Bachelor thesis BSC^2 Econometrics & Economics
Erasmus University Rotterdam - Erasmus School of Economics

Stochastic capacity constraints in the flow-refueling location problem with variable driving range

B.T.C. van Rossum
435111

Supervisor: T. Breugem
Second Assessor: Dr. L.E. Westerink-Duijzer

July 7, 2019

Abstract

Environmental concerns have led to rapid developments regarding the flow-refueling location problem (FRLP), the problem of positioning refueling stations in a road network as to maximise the volume of traffic that is able to successfully complete its trip. Though several stochastic versions of the FRLP have been proposed, increasing model realism by accounting for a variable driving range, so far no probabilistic models with capacitated stations exist. In this paper we argue why ordinary capacity constraints are overly restrictive in a stochastic setting, and derive a novel set of constraints that overcomes this issue by modelling the expected traffic volume that refuels at each station. As the resulting model is highly complex, an intuitive heuristic approach is presented. Numerical experiments on random networks suggest that using stochastic capacity constraints as compared to deterministic ones alters the optimal location choice and leads to substantial gains in covered traffic volume. The heuristic is shown to achieve near-optimal solutions in reduced time on small instances, whilst slightly outperforming the exact solution approach on large instances.

Contents

1. Introduction

Concerns regarding energy security, fuel economy, and environmental impact have led to the increased popularity of alternative-fuel vehicles (AFVs), i.e. vehicles that do not run on traditional petroleum fuels but on other sources of energy, such as hydrogen, natural gas, or electricity. While there are rapid developments in this area, most of these AFVs still offer a limited driving range and thus require to be refueled at alternative-fuel stations (AFSs), perhaps multiple times during a trip. To this day, however, the limited number of such stations leads to range anxiety among drivers and poses a main barrier to the widespread adoption of AFVs (World Economic Forum, 2018).

To overcome this issue, the construction of a carefully designed network of refueling facilities is required, allowing drivers to reach their destinations without running out of fuel. Since current investment in these facilities is still scarce, AFSs can only be built at a small number of locations out of a large candidate set. Since an optimal positioning of facilities requires coordination, non-sequential decision-making is preferred over the sequential positioning of facilities. In the case when multiple stakeholders are involved, for example, a joint location procedure is needed to prevent cannibalization of demand or inferior outcomes due to path-dependency.

One of the first mathematical formulations of this problem is the flow-refueling location model (FRLM) developed by Kuby & Lim (2005), aimed at locating p refueling facilities in a road network as to maximise the total volume of traffic that is able to successfully complete their round trip. This traffic is then represented by a set of flows between origin-destination (O-D) pairs, each flow moving along the shortest path connecting the O-D nodes and having a specific flow volume. This problem has received considerable attention in recent years (see Ko et al. (2017) for an overview), and model realism has been significantly improved, e.g. by the incorporation of capacitated facilities and by allowing drivers to deviate from their predefined shortest paths in order to pass by additional refueling stations.

Relatively little attention has been paid, however, to the stochastic nature of the vehicle driving range. It is well known that the driving range of AFVs is highly variable, as it depends on uncertain factors as the weather, the driving style of the driver, the level of road congestion, and the temperature (Ehsani et al., 2018). Two recently proposed models, the expected flow-refueling location problem and chance-constrainted flow-refueling location problem proposed by de Vries & Duijzer (2017) and later reformulated by Boujelben & Gicquel (2019), do successfully incorporate this uncertainty, yet do not include any of the aforementioned model developments. One of the goals of this paper is therefore to bridge this gap in literature by incorporating capacitated facilities in these existing stochastic models. Moreover, we will show how the capacity constraints commonly used in literature yield an overly restricted model when working in a stochastic setting, and solve this issue by introducing the notion of stochastic capacity and embedding it into a novel set of constraints. As the resulting model is highly complex, an efficient two-stage heuristic approach to this model is presented, its first stage focusing on positioning facilities and its second stage aiming to attribute traffic to these stations.

In extensive computational experiments we find that replacing the restrictive deterministic constraints

by their stochastic counterparts yield significant gains in terms of flow volume covered, gains which partly arise due to a different positioning of facilities. Moreover, due to its efficient distribution of running time, the heuristic is able to find near-optimal solutions in reduced running time on small instances, and outperform an exact solution approach on larger ones. Finally, we find that relaxing one of the key assumptions behind our capacitated model does not have major consequences.

The remainder of this paper is structured as follows. Section 2 presents a comprehensive literature review on the FRLM and discusses the contributions of this paper. Section 3 formally states the flow-refueling location problem and presents the deterministic flow-refueling location problem. In Section 4 the expected flow-refueling location problem and chance-constrained flow-refueling location problem are introduced, whereas Section 5 presents alternative formulations of these models that are more appropriate when incorporating capacity constraints. The notion of stochastic capacity and a corresponding novel set of constraints is discussed in Section 6. All models are extensively evaluated in numerical experiments in Section 7. Section 8 concludes and presents directions for future work.

2. Literature Review

2.1. The FRLM and its Extensions

The literature stream on flow-refueling location models originated in the early 90s, when Hodgson (1990) and Berman et al. (1992) proposed the first flow-capturing location models (FCLM). As compared to set covering location models, which assume that demand in a network is located at its nodes, they realised that demand is sometimes exerted by traffic flows. A significant fraction of customers of discretionary service facilities, such as convenience stores, gasoline stations, or billboards, consumes this service on an otherwise pre-planned trip, e.g. the daily commute to and from work. The goal of the FCLM is to maximise the flow between a set of OD pairs that is captured or intercepted by facilities located at nodes.

Kuby & Lim (2005) extended this logic to the problem of refueling alternative-fuel vehicles along their trips, and were the first to formulate the flow-refueling location problem. Solving this FRLM required pre-generating all possible combinations of locations capable of serving each round trip, limiting its application to small instances. To overcome this issue, Kuby et al. (2009) proposed a greedy heuristic and a genetic algorithm that did not require this pre-processing, and solved the FRLM for the large Florida state-wide road network.

Capar & Kuby (2012) proposed a reformulation of the FRLM that directly incorporates the logic of combinations of locations refuelling a trip, thereby eliminating the costly pre-generation of these combinations. The arc cover-path cover formulation of Capar et al. (2013) uses the idea that a round trip is feasible when each of its arcs is within driving range of a refueling station preceding it, and that there is a set of locations capable of refueling each single arc. This model is even more compact and allows for solving the large aforementioned Florida instance to optimality within a minute. Another alternative formulation was proposed by MirHassani

& Ebrazi (2012), who employ mass balance constraints to model the flow from origin to destination, and show that this also leads to a huge reduction in solution time.

Whereas the FRLM does not allow drivers to deviate from their pre-defined paths in order to refuel, the deviation-flow refueling location model (DFRLM) does incorporate this flexibility. It was initially proposed by Kim & Kuby (2012), who imposed a maximum deviation distance from the shortest path and penalised the flow volume of deviation routes to ensure that trips without deviation are preferred. The symmetric deviation routes were then generated using the k-shortest path (k-SP) algorithm by Hoffman & Pavley (1959), including routes that contain loops since those might pass by an additional refueling station and thereby enable feasibility of a round trip. As the huge number of possible deviation routes led to large solution times even on small networks, Kim & Kuby (2013) developed a network transformation heuristic for the DFRLM. The heuristic selects station in a greedy fashion, iteratively adding the ASS capable of refueling the highest additional flow volume. This volume is determined by finding the shortest feasible paths on the network of OD nodes and refueling stations, thereby allowing for deviation routing without explicitly pre-generating all possible deviation routes.

The FRLM has also been extended to include capacitated facilities. The capacitated flow-refueling location model (CFRLM) initially proposed by Upchurch et al. (2009) relaxes the assumption of unlimited refueling capacity, and limits the vehicle volume that can be refueled by a single station. To this extent, the binary trip variables in the formulation of Kuby et al. (2009) are modified to continuous ones, allowing OD trips to be partially covered. Deviation routing can also be included in the CFRLM, leading to the capacitated deviation-flow refueling location model (CDFRLM) (Hosseini et al., 2017; Hosseini & MirHassani, 2017).

Finally, several authors have attempted to incorporate stochasticity in the FRLM, either by including uncertainty in traffic flow volume or by assuming variability in the driving range. Hosseini & MirHassani (2015) propose a two-stage model which first locates permanent facilities based on uncertain flow forecasts, and subsequently locates portable stations to tune to realised demand. Wu & Sioshansi (2017) develop the stochastic flow-capturing location model (SFCLM), which aims to maximise the expected flow volume covered. Yıldız et al. (2019) formulate a model with capacitated facilities and stochastic demand, and solve this using a tailored branch and cut algorithm. Lee & Han (2017) assume a probabilistic travel range and locate stations such that the expected number of electric vehicles successfully reaching their destination is maximised. Since the resulting objective function is nonlinear and non-convex, a Benders decomposition with column generation techniques is used as solution approach. de Vries & Duijzer (2017) resolve this problem by making several assumptions on the driving range distributions, and propose the expected flow-refueling location problem (EFRLP) and chance-constrained flow-refueling location problem (CCFRLP). The former aims to maximise expected flow volume covered, whereas the goal of the latter is to maximise flow volume when a flow is only considered covered if the probability of successfully completing the round trip is above a certain threshold. Boujelben & Gicquel (2019) build on their work by reformulating the EFRLP and CCFRLP in a more efficient manner, using the arc cover-path cover idea by Capar et al. (2013), and propose

a fast tabu-search heuristic. As we wish to incorporate capacity constraints in models with stochastic driving range, the aforementioned EFRLP and CCFRLP will form the starting point of this paper.

The large number of possible deviation paths or capacity constraints in the extensions of the FRLM lead to large computation times when using commercial solvers. Limiting this solution approach to small- and medium-sized instances, it has led to the development of several heuristic algorithms. Greedy heuristics are the most popular choice, of which versions with and without neighbourhood searches exist (Lim & Kuby, 2010; Li & Huang, 2014; You & Hsieh, 2014; Kuby et al., 2009). Genetic algorithms are also widespread, often outperforming greedy approaches (Lim & Kuby, 2010; You & Hsieh, 2014; Kuby et al., 2009). It is unclear, however, how these methods can be applied to capacitated models. An alternative approach uses an LP relaxation of the FRLM to find a promising set of facility locations, after which facility nodes are selected from this set in an iterative fashion. This idea has been applied to the ordinary FRLM (Tran et al., 2018), as well as the CDFRLM (Hosseini et al., 2017; Hosseini & MirHassani, 2017).

While the ideas underlying the above methods are relatively straightforward, more sophisticated techniques exist. Both Arslan & Karaşan (2016) and Lee & Han (2017) employ a Benders decomposition approach, yet have not extended its application beyond the regular FRLM. Finally, the branch-and-price algorithm developed by Yıldız et al. (2016) is able to find near-optimal solutions for the DFRLM in reasonable time.

