In *The Fourth Workshop on Agents in Traffic and Transportation* (ATT 06), Hakodate, Japan, May 2006.

Human-Usable and Emergency Vehicle–Aware Control Policies for Autonomous Intersection Management

Kurt Dresner and Peter Stone
University of Texas at Austin
Department of Computer Sciences
Austin, TX 78712 USA
{kdresner,pstone}@cs.utexas.edu

ABSTRACT

Traffic congestion and automobile accidents are two of the leading causes of decreased standard of living and lost productivity in urban settings. Recent advances in artificial intelligence and, specifically, intelligent vehicle technology suggest that vehicles driven entirely by autonomous agents will be possible in the near future. In previous work, we presented a novel reservation-based approach for governing interactions of multiple autonomous vehicles, specifically at intersections. This approach alleviated many traditional problems associated with intersections, in terms of both safety and efficiency. However, such a system relies on all vehicles being equipped with the requisite technology — a restriction that would make implementing such a system in the real world extremely difficult. In this paper, we augment the system such that it is able to accomodate traditional human-operated vehicles using existing infrastructure. Furthermore, we show that as the number of autonomous vehicles on the road increases, traffic delays decrease monotonically toward the levels exhibited in the system involving only autonomous vehicles. Additionally, we demonstrate how the system can be extended to allow high-priority vehicles such as ambulances, police cars, or fire trucks through more quickly without placing undue burden on other vehicles. Both augmentations are fully implemented and tested in our custom simulator, and we present detailed experimental results attesting to their effectiveness.

1. INTRODUCTION

Traffic congestion and automobile accidents are two of the leading causes of decreased standard of living and lost productivity in urban settings. According to a recent study of 85 U.S. cities [20], annual time spent waiting in traffic has increased from 16 hours per capita to 46 hours per capita since 1982. In the same period, the annual financial cost of traffic congestion has swollen from $14 billion to more than $63 billion (in 2002 US dollars). Each year, Americans burn approximately 5.6 billion gallons of fuel while idling in heavy traffic. Furthermore, while vehicle safety has historically made gradual improvements each year, collisions cost the United States over $230 billion annually [10]. Globally, automobile accidents account for 2.1% of all deaths, which makes them the 11th overall cause of death [2]. Recent advances in artificial intelligence suggest that autonomous vehicle navigation will be possible in the near future. Individual cars can now be equipped with features of autonomy such as adaptive cruise control, GPS-based route planning [16, 18], and autonomous steering [12, 14]. In fact, in early 2006, DaimlerChrysler began selling the Mercedes-Benz S-Class, which comes with with radar-assisted braking that automatically applies the correct amount of braking force, even if the driver does not. Once individual cars become autonomous, many of the cars on the road will have such capabilities, thus opening up the possibility of autonomous interactions among multiple vehicles.

Multiagent Systems (MAS) is the subfield of AI that aims to provide both principles for construction of complex systems involving multiple agents and mechanisms for coordination of independent agents' behaviors [19]. In earlier work, we proposed a novel MAS-based approach to alleviating traffic congestion and collisions, specifically at intersections [4].

In this paper, we make three main contributions. First, we show how to augment our existing intersection control mechanism to allow use by human drivers with minimal additional infrastructure. Second, we show that this hybrid intersection control mechanism offers performance and safety benefits over traditional traffic light systems. Thus, implementing our system over an extended time frame will not adversely affect overall traffic conditions at any stage. Furthermore, we show that at each stage there exists an incentive for individuals to use autonomous driver agent-equipped vehicles. Historically, many technologies and transit systems aimed at improving safety and decreasing congestion have suffered from a lack of incentive for early adopters. For example, if everyone used mass transit, traffic would be reduced to an extent that the bus or light rail would be cheaper, faster, and safer than driving a personal vehicle is currently. However, given the current state of affairs, it is not in any one person's interest to make the switch. Our third contribution is a separate augmentation that allows the system to give preference to emergency vehicles such as ambulances, police cruisers, and fire trucks. We demonstrate that this is not overly detrimental to the rest of the vehicles. Both augmentations are fully (though separately) implemented and tested in our custom simulator and complete experimental results are presented.

The rest of this paper is organized as follows. In Section 2, we briefly review the reservation system as described in previous work. In Section 3 we explain how our original reservation-based intersection control mechanism can be augmented to allow for human drivers (or cyclists or pedestrians). In Section 4, we describe additions to the system and communication protocol that give further benefits to emergency vehicles without causing excessive delays to civilian traffic. We present the experimental results of these fully-implemented augmentations in Section 5. In Section 6, we discuss the experimental results in the context of related work. Section 7 describes where we plan to take this line of research in the near future, and we conclude in Section 8.

2. RESERVATION SYSTEM

Previously, we proposed a novel reservation-based multi-agent approach to alleviating traffic, specifically at intersections [4]. This system consists of two types of agents: *intersection managers* and *driver agents*. For each intersection, there is a corresponding intersection manager, and for each vehicle, a driver agent. Intersection managers are responsible for directing the vehicles through the intersection, while the driver agents are responsible for controlling the vehicles 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 space-time in the intersection. The intersection manager then determines whether or not these requests can be met based on an *intersection control policy*. Depending on the decision (and subsequent response) the intersection manager makes, the driver agent either records the parameters of the response message (the *reservation*) and attempts to meet them, or it receives a rejection message and makes another request at a later time. If a vehicle has a reservation, it can request that its reservation be changed or can cancel the reservation. It also sends a special message when it finishes crossing the intersection indicating to the intersection manager that it has done so.

