

Optimal Vehicle Routing with Real-Time Traffic Information

Seongmoon Kim¹

*Department of Industrial and Systems Engineering
Florida International University
10555 W. Flagler Street / (EC 3100), Miami, FL 33174
Seongmoon.Kim@FIU.edu, 305-348-6912*

Mark E. Lewis

*Department of Industrial and Operations Engineering
University of Michigan
1205 Beal Avenue, Ann Arbor, MI 48109-2117
melewis@engin.umich.edu, 734-763-0519*

Chelsea C. White III

*School of Industrial and Systems Engineering
Georgia Institute of Technology
765 Ferst Drive, Atlanta, GA 30332
cwhite@isye.gatech.edu, 404-894-0235*

revised, September 14, 2004

¹Corresponding author.

Abstract

This paper examines the value of real-time traffic information to optimal vehicle routing in a non-stationary stochastic network. We present a systematic approach to aid in the implementation of transportation systems integrated with real time information technology. We develop decision-making procedures for determining the optimal driver attendance time, optimal departure times, and optimal routing policies under time-varying traffic flows based on a Markov decision process formulation. With a numerical study carried out on an urban road network in Southeast Michigan, we demonstrate significant advantages when using this information in terms of total costs savings and vehicle usage reduction while satisfying or improving service levels for just-in-time delivery.

Introduction

In order to compensate for uncertainties in travel time due to accidents, bad weather, traffic congestion, etc., trucks hauling time-sensitive freight build “buffer time” into their routes in order to help ensure that deliveries will be made on time. Building buffer time into routes tends to increase the likelihood of on-time delivery, an important measure of service. However, buffer time also tends to reduce measures of productivity associated with cost, such as driver and equipment idle time and the number of miles traveled per hour.

One goal of this paper is to show that real-time traffic information combined with historical traffic data can be used to develop routing strategies that tend to improve both cost and service productivity measures. More specifically, motivated by situations where time-sensitive delivery is required, we examine the value of a real-time traffic information technology (IT) on vehicle routing. We present a systematic approach to aid in the implementation of transportation systems integrated with this real-time IT. To this end, we consider a stochastic shortest path problem on a road network composed of links having non-stationary travel times, where a subset of these links are observed in real-time. We assume that each observed link can be in one of two states (congested or uncongested) that determines the travel time distribution used. Fundamental questions to be addressed are:

- When should the driver of a commercial vehicle be made available to leave the origin?
- Once made available, when should the driver actually depart the origin?
- How should a vehicle be routed through the network, based on real-time traffic information, to reduce time and cost?

A successful analysis of the first question can lead to less time spent idling at the origin resulting in savings to the company. The second question addresses that we can further reduce total vehicle usage and, hence, total cost. Indeed a truck parked at the origin is less expensive than one that is in use. Moreover, if enough time can be saved, the vehicle can be used to make other deliveries

so that the fleet size can be reduced. The last question has the obvious benefits but is the most difficult to answer. It requires that we model a heavily congested road network that is equipped with a real-time traffic information technology.

We assume the following decision-making scenario. We decide to bring the driver in to the origin at a predetermined time called the driver attendance time. We assume the driver is available from the driver attendance time and the driver begins to be paid from that moment. In practice the driver attendance time is set in advance since the driver may not be available immediately when requested. After bringing in the driver, we continuously observe the congestion status of each observed link and at some point we set the actual departure time. Note that the driver attendance time may be (and often is) different than the actual departure time because we may observe traffic conditions that make it advantageous to delay departure. After departing the origin, the decision-maker must choose the next intersection to visit given the current congestion status of each observed link. Upon reaching the next intersection, or perhaps immediately prior to reaching it, the status of the observed links is again observed and the trip continues.

The purpose of this paper is to develop decision-making procedures for determining the optimal driver attendance time, optimal departure times, and optimal routing policies under time varying traffic flows based on a Markov decision process (MDP) formulation. To our knowledge this is the first research that evaluates the benefits of historical and real-time traffic information by using actual traffic data. Real traffic data were collected in Southeast Michigan from the Michigan Intelligent Transportation Systems Center (MITSC) for 30 days in July and August 2000. We demonstrate significant advantages when using this information, relative to the case without it, from the following performance measures

- Percentage savings in total costs,
- Percentage reduction in vehicle usage.

In order to obtain approximate average travel speeds each minute, we use a linear interpolation scheme for smoothing the originally collected data and a Markov chain model for system dynamics

that we claim can be implemented and updated in real-time.

The classic shortest path problem has been extensively examined in the literature. Eiger et al. (1985) demonstrated that when the utility function is linear or exponential, an efficient Dijkstra-type algorithm can be used to compute the minimum cost route in a stationary stochastic network. Hall (1986) showed that standard shortest path algorithms (such as Dijkstra's algorithm) do not find the minimum expected cost path on a non-stationary stochastic network and that the optimal route choice is not a simple path but a policy. The best route from any given node to the final destination depends not only on the node, but also on the arrival time at that node. Dynamic programming was proposed for finding the optimal policy.

When a road map is characterized by a tightly interconnected network of nodes, the complexity of finding the shortest path was estimated in Section 5.3 of Pearl (1984). Using the A^* algorithm, guided by the air-distance heuristic, it was shown that a fraction of the nodes expanded under breadth-first search would also be expanded by A^* . AO^* for a non-stationary stochastic shortest path problem with terminal cost was investigated by Bander and White III (2002). It was shown that AO^* is more computationally efficient than dynamic programming when lower bounds on the value function are available. Similar formulations to our model without real-time traffic congestion information were also presented.

For the non-stationary deterministic case, Kaufman and Smith (1993) have shown that the standard shortest-path algorithm is indeed sound as long as the network satisfies a deterministic consistency condition. For the non-stationary stochastic case, Wellman et al. (1995) developed a revised path-planning algorithm where a stochastic consistency condition holds. Miller-Hooks and Mahmassani (1998) showed and compared algorithms for the least possible cost path in the discrete-time non-stationary stochastic case. An algorithm for finding the least expected cost path in that case was provided in Miller-Hooks and Mahmassani (2000).

Of particular relevance to this paper is the *stochastic shortest path problem with recourse* originally studied by Croucher (1978). In this work, when a driver chooses a link to traverse, there is a fixed probability that it will actually traverse an adjacent link as opposed to the one chosen.

