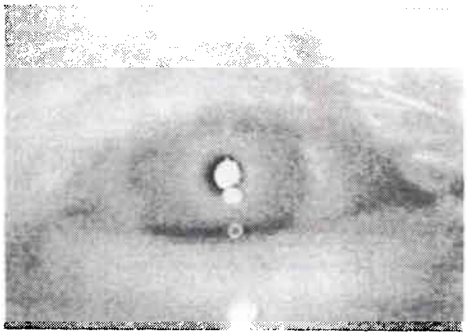

Figure A-2.7-5 Reflection from eyeglass Mistaken as the Corneal Reflection

A sun shield was used that mounts on the brim of the cap, and hangs straight down, wraping around the entire face of the subject. The purpose of the sun shield is to cut down the ambient illumination that may light up areas around the eye. In addition, since the shield acts like a sunglass, subjects do not squint as much when the shield is used. The thresholds for the pupil and corneal reflection are adjusted carefully as part of the eye-tracker calibration procedure at the beginning of each test run. The thresholds are set such that dark areas besides the pupil and bright areas besides the corneal reflection are excluded. These levels are re-adjusted during the run as the ambient lighting condition changes. During the calibration procedure, subjects are also asked to adjust the headgear and the angle of the dichroic mirror such that dark and bright areas are moved out of the eye image. Despite these precautions, not all problems discussed above can be avoided and a software fix has to be implemented as part of the data processing procedure.

A bug has also been discovered in the software provided by the eye-tracker vendor. It may sometimes incorrectly return a plane of regard opposite to the subject's line of sight. For example, the software may indicate that the subject is looking at the left window (plane 12) when, in fact, he is looking at the right window (plane 9). Further testing of the software shows that errors with the other planes are rare, but returning plane 12 when the subject is looking at plane 9 happens quite often. Later versions of this software from the vendor have corrected this problem.

When any of the problems described above occurs, the plane number returned by the software will not be correct. The first task of the correction algorithm is to separate good from bad data. The algorithm first checks the head azimuth measured reliably by the magnetic sensor against the eye plane number returned by the eye-tracker software. If the algorithm detects a gross inconsistency between the head azimuth and the eye plane number, the eye plane data will be marked as bad for subsequent recovery. The data are classified as bad if the head azimuth, measured in degrees, is outside the range listed in Table A-2.7-1. The head azimuth should be roughly zero when the subject looks straight ahead. The angle turns positive to the left and negative to the right. This test is mainly designed to pick out the gross errors when the eye-tracker software returns a plane number directly opposite to the true line-of-sight of the subject.

Planes	Head Azimuth (deg.)
1, 2	-50 to 50
3, 4, 5	-70 to 30
6, 7	-80 to 20
8, 9	-100 to 0
10	-170 to -70
11, 12	0 to 100
13	70 to 170
14	-180 to –130 or 130 to 180

Table A-2.7-1 Valid Head Azimuth Ranges

The algorithm then proceeds to identify bad eye plane data caused by the eye-tracker's failure to track either the pupil or the corneal reflection. The eye-tracker software combines the head position and direction measured by the magnetic head tracking subsystem with the eye azimuth and elevation measured by the eye-tracker to derive the subject's line-of-sight and the plane of regard. Failure in tracking either the pupil or the corneal reflection will result in incorrect eye azimuth and elevation, thus giving an incorrect plane of regard. Fortunately, the bad portion of the data can usually be identified through examining the eye azimuth, eye elevation and pupil diameter data recorded by the eye-tracker software. Fifteen seconds worth of such data are plotted in Figure A-2.7-6. The pupil diameter is measured in unit of pixels in the eye image. Typical values are around 35 depending on the subject and the pupil's response to the ambient lighting conditions. When the eye-tracker fails to locate any region that satisfies the pupil threshold, it returns 1 as the value for the pupil size. This happens when the subject blinks or when the ambient light level rises enough to light up the pupil. As shown in Figure A-2.7-2, the eye-tracker may misidentify a dark region as the pupil. The reported pupil size may indicate the problem if the size of the dark area is significantly larger or smaller than the typical values. With all these taken into consideration, data with pupil diameter falling outside the range of 20 to 50 are marked as bad data for subsequent recovery. When the eye-tracker system misidentifies a bright area as the corneal reflection or a dark area as the pupil, the area is often relatively far from the true location of the corneal reflection or the pupil. Large values for the eye azimuth and/or the eye elevation are generated as the result. As shown in Figure A-2.7-6, portions of the eye azimuth and eye elevation data exceed 100 degrees which is not physically attainable by real eye movement. Cutoff values of ±40 degrees for the eye azimuth and ±30 degrees for the eye elevation are used in the screening process. Data with values falling outside the range are marked as bad.

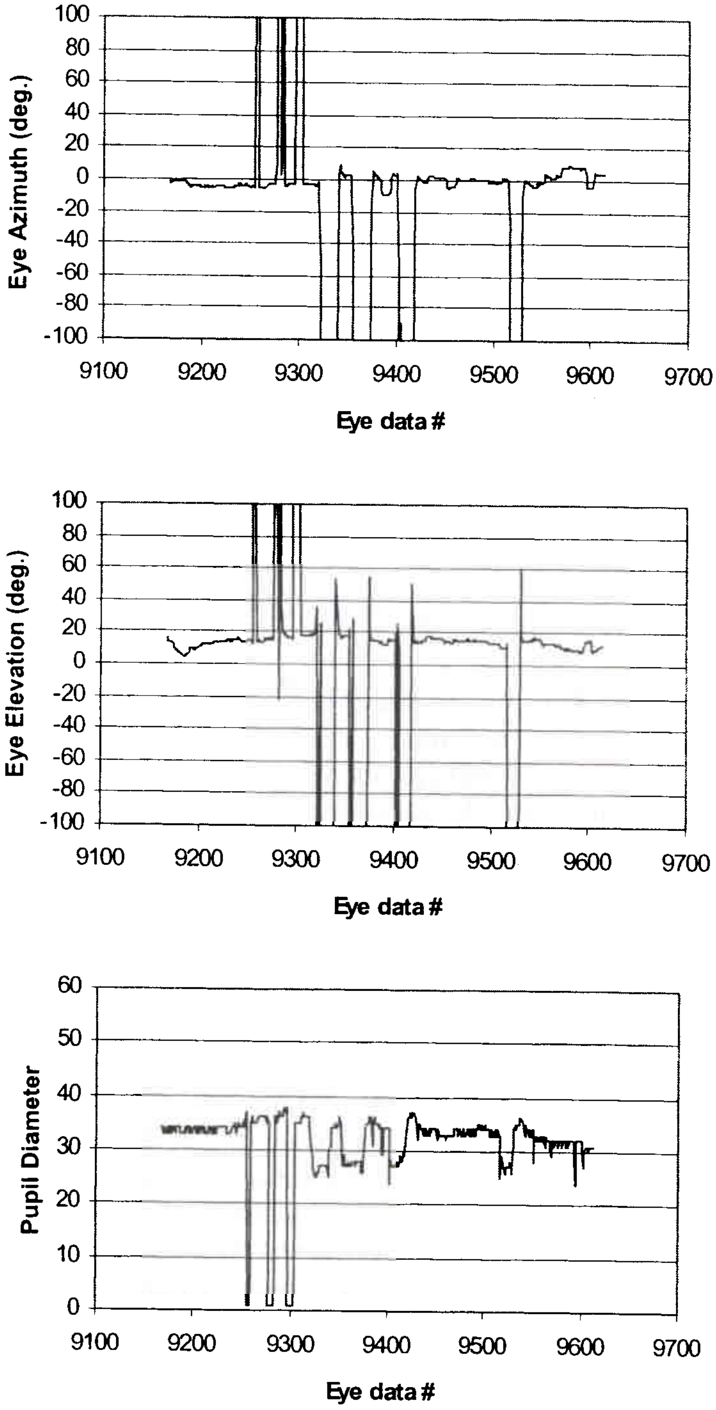

Figure A-2.7-6: Eye Azimuth, Eye Elevation and Pupil Diameter

When the filtering process is applied to the road test data, it is typical to see 10 % to 20 % of the data marked as bad. The error rate for planes other than plane 1 may be significantly higher.

The data rate for the eye-tracker is 30 Hz. To fill a short span of bad data that lasts for less than 0.2 second (i.e., 1 or 2 epochs in duration), the algorithm checks the eye plane number right before the span and after the span. If the numbers are the same, the bad data will be replaced using that plane number. In general, such short spans of bad data are results of blinking. To recover longer spans of bad data, the algorithm relies on the head direction measured by the magnetic sensor to deduce the plane of regard that the subject is looking at. Distributions of head azimuth for subject F61 are plotted in Figure A-2.7-7. Since we are particularly interested in eye gaze direction right before and after lane changes, only data during those periods are included in the plot. The lane change period is defined to start 10 seconds before a lane change mark and to end 5 seconds afterward. An offset equal to the mean head azimuth for plane 1 for each run is subtracted from all the data for that run so that the distribution for plane 1 always centers around 0. Data from all runs are then summed to produce the figure which shows the head azimuth distributions for the various planes based on the good data that pass the filtering process. The top curve includes all the data, both good and bad, for all the planes. The difference between the top curve and the other curves represents the bad data that has been filtered out. Based on the distributions shown, ranges of head azimuth are defined for a set of planes. During the lane change periods, the subject should spend most of the time looking at plane 1 (straight ahead), plane 5 (rearview mirror), plane 8 (right outside mirror), plane 9 (right window), plane 10 (right over-the-shoulder), plane 11 (left outside mirror), plane 12 (left window) and plane 13 (left over-the-shoulder). As a result, all bad data during the lane change periods will be assigned to one of these planes. The other planes are not assigned in the recovery process. Table A-2.7-2 lists the head azimuth ranges defined for subject F61. A plane of regard is assigned when the head azimuth angle falls inside its range defined in the table. The boundaries between planes are chosen manually by examining the distributions. Since the distributions overlap, one cannot be sure that a plane assignment near a plane boundary is correct. However, by choosing the boundaries at the intersections between adjacent distributions, the chance of making a correct plane assignment is higher than the chance of an incorrect assignment. Such a table is defined for each subject based on his or her good data.

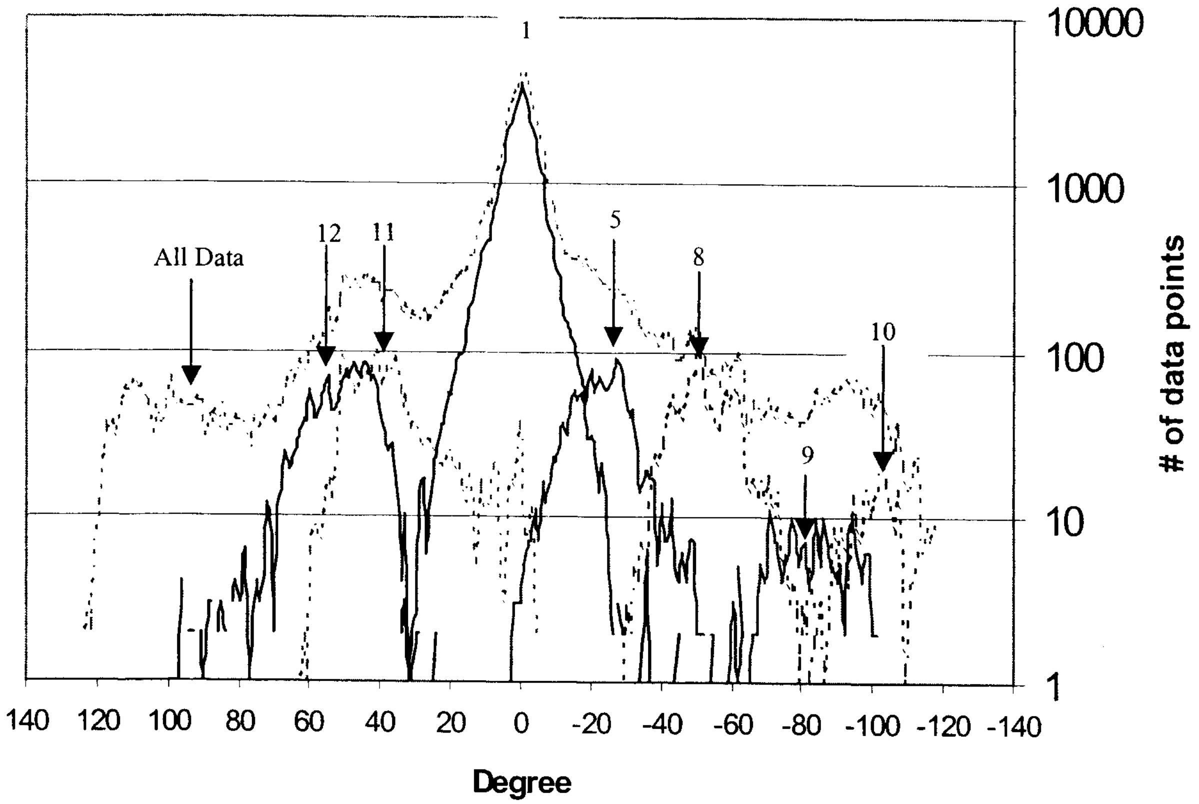

Figure A-2.7-7: Head Azimuth Distributions for Subject F61 During Lane Change Periods

Plane	Head Azimuth
10	< -85
9	-85 to –70
8	-70 to –35
5	-35 to –15
1	-15 to 25
11	25 to 55
12	55 to 80
13	> 80

Table A-2.7-2: Head Azimuth Ranges used in the Recovery Process for Subject F61

Figure A-2.7-8 shows the eye plane data output by the eye-tracker, the eye plane data recovered by the correction algorithm and the head azimuth data for the 15 seconds period around a right lane change during a run by subject F61. The beginning of the lane change (not shown in the figure) was marked by the operator at eye data number 23265. The plot of recovered eye plane data clearly shows the subject checking the rearview mirror (plane 5), the right outside mirror (plane 8), over-the-shoulder (plane 10) and the rearview mirror (plane 5) before the lane change. She then checks the right outside mirror again when she is initiating the lane change.

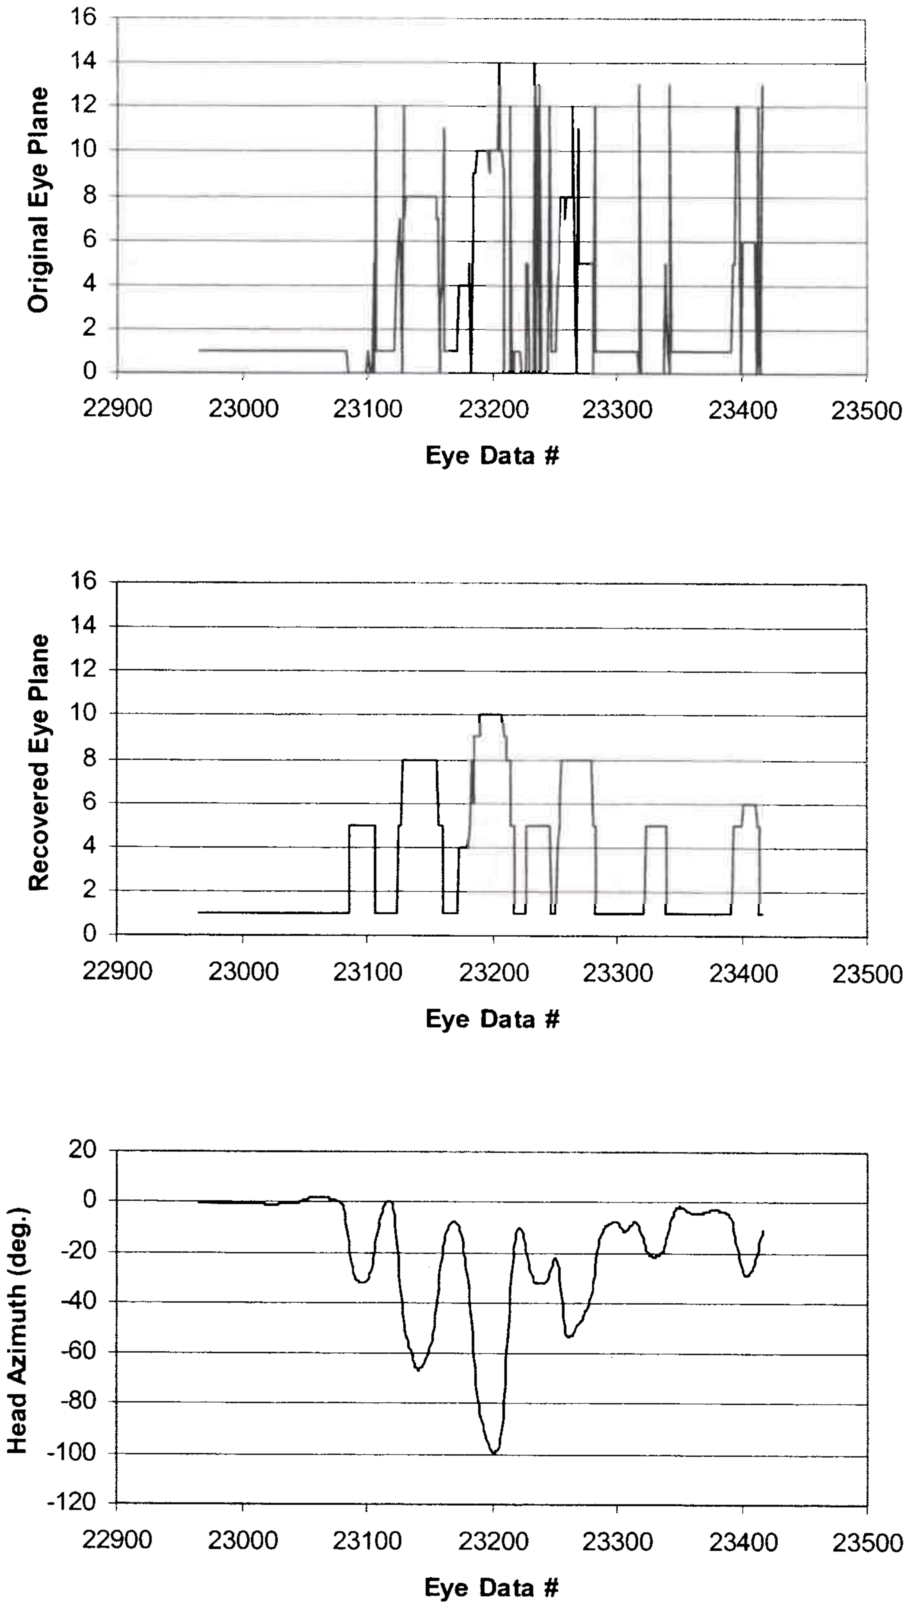

Figure A-2.7-8: Eye Plane Data Output by the Eye-Tracker, Recovered Eye Plane Data and Head Azimuth

By examining the data, we find that the head azimuth distributions outside the lane change periods are not as well separated as those within the lane change periods. The head azimuth distributions for subject F61 outside the lane change periods are plotted in Figure A-2.7-9. In general, the distributions have broader peaks than those within the lane change periods. The differences are especially pronounced for planes 5 (the rearview mirror) and 11 (the left outside mirror). In both Figures A-2.7-7 and A-2.7-9, plane 11 shows a double peak distribution. However, the right peak in Figure A-2.7-9 is much more prominent than the one in Figure A-2.7-7, indicating that the subject tends to use less head movement to look at the left outside mirror when she is not changing lane. The same conclusion can be drawn for the rearview mirror by comparing the distributions in the two figures.

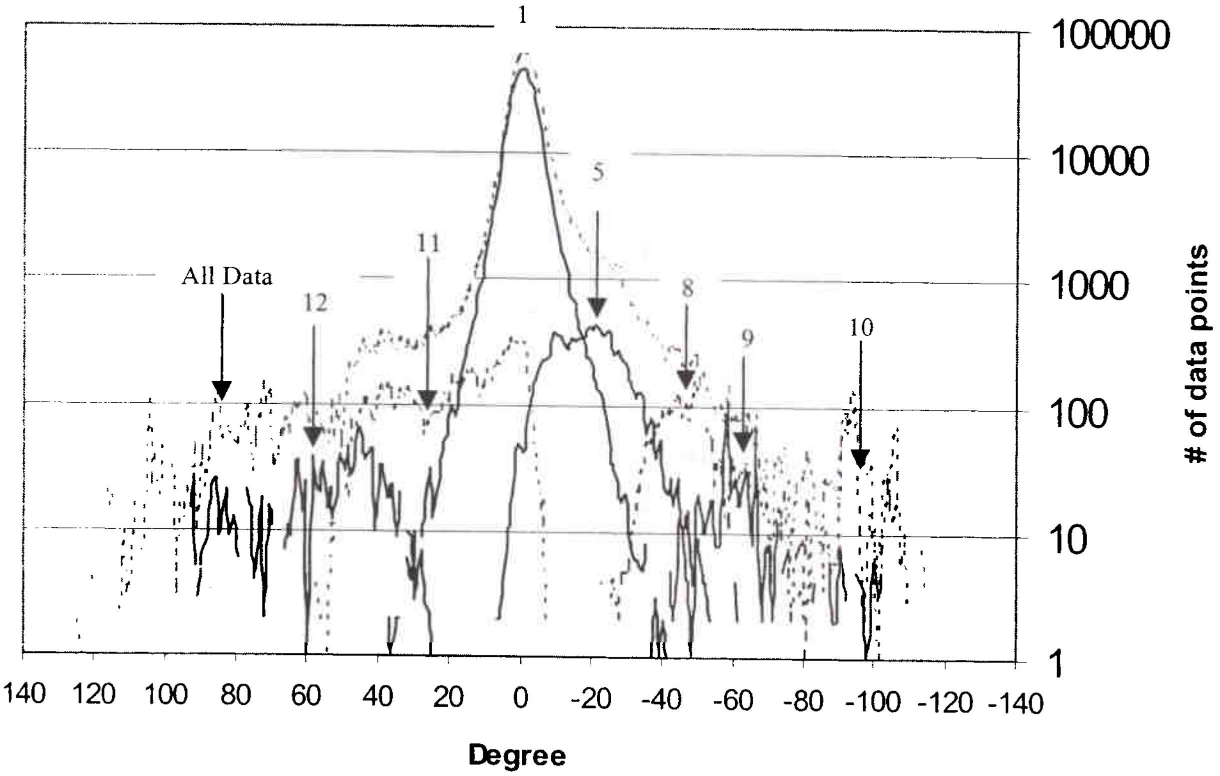

Figure A-2.7-9: Head Azimuth Distributions for Subject F61 Outside Lane Change Periods

Since the distributions are not as well separated, plane recovery based on head azimuth alone becomes more difficult. Outside of the lane change periods, a maximum likelihood method based on both the head azimuth and the head elevation is used. For each run, statistics (mean and standard deviation) for the head azimuth and the head elevation for each plane are compiled using good data that pass the filtering process. To avoid corruption by data taken when the car is not moving, the algorithm checks the velocity of the testbed and uses data taken only when the testbed is in motion. Assuming Gaussian distributions, the likelihood for each plane is defined as

$$P_i = \frac{1}{\sigma_{az,i}} e^{\frac{-(\theta_{az}-\mu_{az,i})^2}{2\sigma_{az,i}^2}} \frac{1}{\sigma_{el,i}} e^{\frac{-(\theta_{el}-\mu_{el,i})^2}{2\sigma_{el,i}^2}}$$

where θ_{az} = head azimuth
$\mu_{az,i}$ = mean head azimuth for plane i
$\sigma_{az,i}$ = head azimuth standard deviation for plane i
θ_{el} = head elevation
$\mu_{el,i}$ = mean head elevation for plane i
$\sigma_{el,i}$ = head elevation standard deviation for plane i

To recover the bad data, the likelihood function is evaluated for planes 1 to 13. The plane with the highest likelihood is assigned.

A-2.7.1 Head-Eye Coordination

Figure A-2.7-10 shows both the head azimuth and the eye azimuth as the subject switches his line-of-sight from plane 1 (straight ahead) to plane 11 (left outside mirror) and then back to plane 1. By closely examining the data, one can see how his eye movement is coordinated with the head movement in the process. Both the head and the eye start out pointing straight ahead between time marks 1 and 2 annotated in the figure. At time mark 2, the subject switches his line-of-sight to the left outside mirror by turning his eye quickly to the left as his head starts to turn slowly in the same direction. Between time marks 2 and 3, his head continues to turn left. Meanwhile, the eye azimuth decreases from its peak in coordination with the head movement such that the subject's line-of-sight remains fixed at the left mirror. At time mark 3, the head azimuth reaches its peak and the subject starts to return his head back towards the forward direction. As the head turns right between time marks 3 and 4, the subject reverses his eye motion in coordination with the head motion such that the line-of-sight remains at the left mirror. At time mark 4, the subject switches his line-of-sight back to plane 1 by quickly turning his eye to the right. Between time marks 4 and 5, both the head and the eye slowly return to the forward direction. In summary, the subject switches his line-of-sight quickly from one plane to another by executing rapid eye motions. After the transition, the eye moves in coordination with the relatively slow but smooth head motion such that the line-of-sight remains fixed. After analyzing the data, all the subjects are found to have this pattern of head-eye coordination.

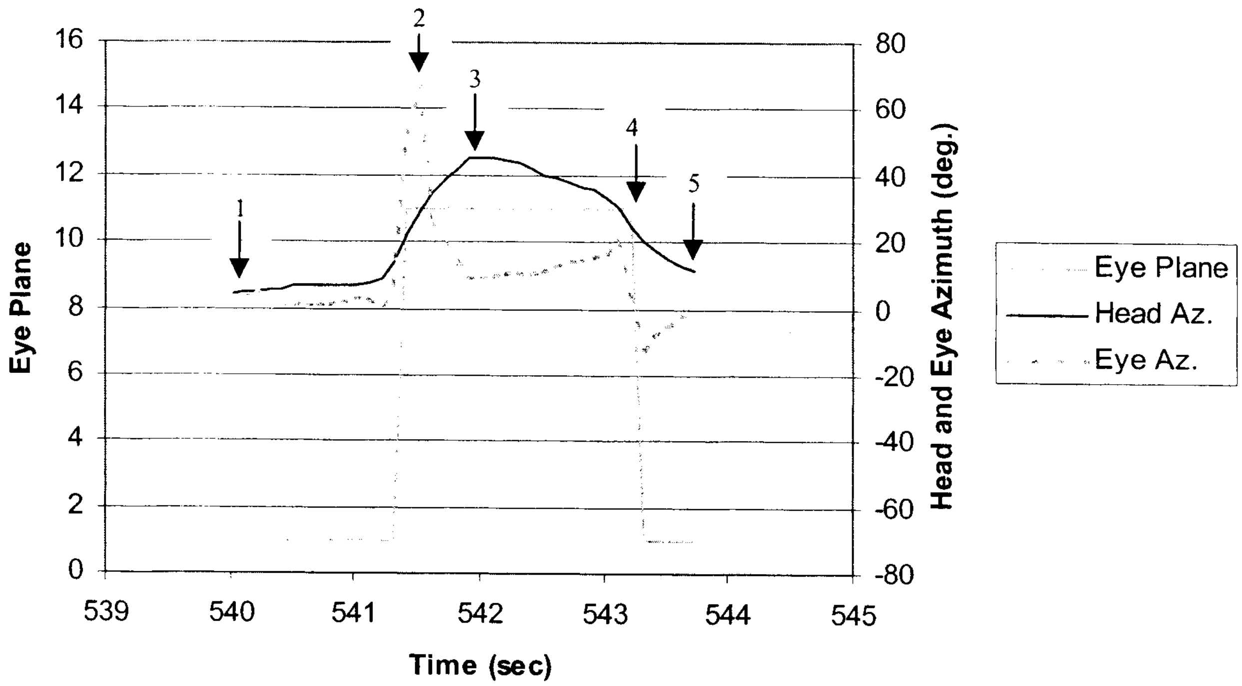

Figure A-2.7-10 Head-Eye Coordination

A-3.0 Static and Semi-Dynamic Tests

These tests verify the operation of the testbed as a collision avoidance system and as a data acquisition system. The parameters tested in this section are essential to the accurate operation of threat detection, target selection and warning activation. Specifically the parameters measured are field of view, sensor range, range measurement accuracy (both statically and dynamically), speed calibration and detection latency.

A-3.1 Detection Area and Range Verification

The purpose of this test is to verify the ability of the laser subsystem to accurately place an object in both range and azimuth. This was done by using the DGPS in a survey like mode, wherein one receiver antenna was always kept fixed on top of the laser scanner. In doing so, all points could be referenced as to this point with delta north and delta east values.

The axis of the car was established by placing the moving antenna at the midway point of the trunk and then the midway point of the hood. The two vectors from the laser scanner to these points would then be subtracted, resulting in a vector that defines the axis of the car. Next the portable antenna was moved to a series of points along the laser zone boundary as defined in Figure A-3.1-1. This line is defined to be parallel to the vehicle axis, at a distance of 11 feet from the side of the car. The numbers in the figure refer to the different zones wherein targets are tracked. The antenna was placed on a vertical pole of approximately 1.5 inch diameter. The placement of this pole was determined by an experimenter to be as coincident as possible with the line labeled laser zone boundary. In each case, the distance between the antennas was measured using a tape measure so as to verify the accuracy of both the GPS and the laser.

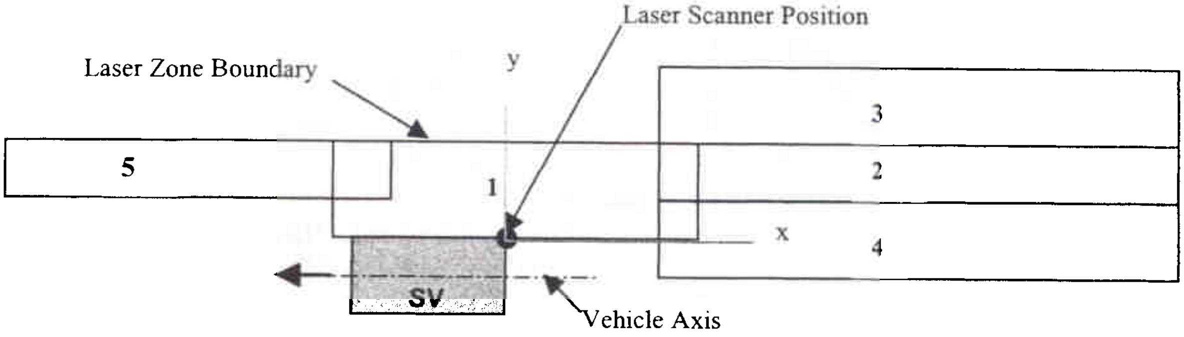

Figure A-3.1-1 Laser subsystem verification

Figure A-3.1-2 shows the car axis and the laser zone boundary plotted in GPS coordinates, i.e. on a north-east grid. The equations for the two lines are as shown, and there is a 0.63259° difference in the slopes. The laser coordinate system was then rotated by exactly this amount so that the laser zones would be parallel to the vehicle axis.

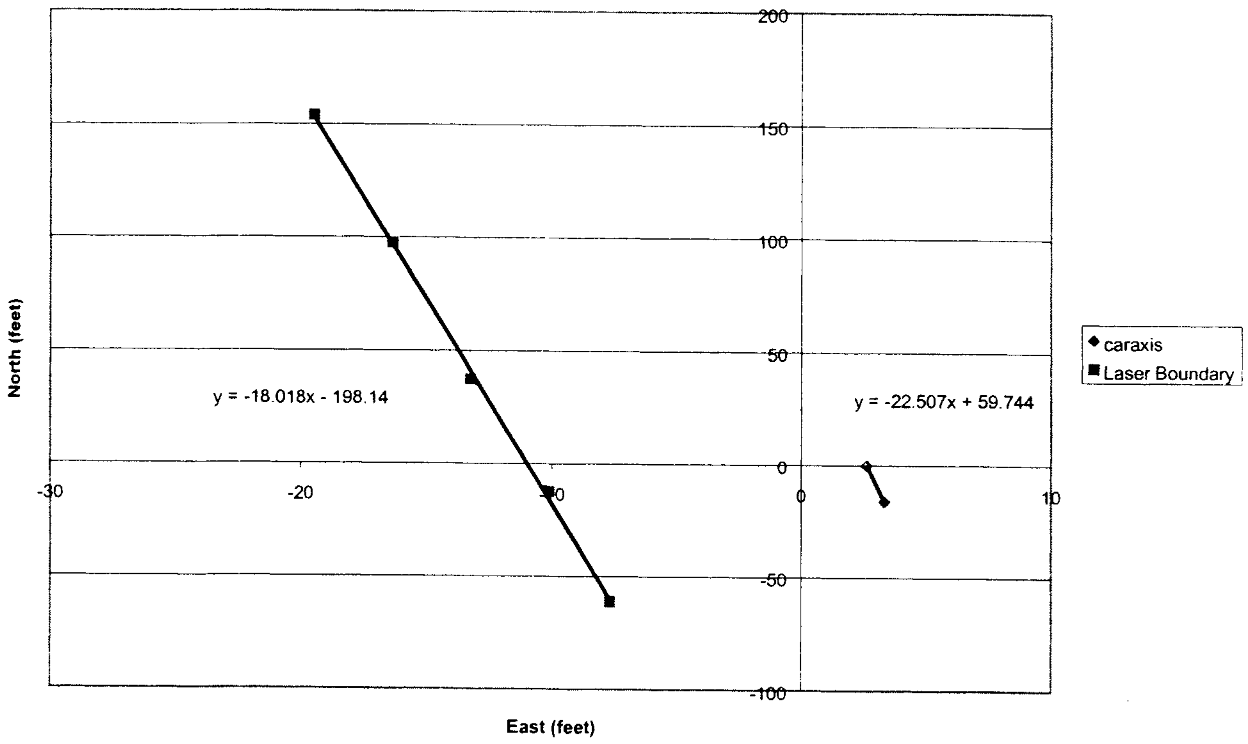

Figure A-3.1-2 Determination of car axis and laser axis by DGPS

Having established the laser coordinate system with the five points shown in the above figure, we can now transform the DGPS coordinates into laser coordinates and compare the accuracy of the measurement. The origins for both systems are coincident at the laser scanner mounted on the right rear corner of the testbed. Therefore the transformation is a simple rotation about the origin. Table A-3.1-1 summarizes the result. The coordinates marked X_{laser}, Y_{laser} are the x and y coordinates in the laser frame. These are derived from the range and azimuth angle outputted by the laser subsystem. The next two columns labelled (DGPS) are the coordinates in the DGPS reference frame. Range (DGPS) is computed from these two values. X_{DGPS}, Y_{DGPS} are the transformed DGPS coordinates in the laser frame. The column marked Tape is of course the distance as determined by the tape measure.

