The complexity of some set-partitioning formulations for the capacitated vehicle routing problem with stochastic demands

by

Joshua Gunter

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics & Optimization

Waterloo, Ontario, Canada, 2020

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The capacitated vehicle routing problem with stochastic demands (CVRPSD) is a variant of the deterministic capacitated vehicle routing problem (CVRP). The CVRP consists of planning routes for vehicles with a given capacity to deliver goods to a set of customers with known demands, with the goal being to find the cheapest such set of routes. In the CVRPSD, rather than being deterministic, customer demands are random variables from a given probability distribution. This creates the possibility of route failures, where the realized demand on a route is greater than the vehicle capacity. In this event, a recourse action following a pre-determined strategy must be taken. The goal is then to minimize the expected cost of the routes, i.e. the sum of the deterministic route lengths and expected additional costs incurred by recourse actions.

A common approach when solving the CVRPSD is to formulate it as a 2-stage stochastic programming problem. In this framework, the first stage is planning a set of routes, while the second stage is computing the expected cost of route failures. While edge-based formulations were the dominant approach originally, the success of set-partitioning formulations for related problems such as the CVRP and the vehicle routing problem with time windows led to research into developing similar formulations for the CVRPSD. These formulations contain an exponential number of variables, necessitating the use of column generation to solve them. In column generation, an additional optimization problem known as the pricing problem needs to be repeatedly solved when solving the current LP relaxation through the branch-and-bound tree.

The pricing problem for the deterministic CVRP is strongly NP-hard, and so to obtain tractable algorithms a relaxed version of the pricing problem which can be solved in pseudo-polynomial time is typically used. Similar methods for relaxing the pricing problem have been explored for the CVRPSD, but these make use of some simplifying assumptions for computing the expected cost of route failures, which needs to be done frequently when solving the pricing problem. In this thesis, we show that using these assumptions results in an “approximate” pricing rather than “exact” pricing, and present results on the hardness of performing exact pricing for set-partitioning formulations for the CVRPSD. Specifically, we show that when customer demands are given by a finite set of demand scenarios, exactly solving the pricing problem for the CVRPSD is strongly NP-hard. Additionally, we show that when customer demands are independent normal, under some assumptions there is a reduction from the Hamiltonian cycle problem to the pricing problem for the CVRPSD. This does not constitute a proof of strong NP-hardness due to the aforementioned assumptions required, but does suggest that even in this case the pricing problem may be harder than currently thought.

Acknowledgements

I would like to thank my supervisor Ricardo for all of his help while writing, and the many suggestions and corrections he contributed.

Additionally I would like to thank my readers for their helpful comments and suggested improvements.

Table of Contents

List of Figures

Chapter 1

Introduction

The capacitated vehicle routing problem (CVRP) is a classic problem in combinatorial optimization, first studied by Dantzig and Ramser in 1959 [11]. It is a generalization of the travelling salesman problem, which is one of the most well-studied discrete optimization problems. The CVRP consists of routing a fleet of vehicles, all starting and ending at a central depot, to service a set of customers, with the goal being to find a set of routes which minimizes the total distance travelled by all of the vehicles. The motivation for this problem is in delivering goods to customers, and as such it is also often assumed that each vehicle in the fleet has a given capacity for how many goods it can carry, with each customer having a demand for goods. This results in an additional constraint on the vehicles' routes: the sum of customer demands along the route must be at most the vehicle capacity, as otherwise some customers will not have their demands met. A comprehensive overview of the CVRP can be found in the book *The Vehicle Routing Problem* by Toth and Vigo [38], as well as the survey by Cordeau, Laporte, Savelsbergh and Vigo [10].

Many real-world applications have more restrictions or different goals than those described in the CVRP. Different variants of the problem accounting for these additional aspects have been developed, such as the vehicle routing problem with time windows [9, 7, 1], where customers have time windows in which they must be serviced, and the vehicle routing problem with pickup and delivery, where customers may have both demands for goods to pickup and goods to deliver rather than one or the other. Both of these are covered by Toth and Vigo [38]. Another major class of variants are *stochastic* vehicle routing problems [10]. These are vehicle routing problems where part of the problem data is stochastic rather than being deterministic. This means that the data is described by a random variable from a given probability distribution instead of having a specific value known ahead of time. Examples of this include stochastic travel times along edges, modelling uncertain travel times of vehicles due to traffic or other physical issues [1], or stochastic customer demands, modelling cases where the demands of customers are not known before routes must be planned [6, 22, 24, 8, 21]. In this thesis, we will focus on the case where the customer demands are stochastic. This problem is known as the capacitated vehicle routing problems with stochastic demands, or CVRPSD. These types of mathematical programming problems involving random variables are known as *stochastic programming problems*. Closely related to these are *robust optimization* problems, where rather than being specified by a probability distribution some of the data lies within a given uncertainty set $\mathcal{U}$. In this paradigm a solution is feasible if and only if it satisfies the constraints of the problem for all given realizations of data in $\mathcal{U}$, with the goal being to optimize the "worst-case" scenario [3].

The introduction of stochastic demands creates new problems when looking to solve a vehicle routing problem. Depending on the distribution of demands, it may no longer be possible to create routes where the cumulative realized demand is less than the vehicle capacity, introducing the possibility of *route failures*. In the case of a route failure, a *recourse action* must be taken. The strategy used to handle route failures and perform recourse actions is known as a *recourse strategy*. The simplest recourse strategy, first introduced by Dror and Trudeau [14], is to make a round trip to the depot whenever a route failure occurs, and continue the route as planned from the customer at which the route failed. Due to its computational simplicity, this strategy is commonly used in the literature, such as in Gendreau, Laporte, and Séguin [22], Hjorring and Holt [24], Christiansen and Lysgaard [8], and Gauvin, Desaulniers and Gendreau [21]. Another available recourse strategy is known as *optimal restocking*. As each customer is visited, this strategy takes into consideration the current remaining vehicle capacity and expected future demand on the route, and computes whether to continue the route or make a preemptive return trip to the depot to empty the vehicle capacity. Yee and Golden [41] were the first to develop this method and provided a dynamic programming algorithm to solve it. Heuristic solutions for formulations with an optimal restocking policy can be found in the papers by Yang, Mathur, and Ballou [40] and Bianchi et al. [16], while more recently Salavanti-Khoshghalb, Gendreau, Jabali and Rei [36] exactly solved such a formulation.

Within the literature, different simplifying assumptions on the distribution of customer demands have been made when developing algorithms. Some common assumptions are that customer demands are independent and that customer demands are drawn from a well-known distribution such as the normal or Poisson distributions [22, 24, 8, 21]. While often necessary for computationally tractable algorithms, these types of assumptions are undesirable, as they may not be suitable in many practical applications. For example, the same factors causing increased demand at one customer may also affect the demands of other customers in the same area, and the true distribution of customer demands may be unknown. For this reason, it is preferable to also be able to solve instances where the distribution of customer demands is specified by a finite set of demand realization scenarios, which may be generated through historical data on customer demands. Using a large number of these scenarios allows one to approximate the true distribution of data with the empirical distribution given by the samples, a technique known as sample average approximation [4, 39].

In addition to the choices of how to handle recourse actions and model the randomness of customer demands, when solving stochastic vehicle routing problems there is also a decision to make regarding whether to explicitly model the cost of route failures. For the CVRPSD, there are two main branches of research: modelling the problem as a 2-stage stochastic program, or modelling the problem as a chance-constrained stochastic program. To distinguish between the two, we will use the term CVRPSD to refer only to models as a 2-stage stochastic program, and use the term CCVRP to refer to chance-constrained models. In 2-stage stochastic programs, the expected cost of recourse of a solution x is explicitly modelled and a solution is found which minimizes the sum of the route lengths and expected recourse cost. In chance-constrained models, the cost of recourse is not explicitly modelled. Instead there is a constraint on the probability of a route failure occurring, with routes failing with probability less than $1 - \epsilon$ for some small parameter ϵ. The effect that failures may have on the expected costs of routes is ignored.

Stewart and Golden [37] solved a chance-constrained model via a reduction to a deterministic VRP under the assumptions of independent demands, and that the cumulative demand on routes has the same distribution as the individual customer demands. They note this is satisfied by many common distributions, such as normal, Poisson, binomial, and chi-squared. They additionally provided heuristics for solving a chance-constrained model with correlated demands. Laporte, Louveaux and Mercure [28] were the first to exactly solve a chance-constrained model using a branch-and-cut algorithm, but required demands to be both independent and normal. More recently, Dinh, Fukasawa and Luedtke [12] exactly solved a chance-constrained model for both joint normal demands and demands given by scenarios with a branch-cut-and-price algorithm.

More work has been devoted to formulations of the CVRPSD as a 2-stage stochastic program with recourse. The first to study this type of formulation was Bertsimas in his Ph.D. thesis [5], which covered stochastic versions of the travelling salesman problem, minimum spanning tree, and vehicle routing problem. For the vehicle routing problem specifically, he provided closed-form expressions and methods to compute the expected cost of a route and heuristics for solving the CVRPSD, as well as asymptotically optimal algorithms. These results were improved further in [6], where he derived upper and lower bounds on the optimal value of the CVRPSD. Gendreau, Laporte and Séguin [22] were the first to successfully solve the CVRPSD as a 2-stage stochastic program exactly. They achieved this by way of the Integer L-shaped algorithm, a branch-and-cut algorithm first developed by Laporte and Louveaux [27] as a general framework for solving stochastic programs with complete recourse. This algorithm models the expected cost of recourse with an additional variable θ, and the sum of the deterministic route costs and θ is minimized. Hjorring and Holt [24] found new optimality cuts for this method for the case of a single vehicle and independent demands. In addition to these, the recent work by Salavanti-Khoshghalb, Gendreau, Jabali and Rei [36] which solves the CVRPSD with optimal recourse also uses the L-shaped algorithm.

While approaches based on the L-shaped algorithm have been widely used for exactly solving the CVRPSD, there are other promising methods, specifically set-partitioning formulations solved via branch-and-price algorithms. In other vehicle routing problems, such as the deterministic CVRP and the vehicle routing problem with time windows (VRPTW), branch-and-price or branch-and-price-and-cut algorithms using set-partitioning formulations are the current state-of-the-art techniques [9, 20]. In these types of formulations, there is a decision variable for each possible route that may be included in a solution, leading to an exponential number of variables in the problem. To solve these, a restricted subset of the variables is initially optimized over, with additional variables being dynamically added over time. This technique is known as *column generation*, and the problem of choosing which variable to add next is known as the *pricing problem*. The pricing problem for these branch-and-price algorithms are NP-hard, but due to pseudo-polynomial time algorithms for pricing they can still be solved efficiently. These algorithms are especially effective in the case of restrictive capacity constraints, a class of problems with which the L-shaped algorithm struggles [8]. Set-partitioning formulations additionally have applications in the design of approximation algorithms for vehicle routing problems, for example in the works by Post and Swamy [33] and Friggstad and Swamy [19], where they are known as configuration LPs.

The efficacy of branch-and-price algorithms in those cases motivated study into whether they could similarly be applied to the CVRPSD, with Christiansen and Lysgaard [8] being the first to develop a branch-and-price algorithm for the CVRPSD using a set partitioning formulation. While this method was not able to solve problems containing as many vertices as the L-shaped method, their results showed it was more effective when problem instances had a higher number of vehicles. This method was further improved by Gauvin, Desaulniers and Gendreau [21]. However, these algorithms rely on the assumption that customer demands are independent.

