Path Planning Under Uncertainty: Complexity and Algorithms (ICAPS Doctoral Consortium: Thesis Abstract)

Evdokia Nikolova
MIT Computer Science & AI Lab
Cambridge MA 02139, USA
enikolova@csail.mit.edu

Introduction

Finding shortest paths is a classic and fundamental problem in theoretical computer science which has influenced a wide array of other fields. Finding stochastic shortest paths has also been extensively studied though it has proven harder to formalize and yield classic results. What is the right way to define stochastic shortest paths, when we only know random distributions for the edge weights? Is it shortest paths on average, or shortest paths minimizing a combination of mean and variance, or minimizing some other specified criterion? Are they found adaptively or non-adaptively? A variety of problem variants have appeared in the literature, and most have ended up minimizing the expected length of paths, or a combination of expected lengths and expected costs such as bicriterion problems (Pallottino & Scutella 1997). 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 nonlinear function of the path length. Some notable work includes that of Loui (Loui 1983) who defines a decision-theoretic framework, where the optimal path maximizes the expected utility of the user for a class of monotonically increasing utility functions. Fan *et al.* (Fan, Kalaba & Moore) 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 the hardness of the problem arises and accumulates from many levels: combinatorial, distributional, analytic, functional, to list a few. We elaborate on these sources below.

We focus on stochastic shortest paths models which can effectively factor the sources of difficulty and whose solution draws from a variety of areas underlying the problem. In addition our solutions contain techniques that may be useful in solving other combinatorial problems and more generally, a number of nonconvex optimization problems.

Related Work

A lot of the related work on shortest paths in stochastic networks has focused on the notion of shortest paths in expectation, e.g., (Bertsekas & Tsitsiklis 1991). Other models have added costs on the edges in addition to travel times (Chabini 2002), (Miller-Hooks & Mahmassani 2000) where the costs depend on the realized travel times and in this way can capture a measure of uncertainty.

Finding the path of smallest expected length trivially reduces to a deterministic shortest path problems and does not take into account risk in predicting the optimal route. Since most real world applications care about a tradeoff between risk and expectation, we consider nonlinear objectives that capture more information about the edge distributions. Closest to this model, Loui (Loui 1983) considered a decision analytic framework for optimal paths under uncertainty, however he only studied monotone increasing cost functions and his algorithm has running time $O(n^n)$ in the worst case. Mirchandani and Soroush (Mirchandani & Soroush 1985) extended his work to a quadratic cost function of the path length, however their algorithm is also an exhaustive search over all potentially optimal paths, and thus exponential in the worst case.

Another branch of the stochastic shortest path literature has focused on adaptive algorithms (Fan, Kalaba & Moore), (Gao & Chabini 2002), (Boyan & Mitzenmacher 2001), which compute the optimal next edge in light of lengths or travel times already realized en route to the current node. Another direction has been to give approximations and heuristics for expected shortest paths in stochastic networks with nonstationary (time-varying) edge length distributions (Miller-Hooks & Mahmassani 2000), (Fu & Rilett 1998), (Hall 1986), to list a few. In this proposal, we only consider stationary edge length distributions, that do not change with time; time-varying distributions will be the subject of future work.

Problem Statement

The offline stochastic shortest path problem takes as input a graph and independent probability distributions for all its edge weights. It asks for the optimal path between a given source and a destination, which minimizes the expectation of a specified objective function. The term offline is used to emphasize that we seek a nonadaptive algorithm for an optimal path, before we observe any of the realized edge weights. When the cost function is linear, the problem becomes equivalent to a deterministic shortest path with edge weights equal to the expectations of their corresponding random variables. Thus, the challenge is when the objective is non-linear, which is also the case that most often occurs in practical applications.

Notation. We denote the graph $G = (V, E)$, with $|V| = n$ and $|E| = m$. Let the source and destination be S and T respectively. Denote the random weight of edge e by X_e. The objective function is $C : \mathcal{R} \to \mathcal{R}$. Strictly speaking, $C(X)$ is a function of the random path length $X = \sum X_e$. Thus, our problem is to solve

$$\min_{\pi} \mathbf{E}[C(\sum_{e\in\pi} X_e)] \tag{1}$$

for paths π between the source and destination.

The meaning of a non-linear cost function of the path length is not as intuitive as the notion of penalty for being late. Thus, we provide an equivalent formulation of the objective function, by including the extra parameter t for the clock time relative to a deadline at time 0. The penalty for arriving at the destination at time t is $\tilde{C}(t)$ (t is negative for early arrivals and positive for late arrivals). The expected cost of a path is then $\int_0^\infty f(x)\tilde{C}(t + x)dx$ where $f(.)$ is the probability density of the length X of the path and t is the departure time. For a fixed departure time t, the cost of the path is $C_t(X) = \tilde{C}(t+X)$, simply a horizontal shift by t units, and the minimization of its expectation over the set of paths is equivalent to the problem (1). When it is clear from the context, the parameter dependence $C_t(X)$ will be suppressed and we will write $C(X)$. This richer framework allows us to solve an additional problem: what is the optimal path and the optimal departure time t? This question is well defined for non-monotone cost functions with a global minimum.

We sometimes distinguish the cost functions by calling $\tilde{C}(t)$ the penalty function (since it explicitly specifies a penalty for being late), and $\mathbf{E}[C(X)]$ the objective function.

Note that we may not expect to solve the problem (1) in full generality for several reasons.

