

MIT Open Access Articles

Stochastic motion planning and applications to traffic

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Lim, Sejoon et al. "Stochastic Motion Planning and Applications to Traffic." Algorithmic Foundation of Robotics VIII. Ed. Gregory S. Chirikjian et al. Vol. 57. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009. 483-500.

As Published: http://dx.doi.org/10.1007/978-3-642-00312-7_30

Publisher: Springer Berlin/Heidelberg

Persistent URL: <http://hdl.handle.net/1721.1/72494>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Stochastic Motion Planning and Applications to Traffic

Sejoon Lim, Hari Balakrishnan, David Gifford, Samuel Madden, and Daniela Rus

CSAIL, MIT, {sjlim,hari,gifford,madden,rus}@csail.mit.edu

Abstract: This paper presents a stochastic motion planning algorithm and its application to traffic navigation. The algorithm copes with the uncertainty of road traffic conditions by stochastic modeling of travel delay on road networks. The algorithm determines paths between two points that optimize a cost function of the delay probability distribution. It can be used to find paths that maximize the probability of reaching a destination within a particular travel deadline. For such problems, standard shortest-path algorithms don't work because the optimal substructure property doesn't hold. We evaluate our algorithm using both simulations and real-world drives, using delay data gathered from a set of taxis equipped with GPS sensors and a wireless network. Our algorithm can be integrated into on-board navigation systems as well as route-finding Web sites, providing drivers with good paths that meet their desired goals.

1 Introduction

This paper presents and evaluates an algorithm for planning the motion of vehicles (autonomous or human-driven) on roadways in the face of traffic delays. Rather than model road delays statically, as in current on-board navigation systems and Web-based mapping services, our algorithm uses past observations of actual delays on road segments to model these delays as probability distributions. The algorithm minimizes a user-specified cost function of the delay distribution. We investigate a few cost functions in detail, particularly one that is equivalent to maximizing the likelihood of reaching a destination within a specified travel deadline.

Our work provides a planning system that can be used by robots as well as human drivers. The system is a useful addition to on-board navigation systems using computer-aided automation to provide good paths that meet desired travel goals (e.g., “when should you leave, and what path should you take, to reach the airport by 8am with high probability?”); it is also a worthwhile addition to Web-based mapping services. We view the incorporation of traffic-aware path computation as an important practical addition in the rapid trend toward computer-assisted driving and autonomous decision-making in vehicles.

Traffic congestion is clearly a serious problem: a recent survey [11] estimates that the annual nationwide cost of traffic congestion is \$78 billion, including 4.2 billion hours in lost time and 2.9 billion gallons in wasted fuel. Drivers today have little knowledge of historic and real-time traffic congestion on the paths they drive, and

even when they do (e.g., from “live” traffic updates), they generally don’t know how to use that information to find good paths. As a result, they often tend to drive sub-optimal routes and often leave well in advance when they need to make an important deadline.

In addition to helping individual cars avoid congested roads, we believe that our work, if deployed widely, can manage traffic flow, reduce congestion, and reduce the fuel consumed by cars on a macroscopic basis by using the under-utilized parts of the road network better than today (thereby reducing load on congested areas). Using our algorithm to investigate this global traffic management question is an area for future work. In this paper, we are concerned with finding good paths for a single car.

The main challenge in planning paths taking traffic delay into account is that these delays are not fixed. The delay on a road segment is best modeled as a probability distribution; in addition, this distribution typically depends on a number of factors, such as time-of-day, whether it’s a working day or not, events such as concerts or sporting events, weather, etc. The shortest-distance path is often not the best path to use if one seeks to minimize the expected travel time or maximize the probability of reaching the destination by a certain time. Our algorithm uses historic observations of travel delays on road segments at different times of day to produce delay distributions (indexed by time-of-day). We posit that this information, together with real-time updates of extraneous conditions (such as accidents), is invaluable (and sufficient) to compute good paths that meet user-specified goals. Given the probability distributions of delays on segments, finding good paths requires more than a shortest-path computation, because the “optimal substructure” property does not hold as explained in [8] (i.e., if the best path from S to T goes through X , it does not follow that the sub-path of this path from S to X is itself the best S - X path).

We have implemented our algorithm and evaluated it by first modeling the historic delays using data from the CarTel vehicular testbed [5], a network of 28 taxis. The data consists of travel times organized by road segment and by time-of-day, yielding statistical profiles for all the road segments. We model the road network as a weighted graph where the nodes represent intersections and the edges represent road segments. An aggregation algorithm combines the road segments into groups to coalesce the important delay characteristics without losing information about alternate paths. Our algorithm has the flavor of searching and pruning the delay statistics on the road network data structure. We evaluate the algorithm and its assumptions using simulation and actual test driving.

1.1 Related Work

Our work is closely related to stochastic planning and stochastic shortest-path algorithms. In [3], an efficient algorithm for a dynamic shortest path with time dependent deterministic edge cost is given. Prior work has considered stochastic shortest paths under the assumption that the travel time follows a known probability distribution [12, 1, 6, 13, 7, 8, 9]. In stochastic shortest path, edge costs are probability distributions rather than deterministic costs and the optimal path depends on drivers’

diverse objectives. When a driver's objective is to minimize the expected travel time, the problem becomes the deterministic shortest-path problem. It is well known that Dijkstra's algorithm is optimal and applicable for deterministic problems. However, for some goals such as maximizing the probability of arriving within a given deadline, the optimal path cannot be found with the standard shortest-path algorithms since the optimal substructure property does not hold. Nikolova et al. [9] developed an algorithm and theoretical bounds by assuming that delays are both Gaussian and independent on different road segments. Inspired by this algorithm, we developed a method that improves performance by removing unnecessary invocations of shortest-path searches.

