REALISTIC SCENARIOS FOR MIXED AUTONOMOUS AND HUMAN VEHICLE INTERSECTION MANAGEMENT

A Thesis

by

AARON PARKS-YOUNG

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Chair of Committee,	Guni Sharon
Committee Members,	Sivakumar Rathinam
	Dylan Shell
Head of Department,	Scott Schaefer

December 2020

Major Subject: Computer Engineering

ABSTRACT

Autonomous vehicles capture people's imagination, but the road to widespread deployment of such vehicles has many challenges ahead. One challenge is the likely eventuality of interaction between autonomous and human operated vehicles at intersections. A large body of work exists addressing this problem, but a common theme among that work is assumption of somewhat artificial intersection conditions and traffic scenarios. The Hybrid Autonomous Intersection Management (H-AIM) protocol is an example of a solution that was created while using artificial conditions. H-AIM is designed to allow human as well as connected and autonomous vehicles (CAVs) to coexist while efficiently traversing intersections.

This thesis discusses an extension to H-AIM to allow the protocol to cope with some uncertainty that is present in modern intersections due to actuated traffic signal control. This thesis also demonstrates the results of simulations of the performance of H-AIM using traffic scenarios created from real traffic data for multiple intersections that are modeled on real intersections in the state of Utah. Sets of simulations show the capabilities of H-AIM with various levels of CAV market penetration. The empirical study performed also demonstrates the effectiveness in simulation of adaptive signal timing, different turning movement profile assignments with fixed signal timing, and signal actuation in terms of delay improvement when used in conjunction with the H-AIM protocol.

The results from simulations in this thesis suggest that mixtures of allowed turning movements by lane that are more permissive for CAVs and less permissive for human operated vehicles are often detrimental in terms of delay. This thesis also demonstrates that a mixture of adaptive signal timing and traffic signal actuation can act cooperatively with the H-AIM protocol in order to further reduce delay when compared to fixed traffic signal timing schemes. Thus, particularly for low CAV market penetration percentages, H-AIM and other substantially similar protocols should be used in conjunction with homogeneous turning movement profiles. Combination with modern optimization and delay reduction schemes such as traffic signal actuation is also beneficial.

DEDICATION

To God, for I am greatly blessed. To my wife, thank you for your patience, understanding, love, support, and occasional proofreading. To my family, your support, love, and encouragement has been invaluable. To my many mentors and teachers, thank you for your time and effort invested. To my friends, you're part of me and thank you for often distracting me when I needed a break.

ACKNOWLEDGMENTS

I'd like to acknowledge and thank Guni Sharon, Stephen Boyles, and Peter Stone whose prior research formed the basis for this thesis. Thanks to the Utah Department of Transportation for much of the data used in this research. And thank you to Texas A&M's High Performance Research Computing group for making it possible to crunch a vast amount of data as well as providing support and resources for special computing needs.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This thesis was supervised by a thesis committee consisting of Professors Guni Sharon and Dylan Shell of the Department of Computer Science and Engineering and Professor Sivakumar Rathinam of the Department of Mechanical Engineering.

All work for the thesis was completed by Aaron Parks-Young, in coordination with and under the advisement of Dr. Guni Sharon of the Department of Computer Science and Engineering.

Funding Sources

There are no outside funding contributions to acknowledge related to the research and compilation of this document.

NOMENCLATURE

AIM	Autonomous Intersection Management
ATSPM	Automated Traffic Signal Performance Measures
CAV	Connected and Autonomous Vehicle
CV	Connected Vehicle
DoT	Department of Transportation
H-AIM	Hybrid Autonomous Intersection Management
HV	Human Operated Vehicle
IM	Intersection Manager
V2V	Vehicle-to-Vehicle
V2I	Vehicle-to-Infrastructure

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

TABLE Page

1. INTRODUCTION AND LITERATURE REVIEW

1.1 Introduction

Personal vehicle communication hardware and protocols are becoming more sophisticated and capable as vehicles transition towards becoming fully autonomous. Though, it's infeasible to expect an instantaneous and complete transition from human driven vehicles to autonomous ones. This means that on top of conquering the technological hurdles of autonomous vehicles themselves, we must also consider the practical concerns of human driven and autonomous vehicles coexisting.

One major such consideration is how vehicles will interact at intersections. Connected autonomous vehicles (CAVs) have been shown to be capable of interacting with an intersection managing agent (or IM for intersection manager) and/or each other to greatly improve efficiency at intersections which service CAVs [1, 2, 3, 4]. While there are numerous approaches to intersection management including centralized [5], decentralized [6], and optimization [7] approaches a foundational work in the area is Dresner and Stone's 2008 paper [1] describing the Autonomous Intersection Management (AIM) protocol.

The AIM protocol is a centralized reservation based protocol. CAVs wishing to traverse the intersection call ahead to an IM with information including arrival time, arrival lane, arrival velocity, vehicle size, vehicle acceleration profile, and destination. The IM then evaluates the reservation request. If a reservation request conflicts with another already existing reservation or is otherwise deemed to be unsafe it will be rejected. Vehicles may not enter the intersection if they do not hold a reservation or are unable to follow their reservation, but may continue to submit subsequent reservation requests.

The AIM protocol was later extended to allow for human driven vehicles (HVs) in a protocol called Hybrid-AIM (H-AIM) [8]. H-AIM necessitates the introduction of some restrictions on both connected vehicles and HVs which wish to traverse an intersection, however. With this extension, the conditions for which a CAV's reservation request may be rejected are widened and HVs may not change lanes within a certain distance of the intersection. HVs follow normal traffic rules and are guided by visual traffic signals. Autonomous vehicles are still required to make reservations to cross an intersection, but reservations will not be approved where they could potentially conflict with any HVs potential paths (assuming all HVs follow applicable traffic laws). Upstream sensors (such as cameras or loop detectors which are commonly used today) can be used to detect vehicles which have not made reservations so that those vehicles may be considered as human for the purposes of evaluating the safety of reservation requests by CAVs. This scheme allows HVs to pass through a managed intersection, in much the same way that they do now, while still allowing CAVs to traverse an intersection in an extremely efficient manner in many cases.

1.2 Hybrid-AIM (H-AIM)

The most relevant work to this thesis is the H-AIM protocol which was first proposed by Sharon and Stone in 2017 [8]. H-AIM, which is a centralized and protocol based approach, is an extension of the research outlined in Dresner and Stone's 2008 work [1] in autonomous intersection management. HV behavior with H-AIM is guided by traffic signals, which is to say HV behaviour is largely unmodified from traditional traffic rules. It is assumed that upstream sensors (such as cameras or loop detectors which are commonly used today) can be used to detect human vehicles by comparing the arrived vehicle count with the number of reservation requests. Autonomous vehicles have a more complex interaction and must converse with an IM, however. The interaction between CAVs, IMs, and HVs is shown in Figure 1.1 and summarized below:

1. A CAV within communication range of the IM sends a reservation request to the IM containing vehicle arrival velocity, vehicle arrival time, arrival lane, destination, vehicle size, and vehicle acceleration capabilities.

2. The IM analyzes the request. The IM must first verify that granting the request does not compromise safety. This is done, within some safety margins, by:

    - Checking that the request does not conflict with any reservation.

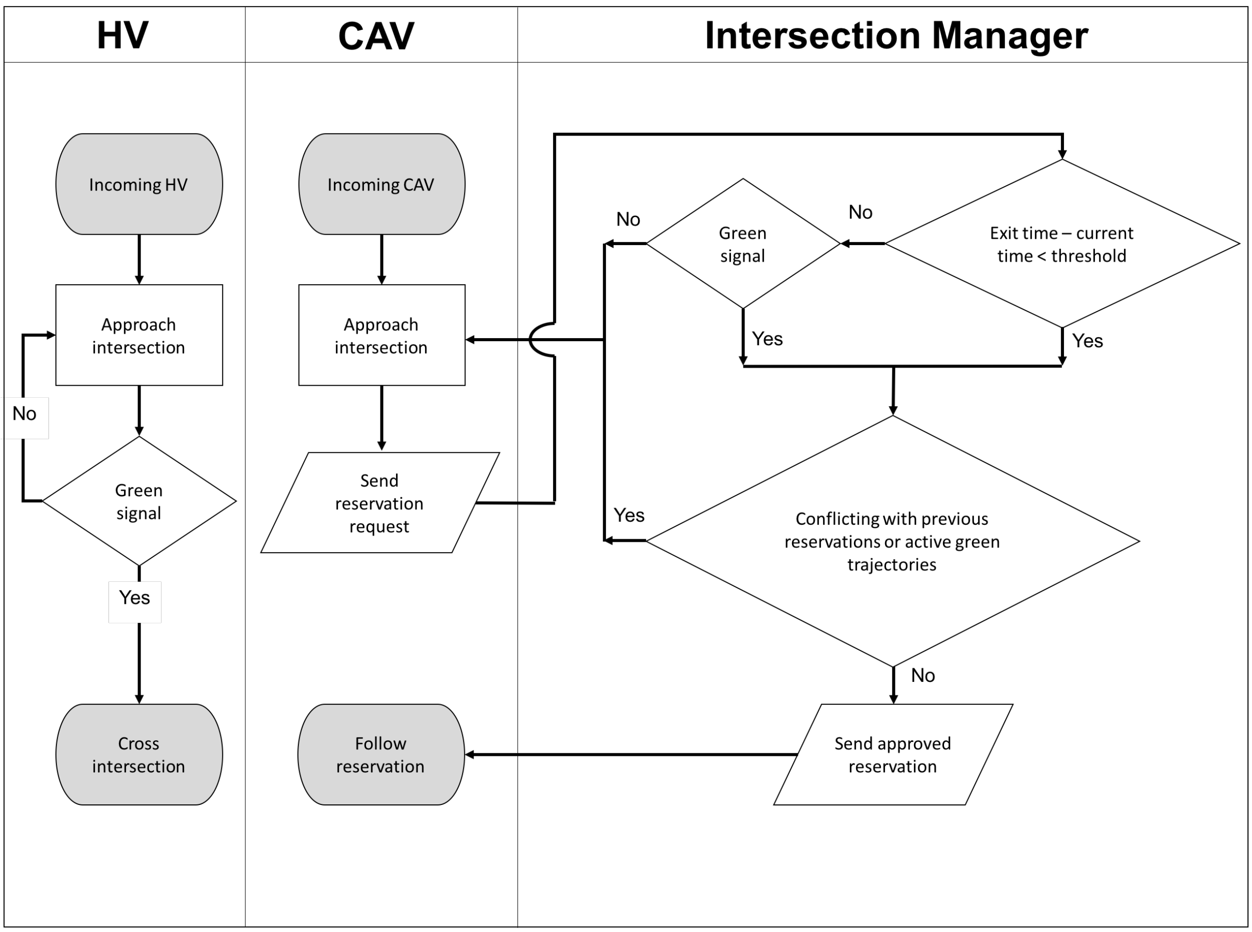

Figure 1.1: A flowchart of the behavior of the H-AIM protocol.

- Verifying that granting the request will not place a CAV's trajectory in conflict with an active green trajectory. Active green trajectories are any path which a detected HV might legally take from its position, provided that lane has a green traffic signal. Human operated vehicles are required not to enter intersections on yellow lights, or to change lanes within a certain distance of the intersection for the purposes of guaranteeing safety.

3. If a reservation request is unsafe, it must be rejected. If a reservation is found to be safe, an IM may decide to approve it. However, an IM is free to utilize any policy for approving or rejecting safe reservation requests. IM policies should be built to improve efficiency while considering edge cases. One edge case is reduction in efficiency due to CAVs which are farther away from the intersection being granted reservations prior to CAVs that are closer, preventing the closer CAVs from making a timely reservation.

4. The IM must then reply to the CAV by either approving or rejecting the request. An IM may include additional information in its message such as advisories, suggestions for attempting subsequent reservations, or the like.

5. In order to preserve safety guarantees CAVs must comply with their reservation in order and are not permitted to choose to follow traffic signals after communicating with the IM. If a communicating CAV is not able to acquire a reservation or is unable to comply with the currently held reservation within some allowable margin, the CAV is not allowed to enter the intersection. CAVs without a reservation must come to a stop at the entrance to the intersection until they are able to acquire a reservation.

6. CAVs notify the IM after leaving the intersection.

1.3 Literature Review

A number of works approach autonomous intersections from a purely control and optimization point of view. Riegger, et al. [7] and Murgovski, et al. [9] formulate the intersection traversal problems as a quadratic program in the space domain. Kim and Kumar [10] formulate a method based on model predictive control which is applicable to vehicle-to-vehicle (V2V) multi-lane road travel and vehicle-to-infrastructure (V2I) intersection traversal. These types of approaches have advantages in their formulation, as they're capable of operating in a continuous domain and provide some nice theoretical guarantees. However, these particular methods often make unworkable assumptions such as drastically limiting the ability of vehicles to follow other vehicles into an intersection or limiting a vehicle's ability to be in an intersection at the same time with another vehicle. Additionally, these methods typically require connected vehicles with high control precision to be applicable, which prevents them from being applied to situations with human operated vehicles. Though some work has begun to emerge which reformulates the optimization approach to account for human operated vehicles, such as in Liu, et al. [11], though Liu, et al. provide no empirical analysis.