A core component of branch-and-price algorithms is the efficient solution of the column generation subproblem. We will refer to this subproblem for the CVRPSD as the minimum cost 2-stage q-route problem. This consists of finding the q-route, or walk starting and ending at the depot, with minimum expected cost that satisfies capacity constraints. The primary contributions in this thesis are two proofs on the hardness of the minimum cost 2-stage q-route problem. First we prove that solving this problem in the case of distributions specified by a finite set of scenarios is strongly NP-hard. This indicates that this problem is very hard in general, and cannot be simply extended to instances with dependent customers demands. Secondly we prove that under some assumptions, namely $O(1)$ elementary operations and computability of the expected number of route failures in polynomial time, there is a reduction from the Hamiltonian cycle problem to the minimum cost 2-stage q-route problem with independent normal demands, even when the means and variances are polynomially bounded in n. This suggests that exactly pricing the 2-stage cost of q-routes may not be feasible in pseudopolynomial time. However this does not constitute a proof of strong NP-hardness, due to the required assumptions and large edge costs in the construction. Additionally, this result is not in contradiction with the results of Christiansen and Lysgaard [8], as they use an alternate pricing method which makes some simplifying assumptions about the probability of route failures.

This thesis is organized as follows. Chapter 2 describes how set-partitioning formulations are applied to the CVRP, the necessity of column generation for this formulation and the structure of the column generation subproblem, and how this subproblem may be solved in pseudopolynomial time via dynamic programming. In Chapter 3, we cover a problem definition and formulation for the CVRPSD which generalizes the one proposed by Christiansen and Lysgaard [8], allowing for probability distributions which may be dependent or specified by a finite set of scenarios. We then examine the differences between the pricing problem for the CVRP and CVRPSD, and take a closer look at the difficulties introduced by stochastic demands. In addition, we contrast our approach to pricing with that taken by Christiansen and Lysgaard [8], and cover the details distinguishing the "approximate" pricing method they used from the "exact" pricing we wish to study. In Chapter 4, we present the results on the hardness of the pricing problem for the CVRPSD for the two cases of finite scenario demands and independent normal demands. Finally, in Chapter 5 we present concluding remarks.

Chapter 2

The deterministic capacitated vehicle routing problem

2.1 Problem Definition

The CVRP is defined on an undirected graph $G = (V, E)$ with vertices $V = \{0, 1, \ldots, n\}$, which may be assumed to be complete. Vertex 0 represents the depot, and $V_+ = \{1, ..., n\}$ represents the set of customers. For each edge $ij \in E$ there is an associated deterministic and metric travel cost ℓ_{ij}. Each customer $i \in V_+$ has an associated demand $d_i > 0$. There is a fleet of m identical vehicles, each with capacity Q. The goal is to find a set of m routes for the vehicles minimizing the total distance travelled, i.e. the sum of the travel costs ℓ_{ij} along each route, such that the total demand on each route does not exceed the vehicle capacity and each customer is visited exactly once. Intuitively, this is because there is no need to re-visit a customer after their demand has been satisfied, and as the edge costs are metric there is no way to find a shorter path between two vertices $u, v \in V$ by re-visiting a previously visited customer. We will define what exactly we mean by a "route" in the following section.

2.1.1 An initial formulation

In this section, we'll be looking at constructing a set-partitioning formulation for the CVRP. This type of formulation requires exponentially many variables, but can give much tighter lower bounds than more compact formulations and may still be solved efficiently through column generation [2]. We will first look at a naive formulation, which will turn out to be too hard to be of practical use, in order to help develop some intuition behind the construction of a standard formulation in the literature which is used in current state-of-the-art exact algorithms for the CVRP. Before this, we first need to define what a *route* is in the context of vehicle routing.

Definition 2.1.1 (Elementary route)**.** An *elementary route* is a sequence of vertices $(0, v_1, ..., v_k, 0)$ such that $k > 0$ and $v_i \neq v_j$ for all $i \neq j$, $1 \leq i, j \leq k$, $v_i \in V_+$ and $\sum_{i=1}^{k} d_i \leq Q$.

As we are primarily concerned with elementary routes, from here on we use the term "route" to refer only to elementary routes, unless otherwise specified. Let $\mathfrak{R}$ denote the set of routes. For a route $r = (v_0 = 0, v_1, v_2, ..., v_k, v_{k+1} = 0) \in \mathfrak{R}$ and vertex $i \in V_+$ let α_{ir} be a parameter defined by:

$$\alpha_{ir} = \begin{cases} 1 & \text{if } i \in r \\ 0 & \text{otherwise} \end{cases}$$

And let:

$$c_r = \sum_{i=0}^{k} \ell_{v_i, v_{i+1}}$$

denote the travel cost associated with r. For each $r \in \mathfrak{R}$, let:

$$x_r = \begin{cases} 1 & \text{if route } r \text{ is used by a vehicle} \\ 0 & \text{otherwise} \end{cases}$$

We can then formulate the CVRP via the following set-partitioning formulation P:

$$(P) \qquad \begin{aligned} \min \sum_{r \in \mathfrak{R}} c_r x_r & & \text{(2.1a)} \\ \text{s.t.} \sum_{r \in \mathfrak{R}} \alpha_{ir} x_r = 1 & \qquad \forall i \in V_+ & \text{(2.1b)} \\ \sum_{r \in \mathfrak{R}} x_r = m & & \text{(2.1c)} \\ x_r \in \{0, 1\} & \qquad \forall r \in \mathfrak{R} & \text{(2.1d)} \end{aligned}$$

The objective function (2.1a) minimizes the total distance travelled by the set of routes chosen. Constraints (2.1b) enforce that each vertex $i \in V_+$ is visited exactly once by the routes chosen. The constraint (2.1c) enforces that we use m vehicles. Finally, the binary constraints (2.1d) enforce that each route is either used or unused.

2.2 Solving the set-partitioning formulation

Difficulties arise when we try to solve the formulation P with a standard IP solving algorithm such as branch-and-bound. To use such an algorithm, we need to be able to solve the LP relaxation of P, denoted hereafter by P^{LP}. As there are exponentially many possible routes in the set $\mathfrak{R}$, there is an exponential number of variables in the problem, which poses a problem even when just solving P^{LP}. To pivot between basic feasible solutions in the simplex algorithm, we need to be able to enumerate all the variables, but this is not feasible to do for P^{LP} due to the cardinality of $\mathfrak{R}$.

Instead, a separate optimization problem known as the *pricing problem* is solved to compute the next variable which should enter the basis, without requiring the enumeration of all possibilities. This is a technique known as *column generation*, where columns of the constraint matrix of the LP are generated dynamically during the solving procedure instead of being available from the start. In this section, we will give a description of column generation algorithms using the LP relaxation of the CVRP as an example.

Since we cannot solve P^{LP} directly due to the aforementioned issues with the huge number of variables, we instead consider a problem with a restricted feasible region. This is known as the restricted master problem, which we denote P^{RM}. Using our example of the CVRP, P^{RM} has the constraints of P^{LP}, with the set of variables restricted to a small subset $\overline{\mathfrak{R}} \subset \mathfrak{R}$, with $|\overline{\mathfrak{R}}| \ll |\mathfrak{R}|$.

$$
\left(P^{RM}\right) \qquad
\begin{aligned}
\min \sum_{r \in \overline{\mathfrak{R}}} c_r x_r & & & \qquad (2.2a)\\
\text{s.t.} \sum_{r \in \overline{\mathfrak{R}}} \alpha_{ir} x_r = 1 & & \forall i \in V_+ & \qquad (2.2b)\\
\sum_{r \in \overline{\mathfrak{R}}} x_r = m & & & \qquad (2.2c)\\
0 \leq x_r & & \forall r \in \overline{\mathfrak{R}} & \qquad (2.2d)
\end{aligned}
$$

Note that an upper bound on the variables x is not needed, as constraint (2.2b) provides an implicit upper bound of 1 on x_r for all $r \in \overline{\mathfrak{R}}$.

Depending on the choice of $\overline{\mathfrak{R}}$, P^{RM} may not be feasible. To get around this problem, similarly to the two-phase simplex method we can add artificial variables s_i to the formulation to ensure we have a feasible solution to start with. These artificial variables are given a large weight M in the objective function such that $M \gg \max\{c_r : r \in \mathfrak{R}\}$, to ensure that they will be eliminated as further routes are added to $\overline{\mathfrak{R}}$. To simplify the following, we will assume that we have already identified a set $\overline{\mathfrak{R}}$ such that P^{RM} is feasible without the need for artificial variables.

Let x^* be the optimal solution to P^{RM}. Observe that x^* can be easily extended to a feasible solution to the original relaxation P^{LP} by setting $x_r^* = 0$ for all $r \in \mathfrak{R} \setminus \overline{\mathfrak{R}}$. The question now is whether x^* is optimal not only for P^{RM}, but also for P^{LP}. Answering this is achieved by solving the pricing problem.

2.2.1 The pricing problem for P^{RM}

Consider the dual D^{LP} for P^{LP}:

$$(D^{LP}) \qquad \max \sum_{i \in V_+} \pi_i + m\pi_0 \tag{2.3a}$$

$$\text{s.t.} \sum_{i \in V_+} \alpha_{ir}\pi_i + \pi_0 \leq c_r \qquad \forall r \in \mathfrak{R} \tag{2.3b}$$

The dual variables π_i, $i \in V_+$ correspond to constraints (2.1b), while the variable π_0 corresponds to the constraint (2.1c). Observe that we can rearrange constraint (2.3b) in the following way:

$$c_r - \sum_{i \in V_+} \alpha_{ir}\pi_i - \pi_0 \geq 0. \tag{2.4}$$

Let D^{RM} denote the dual to P^{RM}. This problem is exactly the same as D^{LP}, except there are constraints (2.3b) for all $r \in \overline{\mathfrak{R}}$ instead of $\mathfrak{R}$. Furthermore, let π^* be the optimal dual solution to D_{RM} corresponding to x^*. Then the *reduced cost* of a route r for the solution x^* is:

$$\bar{c}_r := c_r - \sum_{i \in V_+} \alpha_{ir}\pi_i^* - \pi_0^*. \tag{2.5}$$

Note that this is the left-hand side of (2.4) with π^* substituted in. This means that the reduced cost of a route can be interpreted as a dual feasibility check for the solution π^*. If $\bar{c}_r \geq 0$ for all $r \in \mathfrak{R}$, then π^* is feasible for D^{LP}, and since x^* is feasible for P^{LP} and $\sum_{r \in \mathfrak{R}} c_r x_r^* = \sum_{i \in V_+} \pi_i^* + m\pi_0^*$ by construction, x^* must be optimal for P^{LP} by strong duality.

We can now state the definition of the pricing problem in more formal terms. The pricing problem consists of identifying the route r^* with minimum reduced cost, or in other words solving the problem:

$$r^* := \operatorname{argmin}\{\bar{c}_r : r \in \mathfrak{R}\}. \tag{2.6}$$

If $\bar{c}_{r^*} \geq 0$, then clearly $\bar{c}_r \geq 0$ for all $r \in \mathfrak{R}$. Therefore by the argument above, the current optimal solution to P^{RM} is also optimal for the LP relaxation of P. On the other hand, if $\bar{c}_{r^*} < 0$, then the corresponding dual solution π^* is infeasible for D^{LP}. However, since π^* is feasible for D^{RM}, $\bar{c}_r \geq 0$ for all $r \in \overline{\mathfrak{R}}$ and so it must be case that $r^* \in \mathfrak{R} \setminus \overline{\mathfrak{R}}$. We then set $\overline{\mathfrak{R}} = \overline{\mathfrak{R}} \cup \{r^*\}$, adding the route r^* to our restricted route set, and re-solve P^{RM}. Since $\mathfrak{R}$ contains a finite number of elements, this algorithm necessarily terminates.

