Robust Adaptive Routing under Uncertainty

Arthur Flajolet
Operations Research Center,
Massachusetts Institute of Technology
Email: flajolet@mit.edu

Patrick Jaillet
Department of Electrical Engineering and Computer Science,
Laboratory for Information and Decision Systems and Operations Research Center,
Massachusetts Institute of Technology

Sebastien Blandin
IBM Research Collaboratory, Singapore

1 Problem Description

We study stochastic shortest path problems with a deadline imposed at the destination. The objective is to find a history-dependent routing strategy, as opposed to an a priori path, minimizing a risk function of the lateness under distributional ambiguity. In the setting we consider, arc travel times are revealed upon crossing.

We restrict our attention to risk functions satisfying natural properties meant to prevent infinite cycling, e.g. maximizing the expected delay is not allowed. To capture distributional ambiguity, we assume that the uncertain, possibly time-dependent, arc travel times are only known through confidence intervals on some statistics, e.g. mean, absolute mean deviation, and support, which together define a set of distributions consistent with the available knowledge.

Our model has applications in transportation networks, e.g. in car navigation systems or in delivery route planning, where mandatory stops at intermediate locations could be further incorporated.

2 Models and Solutions

Let $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ be a finite directed graph where each arc $(i, j) \in \mathcal{A}$ is assigned a collection of random travel times $(c_{ij}^{\tau})_{\tau \in \mathbb{R}}$. Crossing arc (i, j) at time τ takes time c_{ij}^{τ}. To simplify the exposition, we assume that the travel times are independent across time and arcs, but the results can be extended to the case of Markov chains of finite order, where the travel time of an arc depends on the past $m \in \mathbb{N}$ experienced travel times. The travel-time distribution of arc (i,j) at time τ is denoted by p_{ij}^{τ}. We consider a user traveling through $\mathcal{G}$ leaving from s at time $\tau = 0$ and willing to reach d by time $\tau = T$. Π is the set of all history-dependent routing strategies, i.e. mappings from the past locations and realizations of the travel times to a neighboring node. For $\pi \in \Pi$, X_π is the random time it takes to reach d when driving in accordance with π. $f(\cdot)$ denotes the risk function of the lateness. Typical examples include $f(t) = 1_{t>0}$ and $f(t) = t \cdot 1_{t\geq 0}$ which correspond to maximizing the probability of on-time arrival and minimizing the expected delay respectively. We impose mild asymptotic restrictions on $f(\cdot)$ to prevent infinite cycling, e.g. $f(\cdot)$ has to be asymptotically non-decreasing.

2.1 Known travel-time distributions

If the travel-time distributions are given, the problem is formulated as:

$$\min_{\pi\in\Pi} \mathbb{E}_{\prod_{\substack{(i,j)\in\mathcal{A}\\ \tau\geq 0}} p_{ij}^{\tau}} [f(X_\pi - T)].$$

Assuming the distributions have compact supports, we prove that, when the driver is already very late, the optimal strategy is to follow the least expected time path. Building on this result, we develop a methodology to compute an approximate optimal strategy using dynamic programming. When $f(\cdot)$ is Lipschitz on compact sets, we show that we obtain an ϵ-approximate optimal strategy in $O(\frac{1}{\epsilon^2})$ running time. The complexity drops to $O(\frac{1}{\epsilon} \cdot \log(\frac{1}{\epsilon}))$ leveraging fast Fourier transforms if the travel-time distributions are time-independent.

2.2 Ambiguity sets

The travel-time distribution of arc (i,j) at τ is only known to lie in an ambiguity set $\mathcal{P}_{ij}^{\tau}$. The routing problem under ambiguity aversion is formulated as:

$$\min_{\pi\in\Pi} \sup_{\forall\tau\geq 0, \forall(i,j)\in\mathcal{A},\, p_{ij}^{\tau}\in\mathcal{P}_{ij}^{\tau}} \mathbb{E}_{\prod_{\substack{(i,j)\in\mathcal{A}\\ \tau\geq 0}} p_{ij}^{\tau}} [f(X_\pi - T)].$$

