

Turning the Corner: Improved Intersection Control for Autonomous Vehicles

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 is one of the leading causes of lost productivity and decreased standard of living in urban settings. Recent advances in artificial intelligence suggest vehicle navigation by autonomous agents will be possible in the near future. In a previous paper, we proposed a reservation-based system for alleviating traffic congestion, specifically at intersections. This paper extends our prototype implementation in several ways with the aim of making it more implementable in the real world. In particular, we 1) add the ability of vehicles to turn, 2) enable them to accelerate while in the intersection, 3) give a better sensor model and communication-efficient heuristic to our driver agent, and 4) augment their interaction capabilities with a detailed protocol such that the vehicles do not need to know anything about the intersection control policy. The use of this protocol limits the interaction of the driver agent and the intersection manager to the extent that it is a reasonable approximation of reliable wireless communication. We then use this protocol to implement a new control policy: the stop sign. All three improvements are fully implemented and tested, and we present detailed empirical results validating their effectiveness.

I. INTRODUCTION

Traffic congestion is one of the leading causes of lost productivity and decreased standard of living in urban settings. According to a recent study of 85 U.S. cities [1], 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. 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 [2], and autonomous steering [3]. Once individual cars become autonomous, it is inevitable that before long 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 [4]. In an earlier paper, we proposed an MAS-based approach to alleviating traffic congestion, specifically at intersections [5]. In this paper, we describe several ways in which we have transformed that

system into a more realistic and implementable system.

Current methods for enabling traffic to flow through intersections include building overpasses and installing traffic lights. However, the former is very expensive and forbids turning, while the latter can be quite inefficient, often requiring cars to remain stopped even when no cars are present on the intersecting road.

At this time, it is possible to create a small-scale system in which all cars are piloted by a central computer. Consider, for example, the task of controlling ten vehicles on an open factory floor. However, growing such a system to handle an intersection in which a city's worth of cars might turn up would involve prohibitively expensive and inefficient communication and control infrastructure. 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 simple wireless communication system and a protocol (which we introduce here) for communicating with oncoming traffic and giving permission for cars to pass.

In our system, vehicles must traverse intersections according to a set of parameters agreed upon by the vehicle and the intersection manager (as they do today by obeying red and green lights), but otherwise are free to decide for themselves how to drive. Each car is an autonomous agent, and in particular need not surrender control to any centralized decision maker.

Given the above assumptions, we have proposed a novel reservation-based system by which cars request and receive time slots from the intersection during which they may pass [5]. While this system showed the potential for a reservation-based system to drastically improve the efficiency of intersections, it required driving agents to maintain a constant velocity in the intersection and forbade turning (a very important part of intersections). Furthermore, it did not adequately specify how they should interact. In this paper, we take three large steps towards making the system implementable in the real world. First, we augment it to allow turning. Second, we make acceleration in the intersection possible, which allows us to subsume the stop sign policy within the reservation framework. Third, we devise a protocol to govern the interactions of the vehicles and the intersection such that the vehicles do not need to know anything about the intersection control policy. The use of this protocol limits the interaction of the driver agent and the intersection manager

to the extent that it is a reasonable approximation of reliable wireless communication.

II. THE ORIGINAL SYSTEM

In our previous paper, we proposed a novel reservation-based multi-agent approach to alleviating traffic, specifically at intersections. This system consisted of two types of agents: *intersection managers* and *driver agents*. Each system consists of an intersection manager for each intersection and a driver agent for each vehicle. 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. Depending on the decision the intersection manager makes, the driver agent either records the parameters of the request (the *reservation*) and attempts to meet them, or it makes another request at a later time.

To determine whether or not a request can be met, the reservation manager simulates the journey of the vehicle across the intersection, which it divides into a grid of $n \times n$ tiles. The parameter n is called the *granularity* of the reservation manager. At each time step of the simulation, it determines which tiles the vehicle occupies. If throughout this simulation, no required tile is occupied by another vehicle (from a previous reservation), the manager reserves the tiles for this vehicle.

