AIM: Autonomous Intersection Management

Kurt Dresner
University of Texas at Austin
Department of Computer Sciences
http://www.cs.utexas.edu/~kdresner
kdresner@cs.utexas.edu

1. INTRODUCTION

Artificial intelligence research is ushering in a new era of sophisticated, mass-market transportation technology. While computers can already fly a passenger jet better than a trained human pilot, people are still faced with the dangerous yet tedious task of driving automobiles. Recent advances in artificial intelligence, multiagent systems (MAS), and Intellgient Transportation Systems (ITS) point to a future in which vehicles themselves handle the vast majority of the driving task. Once autonomous vehicles become popular, autonomous interactions amongst *multiple* vehicles will be possible. Current methods of vehicle coordination, which are all designed to work with human drivers, will be outdated. The bottleneck for roadway efficiency will no longer be the drivers, but rather the mechanism by which those drivers' actions are coordinated. While open-road driving is a well-studied and more-or-less-solved problem, urban traffic scenarios, especially intersections, are much more challenging.

This thesis addresses the question: **"To what extent and how can a multiagent intersection control mechanism take advantage of the capabilities of autonomous vehicles in order to make automobile travel safer and faster?"** First, I introduce and specify the problem of intersection management as a multiagent system and define a metric by which solutions can be evaluated. Next, I propose a novel multiagent intersection control mechanism in which autonomous *driver agents* "call ahead" and reserve space-time in the intersection, pending the approval of an *arbiter* agent called an *intersection manager*, which is located at the intersection. I also define a detailed protocol by which the driver agents and intersection managers communicate. The protocol, which consists of message types and associated semantics, defines a space of behaviors for the two types of agents. This space is extensively explored, and algorithms are presented for the agents that enable the system as a whole to achieve performance levels far beyond those of current intersection control mechanisms like stop signs and traffic lights. To demonstrate the superiority of the mechanism, I present empirical data from a custom simulator developed specifically for this work. Finally, I discuss the feasibility of implementation and deployment of such a system in the real world.

Cite as: Extended Thesis Abstract for Doctoral Mentoring Program, K. Dresner, Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008), Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1732-1733.

Though motivated by the specific problem of intersection management, these contributions generalize to the larger problem of coordinating multiple autonomous, heterogeneous, physical agents that desire access to the same exclusive physical space. By associating an arbiter agent with each resource, a scalable, distributed control mechanism is created that grants more autonomy and is more failure-resistant than a centralized system, while providing fine-grained control that creates a more efficient system than current, coarse-grained methods.

2. PROBLEM STATEMENT AND SOLUTION FRAMEWORK

A major contribution of this thesis is to define the problem of autonomous intersection management. Intersections are one of the most dangerous aspects of automobile traffic, host to a staggeringly disproportionate number of collisions and fatalities. Additionally, they are a bottleneck and a major source of delay and frustration to drivers. By exploiting the capabilities of autonomous vehicles, intersections can be made both more efficient and safer.

2.1 Desiderata

In order to choose directions in which to focus my research as well as establish a set of criteria by which to judge such a framework, I have enumerated an important set of properties I believe an intersection control mechanism for autonomous vehicles should have: autonomy for vehicles, low communication complexity, sensor model realism, protocol standardization, deadlock and starvation avoidance, incremental deployment, safety, and efficiency.

2.2 The Reservation Idea

Of the desiderata, modern-day traffic lights and stop signs completely satisfy all but the last one. However, I show that these mechanisms are terribly inefficient, when autonomous vehicles are involved. With the desiderata in mind, we developed a multiagent approach to get vehicles through the intersection more efficiently. In this approach, *driver agents* controlling vehicles "call ahead" to *interseciton managers* at the intersection 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*. One example of an intersection control policy is the "First Come, First Served" or FCFS policy, which works by dividing the intersection into an $n \times n$ grid of *reservation tiles*. When it receives a reservation request from a vehicle, FCFS simulates the trajectory of that vehicle according to the parameters in the request. By calculating which reservation tiles will be occupied at various points along the vehicle's journey, FCFS can determine whether or not the vehicle can cross safely, or whether any of the reservation tiles are already reserved by another vehicle. In this way, FCFS can orchestrate "close calls" in the intersection that allow more vehicles in the intersection simultaneously, moving faster, and thus increasing efficiency.

