



Improved Approximations for Capacitated Vehicle Routing with Unsplittable Client Demands

Zachary Friggstad, Ramin Mousavi^(✉), Mirmahdi Rahgoshay,
and Mohammad R. Salavatipour

Department of Computing Science, University of Alberta, Edmonton, Canada
{zacharyf, mousavih, rahgosha, mrs}@ualberta.ca

Abstract. In this paper, we present improved approximation algorithms for the (unsplittable) Capacitated Vehicle Routing Problem (CVRP) in general metrics. In CVRP, introduced by Dantzig and Ramser (1959), we are given a set of points (clients) V together with a depot r in a metric space, with each $v \in V$ having a demand $d_v > 0$, and a vehicle of bounded capacity Q . The goal is to find a minimum cost collection of tours for the vehicle, each starting and ending at the depot, such that each client is visited at least once and the total demands of the clients in each tour is at most Q . In the unsplittable variant we study, the demand of a node must be served entirely by one tour. We present two approximation algorithms for unsplittable CVRP: a combinatorial $(\alpha + 1.75)$ -approximation, where α is the approximation factor for the Traveling Salesman Problem, and an approximation algorithm based on LP rounding with approximation guarantee $\alpha + \ln(2) + \delta \approx 3.194 + \delta$ in $n^{O(1/\delta)}$ time. Both approximations can further be improved by a small amount when combined with recent work by Blauth, Traub, and Vygen (2021), who obtained an $(\alpha + 2 \cdot (1 - \varepsilon))$ -approximation for unsplittable CVRP for some constant ε depending on α ($\varepsilon > 1/3000$ for $\alpha = 1.5$).

Keywords: Capacitated vehicle routing · Combinatorial optimization · Approximation algorithm

1 Introduction

Vehicle routing problems are among the most well known and well studied problems in Combinatorial Optimization. The goal is generally to find cost-efficient delivery routes for delivering items from depots to clients in a network using vehicles. The Capacitated Vehicle Routing Problem (CVRP), introduced by Dantzig and Ramser in 1959 [12], generalizes the classic Traveling Salesman Problem and has numerous applications. In CVRP, we are given as input a complete graph $G = (V, E)$ with metric edge weights (also referred to as costs)

Z. Friggstad—Supported by an NSERC Discovery Grant and Accelerator Supplement
M.R. Salavatipour—Supported by NSERC

$c(e) \in \mathbb{R}_{\geq 0}$, a depot $r \in V$, and a vehicle with capacity $Q > 0$, and wish to compute a minimum weight/cost collection of tours, each starting and ending at the depot and visiting at most Q customers, whose union covers all the customers. In the more general setting, each node v is given along with a demand $d(v) \in \mathbb{Z}_{\geq 1}$ and the goal is to find a set of tours of the minimum total cost, each of which includes r , such that the union of the tours covers the demand at every client and every tour serves at most Q demand.

There are three common versions of CVRP: *unit*, *splittable*, and *unsplittable*. In the splittable variant, the demand of a node can be delivered using multiple tours so each tour must also specify how much demand it serves at each client¹. However, in the unsplittable variant the entire demand of a client must be delivered by a single tour (e.g. each demand is an indivisible good of a certain size). This obviously requires that $d_v \leq Q$ for all clients v . The unit demand case is a special case of the unsplittable case where every node has a unit demand, and the demand of a client must be delivered by a single tour. It is easy to see that the splittable demand case can be reduced to the unit demand case in pseudo-polynomial time using multiple collocated clients of unit demands. However, the unsplittable version is more challenging. For example, it contains the bin-packing problem as a special case; when all clients are have distance 1 from r and distance 0 from each other.

