In *The Twentieth International Joint Conference on Artificial Intelligence* (IJCAI 07), pp. 1263-1268, Hyderabad, India, January 2007.

Sharing the Road: Autonomous Vehicles Meet Human Drivers

Kurt Dresner and **Peter Stone**
Department of Computer Sciences
The University of Texas at Austin
1 University Station C0500
Austin, Texas 78712-0233
{kdresner,pstone}@cs.utexas.edu
http://www.cs.utexas.edu/~{kdresner,pstone}

Abstract

In modern urban settings, automobile traffic and collisions lead to endless frustration as well as significant loss of life, property, and productivity. Recent advances in artificial intelligence suggest that autonomous vehicle navigation may soon be a reality. In previous work, we have demonstrated that a reservation-based approach can efficiently and safely govern interactions of multiple autonomous vehicles at intersections. Such an approach alleviates many traditional problems associated with intersections, in terms of both safety and efficiency. However, the 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 extend this system to allow for incremental deployability. The modified system is able to accommodate 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 our previous work. Finally, we develop a method for switching between various human-usable configurations while the system is running, in order to facilitate an even smoother transition. The work is fully implemented and tested in our custom simulator, and we present detailed experimental results attesting to its effectiveness.

1 Introduction

In modern urban settings, automobile traffic and collisions lead to endless frustration as well as significant loss of life, property, and productivity. A recent study of 85 U.S. cities [Texas Transportation Institute, 2004] put the annual time spent waiting in traffic at 46 hours per capita, up from 16 hours in 1982. A recent report puts the annual societal cost of automobile collisions in the U.S. at $230 billion [National Highway Traffic Safety Administration, 2002]. Meanwhile, advances in artificial intelligence suggest that autonomous vehicle navigation may soon be a reality. Cars can now be equipped with features such as adaptive cruise control, GPS-based route planning [Rogers *et al.*, 1999; Schonberg *et al.*, 1995], and autonomous steering [Pomerleau, 1993]. As more and more cars become autonomous, the possibility of autonomous interactions among multiple vehicles becomes a more realistic possibility.

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 [Wooldridge, 2002]. In previous work, we demonstrated a novel MAS-based approach to alleviating traffic congestion and collisions, specifically at intersections [Dresner and Stone, 2005]. However, the system relied on all vehicles being equipped with the requisite technology — a restriction that would make implementing such a system in the real world extremely difficult.

This paper makes two main contributions. First, we show how to augment this 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 the mechanism offers an incentive for individuals to use autonomous-driver-agent-equipped vehicles. Our work is fully implemented and tested in a custom simulator and detailed experimental results are presented.

2 Reservation System

Previously, we proposed a reservation-based multi-agent approach to alleviating traffic, specifically at intersections [Dresner and Stone, 2005]. 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 [Dresner and Stone, 2004a]. In addition to message types (e.g. REQUEST, CONFIRM, and CANCEL), this protocol includes some rules, the most important of which are (1) a vehicle may not enter the intersection unless it is within the parameters of a reservation made by that vehicle's driver agent, (2) 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 its 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. This system is a true multiagent system, integrating heterogeneous, complete, autonomous agents.

2.1 First Come, First Served (FCFS)

To determine if a request can be met, the intersection manager uses what we call the "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.
- If throughout this simulation, no required tile is already reserved, the policy reserves the tiles for the vehicle and confirms the reservation. Otherwise, it rejects it.

We name the policy based on the fact that it 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 is given the reservation (provided there are no conflicts in the rest of the required space-time). Figure 1 shows a successful reservation (confirmed) followed by an unsuccessful reservation (rejected).

Our previous work demonstrated that an intersection manager using the FCFS policy enabled vehicles to cross the intersection while experiencing almost no *delay* (increase in travel time over optimal) [Dresner and Stone, 2004b; 2005]. When compared with a standard traffic light, the FCFS policy not only decreased delay, but also tremendously increased the throughput of the intersection. For any realistic

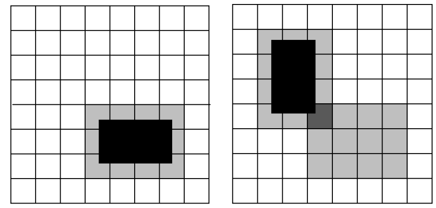

(a) Successful (b) Rejected

Figure 1: A granularity-8 FCFS policy. In 1(a), vehicle A's request reserves tiles at time t. In 1(b), vehicle B's request is rejected because it requires a tile used by A at t.

