Intersection Management Protocol for Mixed Autonomous and Human-Operated Vehicles

Aaron Parks-Young and Guni Sharon

***Abstract*— This paper presents a novel embedding protocol that allows for safe and efficient operation of the Hybrid Autonomous Intersection Management (H-AIM) protocol concurrently with actuated and adaptive signal controllers. The proposed protocol extends H-AIM to allow it to cope with some operational uncertainty that is common in actuated signal controllers. A novel approach for computing safety bounds on signal timing is presented as a way of insuring safety in the face of demand uncertainty. Experimental results show the feasibility and effectiveness of combining H-AIM with actuated controllers for various levels of connected and autonomous vehicle (CAV) market penetration and different combinations of common signal control schemes, namely, adaptive signal timing, fixed signal timing, and signal actuation. The benefits are presented in terms of delay improvement when common actuation protocols are used in conjunction with the H-AIM protocol. In contrast to previous reports, the results presented in this paper suggest that mixtures of turning movement assignments that are more permissive for CAVs and less permissive for human operated vehicles are often detrimental in terms of traffic delay. Nonetheless, when implemented on top of an actuated and adaptive controller, the extended H-AIM protocol is shown to never be detrimental while presenting statistically significant reductions in total delay when more than 15% of the traffic is composed of CAVs.**

***Index Terms*— Autonomous systems, road traffic control, human in the loop.**

I. Introduction

AUTONOMOUS driving capabilities are becoming increasingly common in 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, researchers are developing reservation-based intersection management protocols [1], [5], [6]. By relying on the fine and accurate control of *connected and autonomous vehicles* (CAVs) along with their communication capabilities, intersection management protocols coordinate multiple vehicles simultaneously across an intersection. Such protocols have been shown to lead to significant reductions in traffic delay when compared to traditional traffic signals. When assuming 100% of the vehicles are CAVs the seminal autonomous intersection management protocol (AIM) [6] was shown to reduce the delay imposed on vehicles by an order of magnitude compared to traffic signals [7]. On the other hand, AIM was shown to be not better than traffic signals when more than 10% of the vehicles are *human operated vehicles* (HVs). Similarly, the efficiency of most previously suggested intersection management protocols rely on traffic being mainly composed of CAVs.

Experts speculate that 90% CAV penetration will not occur any time before 2045 [3] deeming many previous protocols to be not relevant in the near future. To this end, [17] extended the AIM protocol and presented Hybrid-AIM (H-AIM) that is suitable for early stages of the transition period. Unlike AIM, H-AIM assumes the capability to sense approaching vehicles which allows the protocol to identify approaching HVs and assign safe reservations to CAVs accordingly. In order to foresee lanes that are potentially occupied by HVs, H-AIM assumes a fixed time signal controller. Such an assumption is not reasonable as many busy intersections are utilizing more advanced actuated and adaptive controllers. This article aims to close this applicability gap and reconcile H-AIM with common advanced signal controllers. This article presents a novel approach for computing time bounds on phase durations. The presented approach is useful for designating and approving safe reservation requests allowing for safe and efficient intersection management. The main contributions of this article are:

1) Extending the H-AIM protocol to be compatible with commonly deployed actuated traffic signal controllers. The proposed extended H-AIM protocol enables relaxation of some assumptions made by the original H-AIM protocol. Specifically, the extended protocol is shown to be compatible with:
   - **Actuated signal control:** A common traffic signal control policy where the amount of time which signals remain green may be extended as vehicles arrive at a green signal, up to predetermined maximums.
   - **Adaptive signal control:** Similar to actuated control, but where the time for signals to remain green is determined both by actuation events caused by vehicles and by recent traffic load at the intersection.
2) Introduction of a publicly available intersection management simulation system along with benchmark traffic scenarios that are based on real-world intersections along with real-world traffic demand.

Manuscript received 21 April 2021; revised 20 December 2021 and 17 March 2022; accepted 30 March 2022. Date of publication 6 May 2022; date of current version 11 October 2022. The Associate Editor for this article was I. Papamichail. *(Corresponding author: Guni Sharon.)*

The authors are with the Computer Science and Engineering Department, Texas A&M University, College Station, TX 77843 USA (e-mail: gunisharon@gmail.com).

This article has supplementary downloadable material available at https://doi.org/10.1109/TITS.2022.3169658, provided by the authors.
Digital Object Identifier 10.1109/TITS.2022.3169658

3) Demonstrating the performance of H-AIM and extended H-AIM in a computer-based simulation. The performance results are analyzed with respect to various turning action profiles (mapping of incoming lanes to possible outgoing ones), traffic signal control strategies (actuated, adaptive, fixed), and CAV market penetration rates.
4) Introducing a general guideline for selection of turning action profiles for extended H-AIM, based on observed experimental trends and analysis.

II. Background

Previous work discussed the potential of connected and autonomous vehicles (CAVs) to communicate with an intersection manager (IM) and/or each other as a way to coordinate movement through- and greatly improve efficiency at- intersections which service CAVs [4]–[6], [20]. Coordination of CAV movement through the intersection is commonly considered from a centralized perspective [15], [16] although decentralized coordination was also examined [13]. Such previous work usually builds on the seminal autonomous intersection management (AIM) protocol [6].

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 the reservation request conflicts with another already existing reservation or is otherwise deemed to be unsafe it will be rejected, else, it may be approved. Vehicles may not enter the intersection if they do not hold a reservation or are unable to follow their previously approved reservation, but may continue to submit subsequent reservation requests.

The AIM protocol was later extended to allow consideration of human driven vehicles (HVs) with no communication capabilities. The resulting, protocol H-AIM [17], requires the introduction of some restrictions on both CAVs and HVs which wish to traverse an intersection. Specifically, 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 are assumed to follow common traffic rules and are guided by visual (traditional) traffic signals. Autonomous vehicles are still required to make reservations to cross the intersection, but reservations will not be approved where they could potentially conflict with paths of HVs (assuming all HVs follow applicable traffic laws). Upstream sensors (such as cameras or loop detectors) are assumed 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. 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 efficient manner in many cases.

A. Preliminaries

Modifications to the H-AIM protocol are necessary to enable its use with adaptive signal timing and actuated signal functionality, which both introduce complications to reservation handling. This section discusses those complications and covers relevant definitions and concepts.

