Stochastic Vehicle Routing with Recourse

Inge Li Gørtz*   Viswanath Nagarajan †   Rishi Saket†

Abstract

We study the classic *Vehicle Routing Problem* in the setting of stochastic optimization with recourse. StochVRP is a two-stage optimization problem, where demand is satisfied using two routes: fixed and recourse. The fixed route is computed using only a demand distribution. Then after observing the demand instantiations, a recourse route is computed – but costs here become more expensive by a factor λ.

We present an $O(\log^2 n \cdot \log(n\lambda))$-approximation algorithm for this stochastic routing problem, under arbitrary distributions. The main idea in this result is relating StochVRP to a special case of *submodular orienteering*, called *knapsack rank-function orienteering*. We also give a better approximation ratio for *knapsack rank-function orienteering* than what follows from prior work. Finally, we provide a Unique Games Conjecture based $\omega(1)$ hardness of approximation for StochVRP, even on star-like metrics on which our algorithm achieves a logarithmic approximation.

1 Introduction

Consider a distribution problem involving a depot location and a set of customer locations. There is a vehicle of capacity Q that is used to distribute items. The demand at customer locations is random with a known (joint) distribution $\mathcal{D}$. The distributor wants to plan a *fixed route* for this capacitated vehicle, that will be employed on a daily basis. However due to the stochastic nature of demands, the fixed route might be insufficient to meet all demands. Therefore the distributor also plans a secondary *recourse strategy*, that satisfies all unmet demands after the fixed route. Each morning the distributor receives the precise demand quantities from all customers (drawn from $\mathcal{D}$). Based on this he/she decides which subset of customers will be satisfied along the fixed route, and then plans a recourse route to satisfy the remaining customers. The goal is to minimize the cost of the fixed route plus the expected cost of the recourse route. Examples of real-world applications are local deposit collection from bank branches, garbage collection, home heating oil delivery, and forklift routing [1, 4].

A solution based on fixed routes is desirable for several reasons, and is commonly used in practice; see [30, 15] for more detailed discussions on this. In our context, there are at least two advantages. First, the driver can get familiar with the road/traffic conditions which results in time savings. Moreover, having fixed routes simplifies the everyday route planning process: the incremental recourse step will typically contain fewer demands.

Fixed-route problems are often modeled in the framework of two-stage stochastic optimization. *A priori* optimization handles some natural but simple recourse strategies: eg., short-cutting over customers without demand in TSP [5, 32], and refill-visits from the depot in the Vehicle Routing Problem (VRP) [4, 19]. Recently, more complex recourse actions have been considered: adding penalty terms in deadline TSP [8], and using backup vehicles in VRP [1].

*Technical University of Denmark, DTU Informatics.
†IBM T.J. Watson Research Center.

In this paper, we penalize the cost of the recourse route by an inflation factor $\lambda \geq 1$. This is also a common approach for two-stage stochastic optimization with recourse. Furthermore, in the stochastic VRP we consider, recourse strategies are non-trivial since it also involves choosing the subset of realized demands served by the fixed route. In this respect it is unlike most previously studied 2-stage stochastic problems (eg. [29, 34, 20]) where the recourse step is just a deterministic instance of the same problem. Before describing the results of this paper, we define the deterministic and two-stage stochastic VRP below.

Vehicle Routing Problem (VRP). There is a vehicle of capacity Q, metric (V, d) with root/depot $r \in V$ and demands $\{q_v \leq Q\}_{v \in V}$. The goal is to find a minimum cost tour of the vehicle that delivers q_v units to each $v \in V$. The demands are "unsplittable", i.e. the demand at any vertex must be satisfied in a single visit. Any VRP solution corresponds to a sequence of round-trips from the depot, where at most Q units of demands are served during each round-trip. It is well-known [2] that an α-approximation ratio for TSP implies an $(\alpha + 2)$-approximation algorithm for VRP.

Two-stage Stochastic VRP (StochVRP). The setting is same as above, with a capacity Q vehicle, metric (V, d) and depot $r \in V$. Here the demands $\{q_v\}_{v \in V}$ are random variables given by a joint demand distribution $\mathcal{D}$ on $\{0, 1, \ldots, Q\}^V$, available as a black-box that can be sampled from. We are also given an inflation parameter $\lambda \geq 1$. The goal is to compute a fixed route solution with a recourse strategy.

- In the first stage the algorithm computes a *fixed tour* τ, without knowledge of the actual demand. The tour τ consists of several round-trips from the depot: each round-trip is a cycle containing r (henceforth called r-tour). We represent τ as a concatenation $\{\tau_1, \ldots, \tau_F\}$ of r-tours. It is important to note that τ only represents the vehicle route, and does not specify demand deliveries (this will be decided after demand instantiations). In particular, a vertex v may appear in multiple r-tours of τ; and even if v appears in τ the instantiated demand at v may not eventually be satisfied by τ.

- In the second stage, the demands $\overline{q}$ are instantiated from $\mathcal{D}$. Knowing this, an algorithm chooses to satisfy subset $\overline{q}_A \subseteq \overline{q}$ of demands using the fixed tour τ, subject to the vehicle capacity of Q. That is, for each r-tour $\{\tau_i\}_{i=1}^F$ the algorithm chooses a subset $S_i \subseteq \tau_i$ of vertices to serve, where $\sum_{v \in S_i} \overline{q}_v \leq Q$; and sets $\overline{q}_A \equiv \{\overline{q}_v : v \in \cup_{i=1}^F S_i\}$. Then the algorithm computes a *recourse tour* σ meeting all residual demands $\overline{q}_B = \overline{q} \setminus \overline{q}_A$. That is, σ is a solution to the deterministic VRP instance with demands $\{\overline{q}_v : v \in V \setminus \cup_{i=1}^F S_i\}$.

Note that the demands $\overline{q}_A$ satisfied by the fixed tour τ differs based on the instantiation $\overline{q}$; however the route taken by the vehicle stays fixed. So the first stage cost is just the length $d(\tau)$ of the fixed tour. The recourse tour σ clearly depends on the demand instantiation. The second stage cost under demand $\overline{q}$ is $\lambda \cdot d(\sigma(\overline{q}))$, the length of the recourse tour inflated by a parameter λ. The objective in StochVRP is to minimize the expected total cost:

$$d(\tau) \quad + \quad \lambda \cdot \mathbb{E}_{\overline{q} \leftarrow \mathcal{D}} \left[d(\sigma(\overline{q})) \right]$$

For any integer $I \geq 0$, we let $[I] := \{1, \ldots, I\}$. For a given StochVRP instance, opt will denote its optimal value. We let $n = |V|$ denote the number of vertices in the metric and $D = \max_{u,v} d(u, v)$ the diameter of the metric.

Our Results, Techniques and Outline. In this paper we show:

Theorem 1.1 *There is a randomized $O(\log^2 n \cdot \log(n\lambda))$-approximation algorithm for* StochVRP *under arbitrary distributions.*

Using a sampling-based reduction [10] we show (in Subsection 3.1) that the objective *value* under any black-box distribution can be well-approximated by another demand distribution having support size $m = poly(n, \lambda)$.

Then, in Section 2 we present an $O(\log^2 n \cdot \log(nm))$-approximation algorithm for StochVRP where m is the support size of the distribution. This is a set-cover type algorithm that uses the *submodular orienteering problem* [12, 7] as a subroutine. In the submodular orienteering problem there is a metric (V, d) with root r, length bound B and monotone submodular function $f : 2^V \to \mathbb{R}_+$; and the goal is to find an r-tour of length at most B visiting some subset $S \subseteq V$ of vertices so as to maximize $f(S)$. Direct use of algorithms from [11, 7] yields an approximation ratio worse than Theorem 1.1 by a factor of $\log^\epsilon n$. Instead we give a better result for submodular orienteering on objective functions of the type encountered in StochVRP, called *knapsack rank-function orienteering* (KnapRankOrient). In particular, we consider the *ratio* KnapRankOrient problem where instead of the length-bound, the objective is to maximize the ratio of function value to the length.

Theorem 1.2 *There is a deterministic $O(\log^2 n)$-approximation algorithm for ratio knapsack rank-function orienteering.*

The main idea here is to use LP rounding techniques for the related group Steiner problem [17, 26], augmented with an *alteration* step (for the analysis). While alteration has been widely used with LP-rounding, eg. [35], we are not aware of an application in context of the group Steiner tree problem. This step only bounds the function-value and length in expectation (separately). In order to bound their ratio, we adapt the group Steiner derandomization from Charikar et al. [9] to our context. We defer further discussion and details on KnapRankOrient to Section 4.

Combined with the sampling-based reduction this suffices to approximate the objective *value* of StochVRP under black-box distributions, satisfying the guarantee in Theorem 1.1. However, more work is required in order to provide an approximate *solution*. This is because the recourse step in StochVRP is quite non-trivial, and a solution must specify an algorithm to construct the recourse tour for *any* possible demand (not merely the m sampled points). It turns out that the recourse step corresponds to solving an "outlier" version of VRP. Although this problem does not admit any true approximation ratio (by a relation to generalized assignment [27]), in Section 3 we give an LP-based $O(1)$ *bicriteria* approximation: this suffices for Theorem 1.1.

Our second main result is a UGC-based hardness of approximation:

Theorem 1.3 *Assuming the Unique Games Conjecture, it is NP-hard to approximate* StochVRP *to within a constant factor, even on star-like metrics.*

This is proved in Section 5 and involves a reduction from the vertex cover problem on k-uniform hypergraphs: we use a result by Bansal and Khot [3] which says that it is UGC-hard to distinguish between the (yes) case when the hypergraph is almost k-partite and the (no) case when any vertex cover is almost the entire vertex-set. We remark that this super-constant hardness holds in star-like metrics, where our algorithm achieves an $O(\log(n\lambda))$-approximation. Our algorithm loses additional log-factors in going from (i) stars to trees, and then (ii) trees to general metrics: these overheads are similar to the best known results for the related *group Steiner tree* problem [17].

Finally, we consider the special case when demands are independent across vertices. Using a different algorithm we obtain a better ratio in Section 6.

Theorem 1.4 *There is a randomized $O(\frac{\log(n\lambda)}{\log\log(n\lambda)})$-approximation algorithm for* StochVRP *under independent demand distributions.*

We show that in this case we can enforce a certain solution structure, while losing an $O(\frac{\log(n\lambda)}{\log\log(n\lambda)})$ factor in the optimal value. Specifically, we show that the demands can be partitioned into two groups: one where each demand is (almost always) served by the fixed tour, and another where each demand is served in the recourse tour. Then we use an LP-based algorithm to find the best such partition, losing another constant factor. We leave open the possibility of a constant approximation in the independent demands case.

Related Work. The VRP [37] is an extensively studied routing problem that combines aspects of both TSP and bin-packing. Several stochastic variants of the basic problem have received attention, eg. [36, 4, 14, 1, 15]. Approximation algorithms for VRP with independent stochastic demands (in the a priori model) were given in [4, 19]. This paper takes a different approach, that of two-stage stochastic optimization with recourse (along the lines of [23, 29, 34, 20] etc). To the best of our knowledge no prior approximation results are known for vehicle routing problems in this model.

Stochastic optimization [6] is a broad area dealing with probabilistic input. Approximation algorithms for two-stage stochastic problems were introduced by Immorlica et al. [23] and Ravi-Sinha [29]. Gupta et al. [20] and Shmoys-Swamy [34] gave general frameworks for approximating a number of stochastic optimization problems; the former result is combinatorial using certain cost-sharing properties, whereas the latter is LP-based. However, these approaches do not seem directly applicable to StochVRP. The results in [20, 34] hold in the most general distribution model, where an algorithm only receives independent samples from a black-box. Charikar et al. [10] showed that any arbitrary distribution can be reduced to one having polynomial support (under certain conditions). We also make use of this result in proving Theorem 1.1. For most other combinatorial optimization problems that have been considered in the two-stage stochastic model (with proportional cost inflation), it has been observed that approximation ratios are the same order of magnitude as the underlying deterministic problem [23, 29, 20, 34, 33]. A notable exception is minimum cost max-matching [24], for which an $\Omega(\log n)$-hardness of approximation was shown. In the case of VRP Theorem 1.3 shows (under UGC) that the stochastic approximation ratio is necessarily worse than its deterministic counterpart, even in very special metrics.