- Combinatorial difficulty. Even in the absence of randomness, when the edge weights are fully deterministic, a wide class of cost functions reduce to finding the longest path in the graph, which is NP-hard and inapproximable within a reasonable factor (Karger, Motwani & Ramkumar 1997).
- Distributional difficulty. The distributional assumptions on the edge lengths may bring a difficulty on their own, to the extent that we cannot even compute the distribution of the total length of a path $X = \sum X_e$, let alone evaluate the function $\mathbf{E}[C(X)]$ and minimize it. For example, Kleinberg *et al.* show that computing the distribution of the sum of n non-identical Bernoulli random variables is #P-hard (Kleinberg, Rabani & Tardos).
- Analytic difficulty. Even with additive edge length distributions such as the Normal distribution, with which we can readily compute the sum $X = \sum X_e$, we might not be able to get a closed analytic form of the objective function $\mathbf{E}[C(X)] = \int f(x)C(x)dx$ and thus cannot optimize it efficiently. This is a common problem in decision theory and related fields, which therefore focus attention on conjugate pairs of function and distribution families (more precisely, conjugate priors), *i.e.,* function-distribution pairs for which the integral can be computed in a closed form and the Expected Cost function lies in the same family as the original Cost function $C(X)$. For example, standard conjugate pairs are (Beta, Binomial) and (Gamma, Exponential).
- Functional difficulty. Having computed the distribution of the path length X and a closed form expression for the objective function $\mathbf{E}[C(X)]$, we are left with an integer optimization problem, to minimize a function over the collection of ST-paths of graph G. Relaxing the integer constraint, we have to optimize the function $\mathbf{E}[C(X)]$ over the path polytope in $\mathcal{R}^m$. The path polytope likely does not have any nice description with fewer than exponentially many linear constraints. Thanks to the separability of a linear objective into the graph edges, the deterministic shortest path problem has an efficient combinatorial solution. However, other than the linear and exponential objectives, no other cost function is separable into the edges (Loui 1983) and thus we might not hope to find exact optimal solutions in the general case. In special cases, convex and quasi-convex objective functions may admit greedy approaches that are equivalent to gradient descent on the path polytope, or they may admit efficient enumeration of a small set of candidate paths, which would contain the optimum. Non-convex functions in the relaxed problem may achieve an optimum anywhere in the path polytope, and as there are no general efficient methods for non-convex programming, it might not be tractable to find the relaxed optimum, nor approximate the integer optimum.

In addition to looking for efficient and approximation algorithms, we would like to understand the degree of difficulty each factor above contributes with.

For the case of general objective, we prove hardness and inapproximability results for objectives with a global minimum. We then describe approximations based on a combination of problem substructure and discretization. This method applies to non-separable objectives which have a separable term, and the solution idea is similar to partial minimization of a multivariate function.

We also study several specific but fundamental cost functions together with several different distributions, and offer hardness results, exact and approximation algorithms based on a variety of techniques.

References

Bertsekas, D. P. and Tsitsiklis, J. N. 1991. An analysis of stochastic shortest path problems. *Math. Oper. Res.* **16** (3): 580-595.

Boyan, J. and Mitzenmacher, M. 2001. Improved Results for Route Planning in Stochastic Transportation Networks. *ACM-SIAM Symposium on Discrete Algorithms*.

Bruno, J.L.; Downey, P.J.; and Fredrickson, G.N. 1981. Sequencing tasks with exponential service time to minimize the expected flowtime or makespan. *Journal of the ACM*, **28**: 100-113.

Chabini, I. 2002. Algorithms for k-Shortest Paths and Other Routing Problems in Time-Dependent Networks. *Transportation Research* Part B: Methodological.

Chvatal, V. 1980. Hard Knapsack Problems. *Operations Research*, **28**: 1402-1411.

Fan, Y.; Kalaba, R.; and Moore, II, J. E. Arriving on Time. *Journal of Optimization Theory and Applications*, forthcoming.

Fu, L. and Rilett, L. R. 1998. Expected Shortest Paths in Dynamic and Stochastic Traffic Networks. *Transportation Research*, Part B, **32**: 499-516.

Gao, S. and Chabini, I. 2002. Optimal Routing Policy Problems in Stochastic Time-Dependent Networks. *Proceedings of the IEEE 5th International Conference on Intelligent, Transportation Systems*, Singapore, pp. 549-559.

Hall, R. W. 1986. The Fastest Path Through a Network with Random Time-Dependent Travel Times, *Transportation Science*, **20**: 182-188.

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

Kleinberg, J.; Rabani. Y.; and Tardos, É. 2000. Allocating Bandwidth for Bursty Connections *SIAM Journal on Computing*, **30**: 191-21.

Loui, R. P. 1983. Optimal Paths in Graphs with Stochastic or Multidimentional Weights. *Communications of the ACM*, **26**: 670-676.

Miller-Hooks, E. D. and Mahmassani, H. S. 2000. Least Expected Time Paths in Stochastic, Time-Varying Transportation Networks. *Transportation Science*, **34**: 198-215.

Mirchandani, P. and Soroush, H. 1985. Optimal paths in probabilistic networks: a case with temporary preferences. *Computers and Operations Research,* **12** (4): 365-381.

Nikolova, E. 2005. Stochastic Shortest Paths. Manuscript.

Pallottino, S. and Scutella M. G. 1997. Shortest Path Algorithms in Transportation Models: Classical and Innovative Aspects. TR-97-06. Universita di Pisa Dipartimento di Informatica, Pisa, Italy.

Papadimitriou, C.H. and Yannakakis, M. 1991. Shortest paths without a map. *Theoretical Computer Science* **84**: 127-150.

Papadimitriou, C. H. and Steiglitz, K. 1998. *Combinatorial optimization: Algorithms and complexity.* Dover Publications, Inc., N.Y.

Vazirani, V. V. 2001. *Approximation Algorithms.* Springer-Verlag Berlin Heidelberg.