CVRP has also been referred to as the k -tours problem [3,4]. Both the splittable and unsplittable versions admit constant factor approximation algorithms in polynomial-time. Haimovich and Kan [17] showed that a heuristic, called iterative partitioning, yields an $(\alpha + 1(1 - 1/Q))$ -approximation for the unit demand case if one uses an α -approximation for the Traveling Salesman Problem (TSP). A similar approach produces a $2 + (1 - 2/Q)\alpha$ -approximation for the unsplittable variant [2]. Despite their simplicity, these remained the best approximations for these two variants for over 35 years. Recently, Blauth et al. [10] improved these approximations giving an $(\alpha + 2 \cdot (1 - \varepsilon))$ -approximation algorithm for unsplittable CVRP and a $(\alpha + 1 - \varepsilon)$ -approximation algorithm for unit demand CVRP and splittable CVRP where ε is a constant depending only on α . For $\alpha = 3/2$, they showed $\varepsilon > 1/3000$. All the variants are APX-hard in general metric spaces [25].

In this paper we make significant progress on improving the approximation guarantee for unsplittable CVRP. More specifically we present a simple combinatorial algorithm with ratio 3.25, and then a 3.194-approximation algorithm based on linear programming (LP). Our algorithms are completely independent of the improvements by Blauth et al. [10]. By incorporating their approach, we can further improve both ratios by a small constant $\varepsilon' > 0$. However, for the sake of simplicity we prefer to present our main results without factoring in this last improvement.

Theorem 1. *There is an approximation algorithm for the unsplittable CVRP with ratio $\alpha + 1.75$, where α is the best approximation ratio for TSP.*

¹ One can show using that restricting the demand served to each client by each tour to integer quantities does not change the optimum solution cost.

The running time of this algorithm is dominated by computing two α -approximate TSP tours and a minimum cost matching. For example, using the simple (combinatorial) Christofides-Serdyukov 1.5-approximation we get a combinatorial 3.25-approximation for unsplittable CVRP whose running time is dominated by computing $O(1)$ perfect matchings in graphs with $O(|V|)$ nodes. Computing a perfect matching in a graph with n nodes can be done in $O(n^3)$ time [15]; hence our algorithm runs in $O(|V|^3)$ time.

If we allow greater running time, we can improve the approximation guarantee further by using linear programming.

Theorem 2. *For any $\delta > 0$, there is an approximation algorithm for unsplittable CVRP with ratio $\ln(2) + \alpha + \frac{1}{1-\delta}$ and running time $n^{O(\frac{1}{\delta})}$, where α is the best approximation ratio for TSP.*

Finally, we show how combining these two results with the approach in [10] actually yields further improvements: a combinatorial $(\alpha + 1.75 - \epsilon')$ -approximation and an LP-based $(\alpha + \ln(2) + \frac{1}{1-\delta} - \epsilon')$ -approximation in time $n^{O(\frac{1}{\delta})}$, where $\epsilon' > 0$ is an absolute constant.

It is worth noting as the classical results on CVRP [2, 17], our results also can be extended to the asymmetric metric where $c(u, v)$ is not necessarily equal to $c(v, u)$. For example, the analogous of Theorem 1 in the asymmetric metric is a $(\beta + 1.75)$ -approximation where β is the best approximation factor for Asymmetric Traveling Salesman Problem.

We discuss these further improvements and the extension to the asymmetric metric in more details in the full version of the paper [14].

1.1 Related Work

CVRP captures classic TSP when Q , the vehicle capacity, is at least the total demand of all clients. For general metrics, Haimovich and Kan [17] considered a simple heuristic, called tour partitioning, which starts from a TSP tour and then splits it into tours of size at most Q by making back-and-forth trips to r at certain points along the TSP tour. They showed this gives a $(1 + (1 - 1/Q)\alpha)$ -approximation for splittable CVRP, where α is the approximation ratio for TSP. Essentially the same algorithm yields a $(2 + (1 - 2/Q)\alpha)$ -approximation for unsplittable CVRP [2]. These stood as the best-known bounds until recently, when Blauth et al. [10] showed that given a TSP approximation α , there is an $\epsilon > 0$ such that there is an $(\alpha + 2 \cdot (1 - \epsilon))$ -approximation algorithm for CVRP. For $\alpha = 3/2$, they showed $\epsilon > 1/3000$. They also describe a $(\alpha + 1 - \epsilon)$ -approximation algorithm for unit demand CVRP and splittable CVRP.