2.2. Contributions

The existing models with stochastic driving range, as proposed by de Vries & Duijzer (2017) and reformulated by Boujelben & Gicquel (2019), rely on the assumption of uncapacitated facilities. The first contribution of this paper is therefore to formulate ordinary capacity constraints for these models, limiting the vehicle volume that can recharge at each facility. Second, we argue how such ordinary capacity constraints are overly restrictive when applied in a stochastic setting. To overcome this issue, we introduce the notion of stochastic capacity and propose a novel set of constraints that incorporate this stochastic capacity in the flow-refueling location problem. Third, as these constraints significantly increase model complexity, we present a heuristic approach that is able to quickly find efficient solutions. Finally, we investigate the impact of the single-path assumption, i.e. the assumption that each vehicle along a given flow follows the same refuelling pattern, which is commonly used in capacitated models.

3. Problem Definition and Formulation

We consider a road network $G = (L, E)$, consisting of a set of locations L and a set of edges E. The set of locations L is defined as the union of the set of potential facility locations K, the set of origins O, and the set of destinations D. A collection of drivers wishes to repeatedly perform symmetric round trips on this network, and we define a flow to be the subset of drivers that travels from the same origin to the same destination. The set F denotes all such flows, and we use O_f and D_f to refer to the origin and destination of flow $f \in F$, respectively. Since our focus is on capacity constraints, for simplicity we assume that there is

no deviation routing and that drivers thus travel along the shortest path connecting origin and destination. Drivers travelling along flow f traverse an ordered sequence of edges $e \in E$, connecting the start location O_f, an ordered set of potential facility locations $K_f \subseteq K$, and the end location D_f, and then pass these edges in opposite order on the way back. The union of all nodes traversed along flow f is denoted by L_f. Moreover, v_f denotes the volume of this flow, i.e. the number of drivers wishing to perform this round trip. As common in literature, we assume drivers start at the origin with half a tank, and that a full tank offers a finite driving range R. Since round trips are repeatedly performed, drivers thus need to recharge or refuel at a charging facility at least once during their trip. More specifically, drivers are only able to complete their trip when the distance between two subsequent facilities along their round trip does not exceed the vehicle range R, and when the facility first passed on the trip is located within a a radius of $R/2$ from the origin. We say flow f is covered when drivers along f can travel from their origin to their destination and back without running out of fuel. The goal of the FRLP is then to locate p charging stations in K such that the total volume of all covered flows is maximised.

Since we assume that drivers perform round trips, we can equivalently see each trip as a cycle connecting the origin and destination. Moreover, each cycle can be said to consist of one or more cycle segments, where a cycle segment is defined as a part of the round trip that a driver traverses without being able to refuel his vehicle, and that cannot be contained in a larger cycle segment. In case no charging stations are located along flow f, the cycle corresponding to f contains only one cycle segment, its distance equalling that of the full round trip. Each additional charging station located on $k \in K_f$ then breaks up one or more existing cycle segments into smaller sub-segments. The positioning of charging stations in the road network can therefore be seen as the partitioning of all cycles into cycle segments such that the volume of flows where no cycle segment has a distance larger than R is maximised. In other words, the use of cycle segments allows us to explicitly incorporate the driving range into a formulation of the FRLP and will subsequently enable the inclusion of a stochastic driving range.

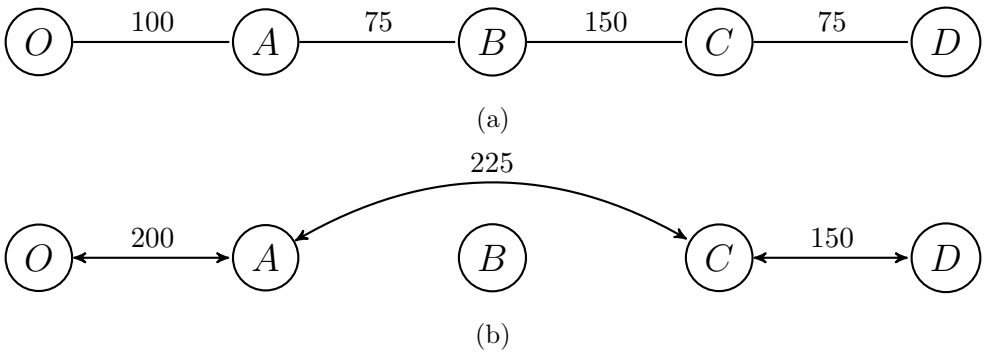

Figure 1: Illustration of cycle segments.

Let us illustrate the concept of cycle segments using the simple single flow depicted in Figure 1a. This flow starts at the origin O, passes by potential facility locations A, B, and C, and reaches destination D.

The distance between each pair of successive nodes is noted on the edge connecting them. We now consider the case where charging facilities are located at both A and C. The arrows in Figure 1b show the cycle segments originating in this situation. These double arrows reflect the fact that each segment is passed in both directions along a trip, once passing from O to A, and vice versa. The segment defined by O and A has length 200, twice the distance between these nodes since a vehicle is unable to recharge at O and we assume that it starts at O with only half a tank. The same logic holds for the segment connecting C and D, whereas the distance of the segment linking A and C is simply the distance of the (shortest) path between them. Since no facility is assigned to B and it is neither an origin nor a destination, there is no cycle segment defined by B. Clearly, B is still passed twice on each round trip.

To formalise the idea of cycle segments, let us first recall that each round trip is symmetric, i.e. the same locations are passed in reverse order when travelling from origin to destination as when travelling from destination to origin. This fact allows us to characterise a cycle segment of flow f by two nodes k and l in L_f, where k is defined to be the starting node of the segment and l to be the node closest to D_f along this cycle segment. Note that by definition $k \in O_f \cup K_f$, whereas $l \in K_f \cup D_f$. The distance t_{kl} of the cycle segment defined by k and l follows directly from this simple characterisation: it is the distance of the shortest path from k to l. Since round trips are symmetric and cycle segments are defined as parts of the cycle where a vehicle cannot recharge, this distance is multiplied by two in case the segment starts or ends at the origin or destination, respectively, and equal to $M \gg R$ in case it directly connects origin and destination. As this latter value largely exceeds the driving range, a vehicle in a covered flow will never be able to travel directly from origin and destination, and is therefore required to recharge at least once along its path. Moreover, note that the multiplication by two enforces the half-tank assumption mentioned earlier, and that as a result drivers return to their starting point with at least half a tank, enabling them to repeatedly preform their round trip.

Using the above notion of cycle segments, a novel formulation of the FRLP can be stated that explicitly incorporates the driving range. To this extent we define the binary variable x_k, which equals one if a facility is placed at potential location k, and zero otherwise. Moreover, the binary variable y_f indicates whether flow f is covered, equalling one if this is the case, and zero otherwise. Finally, the binary variable i_{klf} equals one when k and l identify a cycle segment in f, and zero otherwise. Defining A_{kf} to be the set of locations passed when travelling from k to D_f, the FRLP can be formulated as the following mixed integer linear program

(MILP):

$$\max \quad \sum_{f \in F} v_f y_f \tag{1}$$

$$\text{s.t.} \quad \sum_{k \in K} x_k = p \tag{2}$$

$$\sum_{l \in L_f} i_{klf} t_{kl} - (1 - y_f) M \leq R \qquad \forall f \in F, k \in \{O_f \cup K_f\} \tag{3}$$

$$\sum_{l \in L_f} i_{klf} = x_k \qquad \forall f \in F, k \in K_f \tag{4}$$

$$\sum_{l \in L_f} i_{klf} = 1 \qquad \forall f \in F, k \in O_f \tag{5}$$

$$\sum_{k \in L_f} i_{klf} = x_k \qquad \forall f \in F, l \in K_f \tag{6}$$

$$\sum_{k \in L_f} i_{klf} = 1 \qquad \forall f \in F, l \in D_f \tag{7}$$

$$i_{klf} \in [0, 1] \qquad \forall f \in F, k \in L_f, l \in A_{kf} \tag{8}$$

$$x_k \in \mathbb{B} \qquad k \in K \tag{9}$$

$$y_f \in \mathbb{B} \qquad \forall f \in F \tag{10}$$

The objective function (1) maximises total covered flow volume, whereas constraint (2) ensures that exactly p facilities are built. The range constraints in (3) guarantee that a flow is only covered when none of its cycle segments exceed the driving range R. Constraints (4) - (7) enforce the partitioning of the flow cycle into segments according to the positioning of facilities. Note that the domain of these cycle segments is continuous in (8), yet de Vries et al. (2014) prove that constraints (4) - (7) ensure that $i_{klf} = 1$ if and only if locations k and l identify a cycle segment of flow f and that $i_{klf} = 0$ otherwise. Since the driving range in (1) - (10) is assumed to be deterministic, we will refer to the above model as the deterministic flow-refueling location problem (DFRLP).

4. Stochastic Models

Whereas so far we have treated the driving range to be deterministic, in reality the distance a vehicle can travel on a single tank depends on a collection of random factors. The weather, driving style of the driver, level of road congestion, and temperature may all affect the actual range (Ehsani et al., 2018). In the following, we provide a formal setting to model this stochasticity, and in Section 4.1 and Section 4.2 we extend this approach to present two new formulations of the FRLP as proposed by de Vries & Duijzer (2017).

Without loss of generality, the driving range of a vehicle can be said to be a function $R(\omega)$ of a random variable ω, which represents the aggregated effect of all factors influencing the driving range. We make the following two assumptions regarding $R(\omega)$:

Assumption 1. *Given a realisation ω, the driving range $R(\omega)$ is the same at each cycle segment traversed by a vehicle in flow f.*

Assumption 2. *$R(\omega)$ is randomly distributed with cumulative density function (CDF) $\mathcal{G} : \mathbb{R}_+ \mapsto [0, 1]$.*

Assumption 1 is a rather strong assumption, yet as we will see later on it allows us to avoid non-linearities when formulating the stochastic versions of the FRLP. Moreover, it can be justified by arguing that many factors affecting the driving range, such as temperature and driving style, do not vary strongly along a flow. Assumption 2 has a purely notational purpose. More specifically, it allows us to compute the probability $g_{kl} = \mathcal{G}(t_{kl})$ that a vehicle runs out of fuel when travelling along the cycle segment connecting k and l.

As a consequence of the variable driving range, the notion of flow coverage now becomes a probabilistic one. A driver can be said to be able to complete his round trip only with a certain probability, depending on the interaction between the distance of the cycle segments he traverses and the distribution of the range $\mathcal{G}$. In this setting, the binary representation of coverage used in the DFRLP becomes flawed, and hence we propose two alternatives that do take into account the stochastic nature of the driving range. In Section 4.1 we model the problem of maximising the expected flow volume covered, whereas in Section 4.2 we formulate a model where a flow is considered covered only if the probability of running out of fuel along the corresponding round trip is below a certain threshold.