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 several times higher than for the traffic light.

2.2 Other Intersection Control Policies

While the reservation system as originally proposed [Dresner and Stone, 2004b] uses only the FCFS policy, it can accommodate any intersection control policy that can make a 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 the most common modern-day control mechanisms, the reservation mechanism can perform as well as current systems, provided it uses a reasonable control policy.

Traffic lights are a ubiquitous mechanism used to control high-traffic intersections. In our previous work, we describe a policy that emulates real-life traffic lights by maintaining a model of how the lights would be changed, were they to exist. We name this policy TRAFFIC-LIGHT. 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. In this paper, we extend this work to include a new policy that covers the area between TRAFFIC-LIGHT and FCFS, thereby enabling interoperability with human drivers as well as providing a smooth transition as autonomous vehicles become more prevalent.

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 accommodate humans, whether they are on a bicycle, walking to the corner store, or driving a "classic" car for entertainment purposes. Allowing both humans and autonomous vehicles in the same intersection at the same time presents an interesting technical challenge. In this section we explain how we have extended the FCFS policy and reservation framework to allow for human drivers.

3.1 Using Existing Infrastructure

To add human drivers, we first need a reliable way to communicate information to them. Fortunately, we can use a system that drivers already know and understand — traffic lights. The required 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 can be used that will give enough time for a person to cross as well as alerting the intersection manager 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 needs 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.

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 north traffic (turning and going straight) is given green lights while the east, west, and south traffic all have red lights. The green lights then cycle through the directions. Figure 2 shows a graphical depiction of this light model.

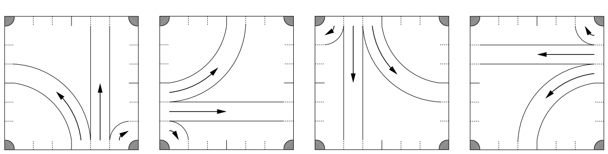

Figure 2: The ALL-LANES light model. Each direction is given all green lights in a cycle: north, east, west, south. During each phase, the only available paths for autonomous vehicles are right turns.

SINGLE-LANE

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 north traffic would have a green light, while all other lanes would have a red light. Next, the straight lane of the north traffic would have a green light, then the right turn. Next, the green light would go through each lane of east traffic, and so forth. The first half of the model's cycle can be seen in Figure 3. This light model does not work very well if most of the vehicles are human-driven. However, we will show that it is very useful for intersections which control mostly autonomous vehicles but also need to handle an occasional human driver.

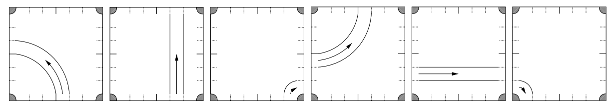

Figure 3: The first half-cycle of the SINGLE-LANE light model. Each lane is given a green light, and this process is repeated for each direction. A small part of the intersection is used by turning vehicles at any given time.

3.3 The FCFS-LIGHT Policy

In order to obtain some of the benefits of the FCFS policy while still accommodating human drivers, a policy needs to do two things. First, if a light is green, it must ensure that it is safe for any vehicle (autonomous or human-driven) to drive through the intersection in the lane the light regulates. Second, it should 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 we have created which does both of these is called FCFS-LIGHT. As with FCFS, the intersection is divided into a grid. When it receives a request, FCFS-LIGHT immediately grants a reservation if the corresponding light will be green at that time. Otherwise, it simulates the vehicle's trajectory as in FCFS. If throughout the simulation, no required tile is reserved by another vehicle or in use by a lane with a green light the policy reserves the tiles and confirms the reservation. Otherwise, the request is rejected.

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 north 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.

