Chapter 11

PERSISTENT HUMAN CONTROL IN A RESERVATION-BASED AUTONOMOUS INTERSECTION PROTOCOL

Karl Bentjen, Scott Graham and Scott Nykl

Abstract Widespread use of fully autonomous vehicles is near. However, the desire of human beings to maintain control of their vehicles – even limited control – is unlikely to ever go away. Several protocols (e.g., AIM, Semi-AIM and H-AIM) have been developed to safely and efficiently manage reservation-based intersections with a mixture of fully autonomous, semi-autonomous and non-autonomous vehicles. However, these protocols do not incorporate the dynamic of a human maintaining control of a semi-autonomous vehicle when approaching and crossing an intersection. This chapter lays the foundation for the extensions required for human-control of semi-autonomous vehicles, the ultimate goal being a protocol that maintains the efficiency of a fully autonomous environment while allowing human control of vehicles when navigating an intersection. This chapter also proposes information feedback mechanisms for human response, such as displays that provide the intersection arrival time, goal velocity, lane maintaining assistance and other warnings. Additionally, it describes a synthetic environment that enables the testing of intersection protocols that support human interaction.

Keywords: Semi-autonomous vehicles, intersections, reservations, human control

1. Introduction

Self-driving vehicles are already on the road, in some cases without backup drivers [4]. Before long, the traffic infrastructure, specifically intersections, will be required to manage autonomous vehicular traffic in an efficient manner. To address this need, Dresner and Stone [8] introduced a reservation-based intersection protocol called Autonomous Intersection Management (AIM), designed for an environment with strictly autonomous vehicles. The AIM protocol was subsequently modified to incorporate semi-autonomous vehicles that allow lim-

J. Staggs and S. Shenoi (Eds.): Critical Infrastructure Protection XII, IFIP AICT 542, pp. 197–212, 2018.
https://doi.org/10.1007/978-3-030-04537-1_11

ited human control [2]. Another modified version of AIM, known as Hybrid-AIM (H-AIM), further accommodates human-operated vehicles without direct communications between vehicles and the intersection [10]. An additional category to be considered is vehicles that can communicate with the intersection manager, but that are driven by humans.

This chapter lays the foundation for the extensions required for human-control of semi-autonomous vehicles, the ultimate goal being a protocol that maintains the efficiency of a fully autonomous environment while allowing human control of vehicles when navigating an intersection. In particular, it attempts to identify how persistent human control can be introduced in an autonomous intersection. It also describes the AFTR Burner synthetic environment [9] and baseline experiments that establish the viability of the reservation-based intersection protocol. Proposed feedback and control mechanisms to enable human control are also detailed, along with a proof-of-concept system that enables humans to maintain vehicular control when navigating autonomous intersections.

2. Background and Motivation

This section provides an overview of autonomous, semi-autonomous and non-autonomous vehicles. Also, it discusses the requirements for the safe and efficient management of a traffic intersection with autonomous and semi-autonomous vehicles, as well as for an environment where all the vehicles are fully autonomous with no human control.

2.1 Autonomous Vehicle Taxonomy

The U.S. National Highway Traffic Safety Administration (NHTSA) and the Society of Automotive Engineers (SAE) International have developed a taxonomy of vehicles and their levels of autonomy [11]. Table 1 lists the five levels of autonomy and provides brief descriptions. Levels 4 and 5 cover autonomous vehicles that are capable of driving themselves; however, these levels provide options for human drivers to assume control of their vehicles.

This research specifically focuses on the ability – or desire – of a human to maintain control of a vehicle, especially when approaching and traversing an intersection. The intersection is designed such that traditional or legacy non-autonomous vehicles (Level 0) are not normally allowed due to the lack of traffic signals at the intersection and/or the vehicles lack vehicle-to-anything (V2X) communications capabilities. If desired, the intersection may be designed to degrade to a standard intersection when a legacy vehicle approaches, but this problem is outside of scope of this research.

2.2 Reservation Concept

Human safety is paramount when designing a protocol for managing autonomous traffic. Dresner and Stone [7] introduced the concept of a reserva-

Table 1. Automation levels set by SAE International [11].

