

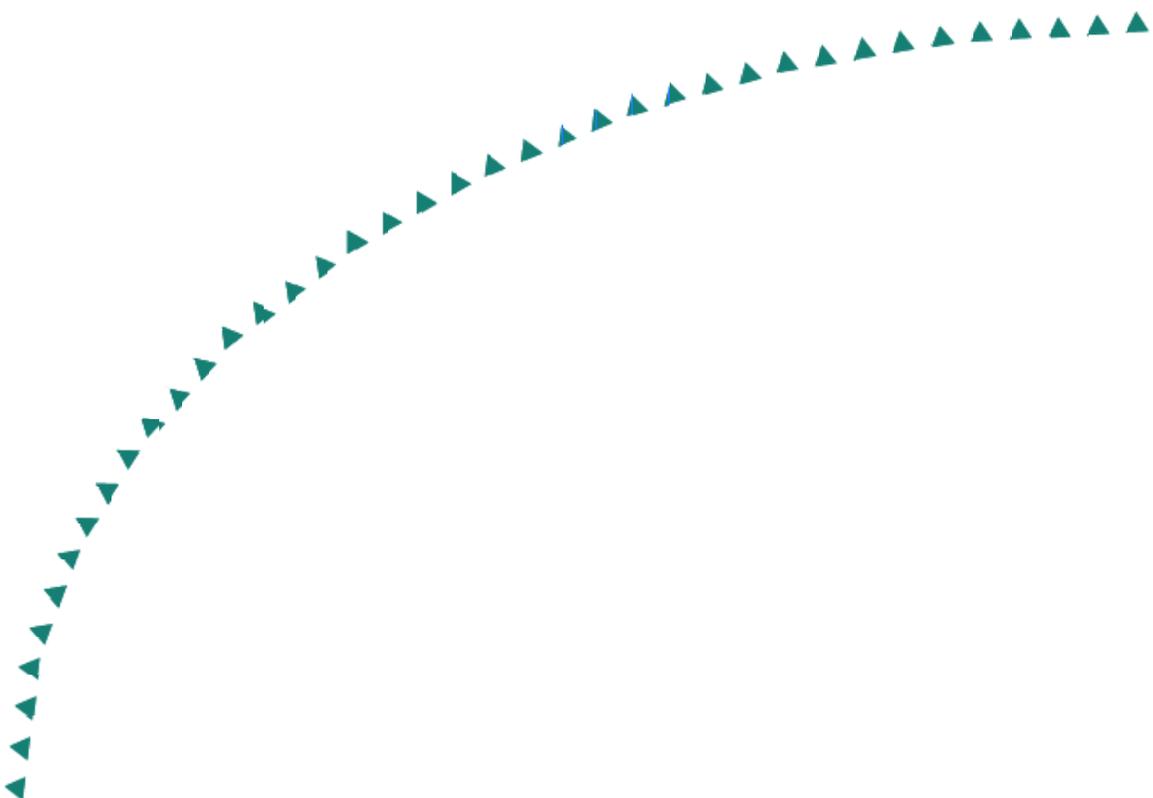
2003-13

Final Report

Part of the Intelligent Vehicle Initiative
**Driver Assistive Systems for
Snowplows**



Research



Technical Report Documentation Page

1. Report No. MN/RC – 2003-13	2.	3. Recipients Accession No.	
4. Title and Subtitle Driver Assistive Systems for Snowplows		5. Report Date March 2003	
		6.	
7. Author(s) Alec Gorjestani, Lee Alexander, Bryan Newstrom, Pi-Ming Cheng, Mike Sergi, Craig Shankwitz, and Max Donath		8. Performing Organization Report No.	
9. Performing Organization Name and Address University of Minnesota Department of Mechanical Engineering ITS Institute Intelligent Vehicles Lab 111 Church St. SE Minneapolis, MN 55455		10. Project/Task/Work Unit No.	
		11. Contract (C) or Grant (G) No. (c) 74708 (wo) 142	
12. Sponsoring Organization Name and Address Minnesota Department of Transportation Office of Research Services 395 John Ireland Boulevard Mail Stop 330 St. Paul, Minnesota 55155		13. Type of Report and Period Covered Final Report 2002	
		14. Sponsoring Agency Code	
15. Supplementary Notes Part of the Intelligent Vehicle Initiative. http://www.lrrb.gen.mn.us/PDF/200313.pdf			
16. Abstract (Limit: 200 words) A comprehensive driver assistive system which utilizes dual frequency, carrier phase real time kinematic (RTK) differential global positioning system (DGPS), high accuracy digital geospatial databases, advanced automotive radar, and a driver interface with visual, haptic, and audible components has been used to assist specialty vehicle operators perform their tasks under these low visibility conditions. The system is able to provide a driver with high fidelity representations of the local geospatial landscape through a custom designed Head Up Display (HUD). Lane boundaries, turn lanes, intersections, mailboxes, and other elements of the geospatial landscape, including those sensed by automotive radar, are projected onto the HUD in the proper perspective. This allows a driver to safely guide his or her vehicle in low to zero visibility conditions in a desired lane while avoiding collisions. Four areas of research, are described herein: driver assistive displays, the integration of a geospatial database for improved radar processing, snowplow dynamics for slippery conditions, and a virtual bumper based collision avoidance/gang plowing system. (Gang plowing is the “flying in formation” of snowplows as a means to rapidly clear multilane roads.) Results from this research have vastly improved the performance and reliability of the driver assistive system.			
17. Document Analysis/Descriptors Head Up Display HUD Collision Avoidance GPS Differential Global Positioning System		18. Availability Statement No restrictions. Document available from: National Technical Information Services, Springfield, Virginia 22161	
19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 72	22. Price

Part of the Intelligent Vehicle Initiative Driver Assistive Systems for Snowplows

Final Report

Prepared by

Alec Gorjestani, Lee Alexander, Bryan Newstrom, Pi-Ming Cheng, Mike Sergi, Craig Shankwitz,
and Max Donath

Intelligent Vehicles Lab
Department of Mechanical Engineering
ITS Institute
University of Minnesota
111 Church St. SE
Minneapolis, MN 55455

March 2003

Published by

Minnesota Department of Transportation
Office of Research Services
395 John Ireland Boulevard, Mail Stop 330
St. Paul, MN 55155

This report represents the results of research conducted by the authors and does not necessarily represent the view or policy of the Minnesota Department of Transportation and/or the Center for Transportation Studies. This report does not contain a standard or specified technique.

The authors and the Minnesota Department of Transportation and/or Center for Transportation Studies do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to this report.

Acknowledgements

The authors would first like to thank John Scharffbillig for his efforts in coordinating our work with Mn/DOT personnel in the field and for keeping the focus of the project on the end user – the snowplow operator. We would also like to thank John Hanzlik and the Maple Grove truck station for their support and maintenance of the SAFELOW research vehicle. Third, as much as our vehicle development work took place at the Mn/ROAD pavement research facility, the authors would like to thank Jack Herndon and the staff at Mn/ROAD for providing access to his facility and help with any problem that arose. Finally, thanks to the staff at the Minnesota Highway Safety Center (MHSC) located at St. Cloud State University in St. Cloud, Minnesota for their help and assistance in the vehicle dynamics testing undertaken at their facility.

TABLE OF CONTENTS

Chapter 1 Introduction	1
Background	1
Work Undertaken and Results	2
<i>Task 1. Driver Assistive Display (DAD)</i>	2
<i>Task 2. Integration of Radar and High Accuracy Geographic Databases</i>	2
<i>Task 3. Snowplow Dynamics Modeling for Slippery Conditions</i>	2
<i>Task 4. Virtual Bumper Based Collision Avoidance – Gang Plowing</i>	3
Chapter 2 Driver Assistive Displays	5
Background	5
Task 1 Results	7
<i>a. Eye Alignment</i>	7
<i>b. HUD Projection Accuracy</i>	8
<i>c. Integration of Radar Information into the HUD</i>	11
<i>d. Information Presented by the HUD</i>	16
<i>e. Preview Amount</i>	16
Chapter 3 Integration Of Radar And High Accuracy Geospatial Databases	17
Background	17
Geospatial Database	18
Range Sensor Processor	19
Experiments.....	20
<i>Database Experiments</i>	20
<i>Range Sensor Filter Experiments</i>	23
Conclusions	25
Chapter 4 Snowplow Dynamics Modeling for Slippery Conditions.....	27
Introduction	27
Kinematic Model.....	29
Icy Surfaces	32
<i>Understeering Slides on Ice</i>	36
<i>Oversteering Slides on Ice</i>	39
Applications.....	39
Conclusions	41
Chapter 5 Virtual Bumper Based Collision Avoidance Applied to Gang Plowing.....	43
Background	43
Overview	43
<i>Longitudinal Controller</i>	45
<i>Speed Controller</i>	47
<i>Lateral Controller</i>	50
Results	52
<i>Lateral Controller Results</i>	54
<i>Longitudinal Controller Results</i>	56
Conclusions	58
Future Work	59

List of Figures

Figure 2-1: Head Up Display mounted in the SAFELOW. The projector is positioned over the driver's right shoulder with the combiner positioned in the driver's forward view.....	5
Figure 2-2: HUD display under bright conditions. Note the lack of glare through the combiner.....	7
Figure 2-3: Reference marks for error analysis. CF (center far) and CN (center near) are points on the pavement lane marking used for analysis.....	9
Figure 2-4: Driving trajectory. The lateral dimension in the graph is amplified to readily show the trajectory.	10
Figure 2-5: HUD error at four different look-head locations (vehicle speed is 40 kph). .	11
Figure 2-6: Data flow diagram for transforming object location relative to vehicle coordinate frame to global coordinate frame.	12
Figure 2-7: Radar, vehicle, and state plane coordinates used to transform range sensor data to global coordinates.....	13
Figure 2-8: Image captured at night behind HUD. Obstacles detected via the radar are represented as clear, square icons. The icon remains white if the time to collision is greater than 3 seconds or if the detected object is more than 50 ft. from the front of the host vehicle. Misalignment of the projected lane boundaries in the front of the image is due to camera misalignment and the difficulties of capturing an image on a rough road with a handheld camera.	14
Figure 2-9: Image captured at night behind HUD. In this shot, the vehicle closest to the SAFELOW is closer than 50 ft., so the icon representing its location has changed color to red. Misalignment of the projected lane boundaries in the front of the image is due to camera misalignment and the difficulties of capturing an image on a rough road with a handheld camera.....	15
Figure 3-1: Host vehicle entering a curve with a road sign directly in the radar's field of view.....	18
Figure 3-2: Illustration of two targets (1 and 2) that were filtered and one target (3) that passed through the filter for further processing.	20
Figure 3-3: Radar target filter query times. For each query update interval the query polygon is increased in size to account for vehicle motion between query updates. Each point corresponds to the mean or maximum query time of 8,385 queries performed along the length of the GDB corresponding to actual vehicle positions. .	22
Figure 3-4: The SAFELOW with DGPS and two Eaton VORAD radar mounted on the front grille and top of the cab.....	23
Figure 3-5: SAFELOW driving southbound on Highway 101. False target filtered off of the right shoulder.....	24
Figure 3-6: SAFELOW driving southbound through an intersection on Highway 101. False targets at road islands were filtered. Trajectories of two vehicles sensed by the onboard radar are drawn in the south bound lane.	25
Figure 4-1: Aerial view of the Minnesota Highway Safety Center. Skid pad on which vehicle dynamic experiments took place is indicated.	28
Figure 4-2: Test data collected from a test run at the MHSC using the SAFELOW on dry pavement. Vehicle speed is approximately 18 miles/hour.....	30

Figure 4-3: Comparison of the computed and measured yaw rates corresponding to the data illustrated in Figure 4-2.	31
Figure 4-4: The difference (residuals) between the computed and the measured yaw rate for four different values of yaw rate lag. The residual is measured in degrees/second.	33
Figure 4-5: A comparison of the measured yaw rate with the output of a kinematic model, including 3% steering compliance and a 0.2 second lag.....	34
Figure 4-6: Test data collected from a test run using the SAFEPLow on an ice covered skid pad. The vehicle speed is approximately 11 miles/hour.....	35
Figure 4-7: A comparison of the measured yaw rate with the output of the kinematic model on an ice covered surface during a gentle weaving maneuver.....	36
Figure 4-8: Test data collected from a test run using the SAFEPLow on an ice covered skid pad at 12 m/s (27 mph) with a steering input strong enough to induce sliding..	37
Figure 4-9: Comparison of predicted and actual yaw rates on ice with induced understeering slides.....	38
Figure 4-10: A plot of the difference between predicted and actual yaw rates makes it evident that sliding is occurring between 25 and 45 seconds.	38
Figure 4-11: This graph shows a series of understeering and oversteering slides induced by strong steering movement on an icy surface. Understeering is indicated when the magnitude of the delayed kinematic yaw rate is greater than the magnitude of the measured yaw rate. Oversteering is indicated when the magnitude of the delayed kinematic yaw rate is less than the magnitude of the measured yaw rate.....	39
Figure 4-12: Snowplow on curve. The dashed line indicates vehicle longitudinal axis. In this situation, to mitigate understeer, the brakes on the left side of the vehicle would be applied. To mitigate oversteer, the right side brakes would be applied.	41
Figure 5-1: The snowplow operator of the following vehicle selects the desired lateral offset and headway distance from the lead vehicle.....	44
Figure 5-2: Gang plow system diagram. The trailing vehicle receives the differential GPS correction from the base station and sends it to the lead vehicle. The lead vehicle sends back its DGPS position.....	45
Figure 5-3: The virtual bumper longitudinal controller consists of a virtual spring with spring stiffness k and a virtual damper with damping ratio b	46
Figure 5-4: Speed controller system diagram.....	47
Figure 5-5: Feed forward speed map showing voltage vs. speed.....	48
Figure 5-6: Brake motor attached to firewall with winch and cable.....	49
Figure 5-7: Cable runs through firewall and attaches to brake pedal.....	50
Figure 5-8: The correction term (δ) is obtained from the lateral error (y), the radius of curvature of the road (r), and the wheel base (W_b).....	51
Figure 5-9: Steering motor attached to steering column.....	52
Figure 5-10:Block diagram of lead vehicle	53
Figure 5-11: Block diagram of trailing vehicle.....	53
Figure 5-12: Lateral position plot for a straight section of road. Initiation of lane change maneuvers is indicated by the spikes in the lateral position error plot in the lower graph.....	55
Figure 5-13: Curved portion of Mn/Road pavement research facility.....	56
Figure 5-14: Lateral position plot for a curved section of road.....	56

Figure 5-15: Longitudinal tracking – speed.....	57
Figure 5-16: Longitudinal tracking - headway.	58

Executive Summary

Poor visibility conditions lead to a significant number of crashes each year. The inability to see road boundaries, obstacles, and other vehicles on the road is of special concern during the winter. In particular, snowplow drivers are affected, because they must drive even when visibility is poor, roads are snow covered and lanes are hard to discern. If snowplows are unable to do their jobs, important safety and economic repercussions may arise. Consequently, focusing attention on technologies that help the snowplow operator benefits the general population, whether through keeping the roads open to the movement of people or freight, or by facilitating the passage of emergency vehicles.

At the University of Minnesota (U of MN), a comprehensive driver assistive system which highly accurate differential GPS (global positioning system), high accuracy digital geospatial databases, advanced automotive radar, and a driver interface with visual, haptic, and audible components has been used to assist specialty vehicle operators perform their tasks under low visibility conditions. The differential GPS used in this system is a dual frequency, carrier phase real time kinetic (RTK) system. This driver assistive system is the subject of this report.

The system is able to provide a driver with high fidelity representations of the local geospatial landscape through a custom designed Head Up Display (HUD). Lane boundaries, turn lanes, intersections, mailboxes, and other elements of the geospatial landscape, including those sensed by automotive radar, are projected onto the HUD in the proper perspective. This allows a driver in low to zero visibility conditions to safely guide his or her vehicle in a desired lane while avoiding collisions.

Four areas of research were addressed in this project: driver assistive displays, integration of a geospatial database for improved radar processing, snowplow dynamics for slippery conditions, and virtual bumper based collision avoidance/gang plowing. (Gang plowing is the “flying in formation” of snowplows used as a means of rapidly clearing multilane roads.)

Each of these tasks, and their respective results, are described below.

Driver Assistive Displays. The objective of this task was to develop a visual based high fidelity system able to provide an accurate visual representation of the roadway and surroundings local to the vehicle. In this task, DGPS data, a high accuracy geospatial database, radar information, and a HUD were integrated. The result is a “see-thru” highly accurate HUD based image displaying (with the proper perspective) road boundaries, relevant geographic features (signs, bridges, end treatments, etc.), and obstacles impeding the motion of the snowplow. Obstacle information provided by vehicle mounted radar is also presented to the driver via the HUD.

This task was completed, resulting in a new HUD hardware and graphical software. The HUD hardware and software developed as part of this task have been used in six

snowplows, an ambulance, a state patrol car, and a transit bus. In addition to these vehicle applications, one HUD has also been sold for a driving simulator application.

Integration of Radar and High Accuracy Geographic Databases. The purpose of this task is to improve the performance of vehicle based radar by integrating a geospatial database into the signal processing task. This integration will allow the radar processor to distinguish radar returns from known elements of the geospatial landscape from unknown moving or fixed obstacles. This simplifies the human interface by “filtering” or rejecting from the display known geospatial objects for which a collision is highly unlikely.

This task has been completed and has been shown to greatly reduce the false positives which arise from radar returns from fixed geospatial elements. This improved signal processing increases the driver confidence associated with the driver assistive system, thereby decreasing the stress associated with driving in low to zero visibility.

Snowplow Dynamics Modeling for Slippery Conditions. The objective of this task is to determine in real time the dynamic capability of a snowplow under changing weather and road conditions. Real time knowledge of the dynamic capability of the snowplow can be used to optimize the collision avoidance algorithms. This optimization will allow the collision avoidance maneuvers to be executed with a minimum probability that the maneuver causes an undesired effect such as a spin or rollover.

This task was completed, and the results were used in the gang plowing research described below.

Virtual Bumper based Collision Avoidance – Gang Plowing. The virtual bumper is a collision avoidance technique that integrates radar, a geospatial database, and throttle, braking, and steering maneuvers. The objective of this task was to develop an interface to the snowplow operator through the steering wheel to assist the operator with the collision avoidance task. The result is a torque applied to the steering wheel that indicates the computed optimal collision avoidance path to the driver.

“Gang” snowplowing using the virtual bumper was implemented. Under gang plowing conditions, a lead snowplow clears one side of the road. A second plow follows closely behind and to the side of the lead plow, moving the snow from the first plow, with that in the second plow’s present lane, to the side of the road or lane. Gangs can be extended *ad infinitum*.

Because of the localized whiteout conditions caused by the lead plow, following plows have difficulty maintaining visual contact and close physical distance with the lead plow.

To overcome the difficulties associated with gang plowing, a virtual bumper based gang plowing system that uses DGPS and inter-vehicle communications was developed and tested. With this approach, the lead plow is equipped with a DGPS receiver and a wireless transceiver. The position of the lead plow is broadcast to the following plow. The driver of the following plow specifies desired lateral and longitudinal offsets from the lead plow. The virtual bumper algorithm uses the lead vehicle position information to

determine the steering, throttle, and brake actuation needed to maintain the lateral and longitudinal offsets specified by the following vehicle operator.

Virtual bumper based gang plowing was demonstrated to Mn/DOT personnel at the Mn/ROAD pavement research facility on 07 July 2002. A video record of its performance can be found on the CD which accompanies this report. This successful demonstration has led to a follow-on project that will test this approach to gang plowing under operational conditions.

Chapter 1 Introduction

Background

In the fall of 1998, the US DOT Intelligent Vehicle Initiative (IVI) Specialty Vehicle pooled fund was established with Minnesota, California, and US DOT as partners. The purpose of this particular pooled fund study was to develop technologies which would aid drivers and maintenance engineers with the task of snow removal in difficult, low visibility conditions. The program was structured so that each partner could focus on technologies appropriate to the environmental and geographical conditions considered most difficult in each locale.

California suffers from extremely wet and heavy snowfalls in their mountainous terrain. Snowfalls of 4 to 5 feet during a single storm are not uncommon. This wet, heavy snow forces snowplows to operate at relatively low speeds. Because of these conditions and the previous experience the California team (CalTrans, University of California at Davis, and the University of California at Berkeley's PATH program) has with magnetic based vehicle guidance, the California team focused on the application of magnetic vehicle guidance to this snow removal problem. In addition to the application of the magnetic technology, the California team developed driver displays designed to convey information provided by the magnetic lane guidance system to the driver.

In contrast to California, issues with snow removal in Minnesota arise because of high winds which can blow light, dry snow at high rates across vast stretches of prairie. This blown snow can create significant drifts which must be continuously cleared in order to keep roads open. Under these conditions, snowplow operations are required to run at relatively high speeds to avoid road closures because of drifting. Because of these weather and geographic conditions, the Minnesota team focused on DGPS based solutions to this particular snow removal problem. The Minnesota team also developed driver displays designed specifically to convey lane keeping and collision avoidance information provided by the on-board systems.

The work undertaken by the Minnesota Team under tasks one and two was implemented and evaluated as part of the FHWA IVI Generation Zero Field Operational Test. Six vehicles were equipped with technology developed under this program: four snowplows, one ambulance, and one state patrol car.

Although the majority of the technical work was carried out independently, it was the intent of the pooled fund partners to jointly develop Methods of Evaluation (MOEs) for emerging snowplow technologies. Unfortunately, contracting problems resulted in the California team completing their study before the Minnesota team was able to begin theirs. This disjointed situation forced a change of plan. Instead of developing MOEs, the Minnesota team requested and was given permission to change the scope of their program. Minnesota increased its efforts under task 4 in the development of a DGPS based gang plowing system. Gang plowing is essentially the operation of snowplows in formation, which improves the efficiency of snow removal operations. The gang plowing

problem addressed in this work is that the snow blown up by the leading plows creates local whiteout conditions for the trailing plows. This forces the trailing plows to follow at a greater distances, providing an opportunity for a “rogue” vehicle to drive between the plows. These rogue vehicles often lead to crashes. A consistently tight formation that does not allow space for rogue vehicles, improves the safety of the snowplow operation.