FCFS-LIGHT Subsumes FCFS

Using a traffic light-like light model (for example ALL-LANES), the FCFS-LIGHT can behave exactly like TRAFFIC-

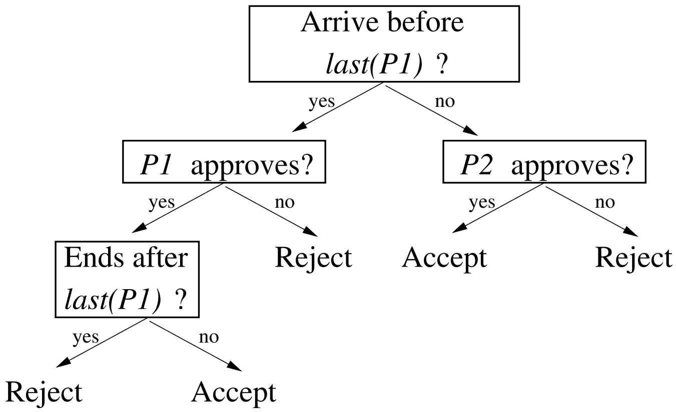

Figure 4: The decision mechanism during a switchover from policy $P1$ to policy $P2$.

LIGHT on all-human driver populations. With a light model that keeps all lights constantly red, FCFS-LIGHT behaves exactly like FCFS. If any human drivers are present it will leave them 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 from TRAFFIC-LIGHT at one extreme to FCFS at the other.

3.4 Policy Switching

Policies based on certain light models may be more effective at governing particular driver populations. In order to further reduce delay, we have developed a way in which the intersection manager can smoothly switch between different policies — the intersection need not bring all the vehicles to a halt and clear out the intersection. The switching mechanism requires every policy to keep track of the last time for which it has authorized a vehicle to be in the intersection. This could be either the last moment of a reservation or the last moment that a vehicle passing through a green light can be in the intersection. Once the intersection manager decides to make the switch, it *freezes* the current policy. When a policy is frozen, it rejects all requests that would cause it to increase the "last reserved" time. Once the current policy is frozen, the intersection manager accesses the last reserved time and henceforth delegates all reservation requests that begin after that time to the new policy. All requests that begin before that time are still processed by the current policy. Because all requests are handled entirely by one policy, if two policies $P1$ and $P2$ are safe (vehicles granted reservations by it are guaranteed not to collide), the same will be true for the period during which the intersection manager is making the switch. A diagram illustrating this compound decision mechanism is shown in Figure 4.

4 Experimental Results

We test the efficacy of the new control policies with a custom-built, time-based simulator. Videos of the simulator in action can be viewed at `http://www.cs.utexas.edu/~kdresner/aim/`. 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, in order to make the comparison between ALL-LANES and SINGLE-LANE more straightforward. 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.

In previous work, we demonstrated that once all vehicles are autonomous, intersection-associated delays can be reduced dramatically. By using the light models presented earlier, we can obtain 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.

4.1 Transition To Full Implementation

The 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 5(a), 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 the effect on the average vehicle, we run our simulator with the FCFS-LIGHT policy, the ALL-LANES light model, and a 100%, 50%, and 10% rate of human drivers: when a vehicle is spawned, it receives a human driver (instead of a driver agent) with probability 1, .5, and .1 respectively. As seen in Figure 5(b), as the proportion of human drivers decreases, the delay experienced by the average driver also decreases. While these decreases are not

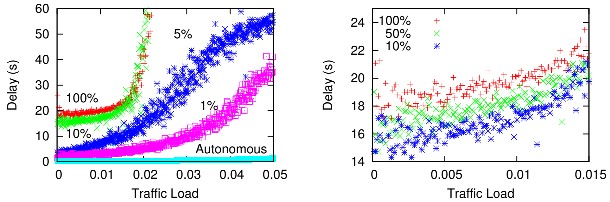

(a) Delays decrease with decreasing proportions of human drivers.

(b) A closeup of the ALL-LANES results.

Figure 5: Average delays for all vehicles as a function of traffic level for FCFS-LIGHT. ALL-LANES is well-suited to high percentages of human drivers (100%-10%), while SINGLE-LANE works well with few humans (10%-0%).

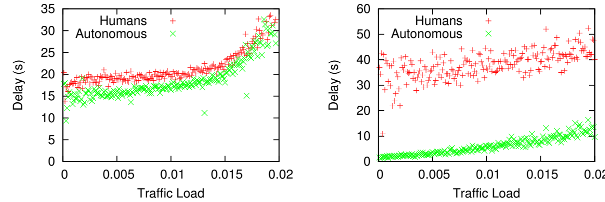

(a)

(b)

Figure 6: Average delay for humans and autonomous vehicles as a function of traffic level for FCFS-LIGHT with 50% human, ALL-LANES(a) and 5% human, SINGLE-LANE(b) .

as large as those brought about by the SINGLE-LANE light model, they are at least possible with significant numbers of human drivers.