4.1. Expected Flow-Refueling Location Problem (EFRLP)

In this stochastic setting, it is natural to consider the problem of maximising the expected flow volume that is covered. To this extent, let us denote by z_f the probability that a vehicle along flow f can complete its round trip. Under Assumption 1, this probability is equal to the probability that the driving range exceeds the distance of the largest cycle segment of flow f. More formally, we have:

$$z_f = 1 - \max_{(k \in L_f,\ l \in A_{kf})} \{i_{klf} \mathcal{G}(t_{kl})\} \tag{11}$$

Note that $\mathcal{G}(t_{O_f D_f}) \approx 1$ as we set $t_{O_f D_f}$ equal to a very large positive value M. We thus obtain the desired result that z_f equals zero when no charging stations are located along a flow and the cycle segment connecting origin and destination is activated. The objective function of expected flow volume is now obtained by multiplying the volume of each flow with its respective coverage probability, i.e. it is equal to $\sum_{f \in F} v_f z_f$. Substituting the new objective function and range constraints into the DFRLP then yields the following model:

$$\max \quad \sum_{f \in F} v_f z_f \tag{12}$$

$$\text{s.t.} \quad (2), (4) - (9)$$

$$z_f \leq 1 - \sum_{l \in L_f} i_{klf} g_{kl} \qquad \forall f \in F, k \in \{O_f \cup K_f\} \tag{13}$$

$$z_f \geq 0 \qquad \forall f \in F \tag{14}$$

where constraint (13) enforces (11). We will refer to model (12) - (14) as the expected flow-refueling location problem (EFRLP).

4.2. Chance-Constrained Flow-Refueling Location Problem

Whereas the EFRLP does incorporate a stochastic driving range, it may lead to solutions where the probability of being able to complete a round trip is below acceptable levels for certain flows. In practice, however, drivers are unlikely to undertake a trip when it is uncertain whether they will reach their destination. Therefore, we now consider the problem of maximising covered flow volume when the probability of running out of fuel on a covered flow must be below a certain threshold α. This condition can be stated as follows:

$$\max_{(k \in L_f,\ l \in A_{kf})} \{i_{klf} \mathcal{G}(t_{kl})\} \leq \begin{cases} \alpha, & \text{if } y_f = 1 \\ 1, & \text{otherwise} \end{cases} \tag{15}$$

Adapting the DFRLP using the above condition yields the following formulation:

$$\max \quad \sum_{f \in F} v_f y_f \tag{16}$$

$$\text{s.t.} \quad (2), (4) - (10)$$

$$\sum_{l \in L_f} i_{klf} g_{kl} \leq \alpha + (1 - y_f) \qquad \forall f \in F, k \in \{O_f \cup K_f\} \tag{17}$$

where the constraints in (17) guarantee that a flow is only considered covered when the probabilistic threshold is not exceeded by any of its cycle segments and thereby enforce (15). We will refer to the above model as the chance-constrained flow-refueling location problem (CCFRLP).

5. Alternative Formulations for the DFRLP, EFRLP, and CCFRLP

The aim of this paper is to incorporate stochastic capacity constraints in the EFRLP and CCFRLP. Clearly, modelling any type of capacity constraint requires the allocation of facilities to flows, i.e. we want to be able to retrieve the vehicle volume that a certain facility covers. In the current formulation, such an assignment is not made. Facilities are not attributed to specific flows, as a result of which cycle segments may be unnecessarily short and vehicles may recharge more often than is desirable in a capacitated set-up. Put differently, it might be impossible to attain the optimal distribution of facility capacity to flows. To overcome this problem, we now present the alternative formulations of the DFRLP, EFRLP, and CCFRLP as proposed by Boujelben & Gicquel (2019) that do allow for such an assignment. Since their formulation is more efficient, we obtain the additional benefit of reduced solution times, as the numerical experiments in Section 7.2 will confirm.

The novel formulations by Boujelben & Gicquel (2019) rely on the node cover-path cover principle, i.e. the observation that a flow is covered if and only if each node along its path is within driving range of a charging facility located on one of the nodes preceding it. In this sense, they very closely resemble the highly

efficient arc cover-path cover formulation first proposed by Capar et al. (2013). Since in this model the location of a facility along a flow does not necessarily imply a node of this flow being assigned to this facility, it is well suited to our goal of incorporating capacity constraints.

To gain some intuition behind this new approach, let us reconsider the example flow of Figure 1. Assuming that facilities are located at nodes A and C, the arrows in Figure 2 show a feasible assignment of nodes to stations preceding them. Though A is assigned to the origin, this simply implies that a driver recharges at A after leaving O. B and C are both assigned to A, indicating that these are within driving range of A and are reached with the tank that is recharged at A. Finally, as D is assigned to C, a driver recharges at C before proceeding to D. Note that the inter-node distances are identical to the ones used in the cycle segment formulation, once again in order to enforce the half-tank assumption.

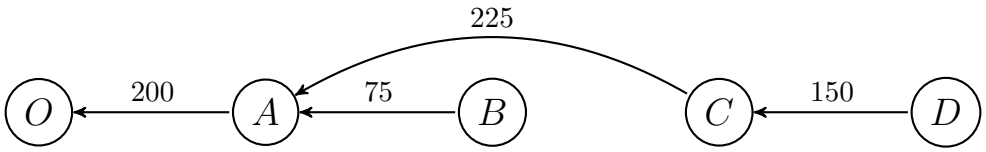

Figure 2: Illustration of the node cover-path cover principle.

Building on the previously used nomenclature, let B_{lf} be the set of locations passed before l when traversing along flow f. Moreover, we introduce the binary variable w_{klf}, which is equal to one if node l along flow f is assigned to a facility located at k, and zero otherwise. The DFRLP can then be reformulated as follows:

$$\max \quad \sum_{f \in F} v_f y_f \tag{18}$$

$$\text{s.t.} \quad \sum_{k \in K} x_k = p \tag{19}$$

$$\sum_{k \in B_{lf}} w_{klf} = y_f \qquad \forall f \in F, l \in L_f \setminus \{O_f\} \tag{20}$$

$$w_{klf} \leq x_k \qquad \forall f \in F, k \in L_f \setminus \{O_f, D_f\} \tag{21}$$

$$\sum_{k \in B_{lf}} w_{klf} t_{kl} \leq R \qquad \forall f \in F, l \in L_f \setminus \{O_f\} \tag{22}$$

$$w_{klf} \in \mathbb{B} \qquad \forall f \in F, l \in L_f \setminus \{O_f\}, k \in B_{lf} \tag{23}$$

$$x_k \in \mathbb{B} \qquad \forall k \in K \tag{24}$$

$$y_f \in \mathbb{B} \qquad \forall f \in F \tag{25}$$

The objective in (18) maximises flow volume covered, and constraint (19) ensures exactly p facilities are located. The constraints in (20) guarantee that each node along a covered flow is assigned to exactly one preceding facility, while those in (21) make sure this is only done when the corresponding facility is actually activated. The range constraints in (22) limit the distance between a node and its assigned facility to the driv-

ing range R, and the remaining constraints define the domain of the binary variables. To distinguish between the new formulation and the previously shown DFRLP, we will refer to this model as the CoverDFRLP.

In a similar fashion, the following model represent the CoverEFRLP:

$$\max \quad \sum_{f \in F} v_f z_f \tag{26}$$

$$\text{s.t.} \quad (19), (21), (23) - (24)$$

$$z_f \leq \sum_{k \in B_{lf}} w_{klf} \qquad \forall f \in F, l \in L_f \setminus \{O_f\} \tag{27}$$

$$z_f \leq 1 - \sum_{k \in B_{lf}} w_{klf} g_{kl} \qquad \forall f \in F, l \in L_f \setminus \{O_f\} \tag{28}$$

$$z_f \in [0, 1] \qquad f \in F \tag{29}$$

Finally, the CoverCCFRLP is given by the model below:

$$\max \quad \sum_{f \in F} v_f y_f \tag{30}$$

$$\text{s.t.} \quad (19) - (21), (23) - (25)$$

$$\sum_{k \in B_{lf}} w_{klf} g_{kl} \leq \alpha \qquad \forall f \in F, l \in L_f \setminus \{O_f\} \tag{31}$$

As indicated before, Boujelben & Gicquel (2019) find that the CoverDFRLP, -EFRLP, and -CCFRLP are solved much more efficiently by the commercial solver CPLEX than their cycle segment-based counterparts. In case of the deterministic model, they attribute this speed-up to the removal of the big-M constraints in (3). Leading to a tighter linear relaxation, the integrality gap is reduced and less branch-and-bound nodes need to be explored. In case of the stochastic models, the authors argue that the additional constraints linking flow coverage variables and assignment variables, i.e. (27) in case of the CoverEFRLP and (21) in case of the CoverCCFRLP, are responsible for the tighter formulation.

6. Stochastic Capacity

6.1. Stochastic Capacity Constraints

Before we introduce the concept of stochastic capacity, it is helpful to consider how ordinary capacity constraints can be modelled. Capacitated versions of the FRLP are plentiful (Upchurch et al., 2009; Kuby et al., 2009; Hosseini et al., 2017; Hosseini & MirHassani, 2017), and incorporating capacity constraints in the node cover-path cover formulations introduced above is straightforward. If we let C be the vehicle volume that a single facility can cover, and a_{kf} denote the location directly succeeding k along the flow f path, this can be done by including the following set of constraints:

$$\sum_{f \in F} w_{ka_{kf}f} v_f \leq x_k C \qquad \forall k \in K \tag{32}$$

It is not restricting to assume that, if a facility is used along a covered flow, the successor node is assigned to recharge at this facility. More specifically, assigning the successor node to recharge at this facility will not reduce leftover capacity, since it was already a charging station along this flow, and can only increase the coverage probability, since it is the station closest to its successor. As a result, the objective function cannot decrease and the assumption is not limiting. This implies that it is valid to say $w_{ka_{kf}f} = 1$ when facility k is used to recharge vehicles on flow f, and hence we use this variable to measure demanded capacity at each station.

Moreover, it is now clear why the cycle segment formulations, i.e. the DFRLP, EFRLP, and CCFRLP, are not suitable when modelling capacitated stations. Since constraints (4) and (6) imply that vehicles recharge at each activated station they pass along their path, stations are more likely to reach their maximum capacity, and as a result only a fraction of the optimal flow volume might be covered. In contrast, the possibility to assign stations to specific flows in the cover formulations yields the required flexibility. Note, however, that the cover formulations still make use of the single-path assumption, i.e. the assumption that all drivers along a single flow have the same charging pattern. As this can lead to sub-optimal outcomes in capacitated models, we investigate a relaxation of this assumption in Section 7.4.

While this approach is sufficient when working with a deterministic driving range, it is overly strict when assuming a stochastic driving range. If we allow for vehicles not to complete their round trip with a certain probability, this implies that there is a chance they might not reach one or several of their allocated charging stations. The capacity constraints in (32), however, implicitly assume that each charging station along the flow is actually reached and used. This leads to an overestimation of required capacity, an overrestrictive model, and possibly erroneous location decisions. Though the overestimation rate is bound by α and may thereby be limited in the CCFRLP, it could be quite substantial in the EFRLP.

Our aim is thus to refine the ordinary capacity constraints above, and account for the fact that a proportion of vehicles may not arrive at their assigned charging locations. In other words, we aim to switch from deterministic capacity to the notion of stochastic capacity, i.e. the realisation that the traffic volume refueling at a facility is the result of a stochastic process and should therefore be modelled as such. To this extent, let us define d_{kf} to be the expected fraction of drivers along flow f that recharges at facility location k on their trip from O_f to D_f, and similarly let e_{kf} be the expected fraction of drivers along flow f recharing at facility location k on the trip back from D_f to O_f. Though there are no stations located at either the origin or destination, we will see that it is necessary for their arrival probabilities to be well-defined, and hence d_{kf} and e_{kf} exist for all $k \in L_f$.