For the case of trees, Labbé et al. [23] showed splittable CVRP is NP-hard, and Golden et al. [16] showed unsplittable version is hard to approximate better than 1.5. This is via a simple reduction from bin packing. For splittable CVRP (again on trees), Hamaguchi et al. [18] defined a lower bound for the cost of the optimal solution and gave a 1.5 approximation with respect to the lower bound. Asano et al. [4] improved the approximation to $(\sqrt{41} - 1)/4$ with respect to

the same lower bound and also showed the existence of instances whose optimal cost is exactly $4/3$ times the lower bound. Later, Becker [5] gave a $4/3$ -approximation with respect to the lower bound. Becker and Paul [9] showed a $(1, 1 + \varepsilon)$ -bicriteria polynomial-time approximation scheme for splittable CVRP in trees, i.e. a PTAS but every tour serves at most $(1 + \varepsilon)Q$ demand. Recently, Jayaprakash and Salavatirpour [19] presented a QPTAS for unit-demand CVRP for trees and more generally graphs of bounded treewidth, bounded doubling metrics, or bounded highway dimension. Even more recently, building upon ideas of [9] and [19], Mathieu and Zhou [24] have presented a PTAS for splittable CVRP on trees.

Das and Mathieu [13] gave a quasi-polynomial-time approximation scheme (QPTAS) for CVRP in the Euclidean plane (\mathbb{R}^2). A PTAS for when Q is $O(\log n / \log \log n)$ or Q is $\Omega(n)$ was shown by Asano et al. [4]. A PTAS for Euclidean plane \mathbb{R}^2 for moderately large values of Q , i.e. $Q \leq 2^{\log^\delta n}$ where $\delta = \delta(\varepsilon)$, was shown by Adamaszek et al. [1], building on the work of Das and Mathieu [13]. For high dimensional Euclidean spaces \mathbb{R}^d , Khachay et al. [20] showed a PTAS when Q is $O(\log \log^{1/d} n)$. For graphs of bounded doubling dimension, Khachay et al. [21] gave a QPTAS when the optimal number of tours is $\text{polylog}(n)$ and Khachay et al. [22] gave a QPTAS when Q is $\text{polylog}(n)$.

The next results we summarize are all for the case $Q = O(1)$. CVRP remains APX-hard in general metrics in this case but is polynomial-time solvable on trees. There exists a PTAS for CVRP in the Euclidean plane (\mathbb{R}^2) (again for when Q is fixed) as shown by Khachay et al. [20]. A PTAS for planar graphs was given by Becker et al. [8] and a QPTAS for planar and bounded-genus graphs was then given by Becker et al. [6]. A PTAS for graphs of bounded highway dimension and an exact algorithm for graphs with treewidth tw with running time $O(n^{tw \cdot Q})$ was shown by Becker et al. [7]. Cohen-Addad et al. [11] showed an efficient PTAS for graphs of bounded-treewidth, an efficient PTAS for bounded highway dimension, an efficient PTAS for bounded genus metrics and a QPTAS for minor-free metrics.

Organization of the Paper: We start with definitions and preliminaries in Sect. 2. The proof of Theorem 1 is presented in Sect. 3 and the proof of Theorem 2 is presented in Sect. 4.

2 Preliminaries

For ease of exposition, we assume we have scaled all the demands and the capacity of the vehicle so that the capacity is 1 and each $d(v) \in (0, 1]$ (so demands can be rational numbers). Also, we treat r as a separate node from the rest of the nodes. Formally:

Definition 1 (CAPACITATED VEHICLE ROUTING). *An instance (V, r, c, d) of CAPACITATED VEHICLE ROUTING (CVRP) consists of:*

- a set of clients V , where $|V| = n$,

- a depot r , not in V ,
- metric travel costs/distances $c : (V \cup \{r\}) \times (V \cup \{r\}) \rightarrow \mathbb{R}_{\geq 0}$,
- a demand $d_v \in (0, 1]$ for each customer $v \in V$.