The interaction among these agents is governed by a shared protocol which we have published in a technical report [**?**]. In addition to message types (e.g. REQUEST, CONFIRM, and CANCEL), this protocol includes some rules, the most important of which are (1) that a vehicle may not enter the intersection unless it is within the parameters of a reservation made by that vehicle's driver agent, (2) that if a vehicle follows its reservation parameters, the intersection manager can guarantee a safe crossing for the vehicle, and (3) a driver agent may have only one reservation at a time. While some may argue that insisting a vehicle adhere to the parameters of such a reservation is too strict a requirement, it is useful to note that vehicles today are already governed by a similar (although much less precise) protocol; if a driver goes through a red light at a busy intersection, a collision may be unavoidable. Aside from this protocol, no agent needs to know how the other agents work — each vehicle manufacturer (or third party) can program a separate driver agent, each city or state can create their own intersection control policies (which can even change on the fly), and as long as each agent adheres to the protocol, the vehicles will move safely through the intersection. A diagram of one type of interaction in the mechanism can be seen in Figure 1.

2.1 First Come, First Served (FCFS)

To determine whether or not a request can be met, our intersection manager uses a "first come, first served" (FCFS) intersection control policy which works as follows:

- The intersection is divided into a grid of $n \times n$ *tiles*, where n is called the *granularity*.

- Upon receiving a request message, the policy uses the parameters in the message to simulate the journey of the vehicle across the intersection. At each time step of the simulation, it determines which tiles the vehicle occupies.

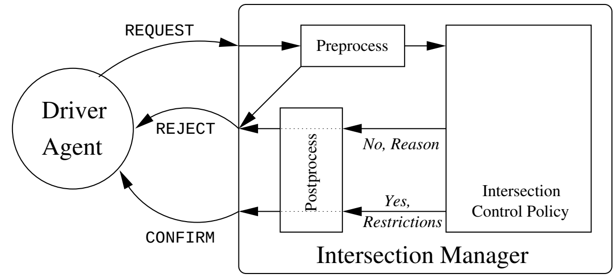

Figure 1: The interaction between the Intersection Manager, Intersection Control Policy, and Driver Agent when a REQUEST message is sent.

- If throughout this simulation, no required tile is reserved by another vehicle, the policy reserves the tiles for the vehicle and confirms the reservation. Otherwise, the request is rejected.

The policy derives its name from the fact that the policy responds to vehicles immediately when they make a request, confirming or rejecting the request based on whether or not the space-time required by the vehicle is already claimed. If two vehicles require some tile at the same time, the vehicle which requests the reservation first will be given the reservation (provided there are no conflicts in the rest of the required space-time). Figure 2 shows a successful reservation (confirmed) followed by an unsuccessful reservation (rejected).

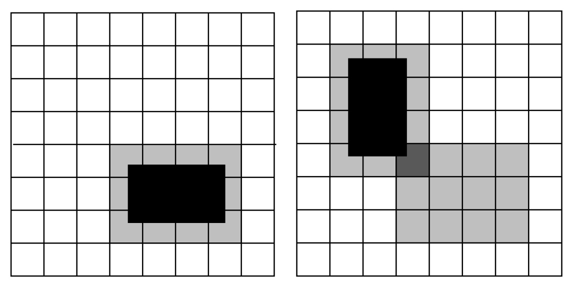

(a) Successful reservation

(b) Failed reservation

Figure 2: The grid for a granularity-8 FCFS policy. In 2(a), the policy is simulating the trajectory of vehicle A and finds that at some time t, all the tiles it requires are available. A's request is confirmed. In 2(b), vehicle B makes a subsequent reservation request. During the simulation of B's trajectory, at time t, the policy finds that a tile required by B is already reserved by A. B's reservation request is thus rejected.

2.2 Other Intersection Control Policies

While the reservation system was designed with the FCFS policy in mind, it can accomodate any intersection control policy that can make a "yes or no" decision based on the parameters in a request message. This includes policies that represent familiar intersection control mechanisms like traffic lights and stop signs. Because the reservation system can behave exactly like our most common modern-day control mechanisms, we can absolutely guarantee that the performance of the reservation mechanism will be no worse than current systems. The descriptions given below are abbreviated; full descriptions (including the STOP-SIGN policy) may be found in our tech report [?].

2.2.1 TRAFFIC-LIGHT

Traffic lights are the most common mechanism used to control high-traffic intersections. The TRAFFIC-LIGHT policy emulates a real-life traffic light by maintaining a model of how the lights would be changed, were they to exist. Then, upon receiving a request message, the policy determines whether the light corresponding to the requesting vehicle's lane would be green. If so, it sends a confirmation, otherwise, it sends a rejection.

2.2.2 OVERPASS

Although called OVERPASS, this policy does not represent a real overpass (or cloverleaf), which are very expensive and built only at the largest and busiest of intersections. Instead, it represents an optimal intersection control policy — one which never rejects a vehicle. This would not be useful in real life as it makes no guarantees regarding the safety of the vehicles, but it does serve as a good lower bound for delays.

2.3 Measuring Performance

After creating a custom simulator (Figure 3 shows a screenshot of the graphical display), we evaluated the performance of the FCFS policy against the OVERPASS and the TRAFFIC-LIGHT policies. Using the simulator, we showed that with the FCFS policy, vehicles crossing an intersection experience much lower *delay* (increase in travel time from the optimal) versus TRAFFIC-LIGHT [3, 4]. The FCFS policy approached OVERPASS in terms of delay, offering safety guarantees that OVERPASS could not. Furthermore, we showed that the FCFS policy increases the throughput of the intersection far beyond that of TRAFFIC-LIGHT. For any realistic (safe) intersection control policy, there exists an amount of traffic for which vehicles arrive at the intersection more frequently than they can leave the intersection. At this point, the average delay experienced by vehicles travelling through the intersection grows without bound — each subsequent vehicle will have to wait longer than all the previous cars. The point for which this occurs in the FCFS policy is five or six times higher than TRAFFIC-LIGHT.