Before proceeding to the formulas governing these probabilities, let us gain some insight into the mechanics of stochastic capacities by reconsidering the example of Figure 2. Once again assuming that facilities are located at A and C, Figure 3 shows the same cover assignment as before, the arrows indicating to which facility nodes are assigned. Moreover, we assume that the probabilities of running out of fuel on the four trajectories of length 200, 75, 225, and 150 are given by 0.15, 0.00, 0.25, and 0.05, respectively.

We can now retrieve the probability d_k of recharging at node k on the way from O to D, omitting the subscript f for simplicity. As a vehicle cannot run out of fuel before arriving at the origin, d_O equals one. The probability of arriving at A is equal to the probability that the driving range exceeds 200, i.e. $d_A = 1 - 0.15 = 0.85$. Since no station is present at B, $d_B = 0$. The case of station C is more elaborate, as a vehicle could be stranded before arriving at C if i) it never reached the station preceding C, or ii) it ran out of fuel on the way from the station preceding C to C. Under Assumption 1, however, we find that d_C equals the minimum of these two odds. In other words, since there is a single realisation ω, the relevant probability is the one corresponding to running out of fuel on the longest distance travelled so far. We thus have that $d_C = \min\{d_A, 1 - 0.25\} = 0.75$. Extending this logic to D and observing that the distance from C to D is shorter than the longest distance travelled so far, we find that $d_D = d_C = 0.75$.

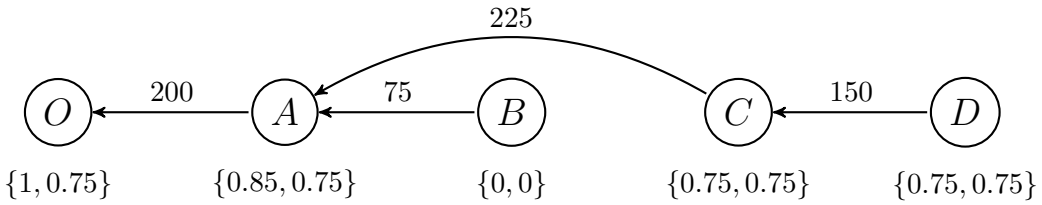

Figure 3: Illustration of stochastic capacity, $\{d_k, e_k\}$ indicated below each location.

Computing the probability e_k of refuelling at station k on the way from D to O is trivial due to Assumption 1: since the trips are symmetric, a vehicle will not run out of fuel on the way back if it has not done so yet. So if a station at k was used along the trip to the destination, we find $e_k = d_D$. If this station is not active along this particular flow, or there is no station located at k, we have $e_k = 0$. These cases are summarised by the final expression $e_k = \min\{d_k, d_D\}$, using the fact that d_k is only smaller than d_D when it equals zero. We see that this yields $e_B = 0$ as it should, since no station is located at B, whereas values of 0.75 are attained at all other nodes. Moreover, it is now clear that arrival probabilities for origin and destination nodes were defined in order to facilitate the recursive relations described above.

The main conclusion from the example in Figure 3 is that d_{kf} should equal the minimum of the probability of reaching the station preceding k, and the probability of crossing the segment connecting this station with k. Special cases are the origin O_f, which is always reached, and nodes that have no stations that are assigned nodes of flow f, for which the probability ought to be zero. The following expression formalises these conditions:

$$d_{kf} = \begin{cases} 0, & \text{if } w_{ka_{kf}f} = 0 \\ 1, & \text{if } k = O_f \\ \min\limits_{\{l: w_{lkf}=1\}} \{d_{lf}, 1 - g_{lk}\}, & \text{otherwise} \end{cases} \qquad \forall f \in F, k \in L_f \tag{33}$$

Note that we use the condition $w_{lkf} = 1$ to identify the station l to which k is assigned. As seen in the

example of Figure 3, computing the value of e_{kf} can be done in the following straightforward manner:

$$e_{kf} = \min\{d_{kf}, d_{D_f f}\} \qquad \forall f \in F, k \in L_f \tag{34}$$

Equation (34) reflects the fact that a driver will not recharge at k on the way back to O_f if he did not do so way on his way to D_f, and that no more drivers will run out of fuel on their way back. The following set of constraints can then be used to incorporate the above expressions for stochastic capacity in the CoverEFRLP and -CCFRLP:

$$d_{kf} \geq d_{lf} - q_{klf} - (1 - w_{ka_{fkf}}) \qquad \forall f \in F, k \in L_f \setminus \{O_f\}, l \in B_{kf} \tag{35}$$

$$d_{kf} \geq 1 - w_{lkf} g_{lk} - (1 - q_{klf}) - (1 - w_{ka_{fkf}}) \qquad \forall f \in F, k \in L_f \setminus \{O_f\}, l \in B_{kf} \tag{36}$$

$$e_{kf} \geq d_{kf} - r_{klf} \qquad \forall f \in F, k \in L_f \tag{37}$$

$$e_{kf} \geq d_{D_f f} - (1 - r_{klf}) \qquad \forall f \in F, k \in L_f \tag{38}$$

$$\sum_{f \in F} (d_{kf} + e_{kf}) v_f \leq x_k C \qquad \forall k \in K \tag{39}$$

$$d_{O_f f} = 1 \qquad \forall f \in F \tag{40}$$

$$d_{kf}, e_{kf} \in [0, 1] \qquad \forall f \in F, k \in L_f \tag{41}$$

$$q_{klf}, r_{klf} \in \mathbb{B} \qquad \forall f \in F, k \in L_f, l \in B_{kf} \tag{42}$$

Note that (33) and (34) are represented by constraints (35) - (36) and (37) - (38), respectively. In both cases we use auxiliary binary variables to model the minimum function. Since our objective is to maximise covered flow volume and capacity is a limiting factor, these constraints ensure that indeed the minimum value is attained in case capacity is scarce. Since the station l to which node k is assigned is not known in advance, in (35) - (36) we are forced to apply the last condition of (33) to all preceding nodes, using the term $1 - w_{ka_{fkf}}$ to activate them when necessary. The constraints in (39) then enforce the actual capacity constraints.

As mentioned before, the benefits of considering stochastic capacity as opposed to ordinary capacity are especially relevant in the context of maximising expected flow volume. In the chance-constrained FRLP, the benefits are bounded by the threshold parameter α, which is why we will focus on the EFRLP. We will refer to the CoverEFRLP in combination with constraints (35) - (35) as the stochastically-capacitated cover expected flow-refueling location problem (SC-CoverEFRLP). Moreover, we will refer to the CoverEFRLP in combination with constraint (32) as the deterministically-capacitated cover flow-refueling location problem (DC-CoverEFRLP).

6.2. A Heuristic Approach to the SC-CoverEFRLP

The large number of binary variables in the SC-CoverEFRLP, in combination with the pairwise constraints modelling the minimum functions, leads to a model that cannot be solved efficiently by commercial solvers. To overcome this issue, we now propose a simple two-stage heuristic that is able to find near-optimal solutions

to the SC-CoverEFRLP in reasonable time. The rationale behind the heuristic is that stochastic capacity constraints increase the flow volume that charging stations can handle. Since the coverage levels z_f are well below one for most flows, the use of stochastic capacity leads to a more efficient use of refuelling capacity. Alternatively, one could say it virtually increases the capacity C of each station. This implies that the facilities activated in the optimal solution to the SC-CoverEFRLP might deviate from those in the optimal solution to the DC-CoverEFRLP solved for the same level of C, yet are less likely to deviate from those in the solution to the DC-CoverEFRLP solved for a (slightly) augmented capacity level.

We make use of this idea by, in the first stage of the heuristic, solving a regular capacited CoverEFRLP, i.e. the CoverEFRLP in combination with the deterministic capacity constraints of (32), for a capacity of $(1+\delta)C$, where δ is a non-negative parameter indicating the virtual increase. This yields a set of selected facilities K^*_δ, whose objective value in the SC-CoverEFRLP we evaluate in the second stage of the heuristic. The idea is that, by selecting a suitable value of δ, this set of facilities is close to the optimal solution to the SC-CoverEFRLP yet is obtained in a more efficient manner. Since finding the optimal assignment of flows to the facilities in K^*_δ is still computationally demanding, the stage-two running time might need to be limited in practice. Note that the set of flows obtained in the first stage cannot be used to provide a warm start in the second stage, since the factor δ might be an overestimate of the actual gains due to stochastic capacity.

7. Numerical Experiments

In this section we perform numerical experiments to analyse the performance of all models presented above. Section 7.1 provides a description of the problem instances and test settings that are used. Section 7.2 investigates the impact of incorporating stochasticity in the FRLP by comparing the DFRLP, EFRLP, and CCFRLP. Moreover, it evaluates the computational advantage of the cover formulations. Section 7.3 then analyses the effect of using stochastic capacity as compared to deterministic capacity in the SC-CoverEFRLP, and evaluates the performance of our heuristic. Finally, Section 7.4 considers the impact of the single-path assumption in capacitated models. All problem instances were solved on a computer with two 2.2 GHz cores and 8 GB memory, using CPLEX 12.6.1 as the solver engine.

7.1. Test Instances

In line with de Vries & Duijzer (2017), all tests are run on instances that are randomly generated using an adapted version of the method used by Capar & Kuby (2012). Section A of the Appendix provides a detailed description of this procedure. We refer to a randomly generated network with X potential facility locations and Y origin-destination nodes as $sXwY$. See Table 1 for an overview of all four generated instances. The parameter settings are as follows. We use a deterministic driving range R of 250, while in all stochastic models we assume $R(\omega) \sim \Gamma(50, 5)$, such that $\mathbb{E}_\omega[R(\omega)] = 250$ in accordance with Capar & Kuby (2012). The coverage threshold of the chance-constrained problem is set to $\alpha = 0.05$.

Table 1: Characteristics of the random test instances.

Instance name	$\|K\|$	$\|F\|$	T	V
s100w50	100	1200	748	10^6
s80w40	80	756	627	10^6
s60w30	60	420	827	10^6
s40w20	40	185	691	10^6

Note: T is the average travel distance from origin to destination, V the total vehicle volume

7.2. Impact of Stochasticity in the FRLP

We now analyse the impact of incorporating stochasticity in the flow-refueling location problem. To measure the benefit of using a stochastic model, we use the so-called value of the stochastic solution (VSS): the gain obtained by using the optimal solution to a stochastic model over using the optimal deterministic allocation. The VSS is then defined as the difference in objective value when evaluating both sets of facilities in a stochastic setting. This value will be zero when the deterministic facility allocation performs as well as the optimal stochastic solution, and above zero when it is outperformed by this stochastic solution. Likewise, the value of the stochastic solution can be seen as the error made when omitting stochasticity from the location problem. In addition to the VSS, we also make a comparison between the EFRLP and CCFRLP solutions.

Table 2: VSS and optimality gaps in terms of expected and chance-constrained flow volume covered for the EFRLP, CCFRLP, and DFRLP solutions.