A feasible solution is a collection of tours \mathcal{T} such that

- every tour $T \in \mathcal{T}$ is a cycle containing r ,
- every client belongs to exactly one tour,
- $\sum_{v \in T} d_v \leq 1$ for all $T \in \mathcal{T}$.

The goal is to find a feasible solution with minimum cost where the cost is the sum of costs of the edges in the solution and denoted by $c(\mathcal{T}) := \sum_{T \in \mathcal{T}} c(T) :=$

$$\sum_{T \in \mathcal{T}} \sum_{(u,v) \in T} c(u,v).$$

Observe we are viewing a tour T as both a set of edges comprising a cycle plus the set of endpoints of these edges, so we may use notation like $v \in T$ for a location v and also $(u,v) \in T$ for a pair of locations (u,v) appearing consecutively along the tour T . It is convenient to view the depot r as having $d_r = 0$, for example when we sum the demand of all locations on a tour.

Fix an unsplittable CVRP instance $\mathcal{I} = (V, r, c, d)$ for the rest of this paper. We use OPT to denote an optimal solution for \mathcal{I} and opt the value of this optimal solution.

Definition 2 (Feasible tours). A tour T that spans r and some clients is called feasible for \mathcal{I} if the total demand of the clients in T is at most 1, i.e., $\sum_{v \in T} d_v \leq 1$.

Clients are partitioned into *small* and *big* clients based on a parameter $\delta \in [0, \frac{1}{2}]$, which will be chosen differently for our two algorithms.

Definition 3 (Small and big clients). For a fixed $\delta \in [0, \frac{1}{2}]$, we say a client v is small if $d_v \in [0, \delta]$, and big otherwise.

Let $\mathcal{D} := \sum_{v \in V} 2 \cdot d_v \cdot c(r, v)$. This is historically referred to as the radial lower bound and the following simple well-known lemma has been used often in previous work.

Lemma 1 (Haimovich and Kan [17]). $\mathcal{D} \leq \text{opt}$.

We also define a similar sum for small and big clients separately, i.e., $\mathcal{D}_{small} := \sum_{\substack{v \in V: \\ v \text{ is small}}} 2 \cdot d_v \cdot c(r, v)$, and $\mathcal{D}_{big} := \sum_{\substack{v \in V: \\ v \text{ is big}}} 2 \cdot d_v \cdot c(r, v)$. Also define

$\mathcal{D}'_{big} := \sum_{\substack{v \in V: \\ v \text{ is big}}} 2 \cdot c(r, v)$, which is the cost of serving all big clients using a separate tour for each client.

Given a TSP tour, the algorithm by Haimovich and Kan has the vehicle begin by randomly filling the “tank” of demand it carries with some value

$\theta \sim (0, 1]$. It then travels about the TSP tour: if the tank has insufficient demand to serve a client it travels to the depot to get enough demand to serve the client, returns to serve the client, and then returns to the depot to refill the tank appropriately before resuming the tour. The probability that such a resupply trip is performed when trying to serve a client v is d_v , so the total cost of performing these round trips is at most $2 \cdot \mathcal{D} \leq 2 \cdot \text{opt}$ in expectation.

One of the main driving forces behind our improvements is the following idea. For a small client, if we think of the vehicle’s tank as only holding $1 - \delta$ demand and keep a reserved tank holding demand δ , then if we cannot serve a client with the demand in the main tank, we can serve it using the reserve tank and only make one round trip to the depot to refill both tanks before proceeding. Both of our main algorithms balance this idea with approaches to handling big clients.

We formalize this notion of using a reserve tank in Lemma 2 below. When $\delta = 0$ this gives the same result as in [2, 17].

Lemma 2 (δ -tank Lemma). *Let \mathcal{A} be a TSP tour on $V \cup \{r\}$ and define small and big clients based on a fixed $\delta \leq 1/2$. There is an algorithm that turns \mathcal{A} into a feasible solution for the CVRP instance with cost*

$$c(\mathcal{A}) + \frac{1}{1 - \delta} \cdot \mathcal{D}_{small} + \frac{2}{1 - \delta} \cdot \mathcal{D}_{big} - \frac{\delta}{1 - \delta} \cdot \mathcal{D}'_{big}, \tag{1}$$

and running time $O(n^2)$.

