SELF-LEARNING INTELLIGENT AGENTS FOR DYNAMIC TRAFFIC ROUTING ON TRANSPORTATION NETWORKS

Adel Sadek

Dept. of Civil and Environmental Engineering
and Dept of Computer Science
University of Vermont
asadek@cems.uvm.edu

Nagi Basha

Dept. of Computer Science
University of Vermont
nbasha@uvm.edu

Abstract

Intelligent Transportation Systems (ITS) are designed to take advantage of recent advances in communications, electronics, and Information Technology in improving the efficiency and safety of transportation systems. Among the several ITS applications is the notion of Dynamic Traffic Routing (DTR), which involves generating "optimal" routing recommendations to drivers with the aim of maximizing network utilizing. In this paper, we demonstrate the feasibility of using a self-learning intelligent agent to solve the DTR problem to achieve traffic user equilibrium in a transportation network. The core idea is to deploy an agent to a simulation model of a highway. The agent then learns by itself by interacting with the simulation model. Once the agent reaches a satisfactory level of performance, it can then be deployed to the real-world, where it would continue to learn how to refine its control policies over time. To test this concept in this paper, the Cell Transmission Model (CTM) developed by Carlos Daganzo of the University of California at Berkeley is used to simulate a simple

highway with two main alternative routes. With the model developed, a Reinforcement Learning Agent (RLA) is developed to learn how to best dynamically route traffic, so as to maximize the utilization of existing capacity. Preliminary results obtained from our experiments are promising. RL, being an adaptive online learning technique, appears to have a great potential for controlling a stochastic dynamic systems such as a transportation system. Furthermore, the approach is highly scalable and applicable to a variety of networks and roadways.

1.0 Introduction

In recent years, there has been a concerted effort aimed at taking advantage of the advances in communications, electronics, and Information Technology in order to improve the efficiency and safety of transportation systems. Within the transportation community, this effort is generally referred to as the Intelligent Transportation Systems (ITS) program. Among the primary ITS applications is the notion of Dynamic Traffic Routing (DTR), which involves routing traffic in real-time so as to maximize the utilization of existing capacity. The solution to the DTR problem involves determining the *time-varying* traffic splits at the different diversion points of the transportation network. These splits could then be communicated to drivers via Dynamic Message Signs or In-vehicle display devices.

Existing approaches to solving the DTR problem have their limitations. This paper proposes a solution for highway dynamic traffic routing based on a self-learning intelligent agent. The core idea is to deploy an agent to a simulation model of a highway. The agent will then learn by itself through interacting with the simulation model. Once the agent reaches a satisfactory level of performance, it could then be deployed to the real world, where it would continue to learn how to refine its control policies over time. The advantages of such approach are quite obvious given the fact that real-world transportation systems are stochastic and ever-changing, and hence are in need of on-line, adaptive agents for their management and control.

1.1. Reinforcement Learning

Among the different paradigms of soft computing and intelligent agents, Reinforcement Learning (RL) appears to be particularly suited to address a number of the challenges of the on-line DTR problem. RL involves learning what to do and how to map situations to actions to maximize a numerical reward signal (Kaelbling, 1996; Kretchmar, 2000; Abdulhai and Kattan, 2003; Russell and Norvig, 2003). A Reinforcement Learner Agent (RLA) must discover on its own which actions to take to get the most reward. The RLA will learn this by trial and error. The agent will learn from its mistakes and come up with a policy based on its experience to maximize the attained reward (Sutton and Barto, 2000).

The field of applying RL to transportation management and control applications is still in its infancy. A very small number of studies could be identified from the literature. A Q-learning algorithm (which is a specific implementation of reinforcement learning) is introduced in Abdulhai et. al. (2003) to study the effect of

deploying a learning agent using Q-learning to control an isolated traffic signal in real-time on a two dimensional road network. Abdulhai and Pringle (2003) extended this work to study the application of Q-learning in a multi-agent context to manage a linear system of traffic signals. The advantage of having a multi-agent control system is to achieve robustness by distributing the control rather than centralize it even in the event of communication problems. Finally, Choy et al. (2003) develop an intelligent agent architecture for coordinated signal control and use RL to cope with the changing dynamics of the complex traffic processes within the network.