Point #	X_{laser}	Y_{laser}	Range (laser)	East (DGPS)	North (DGPS)	Range (DGPS)	X_{DGPS}	Y_{DGPS}	Tape
1	-60.89	11.35	61.94	-7.66	-60.99	61.47	-60.48	11.03	61.5
2	-11.80	10.94	16.09	-10.13	-12.22	15.87	-11.64	10.79	15.9
3	38.52	11.37	40.17	-13.29	37.49	39.77	38.16	11.20	39.75
4	98.35	11.05	98.97	-16.37	97.50	98.87	98.26	10.94	98.7
5	-50.97	7.35	51.50	-4.31	-50.47	50.65	-50.15	7.10	50.6

Table A-3.1-1 Comparison of DGPS and Laser coordinates (all units in feet)

Table A-3.1-2 summarizes the accuracy of the DGPS and laser measurements. In comparing the DGPS measurements with the tape measure, we see that the agreement is excellent, with a difference of less than one inch. Comparing the laser with the DGPS we see a secular offset of about 0.4 feet with a deviation of 0.25 feet. This latter figure falls within the specified accuracy regime of the laser rangefinder of 0.15 to 0.33 feet. Most of the inaccuracy seems to be in the determination of the X coordinate.

Point #	R(DGPS) – Tape	R(DGPS)-R(Laser)	$X_{DGPS} - X_{laser}$	$Y_{DGPS} - Y_{laser}$
1	-.0270	-0470	0.418	-0.324
2	-0.031	-0.226	0.165	-0.155
3	0.023	-0.393	-0.359	-0.171
4	0.166	-0.101	-0.089	-0.114
5	0.049	-0.848	0.821	-0.249
Mean	0.036	-0.408	0.191	-0.202
Std Deviation	0.072	0.255	0.408	0.075

Table A-3.1-2 Accuracy of the measurements (all units in feet)

The general conclusion that can be drawn is that the laser system is quite accurate in determining the positions of objects with respect to the host vehicle. It is within the accuracy requirement of .5 ft., as detailed in reference [1]. It is also close to the manufacturer's stated accuracy of .33 ft.

The final measurement in this section was to refine the measurement of the detection area. The laser has a azimuthal FOV (field of view) of approximately 270°. However part of this FOV is obstructed because the back and side of the SV have a slight bow to them. Therefore a test was done to determine how much of the FOV is obstructed. The pole with the mobile DGPS antenna on it was positioned about 25 feet in front of the car. The pole was moved in incremental steps in the direction away from the side of the car until it moved out of the shadow of the side of the car and finally could be seen by the laser. A similar measurement was done along the rear of the car, again marking the point at which the pole could first be seen. The results for the angles indicated in Figure A-3.1-3 are a value for α of 7.5°, and a value for β of 5.7°.

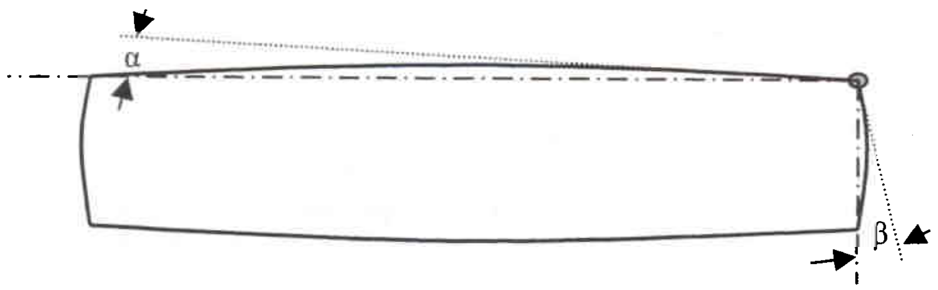

Figure A-3.1-3 Illustration of shadowing of laser by vehicle side and rear

A-3.2 Detection Latency

Detection latency is a key parameter of any remote sensing system. Latency is defined as the delay in system response to an instantaneously introduced object. The only way to introduce an object instantaneously to the testbed is from the side into the proximity zone as illustrated in Figure A-3.2-1. The driver warning algorithm is designed to be sensitive only to presence in the proximity zone, and independent of velocity. Entering the testbed longitudinally (i.e. in a direction parallel to the axis of the testbed) would activate the fast approach warning which is velocity sensitive. Entering the zone from the side at increasing speeds will increase the penetration into that zone. Taking the slope of the curve of penetration depth versus speed yields the detection latency.

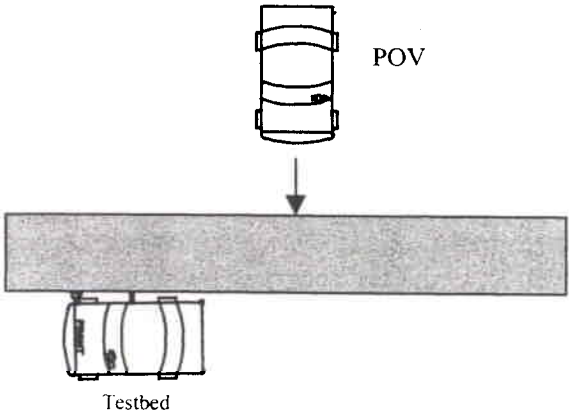

Figure A-3.2-1: Test Arrangement for Detection Latency

The procedure for this test is as follows. The key to this test is the use of DGPS as the yardstick by which the CAS system is measured. The DGPS will be used to measure the relative positions of the testbed and POV at the time when the POV crosses into the proximity zone and initiates a warning. First step is to establish the axis of the testbed in GPS coordinates. This was done, as in section A-3.1, by stationing one of the GPS antennas on top of the laser scan head and the other (mobile) one at the midpoint of the trunk and then the midpoint of the hood of the car. Once the axis is established we then have a simple transformation between the GPS coordinate system and the laser system. We will use the GPS coordinates transformed to the laser coordinate system in plotting our results.

The latency tests were performed by having the POV traverse the proximity zone, approaching from the passenger side, as indicated in Figure A-3.1-1, and also from the driver side. The speed at which the POV traversed the proximity zone was varied from 5 to 35 miles per hour. For these tests, the GPS antenna was attached to the POV at the midpoint of the hood of the car. In operation, the CAS warning system will not turn on unless the car is in motion. Since the testbed cannot be in motion for this test, a software switch was inserted into the code to fool the CAS processor into thinking that the vehicle is in motion. The exact time at which the proximity warning is energized is recorded and in post-processing, the GPS position of the POV with respect to the antenna fixed to the laser scan head can be calculated. In all the plots that follow, distances are referenced to the bumper of the POV as the forward most point of that vehicle.

An example of the results obtained is shown in Figures 3.2-2 and 3.2-3, where the driver of the POV approached the proximity zone from the testbed driver's side and passenger side respectively. Both curves plot the region out to 11 feet perpendicular to the testbed axis , which is the lateral extent of the proximity zone.

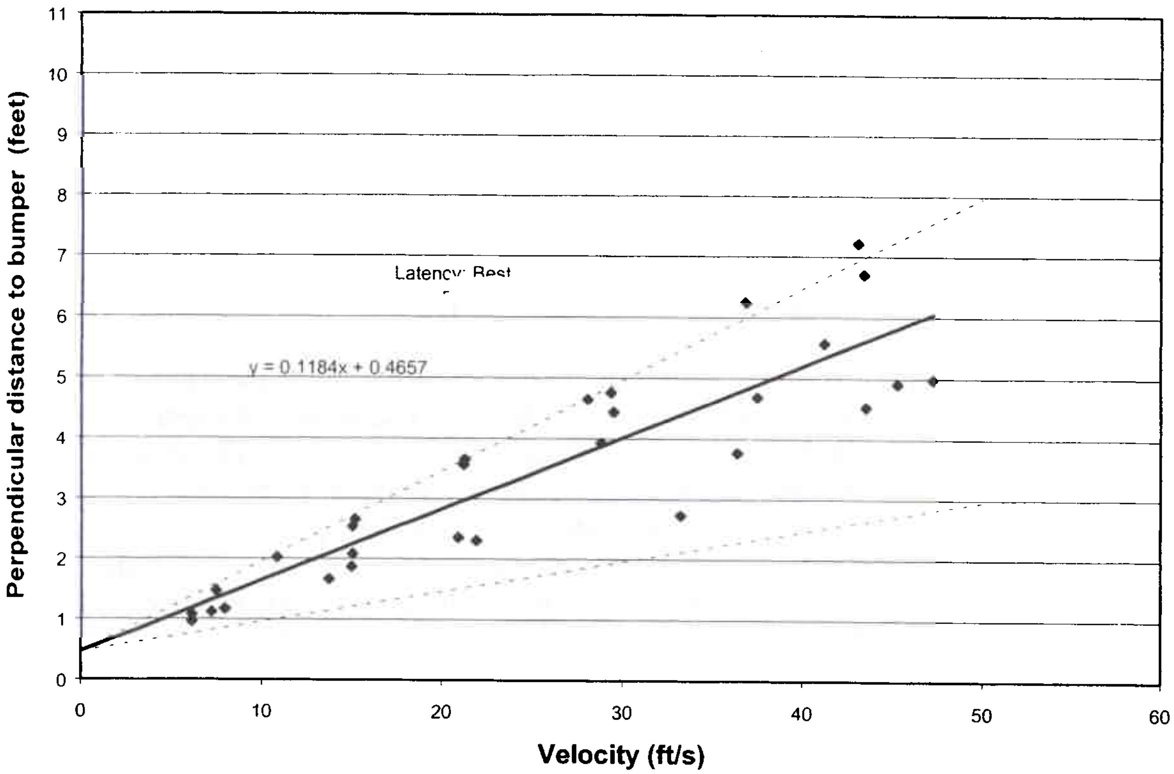

Figure A-3.2-2: Latency: Approach from Driver Side

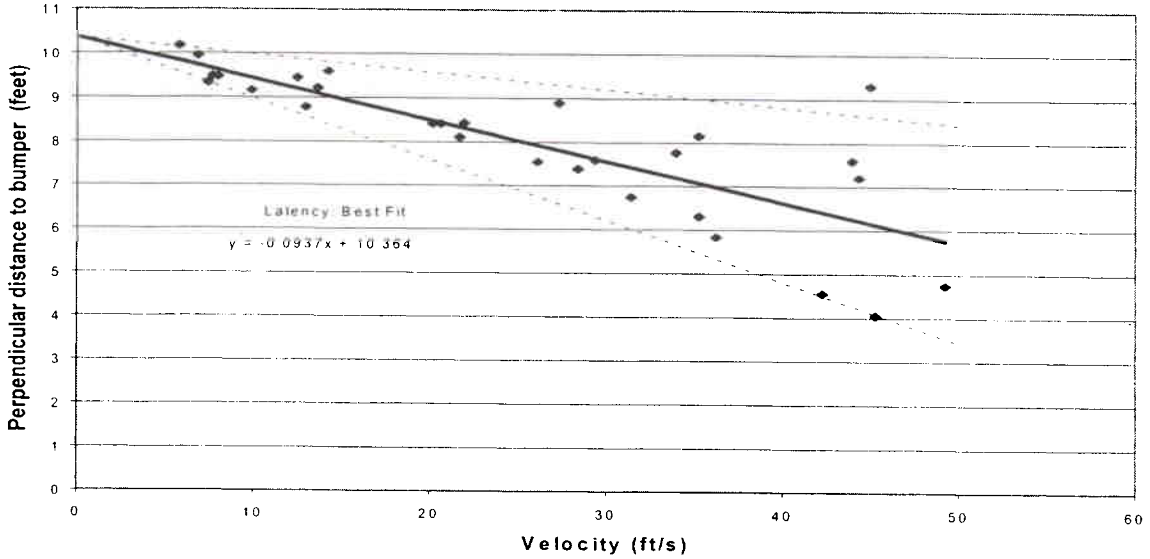

Figure A-3.2-3: Latency: Approach from Passenger Side

From the slopes of the lines, we calculate a latency time of 118 msec and 94 msec. The explanation for the difference can be understood from the operation of the scanning laser in conjunction with the DSP processor. The scanner rotates in a counterclockwise direction, when view from above. The scan cycle begins when the laser beam is pointed along the rear of the car. The first 270° is useful data that will be processed. The next 90° of the cycle represents data taken while the laser is pointing at the testbed car body. This data will be discarded. Consider now a vehicle approaching from the driver's side of the testbed. If the body of the car is just over the boundary, the laser scanner must still traverse one-half a cycle (50 msec) before the data can be processed, and a decision is to be made as to whether the system has detected a presence in the proximity zone. This is one extreme. The other extreme is for the approaching vehicle to be just behind the boundary when the laser passes the first time, so that it is not until the laser makes a second complete revolution before its presence is detected. This other extreme case would take about 150 msec before reporting a warning to the driver. Assuming a flat distribution of arrival times, this would average out to 100 msec delay. We measure a delay of 118 msec. Using similar arguments for a target vehicle approaching from the testbed passenger side, we see that vehicle as being 45° "closer" to the cycle end point, so the expected latency time should be 89 $\underline{+}$ 50 msec. This is much closer to the measured value of 94 msec. Figures A-3.2-2 and A-3.2-3 contain the limit lines corresponding to the extremes in phase differences possible between the target car and the laser scanner. It is seen that nearly all the data falls within these limits.

The second thing to note is that the intercepts are not at 0 or 11 feet, corresponding to an approach from the driver or passenger side of the testbed respectively. The differences of approximately .5 and .6 feet suggest that the laser is not hitting the front bumper of the target car, but in fact is hitting some point on the front grill. This distance corresponds to a longitudinal distance of about 0.5 feet. This is a common difficulty with a laser radar sensor. Because of the small spot size, the reported range is very sensitive to pointing angle. In the present case where the testbed is stationary, the difference is small. As we shall see later on, this difference can be significant when both the testbed and target vehicle are in motion. The laser spot can hit the target vehicle anywhere from the front license plate to high up on the hood, a difference of several feet.

A-3.3 Range and Speed Verification

During these tests the accuracy of the target range and closing speed was verified. The tests were conducted while the testbed was parked and the POV driven toward it at speeds ranging up to 45 mph. The target POV had a magnetically mounted GPS antenna fixed to the left front portion of the hood. This is close to the point that the laser tracking algorithm would latch on to as the point on the POV that was closest to the SV. In plotting the results, all the distances are referenced to the front bumper of the POV, which is the area that would first contact the SV in any collision. The second GPS antenna was placed directly on top of the laser scan head.

For the data to be properly analyzed a discussion of timing is necessarily in order. Referring to Figure A-3.3-1, we see two lines marked scan begins and scan ends. Viewed from above the laser scanhead, the scanner rotates counterclockwise and acquires data over slightly more than three-fourths of a circle. At the point where the scan ends the computer marks the time as counted by the master clock in the DSP board. It is important to note that the tracked point on the POV was actually taken at some earlier time in the scan rotation. Meanwhile the GPS is recording data every tenth of a second exactly, as measured by the GPS clock. Therefore in order to compare the laser measurement to the GPS standard, we must perform the appropriate interpolations to insure that both systems are measuring the same quantity, at the same time.

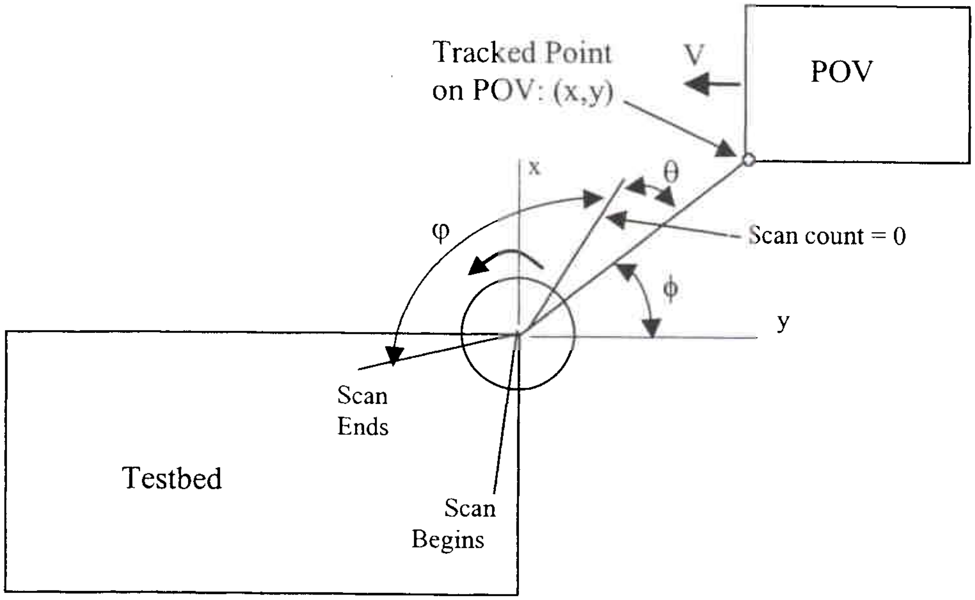

Figure A-3.3-1: Timing sequence of scanning laser

Since the DSP master clock was synched to the GPS clock at the beginning of the run, and updated every five minutes, it is only necessary to work backwards in time from the scan end to the time at which the tracked point on the POV was recorded. The laser scanner has its own internal coordinate system, whose azimuthal origin is labeled scan count = 0. This axis makes an angle of 66.38° with the laser y-axis. This was determined as part of section A-3.1, which established the laser axis as being parallel to the vehicle axis. This angle of 66.38° is equal to $\theta + \phi$ in the above figure. The angle φ is a direct measurement from the laser scanner encoder, while angle $\theta = 66.38° - \tan^{-1}(x/y)$. Therefore the time that must be subtracted from the time at the scan end is 0.1 sec * $(\varphi + 66.38° - \tan^{-1}(x/y))$, where the 0.1 sec comes from the rotation frequency of the laser scanner.

The actual tests of the accuracy of the laser subsystem were performed by having the POV drive past the parked testbed at speeds ranging from 5 to 25 mph. For this test the POV was a 1997 full size Dodge pickup truck. A useful measure of the accuracy is to plot the quantity (GPS range – Laser range) as a function of GPS range. In principle this quantity should be, and is, close to zero. Selected runs are shown in Figures A-3.3-2-5. The spread in the data shows a general increasing tendency to greater deltas at the larger distances. Still in this regard, the deviation from zero is less than 2 feet, which is less than 1% at the extreme distances of about 270 feet. The second set of data in Figure A-3.3-5 marked buffer zone, is the result of tracking the vehicle in the buffer zone because its lateral position was more than one lane to the right of the testbed. When the testbed enters the fast approach zone, it is quickly picked up by the tracking algorithm and then becomes the more accurate value because it is closer laterally to the testbed.

A graphical summary of an entire data set is illustrated in Figure A-3.3-6. For this data set, the POV was a 1989 Honda Accord. The Honda was driven past the stationary testbed at speeds up to 45 mph. The data points represent the speed and velocity of the POV, as measured by the GPS, at the time when the system issued a warning. The solid line represents the driver warning algorithm, with the time to enter the proximity zone set to 2.5 sec. The data is fairly closely clustered at or just below the line, as is expected. The vertical distance below the line is due to the lack of an adequate target return to the laser rangefinder. This is accentuated by those data points marked with ☐, which were taken with the pop-up headlights in the down position. Due to the sloping nature of the hood and the small vertical extent of the grill, the cross section for laser return is quite small. Therefore the car is not seen very well beyond 130 feet. When the headlights are in the up position, the retro-reflective nature of the lamp housing increases the cross section dramatically, and the car is seen out to almost 200 feet. The conclusion from this set of tests is that the performance of the sensing system has been validated.

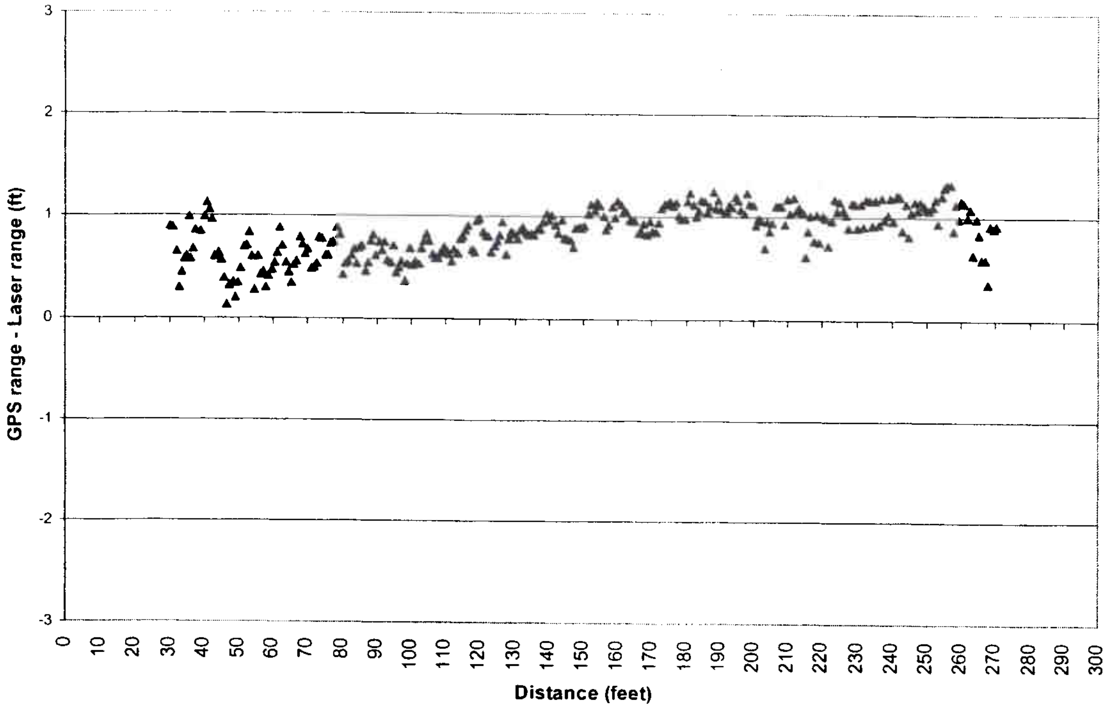

Figure A-3.3-2: Delta Range at 5 mph

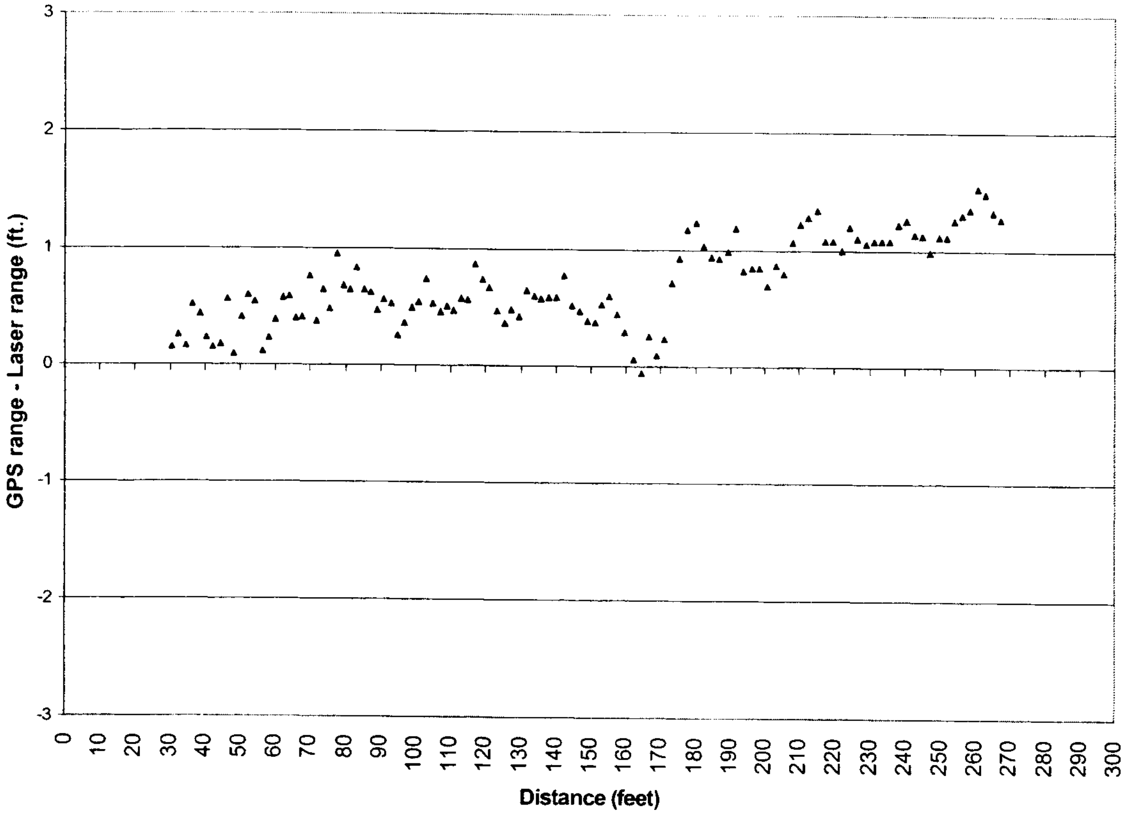

Figure A-3.3-3 Delta Range for ~10 mph

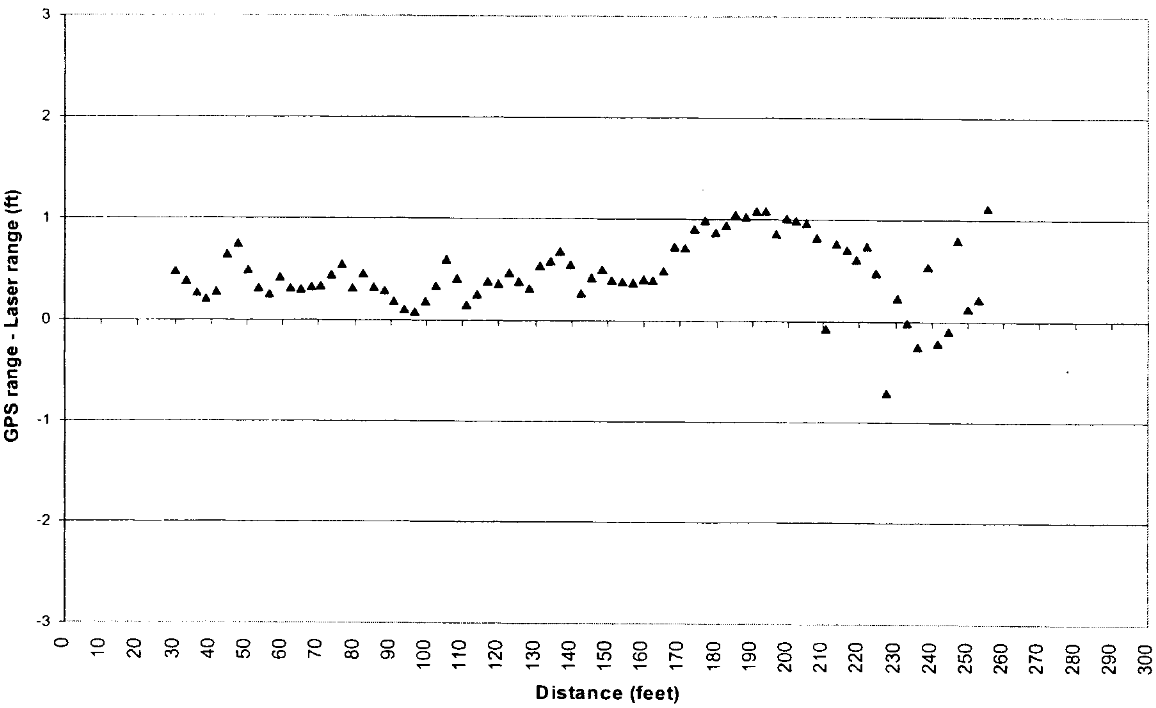

Figure A-3.3 –4: Delta Range for ~20 mph

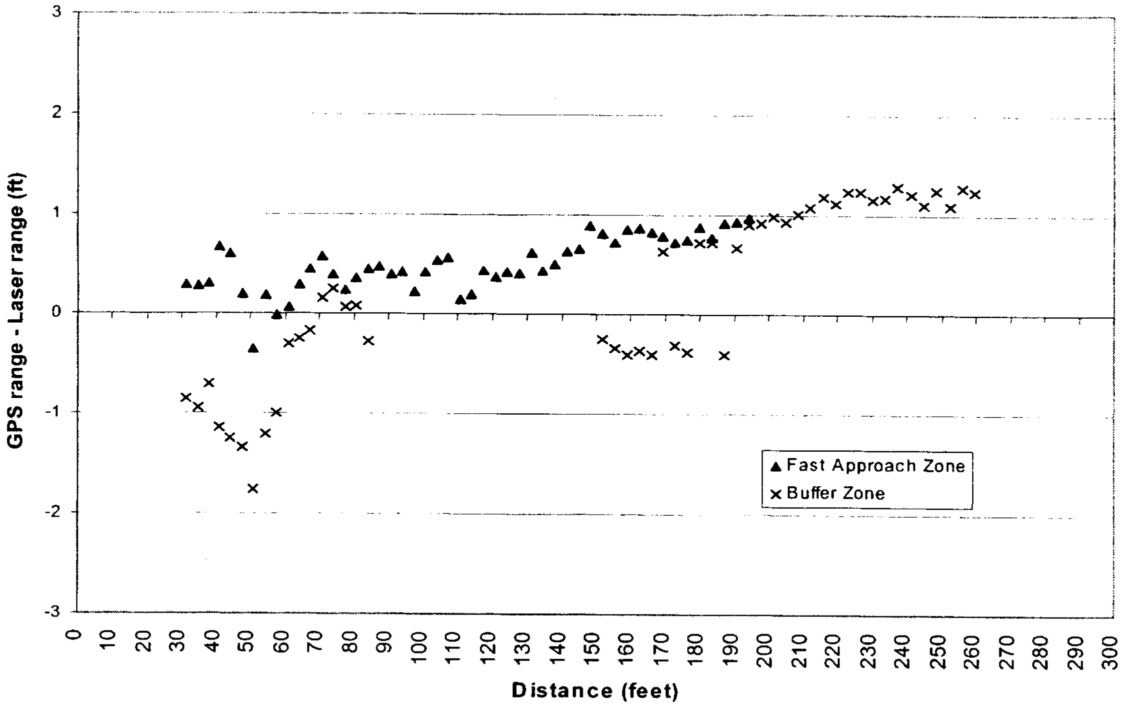

Figure A-3.3-5: Delta Range for ~25 mph

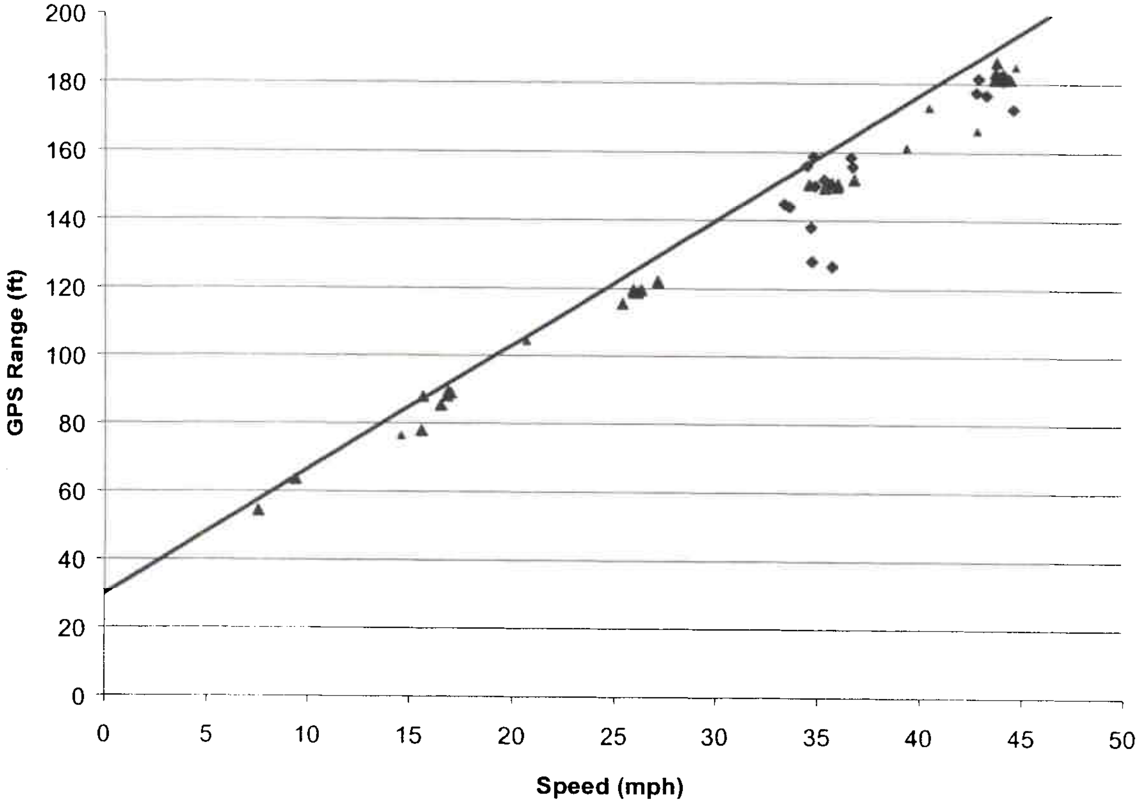

Figure A-3.3-6: GPS range at which fast approach warning is given vs. speed. Data marked with ◆ was taken with the headlights down; data marked with ☐ was taken with the headlights up.