1) Phase Progression: Figure 1 (top) presents a simple ring and barrier signal diagram. For the purposes of this paper, each group of phases within a ring, consists of conflicting turn movements which, if signaled green together ("being active"), would create unsafe traffic conditions. Thus, within a single ring, only one phase may be active at a time, though multiple rings will have an active phase bounded by the same barriers. As an example, prior to the first barrier in the top portion of Figure 1, 1a or 1b may be active in concurrence with either 2a or 2b. However, 1a may not be active while 1b is active. As the associated timer expires, an active phase within a ring will transition to the next phase (to the right) in the current ring by appropriately signaling yellow and then red as the phase transitions. From the right most phase, the ring will transition back to the left most phase. In the case of actuated control, arrival of vehicles may extend the time the light will remain green in increments of up to the preset *gap extension* amount until the phase reaches a preset *maximum time*. Phase transitions which occur in different rings may be coordinated but are not necessarily so, except for when safety is a concern for particular phases (e.g., when left turns of opposite directions could lead to a collision or when transitioning across a barrier).

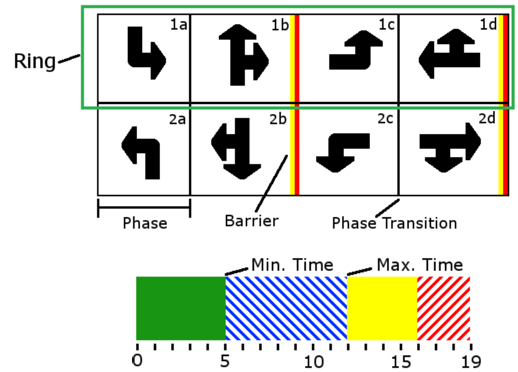

Fig. 1. Ring and barrier diagram (top) and an example of the timing for a simple actuated signal phase with a 19 second time scale (bottom).

The bottom of Figure 1 shows a simple diagram of the timing for a signal phase as used in this article. 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 shown in the example on the bottom of Figure 1 is 5 seconds. The striped portion with a length of 7 seconds represents the additional time the signal may remain green if more vehicles are approaching. The extension is not a requirement and the signal color may transition to yellow immediately after the minimum time has elapsed, this is usually the case if no approaching vehicle is detected. If, however, a vehicle is detected through some sensor (e.g., an inductive loop embedded in the roadway), the green time duration will usually be extended. Importantly, these extensions may not continue indefinitely as they may starve vehicles with a conflicting

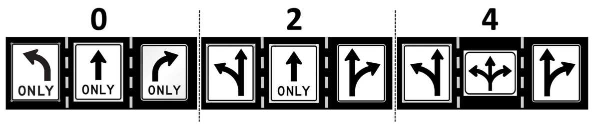

Fig. 2. Three turning assignment policies for a three-lane road approaching a four-way intersection with varying degrees of freedom.

direction of travel. Thus, a phase will typically only be extended at most up to a **maximum green time** (a total green time of 12 seconds in Figure 1), at which point no further activation of sensors will extend the signal for this phase. The exception to this behavior in this article is when the green time of a phase in one ring is complete but is "waiting" on another phase's green time to complete (e.g., before all rings transition through a barrier or preventing conflicting concurrent left turns). 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 1), then red for a fixed time (a duration of 3 seconds in the example in Figure 1) and then the right of passage will transition to the next phase.

2) Adaptive Signal Timing: This article considers a common actuated signal timing scheme that is adopted by the U.S. Department of Transportation Federal Highway Administration's 2008 Traffic Signal Timing Manual [9]. This article bases schedule changes in the actuated controller on Tables 5-6 and 5-10 in the manual.[1]

Two adaptive timing parameters are considered (1) the maximum green time for a phase and (2) the gap extension times. In this paper, these parameters are set according to Tables 5-6 and 5-10 in the Traffic Signal Timing Manual [9].

3) Green and Active Green Trajectories: As HVs are a major consideration in the H-AIM protocol, [17] define trajectories that must be reserved for HVs by IMs as "green" or "active green". 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 occupied by an HV or with an incoming HV 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 active green trajectory.

4) Turning Assignment Policy: As reported in [17], 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}. The sets of turning actions which are allowed on each incoming lane are denoted as the **turning assignment policy** hereafter. Following Sharon and Stone we define **degree of freedom** for a lane as follows.

Definition 1 (Turning Policy Degree of Freedom): Define degree of freedom for a lane as the number of turning options minus one. Define degree of freedom for a turning assignment policy as the sum of degrees of freedom over all lanes.

[1]Other dynamic actuation schemes are available [2], [8], [12] and can potentially work alongside H-AIM. However, adapting H-AIM to fit each of them is left for future work.

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 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.

Definition 2 (Consistent Turning Policy): A turning assignment policy is said to be consistent if trajectories originating from the same road never cross each other.

In the representative policy set, turning policy 4 is not consistent while policies 0 and 2 are. When considering more than one type of vehicle, different turning policy combinations might be considered. For instance, we might choose to assign one turning policy for HVs and a different one to CAVs.

Definition 3 (Consistent Turning Policy Combination): A set of turning assignment policies are said to be a consistent combination if no trajectory from one policy crosses any trajectory from any other policy when both originate from the same road.

In the representative policy set, {0, 4} is a consistent turning policy combination (even though 4 is not a consistent policy on its own). By contrast, {2, 4} is not a consistent turning policy combination. As [17] note, for safety reasons one shouldn't assign an inconsistent policy to HVs. On the other hand, assigning such a policy to CAVs is reasonable since conflicting reservations are monitored and prevented by an IM. Sharon and Stone also point out that assigning inconsistent policy combinations for CAVs and HVs is safe yet counterproductive from an efficiency standpoint and should thus be avoided.

