

3-23-2018

Mitigating the Effects of Cyber Attacks and Human Control in an Autonomous Intersection

Karl C. Bentjen

Follow this and additional works at: <https://scholar.afit.edu/etd>

Part of the [Graphics and Human Computer Interfaces Commons](#), and the [Information Security Commons](#)

Recommended Citation

Bentjen, Karl C., "Mitigating the Effects of Cyber Attacks and Human Control in an Autonomous Intersection" (2018). *Theses and Dissertations*. 1792.
<https://scholar.afit.edu/etd/1792>

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact richard.mansfield@afit.edu.



**MITIGATING THE EFFECTS OF CYBER
ATTACKS AND HUMAN CONTROL IN AN
AUTONOMOUS INTERSECTION**

THESIS

Karl C. Bentjen, Captain, USAF
AFIT-ENG-MS-18-M-008

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENG-MS-18-M-008

MITIGATING THE EFFECTS OF CYBER ATTACKS AND HUMAN CONTROL
IN AN AUTONOMOUS INTERSECTION

THESIS

Presented to the Faculty
Department of Electrical and Computer Engineering
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Computer Engineering

Karl C. Bentjen, BS
Captain, USAF

March 2018

DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED.

AFIT-ENG-MS-18-M-008

MITIGATING THE EFFECTS OF CYBER ATTACKS AND HUMAN CONTROL
IN AN AUTONOMOUS INTERSECTION

Karl C. Bentjen, BS
Captain, USAF

Committee Membership:

Scott R. Graham, PhD
Chair

Scott L. Nykl, PhD
Member

Lt Col Logan O. Mailloux, PhD
Member

Abstract

Widespread use of fully autonomous vehicles is near. However, the desire for a human to maintain control, even if limited, of a vehicle will likely never fully subside. Protocols to safely and efficiently manage reservation-based intersections with a mixture of fully autonomous, semi-autonomous, and non-autonomous vehicles exist such as AIM, SemiAIM, and H-AIM. Missing from these protocols is persistent human control of semi-autonomous vehicles in approaching and navigating autonomous intersections without the use of traditional signals. This thesis offers a proof-of-concept of a reservation-based protocol with necessary extensions required for human control in semi-autonomous vehicles. Desired is a protocol that maintains the benefits in efficiency of a fully autonomous environment, such as AIM, while allowing persistent human control of a vehicle. Proposed are possible feedback mechanisms for human response such as displays detailing intersection arrival time, goal velocity, lane keeping assistance, and other warnings. Also developed is a synthetic environment able to demonstrate cyber attacks, their mitigations, and aid in designing a protocol introducing persistent human control. The AFTR Burner three-dimensional virtual world offers the ability to model this physics based environment in a highly predictable and realistic manner. The reservation-based protocol used in the synthetic environment is first verified and validated against both an established reservation-based protocol, such as AIM, and also use case scenarios to determine if the expected behavior is exhibited. Preliminary observations suggest that persistent human control is a possibility among reservation-based autonomous intersections, but further research must be done to determine its viability.

AFIT-ENG-MS-18-M-008

To my Wife and Daughter

Acknowledgments

I would like to express my appreciation to my research advisor, Dr. Scott Graham, for the support and direction provided throughout the course of this thesis effort. I would also like to thank Dr. Scott Nykl and Lt Col Logan Mailloux, my thesis committee members, for their guidance. Finally, I would like to thank the Center for Cyberspace Research (CCR) laboratory staff for their technical support.

Karl C. Bentjen

Table of Contents

	Page
Abstract	iv
Acknowledgments	vi
List of Figures	x
List of Tables	xii
List of Acronyms	xiv
I. Introduction	1
1.1 Background and Motivation	1
1.2 Problem Statement	2
1.3 Research Objectives	3
1.4 Approach	4
1.5 Organization	5
II. Background	7
2.1 Taxonomy of Vehicle Autonomy	7
2.2 Cyber-Physical Systems	9
2.3 Existing Simulators	9
2.4 Reservation-based Protocols	11
2.5 ECU and CAN Network Cybersecurity	14
2.6 V2X Cybersecurity	15
2.6.1 Availability	16
2.6.2 Authentication	16
2.6.3 Confidentiality	17
2.7 Summary	18
III. Synthetic Environment Development	19
3.1 Objective	19
3.2 Design Decisions and Constraints	20
3.3 Synthetic Environment	22
3.3.1 Simulation Clock	22
3.3.2 Intersection Structure	23
3.3.3 Reservation Grid	25
3.3.4 Reservation Protocol	26
3.3.5 Driver Agent	28
3.3.6 Collision Detection	28
3.3.7 Sensors	30

	Page
3.3.8 Human Controls and Feedback	31
3.4 Assumptions	34
3.5 Measures	35
3.5.1 Independent Variables	35
3.5.2 Dependent Variables	36
3.5.3 Calculation of Dependent Variables	37
3.5.4 Control Variables	38
3.6 Baseline of Synthetic Environment	40
3.6.1 Baseline Objective	40
3.6.2 AIM Simulator Setup	41
3.6.3 Synthetic Environment Setup	41
3.6.4 Variables and Constants	42
3.6.5 Delay Calculations	43
3.7 Scenario 1: Rogue non-Autonomous Vehicle	43
3.7.1 Rogue Vehicle Approach	44
3.7.2 Mitigation	46
3.7.3 Expected Results	48
3.8 Scenario 2: Sybil Attack	49
3.8.1 Sybil Attacker Approach - Serial Reservations	53
3.8.2 Sybil Attacker Approach - Parallel Reservations	54
3.8.3 Mitigations	56
3.8.4 Expected Results	58
3.9 Scenario 3: Squatter Attack	59
3.9.1 Squatter Approach	60
3.9.2 Mitigation	60
3.9.3 Expected Results	61
3.10 Human Control in a Fully Automated Environment	62
IV. Analysis	65
4.1 Verification	65
4.1.1 Baseline Results	66
4.1.2 Baseline Analysis	69
4.2 Validation	70
4.2.1 Rogue non-Autonomous Vehicle	71
4.2.2 Sybil Attack	78
4.2.3 Squatter Attack	86
4.3 Human Controls	91
4.4 Summary	93

	Page
V. Conclusions and Observations	94
5.1 Verification of Baseline Environment	94
5.2 Validation of Proposed Protocol	94
5.3 Introducing Human Controls	95
5.4 Implications of Synthetic Environment	95
5.5 Research Contributions	96
5.6 Future Work	96
Appendix A. Collision Matrix	99
Bibliography	100

List of Figures

Figure		Page
1	Small-scale Gulliver model vehicle	10
2	Screen capture of AIM Simulator	12
3	Example lane numbers for four lane intersection	24
4	Reservation grid at a single moment in time	25
5	High level view of reservation messaging process	27
6	Bounding box of sedan in synthetic environment	30
7	Logitech G920 racing wheel	31
8	BeagleBone Black	32
9	CANtact v1.0 board	32
10	Standard CAN frame	32
11	CAN network diagram	33
12	Single four-lane intersection setup for baseline scenarios	39
13	Reservation grid visualization for $n = 34$ granularity	40
14	Rogue vehicle in center of intersection facing southwest	45
15	Rogue vehicle in southeast quadrant facing north	45
16	Rogue vehicle in center of southeast quadrant facing northwest	46
17	Mitigation in action allowing passage to non-obstructed autonomous vehicles	48
18	Sybil attacker shown on shoulder making false reservation requests	50
19	Legal turn direction options from a single lane with lane numbers	51
20	Parallel reservations from lane 0 to lane 3 made by Sybil attacker	55

Figure	Page
21	Alternating parallel reservations made by Sybil attacker..... 56
22	Average delay observed in Simulation 1 68
23	Average delay observed in Simulation 2 68
24	Average delay observed in Simulation 3 69
25	Queues formed during successful rogue vehicle misbehavior mitigation 72
26	Approaching lanes stopped by rogue vehicle in center of intersection using misbehavior mitigation 73
27	Steady state reached with rogue vehicle in northbound lane with misbehavior mitigation 74
28	Steady state gridlock with rogue vehicle in southeast quadrant with misbehavior mitigation 76
29	Queues formed during successful Sybil attack 79
30	Queues formed as a result of parallel reservation Sybil attack 82
31	Sybil attack using alternating weaved left turns 83
32	Left turning squatter occupies space in intersection 86
33	Queues formed while squatter occupies space in intersection 88
34	Right turning squatter occupies space in intersection 90
35	Screen shot of synthetic environment with human control 92

List of Tables

Table		Page
1	Automation Levels set by SAE	8
2	Lane numbering scheme for intersections	24
3	Bounding volumes for collision detection	29
4	Human controls and feedback devices required for persistent human control	33
5	Excerpt from arrival and departure data log	37
6	Baseline scenarios setup	42
7	Collision matrix for four-lane autonomous intersection	52
8	Human controls and feedback devices required for persistent human control	64
9	Baseline scenario 1 results	66
10	Baseline scenario 2 results	67
11	Baseline scenario 3 results	67
12	Results analysis for all baseline scenarios	70
13	Results of rogue vehicle in center of intersection with mitigation	73
14	Results of rogue vehicle in northbound lane of intersection with mitigation	75
15	Results of rogue vehicle in southeast quadrant of intersection with mitigation	77
16	Results of serial reservations in Sybil attack	80
17	Results of serial reservations in Sybil attack using minimum time mitigation	81
18	Results of parallel reservations in Sybil attack	83
19	Results of parallel reservations in Sybil attack using minimum time mitigation	84

Table		Page
20	Results of parallel reservations in Sybil attack using VIN verification mitigation	85
21	Results of squatter turning left without mitigation	87
22	Results of squatter traveling straight without mitigation	89
23	Results of squatter turning right without mitigation	90

List of Acronyms

- 3D** three-dimensional
- AABB** axis-aligned bounding box
- AIM** Autonomous Intersection Management
- CAN** controller area network
- CCR** Center for Cyberspace Research
- CIP** Critical Infrastructure Protection
- CPS** cyber-physical system
- DARPA** Defense Advanced Research Projects Agency
- DoS** denial of service
- DSRC** Dedicated Short Range Communications
- ECU** electronic control unit
- GPS** Global Positioning System
- H-AIM** Hybrid Autonomous Intersection Management
- HUD** heads-up-display
- ICCWS** International Conference on Cyber Warfare and Security
- IFIP** International Federation for Information Processing
- LiDAR** Light Detection and Ranging
- NHTSA** National Highway Traffic Safety Administration
- OBB** oriented bounding box
- OBD-II** on-board diagnostics
- OBU** on-board unit
- ODE** Open Dynamics Engine
- RFID** radio frequency identification
- RSU** road side unit

SAE Society of Automotive Engineers

SCADA Supervisory Control and Data Acquisition

SemiAIM Semi-Autonomous Intersection Management

V2I vehicle-to-infrastructure

V2V vehicle-to-vehicle

V2X vehicle-to-anything

VCSE Virtual Control System Environments

VIN vehicle identification number

WAVE Wireless Access in Vehicular Environments

WG working group

MITIGATING THE EFFECTS OF CYBER ATTACKS AND HUMAN CONTROL IN AN AUTONOMOUS INTERSECTION

I. Introduction

1.1 Background and Motivation

The time when vehicles communicate with each other and infrastructure is fast approaching. In fact, vehicle-to-vehicle (V2V) communications comes standard on Cadillac's 2017 CTS [1]. Modern infrastructure will aid autonomous vehicles by enabling them to safely and efficiently travel. Vehicles will be able to guide themselves around obstacles, interact with nearby vehicles, and navigate by communicating with infrastructure and each other. Of particular interest are the safety and efficiency implications of managing autonomous intersections. By definition, intersections represent those places where vehicles traveling in different directions must share access to a physical location. In contrast, typical travel lanes, in which vehicles also must share access, have vehicles traveling in the same direction, and with similar velocities.

Reservation-based protocols exist to handle autonomous vehicle traffic, as well as mixtures of semi-autonomous and traditional non-autonomous vehicles. Reservations, analogous to hotel reservations for rooms and guests, are a pairing of times and locations which are used to ensure no two vehicles occupy the same space at the same time. In addition to safety, these protocols promise the benefit of drastically reduced average delay experienced by all vehicles when compared to traditional intersections. However, none of the currently available protocols provide the opportunity for the human to maintain control of a vehicle without traditional traffic signals which un-

avoidably increase average delay. Removing the traditional signals at an autonomous intersection while allowing human control of a semi-autonomous vehicle is the focus of this research.

1.2 Problem Statement

While properly handling the safety and efficiency aspects of managing intersections with autonomous vehicles, V2V communications and reservation-based intersection management protocols introduce aspects that do not exist in traditional intersections. It is paramount to prepare for new attack vectors and failure sources posed by the networks formed among infrastructure and autonomous vehicles. Unlike many cyber attacks today, attacks within this domain may result in property damage, physical injury, and even death. This research uses a synthetic environment to explore the possible misbehaviors of autonomous vehicles in vehicle-to-everything (V2X) communications to avoid costs and personal injury in the real world. Misbehavior includes explicit cyber attacks as well as actions that may not in the real world be malicious in nature.

In addition to exploring the impact of misbehavior, this research also examines the implications of allowing humans to retain control over semi-autonomous vehicles while navigating autonomous intersections without the use of traditional traffic signaling systems. Autonomous intersections communicate with vehicles to schedule arrival times and safe routes rather than using traditional traffic signals or stop signs. Current research in the area of reservation-based protocols has not explored the possibility of humans maintaining control of a semi-autonomous vehicle while navigating an autonomous intersection without the use of the traditional traffic signal.

The ultimate goal of such a persistent human control intersection protocol is two fold: i) allow humans to retain control of a semi-autonomous vehicle using the protocol

and ii) retain as much benefit as possible, in terms of decreased delay, from the use of a reservation-based intersection management protocol for fully autonomous vehicles.

1.3 Research Objectives

It has been shown that attackers are able to gain control of a vehicle via the Internet, given the right circumstances [2]. However, the cybersecurity aspect of this research is focused on misbehavior between vehicles and intersections. The synthetic environment provides the realistic experimental sandbox to visualize the effects of the misbehavior scenarios presented. It also offers a similar environment to safely and quickly test the human ability to safely navigate an autonomous intersection while maintaining control of the vehicle. The research questions are outlined below:

- Does the reservation-based protocol within the synthetic environment guarantee safety by preventing collisions between vehicles?
- Is the total average delay experienced by the autonomous vehicles roughly equivalent to an established reservation-based protocol?
- What scenarios, unique to this domain, can be used to validate the reservation-based protocol within the synthetic environment?
- Are there any effective mitigating strategies to defend against the misbehavior scenarios presented?
- Are the misbehaviors completely mitigated?
- Is it feasible to introduce persistent human control into the reservation-based protocol?
- What feedback mechanisms are necessary to maintain safety while navigating an intersection?

- What precautions need to be taken to ensure the safety of the intersection?

To answer the above research questions, the below objectives are outlined for this thesis:

- Build a synthetic environment capable of hosting autonomous vehicles, semi-autonomous vehicles, and autonomous infrastructure.
- Develop a reservation-based protocol to safely and efficiently manage an autonomous intersection within the synthetic environment.
- Determine scenarios of possible misbehavior within V2X communications with respect to autonomous intersections.
- Develop mitigation strategies for defending against the misbehavior sources determined.
- Introduce persistent human control to the reservation-based protocol within a semi-autonomous vehicle.
- Determine and implement human feedback mechanisms to aid drivers in navigating an autonomous intersection.

1.4 Approach

The approach to answer the research questions mentioned is to make use of a physics based synthetic environment which lends itself to test cyber-physical systems (CPSs) such as autonomous vehicles. The environment hosts the autonomous intersection and its reservation-based protocol as well as fully and semi-autonomous vehicles.

Verification and validation scenarios are designed to establish the developed protocol as roughly equivalent to Autonomous Intersection Management (AIM), an established protocol designed for safely and efficiently managing an autonomous intersection, and that the protocol behaves as expected when misbehavior is introduced [3]. The verification scenarios include trials using the AIM protocol and the synthetic environment's protocol. Once the verification scenarios are executed and the response variables are proven statistically equivalent, then the validation scenarios is executed. The misbehavior scenarios provide validation of the reservation-based protocol within the synthetic environment when expected results are attained.

Finally, once the reservation-based protocol is established as statistically equivalent to AIM and behaves as expected in the misbehavior scenarios, the persistent human controls are introduced. The scenarios executed to test the human controls are notional and do not contain human subjects testing.

1.5 Organization

The organization of this thesis is as follows:

Chapter II details background information and related research. First, it introduces terminology used in this thesis concerning autonomous vehicles and infrastructure, including a taxonomy of autonomy in vehicles. Then, it outlines leading simulation environments for vehicle cybersecurity research as well as current reservation-based autonomous intersection management protocols developed for autonomous and semi-autonomous vehicles.

