

Stochastic Shortest Paths Via Quasi-convex Maximization

Nikolova, E.; Kelner, J.; Brand, M.; Mitzenmacher, M.

TR2006-128 September 2006

Abstract

We consider the problem of finding shortest paths in a graph with independent randomly distributed edge lengths. Our goal is to maximize the probability that the path length does not exceed a given threshold value (deadline). We give a surprising exact $n^{\theta(\log n)}$ algorithm for the case of normally distributed edge lengths, which is based on quasi-convex maximization. We then prove average and smoothed polynomial bounds for this algorithm, which also translate to average and smoothed bounds for the parametric shortest path problem, and extend to a more general non-convex optimization setting. We also consider a number other edge length distributions, giving a range of exact and approximation schemes.

ESA 2006

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Stochastic Shortest Paths via Quasi-Convex Maximization

Evdokia Nikolova*, Michael Mitzenmacher, John A. Kelner, Matthew Brand

MIT, Harvard, MIT, MERL
Cambridge MA, 02139

Abstract. We consider the problem of finding shortest paths in a graph with independent randomly distributed edge lengths. Our goal is to maximize the probability that the path length does not exceed a given threshold value (deadline). We give a surprising exact $n^{\Theta(\log n)}$ algorithm for the case of normally distributed edge lengths, which is based on quasi-convex maximization. We then prove average and smoothed polynomial bounds for this algorithm, which also translate to average and smoothed bounds for the parametric shortest path problem, and extend to a more general non-convex optimization setting. We also consider a number other edge length distributions, giving a range of exact and approximation schemes.

1 Introduction

Finding shortest paths between a given source and destination is a classic and fundamental problem in theoretical computer science which has influenced a wide array of other fields. It is less clear what a *stochastic* shortest path would mean, when the edge lengths are random with given distributions. Is it the shortest path on average, or the path minimizing a combination of mean and variance, or minimizing some other criterion? Is it found adaptively or non-adaptively? A variety of problem variants have appeared in the literature, most minimizing the expected length of a path, or a combination of expected length and expected cost such as bicriterion problems [15], [19]. Adaptive formulations have prevailed, perhaps because a non-adaptive minimization of the expected path length trivially reduces to the deterministic shortest path problem.

Few researchers have considered optimizing a non-linear function of the (random) path length. Some notable work includes that of Loui [12] who seeks the path maximizing an expected utility of the path length for a class of monotone decreasing utility functions. Fan *et al.* [6] present an adaptive heuristic for paths that maximize the probability of arriving on time. Formulations of this type with nonlinear objective, though perhaps most useful in practice, have been sparse, because different sources of hardness arise from many levels: combinatorial, distributional, analytic, functional, to list a few. For example, in the absence of randomness, the combinatorial nature of the problem may be hard to approximate for certain objective functions (*e.g.*, longest path [9]). In the absence of graph structure, the objective function in itself may be difficult to

* enikolova@csail.mit.edu. Part of this work was done while the author was at Mitsubishi Electric Research Labs. Supported in part by NSF grants ANI-0225660 and ITR-0219018.

optimize: we can solve efficiently linear programming but not even quadratic or more generally non-convex programming. The distributions may be hard to work with: calculating values of the cumulative distribution function of the sum of n Bernoulli random variables is #P-hard as it corresponds to counting knapsack solutions [11]. Computing the expectation $\mathbf{E}[u(X)] = \int u(x)f(x)dx$ of the non-linear utility function $u(\cdot)$ of the random path length X with probability density function $f(\cdot)$ may not even have a closed form, thus making the standard notion of computational hardness inapplicable. Superimposing these sources of difficulty may ultimately lead to a problem that has no hope of even being categorized as to what level of hardness it has—partly because we do not understand to what extent each source contributes to the overall complexity.

We thus focus on a stochastic shortest paths model which can effectively factor the sources of difficulty above and at the same time has an innovative solution drawing from a variety of areas. Inspired by recent formulations of the stochastic knapsack and other classic problems turned stochastic [5], [7] our goal is to maximize the probability that the path length would not exceed some threshold value. This is a natural formulation which is also very practical: For example, this is our objective when we are going to the airport and want to pick the path that would maximize our probability of arriving on time for our flight. We consider a pre-planning (nonadaptive) scenario and note that it can easily be converted to an adaptive one by rerunning our algorithms on the fly with updated information.