There have been several approaches to acquiring traffic data. The most prevalent one uses inductor loops installed beneath roads [14, 4]. This is adequate for counting the number of cars that pass a specific location, but it is not suitable for measuring travel time and measurements are possible only on instrumented roads. Recently, GPS sensors installed in probe vehicles have been used [10, 5, 15]. The travel time of vehicles can be measured and recorded for each route segment. In [5], the researchers developed a system called CarTel Network nodes that include GPS and wireless communication. This system was used to study routing and data delivery from cars.

1.2 Outline

This paper is organized as follows. Section 2 gives an overview of a route planning system for traffic. Section 3 gives the stochastic motion planning formulation, and Section 4 presents our algorithm and gives its correctness and performance bounds. Section 5 evaluates the algorithm in simulation and also using physical data from the deployment of GPS sensor nodes in taxis.

2 Transportation System Context

Our research objective is to provide an effective navigation system for autonomous or human-piloted cars that uses historical and real-time traffic data to determine optimal driving directions and traffic estimates. Our intelligent navigation system consists of:

1. a data gathering system included in cars that move in traffic frequently;
2. a data analysis system to compile a historical database of traffic conditions;
3. an algorithm for route planning that uses both historical data and current information;
4. a traffic information system implementation with an appropriate interface.

In [5] a system called CarTel was developed that uses GPS and wireless communication to collect position and time data from cars. CarTel nodes deployed in 28 taxis since January 2007 collected travel data. This data was organized per road segment to create a historical database of traffic delays.

A stochastic route planning algorithm was developed and implemented. The best route depends on the drivers' goals and is a combination of speed and reliability. A

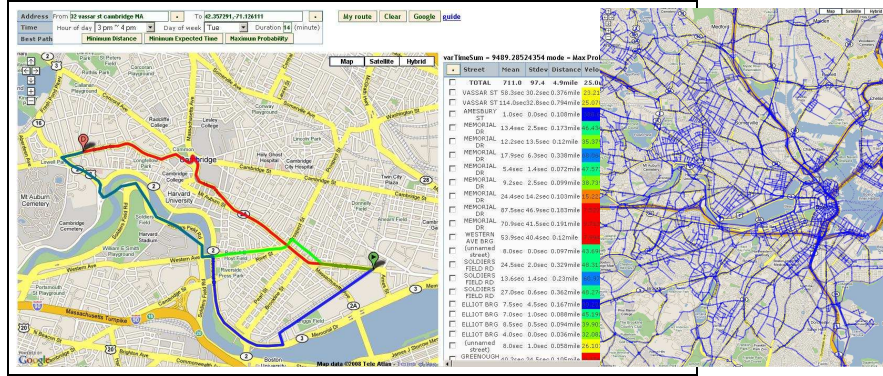


Fig. 1. The user interface to the traffic system with highlighted paths and travel times found by the algorithm described in this paper. Roads with traffic travel time data in our database are also shown in the right.

Web-based interface (Fig. 1) allows users to query the system for traffic conditions and for optimal paths given historical data. The rest of this paper describes the motion planning component 3. of this intelligent traffic system.

3 Problem Formulation

3.1 Road Network Modeling

The road network is a graph (called the Geographic Map), where nodes represent intersections and edges represent road segments. We associate a road delay distribution with each road segment (Fig. 1). This per-intersection granularity road map leads to a large graph for small road segments with related travel statistics. We combine statistically related road segments into groups so that they can capture important delay characteristics without losing information about alternate path segments. This data structure is the Delay Statistics Map. The Geographic Map is used for matching GPS traces onto real road segments, while the Delay Statistics Map is used for statistical delay sensitive routing. We assume that

1. the delay of each edge follows a Gaussian distribution;
2. the delay of each edge is independent of every other edge.

In Section 5, we provide evidence for these assumptions. We formulate stochastic motion planning as a graph search problem over a graph with an origin O and a destination D , where the travel time of each edge is an independent Gaussian random variable. Since the sum of independent Gaussian random variables is also a Gaussian random variable, we can denote the travel time for a path π consisting of edges e of mean m_e and variance v_e as follows:

$$t_\pi \sim \mathcal{N}(m_\pi, v_\pi), \text{ where } m_\pi = \sum_{e \in \pi} m_e \text{ and } v_\pi = \sum_{e \in \pi} v_e.$$

3.2 Cost Functions

Our objective is to find a path that minimizes an expected cost when the cost function models a user's goal. We call this the "optimal" path for the given cost function. We consider several cost functions including:

Linear cost: Here, the cost increases linearly with the travel time. When the cost of arriving at the destination in time t is $C(t) = t$ and the delay PDF of a path π is $f_\pi(t)$, the expected cost of traveling through π is $EC_\pi = \int_{-\infty}^{\infty} t f_\pi(t) dt = m_\pi$. Linear cost models the path with minimum expected time.

