

Replacing the Stop Sign: Unmanaged Intersection Control for Autonomous Vehicles

(Short Paper)

Mark VanMiddlesworth
Harvard University
Elec. Eng. and Comp. Sci. Department
mvandmidd@fas.harvard.edu

Kurt Dresner Peter Stone
University of Texas at Austin
Department of Computer Sciences
{kdresner, pstone}@cs.utexas.edu

ABSTRACT

As computers replace humans as the drivers of automobiles, our current traffic management mechanisms will give way to hyper-efficient protocols designed to exploit the capabilities of fully autonomous vehicles. We have introduced such a system for coordinating large numbers of autonomous vehicles at intersections [2, 3]. Our experiments suggest that this system could alleviate many of the dangers and delays associated with intersections by allowing vehicles to “call ahead” to an agent stationed at the intersection and reserve time and space for their traversal. Unfortunately, such a system is not cost-effective at small intersections. In this paper, we propose an intersection control mechanism for autonomous vehicles designed specifically for low-traffic intersections where the previous system would not be practical. Our mechanism is based on purely peer-to-peer communication and thus requires no infrastructure at the intersection. We present experimental results demonstrating that our system, while not suited to large, busy intersections, can significantly outperform traditional stop signs at small intersections.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Miscellaneous

Keywords

multiagent systems, intelligent transportation systems, intersection control, autonomous vehicles

1. INTRODUCTION

Recent advances in technology suggest that, in the near future, many vehicles will be controlled without direct human involvement [1]. More efficient mechanisms will take advantage of the precision control of autonomous vehicles as well as recent research in the field of Multiagent Systems (MAS). Previously, we created a MAS-based traffic management system that has the potential to vastly outperform current traffic signals [2, 3]. However, the high infrastructure costs associated with this system make it uneconomical at

Cite as: Replacing the Stop Sign: Unmanaged Intersection Control for Autonomous Vehicles (Short Paper), M. VanMiddlesworth, K. Dresner and P. Stone, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX. Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

low-traffic intersections. We thus propose a control mechanism based on peer-to-peer (vehicle-to-vehicle or “V2V”) interaction, that requires no specialized infrastructure.

1.1 Managed Intersection Control

Our previously proposed system [3] involves two classes of agents: *intersection managers* and *driver agents*. In that system, driver agents “call ahead” to an intersection manager at the intersection to reserve the space-time needed to cross the intersection safely. This system offers substantial safety and efficiency benefits over existing mechanisms, such as traffic lights and stop signs. Vehicles traverse the intersection faster, and congestion is reduced.

At the city level, the system is decentralized, but at each intersection, traffic is coordinated by an *arbiter* agent, the intersection manager. We therefore designate this system a *managed* intersection control mechanism. An intersection controlled by a traffic light is also managed—the traffic light is the arbiter agent. Conversely, an *unmanaged* intersection control mechanism, such as a stop sign, has no arbiter agent.

1.2 One Size Does Not Fit All

Managed intersection control mechanisms have a major drawback: cost. While the throughput benefits at large intersections warrant the cost of an arbiter agent, such systems are uneconomical for small intersections. Stop signs are low-overhead, unmanaged systems for low-traffic intersections that complement traffic lights at larger intersections. In this paper, we propose an unmanaged intersection control mechanism for autonomous vehicles, specifically for low-traffic intersections and requiring no specialized infrastructure. This system complements our previously proposed managed system as stop signs complement more expensive traffic light installations. By using the same assumptions about vehicle capabilities, a driver agent can use both systems seamlessly.

2. AN UNMANAGED MECHANISM

In this section, we introduce our unmanaged autonomous intersection control mechanism. First, we specify our system’s goals. Next, we describe our assumptions about driver agents. We then outline the protocol for vehicle communication and describe the vehicles’ required actions.

2.1 Goals Of The System

To be effective and economical, an unmanaged intersection control mechanism for autonomous vehicles should have the following properties:

- Vehicles should traverse intersections more quickly than with current mechanisms (i.e. stop signs).
- The protocol should have minimal (ideally none) per-intersection infrastructure costs.
- If all vehicles follow the protocol correctly, no collisions should result.

2.2 Assumptions

We assume driver agents have access to any information needed to navigate an intersection safely: the layout and location of the intersection, any speed limits, and any other relevant details. We assume each vehicle is outfitted with low-latency (<20ms), medium-range (150m) wireless communication, although the parameters of the protocol can be adjusted to suit the environment. We also assume vehicles have the technology required for autonomous open-road driving, including lidar or short-wave radar capable of reliably sensing objects in their immediate vicinity. Finally, we assume driver agents have access to information about the vehicle they are controlling, including velocity, position, and heading. This information can easily be provided by speedometers, GPS devices, and digital compasses.

