AAAI Workshop on AI in Intelligent-Vehicle Highway Systems, July 1993

Transportation Applications of Artificial Intelligence (Extended Abstract)

Michael P. Wellman
University of Michigan
Artificial Intelligence Laboratory
1101 Beal Avenue
Ann Arbor MI 48109
wellman@engin.umich.edu

I present an overview of some ongoing research on IVHS-related problems at the University of Michigan Artificial Intelligence Laboratory. Our work covers three principal areas: (1) individual route planning under time-dependent uncertainty, (2) decentralized computation of network equilibria using market-price mechanisms, and (3) dynamic traffic modeling for routing and situation assessment.

1. Introduction

In this extended abstract, I present an overview of IVHS-related research recently initiated at the University of Michigan Artificial Intelligence Laboratory. Our ultimate goals in this effort are to further the technology underlying decision support for both individual drivers and traffic management centers within a variety of decision contexts, including routing, scheduling, and control of signals. Part of this technology involves new modeling tools for describing transportation networks and traffic flow, and reasoning mechanisms to support forecasting and situation assessment.

Of particular research interest are problems that involve significant uncertainty, dynamic information-gathering over time from heterogeneous sources, and distributed decision-making processes. The remainder of this abstract describes preliminary efforts in each of these areas. In some of this work, we make significant use of existing techniques from AI, albeit with important extensions. In the following presentation, we describe the extensions to current methods, as well as some new techniques for distributed problem solving applicable to a wide variety of decentralized decision-making tasks.

2. Route Planning under Time-Dependent Uncertainty

Consider a transportation network with nodes denoting locations and edges denoting possible transportation operations between the locations connected. If travel times are static (that is, the duration of a trip from a to b does not depend on departure time), then we can compute the fastest route from any given origin to all possible destinations using Dijkstra's well-known shortest-path algorithm, where the costs on each link are the travel times. This algorithm has a worst-case complexity of $O(N^2)$, where N is the number of nodes in the network. If the travel times are stochastic but independent (that is, the distribution of travel times for one link does not depend on the actual travel time on others), then the route with the fastest *expected* total travel time can be found similarly with Dijkstra's algorithm, where the costs on each link correspond to expected travel times.

Unfortunately, this shortest-path algorithm is not valid when the travel times are time-dependent. This sort of situation should be expected in realistic highway networks, where traffic patterns vary throughout the day, as well as in other transportation networks (e.g., bus routes), where transfer times depend on fixed schedules. For the deterministic case, however, it has been shown (Kaufman & Smith, 1993) that the standard shortest-path algorithm is indeed sound as long as the network satisfies the following reasonable consistency condition. Let s and t be departure times such that $s \le t$, and let $c_{ij}(x)$ denote the time-dependent cost (travel time) of traveling from location i to location j at time x. The network is *consistent* iff

$$\text{for all } i, j.\ s + c_{ij}(s) \le t + c_{ij}(t).$$

This condition seems quite reasonable for time-dependent transportation networks. It merely says that although leaving later can perhaps reduce the duration of the trip, it cannot decrease the ultimate arrival time. Given this condition, the principle of optimality underlying Dijkstra's algorithm applies, and the shortest-path problem can be solved relatively efficiently.

A stochastic version of this condition, with the times replaced by expectations, would similarly validate the use of the standard algorithm with expectations. However, this version would not be reasonable, as demonstrated by some simple examples (Hall, 1986). (Moreover, the expectation version of the shortest-path algorithm would not produce correct results for these examples.) As an alternative, I propose the following condition. Let $c_{ij}(x)$ denote the time-dependent travel time (a random variable) from location i to location j given departure at time x. Let us say the network is *stochastically consistent* iff for all i, j, and z,

$$\Pr(\mathrm{s} + c_{ij}(s) \le z) \ge \Pr(\mathrm{t} + c_{ij}(t) \le z).$$

In other words, the probability of arriving by any given time cannot be increased by leaving later. This appears to be the most natural (and most benign) generalization of the deterministic consistency condition above. It is based on the concept of stochastic dominance, a common way to extend an ordering relation to random variables.