One might think of the original link as being one on which an accident or unexpected road block has occurred. Interesting extensions of this work include that of Andreatta and Romeo (1988) and Polychronopoulos and Tsitsiklis (1996). In the prior, a path from the source to the destination is chosen *a priori*. There is then a fixed probability upon arriving to one of the nodes in the network that the next link is “inactive” and an alternate route must be chosen. In the latter work, the authors assume that the costs to traverse arcs are random variables and that travelling along the network reduces the possibilities for “states” of the network. The problem we consider has a less general cost structure than Polychronopoulos and Tsitsiklis (1996) but allows for both the costs and travel times to be non-stationary in time. A more theoretical discussion of similar problems can be found in Bertsekas and Tsitsiklis (1991) where the authors show the existence of optimal policies without the non-negative arc-cost assumption. The recent paper by Waller and Ziliaskopoulos (2002) addresses the question with random arc costs but with only local arc and time dependencies. In Gao and Chabini (2002) the authors discuss several approximations of the time dependent stochastic shortest path problem where information about the network is learned by the driver as the trip progresses. Psaraftis and Tsitsiklis (1993) considered a problem similar to the one discussed in this paper in which the travel time distributions on links evolve over time according to a Markov process except that the changes in the status of the link are not observed until the vehicle arrives at the link. The work of Fu and Rilett (1998) and Fu (2001) discusses implementations of real-time vehicle routing based on estimation of mean and variance travel times.

This paper is an extension of Kim et al. (2002), which examines more comprehensive issues integrating vehicle routing with real-time information technology. In subsequent work (see Kim et al. (2004)), we examine computational issues of this study. In particular, since the problem presented includes time dependence, the state space of the Markov decision process can be quite large. In Kim et al. (2004) we provide a systematic methodology to delete many of the observed links from the state space before the trip begins. We also discuss how the state space can be further reduced during a trip. In both cases this is done without loss of optimality and stands to make the problem more tractable. For the network studied in this paper the CPU time was reduced on

average to below 3% of the original model.

The rest of the paper is organized as follows. Section 1 presents a formulation of the problem using a Markov decision process and states the dynamic programming optimality equations. Section 2 describes our approach to applying the introduced MDP model and focuses on the implementation procedures integrating vehicle routing with real-time traffic information systems. Section 3 presents how to identify the optimal driver attendance time and provides algorithms determining optimal departure times based on real-time traffic information. Section 4 estimates the value of historical and real-time traffic information with actual traffic data in Southeast Michigan. In Section 5, we discuss conclusions and future research directions.

1 Problem Statement and Optimality Equations

We now formulate the non-stationary stochastic shortest path problem with real time traffic information as a discrete-time, finite horizon Markov decision process (MDP). Consider an underlying network $G \equiv (N, A)$, where the finite set N represents the set of nodes and $A \subseteq N \times N$ is the set of directed links in the network. This network serves as a model of a network of roads (links) and intersections (nodes). By this, we mean that $(n, n') \in A$ if and only if there is a road segment that permits traffic to travel from intersection n to intersection n' . Let $n_0 \in N$ be the start node and the set $\Gamma \subseteq N$ be the goal node set. For each element $n \in N$ define the *successor set* of n (denoted $SCS(n)$) to be the set of nodes that have an incoming link emanating from n . That is, $SCS(n) \equiv \{n' : (n, n') \in A\}$. A path $p = (n_0, \dots, n_K)$ from $n_0 \in N$ is a sequence of nodes such that $n_{k+1} \in SCS(n_k)$ for $k = 0, 1, \dots, K - 1$. We assume the existence of a path from any node in the network to the goal node set.

A link $(n, n') \in A$ is said to be observed if real-time traffic is measured and reported on (n, n') . Suppose that there are Q observed links in A . Define the (random) road congestion status vector

at time t to be $Z(t) = \{Z^1(t), \dots, Z^Q(t)\}$ so that the random variable $Z^q(t)$ is

$$Z^q(t) = \begin{cases} 0 & \text{if the } q^{\text{th}} \text{ link is uncongested} \\ 1 & \text{if the } q^{\text{th}} \text{ link is congested} \end{cases}$$

for $q = 1, 2, \dots, Q$. Denote a realization of $Z(t)$ by z . Thus, $z \in H \equiv \{0, 1\}^Q$. We assume that $\{Z^i(t), t = t_0, t_0 + 1, \dots\}$ and $\{Z^j(t), t = t_0, t_0 + 1, \dots\}$ are independent Markov chains for $i \neq j$. For each $q = 1, 2, \dots, Q$, we assume that the dynamics of the corresponding Markov chain are described by the one-step transition matrix

$$R_q^{(t, t+1)} = \begin{bmatrix} \alpha_t^q & 1 - \alpha_t^q \\ 1 - \beta_t^q & \beta_t^q \end{bmatrix}.$$

Let $P(z'|t, z, t')$ be the probability of a transition occurring from z at time t to z' at time t' .

By our assumption of independence,

$$P(z'|t, z, t') \equiv P[Z(t') = z' | Z(t) = z] = \prod_{q=1}^Q P[Z^q(t') = (z^q)' | Z^q(t) = z^q]$$

where $(z^q)'$ is the q^{th} element of z' and each term satisfies a simple extension of Kolmogorov equations for the non-stationary case (cf. Theorem 5.4.3 of Ross (1996))

$$R_q^{(t, t')} = \begin{bmatrix} \alpha_t^q & 1 - \alpha_t^q \\ 1 - \beta_t^q & \beta_t^q \end{bmatrix} \times \begin{bmatrix} \alpha_{t+1}^q & 1 - \alpha_{t+1}^q \\ 1 - \beta_{t+1}^q & \beta_{t+1}^q \end{bmatrix} \times \dots \times \begin{bmatrix} \alpha_{t'}^q & 1 - \alpha_{t'}^q \\ 1 - \beta_{t'}^q & \beta_{t'}^q \end{bmatrix}.$$

Thus, the congestion status of the network is modeled by a Q -dimensional, non-homogeneous Markov chain.

Let $P(t'|n, t, z, n')$ denote the probability of arriving at node n' at time t' , given that the vehicle travels from node n to n' , departing node n at time t with congestion status vector z . This probability distribution is assumed to be discrete, and the travel time between any two nodes is assumed to be bounded, i.e., $P(t'|n, t, z, n') = 0$ for all $t' \leq t$ and for all $t' > t + \zeta$ for a given, finite value ζ . Notice that we have not precluded n from being a member of its own successor

set. When this situation occurs, the driver is allowed to wait at node n , and we assume that $P(t+1|n, t, z, n) = 1$, for all t and z .

Denote the k^{th} node visited by the vehicle as n_k , and let t_k be the time that node n_k is visited. Thus, n_0 is the start node and t_0 the start time. We assume the existence of a fixed (and known) time T after which no other decisions are made. Thus, if n_k is the k^{th} node visited and t_k is the time of the k^{th} decision epoch, upon arrival to the next node, say n_{k+1} at time $t_{k+1} > T$, a terminal cost $\bar{c}(n_{k+1}, t_{k+1})$ is accrued. We assume that $\bar{c}(n, t) = \infty$ for all $t > T$ and $n \notin \Gamma$.