Via either a street map database or a standard analysis of the intersection layout, we assume agents can determine which paths through it are *compatible*, meaning they can be followed simultaneously without the risk of collision. For example, right turns from the rightmost lanes in any direction are always compatible. Any paths that intersect are not. Because driver agents will use this information to plan their trajectories through the intersection, possibly allowing two vehicles to cross simultaneously, it is important that each agent have the same notion of which paths are compatible.

2.3 Communication Protocol

Driver agents must maintain up-to-date information about the vehicles approaching the intersection. Because connectivity is unpredictable in an ad-hoc wireless network of mobile agents, our communication protocol cannot rely on a dialogue between agents. It consists of two messages, CLAIM and CANCEL.

2.3.1 Claim

Agents transmit a CLAIM message in order to announce the intention to use a specific space and time in the intersection. The message contains the vehicle’s intended arrival lane (`lane`), arrival time (`arrival_time`), exit time (`exit_time`), and any planned turns (`turn`). It also contains the Vehicle Identification Number, or VIN (`vin`), and a message ID (`message_id`), a monotonically increasing counter specific to the message. The message ID is only changed when a vehicle generates a *new* message to broadcast, not when a message is rebroadcast. This allows agents to recognize stale information. Finally, the CLAIM contains a boolean value indicating whether the sending vehicle is stopped at the intersection (`stopped`).

2.3.2 CANCEL

Agents send a CANCEL message to invalidate any previous CLAIM messages. This message has two fields: the VIN, and a message ID.

2.3.3 Message Broadcast

Because each message contains the latest relevant infor-

mation about the sending vehicle, agents need only consider the most recent message from any other vehicle. Each message is broadcast repeatedly to ensure its eventual delivery, should a new vehicle enter transmission range. Although occasional dropped messages may increase the delay in communications between vehicles, they should not compromise the safety of the system. For security purposes, we assume each message is digitally signed, ensuring that driver agents cannot falsify the `vin` field of their messages. Malformed or unsigned messages are ignored.

2.3.4 Conflict, Priority, and Dominance

To facilitate the discussion of agent behavior and protocol analysis, we define three relations on CLAIM messages.

Two CLAIM messages are said to *conflict* if the paths determined by the `lane` and `turn` parameters of the CLAIM messages are not compatible and the time intervals specified in the CLAIM messages are not disjoint.

We define relative *priority* of two CLAIM messages using the following rules, from most to least significant:

1. If both `stopped` fields are `false`, the CLAIM with the earliest `exit_time` has priority.
2. If both `stopped` fields are `true`, the CLAIM whose `lane` is “on the right” has priority. “On the right” is a globally available binary relation, defined similarly to current traffic laws regarding four-way stop signs.
3. If neither message’s `lane` is “on the right,” the CLAIM whose `turn` parameter indicates no turn has priority.
4. If priority cannot be established by the previous rules (including a cycle of “on the right” relations), the CLAIM with the lowest VIN has priority.

For claims c_1 and c_2 , we say that c_1 *dominates* c_2 if the `stopped` field of c_1 is `true` and the `stopped` field of c_2 is `false`, or if the `stopped` fields of c_1 and c_2 are identical, c_1 and c_2 conflict, and c_1 has priority over c_2 . Under the vast majority of circumstances, dominance is a total order and thus driver agents need only reason about whether their own claim is dominated¹.

2.4 Required Agent Actions

A rigid set of rules governing the interaction of autonomous vehicles is required to prevent potentially disastrous failures. Our multiagent system relies on a set of rules analogous to traffic laws. While nothing can prevent an agent from disobeying (just as a human driver can drive through a red light), the safety of each vehicle is guaranteed only if the driver agent follows the rules. Note that the following rules restrict only how the agent behaves in the intersection; driver agents have full autonomy everywhere else:

1. A vehicle may not enter the intersection if its CLAIM is dominated by any other current CLAIM.
2. A vehicle may not enter the intersection without first broadcasting a CLAIM for at least T_p seconds. In our implementation, $T_p = 0.4$.
3. A vehicle must vacate the intersection by the `exit_time` in its most recent CLAIM message.
4. Within the intersection, vehicles must remain in their lanes from entry to departure.
5. The `stopped` field of a CLAIM must be `true` only if the vehicle is stopped at the intersection.

¹In an exceedingly remote case, an “on the right” cycle could lead to a cycle of dominance for 4 stopped cars. In this case, driver agents break the cycle by lowest VIN.