III. Combining H-AIM With Adaptive and Actuated Signal Control

Part of the H-AIM protocol requires that CAVs call ahead to make reservations. A consideration of whether to approve a reservation is based on the status of traffic signals and whether HVs might enter or be within the intersection during the requested reservation. When fixed signal timing is considered, it's simple and efficient for an IM to "look ahead" to determine if a CAV's trajectory will cross any active green trajectory. However, adaptive and actuated controllers complicate this process. Specifically, this article assumes that an adaptive and actuated signal time controller is deployed and used for coordinating HVs while H-AIM is used concurrently for coordinating CAVs. In such cases predicting the set of future green trajectories is challenging due to potentially unforeseen actuation events. To make matters worse, use of the adaptive timing scheme means that maximum green times and gap extension times for different phases may change between subsequent look ups. Thus, without further modification to the H-AIM protocol, an IM naively looking ahead based on the current status of phases might approve reservations which could lead to a collision. Thus, the behavior of the H-AIM protocol must be extended to enable compatibility with actuated and adaptive signal controllers. Like the original H-AIM protocol, our extended H-AIM protocol makes the following assumptions.

1) Human vehicles follow traffic laws, but their trajectory through an intersection is unknown. That is, a human vehicle could choose any legal trajectory through the intersection without informing the IM.
2) An IM running the extended H-AIM protocol has adequate sensing technology allowing it to detect approaching vehicles per lane.

A. Protocol Extension for Actuated Signal Control

Actuated signal control adds uncertainty to phases' green interval duration. In practice, the green interval duration can span any value between the min and max time parameters affiliated with each phase. Due to this uncertainty and in order to guarantee safety, a phase must be assumed active, along with its affiliated green trajectories, at any time step that might allow it. As a result, when actuated control is considered, time bound intervals are computed for each phase and its accompanied green trajectories. For safety reasons, any trajectory is considered green for the entire time interval. Algorithm 1 describes how the relevant time intervals are computed per phase. This algorithm can be used to determine if a specific phase (and affiliated trajectories) is potentially active at a given future time. Consequently, multiple calls to this algorithm can be used to identify future green trajectories coinciding with a given reservation time. The *Compute-Phase*

Algorithm 1 Compute-Phase

Input: phase ϕ; boolean $entry$, True on entry call.
Output: Early end of the phase, late end of the phase.

1: **if** phase immediately precedes a barrier **and** $entry$ **then**
2:   **return** Compute-Barrier($suc(\phi)$)
3: **else**
4:   max $\leftarrow$ max phase ϕ duration
5:   min $\leftarrow$ min phase ϕ duration
6:   **if** phase ϕ is currently active **then**
7:     $t_e, t_l \leftarrow$ start of phase
8:   **else if** ϕ is after a barrier **then**
9:     $t_e, t_l \leftarrow$ Compute-Barrier($pre(\phi)$)
10:   **else**
11:     $t_e, t_l \leftarrow$ Compute-Phase($pre(\phi)$, False)
12:   **return** t_e + min + clearance, t_l + max + clearance

function (Algorithm 1) takes a specific phase and a boolean designating an entry call to the recursive function. It returns the earliest end time (number of seconds past now) and latest end time possible for the given phase following the current cycle of phase progression. The result combined with knowledge about the phase (such as phase turning assignment as well as maximum and minimum phase lengths) can be used to determine whether a phase is potentially green at a particular future time.

Proceeding to the line-by-line description, the first line is a check to determine if the specified phase immediately precedes a barrier and is the target phase (initial entry call). If the check returns true, a call is made on line 2 to the *Compute-Barrier* function (Algorithm 2) for the succeeding barrier ($suc(\phi)$) in a forward order and the result of this call is returned. This is because the end time for this phase is potentially dependent on the end times of the phases in other rings which lead into the same barrier. If either the phase which is currently being examined does not precede a barrier or this is not the entry call, execution proceeds to line 3. Lines 4 and 5 cache the maximum amount of time the phase (ϕ) could last and the minimum amount of time the phase could last. Line 6 checks for the stopping condition of the recursion, which is that the phase being examined (ϕ) is currently active (assigned right of passage at the current time). In such cases, line 7 sets t_e and t_l to the beginning of the phase in preparation for the return statement on line 12. Variables t_e and t_l represent the earliest and latest time the phase in question might begin. If the phase in question (ϕ) is not currently active, line 8 checks to see if it falls directly after a barrier. In such cases, Line 9 makes a call to the *Compute-Barrier* function (Algorithm 2) and stores the result in t_e and t_l. Lines 10 and 11 simply handle the remaining case where a phase precedes the current phase (ϕ). Line 13 calculates the earliest and latest end times for the preceding phase ($pre(\phi)$) and returns them. The clearance variable is the time of any leg (all red, yellow) between the end of the green signal of the current phase and beginning of the next phase.

Algorithm 2 Compute-Barrier

Input: barrier b.
Output: Early end of the barrier, late end of the barrier.

1: t_{e1}, t_{l1} = Compute-Phase($pred(b)$, 1, False)
2: t_{e2}, t_{l2} = Compute-Phase($pred(b)$, 2, False)
3: **return** max(t_{e1}, t_{e2}), max(t_{l1}, t_{l2})

Note: This function assumes there are only 2 rings (ring 1 and 2), though a more generalized formulation is possible.

The *Compute-Barrier* function (Algorithm 2) takes a barrier object as input. It returns the earliest and latest potential times the barrier would be crossed, which coincides with the earliest and latest times the preceding phases (one per ring) would end. On lines 2 and 3, the earliest and latest times the phases preceding the barrier would end if each were considered independently. The presented pseudocode assumes that there are only 2 rings, though a more generic formulation is possible by iterating over an arbitrary number of rings. Line 5 then handles the possibility of one preceding phase being artificially extended due to the other preceding phase (on a different ring) lasting longer.