Proof We sketch the high level idea behind the proof. See the full version of the paper [14] for the complete proof.

The idea is to reserve δ portion of the vehicle’s tank and fill out the rest with a random amount. Then, the vehicle visits the vertices in the same order as they appear in \mathcal{A} . As the vehicle visits the clients (vertices), it serves their demand. However, the vehicle might need to make some round trips to the depot to refill. Using the reserved tank and the initial random filling, we bound the cost of these round trips to the depot against different parts of \mathcal{D} as shown in (1). \square

3 A Combinatorial 3.25-Approximation

In this section, we set $\delta := \frac{1}{3}$. So v is a small client if $d_v \leq \frac{1}{3}$ and big if $d_v > \frac{1}{3}$. Note that in any feasible solution, there are at most two big clients in any single tour. Our algorithm tries two things: the first serves only big clients by pairing them up optimally to form these tours and then runs the classic 3.5-approximation on the small clients but using our δ -tank procedure (see Lemma 2) for performing the tour splitting. The other simply runs the 3.5-approximation using δ -tank tour splitting on all clients.

Let us first explain how we use matching. Consider an auxiliary graph $G_{aux} = (V_{big}, E_{aux})$ where $V_{big} \subseteq V$ is the set of all big clients and E_{aux} constructed as follows: for any pair of big clients u, v where $d_u + d_v \leq 1$ we add

and edge between u and v with cost $c(r, u) + c(u, v) + c(v, r)$. Furthermore, for every big client v there is a loop in G_{aux} with cost equal to $2 \cdot c(r, v)$. We compute a min-cost perfect matching² which corresponds to the cheapest way to select tours to serve only the big clients. The precise details are presented in Algorithm 1.

Algorithm 1. $(\alpha + 1.75)$ -approximation

- 1: The first solution is constructed as follows:
 - 2: Compute a min-cost perfect matching M on G_{aux} . Let \mathcal{T}' be the tours corresponding to the edges in M .
 - 3: Compute a TSP tour \mathcal{A} on small clients and r .
 - 4: Apply Lemma 2 to \mathcal{A} with $\delta = \frac{1}{3}$ and let \mathcal{T}'' be the resulting solution.
 - 5: $\mathcal{T} \leftarrow \mathcal{T}' \cup \mathcal{T}''$
 - 6: The second solution is constructed as follows:
 - 7: Compute a TSP tour \mathcal{A} on $V \cup \{r\}$.
 - 8: Apply Lemma 2 to \mathcal{A} with $\delta = \frac{1}{3}$ and let \mathcal{F} be the resulting solution.
 - 9: Return the cheaper of the two solutions \mathcal{T} and \mathcal{F} .
-

3.1 Analysis

We begin with two simple observations.

Lemma 3 $cost(M) \leq \text{opt}$.

Proof Each tour in any feasible solution contains at most two big clients. So, after shortcutting all tours in OPT past small clients, we get tours corresponding to a perfect matching G_{aux} with cost at most opt. \square

Lemma 4 $cost(M) \leq \mathcal{D}'_{\text{big}}$.

Proof Consider all the loops in G_{aux} . The cost of all the loops is exactly $\mathcal{D}'_{\text{big}}$ and this is a matching so it is an upper bound on the minimum cost of a perfect matching. \square

Next, we compute the cost of the first solution in the algorithm. Note that $c(\mathcal{T}') = cost(M)$. Using an α -approximation for TSP, the cost of \mathcal{A} is at most $\alpha \cdot \text{opt}$: again we are using the metric property which shows opt upper bounds the optimum TSP tour since the union of all tours in OPT is connected and Eulerian. Finally, applying the δ -tank lemma to \mathcal{A} results in a solution of cost at most $c(\mathcal{A}) + \frac{1}{1-\delta} \cdot \mathcal{D}_{\text{small}}$ since there is no big client on \mathcal{A} . Overall, we have

$$c(\mathcal{T}) = c(\mathcal{T}') + c(\mathcal{T}'') \leq cost(M) + \alpha \cdot \text{opt} + \frac{3}{2} \cdot \mathcal{D}_{\text{small}}. \tag{2}$$

² A set of edges M that may contain loops is a perfect matching if each node lies in precisely one edge: so a node is either matched with another node via a normal edge or with itself via a loop.

