

Approximation Algorithms for Route Planning with Nonlinear Objectives

Ger Yang

Electrical and Computer Engineering
University of Texas at Austin
geryang@utexas.edu

Evdokia Nikolova

Electrical and Computer Engineering
University of Texas at Austin
nikolova@austin.utexas.edu

Abstract

We consider optimal route planning when the objective function is a general nonlinear and non-monotonic function. Such an objective models user behavior more accurately, for example, when a user is risk-averse, or the utility function needs to capture a penalty for early arrival. It is known that as nonlinearity arises, the problem becomes NP-hard and little is known about computing optimal solutions when in addition there is no monotonicity guarantee. We show that an approximately optimal non-simple path can be efficiently computed under some natural constraints. In particular, we provide a fully polynomial approximation scheme under hop constraints. Our approximation algorithm can extend to run in pseudo-polynomial time under a more general linear constraint that sometimes is useful. As a by-product, we show that our algorithm can be applied to the problem of finding a path that is most likely to be on time for a given deadline.

Introduction

In this paper, we present approximation algorithms for route planning with a nonlinear objective function. Traditionally, route planning problems are modeled as linear shortest path problems and there are numerous algorithms, such as Bellman-Ford or Dijkstra's algorithm, that are able to find the optimal paths efficiently. However, nonlinearity arises when we would like to make more complicated decision-making rules. For example, in road networks, the objective might be a nonlinear function that represents the trade-off between traveling time and cost. Unfortunately, dynamic programming techniques used in linear shortest path algorithms no longer apply as when nonlinearity appears, sub-paths of optimal paths may no longer be optimal.

Applications include risk-averse routing or routing with a nonlinear objective when we are uncertain about the traffic conditions (Nikolova, Brand, and Karger 2006; Nikolova et al. 2006). A risk-averse user tends to choose a route based on both the speediness and reliability. Such risk-averse attitude can often be captured by a nonlinear objective, e.g. a mean-risk objective (Nikolova 2010). On the other hand, the user might seek a path that maximizes the probability for him to be on time given a deadline. This objective was considered by Nikolova (2010) to capture the risk-averse behavior. For

Gaussian travel times, the latter paper only considered the case where at least one path has mean travel time less than the deadline, and consequently the objective function can be modeled as a monotonic function.

The problem of finding a path that minimizes a monotonic increasing objective has been studied in the literature (Goyal and Ravi 2013; Nikolova 2010; Tsaggouris and Zaroliagis 2009). What if there is no monotonicity guarantee? For example, consider the deadline problem when we know that all paths have mean travel time longer than the deadline, or consider the well-known cost-to-time ratio objective (Megiddo 1979; Guérin and Orda 1999) in various combinatorial optimization problems. Without the monotonicity assumption, the problem usually becomes very hard. It has been shown (Nikolova, Brand, and Karger 2006) that for a general objective function, finding the optimal *simple* path is NP-hard. This is because such a problem often involves finding the longest path for general graphs, which is known to be a strongly NP-hard problem (Karger, Motwani, and Ramkumar 1997), namely finding a constant-factor approximation is also NP-hard. This suggests that for a very general class of functions, including the cost-to-time ratio objective (Ahuja, Batra, and Gupta 1983), not only is the problem itself NP-hard but so is also its approximation counterpart.

Therefore, in this paper, we focus on the problem that accepts *non-simple* paths, in which it is allowed to visit nodes more than once. To make this problem well-defined, we consider the problem under either a hop constraint or an additional linear constraint, where we say a path is of γ hops if it contains γ edges. We design a fully polynomial approximation scheme under hop constraints, and show that the algorithm can extend to run in pseudo-polynomial time under an additional linear constraint. Further, we show how our algorithm can be applied to the cost-to-time ratio objective and the deadline problem.

Related Work

The route planning problem with nonlinear objective we consider here is related to multi-criteria optimization (e.g., Papadimitriou and Yannakakis 2000). The solutions to these problems typically look for approximate Pareto sets, namely a short list of feasible solutions that provide a good approximation to any non-dominated or optimal solution. A fully polynomial approximation scheme (FPTAS) for route plan-

ning with general monotonic and smoothed objective functions was presented by Tsaggouris and Zaroliagis (2009), based on finding an approximate Pareto set for the feasible paths. Later, Mittal and Schulz (2013) provided a general approach for monotonic and smoothed combinatorial optimization problems and showed how to reduce these problems to multi-criteria optimization.

A special class of nonlinear optimization is *quasi-concave minimization*. This class of problems has the property that the optimal solution is an extreme point of the feasible set (Horst, Pardalos, and Van Thoai 2000). For quasi-concave combinatorial minimization problems, Kelner and Nikolova (2007) proved $(\log n)$ -hardness of approximation and they also gave an additive approximation algorithm for low-rank quasi-concave functions based on smoothed analysis. Later, Nikolova (2010) presented an FPTAS for risk-averse quasi-concave objective functions. In recent work, Goyal and Ravi (2013) gave an FPTAS for general monotone and smooth quasi-concave functions.

Without assuming monotonicity, the problem was studied by Nikolova, Brand, and Karger (2006), who proved hardness results and gave a pseudo-polynomial approximation algorithm with integer edge weights for some specific objective functions.