Automation Level	Description
Human Driver Required	
Level 0: No Automation	The vehicle is completely non-autonomous; the driver performs all the driving tasks.
Level 1: Driver Assistance	The vehicle has some driving assist features such as traditional cruise control, but the driver controls the vehicle.
Level 2: Partial Automation	The vehicle has combined automated functions such as acceleration and steering, but the driver must remain engaged with the driving task and monitor the environment at all times.
Level 3: Conditional Automation	The 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 Not Required	
Level 4: High Automation	The vehicle can perform all the driving functions under certain conditions, including limitations on locations and environments; the driver may have the option to control the vehicle.
Level 5: Full Automation	The vehicle can perform all the driving functions under all conditions; the driver may or may not have the option to control the vehicle.

tion to address the issue of safely scheduling the passage of autonomous vehicles through an intersection. They used the reservation concept to develop the AIM protocol that can manage an autonomous intersection in a safe and efficient manner. This is accomplished by ensuring that vehicles do not collide and by reducing the delays experienced by vehicles at the intersection compared with a traditional intersection with traffic signals [8].

The AIM protocol uses the reservation concept to safely eliminate traffic signals as long as all the vehicles are fully autonomous with V2X capabilities. The AIM protocol works well in an environment comprising only fully autonomous vehicles. However, while such an environment will surely be realized in the future, there will be a long transition period during which vehicles with all levels of autonomy will have to be integrated safely and efficiently.

The AIM protocol is designed for an environment where at least 90% of the vehicles are fully autonomous and operate without human control. It is speculated that autonomous vehicles with V2X capabilities will not exceed 90% of the vehicular population until at least the year 2045 [3]. Until this time, protocols will be implemented to handle the integration of all levels of autonomous vehicles. All the autonomous vehicles will incorporate human control to some extent.

2.3 Other Intersection Protocols

In 2015, Au et al. [2] published the SemiAIM protocol, an extension of the AIM protocol that incorporates semi-autonomous vehicles. In the SemiAIM protocol, human drivers relinquish control of their vehicles before entering an intersection. However, the protocol requires the use of traffic signals. Semi-autonomous vehicles that fail to receive confirmed reservations must come to a stop and treat the intersection as a traditional traffic signal intersection. The traffic signals are also used by non-autonomous vehicles, which are allowed in the SemiAIM protocol.

Sharon and Stone [10] developed the H-AIM protocol, which is more efficient than the AIM protocol when there is a low concentration of autonomous and semi-autonomous vehicles. The enhanced protocol assumes that the intersection can detect incoming non-autonomous vehicles and enables autonomous vehicles to receive reservations that do not conflict with the possible paths of non-autonomous vehicles. The protocol also depends on traffic signals for human-driven vehicles that do not have V2X communications capabilities.

The SemiAIM and H-AIM protocols do not provide the option for a human-driven vehicle with V2X capabilities to request and receive a reservation, but they do allow a human to maintain persistent control over his/her vehicle. However, a vehicle at Level 2 or higher automation level permits a human to control the steering and/or velocity while navigating through an autonomous intersection without traffic signals.

2.4 Persistent Human Control

It is safe to assume that there will always be humans who want to drive their vehicles and be in control. Indeed, a recent study by Abraham et al. [1] revealed that 48% of the people surveyed would never purchase a car that completely drives itself. Therefore, it is necessary to consider a future environment where all the vehicles are at least semi-autonomous (whether they require a human driver or not), all the vehicles have vehicle-to-vehicle (V2V) and V2X communications capabilities, and autonomous intersections do not have traditional traffic signals. Some sort of backup signal capability may exist, but not for managing traffic on a regular basis.

A protocol such as AIM could prove to be the protocol of choice in such an environment, especially if, like the AIM protocol, it is already shown to be safe and efficient. However, the protocol would have to be modified to enable the human behind the wheel to maintain control over the steering and/or velocity of the vehicle. The majority of the protocol changes would occur at the vehicle side of transactions instead of at the intersection side. Au et al. [2] discuss some of the feedback and control features that would be necessary to implement the modifications. In fact, they recommend the use of a "button" to make a reservation request and an "OK" indicator that would tell the driver to relinquish control of the semi-autonomous vehicle before it enters an intersection.