Another potential improvement to intersection management schemes are approaches known as platooning and batch reservations. Batch reservations, such as those proposed by Au et al. [12], are when individual reservation requests are considered in groups within certain time windows and are prioritized by some cost function within the group of which they are a member. Platooning is a similar concept in that reservation requests are batched in some way, but this may be accomplished considering multiple vehicles within a single platoon-based reservation request. Batch reservations [12] and platooning [4, 13] both have potential for increasing efficiency of intersections. Platooning also has an added advantage of reducing V2I communication overhead due to the fact that the coordination of a platoon can be offloaded to V2V communications requiring only a representative vehicle from a platoon to communicate with infrastructure such as IMs. However, the degree to which platooning and batch reservations are advantageous is somewhat dependent on the ratio of CAVs to HVs on the road. So while batch reservations and platooning support may be helpful as components of an IM, the major advantages of these methods will not be observable in implementation for some time.

Some game theoretic/incentive based [5, 14] or auction based [15, 16, 17] techniques have also been proposed. Some of these approaches depart slightly from the typical objective of reducing overall delay and instead seek to maximize generic cost or utility functions. Many proposals in this vein allow for agents, sometimes including IMs, to have different utilities associated with receiving their preferred reservations or requests. This means that some of these proposals are formulated as inherently competitive games where agents may choose to act selfishly, though not unsafely. By relying on mechanisms that provide optimality guarantees, overall objectives such as social welfare can be optimized. These methods rely on connected vehicles' abilities to communicate and follow granted reservations, which means this class of approach is not suitable for vehicles with no communication capabilities.

Some areas of work forgo centralized intersection management architecture entirely and focus on employing distributed methods. VanMiddlesworth et al. [18] describe a system where CAVs negotiate to determine a schedule of intersection crossings. This is done in a similar way to a centralized reservation based system, but done solely through V2V communication while incorporating rules for determining priorities of reservations. Mladenović and Abbas [6] propose a similar rule based approach where uniform priority rules based on vehicle class, occupancy, and the like are used for CAVs to determine adjustments in velocity in order to allow for safe passage through an intersection. Results of both of these works are only demonstrated with low traffic volume or as feasible with low traffic volume. Additionally, distributed approaches face significant safety and applicability problems when HVs are considered.

Work has been done incorporating reinforcement learning as part of intersection management systems. Some examples of this are Mirzaei and Givargis' [19] work with a trust region policy optimization based scheme and Wu et al. [3] with a decentralized coordination learning approach. Though, neither of these methods incorporate human operated vehicles. Mirzaei and Givargis primarily demonstrate the feasibility of reinforcement learning in the problem space. Wu et al. show results where they claim their method to be superior to AIM in terms of delay reduction, but with significant assumptions such as a fully observable environment in terms of movement intention of vehicles, vehicle speed, vehicle position, and lane queue lengths.

Work such as in Yang et al. [20], Guler et al. [21], Feng et al. [22], Lee et al. [23], and He et al. [24] focuses on the optimization of traffic signal timing. The underlying theory behind these types of approaches is that connected vehicles (CVs, which may or may not be autonomous) can be utilized as upstream sensors for traffic signal controllers, potentially supplementing more typical sensor technologies. Traffic signal controllers would then have the ability to look far into the future in order to determine what signal phasing adjustments would be most likely to be advantageous given the velocity, density, distance, and headings of approaching vehicles given the data that CVs report as well as the ability to suggest variations in speed to connected vehicles. However, this type of work does not consider the fine control and communication abilities inherent to CAVs that can be used to further optimize intersection operation through coordination of conflicting traffic movements in a safe manner.

Bento et al. [2] present a method called Legacy Early Method for Intelligent Traffic Management (LEMITM) which is somewhat similar to H-AIM in concept. In this system, connected vehicles' capabilities are used by IMs to coordinate travel through the intersection while the IM reserves trajectories on behalf of vehicles that are incapable of communicating. These vehicles that are incapable of communicating, whether due to equipment failure or lack of equipment, are controlled by the IM with traffic lights. LEMITM's effectiveness is shown only in experiments with 90% or more CAV market penetration rate. This thesis is concerned with CAV market penetration rates ranging from 100% human to 100% autonomous.

Since Dresner and Stone's 2008 work, the research field of autonomous vehicle coordination and autonomous intersection management has been very active. The contents of this literature review are just an overview of some of the vast amount of work that's been done. For interested readers, significant reviews of work regarding intelligent intersection management systems have been performed by Namazi, et al. [25] as well as Chen and Englund [26], among others. Many of the methodologies in the specific area of autonomous intersection management show promise, but with unrealistic, cherry-picked, or simplistic simulations it's difficult to imagine fielding many of the so far proposed approaches without significant adjustment and fine tuning.

2. SIMULATOR ADJUSTMENTS AND PROTOCOL CONSIDERATIONS *

2.1 AIM Simulator Modifications

The AIM Simulator was originally built for the work leading up to and including Dresner and Stone's 2008 work [1] with the AIM protocol. It was modified for use for Sharon and Stone's [8] work in extending the AIM protocol to become the H-AIM protocol. Further extensions to the simulator were necessary to create more realistic simulation conditions.

2.1.1 Demand, Lanes, and Turning Profiles

Additions to the AIM Simulator for this thesis allow the simulator to use a comma delimited file generated from data from the Utah Department of Transportation's (UDoT) Automated Signal Performance Measures (ATSPM) system [27] to schedule vehicle arrivals for a single intersection within a single day. This allows for finely controllable and realistic arrival rates in simulation and makes the simulator capable of simulating peaks and valleys in traffic demand over a period of time. The ability to generate demand based on real historical data or even based on finely detailed artificial data makes the simulator capable of modeling more realistic scenarios as compared to earlier versions.

Prior to work performed for this thesis, the AIM Simulator assumed that a modeled intersection would be symmetric. That is, all directions of travel have the same number of incoming and outgoing lanes and all roads have the same speed limits. In order to model single intersections that are not totally symmetric, the simulator was modified so that certain lanes on roads may be "closed" by specifying the number of inbound and outbound lanes as well as some restrictions regarding what vehicles may do from each lane. No vehicle in simulation will spawn on or proceed to a "closed" lane. This specification and the assignment of speed limits is done in a configuration file. An example of an intersection in the AIM Simulator with closed lanes may be seen in Figure

*Some content in this section is reprinted by permission from Springer Nature Customer Service Centre GmbH: Springer Nature. *Springer eBook*. "A protocol for mixed autonomous and human-operated vehicles at intersections," by G. Sharon and P. Stone, © 2017.

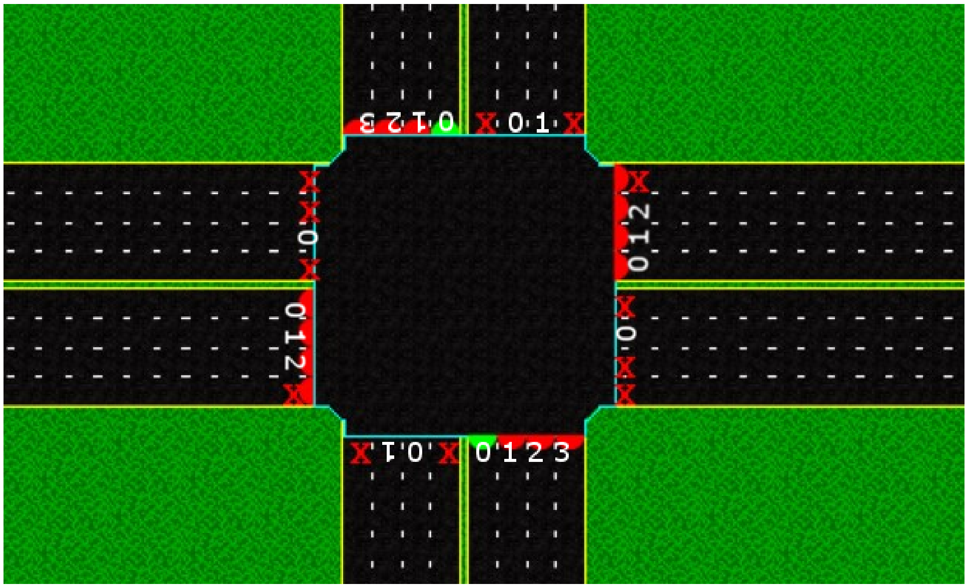

Figure 2.1: Example intersection in the AIM Simulator showing blocked lanes with relative lane indices superimposed.

2.1. Note that the intersection in Figure 2.1 is not totally symmetric.

Another consideration in modeling more realistic scenarios is the actions vehicles may take when entering an intersection. Modern intersections specify allowed turning movements (left, right, straight) for vehicles in each lane. This thesis adopts a similar system to assign such turning profiles to vehicles. Though, turning profiles only specify the general actions a vehicle of a specific type may take when proceeding through an intersection from particular lanes. They are not to be confused with conditions dictating whether a vehicle may enter the intersection at all, such as a requirement for a green light for the desired turning action or a granted reservation. The specifications of allowed movements for vehicles in the modified AIM Simulator used in this thesis are made in the same configuration file that determines the number of lanes and speed limits for roads serviced by the simulated intersection. As an example, in Figure 2.1, to allow vehicles to travel on the northbound road through the intersection while continuing straight the following 2 pairs would be specified in the appropriate place within the configuration file: (1,0) and (2,1). These pairs are of the form (relative lane index inbound, relative lane index outbound).

An example of a configuration file for an intersection, which includes specification of speed limits for roads, number of lanes, and turning profiles can be found in Appendix A.

2.1.2 Signal Timing for Human Operated Vehicles

The AIM Simulator's ability to model traffic signal progression has also been extended for this thesis. Rather than an only fixed signal timing scheme as before, the AIM simulator now allows for a ring and barrier style signal phase specification by configuration file. This modification also incorporates simple signal actuation functionality. Specification of signal phase progression and timing in a configuration file facilitates the use of signal actuation functionality. Users may easily specify the **gap extension** (time a vehicle detection or "actuation" may extend a green signal), maximum, and minimum signal times appropriate for simulation. Users may choose to disable actuation detection altogether or to use an adaptive signal timing scheme if desired. An example of the configuration file can be found in Appendix A, though the data tables for the adaptive timing are hard coded into the simulator (accessible in the "GapExtensionTable" and "MaximumGreenTable" Java files in the modified simulator's code base).

2.1.2.1 Phase Progression and Actuation

The top of Figure 2.2 shows a graphical example of a ring and barrier signal representation. Each group of phases within a ring, separated by barriers, consists of conflicting turn movements which if signaled green together ("being active") would create unsafe traffic conditions. Thus within a group of phases for a ring, only one phase may be active at a time. Though, multiple rings will have an active phase bounded by the same barriers. As timers expire, an active phase within a ring will transition to the next phase in the current group, usually by appropriately signalling yellow and then red as the phase transitions. In the case of actuated control, arrival of vehicles may extend the time the light will may remain green in increments of up to the gap extension amount until the phase reaches a maximum time. Phase transitions which occur in different rings may be coordinated but are not necessarily so, except for when safety is a concern between particular phases (e.g. when left turns of opposite directions could lead to a collision) or when transitioning across a barrier. All rings must transition the active phase across the barrier at the same time in order to maintain safe operation of the intersection.

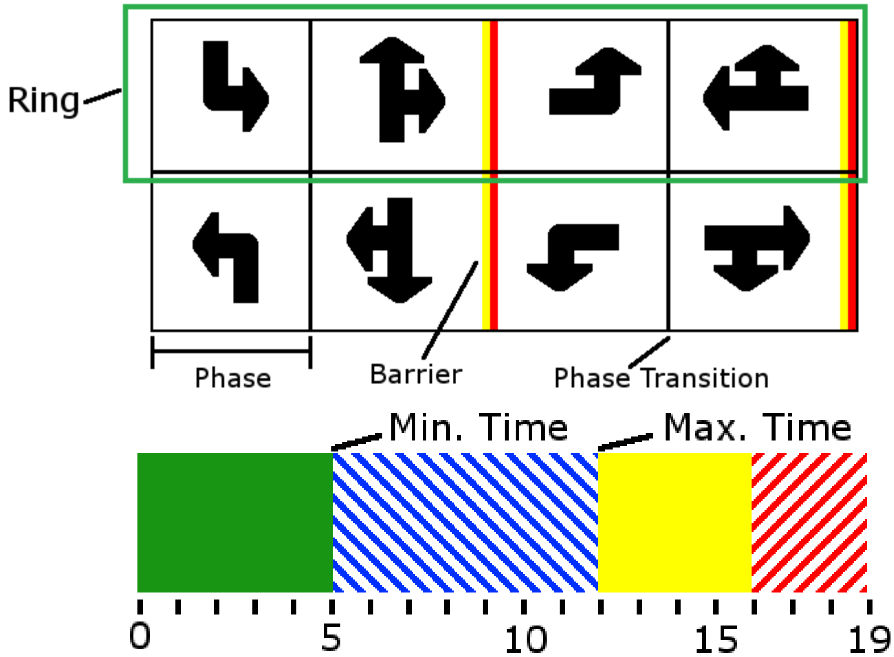

Figure 2.2: Ring and barrier diagram (top) and an example of the timing for a simple actuated signal phase with a 19 second time scale (bottom).