Exponential cost: Exponential cost models a cost function that increases sharply as the arrival time increases. When the cost function is $C(t) = e^{kt}$, where k is the steepness of the cost increase, the expected cost can be written as $EC_\pi = \int_{-\infty}^{\infty} e^{kt} f_\pi(t) dt = \{e^{k(m_\pi + \frac{kv_\pi}{2})}\}$. This exponential cost function minimizes a linear combination of mean and variance determined by k .

Step cost: Step cost models a cost that only penalizes the late arrival after a given deadline. The cost function is $C(t, d) = u(t - d)$, where $u(\cdot)$ is the unit step function and d is the deadline. The expected cost is $EC_\pi(d) = \int_{-\infty}^{\infty} u(t - d) f_\pi(t) dt = \int_d^{\infty} f_\pi(t) dt = \{1 - \Phi(\frac{d - m_\pi}{\sqrt{v_\pi}})\}$, where $\Phi(\cdot)$ is the CDF of the Standard Normal distribution. Thus, when Π is a set of all paths from O to D , the minimum expected cost path is $\operatorname{argmax}_{\pi \in \Pi} \Phi(\frac{d - m_\pi}{\sqrt{v_\pi}})$, which turns out to be the path that maximizes arrival probability. Since $\Phi(\cdot)$ is monotonically increasing, maximizing $\Phi(\cdot)$ is equivalent to maximizing

$$\varphi_d(\pi) = \frac{d - m_\pi}{\sqrt{v_\pi}}. \quad (1)$$

The minimum expected cost path for the linear and exponential cost cases can be found by a deterministic shortest-path algorithm, such as Dijkstra's or A^* search algorithm since the cost of a path can be expressed as the sum of the cost of each edge in the path. However, when the cost is a step function, these algorithms cannot be used since the objective, (1), is nonlinear. Our goal for the rest of the paper is to develop an efficient algorithm for finding the maximum probability path given a deadline.

4 Stochastic Path Planning by Parametric Optimization

In [9], an algorithm for the case of normally distributed edge costs was given based on quasi-convex maximization. It finds the path with the maximum arrival probability by standard shortest-path runs with different edge costs corresponding to varying parameters. We now give a graphical interpretation of the optimal path and show a connection to a parametric optimization problem, which will ultimately lead to a new algorithm that reduces unnecessary runs over [9].

4.1 Transforming the Cost Function into Parametric Form

Let path π be denoted by a point (m_π, v_π) in a rectangular coordinate system, called the $m - v$ plane, where the horizontal axis represents the mean and the vertical

axis represents the variance. The objective of the optimization problem, (1), can be rewritten to show the relation between m_π and v_π as

$$v_\pi = \frac{1}{\varphi_d(\pi)^2} (m_\pi - d)^2, \quad (2)$$

which is a parabola in the $m - v$ plane with apex at d , where $\varphi_d(\pi)$ is determined by the curvature of the parabola. Thus, the optimal path is the path that lies on the parabola of the smallest curvature. Intuitively, the optimization problem is to find the first path that intersects the parabola while we lift up the parabola starting from the horizontal line (see Fig. 2 (Left)). This suggests finding the optimal path using linear optimization with various combinations of cost coefficients.¹

Consider setting the cost of an edge to be linear combinations of mean and variance, $m_e + \lambda v_e$, for an arbitrary non-negative λ . We call the solution for this edge cost the λ -optimal solution. This edge cost follows the optimal substructure property and has the property described in Lemma 1, which was also stated in [9].

Lemma 1. *An optimal path occurs among the extreme points of the convex hull for all the O to D paths in the $m - v$ plane if there exists a path that has a mean travel time smaller than the deadline.*

Proof. Let point P on the $m - v$ plane represent the optimal path. Then, there is no path point that has the φ value larger than that of P . Therefore, every other path point must be inside the parabola. Since the parabola is convex, P must be an extreme point.

With Lemma 1 we can find the optimal solution from λ -optimal solution for a given λ . Since λ -optimal cost satisfies the optimal substructure property, any deterministic shortest-path algorithm (e.g., Dijkstra’s algorithm or A^* search algorithm) will find λ -optimal paths.

4.2 Exhaustive Enumeration

In [9], a method for stochastic motion planning was proposed that exhaustively enumerates all the extreme points of the path convex hull. A brief description of the algorithm is as follows: First, find the λ -optimal paths for $\lambda = 0$ and $\lambda = \infty$. If they are the same, it must be the optimal solution. Otherwise, find the λ -optimal path using $\lambda = -\frac{m_0 - m_\infty}{v_0 - v_\infty}$ since this λ value will cause the algorithm to search the entire region completely unless it finds a new path, as illustrated in Fig. 2 (Middle). If no new path is found, the algorithm terminates with the optimal solution being the one with the largest φ value. Otherwise, the newly found path divides the search region into two parts. Then, the λ -optimal search is executed for each region using λ values determined to search each region completely. In this approach, when the number of extreme points is N_e , there will be N_e searches to guarantee that all the extreme

¹ Linear optimization finds a path that first intersects a straight line when the line is moved in a direction determined by cost parameters.