Work Undertaken and Results

The work performed for this pooled fund consisted of four specific tasks:

Task 1. Driver Assistive Display (DAD)

The objective of this task was to develop a visual based high fidelity system able to provide an accurate visual representation of the roadway and surroundings local to the vehicle. In this task, DGPS data, a high accuracy geospatial database, radar information, and a HUD were integrated. The result was a “see-thru” highly accurate HUD based image, displaying (with the proper perspective) road boundaries, relevant geographic features (signs, bridges, end treatments, etc.), and obstacles impeding the motion of the snowplow. (The University of Minnesota has previously used LCD flat panel technology and has found that HUDs offer better performance.)

This task was completed, and resulted in two published papers [1] [2] and two US Patents pending [3] [4]. Over the course of this project, the design of the HUD evolved continually. What could be considered to be the third generation of this HUD has also been deployed on four Minnesota snowplows, one ambulance, and one state patrol vehicle as part of the FHWA IVI Generation Zero Field Operational test. Two more HUD systems are scheduled to be delivered to two county snowplows during the fall of 2002.

The results of this task are described in Chapter 2.

Task 2. Integration of Radar and High Accuracy Geographic Databases

The objective here was to improve the performance of vehicle based radar by integrating a geospatial database into the signal processing task. This integration allows the radar processor to distinguish between radar returns from known elements of the geospatial landscape and from unknown obstacles. This simplifies the human interface by “filtering”, or rejecting, from the display known geospatial objects for which a collision is highly unlikely.

This task was completed, and the results thereof were published in [5]. This technique is also the subject of a US Patent Pending [3]. The results of this task are described in Chapter 3.

Task 3. Snowplow Dynamics Modeling for Slippery Conditions

The objective of this task is to determine in real time the dynamic capability of a snowplow under changing weather and road conditions. Real time knowledge of the dynamic capability of the snowplow can be used to optimize the collision avoidance

algorithms. This optimization will allow the collision avoidance maneuvers to be executed with a minimum probability that the maneuver causes an undesired effect such as a spin or rollover.

This task was completed, and the results are described in Chapter 4. The dynamic snowplow models are used to implement the virtual bumper applied to gang plowing operations in task 4.

Task 4. Virtual Bumper Based Collision Avoidance – Gang Plowing.

The virtual bumper is a collision avoidance technique which integrates radar, a geospatial database, and throttle, braking, and steering maneuvers. The objective of this task was to develop an interface to the snowplow operator through the steering wheel which will assist the operator with the collision avoidance task. The result is a torque applied to the steering wheel that indicates the computed optimal collision avoidance path to the driver.

Gang snowplowing using the virtual bumper was implemented. Under gang plowing conditions, a lead snowplow clears one side of the road. A second plow follows closely behind and to the side of the lead plow, moving the snow from the first plow with that in the second plow's present lane, to the side of the road or lane. Gangs can be extended *ad infinitum*.

Because of the localized whiteout conditions caused by the lead plow, the following plows have difficulty maintaining visual contact and close physical distance with the lead plow.

To overcome the difficulties associated with gang plowing, a virtual bumper based gang plowing system which uses DGPS and inter-vehicle communications was developed and tested. With this approach, the lead plow is equipped with a DGPS receiver and a wireless transceiver. The position of the lead plow is broadcast to the following plow. The driver of the following plow specifies desired lateral and longitudinal offsets from the lead plow. The virtual bumper algorithm uses the lead vehicle position information to determine the steering, throttle, and brake actuation needed to maintain the lateral and longitudinal offsets specified by the following vehicle operator.

Virtual bumper based gang plowing was demonstrated to Mn/DOT personnel at the Mn/ROAD pavement research facility on 07 July 2002. A video record of its performance can be found on the CD accompanying this report.

Chapter 2 Driver Assistive Displays

Background

At the time the proposal to participate in this pooled fund study was submitted, the Minnesota team had just verified that it was possible to project lane boundaries and other elements of the geospatial landscape onto a HUD using geospatial information provided by a geospatial database and position information provided by a DGPS system. This work was performed at the Mn/DOT Mn/ROAD pavement testing facility and used a small, dim HUD. The HUD used in that research was of sufficient quality to prove a concept, but neither large enough nor bright enough for use in snowplowing operations.

Under a separate Mn/DOT funded project (CUFS 530-6421), the Intelligent Vehicles Lab designed its first large, bright HUD with the performance needed to satisfy the requirements for snowplow operations. Under this contract, the HUD design underwent two required iterations and resulted in a HUD that has been used in snowplows, ambulances, state patrol vehicles, and transit buses. The current generation HUD is shown in Figure 2-1.

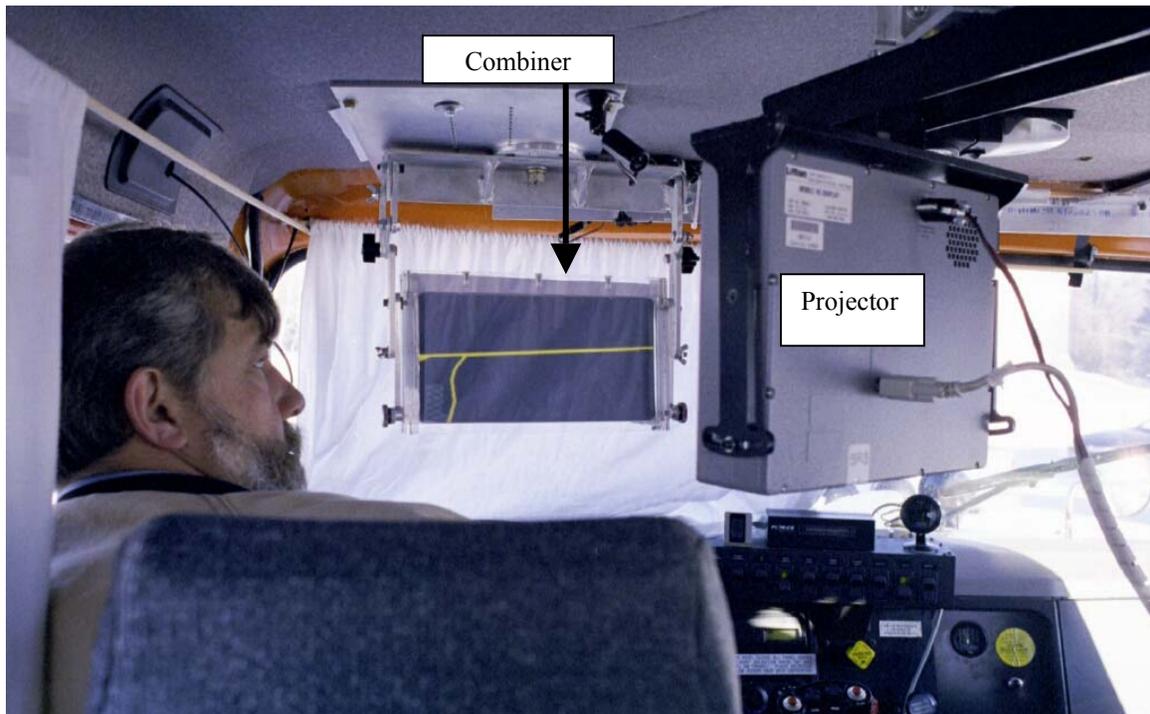


Figure 2-1: Head Up Display mounted in the SAFEFLOW. The projector is positioned over the driver's right shoulder with the combiner positioned in the driver's forward view.

The driver, during low visibility conditions, moves the combiner down to its operational position which is approximately 18 inches in front of the driver's face. The projector then provides the images used by the driver to guide the vehicle during low visibility conditions. The driver sees through the combiner, and also sees reflections from the projector in the combiner.

The combiner (also commonly referred to in the optical world as a beam splitter) is a partially reflective, partially transmissive optical device designed specifically for this purpose. The combiners used in the field operational test are made of chemically tempered optically ground glass, ground and coated to University of Minnesota specifications. (Because of the small number of systems made for this FOT, economics dictate that the combiners be made of glass. In production quantities, the combiners would be made from an injection molded and optically coated piece of acrylic or polycarbonate. These combiners would be much thinner, lighter, and less costly than the ground glass units used in this project.) Chemical tempering strengthens the optical substrate and, in the unlikely event the combiner is broken, allows the broken glass to come apart in smooth pieces. The coating on the concave surface is reflective; the reflective properties are designed such that all colors from the projector are reflected with nearly equal intensity. The level of reflectivity has been optimized for bright to zero ambient light conditions. For extremely bright conditions, (i.e., sunny days with new snowfall), a contrast "enhancer" is overlaid on the combiner frame. This contrast enhancer acts like a pair of sunglasses and reduces the amount of light passing to the driver through the combiner.

An anti-reflective coating covers the convex side of the combiner. This anti-reflective coating improves contrast and eliminates glare. The effectiveness of the anti-reflective coating is best illustrated in Figure 2-2 below. (Note the glare in the windshield, and the lack of glare in the view through the combiner.)

One concern expressed by the California team was that an optical combiner used in a snowplow might be subject to "fogging." This is not a problem for two reasons. First, the entire combiner is inside the vehicle, thereby creating a uniform temperature distribution through the optical glass with the temperature (in steady state conditions) equal to that of the ambient air inside the cab. Without localized "cold spots," there is no impetus for condensation. Second, Mn/DOT snowplows are equipped with Air Conditioning for the purpose of drying air inside the cab during snow events. These two factors result in no combiner fogging.



Figure 2-2: HUD display under bright conditions. Note the lack of glare through the combiner.

Task 1 Results

As proposed, task 1 was to address the following issues:

- a. Determine the alignment of the image with respect to the driver's eyes to assure the proper perspective of the projected image.
- b. Determine the accuracy with which images can be rendered.
- c. Determine how to integrate obstacles sensed by the radar into the HUD.
- d. Investigate those elements of the geospatial landscape that need to be included in the HUD's view.
- e. Determine the amount of preview needed to give the snowplow operator a reasonable comfort level while navigating.

These are addressed below.

a. Eye Alignment

When the proposal was written, a low quality HUD with a small combiner was used for proof of concept work. With that HUD, eye position was critical in order to maintain an accurate projection.

With the redesign of both the projection and combiner components, sensitivity to eye position has been significantly reduced (because of issues involving the pending patents, specific details of the HUD design are not provided). Because of the improvements

resulting from this redesign, the only adjustment the driver needs to make is that of the pitch angle of the combiner. Adjusting the pitch compensates for different driver heights.

b. HUD Projection Accuracy

A series of experiments were performed to document the accuracy of the image projected by the HUD. These experiments and their results are provided below.

Experimental Setup. The SAFEFLOW (an International Model 2540 four door, single axle snowplow used for advanced vehicle research since 1999) was used as the test bed for the experiments. The heading angle of the SAFEFLOW was estimated based only on the trajectory of recent DGPS position values. A simple difference method that calculates a vector angle based on the current and the past positions of the vehicle was used to estimate vehicle heading. To reduce the effect of GPS error, a 5 sample moving average was used for the calculation.

All HUD experiments were performed at Mn/ROAD, a facility maintained by the Minnesota Department of Transportation (Mn/DOT) to test road construction materials and methods. The road is isolated from public traffic, allowing a driver to avoid interactions with other vehicles.

Live video images were captured while driving the SAFEFLOW on a straight section of the low volume test road at Mn/ROAD. Position data coming from DGPS was synchronized with video data and simultaneously recorded. A Canon Optura digital video (640 x 480 resolution, 30 frames per second) camcorder was used to record the projected HUD screens during actual driving situations. The camcorder was mounted at the passenger's right eye position using custom mounting brackets. The optical image stabilizer in the camcorder was enabled while taking video images. The digitally stored images were transferred to a computer, then post processed and analyzed. To synchronize the beginning of the video image stream, a calibration mark was put on the video screen by the HUD software when recording was started.

Error analysis was done by measuring the mismatch between the real centerline and HUD centerline at 2-second intervals. In Figure 2-3, the four linear segments perpendicular to (and on both sides of) the centerline are reference marks used to measure lateral mismatch error. Each horizontal mark is 0.5m long, and the gap between the horizontal reference marks and the projected centerline is 0.5m.

Errors associated with the projected lane boundaries are computed by comparing the "distance" the projected lines are displaced from the actual lane boundary. The lateral displacement of the projected lane boundary from the actual lane boundary at any of the four marks is computed knowing that the length of the reference mark is 0.5m.

The objective of the conformal HUD is to construct and project road characteristics onto the display screen (or combiner) that precisely matches the physical road characteristics.

To quantify how well this is done, VSA is used to describe the mismatch error. The visual sight angle is defined by the ratio of the actual lateral error associated with the lane projection and distance to the eye point as shown below:

$$vsa = \frac{\text{lateral error at distance } x}{\text{distance } x}$$

The visual sight angle normalizes error along the depth of the visual field and captures the error that is perceived by the driver. All our experimental results are described in terms of this visual sight angle based error.

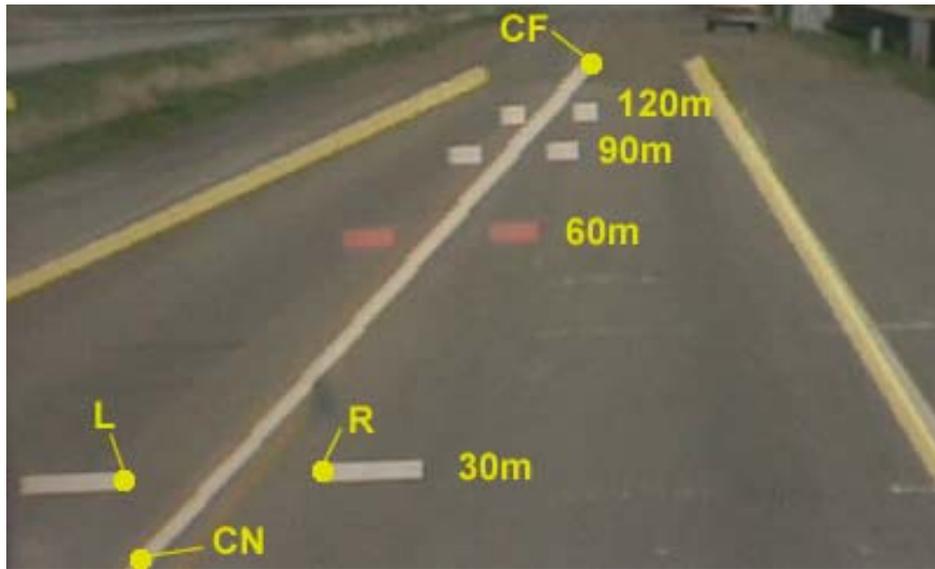


Figure 2-3: Reference marks for error analysis. CF (center far) and CN (center near) are points on the pavement lane marking used for analysis.

The VSA error was measured at four different “look-ahead” distances: 30 m (98.4 ft.), 60 m (196.8 ft.), 90 m (295.3 ft.), and 120 m (393.7 ft.), as measured from the driver’s eye. The topmost horizontal grid mark, i.e., the furthest one in Figure 2-3, is 120 m ahead.

To measure the effect that the heading angle estimation error had on the accuracy of the projection of lane boundaries onto the HUD, a series of experiments was performed. In these experiments, the test truck was driven as shown in Figure 2-4. After full acceleration, the vehicle was driven to the left of the road centerline, and then driven back into the normal (right) driving lane.

Experimental Results. In Figure 2-5, errors for four different positions in front of the driver (30 m, 60 m, 90 m, and 120 m ahead; see Figure 2-3) are drawn.

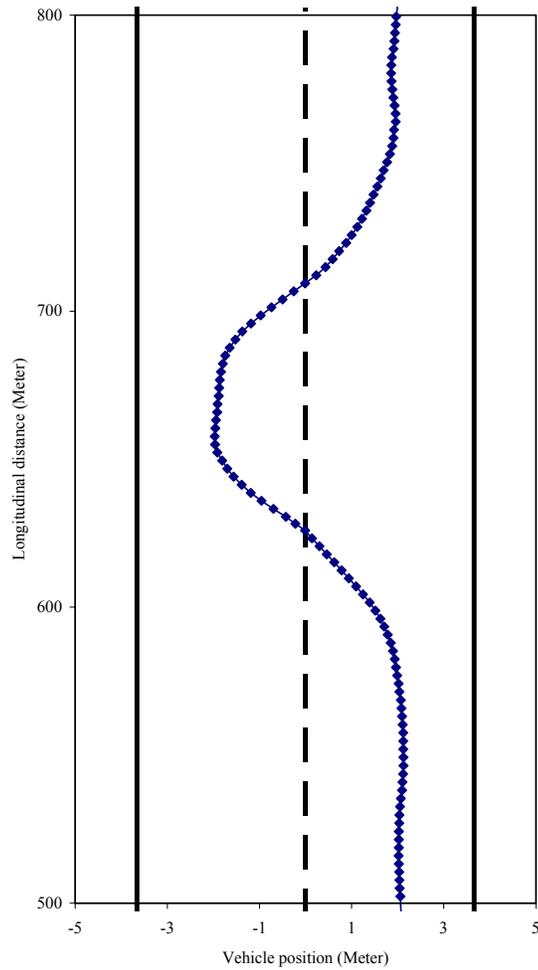


Figure 2-4: Driving trajectory. The lateral dimension in the graph is amplified to readily show the trajectory.

The HUD system was able to estimate the heading angle to within 0.5-degrees even when the vehicle changed lanes.

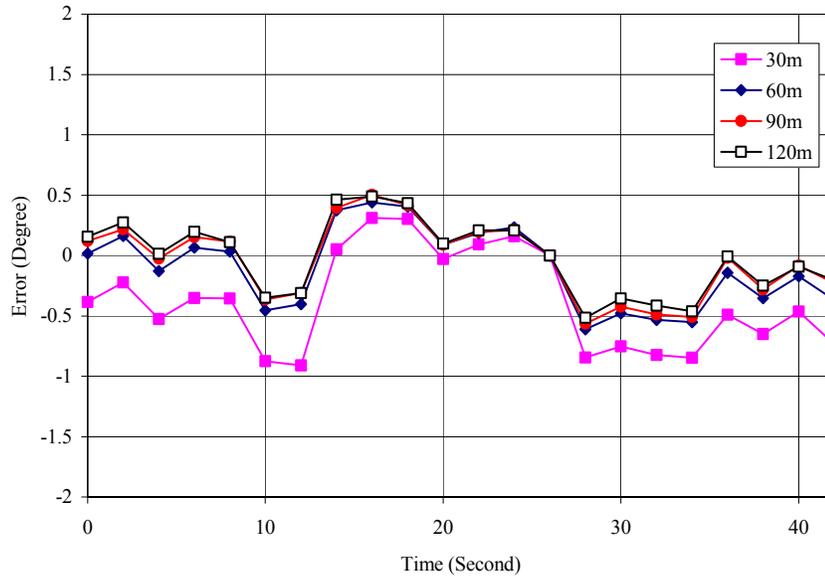


Figure 2-5: HUD error at four different look-ahead locations (vehicle speed is 40 kph).

A large visual field (17.2 degrees horizontal and 15.5 degrees vertical) ensures that the lane markings were within the HUD’s field of view. In Figure 2-5, the HUD vsa error was always less than 0.5 degrees for 60 m, 90 m, and 120 m ahead. The significantly larger errors at 30 m ahead resulted from a misalignment during the calibration performed before the experiment. This can and has been corrected.

c. Integration of Radar Information into the HUD

With the advent of the Eaton Vorad EVT-300 radar system, the projection of the *location* of obstacles relative to the host vehicle became a possibility. The EVT-300 radar provides not only range and range rate information to targets, but *azimuth* information as well. Previous generations of vehicle radar did not provide azimuth information, minimizing their utility as a source of information for a conformal driver display.

With this azimuth information, not only is the presence of an object available, but its location is as well. Given the location relative to the host vehicle, it is a straightforward process to present the location of an obstacle into the HUD. The algorithm dataflow diagram is presented in Figure 2-6 below:

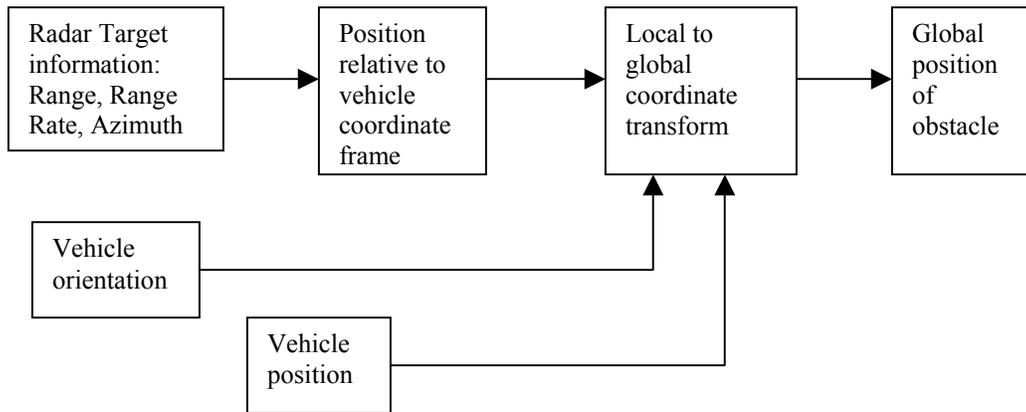


Figure 2-6: Data flow diagram for transforming object location relative to vehicle coordinate frame to global coordinate frame.