Since the goal of a decision-maker is to find the minimum cost path, T is simply for modelling convenience and can be chosen to be arbitrarily large. We choose T such that if t_0 is less than some $\hat{T} \ll T$ there exists a path from any node to Γ such that the time of arrival to Γ is (guaranteed) to be less than T . This is of course possible as a consequence of the bounded travel times assumption. For example, suppose ζ bounds the travel times on each link. Recall that we have assumed the existence of a path from any node to Γ . Since $|N| < \infty$ these paths can be chosen to be finite. Choosing T and \hat{T} such that $T - \hat{T} > |N|\zeta$ guarantees our assertion.

Let $U = \{0, 1, \dots, T\}$ be the set of possible times that decisions are made. Define the state space of our decision problem to be

$$\Omega \equiv \{(n, t, z) : n \in N, t \in U, z \in H\}.$$

A (deterministic, Markov) policy π is a function such that $\pi : \Omega \rightarrow N$ and prescribes which node should be visited next for each node, time, and congestion status; that is, $\pi(n, t, z) \in SCS(n)$. Let z_k be the congestion status at time t_k . We assume for $n_k \in \Gamma$ that $\pi(n_k, t_k, z_k) \in \Gamma$ for all π ; i.e., once the goal node set is reached, it is never left. When the decision is made to travel from n_k to $n_{k+1} = \pi(n_k, t_k, z_k)$, the (finite) random variable T_{k+1} is determined according to the discrete probability distribution $P[T_{k+1} = t_{k+1} | n_k, t_k, z_k, n_{k+1}]$.

Let $\tilde{c}(n, t, z, n', t')$ be the cost accrued by traversing road segment $(n, n') \in A$, given that travel begins at time t and ends at time t' with the congestion status z at time t . We assume the congestion of other links does not affect the cost of traversing, or the time to traverse, the current

link. For $n \in \Gamma$ assume that $\tilde{c}(n, t, z, n', t') = \bar{c}(n, t) = 0$. Thus, once the goal node set is reached, no other costs are accrued. Define the expected cost of traveling from n to n' , starting at time t given congestion status z , by

$$c(n, t, z, n') = \sum_{t'} P(t'|n, t, z, n') \tilde{c}(n, t, z, n', t').$$

Note that in many practical applications, there are “ideal” and “latest acceptable” arrival time windows to the goal node set. Thus, for a given link (n, n') such that $n \notin \Gamma$ and $n' \in \Gamma$ there may be two components of the cost function $\tilde{c}(n, t, z, n', t')$, a part corresponding to the cost of traversing the link (n, n') and a part corresponding to arriving at the goal node set at time t' that represents the penalty for late or early arrival. Assume for all $(n, t, z) \in \Omega$ and n' such that $n \notin \Gamma$ and $t' \leq T$, there exists $b, B \in \mathbb{R}^+$ such that $0 < b \leq c(n, t, z, n', t') \leq B < \infty$. We can now define the total expected cost of a policy.

Let the random variable M be the number of decision epochs before time T . The total expected cost accrued under policy π is

$$v^\pi(n_0, t_0, z_0) = \mathbb{E}_{n_0, t_0, z_0}^\pi \left\{ \sum_{k=0}^M c[n_k, t_k, z_k, \pi(n_k, t_k, z_k)] + \bar{c}(n_{M+1}, t_{M+1}) \right\}$$

where $\mathbb{E}_{n_0, t_0, z_0}^\pi$ is the expectation operator, conditioned on beginning at state $(n_0, t_0, z_0) \in \Omega$. We seek an optimal policy π^* such that

$$v^*(n_0, t_0, z_0) \equiv v^{\pi^*}(n_0, t_0, z_0) \leq v^\pi(n_0, t_0, z_0),$$

for all policies π .

For any function $f : \Omega \rightarrow \mathbb{R}^+$, let

$$h[n, t, z, n', f] \equiv c[n, t, z, n'] + \sum_{t'} P[t'|n, t, z, n'] \sum_{z'} P[z'|t, z, t'] f[n', t', z'].$$

Thus, h represents the total expected cost when in state (n, t, z) , n' is chosen as the next node to visit and a terminal reward f is accrued after moving to n' .

For any $(n, t, z) \in \Omega$ we call the following equation the *optimality equation*

$$v^*(n, t, z) = \min_{n' \in SCS(n)} \{h(n, t, z, n', v^*)\} \quad (1.1)$$

subject to the boundary condition $v^*(n, t, z) = \bar{c}(n, t)$ for $t > T$. Furthermore, a policy, π^* , is optimal if and only if, for all n, t, z ,

$$\pi^*(n, t, z) \in \arg \min_{n' \in SCS(n)} \{h(n, t, z, n', v^*)\}. \quad (1.2)$$

When considered together, the boundedness of the cost function $c(n, t, z, n', t')$ (for $t' \leq T$), the existence of a path from each node to Γ , the fact that $|N| < \infty$, and the definition of \hat{T} guarantee the existence of an optimal policy with finite expected cost.

Since the problem posed is a finite state and action space, finite horizon MDP, a simple extension of the results of Chapter 4 of Puterman (1994) yields that the value function related to the non-stationary stochastic shortest path problem with real time congestion information satisfies the optimality equation (1.1), that (1.2) characterizes an optimal policy, and that we can use (1.1) to compute v^* via backward induction. We note here that computing optimal values and policies may not be a trivial task since the state space of the problem may be quite large. However, state space reduction techniques are discussed in the subsequent paper Kim et al. (2004).

2 Model Application

We now describe our approach to applying the MDP model introduced in Section 1. In particular, we focus on the implementation of an integrated vehicle routing system with real-time traffic information. This system includes a methodology for estimating probability distributions for travel time, transition probabilities of congestion status on each link, and defining costs.

2.1 Data collection and transformation

For 30 days through July and August 2000, the Michigan Intelligent Transportation Systems Center (MITSC) in Detroit collected and provided traffic data, obtained from more than 100 loops on major




Figure 1: A commercial traffic volume map in Southeast Michigan

freeways in Southeast Michigan. Based on the commercial traffic volume map in Figure 1, provided by the Michigan Department of Transportation (MDOT), each intersection was examined. The data from loops provided by MITSC were the average vehicle speed in each 15-minute interval throughout a day as in Table 1. For the road segments of freeways where no traffic data were provided, linear interpolation and extrapolation methods were used to generate pertinent speed data. For each arterial road segment where the speed limit is much lower than that on the freeway, the speed limit was assigned as the vehicle speed (deterministically).