This condition justifies a modified version of the shortest-path algorithm, where instead of maintaining the shortest path found to all intermediate nodes (in the uncertain case, a probability distribution of travel times), we maintain all *undominated* paths. If one path to a node dominates another (in the sense of stochastic ordering), then the stochastic consistency condition ensures that the latter cannot be part of an overall shortest path. This generalized use of the optimum principle can lead to substantial savings if the network contains many dominated paths, as we would expect.

Although I believe this algorithm to possess advantages over existing methods for computing fastest path in stochastic time-dependent networks (e.g., (Hall, 1986)), definitive statements await the result of formal and empirical analysis, which has just begun. This approach should also be applicable to problems with non-additive (but monotonic) costs, generalizing the scope of existing search-based optimization algorithms in OR and AI. For example, with a heuristic estimate of remaining distance, this algorithm is a stochastic version of A*.

3. Network Equilibrium Problems

Network equilibrium analysis is an increasingly popular framework for studying the allocation of traffic flows on transportation networks. Equilibria are typically calculated via optimization algorithms (e.g., variational inequality formulations), which require a centralized or global analysis. Since the ultimate decision-making in traffic applications is generally decentralized, we seek computational methods mirroring the distributed structure of the decision-support environment.

We have studied network equilibrium problems as a special type of distributed resource-allocation task. Our general approach, called "market-oriented programming" (Wellman, 1992)., works by casting a problem as a computational economy and deriving its competitive equilibrium. We formulate the given problem as a general-equilibrium system, with agents corresponding to the decision-making entities, and goods corresponding to relevant resources and outcome attributes. The agents interact exclusively via consumption and production of goods, and communicate solely by submitting bids to auctions for each of the commodities. The protocol implements a variant of tatonnement, converging on equilibrium prices given some restrictions on the form of preferences and technologies.

To test the application of computational markets to transportation network tasks, we have explored a distributed version of the multicommodity flow problem. In a multicommodity flow problem, the task is to allocate a given set of cargo movements over a given transportation network. The transportation network is a collection of locations, with links (directed edges) identifying feasible transportation operations. Associated with each link is a specification of the cost of moving cargo along it. A movement requirement associates an amount of cargo with an origin-destination pair. The planning problem is to determine the amount to transport on each link in order to move all the cargo at minimum cost. In the distributed version, we require that flow allocation decisions are made separately for each movement requirement. The following figure depicts a simple transportation network (from (Harker, 1988)).

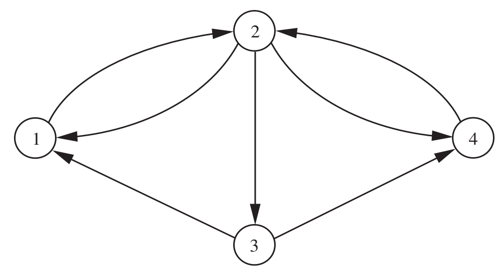

We have investigated three market configurations for solving multicommodity flow problems. In model **S**, the economy consists of a set of shipper agents (consumers), one for each movement requirement. The shippers bid for links serving their origin-destination pair, splitting cost proportionately on shared links. In model **SC** (**S** plus carriers), we augment the economy with carrier agents (producers), one for each link on the network. The carriers "own" the links, and sell transportation services in return for basic resources (fuel, vehicles, etc.), so as to maximize profit. The economic structure of model **SC** is depicted below for the simple transportation network shown above.

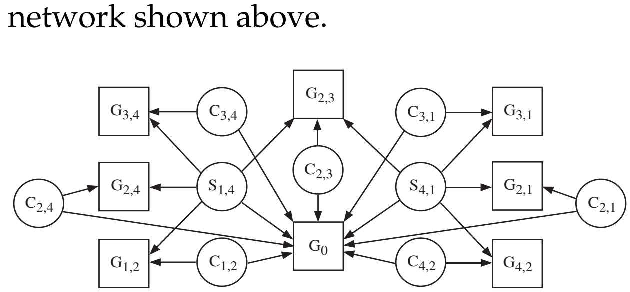

Finally, in model **SCA**, we add arbitrageur agents (producers), whose role is to examine isolated network fragments and try to derive profits from any transient price imbalances.