After creating a custom simulator, we evaluated the performance of the reservation system against two other *intersection control policies* - the overpass and the traffic light. An intersection control policy is a method the intersection managers use to determine when specific vehicles are allowed in the intersection. Using the simulator, we showed that using the reservation-based policy, vehicles crossing an intersection experience much lower *delay* (increase in travel time from the optimal) versus the traffic light. Furthermore, we showed that the reservation-based policy drastically increases the throughput of the intersection. For any realistic intersection control policy, there exists an amount of traffic above which vehicles arrive at the intersection more frequently than they can go through the intersection. At this point, the average delay experienced by vehicles travelling through the intersection grows without bound. We demonstrated that compared to the traffic light, this amount of traffic is much higher for the reservation system. In addition to our original simulator applets¹ Garcia and Vidal have implemented applets reproducing the results².

III. IMPROVING THE ORIGINAL MODEL

The results described in the previous section are very encouraging. In this section, we offer several ways to improve the system with regard to flexibility, efficiency, and making it implementable in the real world.

¹<http://www.cs.utexas.edu/users/kdresner/2004aamas>

²<http://jmvidal.cse.sc.edu/netlogomas/TrafficManagementMendoza.html>

A. Desirable Properties

In order for the reservation-based mechanism to be both realistic and practical, we believe that the following properties ought to hold.

- 1) The agents should only communicate information which is necessary for the system to function properly.
- 2) The agents should only have access to information that can be reliably obtained with current technology.
- 3) Communication failure (dropped messages) should not violate the system’s safety properties.
- 4) The vehicles should be treated as individual agents, and no centralized controller should have any more control over them than necessary.
- 5) The system should incorporate a simple communication protocol that allows agents to know only a minimal amount about each other. As long as agents obey and understand the protocol, no extra information exchange or other interaction should be required.
- 6) Every vehicle should eventually make it through the intersection (i.e. no deadlocks or starvation).

B. Acceleration in the Intersection

Our original reservation system made reservations for vehicles only at a constant velocity. This property is partly responsible (along with others discussed in Section IV) for the deadlocks their system experienced. With this restriction, if a vehicle made a reservation at a low velocity, it would consume a large amount of space-time in the intersection. This, in turn, would cause other vehicles to be delayed making their reservations (which would also be at low velocities). These slow-downs often led to permanent deadlocks. By allowing acceleration in the intersection, our system always eventually recovers from slowdowns caused by heavy traffic.

Because the reservation manager can now return reservations with accelerations, the problem becomes determining what those accelerations should be. By varying its accelerations just right, a vehicle may be able to fit through a small opening in the intersection. Somehow, the intersection manager must choose the correct accelerations. We chose to use a very simple heuristic: the intersection manager first tries to have the entering vehicle accelerate to the maximum allowed velocity. If such a reservation is not possible, it attempts to make a constant-velocity reservation. If the constant-velocity reservation also fails, the request is rejected. Using acceleration in the intersection, along with our protocol, allows us to implement the stop sign policy within this reservation framework.

C. Excess Information

knew each others’ positions and reservation statuses at all times. However, it is not immediately obvious how any vehicle would get this information in the real world. While exact position information would be hard to come by, there is no reason to believe that vehicles would have any access at all to the internal state of other vehicles around it (even ones in close proximity). An older model vehicle interacting with

a new model vehicle can not be expected to understand the newer model's inner workings. Additionally, the manufacturer of the driver agent may not want other vehicles to know what goes on "under the hood."

D. Unspecified Communication Between Driver Agents and Intersection Managers

In our previous paper [5], we specify which agents govern which aspects of their system, but we do not specify exactly *how* the agents coordinate their efforts. Additionally, in our original system, any driver agent would have to understand what kind of intersection control policy the intersection manager was using in order to interact with it. To address these issues, we created a detailed communication protocol to govern and restrict the interactions of driver agents and intersection managers. The complete specification of the protocol as well as how to use it to implement many intersection control policies can be found in our tech report [6].

