

A RESERVATION-BASED MULTIAGENT SYSTEM FOR INTERSECTION CONTROL

Kurt Dresner and Peter Stone¹

*Department of Computer Sciences
The University of Texas at Austin
1 University Station, C0500
Austin, TX 78712-0233
{kdresner,pstone}@cs.utexas.edu*

Abstract: Traffic congestion is one of the leading causes of lost productivity and decreased standard of living in urban settings. Recent advances in artificial intelligence suggest that autonomous vehicle navigation will be possible in the near future. In this paper, we propose a reservation-based system for alleviating traffic congestion, specifically at intersections. First, we describe a custom simulator created to measure the different delays associated with conducting traffic through an intersection. Second, we specify a precise metric for evaluating the quality of traffic control at an intersection. Using this simulator and this metric, we show that the reservation-based system can perform two to three hundred times better than traffic lights. As a result, it can smoothly handle much heavier traffic conditions. We demonstrate that this system very closely approximates an overpass, which is the optimal solution for problem with which we are dealing.

Keywords: autonomous vehicles, traffic control, multiagent systems.

1. INTRODUCTION

Traffic congestion is one of the leading causes of lost productivity and decreased standard of living in urban settings. 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 cruise control, GPS-based route planning (Rogers *et al.*, 1999; Schonberg *et al.*, 1995), and autonomous steering (Pormerleau, 1993; Hakim, 2004). Once individual cars become autonomous, it is inevitable that before long all, or most, of the cars on the road will have such capabilities, thus opening up the possibility of considering 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 (Stone and Veloso, 2000). In this paper, we propose an MAS-based approach to alleviating traffic congestion, specifically at intersections.

Current methods for enabling traffic to flow through intersections include building overpasses and installing traffic lights. However, the former is only worth the cost at the most congested intersections, and the latter can be quite inefficient, often requiring cars to remain stopped even when no cars are present on the intersecting road.

Here we aim to maximize the efficiency of moving cars through intersections with minimal centralized infrastructure. We assume that intersections can be outfitted with a wireless communication system and a protocol for

communicating with oncoming traffic and giving permission for cars to pass.

Cars must only traverse intersections when allowed to by the protocol (as they do today by obeying red and green lights), but otherwise are free to decide for themselves how to drive. That is, each car is treated as an autonomous agent, and in particular need not surrender control to any centralized decision maker.

Given the above assumptions, we propose a novel reservation-based system by which cars request and receive time slots from the intersection during which they may pass. In this paper, we present theoretical and empirical results relating the overpass, traffic signal, and reservation-based intersection control mechanisms. Each of these mechanisms is fully implemented in a custom traffic simulator, which we introduce here. Our results indicate that the reservation-based system allows traffic to flow through the intersection much more efficiently than the traffic light mechanism, and with efficiency approaching that of the overpass mechanism.

2. THE MODEL

Our model of intersection traffic is a somewhat simplified version of real-world intersection traffic. We do not allow cars to turn, and all cars begin travelling roughly the same speed. These simplifications make the analysis and implementation a lot easier, but do not detract from the fundamental challenges of the problem. Since we are trying to create a system where vehicles reach their destinations in a minimum amount of time, it is not unreasonable to expect that each vehicle is always trying to travel at the speed limit. Additionally, almost every vehicle in production is capable of attaining the speed limit of any road, so assuming they all travel at the speed limit is not a far stretch.

¹ This research was supported in part by NSF CAREER award IIS-0237699.

Once we make these assumptions, we next consider the question: How do we measure whether or not one intersection is better than another? Certainly, safety is a primary concern. One of the primary reasons for pursuing automated driving agents is to make car travel safer. However, even at this time, it is possible to create a system in which cars drive themselves, and humans are merely passengers. Such a system would unfortunately be prohibitively expensive. Modifications would need to be made to the entire road system of the country, and quite possibly to every vehicle on that road system. Such a system would also likely be quite inefficient. To measure intersection efficiency we consider specific metrics related to throughput and delay. Throughout this study we consider safety to be a prerequisite — any intersection control policy used in the real world must have 0 probability of collisions, provided the agents follow the protocol correctly.