	Opt. gap in terms of expected flow volume covered (%)			Opt. gap in terms of chance constr. flow volume covered (%)			
Instance	EFRLP solution	CCFRLP solution	DFRLP solution	EFRLP solution	CCFRLP solution	DFRLP solution	VSS
s100w50	0.00	22.08	6.22	37.76	0.00	47.76	42,940
s80w40	0.00	16.92	3.50	43.11	0.00	62.80	61,399
s60w30	0.00	32.46	4.22	38.88	0.00	46.28	50,930
s40w20	0.00	16.79	1.42	3.96	0.00	9.06	25,226

Table 2 shows the optimality gaps, averaged over all values of $p \in \{1, 2, 3, 4, 5, 10, 15, 20, 25\}$, in terms of expected and chance-constrained flow volume covered of the EFRLP, CCFRLP, and DFRLP solutions, as well as VSS averaged over the two stochastic objectives. When evaluating expected flow volume covered, clearly we find that the optimality gap of the EFRLP solution equals zero. The DFRLP solution performs reasonably well, which might be explained by the fact that the range constraint in (3) limits the cycle segment distance and thereby the coverage probability. In contrast, the optimality gaps of the CCFRLP solution are relatively large. Apparently, the strict threshold parameter α implies that facilities be located to enable high coverage levels on a selected set of flows, yet this conservative approach is far from optimal in a setting where

flows with low coverage levels also contribute to the objective function.

A different story is portrayed when considering the chance-constrained flow volume covered. Clearly, the CCFRLP solution attains the optimal solution, yet now we find that both the EFRLP and DFRLP solutions perform poorly. The latter two did not apply strict coverage criteria when positioning facilities, and hence cover only a small fraction of flows in the chance-constrained manner. The EFRLP solution performs slightly better than the DFRLP, indicating that it is better to take into account some form of stochasticity than none at all. Moreover, flows that attain coverage of over $1 - \alpha$ also contribute significantly to the objective of the EFRLP, and hence we can expect more of such flows to be present in the EFRLP solution. Finally, for all four instances we find considerable levels of the VSS, showing that, averaged over all levels of p, as much as 61,399 additional flow volume, i.e. 6.1% of the maximum coverage, can be covered when considering stochasticity in the decision-making process. The value of the stochastic solution is largely driven by the chance-constrained problem, since the optimality gaps of the DFRLP are especially large in terms of chance-constrained flow volume.

The results for instance $s80w40$ displayed in Figure 4 confirm the observations above, as the difference between the three solutions is larger in terms of chance-constrained flow volume than in terms of expected flow volume. In addition, we find that the objective in the CCFRLP is generally lower than that in the EFRLP, which can be attributed to the tighter constraints on covered flows. Moreover, we find that the EFRLP and DFRLP solutions cover (nearly) no chance-constrained flow when only a small number of facilities is to be located. In other words, incorporating stochasticity in the location problem is especially relevant when resources are scarce.

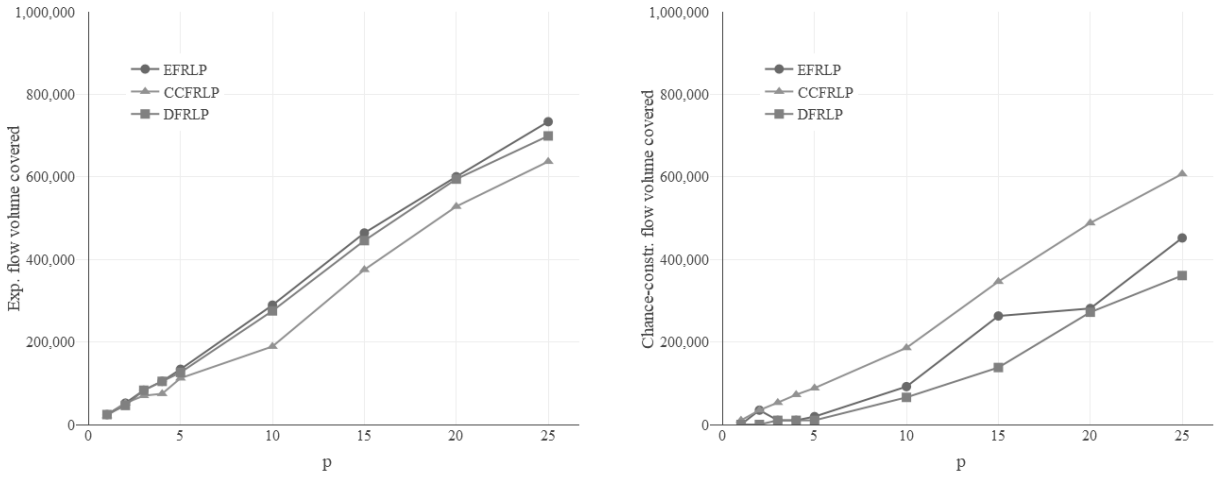

(a) Expected flow volume covered

(b) Chance-constrained flow volume covered

Figure 4: Expected and chance-constrained flow volume covered for the DFRLP, EFRLP, and CCFRLP solutions to the instance $s80w40$.

As indicated before, the rationale behind switching to the cover formulations is that they i) are more

appropriate when modelling capacitated facilities and ii) are significantly more efficient. To support the latter statement, Table 3 displays the computation times of both formulations for all three location problems on the instance $s80w40$. The cover formulations outperform the cycle segment formulations in all three models, though the largest differences in performance are observed in the DFRLP and CCFRLP. Moreover, the solution times of the cover formulations seem to be less sensitive to increases in problem size, as measured by p, as are the cycle segment formulations.

Table 3: CPU times (s) on instance $s80w40$.

	DFRLP		EFRLP		CCFRLP	
p	Segment	Cover	Segment	Cover	Segment	Cover
1	6	1	4	15	5	1
2	55	1	34	51	14	1
3	177	1	66	54	28	1
4	246	1	126	78	80	1
5	341	1	176	70	102	1
10	517	3	264	79	135	1
15	380	2	225	79	139	2
20	279	2	230	107	131	2
25	273	2	172	61	159	2
Average	**253**	**2**	**144**	**66**	**88**	**1**

7.3. Impact of Stochastic Capacity

In order to analyse the impact of stochastic capacity, we now solve the DC-CoverEFRLP and SC-CoverEFRLP on instance $s40w20$ for various capacity levels $C \in \{50, 100, 150, 200, 250\} \times 1,000$. Moreover, we test the heuristic solution approach to the SC-CoverEFRLP presented in Section 6.2 using a range of parameters $\delta \in \{0,\ 0.125,\ 0.250\}$. To enable a wide range of analysis and investigate the practical relevance of our models, we limit all solution times to one hour. As experience suggests that solving the DC-CoverEFRLP is more efficient than finding the optimal set of flows for a given set of facilities in the SC-CoverEFRLP, we attribute fifteen minutes of computation time to the first stage of the heuristic and fifteen minutes to the second. The total solution time of the heuristic thus never exceeds one hour, yielding a fair comparison between the exact and heuristic approach. Finally, we provide the facilities selected in the DC-CoverEFRLP as a warm start to the exact SC-CoverEFRLP solver. Complete results of this experiment are given in Table B.5 - Table B.9 of the Appendix.

Before turning to the gains of stochastic capacity, we consider solutions of the deterministically capaci-

tated model in more detail to get an idea of the role of capacity in the EFRLP. Figure 5 shows the expected flow volume covered for a range of values of p and three capacity levels. As covered flow volume increases significantly more when increasing C from 50,000 to 150,000 than when increasing capacity by an additional 100,000, we find that capacity is a limiting factor especially for lower values. Furthermore, we find evidence that capacity constraints mainly act on certain key nodes in the network. Firstly, increasing p from 20 to 25 barely affects covered flow volume. Secondly, the objective function never attains that of the uncapacitated EFRLP (see Table B.4). These observations suggest that capacity constraints limit the role of certain key nodes in the network, nodes that would normally facilitate a majority of the covered flows. Given the tree-like structure of the network (see Section A), this is not surprising.

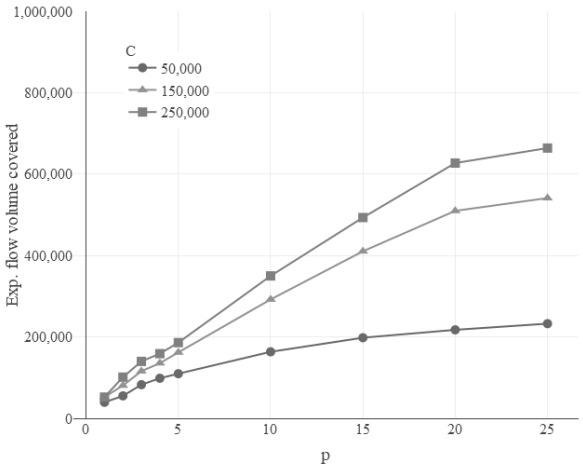

Figure 5: Expected flow volume covered in the DC-CoverEFRLP on instance $s40w20$ for various levels of C.

We now consider the gains in terms of expected flow volume covered when adhering to stochastic rather than deterministic capacity constraints. The first conclusion that can be drawn based on Table B.5 - Table B.9 is that the SC-CoverEFRLP is indeed a large model that cannot be solved efficiently by commercial solvers. As a result, an optimal solution is often not attained within the time limit of an hour, especially for the larger instances where $p \geq 10$. On problem instances where this is the case, we find that optimality gaps above 1% are not uncommon. So despite the fact that the SC-CoverEFRLP is a relaxed version of the DC-CoverEFRLP, we thus find that the objective function of the former does not always exceed that of the latter. In other words, the results for $p \leq 5$ are valid and shed a light on the gains of using stochastic capacity constraints. The results for $p \geq 10$ cannot be used to argue that these gains are non-existent for larger facility numbers, but merely show the huge difference in effiency between the SC-CoverEFRLP and DC-CoverEFRLP.

Figure 6a shows the gains in terms of expected flow volume covered of stochastic capacity as compared to deterministic capacity for various values of p, averaged over all tested capacity levels. The above conclusion is supported, as we find substantial gains of the order of 2.5% when only a few facilities are located, whereas

these gains decrease and even turn negative when more facilities are placed and optimal solutions are not attained anymore. Special attention need be paid to the case where $p = 1$, as here we find no benefits. This is not surprising, as one expects shorter routes to be selected when only a single facility is placed, and the odds of vehicles running out of fuel along such shorter routes are unlikely to be significant. Nonetheless, when facilities are relatively scarce we find that using stochastic capacity constraints can strongly boost the covered flow volume.