Problem Statement and Preliminaries

Consider a directed graph $G = (V, E)$. Denote the number of vertices as $|V| = n$, and the number of edges as $|E| = m$. Suppose we are given a source node s and a destination node t . In the *nonlinear objective shortest path* problem, we seek an $s - t$ path p , possibly containing loops, that minimizes some nonlinear objective function $f(p)$.

More precisely, denote the feasible set of all $s - t$ paths by \mathcal{P}_t , and let $\mathcal{P} = \bigcup_{v \in V} \mathcal{P}_v$ be the set of paths that start from s . Here we have not yet imposed assumptions that \mathcal{P} contains only simple paths. Let $f : \mathcal{P} \rightarrow \mathbb{R}_+$ be a given objective function. Then we can write the problem as finding the path p^* that minimizes this objective function:

$$p^* = \arg \min_{p \in \mathcal{P}_t} f(p) \quad (1)$$

We assume that f is a low-rank nonlinear function with some smoothness condition. Specifically, we say that f is of rank d if there exists a function $g : \mathbb{R}_+^d \rightarrow \mathbb{R}_+$ and an $m \times d$ weight matrix $\mathbf{W} = (w_{e,k})_{m \times d}$, with positive entries, such that for any path $p \in \mathcal{P}$, $f(p) = g(\sum_{e \in p} w_{e,1}, \sum_{e \in p} w_{e,2}, \dots, \sum_{e \in p} w_{e,d})$.

We interpret this as having d criteria, such as for example length in miles, cost, travel time, etc. so that each edge is associated with d additive weights, where the k -th weight for edge e is denoted by $w_{e,k}$. Then, the k -th weights of all edges can be represented by the column vector $\mathbf{w}_k = (w_{1,k}, w_{2,k}, \dots, w_{m,k})^T$. For simplicity, we denote the sum of the k -th weights of the edges along path p as $l_k(p) = \sum_{e \in p} w_{e,k}$, and we call it *the k -th criterion of path p* . Hence, the objective function f can be rewritten as a function of all criteria:

$$f(p) = g(l_1(p), l_2(p), \dots, l_d(p)) \quad (2)$$




Figure 1: Illustration of a path polytope on a doubly weighted graph with a hop constraint. Each blue cross and red dot correspond to an $s - t$ path projected onto \mathcal{W} . We can use polynomially many red dots to approximate the path polytope and consequently to approximate the optimal path.

Consider the linear space spanned by the d criteria $\mathcal{W} = \text{span}(\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_d)$. Then, each feasible $s - t$ path $p \in \mathcal{P}_t$ has an associated point in this d -dimensional space \mathcal{W} given by $\mathbf{l}(p) = (l_1(p), l_2(p), \dots, l_d(p))$, providing the values of the d criteria for that path. We shall call the convex hull of such points the *path polytope*.¹ We assume that the weights matrix \mathbf{W} contains only positive entries, and thus the path polytope is bounded from below: for any path $p \in \mathcal{P}_t$ and criterion $k \in \{1, \dots, d\}$, we have $l_k(p) \geq \min_{p' \in \mathcal{P}_t} l_k(p')$. This lower-bound corresponds to the shortest path with respect to the k -th criterion. For upper bounds of the polytope, we defer the discussion to the later sections.

The size of the path polytope is in worst case exponential in the problem size. The main idea of this paper is to approximate the path polytope with polynomially many paths, as illustrated in Figure 1. As we will show in this paper, the optimal path can be approximated if there are some smoothness conditions on the objective function. The smoothness assumption we adapt here is as follows: g is assumed to be β -Lipschitz on the log-log scale with respect to the L_1 -norm: that is, for any \mathbf{y}, \mathbf{y}' in the path polytope,

$$|\log g(\mathbf{y}) - \log g(\mathbf{y}')| \leq \beta \|\log \mathbf{y} - \log \mathbf{y}'\|_1 \quad (3)$$

where $\log \mathbf{y}$ is defined as a column vector: $\log \mathbf{y} = (\log y_1, \log y_2, \dots, \log y_d)^T$.

In this paper, we design approximation algorithms for this problem, formally defined as follows:

Definition 1. An α -approximation algorithm for a minimization problem with optimal solution value OPT and $\alpha > 1$ is a polynomial-time algorithm that returns a solution of value at most αOPT .

¹Strictly speaking, it is the projection of the feasible path polytope in \mathbb{R}^m onto this d -dimensional subspace but for simplicity we will just call it the path polytope in this paper.

The best possible approximation algorithms provide solutions that are arbitrarily close to the optimal solution:

Definition 2. A fully-polynomial approximation scheme (FPTAS) is an algorithm for an optimization problem that, given desired accuracy $\epsilon > 0$, finds in time polynomial in $\frac{1}{\epsilon}$ and the input size, a solution of value OPT' that satisfies $|OPT - OPT'| \leq \epsilon OPT$, for any input, where OPT is the optimal solution value.

Hardness Results

We can categorize the problem on general directed graphs into two types: one is the problem that accepts only simple paths, by which we mean each vertex can only be visited once in any feasible path; the other is the problem that accepts non-simple paths. Unfortunately, both types of problems turn out to be hard. In fact, the problem that only accepts simple paths in general involves finding the longest simple path, which is strongly NP-hard, namely it is NP-hard to approximate to within any constant factor (Karger, Motwani, and Ramkumar 1997).