2 Algorithm for Polynomial Scenarios

Here we consider the case when the demand distribution $\mathcal{D}$ is specified as a list of possible outcomes. Later on we show how the general case of a black-box distribution can be reduced to this case. Formally $\mathcal{D}$ is a multiset $\{\overline{q}^1, \ldots, \overline{q}^m\}$ where the actual demand $\overline{q} = \overline{q}^i$ (for some $i \in [m]$) with probability $1/m$.

The main idea of our algorithm is to recast the problem as an instance of set-cover with an exponential number of sets. Then we show that the greedy subproblem is an instance of *submodular orienteering* (SOP) for which a poly-logarithmic approximation is known [7, 12]. In fact, for the type of SOP instances obtained from StochVRP we give a better approximation ratio in Section 4. Altogether, this implies Theorem 1.1 for polynomial scenarios.

Set cover instance $\mathcal{I}$. The groundset U consists of tuples $\langle i, v\rangle$ for all scenarios $i \in [m]$ and vertices $v \in V$, which denotes $\overline{q}^i(v)$ demand units at v under scenario i. For any $\langle i, v\rangle \in U$ we use $q(\langle i, v\rangle) := \overline{q}^i(v)$, and for any subset $S \subseteq U$, $q(S) := \sum_{t \in S} q(t)$. Instance $\mathcal{I}$ has the following two types of sets:

1. $S := \cup_{i=1}^m S_i$ is a **first stage set** iff $S_i \subseteq \{\langle i, v\rangle : v \in V\}$ and $q(S_i) \leq Q$ for all $i \in [m]$. The cost of this set S is the minimum length of an r-tour that contains all the vertices represented in S.
2. For any scenario $i \in [m]$, $T \subseteq \{\langle i, v\rangle : v \in V\}$ is a **second stage set** iff $q(T) \leq Q$. The cost of set T is λ/m times the minimum length of an r-tour containing all vertices of T.

Lemma 2.1 *The set cover instance $\mathcal{I}$ is equivalent to* StochVRP.

Proof: Recall that any feasible StochVRP solution is specified by:

- *The fixed tour τ.* It will be convenient to view this as a collection $\{\tau_1, \ldots, \tau_F\}$ of r-tours, each of which is a round-trip from the depot.

- *For each scenario $i \in [m]$, the demands $\overline{q}^i_A \subseteq \overline{q}^i$ satisfied by the fixed tour.* Again this is viewed as follows: for each r-tour $\{\tau_j\}_{j=1}^F$, $S_{i,j} \subseteq \{\langle i, v\rangle : v \in V\}$ denotes the demands satisfied in τ_j. Note that by definition, $\bigcup_{j\in[F]} S_{i,j} \equiv \overline{q}^i_A$. Also due to the capacity constraint, $q(S_{i,j}) \leq Q$ for each $j \in [F]$.

- *For each scenario $i \in [m]$, the recourse tour σ_i which satisfies residual demands $\overline{q}^i \setminus \overline{q}^i_A$.* Again we view this as a collection $\{\sigma_{i,1}, \ldots, \sigma_{i,L_i}\}$ of r-tours. For $k \in [L_i]$ let $T_{i,k} \subseteq \{\langle i, v\rangle : v \in V\}$ denote the demands satisfied in $\sigma_{i,k}$. Clearly, $\bigcup_{k\in[L_i]} T_{i,k} \equiv \overline{q}^i \setminus \overline{q}^i_A$. Again $q(T_{i,k}) \leq Q$ for all $k \in [L_i]$.

Note that corresponding to each first stage r-tour τ_j, the set $\bigcup_{i=1}^m S_{i,j}$ is a valid *first stage set* in $\mathcal{I}$ since for all $i \in [m]$ (a) $S_{i,j} \subseteq \{\langle i, v\rangle : v \in V\}$ and (b) $q(S_{i,j}) \leq Q$. Moreover the cost of this set in $\mathcal{I}$ is at most $d(\tau_j)$.

Similarly, for each scenario $i \in [m]$ and second stage r-tour $\sigma_{i,k}$ ($k \in [L_i]$), set $T_{i,k}$ is a valid *second stage* set. The cost of this set in $\mathcal{I}$ is at most $\frac{\lambda}{m} \cdot d(\sigma_{i,k})$.

Finally, these sets cover U in $\mathcal{I}$ since for each scenario $i \in [m]$, we have:

$$\left(\cup_{j=1}^F S_{i,j}\right) \bigcup \left(\cup_{k\in[L_i]} T_{i,k}\right) \quad = \quad \{\langle i, v\rangle : v \in V\}$$

The total cost of this solution to $\mathcal{I}$ is at most:

$$\sum_{j=1}^F d(\tau_j) \;+\; \frac{\lambda}{m} \cdot \sum_{i=1}^m \sum_{k=1}^{L_i} d(\sigma_{i,k}) \quad = \quad d(\tau) \;+\; \lambda \cdot \mathbb{E}_{\overline{q} \leftarrow \mathcal{D}}\left[d(\sigma(\overline{q}))\right],$$

which is just the StochVRP objective value. The reverse relation (from $\mathcal{I}$ to StochVRP) can be shown in a similar manner, and the lemma follows. ■

Thus it suffices to solve the set cover instance $\mathcal{I}$. We use the greedy algorithm for set cover which requires solving the following *max-coverage subproblem*: given $U' \subseteq U$ find a set (of either first/second type) that maximizes the ratio of the number of U'-elements it covers to its cost. We give separate algorithms for this problem, under the two types of sets.

Max-coverage for second stage sets. We give a constant approximation in this case. Assume that the algorithm knows by enumeration (i) the cost B of the best ratio set (up to a factor two), and (ii) the scenario $i \in [m]$ corresponding to it. Then it suffices to find a set $T \subseteq U' \bigcap \{\langle i, v\rangle : v \in V\}$ maximizing $|T|$ such that $q(T) \leq Q$ and $\text{cost}(T) \leq B$. By the definition of second stage sets, this reduces to finding an r-tour visiting the maximum vertices $W \subseteq \{u \in V : \langle i, u\rangle \in U'\}$, having length at most $\frac{m}{\lambda} \cdot B$ *and* with $\sum_{u\in W} \overline{q}^i(u) \leq Q$. This is just an instance of the *knapsack-orienteering* problem, for which a constant-factor approximation is known [18].

Max-coverage for first stage sets. In this case, we obtain a poly-logarithmic approximation. Again, we assume that the algorithm knows the cost B of the best ratio set (up to a factor two). Recall that unlike the previous case, one first stage set can cover elements from several scenarios. By definition, each first stage set S corresponds to an r-tour visiting vertices $W \subseteq V$ *and* subsets $S_i \subseteq \{\langle i, v\rangle : v \in W\}$

for each $i \in [m]$ such that $\{q(S_i) \le Q\}_{i=1}^m$ and $S = \bigcup_{i=1}^m S_i$. Among all first stage sets visiting a *fixed vertex-set* $W \subseteq V$, the maximum coverage of U' equals:

$$f(W) \quad := \quad \sum_{i=1}^{m} \max\left\{ |S_i| \ : \ S_i \subseteq \{u \in W : \langle i, u\rangle \in U'\}, \ \sum_{v \in S_i} \overline{q}_v^i \le Q \right\}$$

For each $i \in [m]$ let $f_i(W)$ denote the term inside the above summation. Recall that the cost of all first stage sets visiting vertices W is the same, namely the minimum TSP on $\{r\} \cup W$. Thus the subproblem we wish to solve is:

$$\max \ f(W) \ : \ \text{ there is an } r\text{-tour visiting } W \subseteq V \text{ of length } \le B. \tag{1}$$

Recall the submodular orienteering problem (SOP) where given metric (V, d) with root r, bound B and submodular function $g : 2^V \to \mathbb{R}_+$, the goal is to find an r-tour visiting some subset $W \subseteq V$ of vertices, having length at most B that maximizes $g(W)$. If f were submodular then we can use the algorithm [7, 12] to solve this. But f is not submodular. Still, we show below that it can be well approximated by a submodular function g.

We approximate each f_i (point-wise) by a submodular function g_i. Let $V_i := \{u \in V : \langle i, u\rangle \in U'\}$ denote the vertices appearing with scenario i in U'. Define:

$$g_i(W) \quad := \quad \max\left\{ \sum_{v \in V_i \cap W} x_v \ : \ \sum_{v \in W} \overline{q}_v^i \cdot x_v \le Q, \ \ 0 \le x_v \le 1, \ \forall\, v \in W \right\}$$

Observe that $g_i(W)$ is just an LP relaxation for a maximization $\{0,1\}$-knapsack problem. So its value is given by the greedy algorithm that increases x_v (up to 1) in increasing order of $\{\overline{q}_v^i : v \in V_i \cap W\}$. On the other hand, $f_i(W)$ is the value of the same *integral* knapsack problem. Now, function g_i can be rewritten as the rank function of a *polymatroid* [31] which is submodular; see eg. [13]. Moreover, the integrality gap of the natural LP for max-knapsack is two. Thus,

Claim 2.2 *g_i is monotone submodular and $\frac{g_i(W)}{2} \le f_i(W) \le g_i(W)$, $\forall W \subseteq V$.*

So if we define $g(W) := \sum_{i=1}^m g_i(W)$ then it is submodular and maximizing g in (1) is equivalent to maximizing f (up to factor two). Hence, assuming a ρ-approximation algorithm for SOP, we obtain a 2ρ-approximation algorithm for (1). This suffices to give an $O(\rho)$-approximation for the max-coverage subproblem. We have $\rho = O(\log^{2+\epsilon} n)$ in polynomial time using the bicriteria approximation in Calinescu-Zelikovsky [7], and $\rho = O(\log n)$ in *quasi-polynomial* time using the true approximation in Chekuri-Pal [12]. In Section 4 we directly consider the *ratio* objective corresponding to (1), called *ratio knapsack rank-function orienteering*, i.e.

$$\max \left\{ \frac{f(V(\tau))}{d(\tau)} \quad : \quad \tau \text{ is an } r\text{-tour visiting vertices } V(\tau) \right\},$$

and give an improved polynomial time $O(\log^2 n)$-approximation algorithm for it.

Finally, we lose an additional $\log |U| = O(\log(mn))$ factor to solve the set cover instance $\mathcal{I}$ (which is equivalent to StochVRP). Thus we obtain:

Theorem 2.3 *There is a polynomial time $O(\log^2 n \cdot \log(nm))$-approximation algorithm for* StochVRP *for a polynomial number m of scenarios and n vertices. This ratio improves to $O(\log n \cdot \log(nm))$ in quasi-polynomial time.*

3 Algorithm for General Distributions

In this section we prove Theorem 1.1 under an arbitrary distribution $\mathcal{D}$ that is accessed by sampling. We denote the input $\mathsf{StochVRP}$ instance by $\mathcal{J}$. In Subsection 3.1 we apply a sampling-based reduction from [10] to obtain an equivalent $\mathsf{StochVRP}$ instance $\mathcal{J}'$ with $m = poly(n, \lambda)$ scenarios. This allows us to apply the algorithm from the previous section to approximate the *optimal value* of instance $\mathcal{J}$. However a solution to $\mathcal{J}$ must also specify a valid recourse strategy for *every* outcome $\overline{q} \in \mathcal{D}$, and not just for the m outcomes in instance $\mathcal{J}'$. It turns out that the recourse step is captured by an "outlier" version of VRP, and we give an LP-based constant-factor bicriteria approximation for it in Subsection 3.2.

3.1 Sampling Based Reduction to Polynomial Scenarios

Here we show that sampling can be used to reduce an arbitrary demand distribution to one having a polynomial number of scenarios.