2.5 Selfish and Malicious Agents

Driver agents are assumed to be self-interested—they may take any legal action in order to ensure they traverse the intersection in as little time possible. Agents have little incentive to lie about their lane, path, or exit time, because this may put the vehicle at risk of collision. However, an agent may have an incentive to falsely claim it is stopped at the intersection, allowing its CLAIM to dominate the CLAIMS of other moving vehicles. This type of behavior is not currently disincentivized by our protocol, but were it to become a problem, could be tested at random intersections to ensure compliance. This is analogous to current traffic enforcement, which relies on sporadic monitoring and associated penalties to decrease rule violations.

In any multiagent system, malicious agents are a concern. In current traffic scenarios, nothing prevents someone from deliberately crashing into another vehicle, or disabling traffic signals. Similarly, a malicious driver agent could flood the network with useless messages, preventing the system from operating properly. While nothing can be done to stop a determined saboteur, the fact that all messages are signed makes it impossible for vehicles to hide their identity.

3. DRIVER AGENT BEHAVIOR

Our proposed unmanaged intersection control mechanism relies not only on the communication protocol defined above, but also on the existence of driver agents that can abide by the protocol. Our prototype driver agent’s behavior is comprised of three phases: lurking, making a reservation, and intersection traversal.

3.1 Lurking

As an agent approaches the intersection, it begins to receive messages from other agents. However, it does not transmit until its distance from the intersection is less than the *lurk distance*. Lurking prevents agents from prematurely broadcasting a CLAIM that will be dominated by an existing CLAIM held by another agent. Lurk distance depends on both transmission range and the period with which agents broadcast. In our simulations, we set lurk distance at 75 meters—a reasonable value given current technology.

3.2 Making a Reservation

When a vehicle reaches the lurk distance its driver agent needs to tell the other driver agents how it intends to cross the intersection. We call this process “making a reservation,” an analogy to our reservation-based system [3].

As an agent approaches the intersection, it generates a CLAIM based on the earliest possible arrival time, predicted velocity of the vehicle at this time, and predicted exit time. If the agent has received no CLAIMS from other vehicles that dominate this CLAIM, the agent will begin to broadcast this CLAIM. Otherwise, the agent generates a new CLAIM for the earliest sufficiently large block of time such that it will not be dominated by any existing CLAIM.

3.3 Intersection Traversal

When a vehicle has made a reservation, it needs only to broadcast the CLAIM continually and to arrive at the intersection in accordance with its reservation. However, sometimes the vehicle may want to change an existing claim in order to take advantage of an unexpected early arrival. On the other hand, congestion may cause a vehicle to arrive late.

If a vehicle predicts that it cannot fulfill the parameters of its CLAIM message, or if a new CLAIM message arrives that dominates the driver agent’s CLAIM, it must either send a CANCEL or a new CLAIM. Once the vehicle reaches the intersection, it crosses in accordance with its CLAIM. While in the intersection, for safety purposes, the vehicle continues to broadcast its CLAIM, however this CLAIM cannot be dominated. After a vehicle has vacated the intersection, it stops transmitting its CLAIM.

3.3.1 Canceling “Bad” Reservations

To ensure that it can meet the parameters of its reservation, the driver agent constantly re-estimates its arrival time. If the predicted arrival time is later than that of its reservation, the driver agent cancels its current reservation and attempts to make a reservation for a later time.

3.3.2 Improving Reservations

If an agent determines it will arrive early and can broadcast and comply with a new CLAIM that will not be dominated, it immediately begins broadcasting this CLAIM. As specified by the protocol, this invalidates any previous CLAIM. If a vehicle does arrive at the intersection early, it can change its CLAIM to reflect that it is stopped. Because the new CLAIM will dominate that of any vehicle not stopped at the intersection, the agent may be able to set an earlier `arrival_time` than it would were it not stopped at the intersection.

4. EMPIRICAL RESULTS

This section presents empirical results comparing our unmanaged autonomous intersection to four-way stop signs and traffic lights. We first describe our metric, *delay*, and our experimental setup. We then use this metric to compare the various intersection control mechanisms.

4.1 Delay

In our analysis, the metric we use is *delay*: the additional time it takes a vehicle to reach its destination due to the presence of the intersection. The baseline for delay is the time it would take a vehicle to traverse a completely empty intersection. Delay is measured as actual trip time minus baseline trip time, which isolates the effect of the intersection control policies and allows us to accurately compare them.

4.2 Experimental Setup

To test these policies, we use the custom simulator described in our earlier work [3], which simulates a four-way intersection with one lane of traffic in each direction. This configuration is representative of a four-way stop, and provides the best test case for unmanaged control mechanisms. We control traffic levels via a Poisson process governed by the probability of creating a new vehicle at each time step. We simulate traffic levels between 0 and 0.5 vehicles per second, with 15% of vehicles turning left and 15% turning right. Each data point represents the average of 20 simulations, with each run consisting of 30 minutes of simulated time. All error bars indicate a 95% confidence interval.

