Routing under Uncertainty

van Ee, M.

2017

Summary

In this thesis, we study the approximability and complexity of routing problems under uncertainty. In routing problems, we have to walk through a graph while optimizing some objective. For example, in the *traveling salesman problem* (TSP) we have to visit all vertices in the graph using a tour of minimum length. Another example is the *traveling repairman problem* (TRP), where we want to minimize the total latency. In this thesis, we consider two types of routing problems under uncertainty: *a priori routing* and *graph search*.

In a priori routing only a subset of the vertices of the graph needs to be visited but this subset is unknown. However, we need to make just one tour on all vertices called the first-stage tour. After that, the set of vertices to be visited, the active set, is revealed. The second-stage tour can now be derived by shortcutting the first-stage tour to the active set. If we are given a probability distribution over the active sets, we can compute the expected cost of the second-stage tour. Now, our goal is to construct a tour minimizing this expectation. We consider two objectives, namely the length (as in TSP) and the total latency (as in TRP) of a tour. We also consider two models for the probability distribution. In the independent decision model, each vertex is active with its own probability independent of other vertices. In the scenario model, we are given a list of possible active sets, called scenarios, with a probability distribution over the scenarios.

The *graph search problem* is defined as follows. There is a target hidden at a vertex in the graph and one has to find it. We are given a probability distribution and our goal is to find a walk that minimizes the expected length until finding the target. When all probabilities are equal, the graph search problem reduces to the TRP. We consider generalizations to multiple targets, failing edges and non-discrete probability distributions.

In Chapter 2 we study the *a priori traveling repairman problem* in the independent decision model. Using a series of reductions to other routing problems, we were able to give a constant-factor approximation for the uniform case, i.e., when all probabilities are equal. Then, in Chapter 3, we discuss the *a priori traveling salesman problem* in the scenario model. We show that there is a constant-factor approximation when the scenarios are small, big or nested or when the number of scenarios is bounded. By linking the problem to the so-called permutation constraint satisfaction problems, we show strong inapproximability results. We also look into approximation algorithms using the master tour lower bound, i.e., the weighted sum of optimal values for the scenarios. However, we show that this lower bound will not help us to improve on prior work. In Chapter 4, we show that determining if a given graph and a set of scenarios has a master tour, a tour such that each second-stage tour has the same length as the optimal tour on just that scenario, is Δ_2^p-complete. We also show this for the master

versions of satisfiability and Steiner tree. The main open question in a priori routing is whether we can improve upon the current best $O(\log n)$-approximation for a priori TSP in the scenario model.

The *Canadian traveler problem* is investigated in Chapter 5. In this problem, we have to walk from s to t in a graph. However, only edges in the active set (again using a probability distribution) are present. Moreover, we only know whether an edge is present after visiting one of its endpoints. The goal is to construct a policy that minimizes the expected length of the walk. We study both the independent decision and the scenario model. For the independent decision model, we show NP-hardness for series-parallel graphs, solving an open question posed in the literature. Moreover, we investigate what it means for an algorithm to be non-adaptive and show results on the implications of being non-adaptive through analyzing the adaptivity gap. For the scenario model, we show NP-completeness on disjoint-path graphs and cactus graphs. Although this chapter is named after the Canadian traveler problem, it mainly concerns a special case called the *multi-target graph search problem*. For this problem, in the independent decision model, we show NP-completeness on trees and give a 14.4-approximation for general metrics. In the scenario model, the problem is already NP-complete on stars.

Chapter 6 deals with the *lost cow problem*. This problem can be seen as a non-discrete generalization of the graph search problem on the line. For symmetric probability distributions, we characterize whether it is optimal not to turn, i.e., to walk to the end of the line on one side and then to the other end. We generalize this result to more than two directions and also consider non-symmetric distribution functions. In Chapter 7, we introduce the graph parameter *starwidth*. It naturally measures to what extend a graph looks like a star graph in the sense that it is the star-equivalent of treewidth. Our main goal in introducing this parameter was to extend the polynomial solvability of TRP on stars to graphs that look like stars. Unfortunately, TRP turns out to be NP-complete on trees of starwidth 4. Nonetheless, this parameter might be interesting in another setting. In this chapter, we consider characterizations of graphs with small starwidth, the complexity of computing the starwidth of a given graph and the relation between other graph parameters. Finally, we show three problems that are NP-complete on trees and remain NP-complete on trees with small starwidth.