Chapter III describes the salient details of the synthetic environment and the reservation-based protocol. This chapter highlights the assumptions, constraints, and limitations within the synthetic environment. Also provided are a number of

scenarios for the verification and validation of the synthetic environment along with the persistent human control scenario.

Chapter IV provides the results, analysis, and other observations of the scenarios presented in Chapter III. Among the results and analyses presented are those of the verification and validation scenarios. The chapter discusses the baseline verification scenario results of the reservation-based protocol and shows the expected responses to the verification misbehavior scenarios. Preliminary observations are made in response to the persistent human control scenario within the synthetic environment.

Finally, Chapter V concludes with a summary of all the research presented in this thesis. Future work areas are captured along with other implications and benefits of the synthetic environment.

II. Background

This chapter presents background knowledge and information to provide an understanding of current state-of-practice for cyber-physical system (CPS) simulation environments and reservation-based autonomous intersection management protocols. It begins with a vehicle autonomy taxonomy, followed by an overview of CPSs, existing simulators available for vehicle and CPS research, an introduction to the reservation-based protocol Autonomous Intersection Management (AIM), and a brief overview of electronic control unit (ECU), controller area network (CAN), vehicle-to-vehicle (V2V), and vehicle-to-infrastructure (V2I) cybersecurity. Finally, the chapter closes detailing the need for a real-time three-dimensional (3D) cybersecurity research platform for vehicular CPSs to develop a reservation-based intersection management protocol that incorporates persistent human control.

2.1 Taxonomy of Vehicle Autonomy

The National Highway Traffic Safety Administration (NHTSA) and Society of Automotive Engineers (SAE) provide a taxonomy of the levels of autonomy in vehicles [4]. Listed are five levels of autonomy and those levels are shown with a brief description in Table 1. Levels 4 and 5 describe autonomous vehicles which are capable of driving themselves without human intervention (in a majority of circumstances). According to these descriptions, Levels 4 and 5 may provide the option of a human driver to assume control of the vehicle. Levels 0 through 2 contain the majority of the current mainstream vehicles on the road today. Levels 0 and 1 include traditional vehicles that may come equipped with cruise control, while Level 2 autonomous vehicles attain partial autonomy while a human driver is required to stay fully engaged.

Examples of Level 2 autonomous vehicles (and perhaps Level 3 arguably soon) are Tesla's Models S, 3, and X with its Autopilot feature [5].

Table 1. Automation Levels set by SAE

Level of Automation	Description
Human driver required	
Level 0 - No Automation	Completely non-autonomous; the driver performs all driving tasks.
Level 1 - Driver Assistance	Vehicle is controlled by the driver, but some driving assist features may be included in the vehicle design such as traditional cruise control.
Level 2 - Partial Automation	Vehicle has combined automated functions, like acceleration and steering, but the driver must remain engaged with the driving task and monitor the environment at all times.
Level 3 - Conditional Automation	Driver is a necessity, but is not required to monitor the environment. The driver must be ready to take control of the vehicle at all times upon request.
Human driver <i>not</i> required	
Level 4 - High Automation	The vehicle is capable of performing all driving functions under certain conditions including limitations on locations and environments. The driver may have the option to control the vehicle.
Level 5 - Full Automatin	The vehicle is capable of performing all driving functions under all conditions. The driver may, or may not, have the option to control the vehicle.

Throughout this thesis the main focus is on autonomy Levels 3 - 5 where a human may be in control of a semi-autonomous vehicle. Regardless of the autonomy level, vehicles require the capability to communicate with surrounding vehicles and infrastructure via vehicle-to-everything (V2X) to take advantage of the reservation-based protocols presented in this work.

2.2 Cyber-Physical Systems

CPS is a term describing systems that are characterized by the integration of computers and physical processes. They can receive information from a physical environment (via a sensor) and translate it to a digital representation for some useful computation. Also, they can produce a physical effect (e.g. engaging brakes) from some computation, possibly from data received by another CPS. CPSs are often resource constrained, utilize embedded systems, and have a limited network bandwidth. They are also typically networked in groups to produce the desired effects [6]. Naturally, ECUs on a vehicle's CAN architecture are CPSs. For instance, a particular ECU may estimate road speed and provide the estimate to a central ECU to determine whether or not to illuminate a warning to fasten a seatbelt.

With the increased use of CPSs, to include the emergence of V2V and V2I, it is vital to be able to conduct cybersecurity tests on vehicles. Checkoway experimentally and systematically studied the externally-facing attack surface of a vehicle. Through this investigation it was shown that a vehicle can be compromised without direct access to that vehicle: an attacker could compromise a vehicle via the infotainment system using a malicious compact disc or even a PassThru device used by maintenance personnel [7].

These attacks present a need for a synthetic environment to safely conduct cybersecurity research on given vehicle components in custom scenarios. A real-time 3D simulator with customizable hardware and/or software ECUs could meet this need.

2.3 Existing Simulators

There are a number of vehicle and CPS simulators that exist today. This section gives a brief outline of some of the existing simulators, their purposes, and how they fall short of the purpose of this research.

Sandia National Laboratories has developed the Virtual Control System Environments (VCSE) simulator to conduct cybersecurity research on CPSs. However, this simulated environment focuses exclusively on CPSs in the realm of Supervisory Control and Data Acquisition (SCADA) systems (e.g. oil refineries and electrical grid systems). The simulator does, however, offer real-time visualization of these SCADA systems from the virtual controls to the 3D models [8].

Gulliver, developed at the University of Gothenburg, is a testbed for vehicular systems with a focus on cyber-physical aspects of designs. The testbed uses small-scale prototypes of cars rather than a virtual world simulation. One of Gulliver's claims to relevance is its low cost of implementation, which is at least one or two orders of magnitude less than a full-scale vehicular prototyping unit [9]. However, for purposes of cyberspace research, the small-scale model, which can cost up to \$2,500 each, would not be cost effective. Of interest in this research is the real time effects of cyber stimuli (malicious or otherwise) on a vehicular system.




Figure 1. Small-scale Gulliver model vehicle [10]

Figueiredo outlines a number of simulators that were developed for use in the Defense Advanced Research Projects Agency (DARPA) Urban Challenge. The DARPA competition simulators were used to test each team's autonomous vehicle prior to the competition. The simulators were developed in-house at their respective universities. Figueiredo also surveys the use of robotic simulators as well as game engines such as Epic Games' open-source Unreal Engine [11]. Most notably, these simulators do not implement cybersecurity testing. They were designed to test autonomous vehicles using sensors such as Global Positioning System (GPS) and cameras to navigate in urban traffic environments. These simulators also lack V2V/V2I communication devices, or the implementation of individual ECU simulation.

The 3D virtual world game engine known as AFTR Burner offers the ability to build the synthetic environment needed. Nykl *et. al.* developed the game engine originally as an educational teaching tool, however, it has the ability to be quickly adapted to an array of custom purposes such as the development of Automated Aerial Refueling [12], [13]. The AFTR Burner engine uses the open source OpenGL graphics library to render the realistic 3D graphics and is compatible with the open source physics engine known as Open Dynamics Engine [14]. In addition to offering realistic graphics and physics, the AFTR Burner engine is capable of integrating with sensors and CAN traffic. This feature is ideal for introducing human controls such as a steering wheel and an accelerator.

2.4 Reservation-based Protocols

Intersections pose a unique challenge to autonomous vehicles. Open roads ideally do not contain pedestrians, cyclists, or large variances in speed and direction between vehicles. On the other hand, intersections have vehicles crossing paths and changing directions constantly. The challenge of managing autonomous vehicles at an inter-

section is exactly what Dresner and Stone designed AIM to do [3]. Figure 2 shows a screen capture from the AIM Simulator designed to test the protocol.




Figure 2. Screen capture of AIM Simulator

AIM is a protocol that autonomous vehicles, in theory, could use to communicate with intersections to safely and efficiently schedule passage. AIM uses V2I communications to operate. Each vehicle uses a driver agent to operate the vehicle and request clearance, known as a *reservation*, to proceed through an intersection. To make this request, vehicle sends a message to the intersection manager. The intersection manager executes an algorithm to determine if the requesting vehicle may proceed as requested with a reservation confirmation message. Otherwise, the intersection manager may reject the requesting vehicle's request with a rejection message. Along with this reject message is a suggested alternate later reservation. If a vehicle's request is rejected, then it must either accelerate or decelerate and then make another request. AIM, however, is designed with the assumption that all vehicles are fully autonomous.

Using a protocol such as AIM, Sharon and Stone note that improvements to intersection delay are not seen until the expected penetration of fully autonomous vehicles with V2X capabilities exceeds 90% [15]. Practically speaking, for successful evolution from coarse time sharing (traditional traffic signals) to fine time sharing (AIM), a transitional period where all levels of autonomous vehicles are present simultaneously is necessary, for which AIM is not designed.

Au, Zhang, and Stone developed Semi-Autonomous Intersection Management (SemiAIM) to handle this transitional period [16]. This extension of AIM incorporates semi-autonomous vehicles. In this protocol, the human drivers relinquish control of their vehicles prior to entering the intersection. Also, the protocol requires the use of traffic signals. Semi-autonomous vehicles failing to get a confirmed reservation must come to a stop and treat the intersection as a traditional traffic signal intersection. The signals are also used for non-autonomous vehicles, which are also allowed in this protocol.

Sharon and Stone developed Hybrid Autonomous Intersection Management (H-AIM) seeking to improve the efficiency of AIM when there is a low concentration of autonomous and semi-autonomous vehicles [15]. This extension to the protocol increases the efficiency of AIM when less than 90% of vehicles on the road are fully autonomous. The key differences in this protocol are: 1) the ability of the intersection to detect incoming non-autonomous vehicles and 2) allowing autonomous vehicles to receive reservations that do not conflict with the incoming non-autonomous vehicles' possible paths through the intersection. This protocol also requires traffic signals for human driven vehicles, which do not have V2X communication capabilities.

Still missing from each of these protocols is an option for a human-driven vehicle with V2X capabilities to request and receive a reservation, yet allow the human to maintain persistent control over the vehicle. Vehicles at a Levels 2 - 5 automation

level can allow a human to control steering, velocity, or both while navigating through an autonomous intersection without traffic signals. Building and testing this type of protocol necessitates a realistic synthetic environment since the human is reacting to simulated stimuli.

In such an environment, cybersecurity aspects of the developed protocols can also be studied. There are a number of attack scenarios worth studying when considering smart intersections using reservation-based protocols. For instance, malicious actors could spoof reservation request messages, reservation confirmation messages, or reservation rejection messages. A denial of service (DoS) attack against the intersection manager could also be studied in such an environment.

2.5 ECU and CAN Network Cybersecurity

The complexity of electronics and the number of ECUs have increased in the last ten years. ECUs are embedded systems that control one or more subsystems of a vehicle. They can control anything from the power windows to cruise control systems and braking systems. One type of ECU is the engine control module which can control things such as air intake, fuel injection, and monitor temperatures and other sensor information [17]. Compromises in security in these CPSs could modify the vehicle's behavior and potentially lead to crashes, resulting in serious injury or death.

Multiple ECUs present on a vehicle are networked together with the use of one or more CANs. This architecture was originally developed by Bosch as a multi-master message broadcast system. Using this broadcast network, many short messages are sent from one transmitter to all other transmitters on the network [18]. Inherent to the CAN bus architecture is a lack of confidentiality, integrity, availability, authentication, and non-repudiation. Messages sent across the CAN can, and are, read by all other nodes residing on the bus. There is no notion of sender and receiver addresses or

digital signatures [19]. It has been shown by Koscher *et. al.* that, given the right accesses, an attacker can directly manipulate safety critical ECUs to create unsafe conditions. Furthermore, Koscher analyzed the security implications of an attacker with prior access to a vehicle's internal network, not the methods used to gain access. With the future introduction and use of V2X within vehicles, the attack surface of a vehicle will be further broadened [2] beyond the currently vulnerable indirect attack surface presented by Checkoway [7].

Woo *et. al.* proved that it is possible to attack a vehicle's CAN bus through the use of a malicious diagnostics application on a smartphone. The attack does not require direct access to a vehicle. It does, however, require a victim to download a seemingly legitimate, but actually malicious, diagnostic application on his or her smartphone. Once the victim's smartphone is connected to the target vehicle via the on-board diagnostics (OBD-II) port, the attacker then has the ability to execute the attack. The smartphone's cellular connection and the attacker's server are used to carry out the attack which could result in the engine stopping or some other safety feature failing [20].

2.6 V2X Cybersecurity

V2X is the communication between a given vehicle and its surroundings. The surroundings could be other vehicles or infrastructure such as intersections. Its purpose is aimed at improving driver and passenger safety as well as improving traffic management. Dedicated Short Range Communications (DSRC) and Wireless Access in Vehicular Environments (WAVE) are designed to provide an architecture for nodes within a vehicle network (V2X). These V2X communications utilize the 5.9 GHz band with 75 MHz of bandwidth. The nodes within the architecture can be station-

ary nodes known as road side units (RSUs) or nodes located on vehicles known as on-board units (OBUs).

In [21], Laurendeau provides an outline of the threats to V2X's availability, authentication, and confidentiality.

2.6.1 Availability.

Threats to V2X's availability include *DoS*, *malware*, and *spamming*.

- *DoS* - DoS attacks cause a network to be unavailable to its legitimate users by means of flooding or jamming.
- *Malware* - Malware is the introduction of malicious software such as viruses and worms. Addition of malware into a vehicular network could be carried out by a rogue insider when pushing out updates to software or firmware to vehicle ECUs.
- *Spamming* - Spamming by marketers is more of a nuisance, but it can cause troubles in a vehicular network due to increased network latency.

2.6.2 Authentication.

Threats to V2X's authentication include masquerading, black hole, replay, GPS spoofing, broadcast tampering, and transaction tampering.

- *Masquerading* - A masquerading attack is one where a rogue node (RSU or OBU) poses as a legitimate node.
- *Black Hole* - A black hole is an in-network node which does not pass along messages as it is intended to do. This attack cause legitimate nodes to not receive critical update messages, therefore possibly paving the way for further attacks.

- *Replay* - Replay attacks are not considered a strong attack vector due to WAVE's ability to defend against it using a cache of recently received messages. Otherwise, an attacker would be able to eavesdrop a message from a legitimate node and reuse it later in time for malicious purposes.
- *GPS Spoofing* - GPS spoofing consists of an attacker being able to simulate GPS signals strong enough to be received by the GPS receiver rather than legitimate GPS signals.
- *Broadcast Tampering* - Broadcast tampering involves an attacker broadcasting false messages. Since messages are signed, this attack would get the attacker's node quickly placed on a black-list. In addition, the difficulty is high for this type of attack.
- *Transaction tampering* - Transaction tampering is a man-in-the-middle attack where an attacker manipulates multiple messages between two other nodes. This attack is technically difficult due to message encryption.

2.6.3 Confidentiality.

Threats to V2X's confidentiality include outsider eavesdropping, insider eavesdropping, and location tracking.

- *Outsider Eavesdropping* - Outsider eavesdropping involves an attacker listening to messages not destined to the attacker. This does not include broadcast messages since everyone is intended to receive those types of messages. This attack is exceptionally difficult.
- *Insider Eavesdropping* - Insider eavesdropping involves a node that legitimately has information collection authority, however, collects the information outside of any agreed windows of time.

- *Location Tracking* - Location tracking is when an attacker tracks the location of a vehicle over time and builds a potentially damaging profile of the victim.

2.7 Summary

Currently there are no efforts to incorporate persistent human control into a reservation-based intersection protocol. A 3D virtual world synthetic environment is ideal for developing such a protocol. The environment lends itself to also test many aspects of cybersecurity unique to CPSs relating to autonomous vehicles. A virtual world environment with the capability to incorporate hardware or software ECUs with a CAN architecture, autonomous vehicles, human drivers, and V2X communications would greatly add realism to vehicular cyber research. Consider an experiment where high speeds are required prior to an exploitation of a vulnerability. Safety concerns would likely preclude the experiment from being run in real life, even on a closed test range with experienced drivers. Such a virtualized environment could save thousands of dollars in the purchasing of vehicles, or scaled models, solely to conduct research when the physical vehicle may be irrelevant. Also, during a designed experiment with some given scenario, multiple runs could be conducted with essentially a push of a button.