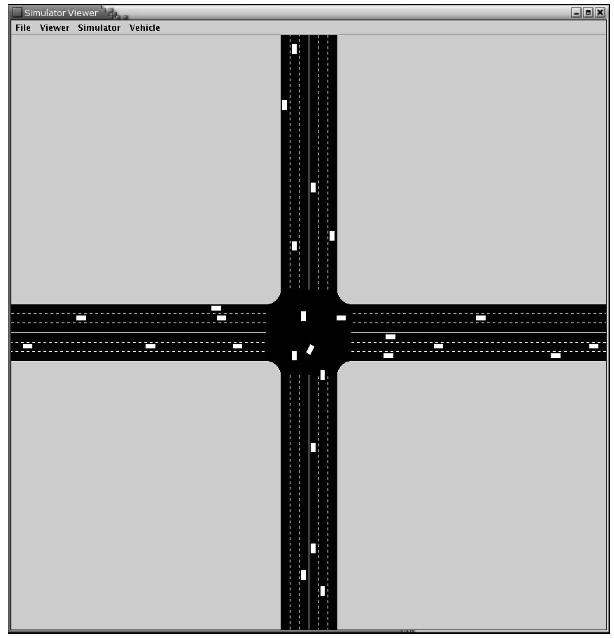

Figure 3: A screenshot of the graphical display of our simulator.

3. INCORPORATING HUMAN USERS

While an intersection control mechanism for autonomous vehicles will someday be very useful, there will always be people who enjoy driving. Additionally, there will be a fairly long transitional period between the current situation (all human drivers) and one in which human drivers are a rarity. Even if switching to a system comprised solely of autonomous vehicles were possible, pedestrians and cyclists must also be able to traverse intersections in a controlled and safe manner. For this reason, it is necessary to create intersection control policies that are aware of and able to accomodate humans, whether they are on a bicycle, walking to the corner store, or driving a "classic" car for entertainment purposes. In this section we explain how we have extended our FCFS policy as well as the reservation framework to incorporate human drivers. Adding pedestrians and cyclists follows naturally and though while we have not actually implemented them in our system, we give brief descriptions of how this would differ from the extensions for human drivers.

3.1 Using Existing Infrastructure

Adding human drivers to the mix means that we need a reliable way to communicate information to the drivers. The best way to do this is to use a system that drivers already know and understand — traffic lights. Traffic light infrastructure is already present at many intersections and the engineering and manufacturing of traffic light systems is well developed. For pedestrians and cyclists, standard "push-button" crossing signals could be used that would give enough time for a person to traverse the intersection. These could also serve to alert the intersection to their presence.

3.2 Light Models

If real traffic lights are going to be used to communicate to human drivers, they will need to be controlled and understood by the intersection manager. Thus, we add a new component to each intersection control policy, called a *light model*. This model controls the actual physical lights as well as providing information to the policy with which it can make decisions. In more complicated scenarios, the light model can be modified by the control policy, for example, in order to adapt to changing traffic conditions. The lights are the same as modern-day lights: red (do not enter), yellow (if possible, do not enter; light will soon be red), and green (enter). Each control policy will need to have a light model so that human users will know what to do. For instance, the light model that would be used with ordinary FCFS would keep all the lights red at all times, informing humans that at no time is it safe to enter. The TRAFFIC-LIGHT policy, on the other hand, would have lights that corresponded exactly to the light system the policy is emulating. Here, we describe a few light models used in our experiments.

3.2.1 ALL-LANES

In this model, which is very similar to some current traffic light systems, each direction is successively given green lights in all lanes. Thus, all northbound traffic (turning and going straight) is given green lights while the eastbound, westbound, and southbound traffic all have red lights. The green lights then cycle through the directions. Figure 4 shows a graphical depiction of this light model.

3.2.2 SINGLE-LANE

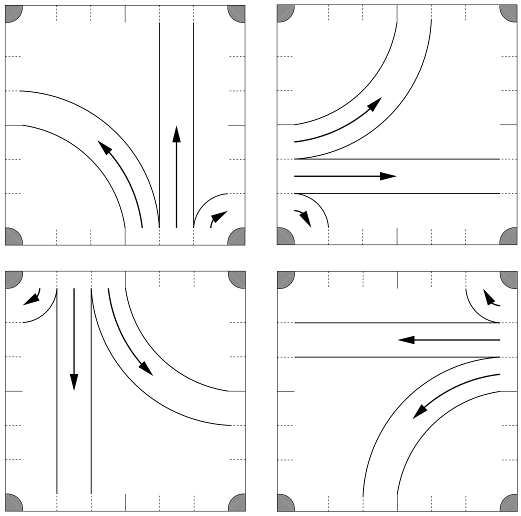

Figure 4: The ALL-LANES light model. Each direction is given all green lights in a cycle: north, east, west, south. During each phase, the only available paths for autonomous vehicles are right turns.

In the SINGLE-LANE light model, the green lane rotates through the lanes one at a time instead of all at once. For example, the left turn lane of the northbound traffic would have a green light, while all other lanes would have a red light. Next, the straight lane of the northbound traffic would have a green light, then the right turn. Next, the green light would go through each lane of eastbound traffic, and so forth. The first half of the model's cycle can be seen in Figure 5. This light model does not work very well if most of the vehicles are human-driven, but as we will show, is very useful for intersections which control mostly autonomous vehicles but need to also handle an occasional human driver.