2.2.2 Formulating the pricing problem as an SPPRC

To solve the pricing problem for the CVRP, we will reformulate it as another combinatorial optimization problem, the shortest path problem with resource constraints (SPPRC).

The SPPRC is defined on a graph $G = (V, E)$. Each vertex $v \in V$ has an associated resource b_v. Despite the name of the problem suggesting otherwise, the objective is to find the shortest walk, rather than path, with respect to a set of edge weights $h \in \mathbb{R}^{|E|}$ from an origin vertex $s \in V$ to a destination $t \in V$ such that the resources picked up along the walk does not exceed a given budget β. This walk may repeat vertices and have edges with negative weights. The SPPRC is known to be NP-hard, but there are pseudo-polynomial time algorithms for solving it when the resources b_v and budget β are bounded by a polynomial in n [25].

For the pricing problem for the CVRP, the resources b_v at each vertex are the customer demands d_v, and the budget β is the vehicle capacity Q. The underlying graph is the same as the graph for the corresponding CVRP instance, with the exception of an additional vertex t_0 which is a copy of the depot. The depot vertex 0 and the copy of the depot t_0 are then used as the origin s and destination t respectively. The last item required is to determine what the edge weights should be. We want to choose edge weights such that the sum of the edge weights along a route r is equal to the reduced cost of that route $\bar{c}_r$. In this next section, we will explore how to compute a set of edge weights satisfying this condition.

Let q_r^e be a parameter denoting the number of times edge e occurs on route r. Observe that the following equality holds:

$$\alpha_{ir} = \sum_{e \in \delta(i)} \frac{q_r^e}{2}$$

By substituting this into (2.5), we can work towards expressing the reduced cost of a route as a sum over the set of edges:

$$\begin{aligned}\bar{c}_r &= c_r - \sum_{i \in V_+} \alpha_{ir} \pi_i^* - \pi_0^* \\ &= \sum_{e \in E} \ell_e q_r^e - \sum_{i \in V_+} \sum_{e \in \delta(i)} \frac{q_r^e}{2} \pi_i^* - \sum_{e \in \delta(0)} \frac{q_r^e}{2} \pi_0^* \\ &= \sum_{e \in E} \ell_e q_r^e - \sum_{i \in V} \sum_{e \in \delta(i)} \frac{q_r^e}{2} \pi_i^*\end{aligned}$$

Note that the sum $\sum_{i \in V} \sum_{e \in \delta(i)} \frac{q_r^e}{2} \pi_i^*$ counts each edge $e = uv$ twice, once for the vertex u and once for the vertex v. We can therefore rewrite this as:

$$\bar{c}_r = \sum_{e = uv \in E} \left(\ell_e - \frac{\pi_u^*}{2} - \frac{\pi_v^*}{2} \right) q_r^e$$

Define the edge weights:

$$h_e := \ell_e - \frac{\pi_u^*}{2} - \frac{\pi_v^*}{2}, \quad \forall e = uv \in E \tag{2.7}$$

We can now restate the pricing problem (2.6) as finding the route of minimum cost with respect to the edge weights $h \in \mathbb{R}^{|E|}$:

$$\min \left\{ \sum_{e \in E} h_e q_r^e : r \in \mathfrak{R} \right\} \tag{2.8}$$

Note that due to the incorporation of the dual variables π, we cannot add any assumptions on the weights h such as non-negativity, or respecting the triangle inequality. Although it may be possible to develop an algorithm for solving (2.8) by using the facts that in equation (2.7) the edge costs ℓ_e are metric and the dual variables π are an extreme point solution to D^{RM}, we are not aware of any such algorithm in the literature. We will therefore focus on the hardness of solving (2.8) for arbitrary edge costs h_e. Solving this problem efficiently is at the core of column generation, as it must be done at every iteration of the simplex algorithm.

Since the routes in the set $\mathfrak{R}$ are elementary, the problem (2.8) is more specifically an elementary shortest path problem with resource constraints (ESPPRC). However, there is one fatal flaw with this: solving the ESPPRC is strongly NP-hard [13]. Thus to solve P^{LP} via column generation, we would need to solve a strongly NP-hard problem at every simplex iteration. There are two approaches which can be taken here. One is to solve the ESPPRC, which gives tighter lower bounds in the branch-and-bound tree, but is much harder to solve and requires more computational tricks. This approach has been taken by several authors, for example Chabrier [7], Feillet, Dejax, Gendreau, and Gueguen [17], and Rousseau, Focacci, Gendreau, and Pesant [35]. Another approach is to relax the set of routes to allow walks with repeated vertices, or q-routes, which means the pricing problem is an SPPRC rather than an ESPPRC.

Definition 2.2.1. (q-route) A q-route is a sequence of vertices $(v_0 = 0, v_1, ..., v_k, v_{k+1} = 0)$ with $k > 0$, such that $v_i \neq v_{i+1}$ and $v_i \in V_+$ for $1 \leq i \leq k$, and satisfying $\sum_{i=1}^{k} d_i \leq Q$.

Note that while the primary difference between q-routes and elementary routes is that q-routes allow for repeated vertices, another implicit difference is that in an elementary route a customer's demand is counted at most once, whereas in a q-route a customer's demand is counted as many times as it is visited. This may seem undesirable, as in practice a customer would no longer have any demand after being visited, but without this property the pricing problem for q-routes would remain strongly NP-hard.

Using q-routes provides weaker lower bounds, but the resulting SPPRC can be solved in pseudo-polynomial time, which will be shown in a further section. This second approach of relaxing the feasible region to simplify the problem is the one we will be using.

We denote the set of all q-routes $\mathfrak{R}_q$. Note that $\mathfrak{R} \subseteq \mathfrak{R}_q$, as every elementary route is also a q-route. Incorporating all of these routes results in a new formulation P_q:

$$(P_q) \qquad \begin{aligned} \min \sum_{r \in \mathfrak{R}_q} c_r x_r & & \qquad (2.9a)\\ \text{s.t.} \sum_{r \in \mathfrak{R}_q} \alpha_{ir} x_r = 1 & \qquad \forall v \in V_+ & \qquad (2.9b)\\ \sum_{r \in \mathfrak{R}_q} x_r = m & & \qquad (2.9c)\\ x_r \in \{0, 1\} & \qquad \forall r \in \mathfrak{R}_q & \qquad (2.9d) \end{aligned}$$

Observe that even though we have relaxed the feasible region of the problem, due to constraint (2.9b) the optimal solution to P and P_q will still be the same. With this change, our pricing problem becomes:

$$\min \left\{ \sum_{e \in E} h_e q_r^e : r \in \mathfrak{R}_q \right\} \qquad (2.10)$$

This pricing problem can be modelled as an SPPRC rather than an ESPPRC, as there is no longer a restriction for vertices to only appear once on each route.

2.2.3 Solving the SPPRC

While the SPPRC is still NP-hard, it can be solved in pseudo-polynomial time via a dynamic programming algorithm [25], so an efficient solution to the pricing problem is in our grasp. However, this algorithm does require the further restriction that the resources $b_v \in \mathbb{Z}$ for all $v \in V$. Furthermore, we assumed that the budget β is an upper bound for b_v for all $v \in V$, as otherwise that vertex cannot be visited without violating the resource consumption constraints.

The particular algorithm used is known as a labelling algorithm. It represents q-routes and partially completed q-routes of minimum cost with respect to the cost of resources by objects called *labels*.

Definition 2.2.2. (partial q-route) A partial q-route is a sequence of vertices $(v_0 = 0, v_1, ..., v_k)$ such that $(v_0 = 0, v_1, ..., v_k, v_{k+1} = 0)$ is a q-route.

A label $L = (u, \gamma)$ is a tuple where u is the final vertex in the partial q-route and $\gamma \geq 0$ is the cumulative resource consumption along the partial q-route. The minimum cost with respect to the edge weights h of a partial q-route corresponding to label L is denoted by $h(L)$. Let $\Pi(L)$ denote the set of predecessor labels that may be extended to L:

$$\Pi(L) := \{L' = (u', \gamma') : \gamma' + \gamma_u = d,\ uu' \in E\}. \tag{2.11}$$

To compute $h(L)$, we apply the following dynamic programming recursion:

$$h(L) := \min\{h(L') + h_{uu'} : L' = (u', \gamma') \in \Pi(L)\}. \tag{2.12}$$

The base case for this recursion is $\bar{c}(L_0) = 0$, where $L_0 = (0, 0)$. Let:

$$\mathcal{L} := \{L = (u, \gamma) : u = 0, 0 < \gamma \leq \beta\}. \tag{2.13}$$

Observe any label in $\mathcal{L}$ corresponds to a complete q-route, as it ends at vertex 0 and has consumed a positive quantity of resources. Therefore the SPPRC can be solved by computing the minimum cost label $L^* = \min\{L : L \in \mathcal{L}\}$. This can done in $O(n^2\beta)$ time, as there are n possible ending vertices for labels, and β possible values for the cumulative resource consumption γ, and therefore $n \times \beta$ possible labels to be computed through dynamic programming, while each label can itself be computed in $O(n)$ time as the set $\Pi(L)$ has cardinality at most n.

Therefore we can use this algorithm to solve the pricing problem in $O(n^2Q)$ time by using the edge weights h as defined in (2.7), the customer demands d for resources, and the vehicle capacity Q for the budget.

Now that we have described the algorithm used to solve the SPPRC, we can elaborate on why we defined q-routes as in Definition 2.2.1, rather than ignoring the demands of customers which have previously been visited. If we ignored the demands of already visited customers, we would have an exponential number of labels, as instead of only needing to track the final vertex visited and the cumulative resource consumption we would also need to track which vertices had been previously visited. This is exactly what is required when pricing elementary routes, and so with such a definition we would not simplify the pricing problem in any way.

For further details about the SPPRC, the ESPPRC and related variants and algorithms, see Righini and Salani [34] or Irnich and Desaulniers [25].

2.2.4 The final restricted master problem

We will now summarize the modifications to the formulation for the CVRP introduced in this section and give a brief recap of how column generation is used. Since we have modified the formulation to use q-routes in order to simplify the pricing problem, the restricted master problem must also be modified accordingly. Let $\overline{\mathfrak{R}}_q \subset \mathfrak{R}_q$ with $|\overline{\mathfrak{R}}_q| \ll |\mathfrak{R}_q|$. We use this in the new restricted master problem P_q^{RM}:

$$(P_q^{RM}) \qquad \begin{aligned} \min \sum_{r \in \overline{\mathfrak{R}}_q} c_r x_r & & \text{(2.14a)} \\ \text{s.t.} \sum_{r \in \overline{\mathfrak{R}}_q} \alpha_{ir} x_r = 1 & \qquad \forall i \in V_+ & \text{(2.14b)} \\ \sum_{r \in \overline{\mathfrak{R}}_q} x_r = m & & \text{(2.14c)} \\ 0 \leq x_r \leq 1 & \qquad \forall r \in \overline{\mathfrak{R}}_q & \text{(2.14d)} \end{aligned}$$

Similarly to P^{RM}, artificial variables may be introduced to ensure feasibility of P_q^{RM}. This formulation is solved via the simplex method to obtain an optimal solution x^* and dual optimal solution π^*. The dual solution π^* is used to compute the edge weights h_e as per

(2.7), and then the pricing problem (2.10) is solved. If the optimal value of (2.10) is negative, then the corresponding q-route is added to the set $\overline{\mathfrak{R}}_q$, P_q^{RM} is re-solved, and the process is repeated. On the other hand, if the optimal value of (2.10) is non-negative, then x^* is optimal for P_q^{LP} and column generation ends. This technique is then applied at every node in the branch-and-bound tree for P_q.