First, we present a natural objective function that has been considered in the literature, that is hard to approximate.

Minimum Cost-to-Time Ratio Path Problem

The cost-to-time ratio objective function is a rank 2 function, defined as follows:

$$f(p) = g\left(\sum_{e \in p} w_{e,1}, \sum_{e \in p} w_{e,2}\right) = \frac{\sum_{e \in p} w_{e,1}}{\sum_{e \in p} w_{e,2}}$$

For simplicity, we assume \mathbf{W} has strictly positive entries. This objective function has been considered previously for different combinatorial optimization problems (Megiddo 1979; Correa, Fernandes, and Wakabayashi 2010). However, finding the minimum cost-to-time ratio simple path is known to be an NP-hard problem (Ahuja, Batra, and Gupta 1983). Here, we show its hardness of approximation:

Theorem 1. *There is no polynomial-time α -approximation algorithm for the minimum cost-to-time ratio simple path problem for any constant factor α , unless $P = NP$.*

Proof. To facilitate the presentation of our hardness of approximation proof, we first show that this problem is NP-hard. We reduce a longest path problem instance to a minimum cost-to-time ratio simple path problem instance. The longest path problem is known to be NP-hard to approximate to within any constant factor (Karger, Motwani, and Ramkumar 1997). First, we choose some $\lambda > 0$. For an instance of the longest path problem, which is some graph G with unweighted edges and a designated source node s and sink node t , we create an instance of the minimum cost-to-time ratio path problem with the same graph G , source s and sink t and the following (cost,time) edge weights:

1. For each edge e incident to s , set $w_e = (\lambda n + 1, 1)$.
2. For the remaining edges e , set $w_e = (1, 1)$.

We can see that a path is optimal under the ratio objective if and only if it is the longest path with respect to the unweighted edges (or, equivalently, with respect to the second weight). This shows that the minimum cost-to-time ratio simple path problem is NP-hard.

To show the hardness of approximation, similarly, we use the fact that the longest path problem cannot be approximated to within any constant factor unless $P = NP$ (Karger, Motwani, and Ramkumar 1997). Suppose the contrary, namely that we can find a $(1 + \epsilon)$ -approximation algorithm for the minimum cost-to-time ratio simple path problem, for some $\epsilon \in (0, \lambda)$. To be precise, since we can choose λ while ϵ comes from the assumed approximation algorithm, we can set $\lambda > \epsilon$. For a given instance of longest path, consider the corresponding instance of minimum cost-to-time ratio path problem above. Let p^* be the optimal min cost-to-time ratio path. Then, the $(1 + \epsilon)$ -approximation algorithm will return a simple path p such that

$$\frac{\lambda n + l_2(p)}{l_2(p)} \leq (1 + \epsilon) \frac{\lambda n + l_2(p^*)}{l_2(p^*)}$$

Rearranging this, we get the inequality:

$$(\lambda n - \epsilon l_2(p))l_2(p^*) \leq \lambda(1 + \epsilon)nl_2(p)$$

Since $l_2(p) \leq n$ and $\epsilon < \lambda$, we get the following:

$$l_2(p^*) \leq \frac{\lambda(1 + \epsilon)}{\lambda - \epsilon} l_2(p)$$

This shows that we can approximate the longest path problem to within a constant factor of $\frac{\lambda(1 + \epsilon)}{\lambda - \epsilon}$, a contradiction. Therefore, the minimum cost-to-time ratio simple path problem cannot be approximated to within any constant factor unless $P = NP$. \square

General Objective Function

For general objective functions, the following hardness results were proven by Nikolova, Brand, and Karger (2006):

Fact 1. *Suppose function g attains a global minimum at $\mathbf{y} > 0$. Then problem (1) of finding an optimal $s - t$ simple path is NP-hard.*

Fact 2. *Consider problem (1) of finding an optimal $s - t$ simple path with f being a rank-1 function. Let p^* be the optimal path, and $l^* = l(p^*)$. If g is any strictly decreasing and positive function with slope of absolute value at least $\lambda > 0$ on $[0, l^*]$, then there does not exist a polynomial-time constant factor approximation algorithm for finding an optimal simple path, unless $P = NP$.*

Although these facts are proven for rank-1 functions, they show that for a large class of objective functions, the problem of finding the optimal simple path is not only NP-hard, but also NP-hard to approximate within any constant factor.

As a result, we turn our attention to the problem that accepts non-simple paths. While this problem is also NP-hard (Nikolova, Brand, and Karger 2006), we are able to design a polynomial-time algorithm to approximate the optimal solution arbitrarily well.

Approximation Algorithms

Consider problem (1) that accepts non-simple paths. For some cases, it is possible that there does not exist an optimal path of finite length. Consider an example of a rank 1 monotonic decreasing function $g(x) = 1/x$. Suppose there is a positive weight loop in any of the $s - t$ paths. Then, we can always find a better solution by visiting such a loop arbitrarily many times: the optimal path here is ill-defined.