Apart from the inherent practicality of the problem, it reveals a deeper theoretic structure intertwining areas such as nonconvex programming, the geometry of path polytopes and combinatorial optimization. As a preview to some of the open questions, we give an exact $n^{\Theta(\log n)}$ algorithm for our main model with edge lengths drawn from normal distributions. It is unknown whether a polynomial exact algorithm exists or whether this problem is complete for the corresponding complexity class LogNP [21]. Our algorithm also reveals a somewhat unexpected connection between Kelner & Spielman’s recent techniques for linear programming [10] and the much more general field of nonconvex optimization. We extend their techniques to get polynomial-time average and smoothed complexity for our superpolynomial algorithms. We stress that these smoothed results are stronger than previous smoothed results in that they do not perturb the feasible set (the path polytope), but just the objective function (the plane on which the polytope is projected). Or, in the terms of the stochastic shortest paths terminology, only the edge means and variances and *not* the solution paths themselves, are slightly perturbed. As an added benefit, we reveal a connection between the stochastic and parametric shortest path problems, which implies new average and smoothed results for the parametric shortest path problem as well. Our results can also generalize to a wide class of non-convex optimization problems, known as low-rank quasiconcave minimization [18].

1.1 Our Results

We define a model for the stochastic shortest path problem in which the edge lengths are independent random variables drawn from known distributions. The optimal path maximizes the probability that the path length does not exceed a given threshold (deadline).

This objective arises naturally in practice where a user wants to maximize the probability of arriving on time to a destination. In an effort to decouple the complexity inherent in this objective from the distributional and analytic complexity of the problem, our first model draws the edges from normal distributions. We show that for a large range of deadlines our problem entails the maximization of a quasi-convex function over the path polytope. Due to the particular form of our quasi-convex objective, the optimal path is attained at an extreme point of the dominant of the projection (shadow) of the path polytope onto a two-dimensional plane. We thus give an exact algorithm for finding the optimal path by walking along extreme points of the shadow dominant. We then establish an equivalence between the shadow dominant and the optimal cost envelope of the parametric shortest path problem. Consequently, this proves that our algorithm has a worst case running time $n^{\Theta(\log n)}$. We give a pseudopolynomial algorithm for the remaining range of deadlines.

In the following section we extend the techniques from Kelner & Spielman [10] to prove linear average and smoothed complexity of the shadow of the path polytope and consequently polynomial running time of our algorithm. These results also imply new polynomial average and smoothed bounds on the complexity of the parametric shortest path problem and hold for a wider class of non-convex optimization problems than the specific stochastic shortest path objective.

Finally we extend our model to distributions other than the normal. For edge lengths coming from a Poisson or a gamma distribution with a fixed second parameter, or more generally distributions which are additive and satisfy stochastic dominance, we show that the problem easily reduces to the deterministic shortest path problem. For the case of exponential and Bernoulli random variables, we give polynomial (PTAS) and quasi-polynomial (QPTAS) approximation schemes respectively based on a discretization of the state space of the random edge lengths.

1.2 Related Work

The majority of the related literature on stochastic shortest paths focuses on adaptive algorithms, which compute the next best hop based on information about realized edge lengths so far [2], [22], [3], [20], [6], [13]. Most of the adaptive formulations focus on minimizing expected path length; few consider minimizing a non-linear function of the length and settle for heuristic algorithms [6].

The most closely related nonadaptive formulation to our model is that of Loui [12]. Loui considers a general utility function of path length which is monotone and nondecreasing, and proves that the expected utility becomes separable into the edge lengths only when the utility function is linear or exponential. In that case the path that maximizes expected utility can be found via traditional shortest path algorithms. For general utility functions he gives an algorithm based on an enumeration of paths, with a very large running time $O(n^n)$. In a consequent paper, Mirchandani and Soroush give exponential algorithms and heuristics for quadratic utility functions [14]. For non-monotone utility functions Nikolova, Brand and Karger [17] give hardness results and pseudopolynomial algorithms. For a separate model on bicriteria shortest paths with monotone objective, Ackerman *et al.* [1] give different average and smoothed analyses.