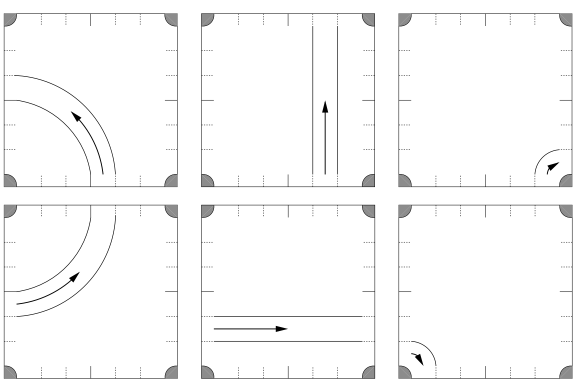

Figure 5: The first half-cycle of the SINGLE-LANE light model. Each individual lane is given a green light (left turn, straight, then right turn), and this process is repeated for each direction. Note how a smaller part of the intersection is used by turning vehicles at any given time. This provides an advantage for autonomous vehicles - there are many available paths through the intersection.

3.3 The FCFS-LIGHT Policy

In order to obtain some of the benefits of the FCFS policy while still accomodating human drivers, a policy needs to do two things:

1. If a light is green, ensure that it is safe for any vehicle (autonomous or human-driven) to drive through the intersection in the lane the light regulates.
2. Grant reservations to driver agents whenever possible. This would allow autonomous vehicles to move through an intersection where a human driver couldn't — similar to a "right on red", but extended much further to other safe situations.

The policy FCFS-LIGHT, which does both of these, is described as follows:

- As with FCFS, the intersection is divided into a grid of $n \times n$ tiles.
- Upon receiving a request message, the policy uses the parameters in the message to establish when the vehicle will arrive at the intersection.
- If the light controlling the lane in which the vehicle will arrive at the intersection would be green at that time, the reservation is confirmed.
- If the light controlling the lane would instead be yellow, the reservation is rejected.
- If the light controlling the lane would instead be red, the journey of the vehicle is simulated as in FCFS (Section 2.1).
- If throughout the simulation, no required tile is reserved by another vehicle or in use by a lane with a green or yellow light, the policy reserves the tiles and confirms the reservation. Otherwise, the request is rejected.

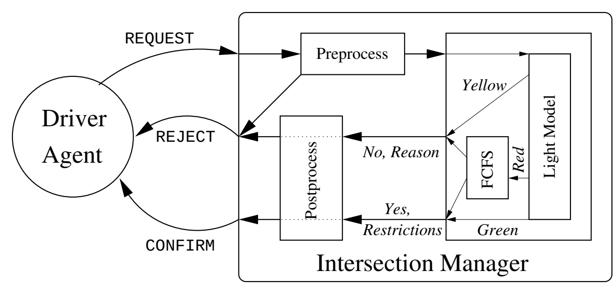

Figure 6: FCFS-LIGHT is the combination of FCFS and a light model. When a request is received, FCFS-LIGHT first checks to see what color the light would be. If it is green, it grants the request. If it is yellow, it rejects. If it is red, it defers to FCFS.

3.3.1 Off-Limits Tiles

Unfortunately, simply deferring to FCFS does not guarantee the safety of the vehicle. If the vehicle were granted a reservation that conflicted with another vehicle following the physical lights, a collision could easily ensue. To determine which tiles are in use by the light system at any given time, we associate a set of *off-limits tiles* with each light. For example, if the light for the northbound left turn lane is green (or yellow), all tiles that could be used by a vehicle turning left from that lane are off-limits. While evaluating a reservation request, FCFS also checks to see if any tiles needed by the requesting vehicle are off limits at the time of the reservation. If so, the reservation is rejected. The length of the yellow light is adjusted so that a vehicle entering the intersection has enough time to clear the intersection before those tiles are no longer off limits.

3.3.2 FCFS-LIGHT *Subsumes* FCFS

Using a traffic light-like light model (for example ALL-LANES), the FCFS-LIGHT can behave exactly like TRAFFIC-LIGHT on all-human driver populations. However, with a light model that kept all lights constantly red, FCFS-LIGHT behaves exactly like FCFS. That is, if any human drivers are present it will fail spectacularly, leaving the humans stuck at the intersection indefinitely. However, in the absence of human drivers, it will perform exceptionally well. FCFS is, in fact, just a special case of FCFS-LIGHT. We can thus alter FCFS-LIGHT's behavior to vary from strictly superior to TRAFFIC-LIGHT to exactly that of FCFS.

4. EMERGENCY VEHICLES

In current traffic laws there are special procedures involving emergency vehicles such as ambulances, fire trucks, and police cars. Vehicles are supposed to pull over to the side of the road and come to a complete stop until the emergency vehicle has passed. This is both because the emergency vehicle may be travelling quickly and because the emergency vehicle must arrive at its destination as quickly as possible — lives may be at stake. Hopefully, once a system such as this is implemented, automobile accidents — a major reason emergency vehicles are dispatched — will be all but eradicated. Nonetheless, emergency vehicles will still be required from time to time as fires, heart attacks, and other emergencies will still be around. While we have proposed other methods for giving priority to emergency vehicles [5], here we present a new, simpler method, which is fully implemented and tested.

4.1 Augmenting The Protocol

In order to accomodate emergency vehicles, the intersection manager must first be aware of their presence. We discovered that the easiest way to accomplish this was simply to add a field to all request messages. In our implementation, this field is simply a flag that indicates to the intersection manager that the requesting vehicle is an emergency vehicle in an emergency situation (i.e. with the siren and the lights on). In practice, however, safeguards would need to be incorporated to prevent normal vehicles from abusing this feature in order to obtain preferential treatment. This could be accomplished using some sort of secret key instead of simply a boolean value, or even some sort of public/private key challenge/response scenario. This level of implementation, however, is beyond the scope of this project and is already a well-studied area of cryptography and computer security.

