The UT Austin Villa 2008 Mixed Reality Team

Peter Stone and Neda Shahidi

Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712-0233
{pstone}@cs.utexas.edu
http://www.cs.utexas.edu/~AustinVilla

Abstract. This paper describes the research focus and ideas incorporated in the UT Austin Villa Simulation League Mixed Reality team entering the RoboCup competitions in 2008.

1 Introduction

This paper describes the research focus and ideas incorporated in the UT Austin Villa Simulation League Mixed Reality team entering the RoboCup competitions in 2008, including:

- Development plans; and
- Proposed research

2 Development Plans

In 2008, our development focus will be on robust and efficient vision for tracking up to 20 robots and on a general multi-robot control platform in support of our proposed research plan as detailed in Section 3.

3 Proposed Research

Our main plan for the EcoBe robots is to use them to support our research on autonomous traffic management [1–3].

Enabling cars to drive autonomously in cities is currently technologically feasible, and will likely be economically feasible within the next 5–10 years. Indeed, The DARPA Urban Challenge was completed successfully by multiple vehicles in November, 2007, and General Motors has announced that it plans to release a nearly autonomous vehicle under its European "Opel" brand. The 2008 Opel Vectra will be able to drive itself at speeds up to 60 miles per hour, even in heavy traffic.

Such autonomous vehicles will change the way we think about transportation, for example enabling people to concentrate on other activities while "driving," and enabling minors and the elderly to be transported on their own. As a result, once there's a single autonomous vehicle, there will likely be many more.

Our research considers the impact of autonomous vehicles on urban traffic infrastructure, specifically at intersections. In particular, it proposes a novel alternative to traffic signals and stop signs that enables cars to coordinate at a much finer granularity while still maintaining robust safety properties. Preliminary results indicate that our method may be able to reduce the time spent waiting at intersections by up to two orders of magnitude.

The key technical idea behind this research is that, taking advantage of their precise controllability, autonomous vehicles will be able to coordinate at intersections by each reserving just the amount of space-time needed within the intersection to enable safe passage. To our knowledge, our approach to this important future problem is both novel and unique.

The central hypothesis of our research is that our proposed form of intersection control can dramatically decrease time wasted at intersections and increase vehicle throughput on roads. We will test this hypothesis in scenarios consisting of all autonomous vehicles and in scenarios mixing autonomous and human-driven vehicles. We will also investigate ways of coordinating multiple such intersections to maximize throughput.

The computer science challenges related to this research are numerous. They include algorithmic design (of intersection controllers and autonomous driver agents); multiagent systems, both pertaining to networks of intersections and vehicle-to-vehicle interactions; theory pertaining to liveness and safety of the proposed algorithms; and computer security for safeguarding against malicious attacks.

Our research promises to have a profound impact on transportation infrastructure by reducing wastefulness that results from traffic congestion. 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. Our research will enable dramatic reductions in this wastefulness via advances that will manifest themselves once autonomous vehicles are feasible.

Our research in this direction has been conducted so far entirely in simulation. We aim to use the *EcoBe* robots as the first physical manifestation of this research. For this purpose, we need to be able to be able to control about 20 robots precisely in a small area so that they can cycle back through an intersection continually. We can also take advantage of the mixed reality nature of this event by injecting simulated traffic into the system to interact with the physical robots.

3.1 Research Overview

Motivated by the preceding discussion, our research aims to answer the following question:

[1] `http://mobility.tamu.edu/ums`

To what extent and how can an autonomous intersection control mechanism take advantage of the capabilities of autonomous vehicles in order to make automobile travel safer and faster?

In order to answer this question, our project is organized according to the following five subgoals, each of which is a contribution of our research.

1. **Novel intersection control mechanism.** Today's intersection control mechanisms were designed to work with humans. This project investigates a solution designed from the ground up to take advantage of the special abilities of autonomous vehicles. The solution is based on a reservation paradigm, in which vehicles "call ahead" to reserve space-time in the intersection.
2. **Detailed protocol.** A multiagent system is defined by the interactions of its agents. Because this project aims to create a multiagent solution, it must specify how the agents will be expected to behave with respect to one another, including exactly how they will communicate. This takes the form of a detailed protocol specification, complete with message types and fields, an intended semantics, and interaction rules governing expected message responses.
3. **Control algorithms.** While agent interactions define a multiagent system, the behaviors of the agents themselves often most directly contribute to the performance of the system as a whole. The protocol makes certain guarantees about the system (e.g. safety, robustness under communication failure), but it also defines very large strategy spaces for the agents. A main technical contribution of this research is an extensive exploration of these spaces. Strategies for all agents in the system are being examined, including adaptive strategies.
4. **Mixed Reality.** To date, all of our research has been conducted in simulation. The EcoBe robots will be the first physical manifestation of this research. We propose to begin by implementing an all-EcoBe scenario with 20 robots continually cycling back through a single intersection in a crossing pattern. We then propose to inject simulated traffic into the system as well, true to the spirit of "mixed reality."
5. **Empirical evaluation.** Before a new mechanism can be considered for deployment in the real world, it must perform quantifiably better than both current methods and ideally all other possible solutions. This project provides detailed empirical results, in a variety of settings, including some in which human drivers are present, and including special attention to failure modes (e.g. what is the worst-case impact of a car getting a flat tire in the intersection).