2.0 Purpose and Scope

The main purpose of this study is to show the feasibility of using RL for solving the problem of providing online Dynamic Route Guidance for motorists, through providing a set of experiments that show how an RL-based agent can provide reasonable guidance for a simple network that has two main routes. Figure 1 shows the simple network used in this study. It should be noted that this network is largely similar to the test network used by Wang et al. (2003) in evaluating predictive feedback control strategies for freeway network traffic.
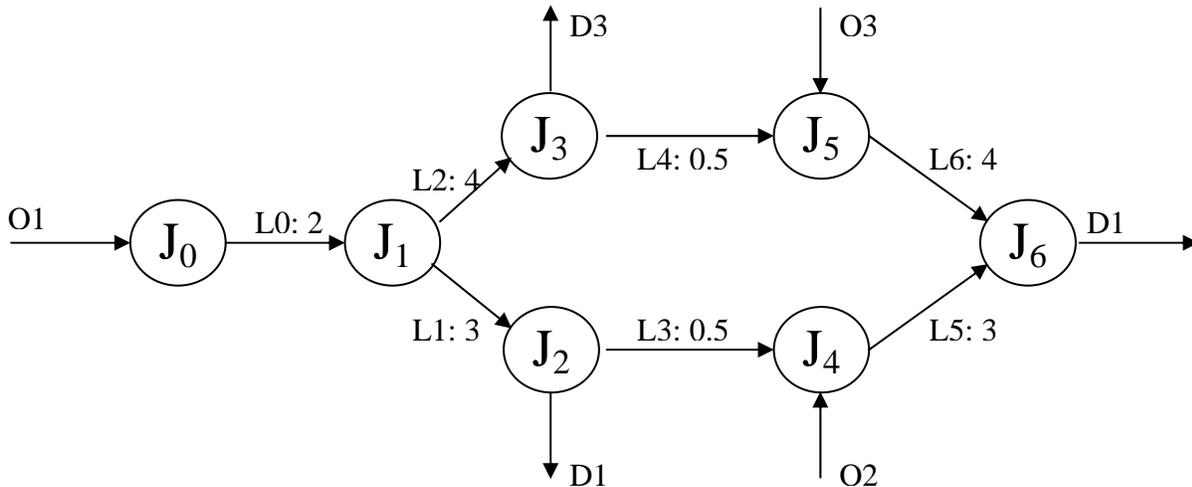

Figure 1: Network Topology

The network has three origins: O1, O2, and O3 and three destinations D1, D2, and D3. Each origin generates a steady flow of traffic. Traffic disappears when it reaches any of the three destinations. The length in miles of each link is indicated on the graph. For example, L0 has a length of 2 miles. The capacity of all links is 4000 veh/h except for L0 that has a capacity of 8000 veh/h. All links have two lanes except L0 that has 4 lanes.

As can be seen from Figure 1, there are two alternate routes connecting origin O1 to destination D1. The primary route (route A) has a total length of 6.50 miles, whereas the secondary route (route B) is 8.50 miles long, is therefore longer than route A. The intelligent RL agent is deployed at the J1 junction. The goal of the agent is to determine an appropriate diversion rate at J1 so as to achieve traffic user equilibrium between the two routes connecting zones O1 and D1 (i.e. so that travel times along routes A and B are as close to each other as possible), taking into consideration the current state of the system.

3.0 Methodology

3.1 Cell Transmission Model

In this study, we selected the Cell Transmission Model (CTM) to build the simulation model, with which the RL agent would interact to learn for itself the best routing strategies. The CTM was developed by Daganzo to provide a simple representation of traffic flow capable of capturing transient phenomena such as the propagation and dissipation of queues (Daganzo, 1994; 1995). The model is macroscopic in nature, and works by dividing each link of the roadway network into smaller, discrete, homogeneous cells. Each cell is appropriately sized to permit a simulated vehicle to transverse the cell in a single time step at free flow traffic conditions. The state of the system at time t is given by the number of vehicles contained in each cell, $n_i(t)$. Daganzo showed that by using an appropriate model for the relationship between flow and density, the cell-transmission model can be used to approximate the kinematic wave model of Lighthill and Whitham (1955). For this study, a C++ implementation of the CTM was developed and used to simulate the test network.

3.2. The Intelligent Agent

As previously mentioned, RL was the paradigm chosen to develop the intelligent, learning agent that will be used for dynamic traffic routing. Specifically, the learning algorithm implemented in the agent is based on the SARSA algorithm, which is an on-policy Temporal Difference (TD) learning implementation of reinforcement learning (Sutton and Barto, 2000). SARSA is a temporal difference algorithm because – like Monte Carlo ideas - it can learn directly from the experience without requiring a model of the dynamics of the environment. Like Dynamic Programming ideas (Bertsekas, 2000), SARSA updates the desirability of its estimates of state-action pairs, based on earlier estimates; i.e. SARSA does bootstrapping. For a complex, unpredictable, and stochastic system like a transportation system, SARSA seemed to be very suitable to adapt with the nature of an ever-changing system.