This protocol solved three problems at once. First, all information between the agents goes through one monitorable channel, which makes it much easier to reason about. Second, by limiting the interactions of the agents to a few message types, we can ensure that no agent has an unrealistic amount of control over another. Third, the agents now have a way to communicate that is identical for any intersection management policy or driver agent policy. A vehicle can cross an intersection using a traffic light without knowing it is a traffic light. The traffic light speaks the same language as a stop sign and a reservation system. The driver agent thus must have a behavior that works with all sorts of intersection control policies — that is, the driver agent must view the intersection as a black box, and vice versa.

IV. NEW DRIVER AGENT

Our protocol is designed to place minimal restrictions on vehicle control. As a result, there remains a lot of freedom in creating driver agents. Though our system does not depend on any specific driver agent implementation, we need at least one concrete instantiation in order to test it empirically. In this section we discuss our extensions to the driver agent we created to work with our original system [5].

In that system, once a driver agent made a successful reservation (at its current velocity), it was forced to maintain that velocity until it reached the intersection. This is a major weakness for the system. If vehicles ever made reservations at very low velocities, not only did they consume a lot of valuable space-time in the intersection, but they also slowed down traffic behind them the rest of the way to the intersection. Repeated iterations of this scenario eventually contribute to deadlocking the system. The other part of this problem (that vehicles cannot accelerate while in the intersection) is addressed via the protocol [6].

A. Optimism and Pessimism

The new driver agent does not calculate its reservation times using only its current velocity. Previously, our driver agent

always made requests by calculating the time to get to the intersection at its current velocity, after which, it maintained that velocity until it was through the intersection. It does not matter *how* the vehicle reaches the intersection, as long as the vehicle arrives as scheduled. The behavior as originally proposed can lead to serious problems when, for example, a vehicle makes a reservation while stuck behind a slower-moving vehicle. If the vehicle in front eventually accelerates, the other vehicle should be able to accelerate as well (possibly switching to an earlier reservation).

To utilize this flexibility, we introduce the notion of an *optimistic* or *pessimistic* driver agent. An optimistic agent makes a reservation assuming it will immediately get to accelerate to full speed. An agent which no longer finds itself stuck behind a slower vehicle will become optimistic and attempt to make a new, earlier reservation. A pessimistic agent assumes it will be stuck at its current velocity until it reaches the intersection. If an agent has to cancel its reservation because there is no way for it to arrive on time, it becomes pessimistic. Due to the relatively infrequent and smooth transitions through these situations, our driver agent can take advantage of improving circumstances without causing it to send excessive numbers of CHANGE-REQUEST messages when things change.

B. Cancellation and Communication Complexity

Another change, very closely related to the previous section, is an improvement in the communication complexity of the model. In the initial model, the agent determined whether or not it could honor a reservation assuming it kept its present velocity for the remainder of the journey to the intersection. While this might keep things more up-to-date, it often caused a decelerating agent to make and cancel new reservations in rapid succession until it stopped decelerating. In order to prevent this, the new agent only cancels a reservation if there is absolutely no physical way it could reach the intersection on time. If a person were a few minutes late in leaving for the airport, that person would not immediately cancel his or her flight entirely. On the contrary, that person would hope to make up lost time at some point before the flight left. Only when there was no hope of making it to the jetway on time would the person actually cancel the reservation.

Reducing the communication complexity of the system is very important for two reasons. First, if fewer total messages are sent, the bandwidth required to send messages is lower; thus, given the available bandwidth, messages are much less likely to be delayed or lost — events which might negatively affect the system's efficiency. Second, many of the messages result in intense computation by the intersection manager. Because the resources of the intersection manager are limited, it can only process these messages at some fixed rate. In order to regulate the driver agents, we envision that some sort of charge (perhaps a micropayment) will be levied for each message. In this case, reducing the number of messages sent will be a priority for driver agents.

C. Sensor Model

In the intersection, the intersection manager is responsible for preventing collisions, but outside the intersection the driver agents are entirely autonomous and must do this themselves. Furthermore, a vehicle approaching an intersection should not be expected to know the exact locations of all other vehicles. Here, we describe a much more limited sensor that is nonetheless sufficient for our purpose.