The transformation outlined in Figure 2-6 shows as an output the location of that obstacle in the global coordinate frame. Once the location of the obstacle is known in the global coordinate frame, the same software used to draw virtual “lines” on the “road” can be adapted to draw representations of obstacles in the proper virtual global location.

Range sensors present the detected object location in a local coordinate system whose origin is located on the sensor's transmitting surface. Most range sensors use a polar coordinate system with a distance (R) and azimuth angle (ϕ) to each target. The first procedure in the range sensor processor is to transform polar sensor measurements into a local Cartesian coordinate system. This coordinate system is the sensor coordinate system ($X_{\text{sensor}}, Y_{\text{sensor}}$) shown in Figure 2-7 and described by

$$\begin{aligned} x_1 &= R \cos(\phi) \\ y_1 &= R \sin(\phi) \end{aligned} \tag{2.1}$$

Each target location is transformed using a translation matrix into the vehicle coordinate frame (X_v, Y_v). The translation distance is simply the distance between the radar and GPS antenna with respect to the vehicle coordinate frame. Thus,

$${}^v \mathbf{x} = {}^{\text{sensor}} \mathbf{T}_v^{-1} {}^{\text{sensor}} \mathbf{x} \tag{2.2}$$

where

${}^{\text{sensor}}\mathbf{x} = (x_1, y_1)$ in the sensor coordinate frame,

${}^{\text{v}}\mathbf{x} = (x_1, y_1)$ in the vehicle coordinate frame, and

${}^{\text{sensor}}\mathbf{T}_v^{-1}$ = Translation matrix from sensor to GPS coordinate frame.

Similarly, the target locations are transformed from the vehicle coordinate system to the global state plane coordinate system. This is done by a transformation matrix consisting of a rotation by the vehicle's heading angle and a translation of the vehicle GPS location.

$${}^{\text{state}}\mathbf{x} = {}^{\text{v}}\mathbf{T}_{\text{state}}^{-1} {}^{\text{v}}\mathbf{x}, \quad (2.3)$$

where

${}^{\text{state}}\mathbf{x} = (x_1, y_1)$ in the state plane coordinate frame,

${}^{\text{v}}\mathbf{T}_{\text{state}}^{-1}$ = transformation matrix consisting of a rotation (heading, z) and translation($x_{\text{gps}}, y_{\text{gps}}$).

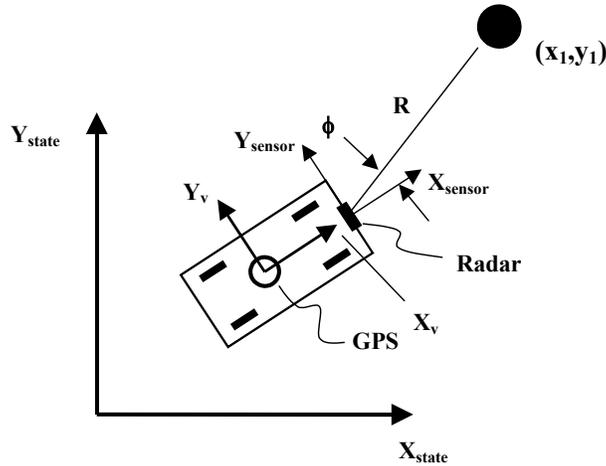


Figure 2-7: Radar, vehicle, and state plane coordinates used to transform range sensor data to global coordinates.

Figure 2-8 shows the system providing a driver advisories by indicating the presence and location of obstacles that are a potential threat to the driver of the host vehicle.

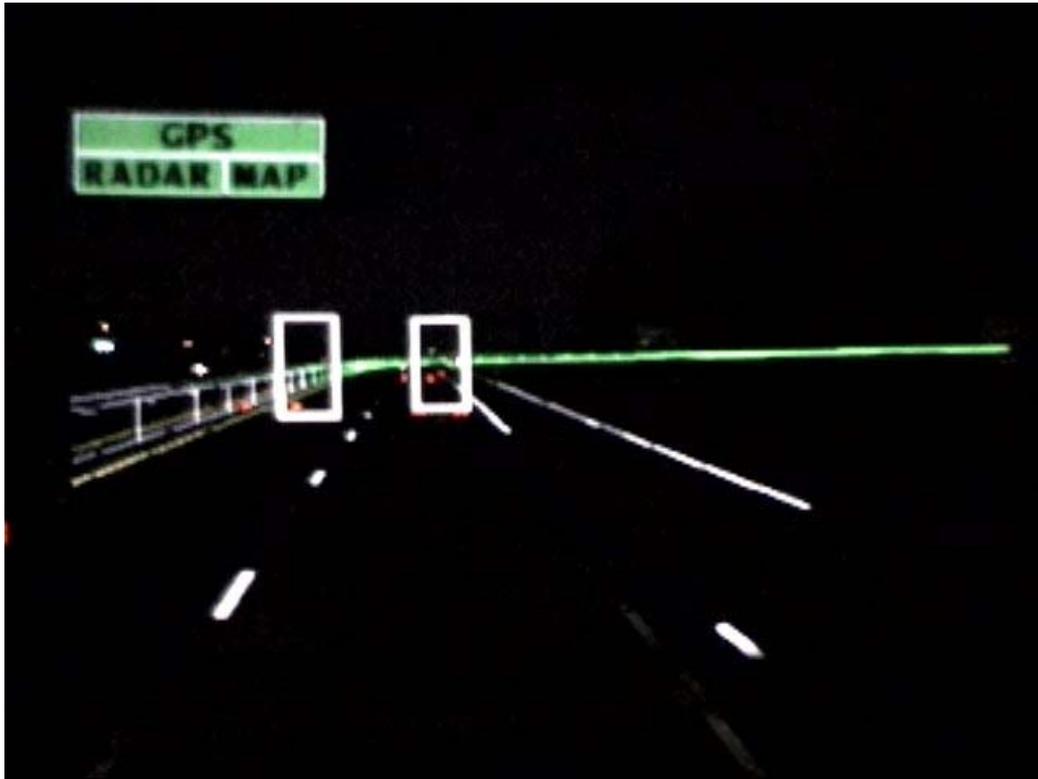


Figure 2-8: Image captured at night behind HUD. Obstacles detected via the radar are represented as clear, square icons. The icon remains white if the time to collision is greater than 3 seconds or if the detected object is more than 50 ft. from the front of the host vehicle. Misalignment of the projected lane boundaries in the front of the image is due to camera misalignment and the difficulties of capturing an image on a rough road with a handheld camera.

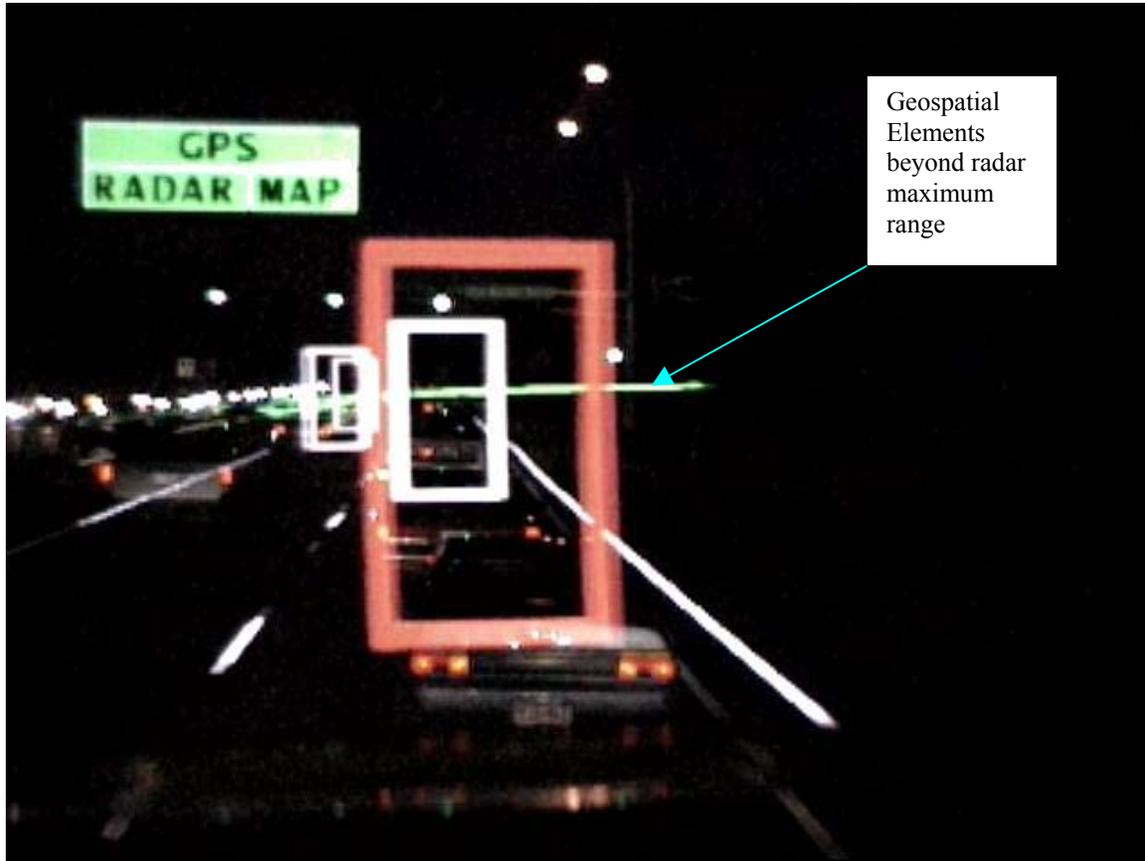


Figure 2-9: Image captured at night behind HUD. In this shot, the vehicle closest to the SAFEFLOW is closer than 50 ft., so the icon representing its location has changed color to red. Misalignment of the projected lane boundaries in the front of the image is due to camera misalignment and the difficulties of capturing an image on a rough road with a handheld camera.

In Figure 2-9, the vehicle closest to the SAFEFLOW is closer than 50 ft., so an icon representing its location has changed color to red.

Two basic representations of targets were used throughout the project. The first icon representing a target is shown in Figure 2-2. The rationale behind this icon was that it represents a silhouette of the rear of a modern, aerodynamic vehicle. However, the icon could potentially block the view of the obstacle itself under conditions of reasonable visibility. Therefore, the icon was changed to the rectangular shape so as to provide a better sight path to the obstacle ahead. The width represents a standard passenger sedan, which is the vehicle most likely to be encountered (and least likely to be seen) by a snowplow in low visibility conditions.

Snowplow operators clearly preferred the open icon as it provided a reasonable opportunity for them to see what was detected by the radar when local visibility conditions improved.

d. Information Presented by the HUD

This question became an issue of personal preference. Some operators preferred to see just the immediate two or three targets detected by the radar sensors; other operators preferred to see all of the targets detected by radar. Clearly, this difference in opinion indicates a need for a system that can be configured by an operator.

With respect to geospatial elements, the preferred amount of data to be provided to the driver varied among drivers. Some operators prefer great detail, including guard rails, jersey barriers, mail boxes, etc. Other drivers prefer to keep the view as free from clutter as possible, and referencing only the lane boundaries to determine vehicle position and desired trajectories.

These results are based on relatively informal development efforts with selected operators. Because the system developed under this program was field tested as part of the FOT, data from that test will provide additional insight into the preferred configuration of the HUD.

e. Preview Amount

Drivers had definite opinions regarding the amount of preview they preferred: more is better. An easy solution would be to provide preview information to the horizon, thereby presenting no limits on how far a driver could “see.” However, there is an issue with providing geospatial information to the horizon. The EVT-300 radar has a maximum effective range of approximately 350 ft. In order to keep the driver aware that the range of the radar is 350 ft., the HUD display was configured so that geospatial information was presented in the proper colors to a distance of 350 ft. from the host vehicle. Beyond the 350 ft. mark, all geospatial elements were rendered in the color green. By splitting colors in this manner, a clear delineation exists between areas covered by radar sensors and those areas not covered. The driver is continually reminded where radar coverage exists, and is, therefore, more likely to drive with greater care in low visibility conditions.

Chapter 3 Integration of Radar and High Accuracy Geospatial Databases

Background

Reducing the severity and frequency of collisions has been a goal of Advanced Vehicle Control and Safety Systems (AVCSS) for some time. To this end, significant research has been conducted with automotive radar to improve the ability to sense obstacles surrounding the host vehicle (the vehicle equipped with the intelligent safety system) [6] [7] [8]. The result of this research has led to the deployment of off-the-shelf products that can sense obstacle position with high accuracy and resolution.

While radar has dramatically improved in quality in recent years, this sensing technology is still prone to detect “false targets”. A false target is defined as any detected object that is not on the roadway, and therefore poses no threat to the host vehicle. Signs, trees, bridge abutments, and light poles detected by radar are all examples of false targets. False targets reduce the effectiveness of AVCSS. In collision warning systems, these false targets may produce false warnings, which when too frequent, cause a reduction in the user's confidence in the system and may lead the user to turn the system off. For collision avoidance systems, false targets may cause the vehicle control/actuation system to produce unneeded maneuvers that may alarm or annoy the driver.

Radar manufacturers go to great lengths to reduce false targets. Internal signal processing is used to attempt to determine whether return echoes originate from stationary or moving targets. Stationary objects can then be identified, rejected, and not reported. This strategy may result in some valid targets (stalled vehicles, for example) not being detected. Radar has traditionally suffered from a significant number of false detections along curved road sections because the radar beam in a front mounted configuration detects objects beyond the outside curved boundary of the road. To reduce this effect, some radar sensors use a gyro to help reduce false targets by estimating the road curvature and ignoring objects outside the curved roadway. This strategy can be effective while the vehicle is traversing the curve and the yaw rate is sufficient enough to be measured by the gyro. However, before entering the curve (when the yaw rate is still zero), this strategy is not effective; please see Figure 3-1.

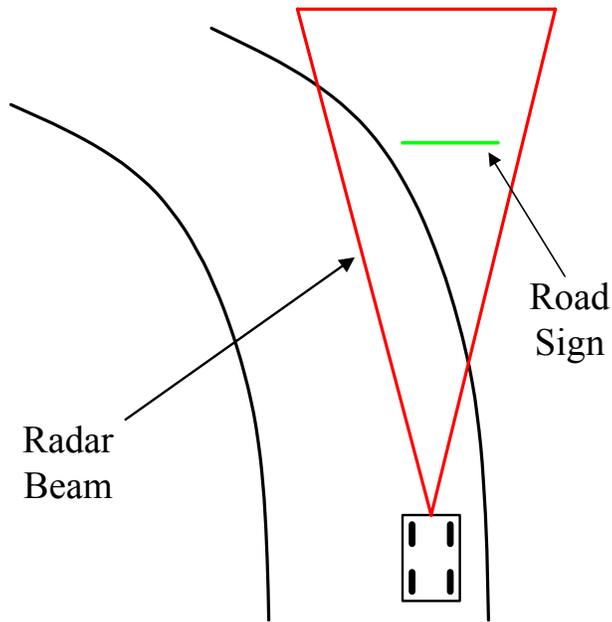


Figure 3-1: Host vehicle entering a curve with a road sign directly in the radar's field of view.

To reduce the number of false targets, an advanced filtering scheme is needed. Such a system has been developed by using a high accuracy DGPS and a highly detailed and accurate Geospatial Database (GDB). Newer DGPS technology provides significantly more accurate dynamic position and orientation of the host vehicle than was possible in the past. The GDB provides the detailed location of not only the road, but also the roadside “furniture”. This chapter documents the development of this advanced range sensor filter and culminates with the presentation of experiments conducted on Minnesota Trunk Highway (TH) 101 between Rogers and Elk River.

Geospatial Database

To satisfy a need for highly accurate and detailed road information, a GDB and real time query processor have been developed as the core around which driver assistive subsystems can be implemented. Details regarding the GDB are fleshed out in [9].

The GDB has been designed to provide the onboard vehicle systems with road information in real time. The onboard systems need specific and timely information about the roadway adjacent to the vehicle as the vehicle moves. The GDB is a collection of spatial objects (points, lines, curves, and polygons) that represent real world geometry. A data model has been developed to satisfy the data needs of multiple onboard driver assistive systems. Each system has specific requirements of the data within the geospatial database. For the range sensor processor, two object types were created that define the extent of the drivable road surface. These two object types have been labeled RoadShoulders and RoadIslands. RoadShoulders are defined as the boundary of any

drivable surface. Usually, this corresponds to the edge of the pavement; however, the edge of dirt and gravel surfaces can also be defined as a RoadShoulder. The range sensor filter uses the road shoulder to determine whether an object detected by the range sensor is on the road surface. RoadIslands are areas contained within RoadShoulders, or within the roadway, that are not drivable surfaces. RoadIslands usually correspond to traffic islands, medians, or the area between divided highway lanes. Once the range sensor processor determines that an object is within the RoadShoulder, the processor compares the location of the detected object against the areas defined as RoadIsland in order to determine whether the detected object lies within the RoadIsland and thus is not a threat.

To query the database, the querying subsystem creates a polygon and a list of data objects with specific attributes. The polygon is the area of interest around or in front of the vehicle. For the range sensor filter, the query polygon consists of the two dimensional cross-section of the radar's horizontal field of view. The query request, which includes the polygon and desired data objects, is then sent to the query processor. The query processor searches through the database to find any data objects that meet the criteria of the query. The query results are the data objects of the specified type with the correct attributes that are located within (or intersect with) the given polygon.

Range Sensor Processor

The primary function of the range sensor processor is to transform the position of the geospatially local objects detected by the range sensor (whose location is relative to the vehicle coordinate system) into a global coordinate system using the DGPS-determined position and orientation of the vehicle. The detected objects positioned in the global coordinate frame can then be compared with roadway features (also in the global coordinate frame) obtained from the database query. Objects located outside the drivable surface of the road are filtered from the valid target pool. Valid target data can then be used by systems requiring this information, such as the HUD referenced in Chapter 2.

Once objects detected by the range sensor are represented in the state plane coordinates (the transformation to state plane coordinates was described in Chapter 2), the geospatial database is queried using the vehicle's DGPS location to extract all the relevant road features local to the vehicle. The database objects of interest are the RoadShoulders and RoadIslands because these objects define the outermost drivable surface of the road. A vehicle to the right of a right shoulder, for example, is off the road and is of no threat to the host vehicle.

The GDB query processor is a server program that accepts queries from either local or remote processes. A query can be specified by the attributes of the database objects, meaning that the range sensor processor is able to query the database for specific objects. For example, RoadShoulder objects with a "right shoulder" attribute were queried to calculate whether the target point was originating to the right or left of the right shoulder. (Left and right shoulders are defined relative to the direction of traffic.) Range sensor detected objects to the right of the right shoulder are filtered, while objects to the left of the right shoulder are identified as valid targets. Target 1 in Figure 3-2 illustrates a target

to the right of the right road shoulder. It would be filtered out at this step. A similar procedure was used for the left shoulder. Finally, a query was made to obtain the RoadIsland objects. The remaining range sensor targets were tested to determine whether they were located within the RoadIslands. If a range sensor target passed through all three filters, it was considered a valid target vehicle (target 3 in Figure 3-2).

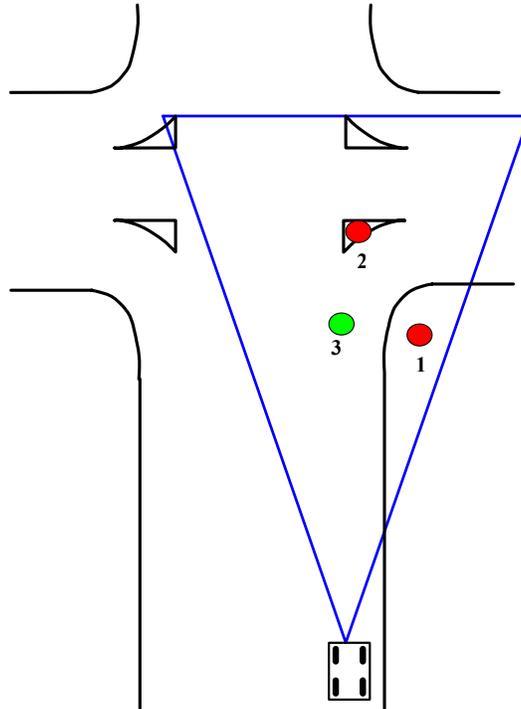


Figure 3-2: Illustration of two targets (1 and 2) that were filtered and one target (3) that passed through the filter for further processing.

Experiments

The system was implemented in the SAFELOW and experiments were conducted in live traffic to determine whether the range sensor processor was capable of filtering unwanted radar returns. Prior to that, however, database experiments were performed to measure the query times and how those times are affected by query area.

Database Experiments

The objects detected by the range sensor system are presented to the vehicle operator in real time by the HUD (see Chapter 2). Any target filtering must be performed in a timely manner so as to minimize any delay or mismatch between the actual detected object and the presented object. One area where processing latency can be minimized is query processing. The query processing routines of the GDB management system were optimized for speed to minimize processing latency.

When the target filter process updates its data set at every new DGPS position measurement, its query polygon normally matches the range sensor's field of view. However, because the vehicle is moving, the query results for one position may not be valid for the next position. Increasing the size of the query polygon will allow the data to be valid for more than one DGPS position. This will increase the local processing performed by each driver assistive system since more information will be returned; however, each system will not have to query the database as often. Experiments were performed to determine the overall query time and the effect of enlarging the query polygon with respect to radar target filtering.

All experiments were conducted on a PC104 computer with a 400 MHz AMD processor running the QNX/Neutrino real time operating system. A decimeter level accurate GDB was compiled using road information provided by the photogrammetry unit of the Minnesota Department of Transportation. The data covers all four lanes, intersections, and turn lanes of the divided Minnesota TH 101 for an 11.6 km segment between Rogers and Elk River. This database consists of 683 separate digital spatial objects that represent lane boundaries, lane centers, and the extents of the drivable surface. The GDB contains 125 RoadIsland and RoadShoulder objects that are used for target filtering.

