

Making Autonomous Intersection Management Backwards-Compatible*

Kurt Dresner and Peter Stone

Department of Computer Sciences

The University of Texas at Austin

1 University Station C0500

Austin, Texas 78712-0233

{kdresner, pstone}@cs.utexas.edu

<http://www.cs.utexas.edu/~{kdresner, pstone}>

Introduction

Traffic congestion and automobile accidents are two of the leading causes of decreased standard of living and lost productivity 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 adaptive cruise control, GPS-based route planning (Rogers, Flechter, & Langley 1999; Schonberg *et al.* 1995), and autonomous steering (Pomerleau 1993). Once individual cars become autonomous, many of the cars on the road will have such capabilities, thus opening up the possibility of autonomous interactions among multiple vehicles.

In earlier work, we proposed a novel Multiagent Systems-based approach to alleviating traffic congestion and collisions, specifically at intersections (Dresner & Stone 2005). In this work, we make three further contributions. First, we augment our existing intersection control mechanism to allow use by human drivers with minimal additional infrastructure. Second, we show that this hybrid mechanism offers performance and safety benefits over traditional traffic light systems. Finally, we show that at each stage, there exists an incentive to use autonomous vehicles over traditional vehicles. All work is fully implemented and tested in our custom simulator and we present experimental results to support its effectiveness.

Reservation System

The reservation system as previously proposed consists of two types of agents: *intersection managers* and *driver agents*. For each intersection, there is a corresponding intersection manager, and for each vehicle, a driver agent. Intersection managers are responsible for directing the vehicles through the intersection, while the driver agents are responsible for controlling the vehicles to which they are assigned.

To improve the throughput and efficiency of the system, the driver agents “call ahead” to the intersection manager and request space-time in the intersection. The intersection manager then determines whether or not these requests can be met based on an *intersection control policy*. Depending

on the decision (and subsequent response) the intersection manager makes, the driver agent either records the parameters of the response message (the *reservation*) and attempts to meet them, or it receives a rejection message and makes another request at a later time. If a vehicle has a reservation, it can request that its reservation be changed or cancel the reservation. It also sends a special message when it finishes crossing the intersection indicating to the intersection manager that it has done so.

The interaction among these agents is governed by a shared protocol which we have published in a technical report (Dresner & Stone 2004). In addition to message types (e.g. REQUEST, CONFIRM, and CANCEL), this protocol includes some rules, the most important of which are (1) that a vehicle may not enter the intersection unless it is within the parameters of a reservation made by that vehicle’s driver agent, (2) that if a vehicle follows its reservation parameters, the intersection manager can guarantee a safe crossing for the vehicle, and (3) a driver agent may have only one reservation at a time. Aside from this protocol, no agent needs to know how the other agents work — each vehicle manufacturer (or third party) can program a separate driver agent, each city or state can create their own intersection control policies (which can even change on the fly), and as long as each agent adheres to the protocol, the vehicles will move safely through the intersection.

Incorporating Human Users

While an intersection control mechanism for autonomous vehicles will someday be very useful, there will always be people who enjoy driving. Additionally, there will be a fairly long transitional period between the current situation (all human drivers) and one in which human drivers are a rarity. Even if switching to a system comprised solely of autonomous vehicles were possible, pedestrians and cyclists must also be able to traverse intersections in a controlled and safe manner. For this reason, it is necessary to create intersection control policies that are aware of and able to accommodate humans, whether they are on a bicycle, walking to the corner store, or driving a “classic” car for entertainment purposes.

Our autonomous intersection control mechanism grants reservations to autonomous vehicles by ensuring that no space in the intersection (which is discretized into *tiles*) is

*This research is supported by NSF CAREER award IIS-0237699.

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

used by more than one vehicle at a time. Vehicles communicate their requirements to the intersection and reservations are granted on a “first come, first served” basis.

To add human drivers to the mix, we first need a reliable way to communicate information to the humans. The best way to do this is to use a system that drivers already know and understand — traffic lights. Traffic light infrastructure is already present at many intersections and the engineering and manufacturing of traffic light systems is well developed. For pedestrians and cyclists, standard “push-button” crossing signals could be used that would give enough time for a person to traverse the intersection. These could also serve to alert the intersection to their presence.