Chapter 3

The capacitated vehicle routing problem with stochastic demands

In this section, we will define the capacitated vehicle routing problem with stochastic demands, and look at the complexities that arise from having stochastic demands instead of deterministic demands.

3.1 Problem Definition & Model Formulation

Most of the definitions for the CVRPSD are shared with the CVRP, the only difference being how the demands are specified. The CVRPSD is defined on an undirected graph $G = (V, E)$, where $V = \{0, ..., n\}$. Vertex 0 represents the depot, and $V_+ = \{1, ..., n\}$ represents the set of customers. For each edge $ij \in E$ there is an associated travel cost ℓ_{ij}. There are m vehicles to be routed, each with capacity Q. Each customer $i \in V_+$ has an associated random variable ξ_i for their demand, with expected value $\mathbb{E}[\xi_i]$ and variance $V[\xi_i]$. These form a customer demand vector ξ. Furthermore, we assume that $\mathbb{E}[\xi_i] \leq Q$ for all $i \in V_+$.

Similar to the CVRP, we will have a decision variable in the formulation for each feasible q-route. However, as customer demands are no longer deterministic, the feasibility characteristic for q-routes is changed. Instead of the sum of demands being less than the vehicle capacity, we require the sum of the *expected* demands on the q-route to be less than the vehicle capacity.

Definition 3.1.1 (Feasible q-route, CVRPSD)**.** A q-route $r = (0, v_1, ..., v_k, 0)$ is *feasible* for a CVRPSD instance if $\sum_{i=1}^{k} \mathbb{E}[\xi_i] \leq Q$.

As before, q-routes are assumed to be feasible unless specified otherwise. Let $\mathfrak{R}_q$ be the set of q-routes. We consider the following set partitioning formulation:

$$(SP) \qquad \begin{aligned} \min \sum_{r \in \mathfrak{R}_q} c_r x_r & & (3.1a) \\ \text{s.t.} \sum_{r \in \mathfrak{R}_q} \alpha_{ir} x_r = 1 \quad \forall i \in V_+ & & (3.1b) \\ \sum_{r \in \mathfrak{R}_q} x_r = m & & (3.1c) \\ x_r \in \{0, 1\} \quad \forall r \in \mathfrak{R}_q & & (3.1d) \end{aligned}$$

The coefficients c_r denotes the expected cost of r, the computation of which is described below; and α_{ir} is a parameter which counts the number of times a q-route r visits vertex i. The objective function (3.1a) minimizes the expected cost of the q-routes. Constraints (3.1b) ensure that each customer is visited exactly once by some q-route, and constraint (3.1c) enforces that exactly m q-routes are used.

To compute the expected cost c_r for a q-route $r = (v_0 = 0, v_1, ..., v_k, v_{k+1} = 0)$, we divide it into two quantities: $\hat{c}_r$, the deterministic cost of travelling along the q-route, and $\tilde{c}_r$, the expected cost incurred by failures along the q-route, with $c_r = \hat{c}_r + \tilde{c}_r$. The value of $\hat{c}_r$ is equal to the sum of the edge costs along the q-route:

$$\hat{c}_r = \sum_{i=0}^{k} \ell_{i,i+1}.$$

We use the simple recourse strategy of making a round-trip to the depot and back if a customer's demand exceeds the vehicle capacity before continuing the q-route as planned. This allows us to compute the expected cost of recourse $\tilde{c}_r$ by summing the *expected failure cost* at each customer v_i along the q-route. Given a q-route $r = (0, v_1, ..., v_k, 0)$, let Ψ_{r,v_i} denote the probability distribution of the cumulative demand of r at vertex v_i, and let ψ_{r,v_i} denote a Ψ_{r,v_i}-distributed random variable. The expected number of failures at customer v_i, $\text{FAIL}_{v_i}(r)$, is then computed by the following:

$$\text{FAIL}_{v_i}(r) = \sum_{u=1}^{\infty} \left[\mathbb{P}\{\psi_{r,v_{i-1}} \leq uQ\} - \mathbb{P}\{\psi_{r,v_i} \leq uQ\} \right] \qquad (3.2)$$

where $\mathbb{P}\{E\}$ denotes the probability of event E occurring. Thus $\mathbb{P}\{\psi_{r,v_{i-1}} \leq uQ\}$ is the probability that the uth failure has not occurred before visiting customer i, and $\mathbb{P}\{\psi_{r,i} \leq uQ\}$ is the probability that the uth failure has still not occurred after visiting customer i. Therefore the difference:

$$\mathbb{P}\{\psi_{r,i-1} \leq uQ\} - \mathbb{P}\{\psi_{r,i} \leq uQ\}$$

is the probability of the uth failure occurring exactly at customer i. The expected failure cost at customer v_i $\mathrm{EFC}_{v_i}(r)$ can then by computed by multiplying this by the distance travelled when a failure occurs:

$$\mathrm{EFC}_{v_i}(r) = 2\ell_{0v_i}\,\mathrm{FAIL}_{v_i}(r). \tag{3.3}$$

Thus:

$$\tilde{c}_r = \sum_{i=1}^{k} \mathrm{EFC}_{v_i}(r). \tag{3.4}$$

3.1.1 The pricing problem for the CVRPSD

Just as with the set-partitioning formulations for the CVRP described in the previous chapter, SP contains an exponential number of variables with respect to the number of customers, and so column generation will be required to solve it. This requires us to construct the restricted master problem SP^{RM}, where the feasible region is restricted to a small set of q-routes $\overline{\mathfrak{R}}_q \subset \mathfrak{R}_q$, and then iteratively add new q-routes to $\overline{\mathfrak{R}}_q$ as they are identified to be useful via column generation.

We refer to the column generation subproblem of finding the q-route with minimum expected reduced cost as the *minimum 2-stage cost q-route problem* to distinguish it from the deterministic version of the problem described in the previous chapter.

Let DSP^{LP} denote the dual of the LP relaxation of SP:

$$\left(DSP^{LP}\right) \qquad \max \sum_{i \in V_+} \pi_i + m\pi_0 \tag{3.5a}$$

$$\text{s.t.} \sum_{i \in V_+} \alpha_{ir}\pi_i + \pi_0 \leq \hat{c}_r + \tilde{c}_r \qquad \forall r \in \mathfrak{R} \tag{3.5b}$$

Here the expected q-route cost c_r has been split into its deterministic and stochastic components. Similarly to the CVRP, the reduced cost of r is defined to be:

$$\bar{c}_r := \hat{c}_r + \tilde{c}_r - \sum_{i \in V_+} \alpha_{ir}\pi_i^* - \pi_0^* \tag{3.6}$$

Where π^* is the optimal solution of the dual to the current restricted master problem. Then the pricing problem is:

$$\min\{\bar{c}_r : r \in \mathfrak{R}_q\} \tag{3.7}$$

However, unlike in the CVRP, due to the stochastic element $\tilde{c}_r$ in $\bar{c}_r$ we can't transform this problem in one only involving edge prices. This is because in our formulation, the distribution of a partial q-route's cumulative demand is dependent on both the q-route and vertex, meaning that two partial q-routes which end at the same vertex and have the same expected cumulative demand may not have the same distribution. Thus, although the deterministic component of $\bar{c}_r$ may be decomposed into a sum of weights on the edges, the stochastic component cannot be.

Despite this, for the later proofs it will still be useful to have this problem formulated as one in which the deterministic component of $\bar{c}_r$ is expressed as a sum of edge weights. This is achieved similarly to our reformulation of the pricing problem for the CVRP: define the parameter q_r^e to denote the number of times an edge $e \in E$ appears on the q-route r, and define:

$$h_e := \ell_e - \frac{\pi_u^*}{2} - \frac{\pi_v^*}{2}, \quad \forall e = uv \in E \tag{3.8}$$

Note again that do cannot have any assumptions on h_e such as non-negativity or integrality. Then (3.7) can be restated as the following formal definition for the minimum cost q-route problem.

Definition 3.1.2. Consider an instance of the CVRPSD on a graph $G = (V, E)$. Given an arbitrary set of edge weights $w_e \in \mathbb{R}$, the minimum cost q-route problem is:

$$\min \left\{ \sum_{e \in E} w_e q_r^e + \tilde{c}_r : r \in \mathfrak{R}_q \right\} \tag{3.9}$$

Where $\tilde{c}_r$ is computed according to (3.4).

3.1.2 Difference between pricing q-routes for CVRP and CVRPSD

As discussed earlier on, q-routes can be priced in pseudo-polynomial time via dynamic programming in the CVRP. The reason this can be done in the deterministic case is because if two q-routes r^1 and r^2 end at the same vertex u and have the same cumulative demand d, they are functionally equivalent with respect to the recursive structure in the pricing problem. By this, we mean that any valid extension that may be made to r^1 may also be made to r^2, and vice versa.

It is clear how dependent customer demands make the CVRPSD very challenging, as in this case it is not necessarily even clear how to evaluate $\tilde{c}_r$. When customer demands are independent, there is more hope for finding a good solution, as the only difference between pricing q-routes for independent stochastic demands and deterministic demands is the addition of variance. However, even this is still more complex than in the deterministic case. This is because in the case of q-routes the demands of vertices at different points along the route will be correlated if the route contains repeated vertices, even if the demands between different customers are independent. Suppose $r = (0, v_1, ..., v_k, 0)$ is a q-route such that $v_i = v_j$ for $i \neq j$, $1 \leq i, j, \leq k$. Then v_i and v_j have perfectly correlated demand.

One problem stemming from this core issue is that the cumulative variance of a q-route is not necessarily additive, a property which is assumed in previous approaches to this problem [8, 21]. If the q-route is elementary, then the cumulative variance is additive, but if the q-route is not elementary then the cumulative variance is also not additive. Consider having k copies of a random variable ξ with expected value $\mathbb{E}[\xi]$ and variance $V[\xi]$. By linearity of expectation, the expected value of $k\xi$ is:

$$\mathbb{E}[k\xi] = \mathbb{E}\left[\sum_{i=1}^{k} \xi\right] = k\mathbb{E}[\xi]$$

However variance is not linear. Using the definition $V[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$, we can compute the variance of $k\xi$ to be:

$$\begin{aligned} V[k\xi] &= \mathbb{E}[(k\xi)^2] - \mathbb{E}[k\xi]^2 \\ &= \mathbb{E}[k^2\xi^2] - (k\mathbb{E}[\xi])^2 \\ &= k^2\mathbb{E}[\xi^2] - k^2\mathbb{E}[\xi]^2 \\ &= k^2(\mathbb{E}[\xi^2] - \mathbb{E}[\xi]^2) \\ &= k^2 V[\xi] \end{aligned}$$

Thus if a customer i is repeated k times in a q-route, it should contribute $k^2V[\xi_i]$ to the total variance of the q-route, rather than $kV[\xi_i]$. This means that given a path $P = (u_1, ..., u_h)$ and a vertex v, to calculate the cumulative variance of the augmented path $P' = (u_1, ..., u_h, v)$ we need to know how many times v has already been visited by P, not simply the cumulative variance of P and the variance of v. This is the core issue that differentiates pricing q-routes for the deterministic CVRP and the CVRPSD with independent demands.