The range sensor filter query was performed at 8,385 locations within the database. These positions correspond to actual vehicle positions determined by DGPS and recorded while driving on TH 101. Using the real time clock available in the QNX/Neutrino real time operating system, the time was measured from when a query was sent to when the results were received. This time includes the communication time, query processing, and any local processing involved in receiving a query.

In this experiment, the query polygon was increased in size by extending the right, left, and total length of the query polygon by 3.12 m per 0.1-second interval. The vehicle's position was updated at 10 Hz. Assuming a maximum speed of 112 kph (70 mph), or 31.19 m per second, the vehicle can move 3.12 m per position update. By enlarging the query polygon, the query frequency can be reduced.

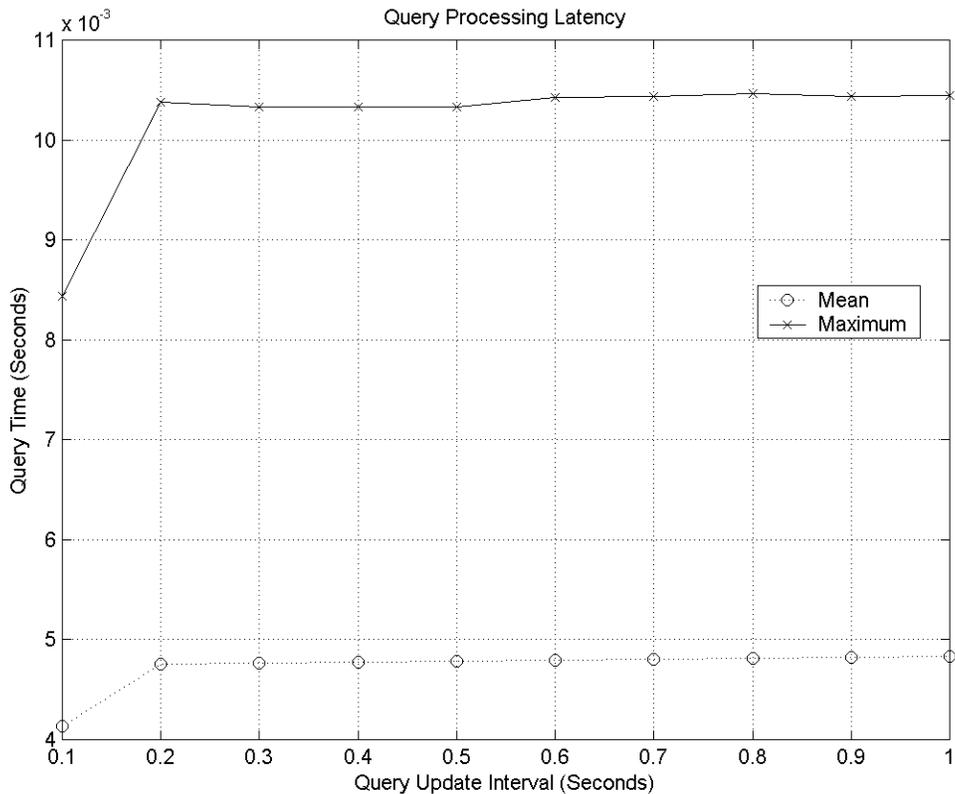


Figure 3-3: Radar target filter query times. For each query update interval the query polygon is increased in size to account for vehicle motion between query updates. Each point corresponds to the mean or maximum query time of 8,385 queries performed along the length of the GDB corresponding to actual vehicle positions.

As can be seen from Figure 3-3, the query time is small compared to the time it takes to update DGPS position and can easily be compensated for by projecting the vehicle's position ahead in time. As the query polygon is increased in size, the corresponding query times do not significantly increase. These results show that the effect of query processing can be further reduced by decreasing the query rate or performing queries when processing load is minimal. The initial jump in query time from 0.1 second query once every position update to a 0.2 second query interval (query every other position update) is due to changing the query polygon from a triangle (3 points) to a trapezoid (4 points). The trapezoid is produced by enlarging the triangle and then clipping off one vertex, as there is no need to query closer to the vehicle. Clearly, all query processing can be preformed under eleven milliseconds, with a mean query time less than five milliseconds.

Range Sensor Filter Experiments

The range sensor filter was implemented in the C programming language on a PC-104 computer running the QNX/Neutrino real time operating system. The computer resides in the SAFEFLOW and uses a Trimble MS 750 GPS receiver and two Eaton VORAD EVT-300 radar units (see Figure 3-4) Experiments were conducted in real traffic on Minnesota TH 101 between Rogers and Elk River.

It should be noted that while the range sensor processor was implemented using the Eaton VORAD radar, the processor is not limited to this sensor. In fact, the processor can be implemented using any range sensor or combination thereof that provides a sufficiently accurate target location (range and azimuth).

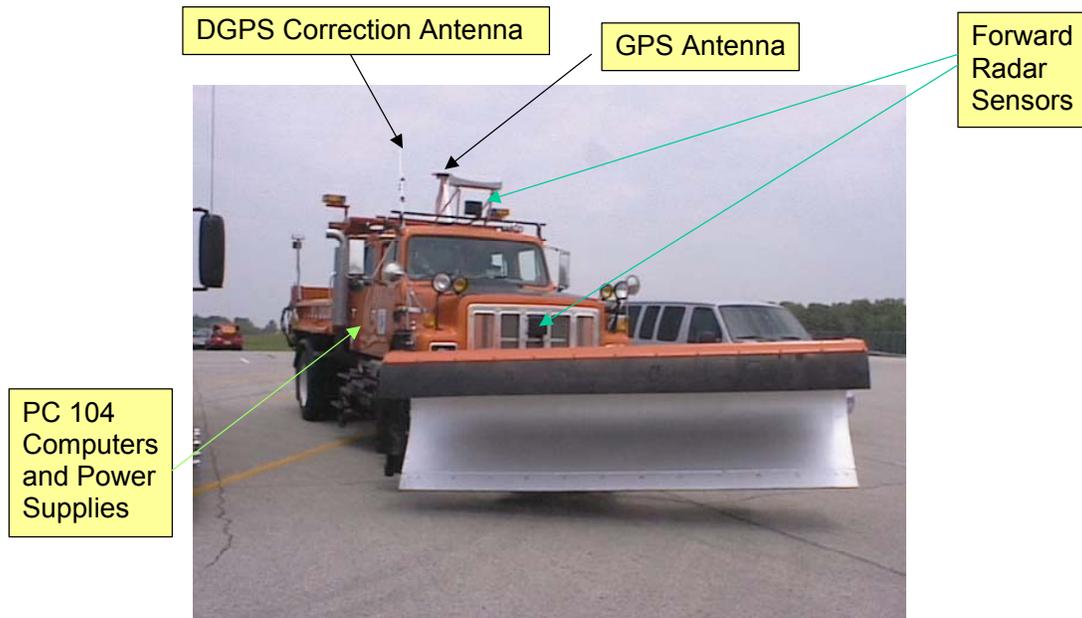


Figure 3-4: The SAFEFLOW with DGPS and two Eaton VORAD radar mounted on the front grille and top of the cab.

The SAFEFLOW was driven North and South on Highway 101; data was collected to test the ability of the range sensor processor to exclude unwanted radar returns. Figure 3-5 shows the results of one such experiment. The SAFEFLOW was traveling southbound on Highway 101 (green “*”) while following a vehicle in the right hand lane. The valid target is shown as a red “x” while filtered targets are shown with a blue “o”. The false target is clearly to the right of the right road shoulder (from the perspective of the SAFEFLOW traveling southbound, as indicated by the arrow in Figure 3-5). The range sensor filter effectively filtered out the radar return emanating from the side of the road. This false target would not be presented to the HUD or a collision avoidance/warning system because it poses no threat to the host vehicle. The roadside object is a road sign.

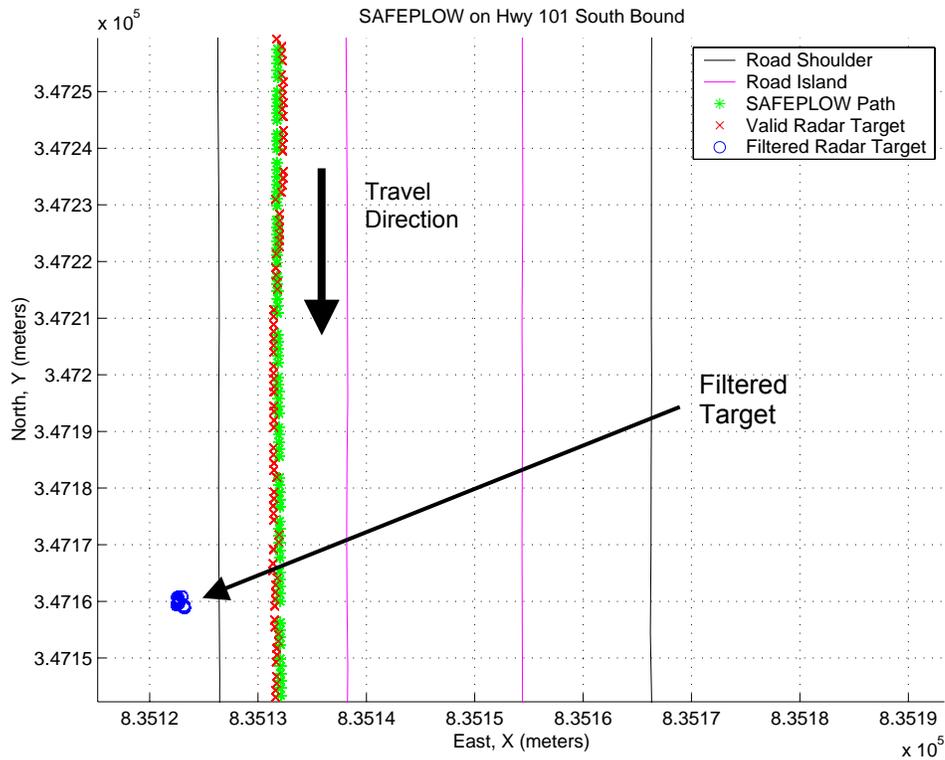


Figure 3-5: SAFEFLOW driving southbound on Highway 101. False target filtered off of the right shoulder.

To demonstrate the ability of the radar processor to filter targets on road islands, this study considered the result of driving the SAFEFLOW through an intersection. The results are shown graphically in Figure 3-6; the symbols in the legend are identical to the previous figure. The SAFEFLOW was driven through the intersection, and the radar tracked several objects. The SAFEFLOW was following a valid target when the target switched lanes inside the intersection. The radar then started tracking the next vehicle in front of the SAFEFLOW. Notice that two separate false targets were successfully filtered from the road islands. In the southern road island, notice that some of the points are located inside the road island and some are located just outside the road island boundary. While the radar returns came from metallic objects (road signs, signal poles) within the road island, inaccuracies in DGPS and the range measurement caused them to be located just outside the island.

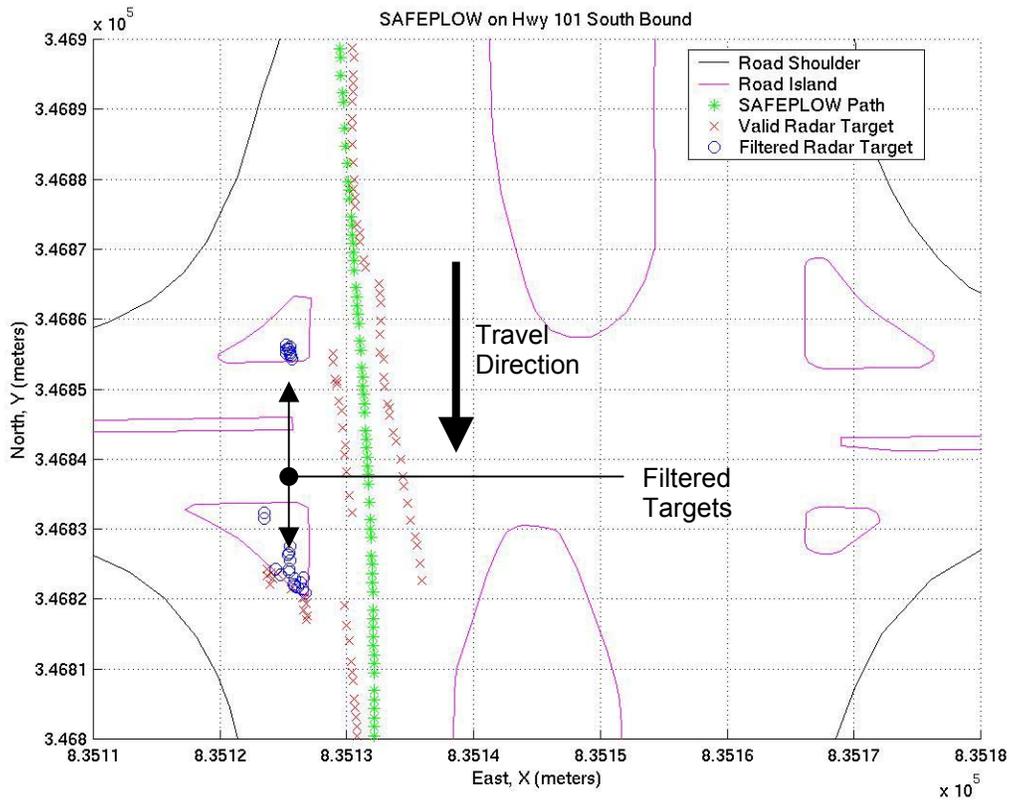


Figure 3-6: SAFEFLOW driving southbound through an intersection on Highway 101. False targets at road islands were filtered. Trajectories of two vehicles sensed by the onboard radar are drawn in the south bound lane.

Conclusions

As a result of this task, a novel range sensor processor that has the ability to improve false target rejection in real time was designed and implemented. The filter locates target objects sensed by radar (or any range sensor that can provide object location) and places them on the road scene using vehicle mounted DGPS and a high accuracy geospatial database. Objects located to the right of right shoulder segments and to the left of left shoulder segments are filtered from the target pool. Additionally, objects located within road islands are filtered. This leaves valid target obstacles that are assured to be located on the roadway. Drivers are then presented warnings based on these remaining radar targets. This significantly reduces the number of irrelevant targets which may lead to driver annoyance in collision avoidance/warning systems.

Chapter 4 Snowplow Dynamics Modeling for Slippery Conditions

Introduction

The goal of task 4 of this project is to develop dynamic models of the motion of a snowplow in various conditions that will aid in the design and implementation of various driver assistive systems. The models developed as part of this task are used to implement virtual bumper-based gang plowing described in Chapter 5.

Tradeoffs exist between model complexity and fidelity. Because of constraints with on-board processors, for real time control and driver assistive systems, the goal is to use the least complex model that provides “reasonable” fidelity. Modern dual frequency, carrier phase DGPS systems provide extremely accurate vehicle position measurements and derived vehicle orientation estimates at relatively high data rates (10 Hz). These measurements, combined with a proper feedback control structure, can compensate for models of dynamic system that provide reasonable, but not full, fidelity. For the present application, simple models offering moderate to good fidelity combined with accurate measurements and solid feedback structures are preferred to complex models offering good to excellent fidelity. The lack of fidelity is compensated for by excellent measurements.

Thus, the models developed herein should be as simple as possible to allow them to run in real time on a mobile computer without using a significant amount of the processing power available. At the same time, the model will have to be robust to handle large variations in vehicle mass and yaw moment of inertia, because it will be used for both loaded and unloaded snowplows. The model should also be accurate enough to detect sliding on slippery roads.

This chapter describes a very simple lateral model of a snowplow that fulfills these requirements. The model was developed based on experimental data collected on the skidpad at the Minnesota Highway Safety Center (MHSC) located at St. Cloud State University in St. Cloud, Minnesota. The skidpad at MCHS is flexible in that it can be used as a high friction surface when dry, a medium friction surface when wet, and a low friction surface (which acts much like ice) when covered with a wetted clay. Experiments were performed under all three conditions during the course of this investigation.

An aerial view of the MHSC in St. Cloud is provided in Figure 4-1. The skid pad is denoted.

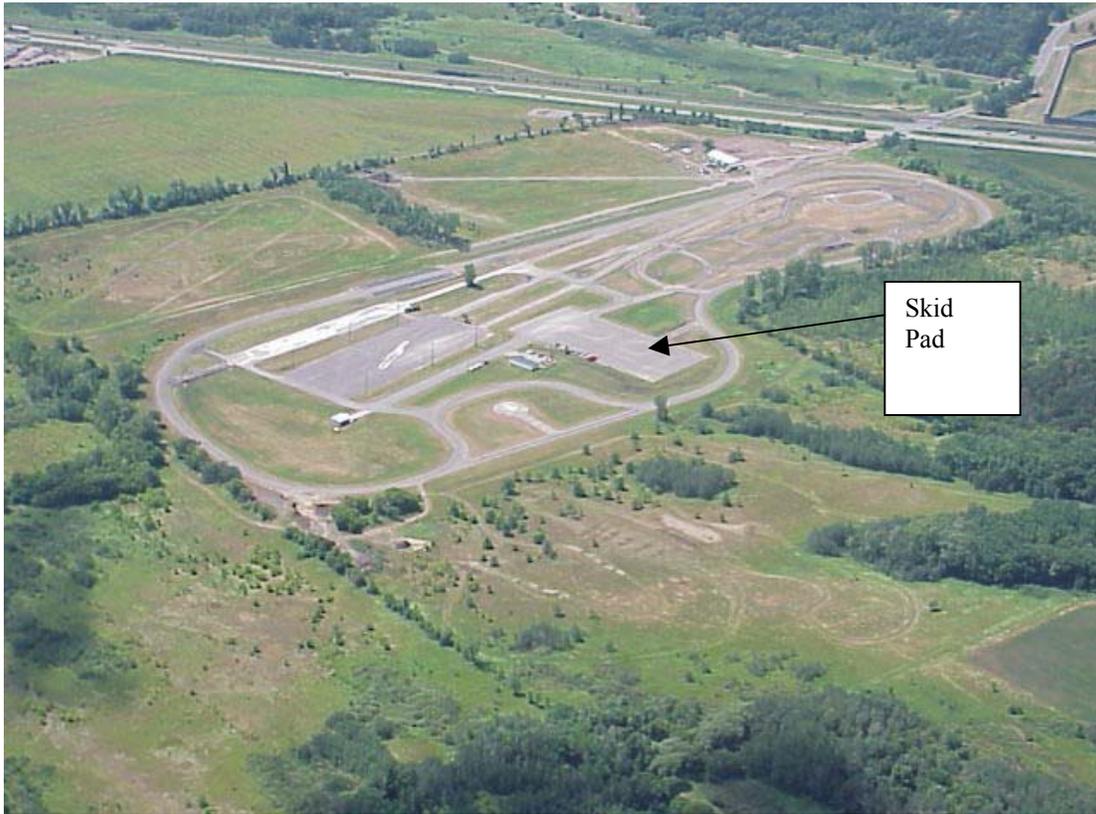


Figure 4-1: Aerial view of the Minnesota Highway Safety Center. Skid pad on which vehicle dynamic experiments took place is indicated.

The vehicle dynamics literature contains a number of mathematical models that attempt to describe the relationship between vehicle control inputs (steering, throttle, and brakes) and the resulting motion of the vehicle. Many of these models are based on the single track (bicycle) model proposed in [10]. The single track model eliminates vehicle roll and pitch motions and computes the motion of a rigid body in a horizontal plane accelerated laterally by the forces acting on a single tire at each end of the vehicle. The lateral tire forces are typically modeled as linearly proportional to their slip angle. This assumption is reasonably accurate when the slip angles are small and the friction coefficient is high, but does not hold when lateral forces are a significant percentage of the traction that is available.

A linearized state space equation for the bicycle model, derived in [11] is:

$$\frac{d}{dt} \begin{bmatrix} \beta \\ \Psi \end{bmatrix} = \begin{bmatrix} -\frac{\mu(c_f + c_r)}{Mv} & -1 + \frac{\mu(c_r l_r + c_f l_f)}{Mv^2} \\ -\frac{\mu(c_r l_r - c_f l_f)}{I_\psi} & -\frac{\mu(c_f l_f^2 + c_r l_r^2)}{I_\psi v} \end{bmatrix} \begin{bmatrix} \beta \\ \Psi \end{bmatrix} + \begin{bmatrix} \frac{\mu c_f}{Mv} \\ \frac{\mu c_f l_f}{I_\psi} \end{bmatrix} \delta_f, \quad (4.1)$$

where β is the side slip angle between the longitudinal axis of the vehicle and the velocity vector at the CG and Ψ is the vehicle's yaw angle. Moreover, the parameters are defined as

- M : vehicle mass,
- I_ψ : vehicle inertia about vertical axis through CG,
- μ : tire/road friction factor,
- l_f and l_r : distance from CG to front and rear axles, and
- c_f and c_r : front and rear tire cornering stiffness.

Except for the mass of the vehicle, all of these parameters are difficult to estimate and change continually.

Although this bicycle model is a simple representation of the vehicle dynamics, integration of the resultant differential equations consumes significant CPU resources. In the sequel, a model of less complexity than the bicycle model is proposed and validated based on the results of a series of experiments performed on the skid pad at the MHSC.

Kinematic Model

A kinematic model computes the motion of a body without regard for the forces that cause that motion. A kinematic model of the lateral dynamics of a snowplow is proposed with the following underlying assumptions:

1. Under normal operating conditions, on dry pavement, the vehicle moves in the direction in which it is steered with small slip angles.
2. Slip angles can be accommodated with a simple factor of proportionality (for a normally understeering vehicle where the front wheel slip angles are greater than those in the rear).