The bottom of Figure 2.2 shows a simple diagram of the timing for a signal phase as used in this thesis. The leftmost solid portion represents the **minimum green time**. This is the minimum time for which the signals of all turning movements associated with the phase must be green. The minimum time for the example on the bottom of Figure 2.2 is 5 seconds. The striped portion with a length of 7 seconds represents the additional time that signal may remain green if more vehicles are approaching. The extension will occur in up to increments of the gap extension amount. The extension is not a requirement and the signal color may transition to yellow immediately after the minimum time has elapsed if no vehicle is detected. If a vehicle is detected through some sensor such as an inductive loop embedded in the roadway, the time the signal is green will be extended by a fixed amount. However, these extensions may not continue indefinitely as they may starve other traffic directions of the opportunity to progress through the intersection. Thus a phase will typically only be extended up to a **maximum green time** (a total green time of 12 seconds in Figure 2.2), at which point no further activation of sensors will extend the signal for this phase. The exception to this behavior in this thesis is when the green time of a phase in one ring is complete but is waiting on another phase's green time to finish before all rings transition through a barrier. In this case, the green time will be artificially extended while other phases complete. Lastly, a phase's signal will transition to yellow for a fixed time (a duration of 4 seconds in the example in Figure 2.2), then red for a fixed time (a duration of 3 seconds in the example in Figure 2.2) and then will transition to the next phase.

2.1.2.2 Adaptive Timing

In the adaptive scheme, the maximum green time for a phase and the gap extension times are determined from IM observed traffic patterns or road parameters. Factors that affect the overall selection these timing parameters include road speed, recent signal cycle length (the time it takes the traffic signal to loop), and recent vehicle arrival rates. Tables 5-6 and 5-10 from the U.S. Department of Transportation Federal Highway Administration's 2008 Traffic Signal Timing Manual [28] are used in this thesis when adaptive timing is enabled in order to determine the values the IM employs at a given time.

2.2 Green and Active Green Trajectories

As HVs are a major consideration in the H-AIM protocol, Sharon and Stone [8] define trajectories that must be reserved for HVs by IMs as "green" or "active green". Figure 2.3 shows green trajectories superimposed on top of an intersection. A green trajectory is a path through the intersection from an incoming lane which is assigned a green signal. An active green trajectory is a green trajectory where a HV is present on it or is near to entering the intersection on the associated incoming lane. The set of green trajectories that exist within an intersection changes as signals progress. Utilizing these definitions as well as knowledge of the schedule for traffic signals, an IM must reject CAV reservation requests which conflict with any potentially active green trajectory.

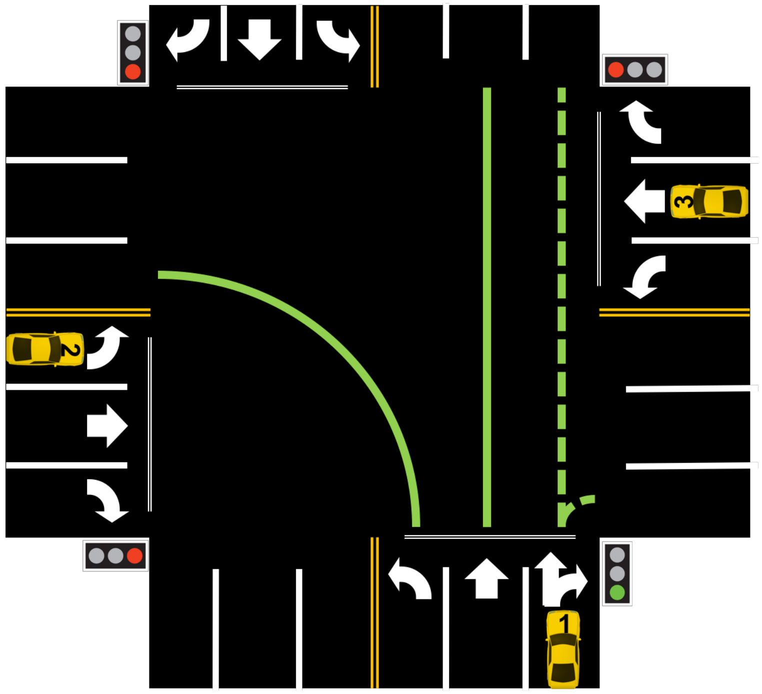

Figure 2.3: Four-way intersection with a green signal assigned only to all northbound lanes. Solid or dashed green lines show green trajectories. Active green trajectories are denoted by dashed green lines.

2.2.1 Reduction of the Number of Green Trajectories

Green trajectories can limit CAVs from obtaining reservations. As such, CAVs benefit from reducing the number of green trajectories to a minimum. On the other hand, HVs cannot cross the intersection unless traveling on a green trajectory. Thus, HVs generally benefit from an increased number of green trajectories. Dresner and Stone [1] presented a one-lane signal policy where only one lane is granted a green light at a time, rather than a collection of lanes. This policy results in green trajectories that originate from a single lane at a time which significantly reduces the number of green trajectories. However, the one-lane signal policy was shown to have a dramatic negative effect on HVs. H-AIM suggests a more conservative approach for reducing the number of green trajectories. Revisiting Figure 2.3, assume vehicle 3 is autonomous and is heading west. When applying H-AIM, vehicle 3 is automatically denied a reservation since the requested reservation crosses an active green trajectory. Currently, the lane on which vehicle 1 approaches the intersection allows continuing straight or turning right. If the turning policy on that lane is changed to "right only", the dashed straight green trajectory will no longer exist allowing vehicle 3 to obtain a reservation.

2.2.2 Turning Assignment Policy Effect on Green Trajectories

As was discussed above, the performance of a managed intersection is affected by the allowed turning options in each lane. When considering a four-way intersection, each incoming lane has between one and three turning options from the set {left, straight, right}. As briefly mentioned in the prior discussion on simulator configuration, **turning assignment policy** defines the sets of turning actions which are allowed on each incoming lane. These actions may also be specified for a specific vehicle type.

For the purpose of demonstrating turning assignment policy concepts, this section assumes three incoming lanes in a perfectly symmetric intersection. Three turning assignment policies which comply with this assumption are depicted in Figure 2.4. The policies are ordered and labeled according to degrees of freedom. The **degrees of freedom** for a lane is the number of turning options minus one. And so the degree of freedom for a turning assignment policy is the sum of degrees of freedom over all lanes. Thus, a restrictive turning policy is one that has a low degree of freedom which, in turn, translates to fewer green trajectories. Policy 0 in Figure 2.4 is an extreme case, representing the most restrictive turning policy (0 degrees of freedom). On the other hand, policy 4 is an extreme case of a liberal turning policy.

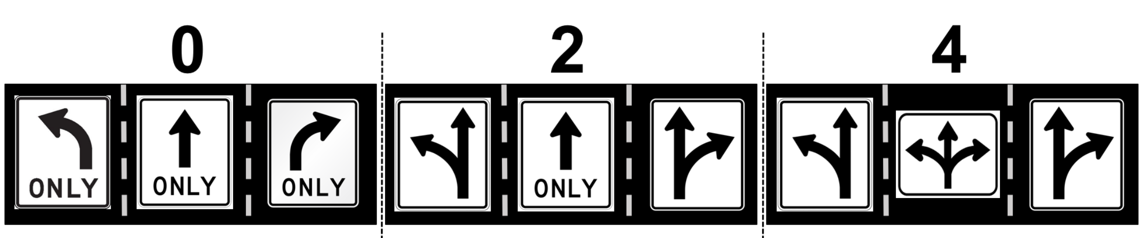

Figure 2.4: Three turning assignment policies for a three lane road approaching a four way intersection with varying degrees of freedom.

Sharon and Stone [8] also discuss the definition of safe and unsafe turning policies. This thesis refers to these as consistent or inconsistent. With the modified moniker, in their formulation a **consistent turning policy** is one where trajectories originating from the same road never cross each other. In in Figure 2.4, turning policy 4 is not consistent while 0 and 2 are. When considering more than one type of vehicle, different turning policy combinations might be employed. For instance, one might choose to assign a turning policy for HVs and a different one to CAVs. Thus a set of turning assignment policies are said to be a **consistent turning policy combination** if no trajectory from any one policy in the set crosses any trajectory which originates from the same road as part of any other policy in the set. In the representative policy set shown in Figure 2.4, {0, 4} is a consistent turning policy combination (even though 4 is not a consistent policy on its own). {2, 4} is not a consistent turning policy combination.

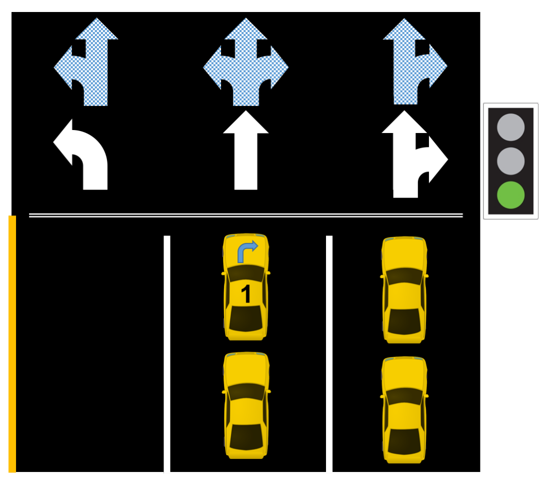

Figure 2.5: An inconsistent policy combination. Top policy (checkerboard texture) for AVs, bottom policy (plain texture) for HVs.

As Sharon and Stone [8] note, for safety reasons one shouldn't assign an inconsistent policy to HVs. Though on the other hand, assigning such a policy to CAVs is reasonable since conflicting reservations will be automatically denied by an intersection manager using H-AIM. Though, Sharon and Stone also note assigning inconsistent policy combinations for CAVs and HVs is counterproductive from an efficiency standpoint and should be avoided. Figure 2.5 demonstrates the inefficiency that stems from an inconsistent turning policy combination. The figure presents a single road approaching a four-way intersection. CAVs are assigned the turning policy shown on the top level (checkerboard texture) while HVs are assigned the bottom turning policy (plain texture). Vehicle 1 is autonomous, is located in the middle lane, and would like to turn right. Assuming a green signal for this incoming road and that HVs are arriving on the rightmost lane, vehicle 1 will not be able to obtain a reservation as it crosses an active green trajectory. Vehicle 1 will thus be stuck and will jam all the vehicles behind it despite having a green signal.

2.3 Protocol Extension

Part of the H-AIM protocol requires that CAVs call ahead to make reservations. A consideration of whether to approve a reservation is the status of traffic signals and whether HVs might enter or be within the intersection while a CAV crosses. With fixed signal timing, it's simple and efficient for an IM to "look ahead" to determine if a CAV's trajectory will cross the active green trajectories of detected HVs. However, actuated controllers complicate this process. When a CAV calls ahead to attempt to make a reservation, an IM supporting actuated control may not know when the current signals will change due to potentially unforeseen actuation events. To make matters worse, use of the adaptive timing scheme mentioned above means that maximum green times and gap extension times for different phases may change between subsequent look ups for signal status in the future. Thus, without further modification to the H-AIM protocol, an IM naively looking ahead based on the current status of phases will likely approve a reservation leading to a collision. Thus the behavior of the H-AIM protocol must be extended in order for it to function with modern and commonly deployed signal management mechanisms.

2.3.1 Adaptive Timing and Signal Actuation Effects on Reservations

In order to combat the shifting phase timing due to use of adaptive timing, an IM may lock the maximum and gap extension times for any phase whose data is used to make a safety critical decision. This lock may be released upon transition out of the particular phase (or phase segment if phases are modeled as segments broken up into green, yellow, and red segments).[1] In order to combat the uncertainty caused by actuated control, a version of the algorithms below are used by IMs during simulations for this thesis when considering future reservations. It should be noted that these algorithms consider phase segments instead of phases. At a high level, these algorithms cooperate to provide a collection of sets of phase segments which could be active at the given future time to an IM. While not precise, this allows IMs to continue to function in actuated environments with human operated vehicles by examining the sets and approving reservations only when there is no possibility that a conflict will occur between HVs and CAVs. Pseudocode and line-by-line descriptions of the algorithms follow.

2.3.2 Recursive Phase Segment Lookup (Algorithm 1) Description

The output for Algorithm 1 is a collection of sets of phase segments which may be active at a provided time for individual rings. [2] The number of sets returned is the number of rings, as each set corresponds to one ring. The input consists of the time relevant to the look up ("futureTime"), the rings of the ring and barrier structure which contain knowledge about the current and future states of signals, and the time at which each of the currently active phase segments within each ring began ("epochCollection").

