In *Autonomous Agents and Multiagent Systems - {AAMAS} 2017 Workshops, Best Papers.*
Lecture Notes in Artificial Intelligence pp. 151--167, Springer International Publishing, 2017.

A Protocol for Mixed Autonomous and Human-Operated Vehicles at Intersections

Guni Sharon[1] and Peter Stone[1]

Department of Computer Science, The University of Texas at Austin, Austin, Texas, U.S.A
gunisharon@gmail.com, pstone@cs.utexas.edu

Abstract. Connected and autonomous vehicle technology has advanced rapidly in recent years. These technologies create possibilities for highly efficient, AI-based, transportation systems. One such system is the Autonomous Intersection Management (AIM), an intersection management protocol designed for the time when all vehicles are fully autonomous and connected. Experts, however, anticipate a long transition period during which human and autonomously operated vehicles will coexist. Unfortunately, AIM has been shown to provide little or no improvement over today's traffic signals when less than 90% of the vehicles are autonomous, making AIM ineffective for a large portion of the transition period. This paper introduces a new protocol denoted Hybrid Autonomous Intersection Management (H-AIM), that is applicable as long as AIM is applicable and the infrastructure is able to sense approaching vehicles. Our experiments show that this protocol can decrease traffic delay for autonomous vehicles even at 1% technology penetration rate.

Keywords: Autonomous Intersection Management, Autonomous vehicles, Multiagent systems

1 Introduction

Autonomous driving capabilities are becoming increasingly common on vehicles. Such capabilities present opportunities for developing safer, cleaner and more efficient road networks. Looking towards a future when most vehicles are autonomous and connected, Dresner and Stone proposed a novel intersection control protocol denoted Autonomous Intersection Management (AIM) [5]. AIM was shown to lead to significant traffic delay reductions when compared to traditional traffic signals.

Connected and autonomous vehicles (CAVs), with the help of advanced sensing devices, are more accurate and predictable compared to *human operated vehicles* (HVs). By relying on the fine and accurate control of CAVs along with communication capabilities, the AIM protocol coordinates multiple vehicles to cross an intersection simultaneously.

The AIM protocol defines two types of autonomous agents: intersection managers, one per intersection, and driver agents, one per vehicle. Intersection managers are responsible for directing the vehicles through the intersections, while the driver agents are responsible for controlling the CAV to which they are assigned.

To improve the throughput and efficiency of the system, the driver agents "call ahead" to the intersection manager and request a path reservation (space-time sequence) within the intersection. The intersection manager then determines whether or not this request can be met by checking whether it conflicts with any previously approved reservation or a potential HV. HVs are assumed to occupy all trajectories that are allowed by the traffic signal i.e., are given a green light. If the intersection manager approves a driver agent's request, the driver agent must follow the assigned path through the intersection. On the other hand, if the intersection manager rejects a driver agent's request, the driver agent may not pass through the intersection but may attempt to request a new reservation.

AIM, assuming 100% of the vehicles are CAVs, was shown to reduce the delay imposed on vehicles by orders of magnitude compared to traffic signals [6]. On the other hand, AIM was shown to be not better than traffic signals when more than 10% of the vehicles are HVs [5].

Given that experts speculate that 90% CAV penetration will not occur anytime before 2045 [3], this paper suggests a new protocol denoted *Hybrid AIM* (H-AIM) that is suitable for the transition period. Unlike AIM, H-AIM assumes sensing of approaching vehicles which allows it to identify approaching HVs. This assumption is reasonable given technological advances allowing vehicle detection using video cameras [4], radar [9], and inductive loop detectors [8]. If no HV is observed on a given lane, then trajectories originating from that lane are assumed to not be occupied by HVs, allowing AVs more flexibility in obtaining reservations.

A single lane entering a four-way intersection can allow three different turning possibilities (turn left, continue straight, turn right) or any combination of the three. The performance of H-AIM is sensitive to the assignment of allowed turns. This paper studies the effect of assigning different turning options to different lanes and different vehicle types (HVs, CAVs).

The main contributions of this paper are:

1. Defining the H-AIM protocol.
2. Presenting a comprehensive empirical study showing that H-AIM improves over traditional traffic signals even for as low as 1% CAV penetration. To the best of our knowledge H-AIM is the first protocol that is shown to be beneficial for low CAV penetration rates. This attribute makes H-AIM relevant for the long transition period expected to take place.
3. Presenting guidelines, potentially useful for practitioners, for assigning allowed turning options from each incoming lane to both autonomous and human operated vehicles such that different traffic measurements are optimized.

2 Background

The work presented in this paper builds on top of the FCFS+Signals policy which is part of the Autonomous Intersection Management (AIM) protocol [5]. This section provides a short overview of both AIM and the FCFS+Signals policy.