To illustrate this, we present the following example. Consider a CVRPSD instance consisting of three customer nodes $1, 2, 3$, each represented by independent normally distributed random variables with expected demand $\mu_i = 100$ and variance $\sigma_i^2 = 100$ for

$i = 1, 2, 3$. Let the vehicle capacity be $Q = 305$. Consider the two q-routes $r^1 = (0, 1, 2, 3, 0)$ and $r^2 = (0, 1, 2, 1, 0)$ (section 3.1.2).

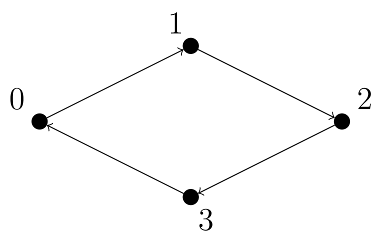

Figure 3.1: Route $r^1 = (0, 1, 2, 3, 0)$.

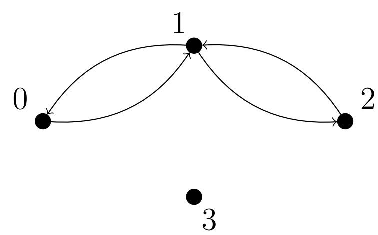

Figure 3.2: Route $r^2 = (0, 1, 2, 1, 0)$.

The expected number of failures on the elementary q-route r^1 is then:

$$\text{FAIL}_1(100, 100) + \text{FAIL}_2(200, 200) + \text{FAIL}_1(300, 300) = 0.386$$

And the expected number of failures on the non-elementary q-route r^2 is:

$$\text{FAIL}_1(100, 100) + \text{FAIL}_2(200, 200) + \text{FAIL}_1(300, 600) = 0.419$$

Thus the non-elementary q-route r^2 has a failure rate significantly higher than that of r^1, even though each q-route visits a total of 3 vertices with each vertex having identical mean and variance.

We additionally computed the expected number of failures empirically by generating 1,000,000 samples of normal demands on such a q-route and checking the number of failures in each sample. This gave us an expected number of failures of 0.412, indicating that increasing the variance when revisiting a node on a q-route more accurately estimates the number of failures than not increasing the variance. Running the same test on q-route r^1 yielded an expected number of failures of 0.386.

Another issue that arises from this is that the distribution of a q-route's cumulative demand can diverge from that of the customers. Consider the case of an instance with n identical Poisson distributed demands, with expected demand and variance λ. Let $r = (0, v_1, v_2, ..., v_h, 0)$ be a non-elementary q-route, with $v_i = v_j$, $1 \leq i < j \leq k$. For each vertex v_k, $k < j$ along this q-route, the cumulative expected demand and variance will be $k\lambda$, and the random variable for the cumulative demand will remain Poisson. But at v_j the cumulative expected demand will be $j\lambda$ while the variance will be $(j + 3)\lambda$, implying that the random variable for the cumulative demand at v_j is no longer Poisson, along with the random variables for cumulative demand at all nodes following v_j. It is then no longer obvious how to even evaluate the expected failure cost of the q-route at each of these nodes.

This fact that we need to know information about the number of times a vertex has occurred so far on the q-route suggests that this problem may be similarly hard as the elementary q-route pricing problem, which is known to be strongly NP-Hard.

3.1.3 Contrast with previous approaches

Previous approaches to this problem such as by Christiansen and Lysgaard [8] and Gauvin et al. [21] solved a very similar formulation with a modified objective function for a restricted family of distributions.

Christiansen and Lysgaard [8] focus on the case where the underlying probability distribution Ψ of the demands is additive, meaning that the sum of independent Ψ-distributed random variables is also Ψ-distributed. This means that there is no need to keep track of different distributions $\Psi_{r,i}$ for each q-route and vertex, which greatly simplifies the calculation of FAIL_i and EFC_i. This property holds for many commonly used distributions, including normal and Poisson distributions [30].

In this case, given a walk with non-repeated vertices $(0, v_1, ..., v_i)$ with cumulative expected mean μ and variance σ^2, the expected number of failures at v_i, $\text{FAIL}_{v_i}(\mu, \sigma^2)$ is computed by the following:

$$\text{FAIL}_{v_i}(\mu, \sigma^2) = \sum_{u=1}^{\infty} \left[\mathbb{P}\{\Psi(\mu - \mathbb{E}[\xi_{v_i}], \sigma^2 - V[\xi_{v_i}]) \leq uQ\} - \mathbb{P}\{\Psi(\mu, \sigma^2) \leq uQ\}\right]$$

Instead of evaluating the expected cost of q-routes exactly, they computed an approximate value obtained by assuming all customer demands on a q-route are independent even if the same customer is repeated. Essentially this means that the demand at a customer is realized once for each time they are visited, rather than once overall. Using this relaxation of the expected cost of q-routes meant that pricing routes could still be done in pseudo-polynomial time using a modified version of the dynamic programming algorithm for pricing q-routes in the CVRP.

Using this approximation of the expected cost still results in an exact solution, as at an integral solution with all elementary q-routes the approximate expected cost and expected cost would coincide. However this method would result in weaker LP bounds throughout the branch-and-price tree than one which exactly prices non-elementary q-routes, while exactly pricing q-routes would give stronger bounds, as the approximate method will underestimate the variance of q-routes, especially those containing many cycles, and by extension their second stage costs. Therefore we want to explore whether it is possible to exactly price routes without significantly increasing the difficulty of solving the pricing problem.

Chapter 4

Complexity of solving the minimum cost 2-stage q-route problem

We now present our results on the complexity of solving the minimum cost 2-stage q-route problem. We first show that in the finite scenario case, pricing q-routes when including the 2-stage costs is strongly NP-hard. Secondly, we show that while it is not necessarily strongly NP-hard in the case of independent normal demands, under the assumption that elementary operations are $O(1)$ and the computability of a function $\textsc{Rfail}(\mu, \sigma^2)$ in polynomial-time, there is a reduction from the Hamiltonian cycle problem to the minimum cost 2-stage q-route problem. This holds even when demands are variances are polynomially bounded in n. While this would seem to imply strong NP-hardness, due to the aforementioned assumptions required, this cannot be concluded. However, it does suggest that this problem may be harder than it is currently thought to be, and that further study of its complexity is warranted.

4.1 Finite Scenarios

In the finite scenario case, we are solving (3.9) with the distribution of customer demands being given as a finite set of sample scenarios $s^j, 1 \leq j \leq S$. Each scenario s^j has a probability $p_j, 1 \leq j \leq S$ of being realized, with $\sum_{j=1}^{S} p_j = 1$. This means that the customer demand vector ξ takes on the value s^j with probability p_j.

Let $G = (V_+ \cup \{0\}, E)$ be a graph, $V_+ = \{1, 2, ..., n\}$, $n \geq 3$. Consider the following instance of the minimum cost 2-stage q-route problem. Let $G' = (V_+ \cup \{0\}, E')$ be the complete graph on the same vertex set as G, and assign edge costs as follows:

$$\begin{aligned} h_{0i} &= n^3 && \forall 0i \in E \\ h_{0i} &= n^3 + 1 && \forall 0i \notin E \\ h_{ij} &= -1 && \forall ij \in E \\ h_{ij} &= 0 && \forall ij \notin E \end{aligned}$$

Let the capacity $Q = 2n - 1$. Construct n scenarios s^v, one for each vertex $v \in V_+$ with vertex v having demand n and all other vertices having demand 1. As in the previous chapter, $\tilde{c}_r$ denotes the expected failure cost of a q-route. In the following lemma, we will show that elementary q-routes in G', i.e. q-routes with non-repeated vertices, will never cause a route failure and therefore have no expected failure cost, while non-elementary q-routes will cause route failures and therefore have a high expected failure cost.

Lemma 4.1.1. *If r is an elementary q-route in G' then $\tilde{c}_r = 0$. If r is non-elementary, then $\tilde{c}_r \geq n^2$.*

Proof. If r is elementary, then the total demand on r in any scenario is at most $2n - 1$. Therefore r will never fail, and so $\tilde{c}_r = 0$.

If r is non-elementary, then there exists some vertex $v \in V_+$ which is visited twice. Therefore r fails in scenario s^v, implying that $\tilde{c}_r \geq \frac{1}{n} n^3 = n^2$. □

Using the fact that non-elementary q-routes have a high failure cost while elementary q-routes do not, we can now show that this positive failure cost will dominate the negative cost of the edges, implying that minimum cost q-route must be elementary.

Lemma 4.1.2. *The minimum cost q-route in G' is elementary.*

Proof. For any vertex $v \in V_+$, the elementary q-route $(0, v, 0)$ has cost $\leq 2n^3 + 2$. Therefore the minimum cost route must have cost less than or equal to $2n^3 + 2$.

If r is a non-elementary q-route, then by Lemma 4.1.1 $\tilde{c}_r \geq n^2$ and so we have $c_r \geq 2n^3 + \tilde{c}_r \geq 2n^3 + n^2$. Since the expected demand at each vertex is $\frac{2n-1}{n}$, any q-route feasible for the capacity constraints can visit at most n nodes, as otherwise the expected demand on the q-route would be greater than the capacity of $2n - 1$. Thus any non-elementary q-route can have its cost reduced by at most n by travelling along negative edges, and so:

$$c_r \geq 2n^3 + n^2 - n$$

Note that for all $n \geq 3$, $2n^3 + n^2 - n > 2n^3 + 2$. Thus any non-elementary q-route r cannot be a minimum cost route. □

We can now prove the hardness result for the finite scenario case.

Theorem 4.1.3. *When demands are given as a finite set of scenarios, solving the minimum cost 2-stage q-route problem is strongly NP-Hard.*

Proof. By Lemma 4.1.2, solving the minimum cost 2-stage q-route problem is equivalent to finding a minimum cost elementary q-route, and G has a Hamiltonian cycle if the minimum cost elementary q-route in G' has cost $2n^3 - n$. All of the parameters for the minimum cost 2-stage q-route problem instance constructed above, such as the vehicle capacity, edge weights, number of scenarios, and demands in each scenario, are polynomial in n. Therefore as the Hamiltonian cycle problem is strongly NP-hard, the minimum cost 2-stage q-route problem is strongly NP-hard. □

4.2 Independent Normal Demands

For the independent normal case we are again solving (3.9) with each customer demand $\xi_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, i.e. the demand at each customer is independent and normally distributed with known means μ_i and variances σ_i^2, $1 \le i \le n$.

In our approach for proving a reduction from the Hamiltonian cycle problem to the minimum cost 2-stage q-route problem, we need to use the intuitive fact that if the cumulative expected demand or cumulative variance of a route is higher, then the route has a higher likelihood of failing, and by extension a higher cost of failure. However, while intuitively true, proving that this is the case using the definitions of FAIL (3.2) and EFC (3.3) is very challenging, as they compute the expected failure rate and failure cost at each vertex independently. While overall the expected failure rate should go up with higher cumulative demand and variance, the expected failure rate at a specific vertex may either increase or decrease. This is because a higher chance to fail at an earlier vertex could decrease the chance of failing at a later vertex, as there would be a higher probability of the vehicle being less full due to the earlier failure emptying all the goods carried by the vehicle.

We therefore first need to define a new function RFAIL, which counts the total expected number of failures on a route. We will first state its definition before explaining how the given expression for RFAIL was derived.