Beginning on line 1, the collections of data which is used to track algorithm progress are

[1] In employing a lock and release strategy for phases/phase segment data when using adaptive timing schemes, one must ensure that look ups do not wrap around the end of the phase/phase segment progression and back to the current phase/phase segment. Depending on implementation detail, this could lead to phases/phase segments potentially being unlocked unsafely and/or could result in permanent locking.

[2] A future signal lookup implementation that uses only the maximum and minimum times that phase segments can persist to determine the set of possible phase segments at a future time is a conservative estimate. This means there may be more phase segments contained in the set than are actually possibly active at the given future time, but there will be no possibly active phase segments missing. This does not violate typical safety constraints, but may harm efficiency.

Algorithm 1 Recursive Phase Segment Lookup

Input: Future time to which to search ("futureTime"), a collection of well formed ring objects through which the current phase segments can be accessed, and a collection of times at which the current phase segments started ("epochCollection").

Output: A collection of sets of phase segments which may be active at the given future time.

1: Let "earlyResults" and "laterResults" be collections used to track algorithm progress.
2: Let "numberDone" be a counter initialized to 0.

3: **while** the "numberDone" counter is less than the number of rings **do**
4:     **for** every ring, with each referenced below as "currentRing" **do**
5:         **if** there is no entry in "earlyResults" for the "currentRing" **then**
6:             Create an entry in "earlyResults" for the "currentRing" which contains a "done" flag as false, a "chooseEarly" flag as true, the beginning time ("epoch") of the current phase segment for "currentRing" taken from "epochCollection", and a set where the most recent addition to the set may be accessed and which contains the current phase segment for "currentRing".

7:         **if** the "done" flag for the entry for "currentRing" in "earlyResults" equals false **then**
8:             Set the entry for "currentRing" in "earlyResults" to the result of a call to Algorithm 2 with futureTime and the entry for "currentRing" as the arguments.
9:             **if** the "done" flag for the entry for "currentRing" in "earlyResults" equals true **then**
10:                 Increment the "numberDone" counter.

11:     **if** any ring's entry in "earlyResults" has the "done" flag equalling false **then**
12:         Let "barrierStart" be the farthest time in the future at which any ring would advance to a barrier from the current group of phase segments assuming the minimum possible time to end each phase segment were used in the look up.
13:         **if** every ring's entry in "earlyResults" has the "done" flag equalling false **then**
14:             For every ring, examine the segment that was most recently added to the phase segment set in the progress tracking entry and then add the segment that immediately follows that one in the ring to the phase segment set.
15:             Set the "epoch" associated with the results tracking for every ring to "barrierStart".
16:         **else**
17:             Set the "done" flag for every ring's entry in "earlyResults" to true.
18:             Set "numberDone" to the number of rings.

19: //Repeat lines 3-18 but set "chooseEarly" on line 6 to false, use the maximum possible time for line 12, use "laterResults" instead of "earlyResults", and reset the "numberDone" counter to 0.

20: **return** A collection where ring's each entry is the intersection of the segment sets within the associated "earlyResults" and "laterResults" data.

defined. Line 2 defines and initializes a counter ("numberDone") to assist with tracking progress. The while loop on line 3 will loop until the "numberDone" counter is equal to the number of rings. In other words, this loop will terminate when the look up for each ring is complete. Multiple passes may be necessary depending on the makeup of the signal progression, signal timing, and the "futureTime" parameter. Line 4 iterates once for every ring. Line 5 checks if an entry used to track progress for the associated ring exists, line 6 creates that entry if it's missing. Next, lines 7 determines the behavior for the current iteration of the loop, depending on whether the search has already been completed for the ring that is currently being examined. If the search is not complete, Algorithm 2 is called on line 9 and the progress tracking entry associated with the current ring is overwritten with the result of the call. Otherwise lines 8-10 are skipped. Algorithm 2 assists Algorithm 1 by looking ahead as far as possible to determine which phase segments may occur at "futureTime". However, Algorithm 2 may block before completion if a barrier is encountered, as all phase segments must progress across a barrier simultaneously and Algorithm 2 only has knowledge about a single ring at a time. Lines 9 and 10 check if Algorithm 2 was able to run to completion and increments the "numberDone" counter if it was. Lines 11-18 handle the cases where Algorithm 2 blocked (as indicated by an entry not being flagged as done). Line 12 determines the earliest time at which all of the phase segments may progress across a barrier, this time is dubbed "barrierStart". Lines 13-15 progress every phase segment to the barrier segment, but only if the search on every ring blocked on the barrier. Lines 16-18 handle the case where the search on at least one ring completed, but others blocked on a barrier. With the implementation of rings and barriers in this thesis, this means that "barrierStart" is greater than "futureTime". Segments in all rings would normally progress across the barrier at "barrierStart". Therefore the search is actually complete and the segments preceding the barrier for the pending searches are the final segments in the search for the respective rings. Lines 17 and 18 set the tracking values so that the search is complete. Line 19 explains the required repetition of lines 3-18 using different parameters in order to obtain the full phase segment range. Line 20 explains that the final step to obtain the result of the algorithm is to take the intersection of the segment sets from the entry of "earlyResults" associated with a particular ring and from the entry of "laterResults" associated with the same ring.

2.3.3 Segment Lookup Helper (Algorithm 2) Description

Algorithm 2 performs the primary search for Algorithm 1 for a single ring (which is effectively a linked list of phase segments). Algorithm 2's inputs are the time to which to search ("futureTime"), a flag to indicate whether to use the earliest or latest possible time for estimation that a segment may end("chooseEarly"), a time for the beginning of a phase segment ("oldEpoch"), and a set where the most recent addition to the set may be accessed and which contains phase segments ("segmentSet"). The output of Algorithm 2 is an object which contains the same types of elements as the algorithm's input except the time to which to search.

Line 1 checks which behavior Algorithm 2 checks whether the earliest possible end time or the latest possible end time should be used when estimating end times of phase segments. Assuming the branch is chosen where the earliest end times should be selected, line 2 assigns a name ("oldSegment") to the most recently added segment in "segmentSet". Line 3 gathers the time at which "oldSegment" would end if it were being considered in a vacuum (i.e., no barriers that may halt progress) and names that time "newEpoch". Line 4 checks if "newEpoch" is less than "futureTime" and that "oldSegment" is not followed immediately by a barrier. This case handles the situation when the search may progress freely and has no dependence on the results of searches for segments on other rings. Assuming the condition on line 4 is true, line 5 adds the segment which follows "oldSegment" in the associated ring to "segmentSet". Line 6 then makes a recursive call, which will in turn return the obtained result, to Algorithm 2. The arguments used are the "chooseEarly" value from the input, the now modified "segmentSet", and "newEpoch". Line 7 handles the case where "newEpoch" is less than "futureTime" but that "oldSegment" is followed immediately by a barrier. To progress any farther, information about other searches is needed. In order for a phase segment to progress across a barrier it must progress at the same time as all other phase segments. Thus, line 8 returns an object with the "done" flag set to false indicating the search is blocked, "oldEpoch" which is the time the currently blocked segment ("oldSegment") began, segmentSet",

Algorithm 2 Segment Lookup Helper

Input: Future time to which to search ("futureTime"), a flag to indicate whether to use the earliest or latest possible times a segment may end for estimation ("chooseEarly"), a time for the beginning of a phase segment ("oldEpoch"), and a set where the most recent addition to the set may be accessed and which contains phase segments ("segmentSet").

Output: An object which contains a "done" flag, a "chooseEarly" flag, the beginning time ("epoch") of a phase segment, and a set where the most recent addition to the set may be accessed and which contains phase segments.

1: **if** "chooseEarly" equals true **then**
2: Let "oldSegment" be the most recently added segment in segmentSet
3: Let "newEpoch" be the minimum possible time that "oldSegment" could end assuming it began at "oldEpoch".
4: **if** "newEpoch" is less than "futureTime" **and** the segment following "oldSegment" is not part of a barrier **then**
5: Add the the segment following "oldSegment" from the appropriate ring to "segmentSet".
6: **return** The result of a recursive call to Algorithm 2 with the arguments of "futureTime", "chooseEarly", "newEpoch", and "segmentSet".
7: **else if** "newEpoch" is less than futureTime **and** the segment following "oldSegment" from the appropriate ring is part of a barrier **then**
8: **return** An object which contains a "done" flag set to false, "oldEpoch", "segmentSet", and "chooseEarly".
9: **else**
10: **return** An object which contains a "done" flag set to true, "oldEpoch", "segmentSet", and "chooseEarly".
11: **else**
12: // Duplicate behavior of lines 2-10, except use the maximum possible time for line 3.

and "chooseEarly". Line 9 and 10 handle the case where where "newEpoch" is greater than "futureTime", and so the search is complete. Line 10 returns the same values as line 8, but with the "done" flag set to true in order to indicate the search on this ring is complete. Finally, lines 11 and 12 deal with the case where the latest possible end times should be used when estimating end times of phase segments. Line 12 alludes to a simple repeating of lines 1-10 with the necessary modification to how segment end times are evaluated.

3. EXPERIMENTAL EVALUATION

This chapter presents results from an empirical study. The goals of these experiments are three-fold:

1. Study the effectiveness of H-AIM for mixed traffic with an emphasis on low CAV ratios while using real traffic data.
2. Indicate which types of turning policy should be assigned to HVs and CAVs in different traffic scenarios.
3. Demonstrate the effectiveness of actuation and adaptive signal timing in combination with H-AIM.

3.1 Intersection Model

Using readily available data from the UDoT's ATSPM system [27], 3 intersections were selected and modeled to use for the experiments to follow. Intersections are modeled based on the number of incoming lanes, the number of outgoing lanes, and the speed limits for each road. Signal progression is approximated as a somewhat standard 8 phase model (similar but not identical to Figure 2.2 for all three intersections, though signal timing and phase order/pairings vary). Incoming and outgoing lane configurations were gathered using Google Maps [29]. Speed limits were gathered through a combination of Google Maps and Utah DoT information [30]. Overhead photos of the intersections may be seen in Figure 3.1.

Traffic demand data for simulations is taken from UDoT's ATSPM system for Friday, March 8, 2019. A snowstorm occurred on this date. The time over which the data spans which was selected are the hours between 5:00 AM inclusive and 8:00 PM exclusive. March 8, 2019 was selected as it demonstrates different traffic patterns for the intersections, which allows for more insightful comparison as opposed to other days where the patterns for all 3 intersections are more similar. The specific time window was chosen with the intent of including peak traffic hours for both the

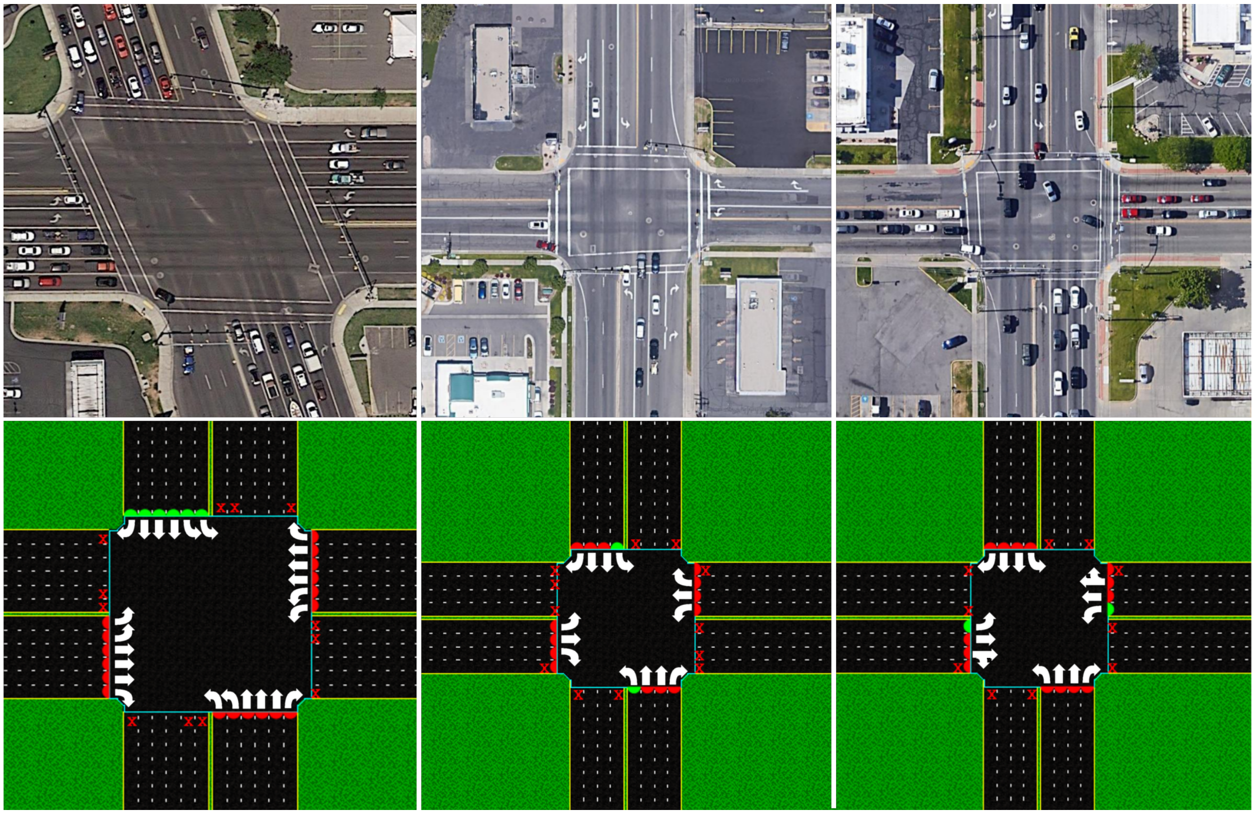

Figure 3.1: Intersection #6303 (top left, State St. at 800 North), #7204 (top center, 900 East at 5600 South), and #7381 (top right, 5600 West at 3500 South) above respective AIM Simulator representations. Current/real turning assignments are superimposed on top of the AIM Simulator representations. Satellite Imagery: Imagery © 2020 Maxar Technologies, State of Utah, Map data © 2020. Courtesy of Google, [29].

morning and the evening. Figure 3.2 shows the demand for all three intersections smoothed across 5 minute intervals.