Given a fixed tour τ and scenario $\overline{q} \in \mathcal{D}$, let $h(\tau, \overline{q})$ denote the minimum cost of a recourse tour. Note that computing $h(\tau, \overline{q})$ involves choosing a subset $\overline{q}_A$ of $\overline{q}$ to be served by τ (at zero cost, but subject to capacity) and then optimally solving the VRP instance with demands $\overline{q} - \overline{q}_A$ and lengths inflated by factor λ. Thus we can express the minimum objective value for a given fixed tour τ as:

$$\mathsf{obj}(\tau) \quad := \quad d(\tau) \; + \; \mathbb{E}_{\overline{q} \leftarrow \mathcal{D}} \left[h(\tau, \overline{q}) \right].$$

The optimal value of $\mathsf{StochVRP}$ instance $\mathcal{J}$ is then $\min_{\tau \in X} \mathsf{obj}(\tau)$, where X denotes the set of all possible fixed tours. Consider drawing m independent samples $\{\overline{q}^1, \ldots, \overline{q}^m\}$ from $\mathcal{D}$, and let $\mathcal{J}'$ denote the (random) instance of $\mathsf{StochVRP}$ with these as explicit scenarios. Define:

$$\widehat{\mathsf{obj}}(\tau) \quad := \quad d(\tau) \; + \; \frac{1}{m} \cdot \sum_{i=1}^{m} h(\tau, \overline{q}^i), \qquad \forall \text{ fixed tour } \tau \in X.$$

It is clear that $\min_{\tau \in X} \widehat{\mathsf{obj}}(\tau)$ is the optimal value of $\mathcal{J}'$. We now use the result of Charikar et al. [10] to relate these two instances. For completeness we give a proof adapted to our context. Let $D = \max_{u,v} d(u, v)$ denote the diameter of the metric; we assume WLOG (by scaling) that all distances are integral.

Theorem 3.1 ([10]) *Using $m = \Theta(\lambda^2 \, D^2 \, n^2 \, \log |X|)$, with probability $1 - o(1)$,*

$$|\mathsf{obj}(\tau) - \widehat{\mathsf{obj}}(\tau)| \le 1, \qquad \textit{for all } \tau \in X.$$

Proof: Fix any fixed tour $\tau \in X$. Define $H = \mathbb{E}_{\overline{q} \leftarrow \mathcal{D}} \, h(\tau, \overline{q})$ and random variables $H_i := h(\tau, \overline{q}^i)$ for $i \in [m]$. Clearly $\mathbb{E} H_i = H$ for all $i \in [m]$. Note that $H_i \le 2\lambda n D$: the worst case recourse action involves a separate round-trip to each vertex. Since $H_i/(2\lambda n D)$ are independent $[0, 1]$ random variables, by Chernoff bound [28],

$$\Pr\left[\left| \frac{1}{m} \cdot \sum_{i=1}^{m} H_i - H \right| > \epsilon \right] \quad \le \quad 2 \exp\left(- \frac{\epsilon^2 \cdot m}{4\lambda^2 n^2 D^2} \right), \qquad \forall \epsilon > 0.$$

Now observe that $\mathsf{obj}(\tau) = d(\tau) + H$ whereas $\widehat{\mathsf{obj}}(\tau) = d(\tau) + \frac{1}{m} \cdot \sum_{i=1}^{m} H_i$. Thus using the above inequality with $\epsilon = 1$ and $m = 8\lambda^2 n^2 D^2 \cdot \log |X|$, we obtain:

$$\Pr\left[|\mathsf{obj}(\tau) - \widehat{\mathsf{obj}}(\tau)| > 1 \right] \quad \le \quad \frac{2}{|X|^2}, \qquad \forall \tau \in X.$$

Finally, a union bound over all X implies the theorem. ■

We now show that $m = poly(n, \lambda)$.

Claim 3.2 *WLOG the number of r-tours in any fixed tour is at most n. Hence the number of edges used in the fixed tour is at most n^2, and $|X| \leq 2^{n^2}$.*

Proof: Consider an arbitrary fixed tour τ consisting of r-tours $\tau_1, \ldots, \tau_F$. Suppose that $F > n$: then we will show there exists another fixed tour τ' with at most n r-tours such that $\mathsf{obj}(\tau') \leq \mathsf{obj}(\tau)$. Let us number vertices so that the depot is numbered 0 and $d(0,1) \leq d(0,2) \leq \cdots \leq d(0,n)$. For each $j \in [F]$ let $M(j) \in [n]$ denote the maximum numbered vertex in r-tour τ_j; note $d(\tau_j) \geq 2 \cdot d(0, M(j))$. Choose $k \in [n]$ as the minimum value so that $|M^{-1}(\{k, \ldots, n\})| \leq n - k$; if there is no such value then set $k = n + 1$.

Let $G = M^{-1}(\{k, \ldots, n\}) \subseteq [F]$; note that $G = \emptyset$ when $k = n + 1$. By choice of k, we have $|G| \leq n - k + 1$.

Also by choice of k, it follows that $|M^{-1}(\{t, \ldots, n\})| \geq n - t + 1$ for all $t \leq k - 1$. Thus there is an injective map $\phi : [k-1] \to [F] \setminus G$ such that $M(\phi(v)) \geq v$ for all $v \in [k-1]$; this can be obtained say by a greedy assignment starting from $k - 1$ (recall $|G| \leq n - k + 1$). Due to the vertex numbering (and definitions of ϕ, M) we have $2 \cdot d(0, v) \leq 2 \cdot d(0, M \circ \phi(v)) \leq d(\tau_{\phi(v)})$ for all $v \in [k-1]$. And since ϕ is injective, $2 \cdot \sum_{v=1}^{k-1} d(0, v) \leq \sum_{j \in [F] \setminus G} d(\tau_j)$.

Set τ' to consist of the following r-tours: (a) all singleton r-tours $\langle 0, v, 0 \rangle$ for $v \in [k-1]$, and (b) $\{\tau_j : j \in G\}$. Using the above inequality, $d(\tau') \leq d(\tau)$. Observe that vertices $\{1, \ldots, k-1\}$ will play no role in the second stage under τ', since they are already individually covered in a first-stage r-tour. Moreover, for any vertex $v \in \{k, \ldots, n\}$, the set of r-tours containing v is identical in both τ and τ'. Hence for *any scenario* $\overline{q} \in \mathcal{D}$, the recourse action (for vertices $\{k, \ldots, n\}$) under τ is also feasible under τ'. This implies $h(\tau', \overline{q}) \leq h(\tau, \overline{q})$ for all $\overline{q} \in \mathcal{D}$; and so $\mathsf{obj}(\tau') \leq \mathsf{obj}(\tau)$. Also by construction, the number of r-tours in τ' is $k - 1 + |G| \leq n$. ■

Claim 3.3 *WLOG, the number of edges used in any recourse tour is at most $2n$.*

Proof: Note that any recourse tour (under any outcome $\overline{q}$) is a solution to some deterministic VRP instance. Since there are at most n demands, and we consider unsplittable routing, it is clear that the number of edges used is at most $2n$. ■

So far we have shown $|X| \leq 2^{n^2}$. Next we will show that $D = poly(n)$, which suffices to prove $m = poly(n, \lambda)$ in Theorem 3.1.

Assume (by enumeration) that the algorithm knows an upper bound B on the optimal value of instance $\mathcal{J}$ (up to factor two), i.e. $B/2 \leq \mathsf{opt}(\mathcal{J}) \leq B$. Let $U \subseteq V$ denote the vertices at distance at most B from r. Clearly, the optimal fixed tour does not visit any vertex outside U (otherwise it incurs cost larger than $2B$). So we may always defer demands at $V \setminus U$ to the second stage (which is what opt does). And by using the $O(1)$-approximation algorithm [2] for deterministic VRP to serve $V \setminus U$, the cost incurred by our algorithm on $V \setminus U$ is at most $O(1) \cdot B$. Now we can focus on the StochVRP instance restricted to vertices U.

Consider the following modification to the metric (U, d) : contract all edges of length smaller than B/n^3 and let (W, ℓ) denote the resulting metric of shortest-path distances. We consider the natural StochVRP instance $\mathcal{J}''$ on (W, ℓ) (where $\mathcal{D}$ induces the demand distribution on W as well). The useful property of metric ℓ is that it has maximum distance $\leq 2B$ and minimum distance $\geq B/n^3$; so by scaling we obtain that it has diameter at most $O(n^3)$. Thus we can apply Theorem 3.1 to instance $\mathcal{J}''$ and $m = poly(n, \lambda)$ samples would suffice. The following lemma relates the $\mathcal{J}''$ to the original instance $\mathcal{J}$.

Lemma 3.4 *The optimal value* $\mathsf{opt}(\mathcal{J}'') \le B$*. Moreover, any solution to* $\mathcal{J}''$ *yields a solution to* $\mathcal{J}$ *with at most a constant factor increase in objective.*

Proof: The first part of the claim is trivial since $\mathcal{J}''$ is obtained from $\mathcal{J}$ by contracting the metric: so $\mathsf{opt}(\mathcal{J}'') \le \mathsf{opt}(\mathcal{J}) \le B$. For the other direction, consider any solution to $\mathcal{J}''$: we now describe the solution corresponding to this in $\mathcal{J}$. Note that each vertex $w \in W$ corresponds to some subset $U_w \subseteq U$ such that there is a spanning tree on U_w in metric d with each edge of length at most B/n^3. So whenever vertex $w \in W$ is visited in $\mathcal{J}''$, we will visit all vertices of U_w along an Euler tour of $MST_d(U_w)$: this results in a cost increase of at most $2|U_w| \cdot B/n^3 = O(B/n^2)$. Note also that each edge e in metric (W, ℓ) corresponds to some path in metric (U, d) of length at most $\ell_e + n \cdot B/n^3$: so each edge traversal causes a cost increase of at most B/n^2. By Claim 3.2 the cost increase in the fixed tour is at most $n^2 \cdot O(B/n^2) = O(B)$. Similarly, using Claim 3.3 the cost increase in the recourse tour (under *any* outcome) is at most $2n \cdot O(B/n^2) \le O(B)$; so the increase in the expected cost of the recourse tour is also $O(B)$. Thus any $\mathcal{J}''$-solution corresponds to a $\mathcal{J}$-solution where the increase in objective is at most $O(B) = O(1) \cdot \mathsf{opt}(\mathcal{J})$. ■

Algorithm 3.1 summarizes the StochVRP algorithm.

Algorithm 3.1 Algorithm for StochVRP under black-box distributions

Input: StochVRP instance $\mathcal{J} = \langle (V, d),\, r,\, Q,\, \lambda,\, \mathcal{D} \rangle$.

1: Guess (by enumeration) value B such that $B/2 \le \mathsf{opt}(\mathcal{J}) \le B$.
2: Restrict instance to vertices $U = \{v \in V : d(r, v) \le B\}$. Vertices $V \setminus U$ are handled separately, always in the recourse tour (costs $O(B)$ in expectation).
3: Modify metric (U, d) to (W, ℓ) by contracting edges shorter than B/n^3 and recomputing shortest paths. By scaling, $D = \mathit{diameter}(W, \ell) \le O(n^3)$. $\mathcal{J}''$ is the induced StochVRP instance on (W, ℓ).
4: Apply Theorem 3.1 to $\mathcal{J}''$ to obtain a (random) instance $\mathcal{J}'$ of StochVRP with explicit scenarios and $m = \mathit{poly}(n, \lambda)$.
5: Obtain fixed tour τ for $\mathcal{J}'$ using the algorithm in Section 2. Theorem 3.1 implies that this is a fixed tour for $\mathcal{J}''$ with increase in objective being at most one w.h.p.
6: By Lemma 3.4, τ is also a fixed tour for $\mathcal{J}$.
7: **Output** the fixed tour for $\mathcal{J}$ containing *five* copies of each r-tour in τ.
8: **Output** recourse action (for any $\overline{q}$) as $\mathsf{Aug}(\tau, \overline{q})$ given in Algorithm 3.2.

3.2 Specifying Recourse Actions

The recourse strategy involves the following *outlier VRP* problem: given a fixed tour τ (as collection $\{\tau_1, \ldots, \tau_F\}$ of r-tours) and outcome $\overline{q} \in \{0, \ldots, Q\}^V$, find

- a subset of vertices whose demands $\overline{q}_A \subseteq \overline{q}$ can be served by the existing route τ, subject to the capacity constraint of Q on its r-tours; *and*
- a minimum cost VRP solution to the residual demands $\overline{q} - \overline{q}_A$.