Next, we compute the cost of the second solution. From the δ -tank lemma,

$$c(\mathcal{F}) = \alpha \cdot \text{opt} + \frac{3}{2} \cdot \mathcal{D}_{small} + 3 \cdot \mathcal{D}_{big} - \frac{1}{2} \cdot \mathcal{D}'_{big}. \tag{3}$$

Combining these, we bound the cost of the solution output by the algorithm as follows:

$$\begin{aligned} \min\{c(\mathcal{T}), c(\mathcal{F})\} &\leq \frac{c(\mathcal{T}) + c(\mathcal{F})}{2} \\ &= \frac{2 \cdot \alpha \cdot \text{opt} + 3 \cdot (\mathcal{D}_{small} + \mathcal{D}_{big}) + \text{cost}(M) - \frac{1}{2} \cdot \mathcal{D}'_{big}}{2} \\ &\leq \frac{2 \cdot \alpha \cdot \text{opt} + 3 \cdot \mathcal{D} + \frac{1}{2} \cdot \text{cost}(M)}{2} \\ &\leq \alpha \cdot \text{opt} + 1.5 \cdot \text{opt} + 0.25 \cdot \text{opt} \\ &= (\alpha + 1.75) \cdot \text{opt}, \end{aligned} \tag{4}$$

where the second inequality follows from Lemma 4 and the last inequality follows from Lemmas 1 & 3. This finishes the proof of Theorem 1.

4 An Improved LP-Based Approximation

In this section let δ be a fixed constant in the range $(0, 1/2]$. Smaller δ lead to better approximations with increased, but still polynomial, running times.

Define the small and big clients for this value δ as in Definition 3. Let V_{big} be the set of big clients. We consider the following configuration LP for big clients: Let \mathcal{J} be the set of all feasible tours where each tour consists of some big clients and the depot. Note $|\mathcal{J}|$ is bounded by $n^{O(\frac{1}{\delta})}$ as there can be at most $\frac{1}{\delta}$ big clients in each tour. For each $T \in \mathcal{J}$ let $c(T)$ be the cost of tour T . For each tour $T \in \mathcal{J}$, we have a variable x_T indicating this tour is chosen by the algorithm.

$$\begin{aligned} \text{minimize:} & \quad \sum_{T \in \mathcal{J}} c(T) \cdot x_T && \text{(Configuration-LP)} \\ \text{subject to:} & \quad \sum_{\substack{T \in \mathcal{J}: \\ v \in T}} x_T \geq 1 \quad \forall v \in V_{big} && (5) \\ & \quad x \geq 0 \end{aligned}$$

By shortcutting all tours in the optimum solution past small clients and discarding tours with no big clients, we see there is an integer solution to (Configuration-LP) with cost at most opt . Thus, the optimum LP value provides a lower bound on opt .

Our algorithm independently samples tours spanning large clients using an optimal LP solution. After this, some large clients and all small clients remain uncovered, we cover them using the classic 3.5-approximation but use the δ -tank tour splitting approach. Algorithm 2 contains the full description of our approach. With foresight, we set $\gamma := \ln(2)$.

Algorithm 2. $(3.194 + \frac{1}{1-\delta})$ -approximation

- 1: $\mathcal{T} \leftarrow \emptyset$. {This will be a collection of tours.}
 - 2: Compute an optimal solution x^* of (**Configuration-LP**).
 - 3: **for** $T \in \mathcal{J}$ **do**
 - 4: with probability $\min\{1, \gamma \cdot x_T\}$ add T to \mathcal{T} .
 - 5: Approximate a TSP tour \mathcal{A} spanning $\{r\} \cup (V \setminus V(\mathcal{T}))$ where $V(\mathcal{T})$ is the vertices covered in \mathcal{T} .
 - 6: Apply the δ -tank lemma to \mathcal{A} and let \mathcal{T}' be the resulting collection of tours.
 - 7: Return $\mathcal{T} \cup \mathcal{T}'$.
-

It could be that some clients lie on multiple tours due to the randomized rounding step. One can shortcut the tours past repeated occurrences of clients so each client lies on exactly one tour.