4.2 Incentives For Individuals

Clearly there are incentives for cities to implement the FCFS-LIGHT policy — the roads will be able to accommodate more traffic and vehicles will experience lower delays. However, we have also shown that these benefits only materialize when a significant portion of the vehicles are autonomous. Here we demonstrate that the system creates incentives for individuals to adopt autonomous vehicles as well, in the form of lower delays for autonomous vehicles. Shown in Figure 6(a) 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 (compare with Figure 5(b)). 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.

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 6(b) 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 5(a)) appears to come at a high price to human drivers.

As shown in Figure 6(b), 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. However, figure 5(a) shows that despite this, at high levels of traffic, the humans get a performance benefit. Because these intersections will be able to handle far more traffic than TRAFFIC-LIGHT, the fact that the human delays are kept more or less constant (as opposed to the skyrocketing delay of Figure 5(a)), means this is actually a win for the human drivers.

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. Because this light model would only be put in operation once human drivers are fairly scarce, the huge benefit to the other 95% or 99% of vehicles far outweighs this cost.

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 initially be a majority), 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 will still benefit as traffic volumes continue to increase.

5 Related Work

Rasche and Naumann have worked extensively on decentralized solutions to intersection collision avoidance problems [Naumann and Rasche, 1997; Rasche *et al.*, 1997]. Others focus on improving current technology (systems of traffic lights). For example, Roozemond allows intersections to act autonomously, sharing the data they gather [Roozemond, 1999]. 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 [Bazzan, 2005].

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 [Hallé and Chaib-draa, 2005]. 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 [Moriarty and Langley, 1998]. 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 [Kolodko and Vlacic, 2003].

Actual systems in practice for traffic light optimization include TRANSYT [Robertson, 1969], which is an off-line system requiring extensive data gathering and analysis, and SCOOT [Hunt *et al.*, 1981], which is an advancement over TRANSYT, responding to changes in traffic loads on-line. However, all methods for controlling automobiles in practice or discussed above still rely on traditional signalling systems.

6 Conclusion

We have extended an extremely efficient method for controlling autonomous vehicles at intersections such that at each phase of implementation, the system offers performance benefits to the average driver. Autonomous drivers benefit above and beyond this average improvement, which creates additional incentives for individuals to adopt autonomous vehicle technology. We also showed how the system can move between different control policies smoothly and safely.

This work opens up the possibility of creating light models that use less of the intersection than ALL-LANES, but don't restrict human drivers as much as SINGLE-LANE. These intermediate models 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. Additionally, adaptive light models could ameliorate the delays for human drivers associated with the SINGLE-LANE light model. If an intersection manager can receive a reward signal based on delays experienced by vehicles, the manager can learn light models and adapt on-line.

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. Efficient, fast, and safe automobile transporation is not a fantasy scenario light-years away, but rather a goal toward which we can make worthwhile step-by-step progress.

References

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

[Dresner and Stone, 2004a] Kurt Dresner and Peter Stone. Multiagent traffic management: A protocol for defining intersection control policies. Technical Report UT-AI-TR-04-315, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, December 2004.

[Dresner and Stone, 2004b] Kurt Dresner and Peter 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.

[Dresner and Stone, 2005] Kurt Dresner and Peter 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.

[Hallé and Chaib-draa, 2005] Simon Hallé and Brahim 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.

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

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

[Moriarty and Langley, 1998] David Moriarty and Pat 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.

[National Highway Traffic Safety Administration, 2002] 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`.

[Naumann and Rasche, 1997] Rolf Naumann and Rainer 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.

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

[Rasche *et al.*, 1997] 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.

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

[Rogers *et al.*, 1999] Seth Rogers, Claude-Nicolas Flechter, and Pat Langley. An adaptive interactive agent for route advice. In Oren Etzioni, Jörg P. Müller, and Jeffrey M. Bradshaw, editors, *Proceedings of the Third International Conference on Autonomous Agents (Agents'99)*, pages 198–205, Seattle, WA, USA, 1999. ACM Press.

[Roozemond, 1999] Danko 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.

[Schonberg *et al.*, 1995] 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.

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

[Wooldridge, 2002] Michael Wooldridge. *Introduction to MultiAgent Systems*. John Wiley & Sons, February 2002.