The optimal value of this instance is exactly function $h(\tau, \overline{q})$ defined Subsection 3.1. We remark that when the capacity $Q = 1$, the outlier VRP problem can be solved exactly using a minimum cost flow formulation. When the fixed tour $\tau = \emptyset$ we obtain the usual VRP, which is NP-hard for $Q \ge 3$. Another special case of outlier VRP is the *restricted assignment* problem [27]. This occurs when V denotes the set of jobs with sizes $\overline{q}$, there are F machines, and potential job-machine assignments are given by τ

(job v can be assigned machine j iff $v \in \tau_j$); there is an assignment of makespan Q iff the outlier VRP optimum is zero. So it is NP-hard to obtain any true approximation ratio for outlier VRP. Instead we give an $(O(1), O(1))$ *bicriteria* approximation algorithm, which suffices to obtain an algorithm for $\mathsf{StochVRP}$ with only constant-factor increase over Theorem 2.3.

The algorithm is based on a natural LP relaxation to outlier VRP. Consider a solution with $S \subseteq V$ as the vertices chosen to be served by τ. Then:

- There is an assignment $\phi : S \to [F]$ such that (1) $v \in \tau_{\phi(v)}$ for all $v \in S$; and (2) for each r-tour $j \in [F]$, the total demand assigned to it $\sum_{v\in\phi^{-1}(j)} \overline{q}_v \leq Q$.
- The objective value is the optimum VRP on metric (V, d), depot r, capacity Q and demands $\{\overline{q}_v : v \in V \setminus S\}$. Using known lower-bounds for VRP [22, 2], at the loss of a constant factor, this is just $\mathsf{MST}(V \setminus S) + \mathsf{Flow}(V \setminus S)$ where for any $T \subseteq V$, $\mathsf{MST}(T)$ = length of minimum spanning tree on $\{r\} \bigcup T$, and $\mathsf{Flow}(T) := \frac{1}{Q} \sum_{v\in T} \overline{q}_v \cdot d(r, v)$.

Thus we can write the following integer programming formulation for outlier VRP, at the loss of $O(1)$-factor.

$$\min \quad \sum_{e\in E} d_e \cdot z_e \;+\; \frac{1}{Q}\sum_{v\in V} d(r,v)\cdot \overline{q}_v \cdot (1-x_v) \tag{2}$$

$$\text{s.t.} \quad \sum_{v\in\tau_j} \overline{q}_v \cdot y_{v,j} \leq Q \qquad \forall j \in [F], \tag{3}$$

$$\sum_{j\in[F]:v\in\tau_j} y_{v,j} = x_v \qquad \forall v \in V, \tag{4}$$

$$\sum_{e\in\delta(U)} z_e \geq 1 - x_v \qquad \forall U \not\ni r,\ \forall v \in U, \tag{5}$$

$$x_v, y_{v,j} \in \{0,1\} \qquad \forall v \in V,\ \forall j \in [F],$$

$$z_e \geq 0 \qquad \forall e \in E.$$

Above x_v is one iff $v \in S$, i.e. served by τ. Variables $y_{v,j}$ denote the assignment $\phi : S \to [F]$. Constraint (4) ensures that each $v \in S$ is assigned to some $\phi(v)$ such that $v \in \tau_{\phi(v)}$. Constraint (3) enforces the total assignment to each r-tour is at most Q. Also $E = \binom{V}{2}$ denotes the edge-set of the metric, and for any $U \subseteq V$, $\delta(U)$ denotes the edges with exactly one vertex in U. Constraint (5) says that $\{z_e : e \in E\}$ is a fractional spanning tree connecting the vertices $\{v : x_v = 0\} = V \setminus S$ to r. In the objective (2), the first term is the length of the fractional spanning tree (corresponding to $\mathsf{MST}(V \setminus S)$), and the second term is $\mathsf{Flow}(V \setminus S)$. Dropping the integrality gives us an LP relaxation $\mathsf{LP}(\tau, \overline{q})$. We can solve this LP in polynomial time, and next we describe a rounding algorithm.

Observe that the solution from Algorithm 3.2 uses *five* copies of the fixed tour τ, whereas we will bound the cost against $\mathsf{LP}(\tau, \overline{q})$.

First we show that our assignment S to the fixed tour is indeed feasible (using 5 copies). It is clear that setting $\alpha_{v,j} = 2\, y_{v,j}$ for $v \in S$, $j \in [F]$ gives a feasible fractional solution to the restricted assignment instance in Step 3 of Algorithm 3.2: this follows from constraints (3) and (4) using $\{x_v \geq \frac{1}{2}\}_{v\in S}$. Thus the rounding algorithm from [27] can be employed on α to obtain an integral solution ϕ having load at most $2Q + \max_v \overline{q}_v \leq 3Q$ for each $j \in [F]$. Then for each $j \in [F]$, we partition $\phi^{-1}(j)$ starting with the trivial partition into singletons and greedily merging parts as long as each part $\leq Q$: this results in at most 5 parts. Thus vertices S can be feasibly assigned to $5\,\tau$.

Algorithm 3.2 Algorithm for computing recourse action $\mathsf{Aug}(\tau, \overline{q})$.

Input: (V, d), r, capacity Q, fixed tour $\tau = \{\tau_1, \ldots, \tau_F\}$, outcome $\overline{q}$.

1: Let (x, y, z) denote the optimal LP solution.

2: Set $S \leftarrow \{v \in V : x_v \geq \frac{1}{2}\}$.

3: Consider the following instance of *restricted assignment*

$$\begin{array}{lll} \sum_{v\in\tau_j} \overline{q}_v \cdot \alpha_{v,j} \ \leq \ 2\cdot Q & & \forall j \in [F], \\ \sum_{j\in[F]:v\in\tau_j} \alpha_{v,j} \ = \ 1 & & \forall v \in S, \\ 0 \ \leq \ \alpha_{v,j} \ \leq \ 1 & & \forall v \in S, \ \forall j \in [F]. \end{array}$$

4: By the rounding algorithm in [27] we can obtain an integral assignment $\phi : S \to [F]$ such that $v \in \tau_{\phi(v)}$ for all $v \in S$ and $\overline{q}\left(\phi^{-1}(j)\right) \leq 3Q$ for all $j \in [F]$.

5: For each $j \in [F]$, partition $\phi^{-1}(j) = \bigsqcup_{l=1}^{5} S_{j,l}$ into at most five parts such that $\{q(S_{j,l}) \leq Q\}_{l=1}^{5}$. This can be done by a greedy algorithm.

6: **Output** for each $j \in [F]$ and $l \in \{1, \ldots, 5\}$, vertices $S_{j,l}$ as served by τ_j.

7: **Output** an $O(1)$-approximate VRP solution [2] on vertices $V \setminus S$ as recourse tour.

Next we bound the cost of our recourse tour by $O(1) \cdot \mathsf{LP}(\tau, \overline{q})$. Observe that $x_v < \frac{1}{2}$ for each $v \in V \setminus S$: so constraint (5) implies that $2 \cdot z$ is a fractional spanning tree on $\{r\} \cup (V \setminus S)$. Hence $\mathsf{MST}(V \setminus S) \leq 4\,(\mathbf{d} \cdot \mathbf{z})$. Moreover, it is clear that $\mathsf{Flow}(V \setminus S) \leq 2 \cdot \frac{1}{Q} \sum_{v\in V} d(r, v) \cdot \overline{q}_v \cdot (1 - x_v)$. Thus the LP objective:

$$\mathsf{LP}(\tau, \overline{q}) \quad \geq \quad \frac{1}{4} \cdot \mathsf{MST}(V \setminus S) + \frac{1}{2} \cdot \mathsf{Flow}(V \setminus S).$$

Since the VRP algorithm (on demands $V \setminus S$) achieves a constant approximation relative to these lower bounds, it follows that the recourse cost is $O(1) \cdot \mathsf{LP}(\tau, \overline{q})$.

Theorem 3.5 *There is an $(O(1), 5)$-bicriteria approximation algorithm for outlier VRP, that uses the fixed tour at most five times.*

4 Algorithm for Ratio Knapsack Rank-function Orienteering

In this section we give an improved result for the ratio version of submodular orienteering, when the objective is a sum of "knapsack rank-functions". This can be used as a subroutine for $\mathsf{StochVRP}$ to yield Theorem 1.1.

An instance of the *knapsack rank-function orienteering* problem ($\mathsf{KnapRankOrient}$) consists of metric (V, d) with root r and length bound B. The objective is a sum of m knapsack rank-functions $f_1, \ldots, f_m : 2^V \to \mathbb{R}_+$. For a solution visiting vertices U, the objective value is $\sum_{i=1}^{m} f_i(U)$. The goal is to find an r-tour of length at most B having maximum objective. Each knapsack rank-function f_i is:

$$f_i(U) \quad := \quad \max\left\{ \sum_{v\in S} w_v^i \ \ : \ \ S \subseteq U, \ \sum_{v\in S} c_v^i \leq 1 \right\}, \quad \forall\, U \subseteq V$$

Above $w^i : V \to \mathbb{R}_+$ and $c^i : V \to [0, 1]$ denote profits and sizes at the vertices; so $f_i(U)$ is the maximum profit in any subset of U having size at most one. (Although a knapsack rank-function may not be submodular, it can always be approximated within factor two by a submodular function, as discussed in the end of Section 2.)

Here we consider the ratio version of this problem, called *ratio* KnapRankOrient, where given metric (V, d) with root r and knapsack rank-functions $f_1, \ldots, f_m$, the goal is:

$$\max \left\{ \frac{\sum_{i=1}^{m} f_i(V(\tau))}{d(\tau)} : \tau \text{ is an } r\text{-tour} \right\}$$

Above $V(\tau) \subseteq V$ are the vertices visited by τ and $d(\tau)$ is its length.

Note that the max-coverage problem for first stage sets corresponds exactly to ratio KnapRankOrient. Here we obtain an $O(\log^2 n)$-approximation algorithm for ratio KnapRankOrient, which combined with the algorithm in Section 2 implies Theorem 1.1.

Related Work. KnapRankOrient is closely related to the group Steiner tree problem, where given metric (V, d) with root r and m groups $G_1, \ldots, G_m \subseteq V$ of vertices, the goal is to find a minimum length tree connecting r to at least one vertex of each group. The best known approximation ratio for this problem is $O(\log^2 n \cdot \log m)$ [17]. Formally, ratio KnapRankOrient generalizes the "density group Steiner" problem [9], which involves finding an r-tour maximizing the ratio of number of groups covered to its length: setting $w_v^i = c_v^i = \mathbf{1}[v \in G_i]$ in ratio KnapRankOrient gives us density group Steiner. Our algorithm builds on [9]; however the previous algorithm is not directly applicable since the natural LP relaxation to ratio KnapRankOrient seems weak. We strengthen the LP relaxation for KnapRankOrient by adding extra constraints (see constraints (11) below), that are motivated from the related *covering* Steiner tree problem [26]. Moreover, as with all LP-based algorithms for group Steiner type problems, the main rounding step is the dependent randomized rounding (called GKR rounding below) from [17]. Even with a good LP relaxation and rounding step, previously used analysis such as [17, 26, 21] is inadequate for bounding the profit in ratio KnapRankOrient as shown next.

Example 1: Consider a star-like tree with a single edge (r, u) from the root and t other edges $(u, v_1), \ldots, (u, v_t)$, all of unit length. There is a single knapsack with zero sizes ($\mathbf{c} = \mathbf{0}$), and profits $\mathbf{w}$ of one at each $V' = \{v_1, \ldots, v_t\}$ (and zero at r, u). Suppose the LP solution sets value $x = 1/t^2$ for all edges. The fractional profit is $\mu = 1/t$. The analysis in all of [17, 26, 21] attempts to upper bound:

$$\Pr[\text{number of } V'\text{-vertices chosen in GKR rounding } < \mu/2]. \tag{6}$$

In this example, the solution (from GKR rounding) is the entire tree with probability $\frac{1}{t^2}$, and empty otherwise. So the probability (6) is $1 - \frac{1}{t^2}$, which by itself would only imply an expected profit of $1/t^2 \ll \mu$ (although the actual expected profit is $\frac{1}{t}$). Such an analysis is sufficient when $\mu = \Omega(1)$ as in [17, 26, 21], but we are not guaranteed this in ratio KnapRankOrient.