    The implementation of the SARSA algorithm is quite simple. Each state action pair (s,a) is assigned an estimate of the desirability of being in state s and doing action a. The desirability of each state action pair can be represented by a function Q(s,a). The idea of SARSA is to keep updating the estimates of Q(s,a) based on earlier estimates of

$Q(s,a)$ for all possible states and all possible actions that can be taken in every single state. Equation 1 shows how the $Q(s,a)$ values are updated.

$$Q(s_t,a_t) = Q(s_t,a_t) + \alpha \: [r_{t+1} + Q(s_{t+1},a_{t+1}) - Q(s_t,a_t)] \qquad \text{[Equation 1]}$$

where α is the step-size parameter or the learning rate. According to Equation [1], at time t, the system was in state s_t, and the agent decided to take action a_t. This resulted in moving the system to state s_{t+1} and obtaining a reward of r_{t+1}. Equation [1] is thus used to find new estimates for the Q-values for a new iteration as a function of the values from the previous iteration. The algorithm typically would go through several iterations until it converges to the optimal values of the Q-estimates.

3.3. Experiment Setup

As was previously mentioned, the objective of the experiment presented in this paper is to have an agent that is capable of recognizing the state of the system and deciding upon a diversion rate at junction J_1. If this diversion rate is followed, the system should eventually move to a state of equilibrium where taking any of the two routes will result in the same travel time. For example, if route A (L1) is totally blocked because of an accident, the agent should guide all motorists to take route B (L2).

3.3.1 State, Action, Reward Definition

State representation is quite important for the agent to learn properly. Ideally, using the CTM, each state is represented by the number of vehicle in each cell. Computationally, it is impossible to use this representation for the states for a RL agent, since this would make the problem state space very large. In this experiment, the state of the system is represented by the difference in the instantaneous time between taking the short route through L1 (route A) and the longer route through L2 (route B). Based on some empirical experiments, the state space was discretized to a finite number of states. Table 1 shows how the state was discretized based upon the difference in instantaneous travel time between the longer route, route B, and the shorter route, route A.

Time difference (dif) in minutes	State code	Time difference (dif) in minutes	State code
0 < dif < 2	0	0 > dif > -2	-10
2 < dif < 4	1	-2 > dif >-4	-1
4 < dif < 7	2	-4 > dif > -7	-2
7 < dif < 10	3	-7 > dif > -10	-3
10 < dif < 15	4	-10 > dif > -15	-4
15 < dif < 20	5	-15 > dif > -20	-5
20 < dif < 30	6	-20 > dif > -30	-6
30 < dif < 40	7	-30 > dif > -40	-7
40 < dif < 60	8	-40 > dif > -60	-8
60 < dif	9	-60 > dif	-9

Table 1: State Space

As can be seen, State 0, for example, refers to the case when both routes are running at free flow speed. For this case, the difference in travel time between routes B and A is in the range of +2 minutes, since route B is 2.0 miles longer than route A. On the other hand, state -9 refers to the case when route A (the shorter route) is extremely congested (totally closed), while the longer route is doing fine. In our experiments, the instantaneous time is determined from speed sensor readings along the two routes.

For actions, ideally the diversion rate is a real number between 0 and 100%; which means an infinite set of actions. In this experiment, the set of actions were reduced to only six actions. Table 2 indicates the six different actions used in this experiment.

Action code	Action meaning
0	Divert 100% of the flow to L1 & 0% to L2
1	Divert 80% of the flow to L1 & 20% to L2
2	Divert 60% of the flow to L1 & 40% to L2
3	Divert 40% of the flow to L1 & 60% to L2
4	Divert 20% of the flow to L1 & 80% to L2
5	Divert 0% of the flow to L1 & 100% to L2

Table 2: Action Set

For the SARSA algorithm, all the values of Q(s,a) are initialized to zero. As the agent experiments with different actions for the sates encountered, the environment responds with a reward, which in our case is equal to the negative of the difference in instantaneous travel time between the two routes. In other words, the goal of the agent is to take the proper action to ensure that the instantaneous difference in time between the two routes does not exceed 2 minutes; i.e. reaching the state of equilibrium.

In the experiment, the authors simulated running the system for around 90 hours of operation. Every five hours and half, an accident is introduced on link L1. The accident lasts for half an hour causing the reduction of L1 capacity to half its original capacity assuming that one of the two lanes on L1 is blocked because of the accident. The purpose of introducing the accident repetitively is to make sure that the agent encounters many different states and lasts in each state long enough to learn the best action for it.

4.0   Results and Discussions

The results of the experiment show that the agent managed to learn the right actions for all the states it has encountered. For example, for the state of complete congestion on link L1 (state -9), the agent learned that the proper action is to divert 100% of the traffic to the longer route (action 5). On the other hand, for the free flow state (state 0), the agent learned that the best action is to divert all the traffic to the shorter path (action 0). The agent also learned the correct actions for the intermediate states. For example, when the travel time on the longer route was 10 minutes less than the shorter route, the agent learned to divert only 20% of the traffic to the shorter route and 80% to the longer one. Once the situation improves and the difference is only 7 minutes instead of 10, the agent adjusts the diversion rate to 40 to the shorter route and 60% to the longer one.