3.2 Desiderata

Replacing modern intersection control with a robust, autonomous control framework is a complex, multi-part problem. In order to establish a set of criteria by which to judge such a framework, here we enumerate a set of desiderata, which simultaneously provides a description of the problem at hand as well as both qualitative and quantitative metrics for judging solutions to that problem.

Autonomy. Each vehicle should be fully autonomous. Were the entire mechanism centrally controlled, it would be susceptible to single point failure, require massive amounts of computational power, and exert unnecessary control over vehicles in situations where they are perfectly capable of controlling themselves.

Low Communication Complexity. By keeping the number of messages and amount of information transmitted to a minimum, the system can afford to put more communication reliability measures in place. Furthermore, each vehicle, as an autonomous agent, may have privacy concerns which should be respected. Keeping the communication complexity low will also make the system more scalable.

Sensor Model Realism. Each agent should have access only to sensors that are available with current-day technology. The mechanism should not rely on fictional sensor technology that may never materialize.

Communication Protocol Standardization. The mechanism should employ a simple, standardized protocol for communication between agents. Without a standardized protocol, each agent would need to understand the internal workings of every agent with which it interacted. An open, standardized protocol would make adoption of the system simpler for vehicle manufacturers.

Deadlock/Starvation Avoidance. Deadlocks and starvation should not occur in the system. That is, any vehicle approaching an intersection should eventually make it through, even if it is better for the rest of the agents to leave that vehicle stranded.

Incremental Deployability. The system should be incrementally deployable, in two senses. First, it should be possible to set up selected intersections to use the system, and then slowly expand to other intersections as needed. Second, the system should function even with few or no autonomous vehicles. At any stage of deployment, be it an increase in the proportion of autonomous vehicles or number of equipped intersections, overall performance of the system should improve, and there should be a benefit to early adopters. At no point should there exist a net disincentive to continue deploying the system.

Safety. Except for gross vehicle malfunction or extraordinary circumstances (natural disasters, etc.), as long as they follow the protocol, vehicles should never collide in the intersection. Note that no stronger guarantee is possible — as with modern mechanisms, a suicidal human driver can always steer a vehicle into oncoming traffic. Furthermore, the system should be safe in the event of total communication failure. If messages are dropped or corrupted, the safety of the system should not be compromised. It is impossible to prevent all negative effects due to communication failures, but those negative effects should be limited to efficiency. If a message gets dropped, it can make someone arrive 10 seconds later at their destination, but it should not cause a collision.

Efficiency. Vehicles should get across the intersection and on their way in as little time as possible. To quantify efficiency, we introduce *delay*, defined as the amount of additional travel time incurred by the vehicle as the result of passing through the intersection.

3.3 The Reservation Idea

Of the above desiderata, modern-day traffic lights and stop signs completely satisfy all but the last one. While many accidents take place at intersections governed by traffic lights, these accidents are rarely, if ever, the fault of the traffic light system itself, but rather that of the human drivers. However, traffic lights and stop signs are terribly inefficient. Not only do vehicles traversing intersections equipped with these mechanisms experience large delays, but the intersections themselves can only manage a limited traffic capacity — much less than that of the roads that feed into them. **The aim of this project is thus to create an intersection control mechanism that exceeds the efficiency of traffic lights and stop signs without sacrificing any of the other properties listed above.**

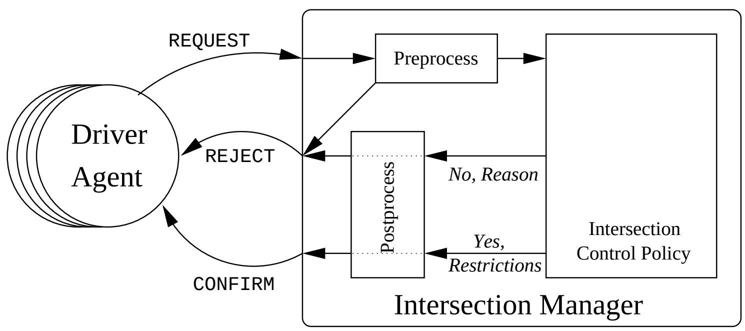

Fig. 1. One of the driver agents attempts to make a reservation. The intersection manager responds based on the decision of an intersection control policy.