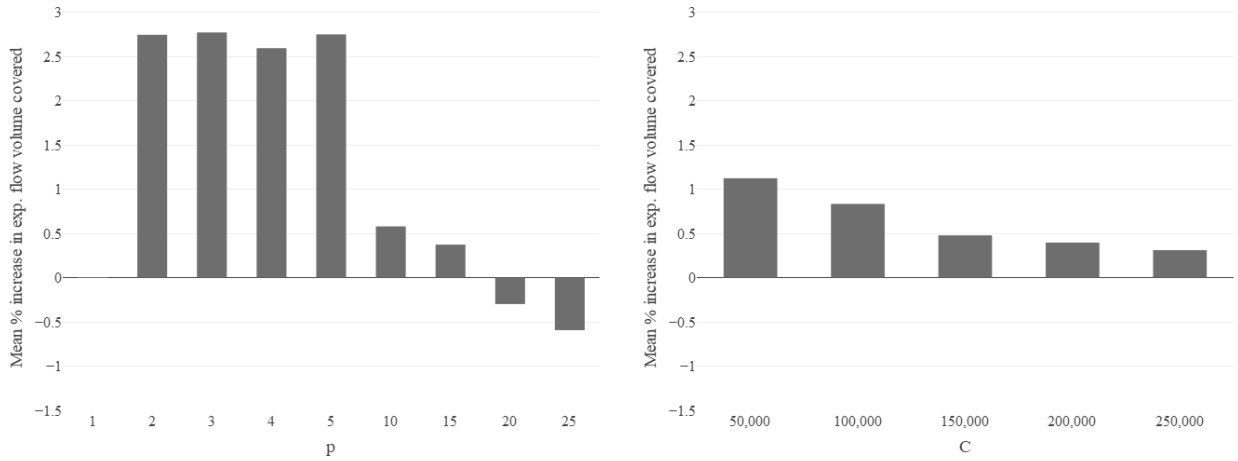

(a) Gains for various facility numbers, averaged over capacity levels

(b) Gains for various capacity levels, averaged over facility numbers

Figure 6: Gains in terms of expected flow volume covered of using stochastic as compared to deterministic capacity constraints on instance $s40w20$.

Similarly, Figure 6b shows the shows the gains in terms of expected flow volume covered for various capacity levels, averaged over all values of p. Firstly, as these values are the averages of both small and large problem instances, we find they are slightly less striking and never exceed 1.5%. Second, we spot a clear trend in the capacity levels, the gains in objective value flattening off as facilities are able to handle larger volumes. This is in accordance with the evidence of Figure 6a, and shows that the gains of stochastic capacity are bigger when capacity is a scarce factor.

We now turn to the performance of the heuristic presented in Section 6.2, which has been tested on the same instances as the exact solution approaches for values of $\delta \in \{0,\ 0.125,\ 0.250\}$. Once again, Figure 7 displays the gains of these solutions in terms of expected flow volume covered as compared to the DC-CoverEFRLP solution. The general pattern of Figure 6 is repeated as i) for large values of p, no improvements upon the deterministic case can be found within the time limit, and ii) the gains become less pronounced as the capacity increases. A new insight is revealed, however, when comparing the heuristic results for $\delta = 0$ with those of the exact solution approach. Looking at the volume increases for $p \leq 5$, values for which both methods achieve their optimal solutions within the time limits, we find that the exact approach outperforms

the heuristic one. Since the facilities selected by the heuristic for $\delta = 0$ correspond to those of the optimal DC-CoverEFRLP solution, this implies that stations are located at different nodes under stochastic capacity constraints than under deterministic ones. In other words, the gains from adhering to stochastic capacity do not purely arise because the covered flows are re-optimised, but originate partly in alternative location choices.

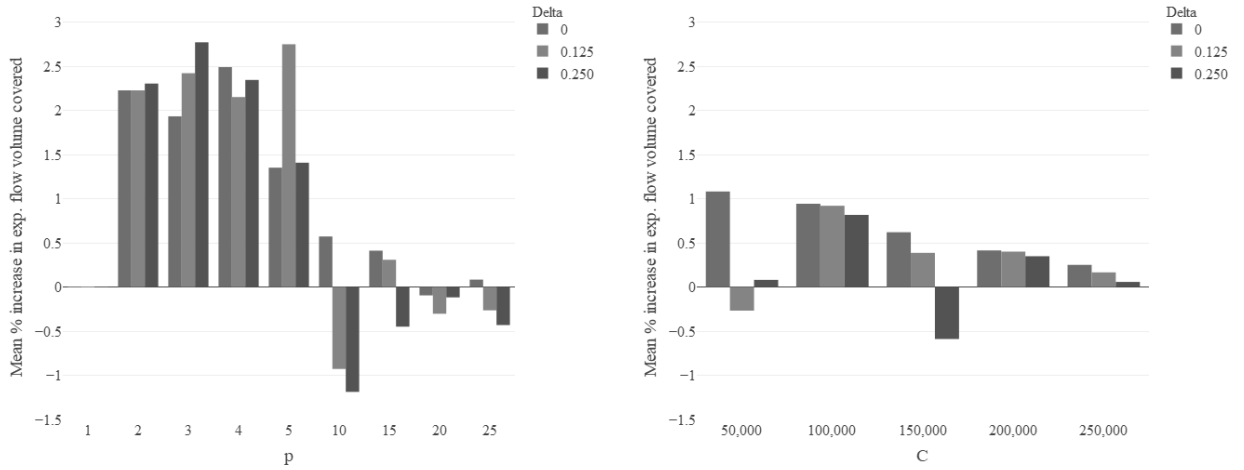

(a) Gains for various facility numbers, averaged over capacity levels

(b) Gains for various capacity levels, averaged over facility numbers

Figure 7: Gains in terms of expected flow volume covered of using stochastic as compared to deterministic capacity constraints on instance $s40w20$ when using heuristic solutions.

Continuing the comparison of the exact and heuristic approach, we find that for $p \leq 5$ the heuristic performance closely tracks that of the exact approach. So despite being a simplification, the idea of virtual capacity underlying the heuristic provides a fruitful basis for tackling the SC-CoverEFRLP. This is even more apparent when analysing the result for larger problem sizes, as we find the exact approach being outperformed by the heuristic multiple times when $p \geq 10$. For $p = 25$, for example, the heuristic with $\delta = 0$ is equal to find an increase upon the deterministic model, whereas the exact approach fails to do so. Clearly, the heuristic provides an efficient distribution of running time between finding facilities and attributing flows to these facilities, whereas the exact solution method attempts the extremely complex task of solving both simultaneously. In case the exact approach would not have been supplied with a warm start, the above analysis would have been even more in favour of our heuristic. The role of the parameter δ in the heuristic performance remains unclear. Whereas $\delta = 0$ seems to provide the best overall results, the difference with $\delta = 0.125$ are of such small magnitude that additional research is desired. Perhaps an intermediate value is most capable of capturing the gains of stochastic capacity in critical scenarios, whilst avoiding to provide an overoptimistic estimate of its benefits.

We conclude the section on stochastic capacity by comparing the solution times of the exact solution

approach with those of the heuristic. Figure 8a displays the solution times of each approach for various levels of p, averaged over all capacity levels. The solution time of the exact approach seems to grow exponentially until it reaches its one hour time limit when $p \geq 10$, in line with the earlier observation of negative increases in flow volume covered when no optimal solution was found. It is not surprising that the CPU time grows so rapidly, given that there are $\frac{\binom{40}{10}}{\binom{40}{5}} \approx 1,288$ times more possible facility combinations for $p = 10$ than for $p = 5$, and that the set of flows which can then be covered grows as well. Based on Figure 8a, we can also note that the performance gains of the heuristic are two-fold. First of all, for $p \leq 10$ we see that the solution times of the heuristic are but a fraction of those of the exact approach, while we have seen before that the solution quality barely differed between the two. This holds for all tested values of δ. Second, for $p \geq 15$ we find that the reduction in solution time is minimal, yet that the heuristic was able to outperform the exact solution approach in terms of flow volume covered. We should note that, since the solution times of the heuristic for $p \geq 15$ hover around forty-five minutes yet never reach one hour, it is likely that the time limit of the second stage was a limiting factor and that performance can be further improved by shifting computation time from stage one to stage two. Analysing the effect of capacity on solution time in Figure 8b, we find a discernible trend among the exact nor heuristic approaches, indicating that the facility capacity is not an indicator of the toughness of a problem.

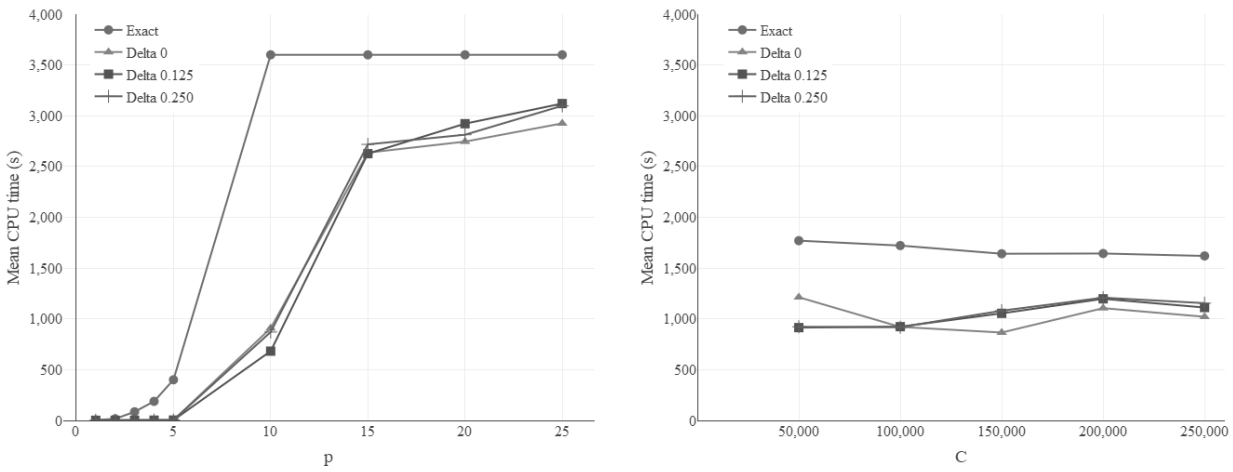

(a) CPU times for various facility numbers, averaged over capacity levels

(b) CPU times for various capacity levels, averaged over facility numbers

Figure 8: CPU times of various solution approaches to the SC-CoverEFRLP on instance $s40w20$.

7.4. Sensitivity to the Simple-Path Assumption

As indicated in Section 6.1, throughout this paper we assume that all vehicles along a given flow have identical charging patterns. This need not be optimal, as coverage in capacitated models might increase when assigning different stations to vehicles along a single flow. As the single-path assumption, in combination

with Assumption 1, enables a stochastic model without non-linearities, we cannot relax this assumption by altering the model. We can, however, obtain the same result by slightly adjusting the set of flows F. If we were to split up each flow f into q identical copies, each having a flow volume that is equal to a fraction $\frac{1}{q}$ of the original, we allow drivers along each flow to follow at most q different charging patterns. Doing this for sufficiently large flow-split factors q is equivalent to relaxing the single-path assumption.

We now solve the DC-CoverEFRLP for all factors $q \in \{2, 5, 10\}$ and a capacity $C = 150 \times 1,000$ on instance $s40w20$. Figure 9 shows the result of this analyis, displaying the gains in covered flow volume as compared to the case $q = 1$. While there is a large gap at $p = 2$, this is likely to be specific to the tested instance, as this magnitude is not observed for any other number of facilities. In contrast, for all other levels of p these gains do not exceed 1%, and increasing q from 5 to 10 barely affects this. The single-path assumption thus does not seem to have much impact on the results of capacitated models. We can, however, expect that the benefits of stochastic capacity are larger when the assumption is relaxed, as small capacity gains are more likely to be exploited.