Modern laser range finders and distance sensors can provide a large amount of distance and angle data to a mobile agent, however, not all this information is entirely relevant. In a real life setting, this information would definitely prove useful in fine-tuning a driver agent. However, in a simple simulation, we must process sensor information for *all* vehicles simultaneously. Thus, we wanted to find a simple, yet pertinent sensor reading which the driver agent could use to control its actions with respect to the other vehicles. A purely straight-ahead sensor suffices when vehicles are travelling only in straight lines. However, when a vehicle turns, it must also take into account what is going on in the direction it is turning. To complicate matters, when a vehicle is turning it must still take into account what is going on directly in front of it because at any point it might straighten out its wheels. Initially, we experimented with a sensor array that points in the same direction as the wheels. After determining that this was not sufficient (for example, vehicles coming out of turns would run into vehicles ahead of them), we settled on a sensor array whose scope widens in the direction of the turn, while narrowing slightly from the other side. Figure 1 shows some images demonstrating this concept. A testament to the sensor's efficacy, vehicles equipped only with the sensor (i.e. no intersection manager was present) were able to avoid many collisions in the intersection, even with moderate amounts of traffic.




Fig. 1. A depiction of the sensor model for the driver agents. Only information inside the gray lines is provided to the driver agent.

V. EMPIRICAL RESULTS

In this section, we evaluate the performance of our improved reservation system for varying amounts of traffic and varying percentages of turning vehicles. Additionally, we show results for the new stop sign control policy as implemented under our protocol. We then compare these to results from an earlier paper regarding standard traffic lights. Finally, we experiment

with allowing vehicles to turn from any lane — something that would be extremely dangerous without the reservation-based mechanism.

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. Figure 2 shows a screenshot of the graphical display. Each time step in the simulator represents .02 seconds of real time. During each time step, a vehicle is spawned with the given probability, each driver is given sensor input and a decision-making phase, the positions of each vehicle are updated based on the decisions of the driver, and finally any vehicles that have left the area of the simulation are removed. Every configuration shown is run for 100,000 steps in the simulator, which corresponds to approximately half an hour. Vehicles that are spawned in any given direction turn both right and left with probability .05. Unless otherwise specified, vehicles turning right are spawned in the right lane, whereas vehicles turning left are spawned in the left lane. Vehicles that are not turning are distributed probabilistically amongst the lanes such that the traffic in each lane is as equal as possible. The reservation system in these simulations has a granularity of 24 and ensures that no two vehicles occupy the same tile within half a second of each other. Videos of the simulator running can be seen at <http://www.cs.utexas.edu/users/kdresner/papers/2005iv/>.

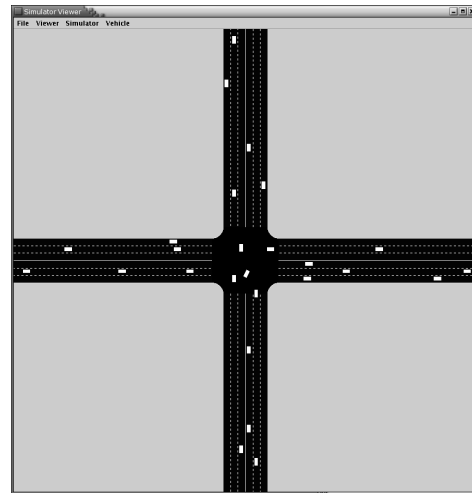


Fig. 2. A screenshot of our simulator in action.

Once turns are allowed, delay does not work very well as a metric. There are many different paths through the intersection and amongst them are several different total distances. In addition, vehicles that are turning must slow down before making their turns, so they may take longer than the minimum time to go through the intersection, even under optimal conditions. Because of this, we have decided to simply measure the average time it takes a vehicle to go from a fixed start point to a fixed destination point. We refer to this time as the *trip time*.