Fig. 2. (Left) Graphical interpretation of the optimal path in the $m - v$ plane. Each square represents a path from the origin to the destination. Equi-probability paths lie on a parabola with an apex at $(d, 0)$ and a curvature of $\frac{1}{\varphi_d(\pi)^2}$. The optimal path is the first point that meets with a parabola as we increase the curvature. (Middle) The result after three executions of λ -optimal searches with $\lambda_1 = 0$, $\lambda_2 = \infty$, and $\lambda_3 = -\frac{m_0 - m_\infty}{v_0 - v_\infty}$. Each square point represents the λ -optimal path for each λ . The gray points represent the paths that are not found yet. The blue regions are guaranteed to contain no path. The white triangles indicate candidate regions for better paths. The round points are the probe points of the regions. (Right) λ -optimal search was done only for the left candidate region. The newly found path turns out to be the new current optimal path and the two round points are the probe points.

points are enumerated. In addition, $N_e - 1$ more searches are needed to conclude that no other paths exist between the extreme points. Thus, the total number of enumerations could be large. Next, we show how to reduce the number of required λ -optimal searches.

4.3 Examining Probe Points

If we know that a certain search region's best possible outcome is worse than the current optimal solution, we do not need to execute the costly λ -optimal search for that region. In this section we formalize this point.

Definition 1. Let the triangular region where a better path can exist be called a candidate region. Let the vertex in the middle of the candidate region be called a probe point. Candidate regions are illustrated as white triangles $\triangle L_i M_i R_i$, and probes as round points M_i in Fig. 2 (Middle).

Theorem 1. If the φ value of the probe point as defined in (1) is smaller than the current optimal value, the candidate region does not contain the optimal path.

Proof. Suppose that a path lies at the probe point. Then, no other point in the candidate region can be an extreme point. The interior points cannot be an optimal solution since the optimal solution occurs at one of the extreme points by Lemma 1. Suppose that a path does not lie at a probe point. Add an imaginary origin-to-destination path that lies on the probe point. The addition of an imaginary path will not make any difference for searching for the optimal solution since it is not better than the current optimal path. The same argument shows that interior points cannot be optimal solutions.

By Theorem 1, we can remove from consideration the candidate region if the region's probe point satisfies the condition in Theorem 1. Fig. 2 (Right) illustrates a case where the right candidate region $\triangle L_2 M_2 R_2$ was removed without any execution of λ -optimal search since the φ value of the probe point M_2 is smaller than that of the current optimal path. The left candidate region $\triangle L_1 M_1 R_1$ was searched since the left probe point gives a larger φ value than the current optimal value, and a new path was found as the λ -optimal path. The same procedure is applied to the new candidate regions built by the newly found λ -optimal path until there is no candidate region remaining.

4.4 Restricting λ by Upper and Lower Bounds

The λ values that should be searched are limited by upper and lower bounds.

Theorem 2. *The optimal path can be found by searching only with the λ values upper bounded by λ_u , the negative inverse of the tangent to the parabola at the intersection of the 0-optimal search line and the ∞ -optimal search line.*

Proof. If the λ -optimal solution is the same as the λ_u -optimal solution for all $\lambda > \lambda_u$, we can trivially find the same path with λ_u instead of $\lambda > \lambda_u$. Suppose that there exists a certain $\lambda > \lambda_u$ for which λ -optimal path (m_λ, v_λ) is different from the λ_u -optimal path $(m_{\lambda_u}, v_{\lambda_u})$. Then, we can say that $m_{\lambda_u} \neq m_\lambda$ and $v_{\lambda_u} \neq v_\lambda$ since $0 < \lambda < \infty$. From the definition of λ -optimal path, $m_{\lambda_u} + \lambda_u v_{\lambda_u} < m_\lambda + \lambda_u v_\lambda$ and $m_{\lambda_u} + \lambda v_{\lambda_u} > m_\lambda + \lambda v_\lambda$. Rewriting these, we get

$$\lambda_u(v_{\lambda_u} - v_\lambda) < m_\lambda - m_{\lambda_u}, \quad \lambda(v_{\lambda_u} - v_\lambda) > m_\lambda - m_{\lambda_u}. \quad (3)$$

From the two inequalities we get $(\lambda - \lambda_u)(v_{\lambda_u} - v_\lambda) > 0$ and since $\lambda > \lambda_u$, it follows that $v_{\lambda_u} > v_\lambda$ and $m_\lambda > m_{\lambda_u}$. We get an expression for λ_u by taking the derivative of (2) and its negative inverse $\lambda_u = -1/\frac{\partial v}{\partial m}|_{m=m_0} = \frac{(d-m_0)^2}{2v_\infty(d-m_0)} = \frac{d-m_0}{2v_\infty}$. From this and (3), and since $d - m_0 > d - m_\lambda$ and $v_\lambda > v_\infty$,

$$\frac{m_\lambda - m_{\lambda_u}}{v_{\lambda_u} - v_\lambda} > \frac{d - m_0}{2v_\infty} > \frac{1}{2} \frac{d - m_\lambda}{v_\lambda}. \quad (4)$$

Since $m_\lambda > m_{\lambda_u}$ and $d - m_\lambda > 0$, $\frac{1}{\frac{m_\lambda - m_{\lambda_u}}{d - m_\lambda} + 2} < \frac{1}{2}$. From this and (4), $\frac{1}{\frac{m_\lambda - m_{\lambda_u}}{d - m_\lambda} + 2} < \frac{m_\lambda - m_{\lambda_u}}{v_{\lambda_u} - v_\lambda} \frac{v_\lambda}{d - m_\lambda}$. We can rewrite this as follows:

$$\begin{aligned} \frac{v_{\lambda_u} - v_\lambda}{v_\lambda} &< \frac{m_\lambda - m_{\lambda_u}}{d - m_\lambda} \left(\frac{m_\lambda - m_{\lambda_u}}{d - m_\lambda} + 2 \right) \rightarrow \frac{v_{\lambda_u}}{v_\lambda} < \left(\frac{m_\lambda - m_{\lambda_u}}{d - m_\lambda} + 1 \right)^2 \\ \rightarrow \frac{v_{\lambda_u}}{v_\lambda} &< \left(\frac{d - m_{\lambda_u}}{d - m_\lambda} \right)^2 \rightarrow \frac{d - m_\lambda}{\sqrt{v_\lambda}} < \frac{d - m_{\lambda_u}}{\sqrt{v_{\lambda_u}}}. \end{aligned}$$