2.1 Autonomous Intersection Management

AIM is a reservation-based protocol in which vehicles request to reserve trajectories crossing an intersection. The AIM protocol assumes that computer-controlled vehicles attempt to obtain a right of passage through the intersection by sending a reservation request message to the *intersection manager* (IM). When using a "First Come, First Served" (FCFS) policy, the IM approves reservation requests that do not conflict with any previously approved reservation or potential HVs. In brief, the protocol proceeds as follows.

1. An approaching CAV, *cav*, sends a message to the IM requesting a reservation. The *request-reservation message* contains data such as the vehicle's size, predicted arrival time, velocity, acceleration, and arrival and departure lanes.
2. The IM processes the request message by simulating the trajectory of *cav* through the intersection, the simulated trajectory is denoted by $path(cav)$.
3. If $path(cav)$ does not conflict with any previously approved reservations or potential HVs then the IM issues a new reservation based on $path(cav)$ and sends an *approve message* containing the new reservation back to *cav*.
4. If $path(cav)$ does conflict with a previously approved reservations or potential HVs then the IM sends a *reject message* to *cav* which, after a predefined time period, may request a new reservation.
5. After receiving an approve message, it is the responsibility of *cav* to arrive at, and travel through, the intersection as specified in $path(cav)$ (within a range of error tolerance).
6. A CAV may not enter the intersection unless it successfully obtained a reservation.
7. Upon leaving the intersection, the CAV informs the IM that its passage through the intersection was successful.

The AIM protocol does not rely on communication capabilities between vehicles (V2V) only between vehicles and the IM (V2I). The protocol is robust to communication failures: if a message is lost, either by the IM or by the CAV, the system's efficiency might be reduced, but safety is not compromised. Safety is guaranteed also when considering a mixed scenario where both HVs and CAVs are present . For such cases Dresner and Stone [5] introduced the FCFS+Signals policy.

2.2 FCFS+Signals

The FCFS+Signal policy [5] is a combination between AIM and traditional traffic signals. Whenever the traffic signal is green for a given lane, all vehicles arriving at that lane have the right of passage (excluding unprotected left turns). However, when the traffic signal shows a red light, only CAVs which were granted a reservation may drive through the intersection.

Since the protocol is not assumed to know the location and trajectory of HVs, such vehicles are assumed to occupy all trajectories that are approved by the traffic signal i.e., have a green light. In this paper we define such trajectories as *green trajectories*. Figure 1 shows an example of green trajectories across an intersection (both the solid and dashed lines represent green trajectories). Note that green trajectories are dynamically changing; once the signal changes, the green trajectories will also change. The signal's timing is assumed to be known so the protocol is able to predict green trajectories in advance.

FCFS+Signals prohibits CAVs from obtaining reservations that conflict with green trajectories. In our example from Figure 1 all reservation requests will be

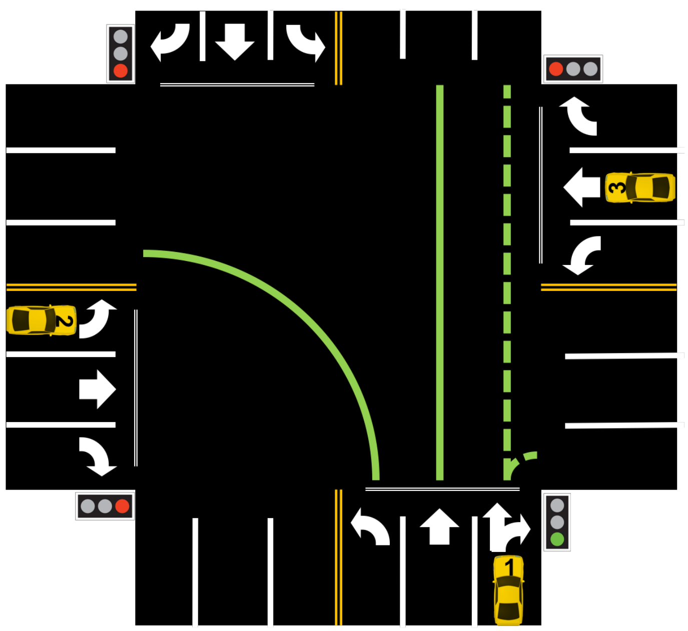

Fig. 1. Four-way intersection. Green light for all lanes originating from the South while all other lanes have a red light. Green trajectories marked with a solid or dashed green lines across the intersection. Active green trajectories marked only by dashed green lines.

automatically denied except those made by CAVs arriving from the south and those arriving from the North or East and request to turn right.[1]

2.3 Experimental Results

Dresner and Stone [5] reported *average delay* for a mixture of CAVs and HVs obtained from the AIM simulator running the FCFS+Signals policy. Delay is defined as the increase in travel time for a vehicle caused by red traffic signals or other vehicles. For CAV penetration of 90% and below, FCFS+Signals yielded a mild improvement. The improvement is attributed to CAVs that make right turns on red lights. If HVs are assumed to be able to turn right on red lights (as in the USA) or turning right has a designated lane bypassing the intersection, then there may be no improvement at all.