Instead of using a bound on the probability (6), we directly lower bound the *expected* profit using a different analysis. In particular our main idea is to use an *alteration* step after GKR rounding (see Lemma 4.1). While alteration has been used with LP-rounding before, eg. [35], we are not aware of an application in context of the group Steiner tree problem; moreover we only use alteration in our analysis and not in the algorithm.

In the next Subsection 4.1, we present the LP relaxation that we use. In Subsection 4.2 we show that the GKR rounding ensures (i) high expected profit and (ii) low expected length (individually). Then in Subsection 4.3 we use the derandomization of GKR rounding [9], and show that it leads to a single (deterministic) solution having a high profit/length ratio. Altogether we obtain an $O(\log^2 n)$-approximation algorithm for ratio KnapRankOrient.

4.1 LP relaxation

At the loss of an $O(\log n)$ factor in the approximation ratio, we can assume that the metric is a tree T (with edge set $E(T)$) rooted at r having $\ell = O(\log n)$ levels [16]. We also enumerate over all choices of the length B of the optimal ratio KnapRankOrient solution.[1] Then we use an LP relaxation similar to the LP for group Steiner tree [17].

First some notation: For any edge e in T, we denote by $\pi(e)$ its parent edge. Similarly $\pi(v)$ is the parent edge of any vertex $v \in V$. For root edges e, the $x_{\pi(e)}$ values are fixed to 1, since the root is always part of the solution. For any $e \in E(T)$, the subtree below edge e is denoted T_e.

$$LP(B) \quad = \quad \max \quad \sum_{i=1}^{m} \sum_{v \in V} w_v^i \cdot z_v^i \tag{7}$$

$$\text{s.t.} \quad x_{\pi(e)} \geq x_e, \qquad \forall\, e \in T \tag{8}$$

$$z_v^i \leq x_{\pi(v)}, \qquad \forall\, v \in V,\, i \in [m] \tag{9}$$

$$\sum_{v \in V} c_v^i \cdot z_v^i \leq 1, \qquad \forall\, i \in [m] \tag{10}$$

$$\sum_{v \in V(T_e)} c_v^i \cdot z_v^i \leq x_e, \qquad \forall\, e \in T,\, i \in [m] \tag{11}$$

$$\sum_{e \in E(T)} d_e \cdot x_e \leq \frac{B}{2} \tag{12}$$

$$\mathbf{0} \leq \mathbf{x}, \mathbf{z} \leq \mathbf{1}$$

Let us show that restricting x and z to integer values gives a valid formulation of KnapRankOrient. In the intended integral solution, x_e is an indicator denoting whether/not edge e is chosen. Constraints (8) ensure monotonicity, that the solution is a subtree rooted at r. Constraints (12) bound the length of the subtree by $B/2$ (so the corresponding Euler tour has length at most B as required). Vertex v is visited by the solution iff $x_{\pi(v)} = 1$; let $U = \{v \in V \,:\, x_{\pi(v)} = 1\}$ denote the vertices visited. For each knapsack rank function $i \in [m]$, variables z^i denote its maximum profit subset $S_i = \{v \in V \,:\, z_v^i = 1\}$. By (9), $S_i \subseteq U$ as required. Moreover (10) ensures that the total size of S_i (in the i^{th} knapsack) is at most one. Finally, the objective (8) is the sum of profits from each knapsack.

Although we do not need constraint (11) to show a valid integer programming formulation, it is crucial in the rounding step. A similar constraint was used in [26] for the related covering Steiner problem. Notice that it indeed holds for integral solutions, so the resulting LP is a valid relaxation of KnapRankOrient. For a fractional solution, (11) says that even *conditional* on edge e being chosen, the total size (in knapsack i) from subtree T_e is at most one.

Algorithm Overview. For each estimate B, we solve the above $LP(B)$ and apply the deterministic rounding algorithm in Subsection 4.3, which guarantees a solution having profit/length ratio at least $\Omega(\frac{1}{\ell}) \cdot \frac{LP(B)}{B}$. Finally we output the best ratio solution amongst all Bs. Note, if B^* denotes the length of the optimal ratio KnapRankOrient solution then $LP(B^*)/B^*$ is at least the optimum ratio. So with $B \approx B^*$ we obtain an $O(\ell)$-approximation to ratio KnapRankOrient.

[1] It suffices to know the length up to a constant factor; so there are only polynomially many choices.

4.2 Expectation guarantee in KnapRankOrient

Here we show that the natural GKR rounding step produces a solution having expected length at most B and expected profit at least $\Omega(LP(B)/\ell)$. Note that this *does not* bound the expectation of profit/length. Still, this property is used by the algorithm in the next subsection to produce a deterministic solution to ratio KnapRankOrient having value $\Omega(\frac{1}{\ell}) \cdot \frac{LP(B)}{B}$.

Algorithm 4.1 LP rounding for KnapRankOrient

1: Solve the LP relaxation $LP(B)$ for KnapRankOrient to obtain (x, z).
2: Perform the (dependent) rounding from [17], i.e. choose each edge $e \in E(T)$ independently with probability $\frac{x_e}{x_{\pi(e)}}$ and retain only subtree $F \subseteq T$ connected to r.
3: **for** $i \in [m]$ **do**
4: For each vertex $v \in V(F)$ choose v into S_i independently with probability $\frac{z_v^i}{x_{\pi(v)}}$.
5: If $c^i(S_i) > 4\ell$ then set $R_i \leftarrow \emptyset$, else $R_i \leftarrow S_i$.
6: **output** r-tour corresponding to F.

In the above algorithm we assume that if e is a root edge then $x_{\pi(e)} = 1$, since the root is always a part of the solution. Steps 2 and 4 are the GKR rounding, and Step 5 is the alteration step. It is clear that in Step 2 we have: $\Pr[e \in F] = x_e$ for each edge e, and so $\Pr[v \in F] = x_{\pi(v)}$ for each vertex v. Note that the expected length of F is:

$$\mathbb{E}\left[\sum_{f\in F} d_f\right] \quad = \quad \sum_e d_e \cdot x_e \quad \leq \quad \frac{B}{2} \tag{13}$$

So taking an Euler tour of F, the expected solution length is at most B. It remains to bound the expected profit. Notice that for each knapsack $i \in [m]$, we have $R_i \subseteq F$ and $c^i(R_i) \leq 4\ell$ with probability one, from Step 5. So a greedy partitioning of R_i yields 8ℓ parts each of size at most one; and by averaging some part has profit at least $w^i(R_i)/(8\ell)$. Thus:

$$\mathbb{E}\left[\sum_{i=1}^m f_i(F)\right] \quad \geq \quad \frac{1}{8\ell}\sum_{i=1}^m \mathbb{E}\left[w^i\left(R_i\right)\right] \quad \geq \quad \frac{1}{8\ell}\cdot\sum_{i=1}^m\sum_v w_v^i \cdot \Pr[v \in R_i] \tag{14}$$

We bound $\Pr[v \in R_i]$ in Lemma 4.1 below. Before doing that, we introduce some notation that will also be useful in the subsequent derandomization step. For any $i \in [m]$ and $u, v \in V$ let I_v^i denote the *indicator* of the event "$v \in S_i$"; and let $I_{u,v}^i$ denote the indicator of event "$u \in S_i$ *and* $v \in S_i$". Also for $i \in [m]$ and $v \in V$ let J_v^i indicate whether "$v \in R_i$". Due to Step 5 it is clear that

$$J_v^i \quad \geq \quad I_v^i - \frac{1}{4\ell}\sum_{u\in V} c_u^i \cdot I_{u,v}^i\,, \qquad \forall i \in [m] \text{ and } v \in V. \tag{15}$$

Lemma 4.1 *For each $i \in [m]$ and $v \in V$,*

$$\Pr\left[v \in R_i\right] \quad = \quad \mathbb{E}[J_v^i] \quad \geq \quad \mathbb{E}[I_v^i] - \frac{1}{4\ell}\sum_{u\in V} c_u^i \cdot \mathbb{E}[I_{u,v}^i] \quad \geq \quad \frac{z_v^i}{4}.$$

Proof: Observe that $\Pr[v \in S_i] = \mathbb{E}[I_v^i] = z_v^i$. Let us now *condition* on $I_v^i = 1$, i.e. $\{v \in S_i\}$. Let $\langle e_1, \ldots, e_\ell\rangle$ denote the edges on the path from r to v; clearly all these edges are in F. For any

$j \in \{1, \ldots, \ell\}$ define $T'_j \subseteq T_{e_j} \setminus \{v\}$ to be those vertices whose least common ancestor with vertex v is edge e_j. Also define T'_0 to be the vertices whose least common ancestor with vertex v is the root r; set $x_{e_0} = 1$. For any $j \in \{0, 1, \ldots, \ell\}$ and $u \in T'_j$ notice that:

$$\mathbb{E}[I_u^i \,|\, I_v^i = 1\,] \quad = \quad \Pr[\,u \in S_i \,|\, v \in S_i\,] \quad = \quad \Pr[\,u \in S_i \,|\, e_j \in F\,] \quad = \quad \frac{z_u^i}{x_{e_j}}$$

Taking expectations, using (11) and $T'_j \subseteq T_{e_j}$ we have for each $j \in [\ell]$,

$$\sum_{u \in T'_j} c_u^i \cdot \mathbb{E}\left[\, I_u^i \,|\, I_v^i = 1\,\right] \quad = \quad \frac{1}{x_{e_j}} \sum_{u \in T'_j} z_u^i \cdot c_u^i \quad \leq \quad 1$$

Observe that $\cup_{j=0}^{\ell} T'_j = V \setminus \{v\}$. So summing the above,

$$\sum_{u \in V \setminus v} c_u^i \cdot \mathbb{E}\left[\, I_u^i \,|\, I_v^i = 1\,\right] \quad = \quad \sum_{j=0}^{\ell} \sum_{u \in T'_j} c_u^i \cdot \mathbb{E}\left[\, I_u^i \,|\, I_v^i = 1\,\right] \quad \leq \quad \ell + 1 \tag{16}$$

By Inequality (15) which is due to the alteration Step 5,

$$\mathbb{E}[\, J_v^i \,|\, I_v^i = 1\,] \;\geq\; 1 - \frac{1}{4\ell} \cdot \sum_{u \in V} c_u^i \cdot \mathbb{E}[I_{u,v}^i \,|\, I_v^i = 1] \;=\; 1 - \frac{1}{4\ell} \cdot \sum_{u \in V} c_u^i \cdot \mathbb{E}[I_u^i \,|\, I_v^i = 1] \;\geq\; 1 - \frac{\ell + 2}{4\ell} \;\geq\; \frac{1}{4}$$

The second last inequality is by (16) and the fact that $c_v^i \leq 1$; the last inequality uses $\ell \geq 1$. The lemma now follows since $\mathbb{E}[I_v^i] = z_v^i$. ∎

Combining (14) and Lemma 4.1 we have:

$$\mathbb{E}\left[\sum_{i=1}^{m} f_i(F)\right] \quad \geq \quad \frac{1}{8\ell} \cdot \sum_{i=1}^{m} \sum_{v} w_v^i \cdot \left(\mathbb{E}[I_v^i] - \frac{1}{4\ell} \sum_{u \in V} c_u^i \cdot \mathbb{E}[I_{u,v}^i]\right) \quad \geq \quad \frac{1}{32\ell} \cdot \sum_{i=1}^{m} \sum_{v} w_v^i \cdot z_v^i \tag{17}$$

Notice that (13) and (17) bound the *expected* length and profit respectively. This is not sufficient for the ratio KnapRankOrient problem. It would suffice to bound the length and profit *simultaneously* (instead of expectation). But this is not possible: Consider the instance in Example 1 and the fractional solution with $x_e = 1/t$ for all edges. This is feasible to the above LP with $B \approx 2$, and has profit $LP(B) \geq 1$. In this example, Algorithm 4.1 produces the following integral solution: the entire tree (profit = length = t), with probability $1/t$, and the empty tree (profit = length = 0) otherwise. Neither of these solutions satisfies bounds on *both* profit and length.

Instead, we show that one can obtain a solution of high ratio "profit/length" by *derandomizing* Algorithm 4.1. This deterministic algorithm uses pessimistic estimators and is similar to [9]; however the details are quite different since we analyze a different random process.