This control dynamic shared by the human and the intersection gives rise to a form of blended control. The intersection dictates when a specific reservation is possible, but the human has ultimate control over the movement of the vehicle. Of course, the vehicle would have to provide feedback such as lane-keeping and velocity warnings to maintain the tight trajectory constraints.

2.5 Synthetic Environment

Developing protocols that blend human control with automated systems requires an environment in which testing can be conducted safely. This research selected the three-dimensional (3D) virtual world called AFTR Burner, the successor to the STEAMiE engine, which utilizes the Open Dynamics Engine (ODE) for physics simulation and collision detection [9]. Incorporating the physics engine in the testing environment enables the human performance introduced by a protocol to be demonstrated and evaluated without endangering human participants and without incurring significant costs.

3. Proposed Design

This section establishes the viability of the proposed synthetic environment for handling an intersection that manages autonomous vehicular traffic. It also details the key protocol components that support human control in the environment. In particular, the section discusses the assumptions and the reservation concept, establishes a baseline and identifies the features required for persistent human control. It includes details about the reservation concept from the AIM protocol for safety, which is the primary goal. Also, it discusses details about V2X communications such as message timing and content, and feedback features needed to guide human drivers safely through an intersection.

3.1 Assumptions

The following assumptions were made in this research:

- **Latency:** The messaging protocol is abstracted to function calls between the vehicle and intersection world object classes. Latency (delay) between a sender and receiver is not modeled and is, therefore, assumed to be zero.

- **Signal Loss:** Signal loss in V2X communications is not modeled in the synthetic environment. While the potential for lost communications is real, this topic is left for future research.

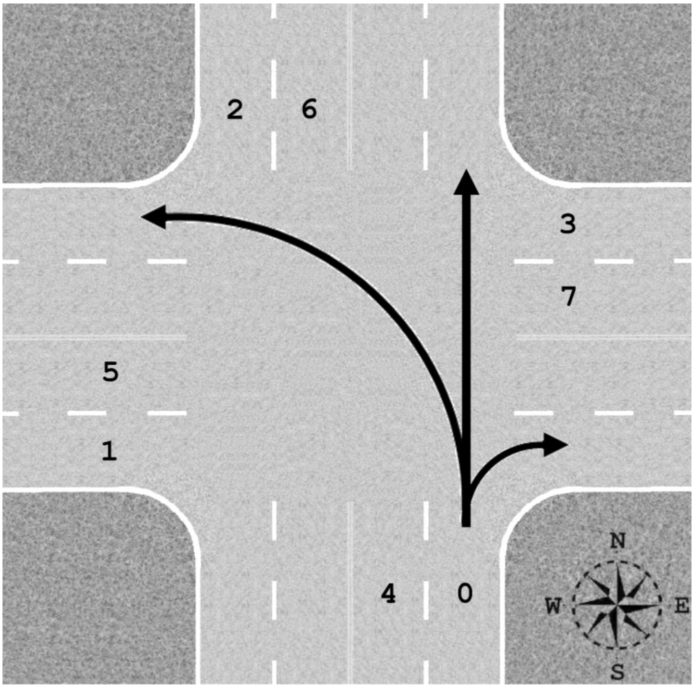

Figure 1. Legal turn direction options from a single lane with lane numbers.

- **Static Lanes:** No lane changes are permitted in an intersection. Turning vehicles move into their respective destination lanes (e.g., left turns from inside lanes terminate at inside lanes). Figure 1 illustrates the possible turning directions from each northbound lane. Note that Lane 0 northbound may turn into Lane 1 eastbound, but not to Lane 5 eastbound.

- **Turning Paths:** Turning paths are smooth or uniform, not abrupt or sharp (Figure 1).

- **Safety Buffer:** In an intersection, the occupied region includes a buffer of approximately 25% of the vehicle length and width for human-operated and autonomous vehicles. This parameter could be the subject of future research that balances safety and efficiency.

- **Stopping Distance:** The stopping distance is set to 25 m from the beginning of the intersection in every direction. This distance represents the beginning of a region where a vehicle must stop if no reservation is confirmed. A vehicle in this region is expected to follow its confirmed reservation.