For AV penetration greater than 90% the *one-lane signal policy* was suggested which yielded a significant reduction in average delay. In the one-lane signal policy, right of passage for HVs (i.e., green light) is given to a single lane at a time instead of an entire road (all lanes arriving from the same direction). The one-lane signal policy results in a significant reduction in green trajectories at the cost of increased delay for HVs. As a result, the one-lane signal policy proved to be inefficient when considering high HV percentage (more than 10%).

3 Intersection Management Protocol for Mixed Traffic

CAVs are expected to penetrate the automobile market gradually over many years. Reaching 90% AV penetration rates will probably not happen in the near future [3]. Hence, a new intersection management protocol is required for managing traffic that is comprised mostly of HVs.

3.1 Assumptions and Desiderata

The new intersection management protocol should provide the following:

- Reduce the average delay suffered by vehicles crossing the intersection. Reduced delay translates into increased social welfare of the passengers.
- Reduce queue length on incoming lanes. Once the vehicle queue is longer than the length of the incoming link, a phenomenon known as queue spillback occurs [1]. Queue spillbacks have a negative cascading effect and should be avoided as much as possible [10].
- Increase throughput. Higher intersection throughput helps reduce congestion accumulated on links leading to the intersection.
- Provide a relative advantage to CAVs over HVs so as to incentivize drivers to transition to CAVs which are assumed to be safer [7] and more efficient [11].

[1] This paper assumes driving on the right side of the road. However, the ideas can trivially be generalized to a left side driving policy.

In contrast to FCFS+Signals we make the following assumptions:

– Humans may turn right on red light if the path is clear. This is a common case in the USA. An alternative assumption is that a right turning lane follows a trajectory outside of the intersection (right turn that bypasses the intersection yields an effect similar to turning right on red).
– A sensor (loop detector, camera or radar) is able to detect approaching vehicles on each lane (sensing speed and heading is not assumed).

3.2 Hybrid AIM

We now present the Hybrid-AIM (H-AIM) protocol for mixed traffic intersection management. Similar to FCFS+Signals, H-AIM grants reservation in a FCFS order. However, while FCFS+Signals automatically rejects reservation requests that conflict with green trajectories, H-AIM rejects reservation requests that conflict with *active green trajectories*. Define an active green trajectory as a green trajectory with a HV present on it or on its incoming lane. Figure 1 illustrates active green trajectories shown as dashed green lines across the intersection (notice vehicle #1 on the incoming lane).

Active green trajectories are a subset of the green trajectories making H-AIM at least as efficient as FCFS+Signals; there can be no reservation that is approved by FCFS+Signals and denied by H-AIM. The other way around, on the other hand, is possible. As an example consider the setting depicted in Figure 1. Assume vehicle #2 is a CAV and is heading North. Under the FCFS+Signals policy vehicle #2 would be automatically denied a reservation as it crosses a green trajectory. H-AIM on the other hand, would consider such a reservation as it doesn't cross an active green trajectory.

Note that the existence of a CAV on an incoming lane does not incur an active green trajectory. This requires the system to be able to identify whether an approaching vehicle is of type CAV or HV. For doing so we suggest the following procedure:

1. v = the number of vehicles detected on a given lane, l.
2. r = the number of reservation requests from unique vehicles seeking to enter the intersection from lane l. Reservations are considered only if the specified exit time is greater than the current time.
3. If $v > r$ then assume a human vehicle on lane l.

Note that the above procedure is safe in the sense that it will never misidentify a HV as a CAV. In the case of faulty communication this procedure might misidentify a CAV as a HV but this does not pose a safety issue. It might, however, hurt efficiency since a green trajectory might, mistakenly, be considered active.

Safety can be compromised if HVs are allowed to change lanes in close proximity to the intersection. For this reason HVs must be prohibited from changing lanes within detection range.

4 Reducing 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 [5] presented the one-lane signal policy (see section 2.3). This policy results in green trajectories that originate from a single lane at a time, which, in turn, leads to a significant reduction in the number of green trajectories. On the other hand, the one-lane signal policy was shown to have a dramatic negative effect on HVs.

We suggest a more conservative approach for reducing the number of green trajectories, which restricts turning options for HVs. Revisiting Figure 1, 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 crossing the intersection by 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.

4.1 Turning Assignment Policy

As was shown in the previous section, the effectiveness of a managed intersection is affected by the allowed turning options in each lane. When considering a three-lane, four-way intersection, each incoming lane has between one and three turning options from the set {left, straight, right}. The turning assignment policy assigns each incoming lane with allowed turns.