4.2 The FCFS-EMERG Policy

Now that the intersection control policy has a way to detect emergency vehicles (in emergency situations), it can process reservation requests giving priority to the emergency vehicles. A first-cut solution is to simply deny reservations to any vehicles that were not emergency vehicles. This, however, is not satisfactory, because if all the traffic comes to a stop due to rejected reservation requests, the emergency vehicle(s) may get stuck in the resulting congestion. The FCFS-EMERG policy prevents this by keeping track of which lanes currently have approaching emergency vehicles. As long as at least one emergency vehicle is approaching the intersection, it only grants reservations to vehicles in those lanes. This ensures that vehicles in front of the emergency vehicles will also receive priority. Due to this increase in priority, even when traffic is fairly congested, lanes with emergency vehicles tend to empty very rapidly, allowing the emergency vehicle to continue on its way relatively unhindered.

5. EXPERIMENTAL RESULTS

We tested the efficacy of our new control policies with our custom-built, time-based simulator. The simulator models one intersection and has a time step of .02 seconds. The traffic level is controlled by changing the spawn probability — the probability that on any given time step, the simulator will attempt to spawn a new vehicle. For each experiment, the simulator simulates 3 lanes in each of the 4 cardinal directions. The total area modelled is a square with sides of 250 meters. The speed limit in all lanes is 25 meters per second. For each intersection control policy with reservation tiles, the granularity is set at 24. We also configured the simulator to spawn all vehicles turning left in the left lane, all vehicles turning right in the right lane, and all vehicles travelling straight in the center lane[1]. During each simulated time step, the simulator spawns vehicles (with the given probability), provides each vehicle with sensor data (simulated laser range finder, velocity, position, etc.), moves all the vehicles, and then removes any vehicles that have completed their journey. Unless otherwise specified, each data point represents 180000 time steps, or one hour of simulated time. Videos of each policy in action (as well as other supplementary material) can be found at `http://www.cs.utexas.edu/users/kdresner/aim/`.

As shown in our earlier work, once all vehicles are autonomous, intersection-associated delays can be reduced dramatically by using the two light models presented in Section 3.2. However, our experiments suggest a stronger result: delays can be reduced at each stage of adoption. Furthermore, at each stage there are additional incentives for drivers to switch to autonomous vehicles. Finally, our experiments verify the efficacy of the FCFS-EMERG policy, reducing emergency vehicle delays across the board.

5.1 Transition To Full Implementation

The whole point of having a hybrid light/autonomous intersection control policy is to confer the benefits of autonomy to passengers with driver-agent controlled vehicles while still allowing human users to participate in the system. Figure 7, which encompasses our main result, shows a smooth and monotonically improving transition from modern day traffic lights (represented by the TRAFFIC-LIGHT policy) to a completely or mostly autonomous vehicle mechanism (FCFS-LIGHT with the SINGLE-LANE light model). In early stages (100%-10% human), the ALL-LANES light model is used. Later on (less than 10% human), the SINGLE-LANE light model is introduced. At each change (both in driver populations and light models), delays are decreased. Notice the rather drastic drop in delay from FCFS-LIGHT with the ALL-LANES light model to FCFS-LIGHT with the SINGLE-LANE light model. Although none of the results is quite as close to the minimum as pure FCFS, the SINGLE-LANE light model allows for greater use of the intersection by the FCFS portion of the FCFS-LIGHT policy, which translates to more efficiency and lower delay.

For systems with a significant proportion of human drivers, the ALL-LANES light model works well — human drivers have the same experience they would with the TRAFFIC-LIGHT policy, but driver agents have extra opportunities to make it through the intersection. A small amount of this benefit is passed on to the human drivers, who may find themselves closer to the front of the lane while waiting for a red light to turn green. To explore how much the average vehicle would benefit, we ran our simulator with the FCFS-LIGHT policy, the ALL-LANES light model, and a 100%, 50%, and 10% rate of human drivers. This means that when a vehi-

[1]This is a constraint we will likely relax in the future. It is included in this work to give the SINGLE-LANE light model more flexibility and for a fair comparison to the FCFS policy, which performs even better in its absence.

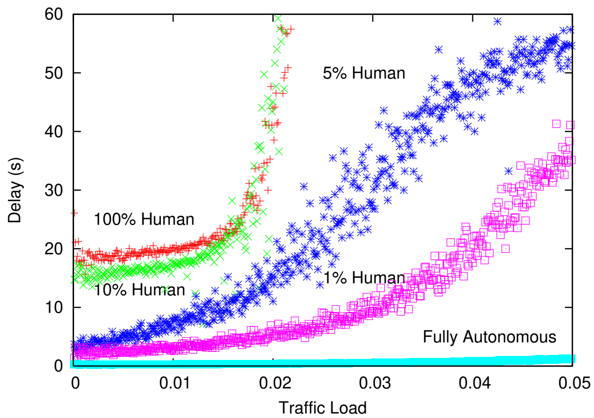

Figure 7: Average delays for all vehicles as a function of traffic level for FCFS-LIGHT with two different light models — the ALL-LANES light model, which is well-suited to high percentages of human-driven vehicles, and the SINGLE-LANE light model, which only works well with relatively few human-driven vehicles. As adoption of autonomous vehicles increases, average delays decrease.

cle is spawned, it receives a human driver (instead of a driver agent) with probability 1, .5, and .1 respectively. As seen in Figure 8, as the proportion of human drivers decreases, the delay experienced by the average driver also decreases. While these decreases are not as large as those brought about by the SINGLE-LANE light model, they are at least possible with significant numbers of human drivers.