We add a new component to each intersection control policy, called a *light model*. This controls the actual physical lights as well as providing information to the policy with which it can make decisions. The lights are the same as modern-day lights: red (do not enter), yellow (if possible, do not enter; light will soon be red), and green (enter). Each control policy will need to have a light model so that human users will know what to do.

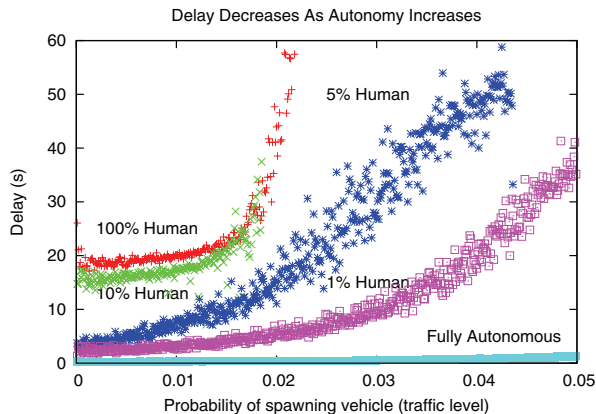


Figure 1: Average delays for all vehicles as a function of traffic level for various proportions of human drivers and incarnations of our intersection control policy. As time progresses and autonomous vehicles become more widely adopted, the policy can be modified to take advantage of the increasing automation, leading to lower delays.

We combined the light model and our original control policy such that the human users obey the lights, but the autonomous vehicles still make reservations. When the intersection manager receives requests from autonomous vehicles, it must first consult the light model to determine whether it can grant safe passage to the vehicle. If the vehicle’s path would coincide with a “green-light” path, it is granted. Otherwise, the tiles in the vehicle’s path are checked as in the fully-autonomous mechanism. Tiles that would be used by human vehicles operating under the control of the lights are considered *off-limits*.

When the system is first introduced and the number of autonomous vehicles is small, the light model can be just like a normal traffic light. Autonomous vehicles, however,

would have the advantage of being able to use the underlying reservation-based framework when they have a red light. Once autonomous vehicles are more prevalent, the door is opened to even more efficient control policies. Shown in Figure 1 are results for various incarnations of the human-usable control mechanism:

1. An all human traffic light
2. A traffic light with 90% autonomous vehicles
3. A traffic light that only makes one lane green at a time, with 95% autonomous vehicles
4. Same as above but with 99% autonomous vehicles
5. A fully-autonomous version with always-red lights.

Notice the rather drastic jump (especially at low amounts of traffic) from the “normal” traffic light paradigm to the “lane-by-lane” version (in which only one lane receives a green light at a time). Each additional autonomous vehicle not only improves efficiency for the vehicle’s owner (over a human-driven vehicle), but also improves the overall efficiency of the intersection. Videos of the mechanism in action, as well as other supplementary material can be found at <http://www.cs.utexas.edu/users/kdresner/aim/>.

Conclusion

As a future of fully autonomous vehicles approaches, it is important that we prepare to exploit the advantages offered by such a scenario. With modern day intersection control technology (e.g. traffic lights), the benefits of autonomous vehicles would be limited to convenience and safety. In this work, we have taken several steps towards a future in which autonomous vehicles not only make our lives easier and safer, but make vehicle travel much more time- and cost-efficient.

References

- Dresner, K., and Stone, P. 2004. Multiagent traffic management: A protocol for defining intersection control policies. Technical Report UT-AI-TR-04-315, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory.
- Dresner, K., and Stone, P. 2005. Multiagent traffic management: An improved intersection control mechanism. In *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 471–477.
- Pomerleau, D. A. 1993. *Neural Network Perception for Mobile Robot Guidance*. Kluwer Academic Publishers.
- Rogers, S.; Flechter, C.-N.; and Langley, P. 1999. An adaptive interactive agent for route advice. In Etzioni, O.; Müller, J. P.; and Bradshaw, J. M., eds., *Proceedings of the Third International Conference on Autonomous Agents (Agents’99)*, 198–205. Seattle, WA, USA: ACM Press.
- Schonberg, T.; Ojala, M.; Suomela, J.; Torpo, A.; and Halme, A. 1995. Positioning an autonomous off-road vehicle by using fused DGPS and inertial navigation. In *2nd IFAC Conference on Intelligent Autonomous Vehicles*, 226–231.