B. Protocol Extension for Adaptive Signal Timing

Adaptive signal timing, when applied in parallel to H-AIM, adds uncertainty regarding the signal timing parameters. Namely, min and max duration per phase (i.e., 2 parameters per phase). Changing these parameter values even for a single phase can obsolete time bounds previously computed for green trajectories by Algorithm 1. As a result, previously approved reservations might no longer be safe to follow. We suggest handling this case by allowing H-AIM to "lock" relevant timing parameters once they are considered for a safety critical decision. These parameters can be "unlocked" and updated as soon as dependent reservations expire. Note that the adaptive timing protocol is orthogonal to the H-AIM protocol. The only caveat is that the H-AIM controller can temporarily limit the ability of the adaptive timing controller to change the timing parameters for specific phases. Though, the "lock" commands can theoretically result in permanent locking of timing parameters if certain restrictions are not imposed. Suppose some phase ϕ is considered when evaluating and approving some future reservation, r_1 (i.e., ϕ will be active sometime between now and r_1). In such a case, the parameters used for computing the bounds for ϕ would be locked until r_1 is cleared. Further assume that prior to the clearing of r_1 a new reservation, r_2, which relies on ϕ in the successive cycle, is approved. As a result, the parameters used for computing the bounds for ϕ will not be unlocked until r_2 is cleared. This process can theoretically continue indefinitely, practically preventing the adaptive parameters used by ϕ from ever being unlocked and updated. A simple solution to this issue would be to bound future reservations to be no further in the future than the minimum cycle duration (given current parameters) minus the (current) maximum phase duration. When applying such a restriction, no reservation that relies on a phase ϕ in a future cycle can be approved until after phase ϕ has terminated, allowing the parameters for ϕ to be unlocked and updated at least once per cycle. This restriction was implemented as part of our experimental setting that is presented in Section V by locking phase parameters as vehicles submit reservations and approach the intersection as well as unlocking phase parameters when a phase completes.

C. Processing Reservation Requests

Given a reservation request, a controller implementing the extended H-AIM protocol populates an *activation bound table* (ABT) with 'early' and 'late' activation bounds per phase. Each entry in the ABT specifies the earliest and latest (relative to the current time, 'Now') that a phase might be active during the current cycle. These values, for a given timestep, are calculated using Algorithm 1. For each phase ϕ, we store the 'early' end time and 'late' end time following 'late end' and 'early end' as returned from *Compute-Phase*$(\phi, True)$. Specifically, 'late' in the ABT is set to be the 'late end' while 'early' in the ABT is set to $\max(e - (m(\phi) + c), 0)$ where e is the 'early end' value, $m(\phi)$ is the max phase duration for ϕ, and c is the clearance time.

After the ABT is populated, the controller follows the original H-AIM protocol to determine whether a reservation can be approved. However, unlike the original protocol, extended H-AIM designates a trajectory as 'green' at time step $t+i$ if and only if it is enabled by a phase ϕ with $\phi[early] \geq i \geq \phi[late]$ according to the ABT.

D. Safety Considerations

The proposed extended H-AIM protocol results in several potential safety issues that must be properly addressed. These issues include the following scenarios.

1) A HV is mistakenly identified as a CAV causing a misclassification of an active green trajectory as not being active.
2) CAV impersonating a HV in order to take advantage of a green light and avoid the need for securing a reservation.
3) Approving a reservation while a potentially conflicting HV is beyond the current sensing range.
4) An approaching HV changes lane within the detection range to trigger a new active green trajectory that was not considered in previously approved reservations.

Note that the extended H-AIM procedure is safe in the sense that it will never misidentify a HV as a CAV (Issue #1) since HVs are not assumed to be able to self-identify as CAVs in front of the IM. In the case of faulty communication or intentional misrepresentation a CAV might be misidentify a HV (Issue #2) but doing so does not pose a safety issue. It might, however, hurt efficiency since a green trajectory might, mistakenly, be considered active. Intentional misrepresentation can be discouraged if approaching CAVs are expected to suffer from less delay compared to HVs. In such cases, a CAV agent seeking to minimize intersection imposed delays would be incentivized to truthfully report its nature.

A safety violation might occur due to delayed sensing of approaching HVs (Issue #3). For example, consider a reservation request, r, which is set to clear the intersection at time step t_{clear}. Further assume an approaching HV which triggers a conflicting active green trajectory. That is, the HV is expected to enter the intersection on a trajectory that intersects with r prior to t_{clear}. In case where r is submitted for approval prior to the detection of the incoming HV, it might be approved leading to a potential conflict. In order to avoid such situations, the extended H-AIM protocol bounds the time horizon for future reservations such that all potentially conflicting HVs are guaranteed to be within sensing range. For instance, when assuming a 250m downstream sensor on every incoming lane and that HVs cannot travel at speeds exceeding 90k/h, then reservations with t_{clear} − 'Now' $\geq 10s$ are considered unsafe and are always rejected. Safety can also be compromised, if HVs are allowed to change lanes in close proximity to the intersection (Issue #4). For this reason HVs must be prohibited from changing lanes within detection range, e.g., by using a physical barrier.

IV. Simulation Environment

Additions to the AIM Simulator (which is the simulator built for the work leading up to and including [6]) were made to facilitate the experiments presented in Section V. The experiments were set to run on real-world data about specific intersections in Utah, USA. This data was gathered from the Utah Department of Transportation's (UDoT or Utah DOT) Automated Signal Performance Measures (ATSPM) system [18], as well as other information available from Utah DOT [19], and GoogleMaps [11]. The modified AIM Simulator along with benchmark scenarios is publicly available at https://github.com/Pi-Star-Lab/Improved-H-AIM. The full technical description along with an example of the relevant input files are provided in a designated technical report [14].

V. Experimental Study

This section presents results from a simulation-based study using the modified version of the AIM Simulator. The goals of these experiments are three-fold:

1) Study the effectiveness of the original H-AIM protocol for mixed traffic with an emphasis on low CAV ratios while using real traffic data.
2) Study the impact of differentiated turning policy assignment (HVs vs. CAVs) in our modeled traffic scenarios.
3) Demonstrate the effectiveness of the extended H-AIM protocol in combination with actuation and adaptive signal timing.

A. Experimental Settings

Similar to the experiments presented in [6], results are reported as averages over 20 simulations per combination of CAV percentage, signal behavior, and assigned turning policy. Unlike Dresner and Stone, speed limits are set according to data gathered about the real intersection being modeled in each experimental setting. Note that Dresner and Stone considered a speed limit of 25 meters/second which is uncommonly high for signaled intersections.