2.1 Throughput

One metric we examine briefly is the amount of traffic that can be handled by an intersection. While this is hard to measure precisely, we make qualitative claims regarding the throughput of several different systems.

2.2 Delay

Delay is the primary metric we consider; what effect does the presence of the intersection have on the overall journey of a vehicle? Additionally, we would like to ensure that no vehicle's travel time is dramatically increased in order to save a few seconds for some other vehicle. We want to ensure that the average case is not bad, but we also want to ensure that the worst case is not *too* bad.

2.2.1. Average Delay Consider the set of all vehicles passing through an intersection in a period of time to be C . Assume that were there no other vehicles on the road, vehicle v_i would have made its trip from point A to point B in time $t_0(i)$. However, due to other cars in the same lane and due to having to go through an intersection involving other vehicles, v_i arrives instead at time $t(i)$. Then we define the *average delay* of an intersection to be:

$$\frac{1}{|C|} \sum_{v_i \in C} t(i) - t_0(i)$$

2.2.2. Maximum Delay The worst case delay is then:

$$\max_{v_i \in C} t(i) - t_0(i)$$

3. OVERPASS, TRAFFIC LIGHT THEORY

Now that we've defined the metrics we are going to use, we can analyze the theoretical performance of some intersection control systems.

First, note that with an overpass the cars never have to slow down at the intersection. Thus, the average and maximum delay are both 0.

Analysis of the traffic light model is somewhat trickier. In this model, a vehicle has a random chance (at an isolated intersection) of making it through the intersection unimpeded. If it does not, the vehicle must come to a complete stop and wait, regardless of the amount of cross-traffic. The average and maximum delays are

complicated functions of the timing of the lights (how long they are green and red), how many cars are on the road, and what the velocities of the other cars are, among other things. In order to analyze this model, we first make some simplifying assumptions.

- (1) *Cars travelling in the same direction do not interact with one another.*

By allowing cars travelling in the same direction to pass through one another, we can remove car-car interactions from the model altogether. We are looking for a theoretical lower bound, and car-car interactions can only increase the delay on the average vehicle.

- (2) *Cars that have to decelerate due to a red light reach the intersection at full speed precisely when the light turns green again.*

Again, as we are looking for a lower bound, we can give the driver agent more information than a normal traffic light would. This allows the driver agent to accelerate at just the right time such that it loses the minimum amount of time due to the traffic light.

Consider the following parameters:

P — the period of the traffic light

α — the fraction of the light's period that the light spends on green in one direction.

From the definitions of these parameters, we can specify two constraints:

$$P > 0 \quad \text{and} \quad 0 < \alpha \leq 1 \tag{1}$$

Because we have assumed the cars are independent, we can find the average delay per car by finding the expected delay for one car. With the assumption that the vehicles always reach the green light at top speed, we remove any dependence on the acceleration or deceleration of the vehicle. To calculate the delay for one vehicle, we find the difference between the time the vehicle would have arrived at the traffic light were it always green and the time that the vehicle actually arrives at the traffic light.

In the end, we get an answer that is dependent only on α and P . For α of the time, the car reaches the green light without having to decelerate and experiences no delay at all. The remaining $(1 - \alpha)$ of the time, the car does not reach the intersection until the beginning of the next cycle. Since the length of time for which the light is red is $(1 - \alpha)P$, on average a car that does not make the green light will wait $\frac{1}{2}(1 - \alpha)P$. Thus, the total expected delay for a car is:

$$\frac{1}{2}(1 - \alpha)^2 P \tag{2}$$

As α approaches 1, the expected delay approaches zero. As α approaches zero, the expected delay approaches $\frac{1}{2}P$. This is slightly counter-intuitive — common sense says the delay should grow without bound — but it is important to remember that at the exact moment the signal turns green (even if only for an instant), the vehicle is allowed to go. As the period of the light increases, our expected delay increases linearly - the same fraction of vehicles miss the light, and those vehicles have to wait longer.