The steering angle and speed of a front-wheel steered vehicle can be related to its yaw rate via the following equation:

$$\dot{\Psi} = v \frac{\delta}{wb}, \quad (4.2)$$

where Ψ is the vehicle's yaw rate, δ is the steering angle of the front wheels, v is the vehicle forward velocity, and wb is the wheelbase.

Figure 4-2 shows a set of test data taken on a particular run of the SAFEFLOW along a dry, paved section of the test track. The graph shows the recorded steering input, vehicle speed, yaw rate, and lateral acceleration. The steering position was calculated from a 16 bit encoder in a Kollmorgen Goldline servo motor attached via toothed belt to the steering column. Vehicle speed was calculated from successive positions sensed by a Trimble MS 750 Differential GPS receiver. Yaw rate and lateral acceleration were measured with a Crossbow HDX six axis inertial measurement unit (IMU).

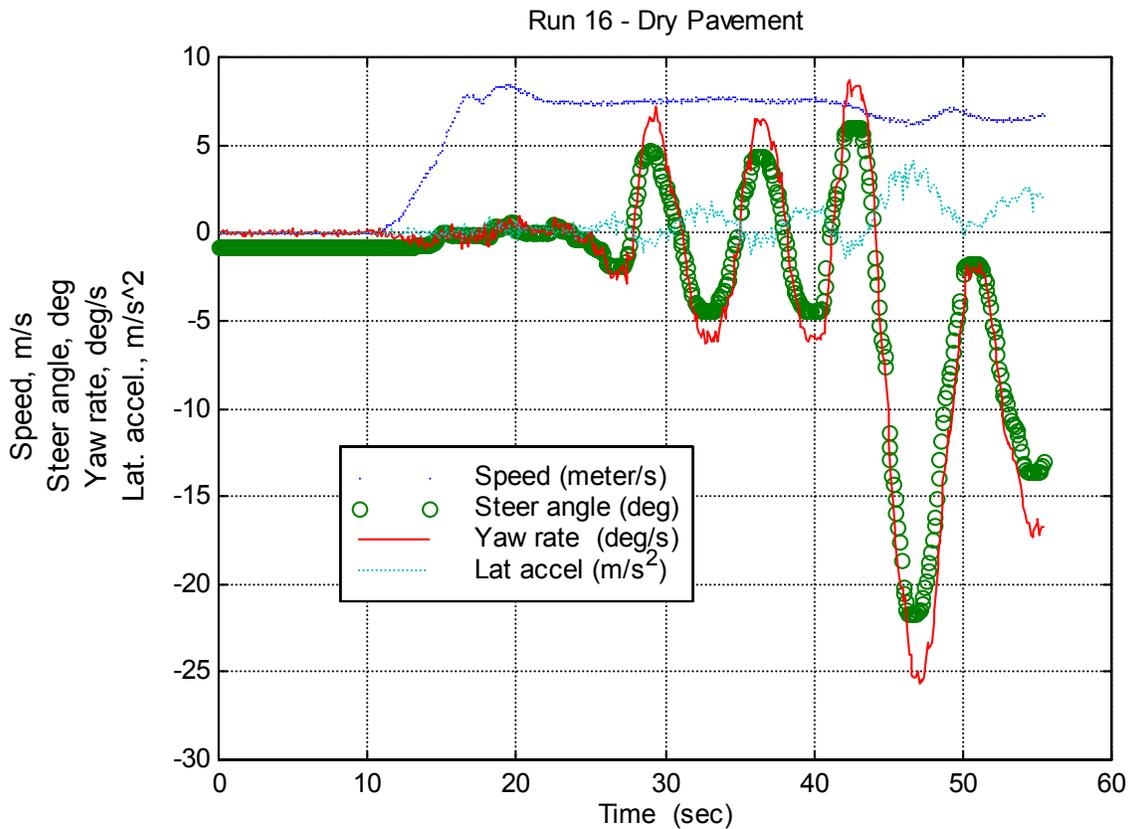


Figure 4-2: Test data collected from a test run at the MHSC using the SAFEFLOW on dry pavement. Vehicle speed is approximately 18 miles/hour.

The yaw rate predicted by equation 4.2 can be compared to the yaw rate measured by the IMU by using the measured vehicle speed and steer angle as input plotting the resulting

estimated yaw rate against the measured yaw rate. Figure 4-3 illustrates this comparison; this simple model well represents vehicle dynamic characteristics at slow speeds and moderate steering angles. Close inspection reveals that the kinematic model usually leads the actual test data and returns an amplitude greater than the measured yaw rate.

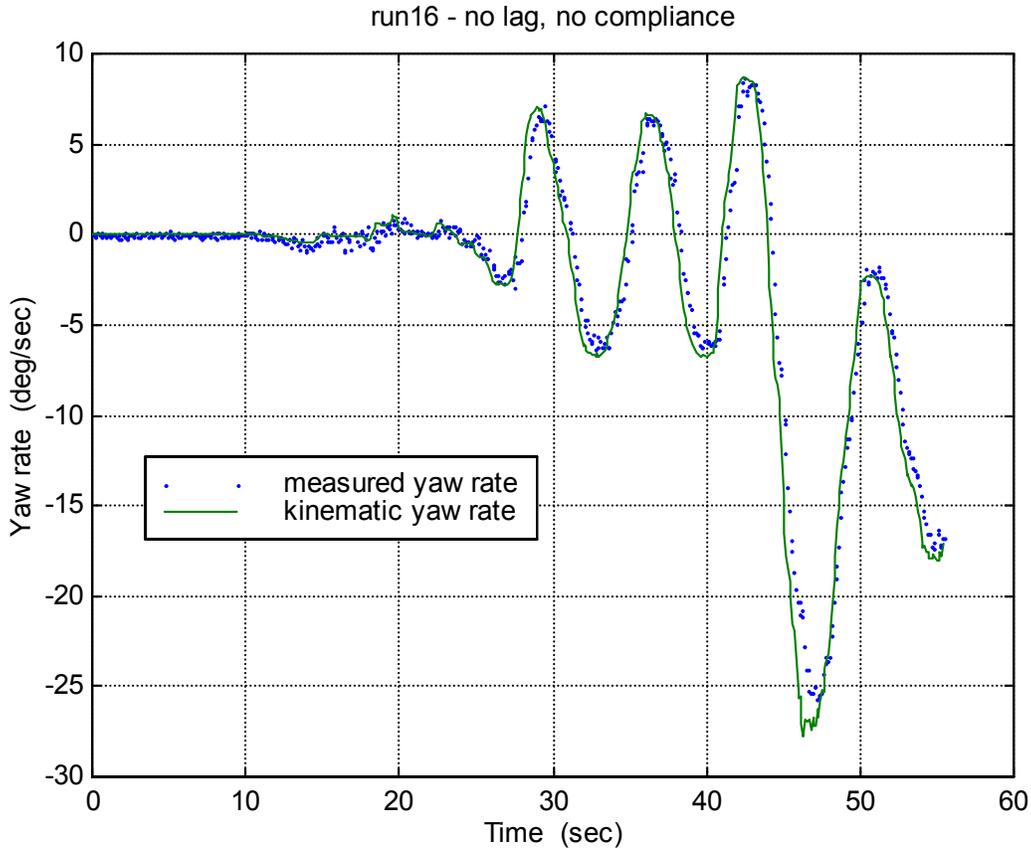


Figure 4-3: Comparison of the computed and measured yaw rates corresponding to the data illustrated in

Figure 4-2.

Two modifications to this simple model can improve its fidelity. First, the front wheel steer angle has been calculated from the position of the steering wheel. The steering ratio relating steering wheel position and the steering angle of the front wheels was calculated by holding the steering wheel in a fixed position and recording the trajectory of the SAFEFLOW using DGPS while moving at very slow speeds. Data from several runs with the steering wheel in various fixed positions were recorded. From that data, steering wheel position was mapped to the radius of the truck trajectory, and a steering ratio of 22:1 was determined.

In this experiment, the slow speed minimized side loading of the suspension and tires. In the higher speed experiments on the test track, side loads are large enough to cause flex in the tire, steering, and suspension. This compliance can be compensated for by slightly reducing the amount of steer angle at the front wheels (3% for the test run shown in

Figure 4-2) from the 22:1 ratio determined with no side loads.

The other modification compensates for the vehicle lateral response delay arising from steering wheel inputs. This delay can be accounted for by adding a lag between the steering input to the kinematic model and the vehicle yaw rate. Figure 4-4 illustrates the error between the computed and the measured yaw rate data for lags of 0, 0.1, 0.2 and 0.3 seconds. The mean of the absolute value of this error is smallest for a lag of 0.2 seconds; a lag of 0.2 seconds is therefore incorporated into the kinematic model.

The modified kinematic model equation becomes

$$\psi = v \frac{0.97 \delta_d}{wb}, \quad (4.3)$$

where δ_d is the steering angle delayed by 0.2 seconds, and 3% compliance is denoted by the weight of 0.97 in the numerator.

Figure 4-5 illustrates the measured yaw rate and the modeled yaw rate modified to include 3% compliance in the steering and a 0.2 second lag in the response. Clearly, the model is a very good predictor of the actual measured yaw rate.

Icy Surfaces

The performance of the kinematic model on the slippery skid pad is also investigated. The cornering stiffness of a tire in contact with the roadway is defined as the ratio of slip angle to the lateral force produced as the slip angle approaches zero. As the slip angle grows, the lateral force increases in a continuous, nonlinear manner until it saturates at a value dependent on the maximum lateral force that the tire can produce with a given vertical load. Tire cornering stiffness should be insensitive to road surfaces because it is measured with slip angles and lateral forces close to zero. This can be verified by examining test results for experiments conducted on an ice covered skid pad at the MHSC. Figure 4-6 shows sensor data for a 6 m/second (13 mph) run with a gentle weave as the steering input. Data was collected using the same sensors as the dry pavement experiments. Using the 0.2 second yaw rate lag and 3% steering compliance derived on dry pavement, Figure 4-7 shows that the computed and measured yaw rates closely agree.

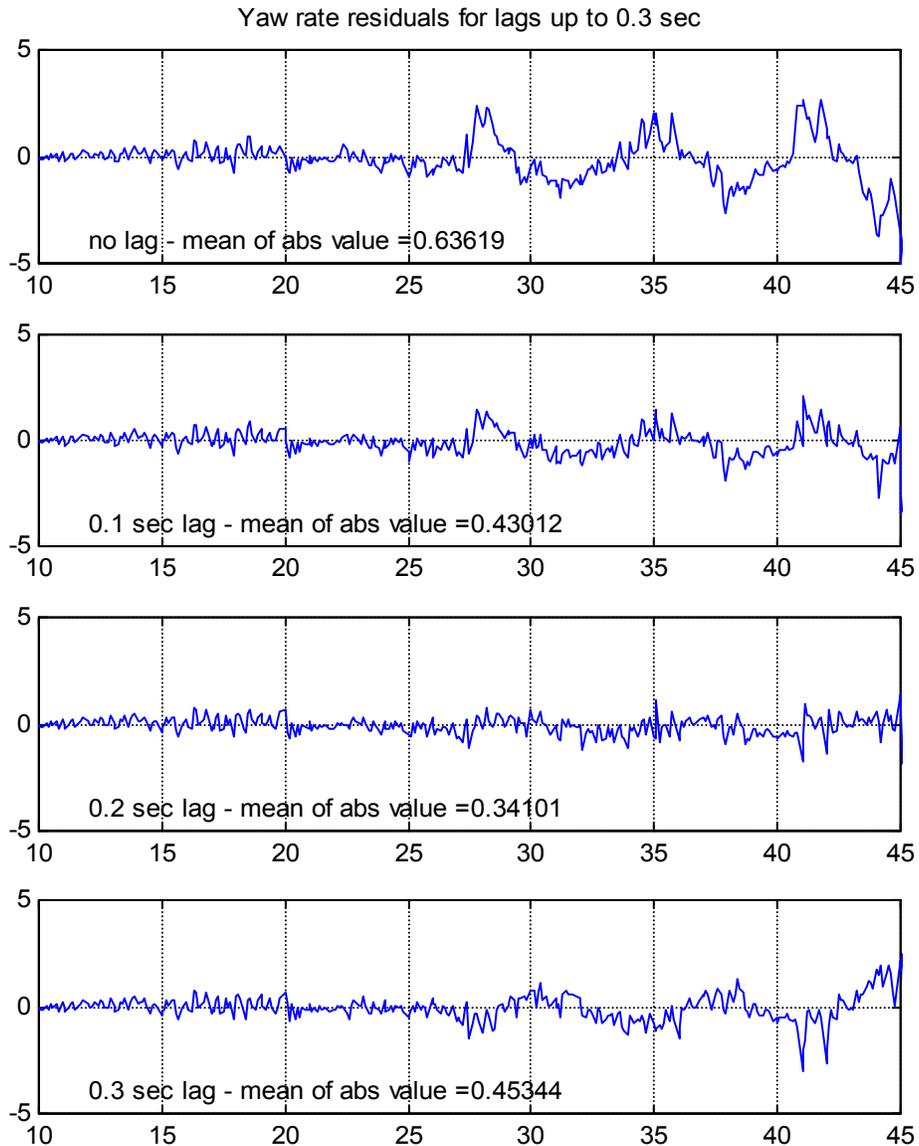


Figure 4-4: The difference (residuals) between the computed and the measured yaw rate for four different values of yaw rate lag. The residual is measured in degrees/second.

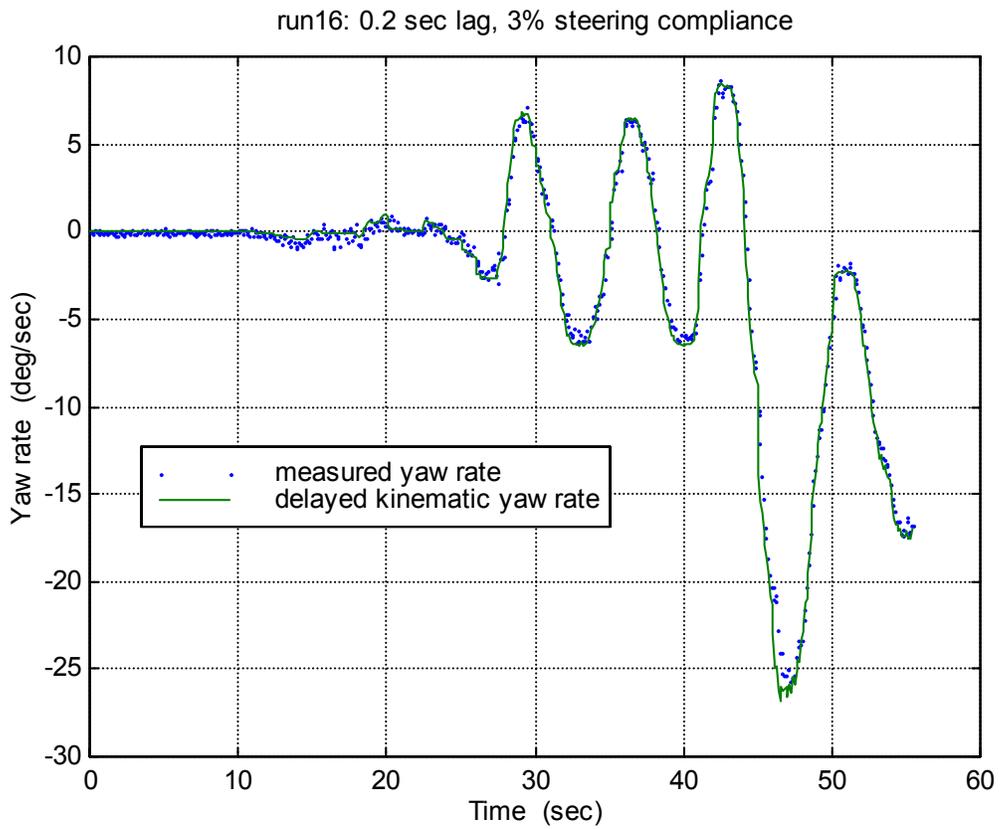


Figure 4-5: A comparison of the measured yaw rate with the output of a kinematic model, including 3% steering compliance and a 0.2 second lag.

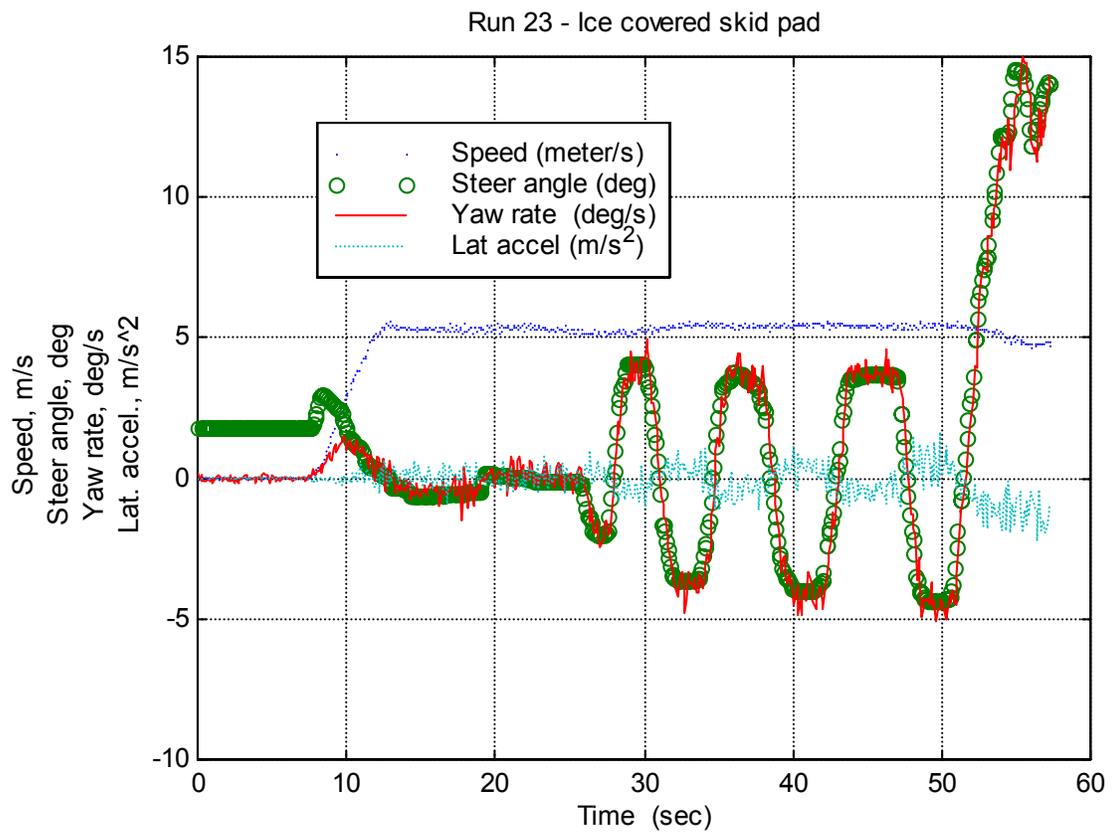


Figure 4-6: Test data collected from a test run using the SAFEFLOW on an ice covered skid pad. The vehicle speed is approximately 11 miles/hour.

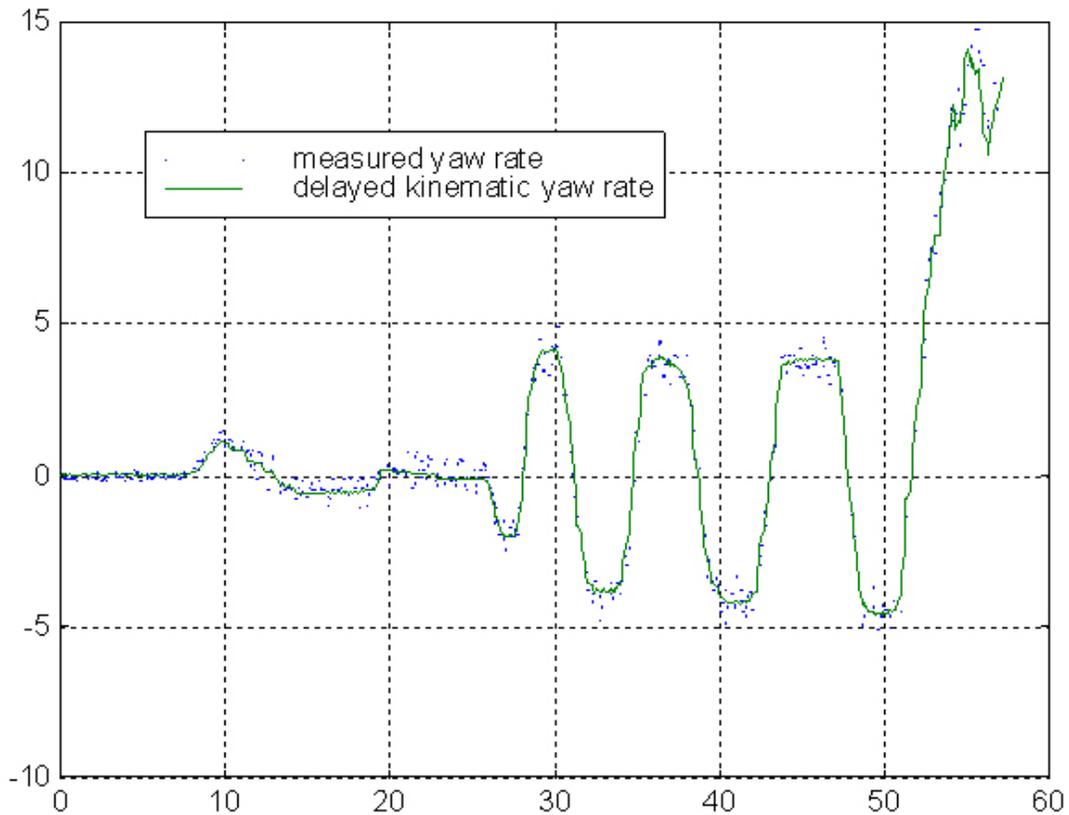


Figure 4-7: A comparison of the measured yaw rate with the output of the kinematic model on an ice covered surface during a gentle weaving maneuver.