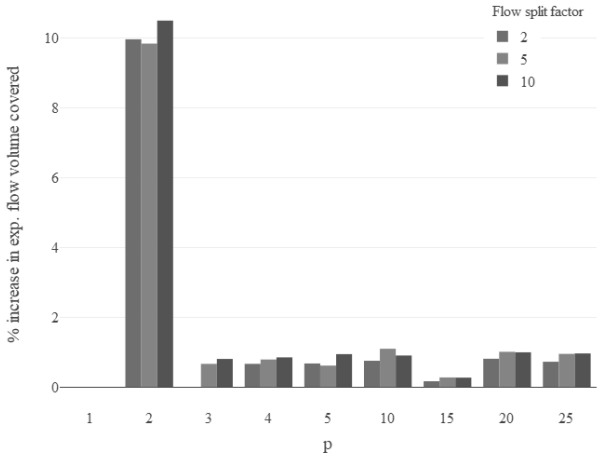

Figure 9: Percentual increase in expected flow volume covered due to flow split factors in the DC-CoverEFRLP on instance $s40w20$ for $C = 150,000$.

8. Conclusions

Existing models on the flow-refueling location problem with variable driving range, such as the EFRLP and CCFRLP, are built on the assumption of uncapacitated refueling stations. In this paper, we propose extended formulations of these models that incorporate capacitated facilities, and show that doing so is only possible when using the arc cover-path cover formulations proposed by Boujelben & Gicquel (2019). An additional benefit of using their formulation is its high efficiency. Building on the probabilistic nature of the problem, we argue that ordinary capacity constraints yield an overly restricted problem, and propose a novel set of stochastic capacity constraints that limit the expected number of vehicles that can recharge

at a given facility. As the resulting SC-CoverEFRLP cannot be solved efficiently, we propose a two-stage heuristic solution approach which relies on the insight that stochastic capacity constraints are a relaxed form of deterministic ones, and thereby virtually increase the maximum capacity of a facility.

In extensive numerical experiments we confirm the relevance of stochastic models, showing that a large additional flow volume is covered when using stochastic solutions rather than deterministic ones. The efficient arc cover-path cover formulation is used to explore the role of capacity in the EFRLP, suggesting the that capacity constraints mainly act by limiting the functioning of several key nodes in the network, nodes that would regularly facilitate the majority of vehicle flows. When investigating the effect of stochastic capacity constraints, we focus on the EFRLP as we expect their role to be limited in the chance-constrained problem. We find that significant increases in objective value arise when modelling capacity constraints in a stochastic rather than deterministic manner. These gains are not merely due to a re-optimisation of covered flows, as the facility location decision is altered after adding the new layer of stochasticity. The notion of stochastic capacity is shown to be especially relevant in case capacity is scarce, i.e. when a small number of facilities is to be located or the capacity of each facility is small. Though further tuning of the parameter δ is required, the heuristic is able to closely track the exact approach on small instances while using only a fraction of its solution time. When a larger set of facilities is to be located, these computational advantages lessen, yet due to a more efficient distribution of running time the heuristic slightly outperforms the exact approach. Relaxing the simple-path assumption yields negligible changes to the objective value under deterministic capacity constraints, though we expect these changes to be more pronounced in the fully stochastic model.

The current experiments are still limited by the complexity of the SC-CoverEFRLP, as in a majority of cases no optimal solution could be obtained within the time limits. This asks for a more efficient solution approach, either in the form of a more sophisticated heuristic, or possibly a Benders decomposition as in Arslan & Karaşan (2016) or Lee & Han (2017). It may also be worthwhile to explore how the tabu-search of Boujelben & Gicquel (2019) can be adapted to capacitated models, since this is currently the only solution approach tailored to the EFRLP and CCFRLP. Another direction for future research could be to investigate the interaction between stochastic capacity and deviation routing, as the positive effect of deviation on coverage is likely to be strengthened under stochastic capacity constraints (and vice versa). When considering capacitated facilities, it is also natural to model the resulting queues at charging stations, as done by Jung et al. (2014). Since both of the above extensions significantly increase model complexity, a feasible implementation strongly relies on the development of a tailored novel solution approach. Similar computational concerns limit the feasibility of incorporating stochastic capacity constraints in the nonlinear model of Lee & Han (2017) where Assumption 1 is relaxed, though we can expect the gains of these constraints to be even more pronounced in such a setting.

References

Arslan, O., & Karaşan, O. E. (2016). A benders decomposition approach for the charging station location problem with plug-in hybrid electric vehicles. *Transportation Research Part B: Methodological*, *93*, 670–695.

Berman, O., Larson, R. C., & Fouska, N. (1992). Optimal location of discretionary service facilities. *Transportation Science*, *26*, 201–211.

Boujelben, M. K., & Gicquel, C. (2019). Efficient solution approaches for locating electric vehicle fast charging stations under driving range uncertainty. *Computers & Operations Research*, .

Capar, I., & Kuby, M. (2012). An efficient formulation of the flow refueling location model for alternative-fuel stations. *IIE Transactions*, *44*, 622–636.

Capar, I., Kuby, M., Leon, V. J., & Tsai, Y.-J. (2013). An arc cover–path-cover formulation and strategic analysis of alternative-fuel station locations. *European Journal of Operational Research*, *227*, 142–151.

Ehsani, M., Gao, Y., Longo, S., & Ebrahimi, K. (2018). *Modern electric, hybrid electric, and fuel cell vehicles*. CRC press.

Hodgson, M. J. (1990). A flow-capturing location-allocation model. *Geographical Analysis*, *22*, 270–279.

Hoffman, W., & Pavley, R. (1959). A method for the solution of the n th best path problem. *Journal of the ACM (JACM)*, *6*, 506–514.

Hosseini, M., & MirHassani, S. (2015). Refueling-station location problem under uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, *84*, 101–116.

Hosseini, M., & MirHassani, S. (2017). A heuristic algorithm for optimal location of flow-refueling capacitated stations. *International Transactions in Operational Research*, *24*, 1377–1403.

Hosseini, M., MirHassani, S., & Hooshmand, F. (2017). Deviation-flow refueling location problem with capacitated facilities: Model and algorithm. *Transportation Research Part D: Transport and Environment*, *54*, 269–281.

Jung, J., Chow, J. Y., Jayakrishnan, R., & Park, J. Y. (2014). Stochastic dynamic itinerary interception refueling location problem with queue delay for electric taxi charging stations. *Transportation Research Part C: Emerging Technologies*, *40*, 123–142.

Kim, J.-G., & Kuby, M. (2012). The deviation-flow refueling location model for optimizing a network of refueling stations. *international journal of hydrogen energy*, *37*, 5406–5420.

Kim, J.-G., & Kuby, M. (2013). A network transformation heuristic approach for the deviation flow refueling location model. *Computers & Operations Research*, *40*, 1122–1131.

Ko, J., Gim, T.-H. T., & Guensler, R. (2017). Locating refuelling stations for alternative fuel vehicles: a review on models and applications. *Transport Reviews*, *37*, 551–570.

Kuby, M., & Lim, S. (2005). The flow-refueling location problem for alternative-fuel vehicles. *Socio-Economic Planning Sciences*, *39*, 125–145.

Kuby, M., Lines, L., Schultz, R., Xie, Z., Kim, J.-G., & Lim, S. (2009). Optimization of hydrogen stations in florida using the flow-refueling location model. *International journal of hydrogen energy*, *34*, 6045–6064.

Lee, C., & Han, J. (2017). Benders-and-price approach for electric vehicle charging station location problem under probabilistic travel range. *Transportation Research Part B: Methodological*, *106*, 130–152.

Li, S., & Huang, Y. (2014). Heuristic approaches for the flow-based set covering problem with deviation paths. *Transportation Research Part E: Logistics and Transportation Review*, *72*, 144–158.

Lim, S., & Kuby, M. (2010). Heuristic algorithms for siting alternative-fuel stations using the flow-refueling location model. *European Journal of Operational Research*, *204*, 51–61.

MirHassani, S., & Ebrazi, R. (2012). A flexible reformulation of the refueling station location problem. *Transportation Science*, *47*, 617–628.

Tran, T. H., Nagy, G., Nguyen, T. B. T., & Wassan, N. A. (2018). An efficient heuristic algorithm for the alternative-fuel station location problem. *European Journal of Operational Research*, *269*, 159–170.

Upchurch, C., Kuby, M., & Lim, S. (2009). A model for location of capacitated alternative-fuel stations. *Geographical Analysis*, *41*, 85–106.

de Vries, H., & Duijzer, E. (2017). Incorporating driving range variability in network design for refueling facilities. *Omega*, *69*, 102–114.

de Vries, H., van de Klundert, J., & Wagelmans, A. (2014). The roadside healthcare facility location problem. URL: http://hdl.handle.net/1765/51315.

World Economic Forum (2018). Electric vehicles for smarter cities: the future of energy and mobility.

Wu, F., & Sioshansi, R. (2017). A stochastic flow-capturing model to optimize the location of fast-charging stations with uncertain electric vehicle flows. *Transportation Research Part D: Transport and Environment*, *53*, 354–376.

Yıldız, B., Arslan, O., & Karaşan, O. E. (2016). A branch and price approach for routing and refueling station location model. *European Journal of Operational Research*, *248*, 815–826.

Yıldız, B., Olcaytu, E., & Şen, A. (2019). The urban recharging infrastructure design problem with stochastic demands and capacitated charging stations. *Transportation Research Part B: Methodological*, *119*, 22–44.

You, P.-S., & Hsieh, Y.-C. (2014). A hybrid heuristic approach to the problem of the location of vehicle charging stations. *Computers & Industrial Engineering*, *70*, 195–204.

A. Instance Generation

The random networks are generated as follows. First, the location of s nodes is drawn from a uniform distribution in the plane $[0, 1000] \times [0, 1000]$. These nodes form the set K of potential facility locations. Subsequently, Prim's algorithm is used to create the minimum spanning tree for the complete graph containing these nodes, where we assume a Eucledian distance measure. We then add s edges connecting the s closest node pairs that were previously unconnected. The resulting graph $G = (K, E)$ thus contains $|K| = s$ nodes and $|E| = 2s - 1$ edges. In order to construct flows on this network, w out of s nodes are duplicated to form the Origin-Destination (O-D) nodes, defining $w(w-1)/2$ O-D pairs. For a given path f, the two O-D nodes are denoted by O_f and D_f. We use the Floyd-Warshall algorithm to find the shortest path connecting each O-D pair, from which follow the set K_f of potential facility locations along this path and the distance t_{kl} between any pair of nodes. Note that, for each flow, both a potential facility node as well as an O-D node are located at both origin and destination.