The maximum delay can be calculated similarly. The longest amount of time for which a driver could wait is the length of time for which there is no green, or:

$$(1 - \alpha)P \quad (3)$$

4. THE SIMULATOR

To test our traffic control policies, we developed a time-based simulator to model each car individually. The simulator (not including the intersection manager) runs with two main parameters:

- (1) Number of lanes travelling in each of the four directions. For all experiments reported in this paper, the number of lanes for each direction was the same.
- (2) Probabilities of attempting to spawn a vehicle in each direction (independently) at the beginning of each time step. For almost all of our experiments, the probability for each direction was the same.

For our simulations, the simulator models an area that is 400 m \times 400 m. Lanes are 3.5 m wide, and each vehicle is 2 m wide by 4 m long. Figure 1 shows a screenshot of the graphical display. Java applets showing the simulator in action can be viewed at <http://www.cs.utexas.edu/users/kdresner/papers/2004iav/>. In each cycle of the simulator (which for our experiments represents one fiftieth of a second) the following events occur:

- (1) In each direction, vehicles are randomly spawned with a predefined probability. Once a vehicle is spawned, it is placed uniformly at random in one of the lanes traveling in that direction. If placing the vehicle in that lane and direction would cause the vehicle to be following another vehicle too closely (within 1 second² or 1 meter), the vehicle is not spawned. Our simulator spawns all vehicles travelling at the speed limit and never spawns a vehicle where it would be in danger of colliding with another vehicle.
- (2) The driver of each vehicle is given a list of the nearby vehicles by which it can access their velocities, positions, etc. This information simulates information that could be sensed with on-board sensors such as cameras and/or range-finders and is used to avoid collisions with cars travelling in the same direction.
- (3) Each driver then takes an action based on this information: ACCELERATE (increase velocity by 3 m/s²), DECELERATE (decrease velocity by 15 m/s²), or COAST (maintain current velocity). The simulator enforces the invariant that $0 \leq \text{speed} \leq \text{top speed}$.
- (4) Each vehicle's position and velocity are updated according to the actions the driver took.
- (5) Any vehicles which have left the domain of the simulator are removed from the simulation. Each vehicle keeps track of how much delay it has experienced, and this value is used to update the global value for total delay and maximum delay. The total number of vehicles which have finished their journey is also updated.

The behavior of the driver agent is completely independent of the simulator, and can be different for every vehicle. For our experiments, we used the same driver agent in each vehicle. It behaves as follows when it receives vehicle information from the simulator.

- COAST;
- If `speed < speed limit`, ACCELERATE;
- If less than 1 second or 1 meter behind the vehicle in front and `speed > 0`, DECELERATE;³
- If not through the intersection already, interact with the intersection as specified separately for each intersection control policy in Section 5;

As an initial test of the simulator, we begin by verifying that the empirical performance matches our theoretical predictions for the cases that can be analyzed completely, namely the overpass and traffic light intersection controllers. We ran some simulations to determine both the average⁴ and maximum delays as functions of both α and P . Figure 2 shows the average and maximum values which we obtained running the simulator for 1,000,000 steps with 1 lane in each direction, spawning cars with probability .002, letting α be .45, and varying P from 10 to 200 in steps of .1. Similarly, Figure 3 shows the average and maximum values obtained by fixing the period at 30, while allowing α to vary from .01 to .89 with steps of .001. Because α is defined as the fraction of the time the light is green in the North/South directions, we only run the simulator with cars spawning in the North/South directions. The graphs for the East/West directions are just a horizontal reflections of these data.




Fig. 1: A screenshot of the graphical display of our simulator with 2 lanes in each direction. Cars travelling in the East/West directions are stopped at a red light.

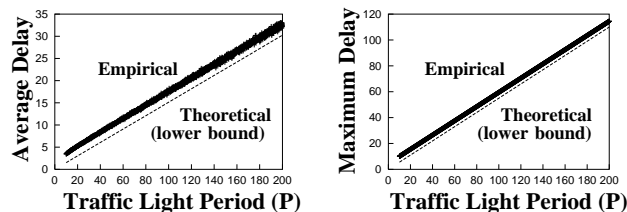


Fig. 2. Average (left) and maximum (right) delay as a function of P for 1,000,000 steps of the simulation, spawning cars with probability .002.