For this study we consider four representative turn assignment policies that are depicted in Figure 2. The policies are ordered and labeled according to degrees of freedom. Define degree of freedom for a lane as the number of turning options minus one. Define degree of freedom for a policy as the summation of degrees of freedom of all lanes.

A restrictive turning policy is one that has a low degree of freedom which, in turn, translates to fewer green trajectories. Policy 0 is an extreme case representing the most restrictive turning policy (0 degrees of freedom). Policy 4 is an extreme case of a liberal turning policy.

Define *safe turning policy* as one in which trajectories originated from the same road never cross each other. Turning policy 4 is not safe while 0, $2a$ and $2b$ are. Define *safe turning policy combination* as two policies in which no trajectory from one policy crosses any trajectory from the other when both originate from the same road. {0, 4} is a safe turning policy combination (even though 4 is not a safe policy on it's own). {2a, 4} is not a safe turning policy combination. A turning policy combination is considered when assigning one turning policy for HVs and a different one to CAVs.

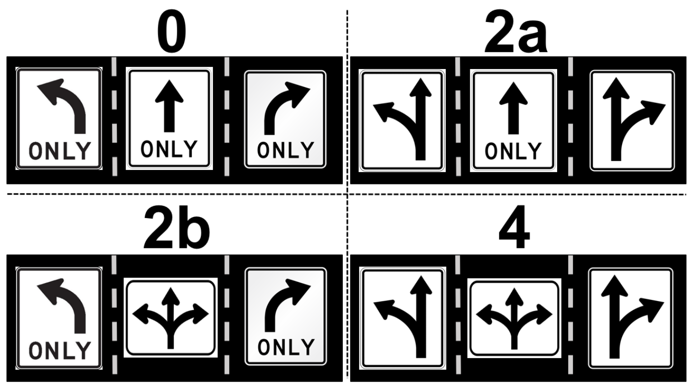

Fig. 2. 4 different turning assignment policies for a 3 lane road approaching a four way intersection.

For safety reasons we don't consider assigning HVs an unsafe policy. During our empirical study, we observed that assigning unsafe policy combinations for CAVs and HVs is counterproductive and should be avoided. Figure 3 demonstrates the inefficiency that stems from an unsafe 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, it is located in the middle lane and would like to turn right. Assuming a green light 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 vehicles behind it despite having a green light.

5 Empirical Study

This section presents results from a comprehensive empirical study. The goals of these experiments are two-fold:

- Study the effectiveness of H-AIM for mixed traffic with an emphasis on low CAV ratios.
- Indicate which turning policy should be assigned to HVs and CAVs in different CAV penetration and traffic levels.

Unless stated otherwise, our experiments used settings identical to those presented by Dresner and Stone [5]:

- Speed limit set to 25 meters/second

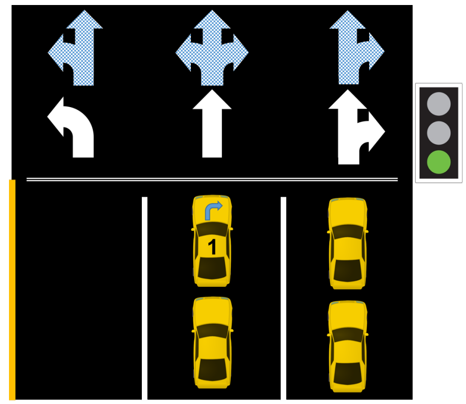

Fig. 3. An unsafe policy combination. Top policy (checkerboard texture) for AVs, bottom policy (plain texture) for HVs.

– CAV may communicate with the IM starting at a distance of 200 meters, which at 25 m/s (approximately 56 miles per hour) is 8 seconds before reaching the intersection.
– One simulated hour per instance. Results present the average over 20 instances per setting.

In line with our desiderata (presented in Section 3.1), we present average results for the following measurements:

– **Average delay -** see definition in Section 2.3.
– **Maximal queue length -** the maximal number of vehicles that simultaneously accommodate a single incoming lane. Note that 32 vehicles is the maximal queue length for any lane in the simulator, no new vehicles will be generated on a lane as long as this limit is reached. When high traffic volumes are considered, the maximal queue length is often reached and queue spillbacks occur. In such cases it is hard to compare different policies as they all return similar results making the maximal queue length measurement less valuable. Hence, we report maximal queue length only for low traffic levels.
– **Throughput -** the number of vehicles that passed the intersection in one hour. When low traffic volumes are considered the maximal throughput is often reached since all approaching vehicles eventually cross the intersection. At high traffic volumes, when queue spillbacks occur, throughput can give evidence on the severity of spillbacks i.e., the magnitude in which the spillbacks block new vehicles from entering the system. Hence, we report throughput only for high traffic levels.

The experiments presented in this section were obtained using the AIM4 simulator (http://www.cs.utexas.edu/ aim/). Several adaptations were required in order to run these experiments.