The traffic light timing is configured such that, in succession, each direction receives a green light for 10 seconds, followed by 3 seconds of yellow. There is a large body of theory and empirical evidence concerning the timing of

traffic lights, but this work is largely irrelevant to our simulated scenario for two reasons. First, much of the theory deals with the timing of lights across multiple intersections, whereas we are examining one intersection in isolation. Second, our simulator generates symmetric traffic, which greatly simplifies timing by eliminating the need to account for higher traffic levels in a particular direction. For these reasons, we established a reasonable timing pattern experimentally by evaluating 10 different candidate patterns and selecting the one that led to the lowest average delay.

4.3 Delay

As shown in Figure 1, our system significantly reduces the average delay experienced by each vehicle. When traffic flow is below 0.35 vehicles per second, the four-way stop is a more effective policy than the traffic light. The unmanaged V2V system outperforms both competing systems over the entire experimental domain, and performs especially well below 0.5 vehicles per second. Note that these results pertain only to low-traffic, isolated intersections. Traffic lights have significant advantages in other situations.




Figure 1: A comparison of average delay of the traffic light, four-way stop, and our unmanaged V2V mechanism. Our system can handle more traffic than the four-way stop and causes negligible delays for low traffic.

Using our unmanaged system, most agents are able to cross the intersection without slowing down to wait for other vehicles when traffic levels are below 0.3 vehicles per second. With the four-way stop sign, each vehicle must stop even if no others are present, resulting in a baseline average delay of approximately 3 seconds. The traffic light system has a higher baseline average delay, around 18 seconds.

When traffic flow is between 0.3 and 0.5 vehicles per second, our system shows a somewhat increased delay. In these cases, cars may slow down to accommodate other vehicles, but only rarely will a vehicle need to stop. With the stop sign, vehicles begin to queue at the intersection, and must often wait for vehicles in front of them to cross.

The stop sign cannot handle traffic levels above 0.35 vehicles per second. At these traffic levels, our system is similar to a four-way stop: because there is almost always at least one vehicle waiting to cross, agents must wait until they are stopped at the intersection to make a reservation. However, the intersection sharing in our system (allowing four simultaneous right turns, for example) provides a noticeable benefit at these traffic levels. Our unmanaged system can handle traffic levels up to approximately 0.7 vehicles per

second (twice that of the stop sign), at which point traffic begins to back up. In these situations, our data suggest that a managed mechanism may be more appropriate.

5. CONCLUSION AND FUTURE WORK

Recent research has already produced fully autonomous, computer-controlled vehicles. As these vehicles become more common, we will be able to phase out human-centric traffic control mechanisms in favor of vastly more efficient computer-controlled systems. This will be especially beneficial at intersections, which are a major cause of delays. For a transition of this magnitude, infrastructure cost will be a central, if not primary, concern. This paper presents a novel, unmanaged intersection control mechanism requiring no specialized infrastructure at the intersection. We have described a protocol for our unmanaged autonomous intersection, and created a prototype driver agent capable of utilizing this protocol. As illustrated by our empirical results, our protocol can significantly reduce delay as compared to a four-way stop. Small intersections far outnumber those that are sufficiently large or busy to warrant the cost of a managed solution. Whereas busier intersections may need to wait for the funding and installation of requisite infrastructure, our proposed mechanism has the potential to open every one of these intersections to be used safely and efficiently by autonomous vehicles.

In the future, we would like to examine other metrics, including a measure of fuel efficiency, test our system’s resilience to communication failures, and extend our protocol to work with both autonomous and human drivers, as we have with our managed system [4]. We would also like to explore the effects of asymmetric traffic flow, including intersections at which a two-way stop is more efficient than a four-way stop.

Acknowledgments

This research is supported in part by NSF CAREER award IIS-0237699 and by the United States Federal Highway Administration under cooperative agreement DTFH61-07-H-00030. Computational resources were provided in large part by NSF grant EIA-0303609.

6. REFERENCES

- [1] DARPA. The DARPA urban challenge, 2007. <http://www.darpa.mil/grandchallenge>.
- [2] K. Dresner and P. Stone. Multiagent traffic management: A reservation-based intersection control mechanism. In *The Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 530–537, New York, NY, USA, July 2004.
- [3] K. Dresner and P. Stone. Multiagent traffic management: An improved intersection control mechanism. In *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 471–477, Utrecht, The Netherlands, July 2005.
- [4] K. Dresner and P. Stone. Sharing the road: Autonomous vehicles meet human drivers. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pages 1263–68, Hyderabad, India, January 2007.