Thus, for any $\lambda > \lambda_u$, the λ -optimal solution is worse than the λ_u -optimal solution. Thus, there is no need to search the area with the λ that is larger than λ_u . Therefore, whether λ -optimal solution is the same with the λ_u -optimal solution or not we can find the optimal solution by searching with $\lambda \leq \lambda_u$.

Algorithm 1: PARAMETRIC-SEARCH

Data: Graph with mean and variance of each edge, origin, and destination
Result: The optimal path

```

1  $bestPath \leftarrow \emptyset$ 
2  $Regions = []$  : FIFO queue containing candidate regions.
3  $path_0 \leftarrow \text{SEARCH-}\lambda\text{-OPTIMAL-PATH}(0)$ 
4  $path_\infty \leftarrow \text{SEARCH-}\lambda\text{-OPTIMAL-PATH}(\infty)$ 
5 if  $path_0 == path_\infty$  then
6   return  $path_0$ 
7  $Regions.push(\text{Region}(l : path_0, r : path_\infty, p : (path_0.mean, path_\infty.var)))$ 
8 calculate  $\lambda_l$  and  $\lambda_u$ 
9 while  $(R \leftarrow Regions.pop()) \neq \emptyset$  do
10   if  $R.probe.\varphi < bestPath.\varphi$  then continue
11    $\lambda \leftarrow -\frac{R.l.mean - R.r.mean}{l.var - r.var}$ 
12   if  $\lambda \geq \lambda_u$  then
13     if  $\lambda_u$  was not searched then  $\lambda \leftarrow \lambda_u$  else continue
14   if  $\lambda \leq \lambda_l$  then
15     if  $\lambda_l$  was not searched then  $\lambda \leftarrow \lambda_l$  else continue
16    $path \leftarrow \text{SEARCH-}\lambda\text{-OPTIMAL-PATH}(\lambda)$ 
17   if  $path \neq R.l$  and  $path \neq R.r$  then
18     if  $path.\varphi > bestPath.\varphi$  then
19        $bestPath \leftarrow path$ , update  $\lambda_l$ 
20     locate  $probe_l$  and  $probe_r$ 
21     if  $probe_l.\varphi > bestPath.\varphi$  then
22        $Regions.push(\text{Region}(l : R.l, r : path, p : probe_l))$ 
23     if  $probe_r.\varphi > bestPath.\varphi$  then
24        $Regions.push(\text{Region}(l : path, r : R.r, p : probe_r))$ 
25 return  $bestPath$ 

```

Theorem 3. *The optimal path can be found by searching only with the λ values lower bounded by λ_l , the negative inverse of the tangent to the current λ -optimal parabola at the intersection of the current λ -optimal parabola and 0-optimal search line.*

Proof. Similar to the proof of Theorem 2.

Theorems 1, 2, and 3 lead to the Parametric Search algorithm (see Algorithm 1) for finding the best route that maximizes the probability of arriving at the destination within a given deadline. In lines 3 and 4, the 0-optimal and ∞ -optimal paths are searched with a shortest-path algorithm (e.g., Dijkstra’s algorithm or A^* search algorithm). If the two found paths are the same, the algorithm terminates. If they are different, the first candidate region consisting of the three points denoted in line 7 is pushed into the queue. Candidate regions are evaluated for searching. The conditions in lines 10, 21, and 23 come from Theorem 1, and those in lines 12 and 14 from Theorems 2 and 3, respectively. If the candidate region does not need to be searched,

the algorithm continues with the next region. Otherwise, the region is searched with the λ value determined by the left and right path of the region (line 11) and possibly modified by the upper and lower bounds (lines 13 and 15) in line 16.

4.5 Correctness

Algorithm 1 finds the optimal solution in a finite number of λ -optimal searches. The paths in the region we exclude from the exhaustive enumeration using the extreme points cannot be optimal by Theorem 1. The paths in the region we excluded using the upper and lower bound of λ cannot be optimal due to Theorems 2 and 3. Since the number of required λ -optimal searches is upper bounded by $2N_e - 1$ as described in Section 4.2, the algorithm finds the optimal solution in a finite number of searches.

4.6 Running Time

As shown by Nikolova, et al. in [9] based on [2], there are total $N^{\Theta(\log N)}$ extreme points in the $m - v$ plane, where N is the number of nodes in the network. Compared to their algorithm that searches every extreme point, our algorithm does not invoke unnecessary searches yielding an average running time of $O(N^2 \log^4 N)$, where N^2 term is due to Dijkstra's runs for each λ -optimal search.