- **Bounding Box:** The vehicular collision detection system uses a rectangular prism bounding box. This type of bounding box simplifies the

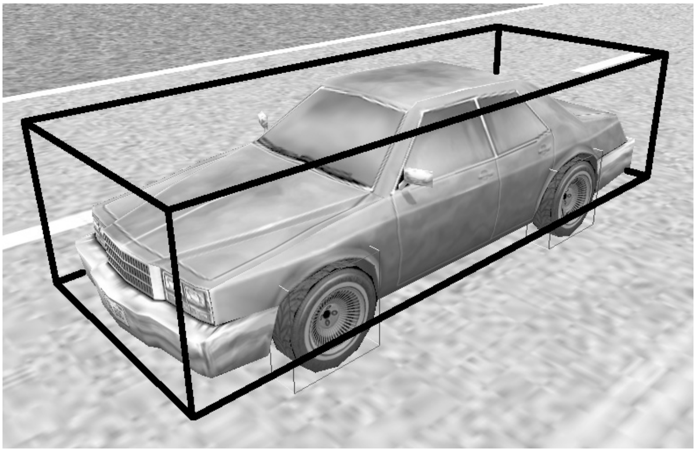

Figure 2. Bounding box of a sedan in the synthetic environment.

computations of the physics engine during collision detection while representing the shapes of vehicles. Figure 2 shows the bounding box of a sedan in the synthetic environment.

- **Velocity:** The maximum velocity before and after an intersection is 8 m/s for fully autonomous vehicles in the synthetic environment. Although this constraint could be relaxed in a future implementation, fully autonomous vehicles are assumed to have a constant velocity of 8 m/s in all directions at an intersection.

- **Single Intersection:** The synthetic environment has a single intersection with two inbound lanes from each cardinal direction.

- **Ambient Environment:** The synthetic environment has no obstructions – visual or otherwise (i.e., the environment is clear with high visibility).

- **Vehicles Only:** The synthetic environment has no obstacles, except for the intersection and other vehicles (i.e., no cyclists, pedestrians, animals or other moving entities).

- **Reservation Order:** A vehicle in a lane may request a reservation if and only if the vehicle directly in front of it already has a reservation. This is determined and enforced via V2V communications.

3.2 Reservations

The primary goal of an autonomous intersection without a traditional signaling system (traffic lights) is safety. Dresner and Stone [6] introduced the

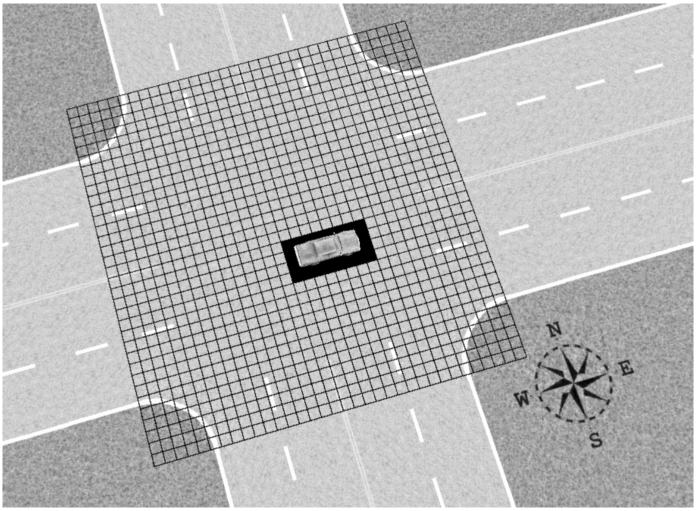

Figure 3. Reservation grid at an instant in time.

reservation concept used here. It divides the intersection into an arbitrary number of squares. The number of squares in each dimension is called the granularity of the reservation grid. A granularity of 34 was employed in this research – an intersection was divided into $34 \times 34 = 1,156$ individual squares.

Figure 3 shows a visual representation of the reservation grid in the synthetic environment at an instant in time when a vehicle was passing through the intersection. The darkened squares represent the space occupied by the vehicle.

The intersection maintains all the reservations for all the vehicles. When a vehicle makes a request, the arrival time, arrival lane, departure lane and velocity are used to determine if, at any instant in time, the proposed reservation overlaps with one or more previously-confirmed requests. If an overlap exists, then the request is denied. Interested readers are referred to [8] for a detailed description of the reservation system.