A graphical depiction of the turning policies used below can be seen in Figures 3.3 and 3.4. The turning policies generally are the most restrictive policy (fewest degrees of freedom), the current (or "real") policy seemingly used by the real intersection, and a permissive policy (most degrees of freedom, but only for CAVs). Note that there are a few oddities in how the turning policies vary by intersection. The intersection model for #6303 is totally symmetric, so only one direction is shown for each policy for this intersection. The restrictive and current policies are identical when considering #6303 or #7204, and so those policies are merged and co-labeled for each inter-

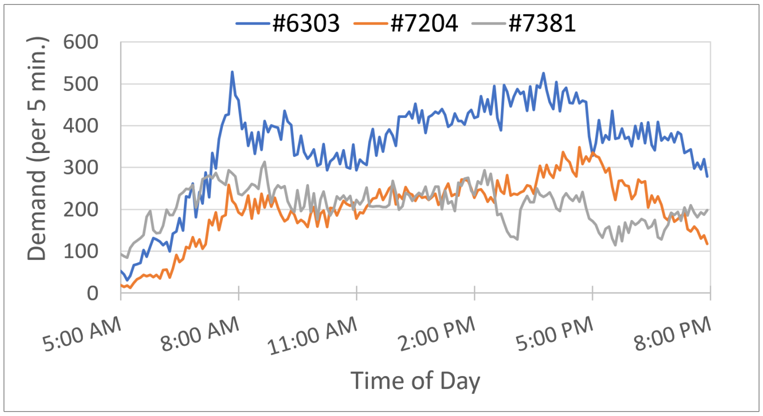

Figure 3.2: Demand in terms of vehicles (aggregated in 5 minute chunks) plotted against time for intersections #6303 (plot on top, starting in the middle of the 3), #7204 (slight upward trend, starting at the bottom of the 3), and #7381 (fairly steady trend, starting at the top of the 3).

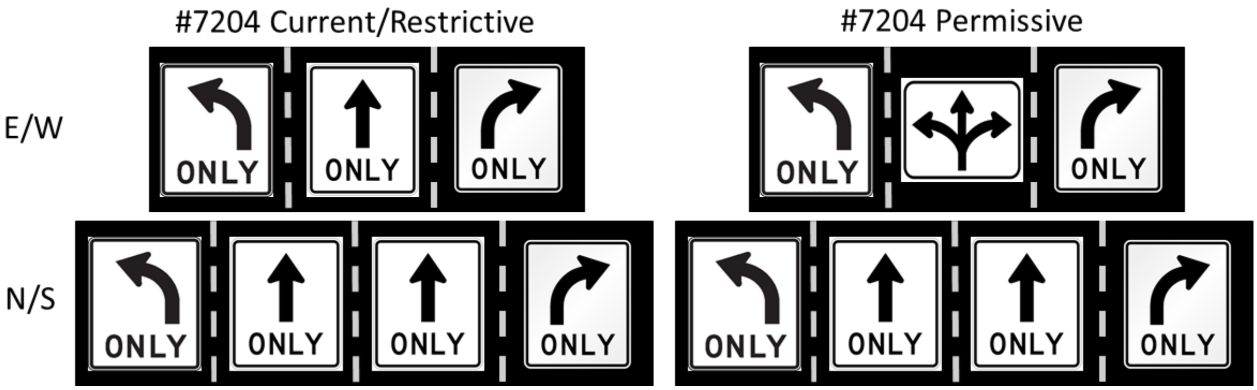

Figure 3.3: Turning policies used in empirical evaluation for intersection #7204. #7204 has symmetry in terms of allowed turning actions between the eastbound and westbound roads and also between the northbound and southbound roads.

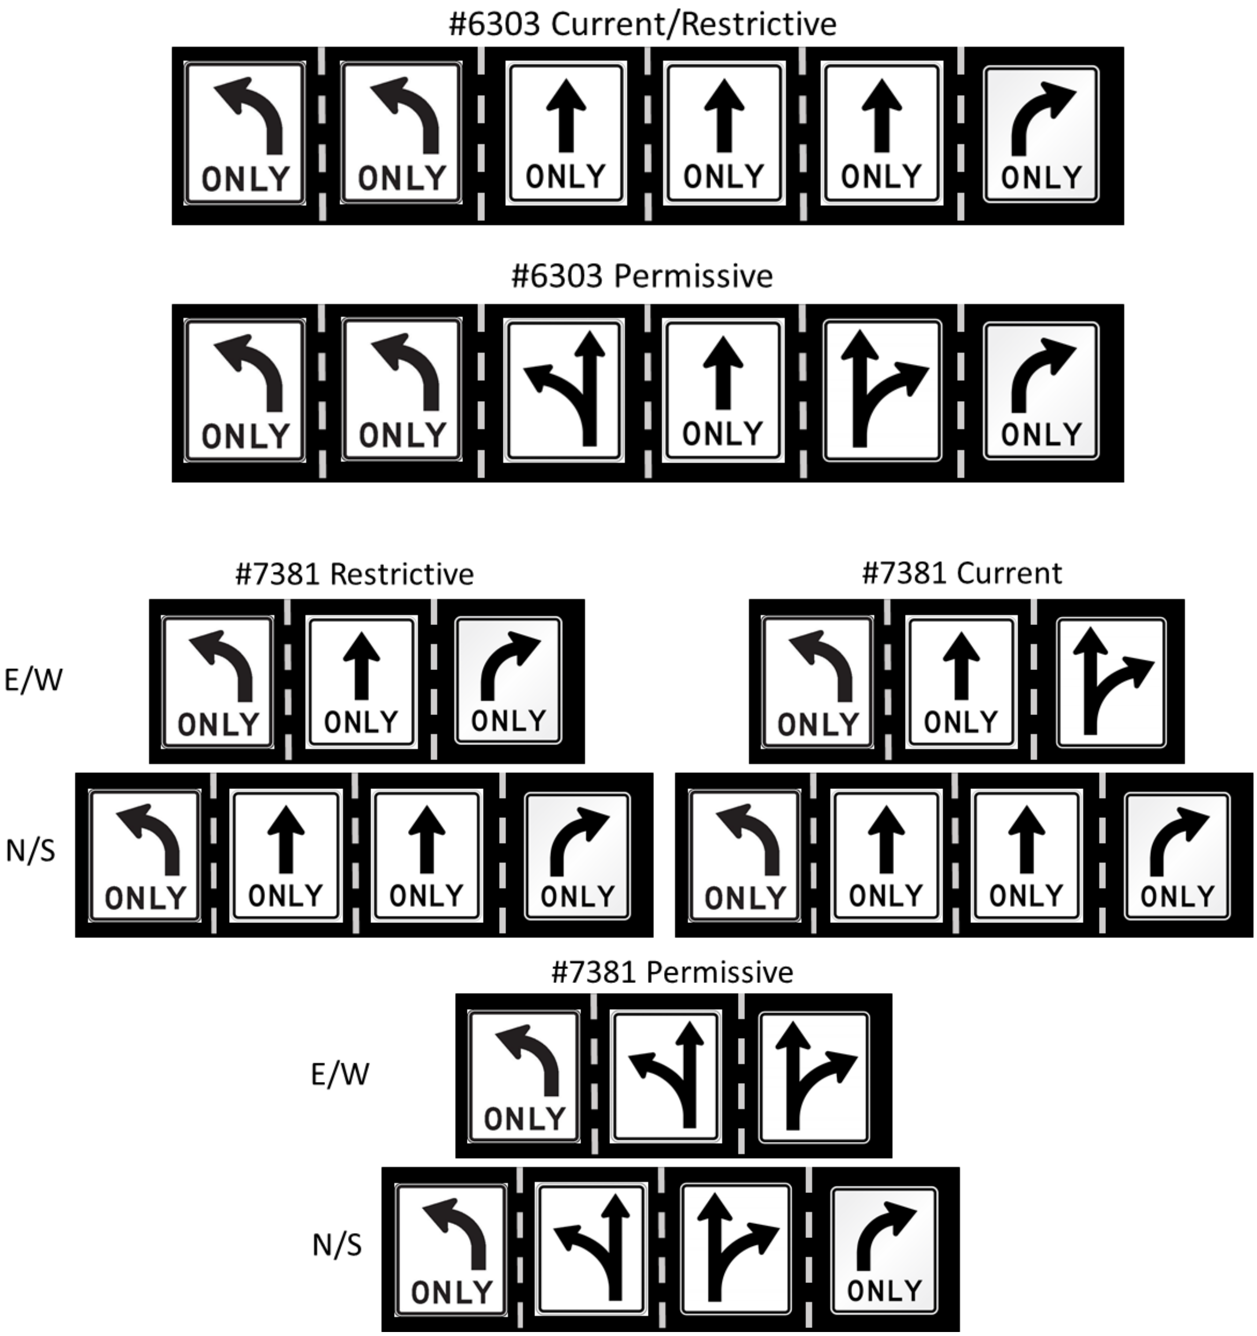

Figure 3.4: Turning policies used in empirical evaluation for intersections #6303 and #7381. #7381 has symmetry in terms of allowed turning actions between the eastbound and westbound roads and also between the northbound and southbound roads. #6303 is entirely symmetrical in terms of turning actions.

section. Lastly, #6303 and #7831 require a slight modification to their permissive policies which makes them slightly less permissive than potentially possible. This is because when given choices about exactly how to assign the turning movements in turning profiles, this thesis prioritizes consistency between experiments, consistency with the actual intersection, and avoidance of assignments where CAVs might be required to turn across potential HV trajectories originating from the same road.

3.2 Experimental Procedure

Similar to the experiments presented by Dresner and Stone [1], the following experiments assume that a CAV may communicate with the intersection manager starting at a distance of 200 meters. Following Dresner and Stone, results are presented as averages over 20 instances per variation of CAV percentage. Standard deviation is shown with bars on the plots. Unlike Dresner and Stone, speed limits are set according to data gathered about the real intersection being modeled in each experimental setting. Dresner and Stone considered a speed limit of 25 meters/second which is uncommonly high for signaled intersections.

A baseline performance using actuation and adaptive timing for each intersection is shown in each plot. The baseline is an average from 20 individual experiments for each intersection with a 0% CAV spawn rate while using currently deployed adaptive timing (which is akin to a manually built schedule by time of day), actuation, and turning policy assignment. Standard deviation of the baseline is shown as a filled region around the representative lines. Note that the baselines for #7204 and #7381 are lower bounds due to the simulator being unable to spawn vehicles on full lanes. The percentages of vehicles which spawn as CAVs are varied from from 0%-100% in all experiments. In order to emphasize the performance evaluation on the early CAV adoption period, this is first done in 11 increments of 1% (starting at 0%) and then the remainder is done in increments of 10%.

The experiments to follow are collected into 2 groups. The first group of experiments uses a fixed timing profile and examines the effects of varying turning policy assignment on **delay** (the time lost due to a vehicle not being in motion at full possible speed). The best and most consistently well performing turning policy combination for each intersection from this set of experiments is carried forward into the later experiments. The next group of experiments demonstrates the performance benefits of using an adaptive signal timing scheme and compares results with and without signal actuation enabled. Right turns on red are permitted for all turning profiles. No permissive left turns on green are permitted for any turning profile. All CAVs in the following simulations act as CAVs and communicate with IMs to obtain reservations.