2 Problem definitions and quasi-convex maximization

2.1 Stochastic Shortest Path definition

Consider a graph $G = (V, E)$, with $|V| = n$ nodes and $|E| = m$ edges. We are given a source S and destination T . Each edge i has an independent random variable length (travel time) X_i . We have a deadline in time t , and we would like to find an ST -path which maximizes the probability that we reach the destination within time t . Thus, we would like to solve

$$\max_{\pi} \Pr \left(\sum_{i \in \pi} X_i \leq t \right) \quad \text{for paths } \pi \text{ between the source and destination.} \quad (1)$$

In the following sections, we see that different distributional assumptions for the edge lengths lead to problem complexity and algorithms of very different nature.

2.2 Parametric Shortest Path definition

Consider a graph G with distinguished source S and destination T . Each edge i has a parameter dependent length $u_i + \lambda w_i$, where u_i, w_i are nonnegative constants, and $\lambda \in [0, \infty)$. The parametric shortest paths problem looks for the parameter values (breakpoints) $\lambda \in (0, \infty)$ at which the shortest path changes. Carstensen [4] proved that the number of breakpoints is at least $n^{\Omega(\log n)}$ in the worst case, and one can easily show a matching upper bound for general graphs (A more involved proof on the upper bound is also available in Carstensen [4]).

In the next section we will establish a connection between the stochastic shortest paths with normal distributions and the parametric shortest paths problem, which will enable us to apply our average and smoothed results for the former to the parametric shortest path setting as well.

2.3 Quasi-convex maximization

In this section we briefly define convex functions and their generalization to quasi-convex functions and state the main property of their global maxima.

Let C be a convex set.

Definition 1. A function $f : C \rightarrow (-\infty, \infty]$ is convex if for all $x, y \in C$ and $\alpha \in [0, 1]$,

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y).$$

A function $f : C \rightarrow (-\infty, \infty]$ is quasi-convex if all its lower level sets $L_{\gamma} = \{x \mid x \in C, f(x) \leq \gamma\}$ are convex.

Informally, quasi-convex functions have a convex cross-section at any height (level).

Definition 2. We say that x is an extreme point of the set C if it cannot be represented as a convex combination of two other points in the set C ,

$$x = \alpha y + (1 - \alpha)z \text{ for } y, z \in C, \alpha \in (0, 1) \quad \Rightarrow \quad y = z = x.$$

The following important property of quasi-convex maximization seems to be attributed to folklore. A statement of the theorem without proof appears in the *Introduction to Global Optimization* [8]; our proof is deferred to the full version of this paper.

Theorem 1. *Let $C \subset \mathcal{R}^m$ be a compact convex set. A quasi-convex function $f : C \rightarrow \mathcal{R}$ that attains a maximum over C , attains the maximum at some extreme point of C .*

We will need a few more definitions. The *shadow* of a convex set in \mathcal{R}^m onto a two-dimensional subspace is the orthogonal projection of the set onto the subspace. The *dominant* of a set C in \mathcal{R}^m is defined as the set of all points that are greater than a point in C , $\{x \in \mathcal{R}^m \mid x \geq y \text{ for some } y \in C\}$.

3 Stochastic shortest paths with normal distributions

In this section we apply quasi-convex maximization to a graph with normally distributed edge lengths, in which we have to select the most certain route to reach a destination by a given time.

Assume each edge i has independent normally distributed length $X_i \sim N(\mu_i, \sigma_i^2)$. Our problem is to

$$\max_{\pi} \Pr\left(\sum_{i \in \pi} X_i \leq t\right) \quad \text{for paths } \pi \text{ between the source and destination.} \quad (2)$$

For any path π , this probability can be computed by

$$\Pr\left(\sum_{i \in \pi} X_i \leq t\right) = \Pr\left(\frac{\sum X_i - \sum \mu_i}{\sqrt{\sum \sigma_i^2}} \leq \frac{t - \sum \mu_i}{\sqrt{\sum \sigma_i^2}}\right) = \Phi\left(\frac{t - \sum \mu_i}{\sqrt{\sum \sigma_i^2}}\right),$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal random variable $N(0, 1)$. Since Φ is monotone increasing, the problem is equivalent to finding the *ST*-path which maximizes its argument,

$$\max_{\pi} \frac{t - \sum_{i \in \pi} \mu_i}{\sqrt{\sum_{i \in \pi} \sigma_i^2}}. \quad (3)$$

The objective in Eq. (3) cannot be separated into edge costs and does not satisfy sub-optimality so a dynamic programming approach based on substructure would fail. To better understand the properties of the objective function, we formulate it as a continuous optimization problem over the path polytope in \mathcal{R}^m , where m is the number of edges.