5.1 Modifications to the AIM simulator

Below is a list of changes introduced to the AIM simulator (on top of the original specifications [5]) for running our experiments.

- vehicles are spawned with equal probability on all roads, and are generated via a Poisson process which is controlled by the probability that a vehicle will be generated at each time step. Each vehicle is randomly assigned a type (HV or CAV) and destination. Given the assigned destination a vehicle is placed on an incoming lane from which it can continue to its destination (the incoming lane must allow turning to the vehicle's destination). If several such lanes exist it will be placed on the lane with the least number of vehicles currently on it. For instance, consider Figure 1, a vehicle arriving at the intersection from the South that is heading North would be assigned the middle lane since the left lane does not allow continuing North and the right lane already has one vehicle.
- CAVs are not granted reservations entering the intersection more than 3.5 seconds in the future. We add this constraint in order to allow the approaching vehicle detector enough time to detect all approaching HVs.
- A reservation is not necessarily denied if it conflicts with a green trajectory.
- A reservation is necessarily denied if it conflicts with an active green trajectory.

5.2 Four-way Intersection

Following Dresner and Stone [5] we start by presenting results from simulating a four way intersection with three lanes in each of the incoming roads (similar to the intersection presented in Figure 1). 0.2 of the vehicles turn right at the intersection, 0.2 turn left and 0.6 continue straight regardless of the incoming road and vehicle type.[2] A six-phase traffic signal timing was used (the signal timing is presented in Table 1).

[2] Dresner and Stone [5] do not report the turning ratios for their mixed traffic experiment. Our turning ratio was chosen since it results in a good balance between the incoming queues when 100% of the vehicles are HV.

Table 1. Six-phase traffic signal timing. Green and yellow duration are given in seconds. Asterisk next to a phase number means that left turns are allowed during that phase.

Phase	Origin	Green	Yellow
1	East-West	30	0
2*	East	15	3
3*	North	15	0
4	North-South	30	0
5*	South	15	3
6*	West	15	0

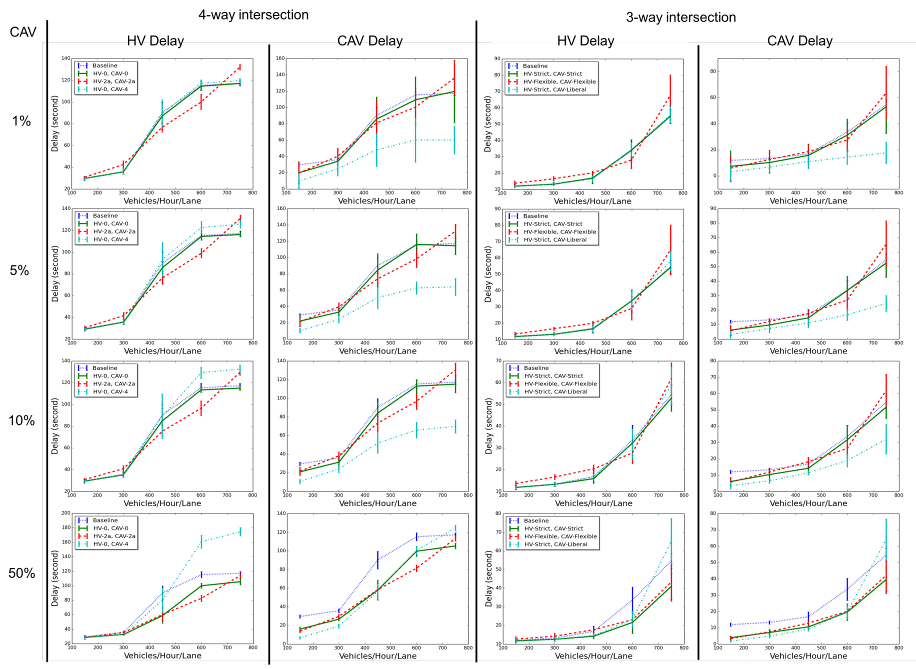

Fig. 4. Average delays for different traffic levels, CAV percentages, vehicle types, and intersection types. Each graph plots the average delay with 95% error intervals as a function of traffic level for three different turning policy combinations as well as the baseline (AIM).

Recall that under our assumption that HVs can turn right on red, the FCFS+Signals protocol has no advantage over traditional traffic signals (unless using the one-lane signal policy, see Section 2.3). Since FCFS+Signals using the one-lane signal policy was stated to be helpful when considering 90% HVs and more, it is not relevant to our current study which focuses on early CAV adoption stages. As a result the baseline for our experiments is the case where all vehicles yield to traffic signals while using turning policy 0 (similar to the turning policy used in [5]).