Date	Loop number	0:00-0:15	0:15-0:30	0:30-0:45	0:45-1:00	...	23:45-0:00
07/21/00	1	73.21	72.19	72.98	73.83		73.98
	2	59.12	60.21	60.89	59.34		58.97
	⋮						
07/22/00	1	72.68	71.49	71.90	72.31		72.99
	2	60.13	61.85	61.94	62.28		61.86
	⋮						
07/23/00	⋮						
⋮							

Table 1: Traffic data: Average of vehicle speed at loops in each 15-min interval. (unit : mph)

To evaluate the total cost by units of one minute, we must estimate the per minute average of vehicle speed from 15-minute average speed. Note that there are ninety-six 15-minute intervals in

a day (from 1440 minutes per day). Let d_k be the average of vehicle speed, the obtained raw traffic data, in each 15-minute interval, and b_{15k+i} be the average of vehicle speed in each 1-minute interval for $k = 0, 1, \dots, 95$ and $i = 0, 1, \dots, 15$. We compute b_{15k+i} for $k = 0, 1, \dots, 95$ and $i = 0, 1, \dots, 15$, satisfying the constraints that the original 15-minute average speed d_k should remain the same after obtaining each 1-minute average speed and that traffic flow should evolve smoothly and linearly in a small increment of time. To this end, we use linear interpolation to estimate the per minute average of vehicle speed from the 15-minute average speeds; See Figure 2.




Figure 2: Inference of 1-minute average speed from 15-minute average speed, where $a_k = b_{15k+7}$ for $k = 0, 1, \dots, 95$ (the midpoint of each interval)

2.2 Transition probability estimation

2.2.1 Estimating probability distributions of travel time

Distances between intersections in the road network are measured by using the commercial software package '*Microsoft Streets & Trips 2001*'. These distances allow us to estimate travel times between node pairs. Applying the central limit theorem, approximate (truncated) normally distributed travel times are obtained for every arc in the network at every minute during a day. If the travel time τ from node n to node n' has a normal distribution, i.e., $\tau \sim N(\mu, \sigma^2)$, we discretize in the

following manner:

$$\begin{aligned}
p\{\tilde{\tau} = t\} &= P\{t - 0.5 \leq \tau < t + 0.5\} \\
&= P\{(t - 0.5 - \mu)/\sigma \leq (\tau - \mu)/\sigma < (t + 0.5 - \mu)/\sigma\} \\
&= \Phi((t + 0.5 - \mu)/\sigma) - \Phi((t - 0.5 - \mu)/\sigma),
\end{aligned}$$

where $\tilde{\tau}$ is the discretized travel time and μ and σ are estimated from the data.

2.2.2 Transition of congestion status

The model requires a prediction of the dynamics of the future traffic congestion status given the current traffic congestion status. The definition of congestion is likely to be problem specific based on the typical travel times along each arc. As suggested by the MITSC, the threshold between congested and uncongested states was set to 50(mph) on freeways. That is, let b_t^q be the 1-minute average vehicle speed at time t for the q^{th} observed arc. Then, for $q = 1, 2, \dots, Q$ and $t = 1, 2, \dots, 1440$, define the random variable $Z^q(t)$ by

$$Z^q(t) = \begin{cases} 0 & \text{if } b_t^q \geq 50 \text{ mph} \\ 1 & \text{if } b_t^q < 50 \text{ mph} , \end{cases}$$

where 0 (1) corresponds to an uncongested (congested) link.

Since the average travel speeds are approximately normally distributed, elements of the transition matrix of $Z^q(t)$ can be approximated using the bivariate normal as follows. Suppose $b_t^q \sim N(\mu_1, \sigma_1^2)$, $b_{t+1}^q \sim N(\mu_2, \sigma_2^2)$, and ρ is the correlation between b_t^q and b_{t+1}^q . Then, if z^q at time t changes to $z^{q'}$ at time $t + 1$, $P(z^{q'} | t, z^q, t + 1)$ is approximated using, for example,

$$P(z^{q'} = 1 \mid t, z^q = 0, t + 1) = P(b_{t+1}^q < 50 \mid b_t^q \geq 50) = \frac{P(b_{t+1}^q < 50, b_t^q \geq 50)}{P(b_t^q \geq 50)}$$

where

$$P(b_{t+1}^q < 50, b_t^q \geq 50) = \int_{-\infty}^{50} \int_{50}^{\infty} \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} \exp\left\{-\frac{1}{2}D(x, y)\right\} dx dy .$$

and

$$D(x, y) \equiv \frac{1}{1 - \rho^2} \left\{ \left(\frac{x - \mu_1}{\sigma_1} \right)^2 - 2\rho \left(\frac{x - \mu_1}{\sigma_1} \right) \left(\frac{y - \mu_2}{\sigma_2} \right) + \left(\frac{y - \mu_2}{\sigma_2} \right)^2 \right\}.$$

Again, each of the parameters is approximated from the data.

2.3 Cost estimation

We let $\tilde{c}(n, t, z, n', t') \equiv \delta(t' - t)$ for $n \neq n'$, $n' \notin \Gamma$, $t' \leq T$, and $\delta > 0$. That is, the cost of traversing the link is proportional to the travel time, with the rate of δ units per unit time. The expected travel cost from n to n' given the start time $t < \hat{T}$ and corresponding congestion status z is then

$$c(n, t, z, n') = \sum_{t'} P(t'|n, t, z, n') \tilde{c}(n, t, z, n', t') = \sum_{t'} P(t'|n, t, z, n') \delta(t' - t).$$

By noting $n \in SCS(n)$ and $P(t' = t + 1|n, t, z, n) = 1$, we define $c(n, t, z, n) \equiv \delta_1 > 0$ for $n \notin \Gamma$ and $t < T$ to be the driver wage rate and we assume it to be less than the rate of $c(n, t, z, n')$ for $n' \neq n$, $n' \notin \Gamma$, and $t < \hat{T}$, i.e.,

$$c(n, t, z, n) \equiv \delta_1 < \delta.$$

Let $[t_e, t_l]$ be an acceptable arrival time window, where t_e is the earliest acceptable arrival time, and t_l is the latest acceptable arrival time at the destination, for $t_e < t_l$ and $t_l - t_e < \infty$. We let penalties for late arrival be $\delta_2(t_K - t_l)^+$ and penalties for early arrival be $\delta_3(t_e - t_K)^+$ where t_K is the first arrival time to the goal node set, and $0 < \delta_3 \leq \delta_2$; it is less costly to arrive early than late. Thus, $\tilde{c}(n_{K-1}, t_{K-1}, z_{K-1}, n_K, t_K) = \delta(t_K - t_{K-1}) + \delta_2(t_K - t_l)^+ + \delta_3(t_e - t_K)^+$ (almost surely) for $t_K \leq T$.

Given this formulation, there are several possible approaches. The first approach reflects the current practice of ignoring real time information and using commercial logistics software and fixed routes. Another approach involves using historical data to develop static policies that, although they disregard real time traffic information, are an improvement over static routes provided by the commercial software. A final approach would use historical data to implement a real-time traffic

information system to develop dynamic routes. In the following section, we demonstrate methods for comparing these alternatives.