Naturally, for the problem that accepts non-simple paths, the path polytope is infinite. Even when the optimal path is finite, for a large class of problems, including the minimum cost-to-time ratio path problem, finding a finite subset of the path polytope that contains the optimal path is difficult, as shown in our Hardness section above. In this section, we will show that, under some additional constraints, this problem becomes well-defined and well-approximable. In what follows, we identify the additional constraints and present the algorithms to solve the corresponding problems.

Hop Constraint

Suppose the path contains at most γ hops, then the set of feasible paths \mathcal{P}_t becomes finite and the problem is well-defined. Denote $\text{hop}(p)$ to be the number of hops in path p , or equivalently, the number of edges in p . Note that an edge can be counted multiple times if path p contains loops.

To approximate the optimal path, we approximate the path polytope with polynomially many paths, and then search for the best one among them. To approximate the path polytope, we first identify a region of weights space \mathcal{W} that contains all paths starting from s , which are within γ edges. We can upper- and lower-bound this region as follows: For each criterion $k = 1, \dots, d$, define $c_k^{\min} = \min_{e \in E} w_{e,k}$ and $c_k^{\max} = \max_{e \in E} w_{e,k}$. Then for any path p that starts from the source s and contains at most γ edges, we have $l_k(p) \in [c_k^{\min}, \gamma c_k^{\max}]$, for each criterion $k \in \{1, \dots, d\}$.

Next, we partition the region $\prod_{k=1}^d [c_k^{\min}, \gamma c_k^{\max}]$ of the weights space \mathcal{W} into polynomially many hypercubes. Specifically, for each criterion $k \in \{1, 2, \dots, d\}$, let $\epsilon > 0$ be a given desired approximation factor, and $C_k = c_k^{\max}/c_k^{\min}$. Define $\Gamma_k = \{0, 1, \dots, \lfloor \log_{(1+\epsilon)}(\gamma C_k) \rfloor\}$ as the set of points that evenly partitions the segment $[c_k^{\min}, \gamma c_k^{\max}]$ on the log scale. Then, we generate a d -dimensional lattice $\Gamma = \Gamma_1 \times \Gamma_2 \times \dots \times \Gamma_d$ according to these sets. After that, we define a hash function $\mathbf{h} = (h_1, h_2, \dots, h_d) : \mathcal{P} \rightarrow \Gamma$ that maps a path to a point on the lattice Γ . Each $h_k : \mathcal{P} \rightarrow \Gamma_k$ is defined as $h_k(p) = \lfloor \log_{(1+\epsilon)} \frac{l_k(p)}{c_k^{\min}} \rfloor$.

The main part of the algorithm iterates over $i \in \{1, 2, \dots, \gamma\}$, sequentially constructing $s - v$ sub-paths containing at most i edges, for each vertex $v \in V$. For each iteration i and each vertex $v \in V$, we maintain at most $|\Gamma|$ configurations as a d -dimensional table Π_v^i , indexed by $\mathbf{y} \in \Gamma$. Each configuration $\Pi_v^i(\mathbf{y})$ can be either an $s - v$ path p of at most i edges and $h(p) = \mathbf{y}$, given there is one, or *null* otherwise. The table can be implemented in various ways.² Here we assume that given index $\mathbf{y} \in \Gamma$, both check-

ing whether the configuration $\Pi_v^i(\mathbf{y})$ is *null*, retrieving the path in $\Pi_v^i(\mathbf{y})$, and inserting/updating a path in $\Pi_v^i(\mathbf{y})$ can be done in $T_\Pi = T_\Pi(n, |\Gamma|)$ time.

At iteration i , for each vertex u , we will construct a set of $s - u$ sub-paths containing i edges based on $(i - 1)$ -edge sub-paths we previously constructed. We loop over each vertex $v \in V$ and each edge $e = (v, u)$ incident to v , and we append the edge to the end of each $(i - 1)$ -edge $s - v$ path in Π_v^{i-1} . However, not all generated sub-paths will be stored. Consider any $s - u$ path p generated with this procedure, it will be stored in table Π_u^i only if $\Pi_u^i(\mathbf{h}(p))$ is empty. With this technique, we can ensure that the number of sub-paths we record in each iteration is at most polynomial, because it is bounded by the size of the lattice Γ times the number of vertices. This procedure is outlined in Algorithm 1.

Algorithm 1 Nonlinear Objective Shortest Path with Hop-constraint

```

1: for each  $\mathbf{y} \in \Gamma$ ,  $v \in V$ , and  $i = 1, 2, \dots, \gamma$  do
2:   Set  $\Pi_v^i(\mathbf{y}) = \text{null}$ 
3: Set  $\Pi_s^0 = \{s\}$ .
4: for  $i = 1$  to  $\gamma$  do
5:   Set  $\Pi_v^i = \Pi_v^{i-1}$  for each vertex  $v \in V$ .
6:   for each vertex  $v \in V$  do
7:     for each path  $p \in \Pi_v^{i-1}$  with  $i - 1$  edges do
8:       for each edge  $e = (v, u) \in E$  do
9:         Set  $p'$  be the path that appends edge  $e$  to the
           tail of  $p$ .
10:        Compute hash  $\mathbf{y} = \mathbf{h}(p')$ .
11:        if  $\Pi_u^i(\mathbf{y}) = \text{null}$  then
12:           $\Pi_u^i(\mathbf{y}) = p'$ .
13: return  $\max_{p \in \Pi_t^\gamma} f(p)$ 