Figure 4 presents eight graphs for the four-way intersection case (left two columns). Each graph presents average delay in seconds (y-axis) versus traffic level in number of vehicles per hour per lane (x-axis). The data is presented for both HVs (first column) and CAVs (second column) and for different CAV penetration levels (with an emphasis on low CAV penetration levels - 1%, 5%, 10%, 50%).

Each graph compares three different safe turning policy combinations based on the policies presented in Figure 2. Note that results for turning policy 2b are not presented. Using the specified experimental settings, policy 2b was inferior to the other policies across all measurements hence it is omitted.

When examining HVs' delay, the results teach us that for low traffic levels (≤ 300) and very high traffic levels (> 700) policy {HV-2a, CAV-2a} is inferior, while for traffic levels in the range (400-700) it is superior. For medium and high traffic levels (> 400) policy {HV-0, CAV-4} is inferior. In 50% CAV penetration levels, policy {HV-0, CAV-4} is particularly inferior. When examining CAVs' delay, we see a clear benefit for policy {HV-0, CAV-4} across all traffic levels and CAV penetration levels except 50% penetration with high traffic levels ($>$ 450). The advantage of this policy is due to its reduction of green trajectories as explained in Section 4.1. Looking at delays of both HV and CAV we see that H-AIM with the base policy {HV-0, CAV-0} was superior to the baseline (FCFS+Signals with policy {HV-0, CAV-0}).

Table 2 presents maximal queue length and throughput for the four-way intersection scenario. At low traffic levels (150, 300) we report maximal queue length. On the other hand, at high traffic levels (450, 600, 750) we report throughput (see Section 5 for reasoning). Similar to Figure 4, results are presented for different traffic levels and different CAV penetration levels. We observe that avoiding congestion (minimizing queue length or maximizing throughput) is best achieved using policy {HV-2a, CAV-2a} regardless of the CAV penetration and traffic levels.

5.3 Three-way Intersection

Next we present results from simulating a three way intersection with two lanes in each of the incoming roads (similar to the intersection presented in Figure 5). 0.6 of the vehicles originating from the East or West continue straight while the rest (0.4) turn (either right or left depending on the incoming road). 0.5 of vehicles originating from the south turn right and the rest (0.5) left. We used a three-phase traffic signal timing presented in Table 3.

Figure 5 depicts three representative turning policies (Strict, Flexible, Liberal). Since a three-way intersection is not symmetrical, each turning policy is broken into three policies (one per origin road: West/East/South). We chose these three policies as they resemble the ones used in the four-way intersection experiment. "Strict" is the most restrictive policy, similar to policy 0 in the four-way case. "Flexible" has the highest degree of freedom among the safe policies, similar to policies 2a and 2b. "Liberal" has the maximal degrees of freedom overall, resembling policy 4. The baseline for our experiments is the case where 100% of the vehicles are HVs (i.e., all vehicles yield to traffic signals) using turning policy "Strict".

Figure 4 presents eight graphs for the three-way intersection case (right two columns). Each graph compares three different safe turning policies combinations based on the policies shown in Figure 5. The results show a picture which is somewhat similar to the one drawn from the four-way intersection scenario. When considering HVs' delay, policy {HV-"Flexible", CAV-"Flexible"} is superior for intermediate traffic levels (600) with the exception of 50% CAV pene-

Table 2. Results for a four-way intersection using different turning assignment policies for each vehicle type (HV, CAV) and different CAV penetration levels (CAV). Reporting maximal queue length for low traffic volumes and throughput for high traffic volumes.

	HV-0, CAV-0	HV-2a, CAV-2a	HV-0, CAV-4
	Maximal queue		
CAV	150 Vehicles/Hour/Lane		
Base	13.2	**8.4**	12.8
1%	13.6	**8.5**	13.2
5%	13.3	**8.4**	12.6
10%	13.6	**8.7**	12.1
50%	12.4	**7.9**	8.4
300 Vehicles/Hour/Lane			
Base	21.8	**16.5**	21.5
1%	22.1	**16.3**	20.8
5%	21.7	**15.3**	21.6
10%	21.2	**16.0**	20.1
50%	20.8	**14.2**	14.7
	Throughput		
CAV	450 Vehicles/Hour/Lane		
Base	4,621	**5,034**	4,630
1%	4,625	**5,039**	4,617
5%	4,639	**5,039**	4,647
10%	4,672	**5,057**	4,702
50%	4,865	**5,155**	5,118
600 Vehicles/Hour/Lane			
Base	4,989	**6,242**	4,983
1%	5,013	**6,239**	4,993
5%	5,029	**6,269**	5,027
10%	5,065	**6,309**	5,075
50%	5,367	**6,514**	5,804
750 Vehicles/Hour/Lane			
Base	5,314	**6,417**	5,328
1%	5,315	**6,429**	5,327
5%	5,361	**6,471**	5,414
10%	5,378	**6,520**	5,500
50%	5,718	**7,004**	5,972

tration levels where {HV-"Strict", CAV-"Strict"} proved most beneficial. Unlike the four-way intersection scenario, policy {HV-"Strict", CAV-"Liberal"} is never significantly inferior to other policies with the exception of 50% CAV penetration levels with traffic level of 750.