3 Optimal Driver Attendance and Departure Times

We now discuss algorithms for computing optimal driver attendance times and departure times when real-time congestion information is available or not available. In essence, we show that the algorithms allow us to compare both scenarios and show that these algorithms converge finitely. Although in the subsequent numerical study we use them to study the value of congestion information, we envision that these algorithms could also be used in practice as a planning aid to schedule drivers and their departure from the origin.

3.1 Value functions without real-time information

We begin by describing how the optimal cost and route would be computed when commercial software is used. This will serve as a baseline for comparison with the dynamic system previously described. Let $d(n, n')$ denote the distance and $l(n, n')$ denote the speed limit between nodes n and n' . Then, it is often the case that commercial logistics software, such as '*Microsoft Streets & Trips 2001*', assumes that the travel time between nodes n and n' , $\tau(n, n') = \frac{d(n, n')}{l(n, n')}$. The software then chooses a (static) route so as to minimize the total route cost from the origin to the destination by solving, for example, the following equations recursively

$$w(n) = \min_{n' \in SCS(n)} \{\tau(n, n') + w(n')\},$$

where $w(\gamma) = 0$ for all $\gamma \in \Gamma$.

Let $P(t'|n, t, n')$ denote the probability of arriving at node n' at time t' , given that the vehicle travels from node n to n' , departing node n at time t . Suppose $c(n, t, n')$ denotes the expected cost of going from n to n' starting at time t . Let $p_c = (n_0, \dots, n_K)$ be the route determined by the commercial software, where n_0 is the origin and $n_K \in \Gamma$. Let $u(n_k, t)$ be the expected cost accrued starting at node n_k on path p_c at time t and travelling, via p_c , to the goal node n_K . Then, $u(n_k, t)$

satisfies the recursive equation

$$u(n_k, t) = c(n_k, t, n_{k+1}) + \sum_{t'} P(t'|n_k, t, n_{k+1})u(n_{k+1}, t'), \quad (3.1)$$

where $u(\gamma, t) = 0$ for $\gamma \in \Gamma$.

For the case with historical traffic data but without real-time traffic information, the minimal expected total cost $v^*(n, t)$ satisfies the optimality equation

$$v^*(n, t) = \min_{n' \in SCS(n)} \{c(n, t, n') + \sum_{t'} P(t'|n, t, n')v^*(n', t')\}, \quad (3.2)$$

where $v^*(\gamma, t) = 0$ for all $\gamma \in \Gamma$. Of course, the routes chosen by the commercial software and by using (3.2) ignore the dynamic nature of traffic and, as we will see, can result in a significant increase in travel costs.

3.2 Definition of optimal driver attendance and departure times

We now formally define the optimal driver attendance times in each case; with commercial logistics software, with historical traffic data, and with both historical and real-time traffic information.

Definition 3.1

(i) *If commercial software is used to solve the shortest path problem as described by equation (3.1), the optimal driver attendance time is the largest t_c^* , such that*

$$u(n_0, t_c^*) \leq u(n_0, t) \quad \text{for all } t.$$

(ii) *If historical traffic data are used as described by equation (3.2), the optimal driver attendance time is the largest t_{h1}^* , such that*

$$v^*(n_0, t_{h1}^*) \leq v^*(n_0, t) \quad \text{for all } t.$$

(iii) *If both historical and real-time traffic data are used as described by equation (1.1), the optimal driver attendance is the largest t_{r1}^* , such that*

$$E\{v^*[n_0, t_{r1}^*, Z(t_{r1}^*)]\} \leq E\{v^*[n_0, t, Z(t)]\} \quad \text{for all } t.$$




Figure 3: The optimal expected total cost as a function of driver attendance time. Each curve is typical of what was found in the numerical results.

Figure 3 illustrates the optimal driver attendance time in each case. Note that it represents the time that minimizes the total cost and is thus the time at which we would like to have a driver available to leave the origin. The obvious questions arise

1. Does each minimum exist?
2. If so, can we compute them easily?

The next result states that the optimal driver attendance time is bounded. Since time is discrete, this then implies that the search for an optimal driver attendance time is over a finite set and can be done exhaustively.

Proposition 3.2 *The optimal driver attendance time is bounded above and below.*

Proof. Recall that the cost for not completing the trip by time T is assumed infinite and that we assume the existence of a path from any node to the goal node set with finite cost when the trip begins at any $t < \hat{T}$. It should be clear that the optimal driver attendance, say t_0^* , is less than T .

Consider any path from n_0 to the goal node set Γ . Let t_w and t_{tr} represent the (policy dependent) cumulative amount of time during the trip that the driver is made to wait at any nodes and travelling

between nodes, respectively. Let $\mu \equiv \min\{\delta, \delta_1, \delta_2, \delta_3\}$. Note $t_K = t_0 + t_w + t_{tr}$, and let $C(t_0)$ denote the cost of a trip from n_0 to Γ that begins at t_0 . Consider the following three cases.

[Case I] : $t_K \in \{\dots, t_e - 2, t_e - 1\}$.

$$\begin{aligned}
C(t_0) &= \delta_1(t_w) + \delta(t_{tr}) + \delta_3(t_e - t_K) \\
&= \delta_1(t_w) + \delta(t_{tr}) + \delta_3(t_e - (t_0 + t_w + t_{tr})) \\
&\geq \min\{\delta, \delta_1, \delta_3\}(t_e - t_0) \\
&\geq \mu(t_e - t_0).
\end{aligned} \tag{3.3}$$

[Case II] : $t_K \in \{t_e, \dots, t_\ell\}$.

$$\begin{aligned}
C(t_0) &= \delta_1(t_w) + \delta(t_{tr}) \\
&\geq \min\{\delta, \delta_1\}(t_w + t_{tr}) \\
&= \min\{\delta, \delta_1\}(t_K - t_0) \\
&\geq \min\{\delta, \delta_1\}(t_e - t_0) \\
&\geq \mu(t_e - t_0).
\end{aligned} \tag{3.4}$$

[Case III] : $t_K \in \{t_\ell + 1, t_\ell + 2, \dots\}$.

$$\begin{aligned}
C(t_0) &= \delta_1(t_w) + \delta(t_{tr}) + \delta_2(t_K - t_\ell) \\
&\geq \delta_1(t_w) + \delta(t_{tr}) \\
&\geq \mu(t_e - t_0).
\end{aligned} \tag{3.5}$$

Note that since the travel times are bounded, t_{tr} is bounded. Moreover, t_e is fixed. Thus, as $t_0 \rightarrow -\infty$, either $\delta_1(t_w) \rightarrow \infty$ (so $C(t_0) \rightarrow \infty$) or $\delta_3(t_e - t_K) \rightarrow \infty$ (so $C(t_0) \rightarrow \infty$). In particular, the optimal cost in each model, $u(n_0, t_0)$, $v^*(n_0, t_0)$ and $v^*(n_0, t_0, z_0)$ approaches ∞ . Suppose we

choose any path with finite cost, say G , and choose t'_0 such that $t'_0 < t_e - \frac{G}{\mu}$. Then, for all $t_0 \leq t'_0$, we have

$$C(t_0) \geq \mu(t_e - t_0) \geq \mu(t_e - t'_0) > G. \quad (3.6)$$

Since $C(t_0)$ is bounded below by $\mu(t_e - t_0)$ as t_0 varies by the aforementioned three cases, we are guaranteed that the search for an optimal driver attendance time is over the finite set, $t_0^* \in \{t'_0, t'_0 + 1, \dots, T - 1\}$, and can be done exhaustively. Therefore, the result is proven. ■

In the next section we show how the results of Proposition 3.2 regarding the optimal driver attendance time can be used to compute the time the driver actually leaves the origin.