The intuition is as follows. Each new path point found with a λ -optimal search yields two candidate regions. Our algorithm only searches the candidate region if its probe point is outside the current optimal parabola. The parabola passing through the newly found path point will almost always divide the two probe points causing one of them to lie outside the parabola. The adverse case where both probe points are outside the parabola happens rarely when the current optimal parabola meets the current λ -optimal search line twice within the interval determined by the two probe points. The decision to remove candidate regions without searching them depends on the distribution of path points in the $m - v$ plane and the deadline. Thus, the running time of our algorithm can be described probabilistically.

Let the probability that both candidate regions are searched be p_2 and the probability that neither candidate region is searched be p_0 . We give the running time bound of our algorithm taking the varying p_2 and p_0 into account in Theorem 4. We use two lemmas: the running time of our algorithm when p_2 and p_0 are given (Lemma 2) and how p_2 and p_0 vary as algorithm iterations proceed (Lemma 3).

Lemma 2. *Let the probability that both candidate regions are searched be p_2 and the probability that neither candidate region is searched be p_0 . If $p_2 \leq p_0$, the average running time of our algorithm is $O(N^2 \log^2 N)$, where N is the number of nodes in the graph.*

Proof. We use a binary tree to represent a hierarchy of candidate regions created by the λ -search sequence. Each node represents a candidate region and its children nodes represent the two candidate regions created by searching the current region. Let h be the height of the tree, and let the function $num(\mathcal{N})$ be defined as the total number of candidate regions to search for a tree with a root node \mathcal{N} and left and right nodes \mathcal{L} and \mathcal{R} . Then,

$$h = O(\log(N^{\log N})) = O(\log^2 N)$$

$$\text{num}(\mathcal{N}) = \begin{cases} 1 + \text{num}(\mathcal{L}) + \text{num}(\mathcal{R}) & \text{with probability } p_2, \\ 1 + \text{num}(\mathcal{L}) & \text{with probability } \frac{1-p_2-p_0}{2}, \\ 1 + \text{num}(\mathcal{R}) & \text{with probability } \frac{1-p_2-p_0}{2}, \\ 1 & \text{with probability } p_0. \end{cases}$$

The conditional expectation of $\text{num}(\mathcal{N})$ given $\text{num}(\mathcal{L})$ and $\text{num}(\mathcal{R})$ is described as follows:

$$E[\text{num}(\mathcal{N}) | \text{num}(\mathcal{L}), \text{num}(\mathcal{R})] = 1 + \frac{1 + p_2 - p_0}{2} (\text{num}(\mathcal{L}) + \text{num}(\mathcal{R}))$$

Then, the expectation of $\text{num}(\mathcal{N})$ can be expressed as follows, taking expectation over $\text{num}(\mathcal{L})$ and $\text{num}(\mathcal{R})$ on the above conditional expectation.

$$\begin{aligned} E[\text{num}(\mathcal{N})] &= E[E[\text{num}(\mathcal{N}) | \text{num}(\mathcal{L}), \text{num}(\mathcal{R})]] \\ &= 1 + \frac{1 + p_2 - p_0}{2} (E[\text{num}(\mathcal{L})] + E[\text{num}(\mathcal{R})]) \end{aligned} \quad (5)$$

Since the expectation of $\text{num}(\mathcal{N})$ depends only on the height of the tree, we can define a function $\text{num}_e(h)$ as the expected number of nodes given the height h . Then, for a node \mathcal{N} whose height is h and whose two children are \mathcal{L} and \mathcal{R} , we have $E[\text{num}(\mathcal{N})] = \text{num}_e(h)$ and $E[\text{num}(\mathcal{L})] = E[\text{num}(\mathcal{R})] = \text{num}_e(h-1)$. From (5), $\text{num}_e(h)$ can be written as:

$$\text{num}_e(h) = \begin{cases} 1 + (1 + p_2 - p_0) \text{num}_e(h-1) & \text{if } h \geq 1, \\ 0 & \text{if } h = 0. \end{cases}$$

Solving the above recursive equations, noting that $(1 + p_2 - p_0) \leq 1$, we get:

$$\text{num}_e(h) = 1 + (1 + p_2 - p_0) + \dots + (1 + p_2 - p_0)^{h-1} \leq h = O(\log^2 N).$$

Thus, the running time is $O(N^2 \log^2 N)$ when we use Dijkstra's algorithm for each λ -optimal search.

Lemma 3. $p_2(i) \leq p_0(i), \forall i > n, n = O(\log^2 N)$, where i is the index of iterations. In other words, p_0 becomes larger than p_2 around $\log^2 N$ iterations.

Proof. First, consider how $p_2(i)$ and $p_0(i)$ vary as iterations proceed. We have three cases: (a) two regions are left in the queue, (b) one region is left in the queue, or (c) no region is left in the queue. Case (a) occurs only when the λ -optimal search line is in the neighborhood of the tangent to the parabola at the point being considered. The size of the neighborhood depends on the distance between the two probe points (e.g., if the distance is large, even a large difference in the slopes can result in the retention of both points.) If the distance is small, only a small difference in the slopes will lead to the retention of both points. Since the distance gets smaller as iterations proceed, $p_2(i)$ decreases monotonically. On the other hand, the probability of case (c) (e.g., both candidate regions are removed) monotonically increases as the distance between the probe points gets smaller. The optimal parabola gets flatter and the lower bound




Fig. 3. p_2 and p_0 was calculated from 5000 random convex hulls with N_e extreme points.

of λ (λ_l) increases as a better λ -optimal path is found, which also causes $p_2(i)$ to decrease and $p_0(i)$ to increase.

