



THE UNIVERSITY OF TEXAS AT AUSTIN
CENTER FOR TRANSPORTATION RESEARCH

TECHNICAL REPORT 0-6877-1

TxDOT PROJECT NUMBER 0-6877

COMMUNICATIONS AND RADAR-SUPPORTED
TRANSPORTATION OPERATIONS AND PLANNING:
FINAL REPORT

Michael Motro
Taewan Kim
Rahi Kalantari
Joydeep Ghosh
Junil Choi
Anum Ali
Preeti Kumari
Enoch Yeh
Yuyang Wang
Duy H. N. Nguyen
Mohammed Eltayeb
Robert Heath W. Jr.
Patricia S. Lavieri
Chandra R. Bhat

CENTER FOR TRANSPORTATION RESEARCH
THE UNIVERSITY OF TEXAS AT AUSTIN

Technical Report Documentation Page

1. Report No. FHWA/TX-16/0-6877-1		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle Communications and Radar-Supported Transportation Operations and Planning: Final Report			5. Report Date May 2016; Published March 2017		
			6. Performing Organization Code		
7. Author(s) Michael Motro, Taewan Kim, Rahi Kalantari, Joydeep Ghosh, Junil Choi, Anum Ali, Preeti Kumari, Enoch Yeh, Yuyang Wang, Duy H. N. Nguyen, Mohammed Eltayeb, Robert W. Heath, Jr., Patricia S. Lavieri, Chandra R. Bhat			8. Performing Organization Report No. 0-6877-1		
9. Performing Organization Name and Address Center for Transportation Research The University of Texas at Austin 1616 Guadalupe Street, Suite 4.202 Austin, TX 78701			10. Work Unit No. (TRAIS)		
			11. Contract or Grant No. 0-6877		
12. Sponsoring Agency Name and Address Texas Department of Transportation Research and Technology Implementation Office P.O. Box 5080 Austin, TX 78763-5080			13. Type of Report and Period Covered Technical Report March 2015–August 2016		
			14. Sponsoring Agency Code		
15. Supplementary Notes Project performed in cooperation with the Texas Department of Transportation and the Federal Highway Administration.					
16. Abstract This project designs a conceptual framework to harness and mature wireless technology to improve transportation safety, with a focus on frontal collision warning/collision avoidance (CW/CA) systems. The framework identifies components of the technology and its capabilities, and how these components can be integrated to improve transportation safety.					
17. Key Words Connected vehicles, traffic safety, machine learning			18. Distribution Statement No restrictions. This document is available to the public through the National Technical Information Service, Springfield, Virginia 22161; www.ntis.gov.		
19. Security Classif. (of report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of pages 162	
				22. Price	



**THE UNIVERSITY OF TEXAS AT AUSTIN
CENTER FOR TRANSPORTATION RESEARCH**

Communications and Radar-Supported Transportation Operations and Planning: Final Report

Michael Motro
Taewan Kim
Rahi Kalantari
Joydeep Ghosh
Junil Choi
Anum Ali
Preeti Kumari
Enoch Yeh
Yuyang Wang
Duy H. N. Nguyen
Mohammed Eltayeb
Robert Heath W. Jr.
Patricia S. Lavieri
Chandra R. Bhat

CTR Technical Report:	0-6877-1
Report Date:	May 31, 2016; Published March 2017
Project:	0-6877
Project Title:	Communications and Radar-Supported Transportation Operations and Planning (CAR-STOP)
Sponsoring Agency:	Texas Department of Transportation
Performing Agency:	Center for Transportation Research at The University of Texas at Austin

Project performed in cooperation with the Texas Department of Transportation and the Federal Highway Administration.

Center for Transportation Research
The University of Texas at Austin
1616 Guadalupe, Suite 4.202
Austin, TX 78701

<http://ctr.utexas.edu/>

Disclaimers

Author's Disclaimer: The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official view or policies of the Federal Highway Administration or the Texas Department of Transportation (TxDOT). This report does not constitute a standard, specification, or regulation.

Patent Disclaimer: There was no invention or discovery conceived or first actually reduced to practice in the course of or under this contract, including any art, method, process, machine manufacture, design or composition of matter, or any new useful improvement thereof, or any variety of plant, which is or may be patentable under the patent laws of the United States of America or any foreign country.

Notice: The United States Government and the State of Texas do not endorse products or manufacturers. If trade or manufacturers' names appear herein, it is solely because they are considered essential to the object of this report.

Engineering Disclaimer

NOT INTENDED FOR CONSTRUCTION, BIDDING, OR PERMIT PURPOSES.

Project Engineer: Chandra Bhat
Professional Engineer License State and Number: Texas No. 88971
P. E. Designation: Research Supervisor

Acknowledgments

The authors would like to thank the undergraduate UT students that helped in the data collection portion of the study: Bryce Arden, Ryan Meek, Phillip Lemons, Jonathon Reynolds, Kevin Rosen, and Robert Syvarth. Under the supervision of Dr. Heath, they developed an assembling system to allow for the equipment (DSRC, radar, and camera) to be connected and powered by any vehicle. They also developed the code to extract and visualize the data from these sensors and provided the images presented in the preliminary data analysis section. We would also like to acknowledge the help from Cooper D. Raterink with all the data collection process.

Table of Contents

Chapter 1. Project Overview	1
1.1 Introduction.....	1
1.2 Report Structure and the Three Main Tasks	3
Chapter 2. Developing Improved Collision Detection and Avoidance Algorithms for Cars	5
2.1 Introduction.....	5
2.2 Predicting Future Motion of Vehicles	5
2.2.1 Vehicle Trajectory Data.....	5
2.2.2 Physical Models	6
2.2.3 Road Context Models	8
2.2.4 Data-based Models.....	10
2.3 Extracting Collisions from Trajectories.....	11
2.3.1 Quickly Approximating Time-to-Collision in Intersections via Collision Zones.....	12
Chapter 3. Develop Improved Collision Detection and Avoidance Algorithms between Cars and Non-Motorized Road Users	13
3.1 Introduction.....	13
3.2 Computer Vision Oriented Methods for Pedestrian and Bicyclist Detection.....	13
3.2.1 Challenges in Pedestrian and Bicyclist Detection	13
3.2.2 Computer Vision for Pedestrian and Bicyclist Detection	14
3.3 Deep-Learning-Based Algorithm with Sensor Fusion.....	16
3.3.1 Fast R-CNN for Pedestrian and Bicyclist Detection.....	16
3.3.2 Fusion of Radar with Fast R-CNN.....	16
3.4 Simulation on KITTI Object Detection Benchmark.....	18
3.4.1 KITTI Object Detection Benchmark	18
3.4.2 Analysis on Region Proposals from Selective Search	19
3.4.3 Fast R-CNN and LIDAR Fusion: Simple ROI Rejection	19
3.4.4 Fast R-CNN and LIDAR Fusion: Supplementary Region Proposal.....	23
Chapter 4. Modeling Uncertainty through Simulations.....	27
4.1 Introduction.....	27
4.2 Highway Simulations.....	28
4.2.1 Simulated Collision Avoidance Algorithm.....	29
4.2.2 Results of Highway Simulations.....	30
4.2.3 Determining the Minimum Required Time for Collision Detection.....	31
4.3 Intersection Simulations	31
4.3.1 Analysis of Desirable Sensor Properties.....	32
4.3.2 Analysis of Radar versus DSRC communication	33
Chapter 5. Harnessing the Power of Collective Intelligence.....	35
5.1 Introduction and Motivation	35
5.2 Intersection Traffic Management using COIN	35
5.2.1 System Evaluation	36
5.2.2 Agent-Based Intersection Traffic Management.....	37

5.2.3 Intersection right-of-way.....	39
5.3 Results of Collective Intelligence Training and Simulations	39
5.3.1 Training Specific Maneuvers.....	39
Chapter 6. Developing a Framework for Joint Millimeter Wave Communication and Radar	43
6.1 Introduction.....	43
6.2 Radar System using Low-Frequency WLAN Signals	44
6.2.1 Forward Collision Vehicular RADAR with IEEE 802.11	44
6.2.2 Prototype	44
6.2.3 Results.....	45
6.3 MmWave Joint Radar and Communication System.....	46
6.3.1 Framework for Joint System.....	46
6.3.2 Proposed Receiver Processing Techniques for Enabling Radar Functions	49
6.4 Summary.....	53
Chapter 7. Optimizing the Joint Waveform to Meet Different Objectives	55
7.1 Introduction.....	55
7.2 Defining Different Traffic Scenarios.....	55
7.3 Optimizing between Data Rate and Radar Performance	58
7.4 Summary.....	60
Chapter 8. Incorporating Antenna Arrays.....	61
8.1 Introduction.....	61
8.2 Alternate Array Design.....	61
8.3 Beamforming Design.....	63
8.3.1 Hybrid Beamforming Architecture and Precoder Design.....	64
8.4 MmWave Channel Acquisition	66
8.4.1 Spatial Congruency in sub-6 GHz and mmWave	67
8.4.2 Correlation Translation from sub-6 GHz to mmWave	68
8.4.3 Performance Evaluation.....	70
8.5 Summary.....	71
Chapter 9. Dealing with Security Issues	73
9.1 Introduction.....	73
9.2 Security Risks of Automotive Radar	73
9.2.1 Jamming.....	73
9.2.2 Spoofing.....	74
9.2.3 Interference	75
9.3 Security Risks of DSRC	75
9.3.1 Jamming.....	75
9.3.2 Spoofing.....	76
9.3.3 Interference	76
9.3.4 Confidentiality	76
9.4 Comparing Automotive Radar and DSRC Security	77
9.5 Hacking a Vehicle.....	77
9.6 Incorporating Antenna Arrays to Enhance Security in mmWave Communication	
Links	78
9.6.1 Proposed Security Technique	78

9.7 Summary.....	81
Chapter 10. Preliminary Field Test of Automotive Radar and Communication Systems.....	83
10.1 Introduction.....	83
10.2 Data Collection Scenarios.....	85
10.2.1 Equipment.....	89
10.3 Preliminary Data Analysis.....	91
10.3.1 Equipment performance and accuracy in uncontrolled environments.....	91
10.3.2 Qualitative Analysis of the joint detection system.....	98
10.3.3 Quantitative analysis of sensor fusion and trajectory prediction.....	100
10.3.4 Qualitative analysis of sensor fusion and trajectory prediction.....	102
10.3.5 System limitations.....	105
Chapter 11. Assessing the performance of the CW/CA system through simulations	107
11.1 Introduction.....	107
11.2 Phases of the Overtaking Maneuver and Definition of Unsafe Maneuvers.....	107
11.3 Simulation Setting.....	109
11.3.1 Distributions of Simulation Variables.....	109
11.3.2 Communication Range.....	110
11.3.3 Packet Error.....	111
11.3.4 Sensor and Estimation Inaccuracy.....	112
11.4 Simulated Data for V2V Communication.....	113
11.5 Performance of the CW/CA system based on V2V communication.....	119
11.6 Comparing V2V and V2I Communication.....	123
Chapter 12. Conclusion	125
References.....	129
Appendix A. Device Properties.....	141
Appendix B. COIN State Space.....	143

List of Figures

Figure 1.1: Architecture of vehicle CW/CA system.....	2
Figure 2.1: Google Maps view of an NGSIM intersection, and plotted trajectories crossing it from the south.....	6
Figure 2.2: Transforming road coordinates	8
Figure 2.3: Seven-second trajectories through an intersection, and motion patterns derived from clustering these trajectories.....	10
Figure 2.4: Visualization of creating collision zone from simulated collisions	12
Figure 3.1: Architecture of pedestrian detection system [Geronimo09].....	15
Figure 3.2: Architecture of Fast R-CNN (Source: [44]).....	16
Figure 3.3: Fusion with mmWave radar (Source: [38]).....	17
Figure 3.4: Architecture of Fast R-CNN with fusion of radar	18
Figure 3.5: Extraction of depth image and edge detection from sparse LIDAR points.....	20
Figure 3.6: Precision-Recall curve of KITTI object detection on validation set trained with CaffeNet (Small size ConvNN). Iteration number 40000. Upper: computer vision only; Lower: computer vision+LIDAR (Simple ROI Rejection)	22
Figure 3.7: Precision-Recall curve of KITTI object detection on validation set trained with VGG_CNN_M_1024 (Medium size ConvNN). Iteration number 40000. Upper: computer vision only; Lower: computer vision+LIDAR (Simple ROI Rejection).....	22
Figure 3.8: Architecture of Fast R-CNN with fusion of LIDAR.....	23
Figure 3.9: Precision-Recall curve of KITTI object detection on validation set trained with CaffeNet (Small size ConvNN). Iteration number 80000. Upper: computer vision only; Lower: computer vision+LIDAR (Supplementary Region Proposal)	24
Figure 3.10: Precision-Recall curve of KITTI object detection on validation set trained with CaffeNet (Small size ConvNN). Iteration number 80000. Upper: computer vision only; Lower: computer vision+LIDAR (Supplementary Region Proposal)	25
Figure 3.11: Pedestrian (upper) and cyclist (lower) detection result examples: ground truth (red), CaffeNet (blue), and CaffeNet+LIDAR (lime). Numbers indicate predicted class probabilities.....	26
Figure 4.1: Three types of collision in a four-way intersection.....	31
Figure 4.2: Labels of Table 4.6 mapped to their corresponding collision scenario.....	33
Figure 5.1: Example of intersection deadlock between three vehicles.....	39
Figure 5.2: Snapshot of left-turn-training simulations.....	40
Figure 5.3: Screenshot of two-entry intersection simulations	42
Figure 6.1: IEEE 802.11 joint radar and communications link setup for measurements.	45
Figure 6.2: RMS range error using IEEE 802.11 packets in a 20 MHz channel with one target.	45

Figure 6.3: Illustration of a traffic scenario for joint automotive radar and vehicular communication systems using IEEE 802.11ad.....	46
Figure 6.4: Illustration of a CPI that consists of M frames, each of K samples.	47
Figure 6.5: Extracted short training field for a SCPHY frame.	47
Figure 6.6: Extracted channel estimation field for a SCPHY frame.....	47
Figure 6.7: (a) The composite ambiguity function of the 128-sample GCP used in the preamble of IEEE 802.11ad. (b) The zero-Doppler cut of the composite ambiguity function of (a).	48
Figure 6.8: The flowchart represents the processing techniques for target detection and range/velocity estimation using IEEE 802.11ad V2V-radar.....	50
Figure 6.9: Probability of detection using different constant false alarm detection rates.....	51
Figure 6.10: MSE of the velocity estimation using the STF of a single and the preamble of the double frames. The numerical results of proposed estimation techniques closely match to the CRLB bounds.	52
Figure 6.11: MSE of the range estimation using the preamble in a single frame based on coarse and fine range estimation algorithms.....	52
Figure 7.1: The passing maneuver on a rural road where the radar signal of the passing car is blocked by the car in front (passed car). In this case, the communication link between the passed car and the passing car can help the latter to detect the incoming traffic.....	57
Figure 7.2: Trade-off between communication data rate and velocity estimation for a fixed size CPI. By increasing the duration of training symbols within a CPI, velocity estimation becomes more accurate with reduced data rate.	59
Figure 8.1: The conceptual figure of mmWave vehicular communications. Multiple mmWave transceivers are deployed on a vehicle to simultaneously establish V2V and V2I communication links.	62
Figure 8.2: Doppler effect in terms of beamwidths	63
Figure 8.3: Channel coherence time vs. beamwidth using directional beams.....	63
Figure 8.4: Proposed hybrid analog/digital architecture with N_T antenna and $N_{RF} \ll N_T$ RF chains.	64
Figure 8.5: An example of resulting beam patterns for a 32-element uniform linear array antenna when using: (a) a fully digital architecture, (b) a hybrid architecture with antenna switches and $N_{RF} = 5$ RF chains, and (c) a hybrid architecture without antenna switches and $N_{RF} = 5$ RF chains.....	66
Figure 8.6: The simulated ray-tracing setup	68
Figure 8.7: The fraction of common paths as a function of frequency.....	68
Figure 8.8: A depiction of collocated Sub-6 GHz and mmWave antenna arrays.....	68
Figure 8.9: CMD as a function of frequency	71
Figure 8.10: EMSE as a function of SNR.....	71
Figure 9.1: Illustration of two different types of attacks as a result of jamming.....	74
Figure 9.2: Illustration of an attack as a result of spoofing.	75

Figure 9.3: Vehicle to vehicle communication with a possible eavesdropper. The transmitting vehicle is communicating with a target vehicle while the eavesdropper tries to intercept the transmitted data.....	79
Figure 9.4: Variance of the interference at an eavesdropper located at a transmit angle of 60 degrees relative to the transmitting vehicle when using different transmission subset sizes M and number of transmit antennas N_T	79
Figure 9.5: Secrecy throughput versus the eavesdropper’s angular direction.	80
Figure 9.6: Secrecy throughput versus the transmission subset size M ; receiver is at a transmission angle of 100 degrees and the eavesdropper is at a transmission angle of 95 degrees relative to the transmitter.....	81
Figure 10.1: Collision avoidance sensor summary. (Source: [159]).	84
Figure 10.2: First data collection scenario—basic car following.	88
Figure 10.3: Second data collection scenario—blind intersection.....	88
Figure 10.4: Delphi Electronically Scanning Radar (2.5) specifications. This equipment can identify up to 64objects simultaneously (Source: Delphi, 2009).	89
Figure 10.5: First vehicle used in the data collection.	90
Figure 10.6: Second vehicle used in the data collection.....	91
Figure 10.7: DSRC reception success rate for the leading vehicle.....	92
Figure 10.8: DSRC reception success rate for the following (tail) vehicle.	93
Figure 10.9: DSRC reception success rate for both vehicles in the blind intersection scenario.	93
Figure 10.10: Distances in which short skips occurred.	94
Figure 10.11: Distances in which long skips occurred.	95
Figure 10.12: Accurate GPS data (blue square) sent through DSRC coincides with radar position detection (green square).	96
Figure 10.13: Inaccurate GPS data (blue square) sent through DSRC shows vehicle in wrong lane.....	96
Figure 10.14: Distances in which long skips occurred.	97
Figure 10.15: Accuracy of radar readings.....	98
Figure 10.16: Blind intersection—DSRC detects approaching vehicle equipped with DSRC (blue square) and not visible to radar system. Radar false positives in green due to light reflex.	99
Figure 10.17: Blind intersection—DSRC detects approaching vehicle equipped with DSRC (blue square) and not visible to radar system. Radar detects vehicle without DSRC equipment (green squares).....	99
Figure 10.18: Blind intersection – DSRC detects approaching vehicle equipped with DSRC. Approaching vehicle is already in the field of vision of the driver but radar does not identify it.	100
Figure 10.19: Pedestrian detection by radar—fails to detect pedestrians that are closer to the vehicle.	100
Figure 10.20: Examples of the vehicles in situation 1 (above) and situation 3 (below). In situation 2, the alternate vehicle is out of view of the camera.....	101

Figure 10.21: Snapshots of video with sensor detections and combined-position estimate.	103
Figure 10.22: Snapshots of video with the future predicted position of the vehicle as a green circle. The snapshots are one second apart from each other.	104
Figure 10.23: Snapshot of video when there is a collision warning. (In the top image the oncoming vehicle is out of camera sight and the triangle on the left shows that the vehicle is off screen to the left).	105
Figure 11.1: Depiction of overtaking maneuver simulations.	108
Figure 11.2: Cumulative distribution of collision maneuvers versus initial distance.	116
Figure 11.3: The actions of the overtaking assistant versus the actual time-to-collision between passing and oncoming vehicles.	117
Figure 11.4: CW/CA system accuracy measures.	118

List of Tables

Table 2.1: Accuracy of physical models on NGSIM dataset 1	7
Table 2.2: Accuracy of physical models on NGSIM dataset 2	8
Table 2.3: Accuracy of road-context model on NGSIM dataset 2	9
Table 2.4: Accuracy of road-context models on NGSIM dataset 1	9
Table 3.1: Analysis on quality of region proposals provided by selective search (with and without LIDAR fusion) on KITTI object detection Train/Validation set.....	19
Table 3.2: Object detection result on KITTI validation set with only computer vision	21
Table 3.3: Object detection result on KITTI validation set with LIDAR information	21
Table 3.4: Cyclist and pedestrian detection AP (%) of Fast R-CNN with different LIDAR support on KITTI validation set.....	24
Table 3.5: Car detection AP (%) of Fast R-CNN with varying levels of LIDAR support on KITTI validation set.....	25
Table 4.1: Details of constant motion simulation	28
Table 4.2: Details of randomized motion simulation.....	29
Table 4.3: Collision detection for constant-motion simulations	30
Table 4.4: Collision detection for random-motion simulations	31
Table 4.5: CW/CA performance with varying sensor properties, intersection	32
Table 4.6: Intersection collision detection accuracy for specific sensors	33
Table 5.1: Results of COIN simulations on a two-way intersection.....	41
Table 5.2: Results of COIN simulations on a two-entry, four-exit intersection	42
Table 7.1: Requirements for joint radar and communication waveforms at urban intersections	56
Table 7.2: Requirements for joint radar and communication waveforms at rural roadways.....	57
Table 7.3: Requirements for joint radar and communication waveforms at locations with substantial pedestrian and bicycle traffic	58
Table 9.1: Security comparison between automotive radar and DSRC.....	77
Table 10.1: List of traffic situations relevant for assessing the developed technologies.....	87
Table 10.2: DSRC message latency	92
Table 10.3: Message skips (periods when communication was unsuccessful)	94
Table 10.4: Worst recorded errors at the three situations	102
Table 11.1: Descriptive statistics of the simulated data.....	115
Table 11.2: Binary probit model estimation results.....	120
Table 11.3: Simulations results for overtaking CW/CA with V2I.....	123

List of Acronyms

AoA	angle of arrival
AP	average precision
CAN	Controller Area Network
CEF	channel estimation field
CMD	correlation matrix distance
ConvNN	Convolutional Neural Network
COIN	Collective Intelligence
CONOPS	Concept of Operations
CPI	coherent processing interval
CRLB	Cramer Rao lower bound
CW/CA	collision warning/collision avoidance
DSRC	dedicated short-range communication
EIRP	equivalent isotropically radiated power
EMSE	excess mean squared error
FMCW	frequency modulated continuous waveform
Gbps	gigabit-per-second
GCP	Golay complementary pair
GCS	Golay complementary sequence
IoU	intersection over union
LRR	long-range radar
LS	least squares
MMSE	minimum mean squared error
mmWave	millimeter wave
MRR	medium-range radar
MSE	mean square error
NGSIM	Next Generation Simulation
PR	perception-reaction
R-CNN	Region-based Convolutional Neural Networks
RF	radio frequency
ROI	region of interest
RSU	road-side unit
SBD	symbol boundary detection
SCNR	signal-to-clutter-plus-noise ratio
SCPHY	single carrier physical layer
SNR	signal-to-noise ratio
SRR	short-range radar

STF	short training field
V2I	vehicle-to-infrastructure
V2V	vehicle-to-vehicle
V2X	vehicle-to-everything
WLAN	wireless local area network

Chapter 1. Project Overview

1.1 Introduction

Crashes represent an enormous cost to society in terms of property damage, productivity loss, injury and even death. Identifying the factors that contribute to crashes and developing countermeasures is therefore a priority for transportation and safety professionals. Roadway traffic accidents in 2012 alone resulted in 33,561 fatalities and 2.36 million injuries in the country as a whole [1], and 3,398 fatalities and 230,506 injuries in Texas [2]. This translates to an average of one fatality every 2 hours and 35 minutes and one reportable crash every 75 seconds in Texas. In fact, motor vehicle crashes continue to be the leading cause of death for people aged 11 through 33 years of age, with no deathless days recorded on Texas roadways in 2012 [2].

Sensors can be a critical component of an active safety system in vehicles. At present, automotive radars and visual cameras are the most common safety sensors found in vehicles [3], [4]. By sensing the existence, position, and velocity of other vehicles or objects, automotive radars make it possible to realize adaptive cruise control, blind spot detection, lane change assistance, parking assistance and more. Cameras make driving safer in several ways: eliminating blind spots, serving as virtual mirrors, and providing better night vision with infrared sensors. A combination of radar and cameras may also effectively improve vehicle safety. For example, a computer vision algorithm may be applied to improve the detection of a pedestrian by a visual camera after preliminary detection from a radar sensor.

The number of sensors on vehicles, along with their data rate capacity, is steadily increasing. On average, there are roughly 100 sensors on deployed vehicles today. That number, however, is expected to double by 2020 as vehicles become smarter [5]. Further, recent developments in autonomous vehicles heavily rely on LIDARs, a detection system that works on the principle of radar, but uses light from a laser to generate high-resolution depth-associated range maps. The amount of data generated by LIDARs is similar to the amount generated by conventional automotive cameras, which will further increase the data generated by a vehicle [6]. Thus, active safety algorithms will need to be designed to process more data sources and higher data volumes.

Many existing sensor technologies exhibit a limited sensing range. For example, radar, cameras, and LIDAR provide information only about objects within the line-of-sight of the sensor and fail to address all potential safety concerns, limiting the capabilities of automation in vehicles. An alternative is to employ wireless communication networks, permitting the exchange of information between vehicles. A vehicle with wireless communication capabilities is called a *connected vehicle* [7]. Various safety-related applications with improved automation capabilities are enabled by vehicular connectivity (despite very limited information being provided by current connected vehicle technologies). These include forward collision warnings, do-not-pass warnings, blind intersection warnings, and red-light violation warnings [8]. The development of collision warning/collision avoidance (CW/CA) systems has a direct connection to the improvement of transportation safety and may be able to reduce more than 80% of all annual car crashes [7]. Fast-growing research on advanced driver assistance systems has resulted in various technologies now available on commercial vehicles. CA also has a symbiotic relationship with partial or full autonomy in vehicles, as both require the same hardware and investments and can either operate separately or cooperate. While most autonomous vehicle research centers around emulating human

driving, CW/CA technologies are meant to deal with cases of unusual or dangerous driving, and therefore require a deeper understanding of the vehicle’s surrounding environment.

The general framework of CW/CA technology can be divided into four phases, visualized in Figure 1.1. First, the vehicle gathers information on itself and its surroundings, using on-vehicle sensors or communication with other vehicles/infrastructure. Secondly, this information is used to predict the future positions of this and other vehicles. Based on the predicted positions, the system then decides whether there is a collision or other dangerous situation imminent. Finally, the vehicle must decide whether to perform an action to prevent impending collisions, such as send an audio-visual warning, force the vehicle to brake, or drive along a different route. The first phase is by far the most heavily researched, partially because gathering information is a difficult task and partially because an actual implementation of the latter phases requires implementation of the first. At the same time, the information-gathering phase of CA is not a goal in and of itself, and is only valuable as input to the next phases.

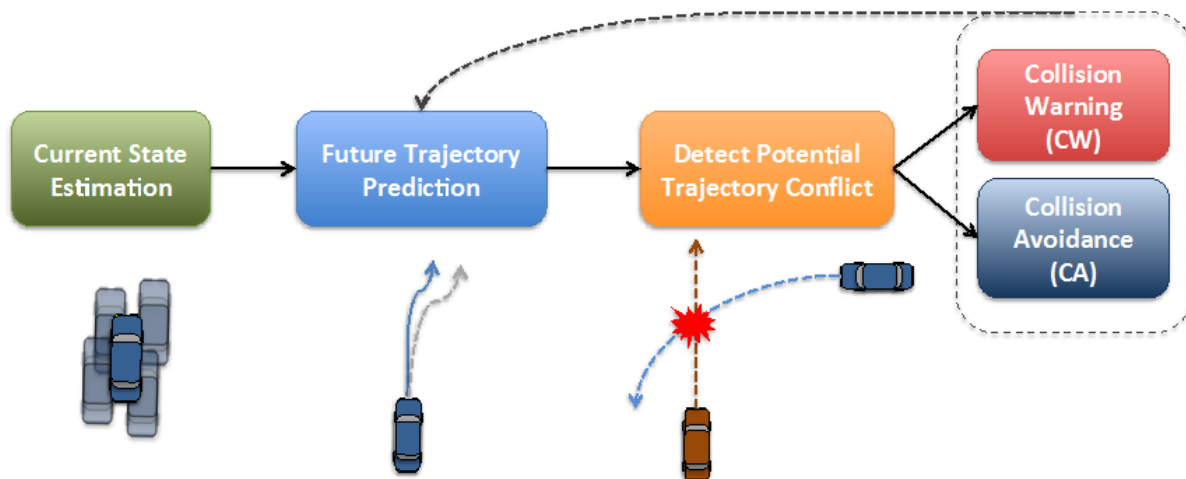


Figure 1.1: Architecture of vehicle CW/CA system

It is natural for future CW/CA technologies to integrate communication with other sensors, as vehicle-to-vehicle (V2V) or vehicle-to-infrastructure (V2I) communication can share detailed road and traffic information to improve the CW/CA algorithm. Vehicle connectivity has two potential benefits. First, if a suitable carrier frequency is chosen, then cars can communicate through non-line-of-sight channels, (e.g., around corners). Second, if a high bandwidth communication link is available, cars can exchange higher raw sensor data rates and achieve full vehicular connectivity. Fully connected vehicles will be beneficial not only for implementing powerful active safety applications but also for cloud-driven, fully automated driving, which has the potential to further improve transportation efficiency and reliability.

The state-of-the-art protocol for connecting vehicles is referred to as *dedicated short-range communication* (DSRC) [9]. Using DSRC, it is possible to implement preliminary V2V, V2I or even vehicle-to-everything (V2X) communication systems. The National Highway Traffic Safety Administration will likely mandate that all new vehicles include DSRC capabilities by 2017 [8]. Although DSRC allows vehicles to exchange messages (including basic sensor information) with a range of up to 1000 meters (ideally), the maximum data rate in practice is only 2–6 Mbps [10].

Fourth generation (4G) cellular systems could be used for V2X communication systems; however, the maximum data rate is still limited to 100 Mbps for high mobility, though much lower rates are typical. Therefore, current technologies cannot sufficiently handle the terabyte-per-hour data rates that can be generated in next-generation vehicles.

Using *millimeter wave* (mmWave) carrier frequencies for communication, it is possible to exploit larger spectral channels and achieve gigabit-per-second (Gbps) data rates (mmWave refers to the spectrum between 30 and 300 GHz). For example, IEEE 802.11ad uses 2.16 GHz of bandwidth in the unlicensed 60 GHz band and supports data rates up to 7 Gbps [11]. Recently, the Federal Communications Commission has proposed authorizing operation in the 28, 37, and 39 GHz of the licensed band and made 64–71 GHz available for unlicensed spectra for mobile use [12], which would facilitate the use of mmWave for various wireless communication scenarios including V2X, leading to the potential for Gbps data rates and realizing raw sensor data exchange among vehicles and infrastructure. Since automotive radars already exploit the mmWave spectrum, it may be also possible to implement joint mmWave radar and communication systems to increase the penetration rate of mmWave V2X capable vehicle in the early deployment stage. Joint systems may also save space within a vehicle and reduce cost and power consumption.

Based on the necessity of improving roadway safety (and thereby also the accessibility and reliability) of the travel experience of Texans, this project presents an interdisciplinary and multidisciplinary effort focusing on research and technology enhancement activities to promote the integration of cutting-edge developments in communications and radar technology with transportation systems. The project is directed toward (a) improving the use of wireless technologies to obtain data from multiple and heterogeneous sources, (b) machine learning and data driven information extraction capabilities that are localized and timely, and (c) bringing information together within traffic contexts to develop safety systems that are effective and customized. More specifically, this project designs a conceptual framework to harness and mature wireless technology to improve transportation safety, with a focus on frontal CW/CA systems. The framework identifies components of the technology and its capabilities, and how these components can be integrated to improve transportation safety.

The project is divided into three main tasks:

- 1) Develop conceptual and functional frameworks for integrated wireless safety systems that incorporate information derived from both communication and radar platforms.
- 2) Advance and develop a new combined communication-radar paradigm for automotive applications using next-generation millimeter wave communication. Consider different performance metrics of radar and communication systems and investigate security issues in vehicular environments.
- 3) Conduct preliminary tests with real data, assess the performance of the proposed system, and develop a Concept of Operations (CONOPS) and requirements for technology deployment.

1.2 Report Structure and the Three Main Tasks

The first task of the project develops conceptual and functional frameworks for integrating communications and radar technologies. It begins in Chapter 2 by analyzing models for predicting the future movements of vehicles. We start with standard algorithms that rely on basic information gathered from sensing or communication. Then we explore more advanced models that rely on

stored or communicated data and we compare the models using real vehicle trajectory data. We also study how to efficiently use these predicted movements for the detection of dangerous driving situations. In Chapter 3, we develop a CW/CA system for non-motorized users—pedestrians and bicyclists—by investigating the sensor fusion approach with a computer-vision-oriented detection algorithm. The use of deep-learning-based detection algorithms and sensors like LIDAR are investigated using a real image dataset. Chapter 4 makes use of computer simulations to analyze the pros and cons of different vehicular sensors. CA systems, relying on various sensor combinations, are simulated in action for several types of scenarios. In the last chapter of Task 1, Chapter 5, a distributed decision-making algorithm, Collective Intelligence, is applied to research cooperative CW/CA systems. This algorithm is tested through simulations of multiple vehicles autonomously navigating an intersection, each following the same algorithm to prevent collisions.

Task 2 starts in Chapter 6 by providing a framework of joint mmWave radar and communication systems for vehicular environments. We first demonstrate the feasibility of joint radar and communication systems using low frequency wireless local area network (WLAN) systems and then demonstrate the feasibility of mmWave joint systems through numerical simulations. Then, in Chapter 7 we identify key scenarios in which radar and communication systems can significantly improve automotive safety mechanisms and demonstrate a method for optimizing the performance of these systems. Chapter 8 details effective methods of antenna and transceiver design for mmWave radar and communication systems. In this chapter we also propose several beamforming and communication channel acquisition techniques. Chapter 9 provides an overview of security issues in vehicular environments and proposes potential ways to improve security in mmWave systems.

In the final task we first assess the performance of the CW/CA system through simulations in Chapter 10. Then, in Chapter 11 we conduct field data collection utilizing vehicles equipped with a joint system of DSRC, radar, and camera. In this report we present preliminary analyses of the collected data. The full data analysis and the CONOPS are left for the final deliverable of the project. Finally, in Chapter 12 we present the overall conclusions of the project and discuss future research.

Chapter 2. Developing Improved Collision Detection and Avoidance Algorithms for Cars

2.1 Introduction

All CW/CA technologies require assumptions about the future. For instance, forward warning systems check for obstacles to the front of the vehicle that are getting closer, with the simple assumption that such obstacles may continue to get closer and collide with the vehicle in a short time. For more complex avoidance or self-driving tasks (such as crossing intersections or safely entering highways), or for vehicles with multiple, simultaneously active CW/CA tasks, it can be useful to view prediction of the future as a distinct and necessary task.

Prediction of the future requires knowledge of the present, but there is currently no clear consensus on what knowledge will be helpful for collision-related prediction. To give two examples, the DSRC basic safety message includes vehicle position, heading, speed, acceleration, and size—in other words, information that is simple to gather using on-vehicle sensors and is also compact enough to communicate easily [13]. On the other hand, a state-of-the-art algorithm that predicts vehicle turns and lane-changes in advance [14] relies on video information of both the outside road and the driver. This understandably adds a lot of prediction power but also requires more hardware, a certain invasion of privacy, and some method of communicating the video or your CW/CA system's predictions to other vehicles on the road.

To establish a consensus on what information vehicles should gather from sensors, cameras, and communication systems, there must be a clearer assessment of the value of each type of information for predicting the future. This in turn requires development of prediction methods that utilize each type of information, and standardized evaluation of each of these methods.

2.2 Predicting Future Motion of Vehicles

Throughout this chapter we refer to a vehicle's changing position over time as its *trajectory*. If the future trajectories of each vehicle are known, then finding any collisions between vehicles in the near future is simply a matter of checking each point in time.

The most common way to predict trajectories is to assume that each vehicle follows a simple physics-based model. This assumption is valid in the short term (because vehicles are bound by the laws of physics) and easily compatible with most sensor fusion methods. Models that incorporate location information and/or previously recorded vehicle paths require far more effort to establish, but might produce more accurate predictions. This section describes and evaluates several physics-based trajectory prediction methods, as well as methods using two additional sources of information—road geometry and previously gathered vehicle data.

2.2.1 Vehicle Trajectory Data

The Next Generation Simulation (NGSIM) Lankershim dataset of vehicle trajectories was used to develop and test the prediction methods [15]. Compared to other driving datasets such as the Strategic Highway Research Program's Naturalistic Driving Study (SHRP2) [16], the NGSIM data is all collected within a specific area. This dataset allows an exact extraction of the geometry of the roads, and also provides multiple observations of different vehicles in the same situation (which is very important for developing and testing prediction methods). Furthermore, the area

covered by the Lankershim dataset contains several types of intersections and high-throughput traffic, which we considered particularly challenging to predict.

Two sets of trajectories from NGSIM are used in this project. One contains 1000 randomly selected 5-second vehicle trajectories from the dataset, most of which are travelling straight along a road. The other contains 5-second cuts of every trajectory that enters a certain intersection from the south (Figure 2.1). The choice of this area of analysis was based on the prediction challenges that occur in the beginning of an intersection or other road entry/exit, as vehicles may change direction as well as brake/accelerate depending on surrounding vehicles or traffic signals. For example, some predictive models perform well on vehicles following a straight and steady path, but fail completely at an intersection. Up to 10 seconds of previous motion are also gathered for each 5-second test trajectory, so that enough prior information can be given to each model.



Figure 2.1: Google Maps view of an NGSIM intersection, and plotted trajectories crossing it from the south

2.2.2 Physical Models

The physics of vehicle motion have been extensively studied for the sake of vehicle design, smooth control, and high-speed driving; there are entire textbooks on the subject [17]. However, trajectory prediction for CW/CA is not intended to capture minute changes or high-speed driving, and therefore is usually achieved with simple models. Following are the three most common models used:

- (i) Constant velocity model:

$$\begin{aligned}x(t + \Delta_t) &= x(t) + v\Delta_t \cos(\theta) \\y(t + \Delta_t) &= y(t) + v\Delta_t \sin(\theta)\end{aligned}$$

- (ii) Constant acceleration model:

$$\begin{aligned}x(t + \Delta_t) &= x(t) + \left(v(t)\Delta_t + \frac{a}{2}\Delta_t^2\right) \cos(\theta) \\y(t + \Delta_t) &= y(t) + \left(v(t) * \Delta_t + \frac{a}{2}\Delta_t^2\right) \sin(\theta) \\v(t + \Delta_t) &= v(t) + a\Delta_t\end{aligned}$$

- (iii) Nearly coordinated turn model:

$$x(t + \Delta_t) = x(t) + \int_t^{t+\Delta_t} v(u) \cos(\theta(u)) du$$

$$y(t + \Delta_t) = y(t) + \int_t^{t+\Delta_t} v(u) \sin(\theta(u)) du$$

$$v(t + \Delta_t) = v(t) + a\Delta_t$$

$$\theta(t + \Delta_t) = \theta(t) + \omega\Delta_t$$

These names were first coined by [18], but the models are based on physics abstractions and have been used for many trajectory prediction tasks [19, 20]. The first two models assume that the vehicle travels in a straight line, while the third assumes that the vehicle follows a steady acceleration and also changes its heading at a steady rate. Intuitively, these parameters correspond to the drivers’ pedaling (gas/brake) and steering wheel angle. As these are the two ways in which the driver controls a vehicle, it is tempting to view this model as the perfect fit for vehicle trajectories. However, physical models do not account for future changes in the vehicle’s motion—for instance, if the road curves or the driver brakes shortly after the present time, these models continue to predict with the same parameters.

In order to ensure that the current position, velocity, etc., of each vehicle are estimated accurately, the Unscented Kalman Filter was applied to the previous 5 seconds of the vehicle’s motion. If the physical model is valid, this filter can near-optimally capture a vehicle’s current state [21]. Moreover, this filter can be easily used to fuse multiple measurements of the same value (such as position) from different sensors, or from sensors and communication.

Table 2.1 shows the predictive power of each model on the randomly gathered trajectories. The errors of the predictions are given in terms of the Euclidean (straight-line) distance in meters between the vehicle’s true position and the predicted position. Errors at three different times are displayed, as predictions farther into the future are expected to be increasingly inaccurate. Note that positional accuracy cannot be immediately translated to the accuracy of a collision detection system, though the relationship between the two is studied in Chapters 2.2 and 4.