4.3 Deterministic Algorithm for Ratio KnapRankOrient

For the randomized algorithm 4.1 recall the indicator variables I_v^is and $I_{u,v}^i$s. Also let K_f for any edge $f \in E(T)$ denote the indicator of event "$f \in F$". Define the following estimators for profit and length:

$$P \;:=\; \sum_{i=1}^{m} \sum_{v \in V} w_v^i \cdot \left(I_v^i - \frac{1}{4\ell} \sum_{u \in V} c_u^i \cdot I_{u,v}^i\right) \qquad \text{and} \qquad D \;:=\; \sum_{f \in E(T)} d_f \cdot K_f \tag{18}$$

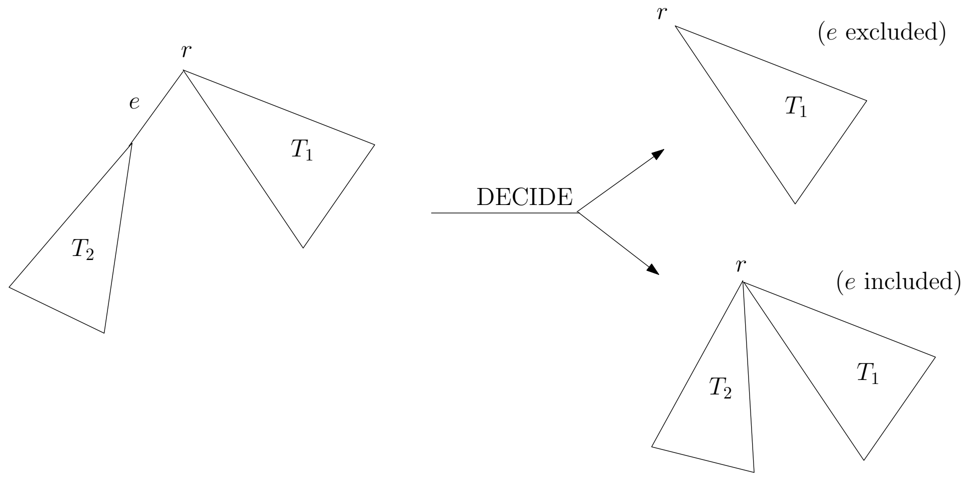

Figure 1: The deterministic edge selection step.

We have $\mathbb{E}[P]/\mathbb{E}[D] \geq \sum_{i=1}^{m} \sum_{v} w_v^i \cdot z_v^i/(4B) = LP(B)/(4B)$ by (17) and (13) which is our initial estimate of the ratio. We will inspect edges e of T one at a time and decide (deterministically) whether/not to include e so that the ratio estimate does not decrease. At the end of the algorithm, we obtain a deterministic subtree F^* with ratio at least $\mathbb{E}[P]/\mathbb{E}[D]$. Details now follow.

To keep notation simple, we extend tree T slightly. For each $v \in V$ add leaves $\{\langle i,v\rangle : i \in [m]\}$ that are adjacent to v (with zero length edges). Let $L_i = \{\langle i,v\rangle : v \in V\}$ denote the leaves corresponding to knapsack $i \in [m]$. Define the following fractional values y^* on edges according to the optimal LP solution (x,z).

$$y_e^* = \begin{cases} x_e & \text{if } e \text{ is an edge in the original tree,} \\ z_v^i & \text{if } e = (v, \langle i,v\rangle) \text{ is a new leaf-edge.} \end{cases}$$

Notice that the GKR rounding steps 2 and 4 in Algorithm 4.1 correspond to choosing each edge e in the modified tree independently w.p. $y_e^*/y_{\pi(e)}^*$ and retaining the subtree F connected to r. (Recall that π maps each edge/vertex to its parent edge.) Moreover, for each $i \in [m]$ subset $S_i \equiv V(F) \cap L_i$.

In our algorithm we will be dealing with trees $\mathcal{T}$ (with vertex set $V(\mathcal{T})$ and edges set $E(\mathcal{T})$) derived from T and edge-weights y (on $\mathcal{T}$) derived from y^*. We will always have the property that y is non-increasing from the root r to any leaf. For any tree $\mathcal{T}$ and edge-weights y, define:

$$P(y,\mathcal{T}) := \sum_{i=1}^{m} \sum_{v \in V(\mathcal{T}) \cap L_i} w_v^i \cdot \left(y_{\pi(v)} - \frac{1}{4\ell} \sum_{u \in V(\mathcal{T}) \cap L_i} c_u^i \cdot \frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}} \right) \quad \text{and} \quad D(y,\mathcal{T}) := \sum_{f \in E(\mathcal{T})} d_f \cdot y_f \quad (19)$$

where $\theta(u,v)$ denotes the least-common-ancestor edge of vertices u and v. For ease of notation, we assume WLOG that there is always a dummy edge above the root r having y-value one. Observe that these values correspond precisely to the expectation of random variables P and D from (18), in tree $\mathcal{T}$ when GKR rounding is performed with edge-values y.

Lemma 4.2 *At any iteration in Algorithm 4.2, after Step 9, we have* $P(y, T \cup F) = y_e \cdot P(y', T') + (1-y_e) \cdot P(y'', T'') = y_e \cdot P_1 + (1-y_e) \cdot P_0$ *and* $D(y, T \cup F) = y_e \cdot D(y', T') + (1-y_e) \cdot D(y'', T'') = y_e \cdot D_1 + (1-y_e) \cdot D_0$.

Proof: Consider first the equation for D. We have:

$$D(y', T') = d_e + \sum_{f \in E(F)} d_f + \sum_{f \in E(T_1)} d_f \cdot y_f + \sum_{f \in E(T_2)} d_f \cdot \frac{y_f}{y_e}$$

Algorithm 4.2 Deterministic algorithm for Ratio $\mathsf{KnapRankOrient}$

Solve the LP relaxation for $\mathsf{KnapRankOrient}$ to obtain (x, z).
Extend tree T by adding leaves $\cup_{i=1}^m L_i$ and define y^* as above.
Initialize tree $F \leftarrow \{r\}$ and $y \leftarrow y^*$.
while there is edge $e \in T$ incident to r **do**
  Let $T_2 = T_e$ denote the subtree below e and $T_1 = T \setminus T_2 \setminus \{e\}$; see Figure 1.
  Set $T'' \leftarrow T_1 \cup F$, and

$$y''_f \leftarrow \begin{cases} y_f & \text{if } f \in T_1 \\ 1 & \text{if } f \in F \end{cases}$$

  Compute estimates $P_0 = P(T'', y'')$ and $D_0 = D(T'', y'')$ upon *excluding* e by (19).
  Set $T' \leftarrow (T \text{ contract } e) \cup (F \cup \{e\})$, and

$$y'_f \leftarrow \begin{cases} \frac{y_f}{y_e} & \text{if } f \in T_2 \cup \{e\} \\ y_f & \text{if } f \in T_1 \\ 1 & \text{if } f \in F \end{cases}$$

  Compute estimates $P_1 = P(T', y')$ and $D_1 = D(T', y')$ upon *including* e by (19).
  if $\frac{P_0}{D_0} \geq \frac{P_1}{D_1}$ **then**
    Set $T \leftarrow T_1$ and $y \leftarrow y''$.
  else
    Set $T \leftarrow (T \text{ contract } e)$ and $y \leftarrow y'$. Also $F \leftarrow F \cup \{e\}$ and $y_e \leftarrow 1$.
output r-tour corresponding to F.

$$D(y'', T'') = \sum_{f \in E(F)} d_f + \sum_{f \in E(T_1)} d_f \cdot y_f.$$

It follows that in the convex combination $y_e \cdot D(y', T') + (1 - y_e) \cdot D(y'', T'')$, each edge $f \in E(F)$ contributes d_f and each edge g in $T = T_1 \cup T_2 \cup \{e\}$ contributes $d_g \cdot y_g$. This is exactly $D(y, T \cup F)$.

Next consider the equation for P. We will show equality term-by-term in the expression (19) for P. Consider the first summation $P_a(y, \mathcal{T}) := \sum_{i=1}^m \sum_{v \in V(\mathcal{T}) \cap L_i} w_v^i \cdot y_{\pi(v)}$. We show that the contribution of any term "$i \in [m]$ and $v \in L_i$" to $P_a(y, T \cup F)$ is the same as to $y_e \cdot P_a(y', T') + (1 - y_e) \cdot P_a(y'', T'')$.

Cases	$P_a(y, T \cup F)$	$P_a(y', T')$	$P_a(y'', T'')$	$y_e \cdot P(y', T') + (1 - y_e) \cdot P(y'', T'')$
$\pi(v) \in E(F)$	w_v^i	w_v^i	w_v^i	w_v^i
$\pi(v) \in E(T_1)$	$w_v^i \cdot y_{\pi(v)}$	$w_v^i \cdot y_{\pi(v)}$	$w_v^i \cdot y_{\pi(v)}$	$w_v^i \cdot y_{\pi(v)}$
$\pi(v) \in E(T_2) \cup \{e\}$	$w_v^i \cdot y_{\pi(v)}$	$w_v^i \cdot \frac{y_{\pi(v)}}{y_e}$	0	$w_v^i \cdot y_{\pi(v)}$

Next consider the second summation $P_b(y, \mathcal{T}) := -\frac{1}{4\ell} \sum_{i=1}^m \sum_{u,v \in V(\mathcal{T}) \cap L_i} w_v^i \cdot c_u^i \cdot \frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$. We again show equality $P_b(y, T \cup F) = y_e \cdot P_b(y', T') + (1 - y_e) \cdot P_b(y'', T'')$ for each term $\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$ corresponding to "$i \in [m]$ and $u, v \in L_i$"; to reduce clutter we drop the multiplier $-\frac{1}{4\ell} w_v^i \cdot c_u^i$.

Cases	$P_b(y, T \cup F)$	$P_b(y', T')$	$P_b(y'', T'')$	$y_e \cdot P_b(y', T')$ $+(1-y_e) \cdot P_b(y'', T'')$
$\pi(u), \pi(v) \in E(F)$	1	1	1	1
$\pi(u) \in E(F)$, $\pi(v) \in E(T_1)$	$y_{\pi(v)}$	$y_{\pi(v)}$	$y_{\pi(v)}$	$y_{\pi(v)}$
$\pi(u) \in E(F)$, $\pi(v) \in E(T_2) \cup \{e\}$	$y_{\pi(v)}$	$\frac{y_{\pi(v)}}{y_e}$	0	$y_{\pi(v)}$
$\pi(u), \pi(v) \in E(T_1)$	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$
$\pi(u), \pi(v) \in E(T_2) \cup \{e\}$	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)} \cdot y_e}$	0	$\frac{y_{\pi(u)} \cdot y_{\pi(v)}}{y_{\theta(u,v)}}$
$\pi(u) \in E(T_1)$, $\pi(v) \in E(T_2) \cup \{e\}$	$y_{\pi(u)} \cdot y_{\pi(v)}$	$y_{\pi(u)} \cdot \frac{y_{\pi(v)}}{y_e}$	0	$y_{\pi(u)} \cdot y_{\pi(v)}$

Since we have checked all cases, it follows that $P(y, T \cup F) = y_e \cdot P(y', T') + (1-y_e) \cdot P(y'', T'')$. ■

This lemma implies that in any iteration,

$$\max \left\{ \frac{P(y', T')}{D(y', T')}, \frac{P(y'', T'')}{D(y'', T'')} \right\} \geq \frac{y_e \cdot P(y', T') + (1-y_e) \cdot P(y'', T'')}{y_e \cdot D(y', T') + (1-y_e) \cdot D(y'', T'')} = \frac{P(y, T \cup F)}{D(y, T \cup F)}$$

So by induction, the ratio $\frac{P(y, T \cup F)}{D(y, T \cup F)}$ is non-decreasing over iterations. It can be seen that the denominator $D(y, T \cup F)$ can always be taken to be non-zero and therefore the final solution F must contain at least one edge. Notice that at the start of Algorithm 4.2, F is empty and the ratio is at least $\rho = LP(B)/(4B)$ by (17) and (13). And at the end of the algorithm, the tree T is empty– so the ratio is exactly:

$$\left(\sum_{i=1}^{m} \sum_{v \in V(F) \cap L_i} w_v^i \cdot \left(1 - \frac{1}{4\ell} \sum_{u \in V(F) \cap L_i} c_u^i \right) \right) \quad / \quad \left(\sum_{f \in E(F)} d_f \right) \quad \geq \quad \rho \tag{20}$$

The denominator is exactly the length of solution tree F. We will show that the numerator is at most $O(\ell)$ times the profit $\sum_{i=1}^m f_i(V(F))$ of F. This would imply that the profit/length ratio of solution F is at least $\Omega(\frac{1}{\ell}) \cdot \rho$ as desired.