Similar to the four-way intersection scenario, when examining CAVs delay, we see a clear benefit for policy {HV-"Strict", CAV-"Liberal"} across all traffic levels and CAV penetration levels except 50% penetration with very high traffic levels (750). Similar to the four-way intersection scenario, H-AIM with base

Table 3. Three-phase traffic signal timing. Green and yellow duration are given in seconds. Asterisk next to a phase number means that left turns are allowed during that phase.

Phase	Origin	Green	Yellow
1	East-West	30	0
2*	East	15	3
3*	South	15	3

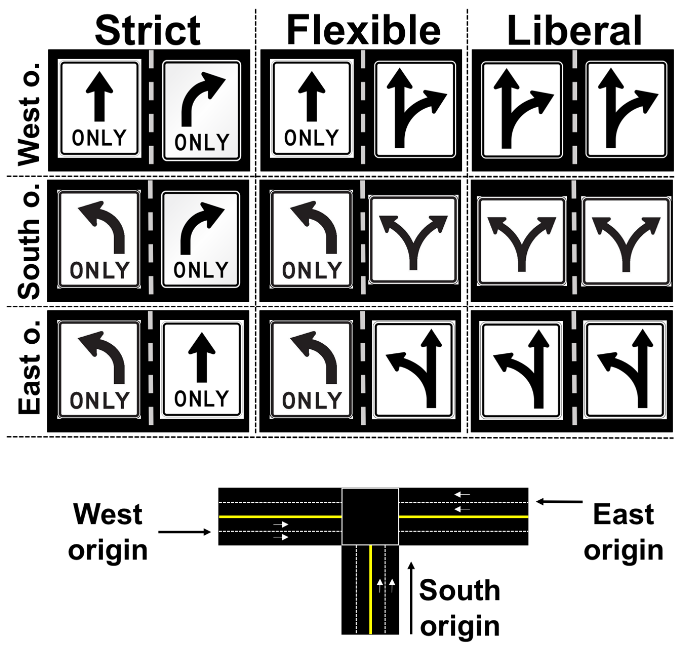

Fig. 5. 3 different turning assignment policies for a 2 lane road approaching a three way intersection.

policy {HV-"Strict", CAV-"Strict"} was superior to the baseline (FCFS+Signals with policy {HV-"Strict", CAV-"Strict"}) when examining delays over both HVs and CAVs. One exception is when considering HVs' delay at 1% CAV penetration, where H-AIM and the baseline performed similarly. Similar to Table 2, Table 4 presents maximal queue length and throughput but for the three-way intersection scenario. Again, we report maximal queue length for scenarios where queue spill back does not occur (traffic levels = {150, 300, 450}) else (traffic levels = {600, 750}) we report throughput. Similar to the four-way intersection scenario, we observe that avoiding congestion (minimizing queue length or maximizing throughput) is best achieved using policy {HV-"Flexible", CAV-"Flexible'}'

Table 4. Results for a three-way intersection using different turning assignment policies ("Strict" - S, "Flexible" - F, "Liberal" - L) for each vehicle type (HV, CAV) and different CAV penetration levels (CAV). Reporting maximal queue length for low traffic volumes and throughput for high traffic volumes.

	HV-S, CAV-S	HV-F, CAV-F	HV-S, CAV-L
	Maximal queue		
CAV	150 Vehicles/Hour/Lane		
Base	7.5	**6.7**	7.2
1%	7.5	**6.5**	7.6
5%	7.3	**7.0**	7.3
10%	7.2	**6.7**	6.9
50%	7.1	6.1	**5.1**
300 Vehicles/Hour/Lane			
Base	11.6	**10.9**	11.5
1%	11.3	**10.6**	11.6
5%	11.0	**10.8**	11.2
10%	11.0	**11.0**	11.3
50%	11.1	10.2	**8.5**
450 Vehicles/Hour/Lane			
Base	16.7	**15.2**	18.5
1%	17.1	**15.5**	16.9
5%	17.3	**14.6**	16.1
10%	17.0	**16.0**	16.3
50%	15.3	14.8	**11.4**
	Throughput		
CAV	600 Vehicles/Hour/Lane		
Base	3,239	**3,377**	3,253
1%	3,273	**3,388**	3,259
5%	3,275	**3,390**	3,301
10%	3,275	**3,391**	3,354
50%	3,355	3,407	**3,446**
750 Vehicles/Hour/Lane			
Base	3,754	**3,909**	3,774
1%	3,755	**3,907**	3,770
5%	3,793	**3,933**	3,862
10%	3,792	**3,941**	3,907
50%	3,942	**4,118**	3,975

with one exception at 50% CAV penetration level where {HV-"Strict", CAV-"Liberal"} was superior.