Table 2.1: Accuracy of physical models on NGSIM dataset 1

Model	Error predicting 1 second ahead (m)	Error at 3 seconds (m)	Error at 5 seconds (m)
Constant Velocity	0.93	3.52	7.28
Constant Acceleration	0.83	6.53	13.28
Nearly Coordinated Turn	1.51	5.46	10.64

Given that a passenger vehicle’s length is roughly 5 m, it is clear that no model can predict well at 5 seconds. Fortunately, looking 3 seconds or less into the future is generally considered sufficient for simple tasks such as avoiding rear-end collisions. The exact amount of time needed for more demanding CW/CA tasks is not known, though we assume 5 seconds is an upper bound. The simplest physical model, the constant-velocity model, noticeably outperformed the others—primarily because erroneously estimated accelerations or heading changes may strongly influence the predicted trajectory. In other words, the constant-velocity model makes a more conservative estimate. To determine which driving maneuvers are predicted well or poorly, the second dataset—of vehicle trajectories from a single intersection—is used. Table 2.2 gives the predictive error on this dataset, categorized by maneuver through the intersection.

Table 2.2: Accuracy of physical models on NGSIM dataset 2

Maneuver	Model	Error at 1 s (m)	Error at 3 s (m)	Error at 5 s (m)
Straight	Constant Velocity	0.8	2.9	6.1
	Constant Acceleration	1.1	8.0	15.9
	Nearly Coordinated Turn	1.7	6.7	13.8
Lane Changes	Constant Velocity	1.1	4.0	7.9
	Constant Acceleration	1.1	8.6	17.1
	Nearly Coordinated Turn	1.8	6.8	13.3
Turns	Constant Velocity	1.3	4.6	9.3
	Constant Acceleration	1.0	7.6	18.9
	Nearly Coordinated Turn	1.5	5.2	11.5
Don't Cross (stop before intersection)	Constant Velocity	2.0	7.9	15.3
	Constant Acceleration	1.0	3.24	7.5
	Nearly Coordinated Turn	1.3	4.6	8.4

The constant-velocity model performs better on all trajectories except for those that stop before the intersection. We can also see that the trajectories that cross the intersection without turning are the easiest to predict. For all other maneuvers, the prediction errors from the intersection are higher than those from the randomized dataset. This supports the hypothesis that in addition to the current physical state of a vehicle, context about its possible future actions is needed to make accurate trajectory predictions. This context is accounted for in the next section.

2.2.3 Road Context Models

In addition to the current position and motion of a vehicle, there are several outside factors that have a clear impact on its future trajectory, which we collectively term *road context*. These include the geometry and other physical characteristics of the road, traffic signs and signals, and the position of other vehicles or objects on the road. We focus on road geometry.

Road Geometry

Petrich et al. [22] developed a trajectory prediction model where vehicles are assumed to move around the centerline of their respective road. This can be viewed as a change in coordinate systems, from one of global position (e.g., longitude versus latitude) to one of the vehicle's movement along the road versus movement across it (Figure 2.2). We use a very similar model to Petrich, with the main difference being that we incorporate the NGSIM dataset's labelling of vehicle location into the method for finding the vehicle's current road.




Figure 2.2: Transforming road coordinates

In this new coordinate system, the vehicle's motion is no longer dependent on the curvature of the road, making directional parameters such as heading and angular velocity unnecessary.

Thus, the predictive models are constant-velocity and constant-acceleration, with independent along-road and across-road components.

Table 2.3 shows the results of a constant-velocity model, in the road-based coordinates, on the second dataset. While it performs worse than the physical models in the short term, by 5 seconds it outperforms all of them, especially for trajectories that change lanes or turn in the intersection. As a caveat, this intersection has turn-only lanes, and thus there is little ambiguity between vehicles turning and going straight. In situations where road location alone is not enough to determine motion, methods to explicitly locate the correct road or motion may be needed (such as those in the next section).

Table 2.3: Accuracy of road-context model on NGSIM dataset 2

Maneuver	Error at 1 second (m)	Error at 3 seconds (m)	Error at 5 seconds (m)
Straight	2.4	2.4	7.4
Lane Change	2.5	2.6	6.0
Turn	1.0	2.1	8.0
Don't Cross (stop before)	1.3	5.6	11.6

Table 2.4 shows the results of several road-context scenarios using the first dataset. The first is a constant-velocity model as seen above, the second is a constant-acceleration model, and the third is a combination of the standard constant-velocity model from the previous section and this constant-velocity model. This combination is simply made by taking the average of the two predictions, but achieves the best results so far—possibly because the road-context models and physics models each perform well in different situations.

Table 2.4: Accuracy of road-context models on NGSIM dataset 1

Model	Error at 1 second (m)	Error at 3 seconds (m)	Error at 5 seconds (m)
Constant Velocity	1.62	4.30	7.09
Constant Acceleration	1.41	4.71	9.45
Physics CV + Road CV	1.07	3.40	6.36

To make a general statement, road geometry seems to be useful information for improving long-term (3 seconds or more) trajectory prediction. This information could be pre-stored in vehicles, estimated by scanning the road with LIDAR or cameras, or made available when needed by means of V2I or vehicle-to-satellite communications.

Lead Vehicle Distance

The distance to the lead vehicle is generally considered an important aspect of driving, with many driving models—known as *car-following models* [23]—based entirely around this figure. However, these models were usually designed for exploratory or simulation purposes, rather than real-time prediction. Incorporating the presence of more than one vehicle in a predictive model is not simple. For instance, if the trajectory prediction for a vehicle depends on the relative position of its lead vehicle (the closest vehicle in front of it), then the position of the lead vehicle must be

predicted first. If this lead vehicle follows the same model, then its lead vehicle’s trajectory must be predicted first, and so on. Additionally, predictions based on the interactions between vehicles are possibly not appropriate when the ultimate goal of our trajectory prediction is to determine when vehicles are interacting in an unsafe way (for instance, ignoring the safe car-following distance).

2.2.4 Data-based Models

Given the complexity of incorporating context in a trajectory prediction model, an alternative is to use previously gathered vehicle trajectories to train a flexible model. Some data-based models work off physics-based models by estimating key parameters with available data. For instance, it is possible to develop multiple physical models and choose the best model for a given trajectory, by comparing the models’ accuracy on similar stored trajectories. Barrios and Motai [24] use this framework with the constant-velocity and constant-acceleration models, as well as two other models. Other data-based models forgo any assumptions about vehicle motion and entirely predict future trajectories by comparing them with similar trajectories.

Motion Patterns

Some predictive models make the explicit assumption that vehicle drivers choose their future motion from a set of likely trajectories, known as *motion patterns* [25, 26]. Figure 2.3 shows trajectories of various vehicles entering an intersection, and the motion patterns that can be gathered from these trajectories. Note that these correspond to individual lanes, as well as motions such as braking or changing lanes.




Figure 2.3: Seven-second trajectories through an intersection, and motion patterns derived from clustering these trajectories.

The task of trajectory prediction is separated into the problems of finding the motion patterns from data, then deciding which motion pattern a vehicle is currently following. We have developed a motion pattern-finding algorithm that groups vehicle trajectories in similar clusters via the Expectation-Maximization algorithm. Higgs and Abbas [27] used a similar clustering of vehicle movements for exploratory analysis of highway driving. Once the clusters are assembled, the average motion of vehicles in any cluster forms a motion pattern. To predict a vehicle’s future

trajectory, compare its past trajectory to the beginning of each motion pattern and choose the most similar motion pattern. Assume the vehicle's future trajectory is equal to the future values of the motion pattern. It is also possible to compromise between the motion pattern and another model's prediction, in the same way that two models were combined in the previous section.

At the moment, the motion patterns gathered mostly correspond to individual lanes on the road, and their prediction accuracy is very similar to that of the road-context model. The current algorithm also has similar weaknesses to the road-context model. In the same way that a road-based model requires information on every lane, motion pattern gathering requires example trajectories from every stretch of road—in comparing the mathematical difference between trajectories, location is by far the most significant factor. For example, the motion pattern of a vehicle that will turn right in 1 second is very different (mathematically) from one that will turn right in 3 seconds—this is a common issue in clustering time data. Similarly, the motion pattern of a vehicle following a sharp curve is quite distinct from one following a shallower curve. When constructing motion patterns from many locations at once, they will capitalize more on location-related differences than driver actions such as lane changes or braking.

It is much more desirable to reframe motion patterns so that a single pattern may encompass a type of motion—say, a certain type of right turn, or braking a certain degree—across multiple location and time references. Decoupling road context from motion patterns is a goal for continuing research in this project. For instance, one technique that we are applying is dynamic time warping [28], which can measure the similarity of trajectories while ignoring shifts in time between them.

Conclusion

Road geometry information was shown to improve the accuracy of vehicle trajectory prediction, even when applied in a simple way. However, no models could reliably provide less than 5 m of error when predicting 5 seconds in the future. Additionally, the simplest model of all, the constant-velocity model, performed relatively well. Any significant improvement in predicting vehicle motion will likely require a combination of advanced modelling techniques that deal with location and time context.

2.3 Extracting Collisions from Trajectories

A CW or CA action typically does not directly depend on the exact positions or trajectories of vehicles. Instead it is formulated as a simple logic statement based on a) the probability of a collision or dangerous situation, and b) the time left until the collision/dangerous situation occurs. However, if only a single trajectory estimate for each vehicle is provided, collisions end up being considered as a deterministic phenomenon. The probability of a collision occurring can only be discussed if multiple trajectory predictions are made, or if the predictions have some inherent probability distribution. While at least one CW/CA project has considered probabilities of future collisions [29], most rely on a single estimate.

Even deterministic collision detection is computationally expensive if done in the most straightforward manner. The potential colliding vehicles must be checked against each other for enough timesteps to be sure that all imminent collisions are caught. This action must be repeated every time new trajectories are calculated, which may be ten times a second or more. A fast computational method for finding collisions is therefore desirable. Several works of research have

shown ways to more quickly calculate the collision time, usually by relying on the simple structure of their prediction method [19, 20].

2.3.1 Quickly Approximating Time-to-Collision in Intersections via Collision Zones

We have developed an approximation specifically for intersections, where vehicles travelling on different roads and lanes may collide.

- For every pair of roads, construct simulations that move a vehicle along each road. Whenever the two vehicles collide, store each vehicle's position (in terms of distance along the road).
- From the stored collision cases, find the smallest and largest distance for each vehicle. These four numbers form the collision zone in which a collision between vehicles at these two roads is possible.
- When determining whether two vehicles will collide, check the times at which they will be within the smallest and largest distance for their respective roads. If these time intervals overlap, then there is a potential collision. The intersection of the two vehicles' time intervals gives the time interval in which the collision may occur.

In essence, the collision-zone method simulates and characterizes all possible collision situations beforehand, so that in real time only a simple check is needed. Figure 2.4 gives a visualization of the first two steps. The method has two drawbacks: one is that each vehicle is assumed to follow a known road or path through the intersection. This is constricting, but cooperates with the road-context and motion-pattern methods in Chapters 2.2.3 and 2.2.4. The other is that the constructed collision zone is an approximation of the actual situations in which a collision will occur. However, as long as enough simulations are gathered, the constructed collision zone will only overestimate the true zone. Thus, this method may cause unnecessary CW/CA actions, but not miss actual collisions (provided the predicted vehicle trajectories are accurate). This method could not be tested on real data, because there are no available records of vehicle trajectories immediately prior to collisions. However, the method is utilized in the many-vehicle simulations in Chapter 5, to make early collision detection more efficient.




Figure 2.4: Visualization of creating collision zone from simulated collisions

Chapter 3. Develop Improved Collision Detection and Avoidance Algorithms between Cars and Non-Motorized Road Users

3.1 Introduction

Two critical differences between motorized and non-motorized users—pedestrians and bicyclists—must be considered in developing collision detection and avoidance algorithms. First, the physical motion of bicyclists and pedestrians are dissimilar to that of other road users, such as cars, trucks, and motorcycles. Furthermore, each type of user is endowed with different degrees of technology. It is not reasonable, at least at the moment, to assume that pedestrians and bicyclists are also equipped with DSRC, radar, speedometer, etc. Therefore, CW/CA algorithms on a vehicle should be developed without relying on direct information from non-motorized road users.

Three main types of sensors can be used for assisting the interaction between automobiles and non-motorized road users. Sensors like radar and LIDAR can detect the position of any object/user relative to the vehicle's surrounding and, therefore, may be the most important source of information in this situation. Communication systems, on the other hand, can be valuable in situations when a vehicle several vehicles away may detect the presence of a bicyclist on the road and send this information to other neighboring vehicles. Then, the information received from other vehicles can be utilized as a source of CW/CA algorithm.

The third on-vehicle sensor to be considered in this scenario is the camera or vision sensor. In the field of computer vision, object detection and tracking has a long history with promising results. Furthermore, a great deal of computer vision research has focused on pedestrian detection [30, 31, 32, 33], while some other techniques consider bicyclists [34, 35]. Fusion of radar and computer vision has also been studied before, as radar can not only provide accurate distance and angle but also robustness against bad weather, which is a limitation of the vision sensors [36, 37, 38]. Additional sensors like radar and LIDAR can be used as a supplementary source of information to enhance computer-vision-oriented algorithms.

To develop improved CW/CA algorithms, we focused on incorporating additional signal information to improve computer vision oriented methods for detecting pedestrians and bicyclists. The baseline computer vision algorithm used in the project is a deep-learning-based approach, which achieves higher performance and has been heavily researched in the last few years. Specifically, we used one of the well-known algorithms, Fast R-CNN [39], which showed successful performance in object detection. To train and evaluate our detection algorithm, we use a popular and recent dataset, the KITTI Benchmark Suite [40, 41], which includes car, pedestrian, and bicyclist detection benchmark. Although data from both radar and cameras are scarce in public datasets, KITTI provides scanned points from LIDAR around car, as well as images. LIDAR is similar to radar in that it provides depth information for surrounding objects with high accuracy, which meets our goal in testing a sensor fusion model.

3.2 Computer Vision Oriented Methods for Pedestrian and Bicyclist Detection

3.2.1 Challenges in Pedestrian and Bicyclist Detection

Computer vision approaches use images as an input to the detection algorithm. Even though a visual source includes plenty of information, pedestrian detection/tracking is not an easy problem because of the following reasons:

- Appearance of pedestrians exhibits high variability
- Outdoor urban scenario
 - Cluttered background
 - Weather condition and illumination vary the quality of sensing
- Highly dynamic scenes
 - Pedestrian and camera are in motion
 - Different view angles
 - Large range of distances
- Requires high performance in terms of speed and robustness

Although the overall process can be similar to pedestrian detection, the detection of a bicyclist can be even more challenging. First, a bicyclist comprises both the bicycle and a human rider, inducing confusion and high variety. Also, the bicyclist's speed is significantly faster than the pedestrian's. High variety in the image space can cause inconsistency of features extracted from a single class. Cho et al. [34] pointed out the integral issues in bicycle detection problems, which center around the bicycle/rider presenting dramatic appearance changes:

- Camera viewpoint (e.g., a bicycle's shape and size appear very different viewed from the front versus the side)
- Class variability (e.g., mountain bikes vs. racing cycles)
- Person riding on the bicycle

3.2.2 Computer Vision for Pedestrian and Bicyclist Detection

In the object detection/tracking area, computer vision focuses more on pedestrian detection than on bicyclist detection. Geronimo et al. [33] investigated pedestrian detection problems for an advanced driver system and proposed a general architecture to describe the pedestrian detection/tracking system for computer vision methods (Figure 3.1). This architecture is divided into six steps and covers the structure of most traditional computer-vision-oriented pedestrian detection systems:

- Preprocessing: Process low-level adjustments like exposure/dynamic range adjustments of the image and calibrate coordinate system with camera.
- Foreground segmentation: Extract the candidate portion of the image to be investigated—the region of interest (ROI)—from the image for the object classifier.
- Object classification: Classify whether ROI contains pedestrian or not.
- Verification/refinement: Filter false positive and refine segmentation.
- Tracking: Track/follow the detected pedestrian over time.
- Application: Apply pedestrian detection in various fields with high-level decisions.

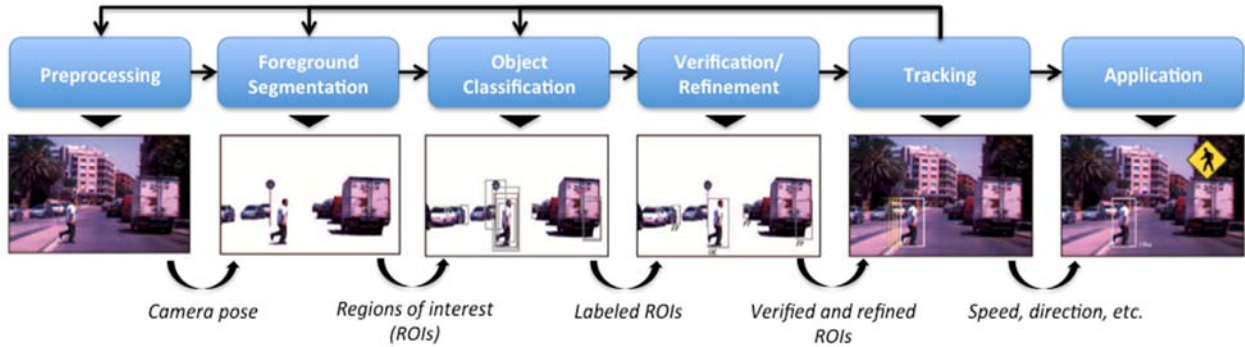


Figure 3.1: Architecture of pedestrian detection system [Geronimo09]

In our study, extracting the ROI and object classification are emphasized to prove the effectiveness of additional information from sensors like radar and LIDAR. Extraction of ROI in foreground segmentation is an integral part of the object detection process. Traditional approaches in computer vision commonly use the sliding window technique for providing ROI. This technique extracts the feature from a partial rectangular region of the image and feeds it to the classifier. Then, it slides the rectangular box to process the algorithm again on a different portion of the image. Since the sliding window process applies the classification algorithm to a large number of rectangular parts, computational efficiency is very low. Also, providing meaningless ROI to the classifier can increase the possibility of erroneously identifying a pedestrian where there is not one. Therefore, a novel algorithm for searching candidate ROI in the image plays an important role in improving speed of detection and decreasing false positive rates.

The most important step of the pedestrian detection is the object classification phase, which includes feature extraction and learning classifier. From the provided ROI, features are extracted to produce sufficient information to recognize the target in the image. Previous research efforts on finding effective features established the following features: Haar wavelet, SIFT (scale-invariant feature transform), HOG (histogram of gradients), edgelet, shapelet, etc. These features are extracted based on pixel difference, shape, edge in object, gradients, and other calculations among close pixels to capture characteristic shapes and appearances in the image. Then, the extracted image features become input for the classifier to learn the parameters, so that it can declare the class of the object in the given ROI. These approaches were later improved by the deformable part-based model approach, which tries to model the object into different parts and build classifiers based on those parts. This was a natural approach, as the human body can be divided into several parts (like head, body, leg, and arms) and these types of algorithms achieved the best performance before the advent of deep-learning-based models.

In modern object detection and image classification tasks, the Convolutional Neural Network (ConvNN) model is widely applied and often achieves state-of-the-art performance. It utilizes convolution to perform nonlinear calculation on neighboring pixels to capture underlying dependencies between the features. Training and testing require a large amount of data and long computation time, but developments in computation power, including multi-core and GPU computing, and plentiful sources of images like ImageNet [42] accelerated the development of deep-learning-based object detection. A unique advantage of the ConvNN approach is that it can extract effective general features from an image without applying a part-based model or other complicated calculation process. In practice, an existing pre-trained model can be applied to the new object detection task as a feature extraction module and extract general features capturing

characteristics of the target in the image. Results and evaluation of pedestrian detection with the ConvNN-based approach are presented in [43].

Details of the history of algorithmic development in pedestrian detection are well explained in the literature [30, 31, 32, 33]. Also, one report focuses on distinguishing bicyclists from pedestrians [35], which is actually an important distinction, as the top part of a bicyclist can be detected erroneously as a pedestrian.

3.3 Deep-Learning-Based Algorithm with Sensor Fusion

3.3.1 Fast R-CNN for Pedestrian and Bicyclist Detection

To achieve reasonable performance in non-motorized-user detection using cutting-edge technology in machine learning, and to investigate the possibility of a real-world application of sensor fusion, we selected the Fast R-CNN algorithm as a baseline model. Region-based convolutional neural networks (R-CNN) were first proposed in [44], combining region proposals and ConvNN. ConvNN has several well-known models that achieve state-of-the-art image classification performance, capturing meaningful features from an image while identifying certain classes (such as person, car, and bicycle). By applying an algorithm that proposes interesting candidate regions in the image regardless of the class of the target, the number of ROIs decreases while meaningful regions remain. Since the original R-CNN were computationally expensive, improved versions were proposed later and termed *Fast R-CNN* and *Faster R-CNN* [45].

In this project, Fast R-CNN is used as a baseline model as it is composed of separate region proposal stages and ConvNN stages, allowing space for additional sensor information. Faster R-CNN computes ROI with deep nets by sharing convolutional layers. The structure of Fast R-CNN is provided in Figure 3.2, extracted from the original paper. Both the original image and ROIs are projected to ConvNN in a fixed scale and output class posterior probabilities and predicted bounding-box offsets per class.




Figure 3.2: Architecture of Fast R-CNN (Source: [44])

3.3.2 Fusion of Radar with Fast R-CNN

Before developing a fusion algorithm for object detection, the characteristics of radar need to be considered. The following points are well-known advantages and disadvantages of radar.

- Advantages of radar
 - Robust against bad weather (e.g., rain and fog)
 - Provide accurate distance and angle
- Disadvantages of radar
 - Cannot provide enough data points to detect obstacle boundaries
 - Reflections from other objects (humans have low reflectance)

Without considering weather, accurate distance information can play an integral role in fusion algorithm. Therefore, it is natural to devise fusion approaches in two ways.

- Give an accurate distance-based estimate to a vision-based system
- Decrease the number of candidate ROIs to decrease false positives and speed.

Wan et al. [38] proposed a three-stage fusion strategy for vision and radar detection system. The first step is to align and calibrate coordinate systems of radar and vision. The system then searches for potential targets from both sensors. Finally, it detects objects on the road from those proposed regions. Figure 3.3 illustrates this three-step strategy.



Figure 3.3: Fusion with mmWave radar (Source: [38])

In Fast R-CNN code, selective search [46] is selected as the region proposal stage, which produces a hierarchical group of pixels based on the predefined similarity. Even though it extracts a number of effective ROIs quickly, it requires about 0.5s per image and includes a number of false positive regions. Therefore, we also tried to improve the overall system by providing additional information from radar in extracting ROIs. As radar can provide accurate distance information with higher accuracy in closer objects, ROIs extracted from selective search are rejected based on the radar data. A simple diagram of the radar fusion with Fast R-CNN is provided in Figure 3.4.

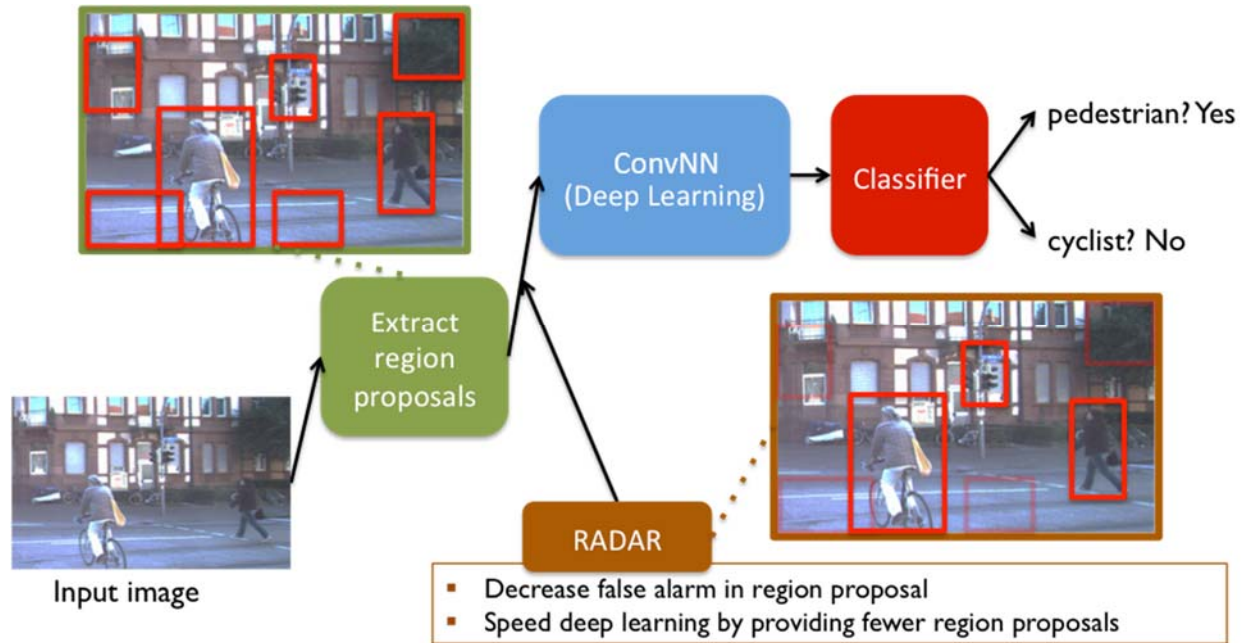


Figure 3.4: Architecture of Fast R-CNN with fusion of radar

3.4 Simulation on KITTI Object Detection Benchmark

3.4.1 KITTI Object Detection Benchmark

The most popular pedestrian detection benchmark is the Caltech Pedestrian Detection Benchmark [47]. However, it doesn't include labels for cyclists. Therefore, we decided to use a more recent dataset, the object detection dataset of KITTI Vision Benchmark [40], which includes 7481 training images and 7518 test images. Seven class labels are provided: car, van, truck, pedestrian, person_sitting, cyclist, and tram. For evaluation, we considered only car, pedestrian, and cyclist, which have a relatively large number of images. One problem in development and evaluation of the model comes from the absence of ground truth label in the testing set. Therefore, we split the KITTI object detection training set into training and validation sets. In the training set, there are annotations of 28,782 cars, 4487 pedestrians, and 1627 cyclists among 7481 images. As the KITTI object detection images are extracted from several videos, the splitting procedure ensured that no image from the same video was present in both the training and the validation sets. The resulting dataset consists of 3471 images in the training set and 3470 images in the validation set with a balanced number of annotations between the two sets.

Finding a public dataset providing both image and radar information is another challenge, but the KITTI benchmark provides depth, including point clouds collected from Velodyne LIDAR. LIDAR and radar operate differently but provide the same information required by our project: accurate distance information. The KITTI benchmark provides the necessary information to project LIDAR points onto the 2D coordinates of images, so the calibration process can be completed prior to the actual fusion step.

3.4.2 Analysis on Region Proposals from Selective Search

To investigate the weak points of Fast R-CNN with selective search on RGB images, we analyzed the ROIs proposed by selected search in several ways. About 1000 ROIs (on average) are proposed per image using selective search on the KITTI object detection dataset. When training the network, ROI is selected as a positive sample if the intersection over union (IoU)—the overlap ratio with a ground truth bounding-box—is at least 0.5. Therefore, the number of images with no IoU exceeding 0.5 is counted to see how many images are provided with no additional positive ROIs. This value is critical in testing because object detection cannot be processed properly without good ROIs; even a perfect object classifier cannot detect an object from the ROI box containing almost nothing. As provided in Table 3.1, 27.7% of images in training set and 17.7% in validation set are proposing ROIs with only IoU less than 0.5, $\#(\text{IoU}_+=0)$. From this statistic, we decided to harness LIDAR in improving the positive ROI proposal ability. By applying selective search on dense LIDAR images, 8.9% and 9.7% of images with no positive ROI sample ($\text{IoU} \geq 0.5$) are recovered in the training and validation sets, respectively.

Table 3.1: Analysis on quality of region proposals provided by selective search (with and without LIDAR fusion) on KITTI object detection Train/Validation set

	Input	$\#(\text{IoU}_+=0)$	$\#(\text{IoU}=0)$	Avg.#ROIs
Train	RGB	962 (27.7%)	672.3	1090.6
	RGB+LIDAR	876 (25.2%)	740.6	1278.1
Validation	RGB	662 (17.7%)	681.6	1138.9
	RGB+LIDAR	598 (16.0%)	733.7	1323.6

In addition, the number of hard ROIs with no overlapping ground truth per image, $\#(\text{IoU}=0)$, and the average number of ROIs—Avg.#ROIs—are calculated to investigate the characteristics and possibility of improvement in Fast R-CNN with LIDAR. About 60% of ROIs extracted using selective search do not overlap with the ground truth objects (pedestrian, cyclist, and car). This finding provides another direction for research, rejecting ROIs with 0 IoU. In the next two sections, different approaches for LIDAR fusion in region proposal stage are suggested with experimental results.

3.4.3 Fast R-CNN and LIDAR Fusion: Simple ROI Rejection

Our first fusion approach focuses on decreasing false positives by rejecting ROIs with no containing object. The LIDAR ‘image’ obtained by projecting LIDAR points into the image plane is initially sparse (only a few points have known values). To extract meaningful depth information from LIDAR points, the procedures described below and shown in Figure 3.5 are applied to the original LIDAR points cloud.

- Project LIDAR points onto image coordinates with depth information
- Transform depth values $d(x, y)$ to have higher weights on closer points

$$\hat{d}(x, y) = 1/\sqrt{d(x, y)}$$

- Rescale depth values $\hat{d}(x, y)$ to 0~255 range (image value range)

- Use inpaint method to impute sparse depth image
(Upper part of image is not considered as it is not covered by LIDAR due to the range of angle)
- Detect edges from the imputed depth image using Canny edge detector

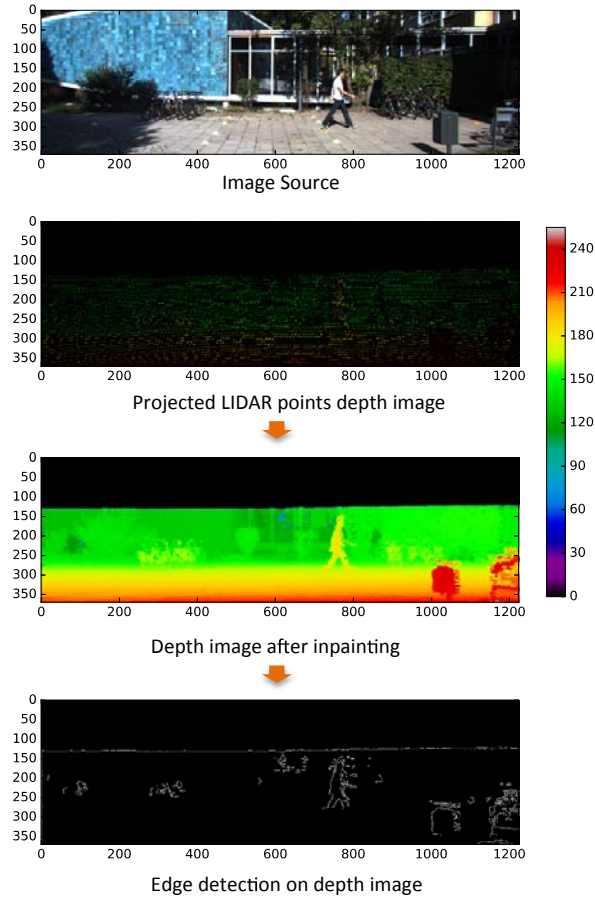


Figure 3.5: Extraction of depth image and edge detection from sparse LIDAR points

Based on the extracted edge detection and depth image, useless ROIs provided by selective search or another general method can be deleted if the following two conditions hold:

- If average depth in ROI box on depth image $>$ threshold value (140)
- If no edge is detected in ROI box on edge detection image

The Fast R-CNN model is trained on the training set and evaluated on validation set with average precision criterion. Training and testing of the Fast R-CNN model was conducted on a machine with K40 GPU and other factors, hosted by the Texas Advanced Computing Center at The University of Texas at Austin. The maximum number of iterations in training is 40,000 and we trained two types of deep ConvNN with different depths of the network: CaffeNet (Small) and VGG_CNN_M_1024 (Medium). The evaluation images have been divided by KITTI into three types: easy, moderate, and hard. Difficulty is defined based on size, occlusion, and other factors.

Both models showed reasonable performance in car detection. Results from computer vision algorithm and sensor fusion algorithm are provided in Tables 3.2 and 3.3 respectively. The evaluation metric chosen was average precision (AP) in percentage, which summarizes the shape of the precision/recall curve, and is defined as the mean precision for a set of eleven equally spaced recall levels. Also, precision/recall curves for each object class and difficulty are provided in Figures 3.6 and 3.7.

Table 3.2: Object detection result on KITTI validation set with only computer vision

		Car			Pedestrian			Cyclist		
Net	Iteration	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
CaffeNet	10000	88.11	87.28	78.04	76.84	66.57	60.24	46.67	46.82	45.79
	20000	87.77	87.86	78.61	81.26	70.32	64.47	47.47	48.82	46.82
	30000	90.76	88.55	79.53	81.18	70.76	64.86	51.42	52.92	51.08
	40000	91.12	88.36	79.31	83.76	74.17	68.38	49.31	51.64	49.53
VGG_CN N_M_1024	10000	92.74	88.75	79.21	78.01	67.71	60.86	47.60	50.03	47.82
	20000	91.42	89.02	79.92	82.41	72.69	66.78	54.51	55.37	53.51
	30000	93.94	89.56	80.41	81.44	74.03	68.74	54.24	55.81	53.37
	40000	91.98	89.42	80.27	83.27	75.13	70.21	55.55	55.92	53.64

Table 3.3: Object detection result on KITTI validation set with LIDAR information

		Car			Pedestrian			Cyclist		
Net	Iteration	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
CaffeNet	10000	87.37	84.69	75.14	74.89	65.96	59.33	42.23	43.83	40.17
	20000	86.95	86.48	77.33	78.43	69.65	64.02	45.13	47.03	45.54
	30000	87.99	87.93	78.93	82.50	72.11	66.28	46.11	48.49	46.58
	40000	87.65	88.06	79.08	82.42	72.54	67.51	47.39	49.71	47.93
VGG_CN N_M_1024	10000	88.68	87.01	77.09	78.69	68.67	61.61	51.43	51.18	49.70
	20000	91.41	88.15	79.19	82.54	71.62	64.92	48.29	49.83	48.50
	30000	91.38	89.08	80.06	84.02	74.11	68.40	53.84	55.70	53.79
	40000	92.03	89.48	80.39	83.38	73.73	68.32	53.17	55.88	54.01

Compared to car detection, the pedestrian and cyclist detection models performed poorly, especially for the cyclist case. There are several possible reasons for the poor performance, but one main reason is an insufficient number of images. The car class had the greatest number of examples to train on. Although improvement in detection process is expected from sensor fusion, the results show similar detection performance with the proposed fusion method. One reason for this is the good performance of the selective search algorithm. In fact, an additional rejection process could only delete approximately 10% of ROIs proposed by selective search algorithm. More advanced fusion algorithms in a different stage, such as the prediction step, can be investigated in the future to improve both accuracy and speed to apply the deep learning approach to an onboard system. Although LIDAR (or radar) didn't show effective influence in object detection process for the CW/CA system with this approach, it can provide further advantages. For instance, position or distance estimation of the pedestrian and bicyclist can be enhanced by LIDAR and radar as they can provide more accurate relative distance. However, we could not evaluate this idea since there is no dataset providing the true position of the pedestrians or bicyclists recorded in the video or image.

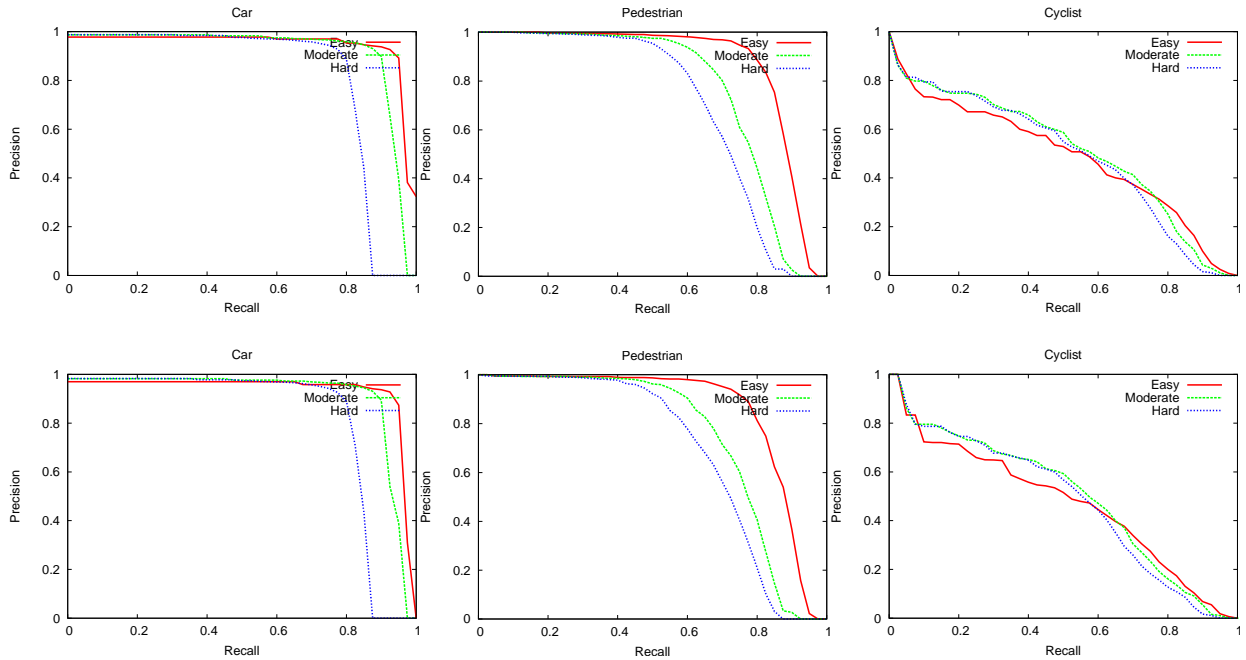


Figure 3.6: Precision-Recall curve of KITTI object detection on validation set trained with CaffeNet (Small size ConvNN). Iteration number 40000. Upper: computer vision only; Lower: computer vision+LIDAR (Simple ROI Rejection)

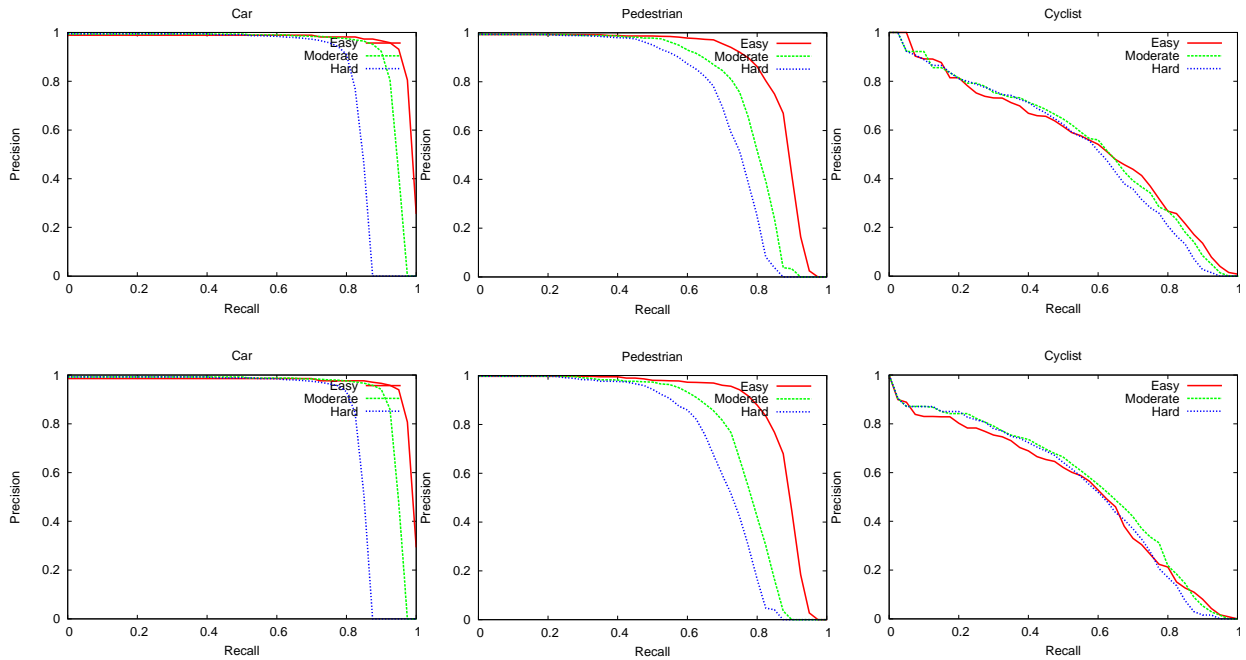


Figure 3.7: Precision-Recall curve of KITTI object detection on validation set trained with VGG_CNN_M_1024 (Medium size ConvNN). Iteration number 40000. Upper: computer vision only; Lower: computer vision+LIDAR (Simple ROI Rejection)

3.4.4 Fast R-CNN and LIDAR Fusion: Supplementary Region Proposal

Our second proposed model using Fast R-CNN and LIDAR image is provided in Figure 3.8. In the region proposal stage, ROIs are extracted from RGB images with selective search's fast mode. In this model, a simple inpainting module provided by OpenCV is used to get a dense LIDAR image. Bounding-box locations of ROIs are extracted from this inpainted LIDAR image using selective search's intensity mode to deal with gray-scale images. Next, ROIs obtained from both RGB and LIDAR images are provided to ConvNN, so that feature vectors can be extracted from the RGB image box at the location of those ROIs. Finally, a fully connected neural network layer provides class both probability scores from a softmax classifier and more precise object locations from a bounding-box regressor.