Consider the point when $p_2(i)$ becomes smaller than $p_0(i)$. When a single child is retained at each step, it takes $\log^2 N$ iterations until the optimal path is found because we have a binary tree. Thus the probability that the search algorithm retains one child decreases fast beyond $\log^2 N$ iterations. This results in $p_0(i)$'s fast increase since its increasing rate is related to the decreasing rate of $1 - p_2(i) - p_0(i)$ by $\frac{\Delta p_0(i)}{\Delta i} = -\frac{\Delta(1-p_2(i)-p_0(i))}{\Delta i} + (-\frac{\Delta p_2(i)}{\Delta i})$. Thus, $p_0(i)$ becomes larger than $p_2(i)$ around $i = \log^2 N$. Fig. 3 illustrates these behaviors of $p_2(i)$ and $p_0(i)$.

Theorem 4. *The average running time of Algorithm 1 is $O(N^2 \log^4 N)$.*

Proof. The total running time is determined by considering the computation before and after $p_2 \leq p_0$ is satisfied separately. By Lemma 3 the computational cost before $p_2 \leq p_0$ is satisfied is $O(N^2 \log^2 N)$. Then, we have $O(\log^2 N)$ candidate regions. In the worst case, all the remaining candidate regions should be searched. Since the height of the sub-trees spanning from those candidate regions is bounded by h and they all satisfy $p_2 \leq p_0$, the computational cost after $p_2 \leq p_0$ is satisfied is $O(N^2 \log^4 N)$ by Lemma 2. Thus, the overall running time is $O(N^2 \log^4 N)$.

5 Algorithm Evaluation

We have evaluated empirically the performance of Algorithm 1 against the exhaustive λ -optimal search proposed in [9] using simulation data and real data from the taxi database.

5.1 Experimental Data

Grid Structure: The simulation data set is a square bidirectional grid structure where each edge has a random mean and variance uniformly distributed between 0 and 1. Grid structures of size from 10×10 to 100×100 are used. The origin and destination are two diagonally opposite points.

Physical Road Network: The Delay Statistics Map built using the taxi database is used as a physical test bed. The map has about 29,000 nodes and 39,000 edges. It is dense around the city area and more sparse in rural areas.




Fig. 4. Running time measured at the square bidirectional grid structure where each edge has a random mean and variance between 0 and 1 with a deadline of half grid size. “exhaustive” is the exhaustive λ -optimal search, “probe” is just applying the candidate-region probing method, “bound” is just applying the bounds of λ and “Alg. 1” is Algorithm 1.

5.2 Running Time

The running time of Algorithm 1 is compared with the exhaustive λ -optimal searches using the two data sets above. Fig. 4 shows the results on the simulation data. The individual effect of each method in Sections 4.3 and 4.4 on the running time is also shown. Algorithm 1 with both methods runs fastest. The speedup is by at least a factor of 10 over the algorithm in [9]. The speedup is due to the reduced number of λ -optimal searches. For the algorithm in [9], the number of λ -optimal searches gradually increases from 17 to 119 as the number of nodes increases from 100 to 10000, whereas it increases from 5 to 7 for Algorithm 1.

The running time of Algorithm 1 on the Delay Statistics Map for a route 144km long is 14 seconds with 5 λ -optimal searches when the deadline is 3 hours. The same query took 178 seconds when we used the exhaustive λ -optimal search yielding 75 λ -optimal searches.

5.3 Path Example

Fig. 1 shows different optimal paths from an origin (the arrow) to a destination (the “D”), according to three different criteria: the minimum distance route (this is the same route recommended by Google Maps and is indicated as the topmost route), the minimum expected time route (the middle route), and the maximum probability route with a deadline of 14 minutes (the bottom route). Our system estimates that minimum distance route (which is 3.1 miles) will take 18 minutes on a Tuesday afternoon. Our system’s minimum expected time route (which is 3.5 miles) takes only 11 minutes and 45 seconds. The maximum probability route (which is 4.1 miles) takes 11 minutes and 51 seconds with 90.3% guarantee of arriving on time. The minimum distance route and the minimum expected time route have lower probabilities at 1% and 88.5%, respectively.

5.4 Overall Path Goodness

Four different routes from an origin to a destination were examined using taxi paths and human test driving. The estimated mean of each path from 7 am to 9 pm was




Fig. 5. The probability of arriving within a given deadline from an origin to a destination. (Left) Estimation. (Right) Measurement.

872 seconds, 899 seconds, 816 seconds, and 795 seconds for route 1 (6.9 km), route 2 (7.2 km), route 3 (6.7 km), and route 4 (6.2 km) respectively. The measured mean was 869 seconds, 895 seconds, 811 seconds, and 799 seconds. Thus, the estimated minimum expected time path, route 4, agrees with the measurement. Fig. 5 gives the maximum probability path. The estimated probability is similar to the measured probability. From both estimation and measurement, for a deadline less than or equal to 12 minutes, route 3 is the best, but for a deadline larger than 12 minutes route 4 is the best. We can observe that route 4, which is the minimum expected time path is the worst path for a deadline less than 12 minutes.

