From: AAAI Technical Report WS-93-04. Compilation copyright © 1993, AAAI (www.aaai.org). All rights reserved.

Improving the Cost-Performance Tradeoff in Traffic Control using Autonomous Intelligent Agents

WAYNE IBA
RECOM Technologies
AI Research Branch, Mail Stop 269-2
NASA Ames Research Center
Moffett Field, CA 94035 USA
IBA@PTOLEMY.ARC.NASA.GOV

1 Introduction

Rapid and concentrated growth in many urban areas has led to a dramatic increase in the number of commuters and a subsequent strain on many cities' traffic signal facilities. Although traffic delays have almost become an accepted part of our culture, significant efforts (and funding) are focussed on improving the overall performance of urban traffic-control systems. Unfortunately, responses to congestion have typically focussed on centralized control systems that are extremely expensive.

Recent developments in overhead infrared traffic sensors enable inexpensive access to more detailed information on traffic conditions [1]. In conjunction with more intelligent controllers that can take advantage of the better information, this new sensor technology can lead to greater efficiency at much lower cost.

The task that we pose is to improve the overall efficiency of traffic flow in urban settings. More specifically, we want to find an ideal tradeoff between the efficiency of traffic flow and the costs of development, acquisition, installation, and maintenance of the necessary physical upgrades. In order to select the appropriate tradeoff, we must identify the significant sources of improvement (over current traffic-control systems), quantify the resulting improvements, and determine the fiscal impact of implementing these improvements. We recognize that current safety requirements must be strictly maintained and assume that they are not subject to any tradeoff.

2 Motivation and Background

There are three very obvious benefits provided by a satisfactory solution to the tradeoff selection problem that we have undertaken. First, by reducing congestion, a more efficient and responsive signal control system would save time for everyone who would otherwise be sitting in traffic. Second, reducing the number of starts and stops would reduce the frequency of individual vehicle accelerations, which is responsible for most of the fuel consumption and auto pollution associated with city driving [2]. We would further argue that reducing the number of starts and stops would reduce the number of minor accidents – a major source of traffic congestion. Third, a low-cost approach such as we suggest would save considerable amounts of money as compared to many current initiatives intended to improve urban traffic flow. As an example of a large centralized project, San Jose's Traffic Signal Management Program currently has a projected total budget over 26 million dollars [3]. Much of the cost involves communication infrastructure, which would be unnecessary in our framework.

Considerable work has been performed towards achieving these benefits. Much of this work has explored the issue of controller *coordination* or integration. A second, controller *responsiveness*, is often confounded with the first. Yet another issue of more recent interest is that of *adaptability*. These three issues can be thought of as dimensions defining a matrix of possible control approaches.

We propose that much of the negative results obtained from experiments with on-line adaptive approaches is a consequence of poor responsiveness and adaptability. A responsive controller should not only react to the current situation, but should operate according to predictions about situations in the very near future. The prediction is a key to efficient traffic control and without it, many useful approaches will appear inferior.

In our current work, we address a very modest instantiation of the problem outlined above. Using techniques for temporal prediction, we hypothesize that intelligent local control with detailed sensory information can approach or exceed the performance of systems using coordinated timing. That is, by improving responsiveness (via better sensory data and machine learning techniques) we can achieve much of the same improvements provided through coordination but at significantly reduced costs.

Naturally, this hypothesis is based on a number of assumptions. First, and most obviously, it requires the availability of low-cost sensors capable of providing rich data about individual vehicles in a timely fashion. This may require some advances in image processing as well as sensor technologies. We also assume (although the hypothesis does is not depend on this) that providing coordination *can* improve overall performance. That is, our approach does not require coordination but should make advantageous use of it when it is provided. We envision a hierarchical control architecture similar in spirit to that described in Cuena, et al [4].

3 Our Traffic Simulator

Because our line of work depends upon detailed sensory information, we need to use a simulator that models individual cars. The outputs of our assumed sensors are individual vehicle locations, velocities, and perhaps sizes. In addition, our research interest is in flexible controllers that utilize the available sensory information in arbitrary ways. That is, the controller should not be limited to adjusting a time period by some factor of the number of cars detected passing or waiting. The simulator should allow the controller to perform more complex reasoning that takes into account current sensor values and predicts likely values in the near future. So, our need is for a microsimulator that decouples the simulation of traffic flow from the behavior of the controller; this can be accomplished through a clean interface between the sensor values and the agent's control actions. The simulator only provides the sensory information and responds to the control actions of the control agent.