To evaluate the effectiveness of newly extracted ROIs from LIDAR in training, results obtained from training only with RGB and testing with additional ROIs from LIDAR are provided in Table 3.4. Figure 3.9 and Figure 3.10 illustrate the precision-recall curves with CaffeNet on object detection, comparing both Fast R-CNN and the proposed model with LIDAR ROIs in Train&Test. Our model with CaffeNet outperformed Fast R-CNN with only vision data in both cyclist and pedestrian detection. In particular, it improved cyclist detection by 3~4 in terms of AP score in percentage and also gave a small improvement in the pedestrian category. Furthermore, the positive influence of LIDAR fusion in training a convolutional network can be identified from the AP score difference between the two LIDAR support levels: Test vs. Train&Test. However, a model based on wider ConvNN architecture, VGG_CNN_M_1024, showed a slight decrease in the performance for the most of tasks in pedestrian and cyclist detection. Since overall performance with a medium-size convolutional network shows better results, this may indicate the biased good performance of ConvNN on vision data. Although more positive ROIs from LIDAR are provided to the ConvNN, it is possible that those ROIs have poor characteristics when viewed in RGB space. Additional ROIs found in depth image space may become a potential confusion factor for well-performing ConvNN architectures.

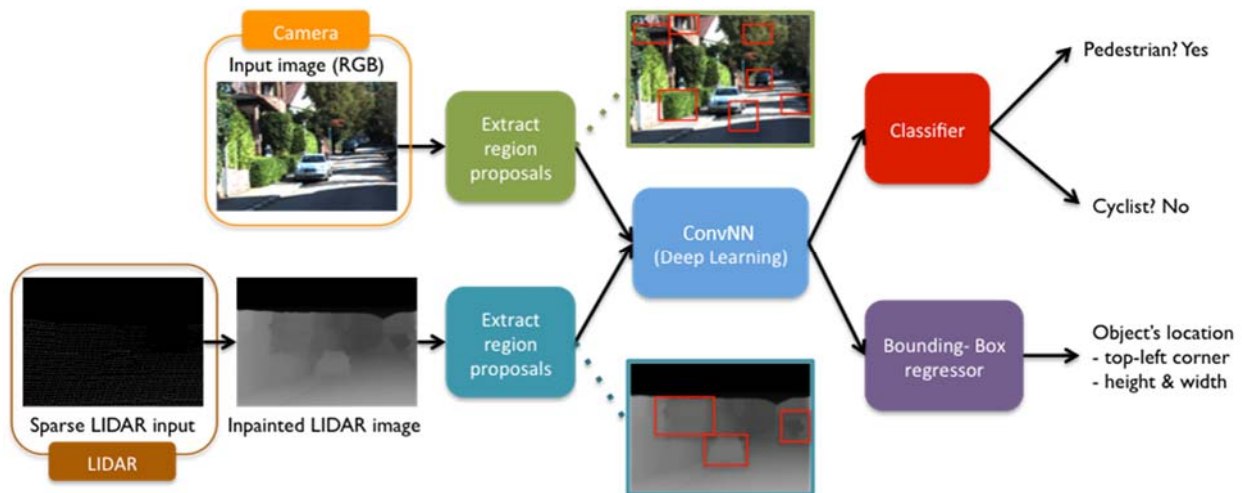


Figure 3.8: Architecture of Fast R-CNN with fusion of LIDAR.

Input data are separated collected from camera and LIDAR to extract region proposals. Convolutional network provides fixed length feature vector from those region proposals to object classifier and bounding-box regressor to declare class and location of object

Table 3.4: Cyclist and pedestrian detection AP (%) of Fast R-CNN with different LIDAR support on KITTI validation set.

Method	Pedestrian			Cyclist		
	Easy	Moderate	Hard	Easy	Moderate	Hard
CaffeNet	82.01	73.78	67.98	38.24	46.04	44.19
CaffeNet+LIDAR (Test)	81.60	72.75	66.99	37.29	44.70	43.06
CaffeNet+LIDAR (Train&Test)	82.73	74.12	67.92	42.69	49.08	47.52
VGG_CNN_M_1024	84.69	76.50	71.49	51.54	54.03	52.58
VGG_CNN_M_1024+LIDAR (Test)	83.90	75.67	70.11	49.91	52.79	51.47
VGG_CNN_M_1024+LIDAR (Train&Test)	84.43	74.76	69.32	52.34	52.49	50.99

Note: (Test) represents additional ROIs from LIDAR only in Testing and (Train&Test) indicates the model trained with those additional ROIs

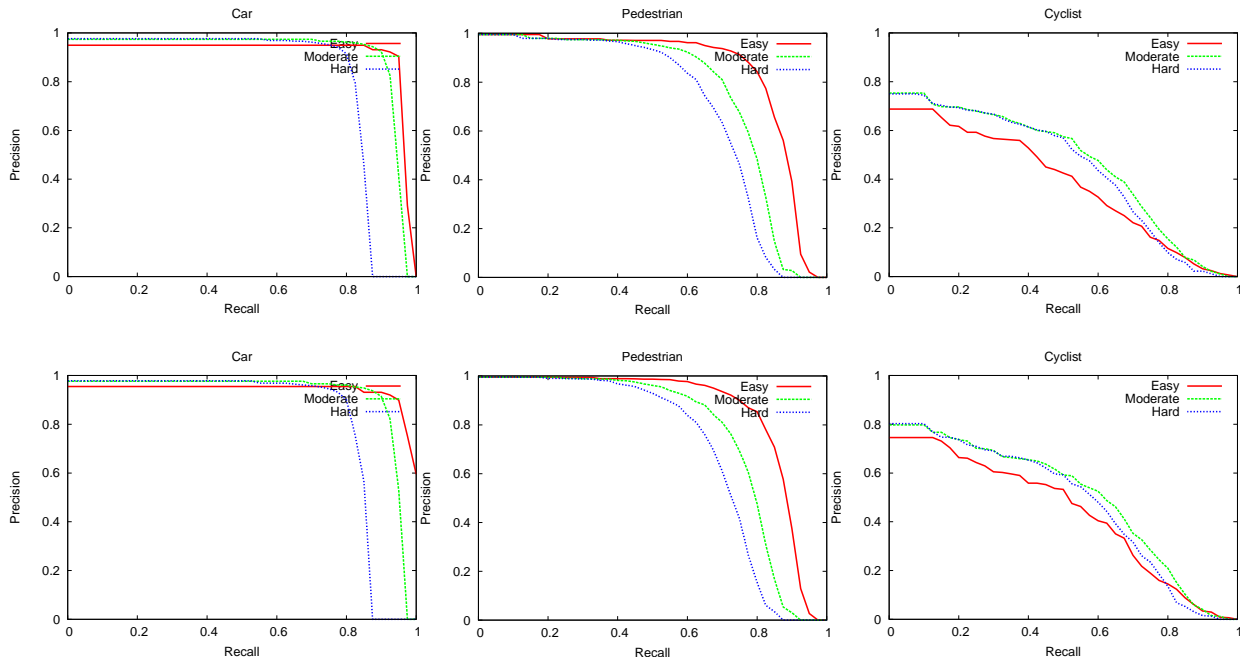


Figure 3.9: Precision-Recall curve of KITTI object detection on validation set trained with CaffeNet (Small size ConvNN). Iteration number 80000. Upper: computer vision only; Lower: computer vision+LIDAR (Supplementary Region Proposal)

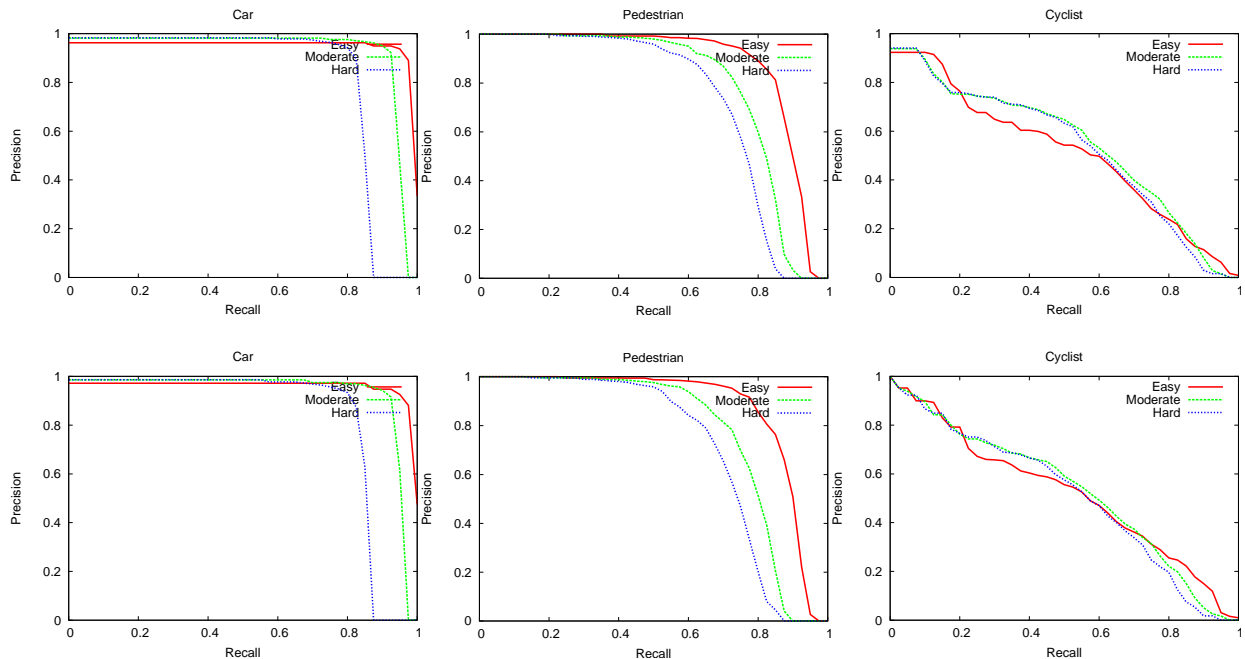


Figure 3.10: Precision-Recall curve of KITTI object detection on validation set trained with CaffeNet (Small size ConvNN). Iteration number 80000. Upper: computer vision only; Lower: computer vision+LIDAR (Supplementary Region Proposal)

When compared to the car detection results in Table 3.5, the pedestrian and cyclist detection models performed poorly, especially for the cyclist case. One main reason is probably an insufficient number of images. The car class has the most examples to train on (approximately 15,000), while the cyclist class contains only 0.8k examples with ground truth data. Nevertheless, the effect of LIDAR fusion in the region proposal stage for deep learning can be discovered using the reliable car detection results. The use of LIDAR fusion increased result accuracy in easy tasks by 2~4 AP, while results for moderate and hard tasks showed almost similar values among three different methods. This outcome may imply the supportive power of the LIDAR sensor if vision algorithms miss objects that can be easily detected by human drivers. Figure 3.11 provides examples of detection results, higher prediction probability or more accurate bounding-box with our proposed fusion method.

Table 3.5: Car detection AP (%) of Fast R-CNN with varying levels of LIDAR support on KITTI validation set

Method	Car		
	Easy	Moderate	Hard
CaffeNet	86.15	88.01	79.14
CaffeNet+LIDAR (Test)	91.04	87.76	78.86
CaffeNet+LIDAR (Train&Test)	92.02	87.98	79.08
VGG_CNN_M_1024	90.46	88.92	79.95
VGG_CNN_M_1024+LIDAR (Test)	93.57	88.75	79.79
VGG_CNN_M_1024+LIDAR (Train&Test)	92.47	88.99	80.01

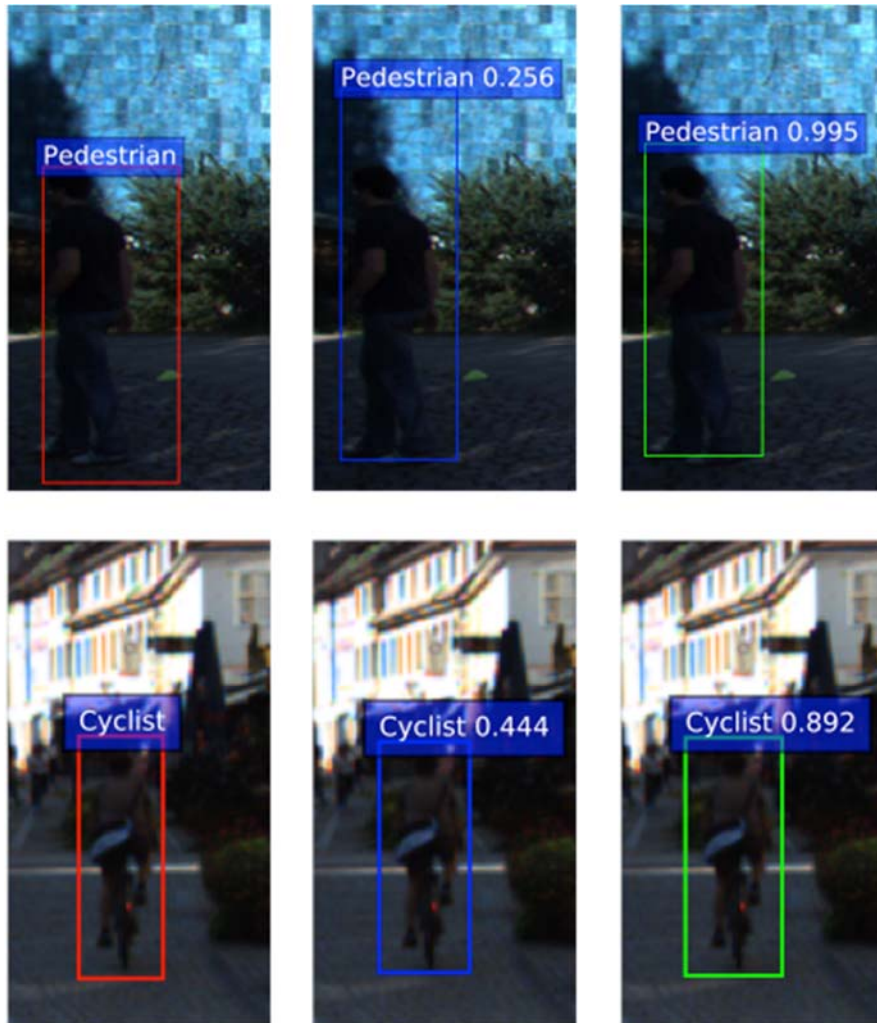


Figure 3.11: Pedestrian (upper) and cyclist (lower) detection result examples: ground truth (red), CaffeNet (blue), and CaffeNet+LIDAR (lime). Numbers indicate predicted class probabilities.

Chapter 4. Modeling Uncertainty through Simulations

4.1 Introduction

Each device—radar/LIDAR, communication, video cameras—has advantages and disadvantages. Furthermore, none will give perfect or complete information. A major open question in CW/CA is how much accuracy and information will be necessary to create an effective CW/CA system. This question is very difficult to address with real data: we cannot feasibly locate collisions among human drivers (before the fact), and prototyping a CW/CA system to test its accuracy is highly expensive/dangerous.

Computer simulation provides an alternative way of measuring the quality of CW/CA systems. Simulation provides specification (each vehicle’s motion can be chosen to address a certain question) and repeatability (the same scenario can be repeated with small changes, for instance with different CW/CA technologies). The disadvantage is that care must be taken to ensure that the simulated driving scenarios are realistic, as well as varied enough to reflect all dangerous situations for a given driving task.

We developed simulations for three main scenarios: a) car-following on highways, b) unregulated intersections, and c) overtaking maneuvers on rural roads. For each of these driving tasks, the goal was to determine which sensor properties or issues can significantly affect the performance of a CW/CA system. Four sensor characteristics were studied:

- Range – the distance and direction at which a communication is capable of detecting a vehicle.
- Packet Loss – this includes any form of temporary communication failure. This is often caused in intersections by competition among vehicles for the wavelength.
- Positioning Accuracy – V2V communication usually relies on GPS to gather the exact positions of each vehicle, while sensors such as radar or LIDAR capture the relative position of nearby vehicles.
- Other information – Speed, heading, and acceleration can augment positioning information to provide a much more accurate prediction of the vehicle’s trajectory. Other potential information is the size of the vehicle.

In addition, sensors with known or measured properties can be compared, to see which provides superior accuracy in any situation.

The popular open-source traffic simulator SUMO [48] was used for initial simulations. However, this simulator is fundamentally intended to simulate normal traffic situations and is not flexible in allowing vehicles to break driving conventions or collide with each other. Therefore, we ultimately created our own simulator, which lacks some of SUMO’s built-in capabilities such as large-scale road creation and car-following models but is better at creating and analyzing collision-prone simulations. This simulator is publicly available at github.com/utexas-ghosh-group/carstop/simulator, and will serve as a valuable resource for future CW/CA analysis.

4.2 Highway Simulations

There has already been a significant amount of research on the detection and prevention of rear-end collisions, particularly in the context of automatic cruise control. A wide array of on-vehicle sensors, including LIDAR, radar, and forward-facing cameras, as well as V2V communication, are easily capable of detecting a vehicle directly in front of them. Thus, the question of whether or not a lead vehicle is detected is less important than how quickly and accurately its distance and motion can be determined. The essence of a rear-end collision system lies in quickly noticing low-speed or decelerating lead vehicles and determining whether braking is necessary.

In order to quantify the link between sensor quality and CW/CA, simulations of one-lane highway driving were constructed. Two separate sets of simulations were conducted: vehicles in the first engage in simple, constant-acceleration motion that can be precisely predicted, while vehicles in the second have periodic, random changes in speed. These may be mistaken for braking or cause misprediction of the vehicle's future trajectory. Tables 4.1 and 4.2 cover the details of each simulation.

Table 4.1: Details of constant motion simulation

Environment		One straight, single-lane, 1-kilometer road with two vehicles	
		Rear Vehicle	Lead Vehicle
Vehicle Shape		5 m length	
Vehicle Initialization	Position	5 m	55 m
	Speed	Uniformly distributed between 55 and 85 mph	
	Acceleration	Uniformly distributed between -1 and 1 m/s ²	
Vehicle Motion		Maintains its initial acceleration	<ul style="list-style-type: none"> - First, maintains initial acceleration - A braking acceleration is also assigned, uniformly distributed between -2 and -9 m/s² - A brake time is randomly assigned with a mean of 35 seconds - After the brake time (follows the braking acceleration)
Simulation Ends		Vehicles collide, or lead vehicle reaches 1 kilometer	

Table 4.2: Details of randomized motion simulation

Environment		One straight, single-lane, 1-kilometer road with two vehicles	
		Rear Vehicle	Lead Vehicle
Vehicle Shape		5 m length	
Vehicle Initialization	Position	5 m	55 m
	Speed	Uniformly distributed between 55 and 85 mph	
Vehicle Motion		Starting with the initial speed, accelerates to a desired speed and stays there for a certain time	
	Desired Speed	Normally distributed around the current speed, with a deviation of 10 mph	
	Time Period	Exponentially distributed with a mean of 3 seconds	
	Acceleration	If desired speed is higher than current speed, uniformly distributed between 0 and 1 m/s ² If desired speed is lower than current speed, distributed between 0 and -1 m/s ²	If desired speed is higher than current speed, distributed between 0 and 1 m/s ² If desired speed is lower than current speed, distributed between 0 and -5 m/s ²
Simulation Ends		Vehicles collide, or lead vehicle reaches 1 kilometer	

4.2.1 Simulated Collision Avoidance Algorithm

Every vehicle’s exact position, speed, and acceleration are available at every timestep in the simulation. To simulate a sensory device, these values can be altered or withheld from the CW/CA system.

- Range: simulate by providing information from a different vehicle only if it is nearby and/or aligned in the right direction
- Inaccuracy: simulate by applying randomized errors to the information. We consider only time-varying and zero-mean errors (a.k.a. white noise), though other types of error may be possible.
- Packet Loss or High Latency: simulate by withholding information randomly, or at a fixed rate

Trajectory prediction was performed using the assumption of constant acceleration, with a Kalman filter for noise mitigation. The algorithm searches for future collisions between the two vehicles and warns the driver (or performs another avoidance action, such as braking) once it detects a collision. There are two requirements imposed on the CW/CA system: it must predict the time of the future collision within 1-second accuracy, and it must predict a collision at least 2.5 seconds in advance. This number was chosen based on the general consensus that up to 3 seconds of braking may be necessary to avoid a collision. Collisions that were predicted to occur more than one second earlier than they actually did are designated as Early Warnings. Collisions that were predicted to occur over 1 second past their actual occurrence, or not detected until less than 2.5 seconds before the collision, were designated as Late Warnings. Note that there were many simulations in which no collision occurred; in these simulations any warning is considered incorrect.

Four hypothetical sensing systems were studied, the first being an ideal noiseless case. Two levels of noise were applied, one double the other. The values for the highest level of noise were chosen from worst-case assumptions about GPS accuracy [49] and speedometer calibration [50]. It is worth noting that the field tests conducted with GPS-equipped vehicles demonstrated a much higher accuracy than this worst-case. In the final system, the sensor had no noise but instead gathered information at 1/5 the rate of the other sensors—a change from 0.1 seconds of latency to 0.5 seconds. This example studies the relative importance of frequency versus accuracy in information retrieval.

4.2.2 Results of Highway Simulations

Results for the constant-motion simulations are available in Table 4.3, and results for the random-motion variant are in Table 4.4. The main takeaway from these simulations is that, for ambitious collision-detection tasks, getting high-accuracy sensor information is much more important than gathering information frequently. Also note that CW/CA systems typically fail by detecting a collision too early. This is perhaps because a single timestep that erroneously points to a collision will immediately cause a warning, while a single timestep that erroneously ignores a collision will only delay the detection of this collision.

The unpredictability of future vehicle actions in the randomized motion simulations strongly limit CW/CA, even when perfect sensing is available. We have not verified with real data/experiments whether our assumptions on the randomness of vehicle motion, but if they are representative of real vehicles' unpredictable actions, then it will be necessary for CW/CA systems to overcome this poor performance. Predicting less far into the future, and accepting imprecise predictions of collision time, will help to improve the overall accuracy of the collision prediction.

Table 4.3: Collision detection for constant-motion simulations

Sensor Setup	254 Collision Sims				146 Safe Sims
	Correct	No Warning	Late Warning	Early Warning	Correct
Ideal (no noise)	244	10	0	0	146
Some noise (2m, 2.5 mph)	66	42	6	140	138
Worst Case (4m, 5mph)	64	18	2	170	102
No noise, 5x latency	170	32	0	32	146

Table 4.4: Collision detection for random-motion simulations

Sensor Setup	248 Collision Sims				252 Safe Sims
	Correct	No Warning	Late Warning	Early Warning	Correct
Ideal (no noise)	116	7	3	122	159
Some noise (2m, 2.5 mph)	81	10	6	155	72
Worst Case (4m, 5mph)	51	0	10	187	52
No noise, 5x latency	112	10	19	107	149

4.2.3 Determining the Minimum Required Time for Collision Detection

As it has been noted that detecting collisions becomes increasingly difficult with higher time constraints, it is worth examining how early, and how precisely, a collision must be detected for a CW/CA algorithm to be successful. While it is possible to look at real drivers' responses, these responses may be overly cautious due to risk adversity or human PR time limits [51]. An alternative is to once again simulate collisions, and focus on finding the time at which an optimal avoidance maneuver, such as braking or swerving, is successful. Hillenbrand et al. [52] have already developed a simulator framework to determine the necessary time-to-avoid metric.

4.3 Intersection Simulations

Avoiding collisions among vehicles entering an intersection is a difficult task for sensors to achieve, and a chief motivation for V2V or V2I communication. A simple four-way, single-lane intersection was simulated to analyze intersection collisions. There are three different ways vehicle routes can conflict in this intersection: straight vehicles can collide from either side, left-turning vehicles can collide with a straight oncoming vehicle, and right-turning vehicles can collide while merging with a straight vehicle (shown in Figure 4.1). Considering both vehicles' points of view, the task of intersection CW/CA really requires detection of six different types of collisions.




Figure 4.1: Three types of collision in a four-way intersection

We assume no traffic signals or a sign, as a collision is most likely to occur in uncontrolled scenarios, such as unprotected left turns, or when one or more vehicles ignore the traffic signals. Each vehicle follows the randomized motion model specified in Chapter 4.2. The vehicles also

travel along fixed paths in the intersection, so the speed of each vehicle is the determining factor in whether or not a collision occurs.

4.3.1 Analysis of Desirable Sensor Properties

As in the case of the rear-end collision simulations, the goal of the intersection simulations is to determine how characteristics of vehicular sensing (communication included) correspond to capabilities in actually detecting collisions. Sensor accuracy is modelled in the same way as in Chapter 4.2. The visibility of each sensor is encapsulated in two parameters: the range, which requires that the absolute Euclidean distance between two vehicles be small enough, and the field of vision, which requires that the detected/sending vehicle be within a certain angle from the front of the detecting/receiving vehicle. Wireless communications are attractive because they typically have a high field of vision, while radar and cameras must usually necessitate a trade-off between vision and range but provide high resolution and accuracy (with the proper processing).

Table 4.5 shows results from 500 simulations of a vehicle crossing the intersection and conflicting with a right-merging vehicle (scenario C in Figure 4.1). Unlike the results from Chapter 4.2, which specified whether collisions were predicted too early or late, all detected collisions are considered correct. Rear-end collisions must be continuously searched for, and whether a CW/CA system should intervene depends on how imminent the predicted collision is. For intersection maneuvering, collisions can only occur within a short timeframe and any collision is worth immediate action (not entering the intersection yet). The timeliness requirement—that a collision must be detected at least 2.5 seconds before it occurs—is still maintained.

Table 4.5: CW/CA performance with varying sensor properties, intersection

Sensor Properties	Collision Scenarios with Warning	Safe Scenarios without Warning
Ideal: 0 noise, 100m range, field of vision 180° either side	84%	96%
0.5 noise	95%	86%
1 noise	100%	67%
50 meter range	19%	99%
25 meter range	0%	100%
Field of vision 90° (half circle facing forward)	84%	96%
Field of vision 45°	80%	94%
Field of vision 22.5°	0%	100%

The most interesting feature of these results is that noise actually improves the rate at which collisions are detected, at the cost of predicting many collisions that do not actually occur. Time-varying noise causes the system to become more paranoid, for reasons explained in the highway simulation discussion. It is possible that the physics-based prediction method, which cannot predict random changes in speed, frequently misses accelerations or decelerations that cause collisions. A method that predicts more cautiously could improve performance overall. The other clear takeaway is that the range of a vehicle sensor must be well over 50 m, and that it must detect vehicles within at least 45 degrees of its heading, in order to detect conflicting vehicles at an intersection.

4.3.2 Analysis of Radar versus DSRC communication

Two devices, representative of in-vehicle sensing and vehicular communication respectively, were simulated to provide a more concrete comparison of the two methods. One is the Delphi ESR module [53], which contains both a short-range, wide-angle radar and a long-range, narrow radar. The other is V2V communication device under the DSRC standard [54]. Both devices were assumed to obtain information about their own state through GPS positioning, speedometers, and accelerometers. The complete properties of each device, as researched and simulated, are available in the appendix. CW/CA was tested using each sensor separately, as well as using a fusion of the values given by each.

The results are shown in Table 4.6 (Figure 4.2 shows which result corresponds to which scenario). The worst performance for each sensor, in terms of both missed collisions and false detections, is highlighted. The Delphi radar system is clearly insufficient for detecting vehicles approaching straight from either side. While it slightly outperforms DSRC in left-turn scenarios, it seems that DSRC communication provides far more reliable accuracy in all cases. A fusion of the two information-gathering methods does not unilaterally improve performance, but it significantly helps in scenario D, the task of collision detection while merging.

Table 4.6: Intersection collision detection accuracy for specific sensors

	Radar (Delphi ESR)		DSRC (+GPS + Speed + Accelerometer)		Radar & DSRC Fusion	
	Crash	Safe	Crash	Safe	Crash	Safe
A	46%	96%	91%	92%	92%	91%
B	34%	93%	96%	81%	96%	82%
C	79%	95%	85%	91%	88%	86%
D	35%	100%	77%	95%	90%	86%
E	92%	89%	89%	89%	91%	89%
F	91%	88%	88%	89%	89%	89%




Figure 4.2: Labels of Table 4.6 mapped to their corresponding collision scenario

Chapter 5. Harnessing the Power of Collective Intelligence

5.1 Introduction and Motivation

In the not-so-distant future, cars will become autonomous entities with the capability to make decisions at the individual-vehicle level to meet a particular passenger's needs. However, autonomous cars will need to work collectively, collaborating so that all the vehicles in a network safely reach their destinations in the shortest possible time. One key research challenge is learning how autonomous vehicles navigate intersections. Intersections, especially in large urban centers, are areas of high risk for all drivers. Traffic lights at these intersections are programmed to control and route traffic safely. However, it is possible that when the entire fleet of operating vehicles becomes equipped with communication equipment, traffic lights will become obsolete. In this chapter, we consider how automobiles could be designed to communicate with one another and sense the appropriate time to execute an intersectional shift, judge the speed at which to cross, and comprehend the proximal distance of other vehicles without the need of traffic control devices or human drivers to make safe decisions.

We adapted the Collective Intelligence (COIN) [55] framework to design an algorithm allowing a fully automated car to navigate safely and timely to their destination through an intersection. Managing a multi-agent environment is often very labor-intensive, as it requires tuning the interaction between the agents to make sure that they are cooperative. These techniques normally result in non-robust systems with limited applicability. The COIN framework is very useful when

- There is little to no centralized communication or control, and
- There is a provided world utility function that rates the possible histories of the full system [55].

It is easy for an agent to learn to optimize its personal utility; the crucial problem is how to design mechanisms that derive benefit from the personal utility functions of the agents so that they “cooperate unintentionally” and optimize the global utility. In large multi-agent systems, it is common to design a global utility function where each agent has its own private utility function and the objective of the system is to optimize the global utility.

The main objective of using COIN for a large system of agents is to autonomously learn a set of actions in specific situations that maximize the individual private utility functions along the global utility, which essentially means that private utility functions need to be designed in such way that their maximization will maximize the global utility function. In this framework, the agents are trained to make independent decisions while automatically ensuring that overall decisions benefit all vehicles.

5.2 Intersection Traffic Management using COIN

Our agents in this framework are fully autonomous cars that know their location and destinations and communicate (perfectly) through DSRC with all nearby vehicles, gathering the approximated location and speed of each vehicle. These agents are trained in a simulation environment. It is possible to simulate realistic, less-than-ideal communication and sensing for

each vehicle, but this has not been explored yet because developing cooperative automation for vehicles is a complex task, even with ideal communication. They ideally should learn to get to their destination by avoiding accidents, keeping a safe distance to the car ahead of them, and also traverse the intersection in a timely manner and without causing issues for the other cars in the environment. COIN is addressing this optimization problem by decomposing it to many sub-problems, each of which can be solved individually (car level instead of intersection level). Using this approach does not mean that those problems are independent of one another, but that each individual solution is a contributor to solving the system-level problem.

5.2.1 System Evaluation

The system performance evaluation function that we use in this section focuses on avoiding collisions, directing vehicles to their destination if it is safe to do so, keeping safe distance to the vehicle ahead, reducing the number of unnecessary maneuvers, and doing all this in a timely manner. The global utility function, $G(z)$, is the sum of these terms. $G(z)$ more precisely can be expressed by its components as follows:

1. Speed measure: A negative reward that penalizes slower vehicles more
 - $D(z) = \sum_{j \in S} D_j(z)$, $D_j(z) = 2(S(Z_{g_j}) - S(Z_g^1)) / S(Z_g^1)$
 - Where $S(Z_{g_j})$ is the speed of car j at the instance of t
 - Where $S(Z_g^1)$ is the max speed allowed on that road

2. Following the correct path to the destination:
 - $F(z) = \sum_{j \in S} F_j(z)$, $F_j(z) = \sum_t -2U(Z_{l_j} - Z_l^1)$
 - Where Z_{l_j} is the path that car j takes at instance t
 - Where Z_l^1 is the correct path to destination for agent j

3. Proximity to the closest agent based on the condition:
 - $P(z) = \sum_{j \in S} P_j(z)$, $P_j(z) = \sum_t -4U(Z_{k_j} - Z_k^0)$
 - Where Z_{k_j} is the proximity to another agent ahead of agent j at time instance t
 - Where Z_k^0 is the safe distance for any agent at certain position with respect to its neighbouring agents

4. Collision reward:
 - $C(z) = \sum_{j \in S} C_j(z)$, $C_j(z) = \sum_t -1000U(Z_{m_j})$
 - Where Z_{m_j} is the collision incident found at instance t for the agent j

5. Intersection priority reward to prevent deadlocks:

- $I(z) = -12$ if the vehicle enters the intersection when there is an immediate conflict with another vehicle already in the intersection, and if this other vehicle's road currently has priority
- $I(z) = 0$ in any other case

The function U in all the equations above is the step function that equals 1 when its argument is greater or equal to zero, and has a value of zero otherwise. S is the set of the all agents and z is vector of state space and action of all the agents. Having defined the components of the global utility above, the global utility function is:

$$G(z) = D(z) + F(z) + P(z) + C(z) + I(z)$$

The global utility function is designed to penalize each outcome differently based on its seriousness, with collisions having a much higher penalty than all other results.

$D(z)$ provides the proximity of the running car to the maximum speed allowed on the road at each instance of time and charges the agent with the fractional amount of deviation from the max speed at that instance. $F(z)$ negatively rewards the cars that are not in the path to their predetermined destinations. $P(z)$ decreases reward of agents that choose to get close to the other cars, and $C(z)$ is a one-time heavy negative reward that the agent is charged at the time of collision.

5.2.2 Agent-Based Intersection Traffic Management

The multi-agent approach to management we present is predicated on agents evolving independent solutions that maximize the system evaluation function. Encompassing a realistic driving environment, and ensuring that the agents can be trained with reasonable computational resources and time, are also concerns. This section describes significant elements of our COIN system.

State Space

One of the biggest challenges for COIN is to pick the state space cleverly enough that it is not so large as to make the learning task intractable, but is inclusive enough as to always give the necessary information to make a decision. One natural choice of the state space is the detailed location of each agent on the road and their speed (agents' trajectories), but it is easy to see that this agent space is far too large. The full state-space environment we have picked for our system is available in the appendix.

Agent Action

The agent's action each step consists of changing speed, changing lane, and turning (at the intersection entrance only).

- Speed = {+10, -10, 0 (Km/h)},
- Lane Change = {Left, Right, None} and
- Turn = {Left, Right, Straight}

Agent Reward Structure Selection

We use the following difference evaluation function:

$$R_i = G(z) - G(z - z_i + c_i)$$

where z_i is the state space/action of agent i . Here z contains all components of all agents' state and action at each time instance, and those z components that are affected by agent i (z_i) are replaced with the fixed constant c_i . In this case, the simplest c_i is taking the agent i "off the grid" and reevaluating the global reward without the presence of agent i . There are two advantages to using a difference function like this. The first effect is on learnability, as the second term removes a significant portion of the impact of other agents in the system, which makes the evaluation function more agent-sensitive. Secondly, it affects the "factoredness": since the second term does not depend on the actions of agent i , any action by agent i that improves R_i also improves G .

Agent Evolution Algorithm Selection

The objective is that each agent learns to make decisions at any given state, which will lead to the best system evaluation G for that specific state. Here we assume that each agent will have a reward function (as described in the previous section) and will aim to maximize its reward using its own reinforcement learner. For complex delayed-reward problems, relatively more sophisticated reinforcement learning systems, such as temporal difference, may have to be used. However, due to our agent selection, the use of state space design and agent action set (described earlier) is needed to utilize immediate rewards. As a consequence, a simple table-based immediate reward reinforcement learning approach is used. Our reinforcement learner is equivalent to an α -greedy Q-learner with a discount rate of $\gamma=0$.

At every time instance an agent takes an action and then receives a reward evaluating that action, with which the Q table is being updated to represent the value of taking that action in that state, as follows:

$$Q_0(s,a) = (1 - \lambda) * Q(s,a) + \lambda * R(s,a)$$

where λ is the learning rate. At every time step the agent chooses the action with the highest table value with probability $1 - \alpha$ and chooses a random action with probability α . In the experiments described in this paper, α is equal to 0.33 and λ is equal to 0.5. The parameters were chosen experimentally, though system performance was not overly sensitive to these parameters. As an alternative, we also have used the following update for Q-Learning:

$$Q_0(s, a) = (1 - \lambda)Q(s, a) + \lambda \left(R(s^+) + \max_{a^+} Q(s^+, a^+) \right)$$

λ is the learning rate

s^+ is the next state, i.e. the state that was reached by performing action a

Based on our experiences for described State-Action space, the former Q-learning update formulation is more effective and converges faster during the training.

5.2.3 Intersection right-of-way

A common challenge in multi-agent systems is that all agents learn to act the same way, when in some scenarios it is necessary for two agents in similar situations to each choose a different action. For the task of autonomous navigation, there is the possibility that two or more vehicles may detect a potential collision with the other(s) and stop immediately. This is especially problematic if each vehicle stops while in the intersection, preventing other vehicles from progressing, as shown in Figure 5.1.



Figure 5.1: Example of intersection deadlock between three vehicles

This issue was addressed by assigning each road a priority, which varied over time. Vehicles that caused dangerous situations when entering from a low-priority road were penalized more harshly than those performing the same actions in a high-priority road. This is somewhat like a ‘soft’ traffic light, by which vehicles are encouraged (but not required) to follow a right-of-way in some situations.

5.3 Results of Collective Intelligence Training and Simulations

5.3.1 Training Specific Maneuvers

In this section, we describe a specific scenario for training and testing, providing some snapshots of simulations as well as statistical results obtained during both training and testing.

The best way to test the learning capabilities of this system is to see if it picks up the correct timing of a given maneuver. Figure 5.2 depicts vehicles performing left turns, in conflict with vehicles coming from the other side—in other words, an unprotected left turn. After less than 100 simulations, vehicles learn to wait when their path would cause a collision with an oncoming vehicle, as well as to speed up or slow down in response to the other vehicles taking the same route.




Figure 5.2: Snapshot of left-turn-training simulations

The results in this section are presented for a symmetrical four-way intersection; training and testing was completed for the cars that are making left turns at the intersection against oncoming traffic. First, the training was performed for a specific scenario with a constant number of cars, fixed time spacing between the cars, and fixed starting speed. The scenario was designed so that every two cars mutually meet at the intersection and learn how to avoid a collision as they proceed to their destination. The training converged quickly for this specific scenario.

The testing results are presented in Table 5.1. Testing results show a learned model in action, unlike training results, which are gathered while the model learns and thus should clearly perform less well. Without the existence of competing algorithms, the training results provide the simplest baseline to check against. For static situations (where the training examples match the testing examples), vehicles learn to totally avoid causing accidents while still following their desired path. Adding some randomness to the starting time between vehicles and without any further training, the number of incidents increases to 5% (rows 3 and 5 of Table 5.1). This is still unacceptable for actual use, but provides clear proof that the algorithm is learning basic collision avoidance. By comparison, the number of collisions during training exceeds 20%.

In the next step, training was completed for two-way traffic in the four-way intersection with a random number of cars, random starting time, and random starting speed. Testing was performed with similar a setup and the same factors being randomized. The number of unwanted incidents under 80 episodes of training increased, but the increased episodes in training ultimately resulted in a decrease in the number of incidents in testing, as more possible scenarios had been trained for (row 7 of Table 5.1).