Using readily available data from UDoT's ATSPM system [18], three real intersections in Utah were selected and modeled for use in each simulation: State St. at 800 North (#6303), 900 East at 5600 South (#7204), and 5600 West at 3500 South (#7381). The reported experiments are collected into two groups. The first group of experiments uses a fixed signal 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 allowed by the road. The best and most consistently well performing turning policy combination for each intersection from this set of experiments is carried forward into the next group of experiments in order to examine the performance benefits of using an adaptive signal timing scheme. Another group of experiments also compares results with and without signal actuation enabled. All experiments permit right turns on red for all turning profiles and disallow permissive left turns on green.

As a baseline for comparison, the average delay was calculated over 20 runs with 0% CAV rate using actuation and adaptive timing (which is akin to a manually built schedule by time of day), and the currently deployed turning policy per intersection (as reported in Figure 3). This baseline, which represents conditions similar to what is found in modern intersections, is contrasted against the average accumulated delay in each experimental setting. The percentages of vehicles which spawn as CAVs are varied from 0%-100% in all experiments. In order to emphasize the performance evaluation on the early CAV adoption period, this is first done in 10 increments of 1% (starting at 0%) and then the remainder is done in increments of 10%. Note that, following [6], [17], we assume the HV and CAV driving patterns to be similar when approaching the intersection with the difference that CAVs can obtain and follow reservations. Under this assumption, the baseline performance is not influenced by CAV penetration rates.

TABLE I

Fixed Traffic Signal Timing Setup (Seconds per Phase) 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

	#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**

A graphical depiction of the turning policies used in the simulations can be seen in Figure 3. The variants of turning policies are chosen as follows.

1) The most restrictive policy (fewest degrees of freedom).
2) The current (or "real") policy seemingly used by the real intersection.
3) A permissive policy (most degrees of freedom, only applicable 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, meaning only one representative direction is shown for each policy. The restrictive and current policies are identical when considering #6303 or #7204, and so those policies are merged and co-labeled for both of those intersections. #7381 requires a slight modification to its permissive policy which makes it slightly less permissive than is potentially possible. This modification is required in order to comply with the "consistent turning policy combination" (Definition 3) when pairing with the 'current' turning policy. Finally, the permissive policy for #6303 is also made slightly less permissive than possible to keep it consistent with the 'current' policy. Thus, by the definitions presented in Section II-A.4, only consistent turning policies and turning policy combinations are considered in the experiments in this article.

The U.S. Department of Transportation Federal Highway Administration's recommendation for a maximum of 120 seconds cycle length for typical intersections [10] was used as a basis to create the fixed timing plan for each intersection in Table I. Green timing per phase was allocated proportional to the phase demand based on the respective March 8, 2019 demand profile gathered from UDoT's ATSPM system [18] for each intersection during the 5:00 AM (inclusive)-8:00 PM (exclusive) 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.

B. Turning Policy Variation

A series of results are shown in Figure 4 for each intersection varying the turning policy assignments. Note that some data points represent a lower bound due to the simulator being

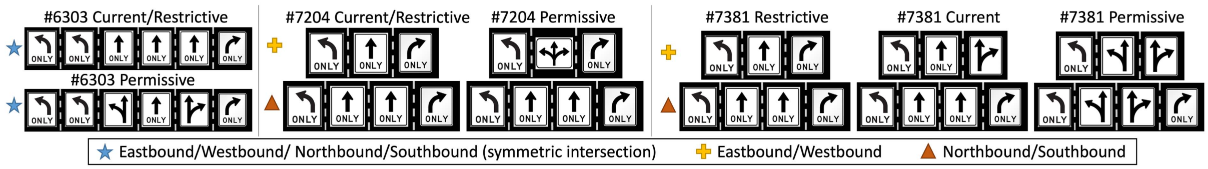

Fig. 3. Turning policies used for evaluating intersections #6303, #7204, and #7381. #6303 is entirely symmetrical in terms of turning actions. #7204 and #7381 both have symmetry in terms of allowed turning actions between the eastbound and westbound roads and also between the northbound and southbound roads.