5.5 Independent Gaussian Assumption

To test the independent Gaussian assumption, we identified a route with a large number of travel samples in the taxi trajectories database. We used a path from an origin to a destination, which is 1.12 km long and has 5 intersections and 6 road segments. From Fig. 6 (Right), we can observe that the empirical data and the estimation by independent Gaussian assumption is very close. More specifically, Fig. 7 shows that the empirical data is very similar to the Gaussian distribution, especially in the probability interval from 0.05 to 0.95. Thus, our assumptions will hold well for the stochastic planning for reaching the destination with the probability in this range. The discrepancies observed over 0.95 and under 0.05 show the limitation of our algorithms due to our assumption. For example, as shown in Fig. 7 (Bottom left), our system will estimate that the users can reach the goal with 99% probability if they leave 230 seconds before the deadline, but the empirical data shows that we will get only a 97% chance, and if users want a 99% guarantee they should leave about 270 seconds before the deadline. The discrepancy over 0.95 is caused by some unusual long delay, which might be due to unexpected events, construction work, or data gathering




Fig. 6. (Left) Delay distribution for one segment. (Right) The comparison of mean and standard deviation between empirical measurement and estimation for various time slots.




Fig. 7. “Empirical” indicates the travel time measurement by driving, “Gaussian” indicates the Gaussian fit for the entire data and “Gaussian(in 270)” indicates the Gaussian fit using only the data in 270 seconds. (Top left) Histogram of empirical travel time data and Gaussian fits for any time for weekdays. (Bottom left) Probability comparison between empirical data and Gaussian fits for any time for weekdays, where the Y axis was scaled to make the Gaussian CDF linear (Right) Corresponding plots to the left plots for 1 ~ 2 pm for weekdays

noise such as taxi drivers’ intentional stops or slow drives. The discrepancy below 0.05 is due to the Gaussian distribution spanning to the negative value whereas the travel time cannot be negative. We observe less discrepancy in case of the right plots, which use only the data from 1 ~ 2 pm whereas the left plots are for entire hours. This result suggests that narrowing the data by conditions that affect the traffic delays such as time of day makes the delay distribution look more like independent Gaussian distribution. Thus, in our ongoing research, we are investigating the proper conditions that constrain the traffic delays.

This observation provides some evidence that the independent Gaussian assumption used for Algorithm 1 holds for the taxi trajectories database, but more testing on road segments with more associated travel data is necessary. As the database grows every day, we plan to continue this validation.

6 Conclusion and Future Work

We developed efficient stochastic shortest-path algorithms that can be used for a practical traffic information system. We evaluated the system with actual measured travel time for selected routes, and observed that our system’s optimal path and travel time estimates are close to reality.

In the future, we are interested in developing path planning algorithms for multiple users. We are also interested in improving our algorithms and system by considering dependencies of each road segment and by using better modeling of delay distributions. We are currently extending the algorithms to integrate current traffic information with historical information to make more accurate estimations and to

predict the future traffic conditions. Considering various conditions affecting traffic like weather, construction work, and events is also part of our current plans. Finally, we plan to integrate this planning system with autonomous vehicles.

7 Acknowledgments

We thank PlanetTran for participating in the taxi data collection. This work has been supported in part by NSF grants EFRI-0710252 and IIS-0426838, by the SWARMS MURI, and by a Samsung Scholarship. We are grateful for this support.

References

1. D. P. Bertsekas and J. N. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595., August 1991.
2. P. Carstensen. The complexity of some problems in parametric linear and combinatorial programming. *Ph.D. Thesis, Mathematics Dept., U. of Michigan, Ann Arbor, Mich.*, January 1983.
3. I. Chabini. Discrete dynamic shortest path problems in transportation applications: Complexity and algorithms with optimal run time. *Transportation Research Record*, 1645:170–175., 1998.
4. R. Chrobok, J. Wahle, and M. Schreckenberg. Traffic forecast using simulations of large scale networks. In *IEEE Conference on Intelligent Transportation Systems*, Oakland, CA, USA, August 2001.
5. B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. K. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In *4th ACM SenSys*, Boulder, CO, November 2006.
6. R. Loui. Optimal paths in graphs with stochastic or multidimensional weights. *Communications of the ACM*, 26(9):670–676, 1983.
7. I. Murthy and S. Sarkar. Exact algorithms for the stochastic shortest path problem with a decreasing deadline utility function. *European Journal of Operational Research*, 103:209–229, 1997.
8. E. Nikolova, M. Brand, and D. Karger. Optimal route planning under uncertainty. In *International Conference on Automated Planning and Scheduling*, 2006.
9. E. Nikolova, J. A. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *ESA*, pages 552–563, 2006.
10. K. K. Sanwal and J. Walrand. Vehicles as probes. *Technical Report UCB-ITS-PWP-95-11, California Partners for Advanced Transit and Highways (PATH)*, January 1995.
11. D. Schrank and T. Lomax. The 2007 urban mobility report. *Annual report, Texas Transportation Institute, The Texas A&M University System*, September 2007.
12. C. E. Sigal, A. A. B. Pritsker, and J. J. Solberg. The stochastic shortest route problem. *Operations Research*, 28(5):1122–1129, 1980.
13. M. P. Wellman, M. Ford, and K. Larson. Path planning under time-dependent uncertainty. In *11th Conference on Uncertainty in Artificial Intelligence*, pages 532–539, August 1995.
14. H. Xu and D. J. Dailey. Real time highway traffic simulation and prediction using inductance loop data. In *Vehicle Navigation and Information Systems Conference*, Seattle, WA, USA, July 1995.
15. J. Yoon, B. Noble, and M. Liu. Surface street traffic estimation. In *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*, pages 220–232, New York, NY, USA, 2007. ACM.