3.2.1 Algorithms determining optimal departure times

In practice it is not possible to predict in advance what $Z(t)$ will be on a specific day. Thus, as has been explained, we minimize the expectation and bring in the driver at the time, t_{r1}^* , minimizing $E\{v^*[n_0, t, Z(t)]\}$ for all t . When the congestion status of the network is observed, we may find it advantageous to delay the actual departure based on the conjectured traffic congestion status in coming periods. In an extreme case, for example, if there is no congestion, say $z_0 \leq z$ for all z , then $v^*[n_0, t_{r1}^*, z_0] \leq E\{v^*[n_0, t_{r1}^*, Z(t_{r1}^*)]\}$ and we may have $\pi^*(n_0, t_{r1}^*, z_0) = n_0$. That is, the actual departure is delayed, perhaps due to the fact that there is a large probability that good road conditions remain the same for several time periods, and the required arrival times will still be met. In the other extreme, if all links are congested, say $z_0 \geq z$ for all z , then $v^*[n_0, t_{r1}^*, z_0] \geq E\{v^*[n_0, t_{r1}^*, Z(t_{r1}^*)]\}$ and we may have $\pi^*(n_0, t_{r1}^*, z_0) \neq n_0$. That is, it is not advantageous to delay the departure, so the optimal action is to leave immediately at t_{r1}^* .

Since the motivation to delay departure is based on the observation of the congestion status, the following result is immediate.

Proposition 3.3 *When real-time traffic information is not available, it is not advantageous to delay the departure after we bring in the driver to the origin. That is, the optimal driver attendance*

time is equal to the optimal departure time when we use only commercial logistics software or historical traffic data.

We now develop algorithms to compute the optimal departure times to minimize total cost. For an arbitrary fixed start time, say \bar{t} , choose the minimum cost $G \equiv E\{v^*[n_0, \bar{t}, Z(\bar{t})]\}$. As is described in the proof of Proposition 3.2, select t'_0 such that $t'_0 < t_e - \frac{G}{\mu}$. The optimal driver attendance time is then between t'_0 and T and can be obtained by finding the largest t_{r1}^* satisfying $E\{v^*[n_0, t_{r1}^*, Z(t_{r1}^*)]\} \leq E\{v^*[n_0, t, Z(t)]\}$ for all $t \in \{t'_0, t'_0 + 1, \dots, T - 1\}$. After the driver has been scheduled and the congestion status z has been observed, we must decide when the vehicle should actually leave the origin. Let t_1^* denote the optimal departure time. It follows that t_1^* is the smallest time greater than or equal to t_{r1}^* satisfying $\pi^*(n_0, t_1^*, z) \neq n_0$. Table 2 summarizes a procedure to dynamically determine an optimal departure time.

In practice, the optimal driver attendance using only commercial logistics software (or t_c^*) or historical traffic data (t_{h1}^*) may be reasonable choices for the arbitrary time \bar{t} in Algorithm 1 since both provide lower bounds on t_{r1}^* .

Theorem 3.4 *The time t_1^* , obtained from ALGORITHM 1, is an optimal departure time. Furthermore, the algorithm terminates after a finite number of iterations.*

Proof. For an arbitrary fixed start time, say \bar{t} , we first choose the minimum cost $G \equiv E\{v^*[n_0, \bar{t}, Z(\bar{t})]\}$, and then select t'_0 such that $t'_0 < t_e - \frac{G}{\mu}$. We have by (3.6)

$$v^*(n_0, t'_0, z_0) \geq \mu(t_e - t'_0) > G,$$

i.e., $v^*(n_0, t'_0, z_0)$ is bounded below for any vector z_0 as t'_0 varies. So $E[v^*(n_0, t, Z(t))]$ is also bounded below for $t \in \{t'_0, t'_0 + 1, \dots, T - 1\}$. Thus, the largest optimal driver attendance time, t_{r1}^* , that satisfies $E\{v^*[n_0, t_{r1}^*, Z(t_{r1}^*)]\} \leq E\{v^*[n_0, t, Z(t)]\}$ for all $t \in \{t'_0, t'_0 + 1, \dots, T - 1\}$, exists and can be found exhaustively (see Proposition 3.2).

By construction $t_{r1}^* \leq t_1^*$ and t_1^* is optimal since it is a solution of (1.2). That is, the actual departure time is later than or equal to the optimal driver attendance time. Moreover, since

ALGORITHM 1.**How to determine an optimal departure time, t_1^* .**

1. For an arbitrary fixed start time, say \bar{t} , let $G \equiv E\{v^*[n_0, \bar{t}, Z(\bar{t})]\}$.
2. Select t'_0 such that $t'_0 < t_e - \frac{G}{\mu}$.
3. Compute the largest $t_{r_1}^*$ satisfying $E\{v^*[n_0, t_{r_1}^*, Z(t_{r_1}^*)]\} \leq E\{v^*[n_0, t, Z(t)]\}$ for all $t \in \{t'_0, t'_0 + 1, \dots, T - 1\}$. That is, $t_{r_1}^*$ is the optimal driver attendance time. Bring in the driver at $t_{r_1}^*$.
4. Observe the congestion status z at time $t_{r_1}^*$.
5. Using the optimality equations (1.1) and (1.2), if $t_{r_1}^*$ satisfies $\pi^*(n_0, t_{r_1}^*, z) \neq n_0$, then start the trip at $t_{r_1}^*$.
6. Otherwise, wait for one more time unit (i.e., $t_{r_1}^* \leftarrow t_{r_1}^* + 1$). Go to 4 until 5 is satisfied.

Table 2: ALGORITHM 1 determining an optimal departure time minimizing total cost

the set $\{t_{r1}^*, t_{r1}^* + 1, \dots, T - 1\}$, where decisions are made, is finite, an optimal departure time that minimizes total cost is determined with certainty and the algorithm terminates after a finite number of iterations. ■

3.3 Optimal vehicle usage times

A second way to see the value of real-time congestion information is to compare the reduced time that the vehicle is actually used when this information is available. We formally define the optimal vehicle usage time as the driver attendance time that achieves at most the same cost as the minimal cost in the case with commercial logistics software; see Figure 3.