```

To prove the correctness of this algorithm, we show the following lemmas:

Lemma 1. *Let p_u and p'_u be two $s - u$ paths such that for each criterion $k \in \{1, 2, \dots, d\}$,*

$$|\log l_k(p_u) - \log l_k(p'_u)| \leq \log(1 + \epsilon). \quad (4)$$

Suppose there is a vertex v and an edge e such that $e = (u, v) \in E$. Let p_v and p'_v be the $s - v$ paths that append (u, v) to the tails of p_u and p'_u , respectively. Then, for each criterion $k \in \{1, 2, \dots, d\}$,

$$|\log l_k(p_v) - \log l_k(p'_v)| \leq \log(1 + \epsilon)$$

Proof. By (4), it directly follows that

$$\begin{aligned} l_k(p_v) &= l_k(p_u) + l_k(e) \leq (1 + \epsilon)l_k(p'_u) + l_k(e) \\ &\leq (1 + \epsilon)(l_k(p'_u) + l_k(e)) = (1 + \epsilon)l_k(p'_v) \\ l_k(p'_v) &= l_k(p'_u) + l_k(e) \leq (1 + \epsilon)l_k(p_u) + l_k(e) \\ &\leq (1 + \epsilon)(l_k(p_u) + l_k(e)) = (1 + \epsilon)l_k(p_v) \end{aligned}$$

Combining the above two inequalities yields the proof. \square

the used entries of the table. According to our experimental results, only a small portion of the table Π_v^i will be used.

²In practice, we can either allocate the memory of the entire table in the beginning, or use another data structure that only stores

Lemma 2. For any $v \in V$, and for any path $s - v$ that is of at most i edges, there exists an $s - v$ path $p' \in \Pi_v^i$ such that for each criterion $k \in \{1, 2, \dots, d\}$,

$$|\log l_k(p) - \log l_k(p')| \leq i \log(1 + \epsilon)$$

Proof. We prove this lemma by induction on the number of edges in a path. The base case is a direct result from the definition of the hash function \mathbf{h} . For the induction step, consider vertex $v \in V$, and we would like to prove the statement is correct for any $s - v$ sub-paths of at most i^* edges, given the statement is true for any $i < i^*$ and any vertices. Let p_v be any $s - v$ path no more than i^* edges. If p_v is of $i' < i^*$ edges, then by the induction hypothesis, there exists an $s - v$ path p'_v in $\Pi_v^{i'}$ such that $|\log(l_k(p_v)/l_k(p'_v))| \leq i' \log(1 + \epsilon)$. The lemma is proved for this case since p'_v is in $\Pi_v^{i^*}$ as well. Next consider p_v to be an $s - v$ path with i^* edges, and suppose p_v is formed by appending the edge (u, v) to an $s - u$ path p_u . By the induction hypothesis, we can see that there is an $s - u$ path $p'_u \in \Pi_u^{i^*-1}$ such that for each criterion $k \in \{1, 2, \dots, d\}$,

$$|\log l_k(p_u) - \log l_k(p'_u)| \leq (i^* - 1) \log(1 + \epsilon)$$

Then, with Lemma 1, as we form $s - v$ sub-paths p_v and p'_v by appending (u, v) to the end of p_u and p'_u respectively, we have the following relationship for each criterion $k \in \{1, 2, \dots, d\}$:

$$|\log l_k(p_v) - \log l_k(p'_v)| \leq (i^* - 1) \log(1 + \epsilon) \quad (5)$$

Consider any $s - v$ path p''_v such that $\mathbf{h}(p''_v) = \mathbf{h}(p'_v)$. We must have $|\log l_k(p'_v) - \log l_k(p''_v)| \leq \log(1 + \epsilon)$ for each criterion $k \in \{1, \dots, d\}$, according to the definition of the hash function \mathbf{h} . Combining this with (5), we can see that even for the case that path p'_v is not stored in the table $\Pi_v^{i^*}$, the path $p''_v = \Pi_v^{i^*}(p'_v)$ must satisfy

$$|\log l_k(p_v) - \log l_k(p''_v)| \leq i^* \log(1 + \epsilon)$$

for each criterion $k \in \{1, \dots, d\}$. This indicates that for any $s - v$ path p_v within i^* edges, there exists a path p''_v in $\Pi_v^{i^*}$ such that their ratio is bounded by $(1 + \epsilon)^{i^*}$. Q.E.D. \square

Now, we are ready to state and prove the main result:

Theorem 2. Suppose g defined in (2) is β -Lipschitz on the log-log scale. Then Algorithm 1 is a $(1 + \epsilon)^{\beta d \gamma}$ -approximation algorithm for the γ -hop-constrained nonlinear objective shortest path problem. Besides, Algorithm 1 has time complexity $O(\gamma m (\frac{\log(\gamma C)}{\log(1 + \epsilon)})^d T_{\Pi})$, where $C = \max_{k \in \{1, 2, \dots, d\}} C_k$.