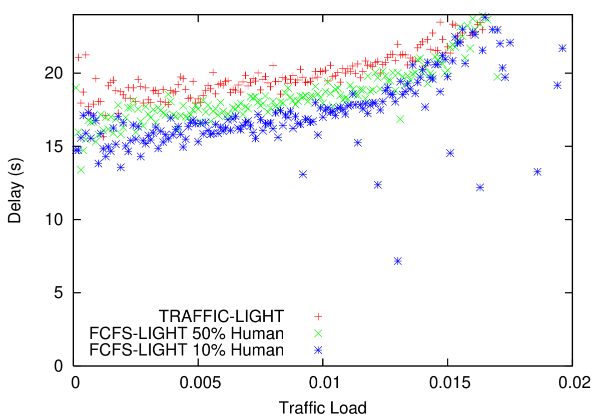

Figure 8: Average delays for all vehicles as a function of traffic level for FCFS-LIGHT with the ALL-LANES light model. Shown are the results for 100%, 50%, and 10% human-driven vehicles. The 100% case is equivalent to the TRAFFIC-LIGHT policy. Note that the average delay decreases as the percentage of human-driven vehicles decreases.

5.2 Incentives For Individuals

Even without any sort of autonomous intersection control mechanism, there are incentives for humans to switch to autonomous vehicles. Not having to do the driving, as well as the myriad safety benefits are strong incentives to promote autonomous vehicles in the marketplace. Our experimental results show additional incentives. Using our reservation system, autonomous vehicles experience lower average delays than human-driven vehicles and this difference increases as autonomous vehicles become more prevalent.

Shown in Figure 9 are the average delays for human drivers as compared to autonomous driver agents for the FCFS-LIGHT policy using the ALL-LANES light model. In this experiment, half of the drivers are human. Humans experience slightly longer delays than autonomous vehicles, but not worse than with the TRAFFIC-LIGHT policy. Thus, by putting some autonomous vehicles on the road, all drivers experience equal or smaller delays as compared to the current situation. This is expected because the autonomous driver can do everything the human driver does and more.

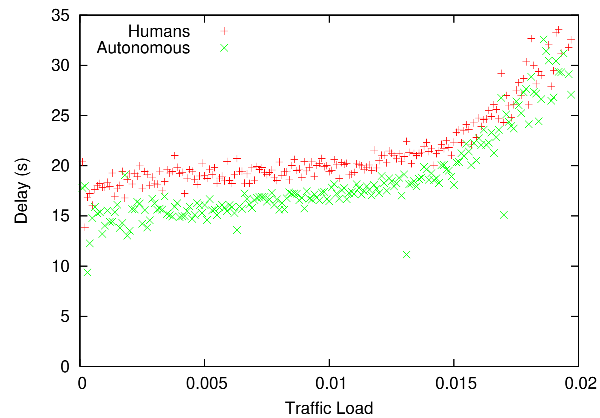

Figure 9: Average delays for human-driven vehicles and all vehicles as a function of traffic level for FCFS-LIGHT with the ALL-LANES light model. In this experiment, 50% of vehicles are human driven. Autonomous vehicles experience slightly lower delays across the board, and human drivers experience delays no worse than the TRAFFIC-LIGHT policy.

Once the reservation system is in widespread use and autonomous vehicles make up a vast majority of those on the road, the door is opened to an even more efficient light model for the FCFS-LIGHT policy. With a very low concentration of human drivers, the SINGLE-LANE light model can drastically reduce delays, even at levels of overall traffic that the TRAFFIC-LIGHT policy can not handle. Using the this light model, autonomous drivers can pass through red lights even more frequently because fewer tiles are off-limits at any given time. In Figure 10 we compare the delays experienced by autonomous drivers to those of human drivers when only 5% of drivers are human and thus the SINGLE-LANE light model can be used. While the improvements using the ALL-LANES light model benefit all drivers to some extent, the SINGLE-LANE light model's sharp decrease in average delays (Figure 7) comes at a high price to human drivers.

As shown in Figure 10, human drivers experience much higher delays than average. For lower traffic levels, the delays are even higher than they would experience with the TRAFFIC-LIGHT policy. Figure 7 shows that despite this, at high levels of traffic, the humans get a performance benefit. Additionally, these intersections will still be able to handle far more traffic than TRAFFIC-LIGHT.

The SINGLE–LANE light model effectively gives the humans a high, but fairly constant delay. Because the green light for any one lane only comes around after each other lane has had a green light, a human-driven vehicle may find itself sitting at a red light for some time before the light changes. However, since this light model would only be put in operation once human drivers are fairly

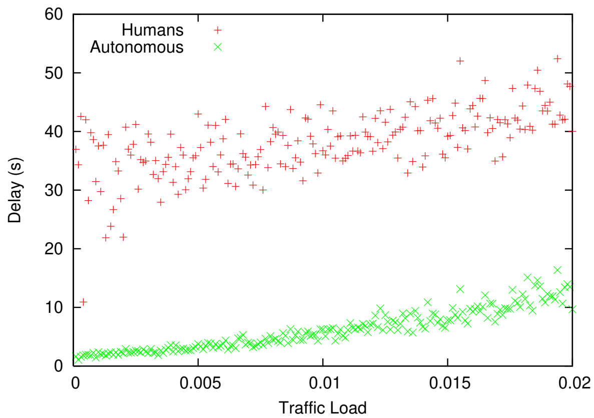

Figure 10: Average delays for human-driven vehicles and all vehicles as a function of traffic level for FCFS-LIGHT with the SINGLE-LANE light model. Humans experience worse delay than with TRAFFIC-LIGHT, but average delay for all vehicles is much lower. In this experiment, 5% of vehicles are human-driven.