To upper bound the numerator in (20), define

$$R_i = \begin{cases} L_i \cap V(F) & \text{if } c^i(L_i \cap V(F)) \leq 4\ell \\ \emptyset & \text{otherwise.} \end{cases} \qquad \forall i \in [m]$$

Note that if $R_i = \emptyset$ then $c^i(L_i \cap V(F)) > 4\ell$, i.e. the contribution of $L_i \cap V(F)$ in the numerator of (20) is negative. On the other hand, if $R_i \neq \emptyset$ then the contribution of $L_i \cap V(F)$ is at most $w^i(L_i \cap V(F)) = w^i(R_i)$. So we obtain that the numerator of (20) is at most $\sum_{i=1}^m w^i(R_i)$. Since each R_i has knapsack-size at most 4ℓ, a greedy partitioning as before implies there is a subset $R_i' \subseteq R_i$ with size $c^i(R_i') \leq 1$ and $w^i(R_i') \geq w^i(R_i)/(8\ell)$; i.e. $f_i(V(F)) \geq w^i(R_i)/(8\ell)$. Rearranging, the numerator of (20) is at most $8\ell \sum_{i=1}^m f_i(V(F))$. Combined with the inequality in (20),

$$\text{Ratio of solution } F \ = \ \frac{f(V(F))}{d(V(F))} \ \geq \ \frac{1}{32\,\ell} \cdot \frac{LP(B)}{B}.$$

Thus we have proved:

Theorem 4.3 *There is a deterministic $O(\ell)$-approximation algorithm for the ratio knapsack orienteering problem on depth ℓ trees. On general metrics there is an $O(\log^2 n)$-approximation algorithm.*

The additional log-factor on general metrics is due to tree embedding [16] which is randomized. This step can also be made deterministic using the algorithm in [9].

5 UGC Hardness of Approximation

In this section we prove a $\omega(1)$ UGC-hardness of approximation for StochVRP even for a very simple star-like metric with a setting of λ that renders the recourse tour trivial. Our hardness result is based on the Unique Games Conjecture (UGC) of Khot [25] , a restatement of which is given below.

Conjecture 5.1 *(Unique Games Conjecture [25]) For any $\varepsilon > 0$, there is a positive integer p such that: given a system of 2-variable linear equations over $\mathbb{Z}_p$, each of the form $x_i - x_j = a_{ij} \mod p$, it is NP-hard to distinguish between the following two cases : (i)* YES CASE: *There is an assignment to the variables that satisfies $1 - \varepsilon$ fraction of equations, (ii)* NO CASE: *Any assigment satisfies at most ε fraction of equations.*

Based on UGC, Bansal and Khot [3] proved the following hardness of approximation result for minimum vertex cover on almost k-partite k-uniform hypergraphs, which shall be the starting point of our reduction.

Theorem 5.2 *[3] Assuming the Unique Games Conjecture, for any $\varepsilon > 0$ and positive integer $k \geq 2$, given a k-uniform hypergraph G with vertex set U and hyperedge set E, it is NP-hard to distinguish between the following two cases:*

YES CASE: *There is a partition of U into $k+1$ disjoint subsets $X, U_1, \dots, U_k$ such that $|X| \leq \varepsilon|U|$ and the hypergraph induced by $U \setminus X$ (consisting of vertex set $U \setminus X$ and hyperedge set $\{e \cap (U \setminus X) \mid e \in E, |e \cap (U \setminus X)| > 0\}$) is k-partite with $U_1, \dots, U_k$ as the k-partition. That is, any hyperedge $e \in E$ has at most one vertex from any U_i. This implies that $X \cup U_i$ is a vertex cover in G for each $i = 1, \dots, k$, and that the minimum vertex cover in G has size at most $(1/k + \varepsilon)|U|$.*

NO CASE: *The size of the maximum independent set in G is at most $\varepsilon|U|$, and therefore the size of the minimum vertex cover in G is at least $(1 - \varepsilon)|U|$.*

In the rest of this section we shall give a hardness reduction from the problem of distinguishing between k-uniform hypergrahs which are almost k-partite (as in the YES case of Theorem 5.2) from those that have a very small maximum independent set (as in the NO case of Theorem 5.2).

5.1 Hardness Reduction

Fix any positive integer $k \geq 2$. Let us suppose we are given a k-uniform hypergraph G on vertex set U and with hyperedge set E as a hard instance from Theorem 5.2, where we shall fix the parameter ε in Theorem 5.2 later. We transform $G(U, E)$ into an instance of StochVRP as follows. For clarity, in this section the nomenclature of "vertices" shall be in context of the hypergraph, while "points" shall be used for corresponding elements in the metric.

Metric (V, d). The set of points V in the metric is $U \cup \{r\}$, where r is the *root*. The distances d are defined as follows. Let $d(r, u) = L$, where $L = (|U|/2k + 1/2)$, for all $u \in U$. Further, for each pair $u, u' \in U$, $u \neq u'$, let $d(u, u') = 1$. It is easy to see that d is a metric. This simple metric can be realized by the shortest paths in a star-like tree of distances as illustrated in Figure 2.

Capacity and Demands. The capacity $Q = 1$ and demands will be $\{0, 1\}$.

Demand Distribution $\mathcal{D}$. There are polynomially many scenarios $m = |E|$, each having uniform probability. Every hyperedge $e \in E$ is a scenario having demand of one at all points in e, and zero demand elsewhere.

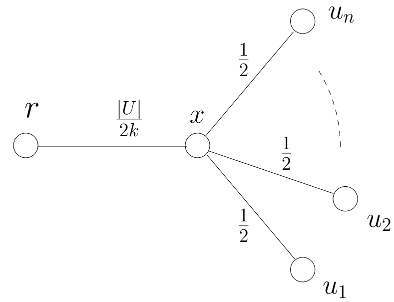

Figure 2: Tree of distances realizing metric d, with intermediate point x and $V = \{r, u_1, \ldots, u_n\}$.

Parameter λ. We set $\lambda = 2m|U|(k+1)$.

Before we proceed to the analysis of this reduction, we note that the cost of the minimum cost r-tour covering points $S \subseteq V \setminus \{r\}$, is simply $|U|/k + |S|$. Also, the optimal value is at most λ/m. Consider the fixed tour consisting of k identical r-tours each covering U: since each scenario has at most k demands, this solution never uses a recourse tour, and has cost $k \cdot (|U|/k + |U|) < \lambda/m$. So we may assume that *the optimal solution has no recourse tour*: if the recourse tour is non-empty in any scenario then its cost is at least λ/m.

5.2 Analysis

We now give the analysis.

YES Case. Suppose that $G(U, E)$ is a YES instance of Theorem 5.2 with $X, U_1, \ldots, U_k$ as the partition of U with the properties as stated in the theorem. Consider the r-tours $\tau_1, \ldots, \tau_k$, where τ_i is an r-tour that covers points $X \cup U_i$ (in addition to r). Since every scenario in our instance of StochVRP corresponds to a hyperedge in G, using the property in the YES case that each hyperedge has at most one vertex from each U_i, we see that the r-tours $\tau_1, \ldots, \tau_k$ satisfy all the scenarios. As noted earlier the cost of each r-tour that covers $S \subseteq V \setminus \{r\}$ is $|U|/k + |S|$. Therefore the total cost of the k r-tours $\tau_1, \ldots, \tau_k$ is,

$$k \cdot (|U|/k) + \sum_{i=1}^{k} |X \cup U_i| \;\leq\; |U| + (1 + k\varepsilon)|U| \;=\; (2 + k\varepsilon)|U|,$$

by the properties of the partition $X, U_1, \ldots, U_k$ of U.

NO Case. Suppose that $G(U, E)$ is a NO instance of Theorem 5.2, so that the maximum independent set in G is of size at most $\varepsilon|U|$. In this case we shall prove that the total cost of any set of r-tours that satisfy all scenarios is at least $k(1 - f_k(\varepsilon))|U|$, where $f_k(\varepsilon) \to 0$ as $\varepsilon \to 0$ for any fixed positive integer $k \geq 2$. We may assume that the number of r-tours in the optimal solution is at most k^2, otherwise the total cost will be at least $k^2(|U|/k) = k|U|$ and we shall be done. Therefore, let $\gamma_1, \ldots, \gamma_T$ be the r-tours in an optimal fixed tour, where $T \leq k^2$. We shall estimate the number of points in U which occur in at most $k-1$ of these r-tours. For any subset $I \subseteq [T]$, let $U(I) \subseteq U$ be the points which *do not* occur in $\{\gamma_i : i \in [T] \setminus I\}$. We have the following simple lemma.

Lemma 5.3 *For any $I \subseteq [T]$ with $|I| = k - 1$, $U(I)$ is an independent set in G.*

Proof: For a contradiction, suppose that e is a hyperedge induced by $U(I)$. Since $|e| = k$, the scenario corresponding to e will not be satisfied by our solution as the k vertices of e appear (as points) in at most $k-1$ of the r-tours, namely those given by $I \subseteq [T]$. Recall that each r-tour can serve only one demand. ■

The total number of points in U that appear in at most $k-1$ of the r-tours is upper bounded by,

$$\sum_{I \subseteq [T], |I|=k-1} |U(I)|.$$

There are $\binom{T}{k-1} \leq 2^T \leq 2^{k^2}$ choices for the subsets I in the above expression. Using the fact that any independent set in G has size at most $\varepsilon|U|$, the fraction of points in U that occur in at most $k-1$ of the r-tours is at most $\varepsilon 2^{k^2} =: f_k(\varepsilon)$. Each of the remaining $(1 - f_k(\varepsilon))|U|$ points appears in at least k of the r-tours; so the total cost of the fixed tour is $k(1 - f_k(\varepsilon))|U|$.

Hardness Factor. In the YES case there is a solution of cost at most $(2 + k\varepsilon)|U|$, whereas in the NO case any solution has cost at least $k(1 - f_k(\varepsilon))|U|$. For any positive integer $k \geq 2$ and arbitrarily small $\delta > 0$, choosing $\varepsilon > 0$ to be small enough in Theorem 5.2, we obtain a hardness factor of $k/2 - \delta$.

6 Independent Demand Distributions

In this section we give an $O(\log(n\lambda)/\operatorname{loglog}(n\lambda))$-approximation for $\mathsf{StochVRP}$ under independent distributions $\mathcal{D}$. That is, the demand $\overline{q}_v$ at each vertex v is independent of all other vertices $V \setminus \{v\}$. The main idea is to show the existence of a near-optimal solution that partitions vertices into two disjoint sets D_1 and D_2 such that: vertices D_1 are served by the fixed tour w.h.p., and vertices D_2 are served in the recourse tour. This step (Lemma 6.1) uses independence. Then we show how an LP-based approach (combined with sampling) yields a constant approximation to the problem of choosing the best partition (D_1, D_2). For any $v \in V$ let $\mu_v := \mathbb{E}[\overline{q}_v]$ the expected demand at v; note $\max_{v \in V} \mu_v \leq Q$ the vehicle capacity. We assume by scaling that the optimal value $\mathsf{opt} \geq 1$.

Lemma 6.1 *Given any instance of* $\mathsf{StochVRP}$ *with independent demands, there exists partition* $D_1 \cup D_2 = V$ *and* $\mathsf{StochVRP}$ *solution with fixed tour* τ *such that:*

- *The total expected demand in each* r*-tour of* τ *is at most* Q.
- *The length of* τ *is* $O(\log(n\lambda)/\operatorname{loglog}(n\lambda)) \cdot \mathsf{opt}$.
- τ *does not visit any* D_2*-vertex; i.e. each* $u \in D_2$ *is served in recourse tour.*
- *Each* $v \in D_1$ *is served by* τ *with probability at least* $1 - 1/(n\lambda)^4$.
- *The recourse cost is at most* $\mathsf{opt} + 1$.