III. Synthetic Environment Development

3.1 Objective

To begin this research effort, a synthetic environment is required. The environment will offer a the ability to realistically model autonomous vehicles and intersections. Incorporating human control is also enabled by the synthetic environment since it keeps costs and safety risks low. The synthetic environment must be able to host cyber attacks against cyber-physical systems (CPSs) and aid in implementing the mitigations and incorporate the use of human controls and realistic feedback mechanisms. The environment must offer real-time visualizations of the attacks and allow a high level of control. This is important for repeatability, given any number of control variables. The specific cyber attacks presented in this research are referred to as misbehaviors within the reservation-based protocol used for managing the intersection.

The validation experiments here are aimed to show the impact of the developed mitigation strategies against a number of misbehavior scenarios. Each scenario has several different control factors be detailed and evaluated in the following sections. However, prior to running experiments in the developed synthetic environment, a proper verification baseline needs to be established to demonstrate the fidelity of the synthetic environment's autonomous behavior. It needs to be established that, given some assumptions, the reservation-based protocol is equivalent to that of a protocol such as Autonomous Intersection Management (AIM) [3].

The synthetic environment will use the AFTR Burner three-dimensional (3D) virtual world game engine [12] and Open Dynamics Engine (ODE) to provide realistic physics. Other simulators, discussed in Chapter II, do not provide an environment conducive to the necessary scenarios without significant modifications. The AFTR

Burner 3D virtual world game engine is a viable option for development of the required synthetic environment. Intersections within the synthetic environment utilize a reservation-based intersection management protocol similar to that of the AIM protocol developed by Dresner [3]. AIM has been proven to provide safe and efficient traffic control for fully autonomous vehicles in a fully autonomous scenario. In addition to a realistic synthetic environment, the protocol and its accompanying two-dimensional simulator also lack the ability to insert misbehavior, execute cyber attack scenarios, and demonstrate the use of mitigations in their defense. The AIM Simulator is also not designed to handle semi-autonomy, namely human controls, such as steering and/or accelerator inputs via controller area network (CAN) traffic or otherwise.

Prior to conducting any misbehavior scenarios, a verification baseline comparison between the synthetic environment and the AIM protocol simulator is needed. This baseline shows the validity of the algorithm used in the synthetic environment. To show this comparison, two comparable simulations are executed with the same control variables. The variables held constant between the two simulations are the number of vehicles generated per lane per hour, number of roads in each direction, and the number of lanes per road. If these two simulations show statistically similar response variables then a baseline will be attained. This baseline shows that the algorithm used in the synthetic environment, given the assumptions made, is sufficiently equivalent to AIM and is fit to simulate a single autonomous intersection for the purposes of this research.

3.2 Design Decisions and Constraints

Level 4 and 5 autonomous vehicles are already occupying public roadways in the United States in addition to the eminent semi-autonomous vehicles mentioned in Chapter I. Google claims its autonomous vehicles in the Waymo project have been

tested on approximately 3.5 million miles of streets since its conception in 2009, including tests without human occupants [22]. It is estimated that over 94% of new cars sold will be fully autonomous in the year 2040 [23]. A synthetic environment offers the ability to model autonomous vehicles and infrastructure well before this time.

Objects within the AFTR Burner 3D virtual world engine are referred to as world objects. The world object chosen to represent the autonomous vehicle is a generic four door sedan and the intersection is a standard intersection without traffic signals. The intersections in this research are limited to the same number of lanes inbound and outbound from all directions (i.e. north, south, east, and west). This limitation on the number of lanes leads to simplicity in implementation.

With the goal of conducting experiments on autonomous vehicles and intersections within the synthetic environment, the need arises for a protocol suited to manage the autonomous intersection. The AIM protocol served as the model reservation-based protocol within the synthetic environment.

The following requirements are to be used to build the synthetic environment:

- **Simulation clock** - simulations shall be timed with a clock that accurately tracks time within the simulation independent of wall-clock time.

- **Intersection**

Reservation grid size - the grid size, discussed in section 3.3.3, can be set to any positive integer.

Lanes - the intersection may be assigned any number of incoming lanes from each direction.

Speed limits - a maximum speed must be set for the road prior to the intersection and a maximum speed within the intersection.

Range - the set maximum distance from the intersection a vehicle can be to communicate with it via vehicle-to-everything (V2X) communications.

Stopping Distance - a set distance from the intersection a vehicle must stop without a confirmed reservation; it may only proceed if it obtains a reservation.

- **Autonomous Vehicles**

Path - the autonomous vehicles shall follow a predetermined path during simulations.

Collision Detection - collisions between vehicles (safety violations) are detected using the ODE, discussed in section 3.3.6.

Human Control - an input device shall be used to control steering and acceleration via CAN traffic.

3.3 Synthetic Environment

This section describes the salient features designed within the synthetic environment. These features include the simulation clock, the intersection reservation grid, and the reservation-based protocol. Also discussed are the algorithm used to control the fully autonomous vehicles, assumptions about the sensors used, and finally the proposed human control and feedback mechanisms unique to this research.

3.3.1 Simulation Clock.

To provide a clock to synchronize the messaging protocol and time-space reservations, a simulation clock class was designed. This class utilizes ODE's built in clock. The physics engine updates the physics of each world object on a single "tick". Each tick represents 1/60th of a second in simulation time. The number of ticks per frame

rendered to the screen is controlled and is typically set to a number between 0 and 10. If the number of ticks per frame is set to 6, then each frame rendered to the screen represents 1/10th second of simulation. The physics engine has the ability to be paused, by setting the engine ticks per frame to 0. This gives the researcher the ability to analyze the behavior of the autonomous vehicles and intersection in real time or compressed time depending on frame rate and the number of physics engine ticks per frame.

The simulation clock class used in the synthetic environment receives the simulation time from the physics engine. Each world object is then able to update the synchronized simulation time from this class. To make reservations within the intersection which is discussed later in this section, the spatial domain must be coupled with the temporal domain. Each moment in time is allocated a complete grid of reservation tiles within the intersection. Moments in time within the synthetic environment are 1/100th of a second. For example, if a reservation is for exactly 1.5 seconds, then there exists exactly 150 reservation tile grids. This is discussed in more detail in section 3.3.3.

3.3.2 Intersection Structure.

The autonomous intersection class includes the algorithm used to operate the intersection autonomously and ensure safety of the approaching vehicles while they cross through the intersection. The intersection class maintains the reservation grid class as one of its member variables. The reservation grid class includes all current reservations and contains the logic to determine if reservations may be safely approved, referred to as *confirmed*, or must otherwise be *denied*. The intersection class handles incoming and outgoing vehicle-to-infrastructure (V2I) request and response messages.

Variables that can be set arbitrarily include the intersections location, road side unit (RSU) communication range distance, intersection identification number, lane width, number of lanes, and grid dimensions. Also included in the autonomous intersection class are the functions that provide the ability to mitigate misbehavior, discussed later.

The lane numbering scheme uses non-negative integer values to identify lanes. Each lane number modulus 4 determines its heading. See Table 2 for the numbering scheme. Figure 3 shows example lane numbers for a four lane intersection.

Table 2. Lane numbering scheme for intersections

(Lane Number) mod 4	Direction
0	North
1	East
2	South
3	West




Figure 3. Example lane numbers for four lane intersection

3.3.3 Reservation Grid.

The idea of the reservation grid was introduced by Dresner and Stone to enable the AIM protocol to guarantee no collisions between vehicles [3]. A reservation grid is broken up into $n \times n$ individual square regions. Each square represents an occupied or vacant state for the corresponding location. Each discrete moment in time (1/100th of a second in simulation time) is represented by a single grid of tiles. Figure 4 shows a 34×34 reservation grid at a single moment in time in the synthetic environment. To ensure no collisions occur, no two vehicles may be granted a reservation that would cause any of the same squares to be occupied at the same moment in time.




Figure 4. Reservation grid at a single moment in time

The reservation grid is its own class within the synthetic environment. This class maintains a collection of the legal turns (or paths) that are available within the intersection, the current confirmed reservations, and all the logic to process reservation requests. These reservation requests may be confirmed or denied based on the confirmed reservations. The reservations are granted on a first come, first served basis.

When a vehicle requests a reservation from the intersection, it first retrieves a copy of the path through the intersection. The vehicle then uses the information to generate a reservation request message based on its estimated arrival time and speed in the intersection. See section 3.3.4 for more details.

Finally, and perhaps most importantly, this class maintains the current reservations made by requests received from incoming autonomous vehicles. Each discrete moment in time requires one reservation grid of tiles which represents the area within the intersection. The reservation grid class also processes each request by checking the current state of the previously confirmed reservations. If there are no conflicts, then the reservation may be confirmed. The next section provides more details on this behavior.

3.3.4 Reservation Protocol.

Reservations are established through messages sent between world objects within the synthetic environment. These messages simulate V2I communication between autonomous vehicles and the intersection. This exchange of information is essential to maintain safety of the autonomous vehicles as they pass through the intersection. The number one priority of the reservation protocol is that no collisions occur between vehicles.

When an autonomous vehicle comes into range of an intersection, it begins to formulate a reservation request message. It first receives the path information from the intersection based on its arrival lane number and the desired turn direction. Once the mapping between times and locations within the intersection has been generated by the autonomous vehicle using the path information, it then sends the request to the intersection. The intersection then determines whether or not the reservation can be granted and replies to the vehicle accordingly. If the reservation is made, then

the car is cleared to proceed. However, if the reservation is denied, then the vehicle is required to decelerate and attempt to make another reservation request. Figure 5 shows a high-level flowchart of the messaging protocol.

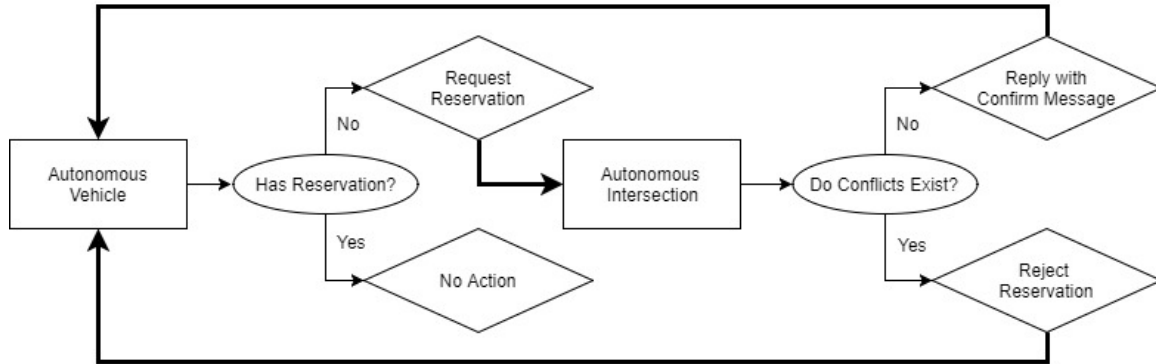


Figure 5. High level view of reservation messaging process

When an intersection receives a message from an autonomous vehicle, it determines if the reservation is in conflict with any reservations that have already been made. It does this by comparing the mapping between times and reservation tile grids that were sent in the message. It compares the reservation tile grid from a given time in the request to the reservation tile grid that already exists. If there are no conflicts found, then the reservations may be made. Once a reservation is made, the intersection maintains it so that subsequent requests do not interfere with it. If there are conflicts found with the requested reservation, then the reservation must be rejected. Otherwise, a collision will likely occur between two autonomously controlled vehicles.

The intersection has the ability to broadcast messages out to all connected vehicles within range. These broadcast messages are used to cancel previously made reservations in the event that cancellation becomes necessary (e.g., in an emergency situation). The cancellation message holds enough information to allow the cancellation of selected lanes and turn directions, rather than simply all reservations.

3.3.5 Driver Agent.

The logic which controls an autonomous vehicle's behavior is the driver agent. The pseudocode for the driver agent used in the AIM Simulator was used as a model for the driver agent in this work [3]. The driver agent controls the velocity of the vehicle based on its location in the synthetic environment, reservation status with the upcoming intersection, and the vehicle's proximity to other vehicles.





The velocity of the vehicle, steering, and messaging functions of the driver agent are executed independently to assist in testing the persistent human control aspect of the protocol. Each can be set to human controlled or autonomously controlled during the human control and autonomous scenarios, respectively.

3.3.6 Collision Detection.

The number one priority for any intersection management protocol is safety. Even traditional intersections focus on safety through the design of traffic lights, stop signs, and traffic laws. Therefore, the ability to detect safety violations is a requirement for the synthetic environment.

The AFTR Burner virtual world uses the ODE physics engine to detect collisions between world objects. Collision detection of a 3D world object requires the physics engine to maintain bounding volumes, a 3D geometry in which to enclose the world object. These volumes are used to detect collisions between world objects. They come in a number of options such as spheres, axis-aligned bounding boxes (AABBs), oriented bounding boxes (OBBs), and convex hulls. These are described in Table 3 in ascending order of precision and descending order of computational speed [24, 25].

Table 3. Bounding volumes for collision detection

3D Geometry	Description	Visualization
Sphere	A spherical geometry of a given radius enveloping a world object. Fast collision detection, low precision bounds.	
AABB	A rectangular prism geometry enveloping the world object aligned with global axes. Changes size when object changes orientation.	
OBB	A rectangular prism geometry enveloping a world object aligned with its local axes. Does not change size when object changes orientation. Used in synthetic environment.	
Convex Hull	A 3D geometry made up of many surfaces that closely approximate the world object's shape without concave surfaces. Slow collision detection, high precision bounds.	

A sphere would over estimate the boundaries above, below, and to the sides of a vehicle. The AABB shape would necessarily change dimensions (become larger) whenever a vehicle turns inside an intersection, which is not ideal since it grossly overestimates the actual volume of the vehicle's world object. The OBB bounding box is ideal for use in the synthetic environment because a typical vehicle, from a top-down view, has a highly rectangular shape and closely matches the world object. The bounding volumes used in the synthetic environment for this research are OBBs, referred to from here on as simply "bounding boxes." A visualization of the bounding box within the synthetic environment is in Figure 6. The red, green, and blue lines represent the x , y , and z axis directions relative to the vehicle's world object, respectively.




Figure 6. Bounding box of sedan in synthetic environment

Simulated collision detection detects safety violations when world objects, or vehicles, come into contact with one another in the 3D virtual world. In addition to this, secondary effects of collisions will also be accounted for in the environment. These secondary effects include collisions between vehicles caused by previous collisions and increased delays due to crashed vehicles. Any cascading effects caused by a single collision could be studied as well using these realistic conditions.

3.3.7 Sensors.

A number of sensors are abstracted in order to introduce mitigations into the synthetic environment. During the rogue human driver scenario, a sensor (perhaps using Light Detection and Ranging (LiDAR) technology) is used to detect and precisely locate the rogue vehicle. Once the sensors obtain the information about the location and size of the vehicle, then the autonomous intersection can take appropriate actions to ensure the safety of the incoming vehicles. Similarly abstracted sensors are used by the driver agent to detect vehicles and their speeds directly in front of the vehicle. Also, the sensors used during the Sybil attack scenario (possibly using radio

frequency identification (RFID) technology) determine the presence of the vehicles that are making reservation requests.

3.3.8 Human Controls and Feedback.

This section describes the implementation of the human controls such as the physical input devices, CAN architecture used, and the feedback mechanisms developed for the human. Real hardware is needed, such as steering wheel and an accelerator, to take input from the human operator and provide as much realism as possible. The feedback mechanisms proposed range from the traditional speedometer to the introduction of a “goal” velocity indicator which is used to communicate to the human how much faster, or slower, he/she must go.

The hardware chosen for human input is the Logitech G920 racing wheel shown in Figure 7. In addition to the this physical hardware, hardware is required to convert data to and from CAN traffic. A BeagleBone Black (Figure 8) simulates of the electronic control units (ECUs) for the accelerator and the steering wheel which converts the output from the wheel and pedals to CAN traffic. The traffic is then placed on the bus from the Beagle Bone via a transceiver. The virtual world receives the CAN traffic via a CANTact v1.0 board which reads the CAN traffic from the bus (see Figure 9).



Figure 7. Logitech G920 racing wheel



Figure 8. BeagleBone Black



Figure 9. CANTact v1.0 board

A quick description of a standard CAN message frame is shown in Figure 10 and the high-level CAN architecture is shown in Figure 11. The CAN frame format used has an 11-bit identifier field and a varying length data field to hold the frame's contents. The message identification field also serves as the priority of the frame within the CAN architecture. Frames in this research, though, are limited to two types of messages: steering and acceleration. It is possible to incorporate other CAN messages into the synthetic environment.



Figure 10. Standard CAN frame




Figure 11. CAN network diagram

Table 8 lists the human feedback mechanisms developed within the synthetic environment. These are in addition to the traditional feedback mechanisms such as a compass, side view mirrors, a rear view mirror, and a speedometer.