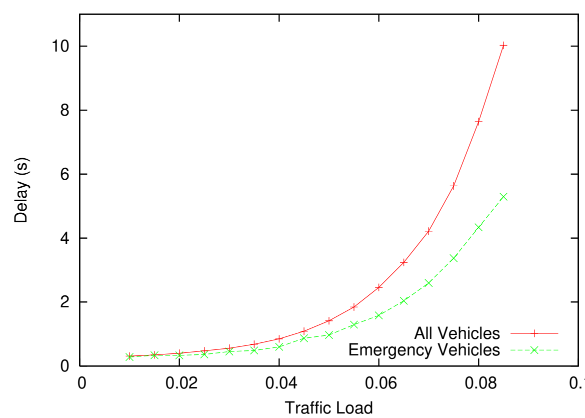

Figure 11: Average delays for all vehicles and emergency vehicles as a function of traffic level for the FCFS-EMERG policy. One out of a thousand vehicles (on average) is an emergency vehicle. Delays for the emergency vehicles are lower for all data points.

scarce, the huge benefit to the other 95% or 99% of vehicles far outweighs this cost. In Section 7, we propose a solution that could ameliorate these long delays for human drivers as well as slightly improving the overall performance of the system.

These data suggest that there will be an incentive to both early adopters (persons purchasing vehicles capable of interacting with the reservation system) and to cities or towns. Those with properly equipped vehicles will get where they are going faster (not to mention more safely). Cities and towns that equip their intersections to utilize the reservation paradigm will also experience fewer traffic jams and more efficient use of the roadways (along with fewer collisions, less wasted gasoline, etc.). Because there is no penalty to the human drivers (which would presumably be a majority at this point), there would be no reason for any party involved to oppose the introduction of such a system. Later, when most drivers have made the transition to autonomous vehicles, and the SINGLE-LANE light model is introduced, the incentive to move to the new technology is increased — both for cities and individuals. By this time, autonomous vehicle owners will far outnumber human drivers, who for high volumes of traffic will still benefit.

5.3 Lower Delays For Emergency Vehicles

While we have already shown that FCFS on its own would significantly reduce average delays for all vehicles, FCFS-EMERG helps reduce delays for such vehicles even further. To demonstrate this improvement, we ran our custom simulator with varying amounts of traffic, while keeping the proportion of emergency vehicles fixed at 0.1% (that is, a spawned vehicle is made into an emergency vehicle with probability 0.001). Because of the very small number of emergency vehicles created with realistically low proportions, we ran each configuration (data point) for 100 hours of simulated time — much longer than the other experiments. As shown in Figure 11, the emergency vehicles on average experienced lower delays than the normal vehicles. The amount by which the emergency vehicles outperformed the normal vehicles increased as the traffic increased, suggesting that as designed, FCFS-EMERG helps most when more traffic is contending for space-time in the intersection.

6. RELATED WORK

Currently, there is a considerable amount of research underway relating to intersection control and efficiency. Rasche and Naumann have worked extensively on decentralized solutions to intersection collision avoidance problems [11, 13]. Many approaches focus on improving current technology (systems of traffic lights). For example, Roozemond allows intersections to act autonomously, sharing the data they gather [17]. The intersections then use this information to make both short- and long-term predictions about the traffic and adjust accordingly. This approach still assumes human-controlled vehicles. Bazzan has used an approach using both MAS and evolutionary game theory which involves multiple intersection managers (agents) that must focus not only on local goals, but also on global goals [1].

Work is also being done with regard to the control of the individual vehicles. Hallé and Chaib-draa have taken a MAS approach to collaborative driving by allowing vehicles to form *platoons*, groups of varying degrees of autonomy, that then coordinate using a hierarchical driving agent architecture [6]. While not focusing on intersections, Moriarty and Langley have shown that reinforcement learning can train efficient driver agents for lane, speed, and route selection during freeway driving [9].

On real autonomous vehicles, Kolodko and Vlacic have created a small-scale system for intersection control which is very similar a reservation system with a granularity-1 FCFS policy [8].

Actual systems in practice (not MAS) for traffic light optimization include TRANSYT [15], which is an off-line system requiring extensive data gathering and analysis, and SCOOT [7], which is an advancement over TRANSYT, responding to changes in traffic loads on-line. However, almost all of the methods in practice or discussed above still rely on traditional signalling systems.

7. FUTURE WORK

Our system as demonstrated can vastly improve the traffic flow and transportation times experienced by all sorts of commuters. In this section, we present some ideas for improving and extending the system further.

7.1 More Intermediate Light Models

In order to smooth the transition further and reap the benefits of autonomous vehicles earlier, we plan to create light models that use less of the intersection than ALL-LANES, but don't restrict human drivers as much as SINGLE-LANE. These would provide the needed flexibility to let autonomous vehicles traverse the intersection using the FCFS portion of FCFS-LIGHT more frequently, decreasing delays relative to ALL-LANES.

7.2 Dynamic Light Models

All the light models presented in this paper have been static — that is they don't change as traffic conditions change. Traffic light systems in use today change throughout the day and week according to pre-programmed patterns created from expensive and time-consuming traffic studies. With the information gathered by the intersection manager and intersection control policy (via messages from the driver agents), the light model could be altered on-line. For example, in a situation with very few human drivers, the light model could keep all lights red until a human vehicle is detected (for example, with a transmitter), at which point the lane or direction from which the human driver is coming could be turned green. Once the human driver is through the intersection, the light(s) could be turned red again. This could offer a two-fold improvement over the SINGLE-LANE light model. First, the human drivers would benefit from not having to wait for the green light to make its way through all the other lanes at the intersection. This would make the system much more equitable to human drivers (who might otherwise have all the fun of driving taken away by extremely long delays at red lights). Secondly, the autonomous vehicles stuck behind the human drivers which would otherwise be stopped at red lights would also benefit. This secondary effect would likely have a much higher influence on the overall average delays, as the scenario assumes human drivers make up only a very small percentage of the total.