Understeering Slides on Ice

A relatively simple model now predicts the motion of a truck when slip angles are small. This model can now be used to detect a slide by comparing yaw rates measured by the IMU to yaw rates estimated from steering position and vehicle speed. High front tire slip angles, indicative of an understeer condition, are indicated when the error between computed and measured yaw rates increases. Figure 4-8 shows the sensor data for a faster run on the icy skidpad with a more aggressive steering input than the previous examples.

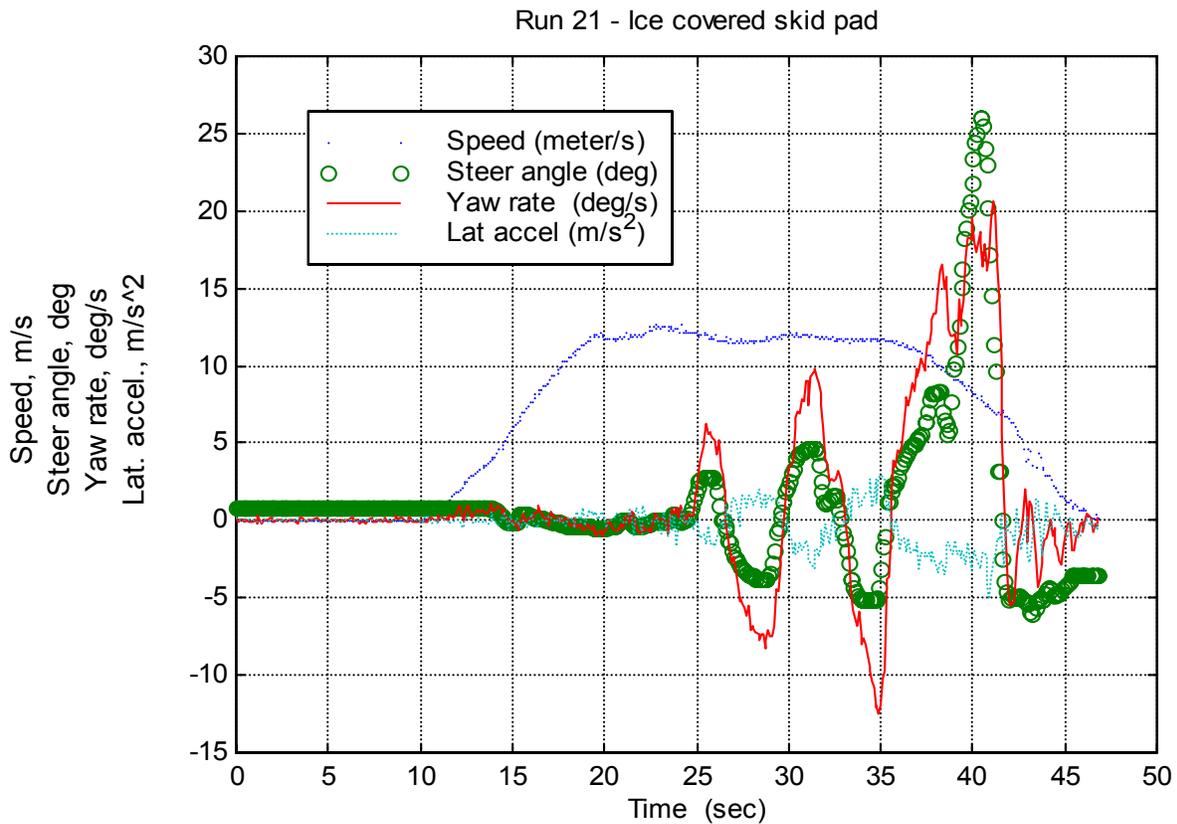


Figure 4-8: Test data collected from a test run using the SAFEFLOW on an ice covered skid pad at 12 m/s (27 mph) with a steering input strong enough to induce sliding.

In Figure 4-9 the kinematic model, which assumes small slip angles, predicts a faster response of higher magnitude than the actual measured yaw rate. This is particularly evident in the response to the large steering input at the 40 second point of the test run where the yaw rate is less than it would be if traction had been maintained. This indicates an understeering condition where the front wheels are sliding more than the rear. The difference between the model and the measurements are illustrated Figure 4-10. This figure shows a significant difference during the smaller steering inputs between 25 and 35 seconds as well as the large difference after 40 seconds. The difference in measured yaw rate and that predicted by equation 4.3 can therefore be used to determine if the front wheels of the truck are sliding.

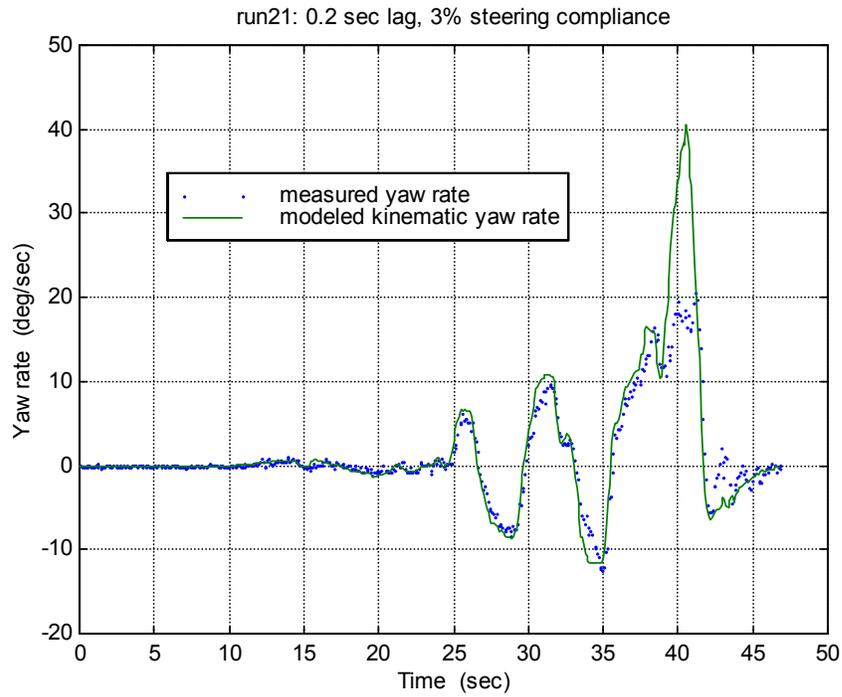


Figure 4-9: Comparison of predicted and actual yaw rates on ice with induced understeering slides.

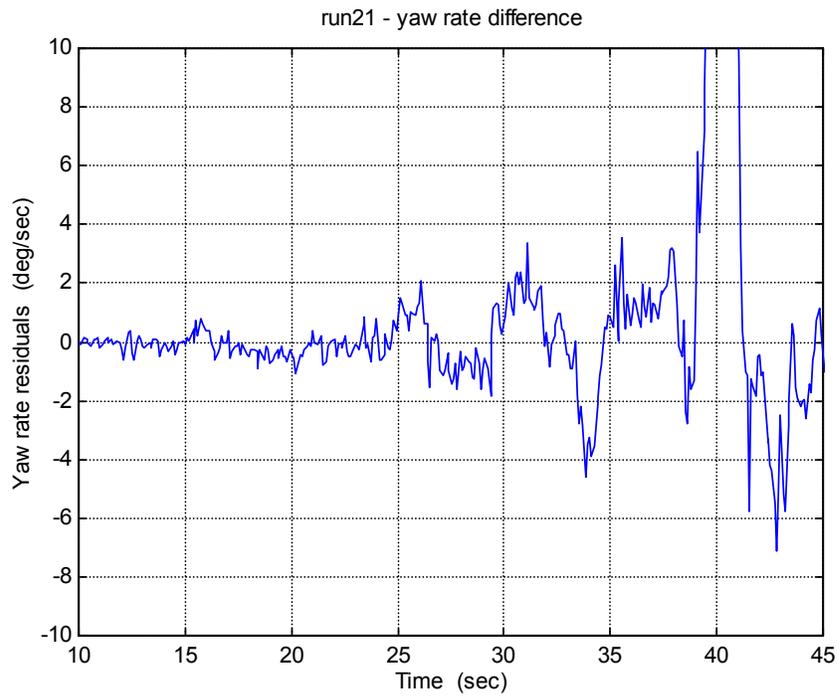


Figure 4-10: A plot of the difference between predicted and actual yaw rates makes it evident that sliding is occurring between 25 and 45 seconds.

Oversteering Slides on Ice

In Figure 4-11 the yaw rate estimates from the kinematic model (equation 4.3) and the yaw rate measurements where sliding was induced with aggressive steering movements at 12 m/second (27 mph) are plotted. In this run the actual yaw rate is greater than the kinematic prediction for the first swerve at 18 seconds, indicating oversteer with the rear wheels sliding. For the next set of swerves up until 30 seconds the actual yaw rate is less than predicted, indicating an understeering condition with the front wheels sliding. In the last large swerve the truck goes into a large oversteering slide.

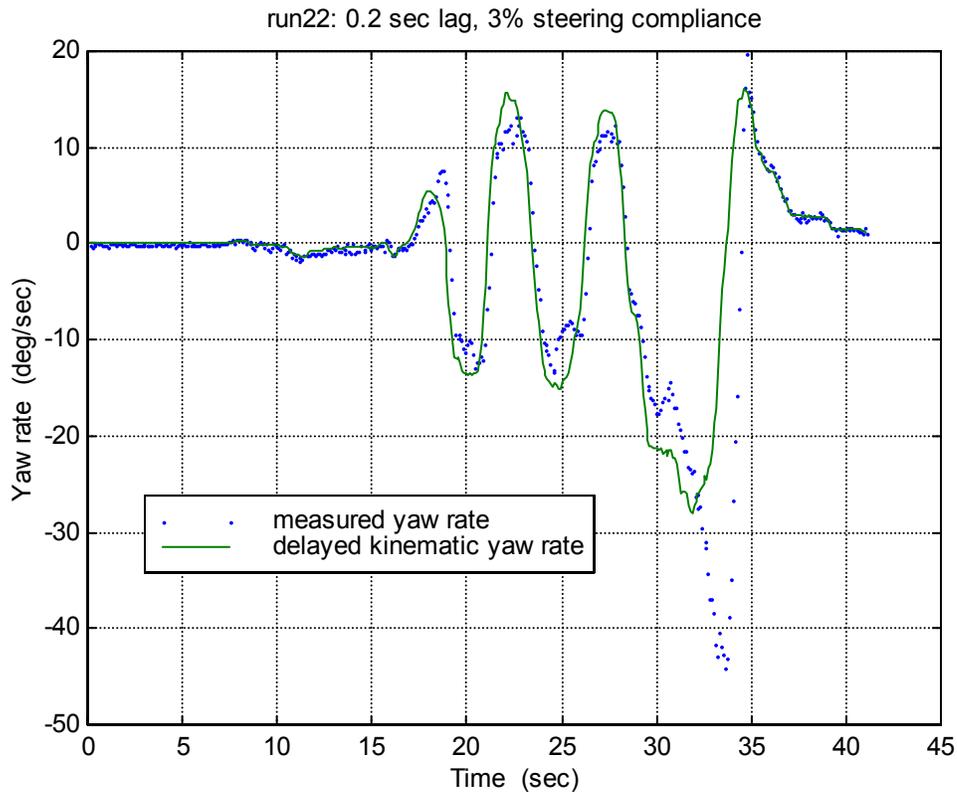


Figure 4-11: This graph shows a series of understeering and oversteering slides induced by strong steering movement on an icy surface. Understeering is indicated when the magnitude of the delayed kinematic yaw rate is greater than the magnitude of the measured yaw rate. Oversteering is indicated when the magnitude of the delayed kinematic yaw rate is less than the magnitude of the measured yaw rate.

Applications

With onboard sensors capable of measuring yaw rate, steering wheel position, and vehicle speed, a means to compare predicted yaw rate (equation 4.3) with measured yaw

rate exists. Understeering is indicated when the magnitude of the predicted kinematic yaw rate from equation 4.3 is significantly greater than the magnitude of the measured yaw rate. Conversely, oversteering is indicated when the magnitude of the predicted kinematic yaw rate from equation 4.3 is significantly less than the magnitude of the measured yaw rate.

The ability to sense understeer and oversteer conditions could be the basis for a snowplow electronic stability control system similar to those found on higher end passenger vehicles. Such a system would require, however, the ability to execute differential braking across the longitudinal axis of the snowplow. For instance, consider the situation shown by figure 4-12 below. If an understeer condition were sensed and differential braking about the vehicle longitudinal axis were available, application of the brakes on the left side of the vehicle would act to reduce understeer. Conversely, if an oversteer condition were sensed, application of the brakes on the right side of the longitudinal axis would act to reduce oversteer. With the anti-lock brake system (ABS) available on most new heavy vehicles, the application of brakes in this manner would be a straightforward endeavor.

The IV Lab has made attempts to gain access to a Bendix (the supplier of the brake system on the SAFEFLOW) brake development kit. However, because of issues with liability, no access to a kit has been provided to the IV Lab.

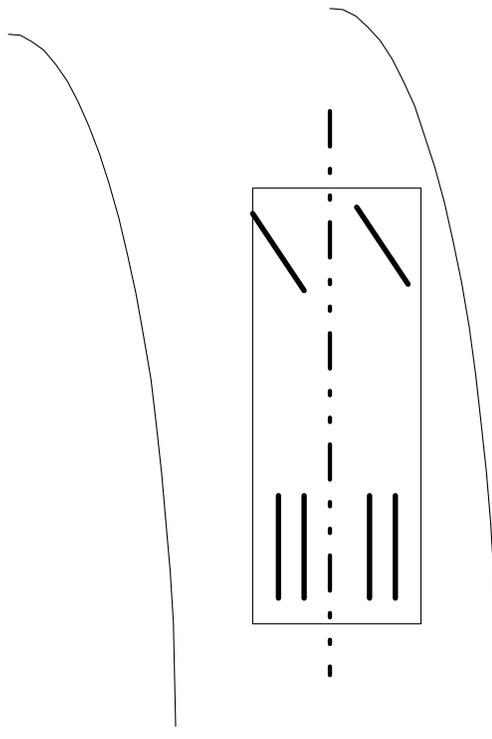


Figure 4-12: Snowplow on curve. The dashed line indicates vehicle longitudinal axis. In this situation, to mitigate understeer, the brakes on the left side of the vehicle would be applied. To mitigate oversteer, the right side brakes would be applied.

Conclusions

It is possible to use a very simple vehicle model to accurately predict the lateral motion of a snowplow on dry pavement. This model can be used in real time with much more computational efficiency than a typical dynamic model that would require the control computer to integrate a set of differential equations. This model can detect when the vehicle is sliding and sense the difference between an understeer and an oversteer condition. This information could be used in a driver warning system, or it could enable an automatic steering or brake system (similar to electronic stability control found on higher end passenger vehicles) to take the proper action when a slide is encountered.

Chapter 5 Virtual Bumper Based Collision Avoidance Applied to Gang Plowing

Background

In order to remove snow on a multilane road in one pass, gang plowing is often employed. Gang plowing consists of multiple snowplows plowing the road in a staggered formation, with the goal of clearing the road in one pass. This improves the traffic throughput, especially for major arteries.

The operators of gang plows face a difficult work environment. The operators of the following vehicles must not only deal with the prevailing weather conditions, but also with the snow from the lead vehicle. Visibility conditions are often poor, and in extreme conditions it is difficult for the operator to locate, much less closely follow, the lead vehicle.

Another difficulty in gang plowing arises because of the motoring public. Hurried motorists will squeeze in between plows in order to get to their destination more quickly. Often the “rogue” vehicle splits the gang, hits the windrow, spins, and is hit by one of the following plows. If the motorist is successful in splitting the gang plows, he finds that the roads aren't in very good condition ahead of the plows and has to slow down anyway. Much is risked for very little gain.

To address these issues, a gang plowing assistive system has been developed. The system provides headway control using both throttle and brakes that maintains a driver selected headway time to the lead vehicle. The system also provides a lateral control system that provides torque through the steering wheel that helps the driver of the following plow at a preselected lateral distance from the lead snowplow. This chapter documents the concept, implementation, and initial testing of the gang plowing assistive system.

Overview

To assist the snowplow operator with the difficult task of gang plowing, a gang plowing assistive system was developed for the trailing snowplow(s). The lead plow in the gang formation has no assistive system, though one could conceivably be added. The system on board the trailing snowplow(s) provides longitudinal control through the throttle and brakes and lateral control through the steering wheel to help the driver accurately follow the lead vehicle. The snowplow operator can select the desired lateral offset (Δy) and headway time (T_h) (Figure 5-1) via a touch screen interface mounted in the cab. A longitudinal control system based on the virtual bumper [12], [13], [14] is used to regulate the driver selected headway time, and therefore headway distance. The headway distance (R_h) is a function of the headway time and the velocity of the lead vehicle (V_{lead}).

$$R_h = T_h * V_{lead} \quad (5.1)$$

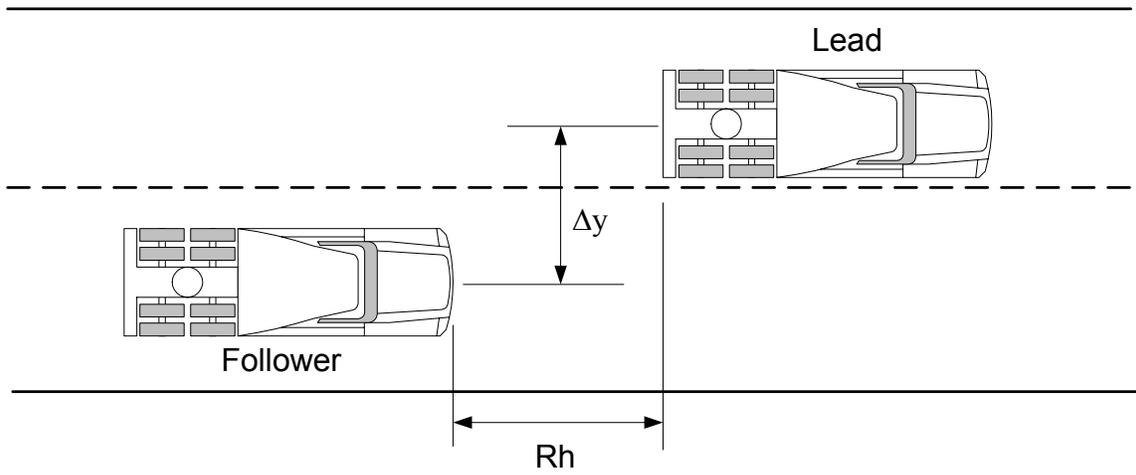


Figure 5-1: The snowplow operator of the following vehicle selects the desired lateral offset and headway distance from the lead vehicle

While the driver assistive system can autonomously follow the lead vehicle, the snowplow operator is in the loop and has the final control of the vehicle. The steering wheel feedback is purposefully limited in torque so that the operator can override it at any moment. The brakes can be applied by the driver at any time to slow the vehicle. In addition to the driver having complete control of the vehicle while the system is active, there are several redundant ways to deactivate the system. First, the touch panel contains a software switch that shuts down the gang plowing system. Furthermore, a hardware emergency switch can be engaged to turn off the motors and deactivate the remote throttle control system. Finally, a master power switch is located within the driver's reach that powers off all computers, motors, and hardware.

The gang plowing assistive system uses DGPS sensors to position the vehicles in the gang plowing platoon. The trailing vehicle is equipped with a radio frequency (RF) modem and receives the differential GPS correction from a base station near the snowplow's route (Figure 5-2). A Wireless Local Area Network (WLAN) was established by installing a WLAN access point in the following vehicle and a station adapter in the lead vehicle. The trailing vehicle sends the differential correction message obtained from the RF modem to the lead vehicle via the WLAN. The lead vehicle then broadcasts its DGPS position back to the following vehicle. A comparison is then made between the lead and trailing vehicle DGPS positions to obtain the lateral offset and headway distance between the vehicles. The relay of DGPS corrections from the trailing vehicle to the lead vehicle is done for two reasons:

1. limited financial resources
2. best use WLAN bandwidth.

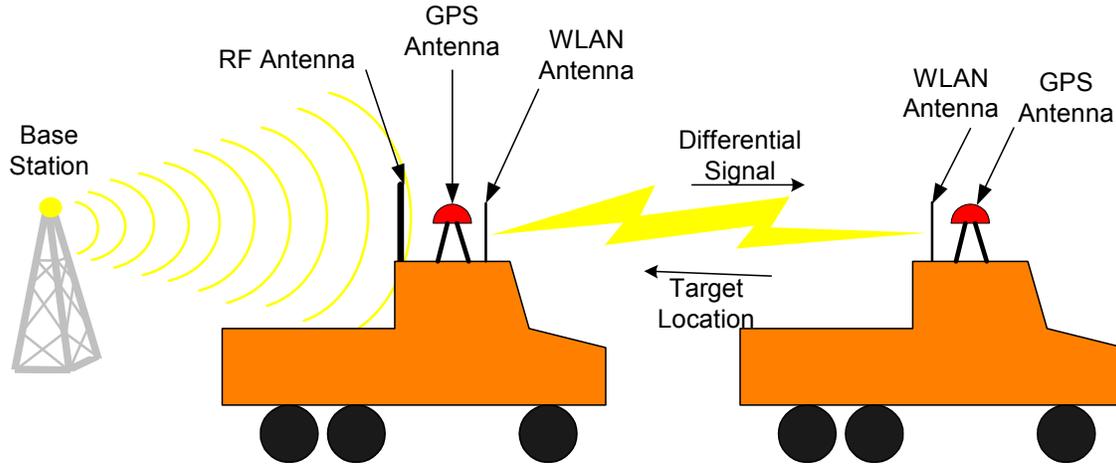


Figure 5-2: Gang plow system diagram. The trailing vehicle receives the differential GPS correction from the base station and sends it to the lead vehicle. The lead vehicle sends back its DGPS position.