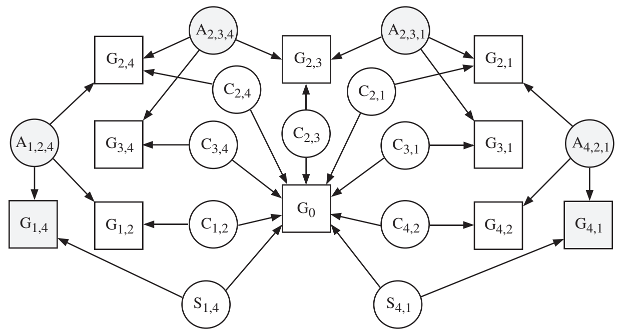

All models converge to an equilibrium. In model **S**, we produce the *user equilibrium*, which is socially suboptimal due to an externality (an interaction outside the market system) among the agents' preferences. **SC** eliminates the externality by delegating the shared links to profit-maximizing carriers, and hence results in a global optimum (the *system equilibrium*). **SCA** retains this global optimum, but with a higher degree of decentralization since the shippers no longer need to perform path analysis. The main conclusions from this exercise are (1) the computational market can derive useful results for a nontrivial transportation problem, (2) different market configurations lead to qualitatively different results and computational behaviors, and (3) these differences can be predicted and analyzed using standard concepts from economics.

4. Dynamic Traffic Interpretation

Finally, we are exploring the use of modern uncertain reasoning techniques from AI for transportation tasks such as dynamic traffic interpretation. We represent our state of information in probabilistic networks, graphical representations of random variables and their interdependencies (Pearl, 1988). The technical advantage of probabilistic networks is that they exploit independence relations among the variables for computation, without imposing any uniform independence restrictions. We have begun preliminary explorations of the problem of representing routing problems under uncertainty in such graphical decision models, and of appropriate structures for intertemporal traffic modeling.

The distinguishing characteristic of dynamic traffic interpretation is that we receive partial and noisy information about a dynamically evolving traffic situation. Observed information may include vehicle sightings, estimated instantaneous loads, gross traffic reports, incident (e.g., accident) notifications, etc. Prior knowledge may include map information, common traffic patterns, etc. The goal is to derive high-level interpretations of the overall situation, including assessments of the overall loads, prediction of short-term future loads, locations of bottlenecks, and identification of prototypical phenomena (e.g., gridlock, accidents, etc.). The particular descriptors of interest will depend dynamically on the control options available (e.g., access control, route guidance, or even commanded diversions) and the nature of information input.

One of the major technical challenges here is how to dynamically modify the model *structure* over time as information increases and traffic evolves, given general knowledge about the behavior of traffic and the impact of important traffic events. This problem of *knowledge-based model construction* is an active area of research in the uncertain reasoning community (Wellman, Breese, & Goldman, 1992).

5. Summary

In summary, we are investigating the application of AI techniques to a variety of transportation and IVHS-related problems. There is no single "AI technique" to be deployed here; rather we are using these application problems to stress and extend the available technology in a variety of ways. The particular methodologies I have cited in this abstract are:

- heuristic search,
- market-oriented programming, and
- probabilistic networks.

Each of these (as well as many others) can make a contribution to IVHS applications, but each will require significant customization and extension to produce useful results in this domain.

References

Hall, R. W. (1986). The fastest path through a network with random time-dependent travel times. *Transportation Science*, **20**(3), 182-188.

Harker, P. T. (1988). Multiple equilibrium behaviors on networks. *Transportation Science*, **22**(1), 39-46.

Kaufman, D. E., & Smith, R. L. (1993). Fastest paths in time-dependent networks for intelligent vehicle-highway systems application. *IVHS Journal*, **1**(1).

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Morgan Kaufmann.

Wellman, M. P. (1992). A general-equilibrium approach to distributed transportation planning. In *Proceedings of the National Conference on Artificial Intelligence*, (pp. 282-289). San Jose, CA: AAAI Press.

Wellman, M. P., Breese, J. S., & Goldman, R. P. (1992). From knowledge bases to decision models. *Knowledge Engineering Review*, **7**, 35-53.