Proof: Consider an optimal fixed tour τ^*. Let $D_1 \subseteq V$ denote the vertices visited at least once in τ^*; note that each vertex might be visited multiple times. Clearly the minimum spanning tree on vertices $D_1 \cup \{r\}$, $\mathsf{MST}(D_1) \leq d(\tau^*) \leq \mathsf{opt}$. Using the "flow lower bound" in VRP [22] it is also clear that:

$$\mathsf{opt} \;\geq\; \frac{1}{Q}\sum_{v \in V} d(r,v)\cdot \mu_v \;\geq\; \frac{1}{Q}\sum_{v \in D_1} d(r,v)\cdot \mu_v \;=\; \mathsf{Flow}(D_1)$$

Recall that each $\mu_v \leq Q$. Thus if we consider a deterministic VRP instance with demands $\{\mu_v : v \in D_1\}$ and capacity Q, then it has a solution τ' of length at most $O(1) \cdot (\mathsf{MST}(D_1) + \mathsf{Flow}(D_1))$ by [22, 2]. From the above, we have $d(\tau') \leq O(1) \cdot \mathsf{opt}$. Let $\tau'_1, \ldots, \tau'_t$ denote the r-tours in τ', each having total μ-value at most Q. We define the fixed tour τ to consist of $\beta := c \cdot \frac{\log(n\lambda)}{\operatorname{loglog}(n\lambda)}$ *copies* of τ', where c is a large enough constant. The first three properties of τ are immediate. For any r-tour τ_i', if the total instantiated demand $\sum_{v\in\tau_i'} \overline{q}_v \le \beta \cdot Q$ then all these demands can be served by τ since it contains β copies of τ_i'. Thus the probability that some $v \in D_1$ (say with $v \in \tau_i'$) is not served by τ is:

$$\Pr[v \text{ not covered by } \tau] \quad \le \quad \Pr\left[\sum_{v\in\tau_i'} \overline{q}_v > \beta \cdot Q\right] \quad \le \quad \frac{1}{(n\lambda)^4},$$

by a Chernoff bound [28] using the fact that $\beta = \Theta\left(\frac{\log(n\lambda)}{\log\log(n\lambda)}\right)$. This proves the fourth property. For the final property, note that vertices D_2 that are never served by the optimal fixed tour τ^*. So the expected VRP value on D_2 (scaled by λ) is at most opt. This is the recourse cost that our solution corresponding to τ pays for D_2. In addition, some D_1-vertices may be uncovered in τ and the recourse cost due to these is at most:

$$\sum_{v\in D_1} \lambda \cdot 2d(r,v) \cdot \Pr[v \text{ not covered by } \tau] \quad \le \quad \frac{2n\lambda D}{(n\lambda)^4} \quad \le \quad 1,$$

where we used the fact that diameter $D = O(n^3)$ from Subsection 3.1. So the total expected recourse cost is at most $\mathsf{opt} + 1$ as claimed. ■

We now find an approximately optimal solution to independent StochVRP that has the above structure. We write an IP formulation to capture the partition (D_1, D_2). For $v \in V$ let $x_v \in \{0,1\}$ denote the indicator that $v \in D_1$. By Lemma 6.1 the fixed tour τ corresponds to a deterministic VRP solution with demands $\{\mu_v \cdot x_v : v \in V\}$. Using MST and flow bounds (as in Section 3) we can express this (losing a constant factor) via linear constraints in x.

We also need to write the expected VRP value (scaled by λ) due to demands D_2. This involves the expected VRP value (equivalently $\mathbb{E}\mathsf{MST} + \mathbb{E}\mathsf{Flow}$) of the random instance where each $v \in V$ has an independent demand of $(1 - x_v) \cdot \overline{q}_v$; recall that $\overline{q}_v$ denotes v's demand in the original StochVRP instance. The expectation $\mathbb{E}\mathsf{Flow}$ is just $\frac{\lambda}{Q}\sum_{v\in V} d(r,v) \cdot \mu_v \cdot (1-x_v)$. Unfortunately it is not clear if one can write linear constraints (in x) for the expectation of MST: this involves the expected MST value when each $v \in V$ is present independently with probability $(1 - x_v) \cdot \Pr[\overline{q}_v > 0]$. Instead we show that sampling can be used to estimate $\mathbb{E}\mathsf{MST}$ within small error, and that the sample expectation can be expressed via linear constraints in x.

For any $x \in \{0,1\}^V$ define $T(x) := \mathbb{E}[\mathsf{MST}(S_x)]$ where S_x contains each vertex $v \in V$ independently w.p. $(1 - x_v) \cdot p_v$ where $p_v := \Pr[\overline{q}_v > 0]$. We now use the result of [10] as in Theorem 3.1. We make $m = \mathit{poly}(n, \lambda)$ independent samples $S^1, \ldots, S^m \subseteq V$ according to $\{p_v\}_{v\in V}$ and set

$$\hat{T}(x) := \frac{1}{m}\sum_{i=1}^{m} \mathsf{MST}\left(\{v \in S^i : x_v = 0\}\right).$$

Then we have $|T(x) - \hat{T}(x)| \le 1$ for all $x \in \{0,1\}^V$ with probability $1 - o(1)$.

We now write the following integer program for finding partition (D_1, D_2).

$$\begin{aligned} \min\ & \sum_{e\in E} d_e \cdot z_e + \sum_{v\in V} \frac{d(r,v)\,\mu_v}{Q} \cdot x_v + \sum_{v\in V} \frac{\lambda\, d(r,v)\,\mu_v}{Q} \cdot (1-x_v) + \frac{\lambda}{m}\sum_{i=1}^{m}\sum_{e\in E} z_e^i \cdot d_e \\ \text{s.t.}\ & \sum_{e\in\delta(R)} z_e \ge x_v \qquad \forall R \subseteq V \setminus \{r\},\ \forall v \in R, \end{aligned}$$

$$\sum_{e\in\delta(R)} z_e^i \geq 1 - x_v \qquad \forall i \in [m],\ \forall R \subseteq V \setminus \{r\},\ \forall v \in R \cap S^i,$$

$$x_v \in \{0,1\} \qquad \forall v \in V,$$

$$z_e,\, z_e^i \geq 0 \qquad \forall e \in E, \forall i \in [m].$$

The last term in the objective captures $\hat{T}(x)$. Based on the preceding discussion, w.h.p. this integer program expresses the objective of *all* partitions (D_1, D_2) up to a constant factor. Relaxing the integrality on x we obtain an LP relaxation that can be solved in polynomial time. The rounding algorithm simply chooses $D_1 = \{v \in V : x_v > \frac{1}{2}\}$ and $D_2 = V \setminus D_1$. We output the fixed tour τ to be $O(\log(n\lambda)/\operatorname{loglog}(n\lambda))$ copies of an approximate VRP on demands $\{\mu_v : v \in D_1\}$. The recourse step involves greedily satisfying the instantiated demands on τ, and then computing an approximate VRP solution on the residual demands. Using Lemma 6.1 it is easy to show that this achieves an $O(\log(n\lambda)/\operatorname{loglog}(n\lambda))$-approximation, i.e. Theorem 1.4.

References

[1] A. Ak and A.L. Erera. A paired-vehicle recourse strategy for the vehicle-routing problem with stochastic demands. *Transportation Science*, 41(2):222–237, 2007.

[2] K. Altinkemer and B. Gavish. Heuristics for unequal weight delivery problems with a fixed error guarantee. *Operations Research Letters*, 6:149–158, 1987.

[3] N. Bansal and S. Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. In *ICALP (1)*, pages 250–261, 2010.

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

[5] D.J. Bertsimas, P. Jaillet, and A.R. Odoni. A priori optimization. *Operations Research*, 38(6):1019–1033, 1990.

[6] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer-Verlag, New York, 1997.

[7] G. Calinescu and A. Zelikovsky. The polymatroid steiner problems. *Journal of Combinatorial Optimization*, 9(3):281–294, 2005.

[8] A.M. Campbell and B.W. Thomas. Probabilistic traveling salesman problem with deadlines. *Transportation Science*, 42(1):1–21, 2008.

[9] M. Charikar, C. Chekuri, A. Goel, and S. Guha. Rounding via trees: Deterministic approximation algorithms for group steiner trees and k-median. In *STOC*, pages 114–123, 1998.

[10] M. Charikar, C. Chekuri, and M. Pál. Sampling bounds for stochastic optimization. In *APPROX-RANDOM*, pages 257–269, 2005.

[11] C. Chekuri, G. Even, and G. Kortsarz. A greedy approximation algorithm for the group steiner problem. *Discrete Applied Mathematics*, 154(1):15–34, 2006.

[12] C. Chekuri and M. Pál. A recursive greedy algorithm for walks in directed graphs. In *FOCS*, pages 245–253, 2005.

[13] B.C. Dean, M.X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008.

[14] M. Dror. Vehicle Routing with Stochastic Demands: Models & Computational Methods. In *Modeling Uncertainty: International Series In Operations Research & Management Science*, volume 46(8), pages 625–649. Springer, 2005.

[15] A.L. Erera, M.W.P. Savelsbergh, and E. Uyar. Fixed routes with backup vehicles for stochastic vehicle routing problems with time constraints. *Networks*, 54(4):270–283, 2009.

[16] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.

[17] N. Garg, G. Konjevod, and R. Ravi. A Polylogarithmic Approximation Algorithm for the Group Steiner Tree Problem. *Journal of Algorithms*, 37(1):66–84, 2000.

[18] A. Gupta, R. Krishnaswamy, V. Nagarajan, and R. Ravi. Approximation Algorithms for Stochastic Orienteering. In *SODA*, pages 245–253, 2012.

[19] A. Gupta, V. Nagarajan, and R. Ravi. Approximation algorithms for vrp with stochastic demands. *Operations Research*, To Appear.

[20] A. Gupta, M. Pál, R. Ravi, and A. Sinha. Sampling and cost-sharing: Approximation algorithms for stochastic optimization problems. *SIAM J. Comput.*, 40(5):1361–1401, 2011.

[21] A. Gupta and A. Srinivasan. An improved approximation ratio for the covering steiner problem. *Theory of Computing*, 2(1):53–64, 2006.

[22] M. Haimovich and A. H. G. Rinnooy Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10(4):527–542, 1985.

[23] N. Immorlica, D.R. Karger, M. Minkoff, and V.S. Mirrokni. On the costs and benefits of procrastination: approximation algorithms for stochastic combinatorial optimization problems. In *SODA*, pages 691–700, 2004.

[24] I. Katriel, C.K. Mathieu, and E. Upfal. Commitment under uncertainty: Two-stage stochastic matching problems. *Theor. Comput. Sci.*, 408(2-3):213–223, 2008.

[25] S. Khot. On the power of unique 2-prover 1-round games. In *STOC*, pages 767–775, 2002.

[26] G. Konjevod, R. Ravi, and A. Srinivasan. Approximation algorithms for the covering steiner problem. *Random Struct. Algorithms*, 20(3):465–482, 2002.

[27] J.K. Lenstra, D.B. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990.

[28] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[29] R. Ravi and A. Sinha. Hedging uncertainty: Approximation algorithms for stochastic optimization problems. *Math. Program.*, 108(1):97–114, 2006.

[30] M.W.P. Savelsbergh and M. Goetschalkx. A comparison of the efficiency of fixed versus variable vehicle routes. *J. Business Logistics*, 16:163–187, 1995.

[31] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Springer-Verlag, Berlin, 2003.

[32] D. Shmoys and K. Talwar. A Constant Approximation Algorithm for the a priori Traveling Salesman Problem. In *IPCO*, pages 331–343, 2008.

[33] D.B. Shmoys and M. Sozio. Approximation algorithms for 2-stage stochastic scheduling problems. In *IPCO*, pages 145–157, 2007.

[34] D.B. Shmoys and C. Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *J. ACM*, 53(6):978–1012, 2006.

[35] A. Srinivasan. New approaches to covering and packing problems. In *SODA*, pages 567–576, 2001.

[36] W. Stewart and B. Golden. Stochastic vehicle routing: A comprehensive approach. *Eur. Jour. Oper. Res.*, 14:371–385, 1983.

[37] P. Toth and D. Vigo. *The vehicle routing problem*. 2001.