7.3 FCFS-LIGHT-EMERG?

This paper begs the question, "What about using both improvements simultaneously?" Unfortunately, making FCFS-LIGHT emergency vehicle-aware requires a dynamic light model as discussed above. However, given a dynamic light model, such an implementation is easy to describe. When the intersection control policy becomes aware of the emergency vehicle, the light model can be changed to one in which the green light rotates through the lanes that contain any approaching emergency vehicles.

7.4 Switching Policies On The Fly

While we have shown that the FCFS-LIGHT policy (with different light models) can span the gamut of scenarios from an all-human to all-autonomous driver population. With dynamic light models, it would seem that any situation could be handled by FCFS-LIGHT. However, should the need arise for a more radical change in intersection control policy (for example, to a stop sign policy in the case of road work or obstacle cleanup in the intersection), the reservation system should have a way to smoothly transition between the policies.

7.5 Learning Light Models/Policy Selection

Once we have a way to change between policies on-line, the next logical step is to get the intersection manager to choose its own policy or light model based on traffic conditions. If vehicles report their delays to the intersection when they finish crossing, the intersection manager will have access to a reinforcement signal that could be used to tune a light model or select a completely different policy altogether.

8. CONCLUSION

A science-fiction future with self-driving cars is becoming more and more believable. As intelligent vehicle research moves forward, it is important that we prepare to take advantage of the high-precision abilities autonomous vehicles have to offer. We have previously proposed an extremely efficient method for controlling autonomous vehicles at intersections. In this work, we have shown that at each phase of implementation, the system offers performance benefits to the average driver. Autonomous drivers benefit above and beyond this average improvement. We have also shown that the reservation system can be adapted to give priority to emergency vehicles, resulting in lower delays. Efficient, fast, and safe automobile transporation is not a fantasy scenario light-years away, but rather a goal toward which we can make worthwhile incremental progress.

Acknowledgements

This research is supported by NSF CAREER award IIS-0237699.

9. REFERENCES

[1] A. L. C. Bazzan. A distributed approach for coordination of traffic signal agents. *Autonomous Agents and Multi-Agent Systems*, 10(2):131–164, March 2005.

[2] R. Bishop. *Intelligent Vehicle Technology and Trends*. Artech House, 2005.

[3] 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, New York, USA, July 2004.

[4] 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.

[5] K. Dresner and P. Stone. Multiagent traffic management: Opportunities for multiagent learning. In K. Tuyls et al., editor, *LAMAS 2005*, volume 3898 of *Lecture Notes In Artificial Intelligence*. Springer Verlag, Berlin, 2006.

[6] S. Hallé and B. Chaib-draa. A collaborative driving system based on multiagent modelling and simulations. *Journal of Transportation Research Part C (TRC-C): Emergent Technologies*, 13:320–345, 2005.

[7] P. B. Hunt, D. I. Robertson, R. D. Bretherton, and R. I. Winton. SCOOT - a traffic responsive method of co-ordinating signals. Technical Report TRRL-LR-1014, Transport and Road Research Laboratory, 1981.

[8] J. Kolodko and L. Vlacic. Cooperative autonomous driving at the intelligent control systems laboratory. *IEEE Intelligent Systems*, 18(4):8–11, July/August 2003.

[9] D. Moriarty and P. Langley. Learning cooperative lane selection strategies for highways. In *Proceedings of the Fifeenth National Conference on Artificial Intelligence*, pages 684–691, Madison, WI, 1998. AAAI Press.

[10] National Highway Traffic Safety Administration. Economic impact of U.S. motor vehicle crashes reaches $230.6 billion, new NHTSA study shows. NHTSA Press Release 38-02, May 2002. Available at `http://www.nhtsa.dot.gov`.

[11] R. Naumann and R. Rasche. Intersection collision avoidance by means of decentralized security and communication

management of autonomous vehicles. In *Proceedings of the 30th ISATA - ATT/IST Conference*, 1997.

[12] D. A. Pomerleau. *Neural Network Perception for Mobile Robot Guidance*. Kluwer Academic Publishers, 1993.

[13] R. Rasche, R. Naumann, J. Tacken, and C. Tahedl. Validation and simulation of decentralized intersection collision avoidance algorithm. In *Proceedings of IEEE Conference on Intelligent Transportation Systems (ITSC 97)*, 1997.

[14] C. W. Reynolds. Steering behaviors for autonomous characters. In *Proceedings of the Game Developers Conference*, pages 763–782, 1999.

[15] D. I. Robertson. TRANSYT — a traffic network study tool. Technical Report TRRL-LR-253, Transport and Road Research Laboratory, 1969.

[16] S. Rogers, C.-N. Flechter, and P. Langley. An adaptive interactive agent for route advice. In O. Etzioni, J. P. Müller, and J. M. Bradshaw, editors, *Proceedings of the Third International Conference on Autonomous Agents (Agents'99)*, pages 198–205, Seattle, WA, USA, 1999. ACM Press.

[17] D. A. Roozemond. Using intelligent agents for urban traffic control control systems. In *Proceedings of the International Conference on Artificial Intelligence in Transportation Systems and Science*, pages 69–79, 1999.

[18] T. Schonberg, M. Ojala, J. Suomela, A. Torpo, and A. Halme. Positioning an autonomous off-road vehicle by using fused DGPS and inertial navigation. In *2nd IFAC Conference on Intelligent Autonomous Vehicles*, pages 226–231, 1995.

[19] P. Stone and M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3):345–383, July 2000.

[20] Texas Transportation Institute. 2004 urban mobility report, September 2004. Accessed at `http://mobility.tamu.edu/ums` in December 2004.