Table 4. Human controls and feedback devices required for persistent human control

Item	Description
In-Range Indicator	A device that informs the driver of an autonomous intersection within range.
Request Reservation Button	A device that initiates V2X communication to request a reservation from the intersection.
Denied Reservation Indicator	A device which warns a driver that the requested reservation is denied.
Granted Reservation Indicator	A device which confirms to the driver a reservation is successful and the assigned velocity while within the intersection.
Goal Velocity Indicator	An active device which informs the driver of the required velocity to maintain to meet keep the reservation. May also be used to maintain the correct velocity within the intersection.
Maintain Path Indicator	An active feedback device which informs the driver of the left/right position correctness based on the lane or planned path within the intersection.

3.4 Assumptions

The assumptions made in this research are related to the synthetic environment and its design decisions.

- **Latency** - The messaging protocol is abstracted to function calls between the vehicle and intersection world object classes. Latency, or delay, between the sender and receiver is not modeled in this research and therefore is assumed to be zero.
- **Signal Loss** - Wireless signal loss is not modeled in the synthetic environment.
- **Static Lanes** - No lane changes are permitted during travel within the intersection. Vehicles turn into the respective destination lanes (e.g. left turns from inside lanes terminate in inside lanes).
- **Turning Paths** - For simplicity, turning paths are uniform in nature (quarter-circles) rather than abrupt 90 degree turns.
- **Safety Buffer** - While within the intersection, the occupied tiles take into account an increased size of the vehicle by 25% in each direction for autonomous vehicles.
- **Bounding Box** - Each vehicle's collision detection makes use of a rectangular prism bounding box. This particular type of bounding box simplifies the physics engine's computation complexity to detect collisions while closely modeling the shape of the vehicles. Refer to section 3.3.6.
- **Velocity** - The maximum velocity before and after an intersection is 8 meters per second for fully autonomous vehicles within the synthetic environment. A constant velocity of 8 meters per second is required within the intersection for all directions for fully autonomous vehicles.

- **Single Intersection** - Inside the synthetic environment, a single intersection with two inbound lanes from each cardinal direction is used.
- **Ambient Environment** - There are no obstructions visual, or otherwise, present in the synthetic environment (i.e. the environment is clear with high visibility).
- **Vehicles Only** - No obstacles, except the intersection and vehicles, are present in the synthetic environment (i.e. no pedestrians or wild animals).
- **Reservation Order** - Within a lane, a given vehicle may request a reservation if and only if the vehicle directly in front of it has a reservation. This is accomplished with vehicle-to-vehicle (V2V) communications, and ensures the road ahead will be clear for this vehicle.

3.5 Measures

The synthetic environment built for this research provides the ability to repeat experiments. Along with repeated experiments comes the ability to measure outputs while setting the desired inputs. This section outlines the high-level independent variables, dependent variables, and constant factors that are common to each experiment presented in this research.

3.5.1 Independent Variables.

Independent variables are those that are varied in order to study their influence on the dependent variables. The independent variables common throughout the presented misbehavior scenarios within the synthetic environment are presented below.

- **Vehicle Generation Rate** - The number of vehicles generated per hour per lane. The AIM Simulator randomly generates vehicles at different time inter-

vals, therefore the actual number of vehicles generated per hour per lane is approximately the generation rate.

- **Misbehavior** - Various misbehaviors take place in their respective scenarios. For example, a non-autonomous vehicle may be stopped within the intersection. These misbehaviors are varied in order to study their influence on the dependent variables.

3.5.2 Dependent Variables.

Dependent variables are those that are being studied in response to varying the independent variables. The dependent variables common throughout the presented misbehavior scenarios within the synthetic environment are presented below.

- **Average Total Delay** - This response variable is measured in seconds. It represents the average amount of extra time taken to travel through the autonomous intersection due to traffic.
- **Maximum Total Delay** - This measure is the largest delay experienced by a single vehicle which completes the route through the intersection.
- **Total Throughput** - This response variable is measured in vehicles. It is the number of vehicles that the intersection was able to safely service during the simulation. The throughput is increased by 1 each time a vehicle safely traverses the intersection from any direction. A vehicle is considered serviced when it departs the intersection.
- **Total Collisions** - Collisions are safety violations that are simply physical contacts between two or more vehicles. One or more collisions in any scenario results in the synthetic environment failing the baseline.

3.5.3 Calculation of Dependent Variables.

This section describes the methods used to collect data to calculate the dependent variables after each experiment in the AIM Simulator and the synthetic environment.

3.5.3.1 Output Logs.

The AIM Simulator and the synthetic environment have the ability to produce an arrival and departure data output log in the format of a comma separated value (CSV) file. This log is generated at the end of an experiment, or any time deemed necessary by the experimenter. Each entry in the log contains the vehicle identification number (VIN), a time stamp (s), event type, type of vehicle, starting lane ID, and destination direction. See Table 5 for example contents of this log data. The lane numbers and destinations have to be translated from the lane numbering scheme used in the synthetic environment to those used in the AIM Simulator so that the same parsing script may be used for both environments. Having these calculations done independent of the synthetic environment allows for processing of the output data well after the experiment ends and the synthetic environment is terminated.

Table 5. Excerpt from arrival and departure data log

VIN	Time Stamp	Location Stamp	Vehicle Type	Start Lane ID	Destination Direction
1002	8.32	East Bound Entrance	Sedan	5	South
1002	38.94	South Bound Exit	Sedan	5	South

3.5.3.2 Processing Output Logs.

A separate custom C++ analysis program parses the arrival and departure data output logs to generate the dependent variables for a given experiment. To calculate delay, a base case time has to be established for each lane and turning direction. A base case time is the time needed for a single vehicle to travel from start to finish without traffic. Once the base case time is recorded for each lane and turning direc-

tion, then the output logs for each experiment can be processed to produce the delay calculations. These base case times must be observed in both the AIM Simulator and the synthetic environment. The script calculates information such as total average delay, maximum delay, total generated vehicles, and total throughput.

The total number of entries including the word “Entrance” represents the total number of vehicles generated. The total number of entries including the word “Exit” represents the total number of vehicles serviced by the intersection, or total throughput. The number of collisions is determined manually by monitoring the simulation for crashes between the vehicles.

Microsoft Excel can then be used to run statistical analyses on the output delay values produced by the C++ analysis program.

3.5.4 Control Variables.

Control variables are those that are held constant in order to study the effects of the independent variables on the dependent variables. Below are the control variables common throughout the presented scenarios within the synthetic environment.

- **Number of Roads** - The number of roads is treated as a constant factor and held at 1 for each direction. Each experiment in this research uses a single north-south road and a single east-west road (see Figure 12).
- **Number of Lanes per Road** - Each road has four lanes, two in each direction. This variable is held constant throughout the scenarios.
- **Grid Size** - The grid size is $n = 34$, which means there are $34 \times 34 = 1,156$ individual reservation tiles within the intersection. See Figure 13 for a visualization of the reservation grid. The AIM Simulator defaults to this value, thus the synthetic environment is also set to this value for each of the baseline scenarios.

- **Maximum Velocity** - The maximum velocity is held constant at 8 meters per second before, within, and after the intersection.
- **Runtime** - This control variable is the amount of simulation time used for any one experiment. This is set to 1,800 seconds of simulation time, or half an hour.

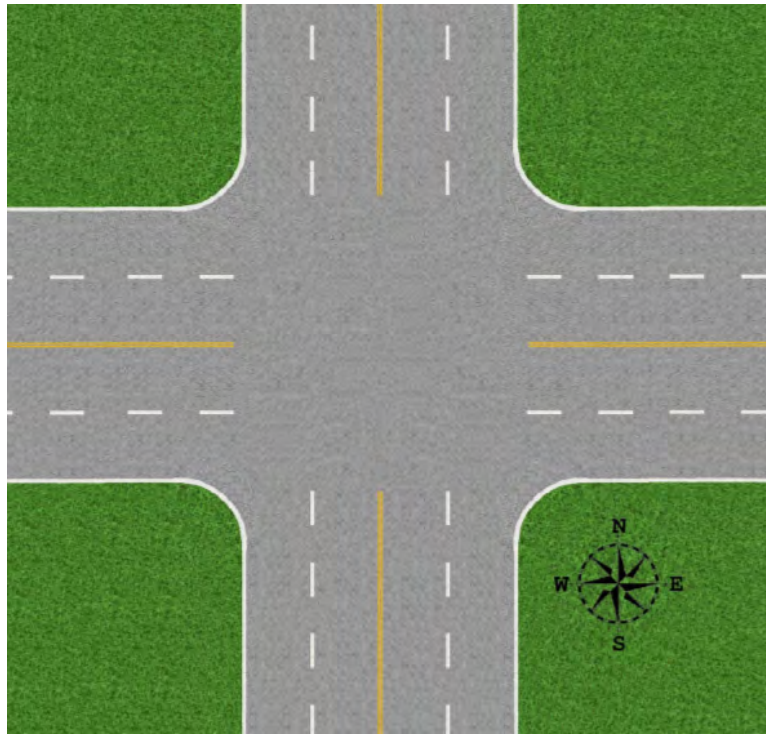


Figure 12. Single four-lane intersection setup for baseline scenarios




Figure 13. Reservation grid visualization for $n = 34$ granularity

3.6 Baseline of Synthetic Environment

This section describes the baseline effort to establish the viability of the autonomous intersection within the synthetic environment. This is also the verification phase of the synthetic environment and its reservation-based protocol.

3.6.1 Baseline Objective.

The baseline is attained when the average total delay experienced by all vehicles traveling through the autonomous intersection in the synthetic environment matches that of the AIM Simulator. The average delay is considered matching if there is no statistical evidence suggesting that the two averages are different. Also the baseline requires zero collisions within the synthetic environment. These results need to show that the algorithm used in the synthetic environment is sufficient to conduct further experiments using the assumptions made of the baseline experiment. These delays are presented in Section 4.1.1.

3.6.2 AIM Simulator Setup.

The baseline experimental setup of the synthetic environment consists of simulations executed between the synthetic environment and the AIM Simulator. Within the AIM Simulator, the number of roads in all directions is set to one. The number of lanes is set to four. That is, two lanes in each direction north, south, east, and west. Finally, the number of vehicles generated per hour per lane is varied. Furthermore, several modifications are made to the AIM source code to closely match that of the synthetic environment, for simplicity. The AIM Simulator has the ability to generate several different types of vehicles, such as coupes, sport utility vehicles, sedans, and vans. For this baseline, the source code is be modified to limit the type of cars to only sedans. Also, the maximum speed of the sedan is set to 8 meters per second, matching that of the synthetic environment. These modifications to the AIM Simulator do not affect the ability of the traffic management protocol to safely and efficiently schedule reservations.

The arrival and departure data output from the AIM Simulator is used in the synthetic environment in order to match the same vehicle generation times. The output data from the AIM Simulator includes VIN, event time stamp, location of event, type of vehicle, starting lane number, and destination road. (See Table 5 above for example data output from the AIM Simulator.) The time stamp from the messages containing entrance data is used to generate vehicles in the synthetic environment along with the starting lane numbers and destination roads to determine their paths.

3.6.3 Synthetic Environment Setup.

The synthetic environment uses the same parameters as the AIM Simulator. There is a single intersection with two inbound and two outbound lanes in each direction.

The distance each vehicle is generated from the intersection closely matches that of the AIM Simulator, approximately 100 meters from the intersection. The synthetic environment parses the AIM output data to match the generation rate and times of vehicles. Each generated car is positioned in the correct lane at the generate distance according to the generate schedule. Once the vehicle passes through the intersection and reaches the same generate distance away from the intersection, it is removed from the synthetic environment. A record is kept at the moment each autonomous vehicle is generated and removed to replicate the AIM Simulator's output data log.

3.6.4 Variables and Constants.

The independent variables and constant factors for the baseline setup are shown in Table 6. Five trial experiments are executed in each scenario with varying generation rates. Dependent variables are calculated at the conclusion of each trial based on arrival and departure data output logs.

Table 6. Baseline scenarios setup

Variable	Scenario 1	Scenario 2	Scenario 3
Generation Rate*	100	200	300
No. North/South Roads	1	1	1
No. East/West Roads	1	1	1
No. Lanes/Road	4	4	4
Grid Size	34	34	34
Speed Limit (m/s)	8	8	8
Runtime (s)	1800	1800	1800
Trials	5	5	5

*Vehicles per hour per lane

3.6.5 Delay Calculations.

The data collected from both the AIM Simulator and the synthetic environment are both parsed to calculate average delay experienced by the autonomous vehicles. The amount of time required for a vehicle to go from start to finish without conflicting traffic is used as a base case. Without modifying AIM's random vehicle generation processes, several simulations must be run to collect all base case times for all turning directions from each lane. A single simulation within the synthetic environment can be designed with no traffic conflicts to get the base case times for each turning direction. These base case times are used to calculate the average delay experienced by each vehicle through their respective simulation environment when traffic is present.

A C++ analysis program is used to parse the arrival and departure times for each vehicle, using their unique VINs, within a single trial. These numbers are then compared to the base case times for that particular type of vehicle's starting lane and turning direction. This number is then averaged for all vehicles within that simulation, becoming the average total delay. In addition to average delay, the program also determines the maximum delay experienced by any one vehicle, total number of generated vehicles, as well as total serviced vehicles by the intersection.

3.7 Scenario 1: Rogue non-Autonomous Vehicle

The first instance of misbehavior introduced into the synthetic environment is that of a stationary rogue non-autonomous vehicle within the autonomous intersection. As expected, this non-autonomous vehicle has no means of communicating information to other vehicles or infrastructure via V2X. Since there is a lack of communication, the vehicle poses a safety concern to any autonomous vehicles traveling through this intersection.

This scenario should not be limited to a rogue vehicle located within the intersection. One could easily imagine a number of other road hazards that could take its place. For instance, especially in rural areas, large wildlife may pose a similar hazard. Items which have fallen off vehicles and rolling boulders could also be dangerous. All of which have no means to communicate their presence via V2X communications.

For the placement of this rogue vehicle, there are three different scenarios. An assumption for these scenarios is that the rogue vehicle is stationary within the intersection.

3.7.1 Rogue Vehicle Approach.

This scenario consists of placing the rogue vehicle in three different locations to show the expected behavior of the environment and its protocol. The experiment has a total simulation runtime of 1,800 seconds. First, the rogue vehicle is placed in the center of the intersection facing southwest. The second position and orientation is in lane 4, just inside the intersection headed north. Finally, the vehicle is placed facing northwest in the center of the southeast quadrant of the intersection. See Figures 14, 15, and 16 for the positions and orientations of the three different test cases.




Figure 14. Rogue vehicle in center of intersection facing southwest



Figure 15. Rogue vehicle in southeast quadrant facing north



Figure 16. Rogue vehicle in center of southeast quadrant facing northwest

3.7.2 Mitigation.

Without any mitigation techniques or strategies, the rogue vehicle (or other hazard) poses a significant safety concern for all autonomous vehicles using this intersection. Within a properly functioning autonomous intersection, near misses are expected and commonplace due to the nature of the multi-directional traffic patterns. Autonomous vehicles may not have the proper amount of time to avoid such an obstacle within an intersection; therefore, the intersection must have some defenses against this type of hazard, otherwise the safety of the intersection is compromised.

The mitigating behaviors of the autonomous intersection should be able to warn incoming traffic of the rogue vehicle hazard. Within the synthetic environment, one can use a virtual sensor to locate these rogue objects and warn the incoming drivers of its location. The strategy to protect against the rogue vehicle involves first locating the hazard, then broadcasting a warning to all incoming vehicles.

The first possible strategy is of a simple detect-and-deny fail safe. The intersection could simply detect the presence of the rogue vehicle and broadcast a warning message to all incoming autonomous within range. Once receiving the warning broadcast, all vehicles would then know not to proceed. This mitigation strategy maintains safety considerations, however, there may not be a reason to block *all* traffic, especially if the intersection is large enough to allow the non-obstructed paths to operate as normal. This alternative method allows the autonomous intersection to stop only some of the traffic, rather than all of the traffic.

In this second possible strategy, the autonomous intersection uses its available sensors to detect the presence of the hazard. Next it approximates the size of the hazard. Following this information gathering, the intersection determines which paths are obstructed by the rogue vehicle. For example, all traffic from the northbound right lane, eastbound traffic in the right lane, and eastbound traffic from the right side southbound lane are obstructed with the rogue vehicle located such as in Figure 15. Once all obstructed lanes are determined, the autonomous intersection broadcasts the warning message to all autonomous vehicles within range. The autonomous vehicles use the list of obstructed lanes to determine if its path is safe. If the autonomous vehicle's path is not safe, then for simplicity sake it stops. This behavior ensures the safety of the autonomous vehicle, the rogue vehicle, and their occupants. Figure 17 demonstrates this mitigation strategy allowing two autonomous vehicles with non-obstructed paths to pass through the intersection. The arrows are added to clarify the direction of the autonomous vehicles. Also notice the reservation tiles highlighted under the rogue vehicle. This strategy enforces the reservation scheme even for the non-autonomous vehicle.