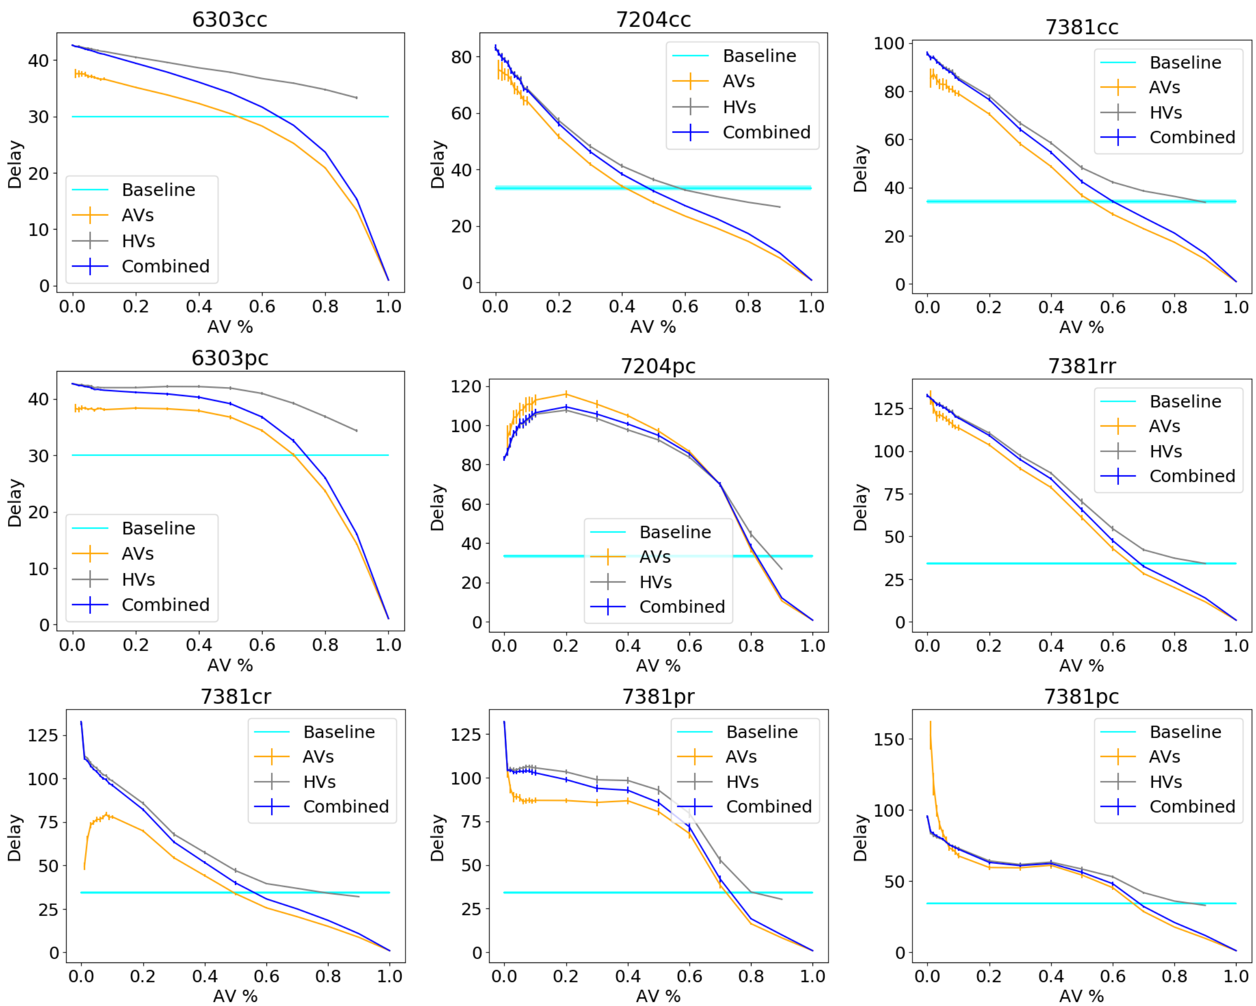

Fig. 4. Delay in seconds by CAV percentage using fixed timing for various turning policies. Each graph is labeled with the intersection number and then 2 letters representing the turning policy combination in the order of CAV policy and then HV policy. 'C' stands for the current policy in use by the real intersection, 'p' represents a permissive policy, and 'r' represents a restrictive policy. Each data point represents the average of the results of 20 trials. Error bars show 95% confidence intervals of those trials. A baseline with a 95% confidence interval is shown as a horizontal line.

potentially unable to spawn vehicles on full lanes (queue spill-back), effectively smoothing the arrival rate of vehicles. Our technical report [14] provides the exact delay values per setting and lists values that represent a lower bound due to queue spillbacks. The policy combinations are denoted as a CAV policy initial then a HV policy initial after the intersection number for each graph. In Figure 4, turning policy assignment pairs for HVs and CAVs are varied between the permutations of restrictive (fewest degrees of freedom) policies; the current policy used by the real intersection; and a permissive policy that has the 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 (as discussed in Section II-A.4).

The {Current, Current} policy combination's performance for all 3 intersections (seen in the first row of Figure 4) 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 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. This result implies that using the original H-AIM protocol (which requires fixed time actuation) is counterproductive for intersections currently applying adaptive and actuated control – at least until most traffic is composed of CAVs.

Noticeably, #6303 shows a different pattern from the other intersections with the {Current, Current} profile combination. Namely, the ‘Combined’ performance trend is concave vs convex and much gain is not observed for H-AIM early on. The {Current, Current} combination for #6303 also seems to be operating closer to its baseline from the start in terms of delay when compared to #7381 and #7204 with the {Current, Current} combination. These discrepancies are 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 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 more green trajectories are likely to be active at any given time. These factors could potentially stifle initial reductions in delay obtained by H-AIM with fixed timing. 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 known behaviors of CAVs which are detrimental in terms of delay that contribute to the odd trends seen in these intersections. Such behaviors, where CAVs block lanes, should be avoided. See Figures 5(a) and 5(b) for representative examples of such scenarios that were observed in simulation. In the example shown in Figure 5(a), HVs are assigned a stricter 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 Eastbound 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 until the blocking trajectory is no longer active. Until Vehicle 1 is able to proceed, Vehicle 3 would be unable to enter

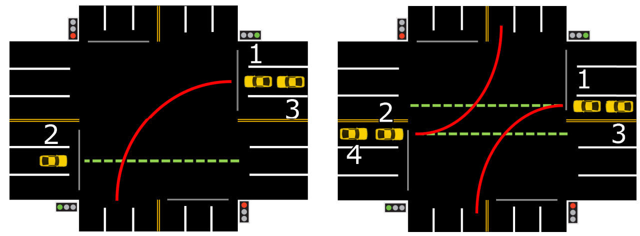

(a) (b)

Fig. 5. (a) An example of a CAV blocking a lane (Vehicle 1) because it is unable to obtain a reservation due to an active green trajectory. The green dashed line represents active green trajectory while the red line represents a path for a reservation which cannot be approved for the CAV. Vehicles 2 and 3 are HVs while Vehicle 1 is a CAV. (b) An example of multiple CAVs deadlocking (Vehicles 1 and 2) because they are unable to obtain reservations due to active green trajectories resulting from an IM's inability to distinguish between vehicle types at an intersection. The green dashed lines represent active green trajectories while the red lines represent paths for reservations which cannot be approved for CAVs. Vehicles 1 and 2 are CAVs while 3 and 4 are HVs.

the intersection. Note that depending on the particular turning profile, the associated green light may or may not apply to the turning action that Vehicle 1 is trying to take. This blockage was observed to be very harmful for an intersection such as #7204 with the {Permissive, Current} policy combination, as the only through lane for the eastbound and westbound roads can be easily blocked by a CAV. This 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 can be seen in Figure 5(b). 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 movements. However, both Vehicles 1 and 2 are trailed by HVs. Because this paper 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 removed 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 5(a), can simply be avoided by assigning identical turning profile combinations to all vehicle types with the provision 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 the one in Figure 5(b) 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 that continuously block the desired reservations for the CAVs. Note that the traffic signal must be green for both directions for this to happen. There are then 4 possible configurations which prevent such a deadlock:

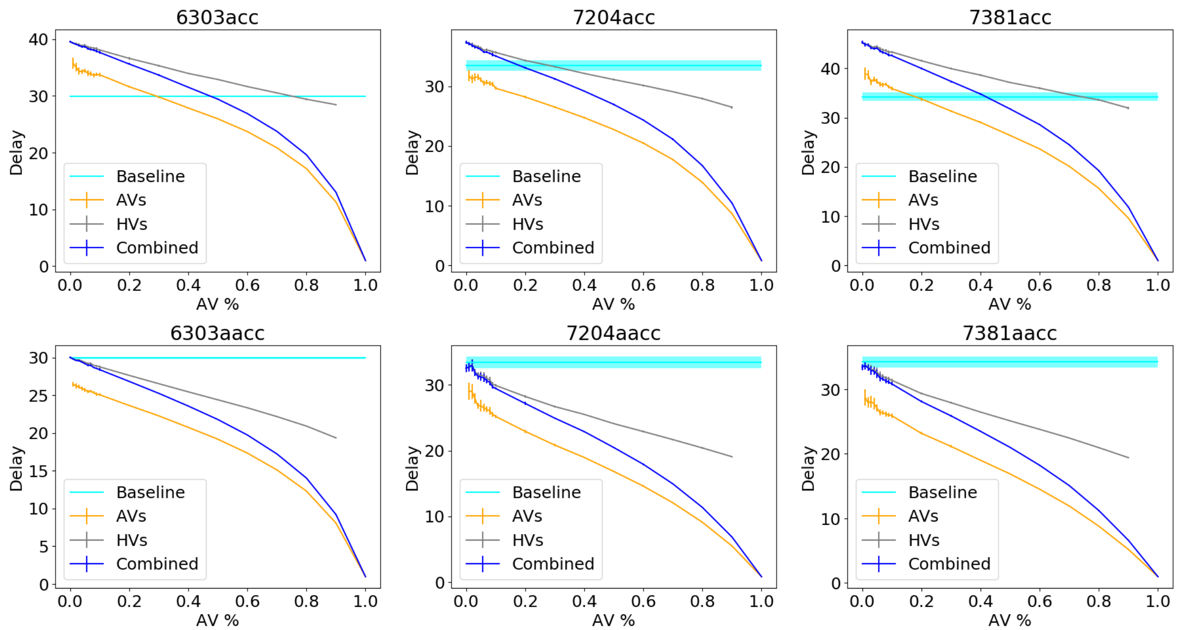

Fig. 6. Graphs of delay in seconds by CAV percentage for both adaptive timing and the combination of adaptive timing and actuation. Each graph is labeled with the intersection number and then 3–4 letters representing the combination of actuated timing and/or adaptive timing and then turning policy combination in the order of CAV policy and then HV policy. A single 'a' represents adaptive timing only, while 2 represents adaptive timing with actuation. '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 of 20 trials. Error bars show 95% confidence intervals of those trials. A baseline with a 95% confidence interval is shown as a horizontal line.

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, 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 IM cannot practically choose the lanes on which HVs arrive or when any vehicle arrives, attempting to create the situations in #1–3 in the list above is impractical. Thus, #4 should be employed as the target condition to prevent a deadlock. Note that in the example in Figure 5(b), 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 (i.e., CAVs and HVs have the same allowed turning actions) 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 simultaneously active green trajectory is considered safe or usable. However, if permissive left turns on green are allowed for a selected turning policy and vehicles may also proceed straight from the same lanes, 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.

C. Adaptive Signal 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 set 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 [9] was used as a basis for the automatic adjustments (gap times, maximum green times) made during the simulations.

For all intersections the actuation detection distance was set to 2 meters (approximately 6 feet). 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 based on their relative demand and table 5-6 in the manual. However as the through demand is more significant than the right turning demand for experimental settings, tallying demand for signals for right movements are merged with through movements when appropriate in this paper. Key values

(values used in conjunction to select a value in the table) are rounded to the nearest key in the appropriate tables. Results can be seen in Figure 6. The left column shows results for all 3 intersections with adaptive signal timing enabled and actuation disabled. This is an improvement in terms of delay for all intersections over the fixed timing scheme seen in the experiments in Section V-B. These results are not surprising since a reasonably configured adaptive signal timing process is better suited to deal with variations in demand throughout a day compared to a fixed timing scheme. The fact that the maximum green times for the phases of traffic signals at an intersection are tailored to recent demand at any given time results in a more efficient signal operation which reduces imposed delays. Nonetheless, all of the extended H-AIM variants in the left column of Figure 6 still initially underperform when compared to the baseline. This is, however, to be expected because the addition of properly configured actuation (which the baseline uses) is not applied for H-AIM here, as opposed to the results in the right column. The right column of Figure 6 shows that adding actuation into the mix results in further improvement of delay when used along with the extended H-AIM protocol. On top of efficiently accommodating varying 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 direction of travel. These experiments suggest that modern signal optimization techniques can efficiently be combined with the extended H-AIM protocol in terms of delay reduction. Moreover, unlike the original H-AIM protocol, these results suggest that applying the extended H-AIM protocol as part of modern intersections is not detrimental for any CAV penetration level and is mostly beneficial.

VI. Summary and Conclusion

Hybrid-AIM (H-AIM), first proposed in [17], is an efficient intersection management protocol capable of operating in early CAV penetration stages. H-AIM builds on top of the Autonomous Intersection Management (AIM) protocol proposed in [6] and can be used in conditions where an IM is able to sense approaching vehicles as well as fulfill the assumptions required by AIM. However, H-AIM, and other subsequent publications, are often demonstrated in simulations that are based on artificial conditions. This paper, by contrast, demonstrates use of readily available historical data combined with information about real intersections to test the efficiency of H-AIM in simulation. The presented results paint a picture where H-AIM underperforms commonly used adaptive and actuated signal controllers for initial/low CAV penetration rates, and thus suggests the original H-AIM protocol should not be used in similar situations. Moreover, this paper presents and discusses uncertainty and safety limitations that prevent a straightforward integration of H-AIM with actuated signal controllers. Next, the extended H-AIM protocol is presented. The extended H-AIM protocol is shown to be a safe and efficient way to reconcile H-AIM with common actuated signal controllers. Specifically, the extended H-AIM protocol enables the determination of when a future signal phase will be potentially active and disallows conflicting trajectories through the intersection. Additionally, a method by which IMs may safely consider adjustments to changing phase times throughout the day (such as is the case with adaptive control) is presented and discussed. Results from the experimental study support the following general conclusions:

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

2. In contrast to the findings in [17], traffic scenarios that are based on real-world data suggest that combinations of more permissive turning profiles for CAVs and more restrictive ones for HVs are not preferable for H-AIM at low CAV market penetration percentages. This may be due to delays brought on by CAVs blocking lanes while waiting for a reservation.