Table 5.1: Results of COIN simulations on a two-way intersection

Scenarios	Number of episodes	Number of passes	Average travel time per car	Number of wrong paths taken by cars	Number of collisions	% of wrong path	% of collision
1- Static Training Scenario	30	180	20.31	39	64	0.217	0.356
2- Static Testing for Exact Same Training Scenario	80	480	16.02	0	0	0	0
3- Static testing + Some Minor Randomness in Scenario	80	492	16.10	19	18	0.039	0.037
4- Randomized Training Scenario	80	532	21.85	84	112	0.158	0.211
5- Randomized Testing Scenario	80	556	18.84	15	32	0.027	0.058
6- Randomized Training Scenario More Episodes	120	1548	27.39	97	445	0.063	0.287
7- Randomized Testing Scenario Using Training with More Episodes	80	717	19.86	28	34	0.039	0.047

The second testing scenario allows traffic to enter from two sides, but travel in any direction, as shown in Figure 5.3. While intersection deadlocks are unlikely, this scenario covers most possible forms of collision, including lane-change conflicts. As shown in Table 5.2, 300 intersection simulations, involving over 6,000 vehicles, were necessary to train the model to a functional level, though there is evidence that further training will further improve the model. Training for a full four-way intersection is currently underway.




Figure 5.3: Screenshot of two-entry intersection simulations

Table 5.2: Results of COIN simulations on a two-entry, four-exit intersection

Number of episodes	Number of vehicles	Average travel time per car	Number of wrong paths taken by cars	Number of collisions	% of wrong path	% of collision
40	1000	11.1	13	24	0.013	0.048

Chapter 6. Developing a Framework for Joint Millimeter Wave Communication and Radar

6.1 Introduction

Vehicular communication and radar sensing are the two primary means of using radio frequency (RF) signals to improve traffic safety and efficiency. So far, we have developed advanced CW/CA techniques assuming the information from automotive sensors and vehicular communication is given. For the rest of sections in this report, we will focus on these two techniques and develop advanced vehicular communication systems using mmWave.

Automotive radars on a source vehicle derive information about the environment (e.g., location and velocity of the surrounding target vehicles) by first sending a special waveform, typically a frequency-modulated continuous waveform (FMCW), and then receiving and processing the reflected echoes [56, 57]. Long-range radar (LRR) operates in the 76–77 GHz mmWave band and is used for adaptive cruise control. In the newer 77–81 GHz mmWave band, medium-range radar (MRR) will be used for cross traffic alert and short-range radar (SRR) will be used for parking aid and pre-crash applications. LRR is used for adaptive cruise control, MRR is used for cross traffic alert, and SRR is used for parking aid and pre-crash applications. The radars provide a high-resolution point-map for continuous automatic detection, but use proprietary waveforms and are expensive. They are already deployed in many luxury vehicles [5].

V2V communication allows vehicles to achieve real-time cooperative detection and ranging using standards-based waveforms and protocols [58]. DSRC is a low-latency wireless communication protocol that operates using a WLAN-based physical layer in the 5.9 GHz microwave band [9]. It is used to send short, low-latency messages about the status of a vehicle for applications such as cooperative forward CW and cooperative adaptive cruise control [58]. DSRC can only support data rates in the range of 3–6 Mbps in practice [9]. The low data rate may restrict novel applications and enhanced active safety functions, which would require sharing raw sensor data at low latency [12, 59]. One such application is vehicular cloud-based autonomous driving, where vehicles generate around 1 Gbps of data from sensors [60].

One potential means to realize the next generation of high-data rate connected vehicles is to exploit the large bandwidths available in the mmWave spectrum. There are several ways to connect vehicles using the mmWave band [61]. A modification of the mmWave consumer WLAN standard, i.e., IEEE 802.11ad, is one of the means for high data rate V2V connectivity. The forthcoming mmWave 5G standard could be another way to provide vehicular connectivity either directly using device-to-device or through future employed cellular infrastructure. The effectiveness of vehicular communication is, however, dependent on its wide-scale deployment.

Although both radar and communication technologies have applications to driver-assist and autonomous driving, they have their own domain-specific challenges and limitations. A joint communication and radar system that shares the same spectrum and hardware will, however, lead to an increase in the penetration rate of communication and radar in commodity vehicles. It will also reduce size and cost of the hardware with efficient spectrum usage and enhanced security. Additionally, using both technologies simultaneously in an integrated unit will allow vehicles to reap the advantage of each technology (e.g., radar for non-communicating traffic and V2V for distances beyond the line-of-sight constraints of radar) and enhance their performance by sharing information with each other.

In the past half-decade, a number of approaches for joint radar and communication that exploit existing radar and communication waveforms have been considered [62]. The approaches can be mainly classified into a joint system, where a single-carrier or a multi-carrier waveform is used for both communication and radar simultaneously, and a time-domain duplex system, where radar and communication will operate in different time cycles. In most of the earlier approaches, the waveforms are designed in a non-standardized fashion. We, however, leverage the mmWave WLAN standard for developing a joint vehicular communication and automotive radar system. In particular, we develop an IEEE 802.11ad V2V-radar, which is an IEEE 802.11ad-based joint V2V communication and LRR paradigm that can be implemented with a low-cost IEEE 802.11ad chip. Our work in [63] is the first to propose the idea of using IEEE 802.11ad for a joint V2V and automotive radar system. This approach motivates a common standard for automotive radar and vehicular communications at the mmWave band. Indeed, the most prevalent V2V standard, DSRC, is based on IEEE 802.11p, which is an evolution of a WLAN standard known as IEEE 802.11a. The most similar work to ours is [64], which analyzed the IEEE 802.11p V2V communication standard for automotive radars. IEEE 802.11p, however, operates at 5.9 GHz and not at the mmWave frequency bands. Using a mmWave standard will provide significant advantages in terms of higher data rates for communication and better accuracy/resolution for radar operation than IEEE 802.11p. The mmWave standard will help simultaneously achieve ultra-low latency and high range of operation for automotive safety applications with minimal hardware size and cost.

6.2 Radar System using Low-Frequency WLAN Signals

6.2.1 Forward Collision Vehicular RADAR with IEEE 802.11

In the beginning of this project, we expected that mmWave equipment would be available during the project. The equipment, however, is not ready for commercialization from any vendor yet. Since mmWave equipment is not yet available, we explored the experimental development of a joint radar and communications framework using low-frequency WLAN signals. Our goal was to develop a radar system that integrates directly into the DSRC standard, enabling joint radar and communication on any vehicle equipped with DSRC. A feasibility demonstration was performed using existing IEEE 802.11 devices with minimal modification through algorithm processing on frequency-domain channel estimates. The results of our work showed that our solution delivered similar accuracy and reliability to mmWave radar devices by increasing allocated spectrum from the 10 MHz DSRC standard to 20 MHz, indicating significant potential for industrial devices with joint vehicular communications and radar capabilities at DSRC frequencies.

6.2.2 Prototype

Initially, we tested the performance of this design via simulation. Results demonstrated accuracy up to 5 m for a 10 MHz channel and meter-level accuracy for a 20 MHz channel. We extended the results of our simulation to a radar and communications framework prototype (Figure 6.1) using a universal software radio peripheral (USRP) reconfigurable I/O transceiver, which enabled us to directly extract IEEE 802.11 channel estimates. We also used two broadband patch antennas for our transmitter and receiver, and developed our code using a desktop machine. Processing was performed in real time on the desktop machine using data from the radio transceiver, enabling single-target range detection.



Figure 6.1: IEEE 802.11 joint radar and communications link setup for measurements. During measurements the antennas were not as close as in the configuration shown. The edges of these antennas were separated by a minimum of 0.5 m.

6.2.3 Results

We tested our prototype for single-target localization at a local parking lot, using a 2002 Toyota Camry as the target vehicle. Initially, we positioned the prototype at a distance of 30 m from the target vehicle. To collect measurements, we transmitted IEEE 802.11 beacon messages every 250 milliseconds. After a sufficient number of channel estimates were obtained at 30 m, we moved the prototype 5 m towards the target. This process was repeated until measurements were obtained at all 5m target distance increments. Our results are shown in Figure 6.2. The results show that the developed radar system using WLAN signals can achieve m-level range estimation accuracy.

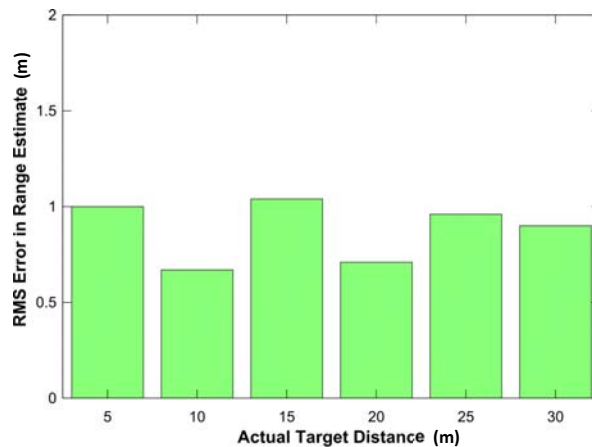


Figure 6.2: RMS range error using IEEE 802.11 packets in a 20 MHz channel with one target. The target has variable range from 5 to 30 m.

6.3 MmWave Joint Radar and Communication System

6.3.1 Framework for Joint System

In this section, we formulate the conceptual framework for the joint automotive radar and V2V communication system based on the IEEE 802.11ad standard. First, we discuss the traffic scenario of interest and the transmit signal model. Then, we model the one-way communication channel and the two-way radar channel to derive the received signal model for both communication and radar systems.

We consider a traffic scenario where a source vehicle sends a waveform, $s(t)$, to a target vehicle using the IEEE 802.11ad-based V2V communication service. The IEEE 802.11ad waveform may get reflected back from the target vehicle and the other surrounding scatters (e.g., trees, road, and the other remote vehicles) as shown in Figure 6.3. Then, the source vehicle receives these reflected echoes from the scatterers and derives information about the target vehicle. We assume a multiple antenna joint communication-radar system with N_T co-located transmit (TX) antennas and N_R co-located receive (RX) antennas mounted on all vehicles under consideration. The TX and the RX antenna arrays on the source vehicle are closely separated such that both arrays will see the same location parameters (e.g., azimuth/elevation angle and range) of a scatterer. We also consider that the TX/RX beams of the source vehicle are pointed towards the target vehicle without blockage and that the 3-dB beamwidth of the TX and RX beams are narrow during mmWave V2V communication [65]. Although a very narrow beamwidth will lead to less clutter interference and long range of operation, it can yield poor performance with vehicle mobility and blockage. Hence, we assume that the TX/RX beams are narrow enough to meet the link budget requirement of V2V communication but are wide enough to illuminate all the scattering centers of a far target vehicle within their resolution. Therefore, we represent the target vehicle as a single point target, as in [66, 67].

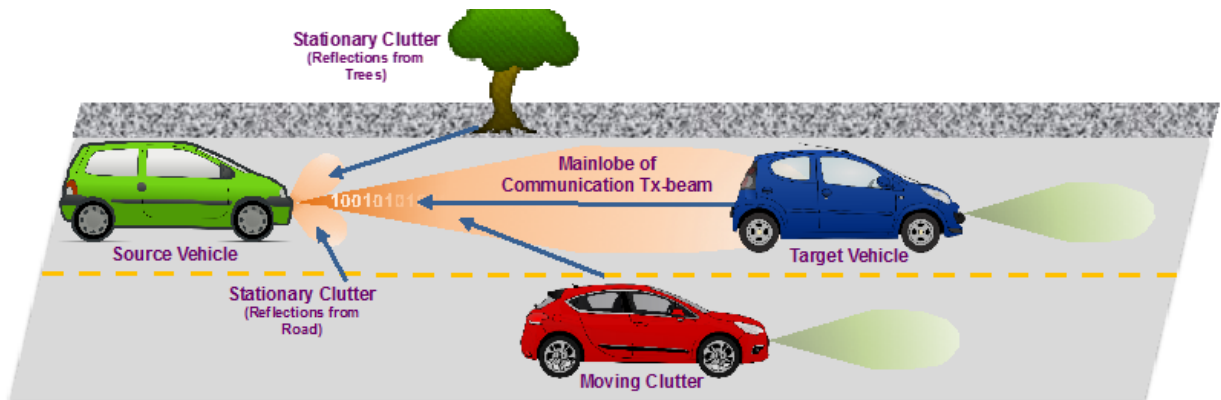


Figure 6.3: Illustration of a traffic scenario for joint automotive radar and vehicular communication systems using IEEE 802.11ad.

We now consider a single coherent processing interval (CPI) of $T=MK T_s$ duration, which is composed of M frames of an IEEE 802.11ad single carrier physical layer (SCPHY), each of K samples with a sampling interval of T_s , as shown in Figure 6.4. An IEEE 802.11ad SCPHY frame is composed of a short training field (STF), a channel estimation field (CEF), a header, and data blocks. The preamble of the IEEE 802.11ad frame is composed of the STF and the CEF and is

generated from a 128-sample Golay complementary pair (GCP) of two 128-chip Golay complementary sequences (GCSs), termed \mathbf{Ga}_{128} and \mathbf{Gb}_{128} , as shown in Figure 6.5 and Figure 6.6. The STF is used in communication for frame synchronization and frequency offset estimation. The CEF is used to estimate the communication channel parameters and to indicate the modulation of the packet (e.g., SCPHY or OFDM [orthogonal frequency division multiplexing] physical layer).

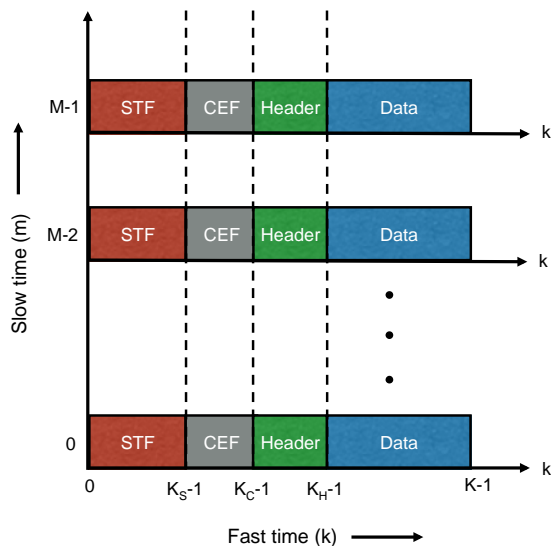


Figure 6.4: Illustration of a CPI that consists of M frames, each of K samples. The end positions of the STF, the CEF, and the header are K_s-1 , K_c-1 , and K_H-1 , respectively.



Figure 6.5: Extracted short training field for a SCPHY frame.




Figure 6.6: Extracted channel estimation field for a SCPHY frame.

The composite ambiguity function of the 128-sample GCP, shown in Figure 6.7, motivates its suitability as a radar waveform [68]. The zero-Doppler cut of this function indicates that the GCP has a perfect auto-correlation with no sidelobe along the zero Doppler axis. This characteristic makes it ideal for target detection in radar applications, which does not exist in FMCW signals typically used in LRR [69]. This figure also shows that the GCP is less tolerant to

large Doppler shifts. These sequences, however, seem to be appropriate for LRR due to its low Doppler shift requirement.




Figure 6.7: (a) The composite ambiguity function of the 128-sample GCP used in the preamble of IEEE 802.11ad. (b) The zero-Doppler cut of the composite ambiguity function of (a).

The mmWave sensing channel during a CPI is composed of a few scattering centers [70] that represent reflections from the target vehicle and the other surrounding objects. The vehicular channel has a key characteristic of temporal variability and inherent non-stationarity [71]. The description of channel in azimuth and elevation directions is also critical for mmWave arrays used in vehicular radar and communication applications [19], [26]. Therefore, we model the mmWave channel for a single CPI as a doubly selective (time- and frequency-selective) mmWave channel with a few N_p dominant paths and 2D TX and RX steering vectors. Each p^{th} path is described by five physical parameters: its azimuth and elevation angle of arrival (AoA) pair $(\phi_{R,p}, \theta_{R,p})$, AoD pair $(\phi_{T,p}, \theta_{T,p})$, delay τ_p , complex gain α_p , and Doppler shift ϑ_p .

Without loss of generality, we assume that the line-of-sight two-way path from the source vehicle to the single point target (the target vehicle that is also the user of the V2V communication link established by the source vehicle) is represented by the u^{th} path. The point target is assumed to be at an arbitrary range $\rho_u(t)$ with a two-way round-trip propagation delay of $\tau_u(t)$ from the reference point on the TX array of the source vehicle to its reference point on the RX array. In vehicular applications, a target vehicle is a slowly moving target compared to the speed of light c and hence, a quasi-stationary assumption can be made. This assumption implies that the range change during the short path of any particular point in the waveform from the transmitter to the receiver is negligible. With this assumption, the two-way round-trip propagation delay at time t is

$$\tau_u(t) = \frac{2\rho_u(t)}{c}. \quad (6.1)$$

We assume that the target velocity relative with respect to the source vehicle is small enough to allow for constant location, that is, constant $\rho_u(t)$ and $\tau_u(t)$ (we drop t from $\rho_u(t)$ and $\tau_u(t)$ because they are assumed to be constant for the time of interest), and azimuth and elevation angle pair (ϕ_u, θ_u) , during the CPI [27]. We also assume that the target vehicle has an arbitrary relative radial velocity of v with respect to the source vehicle. It remains constant within the CPI because of small acceleration. The Doppler shift, therefore, can be represented by

$$\nu_u = 2\frac{v}{\lambda}, \quad (6.2)$$

where λ is the carrier wavelength. In the channel model, we only consider far targets whose $\rho_u(t)$ is large compared to the distance change during the CPI. Hence, we assume a constant complex gain α_u for the target vehicle [73].

IEEE 802.11ad supports multiple antenna communication with a single data stream. Spatial multiplexing as found in IEEE 802.11n/ac is not supported. Therefore, we incorporate the TX/RX analog beamforming vectors into the baseband model even though the actual beamforming may happen at an intermediate frequency or RF.

We assume that the source vehicle attempts to align its TX/RX beams towards the target vehicle using the IEEE 802.11ad beam training approach while establishing the communication link between them. Hence, once the link has been established, the TX and RX beams of the source vehicle are assumed to be pointing towards the (ϕ_u, θ_u) direction with a small beam alignment error. At the same time, the RX beam of the target vehicle will also point towards the (ϕ_u, θ_u) direction with some beam misalignment to receive the V2V communication signal from the source vehicle. Assuming same IEEE 802.11ad-based beamforming codebook at the source and the target vehicles, \mathbf{w}_T and \mathbf{w}_R denote the TX and RX analog beamforming vectors both at the source vehicle and also at the target vehicle that depends on (ϕ_u, θ_u) direction.

The received communication signal at the target vehicle due to the one-way path depends on one-way target delay $\tau_u/2$, target range ρ_u , Doppler shift corresponding to the relative velocity of the source vehicle with respect to the target vehicle $-\mathcal{D}_u$, and receiver noise. The echo reflected from the target vehicle, $y(t)$, when received at the source vehicle, depends on two-way target delay $\tau_u/2$, target round trip distance $2\rho_u$, Doppler shift corresponding to the relative velocity of the target vehicle with respect to the source vehicle \mathcal{D}_u , and receiver noise. This received radar signal, $y(t)$, also depends on the radar cross section of the target vehicle and is affected by the unwanted clutter from the surrounding environment and self-interference factor due to the inherent full-duplex assumption at the source vehicle. We assume that the full-duplex system we are using has a good enough self-interference cancellation mechanism; therefore, we can ignore self-interference effects inherent in the joint system. Due to variability of small-scale fading with scatter distribution, location and orientation, it is reasonable to assume that the small-scale fading corresponding to the communication channel is independent of the small-scale fading corresponding to the radar channel and the receiver noise. For each received frame, the phase shift corresponding to the training sequence is very small due to small relative velocity of the target vehicle; therefore, we can assume the channel to be time invariant within a single frame.

6.3.2 Proposed Receiver Processing Techniques for Enabling Radar Functions

We consider three primary types of radar processing: 1) *vehicle detection* using a constant false alarm rate algorithm; 2) *range estimation* using a time synchronization technique; and 3) *velocity estimation* using a frequency synchronization technique.

Radar processing exploits the special structure of GCP/GCS in the preamble of 802.11ad frames and leverages the communication preprocessing to detect and estimate its parameters of interest, as shown in Figure 6.8. Indeed, the algorithms used in the radar module are based on the pulse-Doppler radar processing and developed by extending the methods used in communication techniques over a single frame to multiple frames [74]. This approach enables the realization of a

joint vehicular communication and radar paradigm using a conventional low-cost IEEE 802.11ad scheme with minimal receiver modifications.

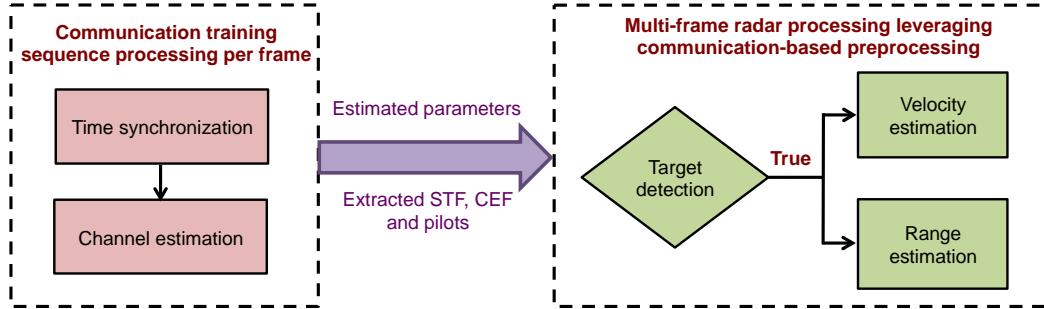


Figure 6.8: The flowchart represents the processing techniques for target detection and range/velocity estimation using IEEE 802.11ad V2V-radar.

The processing techniques leverage the special structure of GCS and GCP present in the STF and the CEF of multiple frames in one CPI for desired automotive radar performance.

In IEEE 802.11ad, the training sequences of a single frame are used for time and frequency synchronization and channel estimation [75]. This is achieved in several steps: 1) *coarse time synchronization* based on preamble detection techniques using the STF; 2) *frequency offset estimation* using the STF; 3) *fine time synchronization* using the CEF symbol boundary detection and the STF/CEF peak detection techniques; and 4) *channel estimation* using the CEF.

The first step of the preprocessing is to detect the IEEE 802.11ad frame using the STF. The frame start detection technique applies a threshold on the normalized auto-correlation of the received STF to coarsely estimate the starting sample of the preamble. The fine range estimate of the time-delay can be obtained either by using an amplitude-based method or a phase-based method. The amplitude-based method estimates the fine time-delay, by detecting the peaks of cross-correlation between $\mathbf{G}a_{128}$ and multiple GCSs in the STF sequence. The amplitude-based fine timing synchronization can be performed by applying the peak detection technique on the CEF instead of the STF in the coarse range estimation. Both the peak detection methods perform well at low signal-to-clutter-plus-noise ratio (SCNR). The timing synchronization can also be fine-tuned by performing phased-based symbol boundary detection (SBD) using the CEF [73]. This method, however, does not perform well in the presence of Doppler shift at low SCNR. After the fine time synchronization, we extract the received CEF signal to estimate the channel using a 512-sample GCP.

Reflected echo from the target vehicle is detected using the constant false alarm detection based on the typical WLAN non-zero channel tap detection [76]. In this technique, the decision is based on a simple thresholding function based on constant false alarm probability P_{FA} and the variance of the clutter-plus-noise received at the source vehicle. Figure 6.9 shows the probability of detection P_D using different false alarm probabilities. It indicates that P_D grows with increasing P_{FA} . For a P_{FA} of 10^{-4} , we achieve radar detection rates greater than 90% above the received SCNR of 0 dB.




Figure 6.9: Probability of detection using different constant false alarm detection rates.

The relative velocity of the target vehicle is estimated by applying (6.2) with the calculated Doppler frequency of the target echo. To estimate the Doppler shift corresponding to the target vehicle, we use the least squares (LS)-based frequency-offset estimation method over single/multiple frames. The resolution and accuracy of frequency-offset estimation will be significantly enhanced when we choose to use multiple frames (similar to pulse-Doppler radar) as compared to a single frame (traditionally used in frequency synchronization algorithms of a standard WLAN receiver) because of larger integration time. The theoretical performance of LS-based frequency-offset estimation can be evaluated using Cramer Rao lower bound (CRLB) of the velocity estimate. The CRLB expresses a lower bound on the variance of velocity estimators based on the preamble of M frames. If the CRLB is above the LRR's desired mean square error (MSE) for velocity estimation, then it indicates that the requirement for LRR velocity accuracy cannot be met in any case.

Figure 6.10 shows MSE of the estimated relative velocity using the STF of a single frame and using the preamble of two frames with $K = 41285$ symbols. The performances of velocity estimation techniques increase linearly (in dB scale) with the SCNR. The LS-based estimation technique that we have used is comparatively better than the one proposed in [75]. The accuracy of LS-based estimation techniques is very close to its CRLB bounds. Using double frames enables us to achieve much better velocity estimation accuracy than using a single frame for all SCNR values. At low SCNR (less than 10 dB), however, even the use of double frames does not enable us to achieve the desired velocity accuracy of 0.1 m/s.




Figure 6.10: MSE of the velocity estimation using the STF of a single and the preamble of the double frames. The numerical results of proposed estimation techniques closely match to the CRLB bounds.

For unambiguous velocity estimation, the accuracy and resolution of velocity estimation grow with the increase in the total training sequence duration and the numbers of frames within a fixed size CPI. The number of communication data symbols, however, decreases with an increase in the training sequence duration. Thus, there is a system trade-off between target velocity estimation accuracy and communication data rate for the number of frames within a CPI.

Once the target vehicle is detected, the range of the target cell from the source vehicle is calculated from (6.1) by estimating the corresponding delay-shift. The range estimation algorithms are applied on the STF and the CEF and can be categorized into two main types: coarse range estimation, using a frame start detection estimate with an error of less than 30 m, and fine range estimation based on SBD boundary detection, and the STF/CEF peak detection technique based on the fine time estimate with an accuracy of 0.1 m, which meets the LRR specifications [77]. The theoretical CRLB bound for velocity estimation signifies that it is possible to achieve cm-level resolution/accuracy using a single frame of IEEE 802.11ad with a bandwidth of 2.16 GHz for SCNR above 0 dB. We can also see from the numerical results in Figure 6.11 that we can achieve the desired range resolution/accuracy of < 0.1 m using the time-delay estimate calculated via communication-based processing.




Figure 6.11: MSE of the range estimation using the preamble in a single frame based on coarse and fine range estimation algorithms.

Using numerical simulations, we compare the performance of various proposed range estimation algorithms and the CRLB bound using a single frame, as shown in Fig. 2.11. The desired range MSE for automotive radars is 0.01m [77]. We observe from Figure 6.11 that the fine range estimation achieves the desired accuracy using the STF/CEF peak detection for SCNR above 0 dB, and using the CEF symbol boundary detection for SCNR above 6 dB. The poor performance of range estimation using the CEF symbol boundary detection at low SCNR can be attributed to the fact the performance of the phase-based estimation gets affected by Doppler shift. The figure also shows that the performance of the frame start detection using the preamble degrades due to a constant threshold, which does not adapt to the increasing SCNR. The amplitude-based peak detection technique using the STF/CEF, however, meets the desired automotive range accuracy of 0.1 m using a single frame without incurring significant complexity.

6.4 Summary

In this chapter, we demonstrated feasibility for IEEE 802.11-based radar in a vehicular environment. Our design can be implemented on existing IEEE 802.11 WLAN devices with minimal modification to the physical layer, supporting a secure and extremely cost-effective design. Although our study was performed using a 20 MHz channel, it is expected that with a more refined algorithm, we can achieve meter-level accuracy with only 10 MHz of bandwidth, the standard for DSRC. This may lead to direct integration of IEEE 802.11p-based radar into the DSRC protocol. Additionally, the results in this study provide insight on various techniques and improvements that can be made for a joint radar and communications framework at mmWave frequencies. The details of this work can be found in [78].

For mmWave spectrum, we developed an automotive radar system using the low-cost IEEE 802.11ad standard, which exploits mmWave. We showed that the proposed radar system can achieve the performance of a commercial automotive LRR with minimal modifications on the IEEE 802.11ad system. To be specific, the proposed radar system can simultaneously achieve 0.1 m range and 0.1 m/s velocity estimation accuracies with Gbps data rates. The details of this work can be found in [79]. Using obtained results, we will build prototypes to demonstrate our joint mmWave radar and communication system.

Chapter 7. Optimizing the Joint Waveform to Meet Different Objectives

7.1 Introduction

Surface transportation safety can be enhanced by the use of wireless technologies, mainly automotive radar and V2V communication. Automotive radar, which requires line-of-sight, provides a high-resolution low-latency approach for automatic detection and ranging of the local area around the vehicle. V2V systems can establish collaborative communication between vehicles that may be hundreds of feet away to achieve a real-time cooperative detection and ranging. By mutually sharing their sensor data, especially lightly processed radar, camera, and LIDAR information, vehicles can overcome the line-of-sight limitation of radar and improve their awareness of the surrounding environment. However, exchanging raw data could easily require hundreds of megabits per second or even gigabits per second, which becomes a challenging task for current V2V communication systems. This is the main thrust of the developed framework on mmWave wireless technologies, which aim to provide a high data rate for communication with low latency as well as better accuracy and resolution for radar operation.

Due to the limited spectrum resources of wireless systems, exclusive allocation of the spectrum resource for either communication or radar may neglect the benefit of the other technology. Combining both functionalities on a common hardware platform brings the advantages of efficient spectrum usage, reduced cost and size, and better performance for the vehicles of tomorrow. A number of joint communication-radar approaches have been considered that exploit the existing radar/communication waveforms. However, these waveforms are usually ad-hoc designed and are not completely integrated as they rely on time [80], frequency [81], or code division ideas [82]. In a fully integrated communication and radar system, optimizing parameters to meet different performance objectives is an important part of having a timely and meaningful wireless system for transportation safety purposes. Depending on the traffic scenarios, we investigate how the resources, such as spectrum and power, can be dynamically allocated to maintain a balanced performance trade-off between the communication and radar functionalities.

7.2 Defining Different Traffic Scenarios

In this section, we present three different traffic scenarios: at urban intersections, at rural roadways (specifically during passing maneuvers), and at university-based or similar locations with substantial pedestrian and bicycle traffic. In all of these instances, the optimal fusion of information from these different sources will be a key issue. These scenarios can be used to benchmark different joint waveform optimizations for communication and radar. The optimization will have to be dynamic based on vehicular and communication traffic. For example, if many vehicles have communication and radar capability, their signals may create interference. This means that additional considerations are needed to provide both communication and radar efficiency. For example, resources might have to be shared among different neighboring vehicles in a distributed fashion. Such functions are common in wireless communications (a key role of the medium access control protocol) but have never been developed for a waveform that also includes radar.

Scenario 1: At urban intersection

Among all traffic accidents in Texas, intersection and intersection-related crashes make up about 40% of total crashes. In addition, 60% of serious crashes (those involving one or more fatalities) occur at urban intersections. It is important for a joint communication and radar system to prevent collisions even if this leads to uncomfortable driving situations. For example, sudden stops or accelerations between two vehicles based purely on radar systems can have a cascading negative safety impact on other vehicle movements. In the context of urban intersections, there are a variety of geometric roadway configurations, such as three-way angular, four-way perpendicular, and four-way angular. Radar system alone will not be able to detect vehicles coming from other traffic directions, if the radar signal and echo are blocked by buildings and other vehicles. This is where communication may play a more important role than radar due to the possible non-line-of-sight propagation of communication signals. In this case, more information can be derived from communication through other vehicles to expand the sensing range and understand the situation several car lengths ahead. Therefore, more resources should be devoted to communication and the joint waveform should be optimized to maintain a certain minimum communication throughput. The radar functionality of the joint waveform would bring additional benefits for collision detection. The comparison between radar and communication systems for this scenario is summarized in Table 7.1.

Table 7.1: Requirements for joint radar and communication waveforms at urban intersections

	Purpose	Required Resources	Required Performance
Radar	Line-of-sight detection of traffic	Moderate	Short range detection of traffic
Communication	Expansion of sensing range with shared radar signal	High	Detection of potential hazards at several car lengths ahead or traffic coming at different directions

Scenario 2: At rural roadways

The rural setting represents an extreme case where there are few vehicles around. Therefore, the majority of system resources might be devoted to radar to look for pedestrians, bicycles, or animals. Nevertheless, communication might play an extremely important role to prevent collisions in overtaking maneuvers on rural roads, which accounts for 82% of head-on fatal collisions in Texas. While radar can prevent collisions between a passing vehicle and other vehicles, the resulting abruptness of preventive maneuvers can cause dangerous downstream safety consequences.

In the context of passing maneuvers as illustrated in Figure 7.1, while the radar functionality of the car in front (passed car) is pivotal to detect the incoming traffic, the follow-up car (passing car) may have its radar signal blocked. Using the long-range communication link between vehicles, the passing car can avoid a collision course with the incoming traffic well before a collision is indicated on its radar alone. Information about different types of vehicles (including trucks) and vehicle kinematics, alternative roadway sight distance configurations, as well as different types of drivers (in terms of age, height, PR times, etc.) can also be transferred via the communication link and considered in the decision-making process of the overtaking maneuver. Although the V2V communication burden may not be as high as in urban intersection scenario,

the rural roadway scenario requires a reliable warning signal mechanism for the passed car. Thus, the joint waveform should be optimized to reserve the communication bandwidth for longer range “talk,” improve reliability, and reduce false alarms. The comparison between radar and communication systems for this scenario is summarized in Table 7.2.

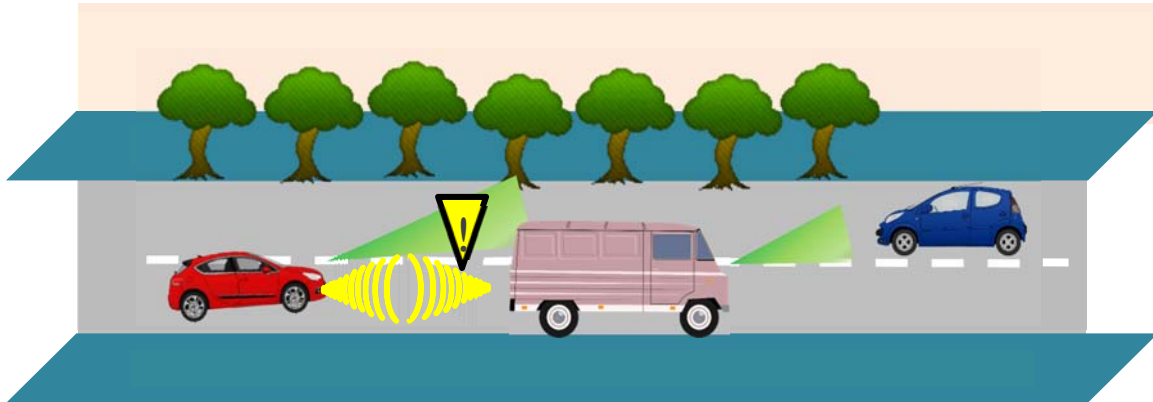


Figure 7.1: The passing maneuver on a rural road where the radar signal of the passing car is blocked by the car in front (passed car). In this case, the communication link between the passed car and the passing car can help the latter to detect the incoming traffic.

Table 7.2: Requirements for joint radar and communication waveforms at rural roadways

	Purpose	Required Resources	Required Performance
Radar	Detection of incoming traffic	High (passed car) Low (passing car)	Long range detection of incoming traffic
Communication	Expansion sensing range with shared radar signal	High (passed car and passing car)	Reliable transmission and reception of warning signals

Scenario 3: At locations with substantial pedestrian and bicycle traffic

Pedestrian and bicyclist fatalities make up about 14% of all Texas roadway fatalities, which is a significant fraction of all fatalities. Locations with high levels of non-motorized traffic should be particular opportunities for reduction of pedestrian and bicyclist fatalities using radar systems. Since pedestrians and bicyclists will not have the communication gadgets installed on vehicles, radar is essential to improve the collision detection and avoidance. Fortunately, mmWave radar provides a high-resolution (range, velocity, and angle), low-latency, compact (small antenna), and single packaging solution to meet the requirements of driver assist functions. In this context, the joint waveform in the mmWave band must be also optimized for multi-target detection with tracking capability. The V2V communication functionality of the joint waveform, however, should not be neglected, as it allows the vehicles to overcome the blockage of radar signal, thus facilitating collaborative detection and tracking of multiple pedestrians and bicycles several car lengths ahead. The comparison between radar and communication systems for this scenario is summarized in Table 7.3.

Table 7.3: Requirements for joint radar and communication waveforms at locations with substantial pedestrian and bicycle traffic

	Purpose	Required Resources	Required Performance
Radar	Detection and tracking of pedestrian and bicycle traffic	High	Multi-target detection with tracking capability
Communication	Expansion sensing range with shared radar signal	Moderate	Reliable sharing of radar signal for collaborative detection and tracking


7.3 Optimizing between Data Rate and Radar Performance

We develop a joint framework of LRR and V2V communication at 60 GHz by exploiting the special preamble structure (repeated GCSs) of an IEEE 802.11ad SCPHY frame. This framework leverages the signal processing algorithms used in the typical WLAN receiver for time and frequency synchronization to perform radar parameter estimation. In particular, we consider a traffic scenario where a source vehicle sends M multiple frames to a target vehicle using the SCPHY layer of IEEE 802.11ad-based V2V communication service within a CPI. The IEEE 802.11ad waveform may get reflected back from the target vehicle and the other surrounding scatterers (e.g., trees, road, and the other remote vehicles). Then, the source vehicle receives these reflected echoes from the scatters and derives information about the target vehicle. In this framework, there is a system trade-off between velocity estimation accuracy and the communication data rate of the target vehicle for the number of frames within a fixed size CPI.


For unambiguous velocity estimation, the accuracy and resolution of velocity estimation grow with the increase in the total training sequence duration and the numbers of frames within a fixed size CPI. This dependence can be evaluated using the MSE of the velocity estimation for the P sample preamble across a variable number of frames, M , within a fixed size coherent processing duration. Increases in M will lead to a longer duration of the training sequence and thus decrease the MSE of the velocity estimation, i.e., better radar performance.

The number of communication data symbols, however, decreases with an increase in the training sequence duration; therefore, it decreases with an increase in M . This dependence can be evaluated using achievable data rate as the performance metric of the communication system.

To evaluate the dependence of velocity estimation on the number of frames within a CPI and investigate its simultaneous effect on the communication data rate of the system, we have performed simulations over several CPI intervals with a varying number of frames at 10 dB received SCNR at the source vehicle, as shown in Figure 7.2.



(a) Data rate of the IEEE 802.11ad system for a fixed duration of CPI



(b) MSE of the velocity estimation using multiple frames for a fixed size CPI

Figure 7.2: Trade-off between communication data rate and velocity estimation for a fixed size CPI. By increasing the duration of training symbols within a CPI, velocity estimation becomes more accurate with reduced data rate.

We consider the radar cross-section of the target vehicle as 10 dBsm (decibel relative to one square meter) [83], receiver noise figure as 6 dB [84], transmit EIRP (equivalent isotropically radiated power) of 82 dBm (the average EIRP for 60 GHz devices with antennas located outdoors [85]), waveform bandwidth of 2.16 GHz [86], and a target range of 60 m (which falls within the typical span of long-range radar range specifications [87]). For a fixed CPI duration, the number of frames is varied from one to the maximum limit within a CPI. We observe from the simulations that as the number of frames increases within a fixed CPI, the communication data rate degrades while enhancing the velocity estimation accuracy. In linear scale, the communication data rate decreases linearly with a linear increase in the number of frames. In logarithmic scale, the velocity estimation decreases linearly with logarithmic increase in the number of frames. The decrease in velocity estimation is, however, steeper than the growth in communication data rate for a given increase in the number of frames. In spite of this trade-off, we observe that it is indeed possible to achieve Gbps communication data rate and cm/s-level accurate target velocity estimation simultaneously.

7.4 Summary

In this chapter, we identified three key traffic scenarios, i.e., urban intersections, rural roads, and at locations with substantial pedestrian and bicycle traffic, where we can harness the full benefit of the joint radar and communication systems. Because the performance metrics of radar and communication systems are different, joint mmWave radar and communication systems should be optimized depending on specific traffic scenarios. The optimization can be performed by dynamically allocating resources, e.g., time, spectrum, and power, to the system that needs better performance. We investigated the impact of frame (i.e., time) allocation on the performance of radar and communication systems in this study. The details of this work can be found in [79].