With our desiderata in mind, we have introduced a novel approach to getting vehicles through intersections more efficiently that is a radical departure from existing traffic signal optimization schemes. In our approach, we assume that computer programs called *driver agents* control the vehicles, while an *arbiter* agent called an *intersection manager* is placed at each intersection. The driver agents "call ahead" and attempt to reserve a block of space-time in the intersection. The intersection manager decides whether to grant or reject requested reservations according to an *intersection control policy*. Figure 1 shows a sample interaction between a driver agent and an intersection manager.

Our prototype intersection control policy divides the intersection into an $n \times n$ grid of *reservation tiles*, where n is referred to as the *granularity* of the

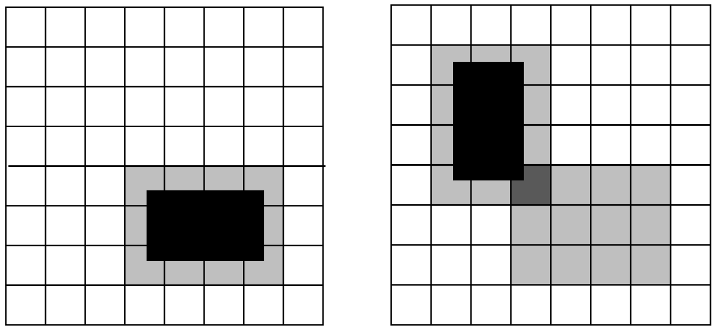

(a) Successful (b) Rejected

Fig. 2. In (a), vehicle A's request is accepted, and reserves a set of tiles at time t. In (b), vehicle B's request is rejected because it requires a tile already reserved by A at time t.

policy. When a vehicle approaches the intersection, it transmits a reservation request, which includes parameters such as time of arrival, velocity of arrival, and vehicle characteristics such as size and acceleration/deceleration capabilities, to the intersection manager. The intersection manager then passes this information to the policy, which simulates the journey of the vehicle across the intersection according to the parameters. At each time step of the simulation, the policy determines which reservation tiles will be occupied by the vehicle. If at any time during the simulation the requesting vehicle occupies a reservation tile that is already reserved by another vehicle, the policy rejects the driver's reservation request, and the intersection manager communicates this to the driver agent. Otherwise, the policy accepts the reservation and reserves the appropriate tiles. The intersection manager then sends a confirmation to the driver.

This control policy allows for much finer-grained coordination than traffic lights: cars traveling in all directions can proceed through the intersection with minimal delay. A graphical depiction of the concept can be seen in Figure 2. A video demonstration[2] shows cars coming from all directions barely slowing down as they traverse the intersection and crossing closely in front of each other, yet with no collisions. Preliminary experiments in simulation indicate that our proposed reservation system may be able to reduce the time spent waiting at intersections by up to two orders of magnitude.

To date, all of our intersection management research has been conducted in simulation [1–3]. The EcoBe robots will be the first physical manifestation of this research. We propose to begin by implementing an all-EcoBe scenario with 20 robots continually cycling back through a single intersection in a crossing pattern. We then propose to inject simulated traffic into the system as well, true to the spirit of "mixed reality."

[2] Available online at `http://www.cs.utexas.edu/~kdresner/aim/?p=video`

4 Conclusion

The UT Austin Villa team is eager to use the EcoBe robots as an integral part of the our research on autonomous intersection management and to participate in the RoboCup 2008 mixed reality competitions.

References

1. 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, July 2004.
2. Kurt Dresner and Peter Stone. Multiagent traffic management: An improved intersection control mechanism. In Frank Dignum, Virginia Dignum, Sven Koenig, Sarit Kraus, Munindar P. Singh, and Michael Wooldridge, editors, *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, New York, NY, July 2005. ACM Press.
3. Kurt Dresner and Peter Stone. Sharing the road: Autonomous vehicles meet human drivers. In *The 20th International Joint Conference on Artificial Intelligence*, pages 1263–68, January 2007.