The reservation system is key to maintaining safety because no reservations are granted if overlaps are detected. The system also eliminates the need for traditional traffic signals because reservations are granted instead of shining green lights. In fact, this approach is highly efficient compared with traditional traffic signals and signs [8].

3.3 Synthetic Environment

The AFTR Burner virtual world was chosen to create the synthetic environment used in this research. Naturally, an algorithm for managing traffic that approaches an intersection is required as well. Although the AIM protocol has not been used in its entirety, many of its concepts are incorporated in a simplified manner in the synthetic environment. For example, the notion of the time-space reservation grid for the intersection manager and the pseudocode for the driver agent managing the autonomous vehicle and messaging protocol are both adapted to the synthetic environment.

A baseline comparison was performed to demonstrate that the intersection management algorithm is viable for handling fully autonomous vehicles. The comparison was conducted between the synthetic environment and the AIM simulator developed by Dresner and Stone [8]. The total average delay experienced and the number of safety violations (i.e., collisions) were selected as response variables in order to determine viability.

The baseline incorporated five trials for each of three traffic levels – 100, 200 and 300 vehicles per lane per hour. The maximum vehicular speed was set to 8 m/s and vehicles in the AIM simulator were limited to sedans. Using the data collection feature of the AIM simulator, the same traffic patterns were used in the synthetic environment to match the response variables in the trials. The relevant data collection items contained in the output included vehicle identification numbers, vehicle generation times, starting lane identifiers, destination identifiers, as well as the simulation exit times for autonomous vehicles.

After the fifteen trials in the AIM simulator and the synthetic environment were completed, the data collection files were compared. A two-tailed z-test with a significance level of $\alpha = 0.05$ was employed.

Table 2 summarizes the results. In every trial, the p-value is at least 0.05. Therefore, the null hypothesis H_o that the total average delays experienced in the AIM simulator and the synthetic environment are the same fails to be rejected. These results suggest that under the assumptions made, the implementation of the intersection manager, which is modeled after the AIM protocol, is roughly equivalent in its operation.

3.4 Messaging

The types of messages exchanged by a vehicle and intersection are modeled closely after those developed by Dresner and Stone [8]. Request messages are sent from a vehicle to the centralized intersection road-side unit (RSU). These messages provide information about the proposed arrival time, starting lane, destination direction and vehicle type. The vehicle type also includes the vehicle size and whether the vehicle is human-controlled. This modification enables the intersection to increase the safety buffer size around a vehicle to provide more flexibility with regard to arrival times and velocities. The road-side unit responds to a vehicle with confirmation messages, rejection messages and acknowledgement messages. A vehicle also can send cancellation messages.

The ovals represent the starting and ending states, rectangles represent processes or actions taken, and diamonds represent decisions. Dashed lines with arrows indicate (wireless) communications between the vehicle and intersection manager.

Table 2. Synthetic environment intersection management baseline results.

H_o: The total average delays experienced in the AIM simulator and the synthetic environment are the same.
H_a: The total average delays experienced in the AIM simulator and the synthetic environment are not the same.

100 Vehicles/Lane/Hour	**Trial 1**	**Trial 2**	**Trial 3**	**Trial 4**	**Trial 5**
AIM Simulator Delay (s)	0.1600	0.1648	0.1562	0.1568	0.1418
Synthetic Environment Delay (s)	0.1846	0.1503	0.1579	0.1933	0.1449
Two-Tailed z-Test p-Value	0.55	0.73	0.97	0.47	0.93
200 Vehicles/Lane/Hour	**Trial 1**	**Trial 2**	**Trial 3**	**Trial 4**	**Trial 5**
AIM Simulator Delay (s)	0.3846	0.3767	0.3690	0.2601	0.2764
Synthetic Environment Delay (s)	0.3327	0.3818	0.2813	0.3152	0.3048
Two-Tailed z-Test p-Value	0.40	0.93	0.05	0.28	0.55
300 Vehicles/Lane/Hour	**Trial 1**	**Trial 2**	**Trial 3**	**Trial 4**	**Trial 5**
AIM Simulator Delay (s)	0.6619	0.5724	0.5479	0.6420	0.6019
Synthetic Environment Delay (s)	0.5698	0.6905	0.5334	0.6782	0.5298
Two-Tailed z-Test p-Value	0.23	0.17	0.83	0.71	0.34

3.5 Human Controls and Feedback Displays

Enabling a human to maintain control of a semi-autonomous vehicle while navigating an intersection without traditional signals is not a trivial problem. In addition to traditional controls such as an accelerator, brake pedal, steering wheel, turn signals, mirrors and speedometer (to list a few), controls and/or displays must be provided to enable persistent human control at an autonomous intersection. This section describes the additional controls and displays that are required.