Note that in our last paper [5], the traffic light was shown to have trip times of at least 5 seconds longer than optimal, even in scenarios with extremely light traffic. The absolute shortest time to go from start to finish in this scenario is 10 seconds, which means that the average trip time for the traffic light would be at least 15 seconds.

A. The Overpass

Previously, we presented the overpass as the optimal solution to the intersection control problem. With the addition of turns, a traditional overpass does not make sense. However, we would like an ideal-case solution in which cross-traffic does not affect the time it takes a vehicle to complete its journey. Thus, although it does not represent a true overpass, we still refer to this solution as “the overpass.” Vehicles are granted reservations at any time and they can pass through one another. The only thing that might hinder a vehicle is slowing for a turn.

Although a lower bound on the trip time of a vehicle is 10 seconds, turning vehicles must slow to make the turn. Thus the average time for the overpass system as shown in Figure 3 is just above 10 seconds.




Fig. 3. Trip times for varying amounts of traffic for three intersection control policies.

B. The Reservation System

The reservation system performs very well, nearly matching the performance of the overpass system. At higher levels of traffic, the average trip time for a vehicle gets as high as 11 seconds, but is never more than 1 second above optimal. Under none of the tested conditions does the reservation system approach the trip times of the original traffic light system.

C. The Stop Sign

Small intersections with slow-moving traffic tend not to be amenable to control by traffic lights. Light traffic can usually regulate itself fairly effectively. For example, consider an intersection with a stop sign - all vehicles must come to a stop, but afterwards may proceed if the intersection is clear. In these situations, a stop sign is often much more efficient than a traffic light, because vehicles are never stuck waiting for a light to change when there is no cross-traffic. Because our new protocol enables us to define such a control policy,

we test how it compares to the other systems as well. Note that this system is much more efficient than an actual stop sign, because once the vehicle has stopped at the intersection, the driver agent and intersection can determine when the car may safely proceed more precisely than a human driver. As shown in Figure 3, the stop sign does not perform as well as the reservation system or the overpass, but for low amounts of traffic, it still performs fairly well, with average trip times only about 3 seconds greater than optimal. As the traffic level increases, however, performance degrades.

D. Allowing Turns from Any Lane

In traditional traffic systems, especially those with traffic lights, vehicles wishing to turn onto the cross street must do so from specially designated turning lanes. This helps prevent cars that want to turn from holding up non-turning traffic. However, with a system like the reservation system, this restriction is no longer necessary. There is nothing inherent in the reservation system that demands vehicles turn from any specific lane, and thus we investigated these effects³. As seen in Figure 4, relaxing this restriction in fact worsens performance. While one might think this allows the vehicles more flexibility, it on average increases the resources used by any one turning vehicle. By making left turns from the left lane and right turns from the right lane, vehicles both travel a shorter distance and use reservation tiles that are less heavily used.




Fig. 4. Comparison of the normal reservation system with turns to one allowing turning from any lane. The system that allows turns from any lane performs marginally worse.

E. Changes to the Driver Agent

As shown in Figure 5, the improvements to the driver agent drastically reduced both the average number of reservations made as well as the average number of messages transmitted. These data were collected using the same simulator settings as the rest of this section, but with a vehicle spawning probability of .02 (approximately 2000 vehicles). For lower amounts of traffic, the effect was less pronounced.

³Videos of this can be seen at <http://www.cppreference.com/aamas05/>.

	Messages	Reservations
Before	560.85	165.89
After	5.97	1.02

Fig. 5. For a moderate amount of traffic, the average number of messages sent and reservations made by driver agents before and after the improvements described in Section IV.

VI. DISCUSSION AND RELATED WORK

We have shown that the reservation system as proposed in our last paper [5] can be extended naturally to incorporate turning and accelerating in the intersection. Furthermore, we have shown that the reservation system can outperform the stop sign, approaching optimal, at a wide range of traffic densities. Our communication protocol, which allows the system to subsume both the stop sign and the traffic light, solves many of the major concerns we had noted.