Definition 4.2.1 (Route Fail)**.** RFAIL(μ, σ^2) counts the expected number of failures on a q-route with a cumulative demand of μ and cumulative variance of σ^2, and can be computed by:

$$\textsc{Rfail}(\mu, \sigma^2) = \frac{1}{\sqrt{2\pi}} \sum_{u=1}^{\infty} \int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}} dt$$

Note that as this function contains an infinite sum, where each term is itself an improper integral. It is therefore not immediately clear that this function is even well-defined. Before proving that it is actually well-defined, we will first explain the logic behind how this function counts the expected number of failures of a q-route.

It is beneficial to consider the case of computing the probability that a route r will fail exactly u times. Let $X_r \sim \mathcal{N}(\mu, \sigma^2)$ be a random variable representing the cumulative demand of a route r with cumulative mean μ and cumulative variance σ^2. As a route experiences a failure whenever the cumulative demand exceeds the vehicle capacity Q, the probability that a route fails exactly u times is the probability:

$$\mathbb{P}\{uQ < X_r \leq (u+1)Q\}$$

It can be seen that this is equal to:

$$\mathbb{P}\{X_r > uQ\} - \mathbb{P}\{X_r > (u+1)Q\}$$

As this is the probability that r fails exactly u times, to count the expected number of failures contributed by r failing exactly u times we simply multiply this quantity by u. Therefore the total expected number of failures of r can be computed by the following infinite series:

$$\sum_{u=1}^{\infty} u[\mathbb{P}\{X_r > uQ\} - \mathbb{P}\{X_r > (u+1)Q\}] \tag{4.1}$$

Observe that as we expand the sum (4.1), the uth term subtracts $u \cdot \mathbb{P}\{X_r > (u+1)Q\}$, while the $(u+1)$th term adds $(u+1) \cdot \mathbb{P}\{X_r > (u+1)Q\}$. We can use this to cancel out the subtracted term, obtaining the significantly simpler series:

$$\sum_{u=1}^{\infty} \mathbb{P}\{X_r > uQ\} \tag{4.2}$$

Let $\Phi(x)$ denote the cumulative distribution function of the standard normal distribution. Given $X \sim \mathcal{N}(0,1)$, this is the probability $\mathbb{P}\{X \leq x\}$. Recall that this is defined to be the integral:

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{t^2}{2}} dt \tag{4.3}$$

Furthermore, let $\Phi^c(x)$ denote the corresponding complementary cumulative distribution function, which gives the probability $\mathbb{P}\{X > x\}$. This is equal to:

$$\Phi^c(x) = 1 - \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{t^2}{2}} dt \tag{4.4}$$

Recalling that $X_r \sim \mathcal{N}(\mu, \sigma^2)$, (4.2) may be expressed as:

$$\sum_{u=1}^{\infty} \Phi^c\left(\frac{uQ - \mu}{\sigma}\right) = \frac{1}{\sqrt{2\pi}} \sum_{u=1}^{\infty} \int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}} dt \tag{4.5}$$

Which is the expression used in the definition for RFAIL.

We can now begin proving of the properties required to use RFAIL for our complexity results. However, before doing this we first require an upper bound on $\Phi^c(x)$.

Lemma 4.2.2 (Upper bound on $\Phi^c(x)$)**.** *The following is an upper bound on the value of $\Phi^c(x)$:*

$$\Phi^c(t) < \frac{1}{\sqrt{2\pi}} \int_x^{\infty} \frac{t}{x} e^{-\frac{t^2}{2}} dt. \tag{4.6}$$

Proof.

$$\begin{aligned}
\Phi^c(t) &= \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{t^2}{2}} dt \\
&< \frac{1}{\sqrt{2\pi}} \int_x^{\infty} \frac{t}{x} e^{-\frac{t^2}{2}} dt \\
&= \frac{1}{\sqrt{2\pi}} \frac{1}{t} e^{-\frac{t^2}{2}}.
\end{aligned}$$

□

Using this upper bound, we will now show that the summation $\frac{1}{\sqrt{2\pi}} \sum_{u=1}^{\infty} \int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}} dt$ converges to a finite value, which in turn implies that RFAIL is well-defined.

Lemma 4.2.3. $\frac{1}{\sqrt{2\pi}} \sum_{u=1}^{\infty} \int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}} dt$ *converges to a finite value.*

Proof.

$$\begin{aligned}\frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sigma}}^{\infty}e^{-\frac{t^2}{2}}dt &= \sum_{u=1}^{\infty}\Phi^c\left(\frac{uQ-\mu}{\sigma}\right)\\ &< \sum_{u=1}^{\infty}\frac{1}{\sqrt{2\pi}}\frac{1}{\left(\frac{uQ-\mu}{\sigma}\right)}e^{-\frac{(uQ-\mu)^2}{2\sigma^2}}\\ &= \sum_{u=1}^{\infty}\frac{\sigma}{\sqrt{2\pi}(uQ-\mu)}e^{-\frac{(uQ-\mu)^2}{2\sigma^2}}.\end{aligned}$$

We will show this series converges by applying the ratio test. This is a criterion for convergence of a series $\sum_{n=1}^{\infty}a_n$, where the value of the limit $L=\lim_{n\to\infty}\left|\frac{a_n}{a_{n+1}}\right|$ to test convergence. Particularly, if $L<1$, then the series $\sum_{n=1}^{\infty}a_n$ converges absolutely.

$$\begin{aligned}\lim_{u\to\infty}\left|\frac{\frac{\sigma}{\sqrt{2\pi}((u+1)Q-\mu)}e^{-\frac{((u+1)Q-\mu)^2}{2\sigma^2}}}{\frac{\sigma}{\sqrt{2\pi}(uQ-\mu)}e^{-\frac{(uQ-\mu)^2}{2\sigma^2}}}\right| &= \lim_{u\to\infty}\left|\frac{(uQ-\mu)e^{-\frac{u^2Q^2+2uQ^2+Q^2-2uQ\mu-2Q\mu+\mu^2}{2\sigma^2}}}{(uQ+Q-\mu)e^{-\frac{u^2Q^2-2uQ\mu+\mu^2}{2\sigma^2}}}\right|\\ &= \lim_{u\to\infty}\left|\frac{uQ-\mu}{uQ+Q-\mu}e^{-\frac{2uQ^2+Q^2-2Q\mu}{2\sigma^2}}\right|\\ &= 0\\ &< 1.\end{aligned}$$

Therefore $\frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sigma}}^{\infty}e^{-\frac{t^2}{2}}dt$ converges, and by extension $\textsc{Rfail}(\mu,\sigma^2)$ is well-defined.

□

Now that we have shown that this expression for RFAIL is well-defined, we can use it to prove the intuitive property discussed at the start of this section, namely that when the cumulative expected demand or cumulative variance of a route is increased, the route has a higher likelihood of failing.

Lemma 4.2.4. RFAIL(μ,σ^2) *is strictly increasing in* μ*, and strictly increasing in* σ^2 *when* $\mu\leq Q$*.*

Proof. First we will show RFAIL is strictly increasing in μ. Let $\epsilon > 0$.

$$\begin{aligned}
\text{RFAIL}(\mu+\epsilon,\sigma^2) - \text{RFAIL}(\mu,\sigma^2) &= \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu-\epsilon}{\sigma}}^{\infty} e^{-\frac{t^2}{2}}\,dt - \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}}\,dt \\
&= \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\left(\int_{\frac{uQ-\mu-\epsilon}{\sigma}}^{\infty} e^{-\frac{t^2}{2}}\,dt - \int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}}\,dt\right) \\
&= \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu-\epsilon}{\sigma}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt
\end{aligned}$$

And for all $u \in \mathbb{Z}, u \geq 1$:

$$\frac{uQ-\mu}{\sigma} > \frac{uQ-\mu-\epsilon}{\sigma} \implies \int_{\frac{uQ-\mu-\epsilon}{\sigma}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt > 0 \implies \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu-\epsilon}{\sigma}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt > 0.$$

Thus RFAIL(μ,σ^2) is strictly increasing in μ.

Now suppose that $\mu \leq Q$, and again let $\epsilon > 0$.

$$\begin{aligned}
\text{RFAIL}(\mu,\sigma^2+\epsilon) - \text{RFAIL}(\mu,\sigma^2) &= \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sqrt{\sigma^2+\epsilon}}}^{\infty} e^{-\frac{t^2}{2}}\,dt - \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}}\,dt \\
&= \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\left(\int_{\frac{uQ-\mu}{\sqrt{\sigma^2+\epsilon}}}^{\infty} e^{-\frac{t^2}{2}}\,dt - \int_{\frac{uQ-\mu}{\sigma}}^{\infty} e^{-\frac{t^2}{2}}\,dt\right) \\
&= \frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sqrt{\sigma^2+\epsilon}}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt
\end{aligned}$$

Since $\mu \leq Q$, for $u = 1$:

$$\int_{\frac{uQ-\mu}{\sqrt{\sigma^2+\epsilon}}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt \geq 0$$

And for all $u \in \mathbb{Z}, u \geq 2$:

$$\int_{\frac{uQ-\mu}{\sqrt{\sigma^2+\epsilon}}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt > 0$$

Thus $\frac{1}{\sqrt{2\pi}}\sum_{u=1}^{\infty}\int_{\frac{uQ-\mu}{\sqrt{\sigma^2+\epsilon}}}^{\frac{uQ-\mu}{\sigma}} e^{-\frac{t^2}{2}}\,dt > 0$, and so RFAIL$(\mu,\sigma^2)$ is strictly increasing in σ^2 when $\mu \leq Q$. □

With these properties, we are now ready to prove the main theorem on the complexity of the 2-stage route pricing problem with independent normal demands. We will first give an informal overview of how we will approach this proof. Similar to the previous proof in the case of finite scenarios, we want to show that there exists a graph such that solving the minimum cost 2-stage q-route problem on that graph gives you a Hamiltonian cycle. Again we construct the graph such that all edges leaving the depot have a high cost, while edges not incident to the depot have a negative cost, in order to force the minimum cost route to be long. We additionally want to use the second stage costs of routes to penalize non-elementary q-routes, in order to force the minimum cost route to be elementary. As the only thing differentiating the second stage costs of elementary and non-elementary q-routes is the slightly increased variance of non-elementary q-routes, this implies that the edge weights we choose must be extremely large. The bulk of the proof is simply showing that the edge weights chosen do result in a distinction between the prices of elementary and non-elementary q-routes.

While most of the assumptions used in the statement of this theorem are standard, we do currently also require a more questionable assumption that the function $\text{RFAIL}(\mu, \sigma^2)$ is computable in polynomial time in n. This is because we need a very precise measurement of the value of RFAIL for many of the inequalities throughout the proof, and by doing this we may be accidentally shifting the complexity in the problem away from the combinatorial aspects and onto precisely computing this function. To develop a stronger theorem, this assumption could be replaced by finding good upper and lower bounds to RFAIL, then altering the theorem to use these bounds instead. However, at this time we were not successful in performing these modifications due to the complexity added to the arguments when using bounds on RFAIL.

Theorem 4.2.5. *Suppose there exists an algorithm to solve the minimum cost 2-stage q-route problem under the following assumptions:*

1. *Demands are independent and identically distributed normal with mean μ and variance σ^2.*
2. *μ and σ are constant integers which do not grow in n.*
3. *The capacity is polynomially bounded in n.*
4. *All elementary operations ($+, -, \times, \div$, assignment, comparison) require $O(1)$ time.*
5. $\text{RFAIL}(\mu, \sigma^2)$ *is computable in polynomial time in n.*

Then there exists an algorithm using polynomially many elementary operations that solves the Hamiltonian cycle problem with polynomially many calls to the algorithm solving the minimum cost 2-stage q-route problem.