The empirical data is consistent with our theoretical predictions modulo the fact that our driver agent does not have the ability to reach every green light at full speed. The discrepancies between the lines are consistent with our expectations — they can be attributed to the time it takes for a vehicle to come to a stop at the light and then accelerate again to full speed. Thus we conclude that the simulator and driving agents operate correctly.

² A vehicle a is considered within t seconds of vehicle b if a is behind b and the distance between the back of b and the front of a is $\text{velocity}_a \cdot t$.

³ This is not done when vehicles are allowed to pass through one another.

⁴ Unless otherwise stated, all averages are over all cars spawned in one run of the simulation.

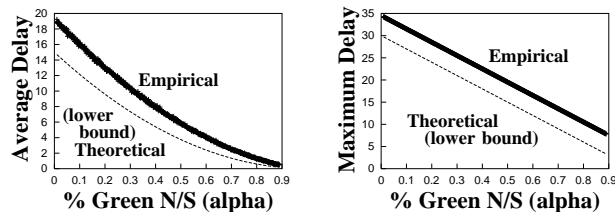


Fig. 3. Average (left) and maximum (right) delay as a function of α for 1,000,000 steps of the simulation, spawning cars with probability .001.

5. INTERSECTION CONTROL POLICIES

Using our simulator, we evaluated the performance of three different intersection control policies: the overpass, the traffic light, and the reservation system.

5.1 Overpass

The overpass is by far the simplest of the three policies: it lets vehicles through all the time. While there is no explicit third dimension in the simulator, by allowing vehicles traveling in orthogonal directions to pass through one another, the same results can be achieved. If the ability to turn is not required (as in our current study) the overpass is an optimal solution. The only delays are caused by cars traveling the same direction, which would happen even if there were not an intersection through which to travel.

5.2 Traffic Light

The traffic light model is a close approximation of actual traffic signals used on real roads. The model has three parameters: the period of the light system, the time between green lights in which the lights in all four directions are red (as a fraction of the period), and the time for which the North/South lights are green (again as a fraction of the period). North and South lights are always identical, as are East and West lights. Yellow lights are not necessary in our system. In real-world traffic signals, yellow lights alert the drivers as to when the next red light is coming. However in our model, the drivers can actually query the light to determine in which state the light will be at a given time. This improves the ability of the traffic light to direct the vehicles through the intersection. It is important to note that simplifying assumption 2 from Section 3 does not apply to this model, because it is specific to the driving agent. Our driver agent does not time the green lights perfectly, although with the traffic light system we implemented, such a driver agent could be created.

The interaction between the driver agent and the traffic light is straightforward. First, the driver calculates when it will reach the traffic light given its current velocity. If the light will be red at that time, the driver decelerates.

5.3 Reservation System

In creating the reservation system we defined both the behavior of the system as well as the behavior of the driver agent when interacting with the system.

5.3.1. Behavior of the System The reservation system, which we introduce here, allows the driver agents to “call ahead” and reserve the spaces they will need. The intersection is divided into an $n \times n$ grid of reservation

tiles, each of which can be reserved by one car per time step. The parameter n is referred to as the *granularity* of the system. A vehicle can ask the reservation system if it will be able to cross the intersection at a certain time, velocity, and starting position. To determine this, the intersection simulates the journey of the vehicle through the intersection with the parameters provided, noting which cells will be occupied by the vehicle at each time step. If any of these cells is already reserved, the intersection responds to the driver’s request with FALSE. Otherwise, it responds to the driver’s request with TRUE.

5.3.2. Behavior of the Driver Agent At the beginning of every cycle, the driver agent determines when it expects to reach the intersection, where it will be when it first enters the intersection, and how fast it will be going.