5.4 Conclusions

Concluding the empirical study we provide the following guidelines:

- H-AIM is superior to FCFS+Signals (baseline) when considering average delay.

– When considering congestion reduction, H-AIM is not superior to FCFS+Signals until more than a 10% CAV technology penetration level is reached.
– If seeking to encourage CAV adoption at early stages (0%-10% CAV penetration levels), one should set turning policies that restrict HVs to the maximum (such as policy 0 and policy "Strict") while allowing maximal flexibility to CAVs (such as policy 4 and policy "Liberal").
– Our experiments showed that setting an unsafe turning policy combination is never worthwhile. These results are not presented in this paper.
– When seeking to reduce congestion, non-restrictive safe turning policies (such as policy 2a and policy "Flexible") should be set for both HVs and CAVs. Note that setting a non-restrictive policy for HVs gives little or no advantage to CAVs and thus, does not encourage CAV adoption.

6 Summary

Though the Autonomous Intersection Management (AIM) protocol was shown to be be extremely efficient in coordinating Connected and Autonomous Vehicles (CAVs) traversing an intersection, it provides no improvement until 90% of the processed vehicles are CAV. This paper aims to enable efficient intersection management for early CAV penetration stages. To this end, we propose a modified AIM protocol denoted Hybrid-AIM (H-AIM). H-AIM is applicable under the assumption that vehicles approaching the intersection can be sensed (along with the assumptions required for AIM).

A comprehensive empirical study shows H-AIM to be superior to AIM when average delay imposed on vehicles is considered. Our study also gives guidelines as to how to assign turning options for each lane and vehicle type. Future work will study the effects of H-AIM when semi-autonomous vehicles [2] are considered and are assigned different turning policies. Future work will also examine restricting entire lanes to one vehicle type.

7 acknowledgments

A portion of this work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (CNS-1330072, CNS-1305287, IIS-1637736, IIS-1651089), ONR (21C184-01), AFOSR (FA9550-14-1-0087), Raytheon, Toyota, AT&T, and Lockheed Martin. The authors would also like to thank the Texas Department of Transportation for supporting this research under project 0-6838, Bringing Smart Transport to Texans: Ensuring the Benefits of a Connected and Autonomous Transport System in Texas. Peter Stone serves on the Board of Directors of, Cogitai, Inc. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

References

1. G. Abu-Lebdeh and R. Benekohal. Development of traffic control and queue management procedures for oversaturated arterials. *Transportation Research Record: Journal of the Transportation Research Board*, (1603):119–127, 1997.
2. T.-C. Au, S. Zhang, and P. Stone. Autonomous intersection management for semi-autonomous vehicles. *Handbook of Transportation. Routledge, Taylor & Francis Group*, 2015.
3. P. Bansal and K. M. Kockelman. Forecasting americans' long-term adoption of connected and autonomous vehicle technologies. In *Transportation Research Board 95th Annual Meeting*, number 16-1871, 2016.
4. B. Coifman, D. Beymer, P. McLauchlan, and J. Malik. A real-time computer vision system for vehicle tracking and traffic surveillance. *Transportation Research Part C: Emerging Technologies*, 6(4):271–288, 1998.
5. K. Dresner and P. Stone. A multiagent approach to autonomous intersection management. *Journal of artificial intelligence research*, 31:591–656, 2008.
6. D. Fajardo, T.-C. Au, S. Waller, P. Stone, and D. Yang. Automated intersection control: Performance of future innovation versus current traffic signal control. *Transportation Research Record: Journal of the Transportation Research Board*, (2259):223–232, 2011.
7. A. Furda and L. Vlacic. Enabling safe autonomous driving in real-world city traffic using multiple criteria decision making. *IEEE Intelligent Transportation Systems Magazine*, 3(1):4–17, 2011.
8. J. Gajda, R. Sroka, M. Stencel, A. Wajda, and T. Zeglen. A vehicle classification based on inductive loop detectors. In *Instrumentation and Measurement Technology Conference, 2001. IMTC 2001. Proceedings of the 18th IEEE*, volume 1, pages 460–464. IEEE, 2001.
9. J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt. Millimeter-wave technology for automotive radar sensors in the 77 ghz frequency band. *IEEE Transactions on Microwave Theory and Techniques*, 60(3):845–860, 2012.
10. Y. Liu and G.-L. Chang. An arterial signal optimization model for intersections experiencing queue spillback and lane blockage. *Transportation research part C: emerging technologies*, 19(1):130–144, 2011.
11. R. Regele. Using ontology-based traffic models for more efficient decision making of autonomous vehicles. In *Fourth International Conference on Autonomic and Autonomous Systems (ICAS'08)*, pages 94–99. IEEE, 2008.