Proof. Consider an instance of the CVRPSD on the graph $G = (V_+ \cup \{0\}, E)$, $V_+ = \{1, ..., n\}$, with $n \geq 3$. Let $\mu_i = 1$ and $\sigma_i^2 = 1$ for all $i \in V_+$, and let $Q = n + \frac{1}{2}$. Furthermore choose some $\epsilon > 0$, and let:

$$M_n = \frac{\textsc{Rfail}(n, n+3) + \epsilon}{\textsc{Rfail}(n, n+3) - \textsc{Rfail}(n, n)}$$

By Lemma 4.2.4 $\textsc{Rfail}(n, n+3) > \textsc{Rfail}(n, n)$ so M_n is well-defined. Additionally, let:

$$\begin{aligned} \alpha_i &= 2(M_n - 1)[\textsc{Rfail}(i, i) - \textsc{Rfail}(i-1, i-1)] \quad \forall i = 2, ..., n \\ \beta_i &= 2(M_n - 1)[\textsc{Rfail}(i, i+3) - \textsc{Rfail}(i-1, i+2)] \quad \forall i = 3, ..., n \\ \gamma &= 2M_n[\textsc{Rfail}(n, n) - \textsc{Rfail}(n-1, n-1)] + 2 \cdot \textsc{Rfail}(n-1, n-1) \end{aligned}$$

And let:

$$W_n = \max(\{\gamma\} \cup \{\alpha_i : 2 \leq i \leq n\} \cup \{\beta_i : 3 \leq i \leq n\}) + \epsilon$$

Define edge costs as follows:

$$\begin{aligned} h_{0i} &= M_n - 1 & \forall 0i \in E \\ h_{0i} &= M_n & \forall 0i \notin E \\ h_{ij} &= -W_n & \forall ij \in E \\ h_{ij} &= nW_n & \forall ij \notin E \end{aligned}$$

To prove the theorem, we will first show that the minimum cost route must only contain edges in G. We will then derive an upper bound UB_H on the cost of a route which is a Hamiltionian cycle in G, along with lower bounds LB_E and LB_Q on the costs of routes which are elementary (but not Hamiltonian in G) and routes which are non-elementary. We will then show that $UB_H < LB_E$ and $UB_H < LB_Q$.

We can use the function Rfail to provide bounds on EFC in the following manner. Since Rfail counts the total expected number of failures on the route, we have $\textsc{Rfail}(\mu_k, \sigma_k^2) = \sum_{i=1}^{k} \text{FAIL}_i(\mu_i, \sigma_i^2)$. Letting $h^+ = \max\{h_{0i} : 1 \leq i \leq k\}$ and $h^- =$

$\min\{d_{0i} : 1 \leq i \leq k\}$, we obtain the following bounds for EFC(r) using RFAIL:

$$\begin{aligned}
\mathrm{EFC}(r) &= \sum_{i=1}^{k} \mathrm{EFC}_i(\mu_i, \sigma_k^i) \\
&= \sum_{i=1}^{k} 2h_{0i}\, \mathrm{FAIL}_i(\mu_i, \sigma_i^2) \\
&\leq \sum_{i=1}^{k} 2h^+\, \mathrm{FAIL}_i(\mu_i, \sigma_i^2) \\
&= 2h^+ \sum_{i=1}^{k} \mathrm{FAIL}_i(\mu_i, \sigma_i^2) \\
&= 2h^+ \textsc{Rfail}(\mu_k, \sigma_k^2)
\end{aligned}$$

By a similar argument, $\mathrm{EFC}(r) \geq 2h^- \textsc{Rfail}(\mu_k, \sigma_k^2)$ as well. Thus:

$$2h^- \textsc{Rfail}(\mu_k, \sigma_k^2) \leq \mathrm{EFC}(r) \leq 2h^+ \textsc{Rfail}(\mu_k, \sigma_k^2).$$

Let v be a vertex such that $0v \in E$, and consider the route $r_v = (0, v, 0)$. This route has cost:

$$c_{r_v} = 2(M_n - 1) + 2(M_n - 1)\textsc{Rfail}(1, 1).$$

Now consider a route $r_u = (0, u_1, u_2, ..., u_k, 0)$ such that $u_i u_{i+1} \notin E$ for some $1 \leq i \leq k-1$. Since the capacity $Q = n + \frac{1}{2}$, we may assume $k \leq n$, as otherwise the route will not be feasible with respect to the capacity constraints. The cost of r_u is:

$$\begin{aligned}
c_{r_u} &= h_{0u_1} + h_{0u_k} + \sum_{i=1}^{k-1} h_{u_i u_{i+1}} + \mathrm{EFC}(R_u) \\
&\geq 2(M_n - 1) + (n - (k-2))W_n + 2(M_n - 1)\textsc{Rfail}(k, k) \\
&> 2(M_n - 1) + 2(M_n - 1)\textsc{Rfail}(1, 1) \\
&= c_{r_v}
\end{aligned}$$

Therefore we may restrict our attention to routes only containing edges in G.

Next we derive an upper bound on the cost of a route which is a Hamiltonian cycle in G. The cost of a route r_H which is a Hamiltonian cycle in G is:

$$c_{r_H} = \underbrace{2(M_n - 1) - (n-1)W_n}_{\text{deterministic travel cost}} + \underbrace{\mathrm{EFC}(r_H)}_{\text{recourse cost}}$$

As the maximum cost of any edge $0i$, $1 \leq i \leq n$, is M_n, we have $2M_n\text{RFAIL}(n,n) \geq \text{EFC}(r_H)$. Thus:

$$UB_H = 2(M_n - 1) - (n-1)W_n + 2M_n\text{RFAIL}(n,n)$$

is an upper bound for c_{R_H}.

Next we will obtain a lower bound LB_E on the cost of an elementary route which is not a Hamiltonian cycle in G. We will show that an elementary route containing $n-1$ non-depot vertices provides the worst-case lower bound by induction.

Let $r_{E_k} = (0, v_1, v_2, ..., v_k, 0)$, $1 \leq k < n-1$, be an elementary route. The cost of r_{E_k} is:

$$\begin{aligned} c_{r_{E_k}} &= h_{0v_1} + h_{0v_k} + \sum_{i=1}^{k-1} h_{v_iv_{i+1}} + \text{EFC}(r_{E_1}) \\ &\geq 2(M_n - 1) - (k-1)W_n + 2(M_n - 1)\text{RFAIL}(k,k) \end{aligned}$$

Let LB_{E_k} be this lower bound on $c_{r_{E_k}}$. Now consider a second elementary route $r_{E_{k+1}} = (0, u_1, ..., u_{k+1}, 0)$. The cost of $r_{E_{k+1}}$ is:

$$\begin{aligned} c(r_{E_{k+1}}) &= h_{0u_1} + h_{0u_{k+1}} + \sum_{i=1}^{k} h_{u_iu_{i+1}} + \text{EFC}(r_{E_2}) \\ &\geq 2(M_n - 1) - kW_n + 2(M_n - 1)\text{RFAIL}(k+1,k+1) \end{aligned}$$

Let $LB_{E_{k+1}}$ be the above lower bound on $c_{r_{E_{k+1}}}$. Then:

$$\begin{aligned} LB_{E_{k+1}} - LB_{E_k} &= 2(M_n - 1)[\text{RFAIL}(k+1,k+1) - \text{RFAIL}(k,k)] - W_n \\ &= \alpha_{k+1} - W_n \\ &< 0. \end{aligned}$$

Therefore $LB_{E_{k+1}} < LB_{E_k}$, and so $r_{E_{k+1}}$ provides the worst-case lower bound. By induction, $r_E = (0, v_1, ..., v_{n-1}, 0)$ is the non-Hamiltonian elementary route providing the strongest lower bound. Thus we set:

$$LB_E = 2(M_n - 1) - (n-2)W_n + 2(M_n - 1)\text{RFAIL}(n-1, n-1).$$

Next we consider non-elementary routes. Similar to the elementary case, we will prove by induction that a non-elementary route with n non-depot vertices provides the worst-case lower bound on the cost of non-elementary routes. We will consider routes $r_{Q_k} =$

$(0, v_1, v_2, ..., v_k, 0)$, $3 \le k < n$, and $r_{Q_{k+1}} = (0, u_1, u_2, ..., u_{k+1}, 0)$. Note that $k \ge 3$, as if $k = 1$ or 2 it is impossible for r_{Q_k} to be non-elementary. These have the lower bounds:

$$LB_{Q_k} = 2(M_n - 1) - (k-1)W_n + 2(M_n - 1)\textsc{Rfail}(k, k+3)$$
$$LB_{Q_{k+1}} = 2(M_n - 1) - kW_n + 2(M_n - 1)\textsc{Rfail}(k+1, k+4)$$

Again we have $LB_{Q_{k+1}} < LB_{Q_k}$:

$$\begin{aligned} LB_{Q_{k+1}} - LB_{Q_k} &= 2(M_n - 1)[\textsc{Rfail}(k+1, k+4) - \textsc{Rfail}(k, k+3)] - W_n \\ &= \beta_{k+1} - W_n \\ &< 0. \end{aligned}$$

Thus $LB_{Q_{k+1}} < LB_{Q_k}$, and so by induction the worst-case lower bound is obtained from a non-elementary route with n non-depot vertices, with corresponding lower bound:

$$LB_Q = 2(M_n - 1) - (n-1)W_n + 2(M_n - 1)\textsc{Rfail}(n, n+3)$$

Finally we will show that $UB_H < LB_E$ and $UB_H < LB_Q$. We will first look at the elementary case:

$$\begin{aligned} UB_H - LB_E &= 2(M_n - 1) - (n-1)W_n + 2M_n\textsc{Rfail}(n, n) - 2(M_n - 1) \\ &\quad + (n-2)W_n - 2(M_n - 1)\textsc{Rfail}(n-1, n-1) \\ &= 2M_n[\textsc{Rfail}(n, n) - \textsc{Rfail}(n-1, n-1)] + 2 \cdot \textsc{Rfail}(n-1, n-1) - W_n \\ &= \gamma - W_n \\ &\le -\epsilon \qquad\qquad < 0. \end{aligned}$$

Therefore $UB_H < LB_E$. Next we consider the non-elementary case:

$$\begin{aligned} UB_H - LB_Q &= 2(M_n - 1) - (n-1)W_n + 2M_n\textsc{Rfail}(n, n) - 2(M_n - 1) \\ &\quad + (n-1)W_n - 2(M_n - 1)\textsc{Rfail}(n, n+3) \\ &= 2M_n[\textsc{Rfail}(n, n) - \textsc{Rfail}(n, n+3)] + 2 \cdot \textsc{Rfail}(n, n+3) \\ &= 2\left(\frac{\textsc{Rfail}(n, n+3) + \epsilon}{\textsc{Rfail}(n, n+3) - \textsc{Rfail}(n, n)}\right)(-1)[\textsc{Rfail}(n, n+3) - \textsc{Rfail}(n, n)] \\ &\quad + 2 \cdot \textsc{Rfail}(n, n+3) \\ &= -2\epsilon \\ &< 0. \end{aligned}$$

Therefore $UB_H < LB_Q$.