The dynamic allocation of radar frames and data communication frames becomes especially beneficial at the urban intersection scenario, where more data frames can be allocated to meet the high performance requirement of communication. Further optimization between radar and communication performance is possible depending on other traffic scenarios. For example, at rural roadways where LRR is needed, dynamic power allocation allows the joint system to transmit radar signal at higher power and expand the radar detection range. At locations with substantial pedestrians and bicycles, the waveform optimization with radar signal, for example, using FMCW, is important to improve the capability of multi-target detection. We will further investigate possible optimization techniques for the joint mmWave radar and communication systems.

Chapter 8. Incorporating Antenna Arrays

8.1 Introduction

Antenna arrays are important in mmWave communication systems. The antennas allow the transmitter and receiver to form sharp beams with high gain. This is useful in providing higher throughput and lower interference in communication, and also aids in determining position in radar. There are many trade-offs to be made in antenna arrays development. For example, a vehicle may point its beam to another vehicle to communicate more data, but scanning that beam around may lead to better detection about the environment. Another important aspect of antenna arrays is the trade-off between power consumption and performance. The use of a mixture of analog and digital beamforming, which is often called *hybrid beamforming*, can dramatically reduce the power consumption associated with beamforming in the communications application.

In this chapter, we proposed possible antenna array designs and developed hybrid beamforming algorithms for the joint mmWave radar and communication system. Since communication systems may require channel state information between the vehicle that is transmitting and the vehicle that is receiving the information, we also developed algorithms to acquire mmWave channel state information using lower frequency channel information.

8.2 Alternate Array Design

In a joint mmWave radar and communication system, it is essential to have a large number of antennas at the transmitter and the receiver (in communication perspective) to form sharp transmit and receive beams and establish good link quality. Due to small wavelengths of mmWave frequencies, it is possible to deploy a large number of antennas in a small form factor. Because of the large number of antennas, however, it may not be cost efficient to have high-quality signal processing components for all antennas. Therefore, analog beamforming (with one RF chain) and hybrid beamforming (with a few RF chains where the number of RF chains is far less than that of antennas) have drawn a significant interest for mmWave communication systems. There has been much work demonstrating that the performance of hybrid beamforming is similar to that of full digital beamforming. We refer to [83] for details about general mmWave communication technologies.

For the joint mmWave radar and communication system, there should be multiple mmWave transceivers to overcome the blockage of mmWave signals by nearby vehicles or even pedestrians, as shown in Figure 8.1. This concept is similar to current automotive radar systems. For example, a vehicle may have mmWave transceivers on front and rear bumpers and sides for V2V and radar, and on its rooftop for V2I communications, because infrastructure will be placed in high positions to ensure good link conditions. The blockage effect in the mmWave V2V scenario might be mitigated by using the gap under vehicles as waveguard, which allows a vehicle to communicate with vehicles other than adjacent ones. Although blockage is usually considered a defect, it can be beneficial for mmWave V2V communications because the effect (combined with narrow mmWave beams) can reduce inter-vehicle interference and enable better spatial packing.



Figure 8.1: The conceptual figure of mmWave vehicular communications. Multiple mmWave transceivers are deployed on a vehicle to simultaneously establish V2V and V2I communication links.

It is well known that antenna array is strongly related to beamwidth in mmWave. With more antennas, it is possible to design beams with small beamwidths and focus transmit and receive power to certain directions. These directional beams can increase the channel coherence time, which is defined as the time duration over which the communication media between the transmitter and the receiver is considered to be not varying [88], [89]. In lower frequencies, the incoming signals arrive uniformly over all the 360-degree angular range, which holds under rich scattering environments with omni-directional reception. To compensate for the increased path loss due to the shrinking antenna size at the mmWave frequencies, transmit and receive beamforming with directional beams, which is often called *beam alignment*, is widely accepted as a necessary component in enabling mmWave communication systems [90]. With beam alignment, the incoming signals at the receiver are limited to a given range of angles. Each angle can be mapped to a Doppler frequency shift, which also means that the Doppler frequency shifts are limited to a certain frequency range with directional reception. Since the average frequency shift can be corrected using standard frequency offset correction methods, this leads to reduced Doppler spread and thus an increased coherence time. The conceptual example of this phenomenon is shown in Figure 8.2. This property has been exploited in [89] and [91] to mitigate the Doppler spread. Small beamwidths, however, will lead to frequent beam misalignment between the transmitter and the receiver due to the movement of vehicles. Therefore, there is a trade-off between the channel coherence time and the received signal strength in vehicular environments.

We analytically modeled the received signal using directional beams and analyzed the trade-off in terms of beamwidths. The results are plotted in Figure 8.3. The plot “Exact” is the result without any approximation, the plot “Approximation” is our derivation, and the plot “No angular difference” is the case without taking any mobility into account. It is clear that “Exact” and “Approximation” follow the same trend of coherence time, i.e., there is an optimal beamwidth that maximizes the coherence time. The result also shows that considering mobility is critical to accurately analyze the coherence time in terms of beamwidths. These results will guide us in designing beams, and eventually antenna arrays, in the joint mmWave radar and communication system. Full derivations and more simulation results can be found in [92].




Figure 8.2: Doppler effect in terms of beamwidths




Figure 8.3: Channel coherence time vs. beamwidth using directional beams.

8.3 Beamforming Design

Phased array antennas consist of multiple-antenna systems where multiple antenna elements can be configured in different ways to, for example, allow beam scanning over a wide area or form sharp beams. The direction of the beam can be electronically steered, thus eliminating the need for any mechanical rotation. These capabilities made phased arrays favorable in a broad range of applications like V2X communications and MIMO radar [9], [93].

Traditionally, each antenna element was allocated a complete RF chain in conventional MIMO systems. The high power consumption of the mixed signal components at the mmWave band, however, makes it difficult to allocate a complete RF chain for each antenna. To reduce the power consumption, hybrid analog/digital precoding architectures, which use a small number of RF chains and divide the precoding/combining processing between RF and baseband domains, have been proposed for MIMO mmWave systems [94–97]. The RF precoding circuit was implemented using networks of variable phase-shifters in [94–96], and switches in [97]. Despite their good performances, these techniques were primarily designed for cellular systems and generally require large number of RF chains to realize good beam patterns. While a larger number of RF chains might be justifiable for cellular systems, the high power consumption and cost of these arrays might limit their use in vehicular systems.

Here, we propose a new hybrid analog/digital architecture that requires fewer RF chains but yields results comparable to those of fully digital solutions, developing an algorithm that designs the beamforming/combining vectors for this architecture.

8.3.1 Hybrid Beamforming Architecture and Precoder Design

The proposed mmWave transceiver architecture consists of N_T antennas and N_{RF} chains, as shown in Figure 8.4. Each RF chain is connected to the N_T antennas via a network of phase-shifters and antenna switches. The antenna switches are used to modify the array geometry and improve the far-field radiation pattern. Unlike the work in [93–96] that uses only a network of phase-shifters, and the work in [97] that only uses one antenna switch per RF chain (without phase-shifters), we will show that the combination of phase-shifters and antenna switches will result in a performance comparable to fully digital solutions with fewer RF chains than required by phase-shifters-only and switches-only based hybrid architectures. To design the hybrid precoder, we first design the optimal digital precoder for a desired range of angles. In the second step, we design the hybrid analog/digital precoder (with a few RF chains) that approximates the fully digital precoder.

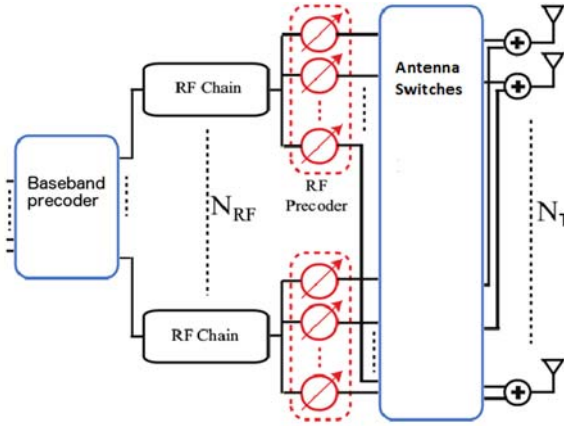


Figure 8.4: Proposed hybrid analog/digital architecture with N_T antenna and $N_{RF} \ll N_T$ RF chains.

Digital precoder design

Let the set \mathcal{R} define the range of target angles. The unconstrained digital beamforming vector (or precoder) is designed such that the beamforming gain is maximized at the desired range of angles and minimized in all other directions. This can be easily achieved using LS techniques, as shown in [94].

Hybrid analog/digital precoder design

Due to hardware constraints, the transmitter may not be able to apply the unconstrained entries of the digital precoder \mathbf{f}_d to form its beams. As discussed in [94–96], the number of RF chains $N_{RF} \ll N_T$, and the RF phase shifters are usually quantized to b bits. One possible solution is to use the limited number of RF chains together with the quantized phase-shifters to form a hybrid analog/digital design. To design the hybrid precoder, we set $\mathbf{f}_h = \mathbf{F}_{RF} \mathbf{f}_{BB}$, where the


subscript h refers to hybrid precoding, \mathbf{F}_{RF} is an $N \times N_{\text{RF}}$ analog RF precoder, and \mathbf{f}_{BB} is an $N_{\text{RF}} \times 1$ digital (baseband) precoder. Consequently, the design of the precoder is accomplished by solving Eq. (8.1),

$$\begin{aligned} \{\mathbf{F}_{\text{RF}}^*, \mathbf{f}_{\text{BB}}^*\} &= \arg \min \|\mathbf{f}_d - \mathbf{F}_{\text{RF}} \mathbf{f}_{\text{BB}}\|_F, \\ \text{s.t. } [\mathbf{F}_{\text{RF}}]_{:,i} &\in \{[\mathbf{A}_{\text{can}}]_{:, \ell} \mid 1 \leq \ell \leq N_{\text{can}}\}, i = 1, \dots, N_{\text{RF}} - 1, \\ \|\mathbf{F}_{\text{RF}} \mathbf{f}_{\text{BB}}\|_F^2 &= 1, \end{aligned} \quad (8.1)$$


where \mathbf{A}_{can} is an $N \times 2^b$ matrix that carries all set of possible analog beamforming vectors due to the angle quantization constraint on the phase shifters. Given the matrix of possible RF beamforming vectors \mathbf{A}_{can} , the above optimization problem can be reformulated as a sparse approximation problem that is solved using matching pursuit algorithms as proposed in [94] to estimate the analog and digital precoders \mathbf{F}_{RF}^* and \mathbf{f}_{BB}^* . One drawback of this solution is that the columns of the matrix \mathbf{A}_{can} are constrained by the number of the phase-shifter quantization bits. This results in spectral spillover by the grid mismatch. The grid mismatch arises because the unconstrained digital precoder \mathbf{f}_d is designed using a matrix \mathbf{A} with an infinite number of columns (or RF beamforming vectors). This grid mismatch destroys the sparsity of the above optimization problem and, as a result, requires the number of RF chains to scale with the sparsity of the system in order to obtain good beam patterns.

One way to minimize the grid mismatch is by letting the analog beamforming matrix \mathbf{A}_{can} carry multiple beamforming vectors that cover wide sector. This can be achieved by modifying the array geometry using antenna switches as shown in Figure 8.4. Based on the target beamforming gain and covering sector, the number of antennas can be selected as $N_s > \min((2\pi / \mu(\vartheta_i)), N)$, where $\mu(\vartheta_i)$ is the length of angle interval $\vartheta_i \in \mathbb{R}$. This choice of N_s ensures that the number of antennas increases for narrow beam applications, and decreases for wide beam applications. Since wider beams do not necessarily require high angular resolution, the number of antennas can be reduced to result in an RF matrix with “virtually grouped” columns. This reduces the sparsity of the optimization problem and as a result, the required number of RF chains.


In Figure 8.5, we plot several beam patterns with different beamwidths when using a fully digital architecture and a hybrid architecture. As shown, the resulting beam pattern when using the hybrid architecture with just five RF chains is similar to that achieved by the fully digital architecture. Figure 8.5(c) shows that the hybrid architecture with fixed array size does not fully cover the desired angular range, especially for wide beams; nonetheless, the pattern can be improved by increasing the number of RF chains. For narrow beams, both hybrid architectures provide similar beam patterns when compared to the fully digital architecture. The reason for this is that only a few columns of the RF matrix could be used to approximate a narrow digital beam. However, multiple columns, and hence more RF chains, are required to approximate wider digital beam patterns.



(a) Digital beamforming covering 180-, 90-, and 15-degree sectors



(b) Hybrid beamforming with antenna switches and phase-shifters covering 180-, 90-, and 15-degree sectors



(c) Hybrid beamforming with phase-shifters only covering 180-, 90-, and 15-degree sectors

Figure 8.5: An example of resulting beam patterns for a 32-element uniform linear array antenna when using: (a) a fully digital architecture, (b) a hybrid architecture with antenna switches and $N_{RF} = 5$ RF chains, and (c) a hybrid architecture without antenna switches and $N_{RF} = 5$ RF chains.

8.4 MmWave Channel Acquisition

In mmWave communications, configuring the large antenna arrays, which may be done through channel estimation or beam training, is a significant source of overhead [83], especially in high mobility V2X applications. The training overhead can be reduced by leveraging the structure in channel, e.g., sparsity [83]. Estimating the mmWave channel or the spatial correlation of mmWave systems is challenging as the channel is not directly accessible due to hardware constraints.

MmWave V2X systems will likely be deployed in conjunction with lower frequency systems, such as DSRC, to provide wide area control signal. We exposed what might be learned about the mmWave channel from sub-6 GHz channel measurements. Specifically, we investigated the possibility of using the *spatial correlation matrix* of a sub-6 GHz system as out-of-band side information about the mmWave channel. The spatial correlation matrix can be used to reduce

training overhead [98], and may help in the precoder and combiner design [99]. Translating statistical information from one frequency to another appears in work on beamforming reciprocity in frequency division duplex systems [100–104]. These studies highlight that although the propagation channels in uplink and downlink are not reciprocal, the spatial information is consistent [100]. Several strategies have been proposed to translate spatial correlation, based on least squares [100], [102], and minimum variance distortion-less response [101]. In [105], a spatio-temporal correlation translation strategy was proposed based on two-dimensional interpolation. These translation strategies work under a key assumption about the congruency (or agreement) of spatial information at the two frequencies under consideration. This is reasonable when the duplex gap is small (e.g., 190 MHz in [103]); however, it was not clear to what extent the spatial information of sub-6 GHz and mmWave agree. Furthermore, sub-6 GHz to mmWave correlation translation involves antenna arrays with different number of elements, requiring a transformation from a smaller to larger spatial correlation matrix.

We proposed to use spatial correlation of a sub-6 GHz channel as a side information for mmWave channel estimation. With the help of ray-tracing, we showed that many but not all paths at sub-6 GHz and mmWave overlap, and hence sub-6 GHz spatial correlation can be leveraged for mmWave channel estimation. Second, we proposed two translation strategies for the specific case of a narrowband single-input multiple-output (SIMO) system model with comparable apertures at sub 6-GHz and mmWave. The first translation strategy extends [55] to the case where the known and translated correlation matrices differ in size. The second translation approach is based on the parametric estimation of the mean AoA and angle spread, and subsequently using them in theoretical expressions of spatial correlation to complete the transformation.

Previously out-of-band measurements had been proposed and validated for mmWave beam steering in indoor 60 GHz WiFi [106]. The directional information was retrieved from legacy WiFi to reduce the beam steering overhead for the 60 GHz WiFi, and the results confirmed the value of out-of-band information. We, however, focused on the congruency in the spatial correlation (targeting specifically the V2I setup), whereas [106] considered primarily the line-of-sight channels.

8.4.1 Spatial Congruency in sub-6 GHz and mmWave

We used 3D ray-tracing to demonstrate spatial congruency in sub-6 GHz and mmWave channels. Ray-tracing has been reasonably successful at predicting site-specific mmWave propagation in the past [107]. We used Wireless Insite[®] software [108], to simulate the V2I communication setup shown in Figure 8.6. There are two lamp posts separated by 30 m, serving as infrastructure. Further, there are two vehicles, a car equipped with five antennas, and a van with four antennas. Multiple antennas (at different locations) on the vehicles are meant to provide diversity to prevent the dramatic effects of blockage.

In Figure 8.7, we show the fraction of paths with same AoA as a function of frequency. We considered 25 strongest paths at each frequency, per transmitter/receiver pair (i.e., from each lamp post to every antenna at car and van). We used five logarithmically spaced test frequencies from 900 MHz to 900 GHz, with 900 MHz as the base case (i.e., the commonality of angles at different frequencies is tested against angles at 900 MHz). We averaged the fraction of common paths (truncated to degree level) across both lamp posts and across all antennas on both vehicles. From the result, we see that as expected, the fraction of common paths decreases with increased frequency separation. That said, the percentage of paths with same angles at frequencies far apart

is still over 90%. This implies good agreement in spatial information, and we exploit this agreement through spatial correlation translation from sub-6 GHz to mmWave.




Figure 8.6: The simulated ray-tracing setup




Figure 8.7: The fraction of common paths as a function of frequency

8.4.2 Correlation Translation from sub-6 GHz to mmWave

Consider a sub-6 GHz communication system, and an mmWave communication system such that the two systems have co-located horizontally aligned receiver antenna arrays, and further that the physical size is same for both arrays, as shown in Figure 8.8. The comparable physical size for both arrays is well motivated—due to smaller wavelengths at mmWave, more antennas can be packed into the same space.




Figure 8.8: A depiction of collocated Sub-6 GHz and mmWave antenna arrays

Based on this configuration and the spatial similarity as demonstrated in Sec. 8.4.1, we proposed to use two correlation translation strategies (discussed next). In the following discussions, we assume that the waves impinging on the arrays have a certain mean AoA and also a certain angle spread around the mean. Note that given the antenna array configuration, the correlation of the channel is completely parameterized by the mean AoA and angle spread.

Non-Parametric Approach

In prior work [100–104], the same array was used for both frequencies; thus, the only source of discrepancy was the slight differences in the wavelength of the received signals in the uplink and the downlink. In our case, the differences come from both the wavelength and differences in the number of elements (more elements in the mmWave array). Along these lines, the first translation approach is a generalization of [105], and we call it the *non-parametric approach*. The idea is to first note the Toeplitz structure (and Hermitian symmetric nature) of the correlation matrix, as this permits us to estimate only a few unknowns. Then, based on the observed values of the correlation via sub-6 GHz spatial correlation matrix, we interpolate and/or extrapolate to get the mmWave correlation matrix. The correlation matrix is typically complex and there are two possibilities to interpolate complex numbers: (i) interpolate the real and imaginary parts and (ii) interpolate the magnitude and phase functions separately. In this work, we interpolate/extrapolate the magnitude and the phase of the correlation separately.

Magnitude: The correlation magnitude is typically smooth and spline interpolation suffices for interpolation. Good quality extrapolation, however, is more challenging. Based on our experiments, we concluded that the quadratic spline provides good extrapolation accuracy. The quadratic spline extrapolation, however, requires end point information, which we propose to estimate by linear extrapolation truncated at zero.

Phase: Note that the phase can be unambiguously determined only in the interval $(-\pi, \pi)$. Note that, however, the phase increases linearly and hence the jumps of 2π can be observed when the phase leaves this interval. For interpolation of phase, the actual linearly increasing phase needs to be reconstructed. Hence, the observed phase is first *unwrapped* and then linearly interpolated (or extrapolated) before it is re-sampled for correlation translation. Based on the phase structure of the correlation, the phase interpolation should have little error if done correctly.

After the magnitude and phase are translated, they can be combined to obtain the translated matrix. The translated matrix, however, may lack positive-semidefiniteness, i.e., a defining characteristic of correlation matrices. Therefore, given the matrix obtained by the aforementioned translation strategy, we project it onto the set of the positive definite matrices, to obtain the final translated covariance matrix.

The Parametric Approach

As a second translation strategy, we proposed a parametric approach. Note that the correlation is a function of only a few parameters. As an example, these parameters could be the mean AoA of the incoming waves, and/or the angular spread of the waves around the mean. Prior work has considered the specific problem of estimating both the AoA and the angle spread jointly from an empirically estimated spatial correlation matrix (e.g., maximum likelihood estimation

[109], covariance matching estimation [110], and spread root-MUSIC estimation [61]). We propose to use spread Root-MUSIC [111] algorithm for mean AoA and angle spread estimation due to its low complexity. We use the estimated mean AoA and angle spread (by the root MUSIC algorithm) to construct the mmWave correlation. Typically, the closed form expressions of the correlation (in terms of AoA and angle spread) result after making some assumptions about the distribution of AoA. As such, it is expected that the performance of the parametric approach will degrade if the actual and postulated distributions did not match.

The correlations obtained by the translation strategies presented can be used for minimum mean square error (MMSE)-based channel estimation and Eigen-beamforming, among other uses. We use the translated correlation for mmWave channel estimation based on linear-MMSE (LMMSE).

8.4.3 Performance Evaluation

To test the performance of the proposed translation schemes, we measure the distance between the true mmWave correlation matrix and the correlation matrices obtained by the proposed translation approaches. This distance is called the *correlation matrix distance* (CMD) and the results for CMD are shown in Figure 8.9. This metric is computed for various values of angle spread σ_θ . As expected the correlation matrix distance increases with frequency. Further, the non-parametric approach seems to perform conversion to a lesser CMD compared to parametric approach.

We perform another experiment to test the error incurred in LMMSE estimation of the channel using translated correlations. Note that even in the presence of perfect correlation knowledge, the channel estimation is not perfect, and there is some MMSE (not to be confused with the estimator LMMSE). Further, when the correlation knowledge is not perfect, there is additional error up on the usual MMSE, and this error is termed excess mean squared error (EMSE). The EMSE for the proposed translation strategies are shown in Figure 8.10. The EMSE is also evaluated for several angle spread values. As expected, the EMSE decreases with signal-to-noise ratio (SNR), and further the non-parametric approach performs better than the parametric approach.

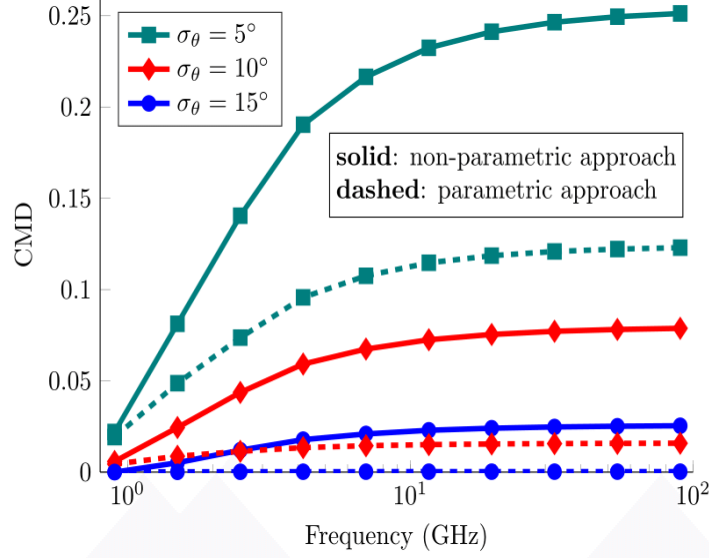


Figure 8.9: CMD as a function of frequency

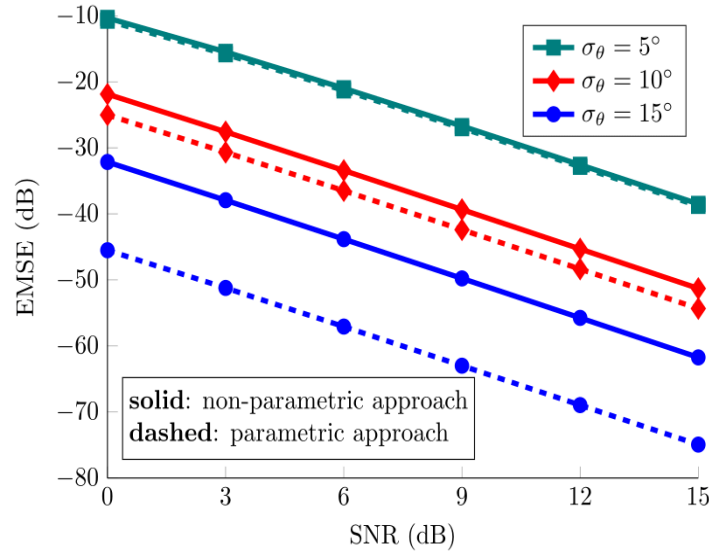


Figure 8.10: EMSE as a function of SNR

8.5 Summary

In this chapter, we identified possible antenna structures for joint mmWave radar and communication systems. Similar to automotive radars, the joint systems should have multiple mmWave transceivers to mitigate blockage. We also proved that there is an optimal mmWave beamwidth that maximizes the performance of joint systems. The result can be exploited to design appropriate mmWave antenna structures for the joint system. This work has been submitted to *IEEE Communications Magazine* [61] and *IEEE Transactions on Vehicular Technology* [92].

Because beamforming, which focuses transmit power to certain directions, is essential in mmWave systems to mitigate pathloss, we also developed a flexible beamforming technique. The proposed solution minimizes the number of RF chains (and consequently, complexity and power

consumption of transceivers) by using antenna switches. Using the proposed solution, it is possible to efficiently create different beamwidths that maximize the performance of joint mmWave radar and communication systems, as discussed in Section 8.2.

To optimize communication performance, it is well known that channel information between the transmitter and receiver is critical. The channel information can be also useful for joint radar and communication systems because it will allow adaptive signal transmission techniques. In this study, we showed that sub-6 GHz channel information can be exploited to acquire mmWave channel information. Because sub-6 GHz channel information can be obtained through conventional DSRC systems, the proposed techniques will significantly lower mmWave channel acquisition overhead and facilitate joint mmWave radar and communication systems.

Chapter 9. Dealing with Security Issues

9.1 Introduction

Many new RF technologies are being deployed to make driving safer and more automated. Automotive radar is one such technology, where RF signals are used for adaptive cruise control, forward CW, or blind spot detection. Going forward, many vehicles will be connected using DSRC for active safety applications. The abundance of bandwidth in the mmWave band could also be exploited to enable much higher data rate communication between vehicles to exchange the raw data from LIDAR, radar, and other sensors to support advanced driver-assisted and safety-related functionalities. Each technology, however, comes with its own security risks. Even isolated security breaches could have a dramatic impact on consumer confidence, resulting in the discontinuation of such technologies.

In this chapter, we present an overview and comparison of security risks associated with both automotive radar and DSRC systems. We make a suggestion about how the industry should respond to these known threats, such as through joint radar and communication. Furthermore, we describe an instance of a past successful attempt to hack a vehicle and speculate on future hacking attempts. Lastly, we suggest several techniques to improve security in both automotive radar and communication.

9.2 Security Risks of Automotive Radar

There are three principal attacks (i.e., intentional disruption of a vehicular system by a third-party) on automotive radar [112], [113].

- **Jamming** is the transmission of RF signals to interfere with a radar signal by saturating its receiver with noise.
- **Spoofing** is the replication and retransmission of radar transmit signals designed to provide false information to a radar to corrupt received data.
- **Interference** is the intentional or unintentional modification or disruption of a radar signal due to unwanted signals, such as signals from different automotive radars.

9.2.1 Jamming

Automotive mmWave radar experiences limited range due to the small wavelength and inability to pass through solid objects consistently [114]. Most radars use a substantial amount of directivity in the system to overcome this effect. This gives automotive mmWave radar more resistance to jamming compared to devices that operate at low frequencies. Additionally, since the purpose of jamming is to deny the victim service, it is moderately difficult to perform an effective jamming attack on automotive radar in a highly mobile environment. If the jammer has a static location, even a successful breach will disrupt the automotive radar for as long as the target is in range, which could be a matter of a few seconds in highly mobile environments (i.e., highways). Although the potential consequences of losing a few seconds of operation are significant (i.e., loss of collision detection for that time frame), it is incredibly difficult for a malicious attacker to predict exactly where and when the jammer needs to operate to cause an accident. As a result, the attacker is limited to jamming in environments with low mobility (i.e., downtown areas) and does

not have the ability to focus an attack on a single radar system. An illustrative example of two different types of attacks as a result of jamming is shown in Figure 9.1.

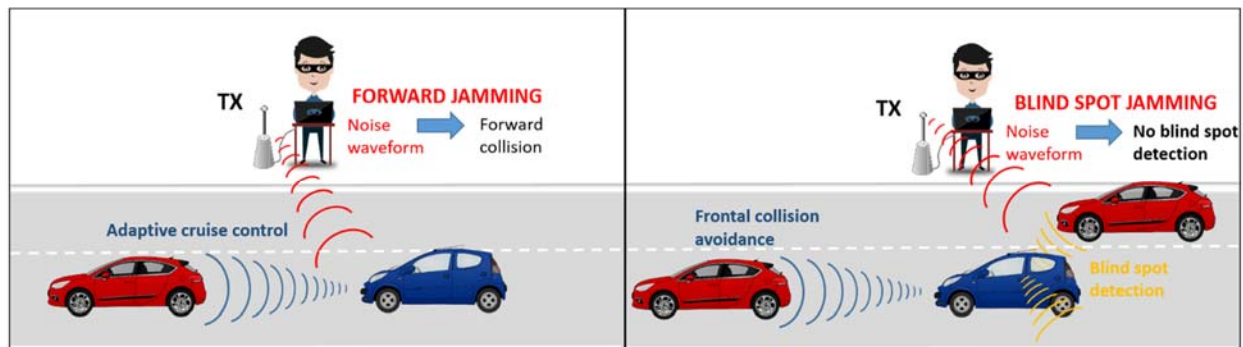


Figure 9.1: Illustration of two different types of attacks as a result of jamming.

If the jammer is mobile, much more damage can be inflicted. A jammer located on a vehicle that is currently following the target may be able to continuously jam the target. Executing a continuous jamming attack requires two major components to be successful. First, the vehicle with the jammer must stay within a certain range of the target vehicle without attracting suspicion to itself. Second, the operation requires a jammer that can accurately scan the wireless channel in a highly mobile environment, which is notably complex. To perform the attack, the jammer must be able to scan the target vehicle from any direction and distinguish the target vehicle's radar signals from any other wireless signal. It must also transmit a strong jamming signal in the direction of the target vehicle. Overall, although jamming attacks have the potential for inducing major collisions in the future, current jammers do not have the necessary adaptability for performing in a highly mobile environment, making it very difficult for malicious attackers to target a single vehicle.

9.2.2 Spoofing

Figure 9.2 provides an illustration of a spoofing attack. Distance- and velocity-falsifying attacks on commercial automotive radars have been shown to be feasible [115, 116]. Automotive radar exploits a specific signal structure that performs well as a radar signal (i.e., has strong autocorrelation properties) but exhibits no inherent authentication, leaving it vulnerable to spoofing attacks. Without a means for checking signal integrity, the receiver is unable to verify the spoofed sequences, making it possible to analyze and replicate the signal. Unlike a jamming attack, a spoofing attack is designed to confuse the target victim. Ideally, a spoofing attack only needs to breach the target radar for a short period of time to severely influence the behavior of the target vehicle, potentially causing it to stop, change direction, or in the worst case, collide. Based on this, a successful spoofing attack can have a devastating effect on automotive radars on the market today. Despite this, there has been no publicized report of a spoofing attack on a vehicle. We believe that this is due to the relatively high implementation complexity of designing an effective and robust spoofing system. Overall, spoofing is the primary security concern for automotive radar due to its potential consequences and feasibility.




Figure 9.2: Illustration of an attack as a result of spoofing.

9.2.3 Interference

FMCW automotive radars can exploit advanced signal cancellation techniques to reduce the effect of interference (including jamming) [117]. There are some forms of interference, however, such as a chirp or sweep signal, that cannot be isolated as detailed in [118], resulting in performance degradation in the presence of heavy interference. Due to the limited use of automotive radar, interference is not a problem in vehicular environments today. As automotive radars become more widespread, however, we predict that interference between automotive radars of different vehicles will become a major issue.

9.3 Security Risks of DSRC

Current implementations of vehicular communication systems are modeled after existing Wi-Fi communication systems (i.e., IEEE 802.11p, the standard used in DSRC, is a subset of the IEEE 802.11 standard). Thus, in general, DSRC technologies are susceptible to similar types of attacks used against traditional Wi-Fi, which include jamming, spoofing, and interference [119]. In addition to these attacks, DSRC technologies are also susceptible to attacks on user confidentiality.

9.3.1 Jamming

In contrast to automotive mmWave radar, DSRC devices operate at relatively low frequencies of 5.9 GHz, improving their maximum range of detection but making them more susceptible to jamming attacks. Research has shown that constant, random, and intelligent jamming attacks can deny service to DSRC applications to the point of disabling their entire functionality [120]. In addition, DSRC may potentially experience denial-of-service attacks designed to jam the system from within the vehicular network, such as malware, spamming, and black hole attacks [121]. All these attacks have the potential to disable vehicular communications for an extended period of time, putting the targeted vehicle and its occupants in danger if the vehicle relies on DSRC for CW.

To combat these potential attacks, considerable research has examined solutions such as implementing additional authentication, physically separating networks within the same vehicle, switching frequencies when denied service, and communicating with legitimate DSRC devices to blacklist rogue devices [120–122]. Despite these efforts, jamming is still a major security concern

for DSRC systems due to the ease of carrying out an attack and the potential consequences it has on targeted systems.

9.3.2 Spoofing

DSRC is a subset of the IEEE 802.11 standard and has a predefined packet sequence that incorporates packet authentication within its packet headers. Therefore, spoofing a DSRC device requires knowledge of the specific sequences used in the packet headers. In addition, the DSRC standard is capable of incorporating public key cryptography during transmission, further improving the security of these devices.

Despite these advantages, DSRC is still vulnerable to specific types of spoofing attacks. These include attacks from within the network itself and attacks that modify the signals sent throughout the network. If the attacker is able to somehow determine or obtain the necessary credentials for authentication, then it may be able to impersonate a legitimate device, enabling the attacker to send false information to the target device [123]. In contrast, spoofing attacks such as replay attacks or man-in-the-middle attacks may allow an adversary to modify signal information by intercepting a transmitted signal and retransmitting a slightly modified version of the signal. Overall, although DSRC technology is ultimately susceptible to spoofing, its inherent robustness due to predefined packet authentication mitigates the severity of this security risk. Furthermore, several supplementary measures can be implemented to provide additional security such as additional authentication.

9.3.3 Interference

DSRC has been allocated a 75 MHz frequency band at 5.9 GHz by the Federal Communications Commission. Due to this allocation, DSRC does not experience any (legal) interference from non-DSRC devices, such as Wi-Fi devices that operate at the 5 GHz band. Currently, there are relatively few DSRC devices implemented in vehicles on the road, rendering interference as a non-issue. In the future, however, when DSRC devices become widespread, interference between mutual devices will be a concern, especially in congested environments such as downtown areas. Although current strategies for reducing mutual interference (such as interference cancellation, power and frequency adaptation, and improved MAC [Medium Access Control] layer protocol design) can decrease the effect of interference on DSRC, mutual interference is still a notable security concern that has yet to be completely addressed [124, 125].

9.3.4 Confidentiality

DSRC devices must maintain information privacy and ensure that unwanted third parties cannot covertly track the location of the device over an extended period of time. Potential threats to confidentiality include eavesdropping, masquerading, and traffic analysis [126]. Maintaining confidentiality is one of the more discussed security topics in vehicular networks. This is due to the exceptionally low complexity of conducting an attack on confidentiality. For naive DSRC technologies, such attacks can be performed by listening to the data transmissions within a network and analyzing the traffic. Furthermore, even if the data itself is encrypted, modern traffic analysis techniques can examine traffic patterns of a specific device and extract location information from the analysis. As a result, DSRC technologies need to be designed intelligently in order to prevent attacks on confidentiality. Currently, there are various measures for preventing attacks on

confidentiality such as device cloaking; however, these solutions introduce considerable complexity to the entire network and are sometimes undesirable.

9.4 Comparing Automotive Radar and DSRC Security

In summary, both automotive radar and DSRC technologies have inherent security flaws, as summarized in Table 9.1. On the one hand, DSRC devices are more susceptible to jamming than automotive radars since they are subject to jamming attacks from within the vehicular network. On the other hand, automotive radars are considerably more susceptible to a spoofing attack than DSRC technologies due to their lack of signal verification. Currently, both automotive radar and DSRC devices are not significantly impacted by interference. In the future, when the technologies become more widespread, interference will become an important security concern that needs to be addressed. In addition, DSRC technology must account for attacks on confidentiality due to its nature as a communications system. Overall, although there are more types of attacks on DSRC systems, DSRC is more secure than automotive radar due to its built-in security mechanisms and its ability to communicate with other legitimate DSRC sources. This does not mean that DSRC equivalents can replace automotive radar, since the functionalities of both technologies are crucial for a variety of vehicular applications.

Table 9.1: Security comparison between automotive radar and DSRC

	Automotive Radar	DSRC
Jamming	Moderate	High
Spoofing	High	Moderate
Interference	Low	Low
Confidentiality	None	Moderate

Note: ‘High’ indicates a security risk with a high potential for major consequences, ‘moderate’ indicates a security risk with a moderate potential for major consequences or a high potential for minor consequences, ‘low’ indicates a security risk with a small potential for both major and minor consequences, and ‘none’ indicates no security risk.

Our research on developing a joint radar and communication framework, as described in the sections above, will address the security concerns detailed above. This framework can improve target localization, increase system reliability, and supplement automotive radar with an extra layer of authentication, dramatically reducing the device’s vulnerability towards a spoofing attack.

9.5 Hacking a Vehicle

In 2014, security researchers published a paper describing a strategy for a remote automotive attack at an international hacker convention [127]. A year later, they took a step further and demonstrated a wireless attack on a Chrysler Jeep being driven on a public highway, posting the footage in a YouTube video [128]. By exploiting a major oversight in Chrysler’s network design, they were able to brute force their way into the system and exploit the Linux operating system. From there, they were able to remotely control steering at low speeds, engine status, the air conditioning system, and radio via the Internet.

Upon the release of the video, the public reacted quite negatively towards this demonstration. Their angry complaints prompted several changes in the automotive industry, one of them being the release of a best practices paper by Intel (McAfee) [129]. This paper outlines all the known ways vehicles can be hacked and the most effective countermeasures, including but not

limited to attacks from wireless V2V and V2I receivers, Bluetooth systems, and the engine control unit.

Despite the paranoia caused by the video, wireless malicious hacking of a vehicle has been virtually nonexistent. Though the idea has been popularized in movies (such as Disney's *Tron*) or video games (such as Ubisoft's *Watch Dogs*), there has only been one documented instance of malicious hacking of a car. In 2010, an angry former employee bricked (rendered inoperable) hundreds of cars at a dealership [130], [131], destroying several million dollars' worth of cars, but injuring no one in the process. Additionally, although [127] provided a substantial list of vehicle models susceptible to the same type of attack they performed, there have been no reported attacks on any of these vehicles.

The public's concern is understandable, given the potentially tragic consequences of allowing vulnerable vehicles to drive on public roads. Although the threat of hacking vehicles is real, with the proper precautions, these threats can be avoided altogether. Like any other networking protocol, vehicular networks will always be subject to attacks. But as long as security concerns are addressed in an ethical, appropriate, and timely manner, there is no reason to prevent or delay the integration of communication networks in vehicles.

9.6 Incorporating Antenna Arrays to Enhance Security in mmWave Communication Links

The security techniques discussed so far for vehicular communications are implemented in higher layer protocols, and are based on digital signature methods that require vehicles to store a large number of public/private key pairs. Additionally, these keys must be regularly exchanged, hence creating, in addition to the processing overhead, an additional communication overhead [8], [132]. Here, we exploit the large dimensional antenna arrays available at mmWave systems to produce direction dependent transmission. This results in coherent transmission to legitimate receivers and artificial noise that jams potential eavesdroppers with sensitive receivers. This technique also provides a low complexity solution that makes it difficult for an eavesdropper to intercept signals from other vehicles or to spoof a vehicle.

9.6.1 Proposed Security Technique

In this section, we propose a security technique for vehicular mmWave communications that does not require the exchange of keys between vehicles. We assume a single transmitter communicating with a single receiver in the presence of one or more potential eavesdroppers as shown in Figure 9.3. The proposed technique uses an antenna architecture with a single RF chain (instead of multiple RF chains), and performs analog beamforming with antenna selection. To transmit an information symbol, a random set (called a *subset*) of M antennas are co-phased to transmit the information symbol to the receiver, while the remaining antennas are co-phased such that a noise like signal is generated in all other directions. The indices of these antennas are randomized after every symbol transmission. Although the target receiver would observe gain reduction, malicious eavesdroppers will observe non-resolvable interference. The variance (power) of this interference term is a function of the subset size M used for data transmission.