Figure 2 shows the convergence to the right action for state –9. The time is represented in seconds*10. Notice that the system got into state –9 after almost six hours of the simulation; 30 minutes after the accident. The agent chose action 0 (diverting all the traffic to the shorter route) the first time the agent encountered state –9. This decision is actually the worst decision that can be made in this situation. But as time passes, the agent decided to switch to action 1 which is still not the proper action to take. By the end of the simulation model, the agent converged to the right action which is diverting all traffic to the long route that exhibits shorter travel time.
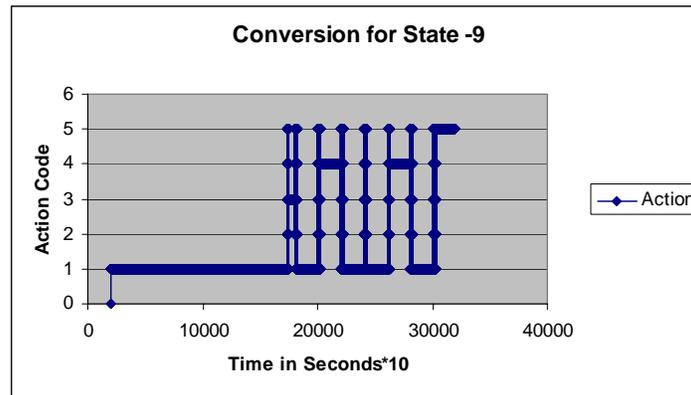

Figure 2: Conversion for State –9

5. Conclusions

The results obtained from this preliminary work are very promising. For our future research, we are planning to use a neural network to augment the state representation of the system. Using a neural network will allow us to deal with a bigger set of states as well as achieving a smother and continuous representation of the state space. Similarly for actions, a neural net would be much more appropriate than the set of 6 different actions we used in this experiment. In the future we are also planning to experiment with a more complex transportation network, and to utilize a community of cooperating intelligent agents for traffic management and control. **.**

6. Acknowledgement

References

Abdulhai, B. and Kattan L.(2003) "Reinforcement learning: Introduction to theory and potential for transport applications" Canadian Journal of Civil Engineering, Volume 30, Number 6 (11), 981-991.

Abdulhai, B., Pringle, P., and Karakoulas, G.J. (2003) "Reinforcement Learning for True Adaptive Traffic Signal Control. ASCE Journal of Transportation Engineering, Vol. 129(3), pp. 278-284.

Abdulhai, B. and Pringle, P. (2003). "Autonomous Multiagent Reinforcement Learning –5 GC Urban Traffic Control." A paper presented at the 2003 Annual Transportation Research Board Meeting, Washington, D.C.

Bertsekas, Dimitri P. (2000). Dynamic Programming and Optimal Control. Athena Scientific, Massachusetts USA, 502 pp.

Choy, M. C., Cheu, R.L., Srinivasan, D., and Logi, F. (2003). Real-Time Coordinated Signal Control Using Agents with Online Reinforcement Learning. A paper presented at the 2003 Annual Transportation Research Board Meeting, Washington, D.C.

Daganzo, C.F. (1994). "The Cell Transmission Model: A Dynamic Representation of Highway Traffic Consistent with the Hydrodynamic." Transportation Research, Vol 28B, No. 4, 269-287.

Daganzo, C.F. (1995). "The Cell Transmission Model, Part II: Network Traffic." Transportation Research, Vol 29B, No. 2, 79-93.

Kaelbling, Leslie P., Littman, Michael L., and Moore, Andrew W. (1996) "Reinforcement Learning: A Survey." Journal of Artificial Intelligence Research, (4) 237-285.

Kretchmar, R. Matthew (2000). A Synthesis Of Reinforcement Learning And Robust Control Theory. Ph.D. Diss., computer Science, Colorado Sate University, Fort Collins, Colorado, 139 pp.

Lighthill, M.J. and Whitham, J.B. (1955) "On Kinematic Waves. I: Flow movement in long rivers; II. A theory of traffic flow on long crowded roads." Proc. Royal Society A, 229, pp. 281-345

Russell, S. and Norving, P.(2003) Artificial Intelligence A Modern Approach. Second Edition, Prentice Hall.

Sutton, Richard S. and Barto, Andrew G. (2000). Reinforcement Learning. MIT Press, Massachusetts USA, 322 pp.

Wang, Y, Papageorgiou, M. and A. Messmer. (2003) "A Predictive Feedback Routing Control Strategy for Freeway Network Traffic." Transportation Research Record 1312, TRB, National Research Council, Washington, D.C, pp. 21-43.