3. 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. This seems to be especially true at low CAV market penetration percentages.

4. The original H-AIM protocol (with fixed time actuation) is mostly counterproductive when replacing a modern, adaptive and actuated signal controller.

5. The extended H-AIM protocol which coexists with an adaptive and actuated signal controller is beneficial in most cases and was not observed to be counterproductive.

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 by the protocol, variations on actuation such as skippable phases, the potential of changing turning profile combinations based on CAV market penetration percentages or other applicable factors, implementation of early gap out functionality which terminates a green light when no HV is close to entering the intersection, studies into different IM reservation approval policies, and evaluation on a physical test bed with preferably real vehicles.

Acknowledgment

The authors would like to thank Dr. Peter Stone (Computer Science at UT-Austin) and Dr. Stephen Boyles (Department of Civil, Architectural and Environmental Engineering at UT-Austin) for sharing the initial code-base with them, providing insightful comments, and supporting this work.

References

[1] T.-Chiu Au, S. Zhang, and P. Stone, "Autonomous intersection management for semi-autonomous vehicles," *Routledge Handbook of Transportation*. Oxfordshire, U.K., 2015, p. 17.

[2] J. Ault, J. Hanna, and G. Sharon, "Learning an interpretable traffic signal control policy," in *Proc. 19th Int. Conf. Auto. Agents MultiAgent Syst.*, 2020, pp. 88–96.

[3] P. Bansal and K. Kockelman, "Forecasting Americans' long-term adoption of connected and autonomous vehicle technologies," in *Proc. 95th Annu. Meeting*, 2016, pp. 1816–1871.

[4] M. Bashiri, H. Jafarzadeh, and C. H. Fleming, "PAIM: Platoon-based autonomous intersection management," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 374–380.

[5] 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 *Proc. 16th Int. IEEE Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2013, pp. 726–731.

[6] K. Dresner and P. Stone, "A multiagent approach to autonomous intersection management," *J. Artif. Intell. Res.*, vol. 31, pp. 591–656, Mar. 2008.

[7] D. Fajardo, T.-C. Au, S. Waller, P. Stone, and D. Yang, "Automated intersection control: Performance of future innovation versus current traffic signal control," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2259, no. 1, pp. 223–232, 2011.

[8] P. Faldu, N. Doshi, and R. Patel, "Real time adaptive traffic control system: A hybrid approach," in *Proc. IEEE 4th Int. Conf. Comput. Commun. Syst. (ICCCS)*, Feb. 2019, pp. 697–701.

[9] FHWA. (2008). *Traffic Signal Timing Manual*. Accessed: Jul. 23, 2020. [Online]. Available: https://ops.fhwa.dot.gov/publications/fhwahop08024/chapter5.htm

[10] FHWA. (Jul. 2013). *Signalized Intersections: An Informational Guide*. Accessed: Jul. 23, 2020. [Online]. Available: https://safety.fhwa.dot.gov/intersection/conventional/signalized/fhwasa13027/ch5.cfm

[11] Google. (2020). *Google Maps*. Accessed: 2020. [Online]. Available: http://maps.google.com/

[12] H. Hu, X. Jia, K. Liu, and B. Sun, "Self-adaptive traffic control model with behavior trees and reinforcement learning for AGV in Industry 4.0," *IEEE Trans. Ind. Informat.*, vol. 17, no. 12, pp. 7968–7979, Dec. 2021.

[13] M. N. Mladenovic and M. M. Abbas, "Self-organizing control framework for driverless vehicles," in *Proc. 16th Int. IEEE Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2013, pp. 2076–2081.

[14] A. Parks-Young and G. Sharon, "Technical report: Hybrid autonomous intersection management," 2022, *arXiv:2204.07704*.

[15] L. Riegger, M. Carlander, N. Lidander, N. Murgovski, and J. Sjoberg, "Centralized MPC for autonomous intersection crossing," in *Proc. IEEE 19th Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2016, pp. 1372–1377.

[16] M. O. Sayin, C.-W. Lin, S. Shiraishi, J. Shen, and T. Başar, "Information-driven autonomous intersection control via incentive compatible mechanisms," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 3, pp. 912–924, Mar. 2019.

[17] G. Sharon and P. Stone, "A protocol for mixed autonomous and human-operated vehicles at intersections," in *Proc. Int. Conf. Auto. Agents Multiagent Syst.* Cham, Switzerland: Springer, 2017, pp. 151–167.

[18] UDOT. (2020). *Udot Automated Traffic Signal Performance Measures—Automated Traffic Signal Performance Metrics*. Accessed: Jul. 23, 2020. [Online]. Available: https://udottraffic.utah.gov/ATSPM

[19] UDOT. (2020). *Udot Data Portal*. Accessed: Jul. 23, 2020. [Online]. Available: http://data-uplan.opendata.arcgis.com/

[20] Y. Wu, H. Chen, and F. Zhu, "DCL-AIM: Decentralized coordination learning of autonomous intersection management for connected and automated vehicles," *Transp. Res. C, Emerg. Technol.*, vol. 103, pp. 246–260, Jun. 2019.

Aaron Parks-Young received the B.S. degree in computer science from The University of Texas at Dallas in 2017 and the M.S. degree in computer engineering from Texas A&M University in 2020. His research interests include intelligent systems, reinforcement and machine learning, as well as multiagent systems. His educational background includes HDL, networking, web development, and facets of artificial intelligence.

Guni Sharon received the Ph.D. degree from the Ben-Gurion University of the Negev in 2015. Following his Ph.D. degree, he held a Post-Doctoral Researcher position for three years with the Computer Science Department, The University of Texas at Austin. Since 2018, he has been serving as an Assistant Professor with the Computer Science and Engineering Department, Texas A&M University. His research interests include artificial intelligence, intelligent transportation systems, reinforcement learning, and combinatorial optimization.