The final test in this section was to check the accuracy of the speed measurements. Again the standard for this test is the GPS. For this test, the POV was parked and the testbed was driven past it. Valid data commenced when the testbed passed the POV. The POV was in this case the 1989 Honda Accord, which was parked in such a way that when the testbed passed it, the viewed target area was the rear of the car. This was done so as to give the laser a larger and more nearly vertical oriented target. For the Honda, the GPS antenna was mounted on the right rear corner, as it was for the testbed. The POV was passed at near constant velocity, and for each pass the velocity from each measurement system (GPS, laser system and the testbed's speedometer) was averaged. The results, plotted as a function of the GPS speed are shown in Figure A-3.3-7. A more instructive plot, in Figure A-3.3-8, shows the difference in velocity between the GPS and the speedometer and the GPS and the laser as a function of GPS velocity. The accuracy of both sensing systems is better than 2 mph at all velocities.

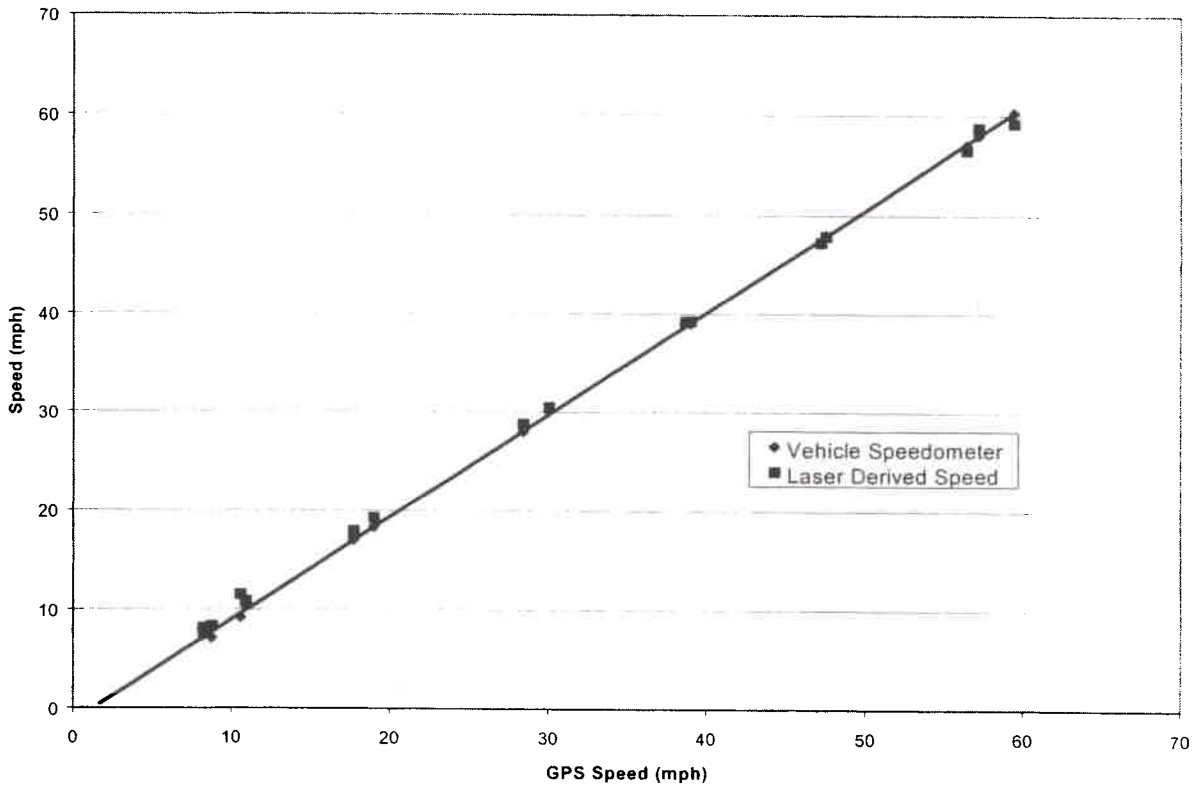

Figure A-3.3-7 Speed Validation

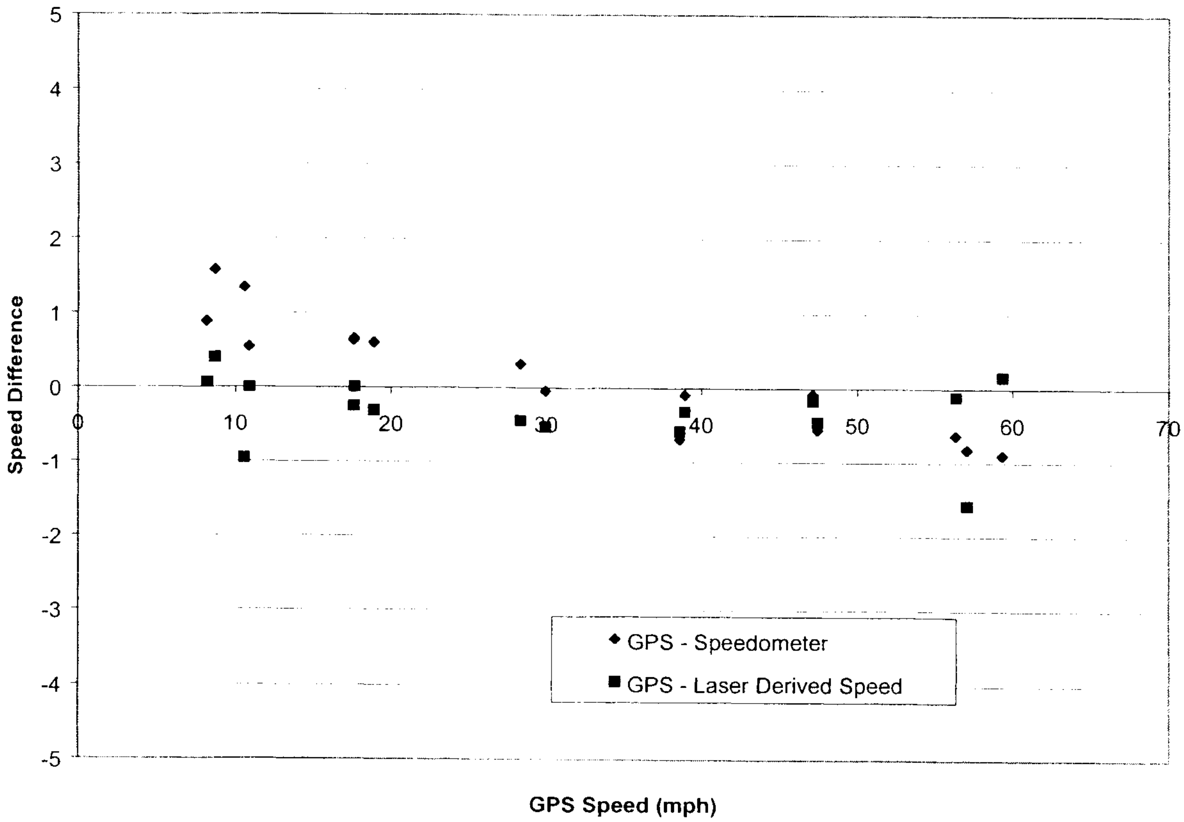

Figure A-3.3-8 Speed Accuracy

A-3.3.1 Motorcycle Tests

A special attempt was made to test the collision warning system on a motorcycle with rider. The test configuration was to park the testbed and drive the motorcycle past it at varying speeds and at two lateral distances – approximately 3 feet and 8 feet. Because the fast approach zone is indented 4 feet laterally from the side of the testbed and because of the narrow extent of the motorcycle, close approaches did not register until 30 feet, where the motorcycle entered the proximity zone. This purpose of this indentation was to reduce false positives coming from vehicles following the testbed in the same lane, but offset to the right enough to extend over into the fast approach zone. For this test, the motorcycle was equipped with a GPS system, with the antenna taped to the helmet of the rider. In this position it was unlikely that any satellite would be blocked during the test period.

The first graph, Figure A-3.3-9, shows the range and speed of the motorcycle, as measured by the GPS, when the fast approach warning first became illuminated. The values are clustered around the driver warning algorithm line, except for a few prominent cases. For this series of tests, the time to enter the proximity zone was set to 3.0 sec. Generally speaking, there are instances which occur when the motorcycle enters the fast approach zone late from one of the buffer zones. It can also be seen that the observable range is limited to less than 120 feet, implying that the motorcycle is a poor reflector.

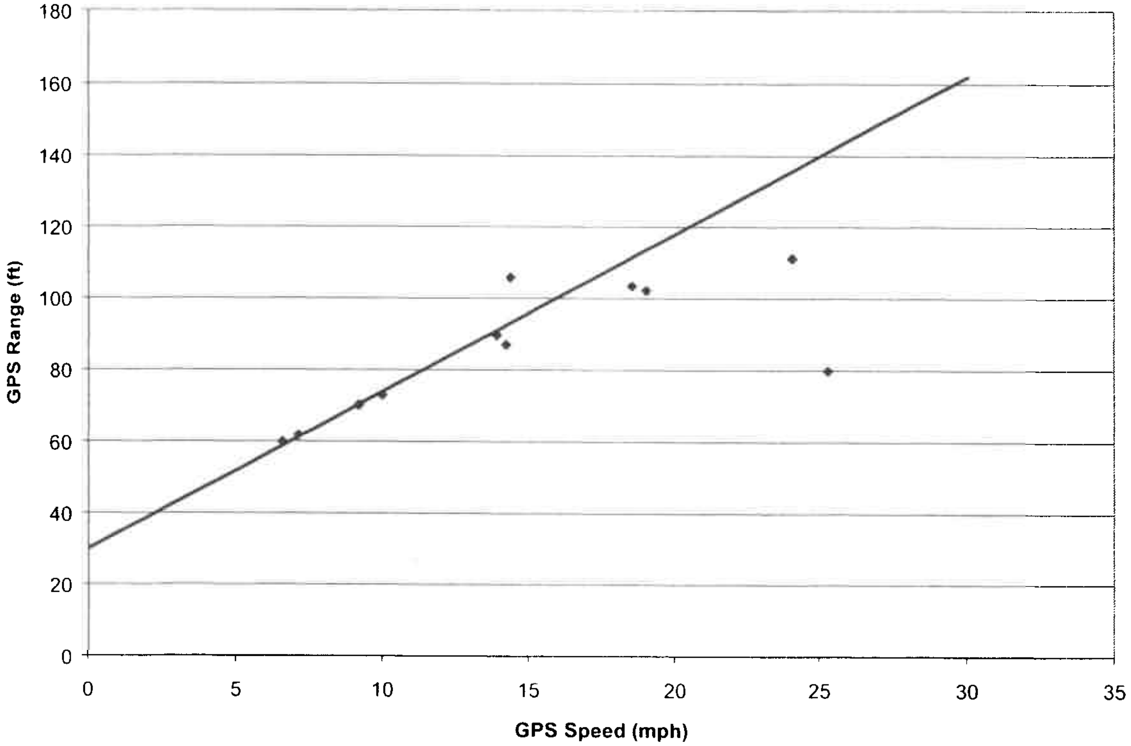

Figure A-3.3-9: Range and speed of motorcycle at fast approach warning onset

Just as was done for other target vehicles, we now plot the difference between the GPS measure of range and the laser measure. Although the GPS antenna is taped onto the helmet of the rider, the reference point, as always, is the most forward point of the vehicle. For the motorcycle this point is the leading edge of the front tire. As can be seen from the following figures, the laser consistently targets points on the motorcycle behind the front tire. At large distances the target points are as much as four feet behind the front tire. Generally speaking the rider's clothes represent a better target than any of the shiny reflective surfaces of the motorcycle itself. This is because the clothes will reflect diffusely while the metal surfaces reflect specularly. At large distances these specular reflections are mostly forward directed and hence are not seen by the receiver optics. The clothes however will reflect energy back toward the laser and therefore are more likely to be seen. At close distances it is more likely that the laser will strike metal surfaces that reflect directly back to the laser, and hence the difference in range will decrease. In all of the following charts, the maximum range observed is generally speaking, less than that of a car and certainly less than a truck. It can also be deduced that the faster the motorcycle enters the fast approach zone, the later it is picked up as a target, and consequently enables the warning at a later time.

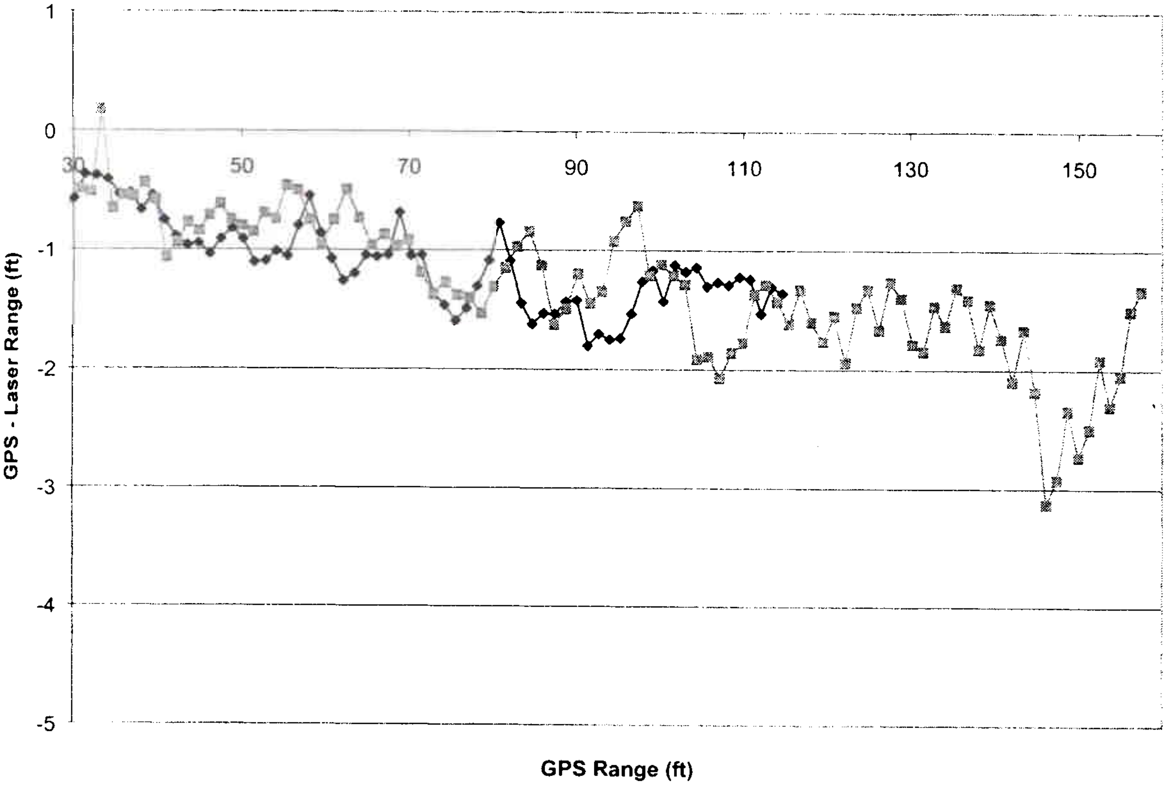

Figure A-3.3-10: Motorcycle at ~ 9 mph; two passes

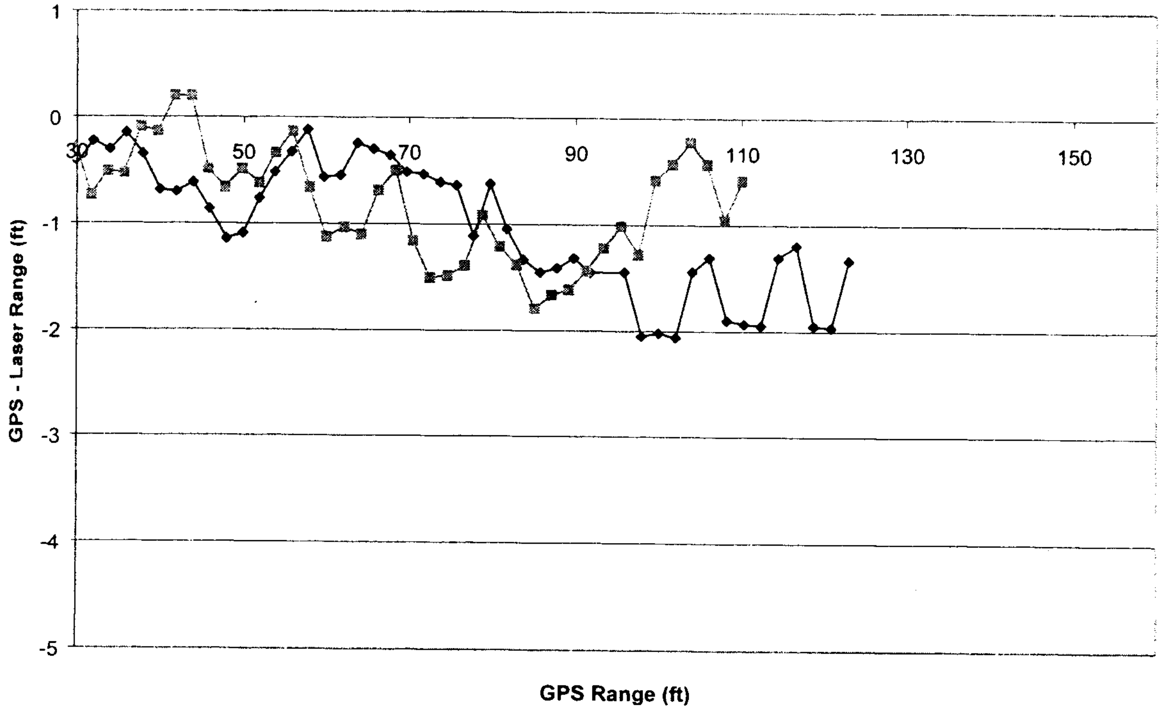

Figure A-3.3-11: Motorcycle at ~ 14 mph; two passes

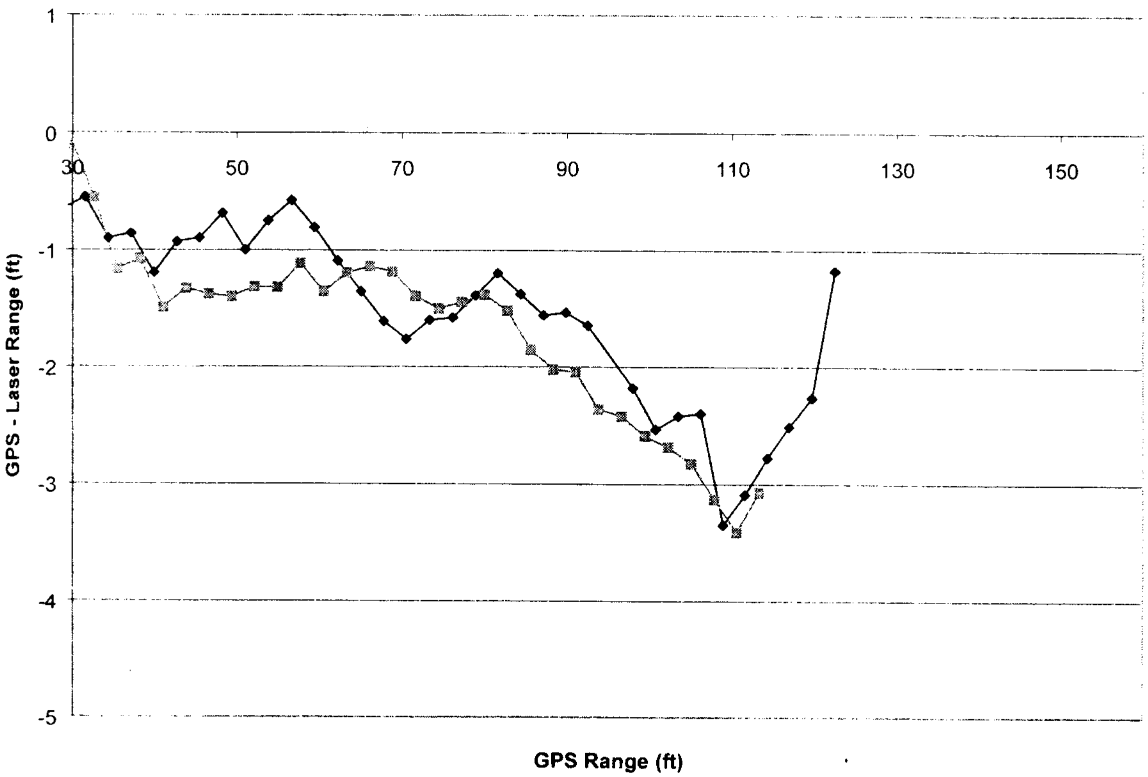

Figure A-3.3-12: Motorcycle at ~ 19 mph; two passes

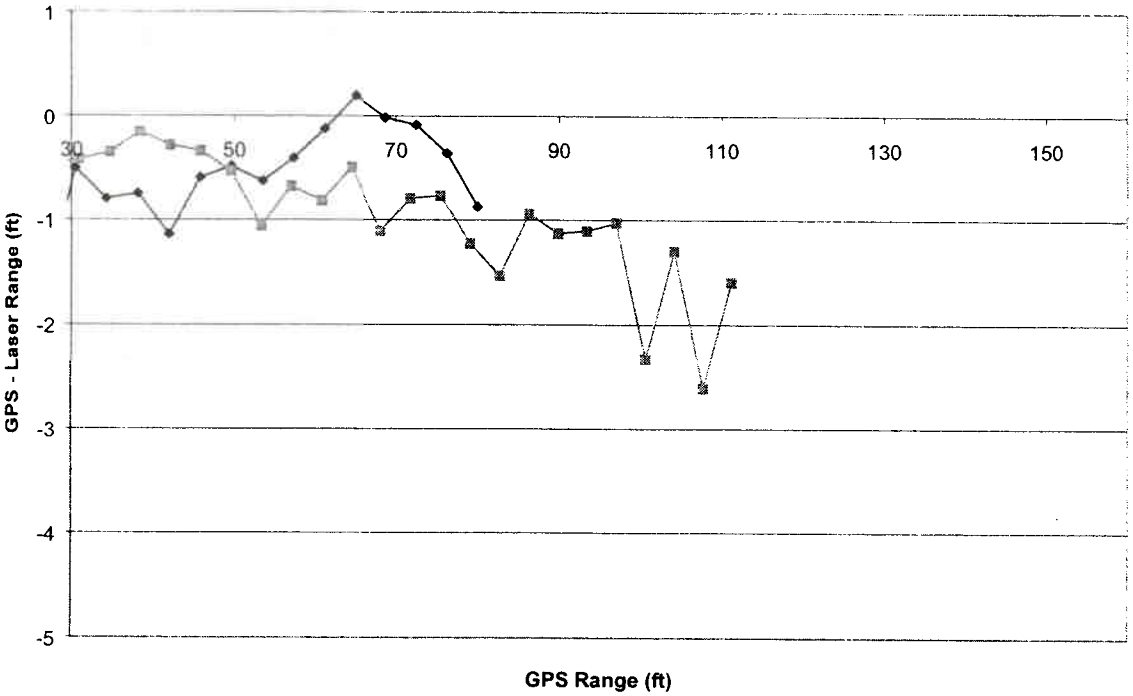

Figure A-3.3-13: Motorcycle at ~ 22mph; two passes

A-4.0 Track Tests

These tests are two-vehicle sequences which validate both the algorithm and testbed sensors at highway speeds. Differential GPS receivers are used to measure relative distances and relative velocities of the vehicles.

The breakdown of lane change crash scenarios into angle, sideswipe and rear end crashes all have in common the fact that the subject vehicle attempted to change lanes without adequate space between itself and the car in the lane to which it desired to move. Furthermore 85% of all lane change crashes occur with the relative speed of the two vehicles being within 15 mph. The test configurations are designed with these points in mind.

The different scenarios test the CAS's ability to handle a variety of POV closing speeds for a variety of different testbed vehicle speeds, a POV move into the adjacent lane from two lanes over, the approach and passing of a POV in the adjacent lane, and the effects of clutter such as parked cars, guardrails, etc. All of these tests were performed on the test track. The final tests in this chapter involve establishing the preferences of the individual team members for the two input parameters that make up the driver warning algorithm. These parameters are the time to enter the proximity zone and the extent of that zone behind the testbed.

All tests in this section were performed on the grounds of the California Speedway in Fontana, CA. A section of the outer perimeter road was cordoned off for these tests. The section of road consisted of a long straightaway (about 1.5 miles) followed by a curve of radius 1800 feet, followed again by a shorter straight section of about 0.75 mile. The procedure for all of the following tests involved taking GPS data with both systems for at least 20 minutes before and after the test runs, so that a forward and backward solution for the position could be compared.

A-4.1 Dynamic Warning Test

The purpose of this test was to verify the performance of the CAS in a classic lane change situation where the POV is approaching the SV from behind in the adjacent lane. The SV traveled at constant speeds of 20, 30 and 40 mph, while the POV passed the SV at speeds of SV +10, SV +20 and SV + 30 mph. Both straight and curved sections of track were used. Each situation was replicated at least five times. A 1989 Honda Accord was used as the POV. The standard of comparison used was of course the DGPS. With one antenna mounted on the right rear corner of the SV and the other mounted on the left front corner of the POV, this was a direct comparison with the laser measurements.

The results are summarized in the following two figures. Figure A-4.1-1 plots the position of the POV as a function of the relative speed, as determined by the GPS, for the straight section of road, while Figure A-4.1-2 plots the same data for the curved sections of road.

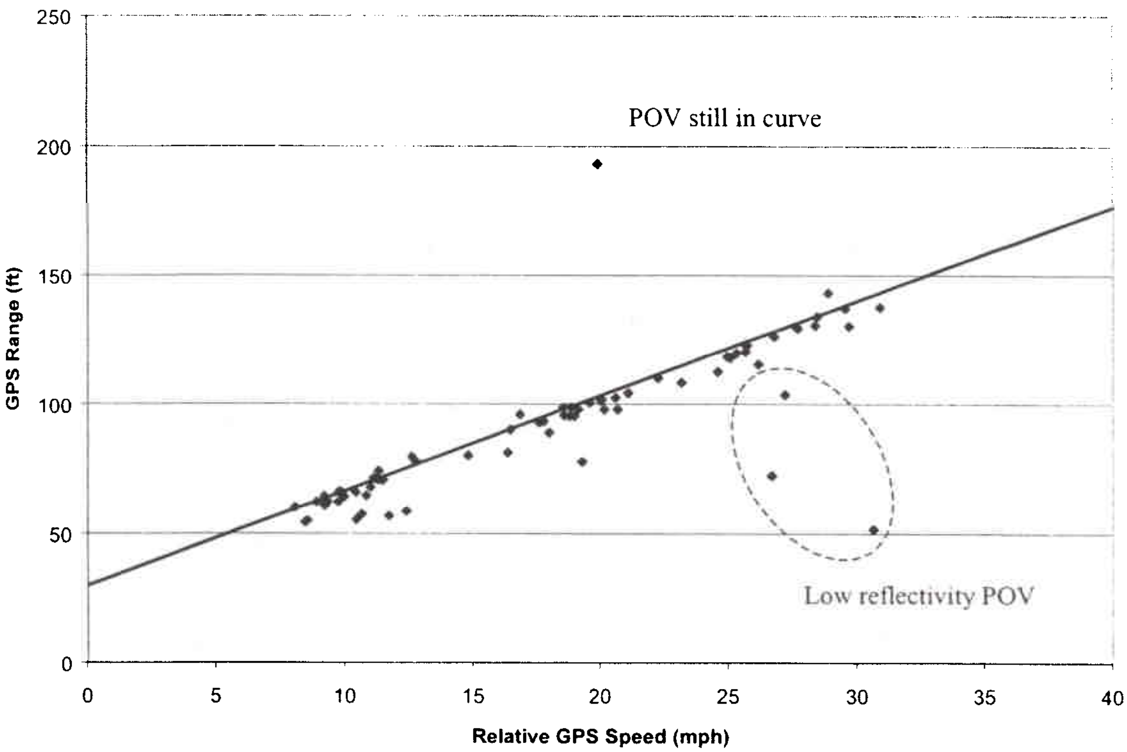

Figure A-4.1-1: Position and Rel. Speed at Onset of Driver Warning on Straight Road

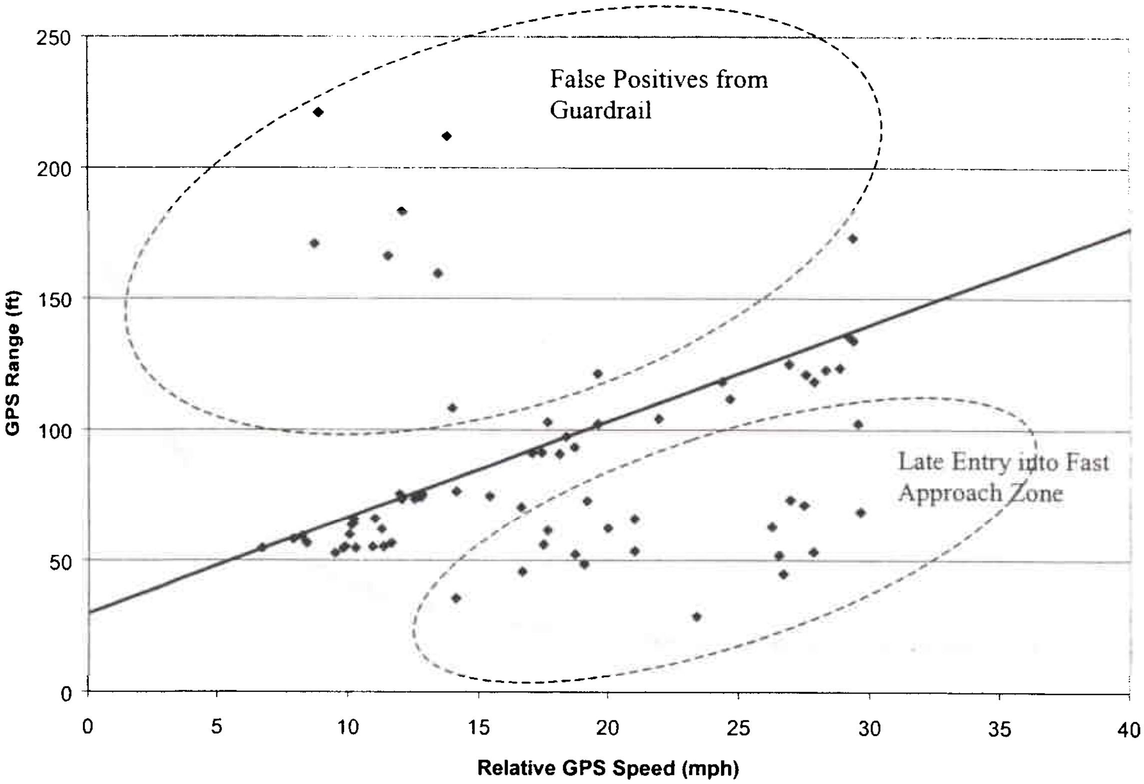

Figure A-4.1-2: Position and Relative Speed at Onset of Driver Warning on Curved Road

The figure for the straight sections of road show a fairly good clustering around the warning algorithm. Small deviations ($\leq$ 10 ft) are typically due to pointing errors. It is almost impossible to train the laser on the front bumper of the POV given the bouncing of the SV and changes in elevation of the road, which redirect the laser either high or low. A number of points are seriously in error. The one point above the line occurred near a curve, which the SV had already left but the POV had not. The laser was picking up on the guardrail which lines both sides of the track. The persistence that is built into the algorithm began tracking that point at a fairly high relative velocity which triggered the warning. Most of the points below the warning line are due to the fact that the laser acquired the target late, probably due to the low reflectivity of the Honda at low incident angles. The point at coordinates of 30 mph, 50 ft. was in fact, the reappearance of the target after it had already triggered the warning. During the testing a dip was noticed in the straight section that on occasion would result in the laser pointing over the POV.