If the driver has not yet made a reservation, it asks the intersection whether or not that combination of parameters is available. If the intersection responds with TRUE, the driver agent notes that a reservation has been made (along with the parameters used), and freezes its velocity. If the intersection responds with FALSE, the driver agent chooses to decelerate.

If the driver has made a reservation, it checks the parameters computed at the beginning of the cycle against the parameters used in its reservation. If it determines that it cannot meet the reservation, it cancels the reservation and the reservation-making process begins again. If it can meet the parameters, it continues at its current velocity.

5.3.3. Constraints on Acceleration Once a reservation is made, the driver must maintain its velocity until it is through the intersection. This is because the velocity and acceleration of a vehicle depend heavily on the vehicles around it, as well as factors like the speed limit or the vehicle’s top speed. Allowing the vehicle to accelerate or decelerate presents a more difficult problem when calculating which tiles the vehicle will need, so for the current implementation this is explicitly disallowed.

6. EMPIRICAL RESULTS

In this section, we evaluate the performance of our reservation system against both the overpass and traffic signal systems for varying amounts of traffic, numbers of lanes, and granularities of the reservation system.

First, it is important to note the amount of traffic that each system can handle. The overpass system can handle as much traffic as we can generate with our simulator, since it does not cause the cars to slow down at all. It is interesting, though to examine the throughput of the traffic signal system. One problem we had was that if we create traffic faster than we can move it through the intersection, it builds up until the simulator can’t simulate it all. For this reason, the simulator stops adding vehicles when it gets full. Qualitatively, the reservation-based system is able to sustain a much higher throughput than any of the traffic signal systems before causing the system to become full. Secondly, and more precisely, we compare the average and maximum delays when using the different intersection control policies.

6.1 Overpass: The Lower Bound

In our simulator, the overpass is as ideal as possible — it never requires any vehicle to slow down at the intersection. Furthermore, since all vehicles travel at the speed limit, vehicles traveling the same direction are never required to slow down for each other. For this reason, every vehicle using the overpass system experiences 0 delay. We ran many trials with the overpass system again to verify that our simulator worked as expected. By varying the top speeds of vehicles, we can produce delays with the overpass system, but in real-world situations, vehicles’ top speeds are much higher than the speed limit.

6.2 Traffic Light

To evaluate the delay of the traffic light model, we set several different periods and ran a million steps of the simulator for an increasing car spawning probability — from .001 to .02. The graph in Figure 4 shows that for lighter traffic, a shorter period is strictly better. This is already known and is used daily in many cities, where late at night the period of traffic lights is reduced. As the traffic increases, higher periods are more efficient. Eventually, however, they cannot handle the traffic, and the delay begins to increase very rapidly. In Figure 4, the delay levels off because at this point the cars are so backed up that the simulation can’t create any more of them.

An interesting thing to note is that even with only one tile, the reservation-based system does not break down until reaching a much higher level of traffic. An overloaded reservation system is chaotic: exactly when and how it breaks down varies wildly with when cars are spawned and which directions they are traveling. For that reason, running the simulation for a fixed number of steps can produce very different amounts of total delay. This explains why the line in Figure 4 representing the single-tile reservation system becomes so jagged.

6.3 Reservation System

To demonstrate the drastic improvement in throughput with the reservation system, we ran the same traffic amounts as in the traffic signal case, but this time used reservation systems with granularity one and two. In Figure 5, we show just how much of an improvement we can obtain over the single-tile case by increasing the granularity to just two. The amount of traffic handled by this system is much higher, and the associated delays are much lower. Note that with a granularity of two, it is permissible to have more than one car in the intersection at the same time: just not in the same quadrant.




Fig. 4: Average delays for traffic light systems with period 10, 30, and 50 seconds plotted against varying traffic levels along with a 1-tiled reservation-based system.




Fig. 5: The average delay for granularity 1 and 2 reservation systems with varying traffic levels. Each data point represents 1,000,000 steps of simulation.

After determining that the reservation system could manage as much traffic as a traffic light, we experimented along two main dimensions. First, we investigated the ability of our system to scale up to wider roads (i.e. multiple lanes going in each direction). Second, we studied the impact of the main parameter in our approach, namely the granularity used on the part of the intersection when assigning reservations. In particular, we characterized the tradeoff in terms of computational complexity vs. performance level.