One of these concerns was allowing the system to work with human drivers, pedestrians, or cyclists. One can imagine a system that shifts to a traffic-light-like control policy (with physical lights) when it detects vehicles or pedestrians that cannot participate in the reservation system. These individuals could then interact with the intersection the way they do currently. Once the traffic consisted only of participating vehicles, the intersection manager could switch back to a more efficient reservation-based policy.

A. Future Work

There are still many challenges and interesting questions to be answered in this domain. For example, we investigated the effects of allowing the vehicle to turn from any lane, but we did not investigate what happens when vehicles are allowed to turn *into* any lane. Furthermore, with the creation of a communication protocol, we can create more interesting driver agents and intersection managers. Both could involve machine learning. The inherent multi-agent nature of the domain makes it a good testbed for multi-agent learning research. The agents can be heterogeneous, and the different types of agents (intersection managers and drivers) have different, but not necessarily opposing, goals.

We also see a large opportunity for more research in designing more intelligent reservation systems and driver agents. Currently both of these use heuristics to find available reservations and reservation times, respectively. Applying machine learning techniques to these issues could increase the efficiency of the system even further.

B. Related Work

Rasche and Naumann have worked extensively on decentralized solutions to intersection collision avoidance problems [7], [8]. Many approaches focus on improving current technology (systems of traffic lights). For example, Roozmond allows intersections to act autonomously, sharing the data they gather [9]. 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.

Work is also being done with regard to the control of individual vehicles. 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 [10].

VII. CONCLUSION

This paper makes four main contributions. First, it augments a proposed intersection control mechanism to allow for more flexible vehicle control, including turning and accelerating while in the intersection. Second, it introduces a detailed protocol by which vehicles and intersection managers can effectively and efficiently communicate and coordinate their actions. Third, it describes a driver agent that makes good use of this protocol. Finally, it demonstrates how this augmented system, using the protocol, can still drastically outperform both the traffic light (as demonstrated in previous work [5]) and the stop sign.

The mechanism is currently limited by the use of straightforward heuristics to calculate reservation parameters, both on the part of the intersection manager and the driver agents. However, this limitation is a focus of our ongoing research. Once autonomous vehicles become common, this mechanism may be useful for controlling real traffic.

REFERENCES

- [1] Texas Transportation Institute, "2004 urban mobility report," September 2004, accessed at <http://mobility.tamu.edu/ums> in December 2004.
- [2] S. Rogers, C.-N. Flechter, and P. Langley, "An adaptive interactive agent for route advice," in *Proceedings of the Third International Conference on Autonomous Agents (Agents '99)*, O. Etzioni, J. P. Müller, and J. M. Bradshaw, Eds. Seattle, WA, USA: ACM Press, 1999, pp. 198–205. [Online]. Available: citeseer.nj.nec.com/rogers95adaptive.html
- [3] C. W. Reynolds, "Steering behaviors for autonomous characters," in *Proceedings of the Game Developers Conference*, 1999, pp. 763–782.
- [4] P. Stone and M. Veloso, "Multiagent systems: A survey from a machine learning perspective," *Autonomous Robots*, vol. 8, no. 3, pp. 345–383, July 2000.
- [5] 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*, New York, New York, USA, July 2004, pp. 530–537.
- [6] —, "Multiagent traffic management: A protocol for defining intersection control policies," The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, Tech. Rep. UT-AI-TR-04-315, December 2004.
- [7] 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.
- [8] R. Rasche, R. Naumann, J. Tacke, and C. Tahedl, "Validation and simulation of decentralized intersection collision avoidance algorithm," in *Proceedings of IEEE Conference on Intelligent Transportation Systems (ITSC 97)*, 1997.
- [9] D. A. Roozmond, "Using intelligent agents for urban traffic control control systems," in *Proceedings of the International Conference on Artificial Intelligence in Transportation Systems and Science*, 1999, pp. 69–79. [Online]. Available: citeseer.nj.nec.com/roozmond99using.html
- [10] D. Moriarty and P. Langley, "Learning cooperative lane selection strategies for highways," in *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. Madison, WI: AAAI Press, 1998, pp. 684–691.