The data for curves shows a much greater spread than the data for the straightaway. The increased variance results from two different configurations. Theses are illustrated in Figures A-4.1-3 and A-4.1-4. In Figure A-4.1-3, the POV is still in the curve while the SV is already coming out of the curve. The fast approach zone however is intersecting the guardrail that stands at the edge of the track. As was discussed in section A-2.6.2, long extended objects that intersect the fast approach zone at an angle can give rise to false alarms. These are the points that lie above the driver warning line in Figure A-4.1-2.

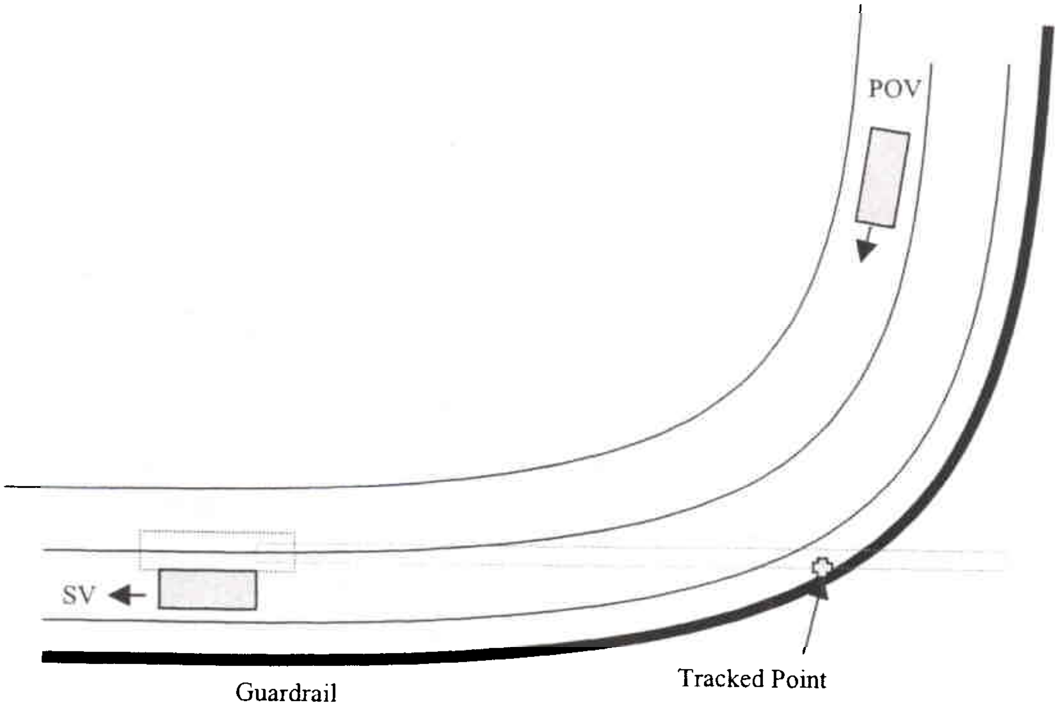

Figure A-4.1-3: Illustration of Improper Warning in Curves

The other situation is one in which the warnings come too late. These are the points in Figure A-4.1-2 that lie well below the driver warning line. This situation is illustrated in Figure A-4.1-4 wherein the POV is simply seen to cut into the fast approach zone at a range closer than that at which it would have triggered a warning, given the speed it was travelling. In principle it would be possible to eliminate these types of late warnings by deforming the fast approach zone to better conform to the curvature of the road. Since for lane change the sensor system is oriented behind the car, we are interested in road curvature that we have already passed. It should be possible to get this information from the time history of the steering or from advanced digital maps plus GPS. In practice this source of false positives from curves in the road was not noticed to be significant. It is therefore believed that the specific geometry of this test (i.e. radius of curvature of 0.6 km) was overstressing the system.

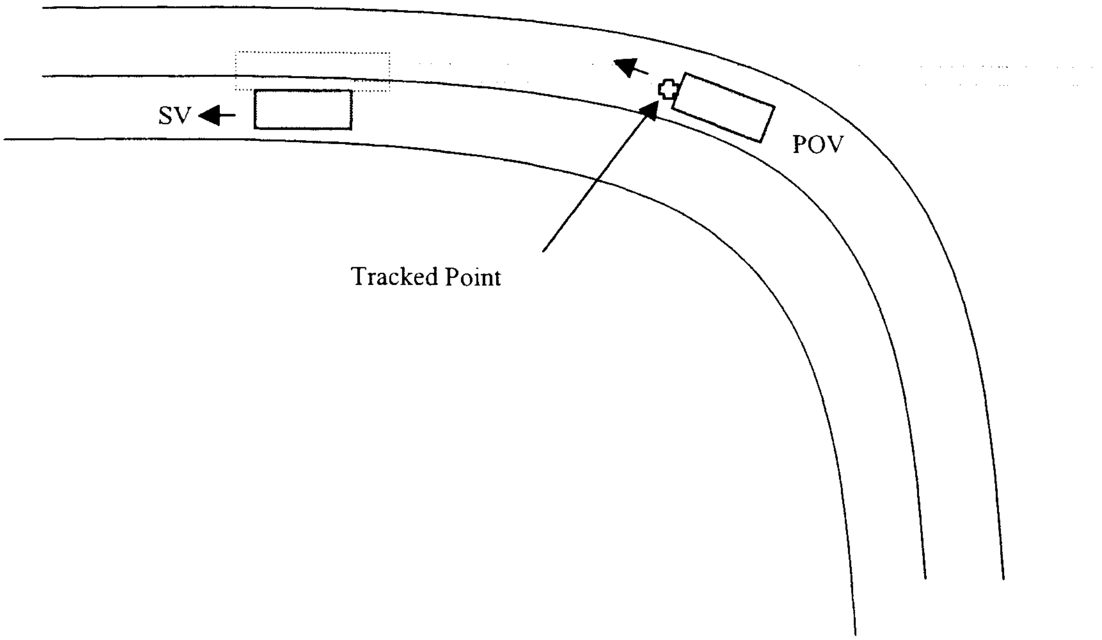

Figure A-4.1-4: Illustration of Late Warning in Curves

The next tests were done with the SV passing the POV. In this case, obviously the POV enters the detection zones from in front of the SV. The first zone it enters is the forward looking zone. This zone determines whether an object is moving or not. If it is moving, then the first opportunity for a warning is when the object enters the proximity zone. We see plotted in Figure A-4.1-5 the coordinates of the closest point of the POV as determined by the GPS. The GPS antenna in this case was placed on the left rear corner of the POV, and referenced to the bumper. The values for the x coordinate are negative because we are considering the region alongside the SV, as opposed to the region behind the SV that we are usually concerned with. The boundary of the proximity zone is shown with a dotted line. The test was carried out by having the POV proceed along the track at constant speeds of 30, 40, and 50 mph. The SV passed at relative speeds of up to 22 mph. The data points are grouped according to relative speed, as indicated in the legend. The spread in the data points in the Y direction is simply a function of the lateral distance between the two vehicles during the encounter. The spread and offset from the boundary edge in the X direction is greater than would be expected due to the latency of the system. Furthermore any spread due to latency would be correlated with relative velocity, of which we see none. The fact that the points are inside the boundary is an indication that the laser system is estimating a greater range that the GPS. This can only mean that the laser is hitting the POV above the rear bumper and along the trunk lid. This is one indication among many that the laser scan head was mounted too high.

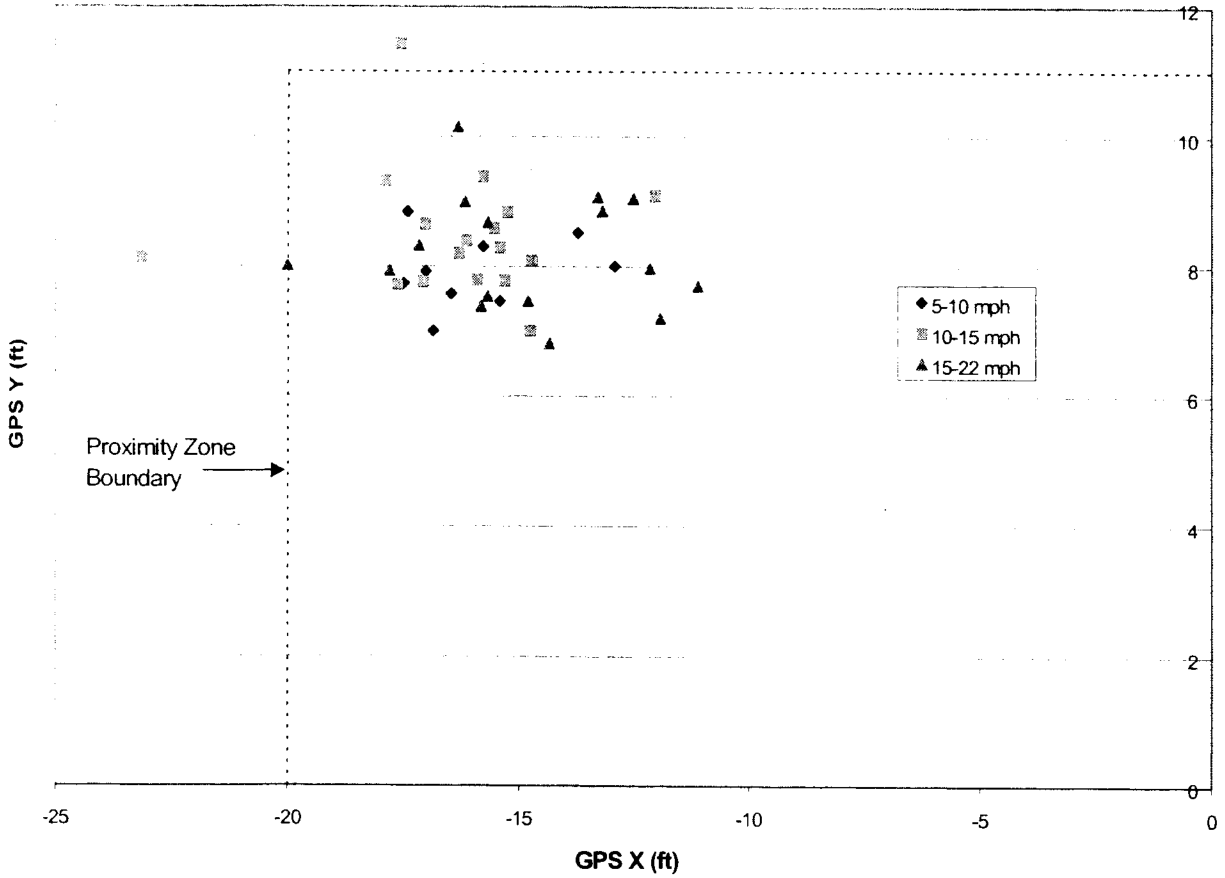

Figure A-4.1-5: X,Y coordinates of POV closest point as determined by the GPS for SV passing POV

The last test in this section was included to demonstrate that the system will recognize a vehicle travelling in the proximity zone at the same speed as the SV. Some systems, notably Doppler radar systems, are designed to detect only relative motion. The laser rangefinder is not one of these systems, and will generate warnings based upon presence alone in the proximity zone. This test was performed by having the SV travel at a constant speed, while bringing the POV alongside at the same speed. When the two vehicles have maintained a constant relative position for at least a few seconds the time is marked for later retrieval of the coordinates of the position of the POV. The test was performed at speeds of 30, 40, and 50 mph. For this test, the GPS was positioned at the left rear corner of the POV. Figure A-4.1-6 shows the results of this test, wherein the POV position was recorded at various locations in the proximity zone, demonstrating that zero relative velocity is not a problem.

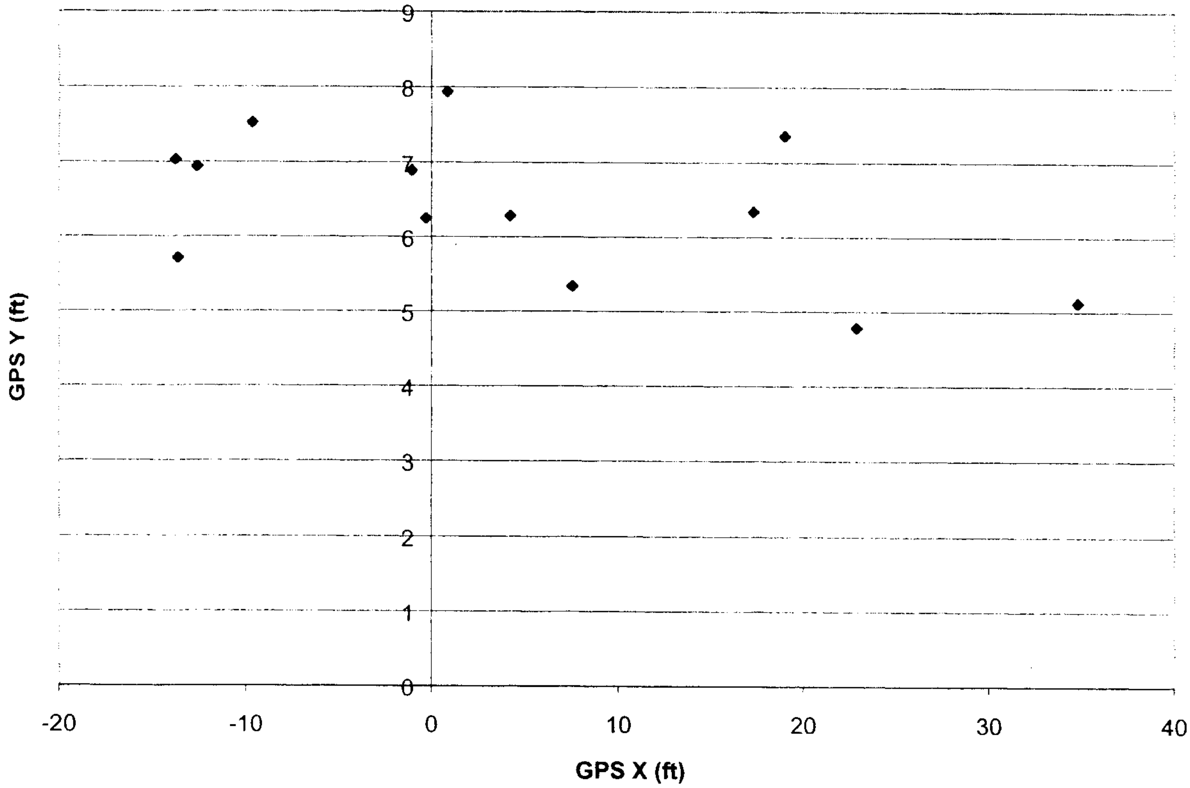

Figure A-4.1-6: X,Y coordinates of the left rear corner of the POV under conditions of zero relative velocity

A-4.2 Dynamic Tests with Interference

This test is an abbreviated form of those in the previous section, with the addition of a third vehicle to follow directly behind the SV. The purpose of this test is to see if this third car can trigger a warning where none is required. For this test, the interference vehicle is to follow behind the SV at distances varying from 20 to 80 feet, at a constant 30 mph. The POV attempts to pass both vehicles on the right. The results of these tests are shown in Figures A-4.2-1 and A-4.2-2. The first figure shows the effect of the interference vehicle on the fast approach warning. We are plotting distance versus relative velocity as measured by the GPS. The bulk of the points are clustered around the driver warning algorithm line, as to be expected. The few that are outside the expected variation of position are those that have been triggered by the third car. The laser has picked up the interference car at distances of between 40 and 80 feet as it intersects the fast approach

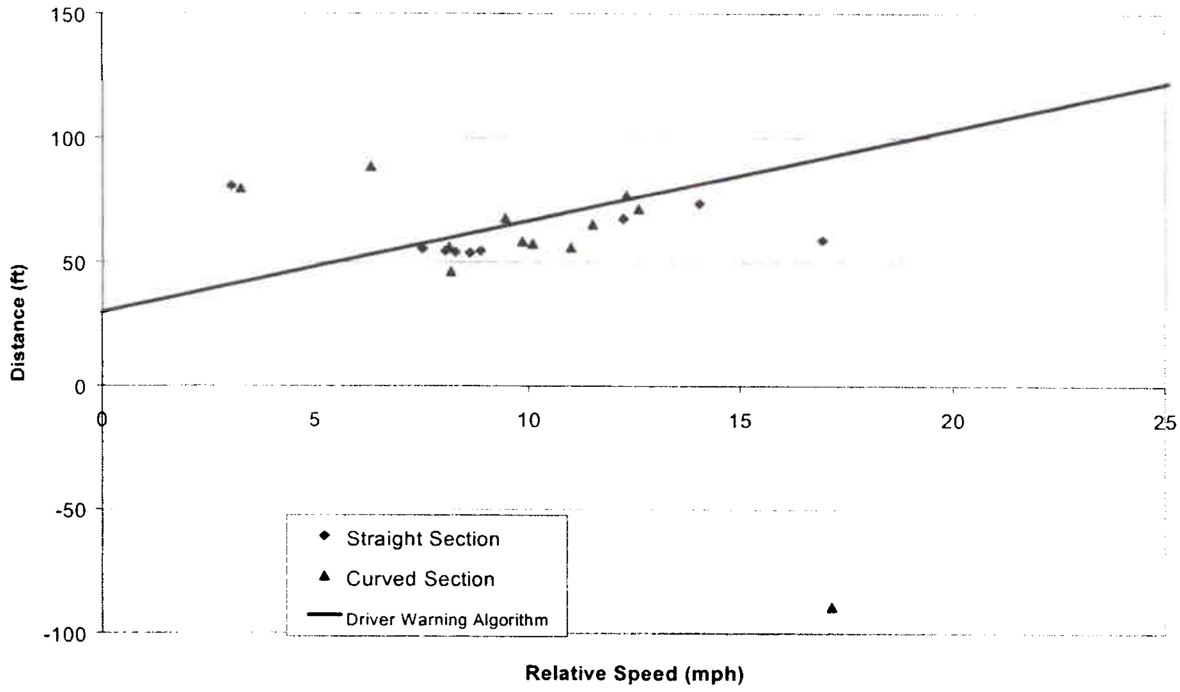

Figure A-4.2-1: Effect of Interference on Fast Approach Warnings

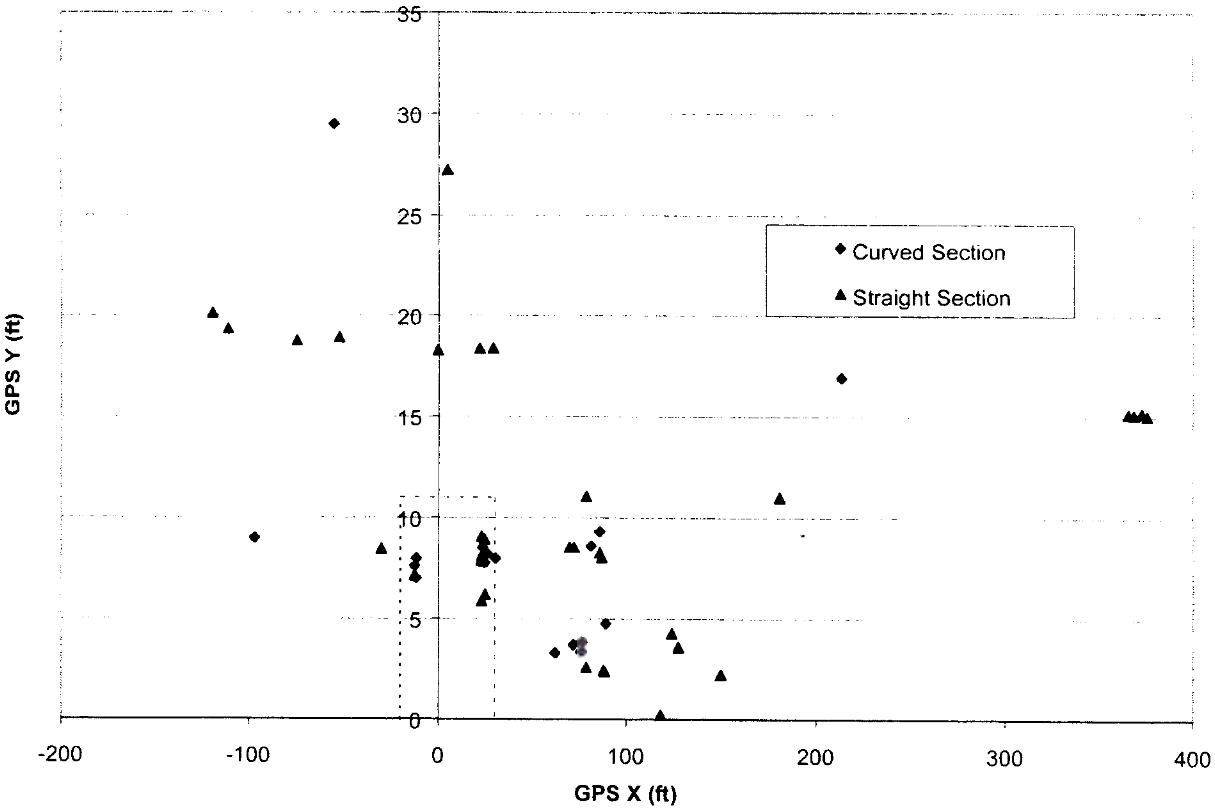

Figure A-4.2-2: Effect of Interference on Proximity Warnings zone, as illustrated in Figure A-4.2-3. Even though this zone is indented 4 feet in the lateral direction from the edge of the SV, it is still possible for a car following the SV to edge into this zone if the SV is closer to the left side of its lane and the interference car is closer to the right side of the same lane. It is possible that a quick excursion into the fast approach zone, could establish a track with enough longitudinal velocity to trigger a fast approach warning. Since this occurs at distances that are relatively close to the SV, they represent a fairly robust return for the scanning laser. Therefore the mitigation of this effect, and that of a fair number of false positives would be to eliminate projected tracking at distances below 80 feet, where the laser is less likely to lose track. This has been previously mentioned as part of the section on tracking (section A-2.6).

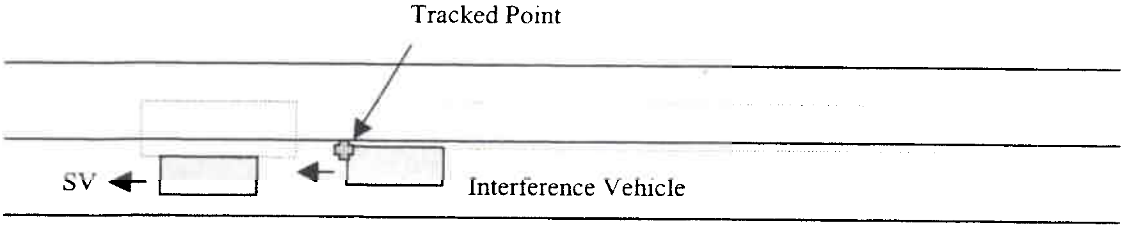

Figure A-4.2-3: Illustration of how an interference vehicle can set off a fast approach warning

Figure A-4.2-2 plots the X,Y position of the POV as determined by the GPS when the proximity warning is displayed. The dotted lines represent the boundary of the proximity zone. As can be seen it is relatively easy to provoke the proximity warning. Because the proximity warning only requires presence of any object, all that is required is that the interference vehicle contact the boundary of the proximity zone anywhere behind the SV. Since the proximity zone extends 30 feet behind the SV, and in a line with the right side with the SV, it is fairly easy for a following vehicle to trigger the proximity warning, as illustrated in Figure A-4.2-4. The proximity zone could have been indented away from the SV, as was the fast approach zone, however it was deemed unwise since motorcycles could slip by undetected until they were right alongside the SV.

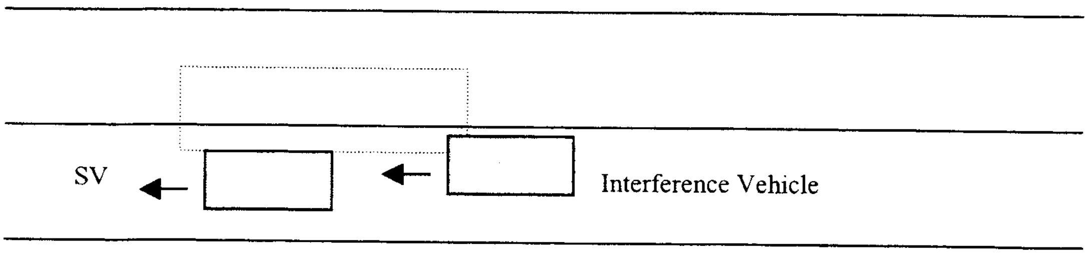

Figure A-4.2-4: Illustration of how an interference vehicle can set off a proximity warning

A-4.3 Three Lane Test

The purpose of this test was to determine the lateral extent of the proximity detection zone in a dynamic situation. The test configuration is shown in Figure A-4.3-1. During this test, both vehicles maintain a constant 40 mph as the POV changes lanes from the top lane to the middle lane. The POV changes lanes while varying the front bumper to front bumper distance S from nose to nose, to nose to midpoint, to nose to tail and finally at a point about 10 ft behind the SV. For this test the GPS antenna on the POV was located on the left front corner of the vehicle.

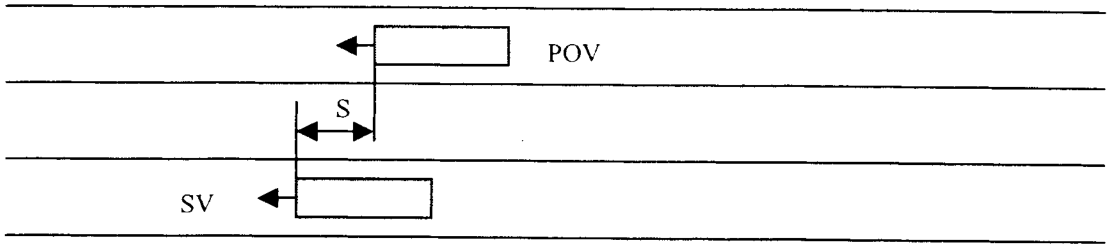

Figure A-4.3-1: Configuration for three lane test

Results from this test are shown in Figure A-4.3-2. This figure plots the X,Y position of the POV as determined by the GPS, when the proximity warning is activated. As usual the dotted lines represent the proximity zone boundary. The points at which the warning are activated are well within the proximity boundary of 11 feet. This is a clear indication that the laser is hitting points that are beyond the edge of the car closest to the SV. The target car in this instance was a 1989 Honda Accord. The height of the laser was such that it was able to enter the side window or graze the hood or trunk lid. This was the second indication that the laser was mounted too high. The first indication came during the tests of the SV passing the POV in section A-4.1.

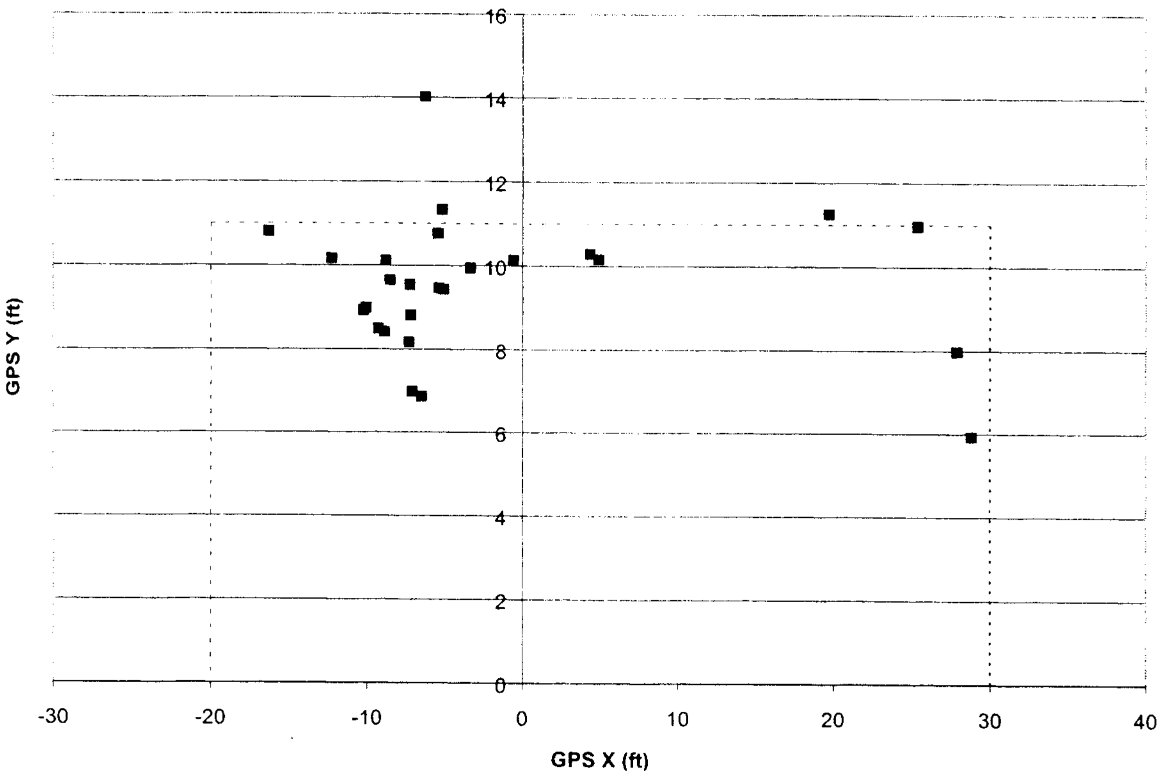

Figure A-4.3-2: Position of POV at proximity warning

A-4.4 Approach and Pass Test

The purpose of this test was to reproduce a common scenario on a multi-lane road, wherein a car will come up directly behind the SV and then suddenly change lanes. Our test configuration is shown in Figure A-4.4-1. The testbed maintained a constant 40 mph. The POV approached the SV in the same lane. The distance of closest approach was whatever the driver felt to be safe. At this point the POV changes lanes and passes the SV. Passes were taken on both the straight and curved sections of track.

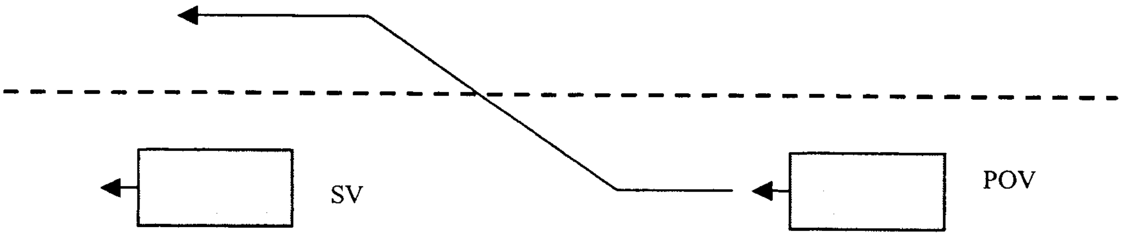

Figure A-4.4-1: Configuration for Approach and Pass Test

The results from that test are shown in Figure A-4.4-2. The most noticeable feature of these results are the fact that there is considerable spread in values away from the driver warning algorithm line. Once again those points considerably above the line are false positives, most likely generated by the laser hitting the guardrail as it has just come out of a curve. The fact that there is a spread of points below the line is to be expected since the POV was entering the fast approach zone from the side, and hence closer than what would ordinarily be warning against. Although the POV could have been tracked in the lane directly behind the SV, it was not done so for this relatively simple driver warning algorithm. This is one loophole that should perhaps be closed in an advanced lane change CAS.