We consider ambiguity sets defined by confidence intervals on piecewise linear statistics:

$$\mathcal{P}_{ij}^{\tau} = \{p \in \mathcal{P}([\delta_{ij}^{\inf,\tau}, \delta_{ij}^{\sup,\tau}]) : \alpha_{ij}^{\tau,q} \leq \mathbb{E}_{X\sim p}[f_{ij}^{\tau,q}(X)] \leq \beta_{ij}^{\tau,q},\ q = 1, \cdots, Q_{ij}^{\tau}\},$$

where $\mathcal{P}([\delta_{ij}^{\inf,\tau}, \delta_{ij}^{\sup,\tau}])$ is the set of probability measures on $[\delta_{ij}^{\inf,\tau}, \delta_{ij}^{\sup,\tau}]$ and the functions $(f_{ij}^{\tau,q}(\cdot))_{q=1,\cdots,Q_{ij}^{\tau}}$ are piecewise linear. This allows us to capture information on the mean and the absolute deviation but also on more general moments since a continuous function can be approximated arbitrarily well by a piecewise linear function. We prove that, when the driver is already very late, the optimal strategy is to follow the shortest path with respect to the worst-case expected travel times. Building on this, we develop a methodology to compute an approximate optimal strategy using dynamic programming, the ellipsoid algorithm and a tailored dynamic convex hull algorithm. When $f(\cdot)$ is Lipschitz on compact sets, we show that we obtain an ϵ-approximate optimal strategy in $O(\frac{1}{\epsilon} \cdot \log(\frac{1}{\epsilon}))$ runtime. In practice, we substitute the ellipsoid algorithm with the simplex algorithm leveraging column generation and warm starting.

2.3 Computational results

We investigate the benefits of the distributionally robust approach when traffic measurements are scarce. The data consists of a 15-day recording of GPS probe vehicle speed samples coming from a combined fleet of over 15,000 taxis in Singapore. We consider that the samples available in our dataset represent the real traffic conditions characterized by the empirical distributions. In our experimental setting, these distributions are not known, instead only a fraction of the samples is available. Based on this limited data, we have to provide an itinerary. We consider three methods: the robust approach with ambiguity sets defined by confidence intervals on the mean and the variance, the non-robust approach with empirical distributions and the least expected time path. We randomly pick the samples in the dataset and repeat the experiment 100 times. We look at the average, over the 100 experiments, probability of arriving on-time as a function of the deadline.

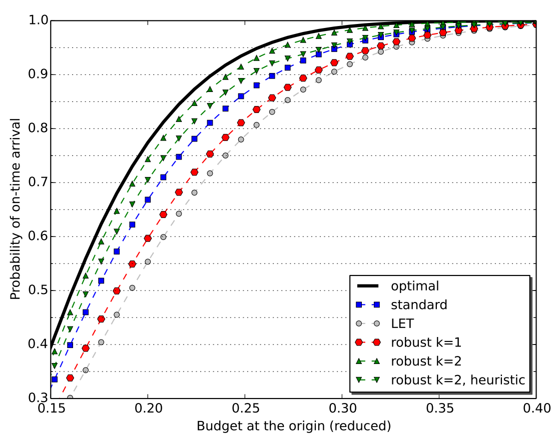

Figure 1: Probability of on-time arrival as a function of the travel-time budget. Average number of samples per arc: ~ 25.

Figure 1 shows that the robust formulation based on confidence intervals can significantly outperform non-robust approaches.

3 Related Work

Stochastic shortest path problems with deadline have gained interest in the transportation research community due to the uncertainty arising in many practical settings as a result of traffic incidents and variability of the travel demand. However, authors have either focused on specific objective functions or made strong assumptions regarding travel-time distributions. Jaillet et al. (2015) develop a polynomial-time algorithm to compute an a priori path minimizing a risk measure, coined as lateness index, that trades off probability of on-time arrival, expected delay and value at risk. The travel-time distributions are assumed to be either given or known to lie in ambiguity sets defined by confidence intervals on moments. Bertsimas and Sim (2003) embrace the robust optimization paradigm and propose a polynomial-time algorithm to find a path minimizing the worst-case total travel time given well-chosen uncertainty sets. Nikolova et al. (2007) tackles the problem of finding a path maximizing the probability of arriving on time with the normal distribution assumption. Our work generalizes the adaptive approach of Samaranayake et al. (2012) formulated for the probability of on-time arrival when the distributions are given.

Our problem can also be cast as a distributionally robust Markov decision process with continuous state space. Nilim and El Ghaoui (2007) investigate the case of a finite state space for ambiguity sets based on a likelihood or entropy constraint. They show that the computational complexity is almost the same as for the non-robust version.

References

D. Bertsimas, and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, volume 98, pages 49-71. Springer, 2003.

P. Jaillet, J. Qi, and M. Sim. Routing Optimization under Uncertainty. Forthcoming in *Operations Research*, 2015.

E. Nikolova, J.A. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. *Proc. 14th Annual Eur. Sympos. Algorithms*, volume 14, pages 552–563. Springer, 2006.

A. Nilim and L. El Ghaoui. Robust Control of Markov Decision Processes with Uncertain Transition Matrices. *Operations Research*, volume 53, pages 780–798. INFORMS, 2005.

S. Samaranayake, S. Blandin, and A. Bayen. A tractable class of algorithms for reliable routing in stochastic networks. *Transportation Science Part C*, volume 20, pages 199–217. INFORMS, 2012.