3.3 Turning Policy Variation

A series of results are shown in Figure 3.5 for each intersection varying the turning policy assignments. Note that some data points represent a lower bound due to the simulator being unable to spawn vehicles on full lanes, effectively smoothing the arrival rate of vehicles. (See delays marked with "*" in Table B.2 in Appendix B.) The policy combinations are denoted as CAV policy initial then HV policy initial after the intersection number for each graph. The policy combinations are listed as {CAV Policy, HV Policy} in the text. Turning policy assignment pairs for HVs and CAVs are varied between the permutations of restrictive (fewest degrees of freedom) policies, the current policy seemingly used by the real intersection, and a permissive policy (most degrees of freedom, but only for CAVs). Combinations of turning policy assignments that are restrictive for CAVs where HVs are not also restricted aren't considered, as they would lead to a loss of efficiency.

The U.S. Department of Transportation Federal Highway Administration's recommendation for a maximum of 120 seconds cycle length for typical intersections [31] was used as a basis to create the fixed timing plan for each intersection in Table B.1 in Appendix B. Green timing per phase was allocated proportional to the phase demand based on the respective March 8, 2019 demand profile for each intersection during the 5:00 AM-8:00 PM time frame. The maximum value of the demand for straight and right movements was used to calculate combined straight/through movement proportions. Yellow and red timing are 4 and 3 seconds, respectively, between all phases and across all barriers.

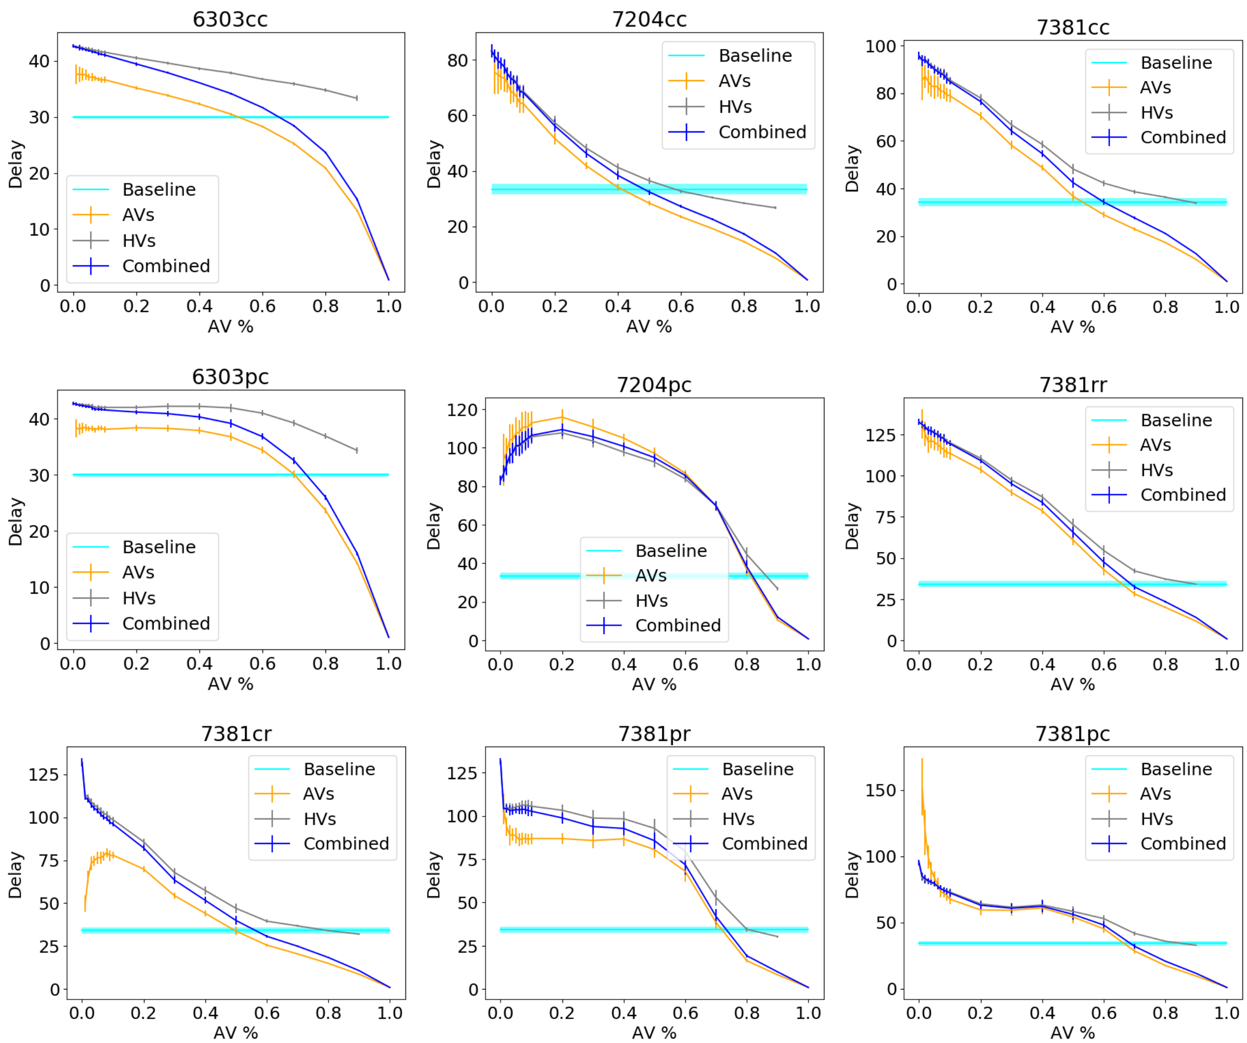

Figure 3.5: Graphs of delay by CAV percentage using fixed timing for various turning policies. Each graph is labeled with the intersection number and then 2 letters representing the turning policy combination in the order of CAV policy and then HV policy. "C" stands for the current policy in use by the real intersection, "p" represents a permissive policy, and "r" represents a restrictive policy. Each data point represents the average of the results 20 trials. Error bars show standard deviation of the results of those trials. A baseline with standard deviation is shown as a horizontal line.

3.3.1 Policy Variation Analysis

The {Current, Current} policy combinations for all 3 intersections seen in the first row of Figure 3.5 are quite consistent for each intersection. CAV delays in these experiments are closely bound to HV delays because CAVs often get trapped behind HVs, making CAV delay highly dependent on HV delay. On the right of the middle row, the {Restrictive, Restrictive} policy combination for #7381 shows a similar pattern to the {Current, Current} combination for the same intersection, but with higher delay. This is seemingly due to decreased throughput caused by removing the through movement in the rightmost lane for the eastbound and westbound directions. Also note that all intersections using fixed timing, including those using the seemingly best performing combination ({Current, Current}), initially underperform as compared to the baseline. This is to be expected, as the baseline employs signal actuation and adaptive signal timing. Fixed timing schemes in practice are inherently inferior to schemes properly employing actuation and adaptive timing.

#6303 shows a noticeably different pattern from the other intersections with the {Current, Current} profile combination. This is possibly due to signal phase order, the number of lanes per road, or some combination of the two. Left and through phases for the same direction of travel in #6303 are more often concurrently active in this intersection than for #7381 and #7204. This means that vehicles entering from a specific direction at #6303 are often given a green light for all turning movements, effectively allowing a single direction of travel to control the intersection for long periods of time. This may also contribute to efficient operation for HVs in high traffic volumes for #6303, limiting the initial improvement brought on by CAVs. The other potentially contributing factor is that #6303 has a large number of lanes per road which means that green trajectories are more likely to be active at any given time. Either of these could potentially stifle initial reductions in delay obtained by integrating CAVs.

The {Current, Restrictive} and {Permissive, Restrictive} policy combinations for #7381; {Permissive, Current} policy combination for #7204; and {Permissive, Current} policy combination for #6303 all show a hump or plateau in terms of delay. The cause of these trends is potentially a complex interaction between various parts of the intersection system. However, there are some behaviors of CAVs which are detrimental in terms of delay that definitely contribute to the odd trends seen for some of these intersections.

CAVs will block lanes in certain conditions. See Figure 3.6 for a representative example. In this example, HVs are assigned a more strict turning policy than CAVs (such as in the case of {Permissive, Restricted}). Vehicle 1 is a CAV and would like to turn left from the middle lane. Assuming that a green signal is assigned to the east and westbound roads, vehicle 1 is blocked from obtaining a reservation due to an active green trajectory. This active green trajectory is caused by continually arriving eastbound HVs (vehicle 2 for instance). Vehicle 1, being unable to obtain a reservation, blocks all vehicles behind it from entering the intersection. Imagine vehicle 3 is a HV and would like to continue straight. As long as vehicle 2 blocks the way, vehicle 1 is unable to cross the intersection. Note that depending on the particular turning profile, the associated green light may or may not apply to the turning action vehicle 1 is trying to take. This blockage very harmful for an intersection such as #7204 with the {Permissive, Current} policy combination, as

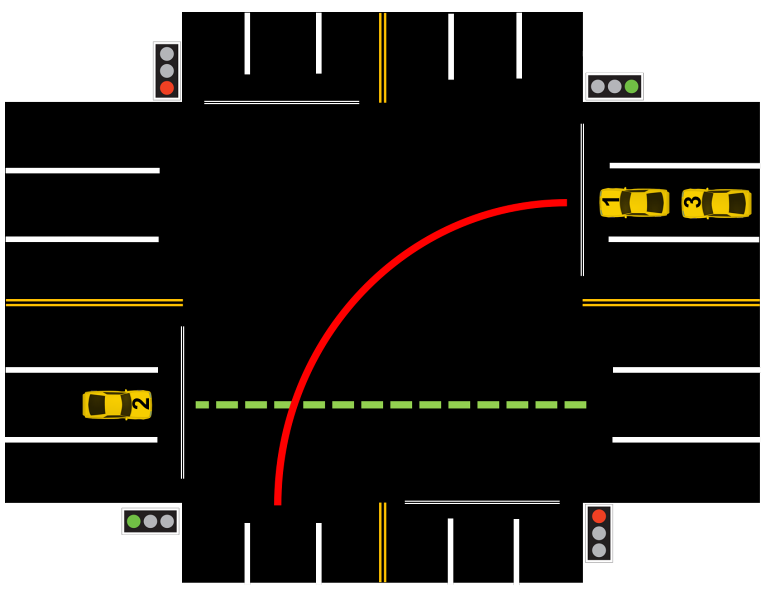

Figure 3.6: An example of a CAV blocking a lane due to being unable to obtain a reservation because of an active green trajectory. The green dashed line represents active green trajectory while the red line represents a path for an unapprovable reservation for the CAV. Vehicles 2 and 3 are HVs while vehicle 1 is a CAV.

the only through lane for the eastbound and westbound roads can be easily blocked by a CAV. This also seems to be a contributing factor to the high CAV delay with low CAV percentages for the {Permissive, Current} experiment for #7381.

A similar, but potentially even more detrimental scenario when it occurs, can be seen in Figure 3.7. Here, both vehicles 1 and 2 are CAVs which wish to make a left turn and have been assigned a more liberal policy than HVs. Thus the current green signals do not apply to the CAVs' desired turning movement. However, both vehicles 1 and 2 are trailed by HVs. Because this thesis assumes IMs are only able to determine if there is an incoming HV on a lane, but not which vehicle is the HV, these two trailing HVs create active green trajectories which prevent the CAV on the opposing road from gaining a reservation. This is a deadlock. Though, the deadlock can be unblocked once the traffic signals change and the active green trajectories are no longer present, provided the intersection is clear. This, as well as the situation shown in Figure 3.6, can simply be avoided by

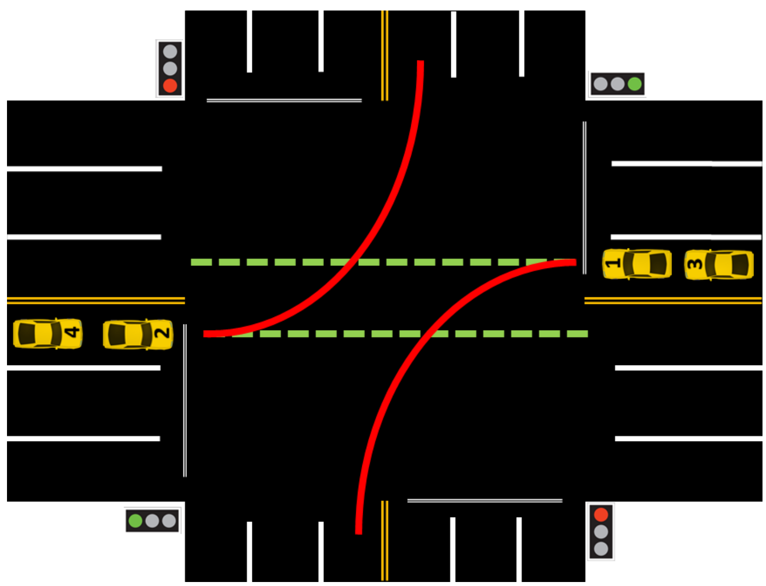

Figure 3.7: An example of multiple CAVs blocking lanes due to being unable to obtain a reservation because of active green trajectories arising from an IM's inability to distinguish between vehicle types at an intersection. The green dashed lines represents active green trajectories while the red lines represents paths for unapprovable reservations for CAVs. Vehicles 1 and 2 are CAVs while 3 and 4 are HVs.

assigning identical turning profile combinations to all vehicle types provided that permissive left turns on green are not permitted. Assigning identical turning profiles also sidesteps potential safety issues related to the exact ordering/method an IM may employ to clear this type of deadlock if the IM were so equipped to be able to identify the deadlock in the first place.