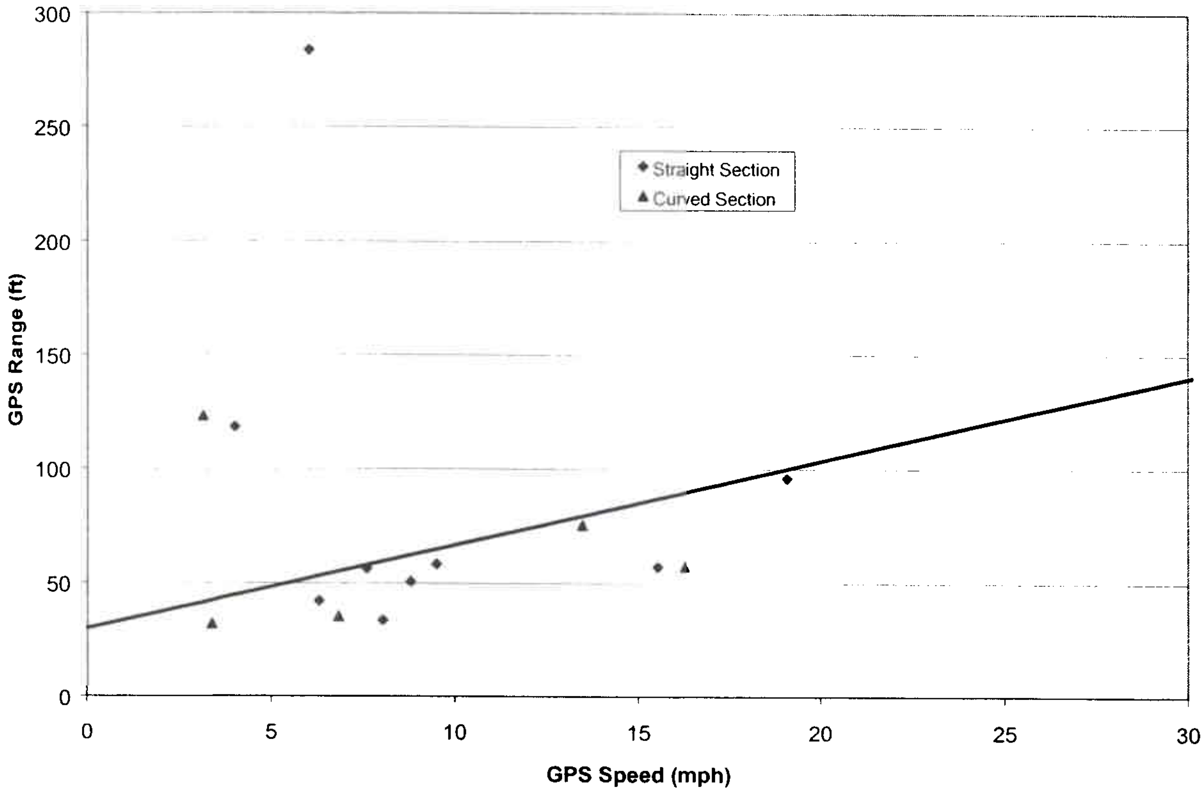

Figure A-4.4-2: Results from Approach and Pass Test

A-4.5 Modifications to SV

As was alluded to in Sections A-4.1 and A-4.3, the vertical position of the laser scan head was deemed too high for optimal performance. The laser was hitting the target cars too high and either generating returns from the interior of the car, skimming the hood and returning a weak signal, or missing the target car entirely. By lowering the laser all of these problems could be lessened. The position of the scanning mirror was lowered from 36.5 inches to 30 inches above the ground. Two of the track tests were repeated to see if there had indeed been any improvement. One of these is representative of a classic lane change situation where the POV passes the SV, while the other was one that showed increased sensitivity to the mounting height of the scan head (three lane test).

Figure A-4.5-1 shows the relation between the range and closing speed, as measured by both the laser and the GPS. This is to be compared to Figure A-4.1-1 for driver warning on a straight section of road. The clustering of points in Figure A-4.5-1 is much closer to the driver warning algorithm line than in Figure A-4.1-1. This is one measure of the quality of the improved system. Figure A-4.1-2 takes the same data and determines the position of the POV when the fast warning turned off. The tracking algorithm turns off at a point 27 feet behind the rear bumper of the SV. When a POV is in the process of passing the SV and advances closer than that point, the tracking algorithm assumes that the target has disappeared and projects the position forward for .8 seconds. The slope of the line connecting these points is a measure of the persistence time, i.e. the time that the warning stays on after the target has gone. This time is measured to be 0.89 seconds which is very close to the time of eight scans (0.8 sec) that the algorithm will project the last known point forward plus the one scan (0.1 sec) where the tracked point was last seen.

The second test performed was that of the three lane test. The x,y position of the front corner of the POV at the point at which the proximity warning turned on is displayed in Figure A-4.5-3. This is to be compared to Figure A-4.3-2. Notice at once the improved tightness of the coordinates, corresponding to a better definition of the edge of the car. The data shows some penetration of the proximity zone by at most 2 feet. The proximity boundary is at 11 feet in the Y direction.

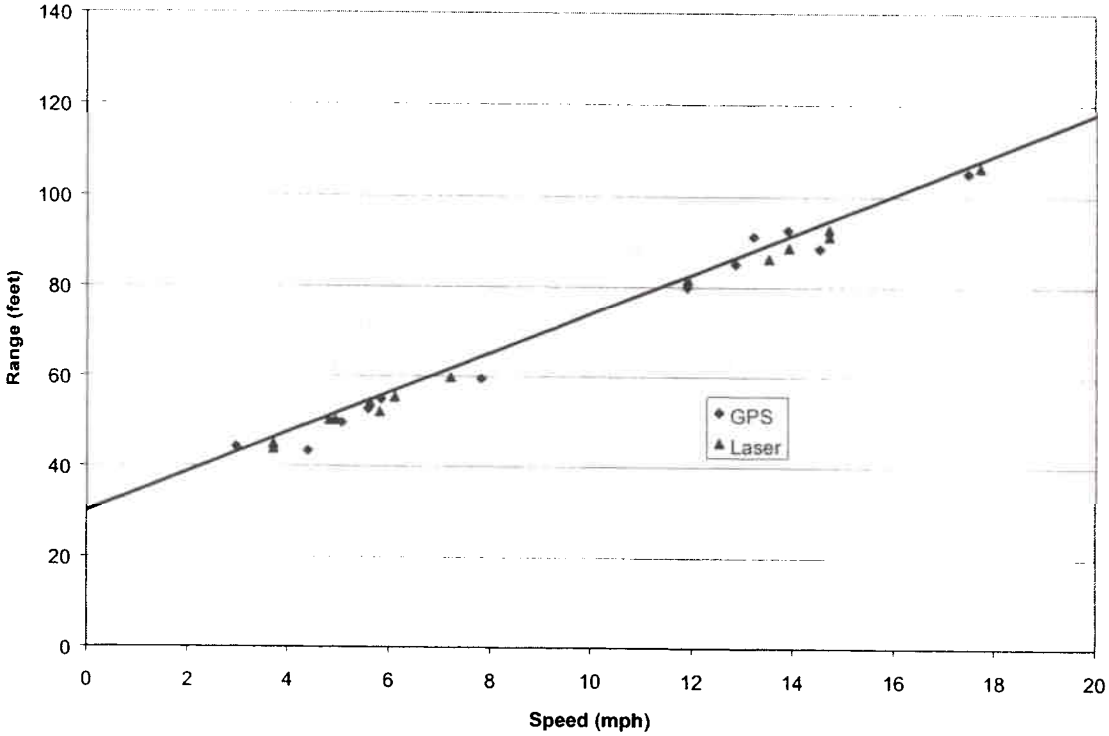

Figure A-4.5-1: Verification of Driver Warning Algorithm after laser modifications

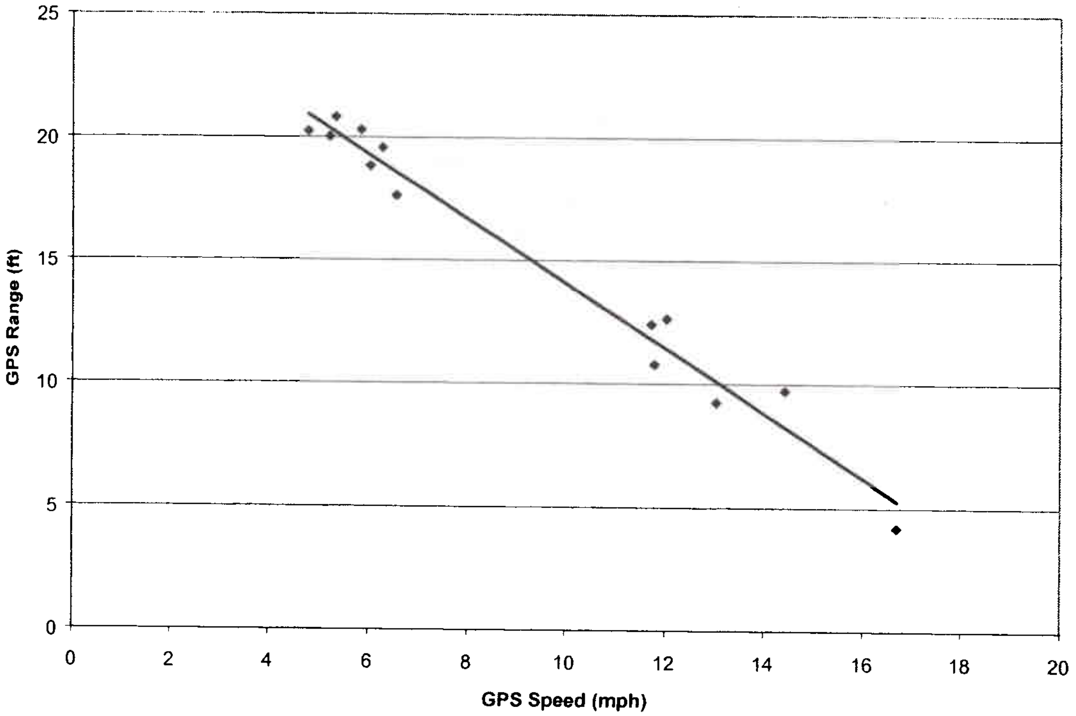

Figure A-4.5-2: Position and Relative Speed of POV at time of fast warning turn off

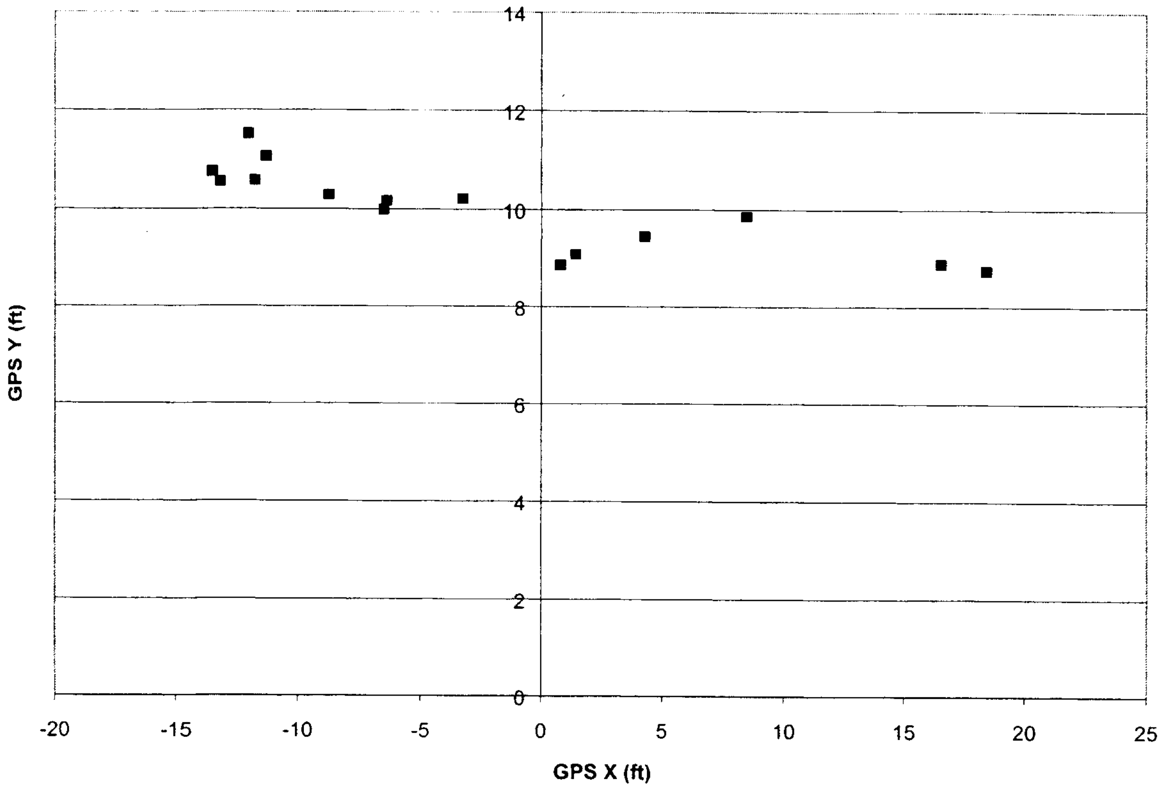

Figure A-4.5-3: POV coordinates at proximity warning turn on after laser modifications

The conclusion reached is that lowering the laser has improved the ability of the laser subsystem to accurately track low reflectivity ("stealthy") targets.

A-4.6 Test Track Preferences

These tests were a systematic attempt to refine the CAS performance parameters by measuring a number of individual preferences as to when they would prefer a warning. The testbed was driven at a steady speed of 20mph, 30mph 40mph, 50 mph and 60 mph. For each of these SV speeds, the POV will attempt to pass the testbed at a relative speed of 5mph, 10mph, 20mph, 30mph and 40 mph on a straight section of track. The actual maximum speeds attained must be consistent with that which the track can safely accommodate. At each run the subject driving the testbed will observe the other vehicle in the side view and center mirrors and indicate the last moment at which he would change lanes by momentarily pressing a switch, which marks that time in the data stream. From this mark one could determine the position and relative velocity of the POV. The instruction "last moment at which you would change lanes" is equivalent to asking that person when they would prefer a warning. The team members rotated the driving tasks so that more than one preference was recorded.

These tests took place in three different venues. The first was on the access road of the California Speedway in Fontana, CA (Track1). The second was a test track operated by the Sheriff's department of San Bernadino County in Devore, CA (Track2). The third venue was on the freeways neighboring TRW in Redondo Beach, CA. In the first two venues, the test procedure was as described in the paragraph above. On the freeway targets of opportunity were used. Four team members were tested, however not all members experienced all three sites.

The raw data results for each team member are plotted in Figures A-4.6-1 to A-4.6-4. For these curves, range (as measured by the laser rangefinder) is plotted as a function of the relative velocity. Each set of points represents a different day, and hence a different venue. The straight line through each set of points is a linear best fit, with the equation of that line shown, along with the coefficient of simple determination, R^2. The R^2 values, with one exception, are all relatively high indicating a strong correlation between range and relative velocity. The implication of this is quite profound, meaning that the desired warning range is independent of the absolute velocity of the subject vehicle. The second point worth noting is that there can be considerable variation on different days (or venues) for a given individual, not to mention the variation between individuals. The form of the equation for each data set has a simple interpretation. The slope of the line represents the warning time in seconds before the POV will enter the proximity zone, while the intercept represents the extent behind the rear bumper of the SV that the proximity zone should extend.

Both the slopes and intercepts have been collected and are shown in Figures A-4.6-5 and A-4.6-6. There is a fair amount of variation in the slopes – from 2.0 to 4.3 seconds. Because we were looking for a one size fits all driver warning algorithm we chose a value of 3 seconds, even though we originally planned to use the most conservative value indicated. In light of the results it was felt that for most drivers a time of 4 seconds would be excessive. The values for the extent of the proximity zone were in fact more consistent. If

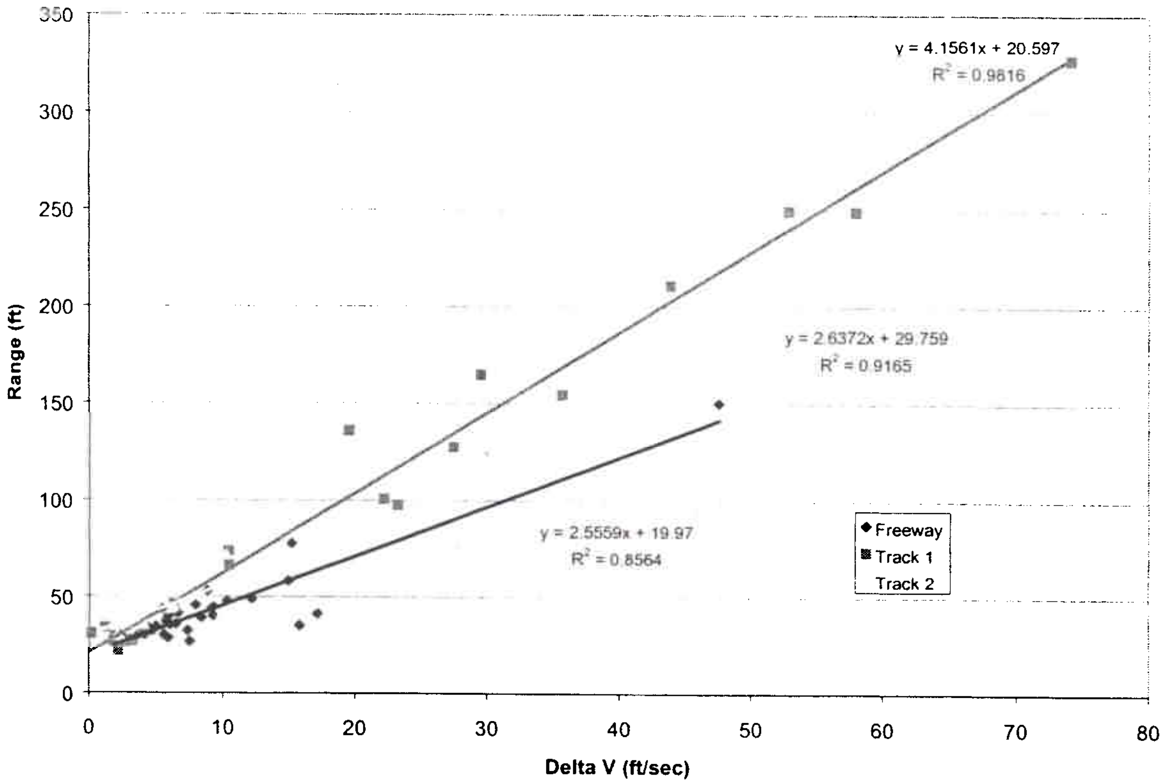

Figure A-4.6-1: Preference data for M25

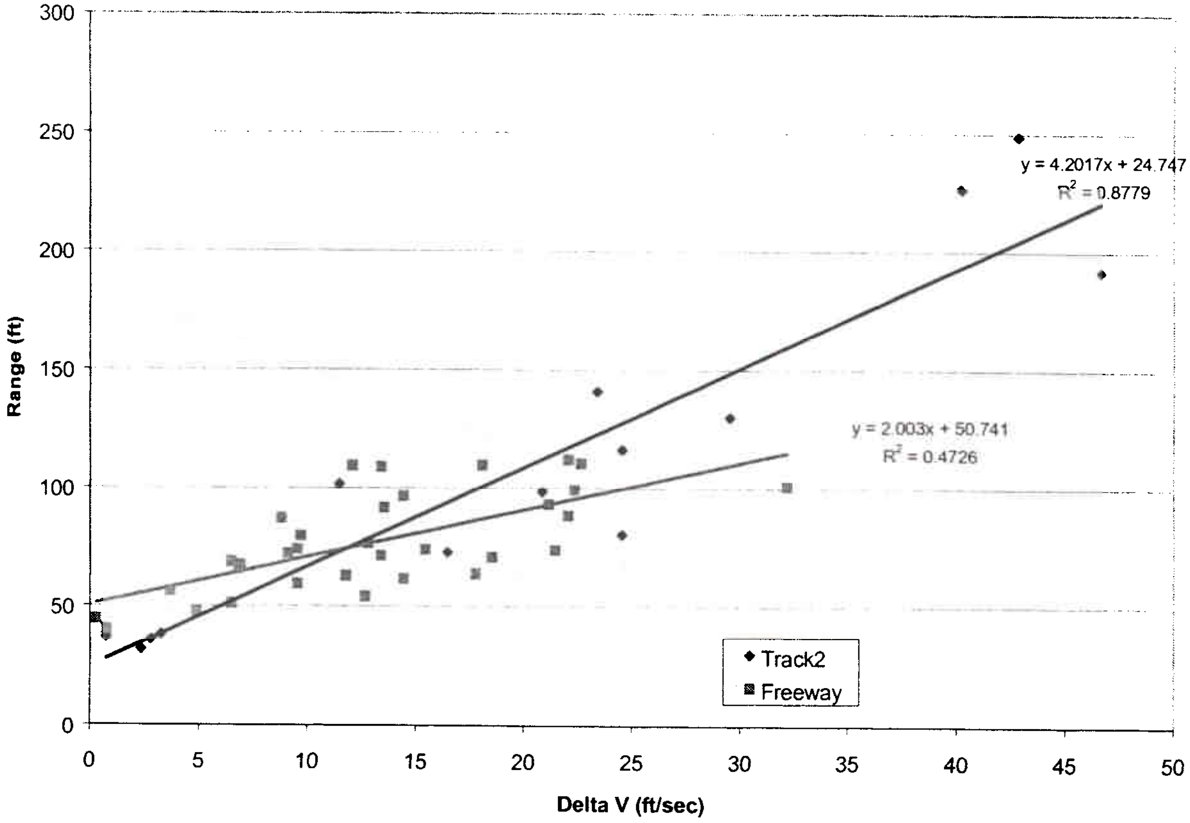

Figure A-4.6-2: Preference data for M37

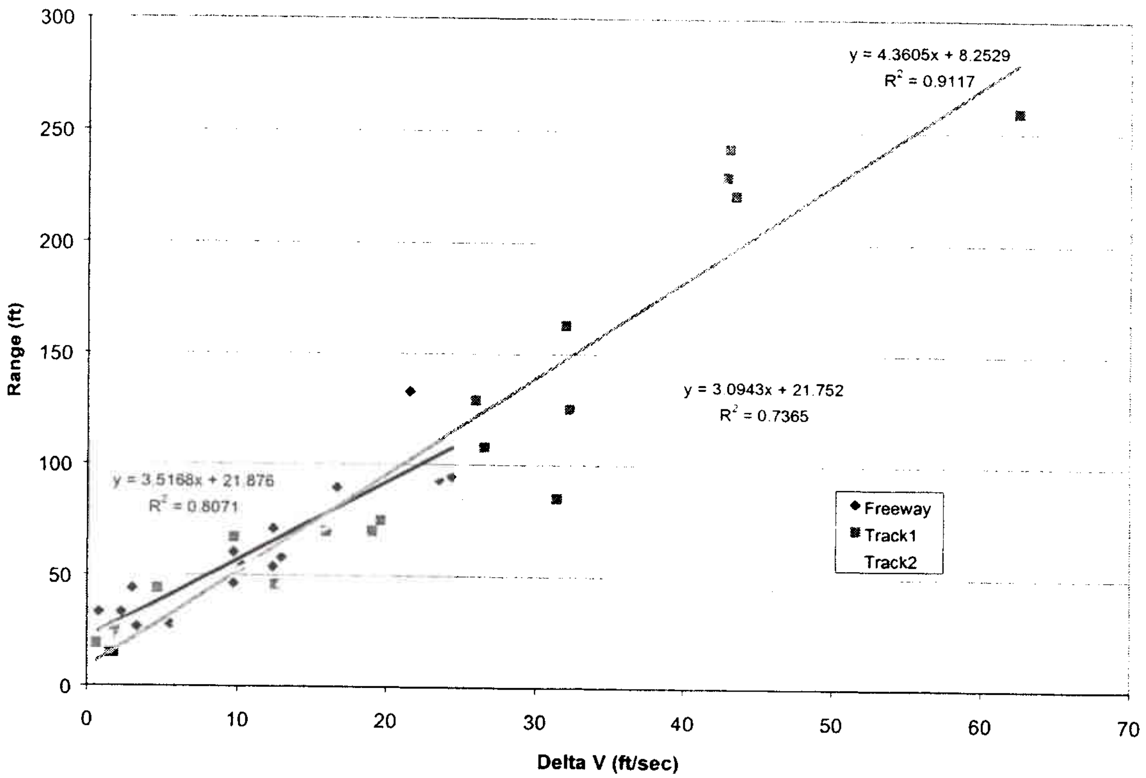

Figure A-4.6-3: Preference data for M48

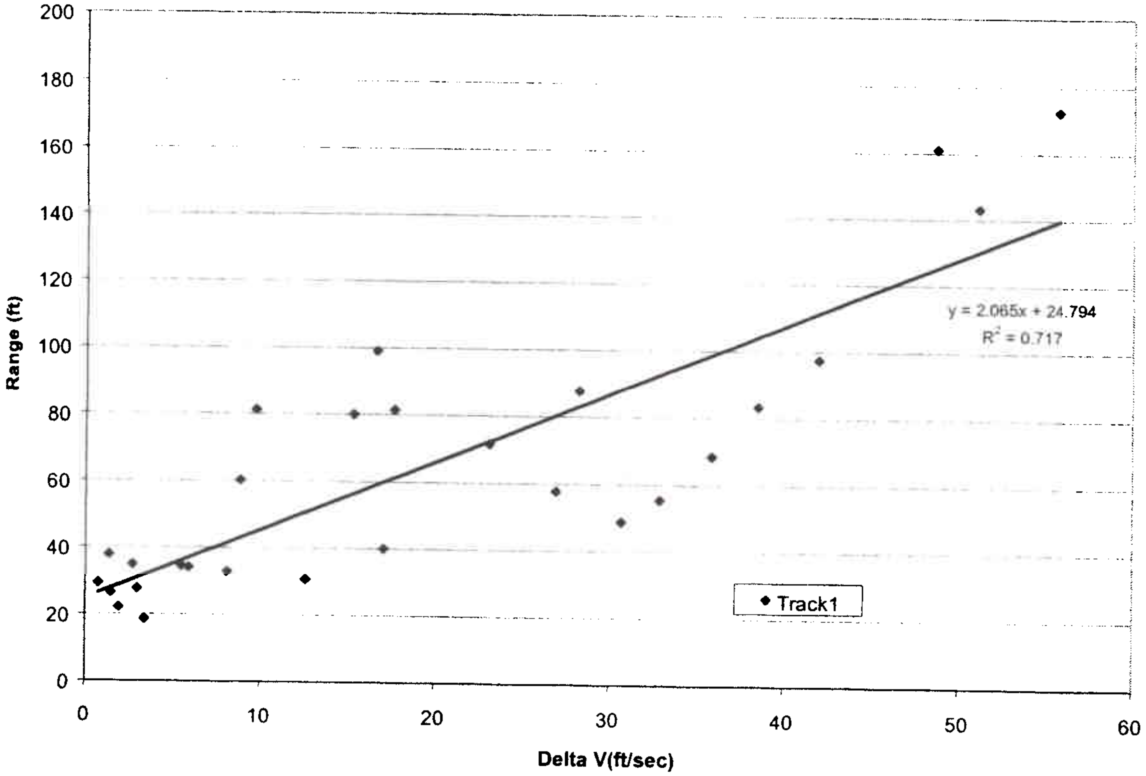

Figure A-4.6-4: Preference data for M55

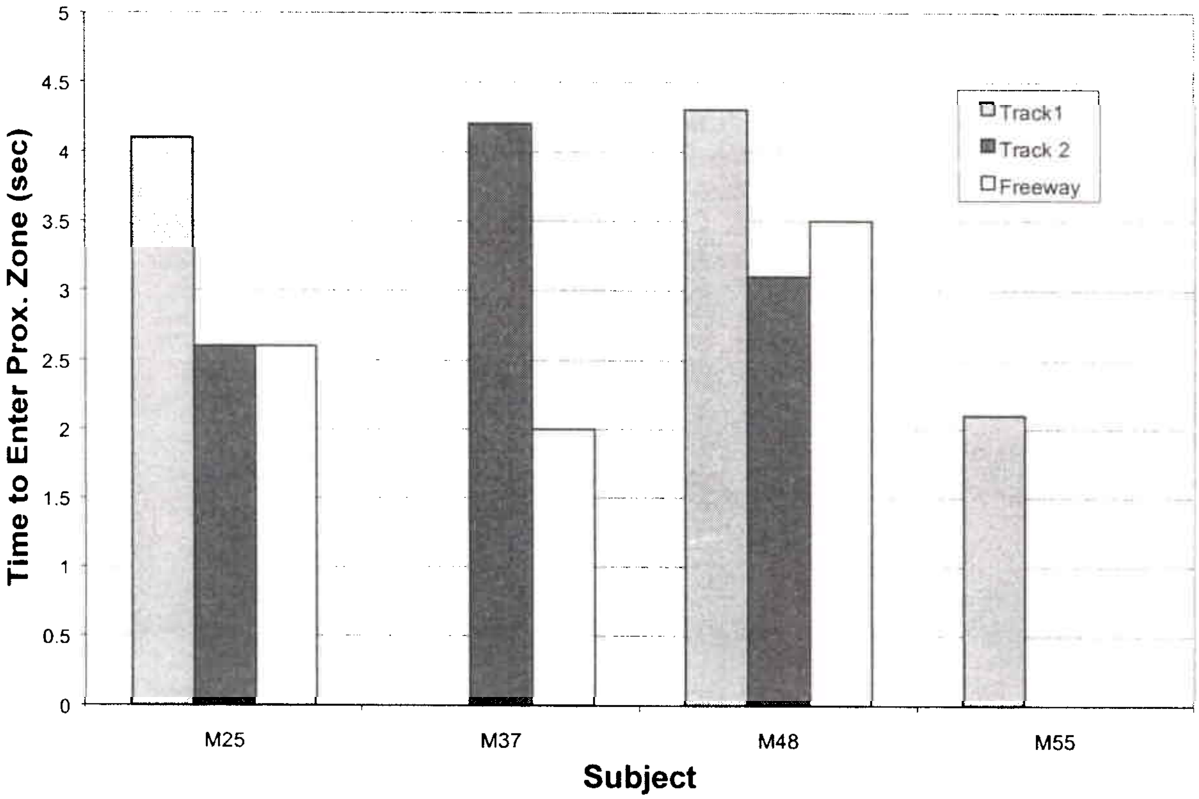

Figure A-4.6-5: Advance warning time for POV to enter proximity zone

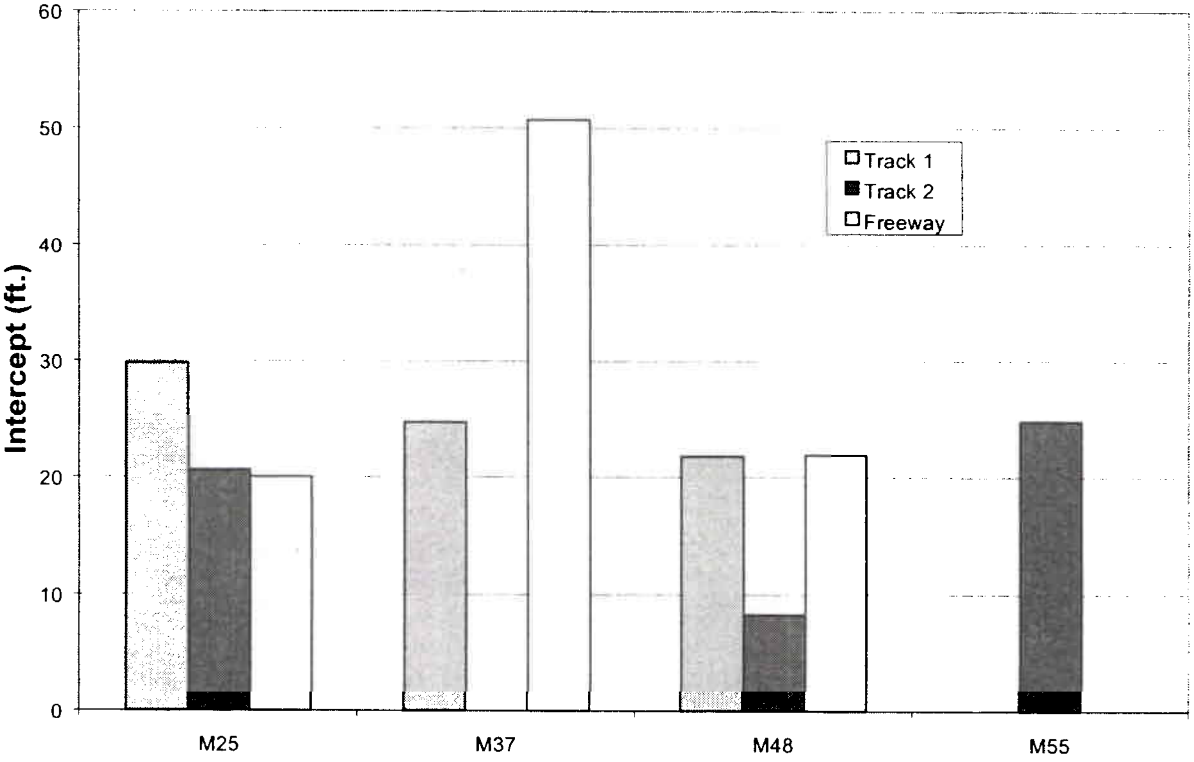

Figure A-4.6-6: Desired rearward extent of proximity zone

one eliminates the shortest and farthest distances, the rest fall consistently between 20 and 30 feet. Wishing to err on the side of conservatism we chose 30 feet. Therefore the final form of the driver warning algorithm that was used in the human use testing is:

$$R_{warn} \leq 30\text{ ft} + 3\text{ sec} * V_{relative}$$

A-5.0 Open Road Testing

The open road tests are designed to learn as much as possible from a restricted sample set of possible drivers and vehicle configurations. This testing required compliance with human use protocols as outlined in NHTSA Order 700-1.