2.3 Communication Protocol

A major part of this thesis is a detailed communication protocol that governs the interactions of driver agents and intersection managers. This communication protocol consists of a set of messages and associated semantics that outline what is expected of each agent. The protocol allows for reasoning about the communication complexity of the system, as well as whether or not agents have incentives or disincentives to be honest in their communications.

2.4 Including Humans

In addition to intersection control policies for managing fully autonomous vehicles, I also present algorithms that allow humans to use the system in a backwards compatibility mode. The policy I introduce, FCFS-Light, combines the reservation tile technique of FCFS with a *light model* that allows the policy to reserve special corridors through the intersection that correspond to green lights. If a vehicle travels through a green light, it can know it is safe because its path through the intersection has been blocked off from use by autonomous vehicles. Areas that are not reserved in this way are allocated in the same way as with FCFS, thus allowing fully autonomous vehicles to go through red lights much in the same way that vehicles today can make a "right on red."

The challenge lies in properly trading off the efficiency of the autonomous vehicles with that of the human drivers. To accomplish this, I demonstrate several policies, each of which allocates a different fraction of the intersection to human vehicles, thereby leaving more or less for the autonomous vehicles. Because different traffic conditions require different intersection control policies, I also show how machine learning can be used to accurately choose the correct policy for the scenario.

3. RESULTS

I empirically evaluate all the features introduced using a custom-built simulator that is also a contribution of the thesis. In order to quantify efficiency, I introduce the metric of *delay*, defined as the amount of additional travel time incurred by the vehicle as the result of passing through the intersection. By comparing the average delay between various mechanisms or intersection control policies, I can establish the relative efficiencies of these policies.

3.1 Fully Autonomous Intersections

When only autonomous vehicles are involved, the improvement over current methods is staggering—delay decreases by as many as two or three orders of magnitude. Instead of spending tens of seconds waiting at intersections, vehicles are delayed by tenths of seconds. This means that vehicles get where they are going faster, using less fuel, and without requiring wider and wider roads.

3.2 Intersections With Humans

As expected, when including human drivers, efficiency is not as high as with only autonomous vehicles. However, I introduced several intersection control policies, each of which is targeted at different proportions of human drivers. By allowing the intersection manager to choose between them automatically and switch between them smoothly, the intersection manager can effect a gradual transition between the current all-human traffic and the almost exclusively autonomous traffic of the future. Furthermore, the system makes efficiency gains at every intermediate step—as the proportion of autonomous vehicles increases, average delay decreases.

3.3 Emergency Vehicles

In addition to extending the system to include humans, I have created a version that gives higher priority to emergency vehicles. The experimental results suggest that not only is this method effective—on average, emergency vehicles experience half the delay of other vehicles—but it also does not put an undue burden on other vehicles, as their average delay sees no appreciable increase.

3.4 Failure Mode Analysis

Before efficiency gains can be realized, a plethora of safety and reliability concerns must be addressed. I demonstrate how the reservation-based mechanism reacts to scenarios in which such malfunctions occur, and also intentionally disable some elements of the system in order to investigate both their necessity and efficacy. Ultimately, I show that even with a few simple safety features, experiments suggest the system will be much safer than current automobile traffic, reducing the number of vehicles involved in accidents by 80% or more.

4. CONCLUSION

My thesis makes five main contributions. First, it defines the problem of autonomous intersection management, including a set of desiderata by which potential solutions can be evaluated. Second, it presents a framework that can meet all of these desiderata, and an algorithm (FCFS) that shows the advantages of the framework over current intersection control methods. Third, it demonstrates how the framework can be extended to allow human-driven (not autonomous) vehicles to use the system, while still exploiting the abilities of the autonomous vehicles to increase throughput and subsequently decrease delays. Fourth, it presents a method by which an intersection manager can automatically choose among several candidate policies to control an intersection with lowest delays given the current traffic conditions. Finally, I provide a detailed failure mode analysis that suggests that in addition to the substantial efficiency benefits, automobile travel will be much safer when a system like mine is in place.