To determine the position of the vehicles with respect to the roadway and to identify road obstacles, a GDB was used. It provides road information to the driver assistive technologies in real time, and consists of a database of road objects and a real time database query manager. The road objects can be selected based on the application, but generally include the lane boundaries, lane centers, and road shoulders. For this gang plowing assistive system, lane boundaries, traffic islands, and road shoulders were mapped and built into the database. Using the DGPS vehicle location, the database can be queried in real time to determine the characteristics of the road and of the road furniture. This information can be used to assure the system does not provide assistance that would go against maintaining a trajectory on the drivable road surface.

Longitudinal Controller

The goal of the longitudinal controller is to maintain the driver selected headway time and distance to the lead vehicle. To achieve this, the virtual bumper system was modified to accommodate gang plowing. The virtual bumper is an impedance based collision avoidance system designed to prevent or mitigate collisions. The algorithm was first implemented on a Navistar 9400 tractor and tested at the Mn/Road pavement research facility in the summer of 1998 [13][14]. That system employed radar and lidar to detect the target vehicle. It was adapted to use DGPS in a gang plowing system.

The virtual bumper longitudinal controller is an impedance controller. It consists of a virtual spring and damper connected to the end of a personal space boundary. Incursions into the personal space boundaries by other vehicles impart a virtual force through deflections in the virtual spring and damper (Figure 5-3). The personal space free length is calculated using the difference in lead/target vehicle speed as shown in equation 5.1. The slope of this line is dependent upon the driver selected headway time (T_H) and the

gains of the controller (b,k). The higher the slope of the line, the larger the personal free space boundary.

$$R_{PS} = -(T_H + b/k)\dot{R} + R_H, \quad (5.2)$$

where

R_{PS} = Personal Space Free Length Range

\dot{R} = Range Rate.

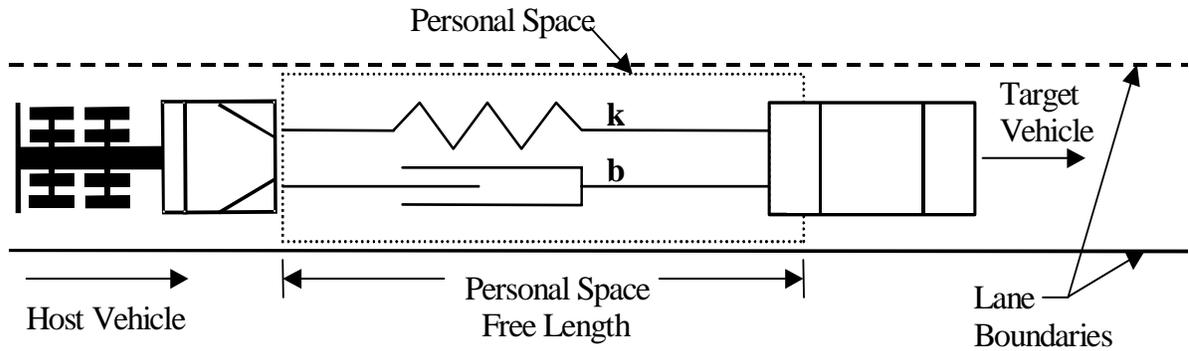


Figure 5-3: The virtual bumper longitudinal controller consists of a virtual spring with spring stiffness k and a virtual damper with damping ratio b .

The virtual force on the truck is calculated using the deflection in the spring and the deflection rate on the damper. This is called the linear force, because it attempts to follow a line in the range – range rate domain.

$$F_{lin} = k(R - R_H) + b\dot{R} \quad (5.3)$$

The admittance controller receives the virtual force as an input and produces a change in desired velocity. This desired change of velocity can then be added to the desired traveling velocity, which is the velocity of the lead vehicle. The virtual force is normalized by the host vehicle mass (m) and then integrated, as shown by the admittance controller transfer function.

$$H_{lin}(s) = \frac{1}{ms} \quad (5.4)$$

As mentioned previously, the linear force controller attempts to follow a linear trajectory in the range – range rate domain. It turns out that constant acceleration in this domain is

a parabolic function [14]. At low accelerations, this behavior produces acceptable ride comfort, but when higher accelerations are needed to achieve the desired headway, the linear force controller can produce very large decelerations that create an uncomfortable situation for the driver. Constant deceleration is more comfortable and natural. Thus, a second force is defined which attempts to produce a constant deceleration when higher forces are needed. This second, nonlinear force is used when larger decelerations are needed to track the lead vehicle. The nonlinear force is computed via

$$F_{nl} = \max \left\{ \begin{array}{l} -\frac{m\dot{R}^2}{2(R - R_H)} \\ -mD_{\max} \end{array} \right\}, \quad (5.5)$$

where

F_{nl} = non-linear force

D_{\max} = Maximum deceleration capability of the host

The nonlinear force attempts to follow a parabolic curve in the range – range rate domain. The bottom term of Equation 5.5 limits the force to the maximum deceleration of the host vehicle.

Speed Controller

The speed controller accepts speed commands from the virtual bumper controller and attempts to follow them. It uses the throttle and brakes of the trailing vehicle to accomplish this goal. The speed controller architecture consists of two sub-controllers: the throttle controller and the brake controller. This is illustrated in Figure 5-4 below.

The throttle controller produces throttle commands to the engine controller through the remote throttle port located in the vehicle cab. The fuel injection “rack” command is proportional to the voltage applied at the throttle port.

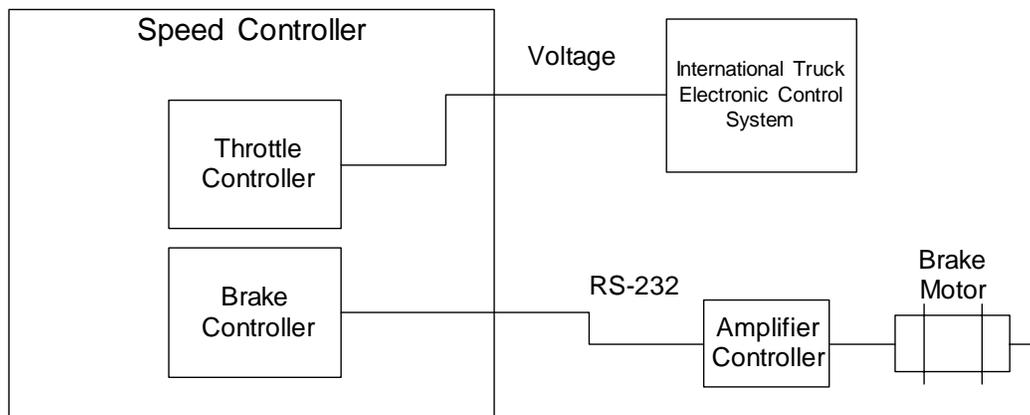


Figure 5-4: Speed controller system diagram.

The throttle controller is a Proportional plus Derivative (PD) controller with a feed forward term. The proportional term acts on the speed error between the actual and desired speed. The derivative term acts on the speed error rate of change. The feed forward term uses a map of speed vs. throttle input to obtain the desired throttle position.

The feed forward map was determined experimentally. The truck electronic control system was given varying fixed voltages. The speed of the truck was measured and the result is shown in Figure 5-5. The scatter plot with the diamond symbol is the experimentally measured speed values for each voltage. A fifth order polynomial function was fitted to the data to arrive at the mapping function (solid line). An R-value of 0.9888 suggests that the polynomial matches well with the experimentally derived map.

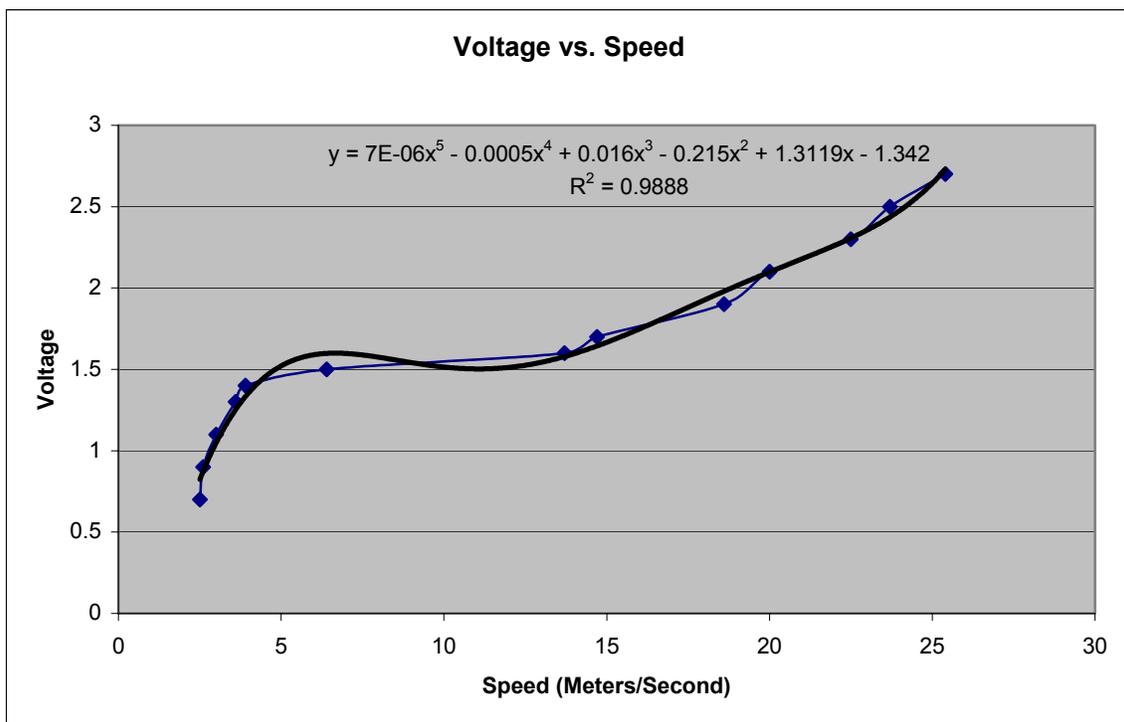


Figure 5-5: Feed forward speed map showing voltage vs. speed.

This controller had one limitation. When the desired speed was just above a speed at which the automatic transmission would typically shift a gear, the output of the throttle controller was insufficient to “force” a gear change. The truck’s electronic control system needed a higher voltage in order to “kick” the transmission into selecting the next gear. In order to solve this limitation, an integral term was added. The integral term is initiated when the speed error reaches a threshold and at low speeds where gear shifts have a greater affect on speed. The integral term integrates the speed error and provides the voltage rise needed to force the engine’s electronic control system to increase the speed of the engine so that the automatic transmission shifts to the next higher gear.

The brake controller controls a motor and winch to depress the brake pedal. (Using a cable allows the driver to override the brake controller and apply more brake pressure as desired.). The motor is mounted in the engine compartment just behind the firewall. At the end of the motor shaft is a winch and cable that runs though the firewall and is attached to the brake pedal (Figure 5-6 and Figure 5-7). Motor armature rotation in the counter clockwise position applies the brake, while reversing the motor direction releases the brake.

The brake controller is a proportional controller that is activated when the speed error rises above a defined threshold.

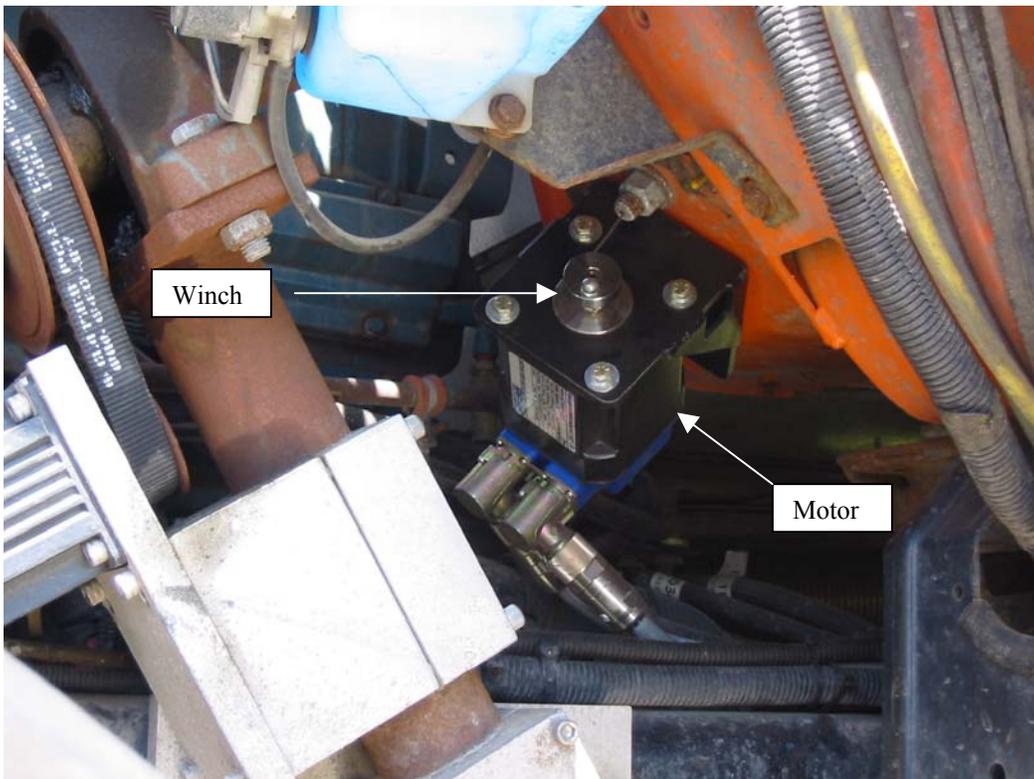


Figure 5-6: Brake motor attached to firewall with winch and cable.



Figure 5-7: Cable runs through firewall and attaches to brake pedal.

Lateral Controller

The lateral controller maintains a desired lateral position with respect to the map. The desired lateral position of the trailing vehicle is obtained by the lead vehicle's position and the driver selectable offset distance. The position of the lead vehicle is compared with the closest LaneCenter object in the GDB to derive its lateral position with respect to its lane of travel. Similarly, the trailing vehicle's position is compared with the map to obtain its lateral position with respect to its lane of travel. Comparison between the lead and trailing vehicles' lateral positions, as well as the driver selected offset, produces the lateral position error.

The lateral controller is a PD controller with a feed forward term (Equation 5.6). The proportional term attempts to reduce the lateral error. The derivative term adds damping compensation in order to reduce oscillations. The feed forward term uses the GDB lane information to point the steering wheel in the direction of the road. The wheelbase (W_b) of the trailing vehicle was measured, and the radius of curvature of the road was obtained from querying the GDB (Figure 5-8). The desired steered wheel angle δ was determined by

$$\delta = k_p y + k_d \dot{y} + k_{ff} \frac{W_b}{r}, \quad (5.6)$$

where y is the lateral error, \dot{y} is the lateral error rate, W_b is the vehicle wheelbase, r is the corner radius, k_p is the proportional gain, k_d is the derivative gain, and k_{ff} is the feed forward gain.

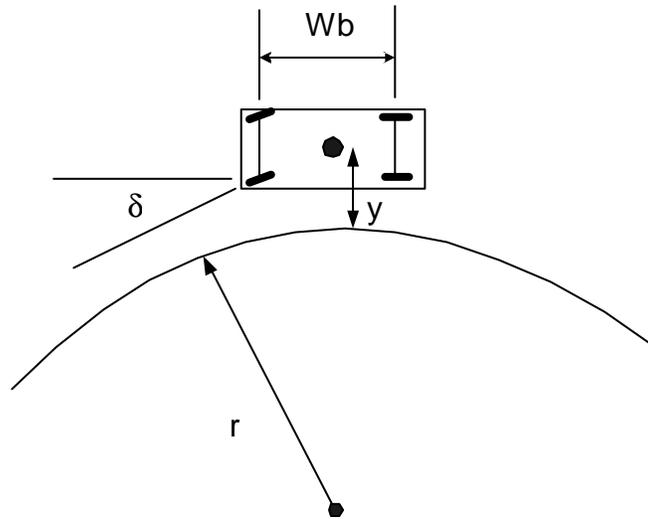


Figure 5-8: The correction term (δ) is obtained from the lateral error (y), the radius of curvature of the road (r), and the wheel base (Wb).

The lateral controller obtains the correct lateral position by sending position commands to the steering motor attached to the steering wheel column (Figure 5-9). The torque of the motor is software limited. The limit is at the point where sufficient torque is applied so that the steering wheel can turn at low vehicle speeds. If necessary, the driver can easily overpower the motor using the steering wheel.

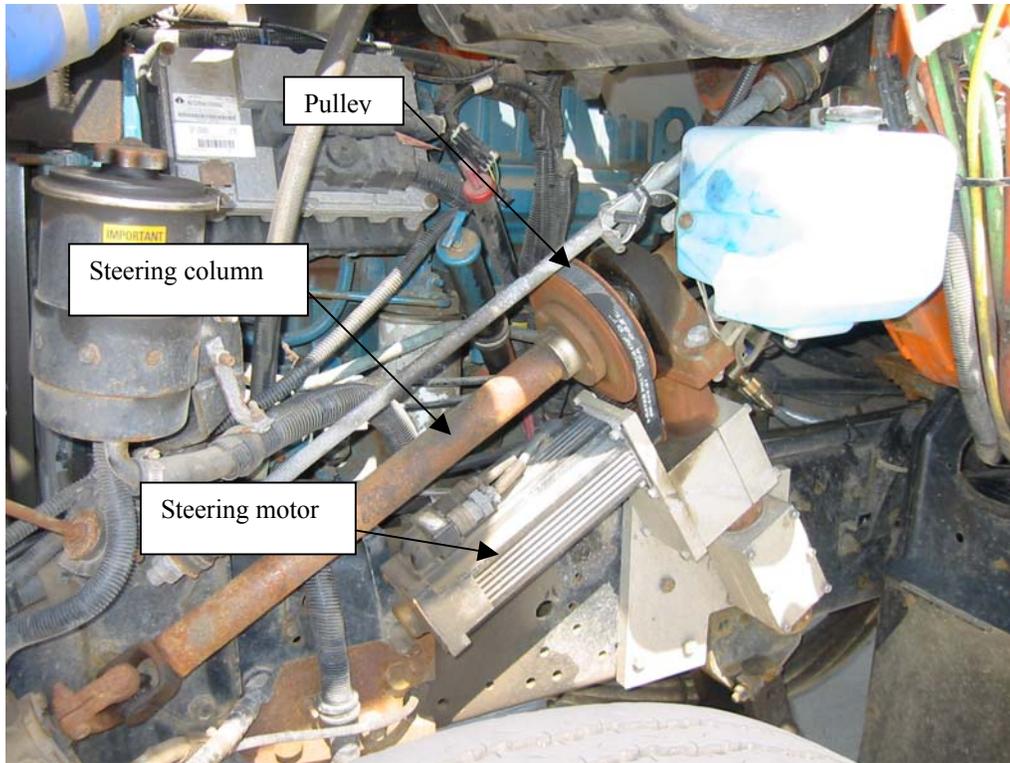


Figure 5-9: Steering motor attached to steering column.

Results

The gang plowing system was implemented on two International snowplows and tested at the Mn/Road research facility near Albertville, Minnesota. The lead vehicle was equipped with DGPS and a wireless Local Area Network (LAN) (implemented using IEEE 802.11b hardware) station adapter which allowed real time communication with the trailing vehicle. The block diagram for the system in the lead vehicle is shown in (Figure 5-10). Two PC-104 computers were used in the trailing vehicle to implement the controllers (Figure 5-11). The first computer received the DGPS correction signal via RF modem. A DGPS receiver on the lead vehicle provided its position via WLAN device to the trailing vehicle. This information was sent via a LAN to the second computer, which performed the control operations. The second computer sent commands to the brake and steering amplifiers which controlled the motors. An analog board was used to provide a voltage to the truck's throttle controller.

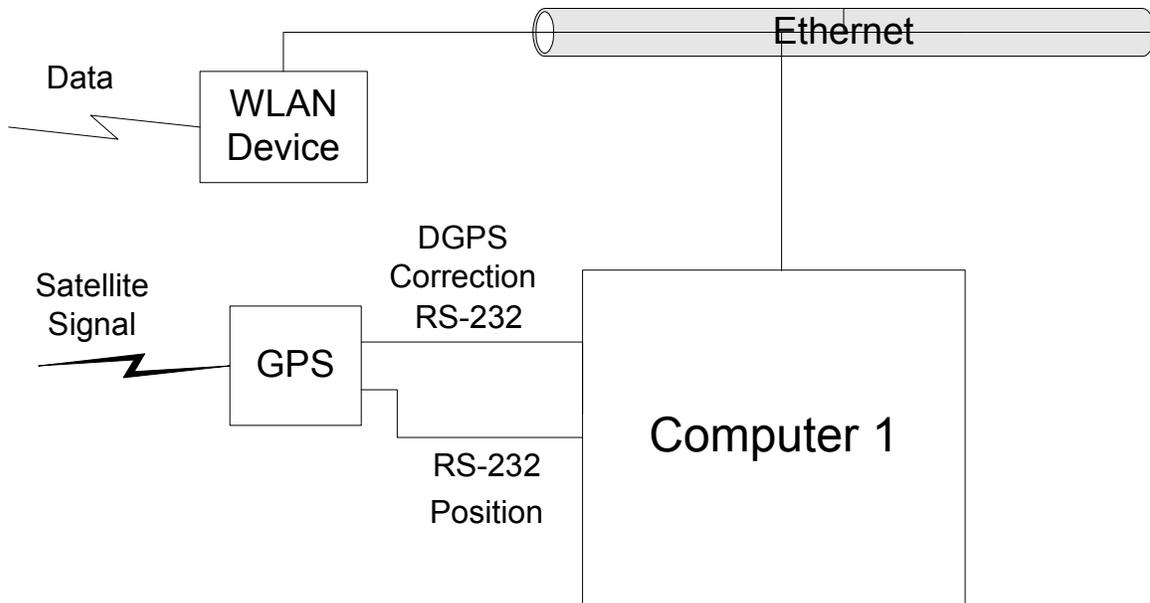


Figure 5-10: Block diagram of lead vehicle.