A-5.1 Experimental Design

This phase of the testing focused on the driver interaction with the CAS. In addition, kinematic information was extracted about the act of making a lane change. Given the relatively small number of test participants (12) it seems appropriate that the number of independent variables be limited to three, with one of the variables having three possible states. In this fashion it is possible to design a balanced randomized blocks partially confounded factorial experiment that reveals the dominant dependencies. The three independent factors are listed below, with +, 0, and - signs indicating their usage.

L (Location of Display) = Side Mirror (-), Center Mirror (0) or Center plus Side Mirrors (+)
M (Mode) = Turn Signal Activated Mode (-) or Monitor Mode (+)
S (System Type) = Proximity (-) or Comprehensive (+)

The display has been discussed before. The proximity type CAS only monitors the proximity zone, whereas the comprehensive type monitors both the proximity zone and the fast approach zone. By monitor mode it is meant that a steady burning light is presented if there is a warning situation without turn signal use. However, turn signal usage in this mode while in a warning situation results in a flashing warning light (i.e., this is the "augmented alert mode"). For turn signal activated mode, there will be no warning unless the turn signal is active. In that case the warning will flash. For all cases, the displays to be energized were the Muth mirrors, and not the added-on LED's.

The revised data collection strategy is known as a balanced randomized blocks partially confounded factorial (RBPF-32^2) design [3], also sometimes referred to as a incomplete blocks design. This strategy reduces the number of combinations of design factors that a test participant experiences during the test to six (6), thus making test execution more manageable. However, this economy of means is bought by partially confounding the three-factor interaction and the Mode x System interaction with blocks. This is thought to be a worthwhile tradeoff because analysis procedures exist that adjust the three-factor interaction and the Mode x System interaction for block effects to provide statistical tests of the influence of those interactions on the variability of the response measures (e.g., turn-signal use, mirror glance frequency, minimum gap accepted for lane change, etc.). A second assumption of the proposed approach is that interactions involving blocks are assumed to equal zero.

The data collection plan is provided in Figure A-5.1-1. This table indicates the need for three groups, each with at least two test participants assigned to it (multiples of two test participants may be assigned to each group as well). This implies a minimum of 6 test participants. Statistical power considerations indicate a minimum of 12 test participants is needed.

In principle the blocks are homogeneous test participants, randomly assigned to groups. For example, T1 and T7 might be younger, while T2 and T8 are older subjects. In practice the assignment of test subjects to blocks was not balanced or randomized, thus confounding the three-factor (LxMxS) and the MxS interaction. Since this is a preliminary study and we are looking for gross effects, not subtleties, this was not deemed a problem. Therefore, in interpreting the results of the statistical analyses below, it is important to keep in mind that high F values for the above two interactions may be misleading.

Reading along a row indicates what six combinations of factors each test participant will experience. Sequential ordering of factor combinations must be counterbalanced or randomized across the ensemble of test participants. They should not be presented in the order given in the table. It is assumed that one baseline run (i.e. no lane change CAS presented) will be collected. Baseline runs are not represented in this table. Baseline runs are used in focused pair comparisons to a standard (i.e. the baseline). Baseline runs were in fact inserted in the middle of a subject's test runs so that the subject would be reasonably acclimated to the handling of the testbed vehicle.

Groups	Blocks	Combinations of L_j M_k S_l Level Codes to be Tested		
Group1	T1, T7	(---) (-++)	(0-+) (0+-)	(+- +) (++-)
	T2, T8	(--+) (-+-)	(0 - -) (0++)	(+--) (+++)
Group 2	T3, T9	(--+) (-+-)	(0- -) (0++)	(+- +) (++-)
	T4, T10	(---) (-++)	(0-+) (0+-)	(+--) (+++)
Group 3	T5, T11	(--+) (-+-)	(0-+) (0+-)	(+--) (+++)
	T6, T12	(---) (-++)	(0 - -) (0++)	(+-+) (++-)

Figure A-5.1-1: Design Matrix

The response measures may be analyzed by means of analysis of variance (ANOVA). Assuming a total of N=12 participants, the ANOVA table provided in Figure A-5.1-2 indicates the effects, degrees of freedom, and appropriate error terms to be included in the analysis. A spreadsheet model, based on the ANOVA in Kirk was set up and used for all the statistical analyses in this report. This model calculates the F-ratio that is referred to in Figure A-5.1-2. The calculated value must be compared to the critical value of F (F_{crit}) which depends on the degrees of freedom of the variable in question (2 for location, and 1 for both mode and system), the degrees of freedom in the error term (49) and the probability level α. This value (usually taken to be 0.05) is the probability of erroneously declaring that there is a correlation when in fact there is not. The F_{crit} values can be looked up in a table. They are $F_{crit} = 3.23$ for location and $F_{crit} = 4.08$ for mode and system. This means that the calculated values for the dependent variables as a function of the independent variables must exceed these numbers. Now, in principle, because this work can be considered exploratory research we might be able to raise the α level to 0.1, and thus lower the threshold for declaring a dependency. However we have chosen for this work the more stringent standard of 0.05.

Line	Source of Variations	Degrees of Freedom	F-ratio (Fraction numerator and denominator denote line numbers for appropriate terms)
1	Location (L)	3-1 = 2	1/10
2	Mode (M)	2-1 = 1	2/10
3	System (S)	2-1 = 1	3/10
4	LxM	(3-1)(2-1) = 2	4/10
5	LxS	(3-1)(2-1) = 2	5/10
6	MxS (adj)	(2-1)(2-1) = 1	6/10
7	LxMxS (adj)	(3-1)(2-1)(2-1) = 2	7/10
8	Groups (G)	(3-1) = 2	not tested
9	Persons within Groups	3(4-1) = 9	not tested
10	Residual (adj)	49	error term
11	Total	(4)(6)(3) –1 = 71	

Figure A-5.1-2: ANOVA table for RBPF-32^2 design

The dependent variables for the evaluation cover several categories of measurement. Examples of these are listed below.

Visual Allocation Measures to be taken
- Mirror glance time and glance frequency per mirror and sequence of glances, including over-the-shoulder
- Eyes-on-Road (ahead) Time (EORT)
- Head turn incidence for over-the-shoulder and right outside mirror

In-vehicle driver Behavior Measures to be taken
- Turn signal use
- Turn signal onset time with respect to lane change start

Driver-Vehicle Measures to be taken
- Lane change completion time
- Range and Range rate of POV at start of lane change
- Adjacent lane vehicle leading range
- Accelerations
- SV speed

A-5.2 Testing Protocols

Test participants were recruited from the general TRW population of employees. A representative sample of participants, with respect to gender and age, with 40 being the dividing line for age, was recruited for this test. Job descriptions of the subjects ranged among secretary, lab technician, business person and scientist/engineer. All subjects were required to possess a valid unrestricted driver's license, have a minimum of two years driving experience and could not be under the influence of alcohol, drugs or any other substances which impair their ability to drive.

All potential recruits were given an information packet (Appendix C) describing the purpose of the tests and the testbed vehicle with its installed instrumentation. Particular attention was paid to the eye-tracker headgear, since that could have been a deciding factor for some people as to whether or not they would participate. Due to the difficulty of using the eyetracker with eyeglasses, only one driver who wore glasses was recruited. All tests were driven by a naïve subject with a CAS Study Team observer in the back seat.

The first test drive was strictly for familiarization only. Insofar as the subject is concerned, it did not differ significantly from any of the other test drives. The subject became familiar with the handling of the testbed and the operation of the displays. The test drive lasted about three-quarters of an hour and encompassed primarily the freeway portion of the designated route, which is described below. As part of the familiarization drive, the subject was fitted with the eye-tracker headgear. The calibration of this device must take place outdoors, so that the subject's pupil size will be stable during the test. This calibration was performed at the beginning of every run. The calibration procedure was described in section A-2.2.

The route for all other test drives, as shown in Figure A-5.2-1, was from TRW to the 405 freeway south, to the 710 freeway north, to the 105 freeway west to Sepulveda Blvd. Then south on Sepulveda Blvd. (PCH) to Hawthorne Blvd going north, turning west on Manhattan Beach Blvd. to the TRW facility. The route contains a mix of freeway and arterial driving. The route typically features moderate density traffic if driven between the hours of 10AM and 3:30PM. Occasionally there were pockets of high density traffic. Only one run was completed on a Friday, given the generally earlier start of the afternoon rush hour on that day. Typically the drive lasted from 1.25 to 1.75 hours. All test runs were completed in fair weather, in order to have a common basis of comparison.

In order to insure that enough lane changes would be executed during a run, the subject was asked to move to the left lane in the freeway when it was clear to do so. At this point the subject was free to change lanes to drive in the lane most comfortable to him. Approximately 3 miles from the freeway exit, as determined from the upcoming exit signs, the driver was warned that he must exit to the right. This procedure was repeated on each of the three freeways of the selected route. The chosen surface streets are 2 and 3 lane roads. On Sepulveda Blvd., the subject was allowed to drive in either lane. After crossing Palos Verde Blvd. the subject was asked to move to the left lane so as to turn left on Hawthorne Blvd. in approximately 2 miles. On Hawthorne Blvd. the subject was asked to move to the right lane and then drive in any lane he felt comfortable. After crossing Artesia Blvd. the subject was asked to move into the left lane in preparation for a left turn onto Manhattan Beach Blvd. At this point the run is nearly over and the subject was instructed to return to TRW.

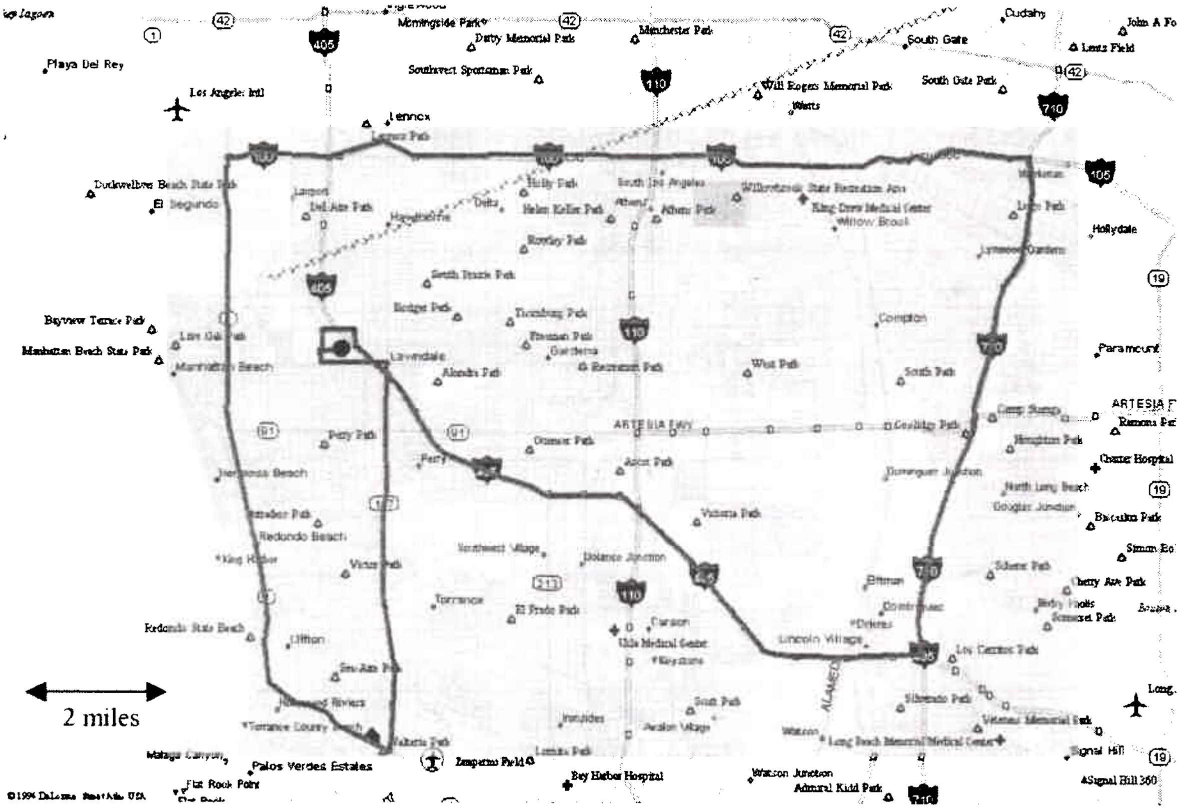

Figure A-5.2-1: Route for Open Road Testing

An assessment of the drivers' subjective comments was made by having the drivers complete a questionnaire, which can be found in Appendix A. The questionnaire has three parts. The first part is completed before the start of the testing. It asks for information about the subject and his driving behavior. The second part is completed after each test run. It assesses the subject experience with that particular day's test configuration. Finally the third part is completed at the end of the testing and asks for an overall assessment.

A-5.3 Lane Change Dynamics

In examining the trajectories of many lane changes it became apparent that they could all be classified into one of five categories. These are illustrated schematically in Figure A-5.3-1. The first is a lane change with no conflicts. In this case the POV is far enough behind the SV and not closing fast or at all. In many instances there may in fact be no POV in sight. The second category is observably the same as category one except for the fact that the SV has recently allowed another vehicle to pass him. The third category is when the SV first passes the POV and then changes lanes in front of him. The fourth category is when the SV cuts in front of the POV and forces him to decelerate. Finally the fifth category is when the SV and POV are going at roughly the same speed. This may occur when the SV is trying to get into a crowded lane and is waiting for enough room to open up so he can proceed.

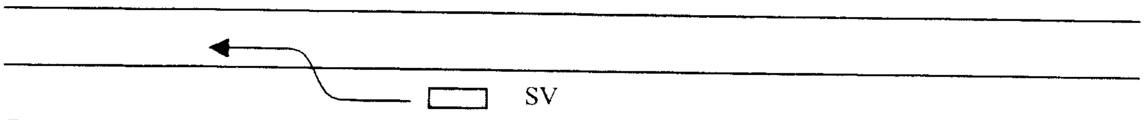

Cat 1: Wide Open

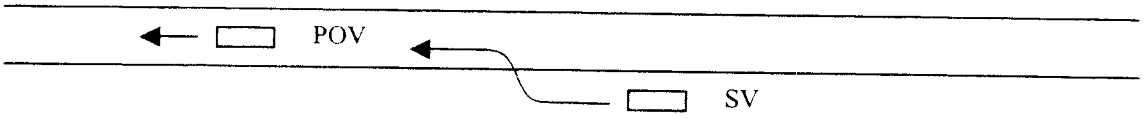

Cat 2: POV passed SV

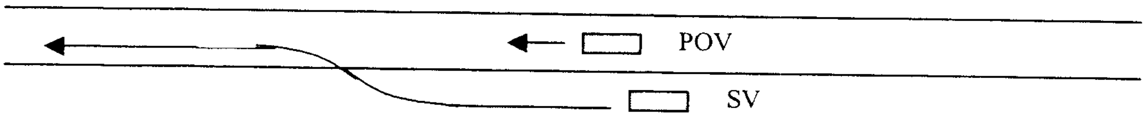

Cat 3: SV passes

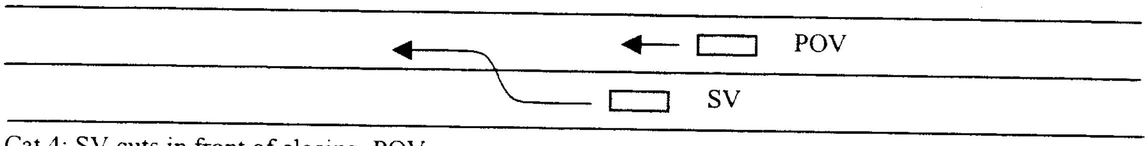

Cat 4: SV cuts in front of closing POV

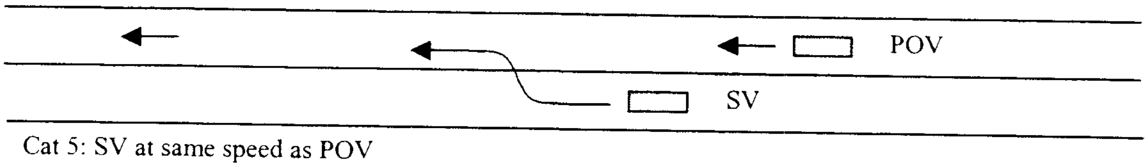

Cat 5: SV at same speed as POV

Figure A-5.3-1: Categories of Lane Change

Figures A-5.3-2 to A-5.3-7 offer examples of trajectories of the last three categories. We are plotting range to the POV, behind the SV versus closing velocity in mph. Positive closing velocity indicates a decreasing gap, while negative indicates an increasing gap. The trajectory itself is broken into four segments, delineated by color. In chronological order the black segment indicates the time before the lane crossing up to a maximum of ten seconds. The red to black change is the point that is one second before the lane crossing. By lane crossing we mean the time at which the wheels of the SV cross the lane markings as determined by the test conductor. The green segment, indicated as "after LC" in the legend indicates the time after the second pair of wheels of the SV cross the lane markings. Finally the blue segment, indicated as " behind SV" in the legend, indicates the trajectory of the POV during the time when it is in the same lane and directly behind the SV. The length of this segment lasts until 5 seconds after the lane crossing. Also shown in the trajectory charts are two sets of lines; one delineating the driver warning algorithm while the other parallel to it but translated further away. The reason for this will be made apparent later in this section.

Figure A-5.3-2 shows a category 3 trajectory at high SV speed, while Figure A-5.3-3 shows the same category at moderate SV speed. In principle the trajectory should start out in negative range, but our tracking software does not easily track in that region. Figure A-5.3-3 shows the SV accelerating away from the POV in the early part of the trajectory which implies that the SV was not going very fast to begin with.

Figure A-5.3-4 and 5.3-5 illustrate two sample trajectories for category 4 where the SV pulls in front of a POV and forces it to decelerate. Figure A-5.3-4 shows a case where the POV came unacceptably close to the SV who made the lane change while the POV was in the warning area. For this case, part of the apparent decrease in the closing velocity is coming from some modest acceleration on the part of the SV (on the order of 1mph/sec). The trajectory in Figure A-5.3-5 represents deceleration by the POV, as the SV is essentially maintaining speed.

Finally Figures A-5.3-6 and A-5.3-7 show trajectories for lane changes with low relative speeds (category 5). These are fairly typical in that they indicate some meandering in phase space (range vs. relative velocity).

A convenient way to characterize lane changes is to plot them in phase space, which has been divided up into zones. We have found it useful to divide the phase space into eight zones as shown in Figure A-5.3-8. The driver warning algorithm that was used for all the drivers is plotted as the lower curve. There is a parallel curve that is displaced 50 feet further away from the SV. The reasoning behind this curve is that there exists a region beyond which the driver of the SV does not consider a POV as in conflict. This is somewhat of a subjective determination, however we believe it to be a reasonable one. Vehicles beyond this zone require only a cursory look before determining whether they are threats. Zones 1 and 2 are regions in the proximity zone behind the SV and in the adjacent lane. The difference between them is solely a function of whether the POV is closing or opening the gap. Zone 3 is the region that will trigger a warning only when the comprehensive system is active. Zones 4 and 5 are intermediary zones with high relative

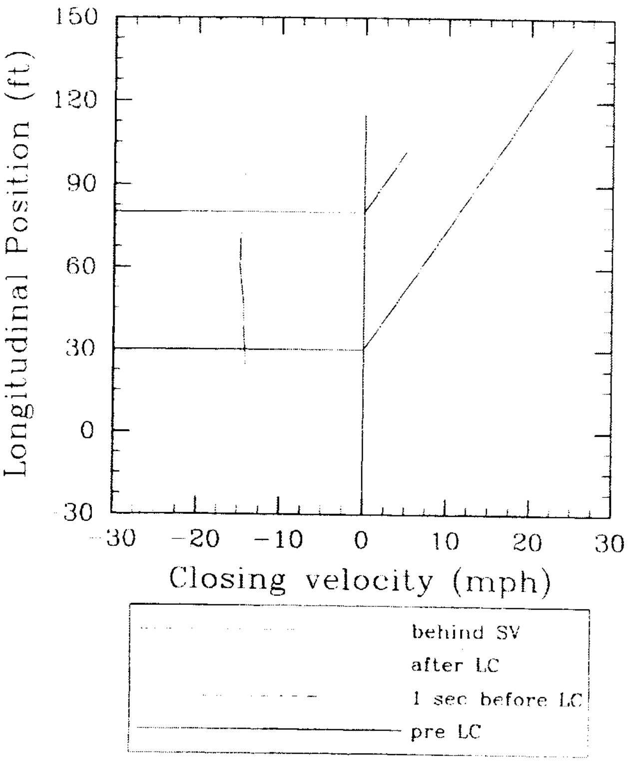

Figure A-5.3 –2: Category 3 passing trajectory at high SV speed

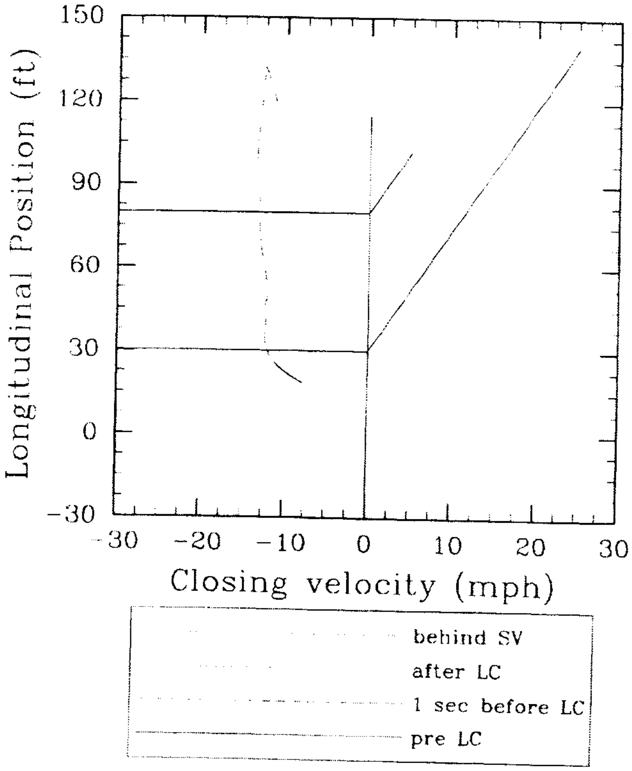

Figure A-5.3-3: Category 3 trajectory at moderate SV speed.

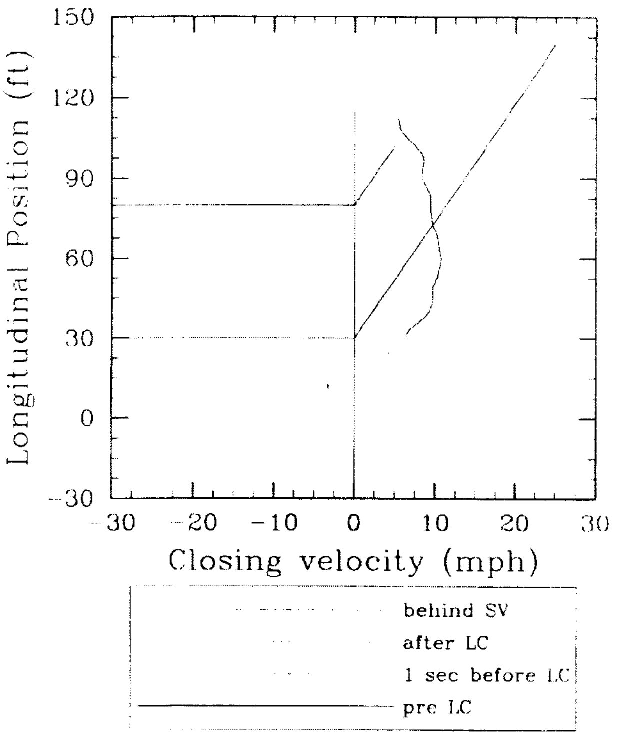

Figure A-5.3-4: Category 4 trajectory with driver warning display triggered

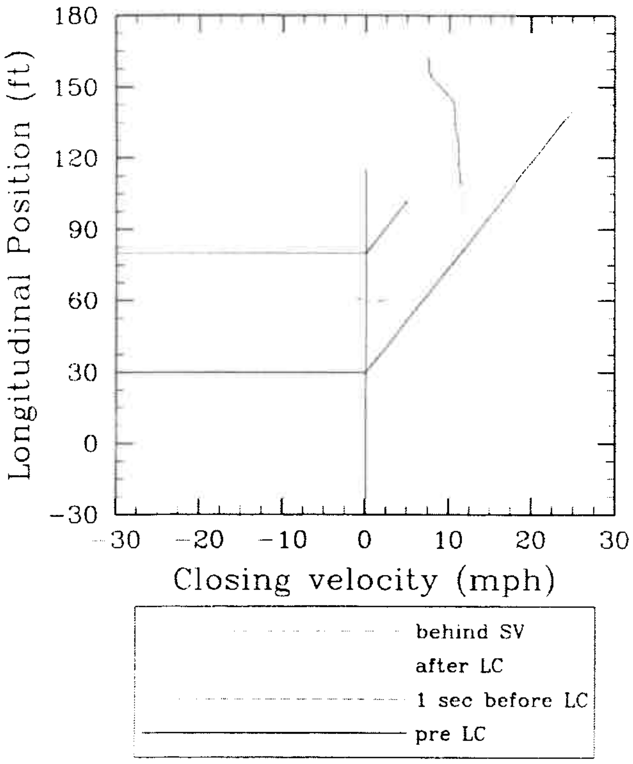

Figure A-5.3-4: Category 4 trajectory

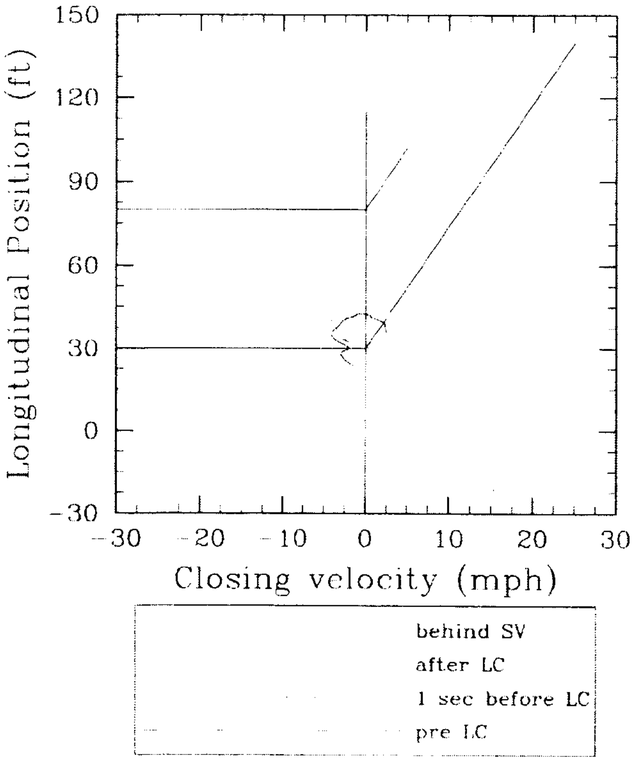

Figure A-5.3-6: Category 5 trajectory

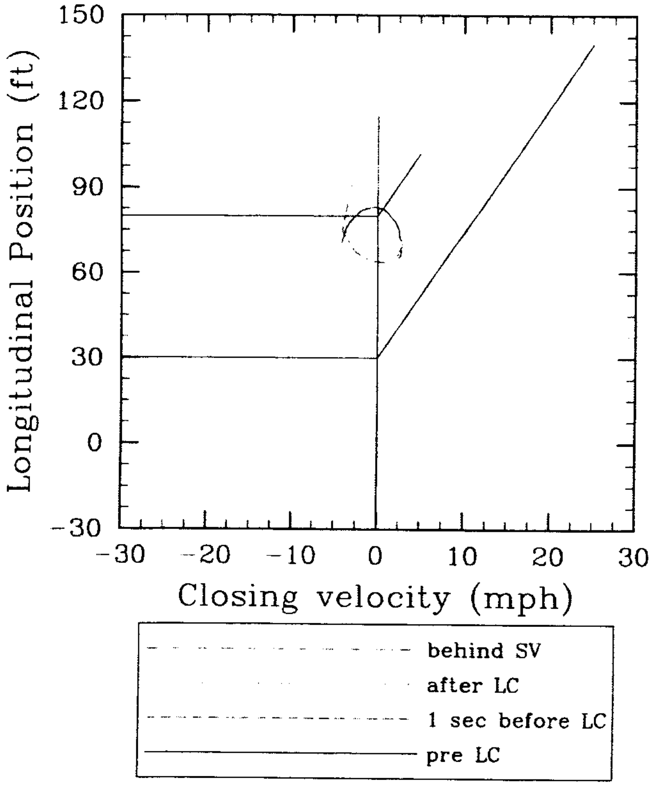

Figure A-5.3-7: Category 5 trajectory

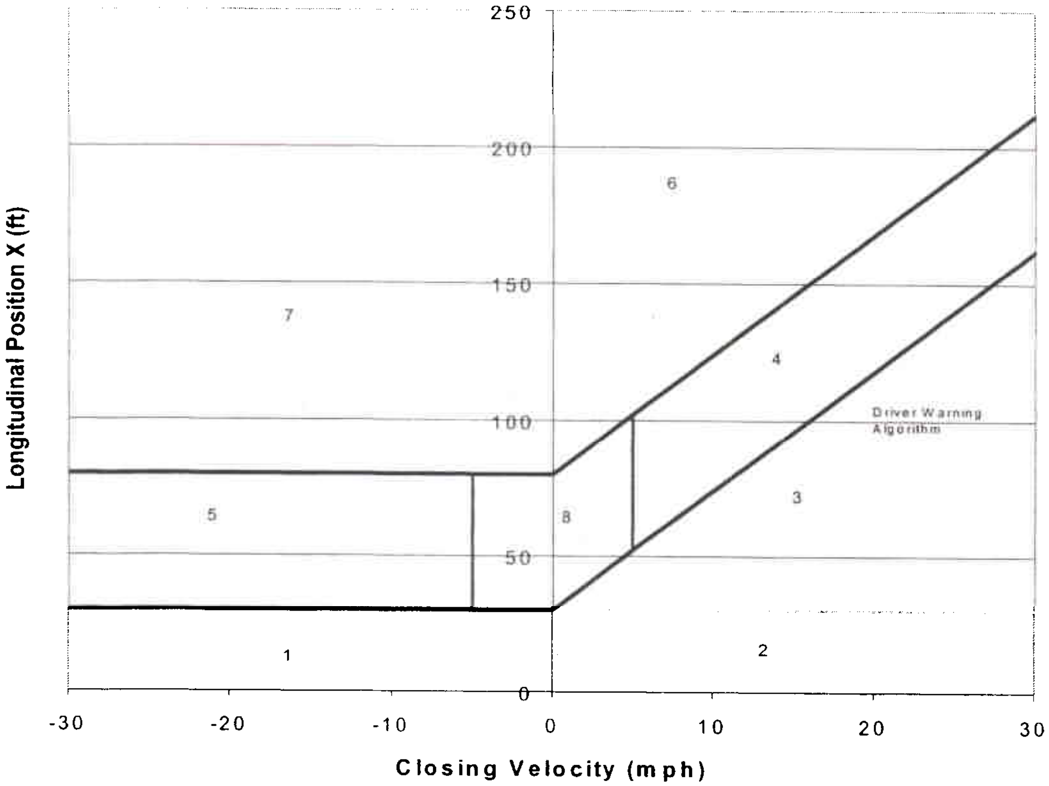

Figure A-5.3-8: Division in Phase Space in Zones

velocity (i.e. |closing velocity| > 5 mph). Zone 8 is the region wherein the SV will make lane changes with low relative velocity. Finally zones 6 and 7 are regions that are fairly distant from the SV.

One way to get a sense of how people make lane changes is to plot the coordinates in phase space of the POV at a given time. Since we define the lane crossing as the moment the SV wheels cross the lane markings it is important to realize that this is in fact a bit late in the decision process. In other words the SV is already in motion and the decision has already been made by this time. Therefore we have found it useful to back up one second and argue that this time is in fact closer to the end of the decision phase. Plotting the POV coordinates at this time for all the baseline, proximity and comprehensive cases gives us the following graphs (Figures A-5.3-9 to A-5.3-11). Not shown on these graphs are those cases wherein there was no POV within sight of the sensor system.