Current methods either help optimize the timing of coordinated signals within a coarse model of vehicular traffic, or provide a fine-grained evaluation of a fixed coordination strategy (e.g., TRANSIT7F and TRAF-NETSIM respectively [5]). However, existing systems assume the only source of sensory information is from a limited number of loop detectors and the control algorithm is constrained to a fixed sequence of modes.

Since a satisfactory simulator was not available off the shelf, we developed our own simple simulator that satisfies our requirements. We modified the NASA Tileworld simulator [6] to create a two dimensional traffic world. Each car occupies a single cell of a grid, and travels along the road according to stochastic operators. An intersection consists of two crossing streams of cars; the common grid-cell can only be occupied by a single car. The controller mitigates the contention for the common cell by stopping traffic on one stream, while permitting the other stream to pass. The sensors provide the controller with information about the state of the tile world. In the current implementation, the controller has access to (but is not required to use) complete information about the simulated world.

With several extensions and modifications that are currently underway, we can provide a framework enabling traffic engineers to experiment with more sophisticated sensors and more intelligent controllers than is possible with currently available simulation frameworks. Furthermore, this can facilitate the rapid prototyping of advanced traffic control agents, which could then be tested on similar, but finer grained, simulators that have yet to be developed.

4 A Design for a Flexible Traffic Control Agent

Our proposed solution to the tradeoff between efficiency and cost is based on a simple testable hypothesis – intelligent local control with detailed sensory information can approach or exceed the performance of systems using coordinated timing. This combination of enhanced sensors and artificial intelligence techniques represents a fresh logic applied to the standard traffic-control problem. If the sensory information can include a limited level of communication between local signals, performance can be further improved. If communication with a central coordinator is provided, even further improvements can be gained. The goal is to provide as much improvement as possible by combining the more detailed sensory information and an intelligent controller, while taking advantage of existing infrastructure.

The intelligent controller is based on developments in integrated cognitive architectures [7]. Our approach views the problem as that of mapping available sensory data onto allowable signal modes, and assumes that more sensor information enables more efficient control. In order to perform the correct mapping, the system treats the sensory data at, and approaching, an intersection as a temporal signature. Previously, we have successfully represented and recognized temporal information as probabilistic concepts that are organized within a polythetic decision tree [8]. The classification of the temporal signature determines a current world state that prescribes a certain mode of operation for the signal. Often, no response by the controller is needed, but when the mode changes the controller will cycle the signal through appropriate states in response to the new conditions.

The particular effectiveness of the approach stems from the ability to predict the state of the world a reasonable distance into the future. These predictions are a function both of the current sensory information and the possible changes to the state of the traffic light. Based on the system's knowledge about world states and the effects of its actions, it can reason about the consequences of various choices. For example, if a platoon of cars is detected approaching the intersection, the system can reason that changing the light at this point will result in greater congestion. However, if other features of the world are also true (such as cross traffic already waiting a certain amount of time), then it may need to change states anyway.

The system's knowledge about the world and about the consequences of its actions can be programmed directly into memory at the beginning. Through experience, this knowledge can be fine-tuned to more closely approximate the actual situations encountered at a particular intersection. However, even when the knowledge base is accurate, the system can search for improvements by evaluating various measures of effectiveness that are continuously available to the system. These measures are the basis for determining how efficiently the system is controlling the flow of traffic.

Although the knowledge base is not static, strict constraints can be imposed to ensure safety and prevent wild divergence from intended or acceptable behavior. For example, the system may discover a modification that would yield a significant overall improvement in efficiency but would prove to be psychologically unacceptable to individual drivers. These constraints, like the initial rules and strategies, can easily be specified ahead of time.

5 Current Status and Future Work

We have implemented the preliminary version of our simulator and it is running on a Sun Sparc station under Allegro Common-Lisp. A graphical interface using X-windows provides a visual display of the movement of traffic in the simulated world and the behavior of the intersection controllers. Traffic streams, controllers, graphics display features, and the evaluation executive all run as separate processes within lisp using the multi-process capabilities provided by Allegro Common-Lisp.

To date, we have implemented two controllers. A pre-timed (non-actuated) controller and a simple "rule-based" controller. The rule-based controller considers the demand on both approaching traffic streams and switches between phases when the demand on the stopped stream exceeds a specified factor of the demand on the permissive stream. Even this primitive controller shows improved performance over the pre-timed one. For the default traffic conditions that we have tested, the rule-based controller improves average speed by 20.6%.