Index all edges by $1, 2, \dots, m$. Represent each edge subset by its incidence vector $x \in \mathcal{R}^m$, with $x_i = 1$ if edge i is in the subset and $x_i = 0$ otherwise. All 2^m subsets of edges correspond to the vertices of the unit hypercube in \mathcal{R}^m . The *ST-path polytope* (or, the *path polytope* for short) is the convex hull of incidence vectors of (simple) *ST*-paths. It is a subset of the unit hypercube in \mathcal{R}^m , and its vertices are a subset of the




Fig. 1. Projection of the unit hypercube (representing all edge subsets) and the path polytope onto the (μ, σ^2) -plane.

vertices of the hypercube. Thus, the optimal ST -path is a solution to

$$\begin{aligned}
 & \text{maximize} && \frac{t - \mu \cdot x}{\sqrt{\sigma^2 \cdot x}} && (4) \\
 & \text{subject to} && x \in \text{path polytope} \\
 & && x \in \{0, 1\}^m,
 \end{aligned}$$

where by $\{0, 1\}^m$ we denote the set of 0–1 vectors of length m . Projecting the path polytope onto the span of vectors $\mu = (\mu_1, \dots, \mu_m)$ and $\sigma^2 = (\sigma_1^2, \dots, \sigma_m^2)$ defines a convex polygon, which we call the path polytope shadow. The objective in Eq. (4) is not separable, far from linear or quadratic and not even convex. This places it in a category of mathematical programming and combinatorial optimization problems, for which there are no general efficient algorithms. Although the integer constraints are what usually causes the main difficulty, in this case it is not clear how to solve even the fractional version.

It turns out our objective has special structure which forces its maximum to lie on the boundary of the feasible set. In particular, it is quasi-convex on a subset of the path polytope and monotone in $\mu \cdot x$ and $\sigma^2 \cdot x$ on the remaining subset of the polytope. This is not automatically good news since we do not have a polynomial description of the boundary of the path polytope or even its shadow. For example computing the rightmost and uppermost vertices of the path polytope shadow corresponds to finding the longest path, in terms of the edge means and edge variances respectively. Thus the




Fig. 2. Correspondence of extreme points P_1, P_2, \dots of the dominant of the shadow path polytope (left) and linear segments from the parametric path cost-function $g(\lambda) = \inf_x \{\mu \cdot x + \lambda(\sigma^2 \cdot x)\}$ (right).

computation of the path convex hull is in general strongly NP-hard [9]. On the other hand, we can efficiently find the extreme points on the dominant of the shadow hull, since they optimize the linear objective $\gamma\mu + (1 - \gamma)\sigma^2$ for $\gamma \in [0, 1]$.

Our main theorem 2 shows that for sufficiently early departure time (eliminating the incidence of a longest path problem), our objective is quasi-convex and we can solve the stochastic shortest paths problem exactly in time $n^{\Theta(\log n)}$. We first state a lemma about the correspondence of the stochastic and parametric shortest paths problems. Its proof is deferred to the full paper version.

Lemma 1. *There is a one-to-one correspondence between the extreme points on the shadow of the path polytope dominant on the plane spanned by vectors $u = (u_1, \dots, u_m)$, $w = (w_1, \dots, w_m)$ and the breakpoints of the parametric shortest path problem with edge weights $u_i + \lambda w_i$.*

By Lemma 1, the results for the complexity of the parametric shortest paths problem [4], [16] imply equivalent bounds for the number of extreme points on the shadow dominant.

Corollary 1. *The dominant of the path polytope shadow has $n^{\Theta(\log n)}$ extreme points in the worst case.*

We now turn to the main result in this section.