Figure 9.3: Vehicle to vehicle communication with a possible eavesdropper. The transmitting vehicle is communicating with a target vehicle while the eavesdropper tries to intercept the transmitted data.

In Figure 9.4 we plot the variance of the interference power of an eavesdropper located at an angle of 60 degrees, while the legitimate receiver is located at 100 degrees relative to the transmitting vehicle. From the figure we observe that variance decreases with increasing transmission subset size M . This is mainly due to the fact that as M increases, the random number of antennas used to generate the interference decreases. Nonetheless, this decrease in the variance results in increased beamforming gain towards the receiver, and hence, there is a trade-off between the interference power at the eavesdropper and the beamforming gain at the target receiver.




Figure 9.4: Variance of the interference at an eavesdropper located at a transmit angle of 60 degrees relative to the transmitting vehicle when using different transmission subset sizes M and number of transmit antennas N_T .

In Figure 9.5, we plot the secrecy throughput achieved by the proposed technique versus the eavesdropper’s angular location. We also plot the secrecy throughput achieved by conventional array transmission methods (when using all antennas). The secrecy throughput is defined as the amount of information that can be sent reliably to a target receiver. From the figure we observe that the secrecy throughput when using the proposed technique is high at all angular locations

except at the target receiver’s angular location (100 degrees). We also observe that conventional array transmission techniques provide poor secrecy throughputs. The reason for this is that conventional array transmission techniques result in a constant radiation pattern at an eavesdropper while the proposed technique randomizes the radiation pattern at an eavesdropper, thereby creating unresolvable interference. For the proposed technique, no randomness is experienced at the target receiver, and the secrecy throughput is minimal when the eavesdropper is located in the same angle with the target receiver, which is not possible in practice.

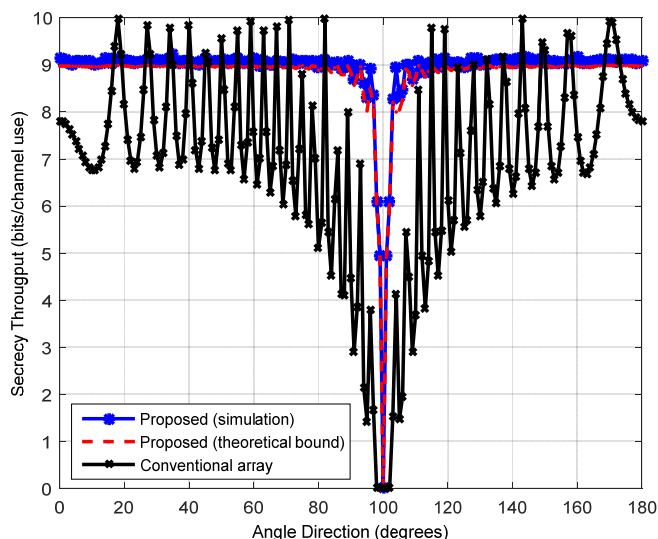


Figure 9.5: Secrecy throughput versus the eavesdropper’s angular direction.

Number of antennas at transmitter is $N_T = 32$, number of antennas at receiver is $N_R = 16$, number of antennas at eavesdropper is $N_E = 64$, and the transmission subset size is $M = 24$. The system operates at 60 GHz with a bandwidth of 50 MHz and an average transmit power of 37 dBm. The distance from the transmitter to all receivers (including the eavesdropper) is set to 30 m, and the path loss exponent is fixed to 2.

Finally, in Figure 9.6, we examine the impact of the transmission subset size M on the secrecy throughput. We observe that as the subset size M increases, the secrecy throughput increases, plateaus, and then decreases. The reason for this is that as M increases, more antennas are co-phased for data transmission. On one hand, this increases the beamforming gain at the target receiver. On the other hand, increasing M decreases the variance of the interference at the eavesdropper. Therefore, there is a trade-off between the beamforming gain at the receiver and the interference power at a potential eavesdropper and there exists an optimum value of M that maximizes the secrecy throughput.




Figure 9.6: Secrecy throughput versus the transmission subset size M ; receiver is at a transmission angle of 100 degrees and the eavesdropper is at a transmission angle of 95 degrees relative to the transmitter.

9.7 Summary

The studies in this chapter showed that security is important for automotive radar and communication systems. Even a small security breach for automotive radar or communication systems could have a dramatic consequence for transportation safety. To improve the security of joint mmWave radar and communication systems, we implemented a beamforming technique that randomizes transmit signals to all directions except to the legitimate receiver. The randomized signals are seen as noise-like signals to possible eavesdroppers, which can significantly improve communication security. This work has been submitted to *Microwave Journal* [134] and the IEEE Vehicular Technology Conference [135].

Chapter 10. Preliminary Field Test of Automotive Radar and Communication Systems


10.1 Introduction

This chapter describes the data collection and the preliminary analysis conducted to evaluate the complementarity of automotive radar and communications. The data was collected with the objective of providing real input to test the architectures and algorithms developed in Task 1, but also with the objective of visualizing the advantages of having a system that joins the capabilities of the radar and the communication devices instead of using single technologies.¹ As discussed in previous tasks, radar is used to estimate the location of objects and motion trajectory of other vehicles relative to the sensing vehicle, while V2V communication relies on on-vehicle sensors such as GPS receivers to determine each vehicle's absolute location and motion. Figure 10.1 presents the capabilities and limitations of these technologies as well as other technologies mentioned in previous technical memoranda, such as cameras and LIDAR. The sensing technologies have the advantage of accuracy but have limited field of view, sensitivity to weather conditions such as heavy rain, fog, and snow, and can simultaneously track only a limited number of objects. In contrast, V2V devices rely on GPS positioning, which usually is less accurate than the sensors, but are not affected by weather and have a 360-degree field of view.

Traffic scenarios and situations have great variability and the isolated technologies may not be enough to ensure a safe CW/CA system. For example, in the context of urban intersections, radars' lines-of-sight may be obstructed by prominent slopes, buildings, walls, and other large objects, so other sources of information would be required to provide the location of oncoming vehicles. Additionally, sudden stops or accelerations between two vehicles based purely on radar systems can have a downstream (in time), cascading negative safety impact on other vehicle movements. This is where communications, in addition to radar, can have a substantial benefit. Radar may also be a good complement to communications, as roadway users such as pedestrians and bicyclists would not be outfitted with vehicular communication devices, and thus radars (and cameras) become essential sensing gadgets.

In the following sections of this chapter, we analyze different scenarios that were considered for the data collection as well as the equipment and software that were designed for this activity. First, we describe the data collection effort and its limitations. Then, we conduct a simplified analysis of the accuracy of the sensing and communication technologies and observe the advantages of using both types of technologies simultaneously. Finally, we perform simplified quantitative and qualitative analyses of the sensor fusion and trajectory prediction techniques proposed in Chapter 2 using the collected data.

¹ Note that in Task 2 the research team developed a framework for joint mmWave radar and communication systems that functions within a single device by using the same waveforms. However, in this task, we explore the use of communication and radar technology through the use of devices that are already available in the market (DSRC and automotive radar units). Therefore, the objective is to analyze the complementarity of these types of technology and not the joint technology per se.



Sensor Type	Radar 24 GHz	Radar 77GHz	Lidar	Mono Camera	Stereo Camera	Radar + Camera	V2V
Field of view	56°	18°	27°	36°	48°	18°/36°	360°
Typical range	60 m	200 m	10 m	(50 m)	(150 m)	200 m / 50 m	300 m
Accuracy	0.2 m	0.2 m	0.2 m	Not defined	Not defined	Not defined	< 1.5 m
Relative reliability in snow, fog, heavy rain	Fair	Fair	Fair	Fair	Fair	Fair	Excellent
Reliability in direct sun and shadows	Excellent	Excellent	Excellent	Fair	Poor	Excellent	Excellent
Reliability in "urban canyons"	Excellent	Excellent	Excellent	Excellent	Excellent	Excellent	Fair
Reliability in tunnels and under heavy foliage	Excellent	Excellent	Excellent	Fair	Poor	Excellent	Fair
Vulnerability to damage or misalignment	Yes	Yes	Yes	No	No	yes	No
Generally considered sufficient to react to fixed objects (by OEMs)	no	No	No	no	yes	yes	Yes
Number of objects (vehicles) that can be tracked/processed at any given time	17	17 to 64	17	Not defined	Not defined	17	TBD >200
Capable of close range, low speed range-rate estimates (city safe capability)	No	No	Yes	No	No	No	for warning applications only
Requires multiple vehicles to be equipped	No	No	No	No	No	No	Yes
Supports pedestrian detection	need multi-sensor system	need multi-sensor system	need multi-sensor system	need multi-sensor system	yes	yes	TBD
Sufficient to support activation of active safety systems	No	No	No	No	yes	yes	TBD

Figure 10.1: Collision avoidance sensor summary. (Source: [159]).

10.2 Data Collection Scenarios

To accomplish the two objectives of the data collection—providing data for testing the developed algorithms and visualizing the advantages of having a system that joins the capabilities of the radar and the communication devices instead of using single technologies—the researchers used a three-stage process to define the scenarios and routes. In the first stage, the research team evaluated the advantages and disadvantages of radar and DSRC and developed a comprehensive list of situations that could be used to test the complementarity of the technologies (Table 10.1). In the second stage, the group selected variables related to transportation infrastructure and traffic that should be addressed by a robust CW/CA system. In the final stage, the group identified streets within the University of Texas campus area that would present characteristics related to the selected variables. The list of variables and locations is presented below:

- a. Traffic Volume and Road User Mix
 - i. Vehicle Types: Cars, Heavy Trucks, Buses, Motorcycles
 1. High volume of cars/buses/motorcycles found on the major streets and high volume of heavy trucks found near construction sites:
 - a. San Jacinto Boulevard
 - b. Red River Street
 - c. Guadalupe Street
 - d. Dean Keeton Street
 - e. W. MLK Boulevard
 2. Low volumes found on minor streets within campus:
 - a. Speedway between W. Dean Keeton Street and 21st Street
 - b. 24th Street
 - c. 21st Street
 - ii. Pedestrians and Bicyclists
 1. Low volumes:
 - a. San Antonio Street
 - b. Red River Street
 - c. Trinity Street
 - d. Manor Road
 2. High volumes:
 - a. Guadalupe Street
 - b. Dean Keeton Street
 - c. All minor streets on within campus
- b. Road Characteristics
 - i. Intersections
 1. Four-way intersections:
 - a. W. Dean Keeton Street & University Avenue
 - b. W. Dean Keeton Street & Speedway
 - c. 24th Street & Speedway
 - d. 24th Street & Guadalupe Street
 - e. 21st Street & Speedway
 - f. Robert Dedman Drive & E. Dean Keeton Street
 2. T-intersections:
 - a. W. Dean Keeton Street & Guadalupe Street

- iii. 45–60 mph (not available in the area of study)
- iv. > 60 mph (not available in the area of study)

Table 10.1: List of traffic situations relevant for assessing the developed technologies

Situation	Relevance
Basic car following—both vehicles equipped with radar and DSRC	Determine which data source is more accurate for position, speed, acceleration, etc., in general circumstances.
Lead vehicle stopped	Determine radar capability of detecting stopped vehicles if the vehicle is stopped before coming into range of the radar.
Blind spots and vehicle out of line of sight	Identify situations in which the radar is not capable of identifying an approaching vehicle or pedestrian/bicyclist due to angle of approximation or distance.
	Test to what extent DSRC can compensate radar blind spots.
Vehicles out of range	Test in which situations DSRC messages stop being exchanged.
Inclement weather	Identify reductions in sensing accuracy of radars in situations of poor visibility.
	Test to what extent DSRC can compensate radar limitations in such situations.
Close range	Identify if radars detect close range objects accurately.
Low speeds	Compare the accuracy of radar and GPS (DSRC) positioning when the vehicles are moving at low speeds.
Left turns and blind intersections	Demonstrate how vehicles entering an intersection to make a permissible left turn may not be able to use radar (limited field of view) to detect oncoming vehicles.
	Evaluate the use of DSRC communication to identify oncoming vehicles.
Interaction with non-connected vehicles	Visualize the information that will be captured by the system when interaction with both connected and non-connected vehicles simultaneously.
	Evaluate the use of communication when a vehicle equipped with DSRC stops many vehicles ahead of another vehicle equipped with DSRC.
Tracking large number of vehicles	Radars can track only 17 objects simultaneously, so this situation will evaluate what is being missed by the radar in situations with more than 17 moving objects.
	DSRC can track over 200 vehicles simultaneously—but for that to happen, all vehicles must be equipped with compatible devices.
Urban canyons, tunnels, and heavy foliage areas	These locations obstruct GPS signal, which affects the accuracy of the location data exchanged through DSRC.
Irrelevant objects	Identify situations that provide erroneous detections by radars that could cause false positive warnings.

For the preliminary data collection, two main scenarios were designed based on the analysis in Table 10.1. Both were planned to be executed with two vehicles equipped with DSRC and at least one vehicle equipped with radar and camera (ideally both vehicles should have all equipment). The first scenario (Figure 10.2) consisted of one vehicle following the other around campus and completing the same loop five times. In the first two laps, the tail vehicle stayed behind the leading vehicle in the same lane without intervening vehicles. In the following laps, the tail vehicle would change lanes, allow different spacing from the leading vehicle, and allow other vehicles to drive between them. This scenario was planned to pass through streets with curves and changes in slope, different types of intersections as well as areas with high density of pedestrians and construction areas subjected to traffic of heavy vehicles. In the second scenario (Figure 10.3), the vehicles met in a blind intersection. Each vehicle approached from one side at least two times. The chosen intersection (27th and Speedway) was the only blind intersection found in the campus area and the vehicle approaching from 27th was on a descending slope.



Figure 10.2: First data collection scenario—basic car following.



Figure 10.3: Second data collection scenario—blind intersection.

10.2.1 Equipment

The equipment used for data collection is divided into three categories—hardware, software, and vehicles—as detailed below.

Hardware

The combined DSRC/radar system contained a Cohda Wireless DSRC On-Board Unit (MK5-OBU), a Delphi Electronically Scanning Radar (see Figure 10.4 for specifications), and a simulated CAN (Controller Area Network) bus made from a Kvaser CAN-to-USB interface adapter, an OBDII connector, and an Ethernet cord. A CAN bus is a specialized internal communications network used to interconnect electronic components and to allow microcontrollers and devices to communicate with each other in applications without a host computer inside a vehicle.

As explained earlier, the function of the DSRC unit is to allow the system to transmit and receive DSRC messages to/from other DSRC-enabled devices (main communication device), while the Delphi radar module is used to sense the surroundings of the vehicle and perceive objects with sub-meter accuracy. The DSRC is accompanied by a GPS antenna that is placed in the vehicle roof. The simulated CAN bus is used to send radar information from the Delphi Radar module across a CAN bus without having to connect the system to the vehicle. The Kvaser CAN-to-USB interface adapter allowed regular laptops to read information from the CAN bus (rate of 8,000 messages per second) and the OBDII connector was used to connect the combined DSRC/radar system to the vehicle. Finally, the Ethernet cable allowed us to create a local network between computers and the Cohda Wireless DSRC unit so that information from the DSRC could be read on the laptops.




Figure 10.4: Delphi Electronically Scanning Radar (2.5) specifications. This equipment can identify up to 64 objects simultaneously (Source: Delphi, 2009).

Software: DSRC/Radar Parsers

The software DSRC/radar parsers took in information coming from the data streams of the DSRC/radar modules and converted it into a Python data structure that could be processed by the algorithms developed. Specifically, the radar parser took in information from the simulated CAN

bus and formed a coherent data structure by using the Kvaser CAN-to-USB reading device. The DSRC parser took information directly from the Cohda module and sent it to the combined DSRC/radar system using a dedicated UDP (User Datagram Protocol) socket. The collection of data coming from both the DSRC and radar modules and the conversion of CAN/DSRC messages into a Python data structure were the first steps for processing the data that was being collected.

The combined DSRC/radar system used a robust logging framework that outputs combined DSRC/radar data into text files. Each process logs its output to its own text file. After the data was collected by the loggers, it was sent through a normalizer to help the combined DSRC/radar system filter out data from the radar or DSRC modules that had been deemed unnecessary. Having each process write logs to a different file avoided the concurrency issues of multiple processes writing to the same file, which could result in information loss. In addition, keeping the raw version of the data and replaying it in real time allows us to use the same data to improve the software that we are developing (to combine DSRC and radar data).

Vehicles

Two vehicles were used for the data collection. The first vehicle was a rental sedan, a 2016 Chevrolet Malibu (Figure 10.5). The vehicle was equipped with DSRC and radar; however, the battery of the vehicle did not provide enough energy for the radar to function properly. The energy provided was enough for the radar to send information to the computer, but not to sense objects. This issue may have been a consequence of the vehicle being a type of hybrid vehicle. The research team did not investigate the issue further, because it was out of the scope of the research, but this experience indicated that the implementation of multiple radars/sensors in hybrid or even electric vehicles may still require further development in both the vehicles' and the radar's technologies.



Figure 10.5: First vehicle used in the data collection.

The second vehicle was an SUV, a 2013 Hyundai Tucson (Figure 10.6). This vehicle had a radar module, a DSRC, and a GoPro camera installed and functioning. The camera was installed in alignment with the radar, such that the horizontal relative position of the identified objects would coincide. Since only the second vehicle had an operating radar, during the basic car following scenario we always kept the first vehicle as the leading vehicle, and the vehicle with the radar as the following (tail) vehicle so that we could compare the GPS position information of the lead vehicle sent by DSRC to the position sensed by the radar.



Figure 10.6: Second vehicle used in the data collection.

10.3 Preliminary Data Analysis

In the basic car following scenarios, the two vehicles completed five laps around the defined loop, yielding 53 minutes of uninterrupted data collection and approximately 15,000 messages exchanged per vehicle through the DSRC equipment. The blind intersection scenario was performed four times (each vehicle approached 27th Street twice) and resulted in 25 minutes of data (including the time necessary for the vehicles to maneuver back to their initial positions) and approximately 7,500 messages per vehicle. In the following sections, some quantitative analyses are made, however in most cases only qualitative analyses could be performed due to the limited amount of data collected and the absence of the ground-truth position of the vehicles.

10.3.1 Equipment performance and accuracy in uncontrolled environments

Table 10.2 shows the DSRC message latency for both scenarios. The results suggest an overall smaller latency for the tail vehicle. Specific testing with each of the DSRC units will be necessary

to determine whether the difference in latency was specific to the equipment unit. If both units present the same latency in a controlled environment, then we will discuss in future research whether the DSRC functioning is affected by the vehicle type.

Figures 10.7, 10.8, and 10.9 are plots of the ratio of messages that were successfully received by the communication equipment over the distance between the vehicles. There is a significant variation in the success rate of the lead and tail vehicle in the basic car following scenario, but high rates were maintained up to 300 meters. For the blind intersection scenario there is a drastic reduction in the success rate when vehicles are more than 100 meters apart, probably due to the many obstacles (e.g., buildings, hills, and trees) in the environment. Again, further data collection, including more controlled environment scenarios, will be necessary to investigate the causes of these results.

Table 10.2: DSRC message latency

Message Latency (seconds)			
Basic car following scenario	Mean	99th percentile	Maximum
Lead vehicle	0.15	0.35	0.5
Tail vehicle	0.02	0.03	0.6
Intersection	Mean	99th percentile	Maximum
Both vehicles	0.08	0.16	0.39




Figure 10.7: DSRC reception success rate for the leading vehicle.




Figure 10.8: DSRC reception success rate for the following (tail) vehicle.




Figure 10.9: DSRC reception success rate for both vehicles in the blind intersection scenario.

The above plots do not allow us to distinguish between occasional packet loss and large gaps where communication has failed. Therefore, we identify “skips,” which are the periods when communication was not successful. They are classified as long skips, which are greater or equal to one second, and short skips, which are less than a second. The skips for both scenarios are presented in Table 10.3. For the basic car following scenario the period of time in which communication is missed is very short and inferior to 7 percent of the total time. In this scenario, the overall total duration of short and long skips are similar. On the other hand, for the intersection scenario, one of the vehicles did not receive messages during approximately 43 percent of the time tested. As mentioned before, in this scenario there were many obstacles between the vehicles and at some points, while returning to the scenario start point, vehicles stayed more than 400 meters

apart, which may have contributed to the long skips. Again, more tests are necessary in order to identify the true limitations of the DSRC equipment in real world environment.

Table 10.3: Message skips (periods when communication was unsuccessful)

Skips									
Scenario	Vehicle	Number of long skips		Total long skips time(s)		Number of short skips		Total short skips time(s)	
Basic car following scenario (53 minutes)	Lead vehicle	16	0.30/min	122	3.84%	154	2.90/min	66.0	2.07%
	Tail vehicle	17	0.32/min	132	4.15%	1388	26.19/min	60.1	1.89%
Intersection (25 minutes)	First vehicle	16	0.64/min	634	42.27%	30	1.20/min	13.8	0.92%
	Second vehicle	11	0.44/min	369	24.60%	38	1.52/min	18.6	1.24%

Figure 10.10 shows scatter plots of short and long skips for the intersection case. The blue points show how far apart the cars were before the skip began, and the red points show their distance once communication was restored. Short skips occurred even when the vehicles were less than 25 meters apart; however, they were more frequent at distances between 100 and 150 meters. The majority of long skips was just above one second and they occurred only when the vehicles were more than 100 meters apart (Figure 10.11).




Figure 10.10: Distances in which short skips occurred.




Figure 10.11: Distances in which long skips occurred.

Besides analyzing the performance of the message exchange, we also checked the GPS accuracy. For this analysis, we picked a 45-second interval in which both cars are waiting in an intersection and remain still for the most of the time. The RMS error of the leading vehicle's GPS, during this period was 0.22 meters and for the tail car's GPS was 0.28 meters. The maximum errors were 0.31 meters for the lead car and 0.76 meters for the tail car. Any error in the sub-meter range should be considered acceptable. However, it is not guaranteed that GPS accuracy will not degrade when vehicles are in motion. Merging the video data with the DSRC and radar data, we were able to capture moments in which the GPS precision was adequate and moments in which it was very poor (indicating that the sensed vehicle was in a different lane than its actual lane). Figure 10.12 and Figure 10.13 depict both situations respectively. The green squares represent the objects identified by the radar, while the blue square represents the GPS position of the leading vehicle. In Figure 10.13, the GPS informs that the vehicle is in the neighboring lane from its actual lane, an error of approximately 3.5 meters.



Figure 10.12: Accurate GPS data (blue square) sent through DSRC coincides with radar position detection (green square).




Figure 10.13: Inaccurate GPS data (blue square) sent through DSRC shows vehicle in wrong lane.

Another aspect that can be evaluated regarding the GPS data is the variation in the logged relative distance and angle of the vehicles during the 45-second period that both cars are waiting in an intersection. In Figure 10.14, we can see that the variation in the angle measurement is about 3 percent (we can assume that the vehicles did not change their true relative angle in this time period), while the distance increases then goes back to roughly its old value (this occurs because both cars move less than a meter after 10 seconds, when the front car starts the movement and the tail car follows it).




Figure 10.14: Distances in which long skips occurred.

While the positioning does show some error induced by motion (otherwise the upper plot would be much more block-like), it does capture the major movements at the right times. The lead vehicle moves forward at 12 seconds, and the tail car begins to move shortly after. At around 22 seconds, the lead car shifts forward again very slightly. The extraneous motion given by the GPS accuracy is on the order of 0.2 meters or less. The relative direction of the vehicles gets a slow shift in measurement, on the order of 0.1 radians (5 degrees), which seems to be acceptable. Again, it is not possible to determine if this angle measurement error increases when the vehicles are in motion.

We also used the same 45-second interval of data to evaluate the accuracy on the radar readings (Figure 10.15). We can observe that the radar data presents more noise than the DSRC data. Instead of a continuous type of error, the radar presents occasional peaks of mismeasurements. In this interval, 17 percent of measurements were “missing,” meaning that the vehicle stopped in front of the radar was not detected and the longest period of time without missing measurements was 0.85 seconds. Filtering out the noise peaks, we can consider that the radar error remained close to 1 and 2 meters during most of the time.

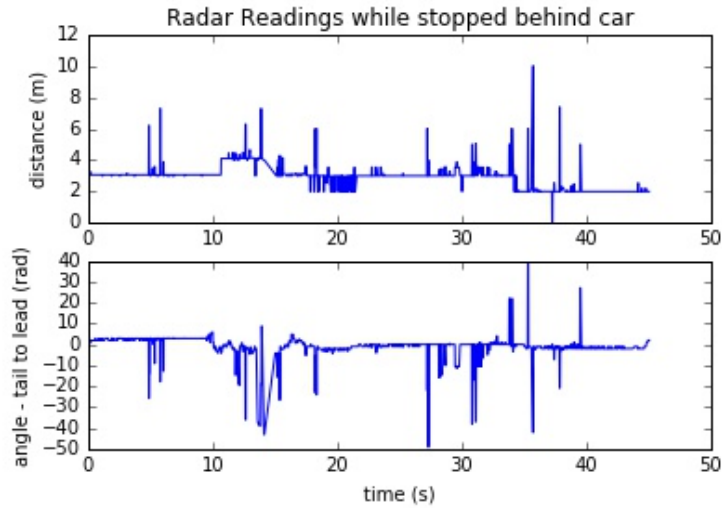


Figure 10.15: Accuracy of radar readings.

10.3.2 Qualitative Analysis of the joint detection system

Using the fully merged data that includes the video, DSRC, and radar data, it is possible to qualitatively analyze the performance of the joint system in identifying moving objects and approaching vehicles in the blind intersection scenario. Figures 10.16, 10.17, and 10.18 show the field of view of the vehicle approaching Speedway from 27th Street. In the first figure of the sequence, we observe some green squares that represent objects identified by the radar. In this case, non-moving objects are identified by the radar. This street approach is a descending slope. The ground in the center of the intersection is in the line-of-sight of the radar and generates erroneous detections (radar false positives). The blue square represents the position of the oncoming vehicle equipped with DSRC. In this situation the DSRC proves to be an important source of information to identify vehicles obscured by buildings or objects. Figure 10.17, shows a vehicle not equipped with DSRC being identified by the radar (green squares) while the vehicle with DSRC is still out of sight but is recognized through DSRC. This is a clear example of the complementarity of both technologies in situations where there is a mix of vehicles using and not using the communications technology. Finally, in Figure 10.18, the approaching vehicle appears in the visual field of the driver; however, due to the angle and slope, the vehicle is still not sensed by the radar but continues to be identified through DSRC (the square is vertically misaligned with the vehicle only because a simplified algorithm was used to merge the three sources of data).

The last situation we looked at in this preliminary data analysis was the recognition of pedestrians by the radar. Figure 10.19 shows that the radar identifies a limited number of pedestrians (as mentioned earlier, this type of automotive radar identifies up to 64 targets but in our data collection it hardly ever identified more than 20 objects simultaneously) and the ones that are more than 10 meters ahead of the vehicle seem to be prioritized (long-range seems to be prioritized over mid-range; see Figure 10.4 for radar specifications), while the ones that are closer are missed. This problem was identified multiple times along the video, proving that using cameras and developing algorithms to assist pedestrian- and bicyclist-sensing is indeed necessary, as shown in Chapter 3. A video containing the images of the situations described in this section is available at <https://www.youtube.com/watch?v=BnujYSEJ26I&feature=youtu.be>.



Figure 10.16: Blind intersection–DSRC detects approaching vehicle equipped with DSRC (blue square) and not visible to radar system. Radar false positives in green due to light reflex.




Figure 10.17: Blind intersection–DSRC detects approaching vehicle equipped with DSRC (blue square) and not visible to radar system. Radar detects vehicle without DSRC equipment (green squares).



Figure 10.18: Blind intersection – DSRC detects approaching vehicle equipped with DSRC. Approaching vehicle is already in the field of vision of the driver but radar does not identify it.



Figure 10.19: Pedestrian detection by radar–fails to detect pedestrians that are closer to the vehicle.

10.3.3 Quantitative analysis of sensor fusion and trajectory prediction

The previous sections examined the consistency of both V2V information and radar, but not their accuracy in terms of the distance between each vehicle’s estimated position and its true position. The accuracy of each sensor, the synthesized position estimate from both sensors, and the predicted future position of vehicles are quantified in this section. Unfortunately, it is not feasible to calculate the vehicles’ true positions for the entire experiment. The vehicles moved through many locations and there are no references (planted placemarks, etc.) with which to automate this calculation. Thus the accuracy is shown for a few locations that were visited several times by the vehicles.

The test vehicle both receiving DSRC communication and using radar is denoted the ego vehicle, while the other is denoted the alternate vehicle. In the first situation for which accuracy is calculated, the ego vehicle is stopped behind the alternate vehicle. In the second situation the vehicles pass each other on different sides of the road. At this point the alternate vehicle is out of the range of the ego vehicle’s radar. In the third situation the vehicles have both stopped at the entry of an intersection, with the ego vehicle turning to the right and the alternate vehicle moving straight from the ego vehicle’s right. The ego vehicle’s camera view for each situation is shown in Figure 10.20.



Figure 10.20: Examples of the vehicles in situation 1 (above) and situation 3 (below). In situation 2, the alternate vehicle is out of view of the camera.

For the data fusion, we adopt the premise that the vehicle equipped with radar and V2V communication to detect vehicles should combine both sensors’ information whenever possible, and rely on one when the other is not available. One difficulty that was not anticipated in Task 1’s study of DSRC/radar fusion—but became evident when analyzing the collected data—was the radar’s propensity to detect extraneous objects (frequent radar false positives). To solve this issue, these objects are removed in two steps during the data processing: objects that are detected for less than 200 milliseconds are removed; then detected objects that are more than 6 meters (distance that is greater than the maximum GPS error) from the position identified by the GPS are also removed².

The accuracy of the sensor fusion in each situation is shown in Table 10.4. The error in position estimation is expressed as the error in the distance between the two vehicles, in meters, and the error in the angle where the alternate vehicle is with regard to the ego vehicle, in radians. In the stopped situation, the fusion of radar and DSRC is seen to provide significant improvements over prediction using either individual information source. This is because radar often makes false detections or misses the object for a short time, and the GPS information sent through V2V is more

² Note that the algorithm being tested is for detection and trajectory prediction of a single object, therefore we are only interested in the radar’s detection of this specific object.

stable but often incorrect by several meters (as previously seen in Figure 10.13). The sensor fusion incorporates a radar estimate if, and only if, it matches closely enough to the V2V estimate, thus using the strengths of each sensor.

Table 10.4: Worst recorded errors at the three situations

		DSRC	Radar	Fused
Situation 1	Worst Error in Angle Est.	20°	10°	6°
Situation 2	Worst Error in Distance Est.	1.3 meters	---	1.6 meters
	Worst Error in Angle Est.	58°	---	54°
Situation 3	Worst Error in Distance Est.	4.7 meters	---	4.5 meters
	Worst Error in Angle Est.	11°	---	11°

Even in the situation wherein radar is available, it gives less worst-case accuracy than the fused result. This is partially because radar jumps a lot and the fusion method incorporates past measurements from both sources as well as present measurements. For similar reasons, the fused method outperforms DSRC in the situations where radar is unavailable, but is less accurate for judging the distance between adjacent (moving) cars. The degree to which to incorporate past measurements is an important parameter of the fusion method that can be adjusted by using more testing data. It is also shown, numerically this time, that the fusion method can combine the information from radar and DSRC to provide a more accurate measurement of vehicle position than either individually. The situation in question is estimation of the angle at which a stopped vehicle is in front of you: while ten degrees of error may be acceptable, a higher degree of error could cause the vehicle to assume that it can pass by the alternate vehicle, or that braking is unnecessary. Because it takes effort to determine the actual significance of a numeric error, and because the situations in which we can calculate numeric errors are limited, the next section describes a different approach to determining sensor fusion effectiveness.

10.3.4 Qualitative analysis of sensor fusion and trajectory prediction

As an alternative analysis, we used the driver's-view video that was recorded during testing. The estimated location of the vehicle is inserted into this video and can be visually compared to the actual position of the vehicle. For this analysis we focus on the blind intersection scenario.

The sensor fusion technique used in this analysis is the simplest trajectory prediction model studied in Task 1—the constant-velocity model—to combine present and past GPS/radar measurements into a cleaner estimate of the oncoming car's position. Figure 10.21 shows several examples of this sensor fusion. The vehicle position detected by DSRC is represented as a blue square and radar-detected objects are represented as green squares, as in the previous section. The final estimate (using the joint information) of the car's position is shown as a cyan circle. This circle is larger when the two vehicles are closer to each other. Note that the vertical displacement of the vehicles is not estimated. GPS and radar do not give accurate vertical information (note that the camera and the radar are not positioned at the same height), and vertical positioning is unlikely to significantly interfere in collision detection.



Figure 10.21: Snapshots of video with sensor detections and combined-position estimate.

The same physical model can be used to predict the vehicles' future positions. It was shown in Task 1 that road-fitted trajectory prediction models are more accurate. However, these experiments were set up so that the paths the vehicles follow do not cross, so there is no chance of collision if both vehicles follow the road. A physical trajectory prediction system is used to highlight dangerous situations if the vehicles followed other paths, or veered out of their lane. Figure 10.22 shows snapshots of increasing time as both vehicles approach an intersection. The cyan circle represents the current estimate of the vehicle's position as before, and the green circle

represents the expected position of the oncoming vehicle two seconds in the future. The future movement is overestimated as the vehicle stops before entering the intersection.



Figure 10.22: Snapshots of video with the future predicted position of the vehicle as a green circle. The snapshots are one second apart from each other.

The subject vehicle's motion is also predicted, and a red bar appears on the bottom of the video if the two vehicles are predicted to be dangerously close in the near future (within two seconds), emulating a collision warning. Figure 10.23 shows this occurring at the same intersection (triangles are used to represent the detected vehicle when it is out of the camera line of sight).




Figure 10.23: Snapshot of video when there is a collision warning. (In the top image the oncoming vehicle is out of camera sight and the triangle on the left shows that the vehicle is off screen to the left).

10.3.5 System limitations

This experiment showed several gaps with the current sensor fusion and trajectory prediction systems. The first is that radar gives very noisy information, and provides more value when present

for a steady period of time such as when detecting another car driving in front. For intersection cases where the vehicle may appear only briefly, vehicle communications are much more reliable. In general, an optimal sensor-fusion method should recognize situations in which radar or communications are more reliable and adjust its operation accordingly. This is a goal for sensor fusion in Phase II. The second phase will also involve more complex vehicle experiments, and this phase has shown that would be extremely valuable to gather true vehicle positions throughout the experiment. For instance, cameras on the side of each vehicle and markers with known positions can give meter-level accuracy. The second finding is that, in some cases, physical predictive models provide greater subjective value than road-based models because they consider only time and distance and do not incorporate rules such as lane division. In the current experiments the vehicles did not cross paths (there was no left turn, for example) and, therefore, collision between the vehicles was only possible from rear-ending, unless the vehicles were to deviate from their lane or direction (then a side collision could happen). Thus a road-based model as shown in Chapter 2 will not detect any dangerous activity (because it will assume that each vehicle stays within its lane limits). However, considering that, especially in intersections, crashes can happen because one vehicle deviates from its lane, the knowledge that another vehicle is approaching the intersection may be appreciated on its own, and a combination of physical predictive models and road-based models might be required.

Chapter 11. Assessing the performance of the CW/CA system through simulations

11.1 Introduction

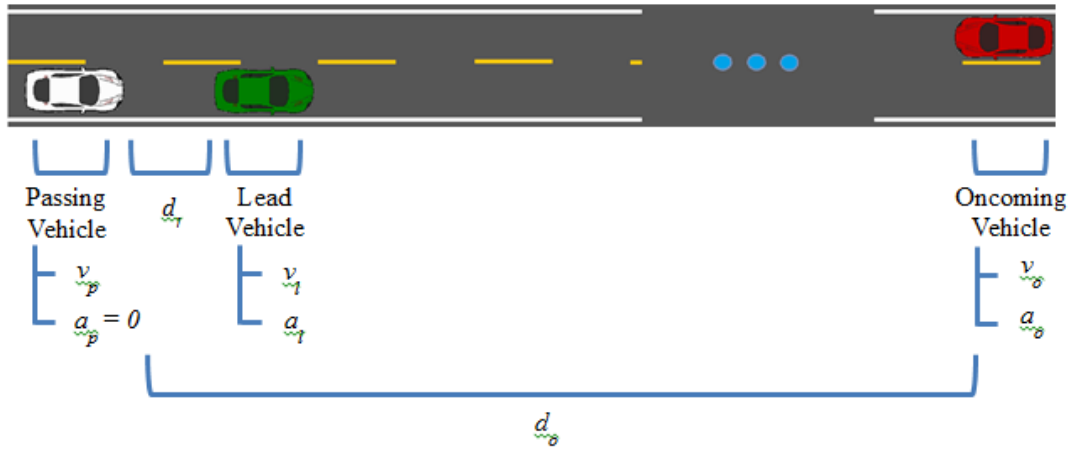
The simulator developed in the initial chapters was used for a detailed analysis of overtaking maneuvers, as performed on two-lane rural roads. These maneuvers are responsible for 54 percent of traffic fatalities on rural highways [136], yet research on preventing such collisions through collision avoidance/warning is still in the beginning stages [137, 138]. Collision warning/avoidance for overtaking maneuvers faces some different challenges than typical CW/CA tasks, such as determining whether a gap is considered safe for overtaking, given the trajectory information of the vehicles in the vicinity. During an overtaking maneuver, oncoming vehicles will be travelling toward the passing vehicle rapidly enough that braking may not be sufficient. In addition, until the lead vehicle has been completely passed it is not possible to switch back to the correct lane. The ultimate effect is that, when overtaking, CW/CA systems cannot expect to safely abort a maneuver midway, and must instead decide beforehand whether or not a maneuver is safe. This requires predicting vehicle paths up to 10 seconds in the future, far more than for other CW/CA algorithms. Additionally, analysis of an overtaking maneuver requires modelling at least three vehicles, while for most scenarios any pair of vehicles can be modelled separately. Finally, sensory requirements are different. In urban areas vehicular sensors must detect many surrounding vehicles simultaneously. On the other hand, overtaking is expected to occur on sparsely populated roads where the number of moving objects to be detected is much lower, but these objects move much faster and may not necessarily be in the overtaking vehicle's line-of-sight. Therefore, determining the location of oncoming traffic (i.e., traffic in the opposite lane) is usually not a task that radars, lasers, or cameras have been able to achieve successfully, mainly because the reported detection ranges of these sensors are shorter than the safe overtaking sight distances (or passing sight distances) recommended in the transportation literature [139, 140, 141, 142]. Thus, V2V and V2I communications are much better suited for this task than standard sensors.

11.2 Phases of the Overtaking Maneuver and Definition of Unsafe Maneuvers


Per the terminology of Hegeman et al. [139], we consider a simple, *normal* overtaking maneuver involving three vehicles on a two-lane rural roadway: passing vehicle, lead vehicle, and oncoming vehicle. In Figure 3.1, the passing, lead, and oncoming vehicles are represented by the white, green, and red colored vehicles, respectively. All three vehicles are considered passenger vehicles, each of length 5.8 meters (19 feet).

The simulation is assumed to begin when the passing vehicle indicates its desire to overtake the lead vehicle traveling ahead of it. At the beginning of the simulation ($t = t_0 = 0$), the passing vehicle is assumed to be traveling behind the lead vehicle at a constant speed (i.e., no acceleration, or $a_p = 0$ as in Figure 11.1) in its travel lane; the speed of the passing vehicle is assumed to remain constant for the duration of its driver's PR time (t_{pr}). (As discussed later, we allow this PR time to be heterogeneous in the population of drivers.) During the perception/reaction time ($0 \leq t < t_{pr}$), the driver is assumed to perceive and process information on the lead vehicle and oncoming vehicle and determine whether the gap available is safe for completing the overtaking maneuver. At the end of the PR time ($t = t_{pr}$), the passing vehicle is assumed to accelerate and move into the opposite lane. This is considered the start of the overtaking maneuver.

$t = 0$: Initialization



$t = t_{pr}$: Overtaking begins



$t = t_{fin}$: Overtaking complete




Figure 11.1: Depiction of overtaking maneuver simulations.