Figure 4 shows a high level view of the message decision making flow between an autonomous vehicle and intersection.

Currently, traditional road signs communicate information to drivers about upcoming hazards and roadway features such as sharp bends, intersections and speed limits. An in-dash indicator that notifies a human driver of an intersection that has come into range would be needed. In addition to this indicator, a button option [2] could be implemented to initiate vehicle communications with the autonomous intersection. After the button is pressed by the driver, the vehicle communicates with the intersection to arrange a reservation for passage, initiating the messaging flow shown in Figure 4.

When there is traffic congestion, the reservation request sent by a vehicle may be denied. In this scenario, there is a requirement to inform the driver

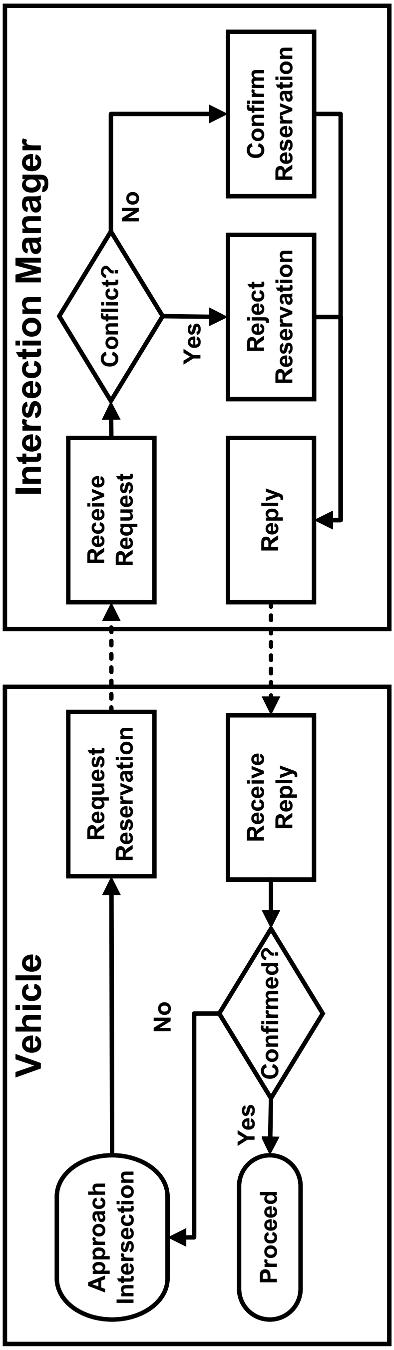

Figure 4. High-level view of a reservation messaging process.

about the reservation denial. For simplicity, a reservation denial would require 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 would continue until a successful reservation is made before entering the intersection.

After a reservation is made, the information supplied to the vehicle must be displayed to the driver. An indicator is needed to show that the reservation has been made for the desired path and the velocity to be maintained in the intersection. This requires a mechanism that communicates to the driver the goal velocity at arrival and/or the velocity needed to arrive at the required time, along with the velocity to be maintained in the intersection. This feedback mechanism is pivotal to ensuring that the vehicle arrives at and traverses through the intersection at the correct times. The indicator must continuously update the goal velocity based on the current time, arrival time and distance to the intersection. The indicator may also be used to communicate the goal velocity to be maintained in the intersection.

Finally, regardless of the vehicle velocity, the human must be able to maintain the correct lateral control of the vehicle, especially in the intersection. The corresponding path maintainer feedback indicator would notify the driver if the vehicle is too far left or right from the center of the current lane, along with the designated path through the intersection.

Table 3 summarizes the human controls and feedback devices required for persistent human control. The next section discusses the manner in which feedback information should be displayed to human drivers. Armed with the human controls and feedback devices, a driver would able to safely enter and navigate an autonomous intersection. Due to the security concerns, the synthetic environment provides the best venue for evaluating the ability of humans to safely traverse an autonomous intersection.