Theorem 2. *When the deadline t is no less than the mean of the smallest-mean path, the solution to Eq. (4) is an extreme point of the dominant of the path polytope shadow and can be found in time $n^{\Theta(\log n)}$.*

Proof. We first consider the relaxed version of Eq. (4). Denoting $z_1 = \mu \cdot x$ and $z_2 = \sigma^2 \cdot x$, the system becomes equivalent to

$$\begin{aligned} & \text{maximize} && \frac{t - z_1}{\sqrt{z_2}} && (5) \\ & \text{subject to} && (z_1, z_2) \in \text{path polytope shadow } S \end{aligned}$$

We first show that the induced objective $f(z_1, z_2) = \frac{t-z_1}{\sqrt{z_2}}$ is quasi-convex on a subset of the feasible set $\bar{S} = S \cap \{z_1 \mid z_1 < t\}$ (which is non-empty assuming there is a path with mean less than t). Since $z_1 = \mu \cdot x < t$, the value of $f(z_1, z_2)$ on this feasible subset is positive, and must contain the maximum. Consider the level set $L_\gamma = \{z \in \mathcal{R}^2 \mid f(z) \leq \gamma\}$. This set consists of points (z_1, z_2) such that

$$\frac{t-z_1}{\sqrt{z_2}} \leq \gamma \iff z_2 \geq \left(\frac{t-z_1}{\gamma}\right)^2,$$

hence for positive γ and $z_1 < t$, the level set L_γ is convex. Therefore $f(z_1, z_2)$ is quasi-convex on \bar{S} , which is the part of path polytope shadow to the left of $z_1 = t$. By Theorem 1, the maximum is attained at an extreme point of \bar{S} . Further, since $f(z_1, z_2)$ is monotone decreasing in both z_1 and z_2 , the solution must be an extreme point of the dominant of the shadow, to the left of $z_1 = t$.

Now, any extreme point of the shadow is the projection of an extreme point of the original path polytope (which has integer coordinates). Hence the optimal solution of the relaxed program (5) is also a solution to the integer program (4).

Next, the extreme points of the dominant of the shadow can be found in time linear in their number, for example with a binary search type enumeration as follows. Each extreme point on the shadow dominant is the solution to a linear program

$$\begin{aligned} \min \quad & c \cdot z \\ \text{subject to} \quad & z \in \text{shadow path polytope} \end{aligned} \tag{6}$$

for some $c = (c_1, c_2) \geq 0$. Equivalently, each extreme point corresponds to a path minimizing $c_1 z_1 + c_2 z_2$ where $z_1 = \mu \cdot x$ is the total mean of the path and z_2 is the total variance so for $c_1, c_2 \geq 0$ it can be found via any shortest path algorithm. To find all extreme points on the shadow dominant, we start with its two endpoints: the leftmost point, which corresponds to the path with smallest mean, and the bottom-most point, which is the path of smallest variance. Denote these $\pi_1 = (m_1, s_1), \pi_2 = (m_2, s_2) \in \mathcal{R}^2$, where m_i is the mean and s_i the variance of path π_i , then solve Eq. (6) with $(c_1, c_2) = \left(-\frac{s_2-s_1}{m_2-m_1}, 1\right)$ if $m_2 - m_1 \neq 0$, otherwise $(c_1, c_2) = (1, 0)$. The new solution is $\pi_3 = (m_3, s_3)$, a vertex between π_1 and π_2 on the shadow boundary. If different from both π_1 and π_2 , we repeat the procedure for finding a vertex between π_1, π_3 and between π_3, π_2 , etc. Clearly in this way we find all vertices on the shadow dominant in time linear in their number, multiplied by the time to solve the auxiliary program (6). Similar enumeration methods for extreme points are discussed in Carstensen [4].

Finally, since there are $N = n^{\Theta(\log n)}$ extreme points of the shadow dominant in the worst case by Corollary 1 and we can find each in polynomial time, the running time of the algorithm is $n^{\Theta(\log n)}$.

When the departure time is closer to the deadline, so that any shortest path has mean greater than t , our objective is no longer quasi-convex, in fact it is increasing in the variance. Since finding the simple path with highest variance is strongly NP-hard [9], we might not expect to find a good polynomial-time approximation. Settling for potentially non-simple paths, we can give a pseudopolynomial dynamic programming solution, which finds the path of smallest mean for every possible value of its variance and then selects the ST-path with optimal objective value.

Theorem 3. For general deadline t , the solution to Eq. (4) can be found in time $O(\sigma^2 nm)$ where σ^2 is the maximum variance of an edge.