In order to outline turning policy assignments for which a deadlock such as in Figure 3.7 cannot occur, consider the cause of the deadlock once more. The deadlock only occurs when at least 2 CAVs arriving from opposite directions are both trailed by HVs in the same lane which trigger active green trajectories. If a deadlock were to occur, these trajectories would be blocking the desired reservations for the CAVs. Note that the traffic signal must be green for both directions for this to occur. There are then 4 possible configurations which do not result in a deadlock:

1. The simultaneous triggering of active green trajectories which block travel by CAVs does not occur due to HVs being absent for at least one direction of travel which has a green light. CAVs in this case may be delayed by HVs crossing the intersection, but will not be deadlocked.

2. HVs do not trail the CAVs at all but are still present. In this situation the HVs will proceed through the intersection on active green trajectories. Thus no deadlock will occur. Though, again, CAVs may be delayed as HVs traverse the intersection.

3. Trivially, CAVs are absent for a direction of travel which has a green light and so no deadlock can occur.

4. HVs are present and behind CAVs arriving from opposite directions as previously mentioned, but no green trajectory from either of the two active phases crosses any green trajectory from the other phase. Thus there is no green trajectory that could become active to block a CAV's reservation.

Because an intersection cannot practically choose the lanes on which HVs arrive or when any vehicle arrives, #1-3 in the in the list above are impractical. Thus #4 must be employed as the target condition to prevent a deadlock. Note that in the example in Figure 3.7, the green trajectories associated with allowed CAV turning actions are not the same as those associated with HV turning actions. This is what gives rise to the deadlock. If permissive left turns on green are not allowed, any homogeneous set of safe turning policy combinations will suffice to fulfill the requirements of #4. This is because no turning policy allowing green trajectories for human vehicles which intersect with any other green trajectory is safe or usable. However, if permissive left turns on green are allowed for a selected turning policy and vehicles my also proceed straight from the same lanes they may turn permissively, this deadlock may be unavoidable without specific tailoring of signal timing or addressing this as a special case for IMs employing the H-AIM protocol.

3.4 Adaptive Timing with Actuation Variation

Next, performance using adaptive signal timing with and without actuation is shown. The best and most consistently well performing turning profile assignments from the previous experiment are used ({Current, Current} for all 3 intersections). Here, all green times have a minimum of 4 seconds, but the maximum green times are varied throughout the day. Gap times are ignored if actuation is disabled. Information on signal optimization from the U.S. Department of Transportation Federal Highway Administration's 2008 Traffic Signal Timing Manual [28] was used as a basis for the automatic adjustments made during the simulations.

For all intersections the actuation detection distance was set to just under 2 meters. The portion of table 5-10 in the 2008 Traffic Signal Timing Manual corresponding to 3.0 seconds was used in order to determine gap extension times per phase when applicable. Maximum green signal times were determined per phase, but are proportioned across all phases within a ring based on their relative demand and table 5-6 in the aforementioned manual. However, signals for right turning movements are always associated with a through movement in this experimental setting and the through demand is more significant than the right turning demand. Thus, evaluations of demand which are used to adjust maximum signal lengths exclude lanes which only permit right turns in order to better estimate the impact of different phases on delay. Key values (values used in conjunction to select a value in the table) are rounded to the nearest key in the appropriate tables.

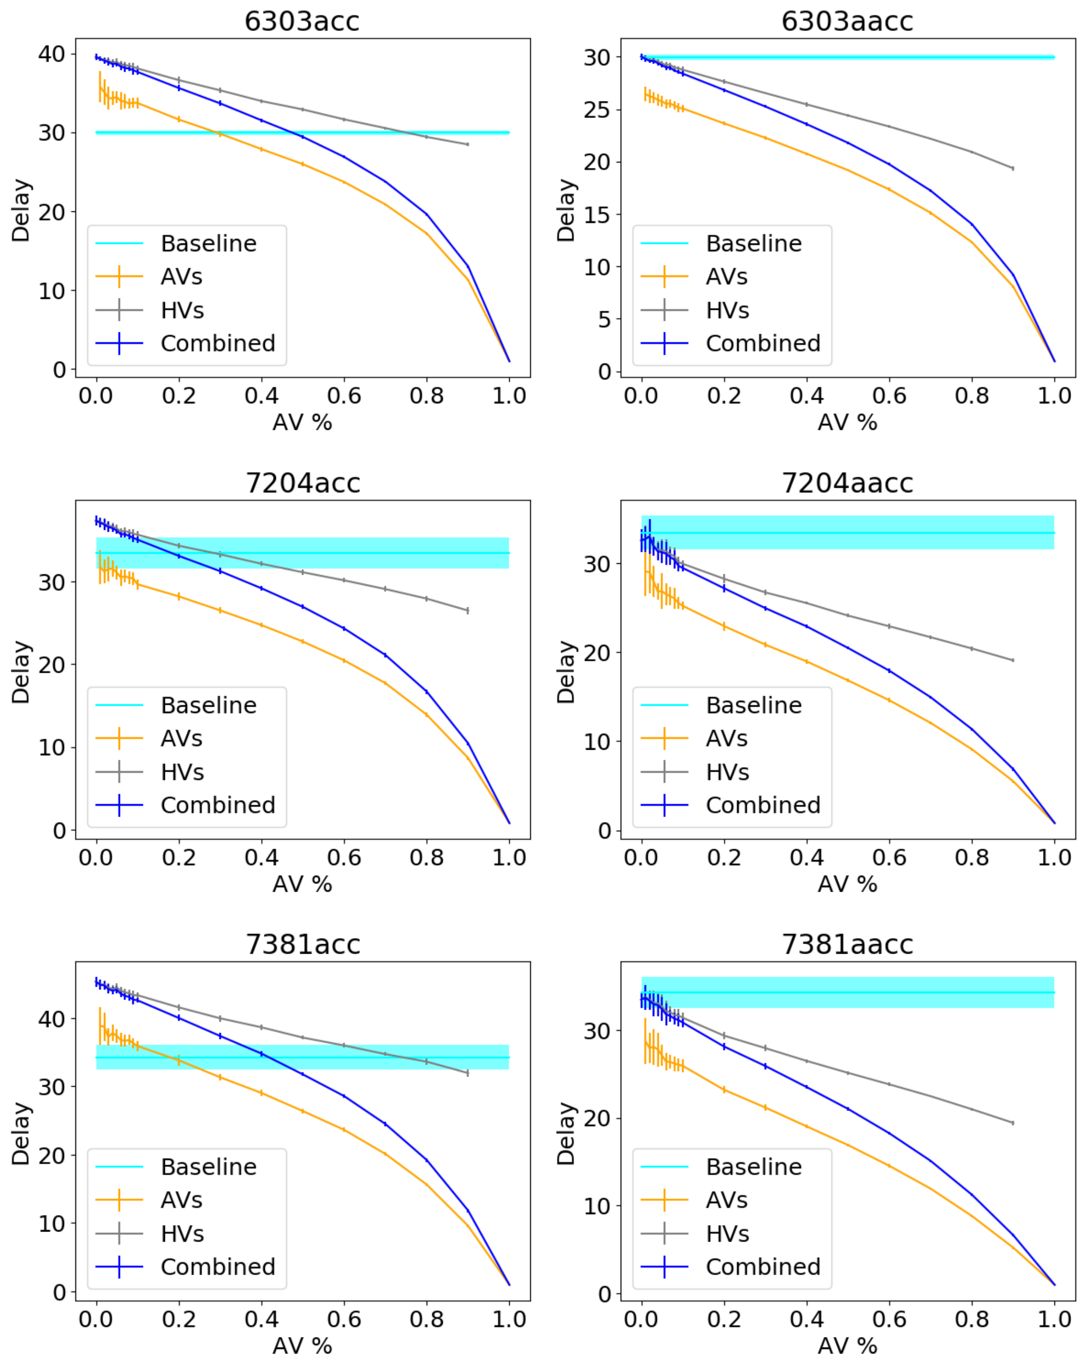

Figure 3.8: Graphs of delay by CAV percentage for both adaptive timing and the combination of adaptive timing and actuation. Each graph is labeled with the intersection number and then 3-4 letters representing the combination of actuated timing, adaptive timing, and turning policy combination in the order of CAV policy and then HV policy. "C" stands for the current policy in use by the real intersection, "p" represents a permissive policy, and "r" represents a restrictive policy. A single "a" represents adaptive timing only, while 2 represents adaptive timing with actuation. Each data point represents the average of the results 20 trials. Error bars show standard deviation of the results of those trials. A baseline with standard deviation is shown as a horizontal line.

3.4.1 Adaptive Timing and Actuation Analysis

Results can be seen in Figure 3.8. The left column shows results for all 3 intersections with adaptive signal timing and actuation disabled. This is an improvement in terms of delay for all intersections over the fixed timing scheme seen previously. An adaptive timing process, whether manually determined or automated as in this experiment, is better suited to deal with variations in demand than a fixed timing scheme. The fact that the maximum green times for the phases of traffic signal at an intersection are tailored to demand at any given time causes delay to be reduced. Though, all of the intersections in the left column initially underperform when compared to the baseline as well. This is still expected because the addition of properly configured actuation is inherently advantageous when considering delay.

The right column shows that experiments adding actuation into the mix further improve delay. On top of efficiently accommodating potentially high demand with the adaptive timing, actuation allows an IM to switch phases prior to waiting the maximum time for each phase in order to avoid wasting time when few vehicles are present for a particular phase. These experiments suggest that modern signal optimization techniques are beneficial to H-AIM in terms of delay reduction.

4. CONCLUSION

Hybrid-AIM (H-AIM) [8] is an efficient intersection management protocol for early CAV penetration stages. H-AIM builds on top of the Autonomous Intersection Management (AIM) protocol [1] and can be used in conditions where an IM may sense approaching vehicles provided the assumptions required by AIM are fulfilled. However H-AIM, and other works like it, are often demonstrated only in simulations with artificial conditions. This thesis demonstrates use of readily available historical data combined with information about real intersections to model simulations for an empirical study using H-AIM.

This thesis also extends the H-AIM protocol to cope with some uncertainty which is present when considering modern traffic signal behaviors. The proposed method to determine the possible set of signal phase segments potentially active at a future time enables the intersection manager to continue to approve autonomous vehicle reservations despite uncertainty introduced by actuated signal control. Additionally, a method by which intersection managers may cope with changing phase times throughout the day is discussed.

Results obtained from the empirical study support the following general conclusions:

- Combinations of similar turning profile policies for CAVs and HVs provide the most consistent and predictable improvement across varied CAV market penetration percentages.

- Combinations of more permissive turning profiles for CAVs and more restrictive ones for HVs are usually not preferable at low CAV market penetration percentages due to delays brought on by CAVs blocking lanes while waiting for a reservation. Future work could investigate the potential of changing turning profile combinations based on CAV market penetration percentages or other applicable factors.

- Modern signal optimization techniques such as varying signal parameters by time of day and signal actuation are seemingly compatible with and beneficial to the H-AIM protocol, especially at low CAV market penetration percentages.

Future extensions to work on the H-AIM protocol include development of a more precise method of dealing with uncertainty of signal timing, consideration of pedestrians within the protocol, variations on actuation such as skippable phases, implementation of early gap out functionality which terminates a green light when no vehicles are close to entering the intersection, studies into different intersection manager reservation approval policies, and eventual evaluation on a physical test bed with preferably real vehicles.

REFERENCES

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

[2] L. C. Bento, R. Parafita, S. Santos, and U. Nunes, "Intelligent traffic management at intersections: Legacy mode for vehicles not equipped with v2v and v2i communications," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 726–731.

[3] Y. Wu, H. Chen, and F. Zhu, "Dcl-aim: Decentralized coordination learning of autonomous intersection management for connected and automated vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 103, pp. 246–260, 2019.

[4] M. Bashiri, H. Jafarzadeh, and C. H. Fleming, "Paim: Platoon-based autonomous intersection management," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 374–380.

[5] M. O. Sayin, C.-W. Lin, S. Shiraishi, J. Shen, and T. Başar, "Information-driven autonomous intersection control via incentive compatible mechanisms," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 3, pp. 912–924, 2018.

[6] M. N. Mladenović and M. M. Abbas, "Self-organizing control framework for driverless vehicles," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 2076–2081.

[7] L. Riegger, M. Carlander, N. Lidander, N. Murgovski, and J. Sjöberg, "Centralized mpc for autonomous intersection crossing," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2016, pp. 1372–1377.