Proof. By Lemma 2, we can see that for any $s - t$ path p of at most γ edges, there exists an $s - t$ path $p' \in \Pi_t^\gamma$ such that

$$|\log l_k(p) - \log l_k(p')| \leq \gamma \log(1 + \epsilon), \quad \forall k = 1, 2, \dots, d$$

It then follows from (3) that our algorithm finds a path p' such that $|\log f(p') - \log f(p^*)| \leq \beta d \gamma \log(1 + \epsilon)$. Then, we have

$$f(p') \leq (1 + \epsilon)^{\beta d \gamma} f(p^*)$$

This shows that p' is an $(1 + \epsilon)^{\beta d \gamma}$ -approximation to p^* .

Next, we show the time complexity. It is easy to see that at each iteration, the algorithm generates at most $m|\Gamma|$ sub-paths. Each sub-path requires at most $O(T_{\Pi})$ time to operate. Since $|\Gamma| = (\frac{\log(\gamma C)}{\log(1 + \epsilon)})^d$, we can conclude that Algorithm 1 runs in time $O(\gamma m (\frac{\log(\gamma C)}{\log(1 + \epsilon)})^d T_{\Pi})$. \square

Note that we can obtain a $(1 + \delta)$ -approximation algorithm by setting $\delta = (1 + \epsilon)^{\beta d \gamma} - 1$. With this setting, then Algorithm 1 has time complexity $O(\gamma m (\frac{\beta d \gamma \log(\gamma C)}{\delta})^d T_{\Pi})$. If we assume $\gamma = O(n)$, then Algorithm 1 is indeed an FPTAS.

Linear Constraint

The hop constraint we considered can be relaxed: we now consider the more general problem that any returned path p must satisfy an additional linear constraint $\sum_{e \in p} w_{e,d+1} \leq b$, for some budget $b > 0$ and some weight vector $\mathbf{w}_{d+1} \in \mathbb{R}_{+}^m$ with strictly positive entries. We define the $(d + 1)$ -th criterion $l_{d+1}(p) = \sum_{e \in p} w_{e,d+1}$, which can be interpreted as having a budget b for a path. To solve this problem, we can find an upper bound of the number of edges for any feasible $s - t$ path. A simple upper bound can be given by

$$\gamma = \left\lceil \frac{b - \sum_{e \in p_{d+1}^*} w_{e,d+1}}{\min_{e \in E} w_{e,d+1}} + \text{hop}_{\min} \right\rceil \quad (6)$$

where $p_{d+1}^* = \arg \min_{p \in \mathcal{P}} \sum_{e \in p} w_{e,d+1}$ is the shortest path with respect to the weight \mathbf{w}_{d+1} , and hop_{\min} is the minimum number of edges for any $s - t$ path. Note that either hop_{\min} and p_{d+1}^* can be solved by Dijkstra's or Bellman-Ford's shortest path algorithms efficiently.

Then, we can adapt the techniques discussed for the hop-constrained problem but with some modifications. Instead of assigning the least hop path to a configuration, we now assign the path with the least $(d + 1)$ -th criterion. Formally, the modification is to change line 11 and line 12 of Algorithm 1 to: Assign p' to $\Pi_u^i(\mathbf{y})$ if p' satisfies both of the following:

- $l_{d+1}(p') < b$
- $\Pi_u^i(\mathbf{y})$ is null or $l_{d+1}(p') < l_{d+1}(\Pi_u^i(\mathbf{y}))$

As a corollary of Theorem 2, we have the following result:

Corollary 1. Algorithm 1 with the modifications stated above is a $(1 + \epsilon)^{\beta d \gamma}$ -approximation algorithm for the linear-constrained nonlinear objective shortest path problem, where γ is defined as in Eq. (6).

The proof of this corollary is presented in the full version of this paper (Yang and Nikolova 2015).

Application: Deadline Problem

With the techniques mentioned in the above section, we are able to solve various problems, including the minimum cost-to-time ratio path problem with either a hop or a linear budget constraint. In this section, we discuss a related problem, which is the *deadline problem* in a stochastic setting.

The deadline problem is to find a path that is the most likely to be on time, which is defined as follows: Suppose each edge is associated with an independent random variable X_e , which we call the travel time of edge e . Given

Table 1: Deadline Problem with Hop Constraint

Network	Hops	Worst-case Theoretical Accuracy (α)	Memory Usage (10^6 paths)	Run time (s)
5×5 Grid	12	(Exhaustive)	1.63	3.441
	12	2.2923	0.67	2.057
	15	2.2926	4.30	20.276
10×10 Grid	15	2.2891	18.05	83.430
Anaheim	15	3.0436	9.60	37.671

a deadline $D > 0$, we are going to find an $s - t$ path such that the probability of being on time is maximized: $\max_{p \in \mathcal{P}_t} \Pr(\sum_{e \in p} X_e < D)$. We assume that the random variables $\{X_e\}$ are Gaussian. Based on this, we denote the mean travel time for edge e as μ_e and variance as σ_e^2 . For simplicity, we assume for each edge $e \in E$, its mean μ_e and variance σ_e^2 are bounded by $\mu_e \in [\mu_{\min}, \mu_{\max}]$, and $\sigma_e^2 \in [\sigma_{\min}^2, \sigma_{\max}^2]$. The problem can hence be written as:

$$\max_{p \in \mathcal{P}_t} \Phi\left(\frac{D - \sum_{e \in p} \mu_e}{\sqrt{\sum_{e \in p} \sigma_e^2}}\right) \quad (7)$$

where $\Phi(\cdot)$ is the cumulative distribution function for the standard normal distribution. The objective function in (7) is a rank 2 function. This problem is studied in (Nikolova et al. 2006) for a risk-averse shortest path problem, where they considered the case that the optimal path has mean travel time less than the deadline: for this case (7) is a monotonic decreasing function. In this paper, we consider the case that all feasible $s - t$ paths have mean travel time strictly greater than the deadline. This means that we assume each path has poor performance. Then, the objective function in (7) becomes monotonic decreasing with respect to the mean but monotonic increasing with respect to the variance, for which known algorithms no longer apply. However, we are able to use techniques presented in this paper to solve this problem under either a hop or an additional linear constraint.