Figure 17. Mitigation in action allowing passage to non-obstructed autonomous vehicles

3.7.3 Expected Results.

Lacking the mitigation policy in place, the expected behavior is that autonomous vehicles collide with the rogue vehicle. Given the assumption that lane changes are not allowed, the expected behavior with mitigation is straight forward. One can analytically determine which paths are obstructed and which are not. For the paths that are not obstructed, the total average delay should be unchanged, or even decreased. However, the first autonomous vehicle that is unable to pass through the intersection from a given lane blocks all subsequent autonomous vehicles in said lane. For example, referring to Figure 17, an eastbound vehicle in the right lane turning right has a clear path. However, if that vehicle is behind another vehicle going straight, the second vehicle is unable to complete its path because lane changing is not available. More complex mitigation techniques, such as lane changes, are out of scope for this thesis.

3.8 Scenario 2: Sybil Attack

Another instance of misbehavior that can be studied in this synthetic environment is the Sybil attack [26]. This scenario portrays a simple Sybil attack taking place at a single autonomous intersection using a number of different strategies. Consider an autonomous vehicle controlled by a malicious actor herein referred to as the Sybil attacker. A Sybil attacker's motives may include:

1. Nuisance: Imposing delays on commuters to cause inconveniences.
2. Herding: Considering multiple intersections over a large area, such an attack may herd vehicles in a given direction.
3. Carjacking: With sufficient delays, vehicles necessarily stop within roadways. This type attack may cause commuters to be more susceptible to crimes such as carjacking.

See Figure 18 for a possible positioning of a Sybil attacker in the synthetic environment. The Sybil attacker carries out this attack with the following goals:

- Carry out the attack with the least cost possible
 - Minimize network traffic (reservation requests)
 - Minimize faulty identities used (VINs)
- Cause disruption for legitimate autonomous vehicles at the intersection in terms of increased delay




Figure 18. Sybil attacker shown on shoulder making false reservation requests

In order to minimize network traffic, the Sybil attacker may investigate possible reservations to fabricate. For instance, from each lane, one could proceed straight or turn one of two directions. Therefore, from each travel lane there are three possible paths. In a four lane highway autonomous intersection (see Figure 19), there are 24 possible paths for which to make reservations (i.e. $3 \text{ paths/lane} \times 8 \text{ lanes} = 24 \text{ paths}$).




Figure 19. Legal turn direction options from a single lane with lane numbers

Consider the northbound left turn beginning in lane 0 and terminating into lane 3 shown in Figure 19. This reservation path will obstruct 19 of the 24 possible paths at some point in time. (Table 7 provides the number of obstructed paths possible given a particular path with Appendix A showing a detailed matrix showing which paths obstruct with each other.) Note that two overlapping paths do not necessarily mean a collision between vehicles. In the matrix in the appendix, an “x” denotes two paths that may intersect, or obstruct, at some point while an “o” denotes two paths that cannot conflict at all. Lanes 0 through 3 are outside lanes, while lanes 4 through 7 are inside lanes. All outside lanes hold the same number of total obstructing paths while all inside lanes respectively do the same. A Sybil attacker would be most interested in the maximum effect with the least amount of effort. Therefore, the attack may choose as few as two turns, namely left turns from lanes 0 and 3. Using these two

lanes, all paths within the intersection are blocked. Note that an “x” appears in all columns given the rows in lane 0 and 3 corresponding to left turns.

Table 7. Collision matrix for four-lane autonomous intersection

Lane	Turn*	Total**
0	L	19
	S	15
	R	5
1	L	19
	S	15
	R	5
2	L	19
	S	15
	R	5
3	L	19
	S	15
	R	5
4	L	14
	S	15
	R	9
5	L	14
	S	15
	R	9
6	L	14
	S	15
	R	9
7	L	14
	S	15
	R	9

*L = left turn, S = straight (no turn), R = right turn

**Number of obstructed paths

Next, consider how the Sybil attacker may request reservations using these two turning directions to maximize the effects. The Sybil attacker does not want to make

requests to the autonomous intersection that the attacker knows will be denied since that would increase unneeded network traffic.

3.8.1 Sybil Attacker Approach - Serial Reservations.

This approach demonstrates the attacker making a single reservation within the intersection at a time. In other words, at any given moment, the attacker only holds a single confirmed reservation. There are multiple reservations, however, they are serial in nature since only one reservation is active within the intersection at any moment.

To disrupt traffic as much as possible with this approach, the arrival and departure times have to be considered. These times are directly related to the velocity within the intersection. With low speeds used by the attacker, it is possible to have a legitimate reservation granted that allows a vehicle to cross the Sybil attacker's paths. For instance, consider a right turn from southbound lane 2 to eastbound lane 3. Despite the reservation for the Sybil attacker being made for a left turn from lane 0 to 3, the southbound legitimate vehicle is still able to proceed, possibly without delay depending on its arrival time. Thus, the arrival time and departure time must be close together for the attacker's request, i.e. the velocity within the intersection must be increased, to require a denial to legitimate vehicles.

The first experiment's setup includes reservation requests for left turns from lane 0 to lane 3 by the attacker. These reservations are made with the arrival times and departure times reflecting a high velocity. The second experiment's setup is extend the previous setup with left turns from lane 3 to lane 2. Reservation requests are sent in an alternating fashion to keep the serial nature of the reservations in tact.

3.8.2 Sybil Attacker Approach - Parallel Reservations.

An attacker may not be able to make successful reservations with high velocities within the intersection due to defenses in place by the autonomous intersection. Introduced in this approach is the concept of parallel, or simultaneous, reservations. Rather than a single reservation at a time as in the previous approach, the attacker is able to make multiple non-conflicting reservation requests. The resulting parallel reservations reflect a platoon of closely spaced vehicles traveling at the appropriate speed through the intersection, however, none actually exits.

This approach demonstrates the attacker's ability to request several reservations within the intersection at a time. Reservations must be timed such that each of them does not conflict with another. Subsequent requests have to be delayed enough to not provide enough space for legitimate drivers' reservation requests. If the attacker makes two requests that start at the same time, the second request is necessarily denied by the autonomous intersection due to the safety violation of the reservations overlapping within the intersection. The attacker must delay the subsequent reservation requests enough to be granted the reservations.

The first experiment's setup includes reservation requests for left turns from lane 0 to 3 by the attacker. The reservations are made such that there is less than one vehicle's length between reservations. The reservations resemble those in Figure 20.




Figure 20. Parallel reservations from lane 0 to lane 3 made by Sybil attacker

The second experiment's setup extends the previous setup with left turns from lane 3 to lane 2 (see Figure 21). These reservations are sent in an alternating fashion maintaining the parallel nature of the reservations. The space between each reservation in a particular path needs to be adjusted appropriately in order to weave the alternating reservation paths together. Each reservation request is made in such a way that the paths are close without overlapping. Recall the attacker intends to not get rejections for reservation request.




Figure 21. Alternating parallel reservations made by Sybil attacker

3.8.3 Mitigations.

The main problems exemplified by this scenario are those of spoofing identities and requesting reservations with characteristics that should not be allowed. It follows that an update in the autonomous intersection policies will suffice. In the event of one of the Sybil-type attacks previously described, a traffic intersection unwittingly accepts reservations that should never be fulfilled since they are not coming from legitimate vehicles. The intersection also unwittingly denies legitimate users' reservation requests, therefore queuing any vehicle whose requested reservation conflict with the Sybil attack. Also affected are any vehicles behind the queued vehicles regardless of their intended direction of travel through the intersection.

The first mitigation technique is to only allow reservations up to a fixed time in the future. Without this mitigation technique, the attacker is able to preemptively request reservations far in advance. Since these reservations are made so far in advance, when a legitimate vehicle finally arrives within range of an intersection, the reservation request based on its estimated time of arrival derived from its current speed will have

already been taken by the attacker. This technique, however, does not guarantee that the attacker does not continue to disrupt the flow of traffic. Since the attacker is constantly making reservations with spoofed identifications, there exists a race condition for the times furthest in the future. If a legitimate user secures one of these later reservations, then it will likely have to slow down, causing significant delay depending on the latest available time to enter the intersection.

To mitigate the reservations made in the serial reservations approach, the intersection could update its policy to not allow for high speeds within the intersection. Establishing a check for a minimum amount of time spent within the intersection will directly mitigate the high speed reservation requests. The high speed reservation requests used will be rejected. In fact, all reservation requests will be rejected when exceeding the maximum velocity allowed within the intersection. Each of the serial reservations experiments will be executed using this mitigation strategy.

Offered here are two techniques to tackle the issue of spoofing identities. First, each vehicle's messages could contain a unique signed certificate by a third party, thereby increasing the difficulty of the Sybil attack. The autonomous intersection will then be able to verify the validity of the vehicle's certificate. If the certificate is not legitimate, then the reservation request is denied. This approach is easily abstracted in the synthetic environment by maintaining a list of valid autonomous vehicle VINs. Each time a reservation request is made, the intersection manager simply has to see if the VIN matches to any of the valid VINs.

A second, more elaborate mitigation technique to combat spoofing identities requires the incorporation of sensor data collected by the intersection. Giving the autonomous intersection the ability to detect vehicles approaching and record the approaching vehicles' identifiers, then a match between the requested reservation and the sensors' collected data could be used to verify the request prior to granting the

requests. When a reservation request arrives at the autonomous intersection, the intersection would compare the arrival time in the message with the estimated arrival time from the sensors of that particular vehicle. If there is a mismatch in the arrival times, or the vehicle is not found, then the intersection could refuse the reservation request. This mitigation technique would require a network of sensors able to detect the vehicles and their identifiers, however, it would mitigate the presented Sybil attack. There is a possibility of hacking the network of sensors which would involve attacks beyond the presented Sybil attack scenario's scope.

The mitigations used in the synthetic environment is a combination of the policy updates and using certificates to verify valid users via a list of legitimate VINs. This ensures legitimate autonomous vehicles do not exceed the intersection's maximum velocity as well.

3.8.4 Expected Results.

Considering no mitigation strategy, a Sybil attacker would be able to forge numerous VINs in order to make false reservation requests. These requests would be granted without proper defenses in place by the intersection. Depending on the false reservations made, some or all autonomous vehicles would be affected. The effects could either be increased delay or inability to get a reservation at all. This instance of misbehavior, left unmitigated, does not result in a safety concern. However, it does present a special kind of denial of service (DoS), denying reservations to legitimate users while maintaining connectivity.

Within the serial reservation Sybil attack approaches, traffic from all directions is expected to come to a complete halt at some point. As soon as the leading autonomous vehicle in each lane is denied its reservation request by the intersection, then all vehicles in that lane are stopped. This will happen immediately in the second

serial reservation experiment since the two alternating left turns block all lanes' paths through the intersection.

The parallel reservations Sybil attack approach is a little bit more interesting. The first experiment only uses left turns from lane 0. The results of this experiment is expected to be the same as the results of the first during the serial reservations Sybil attack. At some point, all directions of travel will be blocked. In the experiment where two left turns are alternated between lane 0 and lane 3, the reservations must be separated enough to ensure the reservations will not be rejected. Given the spacing of reservations, vehicles taking certain paths will be able to attain reservations. Consider a right turn from lane 1. It is expected that this right turn would be reserved given the space between the Sybil attacker's requests. The final steady state of stopped traffic in every lane is expected, however, not as quickly as the other experiments.

When the mitigation strategies are in place though, the attacker should have no effect on the delay experienced by the legitimate autonomous vehicles. The mitigation strategy used to defend against the serial reservations (maximum velocity within intersection) should show to be ineffective with the parallel reservations. The mitigation strategy used in the parallel reservations attack should work in both approaches since it is checking identities rather than intersection velocities.

3.9 Scenario 3: Squatter Attack

In this instance of misbehavior, a legitimate autonomous vehicle fabricates a custom atypical reservation request. This custom reservation request allows for the misbehaving autonomous vehicle to come to a complete stop for a given period of time inside the intersection. At some point, the vehicle leaves the intersection, all while being compliant with its requested and confirmed reservation.

3.9.1 Squatter Approach.

An autonomous vehicle with an occupant with the intent to inject misbehavior, known as the squatter, approaches an autonomous intersection. The squatter fabricates a custom reservation request. Within this request, he proposes to come to a complete stop within the intersection. Unlike the typical reservation request, this request has an arrival time as usual, however, the reservation has him coming to a complete stop in the center of the intersection, and then continue on. This unique reservation request is sent to the autonomous intersection and the reply is sent. The autonomous vehicle receives a confirmation message so long as the requested reservation does not conflict with other confirmed reservations. The vehicle is now able to enter the intersection and come to a complete stop in the middle of the intersection (or wherever the custom reservation placed the stopping point).

Within the synthetic environment, the squatter attempts to stop within the intersection for approximately 180 seconds in simulation. The squatter stops at the middle point in its path through the intersection. Three separate paths are used to validate the expected behavior of the autonomous intersection. The squatter is headed northbound in lane number 1 (inside lane) and have 3 different paths: left turn, straight, and right turn. Each path taken is studied at the 3 traffic levels of 100, 200, and 300 vehicles generated per hour per lane.

The first generated vehicle in each generation schedule is modified to fit the squatter profile described above. Also, since the squatter's goal is to cause delays from within the intersection, only the first generated car is assigned the squatter role.

3.9.2 Mitigation.

The mitigation for this type of misbehavior is similar to that of the Sybil attack scenario. The autonomous intersection's reservation policy requires an update. This

update differs from that of the Sybil attack, since it is quite the opposite. Rather than requiring the proposed reservations to have a maximum acceptable velocity, this policy would require a maximum acceptable amount of time within the intersection or a minimum acceptable velocity.

The policy checks each request to ensure the amount of time between the arrival and departure times is reasonable. In the case of the squatter attack, any stopping within the intersection quickly adds to the amount of time spent within the intersection. Within the synthetic environment, the autonomous intersection checks to ensure the amount of time requested does not exceed the time needed to pass through the intersection at a given velocity. The speed limit within the intersection is set to be 8 meters per second.

3.9.3 Expected Results.

The autonomous vehicle squatter attack, without mitigation, causes increased delays. Since the squatter requested, and was granted, a legitimate reservation made prior to its arrival, there is no safety concern since no conflicting reservations would have been made. However, there will still be the issue of increased delays while this misbehavior is taking place. Considering lanes and paths affected by the squatter, the average delays will be greater than without the squatter. However, lanes with paths which are not affected by the squatter will experience no additional delay. Also, for the unaffected lanes and paths, the average delay may decrease due to the lower amount of traffic competing for the same time-space within the intersection.

With the mitigation in place, the overall average delays per lane and total should be unchanged when compared to the baseline scenarios. This is expected since the mitigation eliminates the squatter's attack completely.

3.10 Human Control in a Fully Automated Environment

Other modifications to the AIM protocol have been made to accommodate human control. However, when humans maintain control of the vehicle in these protocols, the intersection reverts to traditional traffic signals to communicate with the human-controlled traffic. Each of the modified versions of AIM discussed in section 2.4 do not allow the human to persistently maintain control of the vehicle while also removing the need for the traditional traffic light.

Can humans safely maintain control at all times when traveling through a *fully autonomous* intersection? A fully autonomous intersection does not use traditional traffic signals in normal operation. Backup traffic signals may be in place for emergency and/or failure situations. The synthetic environment is the only reasonable place to test such modification to a reservation-based protocol for to many reasons including safety.

Incorporating the ability of a human to maintain control of a semi-autonomous vehicle while navigating an intersection without traditional signals is not trivial. In addition to the traditional controls such as an accelerator, brake pedal, steering wheel, turn signals, mirrors, and speedometer (to only list a few), there must be some additional controls and/or displays available to accommodate human control through an autonomous intersection. Described below are these additional controls and displays that are required.

Currently, traditional road signs communicate to drivers information about upcoming hazards and changes in traffic patterns such as sharp bends, speed limits, and intersections. An in-dash indicator to notify the human of an intersection that has come into range would be needed. In addition to this indicator, a button could be used to initiate communication with the autonomous intersection [16]. Once this button is pressed, the vehicle then communicates with the intersection to arrange

a reservation for passage, initiating the messaging flow starting in the autonomous vehicle entity in Figure 5 (see above).

During times of congestion, the reservation request sent may be denied. In this scenario, there is a requirement to inform the driver of this fact. For simplicity, any denial of a reservation requires the driver to slow the vehicle. At this point, the driver may press the button again or the vehicle may make another request automatically. This could continue until a point is reached prior to entering the intersection until a reservation is successful.