Definition 3.5

(i) *If historical traffic data are used to compute the shortest path as described by equation (3.2), the optimal vehicle usage time is the largest t_{h2}^* , such that*

$$v^*(n_0, t_{h2}^*) \leq u(n_0, t_c^*).$$

(ii) *If both historical and real-time traffic data are used as described by equation (1.1), the optimal vehicle usage time is the largest t_{r2}^* , such that*

$$E\{v^*[n_0, t_{r2}^*, Z(t_{r2}^*)]\} \leq u(n_0, t_c^*).$$

Since $E\{v^*[n_0, t, Z(t)]\} \leq u(n_0, t)$ for all t when real-time traffic information is available, the optimal vehicle usage time is between t_c^* and T and can be obtained by finding the largest t_{r2}^* satisfying

$$E\{v^*[n_0, t_{r2}^*, Z(t_{r2}^*)]\} \leq u(n_0, t_c^*).$$

where $t_{r2}^* \in \{t_c^*, t_c^* + 1, \dots, T - 1\}$. This is illustrated in Figure 3. In the next section we present algorithms for computing the actual departure time that minimizes vehicle usage.

3.3.1 Algorithms determining the optimal departure time minimizing vehicle usage

After the driver has been scheduled, we must decide when the vehicle should actually leave the origin. Note in this case, the driver attendance time has been computed with the objective of

maintaining the same cost as the case with only commercial logistics software. Let t_2^* denote the optimal departure time. It follows that t_2^* is the smallest time greater than or equal to t_{r2}^* satisfying $\pi^*(n_0, t_2^*, z) \neq n_0$ for given z at time t_2^* . Table 3 summarizes a procedure to dynamically determine an optimal departure time after we bring in the driver at t_{r2}^* .

ALGORITHM 2.

How to determine an optimal departure time, t_2^* , minimizing vehicle usage

1. Search for the largest t_{r2}^* satisfying $E\{v^*[n_0, t_{r2}^*, Z(t_{r2}^*)]\} \leq u(n_0, t_c^*)$ where $t_{r2}^* \in \{t_c^*, t_c^* + 1, \dots, T - 1\}$. Bring in the driver at t_{r2}^* .
2. Observe the congestion status z at time t_{r2}^* .
3. Using the optimality equations (1.1) and (1.2), if $\pi^*(n_0, t_{r2}^*, z) \neq n_0$, then start the trip at t_{r2}^* .
4. Otherwise, wait for one more time unit (i.e., $t_{r2}^* \leftarrow t_{r2}^* + 1$). Go to 2 until 3 is satisfied.

Table 3: ALGORITHM 2 determining an optimal departure time minimizing vehicle usage

Theorem 3.6 *The time t_2^* , obtained from ALGORITHM 2, is an optimal departure time after we bring in the driver at the optimal vehicle usage time. Furthermore, the algorithm terminates after a finite number of iterations.*

Proof. The result follows from a similar argument as that in Theorem 3.4, so the proof is omitted.

■

In what follows we use the results of this section to compare the value of real time IT on a network in southeast Michigan.

4 Numerical Evaluation

We investigated 10 origin and destination pairs in Southeast Michigan to calculate the total cost savings and the reduction in vehicle usage by using historical and real-time traffic information. Figure 4 shows one such road network. In the road network, the highlighted path connecting the origin (1) and destination (2) is the best route generated by the commercial logistics software.



Figure 4: An example of the origin and destination pair analyzed

In order to compare the cost computed using commercial software to the dynamic model that we propose, we examined 5 shipping time slots (6am-9am, 9am-12pm, 12pm-3pm, 3pm-6pm, 6pm-6am) and randomly selected the desired arrival time window in each shipping time zone. We have considered 3 ratios of the late penalty to the travel cost (1:1, 10:1, 100:1). These ratios are likely to be dependent on the specific contracts based on the typical field of industry. We assume these ratios are uniformly distributed in Southeast Michigan. In order to evaluate the benefits of historical and real-time traffic information, 30 different cases were analyzed in each of 5 shipping time slots (i.e., 10 origin and destination pairs \times the 3 aforementioned late penalty to travel cost ratios).

4.1 Benefits of historical and real-time traffic information

4.1.1 Cost savings

Table 4 shows cost savings (%) by using historical and real-time traffic information compared with the case using commercial logistics software over different delivery time zones.

	Cost Savings (%)		
	Historical Data	Real-Time Data	Total
6am - 9am	4.39	2.57	6.96
9am - 12pm	8.20	2.35	10.55
12pm - 3pm	7.16	1.45	8.61
3pm - 6pm	4.07	3.65	7.72
6pm - 6am	5.30	0.61	5.91

Table 4: Cost savings by historical and real-time traffic information over time

For example, between 6am to 9am the percentage savings in total cost by using historical traffic data compared to the base case with commercial logistics software is 4.39(%). An additional savings of 2.57(%) can be achieved by using real-time traffic information together with historical traffic data. Figure 5 is a graph of the results in Table 4. The cost savings due to real-time traffic information make up about 37(%) of the total cost savings during rush hours in the morning and about 47(%) of total cost savings during rush hours in the afternoon. However, when the traffic volume is relatively low, for example, between 6pm and 6am, cost savings due to real-time traffic information is only 10(%).

This exhibits the intuitive idea that real time information can be quite useful during times of potential heavy congestion like during rush hour times and less useful when the traffic volumes are low. By analyzing the results, we suggest that an appropriate level of traffic information should be provided for each trucking company. For example, if a package delivery company usually ships packages at night, expensive real-time traffic information may not be warranted. However, for a trucking company, that is responsible for the just-in-time delivery of products to automobile assembly plants arriving in the morning or afternoon rush hours, a real-time traffic information

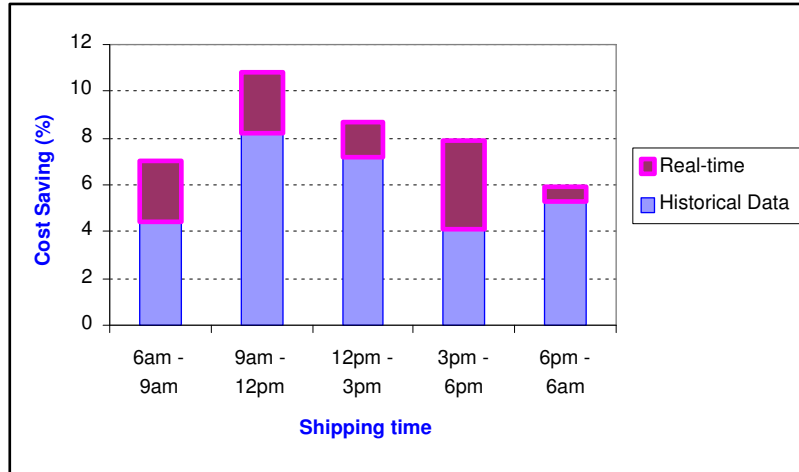


Figure 5: Visualization of Table 4

system can provide a significant payoff.