With reference to Figures A-5.3-10 and A-5.4-11 we see numerous lane changes “below the line”, where the driver changed lanes in spite of the warning. One of the most important effects we are looking for in this study is to see whether the use of the CAS is having an influence on the drivers. If drivers can be induced to make safer lane changes, which in effect means taking those below the line lane changes and moving them above the line, then one might infer that when deployed in the field, the number of lane change accidents will decrease. A complete discussion and derivation of the expression for effectiveness can be found in section 6.4 of this report.

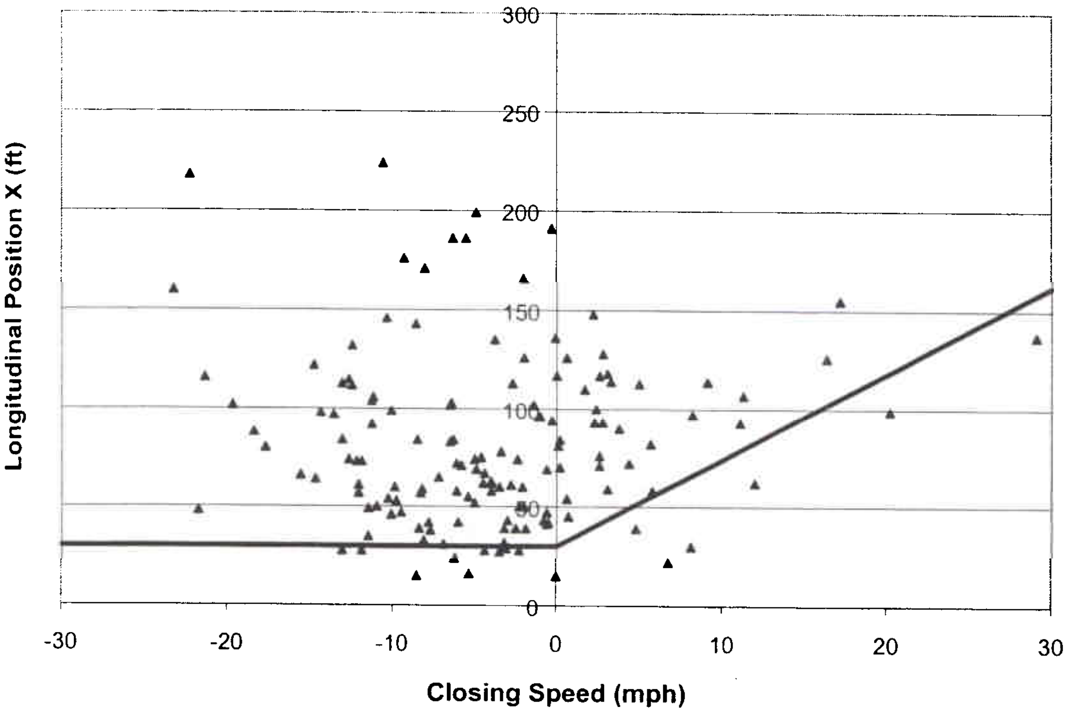

Figure A-5.3-9: Phase Space coordinates at 1 sec before lane crossing for all baseline test runs

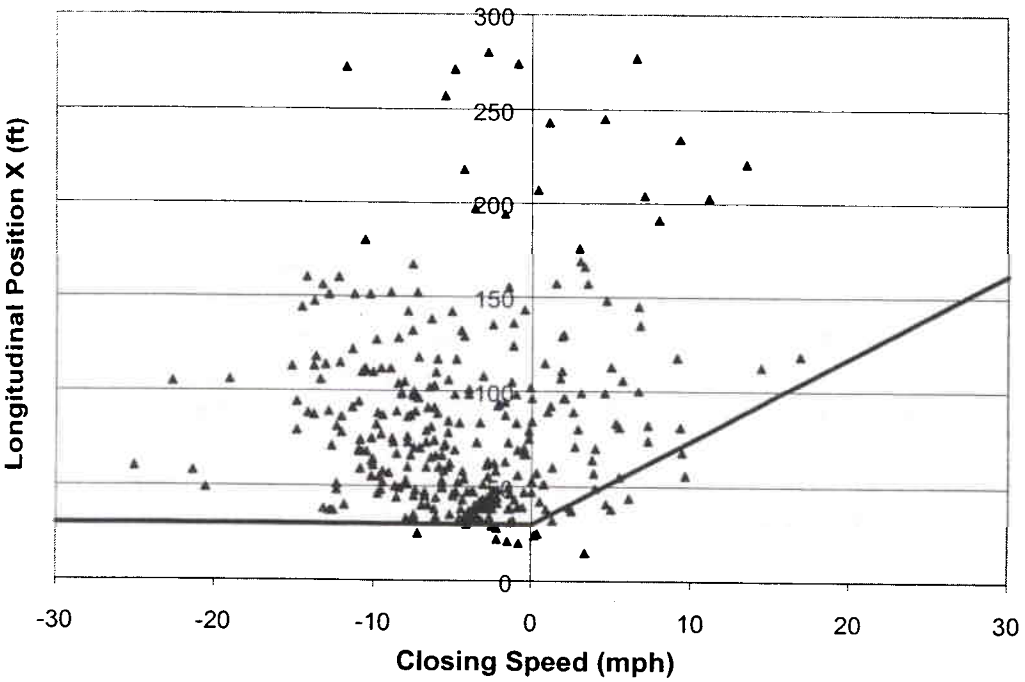

Figure A-5.3-10: Phase space coordinates at 1 sec before lane crossing for all comprehensive test runs

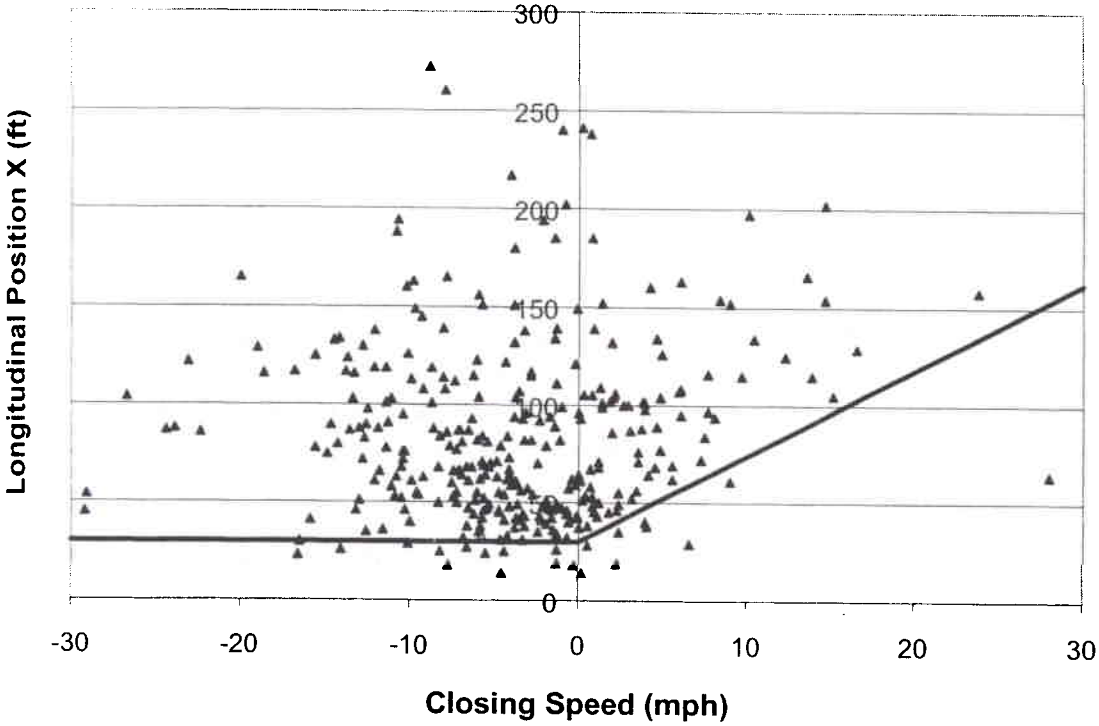

Figure A-5.3-11: Phase space coordinates at 1 sec before lane crossing for all proximity test runs

The final expression for effectiveness is:

$$E = 1 - \frac{fraction \;\; of \;\; lane \;\; changes \;\; in \; conflict \;\; zone \;\; with \;\; CAS}{fraction \;\; of \;\; lane \;\; changes \;\; in \; conflict \;\; zone \;\; without \;\; CAS}$$

The conflict zone is taken to be identical to the warning zone of the comprehensive system.

The number of lane changes under conflict for baseline, comprehensive and proximity runs are shown in Table A-5.3-1. Note that 5 of the 22 lane changes under conflict with the proximity lane change CAS involved fast approaching vehicles outside of the proximity warning zone which extends 30 feet behind the SV. In those cases the drivers were not warned by the proximity lane change CAS.

	Baseline runs	CAS runs	
		Comprehensive	Proximity
# conflict LC's	16	18	22
Total # LC's	273	541	552
Fraction of LC's under conflict	0.059	0.033	0.040

Table A-5.3-1: Tabulation of Conflict/Non-conflict totals for all Right Lane Changes

Substituting into the expression for effectiveness we get for the two systems:

$$E_{comp} = 0.43$$
$$E_{prox} = 0.32$$

In our study the errors in these effectiveness numbers were dominated by the relatively small number of lane changes that fell below the line. These numbers can vary by as much as 50%. However, the methodology is sound and future studies of longer duration should yield more reliable estimates. It is interesting to note that the comprehensive system has a higher effectiveness as one would expect intuitively since it applies in more situations. Also interesting is the fact that the effectiveness for the proximity system is seen to be positive, so that it appears to be able to make a contribution to the reduction of lane change crashes.

As a check on whether these values have any reliability one can perform the identical calculations for lane changes made to the left. This side of course has no warnings and therefore should show no differences. In place of the CAS runs we simply use the measured values from the left side of the testbed for whatever CAS configuration was used on the right side. For example, if the configuration on the right side was center mirror/monitor mode/comprehensive, then we simply tallied the left side values as if there was a comprehensive system being tested on that side. Performing the same substitutions we get for effectiveness:

$$E_{comp} = .15$$
$$E_{prox} = .19$$

This is larger than we expected but on a speculative note we cannot rule out the fact that since the CAS seems to be influencing more conservative behavior on the right side, that it may also be spilling over to the left side.

It has been suggested that one qualitative measure of the effect of the CAS might be the increase in category 2 trajectories (see Figure A-5.3-1). These are the situation where the SV allows a POV to pass before making a lane change. Using our statistical analysis package we were able to find that there is a weak dependence of the fraction of category 2 lane changes on the location of the display, but not the system type. The F value for location is 3.22, while the critical F value for a confidence level of .05 is 3.23. The means and standard deviations for display location are graphed in Figure A-5.3-12, and for system type are graphed in Figure A-5.3-13 for comparison. At present there is no clear interpretation of the reason for the location dependency. Nor is there one for the lack of any clear effect of CAS vs. baseline. It may simply be an artifact of the poor statistics or there may in fact be no causal relationship.

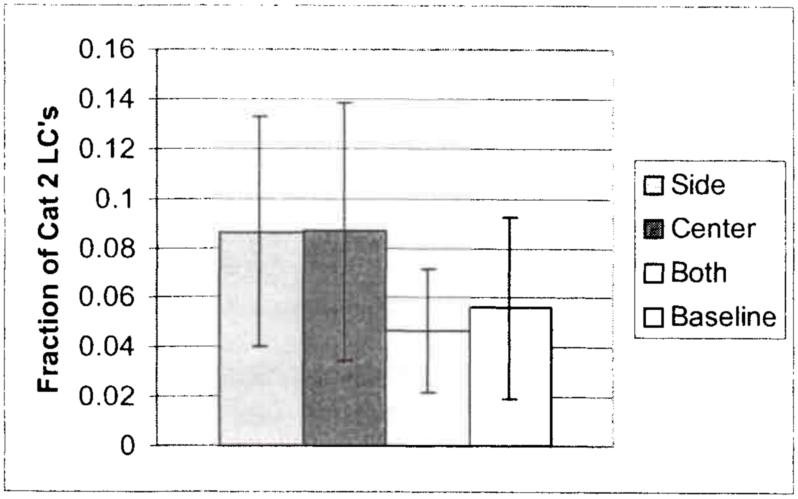

Figure A-5.3-12: Fraction of Cat2 lane changes as a function of display location

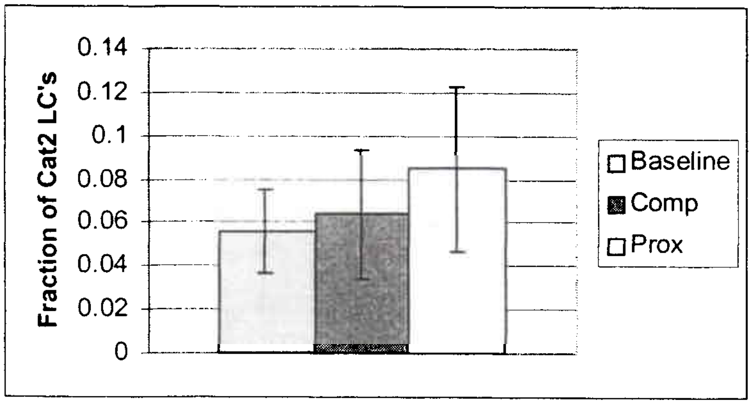

Figure A-5.3-13: Fraction of Cat2 lane changes as function of system type

It is interesting to collect the data concerning the velocity distribution of POV's that have passed the SV. We have simply summed the passing vehicles over all runs and all drivers, but segregated according to right side/left side and above and below 50 mph, SV speed. This last constraint is to separate freeway from city street driving. The data was taken only when the SV was in motion (specifically speed> 10 mph). Figure A-5.3-14 shows the results. Generally speaking the slower SV velocities show a longer high passing velocity tail. In addition the left side distribution is slightly more skewed to the higher speeds. Figure A-5.3-15 is the distribution for all passing vehicles and Figure A-5.3-16 is the cumulative fraction of passing vehicles as a function of relative passing speed. An important point to note is that accommodating relative passing speeds of 30 mph as opposed to 20 mph increase the cumulative percentage from 95% to 99%. This may not seem like much, but it is precisely these high velocity impacts that cause the most injuries and damage. This has an implication on the longitudinal extent of the fast approach zone.

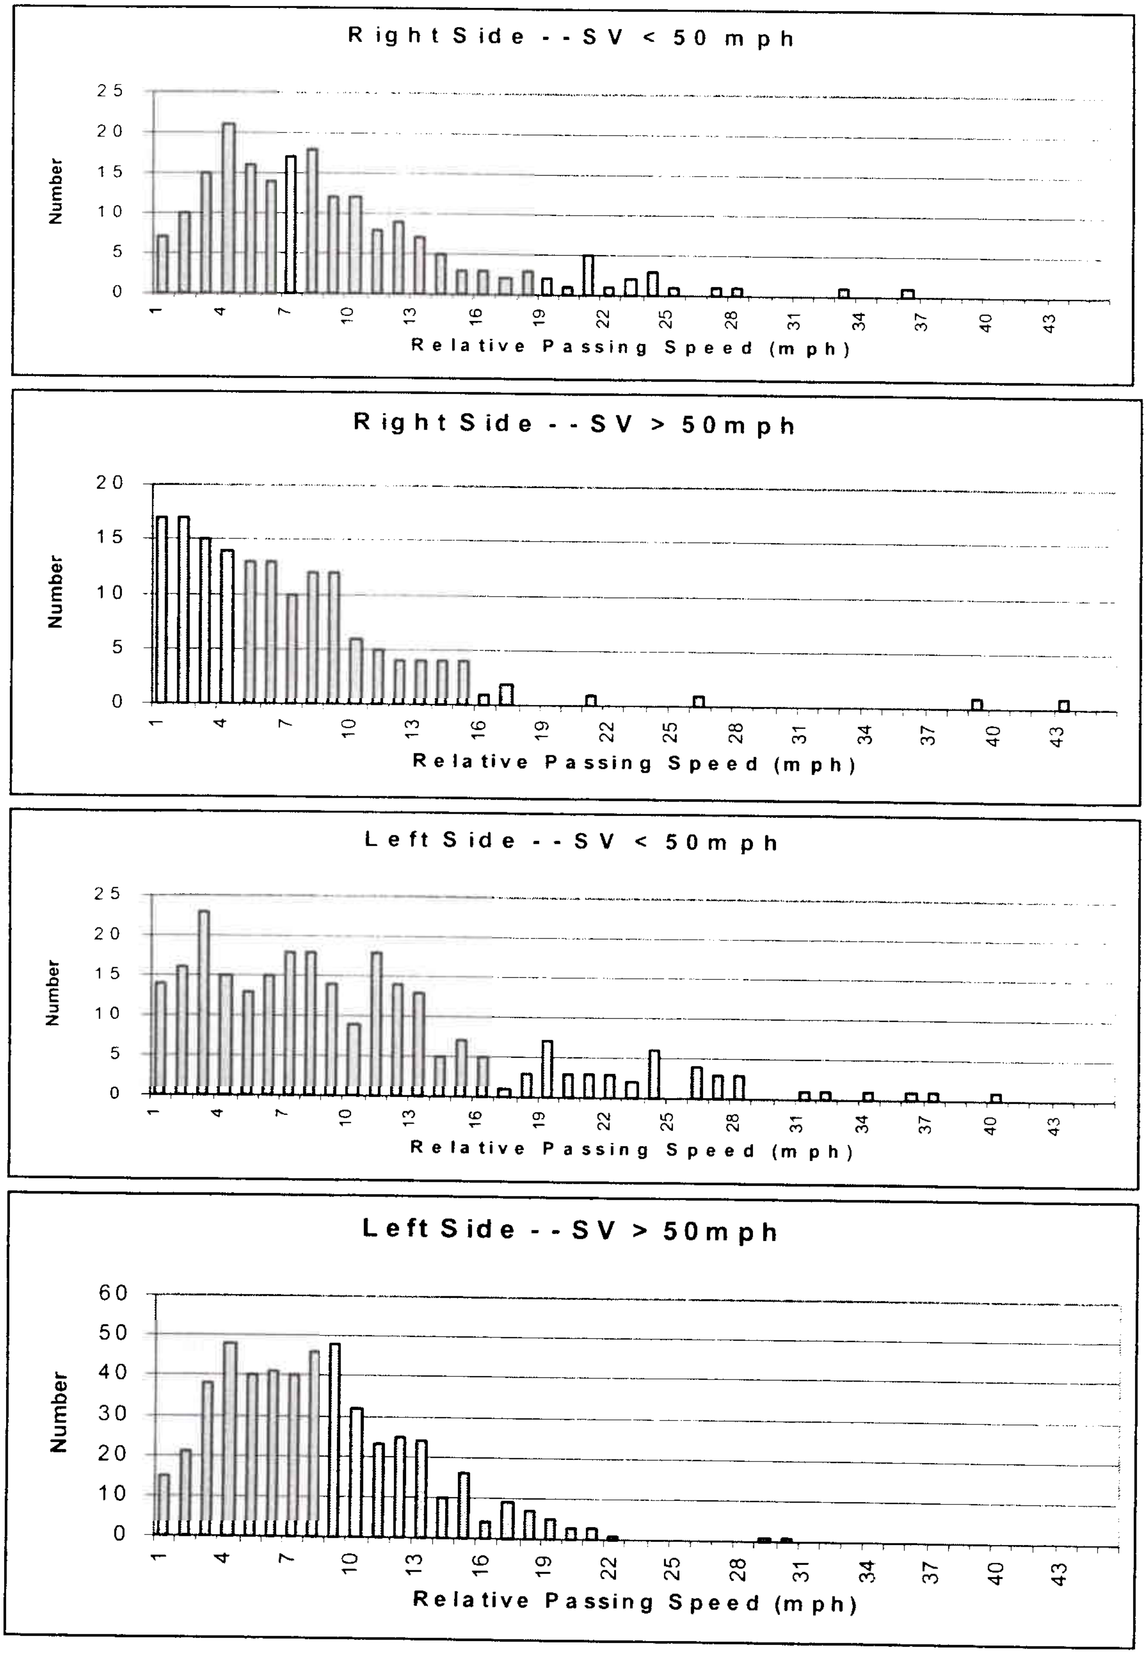

Figure A-5.3-14: POV relative passing speeds

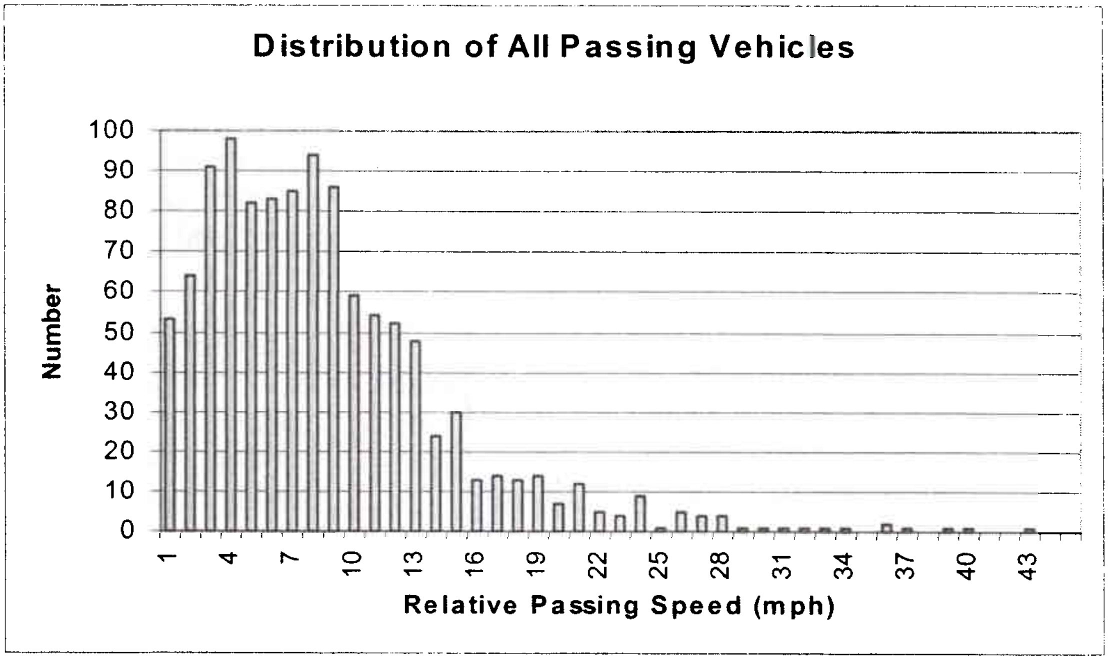

Figure A-5.3-15: Distribution of all passing POV's

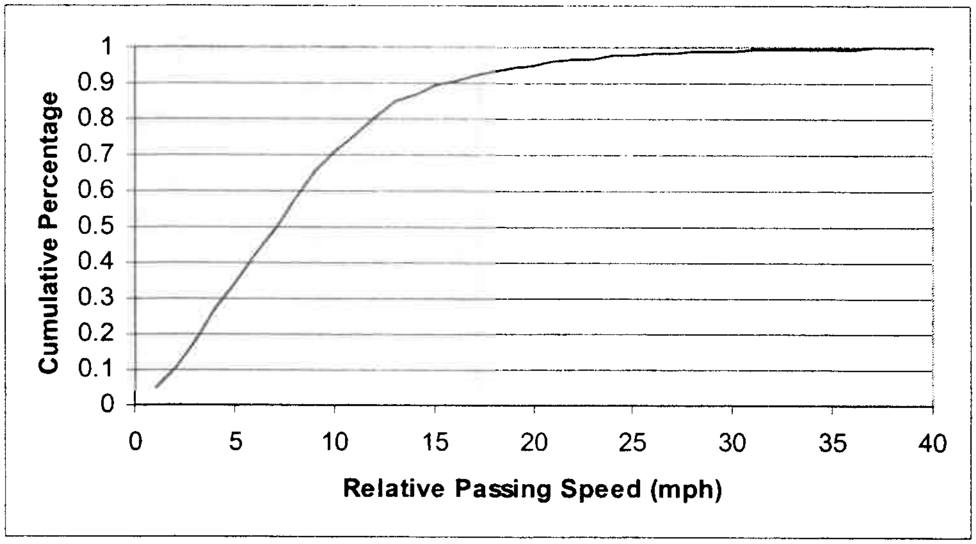

Figure A-5.3-16: Cumulative Percentage of all passing POV's

A-5.4 Eye Tracker Results

If there is evidence to suggest that the CAS is having an effect on the subject's driving behavior, then there must be some evidence of that in the eye glance behavior. First it is useful to get some idea of some of the eye glance characteristics during baseline driving. Table A-5.4-1 is such a table for two different drivers. The values here represent average values for one baseline run for each driver. The lane change period is defined as starting ten seconds before the lane crossing to 5 seconds after the lane crossing.

	Right Lane Change			Left Lane Change		
Glance Location	Dwell time per glance (sec)	No. of glances per lane change	% time during lane change	Dwell time per glance (sec)	No. of glances per lane change	% time during lane change
M21						
Straight Ahead	5.2	2.42	93.51	1.91	5.04	74.67
Rear View Mirror	.3	1.33	2.94	.33	.45	1.14
Right Mirror	.29	.33	.72	0	0	0
Right Window	0	0	0	0	0	0
Over Right Shoulder	0	0	0	0	0	0
Left Mirror	.26	1.33	2.52	.57	4.94	21.77
Left Window	.47	.03	.1	.51	.49	1.96
Over Left Shoulder	0	0	0	0	0	0
F61						
Straight Ahead	2.98	3.62	73.07	2.25	3.68	66.68
Rear View Mirror	.4	4.29	11.69	.5	.36	2.08
Right Mirror	.35	2.9	6.87	0	0	0
Right Window	.16	1.05	1.13	0	0	0
Over Right Shoulder	.75	1.05	5.33	0	0	0
Left Mirror	1.03	.05	.33	.58	4.57	21.34
Left Window	0	0	0	.31	1.82	4.51
Over Left Shoulder	0	0	0	.78	.64	4.05

Table A-5.4-1: Eye glance characteristics for baseline driving during lane change averaged over a run

These two drivers represent polar extremes in age and gender. M21 only uses his left and center mirrors and never looks over his shoulders, whereas F61 makes use of all her mirrors and looks over the relevant shoulder. Generally speaking, those who look over their shoulders represent slightly more than half the subjects tested. Most drivers use both mirrors when making a lane change to the right. By way of comparison, the same numbers for the far greater periods of time when there were no lane changes are presented in Table A-5.4-2. The differences between the two drivers narrow when not making lane changes.

	Non-Lane Change Period	
Glance Location	Dwell time per glance (sec)	% time during lane change
M21		
Straight Ahead	3.14	85.93
Rear View Mirror	.37	3.76
Right Mirror	0	0
Right Window	0	0
Over Right Shoulder	0	0
Left Mirror	.56	2.45
Left Window	1.18	.81
Over Left Shoulder	0	0
F61		
Straight Ahead	2.59	81.24
Rear View Mirror	.5	2.82
Right Mirror	.65	.86
Right Window	.57	.05
Over Right Shoulder	.78	.06
Left Mirror	.64	2.99
Left Window	.78	.6
Over Left Shoulder	0	0

Table A-5.4-2: Eye glance characteristics for baseline driving averaged over a run

The key eye planes for determining whether the CAS has an effect on the driver are straight ahead, center mirror and right side mirror. By evaluating the dwell time per glance, the number of glances per lane change and the fraction of time spent looking at that plane during the lane change period, we can determine whether there has been any effect. Figure A-5.4-1 shows these three quantities for the right side mirror as a function of location, mode and system type. They are also classed according to whether or not there was a warning during the lane change or not. The term baseline with or without conflicts is an attempt to correlate better with situations with the CAS energized wherein there was a warning. As it turns out it made little difference, so it does not appear in every plot. The descriptions in the legend take the form location-mode-display on or off. Error bars are plotted for each value to give a sense of the variability. Looking at the % time during lane change, we see a definite increase over baseline and some variation within the subset of display on. Running our statistical analysis (ANOVA) on the case when the warning is displayed, we get a strong correlation with location ($F = 5.4$, where $F_{crit}=3.23$). We can see that the longest relative time spent looking at the side mirror is when that mirror alone is energized. The fraction of time spent looking at a particular eye plane is really the product of the number of glances per lane change and the dwell time per glance. For the right side mirror the source of most of this increase in fractional time seems to come from the dwell time per glance. To determine the statistical significance of

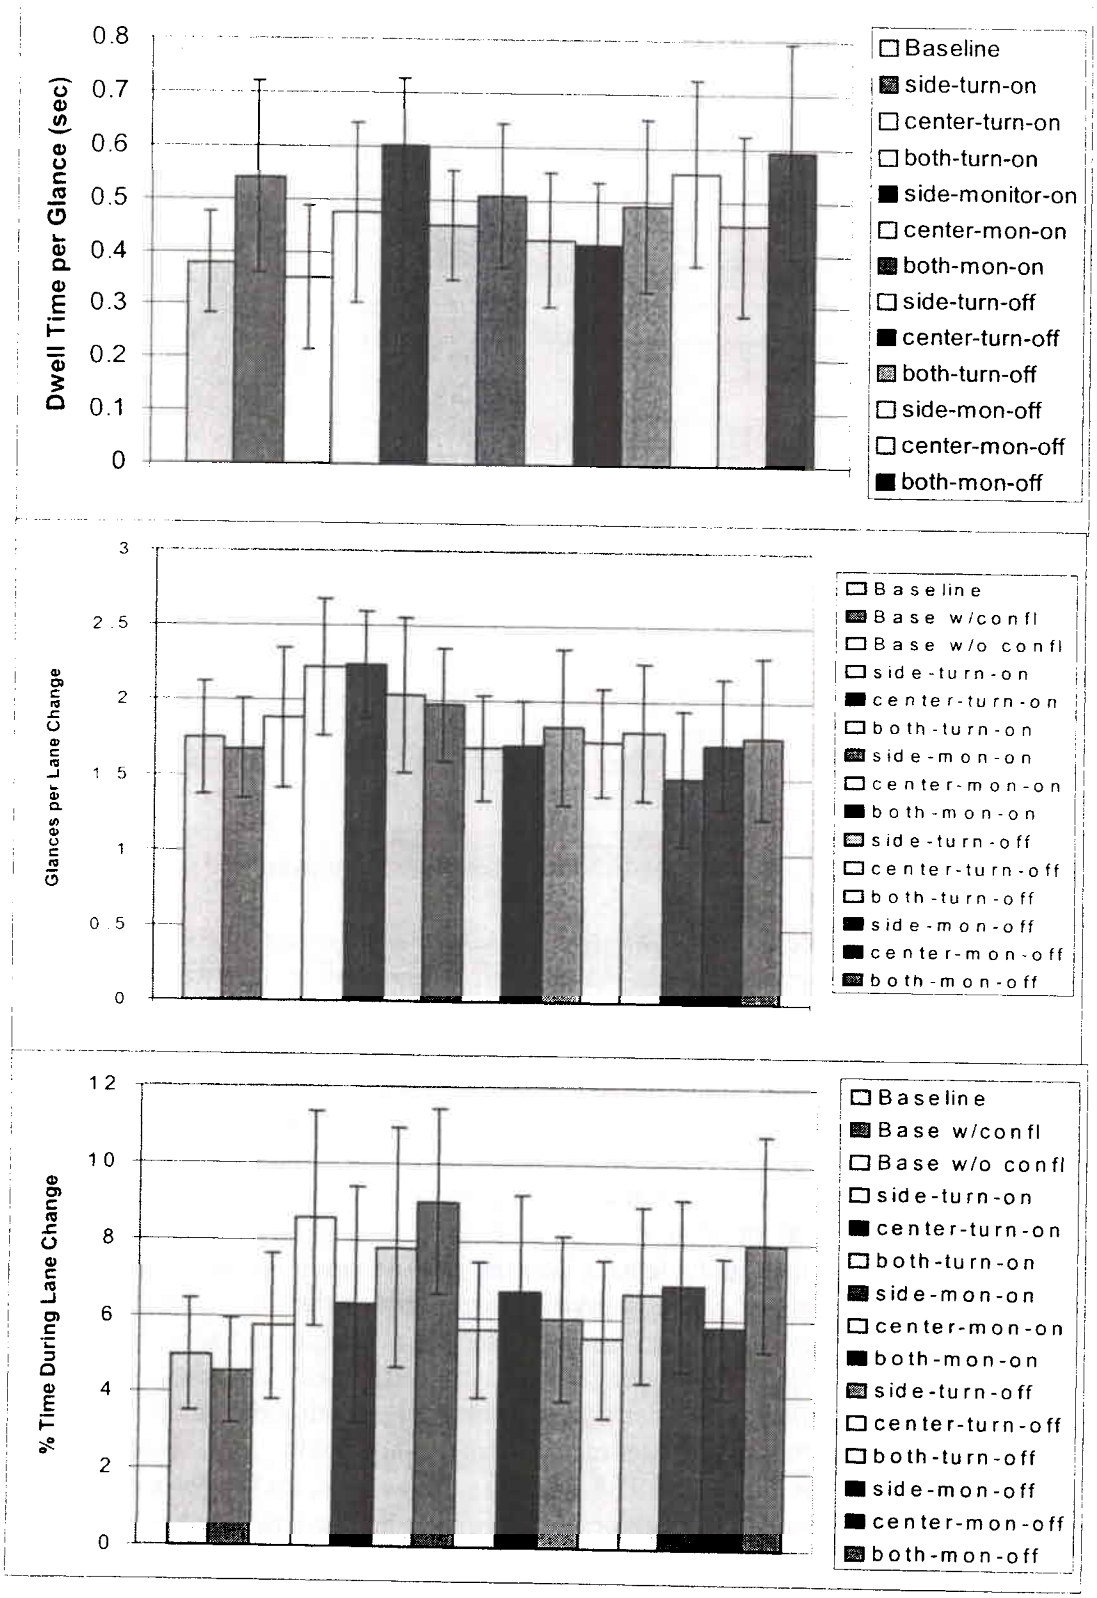

Figure A-5.4-1: Eye statistics for Right Side Mirror the increase in time spent looking at the side view mirror when only that mirror is lit we perform a t-test comparing the baseline values with those lane changes when the display was lit during the lane change period and then the baseline with those changes where it did not come on at all during the lane change period. A t-test in the former case yields a p-value of 0.023, which means that there is a moderate evidence for the fact that the mean for the baseline case is statistically different than that for the case where the side mirror is lit. The same test comparing baseline with the case where the display was not lit yields no difference in the mean time spent looking at the side mirror.