Once a reservation is successful, this information is fed back to the vehicle and needs to be displayed to the driver. An indicator would show that a reservation is made for the desired path as well as the velocity to maintain while inside the intersection. The confirmed reservation presents the need for a mechanism to communicate to the driver the goal velocity to arrive at the required time and/or the velocity required to travel within the intersection. This feedback is pivotal in ensuring the human-controlled vehicle arrives and passes through the intersection at the correct time. It must be able to constantly update the goal velocity based on the current time, arrival time, and distance to the intersection. This indicator may also be used to communicate the goal velocity while within the intersection.

Finally, independent of the velocity of the vehicle, the human must be able to maintain the correct lateral control of the vehicle, especially within the intersection. This path maintainer feedback indicator would notify the driver if the vehicle is too far left or right from the center of the current lane and also the dictated path through the intersection.

Table 8 shows a summary list of the possible controls and feedback devices discussed. The manner at which these devices are displayed to the driver are discussed in section 4.3.

Table 8. Human controls and feedback devices required for persistent human control

Item	Description
In-Range Indicator	A device that informs the driver of an autonomous intersection within range.
Request Reservation Button	A device that initiates V2X communication to request a reservation from the intersection.
Denied Reservation Indicator	A device which warns a driver that the requested reservation is denied.
Granted Reservation Indicator	A device which confirms to the driver a reservation is successful and the assigned velocity while within the intersection.
Goal Velocity Indicator	An active device which informs the driver of the required velocity to maintain to meet keep the reservation. May also be used to maintain the correct velocity within the intersection.
Maintain Path Indicator	An active feedback device which informs the driver of the left/right position correctness based on the lane or planned path within the intersection.

Armed with these controls and feedback devices, the human may now be able to safely enter the fully autonomous intersection. The synthetic environment is the best place to test the ability of the human to safely pass through the autonomous intersection due to obvious safety concerns.

IV. Analysis

This chapter presents the analysis results from the experiments executed for the verification and validation of the synthetic environment. The first section discusses the baseline scenarios which were used to establish the viability of the synthetic environment's capability to manage an autonomous intersection. Next, the misbehavior scenarios and are discussed in detail along with the persistent human control observations made.

4.1 Verification

This section describes the analysis of the baseline scenarios. These scenarios show that the synthetic environment's autonomous control management system functions properly and is a viable intersection management system. To establish the autonomous intersection's legitimacy in the synthetic environment, three scenarios were executed with five separate trials for each scenario. To consider the autonomous intersection management system viable, the dependent variables of average delay and number of collisions were considered. These variables have two conditions that must be met. First, there must be no physical collisions between any two autonomous vehicles. Second, there must not be a significant difference between the average delays between the Autonomous Intersection Management (AIM) Simulator and the synthetic environment's autonomous intersection manager.

To show there is no significant difference between the mean delays, a two sample z -test is run for each trial in each scenario since the sample sizes are large. The null hypothesis for each test is that there is no difference, while the alternate hypothesis in this two-tailed test is that there is a difference in the mean delays. If, for any test, the p -value is less than $\alpha = 0.05$, then the null hypothesis is rejected, meaning

that there is statistical evidence to suggest there is a difference in the mean delays between the AIM Simulator and the synthetic environment.

The first scenario set the traffic volume at 100 cars generated per lane per hour. Each of the five trials for this experiment were executed for a total time of 30 minutes, or 1,800 seconds, in simulation time.

4.1.1 Baseline Results.

Each scenario is executed with all five trials running concurrently in five separate instances of the synthetic environment. Tables 9, 10, and 11 show the results from the baseline scenarios 1, 2, and 3 respectively. Figures 22, 23, and 24 show the comparison between the average delays between the AIM Simulator and the synthetic environment. The circles and the plus signs on the graphs represent the total average delay experienced by the autonomous vehicles in the AIM Simulator and the synthetic environment, respectively.

Table 9. Baseline scenario 1 results

Response Variable	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
<i>AIM Simulator Output</i>					
Throughput	404	380	394	380	392
Average Delay (s)	0.16	0.1648	0.1562	0.1568	0.1418
Standard Deviation (s)	0.26	0.65	0.60	0.61	0.52
Maximum Delay (s)	3.94	5.86	6.24	6.02	3.94
Collisions	0	0	0	0	0
<i>Synthetic Environment Output</i>					
Throughput	404	381	394	381	393
Average Delay (s)	0.1846	0.1503	0.1579	0.1933	0.1449
Standard Deviation (s)	0.65	0.50	0.50	0.77	0.44
Maximum Delay (s)	7.59	5.05	4.78	9.99	4.51
Collisions	0	0	0	0	0

Table 10. Baseline scenario 2 results

Response Variable	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
<i>AIM Simulator Output</i>					
Throughput	703	722	711	640	681
Average Delay (s)	0.3846	0.3767	0.3690	0.2601	0.2764
Standard Deviation (s)	1.23	1.04	0.97	0.80	0.82
Maximum Delay (s)	14.56	9.06	7.26	6.84	5.98
Collisions	0	0	0	0	0
<i>Synthetic Environment Output</i>					
Throughput	705	726	715	640	688
Average Delay (s)	0.3327	0.3818	0.2813	0.3152	0.3048
Standard Deviation (s)	1.07	1.15	0.73	1.00	0.95
Maximum Delay (s)	9.59	11.40	6.66	12.48	9.80
Collisions	0	0	0	0	0

Table 11. Baseline scenario 3 results

Response Variable	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
<i>AIM Simulator Output</i>					
Throughput	935	935	929	932	957
Average Delay (s)	0.6619	0.5724	0.5479	0.6420	0.6019
Standard Deviation (s)	1.73	1.45	1.41	1.68	1.67
Maximum Delay (s)	16.62	12.28	16.06	17.56	16.64
Collisions	0	0	0	0	0
<i>Synthetic Environment Output</i>					
Throughput	937	939	933	933	957
Average Delay (s)	0.5698	0.6905	0.5334	0.6782	0.5298
Standard Deviation (s)	1.55	2.20	1.55	2.50	1.64
Maximum Delay (s)	14.62	31.77	18.59	43.91	17.12
Collisions	0	0	0	0	0




Figure 22. Average delay observed in Simulation 1

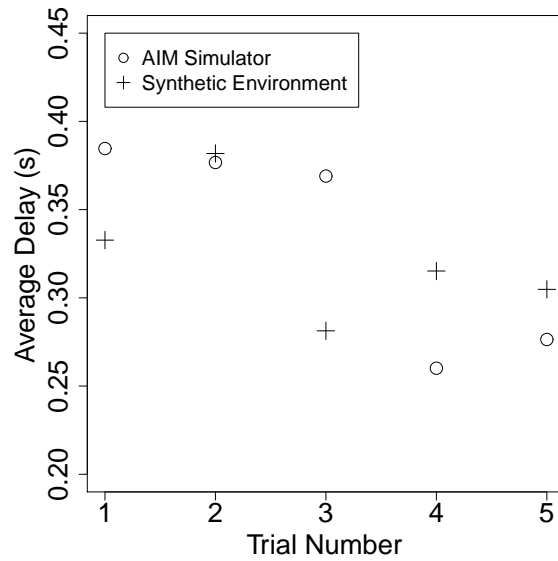


Figure 23. Average delay observed in Simulation 2




Figure 24. Average delay observed in Simulation 3

4.1.2 Baseline Analysis.

The two sample z -tests results from each of the scenarios are in Table 12. The p -values are all greater than or equal to the significance level of $\alpha = 0.05$. This gives statistical evidence to support that the total average delays between the AIM Simulator and the synthetic environment do not differ with strong statistical evidence. With this establishment of the viability of the synthetic environment, it can be used to conduct further experiments. These experiments are conducted without comparing the delays from the AIM Simulator since it is not able to be used to introduce any sort of misbehavior or cyber attack scenario. These experiments are also used to validate the expected behavior of the synthetic environment.

Table 12. Results analysis for all baseline scenarios

	Scenario 1				
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
AIM* Average Delay (s)	0.16	0.1648	0.1562	0.1568	0.1418
AIM Standard Deviation (s)	0.26	0.65	0.60	0.61	0.52
SE** Average Delay (s)	0.1846	0.1503	0.1579	0.1933	0.1449
SE Standard Deviation (s)	0.65	0.50	0.50	0.77	0.44
Two-tailed <i>z</i> -test <i>p</i> -value	0.55	0.73	0.97	0.47	0.93
	Scenario 2				
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
AIM Average Delay (s)	0.3846	0.3767	0.3690	0.2601	0.2764
AIM Standard Deviation (s)	1.23	1.04	0.97	0.80	0.82
SE Average Delay (s)	0.3327	0.3818	0.2813	0.3152	0.3048
SE Standard Deviation (s)	1.07	1.15	0.73	1.00	0.95
Two-tailed <i>z</i> -test <i>p</i> -value	0.40	0.93	0.05	0.28	0.55
	Scenario 3				
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
AIM Average Delay (s)	0.6619	0.5724	0.5479	0.6420	0.6019
AIM Standard Deviation (s)	1.73	1.45	1.41	1.68	1.67
SE Average Delay (s)	0.5698	0.6905	0.5334	0.6782	0.5298
SE Standard Deviation (s)	1.55	2.20	1.55	2.50	1.64
Two-tailed <i>z</i> -test <i>p</i> -value	0.23	0.17	0.83	0.71	0.34

*AIM = AIM Simulator

**SE = Synthetic Environment

4.2 Validation

This section discusses the results obtained from running the validation scenarios. These scenarios introduce a number of different misbehaviors into the synthetic envi-

ronment. As discussed previously, these misbehaviors cannot be modeled in the AIM Simulator.

4.2.1 Rogue non-Autonomous Vehicle.

The experiments run in this misbehavior scenario were executed without the mitigating techniques in place. The first autonomous vehicle whose path was obstructed by the rogue vehicle collided with it. This result was expected due to the intersection having no ability to locate and account for the presence of the rogue vehicle.

Using the mitigation technique discussed in section 3.7.2, the autonomous vehicles were informed to not approach the intersection. The broadcast warning to each of the autonomous vehicles contained which starting lanes and turning directions that are unsafe for passage. Additionally, when an autonomous vehicle is queued behind another vehicle who is blocked by the rogue vehicle, the trailing vehicle is notified that the leading vehicle is blocked by a rogue vehicle. This is important in that the trailing vehicle's path may not be obstructed, however, since the trailing vehicle is queued behind one that is, then it is also blocked indirectly by the rogue vehicle.

Each experiment below consists of three separate traffic volume levels. The first trial's generation schedule from each of the scenarios in the baseline experiment were used to conduct the experiments. The traffic volumes used here are 100, 200, and 300 vehicles per hour per lane.

4.2.1.1 Rogue Vehicle Location: Center.

The location of the rogue vehicle in this experiment blocks at least one path from every approaching lane. In Figure 25, the mitigation strategy is in place, stopping vehicles which have an obstructed path by the rogue vehicle. Also, it was observed that the vehicles approaching which are able to pass still may, provided they are

not queued behind another vehicle being denied passage. Note the rogue vehicle's reservation is showing in the figure while the other autonomous vehicle is passing through the intersection. This shows that the intersection is able to locate the vehicle accurately. Each new reservation request is made with this rogue vehicle's reservation set.




Figure 25. Queues formed during successful rogue vehicle misbehavior mitigation

Once the experiment reaches a steady state, all directions are queued and all approaching vehicles are stopped. This steady state gridlock occurred within a few minutes into the simulation. This is the expected behavior given that at some point, all starting lanes have a possible path that is obstructed by the rogue vehicle. Once an autonomous vehicle is stopped in a given lane, all traffic behind it is stopped as well. This steady state level of gridlock can be seen in Figure 26. This figure shows the southbound and westbound lanes. Table 13 shows the total vehicles generated, throughput, and delay for this experiment. (Note if no vehicles pass through an intersection from a given lane, then the delay is listed as not applicable (N/A) due to

the way delay is calculated (see Section 3.5.3.) Recall that the total vehicles generated per lane is low due to the queues forming.




Figure 26. Approaching lanes stopped by rogue vehicle in center of intersection using misbehavior mitigation

Table 13. Results of rogue vehicle in center of intersection with mitigation

100 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	13	10	14	8	8
Throughput	0	0	0	5	2	6	0	0
Delay (s)	N/A	N/A	N/A	0.01	0	0	N/A	N/A

200 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	8	8	8	15	23	11	8	8
Throughput	0	0	0	8	15	3	0	0
Delay (s)	N/A	N/A	N/A	0	0.08	0	N/A	N/A

300 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	11	7	12	8	8	10	9
Throughput	0	3	0	4	0	0	2	1
Delay (s)	N/A	0	N/A	0.01	N/A	N/A	-0.01	1.33

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.1.2 Rogue Vehicle Location: Northbound Outside Lane.

This experiment showed that the expected results for the rogue vehicle placement in the southern portion of the northbound outside lane of the intersection were attained. Since the rogue vehicle is placed in such a way in the intersection, some approaching lanes, no matter the path taken, were unobstructed. However, as soon as the first vehicle approaches with an obstructed path, all vehicles behind it are queued up behind it. Figure 27 shows the steady state behavior of the intersection with the rogue vehicle mitigation in action. Both northbound lanes are blocked since the vehicle in the inside lane is not able to take a right turn to the east. These are the expected results given the placement of the rogue vehicle. Table 14 shows the expected results attained. The steady state gridlock, for the lanes that become blocked, was achieved within the first few minutes of simulation.




Figure 27. Steady state reached with rogue vehicle in northbound lane with misbehavior mitigation

Table 14. Results of rogue vehicle in northbound lane of intersection with mitigation

100 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	11	7	57	13	37	8	56	55
Throughput	3	0	55	5	37	1	54	54
Delay (s)	0	N/A	0.08	0.01	0.39	0	0.10	0.05

200 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	9	7	81	15	86	8	90	88
Throughput	1	0	79	8	85	0	88	85
Delay (s)	0	N/A	0.24	0	0.11	N/A	0.17	0.04

300 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	11	7	119	12	127	9	132	116
Throughput	3	0	116	4	124	2	129	114
Delay (s)	0.01	N/A	0.41	0.58	0.20	-0.01	0.15	0.07

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.1.3 Rogue Vehicle Location: Southeast Quadrant.

This experiment showed that the expected results for the rogue vehicle placement in the southeast quadrant of the intersection were attained. Similar to the previous placement of the rogue vehicle, some approaching lanes no matter the path taken were unobstructed. However, as soon as the first vehicle approaches with an obstructed path, all vehicles behind it are queued up behind it. Figure 28 shows the steady state behavior of the intersection with mitigation. The northbound and eastbound lanes were completely blocked while the westbound lanes were unobstructed. The

southbound inside lane was not block as well. These are the expected results given the placement of the rogue vehicle. Table 15 shows the expected results after reaching the steady state gridlock within a few minutes in simulation.




Figure 28. Steady state gridlock with rogue vehicle in southeast quadrant with misbehavior mitigation

Table 15. Results of rogue vehicle in southeast quadrant of intersection with mitigation

100 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	8	7	57	13	10	8	9	55
Throughput	0	0	55	5	2	1	1	54
Delay (s)	N/A	N/A	0.05	0.01	0	0	0	0.03

200 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	8	7	81	15	23	8	13	88
Throughput	1	0	79	8	15	0	5	85
Delay (s)	0	N/A	0.18	0	0.16	N/A	0.39	0.03

300 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	8	7	119	12	8	9	11	116
Throughput	0	0	116	4	0	2	4	114
Delay (s)	N/A	N/A	0.14	0.58	N/A	-0.01	0.44	0.04

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.1.4 Rogue Vehicle Conclusion.

These experiments validated the expected behavior of the synthetic environment given this type of misbehavior and mitigation strategy. Three different rogue vehicle positions along with three different traffic levels each were used in the experiments. The synthetic environment produced results that are in line with the expectations.

4.2.2 Sybil Attack.

The Sybil attacker in this research is a malicious actor with the ability to pose as multiple false identities. The identities used here are the autonomous vehicles' vehicle identification numbers (VINs). The Sybil attacker uses two approaches to cause disruption at the intersection. First, the attacker sends a series of reservation request, each of which do not overlap in the time domain. That is, each confirmed reservation's arrival time is after the previous reservation's departure time. This approach is called the serial reservations approach. Secondly, the attacker sends a chain of reservation requests, each of which may overlap with previous or subsequent reservation requests. Since these reservations overlap in the time domain, this approach is called the parallel reservations approach. In each of the approaches and experiments, the attacker is stationary just south of the intersection on the northbound shoulder.