In a second experiment with similar traffic conditions but with two traffic streams crossing a main street, we compared a pair of pretimed controllers in coordination to two independent rule-based controllers. The rule-based controllers had no information about the actions or states of the other controller. The independent controllers achieved a 26.2% improvement over the coordinated scheme. These results must be taken with a grain of salt; they represent a single set of traffic conditions and measure only distance and travel time (speed). Clearly, more extensive experiments are necessary (and are underway). However, we can interpret the results as weak evidence in support of our primary hypothesis and as an encouragement to continue this line of work.

Above, we outlined a possible design for a flexible intelligent controller that is based on earlier work with autonomous agents. Most notably, we need to ground this design in an implementation that can be tested against more conventional methods. This work is currently underway. We are also extending the simulator to handle more sophisticated and realistic driving behavior in individual cars. In conjunction with these developments, we are experimenting with other independent variables, such as number of starts and stops, and expanding the number of dependent measures we can use to evaluate a controller's performance.

These research directions are straightforward extensions of the work described in this paper. However, we also intend to characterize the improvements that are possible in order to better assess the utility of expenditures towards new and more complex traffic control systems. In some cases, cities are spending many millions of dollars to improve traffic flow. Although an improvement may be measurable, it may not be subjectively significant to individual drivers. Traffic improvement must be measurable and significant, but must also be individually relevant; an individual driver should notice the difference. This essentially introduces another dimension or dependent variable along which proposed (and implemented) control systems should be evaluated.

6 Conclusions

This research is strongly motivated by at least three factors. First, growth in urban areas has caused considerable strain on existing systems, causing significant congestion and delays. Second, progress in infrared detectors suggests that a low-cost traffic detector with sufficient resolution can be produced in the near term. These sensors will provide more extensive information (e.g., traffic locations and velocities over a significant range of distances) at a lower initial cost, and will operate with much less maintenance than loop detectors. Finally, advances in autonomous agents and adaptive control mechanisms have paved the way for more responsive and safer traffic signals. We believe our experience with autonomous agents can provide a solid foundation on which to explore significant improvements to existing traffic control systems.

Currently, we have implemented and tested an initial version of a micro traffic simulator that does not constrain the research to a particular sensor or control paradigm. Preliminary experiments with the simulator and a very simple rule-based agent have demonstrated the expected advantages of using actuated control to mitigate traffic contention, and more importantly, have provided supporting evidence to our hypothesis that intelligent local control can compete favorably with centralized control. We are currently designing and running more extensive experiments that test our hypotheses over a broader range of conditions and evaluate a more complete range of measures of effectiveness. The primary focus of our continuing work is the design and implementation of flexible intelligent controllers.

Acknowledgments

We thank Al Harbury, Rajeev Saraf, and Cindy Mason for lengthy and helpful discussions on the ideas presented in this paper. Mark Drummond helped focus the ideas for the traffic simulator. Peter Friedland, Kevin Thompson, and John Allen all provided helpful comments on earlier drafts of this paper.

References

[1] Wood, R. A., et al. (1992). Uncooled Monolithic Silicon Focal Plane Development. Presented at *Infrared Information Symposium, Detector Specialty Group.*

[2] Fuel Efficient Transportation Simulation contract. Institute of Transportation Studies, University of California, Berkeley

[3] Ron Northouse (personal communication), City of San Jose's Traffic Signal Management Program.

[4] Cuena, J., et al. (1992). A general knowledge-based architecture for traffic control: The KITS approach. Proceedings of the Second OECD Workshop on Knowledge-based Expert Systems in Transportation.
[5] McTrans, Center for Microcomputers in Transportation, Gainsville, FL.
[6] Phillips, A. B., & Bresina, J. L. (1991). *NASA TileWorld Manual.* Artificial Intelligence Research Branch Technical Report, FIA-91-11. NASA Ames Research Center.
[7] Langley, P., McKusick, K. B., Allen, J. A., Iba, W. F., & Thompson, K. (1991). A design for the ICARUS architecture. *SIGART Bulletin, Vol. 2, No. 4*, 104–109.
[8] Iba, W. (1991). Learning to classify observed motor behavior. *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence* (pp. 732–738). Sydney: Morgan Kaufmann.