The vehicle volumes v_f are generated by first assigning to each of the w O-D nodes a number drawn from a standard uniform distribution $\mathcal{U}(0, 1)$. Let e_f^O and e_f^D denote these numbers for the origin and destination, respectively, and let T_f be the travel distance of the shortest path between O_f and D_f. The non-normalized vehicle volume v_f^* is then computed as follows:

$$v_f^* = \frac{e_f^O e_f^D}{T_f} 1_{T_f \geq 100} \tag{A.1}$$

Note that paths of length below one hundred are excluded, in order to limit the impact of short routes on location decisions. Finally, vehicle volumes v_f are obtained by normalizing v_f^* such that the total volume adds up to 10^6.

B. Full Results

Table B.1: Expected volume covered and chance-constrained flow volume covered in the DFRLP, EFRLP, and CCFRLP solutions to instance $s100w50$.

	Expected flow volume covered			Chance-constrained flow volume covered			
p	EFRLP solution	CCFRLP solution	DFRLP solution	EFRLP solution	CCFRLP solution	DFRLP solution	VSS
1	17,761	10,723	17,224	0	7,696	4,664	1,785
2	39,141	21,934	32,713	0	19,784	8,306	8,953
3	57,225	48,828	50,344	10,222	38,941	8,306	18,757
4	77,377	71,438	71,854	48,402	58,512	8,306	27,864
5	102,095	85,828	97,725	69,247	73,865	52,045	13,095
10	237,976	179,509	224,513	111,614	163,478	111,614	32,664
15	351,543	285,588	348,149	121,835	259,695	128,609	67,240
20	479,211	393,292	450,028	271,116	380,112	162,133	123,581
25	594,830	501,930	592,328	338,419	476,964	294,419	92,524

Table B.2: Expected volume covered and chance-constrained flow volume covered in the DFRLP, EFRLP, and CCFRLP solutions to instance $s80w40$.

	Expected flow volume covered			Chance-constrained flow volume covered			
p	EFRLP solution	CCFRLP solution	DFRLP solution	EFRLP solution	CCFRLP solution	DFRLP solution	VSS
1	24,038	21,059	24,038	0	10,076	0	5,038
2	51,842	51,842	46,127	34,731	34,731	0	20,223
3	83,027	69,761	83,027	10,076	53,012	10,076	21,468
4	104,638	74,601	104,507	10,076	72,248	10,076	31,151
5	133,900	11,2271	126,118	19,125	88,208	10,076	42,957
10	289,238	189,046	275,248	91,805	185,784	65,778	66,998
15	464,075	375,037	445,543	262,940	346,243	138,066	113,354
20	600,466	528,250	594,616	281,105	488,281	272,060	111,036
25	733,642	637,087	699,179	452,049	607,170	360,903	140,365

Table B.3: Expected volume covered and chance-constrained flow volume covered in the DFRLP, EFRLP, and CCFRLP solutions to instance $s60w30$.

	Expected flow volume covered			Chance-constrained flow volume covered			
p	EFRLP solution	CCFRLP solution	DFRLP solution	EFRLP solution	CCFRLP solution	DFRLP solution	VSS
1	48,068	0	48,068	0	0	0	0
2	93,070	58,203	66,597	0	31,630	0	29,051
3	132,140	69,159	132,140	31,630	68,660	31,630	18,515
4	172,348	83,930	172,348	31,630	81,477	31,630	24,923
5	209,999	151,548	209,999	68,660	114,936	68,660	23,138
10	312,481	278,177	312,458	150,137	277,506	150,137	63,696
15	419,316	405,592	403,271	315,948	390,015	182,337	111,861
20	560,135	502,572	534,090	419,152	497,155	278,561	122,319
25	706,483	683,096	699,229	583,067	677,773	555,287	64,870

Table B.4: Expected volume covered and chance-constrained flow volume covered in the DFRLP, EFRLP, and CCFRLP solutions to instance $s40w20$.

	Expected flow volume covered			Chance-constrained flow volume covered			
p	EFRLP solution	CCFRLP solution	DFRLP solution	EFRLP solution	CCFRLP solution	DFRLP solution	VSS
1	52,166	0	52,166	0	0	0	0
2	101,913	101,913	101,913	66,590	66,590	66,590	0
3	153,346	153,346	153,346	116,952	116,952	116,952	0
4	178,880	157,688	169,582	129,830	143,874	116,952	18,109
5	214,501	166,789	214,501	164,096	166,734	164,096	1,319
10	429,023	394,127	405,283	384,467	395,459	292,578	63,311
15	662,298	611,476	662,298	515,134	613,137	515,134	49,001
20	871,545	861,246	859,168	817,047	864,494	748,106	64,382
25	963,753	963,753	957,940	962,631	962,631	906,628	30,909

Table B.5: Expected flow volume covered in the DC-CoverEFRLP, SC-CoverEFRLP, and heuristic solutions to instance $s40w20$ with $C = 50,000$.

	Exact		Heuristic		
p	DC-CoverEFRLP	SC-CoverEFRLP	$\delta = 0$	$\delta = 0.125$	$\delta = 0.250$
1	39,321	39,321	39,321	39,321	39,321
2	55,142	57,686	55,557	55,557	55,873
3	82,630	83,307	83,306	83,306	83,306
4	98,451	99,543	99,543	99,543	99,543
5	109,585	116,313	111,146	116,313	116,313
10	163,393	165,589[1]	165,589[1]	149,917	149,917
15	197,988	199,842[1]	199,38[1]	199,283[1]	199,283[1]
20	217,365	214,508[1]	219,466[1]	214,531[1]	219,633[1]
25	232,574	233,774[1]	236,415[1]	235,470[1]	234,220[1]

[1] No optimal solution was attained within the time limit

Table B.6: Expected flow volume covered in the DC-CoverEFRLP, SC-CoverEFRLP, and heuristic solutions to instance $s40w20$ with $C = 100,000$.

	Exact		Heuristic		
p	DC-CoverEFRLP	SC-CoverEFRLP	$\delta = 0$	$\delta = 0.125$	$\delta = 0.250$
1	52,166	52,166	52,166	52,166	52,166
2	74,717	79,831	79,831	79,831	79,831
3	95,455	100,159	95,455	98,196	100,159
4	120,432	125,651	124,976	125,651	125,651
5	147,104	148,956	148,956	148,956	148,956
10	248,952	250,890[1]	250,803	248,726	248,726
15	332,992	336,044[1]	335,444[1]	335,444[1]	335,444[1]
20	394,601	395,788[1]	394,32[1]	394,958[1]	394,274[1]
25	412,421	405,000[1]	414,878[1]	412,186[1]	408,967[1]

[1] No optimal solution was attained within the time limit

Table B.7: Expected flow volume covered in the DC-CoverEFRLP, SC-CoverEFRLP, and heuristic solutions to instance $s40w20$ with $C = 150,000$.

	Exact		Heuristic		
p	DC-CoverEFRLP	SC-CoverEFRLP	$\delta = 0$	$\delta = 0.125$	$\delta = 0.250$
1	52,166	52,166	52,166	52,166	52,166
2	80,602	81,720	81,720	81,720	81,720
3	115,786	119,255	119,255	119,255	119,255
4	135,474	138,900	138,900	137,621	138,900
5	161,902	165,338	165,338	165,338	156,084
10	291,997	295,741[1]	295,741	292,798	292,798
15	410,315	410,404[1]	412,085	412,085	397,801[1]
20	509,659	508,442[1]	508,260[1]	508,260[1]	508,415[1]
25	541,059	537,969[1]	539,743[1]	538,596[1]	538,270[1]

[1] No optimal solution was attained within the time limit

Table B.8: Expected flow volume covered in the DC-CoverEFRLP, SC-CoverEFRLP, and heuristic solutions to instance $s40w20$ with $C = 200,000$.

	Exact		Heuristic		
p	DC-CoverEFRLP	SC-CoverEFRLP	$\delta = 0$	$\delta = 0.125$	$\delta = 0.250$
1	52,166	52,166	52,166	52,166	52,166
2	99,898	101,615	101,615	101,615	101,615
3	127,622	131,930	131,930	131,930	131,930
4	146,970	150,920	150,918	150,918	150,918
5	174,264	179,246	175,979	179,246	178,055
10	325,178	325,208[1]	325,208	325,208	325,208
15	451,848	453,535[1]	453,535[1]	453,535[1]	453,535[1]
20	576,333	573,790[1]	575,597[1]	575,597[1]	575,33[1]
25	610,074	606,085[1]	608,017[1]	604,400[1]	604,814[1]

[1] No optimal solution was attained within the time limit

Table B.9: Expected flow volume covered in the DC-CoverEFRLP, SC-CoverEFRLP, and heuristic solutions to instance $s40w20$ with $C = 250,000$.

	Exact		Heuristic		
p	DC-CoverEFRLP	SC-CoverEFRLP	$\delta = 0$	$\delta = 0.125$	$\delta = 0.250$
1	52,166	52,166	52,166	52,166	52,166
2	101,117	101,913	101,913	101,913	101,913
3	139,842	142,233	142,233	142,233	142,233
4	158,808	162,234	162,234	160,598	160,598
5	185,955	190,359	187,911	190,359	190,359
10	349,842	349,908[1]	349,908	349,908	346,325
15	493,213	493,555[1]	494,006[1]	491,817[1]	491,817[1]
20	627,078	625,556[1]	625,463[1]	624,636[1]	624,944[1]
25	663,889	662,574[1]	663,008[1]	662,855[1]	663,132[1]

[1] No optimal solution was attained within the time limit

C. List of Files Contained in Code Attachment

The following programs are present in the code attachment:

- `graph` folder:
  - `Dijkstra.java`: Implementation of the Dijkstra algorithm
  - `FloydWarshall.java`: Implementation of the Floyd-Warshall algorithm
  - `PathFind.java`: Implementation of a k-SP algorithm
  - `Prim.java`: Implementation of Prim's algorithm for minimum spanning trees
  - `Test.java`: Test class for all tree methods
  - `Tree.java`: Models a tree
  - `UndirectedGraph.java`: Models an undirected graph
- `instance` folder:
  - `CycleSegment.java`: Models a cycle segment
  - `Flow.java`: Models a flow
  - `Instance.java`: Representation of a problem instance
  - `Location.java`: Models a node in the road network
  - `Main.java`: Main class for DFRLP, EFRLP, and CCFRLP experiments
  - `MainCover.java`: Main class for CoverDFRLP, -EFRLP, and -CCFRLP experiments
  - `MainHeuristic.java`: Main class for heuristic experiments
  - `Route.java`: Models a route from origin to destination
- `models` folder:
  - `CCFRLP.java`: CCFRLP implementation in CPLEX
  - `CoverCCFRLP.java`: CoverCCFRLP implementation in CPLEX, extends CoverFRLP
  - `CoverDFRLP.java`: CoverDFRLP implementation in CPLEX, extends CoverFRLP
  - `CoverEFRLP.java`: CoverEFRLP implementation in CPLEX, extends CoverFRLP
  - `CoverFRLP.java`: Abstract class implementing various FRLPs in CPLEX
  - `DFRLP.java`: DFRLP implementation in CPLEX
  - `EFRLP.java`: EFRLP implementation in CPLEX
  - `SCCoverEFRLP.java`: SCCoverEFRLP implementation in CPLEX, extends CoverEFRLP
- `utils` folder:
  - `Timer.java`: Used to track CPU times