The same plot for the center mirror is shown in Figure A-5.4-2. The source of the increased time spent in looking at this mirror when there are warnings is clearly located in the number of glances per lane change. Performing the ANOVA test for those cases with the display lit yields $F = 3.6$ for location ($F_{crit}=3.23$) and $F=4.7$ for mode ($F_{crit}=4.08$) for the number of glances per lane change. These are mild correlations, and they disappear when the ANOVA test is performed on the fraction of time spent looking at the rear view mirror. Interestingly, the general trend is that the lowest fractional time spent looking at this mirror is during the baseline runs. The highest is during those lane changes during which there was at some time a warning. The in-between case occurs for those lane changes wherein there was no warning. Because the ANOVA test yielded no correlation between the time spent looking in the rearview mirror and any of the independent variables when the display was both lit and unlit during a lane change, we can simply group all the lane changes when the display was lit and all the lane changes when the display was not lit and compare them to the baseline case via the t-test. When we compare display on to baseline we get a p-value of 0.0098, which is evidence for a real effect. Comparing display off to baseline yields no such evidence.

Finally we plot the fractional time spent looking straight ahead in Figure A-5.4-3. Clearly there is a drop in eyes on the road time with the use of the CAS. ANOVA tells us that there is no correlation between time spent looking straight ahead and any of the three independent variables. Therefore as in the case for the center mirror we can compare the baseline case with the sum of all lane changes made when the display was lit and not lit. The p-value for display on versus baseline equals 0.003 which is strong evidence for a real difference, whereas there is no evidence for a real difference between display off and the baseline.

Although it has often been posited that a lane change CAS would aid the driver by allowing him to spend more time looking straight ahead, ironically we see in these results that the average driver spends less time looking straight ahead. Since the structure of this study was such that it was impossible to study long term adaptation, it is entirely possible that all of this eye data may be a form of the novelty effect. If it is real, and persists over time, it may be something to be concerned about. In any case, at the present we do see clear evidence that the drivers are noticing the displays.

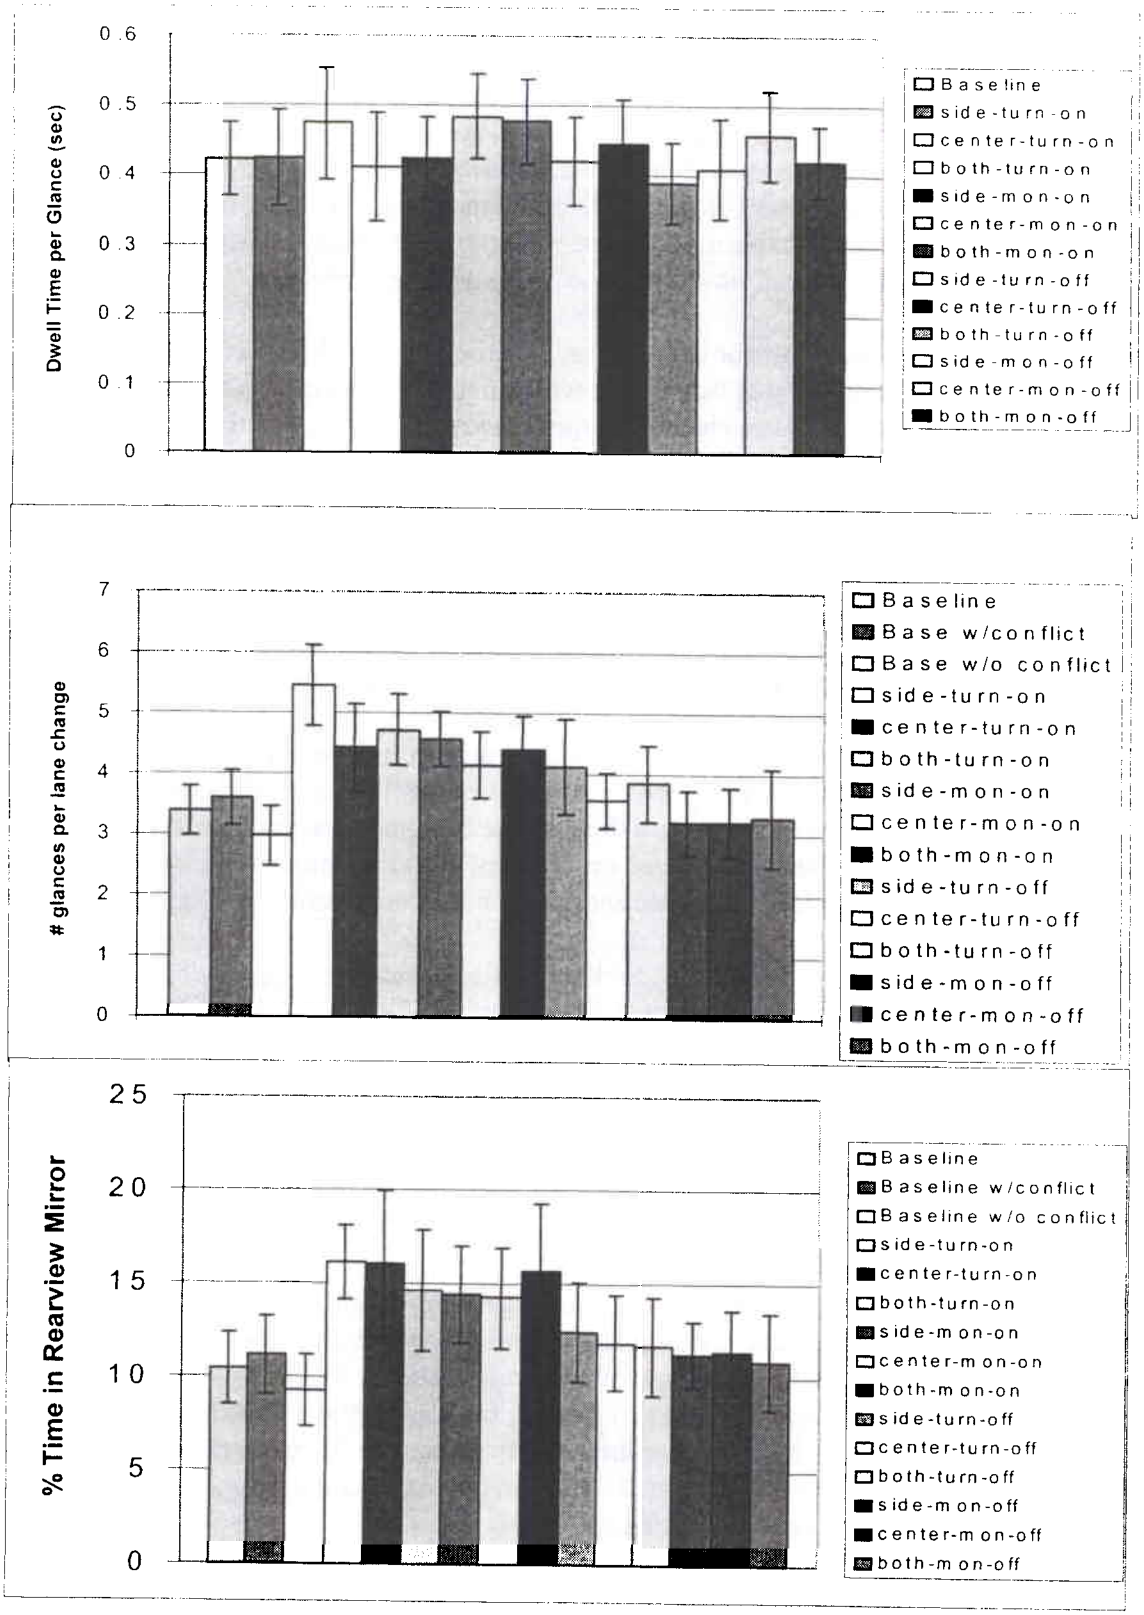

Figure A-5.4-2: Eye statistics for Center Mirror

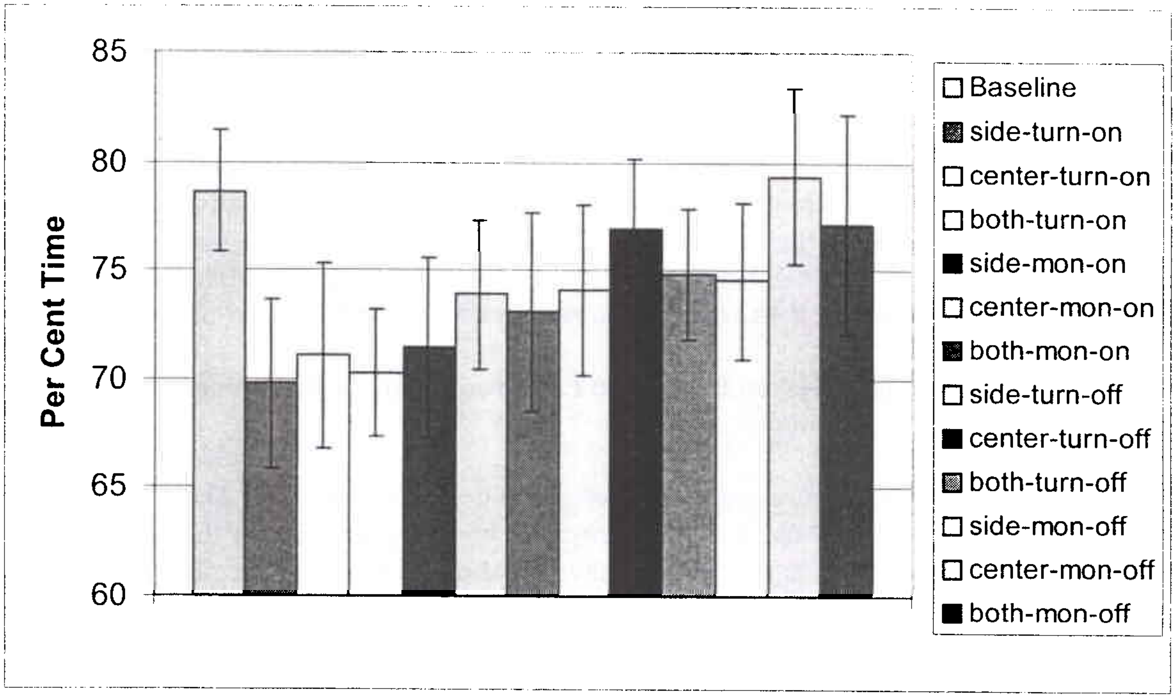

Figure A-5.4-3: Fractional time spent looking straight ahead during lane changes

A-5.5 System Performance and Driver Reactions

Through thorough analysis of a selected sample of runs, one can compile statistics concerning the base performance of the lane change system. Each run contains literally many hundreds of interactions by which to judge the performance of the lane change CAS. The criteria for choosing the runs for analysis were those runs wherein the test conductor noted the possibility of a missed detection. The following performance quantities listed in Table A-5.5-1, will first be defined and then evaluated. In that light, driver responses as detailed in the questionnaires have been compiled and presented.

True Positive: Warnings given on targets that can be considered threats; namely other moving vehicles.

False Positive: Warnings generated when there is no true threat present. The overwhelming majority of these alerts are generated by stationary objects that have evaded detection as a stationary object and have therefore triggered a warning. In addition warnings can be generated with no true object present as explained in section A-2.6.2.

False Negative: Warnings not generated when there is a moving vehicle present.

Detection Probability (P_D): Ratio of True Positive to the sum of True Positive plus False Negative.

Inappropriate warnings: Warnings given for stationary objects which the algorithm is supposed to eliminate. These include guardrails, bushes, poles and parked cars. Generally speaking these warnings are assumed to be a major contributer to the nuisance alarms.

Rejection Ratio: Ratio of the number of inappropriate targets that do not generate a warning to the total number of those (stationary) targets.

Nuisance Alarm: Warnings given that the driver considers annoying.

P_D	99.3%
Rejection Ratio	90.7%
False Positive Alarm Rate	42 / hr

Table A-5.5-1: Performance Statistics for CAS

Two comments are in order. First, we believe that the detection probability is a lower bound. The particular runs that were selected for evaluation were those in which it was noted by the operator that there might have been an instance of a false negative. The overwhelming majority of the runs did not show any evidence of false negatives. It is important to note that upon careful examination, it was determined that it was not a failure of the sensor that resulted in a lack of warning, but rather a failure of the algorithm to correctly anticipate the situation. The situation was that the proximity warning was disabled by the presence of stopped traffic in the right adjacent lane. The system would reset itself when the SV slows to below 10 mph or the zone is clear for 3 cycles (0.3 sec). Since neither of these events happened before the stopped traffic in the adjacent lane starting moving again, the system failed to warn on the now moving cars. This is clearly correctable in software if it had been properly anticipated.

The second comment has to do with inappropriate alarms. The vast majority of these alarms are triggered by stationary objects that enter the proximity zone on a slight angle and thus avoid being tracked and evaluated in the forward looking zone. Typically, passing alongside a row of bushes in such a way that the bushes lie just on the boundary of the proximity zone will trigger a number of alarms. There will always be some branches that will not be observable until they are already in the proximity zone. Guardrails might also trigger alarms if they are close to the zone boundary in much the same way. Depending on how the laser beam intersects the guardrail, it might show up as a solid line or as a row of discrete points if the laser strikes the supports.

The driver questionnaires had three sections to them. The first section dealt with personal information as relates to driving. Eleven out of the twelve drivers have had accidents in the last three years. Out of these eleven, six involved changing lanes. This is clearly way out of the norm, considering that only about 5% of the total number of accidents nationally involve changing lanes. Also in the same section, all the respondents said that they used their turn signals either often or constantly. In fact the lowest individual usage of turn signals in these runs was 98%. It is conceivable that having a ride along observer influenced turn signal usage.

The questions pertaining to the individual drive configurations are presented in the following table. A copy of the actual questionnaire can be found in Appendix B. In this table, yes or no response are scored as zero or one, while questions asking for a range of responses according to the levels of choice given. For example, when asked how meaningful was the information when the display was turned on (Appendix B, question 7, part 2) the choices presented to the subjects and the numerical values attributed are shown below.

Level	Numerical Score
Very important	5
Important	4
Somewhat important	3
Slightly important	2
Not at all important	1

Question	Scale	Score	Analyst's Comments
Was system useful in conflicts?	1 – 4	3.23	Closer to slightly useful than very useful
Did you notice any alarms due to things other than cars?	0,1	0.32	Noticed nuisance alarms in about 1 out of 3 rides
How meaningful was the information?	1 – 5	3.47	Halfway between somewhat important and important
Were there too many inappropriate alarms?	0,1	0.03	Occurred on 2 out 72 rides
Was alert noticeable?	1 – 3	2	Noticed when needed, could ignore it when not needed
Did the alert annoy you?	0,1	0	No one found the display warnings annoying
Was the alert helpful in notifying you of potential conflicts?	1 – 5	3.65	Halfway between somewhat helpful and helpful
Did it provide you with adequate notice of hazards?	0,1	0.88	Provided adequate notice of hazards in almost 9 out 10 rides.
How did the warnings affect your driving?	1 – 5	4.10	Just above slightly positively affected
Did the blinking display attract your attention?	0,1	0.83	Attracted attention is just over 8 out of 10 rides
How did blinking display affect your driving performance?	1 – 5	4.06	Slightly positive

Table A-5.5-2: Summary of Driver Questionnaires

Finally we summarize the drivers recommendations as to the type of variable combinations that they preferred.

Variable	Options	% choosing that option
Location	Center Only	13
	Side Only	37
	Center + Side	50
Mode	Monitor	58
	Turn Signal	42
System	Proximity	42
	Comprehensive	58
Timing of Comprehensive System	Too Early	33
	Too Late	0
	About Right	67

Table A-5.5-3: Summary of Driver's Preferences

In discussing the responses it is perhaps more instructive to start with those questions that elicited a strong response. For example no one thought that the display warnings were annoying and that they would notice it when they needed it and ignore it when it was not needed. These responses tend to support the conclusion that Muth mirrors are user friendly. In the case of false positive alarms, the drivers noticed them on approximately 1 out of 3 rides, and on only 2 out of 72 rides did anyone find them annoying. Of those runs wherein the driver noticed false positive alarms, 57% were with monitor mode while 43% were with turn signal mode. Since the purpose of turn signal mode was to reduce the noticeability of false positive warnings this is not as strong a split as might have been expected. However both of the rides wherein the drivers found them to be annoying occurred during monitor mode.

A question arises as to whether there is a contradiction between the fact that on such a small percentage of the drives were the false positive alarms found to be annoying, but yet the average false positive rate is 42 per hour. There are two general explanations for this result. First, the displays themselves are unobtrusive and can be ignored when not needed. In addition, comments from the drivers generally find the information conveyed to be of use. In a qualitative sense, one can therefore consider the annoyance to benefit ratio as being fairly low. The second explanation has to do with the character of the false positive alarms. Since the overwhelming majority of these alarms are generated by stationary objects entering the proximity zone without being tracked in the forward looking zone, this rate is not likely to vary as a function of system type. However it is worth noting that this is an average rate, while in practice these alarms are clustered. They occur primarily while the driver is driving in the right lane next to walls, bushes, etc. Clearly the driver has no interest in making a right lane change and therefore is capable of ignoring or not even noticing the displays. When driving toward the left in a multi-lane road, false positive alarms are very rare. In addition, it is important to note that false positive warnings almost invariably last for a fraction of a second. However true positive warnings last as long as there is a threat. This can last from a few seconds to many tens of seconds if a vehicle is traveling alongside at roughly the same speed as the SV. The true nuisance alarm rate can be calculated by multiplying the false positive alarm rate by the fraction of test drives that the display was found to be annoying (0.03). This yields a nuisance alarm rate of 1.2/hr. This is a very reasonable and acceptable level of nuisance.

In 88% of the rides, the drivers stated that the system provided them with adequate notice of potential hazards. Of those rides where the driver felt that the system did not provide adequate notice the breakdown between turn signal vs. monitor mode and proximity vs. comprehensive system is as follows. There was almost an even split between the proximity (56%) and comprehensive (44%) system. However between the modes there was more of a difference with 22% of the negative responses coming from monitor mode runs, while 78% of the negative responses coming from the turn signal mode runs. Statistical analysis shows that this is a significant difference if we relax the confidence level from 0.05 to 0.10. This is acceptable for experimental research of this type. The F value for the mode variable is 3.9 while the critical F is 2.84. All the other variables did not come close to being statistically significant. Even though a large majority of the cases where the system did not provide adequate notice were turn signal cases, slightly more than 40% of the respondents preferred the turn signal mode over the monitor mode. It would seem that there is a clear reason for overruling this preference in the establishment of design guidelines, in light of the fact that the turn signal mode does not always provide adequate notice of hazards. However, given the limited extent of these tests, it might be premature to draw such a conclusion. It is definitely something to be alert for in future studies.

In terms of driver acceptance, the relevant questions such as usefulness in conflicts, helpfulness in notification of conflicts and how did the warnings affect your driving, the responses clearly indicate a positive response although not overwhelmingly so. However some of the questions had a more polarized response. These included questions as to how meaningful was the information presented and how helpful was the alert in notifying you of potential conflicts. The impression of the test observer was that some people were quite enthusiastic about the system while others commented that they basically ignored it. The spread of responses to the question about how the CAS affected one's driving elicited a generally a positive response with not one negative response and only a few neutral responses.

It is interesting to note that as regards the timing of the comprehensive system, no one thought that the system warned too late. Most thought it about right while 1/3 thought that it warned too early. Our bias toward a conservative system was evidently close to the mark.

A few last observations. There were only three close calls during the entire test series. All of these involved another vehicle changing lanes into the same lane as the SV. In one of these situations the system was configured to only display a warning in the center mirror. The driver informed the observer that while he was looking at the side mirror to check for conflicts to the rear of the SV, he caught the warning display in the center mirror in his peripheral vision and then looked forward to see the threatening vehicle. On this basis it seems prudent to have the display in both mirrors to accommodate all possible situations and driver habits.

A-6.0 Summary and Conclusions

- Built an instrumented testbed that was highly effective in collecting new and interesting data in an extremely detailed fashion.
- Developed a "prototype" CAS that demonstrated good performance and driver acceptance.
- Applied the data collected to projecting a rough estimate for probable effectiveness
    - $E_{comp} = .43$
    - $E_{Prox} = .32$
- Acquired significant data on driver eye behavior
- Showed that the displays drew the driver's attention
- Data implies that turn signal mode may not always give adequate notice of potential hazards
- Drivers did not consider the relatively high (42/hour) rate of false positive alarms as annoying
    - Only noticed it about 1/3 of the time
    - Felt it was annoying only about 3% of the time
    - Resulting nuisance alarms rate = 1.3/hour
- This study points out the need for more extensive work, collecting data of this kind to assess the issue of driver long-term adaptability and to collect better statistics on the effects already noted.

A-7.0 References

1. S. Talmadge, D. Dixon and B. Quon, “Development of Performance Specifications for Collision Avoidance Systems for Lane Change Crashes, Task 6 Interim Report: Testbed Systems Design and Associated Facilities,” NHTSA Technical Report DOT HS 808 369, November 2001.
2. S. S. Blackman, “Multiple Target Tracking with Radar Applications,” Artech House, Inc., 1986.
3. R.R. Kirk, “Experimental Design: Procedures for the Behavioral Sciences,” Brooks/Cole Publishing Company, 1985.

Appendix B - Display Modality Questionnaire

PART I - Personal Information (to be answered before Introduction session)

Please answer the following questions as honestly as possible.

Date:__________________________ Name:____________________________

1) How old are you? ____

2) What is your profession? ____

3) How many years have you had your driver's license?

____

4) What type of vehicle(s) do you regularly
drive?___________________________

5) How many tickets have you had in the last 3 years?_______________________
For what offenses?____

6) How many accidents have you been involved in (regardless if
they were your fault or not)? _______________________________________

7) How many of those accidents involved you or someone changing or
drifting into another lane? _______________________________________

8) To what degree do you use your center mirror when driving a passenger car?

_____ constantly (skip to question 10)
_____ often (skip to question 10)
_____ sometimes
_____ seldom
_____ rarely

9) Can you explain why you do not use your center mirror often or constantly
(for example, do you just look over your shoulders more)?

____
____
____

10) To what degree do you use your right outside mirror when driving a passenger car?

_____ constantly (skip to question 12)
_____ often (skip to question 12)
_____ sometimes
_____ seldom
_____ rarely

11) Can you explain why you do not use your right outside mirror often or constantly

____
____
____

12) To what degree do you use your turn signal when it would be appropriate for you to do so?

_____ always
_____ most of the time
_____ sometimes
_____ seldom
_____ rarely

END OF PART I

PART II - After Each Driving Session

Session Number:____________________

Conditions:________________________

Date:_________________________

Subject Name :_____________________

1) Overall, how tense or relaxed were you while making lane changes during this driving session?

_____ very tense
_____ tense
_____ in-between
_____ relaxed
_____ very relaxed

2) How confident or unsure were you in your ability to avoid getting into an a crash while making a lane change?

_____ very confident
_____ confident
_____ neither confident nor unsure
_____ unsure
_____ very unsure

3) How many conflict situations, where you considered changing lanes in front of another vehicle, did you experience? (If none, go to question 5.)

____

4) For the worst of these, was your warning useful?

_________ very useful
_________ slightly useful
_________ not sure
_________ not useful

5) Were there any instances in which you noticed other things besides cars that turned the display on?

_____ no (go to question 7)
_____ yes – what were they? ______________________________________

6) What kind of effect, if any, did this have on your willingness to rely on the system?

_____ positive
_____ slightly positive
_____ none
_____ slightly negative
_____ negative

7) When the display turned on, how meaningful was that information to you (that is, did it matter much to you)?

_____ very important
_____ important
_____ somewhat important
_____ slightly important
_____ not at all important

8) Did you feel that there were too many inappropriate alarms?

_____ no (skip to next question)
_____ yes (answer below)
_____ far too much
_____ too much
_____ somewhat too much
_____ slightly too much

9) To what extent did an alert <u>distract</u> you from you driving?

___ It would distract me from my driving almost every time it came on.
___ It would sometimes distract me from my driving.
___ It would rarely distract me from my driving.

10) To what extent was an alert <u>noticeable</u> to you while driving?

___ It was not noticeable when I needed it.
___ I could notice it when I needed it, ignore it when I didn't need it.
___ I noticed it even when I didn't need it.

11) To what extent did an alert <u>annoy</u> you while driving?

___ I wanted to disconnect the display.
___ I wanted to turn the display brightness down.
___ I was not annoyed.
___ I appreciated the alert to warn me of other vehicles.

12) Did you ever confuse an alert with something else while driving?

___ No
___ Yes___Describe

13) How did you interpret the alert signal?

___ Go in the direction indicated if you need to change lanes.
___ Do not go in the direction indicated if you need to change lanes.
___ Put turn signal on in the direction indicated.
___ Turn signal is "on" in the direction indicated
___ The alert had no meaning to me at all
___ Other (Specify)

14) How helpful was the alert in notifying you of conflicting vehicles?

_____ very helpful
_____ helpful
_____ somewhat helpful
_____ slightly helpful
_____ not at all helpful

15) Did you feel the situational awareness display gave you adequate notice of potential hazards?

_____ yes
_____ no

16) Do you feel that the alert positively or negatively affected your driving performance?

_____ positively affected
_____ slightly positively affected
_____ not affected
_____ slightly negatively affected
_____ negatively affected

17) Do you feel that the augmented alert display (i.e., the one that blinked on and off after the turn signal was turned on) helped attract your attention?

_____ no
_____ yes,

18) What effect, if any, did this augmented alert display (turned on after the turn signal was turned on) have on your driving performance?

_____ positive effect
_____ slight positive effect
_____ no effect
_____ slight negative effect
_____ negative effect

19) Any other comments? ____
____
____
____
____
____
____
____

END OF PART II

PART III – At the end of the entire test

Date:____________________________ Subject Number:______________________

1) Which of the display locations did you prefer?

___ Center mirror only
___ Side mirror only
___ Center mirror and side mirror
___ Like all about the same

2) Which mode of operation did you prefer?

___ Monitor mode (i.e., alerts came on whenever an object was in sensor coverage zone; alerts blinked if turn signal on and object was in sensor coverage zone)
___ Turn-signal activated mode (i.e., no alerts presented unless the turn signal was in use)
___ Like both about the same

3) Which of the two systems did you prefer?

___ Proximity system (looked for objects in the immediate area only)
___ Rearward looking plus proximity Comprehensive System (looked backward to objects approaching from rear as well as objects in the immediate area)

Recommendations:____

4) When using the comprehensive system, did you feel that the warnings came:

_____ too early
_____ too late
_____ about right

5) What would you recommend as the "best" combinations of lane change crash avoidance system "features" for further testing?

Location:_______________
Mode:_________________
System:________________

6) Do you have any further recommendations or observations you would like to share with us?
(use additional paper)

This page intentionally left blank

Appendix C – Driver Information Letter

INFORMATION LETTER -
LANE CHANGE COLLISION AVOIDANCE STUDY

Dear Driver,

TRW and the National Highway Traffic Safety Administration are conducting a study of a new lane change collision avoidance warning system. We are examining the impact of these devices on driving safety, comfort, and convenience. We believe this is important research that will contribute to enhancing automobile safety and comfort, but want to ensure that these devices are designed with the drivers' needs and habits in mind.

You have been asked to participate in this study, and as such you will be driving a car on local streets and freeways that is equipped with this new form of lane change Collision Avoidance System (CAS). The CAS has a number of different modes of operation, but also has some features that are always present. It will always warn the driver when a vehicle pulls alongside the car. This "proximity" zone extends along the right side of the car, 12 feet from the side, up to the front bumper and approximately 30 feet behind the rear bumper. The system will not turn on until the vehicle reaches 10 mph and will also not warn against stationary objects such as parked cars and poles. Among the conditions to be varied will be the ability to warn against fast approaching cars in the right adjacent lane. The warning distance will vary with the speed of the approaching vehicle. Also variable will be the type of display this is lit, as well as whether the display will be activated by the turn signal.

You will also be asked to complete some questionnaires. Part I collects general information as it is completed before the start of testing. Part II collects your response to the system configuration just tested and is completed at the end of each test run. Part III is completed at the end of testing and asks for your general reactions.

During testing an experimenter will always be present in the car with you. The experimenter will instruct you where to drive, and establish the conditions under which you will be driving that day. Typically this involves variations in the display modes and/or the functionality of the warning system. The experimenter will also be present to answer any questions you many have during the course of the study. As you drive, the experimenter will operate a computer that records specific information about how the car is being operated. In addition, video cameras will be used to record images of the road and other traffic near the experimental car. No pictures of you will be included in the video.

At no time during this study will you be asked to perform any unsafe driving actions. You are required to obey all traffic laws during this study. You must possess a valid, unrestricted, driver's license. You must have a minimum of two years driving experience. You may not be under the influence of alcohol, drugs, or any other substances which impair your ability to drive (you are asked to refrain from the use of alcohol, drugs, or substances which impair their ability to drive for a period of no less than 12 hours prior to participating).

RISKS: While participating in this study, you will be subject to all the risks that are normally present when driving a passenger car on streets and freeways. You will be required to wear an eye-tracking device, which determines at all times the direction in which you are looking. This device is basically a baseball cap with an infrared illuminator and a camera mounted on the brim and a transparent piece of glass suspended in front of your left eye. Experience has shown that people adapt to this device in a fairly short amount of time. If this makes you uncomfortable then you may stop the testing at any time. Use of the lane change collision warning device being studied should not make driving any more hazardous than normal. It is however an experimental device, and the absence of a warning should not be construed as green light to change lanes. In addition caution should be used when operating a vehicle with which you are not familiar.

Be aware that accidents can happen at any time when driving, and that you can not rely on any device being studied to prevent an accident. You remain responsible for your driving during this testing. Nothing in this study should affect the manner in which you accomplish normal driving tasks. **It is very important to always remember that you, as the driver, are in control of the vehicle and are fully responsible for driving safely at all times.** In the unlikely event that an accident occurred, you, the experimenter, the test vehicle, as well as any other persons or property involved, would be covered under an insurance policy held by TRW.

BENEFITS: The results of this study will provide valuable guidance for the development of lane change warning systems for passenger cars. By participating in this study, you will be lending your experience and expertise to support highway safety research.

PAYMENT: You will be given a job number for your time spent in this study. The study requires that you drive the car a total of eight times over a period of approximately two weeks and for about 1.5 to 2 hours duration each time.

CONFIDENTIALITY: TRW and the National Highway Traffic Safety Administration are gathering information on the use of lane change collision warning devices in passenger cars. We are not testing you or your skills. If you agree to participate in this study, your name will not be voluntarily released to anyone who does not work on this project. Your name will not appear in any reports or papers written about the project.

TRW and the National Highway Traffic Safety Administration hope that you will agree to participate in this study. If you have any questions, please feel free at any time to ask the experimenter.

Once you have had your questions answered, please let the experimenter know whether you are interested in participating in this study. If you are willing to participate, the experimenter will ask you some questions to ensure that your skills and experience match our research needs. If it is determined that you qualify to participate, you will be asked to read and sign an Informed Consent Form before you can actually participate in the study.

This page intentionally left blank

This page intentionally left blank

DOT HS 809 414
August 2000

U.S. Department
of Transportation

**National Highway
Traffic Safety
Administration**

NHTSA
People Saving People
www.nhtsa.dot.gov