[8] G. Sharon and P. Stone, "A protocol for mixed autonomous and human-operated vehicles at intersections," in *International Conference on Autonomous Agents and Multiagent Systems*.

Springer, 2017, pp. 151–167.

[9] N. Murgovski, G. R. de Campos, and J. Sjöberg, “Convex modeling of conflict resolution at traffic intersections,” in *2015 54th IEEE Conference on Decision and Control (CDC)*. IEEE, 2015, pp. 4708–4713.

[10] K.-D. Kim and P. R. Kumar, “An mpc-based approach to provable system-wide safety and liveness of autonomous ground traffic,” *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3341–3356, 2014.

[11] X. Liu, P.-C. Hsieh, and P. Kumar, “Safe intersection management for mixed transportation systems with human-driven and autonomous vehicles,” in *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2018, pp. 834–841.

[12] T.-C. Au, N. Shahidi, and P. Stone, “Enforcing liveness in autonomous traffic management,” in *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, ser. AAAI’11. AAAI Press, 2011, p. 1317–1322.

[13] Q. Jin, G. Wu, K. Boriboonsomsin, and M. Barth, “Platoon-based multi-agent intersection management for connected vehicle,” in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 1462–1467.

[14] H. Wei, L. Mashayekhy, and J. Papineau, “Intersection management for connected autonomous vehicles: A game theoretic framework,” in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 583–588.

[15] M. Vasirani and S. Ossowski, “A market-inspired approach for intersection management in urban road traffic networks,” *Journal of Artificial Intelligence Research*, vol. 43, pp. 621–659, 2012.

[16] D. Carlino, S. D. Boyles, and P. Stone, “Auction-based autonomous intersection management,” in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 529–534.

[17] M. W. Levin and S. D. Boyles, "Intersection auctions and reservation-based control in dynamic traffic assignment," *Transportation Research Record*, vol. 2497, no. 1, pp. 35–44, 2015.

[18] M. VanMiddlesworth, K. Dresner, and P. Stone, "Replacing the stop sign: Unmanaged intersection control for autonomous vehicles," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 3*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 1413–1416.

[19] H. Mirzaei and T. Givargis, "Fine-grained acceleration control for autonomous intersection management using deep reinforcement learning," in *2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*. IEEE, 2017, pp. 1–8.

[20] K. Yang, S. I. Guler, and M. Menendez, "Isolated intersection control for various levels of vehicle technology: Conventional, connected, and automated vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 72, pp. 109–129, 2016.

[21] S. I. Guler, M. Menendez, and L. Meier, "Using connected vehicle technology to improve the efficiency of intersections," *Transportation Research Part C: Emerging Technologies*, vol. 46, pp. 121–131, 2014.

[22] Y. Feng, K. L. Head, S. Khoshmagham, and M. Zamanipour, "A real-time adaptive signal control in a connected vehicle environment," *Transportation Research Part C: Emerging Technologies*, vol. 55, pp. 460–473, 2015.

[23] J. Lee, B. Park, and I. Yun, "Cumulative travel-time responsive real-time intersection control algorithm in the connected vehicle environment," *Journal of Transportation Engineering*, vol. 139, no. 10, pp. 1020–1029, 2013.

[24] Q. He, K. L. Head, and J. Ding, "Pamscod: Platoon-based arterial multi-modal signal control with online data," *Transportation Research Part C: Emerging Technologies*, vol. 20, no. 1,

pp. 164–184, 2012.

[25] E. Namazi, J. Li, and C. Lu, "Intelligent intersection management systems considering autonomous vehicles: a systematic literature review," *IEEE Access*, vol. 7, pp. 91 946–91 965, 2019.

[26] L. Chen and C. Englund, "Cooperative intersection management: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 2, pp. 570–586, 2015.

[27] "Udot automated traffic signal performance measures - automated traffic signal performance metrics," accessed July 23, 2020. [Online]. Available: https://udottraffic.utah.gov/ATSPM

[28] "Traffic signal timing manual," accessed July 23, 2020. [Online]. Available: https://ops.fhwa.dot.gov/publications/fhwahop08024/chapter5.htm

[29] "Google maps," accessed 2020. [Online]. Available: http://maps.google.com/

[30] "Udot data portal," accessed July 23, 2020. [Online]. Available: http://data-uplan.opendata.arcgis.com/

[31] "Signalized intersections: An informational guide," accessed July 23, 2020. [Online]. Available: https://safety.fhwa.dot.gov/intersection/conventional/signalized/fhwasa13027/ch5.cfm

APPENDIX A

CONFIGURATION DATA

A.1 XML Intersection Configuration Example

```xml
<?xml version="1.0" encoding="UTF-8"?>
<intersection>
    <!--Direction, incoming, outgoing, speed in m/s-->
    <road>EAST, 3, 1, 13.4</road>
    <road>SOUTH, 4, 2, 20.1</road>
    <road>WEST, 3, 1, 15.6</road>
    <road>NORTH, 4, 2, 20.1</road>
    <!--Connections-->
    <direction>
        <from_to>EAST, EAST</from_to>
        <vehicle type="HUMAN">(1,0)</vehicle>
        <vehicle type="AUTO">(1,0)</vehicle>
    </direction>
    <direction>
        <from_to>EAST, NORTH</from_to>
        <vehicle type="HUMAN">(0,0)</vehicle>
        <vehicle type="AUTO">(0,0)</vehicle>
    </direction>
    <direction>
        <from_to>EAST, SOUTH</from_to>
        <vehicle type="HUMAN">(2,1)</vehicle>
        <vehicle type="AUTO">(2,1)</vehicle>
```

```
</direction>
<direction>
    <from_to>WEST, WEST</from_to>
    <vehicle type="HUMAN">(1,0)</vehicle>
    <vehicle type="AUTO">(1,0)</vehicle>
</direction>
<direction>
    <from_to>WEST, NORTH</from_to>
    <vehicle type="HUMAN">(2,1)</vehicle>
    <vehicle type="AUTO">(2,1)</vehicle>
</direction>
<direction>
    <from_to>WEST, SOUTH</from_to>
    <vehicle type="HUMAN">(0,0)</vehicle>
    <vehicle type="AUTO">(0,0)</vehicle>
</direction>
<direction>
    <from_to>NORTH, NORTH</from_to>
    <vehicle type="HUMAN">(1, 0), (2, 1)</vehicle>
    <vehicle type="AUTO">(1, 0), (2, 1)</vehicle>
</direction>
<direction>
    <from_to>NORTH, EAST</from_to>
    <vehicle type="HUMAN">(3, 0)</vehicle>
    <vehicle type="AUTO">(3, 0)</vehicle>
</direction>
<direction>
```

```
        <from_to>NORTH, WEST</from_to>
        <vehicle type="HUMAN">(0,0)</vehicle>
        <vehicle type="AUTO">(0,0)</vehicle>
    </direction>
    <direction>
        <from_to>SOUTH, SOUTH</from_to>
        <vehicle type="HUMAN">(1, 0), (2, 1)</vehicle>
        <vehicle type="AUTO">(1, 0), (2, 1)</vehicle>
    </direction>
    <direction>
        <from_to>SOUTH, WEST</from_to>
        <vehicle type="HUMAN">(3, 0)</vehicle>
        <vehicle type="AUTO">(3, 0)</vehicle>
    </direction>
    <direction>
        <from_to>SOUTH, EAST</from_to>
        <vehicle type="HUMAN">(0,0)</vehicle>
        <vehicle type="AUTO">(0,0)</vehicle>
    </direction>
</intersection>
```

A.2 XML Signal Phase Configuration Example

```
<?xml version="1.0" encoding="UTF-8"?>
<root>
    <ring>
            <!--Road direction outbound; -->
            <!-- cross ("c") turn (left in US) turn -->
            <!--or through ("t", and right);-->
```

```
        <!--gap timeout/extension; min time;-->
        <!-- max time-->

        <green>W, c, 5, 4, 35</green>
        <!--Alternative: Road direction outbound; -->
        <!--actions; time (min time == max time, -->
        <!-- no gap out) -->
        <yellow>W, c, 4</yellow>
        <red>W, c, 3</red>

        <green>E, t, 5, 4, 35</green>
        <barrier id="b1"></barrier>

        <green>N, c, 1, 4, 15</green>
        <yellow>N, c, 4</yellow>
        <red>N, c, 3</red>

        <green>S, t, 1, 4, 15</green>
        <barrier id="b2"></barrier>
</ring>
<ring>
        <green>E, c, 1, 4, 15</green>
        <yellow>E, c, 4</yellow>
        <red>E, c, 3</red>

        <green>W, t, 1, 4, 15</green>
        <barrier id="b1"></barrier>
```

```
            <green>S, c, 5, 4, 35</green>
            <yellow>S, c, 4</yellow>
            <red>S, c, 3</red>

            <green>N, t, 1, 4, 15</green>
            <barrier id="b2"></barrier>
    </ring>
    <!--yellow time, red time-->
    <barrier id="b1">4, 3</barrier>
    <barrier id="b2">4, 3</barrier>
</root>
```

APPENDIX B

ADDITIONAL EXPERIMENT DETAILS

	#6303	#7204	#7381
Eastbound Left	**11.02**	10.80	13.04
Eastbound Through/Right	36.84	**14.54**	**21.96**
Westbound Left	10.21	**8.13**	**13.40**
Westbound Through/Right	**36.34**	15.73	17.66
Northbound Left	11.44	**6.33**	**7.62**
Northbound Through/Right	**31.79**	57.68	45.74
Southbound Left	**12.51**	3.90	10.88
Southbound Through/Right	33.18	**59.11**	**44.34**

Table B.1: Fixed traffic signal timing selection in seconds for each intersection based on demand percentages, split and scaled across the available green time for a 120 second cycle. Values are rounded down to the nearest hundredth second from the actual calculated value. Note that values in bold are situated in one ring while non-bolded values are situated in another ring for the same intersection.

CAV Ratio	#6303		#7204		#7381				
	Turning Policy Variation								
	C, C	P, C	C, C	P, C	C, C	P, C	C, R	P, R	R, R
0	42.6	42.7	*83.1	*83.1	*95.6	*95.2	*132.1	*131.7	*132.5
0.01	42.4	42.5	*81.3	*86.6	*93.7	*84.8	*111.5	*104.0	*131.0
0.05	41.8	42.1	*74.9	*101.2	*89.9	*79.4	*103.9	*103.7	*125.7
0.1	41.0	41.5	*68.2	*106.4	*84.9	*72.3	*96.1	*102.7	*119.4
0.2	39.4	41.1	*56.2	*109.4	*76.5	*63.0	*82.2	*98.8	*109.1
0.3	37.9	40.9	*46.3	*105.7	*64.1	*60.8	*63.5	*93.8	*95.1
0.4	36.1	*40.3	*38.4	*100.7	*54.6	*62.3	*51.7	*92.8	*83.8
0.5	34.2	*39.1	32.4	*94.8	*42.4	*56.1	*39.9	*85.7	*65.7
0.6	31.6	*36.8	27.2	*85.6	34.2	*48.1	30.7	*71.8	*47.6
0.7	28.4	*32.5	22.6	*69.8	27.6	*32.1	25.0	*41.9	*32.5
0.8	23.6	26.0	17.3	*38.2	21.1	20.8	18.4	19.3	*23.5
0.9	15.3	16.0	10.4	*12.1	12.6	11.6	10.7	10.0	13.9
1	0.9	1.1	0.8	0.8	0.9	1.0	0.9	1.0	1.0
	Adaptive Timing with Actuation Variation								
	A	AA	A	AA	A	AA			
0	*39.5	30.0	37.3	*32.5	45.3	*33.5			
0.01	*39.3	29.8	37.1	32.7	44.9	*33.7			
0.05	*38.7	29.2	36.1	*31.2	44.1	*32.4			
0.1	*37.7	28.4	35.1	29.5	42.6	30.9			
0.2	*35.6	26.8	33.1	27.2	40.0	28.1			
0.3	*33.7	25.2	31.2	24.9	37.4	25.9			
0.4	31.5	23.6	29.2	22.9	34.8	23.5			
0.5	29.4	21.8	27.0	20.5	31.8	21.0			
0.6	26.9	19.7	24.3	17.9	28.6	18.2			
0.7	23.7	17.2	21.1	14.9	24.5	15.1			
0.8	19.6	14.0	16.7	11.3	19.2	11.2			
0.9	13.0	9.2	10.4	6.9	11.8	6.6			
1	0.9	0.9	0.8	0.8	0.9	0.9			

Table B.2: Data table showing the aggregated average results for delay for select CAV percentages in seconds rounded to the nearest tenth second. "R" is the restrictive policy, "P" is the permissive policy, and "C" is the current/real policy. Policies are listed as "CAV Policy, HV Policy". "A" represents adaptive timing while "AA" represents adaptive timing with actuation. "*" represents that data from AV percentages at this point are a lower bound, as the simulator may have had difficulty spawning vehicles at scheduled times in at least 1 simulation due to a lane being full.