Theorem 3. *Suppose each $s - t$ path $p \in \mathcal{P}_t$ satisfies $\sum_{e \in p} \mu_e > D$ and $\sum_{e \in p} \sigma_e^2 > S$. For the deadline problem with a hop constraint γ , if Algorithm 1 returns some path p' such that $f(p') > \Phi(-3) \approx 0.0013$, then p' is an α -approximation for the deadline problem, where $\alpha = \min\{384.62, (1 + \epsilon)^{6.568(3 + \frac{D}{\sqrt{S}})\gamma}\}$. Similarly, if p' is such that $f(p') > \Phi(-2) \approx 0.023$, then $\alpha = \min\{21.93, (1 + \epsilon)^{4.745(2 + \frac{D}{\sqrt{S}})\gamma}\}$.*

Note that this theorem does not guarantee Algorithm 1 admits FPTAS for the deadline problem, since given α , the running time is polynomial in the factor D/\sqrt{S} .

Experimental Evaluations

We tested our algorithm on a grid graph and a real transportation network. The algorithm is implemented in C++. Experiments are run on an Intel Core i7 1.7 GHz (I7-4650U) processor. Due to memory constraints, the table $\{\Pi_v^i | v \in V\}$ is implemented as a binary search tree, where query, insertion, and update operations can be done in time $T_{\Pi} =$




Figure 2: Memory usage for running our algorithm on a 5×5 grid graph to the number of iterations.

$O(\log(n|\Gamma|))$. Memory turns out to be the limiting factor for the problem size and accuracy available to our experiments. In this section, *memory usage* refers to the total number of sub-paths maintained in the table $\{\Pi_v^i | v \in V\}$.

We tested the deadline problem with a hop constraint on a 5×5 grid graph, 10×10 grid graph, and the Anaheim network (416 vertices, 914 edges) from the real world data set (Bar-Gera 2002). The grid graphs are bi-directional, and the mean and the variance for each edge are randomly generated from $[0.1, 5]$. We routed from node $(0, 0)$ to node $(4, 4)$ on the 5×5 grid, and from $(1, 1)$ to $(8, 8)$ on the 10×10 grid. The Anaheim dataset provides the mean travel time. The variance for each edge is randomly generated from 0 to the mean. The source and destination are randomly chosen.

Table 1 summarizes the results on the datasets, which are the average of 20 experiments. Even though theoretically the algorithm can be designed for arbitrary desired accuracy, it is limited by the memory of the environment.³ However, comparing with exhaustive search, we found that the practical accuracy is 100% on 5×5 grid graphs.⁴ This means that our algorithm is far more accurate than its theoretical worst case bound. We explain this phenomenon through Figure 2, which shows the number of sub-paths generated up to each iteration. We see that the number of sub-paths generated in the initial iterations grows exponentially, just as in exhaustive search, but levels off eventually. Recall Lemma 2, which tells us that an error is induced when a sub-path is eliminated due to similarity to a saved sub-path, and Figure 2 suggests that this rarely happens in the initial iterations.

Conclusion

We investigated the shortest path problem with a general nonlinear and non-monotonic objective function. We proved that this problem is hard to approximate for simple paths and

³It takes about 3GB physical memory usage for the 18M paths in the 10×10 grid network.

⁴Exhaustive search becomes prohibitive for paths with more than 12 hops on 5×5 grid graphs.

consequently, we designed and analyzed a fully polynomial approximation scheme for finding the approximately optimal non-simple path under either a hop constraint or an additional linear constraint. We showed that this algorithm can be applied to the cost-to-time ratio problem and the deadline problem. Our experiments show that our algorithm is capable of finding good approximations efficiently.

Acknowledgement

This work was supported in part by NSF grant numbers CCF-1216103, CCF-1350823, CCF-1331863, and a Google Faculty Research Award.

References

Ahuja, R.; Batra, J.; and Gupta, S. 1983. Combinatorial optimization with rational objective functions: A communication. *Mathematics of Operations Research* 8(2):314–314.

Bar-Gera, H. 2002. Transportation network test problems.

Correa, J.; Fernandes, C. G.; and Wakabayashi, Y. 2010. Approximating a class of combinatorial problems with rational objective function. *Mathematical programming* 124(1-2):255–269.

Goyal, V., and Ravi, R. 2013. An fptas for minimizing a class of low-rank quasi-concave functions over a convex set. *Operations Research Letters* 41(2):191–196.