4. Experimental Observations

This section discusses the observations made when testing the proposed protocol that leverages additional human control and feedback devices. The experiments described in this section are notional and serve as proofs-of-concept instead of actual tests involving human subjects.

Figure 5 shows a screenshot of the synthetic environment with the human feedback mechanisms mentioned above. The screens outlined with thick black borders mimic the side-view and rear-view mirrors. The remaining displays present feedback information. On the left-hand side and moving from top to bottom are: the current time in the simulation, the arrival time of the confirmed reservation, the current velocity (m/s) and the goal velocity (m/s). The compass in the upper center of the screen displays one of the eight cardinal or intercardinal headings (i.e., N, NE, E, etc.). On the right-hand side of the screen and moving from top to bottom are: the current simulation name (used for reference purposes), the reservation status indicator and the digital lateral offset.

Table 3. Human controls and feedback devices required for persistent human control.

Item	**Description**
In-Range Indicator	This device informs the driver that an autonomous intersection is within range.
Request Reservation Button	This device initiates V2X communications to request a reservation from the intersection.
Denied Reservation Indicator	This device informs the driver that the requested reservation was denied.
Granted Reservation Indicator	This device informs the driver that the requested reservation was successful and provides the assigned velocity in the intersection.
Goal Velocity Indicator	This active device informs the driver of the velocity to be maintained to keep the reservation; the device may also be used to maintain the correct velocity in the intersection.
Maintain Path Indicator	This active feedback device informs the driver about the left/right position correctness based on the lane or planned path in the intersection.

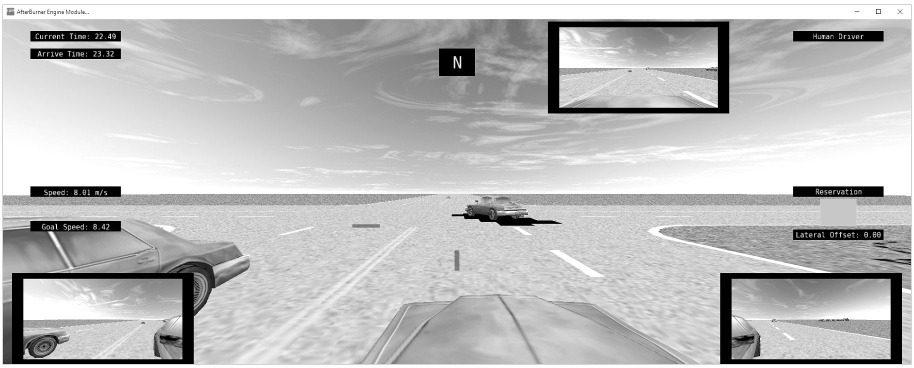

Figure 5. Screenshot of the synthetic environment with human control.

Table 4. Reservation feedback mechanism states.

Vehicle Location	**Reservation Status**	**Indicator Color**
Out of Range	N/A	Clear
In Range	Unconfirmed	Yellow
In Range	Confirmed	Green
Within Stopping Distance	Unconfirmed	Red

Table 4 presents the reservation feedback mechanism states. The driver in the experiment with a current speed of 8.01 m/s was required to increase the velocity slightly to arrive at the intersection on time, as indicated by the goal speed (8.42 m/s) in the feedback display.

Analog versions of the goal speed and lateral feedback indicators are presented on a heads-up-display (HUD) in the direct line of sight of the human driver. The goal velocity is indicated by a green box that hovers around a horizontal black line. If the goal velocity is higher than the current velocity, then the green box hovers above the line; the green box hovers below the line if the goal velocity is lower than the current velocity.

The analog lateral feedback operates similarly. If the human veers to the right or left of the planned path, then the green box hovers horizontally to the left or right of the vertical black line, respectively.

Extreme deviations from the goal velocity and vehicle path turn the green box to a red box. The mechanisms in the heads-up-display are translucent to minimize obstructions to the driver's view.

The design and placement of feedback devices are important. The digital speedometer and goal speed indicator should be close to each other. In fact, an analog display may be better than a digital display. An analog speedometer could have the goal velocity indicated in a separate colored dial located directly above the current velocity dial. Drivers may prefer to have the option of choosing digital versus analog as well. Extensive testing is required to determine the optimal design and placement of the feedback devices.