6.3.1. More lanes One question that arises is whether or not this system scales up to a larger number of lanes. We tested it up through a 6×6 intersection (6 lanes travelling in each direction), and it performed well in all the cases, averaging less than one half of one percent the delay with the traffic light system. We present these results in Figure 6. Recall that the overpass system, representing the lower bound, leads to 0 delay.

Varying granularity represents another tradeoff - computation complexity and fault-tolerance versus efficiency. With fewer

Lanes	Reservation		Light	
	Avg	Max	Avg	Max
1	0.016	0.912	5.847	15.526
2	0.017	0.925	5.488	15.536
3	0.023	1.435	5.482	15.506
4	0.019	1.590	5.351	15.536
5	0.031	1.902	5.439	15.506
6	0.025	1.926	5.378	15.517

reservation tiles, the intersection is more exclusive. With a traffic light system along with numbers from a more exclusive. With a traffic light system along with numbers from a more exclusive. With a traffic light system along with numbers from a more exclusive.

Fig. 6: Statistics for different numbers of lanes and granularities of the reservation-based system along with numbers from a traffic light system with period 20 seconds. The simulation was run for 1,000,000 steps with a car-spawning probability of .001.

only 1 tile, for example, only one car may be in the intersection at a time, however with an extremely large number of tiles, the intersection reserves only enough space for the car, leaving little room for mistakes. As we will show, having only one tile can also lead to deadlock problems, whereas having a large number of tiles gives the intersection more flexibility with which to guide the cars. As for computation complexity, both the amount of calculation needed to process a single reservation and the space needed to hold the reservation grid increase as the square of the granularity.

One problem we discovered occurs whenever cars travelling in the inside lanes, but in opposite directions are competing for the same tiles. This happens usually when the granularity is an odd

Gran.	Delay		
	Avg.	Max	Std Dev.
2	0.109	1.698	0.239
3	0.131	2.750	0.274
4	0.043	1.026	0.134
5	0.071	1.812	0.185

Fig. 7: Simulation statistics for a reservation system with varying granularity. There are 2 lanes in each direction and the car-spawning probability is .001.

number. If one car has to slow down because it can’t get a reservation, when it finally does get a reservation it will occupy those tiles for a longer period of time. This makes it more likely that the next car coming in the opposite direction will have to slow down even more. This process eventually slows the cars down more and more. For small to average amounts of traffic, this causes larger delays. For very heavy traffic, it will eventually deadlock the intersection. As shown in Figure 7, with 2 lanes going in each direction it is better to have a 2×2 grid rather than a 3×3 grid. Increasing to a 4×4 grid is better than 2×2 ,

but increasing it to 5×5 is again worse. The deadlocking effect is hard to measure quantitatively, because if the cars start to deadlock, they begin to make reservations for very long periods of time — so long, in fact, that they overflow the memory of the computer system.

Because of this last problem, the reservation system should always be run with a granularity at least as high as the number of lanes, and ideally a multiple of the number of lanes. To test the effect of higher granularities, we ran the simulator with varying numbers of lanes and granularities. As Figure 8 demonstrates, more lanes require a higher granularity (though even with low granularity, this system out-performs the traffic light system). However, once the granularity is greater than the number of lanes travelling in each direction, subsequent increases in granularity provide only a marginal benefit, if any. Figure 9 shows a typical moment during a simulation with 6 lanes going in each direction and a granularity of 48.

Lanes	Granularity				
	1	2	6	12	48
1	0.016	0.005	0.007	0.005	0.007
2	0.059	0.017	0.010	0.008	0.003
3	0.142	0.038	0.009	0.011	0.013
6		0.132	0.025	0.011	0.006

Fig. 8. Average delays for the reservation system with independently varying numbers of lanes and granularity. All simulations run for at least 500,000 steps. 6 lanes with 1 tile breaks down and overflows the system memory before 500,000 steps can complete.