Guérin, R. A., and Orda, A. 1999. Qos routing in networks with inaccurate information: theory and algorithms. *IEEE/ACM Transactions on Networking (TON)* 7(3):350–364.

Horst, R.; Pardalos, P. M.; and Van Thoai, N. 2000. *Introduction to global optimization*. Springer Science & Business Media.

Karger, D.; Motwani, R.; and Ramkumar, G. 1997. On approximating the longest path in a graph. *Algorithmica* 18(1):82–98.

Kelner, J. A., and Nikolova, E. 2007. On the hardness and smoothed complexity of quasi-concave minimization. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, 472–482. IEEE.

Megiddo, N. 1979. Combinatorial optimization with rational objective functions. *Mathematics of Operations Research* 4(4):414–424.

Mittal, S., and Schulz, A. S. 2013. An fptas for optimizing a class of low-rank functions over a polytope. *Mathematical Programming* 141(1-2):103–120.

Nikolova, E.; Kelner, J. A.; Brand, M.; and Mitzenmacher, M. 2006. Stochastic shortest paths via quasi-convex maximization. In *Algorithms–ESA 2006*, 552–563. Springer.

Nikolova, E.; Brand, M.; and Karger, D. R. 2006. Optimal route planning under uncertainty. In *ICAPS*, volume 6, 131–141.

Nikolova, E. 2010. Approximation algorithms for offline risk-averse combinatorial optimization.

Papadimitriou, C. H., and Yannakakis, M. 2000. On the approximability of trade-offs and optimal access of web

sources. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, 86–92. IEEE.

Tsaggouris, G., and Zaroliagis, C. 2009. Multiobjective optimization: Improved fptas for shortest paths and non-linear objectives with applications. *Theory of Computing Systems* 45(1):162–186.

Yang, G., and Nikolova, E. 2015. Approximation algorithms for route planning with nonlinear objectives. *arXiv preprint arXiv:1511.07412*.

Appendix: Proof of Theorem 3

In order to prove Theorem 3, we need the following lemma:

Lemma 3. *The function $\Phi(\frac{D-x}{\sqrt{y}})$ is $3.284(3 + \frac{D}{\sqrt{S}})$ -Lipschitz on log-log scale with respect to L1 norm on the region $\mathcal{C} = \{(x, y) | x > D, y > S, x - D \leq 3\sqrt{y}\}$.*

The proof of the lemma is omitted here and will be presented in the full version of this paper. With this lemma, we are ready to prove Theorem 3:

Proof of Theorem 3. Let $l_1(p)$ denote the mean traveling time for path p , and $l_2(p)$ denote the variance for path p . Further denote $\mathbf{l}(p) = (l_1(p), l_2(p))$. By Lemma 2, we can see that for the optimal $s - t$ path p^* , there exists an $s - t$ path $p'' \in \Pi_t^s$ such that

$$|\log l_k(p^*) - \log l_k(p'')| \leq \gamma \log(1 + \epsilon)$$

Denote the region $\mathcal{C} = \{(x, y) | x > D, y > S, x - D \leq 3\sqrt{y}\}$. We can see that $(x, y) \in \mathcal{C}$ if and only if $\Phi(\frac{D-x}{\sqrt{y}}) \in [\Phi(-3), 0.5]$ and $y > S$. The optimal path must satisfy $\mathbf{l}(p^*) \in \mathcal{C}$ since $f(p^*) \geq f(p') > \Phi(-3)$.

First consider the case that $\mathbf{l}(p'')$ is in the region \mathcal{C} . Then we can adapt the same technique as in the proof of Theorem 2 and the smoothness condition we showed in Lemma 3 to claim that the path p' returned by the algorithm satisfies

$$f(p^*) \leq (1 + \epsilon)^{2 \times 3.284 \times (3 + \frac{D}{\sqrt{S}}) \gamma} f(p')$$

Now consider the case that $\mathbf{l}(p'')$ is not in the region \mathcal{C} , then there must exist a point $(x_0, y_0) \in \mathcal{C}$ on the line $L = \{x - D = 3\sqrt{y}\}$ such that

$$\begin{aligned} |\log x_0 - \log l_1(p^*)| &\leq \gamma \log(1 + \epsilon) \\ |\log y_0 - \log l_2(p^*)| &\leq \gamma \log(1 + \epsilon) \end{aligned}$$

It then follows from Lemma 3 that

$$\left| \log \frac{\Phi(\frac{D-x_0}{\sqrt{y_0}})}{f(p^*)} \right| \leq 2 \times 3.284 \times \left(3 + \frac{D}{\sqrt{S}} \right) \gamma \log(1 + \epsilon)$$

Since the algorithm returns some other path p' such that $f(p') > \Phi(-3) = \Phi(\frac{D-x_0}{\sqrt{y_0}})$, we must have

$$f(p^*) < (1 + \epsilon)^{6.568(3 + \frac{D}{\sqrt{S}}) \gamma} f(p')$$

However, the approximation ratio is upper-bounded by $0.5/\Phi(-3) \approx 384.62$, so this shows that p' is an $\min\{384.62, (1 + \epsilon)^{6.568(3 + \frac{D}{\sqrt{S}}) \gamma}\}$ -approximation to p^* . We can apply the same technique for the case that $f(p') > \Phi(-2)$ or any other values. \square