Thus one can solve the Hamiltonian cycle problem in G by solving the minimum 2-stage feasible q-route problem in G' and checking whether the output solution has cost $\le UB_H$. □

Chapter 5

Concluding Remarks

In this thesis we identified some difficulties that stochastic demands introduce to the pricing problem in set partitioning formulations for the CVRPSD. Specifically, we show that even when customer demands are independent, vertices along q-routes may still have correlated demands. One issue that this introduces is the cumulative variance of a q-route may not be the sum of the variances of all the customers on the q-route. This in turn means that we cannot as easily reuse the dynamic programming algorithm for pricing q-routes in the CVRP if we wish to exactly price q-routes in the CVRPSD. However, the pricing algorithm for the CVRP still suffices if we only wish to perform approximate route pricing such as in the approach by Christiansen and Lysgaard [8], where the price of a q-route during column generation is not computed exactly, but with a heuristic approach which assumes the demand of repeated vertices is re-evaluated independently of previous visits along the q-route.

In addition, we proved that when the distribution is specified as a finite set of scenarios, solving the minimum 2-stage cost q-route pricing problem is strongly NP-hard. We also provide a proof that, under the assumptions that elementary operations can be performed in $O(1)$ time and the expected number of failures of q-routes can be computed in polynomial time with respect to the number of vertices n, there is a reduction from the minimum 2-stage cost q-route problem to the Hamiltonian cycle problem, and that this reduction is valid even when demands and variances are bounded by a constant. While this does not constitute a proof of strong NP-hardness, both due to the aforementioned assumptions as well as the construction in the proof requiring extremely large edge costs, it does suggest that pricing q-routes for the CVRPSD may be harder than for the CVRP.

Further questions on this topic would first be whether the minimum cost 2-stage q-

route pricing problem is strongly NP-hard in the case of independent normal demands. While our results suggest this may be case, further work needs to be done for a complete proof. To eliminate the need to compute $\textsc{Rfail}(\mu, \sigma^2)$, one approach could be to find easily computable upper and lower bounds for this function and alter the proof to use those. However there are challenges here both in finding bounds as well as in reworking the proof, due to the dependence of many of the parameters on Rfail. Since this would still not eliminate all of the issues, such as the huge edge weights M_n and W_n, it may be more productive to try to find a completely different approach to proving this complexity result. Furthermore, while the pricing problem we considered may prove to be hard, it may still be possible to develop efficient algorithms for pricing by exploiting the additional structure to the edge weights, as they are derived from the dual variables to a specific LP, and are not actually arbitrary real numbers.

Secondly, given that pricing is strongly NP-hard in the scenarios case, other avenues must be explored to solve this class of problems. One possibility is the use of Monte Carlo methods to verify route feasibility, an idea very recently applied to the vehicle routing with stochastic demands and probabilistic duration constraints by Florio, Hartl, Minner, and Salazar-González [18]. For that application, they also could not simplify the pricing problem by relaxing the feasible region, and this technique allowed them to efficiently price elementary routes. This could solve the similar issue for the CVRPSD where relaxing the feasible region does not necessarily simplify the pricing problem.

Another possibility worth exploring is combining the set-partitioning formulation for the CVRPSD with an edge-based formulation such as in Dinh et al. [12], where they combined edge-based and set partitioning formulations for the chance-constrained vehicle routing problem (CCVRP). Pricing q-routes for the CCVRP is also strongly NP-hard, which they addressed by further relaxing the constraints involved in defining the set. The constraints of edge-based formulations guide the solutions away from containing many cycles in the routes, but on the other hand struggle more with capacity constraints and finding strong lower bounds on the expected failure cost of routes. In contrast, set-partitioning formulations can incorporate the expected failure cost of routes directly in the objective function coefficients for variables, but the q-routes generated will tend to have more cycles due to the structure of the pricing problem. The two formulations therefore complement each other well when combined, each addressing the weaknesses of the other.

References

[1] Agostinho Agra, Marielle Christiansen, Rosa Figueiredo, Lars Magnus Hvattum, Michael Poss, and Cristina Requejo. The robust vehicle routing problem with time windows. *Computers & Operations Research*, 40(3):856–866, 2013.

[2] Cynthia Barnhart, Ellis L. Johnson, George L. Nemhauser, Martin W. P. Savelsbergh, and Pamela H. Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46(3):316–329, 1998.

[3] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

[4] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Robust sample average approximation. *Mathematical Programming*, 171(1-2):217–282, 2018.

[5] Dimitris J. Bertsimas. *Probabilistic combinatorial optimization problems*. PhD thesis, Massachusetts Institute of Technology, 1988.

[6] Dimitris J. Bertsimas. A vehicle routing problem with stochastic demand. *Operations Research*, 40(3):574–585, 1992.

[7] Alain Chabrier. Vehicle routing problem with elementary shortest path based column generation. *Computers & Operations Research*, 33(10):2972–2990, 2006.

[8] Christian H. Christiansen and Jens Lysgaard. A branch-cut-and-price algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research Letters*, 35:773–781, 2007.

[9] William Cook and Jennifer L. Rich. A parallel cutting-plane algorithm for the vehicle routing problem with time windows. Technical report, Rice University, 1999. Source: https://www.math.uwaterloo.ca/~bico/papers/vrptw.pdf.

[10] Jean-François Cordeau, Gilbert Laporte, Martin Savelsbergh, and Daniele Vigo. Vehicle routing. In Cynthia Barnhart and Gilbert Laporte, editors, *Transportation, handbooks in operations research and management science*, volume 14, pages 195–224. North Holland, 2007.

[11] George B. Dantzig and John H. Ramser. The truck dispatching problem. *Management Science*, 6(1):80–91, 1959.

[12] Thai Dinh, Ricardo Fukasawa, and James Luedtke. Exact algorithms for the chance-constrained vehicle routing problem. *Mathematical Programming*, 172(1-2):105–138, 2018.

[13] Moshe Dror. Note on the complexity of the shortest path models for column generation in vrptw. *Operations Research*, 42(5):977–978, 1994.

[14] Moshe Dror and Pierre Trudeau. Stochastic vehicle routing with modified savings algorithm. *European Journal of Operations Research*, 23(2):228–235, 1986.

[15] Alan L. Erera, Juan C. Morales, and Martin Savelsbergh. The vehicle routing problem with stochastic demand and duration constraints. *Transportation Science*, 44(4):474–492, 2010.

[16] Yao X. et al., editor. *Metaheuristics for the Vehicle Routing Problem with Stochastic Demands*, Berlin, Heidelberg, 2004. Springer.

[17] Dominique Feillet, Pierre Dejax, Michel Gendreau, and Cyrille Gueguen. An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems. *Networks*, 44(3):216–229, 2004.

[18] Alexandre M. Florio, Richard F. Hartl, Stefan Minner, and Juan-José Salazar-González. A branch-and-price algorithm for the vehicle routing problem with stochastic demands and probabilistic duration constraints. Technical report, University of Vienna, 2020. Source: http://www.optimization-online.org/DB_HTML/2020/03/7713.html.

[19] Zachary Friggstad and Chaitanya Swamy. Approximation algorithms for regret-bounded vehicle routing and applications to distance-constrained vehicle routing. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 744–753, New York, NY, USA, 2014. Association for Computing Machinery.

[20] Ricardo Fukasawa, Humberto Longo, Jens Lysgaard, Marcus Poggi de Aragão, Marcelo Reis, Eduardo Uchoa, and Renato F. Werneck. Robust branch-and-cut-and-price for the capacitated vehicle routing problem. *Mathematical Programming*, 108:491–511, 2006.

[21] Charles Gauvin, Guy Desaulniers, and Michel Gendreau. A branch-cut-and-price algorithm for the vehicle routing problem with stochastic demands. *Computers & Operations Research*, 50:141–153, 2014.

[22] Michel Gendreau, Gilbert Laport, and René Séguin. An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Science*, 29:143–155, 1995.

[23] Michel Gendreau, Gilbert Laport, and René Séguin. Stochastic vehicle routing. *European Journal of Operational Research*, 88:3–12, 1996.

[24] Curt Hjorring and John Holt. New optimality cuts for a single-vehicle stochastic routing problem. *Annals of Operations Research*, 86:569–584, 1999.

[25] Stefan Irnich and Guy Desaulniers. Shortest path problems with resource constraints. In Guy Desaulniers, Jacques Desrosiers, and Marius M. Solomon, editors, *Column Generation*, chapter 2, pages 33–65. Springer, Boston, MA, 2005.

[26] Gilbert Laporte. The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 59(3):345–358, 1992.

[27] Gilbert Laporte and François Louveaux. The Integer L-shaped method for stochastic integer programs with complete recourse. *Operation Research Letters*, 13:133–142, 1993.

[28] Gilbert Laporte, François Louveaux, and Hélène Mercure. Models and exact solutions for a class of stochastic location-routing problems. *European Journal of Operational Research*, 39:71–78, 1989.

[29] Gilbert Laporte, François Louveaux, and Luc Van Hamme. An Integer L-shaped method for the capacitated vehicle routing problem with stochastic demands. *Operations Research*, 50(3):415–423, 2002.

[30] Averill M. Law and W. David Kelton, editors. *Simulation Modeling and Analysis*. McGraw-Hill, Singapore, 2000.

[31] Hongtao Lei, Gilbert Laporte, and Bo Guo. The capacitated vehicle routing problem with stochastic demands and time windows. *Computers & Operations Research*, 38:1775–1783, 2011.

[32] Jorge E. Mendoza, Louis-Martin Rousseau, and Juan G. Villegas. A hybrid metaheuristic for the vehicle routing problem with stochastic demand and duration constraints. *Journal of Heuristics*, 22:539–566, 2015.

[33] Ian Post and Chaitanya Swamy. Linear programming-based approximation algorithms for multi-vehicle minimum latency problems: Extended abstract. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, page 512–531, USA, 2015. Society for Industrial and Applied Mathematics.

[34] Giovanni Righini and Matteo Salani. Symmetry helps: Bounded bi-directional dynamic programming for the elementary shortest path problem with resource constraints. *Discrete Optimization*, 3:255–273, 2006.

[35] Louis-Martin Rousseau, Michel Gendreau, Gilles Pesant, and Filippo Focacci. Solving vrptws with constraint programming based column generation. *Annals of Operations Research*, 130(1-4):199–216, 2004.

[36] Majid Salavati-Khoshghalb, Michel Gendreau, Ola Jabali, and Walter Rei. An exact algorithm to solve the vehicle routing problem with stochastic demands under an optimal restocking policy. *European Journal of Operations Research*, 273(1):175–189, 2019.

[37] William R. Stewart and Bruce L. Golden. Stochastic vehicle routing: a comprehensive approach. *European Journal of Operational Research*, 14:371–385, 1983.

[38] Paolo Toth and Daniele Vigo, editors. *The Vehicle Routing Problem.* Society for Industrial and Applied Mathematics, 2001.

[39] Bram Verweij, Shabbir Ahmed, Anton J. Kleywegt, George Nemhauser, and Alexander Shapiro. The sample average approximation method applied to stochastic routing problems: A computational study. *Computational Optimization and Applications*, 24(2):289–333, 2003.

[40] Wen-Huei Yang, Kamlesh Mathur, and Ronald H. Ballou. Stochastic vehicle routing problem with restocking. *Transportation Science*, 34(1):99–112, 2000.

[41] James R. Yee and Bruce L. Golden. A note on determining operating strategies for probabilistic vehicle routing. *Naval Research Logistics Quarterly*, 27(1):159–163, 1980.