4 Average and Smoothed Complexity

In this section we show that if the edge weight vectors $u, w \in \mathcal{R}^m$ are uniformly random unit vectors or fixed vectors which are slightly perturbed, then the expected number of extreme points on the path polytope shadow is linear and consequently our $n^{\Theta(\log n)}$ algorithm from the previous section will have a low expected polynomial running time. The techniques in this section are motivated by the recent techniques of Kelner and Spielman [10] for the polynomial simplex algorithm for linear programming.

Note that the vertices of the path polytope P are a subset of the vertices of the unit hypercube, in particular:

Fact 1 Each edge of the polytope P has length at least 1.

Fact 2 The polytope P is contained in the unit hypercube, which in turn is contained in a ball with radius $\sqrt{m}/2$.

4.1 Average bounds

Theorem 4. Let $u, w \in \mathcal{R}^m$ be uniformly random unit vectors and let V be their span. Then the expectation of the number of edges on the projection of P onto V is at most $2\sqrt{2}\pi m$.

Proof. By Fact 2, the perimeter of the shadow of P onto V is bounded above by $\pi\sqrt{m}$. Next, for each edge I of the polytope P , denote by $S_I(V)$ the event that edge I appears in the shadow, and let $l(I)$ be the length of the edge in the shadow. The sum of expected edge lengths in the shadow is at most equal to the biggest possible perimeter:

$$\sum_I \mathbf{E}[l(I)] = \sum_I \mathbf{E}[l(I)|S_I(V)] \Pr[S_I(V)] \leq \pi\sqrt{m}.$$

By Lemma 2 below, $\mathbf{E}[l(I)|S_I(V)] \geq \frac{1}{2\sqrt{2m}}$. Therefore,

$$\mathbf{E}[\text{number of shadow edges}] = \sum_I \Pr[S_I(V)] \leq 2\sqrt{2}\pi m,$$

where m is the dimension of the polytope P , in our case it is the number of edges of the original graph.

Lemma 2. For all edges I of the polytope P , $\mathbf{E}[l(I)|S_I(V)] \geq \frac{1}{2\sqrt{2m}}$.

Proof. We first note a direct corollary from Lemma 3.8 in Kelner & Spielman [10], namely that if an edge I of the polytope appears in the shadow, it must make a small angle $\theta_I(V)$ with the projection plane V , $\Pr_V [\cos(\theta_I(V)) \geq \frac{1}{\sqrt{2m}} | S_I(V)] \geq \frac{1}{2}$.

Now, since any edge in the polytope P has length at least 1 (by Fact 1 above), the length of the edge in the shadow would be at least $\cos(\theta_I(V))$ and its expectation provided it appears in the shadow is

$$\mathbf{E}[l(I)|S_I(V)] \geq \frac{1}{\sqrt{2m}} \frac{1}{2}.$$

4.2 Smoothed bounds

We now provide smoothed results for the maximization of our quasi-convex objective. In particular, we show that the expected number of extreme points (equivalently edges) on the projection of a general 0–1 vertex polytope onto a perturbed plane is polynomial in m and $1/\rho$, the inverse of our perturbation.

We first define a ρ -perturbation of the vector u , for $\rho > 0$. Choose an angle $\theta \in [0, \pi]$ at random from an exponential distribution with mean ρ , restricted to the range $[0, \pi]$. Set the ρ -perturbation of u to be a unit vector chosen uniformly at random at an angle θ to u . The following theorem states that the expected number of edges on the polytope shadow is polynomial.

Theorem 5. *Let $u_1, u_2 \in \mathcal{R}^m$ be given vectors and let v_1 and v_2 be their respective ρ -perturbations. Denote $V = \text{span}(v_1, v_2)$. The expected number of edges of the projection of P onto V is at most $4\pi\sqrt{2m}/\rho$, for $\rho < 1/\sqrt{m}$.*

The theorem follows similarly to the argument in Section 4.1 from the next lemma.

Lemma 3. *With the variables above, $\Pr_{v_1, v_2}[\cos(\theta_I(V)) \leq \epsilon \mid S_I(V)] \leq 4(\epsilon/\rho)^2$.*

This lemma generalizes the lemma of Kelner and Spielman [10] by allowing both v_1 and v_2 to be drawn from ρ -perturbed distributions, as opposed to requiring one of them to be uniformly random. Its proof is deferred to the full version of this paper.