4.1.2 Vehicle usage reduction

Similarly, Table 5 shows the reduction of vehicle usage (%) by using historical and real-time traffic information over different time zones achieving at most the same cost as the minimal cost in the case with commercial logistics software.

	Reduction in vehicle usage (%)		
	Historical Data	Real-Time Data	Total
6am - 9am	7.42	2.92	10.34
9am - 12pm	11.78	4.41	16.19
12pm - 3pm	10.18	1.66	11.84
3pm - 6pm	5.11	6.88	11.99
6pm - 6am	7.65	2.17	9.82

Table 5: Reduction in vehicle usage by historical and real-time traffic information over time

Figure 6 is a graph of the results in Table 5. The results show that the vehicle usage reduction due to real-time traffic information is about 28(%) of the total reduction in vehicle usage during rush hours in the morning. During rush hours in the afternoon the vehicle usage reduction due to real-time traffic information constitutes about 58(%) of total. Even when the traffic volume

is relatively low, for example, between 6pm and 6am, the reduction in vehicle usage due to real-time traffic information is approximately 22(%). This implies that no matter what time of day (during rush hours or not), the real-time traffic information may play a major role in vehicle usage reduction.




Figure 6: Visualization of Table 6

5 Conclusions

This paper provides a systematic approach to aid in the implementation of transportation systems integrated with real time information technology and develops decision-making procedures for determining the optimal driver attendance time, optimal departure times, and optimal routing policies under time-varying traffic flows. Our primary conclusion is that real-time traffic information incorporated with historical traffic data can significantly reduce expected total costs and vehicle usage during times of potential heavy congestion while satisfying or improving service levels for just-in-time delivery. We remark that implementation of only historical traffic data is easier than incorporating real-time IT. However, only a small amount of additional effort and investment may be required to achieve a complete implementation of real-time vehicle routing and scheduling. All required calculations such as the optimal driver attendance time, optimal departure times, and routing policies can be done off-line in advance. A central dispatcher or a logistics manager would

observe real-time traffic information and communicate to the commercial vehicle updated routing instructions. Of course, to fully reflect recent traffic flows, one would need to periodically update the database of historical traffic information.

Although we have formulated the problem with the congestion status of each link being modelled as a two-state Markov chain, the methods are extendable to the case when the congestion status of each link can be in one of $N \geq 2$ states. The definitions of congestion status are likely to be problem specific based on the typical travel times along each link. Moreover, in order to make the problem more tractable we have made the assumption that the link travel times are independent. Although this may not perfectly reflect reality, we believe that this is much needed first step in the valuation of real-time traffic information. Relaxing this assumption is an interesting (and potentially quite difficult) extension.

We also remark that the assumption that the unobserved links have deterministic cost functions is only a restriction in the sense that we require that the cost functions are stationary. If the cost is to be a random amount, we would then use the expected cost along each link as the deterministic cost function. Finally, we mention that when the number of observed links with real-time traffic information increases, the offline calculations can be computationally intractable. This computational challenge is in large part due to the amount of data available that may be useful for optimal route selection. As has been mentioned, developing algorithms for efficiently reducing the amount of data to improve computation of a Markov decision process is the subject of our subsequent work (see Kim et al. (2004)).

References

- G. Andreatta and L. Romeo, “Stochastic Shortest Paths with Recourse,” *Networks* 18, 193–204 (1988).
- J. Bander and C. White III, “A Heuristic search approach for a nonstationary stochastic shortest path problem with terminal cost,” *Transportation Science* 36, 218–230 (2002).

- D. P. Bertsekas and Tsitsiklis, “An Analysis of Stochastic Shortest Path Problems,” *Mathematics of Operations Research* 16, 580–595 (1991).
- J. S. Croucher, “A Note on the Stochastic Shortest-Route Problem,” *Naval Research Logistics Quarterly* 25, 729–732 (1978).
- A. Eiger, P. Mirchandani, and H. Soroush, “Path preferences and optimal paths in probabilistic networks,” *Transportation Science* 19, 75–84 (1985).
- L. Fu, “An adaptive routing algorithm for in vehicle route guidance systems with real-time information,” *Transportation Research Part B* 35, 749–765 (2001).
- L. Fu and L. Rilett, “Expected Shortest Paths in Dynamic and Stochastic Traffic Networks,” *Transportation Research Part B* 32, 499–516 (1998).
- S. Gao and I. Chabini, “Best routing policy problem in stochastic time-dependent networks,” *Transportation Research Record* 1783, 188–196 (2002).
- R. Hall, “The fastest path through a network with random time-dependent travel times,” *Transportation Science* 20, 182–188 (1986).
- D. Kaufman and R. Smith, “Fastest paths in time-dependent networks for intelligent vehicle highway system application,” *IVHS Journal* 1 (1993).
- S. Kim, M. Lewis, and C. White III, “The impact of real-time traffic congestion information on the trucking industry,” in *Proceedings of the 9th World Congress on Intelligent Transport Systems*, 2002.
- S. Kim, M. Lewis, and C. White III, “State Space Reduction for Non-stationary Stochastic Shortest Path Problems with Real-Time Traffic Information,” (2004), submitted for publication.
- E. Miller-Hooks and H. Mahmassani, “Least possible time paths in stochastic, time-varying transportation networks,” *Computers and Operations Research* 25, 1107–1125 (1998).

- E. Miller-Hooks and H. Mahmassani, “Least expected time paths in stochastic, time-varying transportation networks,” *Transportation Science* 34, 198–215 (2000).
- J. Pearl, *Heuristics : Intelligent search strategy for computer problem solving*, Addison-Wesley, Massachusetts (1984).
- G. H. Polychronopoulos and J. N. Tsitsiklis, “Stochastic Shortest Path Problems with Recourse,” *Networks* 27, 133–143 (1996).
- H. Psaraftis and J. Tsitsiklis, “Dynamic shortest path in acyclic network with Markovian arc costs,” *Operations Research* 41, 91–101 (1993).
- M. Puterman, *Markov decision processes : Discrete stochastic dynamic programming*, Wiley, New York (1994).
- S. Ross, *Stochastic processes*, Wiley, New York (1996).
- S. T. Waller and A. K. Ziliaskopoulos, “On the online shortest path problem with limited arc cost dependencies,” *Networks* 40, 216–227 (2002).
- M. P. Wellman, M. Ford, and K. Larson, “Path planning under time-dependent uncertainty,” in *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, 532–539, 1995.