Once in the opposite lane, the passing vehicle is assumed to travel at a constant acceleration ($a_p > 0$) until it overtakes the lead vehicle and gains a one second *headway* ahead of the lead vehicle. In this context, the term *headway* refers to the time the lead vehicle will require to traverse the gap between the front of the lead vehicle and the back of the passing vehicle (i.e., the time required for the lead vehicle to travel d_l distance shown in the bottom part of Figure 11.1). This is equivalent to the time until collision between the lead vehicle and a (hypothetical) stationary object at the rear end position of the passing vehicle in the bottom part of Figure 11.1. At the time instant that the passing vehicle's *headway* becomes one second ahead of the lead vehicle, the passing

vehicle is assumed to have returned to the original lane to complete the overtaking maneuver, *if* the maneuver were a successful one.

At the moment the passing vehicle's headway becomes one second ahead of the lead vehicle ($t = t_{fin}$), the *time-to-collision* may be calculated between the passing vehicle and the oncoming vehicle. The term *time-to-collision* refers to the amount of time in which the passing vehicle would collide with the oncoming vehicle had it continued traveling in the opposite lane. To be precise, if the passing and oncoming vehicles maintain their speeds and accelerations at time $t = t_{fin}$, time-to-collision is the time in which the two vehicles would together travel the distance between their front bumpers (denoted by d_o in the last part of Figure 11.1). If the *time-to-collision* is less than one second, we deem this overtaking maneuver as unsafe, as in Harwood et al. (2008), and label it as resulting in a collision.³ On the other hand, if the *time-to-collision* is greater than one second, we deem the overtaking maneuver as safe (and resulting in no collision).

Throughout the discussed duration (i.e., from the beginning to the end of the simulation), the lead vehicle and the oncoming vehicle are assumed to travel at their respective constant acceleration rates in their respective lanes, regardless of the position, speed, and acceleration of the passing vehicle. While it may be considered a bit too conservative, we did not want the assistant to rely on the oncoming vehicle's braking because the maneuver may err toward a collision if the oncoming vehicle does not brake (but the assistance system assumed it would). The crash statistics mentioned in the introduction make it clear that other vehicles will not always notice or avoid an overtaking vehicle in time.

11.3 Simulation Setting

We performed simulations to evaluate the accuracy of the CW/CA system based on V2V communication. Two thousand randomized simulations of overtaking maneuvers were conducted. After obtaining the datasets of maneuvers, we tested the effect of different levels of equipment accuracy on the performance of the CW/CA system.

11.3.1 Distributions of Simulation Variables

The passing vehicle's driver PR time after he/she indicates a desire to overtake is drawn from a triangular distribution between one to four seconds with a mode of 2.5 seconds. Since PR times vary depending on the driver's state (e.g., alertness or fatigue), complexity of the driving situation, and the type of highway [143], assuming a maximum of four seconds captures that drivers might need longer PR times in rural settings than in urban settings, and in passing maneuvers than in simpler driving tasks.

The initial speeds (i.e., speeds at the beginning of each simulation; denoted by v_p , v_l , and v_o in Figure 11.1) for the three vehicles are generated from a truncated normal distribution with a mean value 70 mph, minimum value 55 mph, and maximum value 90 mph. Typical speed limits for rural interstates in the US range from 55 to 80 mph [152]. We simulated scenarios over the typical speed limits on rural two-lane highways to capture excessive speeding situations.

The passing vehicle's acceleration is assumed to be zero at the beginning of the simulation. After the perception/reaction time, the passing vehicle is assumed to accelerate at a constant rate

³ Of course, not all situations where the *time-to-collision* is less than one second may result in collisions. To be precise, a collision happens only when the time gap goes to zero or beyond. However, since it is not safe to be within such a small *time-to-collision* window, we deem all such collision-prone situations (with less than one second *time-to-collision*) as collisions.

and move into the opposite lane. This acceleration is drawn from a truncated normal distribution with mean 3.6 ft/sec^2 and truncated at 1 ft/sec^2 and 8.2 ft/sec^2 (see [144] for empirical data on accelerations in rural roads). The accelerations for the lead and oncoming vehicles (denoted by a_l and a_o in Figure 11.1) are drawn from another normal distribution with mean zero and truncated at $\pm 3.2 \text{ ft/sec}^2$ on both sides of the distribution [144]. Deceleration was allowed only for the lead and oncoming vehicles because the passing vehicle cannot typically overtake the lead vehicle while decelerating.

The vehicular dynamics in the simulation begin with positioning the passing vehicle in the right lane at initialization (t_0). Subsequently, the lead vehicle is positioned in the right lane at an arbitrary location (drawn from uniform distribution) ahead of the passing vehicle's initial location as long as its position is within 15 feet of a one second headway in front of the passing vehicle. The oncoming vehicle's initial position is set to be uniformly distributed between a lower bound and an upper bound such that the passing and oncoming vehicles are neither too close at the beginning of the overtaking maneuver nor very distant at the end of the maneuver. The lower bound of the allowed distance between the passing and oncoming vehicles was taken as the minimum distance needed for a vehicle (taking the fastest possible maneuver) to successfully overtake, minus one second of headway. In other words, a scenario with an initial passing-oncoming distance at or below the lower bound would never result in a safe maneuver. The upper bound was obtained from the speed-dependent passing sight distance (PSD) guidelines from AASHTO [146].⁴ It is worth noting here that the initial vehicle-to-vehicle spacing and other parameters were set such that a considerable proportion of simulated overtaking maneuvers are difficult (but not unrealistic) to complete, since one of the objectives of this research was to assess the usefulness of communications in preventing overtaking crashes. Since the focus of this research study is to evaluate DSRC's effectiveness in an overtaking safety application, some outliers were excluded from the simulated data. Scenarios where the lead vehicle is travelling more than 10 mph faster than the passing vehicle at PR time were discarded, as an overtaking maneuver is very unlikely to occur in such circumstances. Scenarios where the overtaking vehicle failed to pass the lead within 0.621 miles (1 kilometer) were also considered unrealistic and discarded. Finally, scenarios in which the oncoming vehicle passes the lead vehicle before the PR time were discarded.

11.3.2 Communication Range

For the CW/CA system to estimate the trajectory of the lead or oncoming vehicles, the two vehicles must be within communication range of the passing vehicle to receive the CAMs containing position, speed, and acceleration information of the lead and oncoming vehicles. The communication range, in turn, depends on the maximum transmit power of the DSRC devices. The Federal Communications Commission defines four classes of DSRC devices depending on their maximum allowed transmit powers as: Class A, Class B, Class C, and Class D. DSRC devices are normally in the Class C category, with a maximum transmit power of 20 decibel-milliwatts (or dBm (dBm is a logarithmic scaled unit of milliwatts) [145]. On the receiving side, devices are only guaranteed to correctly receive messages above a certain power, which is referred to as the minimum sensitivity. IEEE requires the minimum sensitivity of VANET (Vehicular Ad-hoc

⁴ The PSD calculations from the AASHTO Green Book are used to set the upper bound on the initial distance between the passing vehicle and the oncoming vehicles, because these PSD values are considered to be very conservative in the literature [140].

Network) systems to be at least -85 dBm. A wireless signal's loss in power over distance is measured by its path loss exponent, which has a value of two in free space. We opted to set the path loss exponent to 2.1 due to the low density of vehicles on rural roads. For these communication strength settings, the communication range in our simulations was approximately 600 meters (2000 feet)⁵. This doesn't necessarily imply that V2V communication is fully present before 600 meters and becomes completely absent right after 600 meters. Rather, the reliability of the communication is likely to taper continuously (but quickly) beyond 600 meters.

If an oncoming vehicle is out of communication range when the overtaking begins, there would be no communication of information between vehicles. In such situations, there would be no warning issued by the overtaking assistant, even if the passing maneuver would lead to a collision. Therefore, to ensure timely onset of communications between vehicles involved in overtaking maneuvers, it is useful that the communication range be more than the design-speed dependent safe passing distances given in AASHTO's Green Book [146]. At the least, the passing and oncoming vehicles must come within the communication range before the passing vehicle driver's PR time. However, increasing the communication range has not been a major focus in the development of DSRC devices since the allocated spectrum is designed to support many other applications [145]; and, widely researched applications such as collision warning at intersections or platooning require a much shorter range [147].

In addition to 20 dBm transmit power, we also simulated scenarios with transmission powers of 17 and 23 dBm, which are close to half and double the power of 20 dBm and roughly equate to maximum communication ranges of 430 and 860 meters (1,400 and 2,800 feet) respectively. Note that other factors such as minimum sensitivity and path loss were kept constant, as they have a very similar effect on communication range.

11.3.3 Packet Error

When the vehicles are within communication range, the receipt of speed and acceleration information may be affected by communication errors called packet errors that lead to the loss of some CAMs without their receipt. One major cause of these errors is latency, or the delay between a message's initial broadcast and complete reception. The DSRC standards for the US specify communication every 100 milliseconds [145], so a message with latency greater than 100 milliseconds will be abandoned as the next message is sent. Latency is not constant and is determined by many factors, such as congestion caused by high vehicle density and the data size of each message. Other miscellaneous issues, including physical interference from precipitation or obstacles and software errors, could also prevent a single message from being received. The term packet error encompasses all the reasons (other than vehicles being outside communication range) why timely communication may not be established between vehicles and why, therefore, potential collisions may not be detected by the CW/CA system.

Communication protocols are generally designed to maintain an acceptable rate of packet errors for a given application. Congestion control methods, for instance, focus on minimizing the bandwidth used by each broadcasting vehicle while ensuring that all important information is transmitted reliably. Advanced DSRC communication protocols are still an active area of research

⁵ The limited data collection performed in this task did not include rural road scenarios. In the urban environment, we observed a reduction in reception rate starting from when vehicles are 150 meters apart. The loss becomes more significant when the distance is greater than 300 meters. However, since our data collection was very limited, the following simulations are based on results available in the literature such as the ones reported by [148].

(for instance, see 149, 150, 151). While Veins is capable of simulating many protocols and error sources, the exact nature of these error sources for overtaking applications is not known. For instance, high-density traffic is uncommon on rural roads and less likely to permit overtaking in the first place. Rather than make arbitrary assumptions on each case, we encompass all errors into a single packet error rate. For each message successfully received by the overtaking vehicle (within the Veins simulator), with a certain probability this message will be removed and not reported to the overtaking assistant. This probability number is termed the packet error rate. Simulations were performed with the following packet error rates: 0 percent, 50 percent, 75 percent, and 87.5 percent.

11.3.4 Sensor and Estimation Inaccuracy

In DSRC-enabled connected vehicles, many in-vehicle sensors are used to determine the position, speed, and acceleration of the vehicles. Such sensor measurements are, of course, subject to sensing error (or inaccuracy), which in turn influences the accuracy of the trajectory prediction. To capture this, each simulated measurement of the vehicle position, speed, and accelerations was subject to random noise to represent sensor error (or inaccuracy) of the variables used for trajectory prediction. That is, while the values of the position, speed, and acceleration variables used for simulating each scenario were assumed as “true” values, the corresponding values used for trajectory prediction were subject to sensor error. This is one reason why the trajectory predictions could differ from the simulated trajectories.

The magnitude of sensor error for all variables was controlled by a single noise parameter η . When η is 0 percent, information used for trajectory prediction is assumed to be known perfectly. That is, the values of the vehicle state variables used for trajectory prediction are exactly the same as the simulated values. For nonzero η , normally distributed noise is added to each value, the magnitude of which depends on η . For a variable X with a measurement x (i.e., a simulated value x), the after-noise measurement \hat{x} , which is used for trajectory prediction, is considered to be normally distributed as:

$$\hat{x} \sim N\left(\mu = x, \sigma = \frac{\eta}{100} X_{range}\right)$$

In the above equation, X_{range} is 2 meters for position variables, 0.5 m/s for speed variables, and 0.25 m/s² for acceleration variables⁶. As vehicle positioning is typically achieved by a combination of GPS location and reckoning/filtering, the X_{range} value position error was taken as half the standard RMS of error for GPS [153]. For the velocity and acceleration sensors used within vehicles, the X_{range} values are chosen such that the sensor error is in the similar range as in standard commercial devices (see [154] for an accelerometer example).⁷

In addition to the above discussed sensor errors, it is important to note that the passing vehicle's behavior variables—driver's PR time and acceleration during overtake—cannot be known with certainty before the beginning of the overtaking maneuver. Therefore, the overtaking assistant has to *estimate* the driver's PR time and acceleration for trajectory prediction purposes. To capture such uncertainty (or errors) in estimation, these two variables were subject to a random

⁶ Again the data collected in this task was insufficient to be used as a reliable source for the simulations. However, what we observed in the field is not significantly different from the assumptions adopted herein.

⁷ The error bounds on each variable are relative to its assumed possible error, not the overall range or significance of its values. A separate study with a different variable to represent error on each of the eight sensed variables is outside the scope of this paper.

noise, using the same control parameter η used for sensor error. The parameters should still follow all previously outlined assumptions on realistic driving parameters (i.e., the maximum and minimum threshold values). Thus, for a variable X with a measurement (or simulated value) x and the threshold values X_{max} and X_{min} , the estimated value \hat{x} is distributed as a truncated normal:

$$\hat{x} \sim N\left(\mu = x, \sigma = 2 \frac{\eta}{100} (X_{max} - X_{min})\right), X_{min} \leq \hat{x} \leq X_{max}$$

In our simulations, multiple settings are tested for the sensor/estimation error rate (η) in conjunction with the packet error rate. These are: 0, 25, 50, and 100 percent for the sensor and estimation error parameters (η); and 0, 50, 75, and 87.5 percent for the packet error rate.

The simulations include nine distinct combinations of transmission power, packet error rate, and sensor/estimation error rates, as itemized in the bottom right of Table 1. Note that some of these settings, particularly those with high packet error (higher than 50 percent) or high sensor inaccuracy rates (higher than 50 percent) may not be realistic vis-à-vis the current performance of DSRC devices, but are considered in the simulations to allow for worst-case communication settings.

11.4 Simulated Data for V2V Communication

The simulated dataset to investigate the use of V2V communication for the CW/CA system includes 2,000 unique overtaking scenarios in terms of vehicle dynamics. Each scenario is used to test nine different settings of the CW/CA system. This results in 18,000 overtaking assistance simulations, with 14,121 collisions (78.8 percent) and 3,879 (21.6 percent) non-collisions. It is worth noting here that we purposely simulated a higher than realistic proportion of collisions to obtain a sufficient sample of collisions to study.

Of the 14,121 collisions, the DSRC-enabled CW/CA system detected collisions in a timely manner (i.e., detected collision before driver's perception reaction time) for 9,496 cases (67 percent successful) but did not detect collisions for the remaining 4,625 cases. Among all the 3,879 simulated successful overtaking maneuvers without a collision, passing vehicles took an average of nine seconds to complete the overtaking maneuver, which is consistent with the overtaking maneuver times reported in previous literature [155, 156]. The warning system detected collisions (i.e., unnecessary or false warnings) for less than 4 percent of the 3,879 successful (or safe) overtaking maneuvers.

Table 11.1 presents the inputs used across all the overtaking maneuver scenarios studied in this research. These include driver behavior and vehicular dynamics (i.e., PR time, initial speed and acceleration of all the three vehicles—passing, lead, and oncoming vehicles) and V2V communication settings. In addition, initial distances between (1) passing and lead vehicles and (2) passing and oncoming vehicles are presented to give a sense of relative positioning of the vehicles in the beginning of the simulation. As can be observed, the descriptive statistics of the driver behavior and vehicular dynamics parameters are consistent with the assumptions made on these parameters in Section 11.2. The parameters defining V2V communication settings include the frequency of CAM messages, power setting parameters (transmission power, minimum sensitivity, and path loss exponent, packer error rate, and sensor/estimation inaccuracy rates. As discussed in Section 3.3, the frequency of CAM messages and some power setting parameters were fixed across all simulations, while the transmission power, packet error rate, and sensor and estimation inaccuracy rates were varied.

Comparison of Table 11.1's descriptive statistics between simulated collisions and non-collisions provides insight into how driver behavior and vehicular dynamics might influence collision and non-collision outcomes. Within driver behavior and vehicular dynamics, a higher proportion of passing vehicles with a longer driver PR time ended up in collisions. This result demonstrates the importance of quick and correct decisions in overtaking maneuvers and highlights the need for V2V technologies that can potentially assist in making quick decisions. It can be seen that passing vehicles in the highest speed category (> 80 mph) show a greater chance of avoiding a collision, despite the notion that fast driving is more dangerous. Yet, this result needs to be interpreted with caution, because in our simulations the maximum allowed distances between the passing and oncoming vehicles are speed-dependent (see Section 11.2.1). So fast passing vehicles often start farther away (from oncoming vehicles) than slower vehicles, and therefore might lead to safer simulated maneuvers. Lead vehicles in the slowest speed category (< 70 mph) are also represented in greater proportions in non-collisions than in collisions. A larger proportion of non-collisions started with a larger initial gap between the passing and oncoming vehicles (> 750 m). A different trend is seen in the case of collisions, where the largest proportion of collision scenarios start with an initial gap of 600 to 750 meters.

One may note that the V2V communication parameters have no influence on simulated collision or non-collision outcomes. This is because the simulations allowed all the overtaking maneuvers to complete despite any warning from the overtaking assistant. Such simulation outcomes are compared with the trajectory prediction outcomes (which depend on the V2V communication settings) to understand the performance of the CW/CA system.

Table 11.1: Descriptive statistics of the simulated data

	All Scenarios	Collisions	Non- Collisions		All Scenarios	Collisions	Non- Collisions
<i>No. of observations</i>	2000	1569	431	<i>No. of observations</i>	2000	1569	431
Driver Behavior & Vehicular Dynamics				Driver Behavior & Vehicular Dynamics (continued)			
Passing Vehicle				Initial Distance between Vehicles			
Perception/Reaction Time (Seconds)				Passing and Lead (m)			
Min	1.01	1.01	1.11	Min	18.9	18.9	25.2
Max	3.98	3.98	3.78	Max	44.6	44.6	44.5
<= 3 seconds	80%	80%	79%	< 30 m	22%	26%	8%
> 3 seconds	20%	20%	21%	30-40 m	71%	68%	78%
Initial Speed (mph)				> 40 m	7%	6%	14%
Min	49.1	49.1	63.5	Passing and Oncoming (m)			
Max	90.0	89.9	90.0	Min	390	390	513
< 70 mph	29%	34%	8%	Max	1153	1101	1153
70-80 mph	44%	44%	41%	< 600 m	20%	24%	5%
> 80 mph	28%	22%	51%	600-750 m	42%	47%	27%
Overtaking Acceleration (m/s ²)				> 750 m	38%	30%	68%
Min	0.306	0.306	0.332	V2V Communication Settings, Fixed			
Max	2.50	2.49	2.50	Frequency of Cooperative Awareness Messages – 100 ms			
< 1 m/s ²	34%	37%	26%	Minimum Sensitivity – -85 dBm			
1-1.5 m/s ²	43%	44%	41%	Path Loss Exponent – 2.1			
> 1.5 m/s ²	28%	26%	37%	V2V Communication Settings, Combinations			
Lead Vehicle				Transmission Power (dBm)	Packet Error Rate	Sensor and Estimation Inaccuracy Rate (or Noise)	
Speed (mph)				20	0%	0%	
Min	55.0	55.0	55.0	20	50%	0%	
Max	89.9	89.9	87.5	20	75%	0%	
< 70 mph	47%	42%	68%	20	87.5%	0%	
70-80 mph	37%	40%	27%	20	0%	25%	
> 80 mph	16%	19%	5%	20	0%	50%	
Acceleration (m/s ²)				20	0%	100%	
Min	-0.998	-0.890	-0.998	23	0%	0%	
Max	0.972	0.972	0.632	17	0%	0%	
<= 0 m/s ²	51%	47%	64%				
> 0 m/s ²	49%	53%	36%				
Oncoming Vehicle							
Speed (mph)							
Min	55.0	55.0	55.1				
Max	89.8	89.6	89.8				
< 70 mph	49%	47%	55%				
70-80 mph	36%	37%	33%				
> 80 mph	15%	16%	12%				
Acceleration (m/s ²)							
Min	-0.979	-0.979	-0.840				
Max	0.931	0.931	0.765				
<= 0 m/s ²	53%	51%	59%				
> 0 m/s ²	48%	49%	41%				

Figure 11.2 shows the cumulative distribution of the distance between passing and oncoming vehicles, at the starting time of an unsafe (ultimately collision-causing) overtaking maneuver. This is the final, and minimum, distance at which these two vehicles may communicate to enable an automated warning. Therefore, the distribution of this value provides insight into the essential range of communication for reliable overtaking assistance: for any given distance, this figure displays the proportion of overtaking maneuvers that could have had sufficient communication at the matching DSRC communication range—excluding other factors such as congestion-related packet error. In order to capture nearly every unsafe overtaking maneuver, vehicular communication will have to operate over roughly 900 meters. This is a tall order for DSRC, as it is usually designed for other goals (see 147, 157, 158 for typical assumptions of the maximum necessary distance for urban ADAS). Using more typical long-range DSRC settings, which achieve less than 700 meters, an overtaking assistant may not detect at least 10 percent of unsafe maneuvers.




Figure 11.2: Cumulative distribution of collision maneuvers versus initial distance.

This insight is matched by simulation results. In total, out of 4,625 undetected collisions, 4,555 (98.5 percent) occurred because communication was not established between the passing and oncoming vehicles. For 4,498 (98.7 percent) of the undetected collisions where communication was not established, the passing and oncoming vehicles had still not come within communication range before the passing vehicle driver’s PR time.⁸ This suggests that communication range is the primary factor in the performance of the overtaking assistant.

Figures 11.3 and 11.4 both categorize the scenarios by four assisted overtaking outcomes—undetected collisions, correctly detected collisions, no-collision scenarios without warning, and no-collision scenarios with a false (or unnecessary) warning. Figure 11.3 shows the distribution of

⁸ As discussed earlier, the three DSRC power settings employed in our simulations imply communication ranges of about 430, 600, and 860 meters respectively. However, this doesn’t necessarily imply that V2V communication is fully present before 430 meters and becomes completely absent right after 430 meters. Rather, the strength of the communication is likely to taper continuously (but quickly) beyond 430 meters.

the actual time-to-collision - i.e., the time it took for the passing vehicle to collide with the oncoming vehicle after the passing vehicle achieved 1 second headway ahead of the leading vehicle. For the majority of simulations with an actual time-to-collision greater than 1 second and a warning from the assistant – in other words, a false warning – the time-to-collision is less than 2, suggesting that the majority of false warnings are issued for scenarios that were relatively close to collision. Therefore, the issue of false warning does not appear to be a severe issue in the context of DSRC-assisted collision warning systems for overtaking scenarios on rural highways. On the other hand, the ratio of warnings for collision scenarios (with time-to-collision less than one second) appears unrelated to the time-to-collision.




Figure 11.3: The actions of the overtaking assistant versus the actual time-to-collision between passing and oncoming vehicles.

Figure 11.4 presents descriptive statistics to understand the influence of packet error and sensor/estimation error rate (or noise) on the performance of the overtaking assistant. The information is presented separately for scenarios that resulted in collisions and scenarios that did not lead to collisions. For the collisions, the figure presents the distribution of the scenarios between undetected collisions and detected collisions for different levels of packet error and sensor and estimation error rates. Similarly, for non-collisions, the figure presents the distribution of the scenarios between cases where no warning was issued and cases where a false warning was issued by the overtaking assistant.

As can be observed from the bars representing “undetected collisions,” when transmission power was set to 20 dBm and both packet error rate and noise were set to 0 percent, 26.8 percent of the collisions were undetected. All of these undetected collisions may be attributed to lack of communication due to vehicles being out of communication range before the passing vehicles’ PR time. An increase or decrease in transmission power critically affects the level of undetected collisions, reducing them as low as 1 percent for 23 dBm communication or as high as 78 percent for 17 dBm.

As the packet error increases from 0 percent to 87.5 percent, the percentage of collisions that were not detected increases and the percentage of collisions that were detected decreases (see the column titled detected collisions). However, the increase in the percentage of undetected collisions is less than 5 percent for packet error rates of up to 75 percent. It is only beyond 75 percent packet error rates that the percentage of scenarios with undetected collisions increases considerably. In reality, as discussed earlier, packet errors of DSRC devices are rarely as high as 75 percent or more. Therefore, these results suggest that the influence of packet error rates on missing the detection of a potential collision is not as strong as that of the vehicles being out of communication range. Next, note from the bars representing “false warning” that increasing packet error rate did not influence whether or not a false warning is issued for overtaking scenarios that did not end up in collisions. This is expected because packet errors influence only whether communication is established or not, not the accuracy of trajectory prediction itself.




Figure 11.4: CW/CA system accuracy measures.

The rightmost segment of the figure corresponds to the influence of sensor and estimation error rates (i.e., the noise parameter) on the performance of the overtaking assistant. As can be observed from the collision bars, increasing the sensor and estimation error rate leads to a small decrease in the ability to detect collisions. Specifically, the percentage of collisions that were not detected increases from 26.8 percent at zero noise to only 28.3 percent at 100 percent noise. On the other hand, the percentage of no-collision scenarios that had a false warning issued by the overtaking assistant rises to 7.2 percent at 25 percent noise and 9.0 percent at 50 percent noise. These trends suggest that the sensor and estimation errors, as simulated, are more likely to cause the overtaking assistant to be overly conservative, leading to false warnings, than to be overly optimistic, leading to undetected collisions or false negatives. This is expected because the sensor and estimation errors simulated in our experiments were symmetric around the true values (i.e., not biased toward the right or left of the true values) and sensor errors varied with every V2V message sent. Furthermore, only a single prediction of collision is needed at any time between the beginning of the scenario and the PR time of the passing vehicle. While the predicted time-to-collision is equally likely to be conservative or optimistic (because sensor and estimation errors are symmetric), the collision warning is issued on the first instance the predicted time-to-collision is less than one second. Therefore, sensor and estimation errors combined with our collision warning protocol primarily increase the likelihood of false warnings. It’s worth noting that symmetric and time-varying noise is likely to have a stronger effect on the overtaking assistant’s

performance than constant or one-sided noise for the same reason: the assistant can overreact to one point in time with exceptional noise.

11.5 Performance of the CW/CA system based on V2V communication

The descriptive analysis of the simulated data provides useful insights on the influence of V2V communication parameters on the effectiveness of DSRC-enabled warning systems in predicting and preventing rural road overtaking collisions. Nevertheless, a univariate descriptive analysis cannot conclusively isolate the influence of different factors on the performance of the DSRC-enabled warning systems. One reason is that the safety of an overtaking maneuver, or even whether a simulated overtaking maneuver is realistic in the first place, depends on multivariate relationships in the vehicle dynamics. As an example, the speed of the passing vehicle might be related to the effectiveness of an overtaking assistant. A univariate analysis requires simulations of varying speed while all other vehicle parameters are kept constant (otherwise correlation effects can lead to false conclusions). However, depending on the distance to and speed of the lead vehicle, a passing vehicle's speed may not be high enough for an overtaking maneuver to occur, or may be so high that the driver must slow down before his perception/reaction time is complete. Thus, it is not possible to fix the environment while analyzing the overtaking assistant. Therefore, the next section provides a multivariate analysis to isolate the influence of each of the above factors while controlling for the influence of vehicular dynamics and driver behavior variables. The simulated data discussed above were used to estimate three binary discrete outcome models.

The first model, called collision occurrence model, was estimated on all 18,000 simulated overtaking maneuvers to examine the influence of driver behavior and vehicular dynamics on collision occurrence (i.e., whether collision occurred or not). The second model, called collision detection model, was on only the subset of simulated overtaking maneuvers that resulted in collisions. This model explores the influence of driver behavior, vehicular dynamics, and V2V communication parameters (packet error rate and sensor/estimation inaccuracy rate) on the ability of the DSRC-enabled system to detect collisions⁹ in a timely manner (i.e., before the passing vehicle driver's PR time). The binary outcomes analyzed in this model are: (1) undetected collision and (2) detected collision. The third model, called false warning model, was on only the subset of simulated overtaking maneuvers that did not result in collisions. This model was used to examine the influence of various factors on the likelihood of the overtaking assistant to provide unnecessary warnings (or the false alarm of a collision). The binary outcomes analyzed in this model are: (1) collision detected but there was no collision (i.e., false warning), and (2) no collision detected and there was no collision. The parameter estimates of all the three models are presented in Table 11.2.

⁹ Recall that a collision would be detected if the estimated time-to-collision (i.e., time-to-collision at the instance the passing vehicle's headway is one second ahead of the lead vehicle) is less than one second.

Table 11.2: Binary probit model estimation results

	Model #1	Model #2	Model #3
	Collision Occurrence (base: Non-Collisions)	Undetected Collisions (base: Collision occurred and warning issued)	False Warnings (base: No collision and warning not issued)
<i>No. of observations</i>	18000 overtaking maneuvers	14121 overtaking maneuvers that lead to collisions	3879 overtaking maneuvers that did not lead to collisions
	Coefficient (t-stat)	Coefficient (t-stat)	Coefficient (t-stat)
Constant	4.4888 (50.96)	-1.0543 (-20.67)	-2.1587 (-21.30)
Passing Vehicle			
Perception/Reaction Time			
<= 3 seconds	base category	base category	base category
> 3 seconds	0.2862 (6.63)	-1.2267 (-25.40)	0.2974 (2.15)
Initial Speed			
< 60 mph	--	--	0.4797 (3.57)
60-70 mph	base category	--	0.4797 (3.57)
70-80 mph	-1.3560 (-26.04)	base category	base category
> 80 mph	-1.4802 (-32.74)	-0.2123 (-3.99)	--
Overtaking Acceleration			
< 3 ft/sec ²	0.5779 (15.07)	0.1068 (2.84)	-0.3055 (-2.14)
3-5 ft/sec ²	base category	base category	base category
> 5 ft/sec ²	-0.4291 (-11.20)	--	--
Lead Vehicle			
Speed			
< 60 mph	-3.8337 (-47.22)	--	--
60-70 mph	-1.9778 (-41.04)	base category	--
70-80 mph	base category	0.1579 (3.85)	base category
> 80 mph	1.4702 (23.58)	0.2985 (5.44)	--
Acceleration			
<= 0 ft/s ²	-1.0270 (-28.08)	-0.1271 (-3.71)	base category
> 0 ft/s ²	base category	--	--
Oncoming Vehicle			
Speed			
< 60 mph	-0.8192 (-15.21)	0.0952 (2.76)	--
60-70 mph	-0.4136 (-11.27)	0.0952 (2.76)	--
70-80 mph	base category	base category	base category
> 80 mph	0.3109 (5.87)	--	0.3636 (2.03)
Acceleration			
<= 0 ft/s ²	base category	base category	base category
> 0 ft/s ²	0.3624 (11.50)	-0.1196 (-3.45)	--
Initial Distance between Vehicles			
Passing and Lead			
< 100ft	--	-0.1492 (-3.49)	--
100-120 ft	base category	base category	base category
> 120 ft	-0.1287 (-3.23)	--	--
Passing and Oncoming			
< 2000 ft	1.4604 (24.35)	-1.5274 (-26.96)	0.5234 (3.67)
2000-2500 ft	base category	base category	base category
> 2500 ft	-1.3674 (-33.32)	2.4363 (50.01)	-1.5317 (-8.74)
V2V Communication Settings			
Packet Error Rate			
0%	**	base category	base category
50%	**	0.1562 (2.11)	--
75%	**	0.3182 (4.33)	--
87.5%	**	1.0795 (16.05)	--
Sensor & Estimation Error Rate			
0%	**	base category	--
25%	**	--	base category
50%	**	--	1.2205 (9.35)
100%	**	--	1.5277 (12.26)
Sensor Power			
17 dBm	**	2.7132 (31.80)	--
20 dBm	**	base category	base category
23 dBm	**	-3.2786 (-21.10)	--
Summary Statistics			
R ²	0.55	0.60	0.37
Restricted Log-Likelihood:	-9546.44	-6350.23	-569.11
Final Log-Likelihood:	-4295.90	-2978.21	-355.39

** Not included in model. V2V communication only warns of a potential collision but does not influence the simulated outcome.

-- Dropped from specification as the coefficient was statistically insignificant (i.e., not different from zero).

Model #1: Collision Occurrence Model

The collision occurrence model parameter estimates are shown in the second column of the table. The positive coefficient on the passing vehicle driver's perception/reaction time suggests that higher PR times increased the likelihood of collisions in our simulations. This is because the distance between the passing and oncoming vehicles diminishes as more time elapses from the beginning of the simulation. Also recall that all our simulations continued to complete the overtaking maneuver despite any potential for collisions because the primary goal of this work is to assess the effectiveness of DSRC-enabled V2V communication systems in predicting and preventing overtaking collisions. In real life situations, however, longer PR times might provide the driver an opportunity to carefully evaluate the situation and abort the overtaking maneuver if necessary. Similarly, as discussed later, in the context of the DSRC-assisted collision detection systems, longer PR times increase the likelihood of timely detection of collisions.

In the context of the vehicular dynamics of the passing vehicle, *ceteris paribus*, greater initial speeds and higher overtaking accelerations decreased the likelihood of collisions—perhaps because such passing vehicles spend less time in the opposite lane. On the other hand, the initial speed and accelerations of the lead and oncoming vehicles had an opposite influence. Greater speeds and higher accelerations of either vehicle increased the likelihood of collisions. This is because the available gap between passing and oncoming vehicles (when the passing vehicle achieves one second headway ahead of the lead vehicle) becomes smaller at higher speeds and accelerations of the lead or oncoming vehicles.

Finally, as expected, smaller initial distance between passing and lead vehicles increased the likelihood of collision, while greater initial distance between passing and oncoming vehicles reduced the likelihood of collision.

Model #2: Collision Detection Model

Model #2 may be used to examine the influence of driver behavior, vehicular dynamics, and V2V communication settings on the likelihood of a missed warning (or undetected collision) for unsafe overtaking maneuvers. Most of the parameter estimates from this model point to the relative importance of the passing and oncoming vehicles coming within communication range. For instance, in the context of driver behavior, longer PR times of passing vehicle drivers decreased the likelihood of missing the detection of a collision, presumably because longer PR times provide a greater opportunity for the passing and oncoming vehicles to come within communication range.¹⁰ In addition, increasing the speed of oncoming vehicles also increased the likelihood of a collision being properly detected by the overtaking assistant. More importantly, as can be observed from the high t-statistic values of the variable “initial distance between passing and oncoming vehicles,” this variable exhibits a significant influence on the ability to detect collisions. Specifically, scenarios that begin with a greater separation between passing and oncoming vehicles and end in collisions are less likely to be detected in a timely manner. Again, this is because a greater initial separation between the two vehicles lowers the likelihood of them coming within communication range in a timely manner (i.e., prior to passing vehicle's PR time). These results suggest that increasing the DSRC power settings to broaden the communication range may be an effective way of increasing the performance of DSRC devices for improving the safety of rural highway overtaking maneuvers.

¹⁰ Longer PR times also result in a higher likelihood for the V2V communication to overcome packet loss, which in turn, increases the likelihood of detecting collisions.

Both the speed and acceleration of the lead vehicle appear to be positively associated with the likelihood of undetected collisions. Increasing the lead vehicle speed increases the amount of time needed for the passing vehicle to complete the overtaking maneuver, thus increasing the likelihood that a distant oncoming vehicle (one outside of communication range) could cause a collision. For the same reason, the acceleration at which the passing vehicle performs the overtaking maneuver is negatively correlated with the likelihood of undetected collisions.

In the context of V2V communication settings, the transmission power has an expectedly high correlation with the detection of collisions. As the packet error rate increases beyond 50 percent, the likelihood of undetected collisions also increases, presumably because it increases the likelihood of missed communication among the three vehicles. However, as discussed earlier, packet error rates of greater than 50 percent are unlikely in DSRC-enabled V2V communication systems. Therefore, in the context of rural highways where the vehicular traffic volumes are not as high as those in urban environments, relieving communication channel congestion is perhaps not a high-priority concern unless packet error rates increase beyond 50 percent.

Sensor and estimation errors were not determined to have significant effect on the detection of collisions. Figure 11.4 shows that higher errors will in fact cause slightly fewer collisions to be detected, but this amount is so small as to be probabilistically insignificant according to a multivariate model.

Model #3: False Warning Model

The parameter estimates of Model #3 may be used to understand which safe overtaking scenarios are associated with an increased likelihood of an unnecessary warning issued by the overtaking assistant. Specifically, safe overtaking scenarios with lower initial speeds of passing vehicles, higher lead vehicle speeds, or those with higher oncoming vehicle speeds are associated with a higher likelihood of a false warning. This is because passing vehicles with lower initial speeds and lead vehicles with higher speeds tend to require a longer time for completing the overtaking maneuver. Long overtaking maneuvers and fast oncoming vehicles may lead to situations that are near collisions but deemed safe (i.e., time-to-collision is higher than one second, but not by much). As seen in Figure 11.3, such maneuvers are common in these simulations and contain a high proportion of false warnings. In such cases, it is perhaps easier for sensor and estimation errors (that influence the trajectory prediction) to cause an under-estimation of the time-to-collision to be below one second, leading to a false warning.

In the context of communication settings, as expected, packet error rates do not significantly impact the likelihood of unnecessary warnings. However, increasing the sensor and estimation inaccuracy rates leads to an increase in the likelihood of unnecessary warnings. As discussed at the end of Section 4, this result may be attributed to the unbiasedness of the simulated sensor and estimation errors combined with our protocol to issue a warning at the first instance of predicted time-to-collision falling below one second. To reduce such incidence of unnecessary warnings, Haas and Hu [157] built in logic to their collision warning model to only issue a warning to the driver if the vehicle predicts a collision two consecutive times. However, given the low incidence rate of false warnings (less than 15 percent at the highest noise setting in our simulations) and that the warnings occurred for scenarios that were near collisions, the issue of false warnings does not appear to be a severe concern for DSRC enabled collision warning systems in rural overtaking settings. Of course, to the extent that sensor and estimation errors in reality might be biased toward being conservative or optimistic, the predictions may also be biased in the same manner.

11.6 Comparing V2V and V2I Communication

In this final section we perform a brief comparison between the use of V2V and V2I communication for a CW/CA system. A hypothetical rural-road V2I infrastructure was developed as a potential alternative to V2V warnings and 4,690 overtaking maneuvers were simulated. This infrastructure consists of periodic radar road-side units (RSUs) positioned to the side of a road, facing down that road. Each RSU can detect vehicles up to 200 meters away, and can transmit the vehicles' presence up to 600 meters. While this type of system could potentially be costly, especially if many radar units are required, it offers several advantages over V2V methods. Most notably, it makes up for DSRC's potentially inadequate range by adding an additional 200 meters of visibility. There is also the potential for each RSU to relay a message from the others, enabling vehicle information to spread much farther.

Table 11.3 shows the comparison of the CW/CA capabilities of this infrastructure system to that of the previously studied V2V system. The sensor and estimation accuracies were set to one-half of their worst-case accuracy, and packet loss was set to zero for this result. If RSUs are spaced at two per kilometer, they can potentially warn of far more collision-causing maneuvers than V2V can. These RSUs also give false warnings at a higher rate than V2V, but that is entirely due to the fact that they can detect more vehicles in the first place, rather than some fundamental difference in their accuracy or predictive algorithm. These results show that V2I may have an important role in complementing V2V for safe collision warning systems. However, the analysis performed herein is not conclusive and cannot guarantee that V2I is the most suitable solution for rural roads.

Table 11.3: Simulations results for overtaking CW/CA with V2I

Method	3,593 collisions	1,097 safe maneuvers
V2V—600m range	70.8 % warned	10.3% false warning
V2I with 200m spacing	96.9% warned	29.4% false warning
V2I with 500m spacing	95.6% warned	26.2% false warning

Chapter 12. Conclusion

Historically, the focus of highway safety has been geared toward implementing passive safety systems (such as airbags and road barriers) that attempt to reduce the severity of crash outcomes. With the advancement of technology, however, efforts have expanded to design advanced driver assistance systems. In this report, we provide four different approaches to utilizing communications and radar sensors to enhance CW/CA systems. The overall goal of the trajectory prediction research was to quantify the value of different types of information in predicting a vehicle's future motion, and how these predictions translate into more effective CW/CA systems. Position and basic dynamic information, such as speed, direction, and acceleration, can contribute to accurate prediction of a vehicle's position for up to 3 seconds. A 3-second time gap is enough for many CW/CA systems, and this information is easily conveyed with V2V communications and/or gathered through radar. For further prediction, information such as road maps and recorded vehicle trajectories can improve accuracy, but our current methods require detailed information for every location.