The vehicle generation schedule used in each of the experiments is that of the baseline scenario 1, trial 1. This generation schedule generates approximately 100 vehicles per lane per hour randomly.

4.2.2.1 Serial Reservations.

The first experiment executed in this approach had the attacker make many high velocity reservation requests to the intersection. These requests for a left turn from lane 0 were all granted. The upper half of Table 16 shows the expected results of this experiment without the mitigation in place. After the first vehicles traveling eastbound, the traffic from that direction stopped due to the attack.

The purpose of this experiment was to show the expected overall behavior of queues formed in all directions where paths intersect with the attacker's requests. Figure 29 shows the queues formed. All lanes were stopped at approximately 100 seconds into the simulation. Each of the eastbound lanes had one vehicle pass through

the intersection before the queues formed for those lanes. These vehicles turned right into the northbound lane unobstructed.

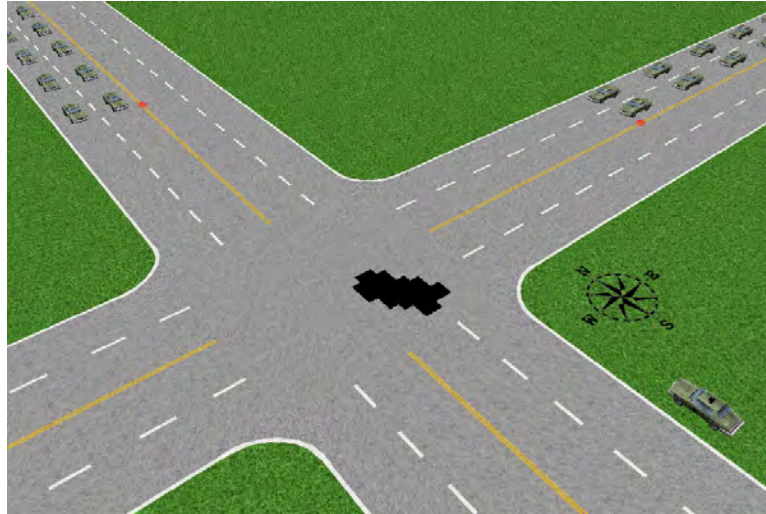


Figure 29. Queues formed during successful Sybil attack

The second experiment in this approach had the attacker make many high speed reservation requests to the intersections for two separate turn directions. These requests for left turns from lane 0 and 3 were all granted. The lower half of Table 16 shows the expected results of this experiment attained without the mitigation in place. No autonomous vehicles were able to pass through the intersection.

The expected behavior once again was attained. Since all paths from all directions were obstructed by the attacker's false reservations, queues began to form from all directions, allowing no vehicles to pass through. The steady state of all directions blocked is similar to that in Figure 29 from the previous experiment.

Table 16 shows the results for total throughput and delay for the experiment. The experiment used the same vehicle generation schedule as that of the baseline scenario 1, trial 1 which generated 100 vehicles per lane per hour. The delays listed are those of the specific lane of travel.

Table 16. Results of serial reservations in Sybil attack

Left turns from NO lane	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	7	8	8	7	7
Throughput	0	0	0	0	1	1	0	0
Delay (s)	N/A	N/A	N/A	N/A	0	-0.06	N/A	N/A

Left turns from NO and WO lanes	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	7	7	7	7	7
Throughput	0	0	0	0	1	0	0	0
Delay (s)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

Introducing the mitigation strategy of requiring a minimum amount of time within the intersection successfully thwarts this type of attack. The attacker was at first able to make multiple reservations with high velocities. Now that the intersection's policy is updated to automatically deny reservations at high velocities, this attack is no longer able to be carried out. See Table 17 for the expected results of this experiment. The two experiments had the same results due to the successful defense against the serial reservation Sybil attack.

Table 17. Results of serial reservations in Sybil attack using minimum time mitigation

Left turns from NO lane	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	51	54	57	54	37	47	56	55
Throughput	50	54	55	52	37	47	54	54
Delay (s)	0.14	0.14	0.10	0.11	0.54	0.22	0.12	0.07
Total Average Delay (s)	0.16							

Left turns from NO and WO lanes	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	51	54	57	54	37	47	56	55
Throughput	50	54	55	52	37	47	54	54
Delay (s)	0.14	0.14	0.10	0.11	0.54	0.22	0.12	0.07
Total Average Delay (s)	0.16							

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.2.2 Parallel Reservations.

The first experiment executed in this approach had the attacker make many normal velocity (8 meters per second) reservation requests to the intersection. The requests were parallel, meaning that more than one reservation from the attacker were confirmed at any given moment in time. See Figure 30 to see the parallel nature of the reservations. These requests for a left turn from lane 0 were all granted. The upper half of Table 18 shows the expected results of this experiment without the mitigation in place. After the first vehicles traveling eastbound in each lane, the traffic from that direction stopped due to the attack. This is identical to the behavior in the first experiment in the serial reservations approach.

The purpose of this experiment was to show the expected overall behavior of queues being formed in all directions where the paths intersect with the attacker's requests. Figure 30 shows the the queues formed. All lanes were stopped at approximately 100 seconds into the simulation. Each of the eastbound lanes had one vehicle pass through the intersection before the queues formed from that direction. Each of these vehicles turned right into the northbound lane unobstructed.




Figure 30. Queues formed as a result of parallel reservation Sybil attack

The second experiment in this approach had the attacker make many normal velocity (8 meters per second) reservation requests to the intersections for two separate turn directions. These requests for left turns from lane 0 and 3 were all granted. The lower half of Table 18 shows the expected results of this experiment without the mitigation in place. Three starting lanes were able to attain a non-zero throughput. These lanes, however, did eventually experience gridlock when the first vehicle whose path was obstructed by the attacker's reservations was denied.

The expected behavior once again was attained. Since all paths from all directions were obstructed by the attacker's false reservations, queues began to form from all directions, allowing no vehicles to pass through. The steady state the gridlock is shown in Figure 31.




Figure 31. Sybil attack using alternating weaved left turns

Table 18 shows the results for total throughput and delay for the experiments. The experiments used the same generate schedule as that of the baseline scenario 1, trial 1 which generated 100 vehicles per lane per hour.

Table 18. Results of parallel reservations in Sybil attack

Left turns from NO lane	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	7	8	8	7	7
Throughput	0	0	0	0	1	1	0	0
Delay (s)	N/A	N/A	N/A	N/A	0	-0.06	N/A	N/A
Total Average Delay (s)	-0.03							

Left turns from NO and WO lanes	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	12	9	8	7	7
Throughput	0	0	0	5	2	1	0	0
Delay (s)	N/A	N/A	N/A	1.69	0.74	2.34	N/A	N/A
Total Average Delay (s)	1.53							

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

Introducing the mitigation technique used in the serial reservation approach (ensuring the reservation request does not use a high velocity within the intersection) is not effective in this approach. Table 19 shows these expected results attained. The results are identical to the experiment run without the mitigation in place. This is the expected behavior since the velocity used to form the reservation requests is 8 meters per second, which passes the check used in the mitigation strategy.

Table 19. Results of parallel reservations in Sybil attack using minimum time mitigation

Left turns from NO lane	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	7	8	8	7	7
Throughput	0	0	0	0	1	1	0	0
Delay (s)	N/A	N/A	N/A	N/A	0	-0.06	N/A	N/A
Total Average Delay (s)	-0.03							

Left turns from NO and WO lanes	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	7	7	7	12	9	8	7	7
Throughput	0	0	0	5	2	1	0	0
Delay (s)	N/A	N/A	N/A	1.69	0.74	2.34	N/A	N/A
Total Average Delay (s)	1.53							

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

Using the mitigation technique of verifying the autonomous vehicle's identity proved to be successful in thwarting the Sybil attacker. The attacker relies heavily on being able to fake identities, when this ability is taken away, the attack is no longer valid. Table 20 shows the expected results attained. These results match that of the baseline scenario 1, trial 1 where 100 vehicles per lane per hour are generated.

Table 20. Results of parallel reservations in Sybil attack using VIN verification mitigation

Left turns from NO lane	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	51	54	57	54	37	47	56	55
Throughput	50	54	55	52	37	47	54	54
Delay (s)	0.14	0.14	0.10	0.11	0.54	0.22	0.12	0.07
Total Average Delay (s)	0.16							

Left turns from NO and WO lanes	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	51	54	57	54	37	47	56	55
Throughput	50	54	55	52	37	47	54	54
Delay (s)	0.14	0.14	0.10	0.11	0.54	0.22	0.12	0.07
Total Average Delay (s)	0.16							

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.2.3 Sybil Attack Scenario Conclusion.

The synthetic environment behaved as expected during the Sybil attack scenario. The mitigation strategy which was designed to handle high velocity reservation requests from the Sybil attacker was successful. However, the defense was shown to be insufficient when the Sybil attacker did not use high speeds within the intersection. Giving the intersection the ability to verify identities was shown to thwart the Sybil attacker's efforts to disrupt normal traffic reservations. Both approaches, the serial reservation approach and the parallel reservation approach which used normal speeds within the intersection, were successfully defended against.

4.2.3 Squatter Attack.

For each of the squatter attack experiments, the squatter formulates a custom reservation. The reservation contains a prolonged stay within the intersection. The amount of time within the intersection, in addition to traveling to and from the stopping point, is set to 3 minutes, or 180 seconds. The squatter stopped midway through the intersection in each experiment as expected and stayed within its reserved space.

4.2.3.1 Left Turn Path.

The first round of experiments placed the squatter in a left turn path. The squatter came to a complete stop halfway through its turn, matching its reservation. The vehicles whose paths are now obstructed are queued up, as shown in Figure 32. The northbound outside lane is the only lane whose paths are not affected by the squatter.

The results for this experiment are in Table 21. Note that each lane experienced high average delays except for the northbound outside lane. This is to be expected since no path is obstructed from this lane.




Figure 32. Left turning squatter occupies space in intersection

Table 21. Results of squatter turning left without mitigation

100 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	52	54	56	54	37	46	56	53
Throughput	51	54	54	52	37	46	54	52
Delay (s)	16.17	0.14	6.76	8.15	7.62	11.91	4.74	11.92

200 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	87	90	79	90	79	94	90	85
Throughput	86	88	77	89	78	93	88	82
Delay (s)	13.51	0.26	10.28	3.47	16.04	11.09	3.21	7.99

300 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	111	111	112	112	118	102	128	110
Throughput	107	108	109	109	115	99	125	108
Delay (s)	8.46	0.68	9.61	8.89	9.63	13.78	6.08	6.33

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.3.2 Straight Path.

The second round of experiments placed the squatter in a straight path inside the intersection. The squatter stopped midway through the intersection as expected. The traffic with conflicting reservations are denied and queues begin to form while the squatter is stopped inside the intersection. Figure 33 shows the queues formed while the squatter is occupying the middle of the intersection.

The results for this experiment are in Table 22. At least one path from every direction is affected by this squatting location. This is evident by the increased

delays experienced by the vehicles in each lane. The low delays seen in the 100 and 200 vehicles per lane per hour charts for this experiment are an artifact of the generation schedule. Since the generation schedule randomly selected which direction to choose, these particular directions falsely show a low average delay. For instance, if no vehicles were scheduled to turn left from the eastbound outside lane, then the average delay for that lane would not be affected by the squatter since there would be no paths blocked by the squatter in that situation.




Figure 33. Queues formed while squatter occupies space in intersection

Table 22. Results of squatter traveling straight without mitigation

100 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	51	54	56	54	37	47	56	53
Throughput	50	54	54	52	37	47	54	52
Delay (s)	20.13	13.97	10.82	0.02	4.66	0.29	8.55	12.73

200 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	89	85	80	90	86	96	87	85
Throughput	88	83	78	89	85	95	85	82
Delay (s)	10.19	11.85	6.84	0.20	0.16	9.19	11.80	8.19

300 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	114	106	111	115	121	103	126	107
Throughput	110	103	108	112	118	100	123	105
Delay (s)	5.09	9.29	10.94	7.26	8.25	12.82	8.71	8.32

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

4.2.3.3 Right Turn Path.

Finally, the squatter's path is a right turn. After stopping midway through the intersection, as expected, conflicting traffic begins to queue. Figure 34 shows both westbound lanes unaffected as well as the southbound inside lane at approximately 145 seconds into simulation. This is the expected behavior since no paths from these lanes are obstructed by the squatter. Table 23 reflects the expected results as well with increased delays on all but the lanes mentioned.




Figure 34. Right turning squatter occupies space in intersection

Table 23. Results of squatter turning right without mitigation

100 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	51	54	56	54	37	46	56	55
Throughput	50	54	54	52	37	46	54	54
Delay (s)	17.65	13.57	0.10	0.11	3.98	11.86	0.12	0.07

200 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	89	85	81	90	86	94	90	88
Throughput	88	83	79	89	85	93	88	85
Delay (s)	11.23	10.08	0.34	0.20	0.17	10.12	0.37	0.42

300 vehicles/lane/hour	Starting Lane*							
	NI	NO	SI	SO	EI	EO	WI	WO
Total Generated	113	103	119	114	121	106	132	116
Throughput	109	100	116	111	118	103	129	114
Delay (s)	5.52	11.71	0.46	7.50	8.22	10.42	0.61	0.14

*{N, S, E, W} = {North, South, East, West}

{I, O} = {Inside Lane, Outside Lane}

Each of the squatter experiments return to normal experienced delays when the mitigation strategy is enabled. Essentially each mitigation behaves identically to

the baseline experiments discussed in section 4.1. Once the squatter is denied his custom reservation, he resorts to requesting a normal reservation and is granted the reservation.

4.3 Human Controls

This section discusses the preliminary observations made while testing the prototype protocol using the additional human control and feedback devices. The experiments conducted are notional and serve as a proof of concept rather than human subjects testing.

Figure 35 shows a sample screen shot of the synthetic environment. The black rectangles in the lower left, lower right, and upper right mimic side view and rear view mirrors. The remaining displays show some feedback information. On the left, top to bottom, are: current simulation time (s), confirmed reservation's arrival time (s), current velocity (m/s), and goal velocity (m/s). In the upper center, a compass displays one of eight cardinal or intercardinal headings (e.g., N, E, NE, etc.). On the right, top to bottom are the current simulation name (used for reference purposes), the reservation status, and a digital lateral offset mechanism.

The reservation status is a solid colored square located to the far right of the display. The square is green when a reservation is confirmed by the intersection. The square is yellow if no reservation is made while the vehicle is within range of the intersection's road side unit (RSU) and it turns red only if the vehicle does not have a reservation and it is within the stopping distance to the intersection. The square is black with the letters "N/A" when no intersection is within range.

In the center of the figure are two heads-up-displays (HUDs). They are each made up of two translucent rectangles; one black and one green. The vertical black rectangle, located just underneath of center screen, marks the desired center of the

path taken by the vehicle. If the vehicle veers left of the correct path, the green rectangle moves to the right of the vertical black rectangle. This indicates to the human driver the vehicle needs to be positioned further right to be placed correctly. The green rectangle turns red in the event the vehicle veers too far in any direction, left or right, of the intended path. Similarly, the second HUD located left of center screen, provides feedback to the human concerning the velocity of the vehicle. The horizontal black rectangle and its corresponding green rectangle indicates to the driver whether he/she should increase or decrease velocity. If the vehicle determines an increase in velocity is necessary, the green dot slides upward (eventually turning red in warning). Likewise, the green rectangle slides downward if the vehicle determines it needs to slow down.



Figure 35. Screen shot of synthetic environment with human control

The design and placement of the feedback devices prove to be important. With digital speedometer and goal speed indicator, having them close in proximity is important since they both contain information about velocity. Furthermore, an analog display (common in traditional vehicles) may be more beneficial than a digital display. An analog speedometer could have the goal velocity indicated in a separately colored

dial overlaid on top of the current velocity dial. Drivers may prefer to have the option of digital versus analog as well. Answers to these questions requires extensive testing.

Maintaining the center of the correct lane appears to be straight forward. However, maintaining the correct position within the intersection proves to be more difficult during turns. The path maintainer feedback device proves useful in maintaining the proper latitudinal placement of the vehicle, but is largely reactive in nature. As a proactive measure, as in many traditional larger intersections, turning lane marks within the intersection may be useful.

4.4 Summary

This chapter presented the results and analysis from the experiments executed for the verification and validation of the synthetic environment. The first section discussed the baseline scenarios which were used to establish the ability of the synthetic environment to manage autonomous vehicles within an autonomous intersection. Next, the misbehavior scenarios and were discussed in detail along with the persistent human control observations made.