Maintaining the center of the correct lane appears to be a straightforward task. However, maintaining the correct position in an intersection is more difficult. The path maintainer feedback device helps keep the proper placement of the vehicle, but it is largely reactive in nature. As a proactive measure, it would be prudent to mark the paths of turning lanes, as is done in many traditional intersections.

5. Conclusions

This chapter has laid the foundation for the extensions required for human-control of semi-autonomous vehicles, the ultimate goal being a protocol that maintains the efficiency of a fully autonomous environment while allowing human control of vehicles when navigating an intersection. The reservation-based autonomous intersection protocol derived from the AIM protocol [8] and implemented in the synthetic environment proved to be roughly equivalent to the AIM protocol given the assumptions made; this result was established by the baseline experiments. The limited feedback mechanisms enable a manually-controlled, semi-autonomous vehicle to safely approach, enter, traverse and exit an autonomous intersection, despite the fact that the intersection does not have traditional traffic control signals. In such a scenario, all the control signals must be transmitted to the vehicle via V2X communications at a rate of up to ten signals per second.

Introducing persistent human control has been shown to be feasible given the feedback mechanisms and controls. The AFTR Burner virtual world provides an appropriate synthetic environment. This highly-configurable synthetic environment supports extensive testing of the reservation-based autonomous intersection protocol as well as the integration of semi-autonomous vehicles.

Future research will attempt to determine the minimum amount of information required for a human driver to safely maintain vehicular control, and the optimal types and placement of the driver interaction and feedback mechanisms. Other research topics include maintaining vehicular velocities and paths, establishing safety buffer zones and integrating autonomous, semi-autonomous and legacy vehicles in a busy intersection while ensuring safe and efficient traffic flow.

Note that the views expressed in this chapter are those of the authors and do not reflect the official policy or position of the U.S. Air Force, U.S. Army, U.S. Department of Defense or U.S. Government.

References

[1] H. Abraham, B. Reimer, B. Seppelt, C. Fitzgerald, B. Mehler and J. Coughlin, Consumer Interest in Automation: Preliminary Observations Exploring a Year's Change, White Paper 2017-2, AgeLab, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2017.

[2] T. Au, S. Zhang and P. Stone, Autonomous intersection management for semi-autonomous vehicles, in *Routledge Handbook of Transportation*, D. Teodorovic (Ed.), Routledge, New York, pp. 88–104, 2016.

[3] P. Bansal and K. Kockelman, Forecasting Americans' long-term adoption of connected and autonomous vehicle technologies, *Transportation Research Part A: Policy and Practice*, vol. 95, pp. 49–63, 2017.

[4] J. Doubek, Study backs getting driverless cars on the road, as Waymo ditches backup drivers, *National Public Radio*, November 10, 2017.

[5] J. Douceur, The Sybil attack, *Proceedings of the First International Workshop on Peer-to-Peer Systems*, pp. 251–260, 2002.

[6] K. Dresner and P. Stone, Multiagent traffic management: A reservation-based intersection control mechanism, *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 530–537, 2004.

[7] K. Dresner and P. Stone, Multiagent traffic management: An improved intersection control mechanism, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 471–477, 2005.

[8] K. Dresner and P. Stone, A multiagent approach to autonomous intersection management, *Journal of Artificial Intelligence Research*, vol. 31, pp. 591–656, 2008.

[9] S. Nykl, C. Mourning, M. Leitch, D. Chelberg, T. Franklin and C. Liu, An overview of the STEAMiE educational game engine, *Proceedings of the Thirty-Eighth Annual Frontiers in Education Conference*, pp. F3B-21–F3B-25, 2008.

[10] G. Sharon and P. Stone, A protocol for mixed autonomous and human-operated vehicles at intersections, in *Autonomous Agents and Multiagent Systems (AAMAS 2017)*, G. Sukthankar and J. Rodriguez-Aguilar (Eds.), Springer, Cham, Switzerland, pp. 151–167, 2017.

[11] Society of Automotive Engineers International, Surface Vehicle Recommended Practice, J3016 SEP2016, Warrendale, Pennsylvania, 2016.