Naturally, the smaller the perturbation, the weaker the bound in the theorem. However setting $\rho = \frac{1}{\sqrt{2m}}$ for example, gives the linear bound $8\pi m$ which is just a little larger than the bound on the number of shadow edges for the average case. Finally note that by Lemma 1, these bounds imply linear (in the number of graph edges) average and smoothed bounds for the number of optimal paths in the parametric shortest paths problem as well.

5 Extensions to other distributions

5.1 Poisson and additive stochastic dominant distributions—exact solution

The probability distribution $\mathcal{D}(\lambda)$ is called *additive* if the sum of two independent random variables with distributions $\mathcal{D}(\lambda_1)$ and $\mathcal{D}(\lambda_2)$ is another random variable with the same distribution and parameter equal to the sum, $\mathcal{D}(\lambda_1 + \lambda_2)$. With a slight abuse of notation we use $\mathcal{D}(\lambda)$ to also denote a random variable with this distribution. Assume in addition that the distribution \mathcal{D} satisfies *stochastic dominance*, that is $\Pr(\mathcal{D}(\lambda_1) \leq t) \geq \Pr(\mathcal{D}(\lambda_2) \leq t)$ whenever $\lambda_1 \leq \lambda_2$. Examples of such distributions are Poisson and *gamma*(a, b) with constant parameter b .

Suppose the random length of edge i is $X_i \sim \mathcal{D}(\lambda_i)$. Now, despite the non-separable objective function, the form of distribution makes the problem separable:

$$\Pr\left(\sum_{i \in \pi} X_i \leq t\right) = \Pr\left(\sum_{i \in \pi} \mathcal{D}(\lambda_i) \leq t\right) = \Pr\left(\mathcal{D}\left(\sum_{i \in \pi} \lambda_i\right) \leq t\right) \geq \Pr\left(\mathcal{D}(\lambda') \leq t\right),$$

where the last inequality follows from the stochastic dominance property of the distribution for all $\lambda' \geq \sum_{i \in \pi} \lambda_i$. With this, the optimal path is the one that has the smallest sum of distribution parameters along its links and can be found exactly with a deterministic shortest path algorithm.

5.2 Exponential PTAS and Bernoulli QPTAS

Unlike the Poisson, the exponential distribution is not additive and we cannot write a simple closed form expression for the objective function. We propose a polynomial-time approximation scheme, based on dynamic programming over a discretization of the distribution parameter space. More precisely, we give a bicriterion approximation, which is given a lateness tolerance p and for $\epsilon > 0$ it finds an ST -path π satisfying

$$\Pr\left(\sum_{i \in \pi} X_i < t(1 + \epsilon)\right) > 1 - \epsilon p.$$

Due to space constraints, we defer the algorithm description to the full paper version and only state its running time.

Theorem 6. *An approximately optimal path π with $\Pr[\sum_{i \in \pi} X_i > (1 + \epsilon)] \leq p(1 + \epsilon)$ can be computed in time $O(n^4 \log n) O\left(\frac{\log(1/\epsilon)}{\epsilon^4} \log \frac{1}{\epsilon p}\right) \gamma^{O(\frac{1}{\epsilon} \log 1/\epsilon)}$.*

A similar discretization of the state space yields a quasi-polynomial approximation scheme for the case of Bernoulli distributions; we omit the details from this version.

6 Conclusion

We have considered a novel framework for stochastic shortest paths with independent random edge lengths. When the edges are normally distributed, we give an exact $n^{\Theta(\log n)}$ algorithm. Several points worth noting are that this is an unusual algorithm (not based on dynamic programming) with an unusual running time for the classic shortest path problem in the presence of uncertainty. Although the problem is inherently discrete, in its core are properties from continuous optimization. One possibility to prove a polynomial worst-case bound on our $n^{\Theta(\log n)}$ algorithm is to restrict the class of graphs under consideration. We conjecture that the number of extreme points on the corresponding shadow dominant of planar graphs is polynomial (linear) in the size of the graph.