V. Conclusions and Observations

This chapter explores the conclusions and observations drawn from the experimental setup and analysis of the data collected in this research. Discussed here are the verification and validation scenarios used to establish the reservation-based autonomous intersection protocol within the synthetic environment. The chapter then discusses the preliminary observations obtained from the introduction of persistent human control in semi-autonomous vehicles. Finally, a future work section enumerates several areas of interest for additional research.

5.1 Verification of Baseline Environment

The verification phase of development was intended to ensure the reservation-based protocol meets the design requirements. The analysis shows that the protocol used in the synthetic environment is roughly equivalent to the protocol used in the Autonomous Intersection Management (AIM) Simulator developed by Dresner, given the assumptions made in section 3.4 [3]. There were no safety violations at any point during the outlined scenarios in section 3.6. The generated autonomous vehicles also experienced comparable delays as outlined in section 4.1.2. With the successful completion of the baseline verification scenarios, the validation scenarios were executed.

5.2 Validation of Proposed Protocol

In the validation phase of development, the focus is on determining if the protocol produced expected results in scenarios unique to the synthetic environment. The validation scenarios are outlined in sections 3.7 - 3.9 and the experimental results from these scenarios are detailed in section 4.2. The results and analyses show that the reservation-based autonomous intersection protocol operates as expected in the

synthetic environment. The results of these validation scenarios reinforce the viability of the reservation-based protocol as built. With the verification and validation scenarios completed successfully, the introduction of the human controls is possible.

5.3 Introducing Human Controls

It was shown that, with the feedback mechanisms presented, a manually controlled semi-autonomous vehicle may in fact safely approach, enter, traverse, and exit an autonomous intersection with no traditional traffic control signals as demonstrated in the developed synthetic environment. All control signals are communicated to the vehicle via vehicle-to-everything (V2X) communications, allowing the protocol to benefit from the increased efficiency of reservation-based protocols such as AIM.

5.4 Implications of Synthetic Environment

A synthetic environment able to host semi-autonomous vehicles has great promise to cybersecurity research. With autonomous vehicles and semi-autonomous vehicles communicating with infrastructure come unique cyber threats. Human-controlled or otherwise, the ability of the vehicles and intersections to ensure the integrity and authenticity of messages sent and received is vital because there are many cybersecurity implications unique to semi-autonomous vehicles. The ability to defend against man-in-the-middle attacks needs to be built in to the system. Sybil-type attacks could cause increased delays where a malicious actor is making false reservation requests and therefore delaying legitimate vehicles. Malicious software present on the vehicles could misrepresent data to the human driver and cause the driver to arrive at the intersection at the wrong speed, time, etc. Petit and Shladover enumerate a list of potential cyber attacks unique to autonomous vehicles in [27]. Many of these attacks involve controller area network (CAN) bus vulnerabilities.

CAN cybersecurity is another area of research that the synthetic environment can aid in exploring. It was shown that the AFTR Burner virtual world was able to communicate with human controls via CAN traffic. CAN research could benefit from the real-time realistic observations provided by the three-dimensional (3D) virtual world.

Further development of the human feedback mechanisms can also be accomplished within the synthetic environment. Determining the type and amount of feedback required, the best location for these mechanisms, and the manner in which a human interacts with the semi-autonomous vehicle are all topics to which the synthetic environment is able to provide answers.

5.5 Research Contributions

Identifying and visualizing the cyber threats facing a reservation-based autonomous intersection presented as validation scenarios are the subject of an accepted International Conference on Cyber Warfare and Security (ICCWS) paper. This work was presented at the 13th annual conference held in Washington, D.C.

A second paper was accepted for presentation at International Federation for Information Processing (IFIP) working group (WG) 11.10 international conference on Critical Infrastructure Protection (CIP). This second paper discussed the introduction of the persistent human control within the reservation-based intersection management protocol. It was presented at the 12th annual conference held in Arlington, VA.

5.6 Future Work

The reservation-based protocol presented in this thesis is a proof of concept. There are many aspects that need to be matured and further developed. Itemized below are

some questions that need to be answered to implement a reliable, safe, and efficient method for persistent human control:

- What is the minimum amount of information required for the driver to safely maintain control of the vehicle?
- Should the current velocity and goal velocity be displayed separately in digital format or together in an analog format?
- What is the optimal placement of these feedback devices?

Should the placement be the traditional in-dash location?

Should the placement be in a heads-up-device location in-line with the driver's view of the road?

- At what distance from the intersection should the vehicle assume control (if at all)...

if the driver does not initiate communications using a manual button?

if the driver does not maintain the appropriate velocity?

if the driver does not maintain the appropriate path?

- What is the ideal safety buffer zone within an intersection for a human-controlled vehicle considering both safety and efficiency?
- If the negative impact on efficiency is substantial, should there be geographical zones and/or times where persistent human control is allowed?
- How should legacy vehicles be handled in such an environment?

Networking together several instances of the virtual world could add fidelity to a more urban-like environment. The environment presented in this thesis consisted of

a single intersection with many assumptions abstracting much of the V2X protocols. Future work could be accomplished to integrate any number of virtual worlds which could communicate in order to simulate a larger scale system. In addition, incorporating tools such as OMNeT++ would add fidelity to the wireless communications between vehicles via V2X capabilities [28]. With this last addition, more realistic experiments could be conducted to test the reservation-based protocol using actual wireless V2X protocols.

Finally, introducing persistent human control has notionally been shown to be feasible in the designed synthetic environment given the feedback mechanisms and controls described. Additionally, the synthetic environment is highly configurable to test many different aspects of the reservation-based autonomous intersection protocol as well as the semi-autonomous and autonomous vehicles modeled within it. The future work topics proposed can each be studied and answered via the synthetic environment well before the rubber meets the road.

Appendix A. Collision Matrix

This matrix shows which paths within a four lane intersection obstruct each other. An “x” denotes a collision is possible, while an “o” represents two paths where no collision is possible. This assumes vehicles follow the correct path within the intersection.

Lane	Turn*	0			1			2			3			4			5			6			7			Total**
		L	S	R	L	S	R	L	S	R	L	S	R	L	S	R	L	S	R	L	S	R	L	S	R	
0	L	x	x	x	x	x	o	x	x	x	x	x	o	o	x	x	x	x	o	x	x	x	x	x	o	19
	S	x	x	x	x	x	o	x	o	o	x	x	x	o	o	x	o	x	o	x	o	o	x	x	x	15
	R	x	x	x	o	x	o	x	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	5
1	L	x	x	o	x	x	x	x	x	o	x	x	x	x	x	x	o	o	x	x	x	o	x	x	x	19
	S	x	x	x	x	x	x	x	x	o	x	o	o	x	x	x	o	o	x	o	x	o	x	o	o	15
	R	o	o	o	x	x	x	o	x	o	x	o	o	o	o	o	o	o	o	o	o	o	o	o	o	5
2	L	x	x	x	x	x	o	x	x	x	x	x	o	x	x	x	x	x	o	o	x	x	x	x	o	19
	S	x	o	o	x	x	x	x	x	x	x	x	o	x	o	o	x	x	x	o	o	x	o	x	o	15
	R	x	o	o	o	o	o	x	x	x	o	x	o	o	o	o	o	o	o	o	o	o	o	o	o	5
3	L	x	x	o	x	x	x	x	x	o	x	x	x	x	x	x	o	x	x	x	x	o	o	x	x	19
	S	x	x	o	x	o	o	x	x	x	x	x	x	o	x	o	x	o	o	x	x	x	o	o	x	15
	R	o	x	o	x	o	o	o	o	o	x	x	x	o	o	o	o	o	o	o	o	o	o	o	o	5
4	L	o	o	o	x	x	o	x	x	o	x	o	o	x	x	x	x	x	o	o	x	x	x	x	o	14
	S	x	o	o	x	x	o	x	o	o	x	x	o	x	x	x	x	x	o	x	o	o	x	x	x	15
	R	x	x	o	o	x	o	x	o	o	o	o	o	o	x	x	x	o	x	o	o	o	o	o	o	9
5	L	x	o	o	o	o	o	x	x	o	x	x	o	x	x	o	x	x	x	x	x	o	o	x	x	14
	S	x	x	o	x	o	o	x	x	o	x	o	o	x	x	x	x	x	x	x	x	o	x	o	o	15
	R	o	o	o	x	x	o	o	x	o	x	o	o	o	o	o	x	x	x	o	x	o	x	o	o	9
6	L	x	x	o	x	o	o	o	o	o	x	x	o	o	x	x	x	x	o	x	x	x	x	x	o	14
	S	x	o	o	x	x	o	x	o	o	x	x	o	x	o	o	x	x	x	x	x	x	x	x	o	15
	R	x	o	o	o	o	o	x	x	o	o	x	o	x	o	o	o	o	o	x	x	x	o	x	o	9
7	L	x	x	o	x	x	o	x	o	o	o	o	o	x	x	o	o	x	x	x	x	o	x	x	x	14
	S	x	x	o	x	o	o	x	x	o	x	o	o	x	x	o	x	o	o	x	x	x	x	x	x	15
	R	o	x	o	x	o	o	o	o	o	x	x	o	o	x	o	x	o	o	o	o	o	x	x	x	9

*L = left turn, S = straight (no turn), R = right turn

**Number of obstructed paths

Bibliography

1. J. Lavina and C. Bonelli. (2017) ‘V2V Safety Technology Now Standard on Cadillac CTS Sedans’. [Online]. Available: <http://media.cadillac.com/media/us/en/cadillac/news.detail.html/content/Pages/news/us/en/2017/mar/0309-v2v.html>
2. K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham *et al.*, “Experimental security analysis of a modern automobile,” in *Security and Privacy (SP), 2010 IEEE Symposium on*. IEEE, 2010, pp. 447–462.
3. K. Dresner and P. Stone, “A multiagent approach to autonomous intersection management,” *Journal of Artificial Intelligence Research*, vol. 31, pp. 591–656, 2008.
4. *Surface Vehicle Recommended Practice*, Society of Automotive Engineers Std. J3016 SEP2016, 2016.
5. M. Matousek. (2017) ‘The most impressive things Tesla’s cars can do in Autopilot’. [Online]. Available: <http://www.businessinsider.com/tesla-autopilot-functions-and-technology-2017-12/#tesla-introduced-its-autopilot-feature-in-2015-it-gives-its-cars-the-ability-to-drive-autonomously-in-some-situations-last-year-the-company-rolled-out-new-hardware-on-all-of>
6. J. Shi, J. Wan, H. Yan, and H. Suo, “A survey of cyber-physical systems,” *2011 International Conference on Wireless Communications and Signal Processing, WCSP 2011*, 2011.
7. S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno *et al.*, “Comprehensive experi-

- mental analyses of automotive attack surfaces,” in *USENIX Security Symposium*. San Francisco, 2011.
8. M. McDonald, J. Mulder, B. Richardson, R. Cassidy, A. Chavez, N. Pattengale, G. Pollock, J. Urrea, M. Schwartz, W. Atkins *et al.*, “Modeling and simulation for cyber-physical system security research, development and applications,” *Sandia National Laboratories, Tech. Rep. Sandia Report SAND2010-0568*, 2010.
 9. M. Pahlavan, M. Papatriantafidou, and E. M. Schiller, “Gulliver: a test-bed for developing, demonstrating and prototyping vehicular systems,” in *Proceedings of the 9th ACM international symposium on Mobility management and wireless access*. ACM, 2011, pp. 1–8.
 10. Gulliver 8. [Online]. Available: <https://www.chalmers.se/hosted/gulliver-en/documents/vehicles>
 11. C. Parsons and S. Nykl, “An approach to simulate autonomous vehicles in urban traffic scenarios,” in *Proceedings of the 12th International IEEE Conference on Intelligent Transportation Systems, St. Louis, MO, USA, October 3-7, 2009*. IEEE, 2009, pp. 322–327.
 12. S. Nykl, C. Mourning, M. Leitch, D. Chelberg, T. Franklin, and C. Liu, “An overview of the steamie educational game engine,” in *Frontiers in Education Conference, 2008. FIE 2008. 38th Annual*. IEEE, 2008, pp. F3B21–25.
 13. C. Parsons and S. Nykl, “Real-time automated aerial refueling using stereo vision,” in *International Symposium on Visual Computing*. Springer, 2016, pp. 605–615.
 14. Khronos Group. (2016) OpenGL Overview. [Online]. Available: <https://www.khronos.org/about/#1>

15. G. Sharon and P. Stone, "A protocol for mixed autonomous and human-operated vehicles at intersections," in *The 2nd International Workshop on Agent-based Modeling of Urban Systems*. Sao Paulo: Springer, May 2017.
16. T. Au, S. Zhang, and P. Stone, "Autonomous intersection management for semi-autonomous vehicles," *The Routledge Handbook of Transportation*, pp. 88–104, 2015.
17. S. Jafarnejad, L. Codeca, W. Bronzi, R. Frank, and T. Engel, "A car hacking experiment: When connectivity meets vulnerability," in *Globecom Workshops (GC Wkshps), 2015 IEEE*. IEEE, 2015, pp. 1–6.
18. S. Corrigan, "Introduction to the controller area network (can)," *Texas Instruments, Application Report*, 2008.
19. P. Kleberger, T. Olovsson, and E. Jonsson, "Security aspects of the in-vehicle network in the connected car," in *Intelligent Vehicles Symposium (IV), 2011 IEEE*. IEEE, 2011, pp. 528–533.
20. S. Woo, H. J. Jo, and D. H. Lee, "A practical wireless attack on the connected car and security protocol for in-vehicle can," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 993–1006, 2015.
21. C. Laurendeau and M. Barbeau, "Threats to security in dsrc/wave," in *International Conference on Ad-Hoc Networks and Wireless*. Springer, 2006, pp. 266–279.
22. J. Doubek. 'Study Backs Getting Driverless Cars On The Road, As Waymo Ditches Backup Drivers'. [Online]. Available: <https://www.npr.org/sections/alltechconsidered/2017/11/10/563002343/study-backs-getting-driverless-cars-on-the-road-as-waymo-ditches-backup-drivers>

23. G. Munster. (2017) Auto Outlook 2040: The Rise of Fully Autonomous Vehicles. [Online]. Available: <http://loupventures.com/auto-outlook-2040-the-rise-of-fully-autonomous-vehicles/>
24. H. Serrano. (2016) Game Engine Development: Tips for developing a Collision Detection System. [Online]. Available: <https://www.haroldserrano.com/blog/tips-for-developing-a-collision-detection-system>
25. R. Smith, *Open Dynamics Engine v0.5 User Guide*, 2006. [Online]. Available: <http://ode.org/ode-latest-userguide.pdf>
26. J. R. Douceur, "The Sybil attack," *Lecture Notes in Computer Science*, vol. 2429, pp. 251–260, 2002.
27. J. Petit and S. E. Shladover, "Potential cyberattacks on automated vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 546–556, 2015.
28. OpenSim Ltd. (2015) What is OMNeT++? [Online]. Available: <https://www.omnetpp.org/intro>

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 22-03-2018		2. REPORT TYPE Master's Thesis		3. DATES COVERED (From — To) Aug 2016 – Mar 2018	
4. TITLE AND SUBTITLE Mitigating the Effects of Cyber Attacks and Human Control in an Autonomous Intersection			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
6. AUTHOR(S) Bentjen, Karl C., Captain, USAF					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENG-MS-18-M-008		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory, Sensors Directorate Attn: Lt Col Patrick J. Sweeney 2241 Avionics Circle Wright-Patterson AFB, OH 45433 (937) 938-4252 (DSN: 713-4252) Patrick.Sweeney@us.af.mil			10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/Rywa		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.					
13. SUPPLEMENTARY NOTES This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.					
14. ABSTRACT Widespread use of fully autonomous vehicles is near. The desire for a human to maintain control, even if limited, of a vehicle will likely never fully subside. Protocols to safely and efficiently manage reservation-based intersections with a mixture of fully autonomous, semi-autonomous, and non-autonomous vehicles exist. Missing from these protocols is persistent human control of semi-autonomous vehicles in approaching and navigating autonomous intersections without the use of traditional signals. This paper attempts to lay the foundation for the necessary extensions required for human control in semi-autonomous vehicles. Desired is a protocol that maintains the benefits in efficiency of a fully autonomous environment, such as AIM, while allowing persistent human control of a vehicle. The AFTR Burner three-dimensional virtual world offers the ability to model this physics based synthetic environment in a highly predictable and realistic manner. The preliminary observations suggest that persistent human control is a possibility among reservation-based autonomous intersections, but further research must be done to determine its viability.					
15. SUBJECT TERMS Autonomous Intersection Management, Semi-autonomous Vehicles, Human Control, V2X Communication, AFTR Burner					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Scott Graham (ENG)
U	U	U	UU	120	19b. TELEPHONE NUMBER (include area code) (937) 255-6565, x4581 scott.graham@afit.edu