CW/CA system design for non-motorized users was investigated by coupling a radar/LIDAR sensor with the recent, state-of-the-art deep-learning approaches of computer vision. Pedestrian detection was achieved, while cyclist detection was less successful—partially due to the lack of example data for cyclists. The possibility of an onboard detection system with deep learning was also investigated. Simulations of intersection maneuvers verified that both communication and radar have strengths and weaknesses that result in different CW/CA performance, in terms of how accurately collision situations can be detected in advance. Sensor fusion of the two information sources achieved the best performance. In addition, the different properties of a vehicular sensor were prioritized, with coverage area being the most important, accuracy being next, and frequency of information last. The COIN simulations showed that automated systems are capable of learning simple driving precepts simply by avoiding collision situations. The model is being continually refined and the goal is to achieve a rule-free and crash-free navigation in a crowded intersection using only information that is readily available from DSRC communication. This work is an unprecedented experiment in the design of futuristic transportation systems.

In this report we also explained why current technologies for vehicular communications, such as DSRC or 4G cellular systems, would be insufficient for future connected vehicles that wish to share raw sensor data on a large scale. To exchange large amounts of data from automotive sensors (e.g., cameras and LIDARs), we proposed the use of mmWave networks to exploit a large amount of bandwidth and enable Gbps data rates for future vehicular communications. Since automotive radars are already using mmWave spectrum, we developed a framework for joint mmWave radar and communication systems that functions within a single device by using the same waveforms. We showed that it is possible to simultaneously achieve desired automotive radar performance and data rates to exchange raw sensor data using a joint mmWave system. We also verified the feasibility of joint radar and communication systems using low frequency WLAN systems.

To optimize the performance of joint radar and communication systems, we first identified the three key scenarios where automotive radar and communication systems would be mostly needed. Then we studied performance trade-off between radar and communication functions of the joint systems. Using the trade-off characteristic of joint systems derived in this report, joint

mmWave radar and communication systems can adaptively optimize radar and communication performance depending on traffic scenarios.

We also investigated mmWave antenna and transceiver structures and identified an optimal beamwidth that maximizes the performance of mmWave systems, which will guide future antenna design. We proposed a beamforming technique to efficiently adapt beamwidths to maximize the performance of joint systems. We further investigated mmWave channel acquisition techniques, which exploit low frequency channel information. The proposed techniques will enable sufficient flexibilities on designing mmWave joint radar and communication systems. Since security is critical in automotive environments, we thoroughly investigated security issues in automotive radar and communication systems. We then proposed potential ways to improve security in mmWave systems exploiting a large number of antennas.

There is still a significant amount of work to in this field of automotive communication technology research. In order to mitigate high velocities in vehicular settings, we need to develop efficient protocols to establish mmWave communication links with minimal overhead. We believe it is possible to exploit side information from auxiliary automotive sensors and leverage this information in the implementation of efficient communication protocols. For example, automotive radars can detect possible vehicles for communication, and automotive cameras can tell where to transmit mmWave beams to establish communication links. The use of out-of-band sensing technologies to support communication links is a novel concept in the wireless communication industry, and requires a substantial amount of work. In addition, prototyping will also be necessary to demonstrate the true benefits of the proposed techniques.

In the final task of this research project, we used simulations to assess the performance of a CW/CA system that uses DSRC to assist overtaking maneuvers. The overtaking assistant predicted collisions successfully for 70% of the simulated collisions and gave false CWs for less than 3% of simulated safe maneuvers. A descriptive analysis followed by a multivariate analysis (using binary discrete outcome models) of the simulated data revealed that the majority of collisions that could not be detected were due to the passing and lead vehicles being out of communication range (2000 ft) when the passing vehicle started the overtaking maneuver (at least for the communication power settings used in the simulation). These results suggest that a promising way forward to enhance the effectiveness of DSRC devices for improving the safety of rural highway overtaking maneuvers is by increasing their power settings to broaden the communication range. Another notable result was that packet errors up to a 50% error rate did not have a significant influence on the ability to detect collisions. This result points to how the communication requirements of rural road overtaking scenarios might differ from those of urban intersection scenarios with large traffic volumes where decreasing latency (or packet errors) and relieving communication channel congestion might be a critical need. While still a factor, channel congestion will not have the same magnitude in rural settings as in urban settings. Furthermore, the rural road overtaking maneuver is very deliberate and allows a large span of time in which communication can occur. However, even in rural road settings, latency may be a key factor for other safety applications such as forward CW or emergency brake warning. Sensor error and estimation inaccuracies were found to increase the rate of false warnings more than that of undetected collisions. However, since the incidence of false alarms was small and a majority of them occurred for scenarios that were near collisions, the issue of false alarms does not appear to be a major concern in the case. It is important to note, however, that any systematic biases in sensor and estimation errors, or systematic errors in the trajectory prediction method, may increase the incidence of false alarms or undetected collisions in ways not covered by this simulation.

The data collection effort performed in the project was very limited but enough to show the importance of merging multiple technologies for sensing the traffic environment in cities. The blind intersection scenario, the situations in which the GPS positioning showed the vehicle in the wrong lane, and the interaction with pedestrians are some examples of situations in which we could use the data to visualize the importance of having multiple technologies backing each other. The difficulty of identifying the causes of errors in the information produced by the sensors and transmitters pointed to the necessity of more controlled experiments. One possible approach could be the use of efficient experimental design to allow the identification of sets of environmental and situational variables that adversely affect the functioning of the system. The identification of these special circumstances can also allow specific programming to enhance the performance of the algorithms and the overall system. Situations in which the radar component outperforms the communication system (and vice versa) can be identified and the algorithms can be adapted in order to give more weight to the information provided by the best sensor in each occasion.

The research described in this report is a solid first step towards the development of connected and even autonomous vehicles. We developed important frameworks that need to be prototyped and further enhanced. Overall, our effort in this report makes clear the huge value of being able to rely on multiple information sources for CW/CA. The different strengths of each in-vehicle device, from cameras to radar to V2X networks, mean that each device plays the main role in at least one important CA task.

References

- [1] National Highway Traffic Safety Administration - NHTSA (2013). Traffic Safety Fact Sheet, 2012. Publication DOT HS 811392, National Center for Statistics and Analysis, National Highway Traffic Safety Administration, Available at: <http://www-nrd.nhtsa.dot.gov/Pubs/811856.pdf>
- [2] Texas Department of Transportation (2013). Texas Motor Vehicle Traffic Crash Highlights, Calendar Year 2012, Available at http://ftp.dot.state.tx.us/pub/txdot-info/trf/crash_statistics/2012/01_2012.pdf
- [3] H. H. Meinel and J. Dickmann, “Automotive radar: From its origins to future directions,” *Microwave Journal*, vol. 56, no. 9, pp. 24–40, Sep. 2013.
- [4] B. Fleming, “New automotive electronics technologies (automotive electronics),” *IEEE Vehicular Technology Magazine*, vol. 7, no. 4, pp. 4–12, Dec. 2012.
- [5] N. Lu et al., “Connected vehicles: Solutions and challenges,” *IEEE Internet of Things Journal*, vol. 1, no. 4, pp. 289–299, Aug. 2014.
- [6] Velodyne, <http://velodynelidar.com>.
- [7] United States Department of Transportation, “Connected vehicle research in the United States,” Aug. 2015. [Online]. Available: http://www.its.dot.gov/connected_vehicle/connected_vehicle_research.htm
- [8] J. Harding et al., “Vehicle-to-vehicle communications: Readiness of V2V technology for application (report no. DOT HS 812 014),” National Highway Traffic Safety Administration (NHTSA), Washington, DC, 2014.
- [9] J. B. Kenney, “Dedicated short-range communications (DSRC) standards in the United States,” *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1162– 1182, Jul. 2011.
- [10] G. Bansal, “Potential 5G communications contributions to automated vehicle applications,” Texas Wireless Summit, Oct. 2015.
- [11] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band, IEEE 802.11ad Std., 2012.
- [12] Federal Communications Commission (FCC), “FCC 15-138,” 2015.
- [13] B. Cronin, “Vehicle Based Data and Availability,” USDOT ITS Program Advisory Committee, 2012.
- [14] A. Jain, H. Koppula, B. Raghavan, S. Soh, A. Saxena, “Car that Knows Before You Do: Anticipating Maneuvers via Learning Temporal Driving Models,” ICCV 2015.

- [15] NGSIM Homepage. FHWA. <http://ngsim.fhwa.dot.gov>.
- [16] Transportation Research Board of the National Academy of Sciences, "The 2nd Strategic Highway Research Program Naturalistic Driving Study Dataset," Available from the SHRP 2 NDS InSight Data Dissemination web site: <https://insight.shrp2nds.us>, 2013.
- [17] R. Rajamani, "Vehicle dynamics and control," Springer Science and Business Media, 2011.
- [18] Y. Bar-Shalom, X. R. Li, T. Kirubarajan, "Estimation with Applications to Tracking and Navigation: theory algorithms and software," John Wiley and Sons, 2004.
- [19] E. Wang, L. Tian, Y. Wang, "Vehicle Collision Warning System and Collision Detection Algorithm Based on Vehicle Infrastructure Integration," The Seventh Advanced Transportation Forum of China, 2015.
- [20] A. Ward, N. Worrall, "Vehicle Collision Probability Calculation for General Traffic Scenarios under Uncertainty," IEEE Intelligent Vehicles Symposium, 2014.
- [21] E. A. Wan, R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," Adaptive Systems for Signal Processing, Communications, and Control Symposium, 2000.
- [22] D. Petrich, T. Dang, D. Kasper, G. Breuel, and C. Stiller, "Map-based long term motion prediction for vehicles in traffic environments," Intelligent Transportation Systems, 2013.
- [23] J. J. Olstam and A. Tapani, "Comparison of Car-following models," No. VTI report 960A, 2004.
- [24] C. Barrios and Y. Motai, "Improving Estimation of Vehicle's Trajectory Using the Latest Global Positioning System with Kalman Filtering," IEEE Transactions on Instrumentation and Measurement, vol 60, 2011.
- [25] L. Aoude, B. D. Luders, R. Joseph, J. P. How, "Probabilistically Safe motion planning to avoid dynamic obstacles with uncertain motion patterns," Autonomous Robots 35(1), 51-76, 2013.
- [26] C. Hermes, C. Wohler, K. Schenk, F. Kummert, "Long-term vehicle motion prediction," IEEE Intelligent Vehicles Symposium, p 652-657, 2009.
- [27] B. Higgs and M. Abbas, "Segmentation and Clustering of Car-Following Behavior: Recognition of Driving Patterns," IEEE Intelligent Transportation Systems, 16(1) 81-90, 2015.
- [28] D. J. Berndt and J. Clifford, "Using Dynamic Time Warping to Find Patterns in Time Series," KDD workshop Vol 10 359-370, 1994.

- [29] G. Weidl, G. Breuel, V. Singhal, "Collision risk prediction and warning at road intersection using an object oriented Bayesian network," *Automotive User Interfaces* 5 270-277, 2013.
- [30] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: An evaluation of the state of the art." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 34.4 (2012): 743-761.
- [31] R. Benenson, M. Omran, J. Hosang, and B. Schiele, "Ten years of pedestrian detection, what have we learned?." *Computer Vision-ECCV 2014 Workshops*. Springer International Publishing, 2014.
- [32] M. Enzweiler, and D.M. Gavrila, "Monocular pedestrian detection: Survey and experiments." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 31.12 (2009): 2179-2195.
- [33] D. Geronimo, A.M. Lopez, A.D. Sappa, and T. Graf, "Survey of pedestrian detection for advanced driver assistance systems." *IEEE Transactions on Pattern Analysis & Machine Intelligence* 7 (2009): 1239-1258.
- [34] H. Cho, P. E. Rybski, and W. Zhang. "Vision-based bicycle detection and tracking using a deformable part model and an EKF algorithm." *Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on*. IEEE, 2010.
- [35] Y. Gu and S. Kamijo. "Recognition and pose estimation of urban road users from on-board camera for collision avoidance." *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*. IEEE, 2014.
- [36] S. Milch, and M. Behrens, "Pedestrian detection with radar and computer vision." [Online]. http://www.smart-microwavesensors.de/Pedestrian_Detection.pdf, 2001
- [37] M. Bertozzi, L. Bombini, P. Cerri, P. Medici, P. C. Antonello and M. Miglietta, "Obstacle detection and classification fusing radar and vision." *Intelligent Vehicles Symposium, 2008 IEEE* (pp. 608-613). June 2008
- [38] T. Wang, N. Zheng, J. Xin, Z. Ma, "Integrating millimeter wave radar with a monocular vision sensor for on-road obstacle detection applications." *Sensors*, 11(9), pp.8992-9008, Sep 2011
- [39] R. Girshick, "Fast r-cnn." *Proceedings of the IEEE International Conference on Computer Vision*. 2015.
- [40] A. Geiger, L. Philip, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite." *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 2012
- [41] A. Geiger, L. Philip, C. Stiller and R. Urtasun, "Vision meets robotics: The KITTI dataset." *The International Journal of Robotics Research* (2013): 0278364913491297

- [42] D. Jia, W. Dong, R. Socher, L.J. Li, K. Li, L. Fei-Fei, "Imagenet: A large-scale hierarchical image database." *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009.
- [43] J. Hosang, M. Omran, R. Benenson, B. Schiele, "Taking a deeper look at pedestrians." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2015.
- [44] R. Girshick, J. Donahue, T. Darrell, J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2014.
- [45] S. Ren, K. He, R. Girshick, J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks." *Advances in Neural Information Processing Systems*. 2015.
- [46] J. R. Uijlings, K. E. van de Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition." *International journal of computer vision* 104.2 (2013): 154-171.
- [47] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: A benchmark." In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on* (pp. 304-311). IEEE. June 2009
- [48] D. Krajzewicz, J. Erdmann, M. Behrisch, L. Bieker, "Recent Development and Applications of SUMO," *International Journal on Advances in Systems and Measurements*, 5(3&4) 128-138, 2012.
- [49] Global Positioning System (GPS) Standard Positioning Service (SPS) Performance Analysis Report (2014). *Federal Aviation Administration, Washington, DC* 410.
- [50] US CFR Title 49 Part 393.82 – Speedometer. Viewed online at <https://www.law.cornell.edu/cfr/text/49/393.82>
- [51] R. Layton and K. Dixon, "Stopping Sight Distance," Discussion Paper for Oregon Department of Transportation, 2012.
- [52] J. Hillenbrand, A. M. Spieker, K. Kroschel, "A Multilevel Collision Mitigation Approach – ITS situation assessment, decision making, and performance tradeoffs," *IEEE Intelligent Transportation Systems*, 7(4) 528-540, 2006.
- [53] Delphi, "Delphi Electronically Scanning Radar (datasheet)", 2009. <http://delphi.com/docs/default-source/old-delphi-files/b87cda8b-468d-4f8e-a7a8-836c370fc2c2-pdf>
- [54] J. B. Kenney, "Dedicated Short-range Communications (DSRC) Standards in the United States," *Proceeding of the IEEE*, 2011.
- [55] Wolpert, David H., and Kagan Tumer. "An introduction to collective intelligence." *arXiv preprint cs/9908014* (1999).

- [56] S. Saponara, M. Greco, E. Ragonese, G. Palmisano, and B. Neri, *Highly Integrated Low Power Radars*. Artech House, 2014.
- [57] J. Wenger, “Automotive radar-status and perspectives,” in *Proceedings of the IEEE Compound Semiconductor Integrated Circuit Symposium*, October 2005, pp. 21–25.
- [58] P. Papadimitratos, A. La Fortelle, K. Evenssen, R. Brignolo, and S. Cosenza, “Vehicular communication systems: Enabling technologies, applications, and future outlook on intelligent transportation,” *IEEE Communications Magazine*, vol. 47, no. 11, pp. 84–95, 2009.
- [59] S. Kumar, S. Gollakota, and D. Katabi, “A cloud-assisted design for autonomous driving,” in *Proceedings of the first edition of the MCC workshop on Mobile Cloud Computing*. ACM, 2012, pp. 41–46.
- [60] K. Zheng, Q. Zheng, H. Yang, L. Zhao, L. Hou, and P. Chatzimisios, “Reliable and efficient autonomous driving: the need for heterogeneous vehicular networks,” *IEEE Communications Magazine*, vol. 53, no. 12, pp. 72–79, 2015.
- [61] J. Choi, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath Jr, “Millimeter wave vehicular communication to support massive automotive sensing,” arXiv preprint arXiv:1602.06456, 2016.
- [62] L. Han and K. Wu, “Joint wireless communication and radar sensing systems—state of the art and future prospects,” *IET Microwaves, Antennas & Propagation*, vol. 7, no. 11, pp. 876–885, 2013.
- [63] P. Kumari, N. Gonzalez-Prelcic, and R. W. Heath Jr, “Investigating the IEEE 802.11ad Standard for Millimeter Wave Automotive Radar,” in *Proceedings of the 82nd IEEE Vehicular Technology Conference*, September 2015.
- [64] L. Reichardt, C. Sturm, F. Grunhaupt, and T. Zwick, “Demonstrating the use of the IEEE 802.11p car-to-car communication standard for automotive radar,” in *Proceedings of the 6th European Conference on Antennas and Propagation (EUCAP)*, 2012, pp. 1576–1580.
- [65] W. Roh, J.-Y. Seol, J. Park, B. Lee, J. Lee, Y. Kim, J. Cho, K. Cheun, and F. Aryanfar, “Millimeter-wave beamforming as an enabling technology for 5G cellular communications: theoretical feasibility and prototype results,” *IEEE Communications Magazine*, vol. 52, no. 2, pp. 106–113, 2014.
- [66] H. Rohling and R. Mende, “OS CFAR performance in a 77 GHz radar sensor for car application,” in *Proceedings of 1996 CIE International Conference of Radar*, 1996, pp. 109–114.
- [67] A. Bazzi, C. Karnfelt, A. Peden, T. Chonavel, P. Galaup, and F. Bodereau, “Estimation techniques and simulation platforms for 77 GHz FMCW ACC radars,” *The European Physical Journal Applied Physics*, vol. 57, no. 01, p. 11001, 2012.

- [68] A. Pezeshki, A. R. Calderbank, W. Moran, and S. D. Howard, "Doppler resilient Golay complementary waveforms," *IEEE Transactions on Information Theory*, vol. 54, no. 9, pp. 4254–4266, 2008.
- [69] C. Karnfelt, A. P'eden, A. Bazzi, G. Shhad'e, M. Abbas, T. Chonavel, and F. Bodereau, "77 GHz ACC radar simulation platform," *Proceedings of the 2009 International Transport System Technology Symposium*, pp. 20–22, 2009.
- [70] M. R. Akdeniz, Y. Liu, M. K. Samimi, S. Sun, S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 6, pp. 1164–1179, 2014.
- [71] A. F. Molisch, F. Tufvesson, J. Karedal, and C. F. Mecklenbr'auker, "A survey on vehicle-to-vehicle propagation channels," *IEEE Wireless Communications*, vol. 16, no. 6, pp. 12–22, 2009.
- [72] W. Menzel and A. Moebius, "Antenna concepts for millimeter-wave automotive radar sensors," *Proceedings of the IEEE*, vol. 100, no. 7, pp. 2372–2379, 2012.
- [73] O. Bar-Ilan and Y. C. Eldar, "Sub-Nyquist radar via Doppler focusing," *IEEE Transactions on Signal Processing*, vol. 62, no. 7, pp. 1796–1811, 2014.
- [74] M. I. Skolnik, *Radar Systems*. McGraw-Hill, NY, 2001.
- [75] W.-C. Liu, T.-C. Wei, Y.-S. Huang, C.-D. Chan, and S.-J. Jou, "All-Digital Synchronization for SC/OFDM Mode of IEEE 802.15. 3c and IEEE 802.11ad," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 62, no. 2, Feb. 2015.
- [76] G. Bo, Z. Changming, J. Depeng, and Z. Lieguang, "Compressed SNR-and-channel estimation for beam tracking in 60-GHz WLAN," *Communications, China*, vol. 12, no. 6, pp. 46–58, 2015.
- [77] J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt, "Millimeter-wave technology for automotive radar sensors in the 77 GHz frequency band," *IEEE Transactions on Microwave Theory and Techniques*, vol. 60, no. 3, pp. 845–860, 2012.
- [78] E. Yeh, R. Daniels, and R. Heath, Jr., "Forward collision vehicular RADAR with IEEE 802.11: feasibility demonstration through measurements," in preparation.
- [79] P. Kumari, J. Choi, N. G. Prelcic, and R. W. Heath Jr., "IEEE 802.11ad V2V-Radar: A Joint Vehicle-to-Vehicle Communication and Automotive Radar System," in preparation.
- [80] L. Han and K. Wu, "Joint wireless communication and radar sensing systems – state of the art and future prospect," *IET Microwaves, Antennas & Propagation*, vol. 7, iss. 11, 2013, pp. 876–885.

- [81] B. J. Donnet and I. D. Longstaff, "Combining MIMO radar with OFDM communications," Proc. Third European Radar Conf., Manchester, U.K., Sep. 2006, pp. 99-106.
- [82] K. Mizui, M. Uchida, and M. Nakagawa, "Vehicle-to-vehicle communication and ranging system using spread spectrum techniques," Proc. IEEE 43rd Vehicular Technology Conf., Secausus, NJ, May 1993, pp. 335-338.
- [83] R. W. Heath Jr. et al., "An overview of signal processing techniques for millimeter wave MIMO systems," IEEE J. Sel. Topics Signal Process., to appear. [Online]. Available: <http://arxiv.org/abs/1512.03007>
- [84] J. Kim and A. F. Molisch, "Enabling gigabit services for IEEE 802.11 ad-capable high-speed train networks," in IEEE Radio and Wireless Symposium (RWS), 2013, pp. 145-147.
- [85] Part 15 Rules for Unlicensed Operation in the 57-64 GHz Band; <http://www.fcc.gov/document/part-15-rules-unlicensed-operation-57-64-ghz-band/>
- [86] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," IEEE Std. 802.11ad, 2012.
- [87] J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt, "Millimeter-wave technology for automotive radar sensors in the 77 GHz frequency band," IEEE Transactions on Microwave Theory and Techniques, vol. 60, no. 3, pp. 845-860, 2012.
- [88] G. D. Durgin and T. S. Rappaport, "Theory of multipath shape factors for small-scale fading wireless channels," IEEE Trans. Antennas Propag., vol. 48, no. 5, pp. 682-693, 2000.
- [89] D. Chizhik, "Slowing the time-fluctuating MIMO channel by beam forming," IEEE Trans. Wireless Commun., vol. 3, no. 5, pp. 1554-1565, 2004.
- [90] T. S. Rappaport, R. W. Heath Jr., R. C. Daniels, and J. N. Murdock, Millimeter Wave Wireless Communications. Pearson, Sep. 2014.
- [91] O. Norklit and R. G. Vaughan, "Angular partitioning to yield equal Doppler contributions," IEEE Trans. Veh. Technol., vol. 48, no. 5, pp. 1437-1442, 1999.
- [92] V. Va, J. Choi, and R. W. Heath Jr., "The Impact of Beamwidth on Temporal Channel Variation in Vehicular Channels and its Implications," submitted to IEEE Transactions on Vehicular Technology, Nov. 2015. [Online]. Available: <http://arxiv.org/abs/1511.02937>.
- [93] M. Rossi, A. Haimovich and Y. C. Eldar, "Spatial Compressive Sensing for MIMO Radar", IEEE Transactions Signal Processing, vol. 62, issue 2, pp. 419-430, Jan. 2014.
- [94] O. El Ayach, S. Rajagopal, S. Abu-Surra, Z. Pi, and R. Heath, "Spatially sparse precoding in millimeter wave MIMO systems," IEEE Trans. on Wireless Commun., vol. 13, no. 3, pp. 1499-1513, Mar. 2014.

- [95] X. Zhang, A. Molisch, and S. Kung, "Variable-phase-shift-based RF baseband codesign for MIMO antenna selection," *IEEE Trans. Signal Process.*, vol. 53, no. 11, pp. 4091-4103, Nov. 2005.
- [96] A. Alkhateeb, O. Ayach, and R. W. Heath Jr, "Channel estimation and hybrid precoding for millimeter wave cellular systems," *IEEE J. on Select. Areas in Commun.*, vol. 8, no. 5, pp. 831-846, Oct. 2014.
- [97] R. Mendez-Rial, C. Rusu, A. Alkhateeb, N. Gonzalez-Prelcic, and R. W. Heath Jr, "Channel estimation and hybrid combining for mmwave: Phase shifters or switches?" in *Info. Theory and App. Workshop*, 2015.
- [98] J. H. Kotecha and A. M. Sayeed, "Transmit signal design for optimal estimation of correlated MIMO channels," *IEEE Trans. Signal Process.*, vol. 52, no. 2, pp. 546-557, 2004.
- [99] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath Jr., "Hybrid precoding for millimeter wave cellular systems with partial channel knowledge," in *Proc. Inf. Theory Appl. (ITA) Wksp*, 2013, pp. 1-5.
- [100] T. Aste, P. Forster, L. F'ety, and S. Mayrargue, "Downlink beamforming avoiding DOA estimation for cellular mobile communications," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 1998, pp. 3313-3316.
- [101] K. Hugl, J. Laurila, and E. Bonek, "Downlink beamforming for frequency division duplex systems," in *Proc. IEEE Glob. Telecom. Conf. (GLOBECOM)*, vol. 4, 1999, pp. 2097-2101.
- [102] Y.-C. Liang and F. P. Chin, "Downlink channel covariance matrix (DCCM) estimation and its applications in wireless DS-CDMA systems," *IEEE J. Sel. Areas Commun.*, vol. 19, no. 2, pp. 222-232, 2001.
- [103] K. Hugl, K. Kalliola, and J. Laurila, "Spatial reciprocity of uplink and downlink radio channels in FDD systems," *COST 273 Technical Document TD(02) 066*, 2002.
- [104] B. K. Chalise, L. Haering, and A. Czylik, "Robust uplink to downlink spatial covariance matrix transformation for downlink beamforming," in *Proc. IEEE Int. Conf. Commun. (ICC)*, vol. 5, 2004, pp. 3010-3014.
- [105] M. Jordan, X. Gong, and G. Ascheid, "Conversion of the spatio-temporal correlation from uplink to downlink in FDD systems," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, 2009, pp. 1-6.
- [106] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, 2015, pp. 2416- 2424.

- [107] B. Neekzad, K. S. Pour, J. Perez, and J. S. Baras, "Comparison of ray tracing simulations and millimeter wave channel sounding measurements," in Proc. IEEE Int. Symp. Pers., Indoor Mobile Radio Commun. (PIMRC), 2007, pp. 1–5.
- [108] "Wireless Insite," <http://www.remcom.com/wireless-insite>.
- [109] T. Trump and B. Ottersten, "Estimation of nominal direction of arrival and angular spread using an array of sensors," *Signal Process.*, vol. 50, no. 1, pp. 57–69, 1996.
- [110] B. Ottersten, P. Stoica, and R. Roy, "Covariance matching estimation techniques for array signal processing applications," *Digital Signal Process.*, vol. 8, no. 3, pp. 185–210, 1998.
- [111] M. Bengtsson and B. Ottersten, "Low-complexity estimators for distributed sources," *IEEE Trans. Signal Process.*, vol. 48, no. 8, pp. 2185–2194, 2000.
- [112] N. Currie and C. Brown, *Principles and Applications of Millimeter-wave Radar*, Artech House, 1987.
- [113] "Millimeter wave radar brings international recognition," <http://www.gtri.gatech.edu/history/innovations/millimeter-wave-radar-brings-international-recognition>, 2011.
- [114] V. Richard, *Millimeter wave radar applications to weapons systems*, USA Ballistic Research Laboratories, 1976.
- [115] R. Chauhan, "A platform for false data injection in frequency modulated continuous wave radar," <http://digitalcommons.usu.edu/cgi/viewcontent.cgi?article=4983&context=etd>, 2014.
- [116] S. Roome, "Digital radio frequency memory," *Electronics & Communication Engineering Journal*, vol. 2, no. 4, pp. 147–153, 1990.
- [117] A. Stove, "Linear FMCW radar techniques," *IEEE Radar and Signal Processing*, vol. 139, pp. 343–350, 1992.
- [118] M. Brooker, "Mutual interference of millimeter-wave radar systems," *IEEE Transactions on Electromagnetic Compatibility*, vol. 49, pp. 170–181, 2007.
- [119] C. Laurendeau and M. Barbeau, "Threats to security in DSRC/WAVE," in Proc. 5th International Conference on Ad-Hoc Networks & Wireless, LNCS 4104, 2006, pp. 266–279.
- [120] A. Serageldin, H. Alturkostani, and A. Krings, "On the reliability of DSRC safety applications: a case of jamming," in *International Conference on Connected Vehicles and Expo*, 2013, pp. 501–506.

- [121] T. Zhang, H. Antunes, and S. Aggarwal, "Defending connected vehicles against malware: challenges and a solution framework," *IEEE Internet of Things Journal*, vol. 1, pp. 10–21, 2014.
- [122] H. Hasbullah, I. Soomro, and J. Ab Manan, "Denial of service (DOS) attack and its possible solutions in VANET," *International Journal of Electrical, Computer, Energetic, Electronic and Communication Engineering*, vol. 4, no. 5, pp. 813–817, 2010.
- [123] E. Hamida, H. Noura, and W. Znaidi, "Security of cooperative intelligent transport systems: standards, threats analysis and cryptographic countermeasures," *Connected Vehicles, V2V Communications, and VANET*, vol. 4, no. 3, pp. 380–423, 2015.
- [124] S. Maddio, A. Cidronali, M. Passafiume, G. Collodi, and G. Manes, "Interference cancellation for the coexistence of 5.8 GHz DSRC and 5.9 GHz ETSI ITS," in *IEEE MTT-S International Conference on Microwaves for Intelligent Mobility*, 2015, pp. 1–4.
- [125] Q. Xu, T. Mak, J. Ko, and R. Sengupta, "Vehicle-to-vehicle safety messaging in DSRC," in *Proceedings of the 1st ACM International Workshop on Vehicular ad hoc Networks*, 2004, pp. 19–28.
- [126] Information Resources Management Association, *Transportation Systems and Engineering: Concepts, Methodologies, Tools, and Applications*, IGI Global, 2015.
- [127] C. Miller and C. Valasek, "A survey of remote automotive attack surfaces," in *DEF CON*, 2014.
- [128] "Hackers remotely kill a jeep on the highway-with me in," <https://www.youtube.com/watch?v=MK0SrxBC1xs>, 2015.
- [129] "Automotive security: best practices: recommendations for security and privacy in the era of the next-generation car," <http://www.mcafee.com/us/resources/white-papers/wp-automotive-security.pdf>, 2015.
- [130] "Hacker disables more than 100 cars remotely," <http://www.wired.com/2010/03/hacker-bricks-cars>, 2010.
- [131] "Hackers can now hitch a ride on car computers," <http://www.latimes.com/business/autos/la-fi-hy-car-hacking-20150914-story.html>, 2015.
- [132] M. Raya, P. Papadimitratos, and J. Hubaux, "Securing vehicular communications," in *IEEE Wireless Commun.*, vol. 13, no. 5, pp. 8-15, Oct. 2006.
- [133] C. Karnfelt, A. Peden, A. Bazzi, G. Shhade, M. Abbas, T. Chonavel, and F. Bodereau, "77 GHz ACC radar simulation platform," *Proceedings of the 2009 International Transport System Technology Symposium*, pp. 20–22, 2009.

- [134] M. E. Eltayeb, J. Choi, T. Y. Al-Naffouri, and R. W. Heath, Jr., "On the Security of Millimeter Wave Vehicular Communication Systems using Random Antenna Subsets," submitted to IEEE Vehicular Technology Conference.
- [135] E. R. Yeh, J. Choi, N. G. Prelcic, C. R. Bhat, and R. W. Heath, Jr., "Security in Automotive Radar and Vehicular Networks," submitted to Microwave Journal.
- [136] Federal Highway Administration (2015). Local and Rural Road Safety Program. Retrieved from http://safety.fhwa.dot.gov/local_rural/, July 30, 2015.
- [137] Olaverri-Monreal, C., Gomes, P., Fernandes, R., Vieira, F., Ferreira, M. (2010). The See-Through System: A VANET-enabled assistant for overtaking maneuvers. In Intelligent Vehicles Symposium (IV), 2010 IEEE (pp. 123-128).
- [138] Sousa Vieira A. S. de, Júnior, J. C., Patel, A., Taghavi, M. (2013). Driver assistance system towards overtaking in vehicular ad hoc networks.
- [139] Hegeman, G., Brookhuis, K., Hoogendoorn, S. (2005). Opportunities of advanced driver assistance systems towards overtaking. *European Journal of Transport and Infrastructure Research EJTIR*, 5(4), 281-296.
- [140] Harwood, D.W., Gilmore, K.D., Richard, R.K., Dunn, M.J., Sun, C. (2008). Passing sight distance criteria. NCHRP Report 605. National Cooperative Highway Research Program, Transportation Research Board, Washington D.C.
- [141] Delphi (2009). Delphi electronically scanning radar (datasheet). <http://delphi.com/docs/default-source/old-delphi-files/b87cda8b-468d-4f8e-a7a8-836c370fc2c2-pdf>.
- [142] Velodyne (2016). HDL-64E high definition real time 3D LIDAR (datasheet). http://velodynelidar.com/docs/datasheet/63-9194_Rev-D_HDL-64E_Data%20Sheet_Web.pdf.
- [143] Layton, R., Dixon, K. (2012). Stopping sight distance. Kiewit Center for Infrastructure and Transportation, Oregon Department of Transportation.
- [144] Brooks, R.M. (2012). Acceleration characteristics of vehicles in rural Pennsylvania. *ANALYSIS*, 1(5), 5-22.
- [145] Kenney, J.B. (2011). Dedicated short-range communications (DSRC) standards in the United States. *Proceedings of the IEEE*, 99(7), 1162-1182.
- [146] AASHTO Green Book (2004). A Policy on Geometric Design of Highways and Street. American Association of State Highway and Transportation Officials, Washington, D.C.
- [147] Rabadi, N.M., Mahmud, S.M. (2007). Performance evaluation of IEEE 802.11a MAC protocol for vehicle intersection collision avoidance system. *Proceedings of 4th IEEE Consumer Communications and Networking Conference*, Las Vegas, NV, January, 54-58.

- [148] Abbas, T., F. Tufvesson, J. Karedal (2012). Measurement based shadow fading model for vehicle-to-vehicle network simulations. arXiv preprint.
- [149] Sepulcre, M., Gozalvez, J., Harri, J., Hartenstein, H. (2011). Contextual communications congestion control for cooperative vehicular networks. *IEEE Transactions on Wireless Communications*, 10(2), 385-389.
- [150] Sepulcre, M., Gozalvez, J. (2011). On the importance of application requirements in cooperative vehicular communications. *Proceedings of 2011 Eighth International Conference on Wireless On-Demand Network Systems and Services (WONS)*, Bardonecchia, Italy, January, 124-131.
- [151] Bansal, G., Kenney, J. B., Rohrs, C. E. (2013). LIMERIC: A linear adaptive message rate algorithm for DSRC congestion control. *IEEE Transactions on Vehicular Technology*, 62(9), 4182-4197.
- [152] National Motorists Association. (2014). State speed limit chart. Retrieved from <https://www.motorists.org/issues/speed-limits/state-chart>, July 30, 2015.
- [153] Global Positioning System (GPS) Standard Positioning Service (SPS) Performance Analysis Report (2014). Federal Aviation Administration, Washington, DC 410.
- [154] Analog Devices Inc. (2009). ADXL345 datasheet. <http://www.analog.com/en/products/mems/mems-accelerometers/adxl345.html#product-overview>
- [155] Polus, A., Livneh, M., Frischer, B. (2000). Evaluation of the passing process on two-lane rural highways. *Transportation Research Record: Journal of the Transportation Research Board*, 1701, 53-60.
- [156] Mocsári, T. (2009). Analysis of the overtaking behaviour of motor vehicle drivers. *Acta Technica Jaurinensis*, 2(1), 97-106.
- [157] Haas, J.J., Hu, Y.C. (2010). Communication requirements for crash avoidance. *Proceedings of the seventh ACM international workshop on VehiculArInterNETworking*, Chicago, IL, September.
- [158] Joerer, S., Segata, M., Bloessl, B., Lo Cigno, R., Sommer, C., Dressler, F. (2014). A vehicular networking perspective on estimating vehicle collision probability at intersections. *IEEE Transactions on Vehicular Technology*, 63(4), 1802-1812.
- [159] Harding, J., Powell, G., R., Yoon, R., Fikentscher, J., Doyle, C., Sade, D., Lukuc, M., Simons, J., Wang, J. (2014, August). Vehicle-to-vehicle communications: Readiness of V2V technology for application. (Report No. DOT HS 812 014). Washington, DC: National Highway Traffic Safety Administration.
- [160] US CFR Title 49 Part 393.82 – Speedometer. Viewed online at <https://www.law.cornell.edu/cfr/text/49/393.82> [50]

Appendix A. Device Properties

These are the assumed properties of certain on-vehicle sensors, as used in Chapter 4.3.

- Global Positioning Service (GPS) – worst case r.m.s. error of 4 meters, as found in the 95th percentile of tests by a US survey [49]. The simulations used normally distributed noise with a standard deviation of 2 meters. This way the 95th percentile of noise has a magnitude of roughly 4 meters.
- Speedometer – the required accuracy is 5 miles per hour [50]. This is surprisingly high for a mechanical sensor, the primary reason being that tires have varying-changing size. Simulations used normally distributed noise with a standard deviation of 2.5 miles per hour.
- Accelerometer – not included in these simulations, but included in the Task 3 overtaking simulations. Their error is usually reported well under 0.25 meters per second squared [142], though numbers have not been reported for vehicle use in particular.
- Delphi ESR [DELPHIESR]
 - Communication range: 174 meters from the long-range radar, 60 meters from the short-range radar
 - Communication field of vision: 10 degrees from the long-range radar, 45 degrees from the short-range radar
 - Positioning error: normally distributed with one meter standard deviation. It was very difficult to find any statements on the accuracy of the radar; the vehicle tests from Task 3 will provide more insight into this. Note that this is the relative position of one vehicle to the other, not the absolute position of the vehicle!
 - Speed error: normally distributed with 1.25 miles-per-hour (roughly 0.5 meters per second) standard deviation. Same as above, a characterization of radar accuracy was not easy to locate. Also same as above, this is speed relative to the reference of the sensing vehicle.
 - Frequency of information: 50 milliseconds. Our simulations only occur in 100-millisecond timesteps, so this is the limit that we could simulate. It was assumed that the increased frequency would have no benefits.
- DSRC communication [DSRC]
 - Communication range: Values vary depending on the implementation, but most studies accept that at least 100 meters will be standard
 - Communication field of vision: complete
 - Positioning and speed errors: same as GPS and speedometer errors. Sending the vehicle's values by communication will add no error.
 - Frequency of information: once per 100 milliseconds. Packet loss may occur in cases of heavy traffic, but is not studied in our simulations.

Appendix B. COIN State Space

This list includes every element of the states that the agents learn on: note that all are discrete variables, and some are heavily rounded to keep the state space small.

- Location = {pre-intersection, entering intersection, within intersection, post intersection}, this specifies the rough position of the car in the intersection.
- Lane={Agent's current lane number depending on road }
- Lane to Destination = {T, F} whether the path that agent is currently in ends up in its desired destination (e.g., the car that wants to go straight should learn not to take the lane which is specifically meant to be used by vehicles which make the left turn at intersection)
- Adjacent Lane is Free = {None, Left, Right, Both}, measures the presence of other vehicles in adjacent lanes
- Agent Speed = {0,10,20,30,40,50,60}(Km/h), Agent speed with resolution of 10 Km/h
- Time to Nearest Vehicle = {[0,1],[2,4],[4,8],[8,]} (seconds), the shortest time-to-collision with any other vehicle, using a constant-velocity model to predict vehicles' trajectories. The model from Chapter 2.2.2 is used to gather this time value.
- Slower Time to Nearest Vehicle = {[0,2],[2,]} (seconds), checks for potential collisions under the assumption that this vehicle moves 10 kph slower than its current speed.
- Faster Time to Nearest Vehicle = {[0,2],[2,]} (seconds), checks for potential collisions under the assumption that this vehicle moves 10 kph faster than its current speed.
- Distance to the car ahead = {[0,2],[2,4],[4,]} (meters), the approximate distance of the car to the car ahead (on the same lane)
- Distance to a crossing car = {[0,2],[2,4],[4,]} (meters), the approximate distance to a car that is blocking this car's way through the intersection
- Intersection Open = {True, False} whether there is a vehicle currently entering or in the intersection with a conflicting path with this vehicle – regardless of whether the two vehicles are expected to collide
- Priority = {True, False} whether this entry road has priority for the intersection at the moment