We present polynomial average and smoothed bounds with respect to the means and variances of the edge length distributions in the case of normal distributions. We note that these bounds hold for the maximization of any quasi-convex function of rank 2 (that is, a function of the form $f(a \cdot x, b \cdot x)$ for vectors $a, b \in \mathcal{R}^m$) over general polytopes with 0–1 vertex coordinates. Our results could be further generalized [18] and apply to diverse other settings as well as serve of independent interest to non-convex optimization.

Other open questions for our stochastic shortest path model include considering correlated as well as dynamically varying edge length distributions.

Acknowledgement. We wish to thank Brian Dean, David Karger, Asu Ozdaglar, and Santosh Vempala for many helpful suggestions.

References

1. H. Ackermann, A. Newman, H. Röglin, and B. Vöcking. Decision making based on approximate and smoothed pareto curves. In *Proc. of 16th ISAAC*, pages 675–684, 2005.
2. D. Bertsekas. *Dynamic Programming and Optimal Control*, volume II, 2nd Edition. Athena Scientific, Belmont, MA, 2001.
3. J. Boyan and M. Mitzenmacher. Improved results for route planning in stochastic transportation networks. In *Proc. of Symposium of Discrete Algorithms*, 2001.
4. P. Carstensen. *The complexity of some problems in parametric linear and combinatorial programming*. Ph.D. Thesis, Mathematics Dept., U. of Michigan, Ann Arbor, Mich., 1983.
5. B. Dean, M. Goemans, and J. Vondrak. Approximating the stochastic knapsack: the benefit of adaptivity. In *Proceedings of FOCS*, pages 208–217, 2004.
6. Y. Fan, R. Kalaba, and I. J. E. Moore. Arriving on time. *Journal of Optimization Theory and Applications*, forthcoming.
7. A. Goel and P. Indyk. Stochastic load balancing and related problems. In *Proceedings of the 40th Symposium on Foundations of Computer Science*, 1999.
8. R. Horst, P. M. Pardalos, and N. V. Thoai. *Introduction to Global Optimization*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2000.
9. D. Karger, R. Motwani, and G. D. S. Ramkumar. On approximating the longest path in a graph. *Algorithmica*, 18:82–98, 1997.
10. J. A. Kelner and D. A. Spielman. A randomized polynomial-time simplex algorithm for linear programming. *Electronic Colloquium on Computational Complexity*, No. 156, 2005.
11. J. Kleinberg, Y. Rabani, and É. Tardos. Allocating bandwidth for bursty connections. *SIAM Journal on Computing*, 30(1):191–217, 2000.
12. R. P. Loui. Optimal paths in graphs with stochastic or multidimensional weights. *Communications of the ACM*, 26:670–676, 1983.
13. E. D. Miller-Hooks and H. S. Mahmassani. Least expected time paths in stochastic, time-varying transportation networks. *Transportation Science*, 34:198–215, 2000.
14. P. Mirchandani and H. Soroush. Optimal paths in probabilistic networks: a case with temporary preferences. *Computers and Operations Research*, 12(4):365–381, 1985.
15. J. Mote, I. Murthy, and D. Olson. A parametric approach to solving bicriterion shortest path problems. *European Journal of Operational Research*, 53:81–92, 1991.
16. K. Mulmuley and P. Shah. A lower bound for the shortest path problem. *Journal of Computer and System Sciences*, 63(2):253–267, 2001.
17. E. Nikolova, M. Brand, and D. R. Karger. Optimal route planning under uncertainty. In *Proceedings of International Conference on Automated Planning and Scheduling*, 2006.
18. E. Nikolova and J. A. Kelner. On the hardness and smoothed complexity of low-rank quasiconcave minimization. *Manuscript*, May, 2006.
19. S. Pallottino and M. G. Scutella. Shortest path algorithms in transportation models: Classical and innovative aspects. Technical Report TR-97-06, Università di Pisa Dipartimento di Informatica, Pisa, Italy, 1997.
20. C. H. Papadimitriou and M. Yannakakis. Shortest paths without a map. *Theoretical Computer Science*, 84:127–150, 1991.
21. C. H. Papadimitriou and M. Yannakakis. On limited nondeterminism and the complexity of the V-C dimension. *Journal of Computer and System Sciences*, 53(2):161–170, 1996.
22. G. H. Polychronopoulos and J. N. Tsitsiklis. Stochastic shortest path problems with recourse. *Networks*, 27(2):133–143, 1996.