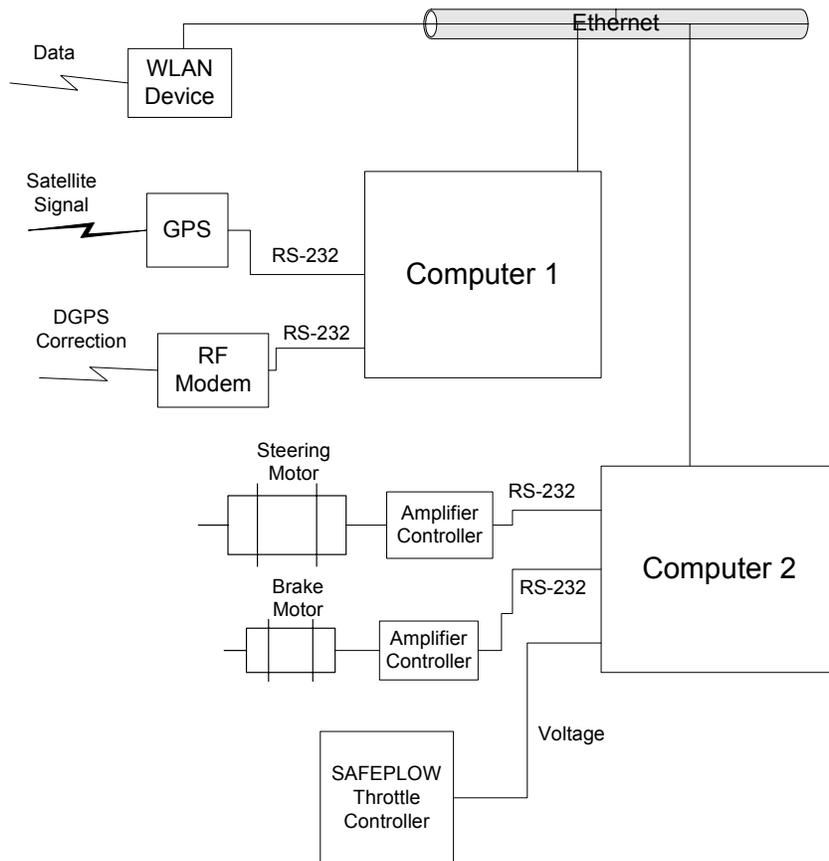


Figure 5-11: Block diagram of trailing vehicle.

Lateral Controller Results

In order to test various plowing situations and to document the ability of the trailing vehicle to follow the lead vehicle, the lead snowplow was driven on the Mn/Road track while continually changing its lateral position and speed. The trailing vehicle was under computer control. The speed of the lead vehicle varied from 5 to 15 m/second (11 mph to 33 mph). The headway time was set to 2.0 seconds and the lateral offset was set to 0. Figure 5-12 shows the results of the lateral tracking capabilities on a straight section of the track. On the upper portion of the plot, the lateral position of the lead and follower were plotted against time. The lateral errors are shown on the lower plot. The discontinuities on both plots happened when the vehicle switched lanes. The lateral position was calculated from the center of each lane, with a positive lateral position to the left of center of the lane. When a vehicle switched from lane 1 (right lane) to lane 2 (left lane), the sign of the lateral error changed. Ignoring the discontinuities, the lateral errors in Figure 5-12 are reasonable, considering that the lead vehicle was swerving constantly between lanes. The trailing vehicle did overshoot when the lead vehicle finished its lane change (i.e., at 7,275, 7,296, and 7,320 seconds), but the lateral controller recovered quickly, and the error decreased within a few seconds. The errors were less than 1 m in all cases.

In this example, this lateral overshoot placed the outside tires of the trailing vehicle slightly over the boundary of the lane. It should be noted, however, that these maneuvers would be considered aggressive for gang plowing operations. With less aggressive maneuvers, the magnitude of the lateral overshoot would be considerably smaller.

It should also be noted that the following vehicle is using DGPS as its only feedback device, and only position signals from the lead truck are sent back to the trailing vehicle. The lateral following performance of the trailing vehicle would be improved with feedforward (steering wheel position from the lead vehicle) and additional feedback (lateral acceleration) from the trailing vehicle. This improved performance would come at a much higher cost, however. Additional feedforward and feedback loops will be investigated in follow-on work.

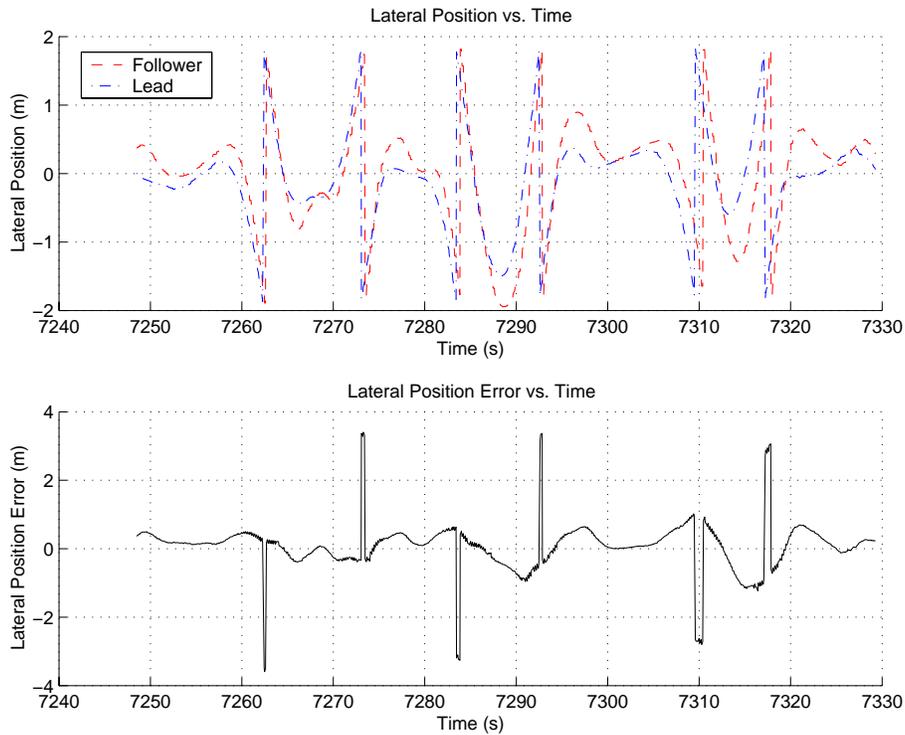


Figure 5-12: Lateral position plot for a straight section of road. Initiation of lane change maneuvers is indicated by the spikes in the lateral position error plot in the lower graph.

The same experiment was repeated using a long sweeping curve with an “S” section. Figure 5-13 shows the curved section of road at the Mn/Road pavement research facility. As can be seen, the radius of curvature is small, and the curvature switches at the end of the curve immediately before the straight section begins. Even with this challenging road geometry and the lead vehicle swerving within the curve, the lateral controller kept the trailing vehicle within 1 m of the lead vehicle (excluding discontinuities) for the entire curved section (Figure 5-14). Again, the trailing vehicle overshoot the lead vehicle but recovered in a timely manner.



Figure 5-13: Curved portion of Mn/Road pavement research facility.

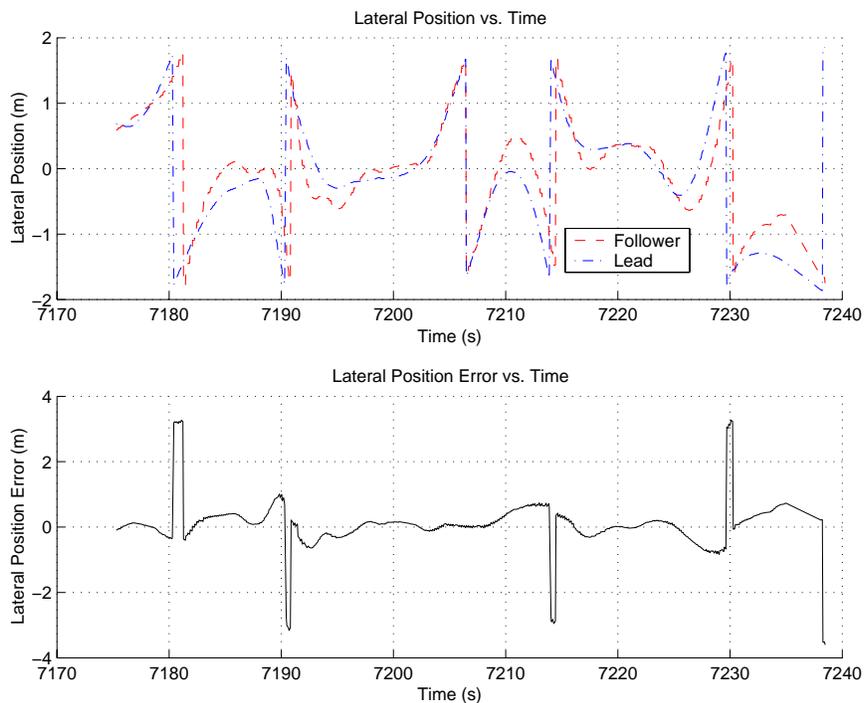


Figure 5-14: Lateral position plot for a curved section of road.

Longitudinal Controller Results

The longitudinal controller was tested by varying the speed of the lead vehicle. Figure 5-15 shows the results of one experiment on the Mn/Road pavement facility. The plot shows the lead and trailing vehicle speed on the upper plot and the speed error on the lower plot. The lead vehicle varied its speed between 0 and 17 m/second (0 to 38 mph). The following vehicle was able to match the speed of the lead. The errors were generally

within 2 m/second (4.5 mph). The speed error plot shows some oscillations when the target maintained a somewhat constant speed at several intervals during the experimental run. These oscillations are less than 1 m/second (2.2 mph) and do not cause significant acceleration that could annoy the driver. The oscillations were observed to be more audible than tactile as the rack controller changed the fuel flow rate to the engine, changing engine RPM. This lack of tactile sensation can be attributed to a “loose” transmission torque converter design that maintains a relatively soft coupling between the engine and transmission.

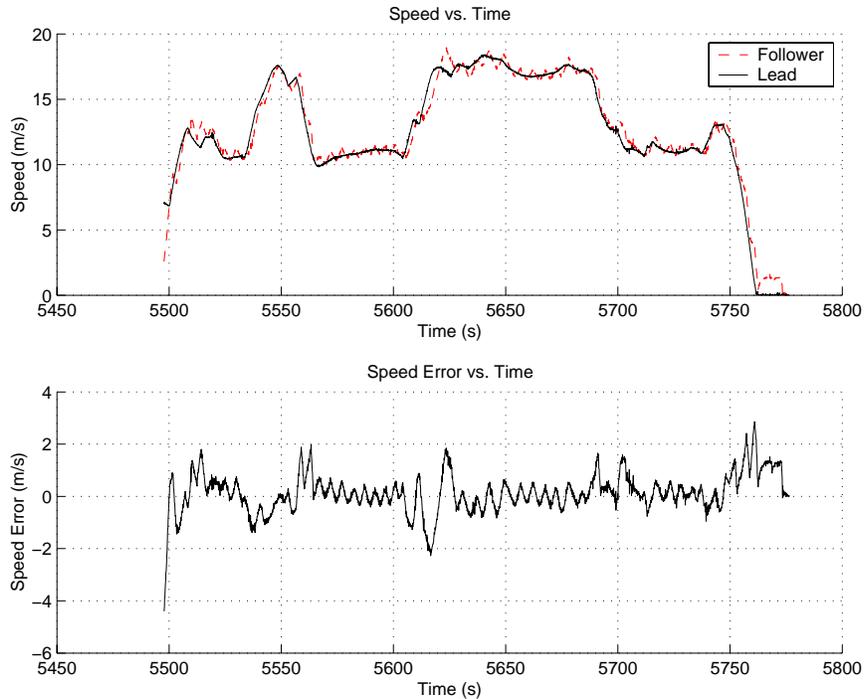


Figure 5-15: Longitudinal tracking – speed.

While the longitudinal controller was good at matching the speed of the lead vehicle, it is actually a headway controller, not a speed controller. Of course, it achieves headway control using speed variations. Figure 5-16 shows the results of the same experiment but the headway (range) and desired headway (R_h) are plotted vs. time. The errors are again shown on the lower plot. For steady state conditions, the longitudinal controller was able to achieve the desired headway. During transitions, the error was generally within 5 m (16 ft.) until the end of the maneuver when the target vehicle came to halt. The desired headway was a function of the lead vehicle velocity. When the lead vehicle’s speed dropped quickly, the desired headway also dropped. It took the longitudinal controller some time to slow the following vehicle.

It is important to note that no feedforward from the lead truck was provided back to the trailing vehicle. Longitudinal performance could be enhanced by providing brake and throttle positions from the lead vehicle to the trailing vehicle via the wireless LAN. This

feedforward control would allow the trailing vehicle to apply throttle and brake commands simultaneously with the lead vehicle, substantially improving longitudinal control. Provision of these feedforward terms from the lead vehicle would come at a greater cost, however. The performance demonstrated by this system is sufficient for gang plowing operations.

The final range at the end of the experiment (after 5,750 seconds) is less than the desired headway range. The headway range actually should go to zero as the lead vehicle's speed goes to zero. Obviously, it is difficult to obtain such a small headway. A buffer headway was added so that the trailing vehicle would stop with a comfortable headway when the lead vehicle stops. However, if the following vehicle failed to achieve the desired final headway when it stopped, it would accelerate and stop again. This caused an unpleasant ride. To improve ride quality, when the lead vehicle came to a stop, the headway controller was turned off, and the trailing vehicle was allowed to coast to a specified headway. Once the headway target was reached, the trailing vehicle brakes were applied to a full stop. This provided the operator with a more pleasant ride.

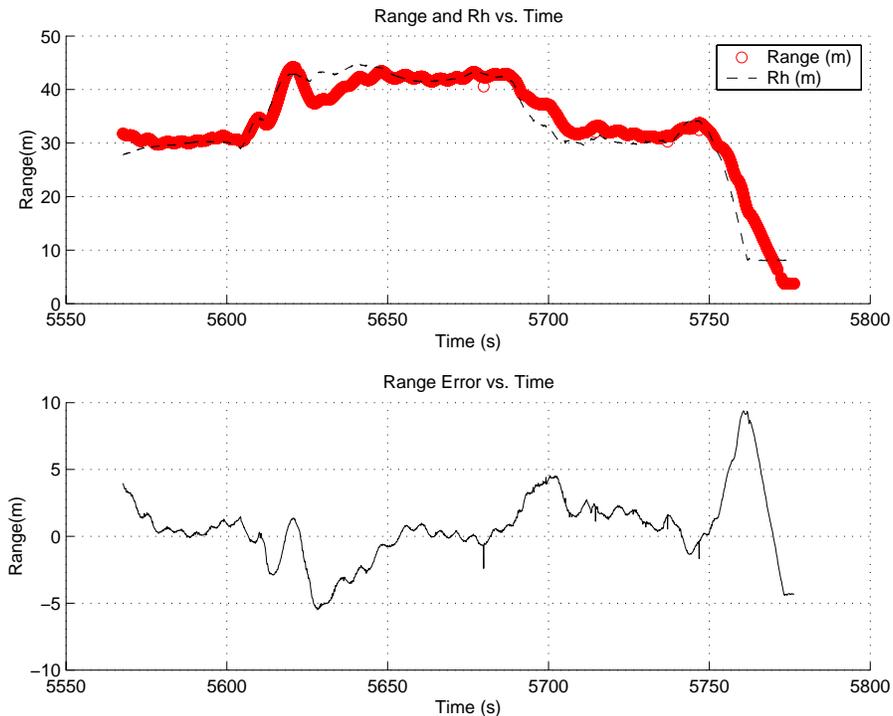


Figure 5-16: Longitudinal tracking - headway.

Conclusions

The gang plow assistive system was implemented on an International snowplow and showed promise as a driver assistive system to help snowplow operators gang plow. The lateral errors between the lead and following vehicle were less than 1 m (3.2 ft.), even

when the lead vehicle was swerving back and forth. There was some overshoot in lateral position, but it settled within a few seconds. The longitudinal controller maintained the desired headway within 5 m (16.2 ft.) most of the time and peaked at 9 m (29.5 ft.) when the lead vehicle came to a quick stop. The speed plot revealed some small oscillation, but the effect on driving comfort was negligible.

The lateral position at speeds over 15 m/second (33.5 mph) showed increased oscillations. The oscillations could cause the controller to become unstable if the steering wheel was left untouched. However, placing a hand on the steering wheel was enough to damp the system. Furthermore, snowplows usually plow under 15 m/second (33.5 mph).

Future Work

Damping the oscillations experienced when traveling over 15 m/second (33.5 mph) should be addressed. Even though snowplows generally plow snow under 15 m/second (33.5 mph), the system should be able to function at higher speed. Potential solutions include adding an Inertial Measurement Unit (IMU) to close a control loop around lateral acceleration. The steering wheel angle and/or lateral acceleration of the lead vehicle could also be used to damp the lateral acceleration.

The system was developed and tested with a two snowplow platoon to prove the concept. Eventually, more than two vehicles are likely to be employed. Errors between each successive vehicle are additive and a lag develops between the lead vehicle and the last vehicle in the formation. The platoon of vehicles acts like an accordion, as the lead vehicle slows and accelerates, and the errors propagate down the platoon. The oscillatory behavior of this platoon increases with additional vehicles. String stability can be achieved, however, by communicating the lead vehicle's position and speed to the rest of the vehicles in the platoon. Since a WLAN was employed to communicate vehicle data, more nodes can easily be added to the network, and all vehicles can communicate with each other. It is then possible to send the lead vehicle data to all vehicles in the platoon.

The GDB contains the location and attributes of various road elements. Using the road shoulder or road island objects, it is possible to make the system "aware" of the various road features so it does not provide feedback that would drive the snowplow off the road. The GDB already has the needed information, but time ran out before this feature could be implemented.

The next logical step would be to test the system on an actual road. The robustness and performance of the system can be measured under real operating conditions. Driver feedback can also be gathered to improve the utility of the driver assistive system.

References

- [1] H.M. Lim, B. Newstrom, C. Shankwitz, and M. Donath, "A Head Up Display Based on a DGPS and Real Time Accessible Geo-Spatial Database for Low Visibility Driving", In *Proceedings of the 12th International Meeting of the Satellite Division of the Institute of Navigation (ION GPS '99)* (Nashville, TN: September 1999).
- [2] H.M. Lim, B. Newstrom, C. Shankwitz, and M. Donath, "A Conformal Augmented Head Up Display for Driving Under Low Visibility Conditions", In *Proceedings of the 5th International Symposium on Advanced Vehicle Control (AVEC'2000)* (Ann Arbor, MI: August 2000).
- [3] Mobility Assist Device, Patent Application #09/618,613.
- [4] Real Time Hi Accuracy Geo-Spatial Database for Onboard Intelligent Vehicle, Patent Application #10/091,182.
- [5] A. Gorjestani, B. Newstrom, C. Shankwitz, and M. Donath, "Advanced Range Sensor Processing Using DGPS and a Geospatial Database", In *Proceedings of the IEEE Conference on Intelligent Transportation Systems, (ITSC '01)* (Oakland, CA: August 2001).
- [6] H. Kuroda, M. Nakamura, K. Takano, and H. Kondoh, "Fully-MMIC 76GHz Radar for ACC". In *2000 IEEE Intelligent Transportation Systems Conference Proceedings* (Dearborn MI: October 2000).
- [7] K. Hatanaka, T. Kashihara, and Y. Yoshida, "An Electrically Scanning 76GHz FM-CW Radar for Adaptive Cruise Control System", In *Proceedings of the 7th World Congress on Intelligent Transport Systems* (Turin, Italy: November 2000).
- [8] T. Harada, Y. Asano, S. Oshima *et al*, "Phase-Comparison Monopulse Radar with Switched Transmit Beams for Automotive Application", In *Proceedings of the 10th Intelligent Transportation Society of America* (Boston MA: May 2000).
- [9] B. Newstrom. "Real Time High Accuracy Geo-Spatial Database for Onboard Intelligent Vehicle Applications," Master's thesis, Mechanical Engineering Department, University of Minnesota, October 2000.
- [10] P. Riekert and T.E. Schunck, "Zur Fahrmechanik des gummibereiften Kraftfahrzeugs (in German)," *Ingenieur Archiv*, vol. 11, (1940) p. 210-224.
- [11] H. Peng and M. Tomizuka, "Vehicle Lateral Control for Highway Automation", In *Proceedings of the American Control Conference* (San Diego, CA: 1990), p788-794.

-
- [12] A. Gorjestani and M. Donath, "Longitudinal Virtual Bumper Collision Avoidance System Implemented on a Truck," In *Proceedings of the 6th ITS World Congress* (Toronto, Canada: November 1999).
- [13] A. Gorjestani, L. Alexander, and M. Donath, "Radar Based Longitudinal Virtual Bumper Collision Avoidance System Implemented on a Truck", Final Report, Minnesota Dept. of Transportation, Report #2000-07 (St. Paul, MN: 1999).
- [14] A. Gorjestani, C. Shankwitz, and M. Donath, " Impedance Control for Truck Collision Avoidance", In *Proceedings of the 2000 American Control Conference* (Chicago, IL: June 2000).