7. RELATED WORK

Most study of intersection management has been from a theoretical perspective, reducing the problem to scheduling jobs that may be competing for mutually exclusive resources. Irani and Leung create what they call “conflict graphs” in which nodes competing for the same resources share an edge. They then search these graphs for independent sets (Irani and Leung, 1996). Later, they present a probabilistic approach to the problem, and report actual simulated results using this approach (Irani and Leung, 1997). Their approach, however does not aim to improve the underlying system of traffic signals, but rather to tune them more effectively.

As far as multiagent approaches, most prior work has focused on increasing the throughput and decreasing delays for traditional traffic light systems. For example, Roozmond allows intersections to act autonomously, sharing the data they gather (Roozmond, 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 — the agents in this case are only the intersection controllers themselves.




Fig. 9: A screenshot of with 6 lanes in each direction and a granularity of 48. Note that there are cars travelling all four directions in the intersection simultaneously.

Coming from the other direction, agent-based controllers have been designed for mechatronic systems, which are defined as “complex technical systems whose motion behaviour is actively controlled with the help of computer technology” (Oberschelp *et al.*, 2002). They point to intersection management as one application of such controllers.

Finally, Kolodko and Vlacic describe a primitive system for intersection control which is very similar to the reservation system with granularity 1 (Kolodko and Vlacic, 2003). This system has been successfully implemented on real autonomous vehicles.

8. CONCLUSION

This paper makes two main contributions. First, it describes and specifies a custom intersection simulator for modeling intersection control policies along with precise metrics for measuring the efficiency of these policies. Second, it proposes a reservation-based policy that dramatically out-performs a traffic light policy and approaches the theoretical optimum in simulation. This policy is fully implemented and tested in simulation.

Current limitations include the inability of vehicles to turn and the constraint that vehicles may not change their velocity while in the intersection. Relaxing these limitations is a part of our on-going research agenda. Once autonomous vehicles are common, this mechanism may be useful for managing real traffic.

REFERENCES

- Hakim, Danny (2004). Robo-cars make cruise control so last century. *The New York Times*. April 4th.
- Irani and Leung (1996). Scheduling with conflicts, and applications to traffic signal control. In: *SODA: ACM-SIAM Symposium on Discrete Algorithms (A Conference on Theoretical and Experimental Analysis of Discrete Algorithms)*.
- Irani and Leung (1997). Probabilistic analysis for scheduling with conflicts. In: *SODA: ACM-SIAM Symposium on Discrete Algorithms (A Conference on Theoretical and Experimental Analysis of Discrete Algorithms)*.
- Kolodko, Julian and Ljubo Vlacic (2003). Cooperative autonomous driving at the intelligent control systems laboratory. *IEEE Intelligent Systems* **18**(4), 8–11.
- Oberschelp, Oliver, Thorsten Hestermeyer, Bernd Kleinjohann and Lisa Kleinjohann (2002). Design of self-optimizing agent-based controllers. In: *Proceedings of the 3rd International Workshop on Agent-Based Simulation*.
- Pormerleau, Dean A. (1993). *Neural Network Perception for Mobile Robot Guidance*. Kluwer Academic Publishers.
- Rogers, Seth, Claude-Nicolas Flechter and Pat Langley (1999). An adaptive interactive agent for route advice. In: *Proceedings of the Third International Conference on Autonomous Agents (Agents’99)* (Oren Etzioni, Jörg P. Müller and Jeffrey M. Bradshaw, Eds.). ACM Press. Seattle, WA, USA. pp. 198–205.
- Roozmond, Danko A. (1999). Using intelligent agents for urban traffic control systems. In: *Proceedings of the International Conference on Artificial Intelligence in Transportation Systems and Science*. pp. 69–79.
- Schonberg, T., M. Ojala, J. Suomela, A. Torpo and A. Halme (1995). Positioning an autonomous off-road vehicle by using fused DGPS and inertial navigation. In: *2nd IFAC Conference on Intelligent Autonomous Vehicles*. pp. 226–231.
- Stone, Peter and Manuela Veloso (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots* **8**(3), 345–383.