4.1 Analysis

We first bound the probability of a big client not being covered in the randomized rounding step of Algorithm 2 (steps 3–4).

Lemma 5 For a $v \in V_{big}$, $\Pr[v \text{ is not covered by } \mathcal{T}] \leq e^{-\gamma}$.

Proof The event that a big client v is not covered is if we do not sample any tour T that contains v in the randomized rounding step. So

$$\Pr[v \text{ is not covered by } \mathcal{T}] = \prod_{T \in \mathcal{T}} (1 - \gamma \cdot x_T) \leq e^{-\gamma \cdot \sum_{T \in \mathcal{T}: v \in T} x_T} \leq e^{-\gamma},$$

where the last bound follows from the constraint in (**Configuration-LP**) for v . \square

Next, we bound the expected costs of \mathcal{T} and \mathcal{T}' , separately. The cost of \mathcal{T} is bounded as follows:

$$\mathbb{E}[\mathcal{T}] = \gamma \cdot \text{cost}(x^*) \leq \gamma \cdot \text{opt}. \tag{6}$$

Using the δ -tank lemma, we bound the expected cost of \mathcal{T}' but with the following changes: in (1), we drop the negative term and we incorporate the fact that a big client is on \mathcal{A} with probability at most $e^{-\gamma}$, see Lemma 5.

$$\begin{aligned} \mathbb{E}[c(\mathcal{T}')] &\leq c(\mathcal{A}) + \frac{1}{1-\delta} \cdot \mathcal{D}_{small} \\ &\quad + 2 \cdot \sum_{v \in V_{big}} \Pr[v \text{ is not covered by } \mathcal{T}] \cdot \frac{d_v}{1-\delta} \cdot 2 \cdot c(r, v) \\ &= c(\mathcal{A}) + \frac{1}{1-\delta} \cdot \mathcal{D}_{small} + e^{-\gamma} \cdot \frac{2}{1-\delta} \cdot \mathcal{D}_{big} \\ &= c(\mathcal{A}) + \frac{1}{1-\delta} \cdot \mathcal{D}_{small} + \frac{1}{1-\delta} \cdot \mathcal{D}_{big} \\ &= c(\mathcal{A}) + \frac{1}{1-\delta} \cdot \mathcal{D} \leq \alpha \cdot \text{opt} + \frac{1}{1-\delta} \cdot \mathcal{D}. \end{aligned} \tag{7}$$

The second equality follows from our choice of $\gamma = \ln 2$. From (6) and (7), the expected cost of the solution returned by Algorithm 2 is at most

$$\begin{aligned} \mathbb{E}[c(\mathcal{T} \cup \mathcal{T}')] &\leq \ln 2 \cdot \text{opt} + \alpha \cdot \text{opt} + \frac{1}{1-\delta} \cdot \mathcal{D} \\ &\leq \left(\ln 2 + \alpha + \frac{1}{1-\delta}\right) \cdot \text{opt}, \end{aligned} \tag{8}$$

where the last inequality follows from Lemma 1. This finishes the proof of Theorem 2. We briefly comment that this algorithm can be derandomized efficiently using the method of conditional expectation since the probability a big client is covered and its expected contribution to the δ -tank upper bound can be computed efficiently even if some tours have been sampled or rejected so far. Note there is a numerical issue in that $\gamma \cdot x_T$ may not be a rational number, but this error can be absorbed in the $\frac{1}{1-\delta}$ part of the guarantee by choosing δ to be slightly smaller.

Acknowledgements. We thank an anonymous reviewer for pointing out that our approaches could be extended to the asymmetric metric setting.

References

1. Adamaszek, A., Czumaj, A., Lingas, A.: PTAS for k -tour cover problem on the plane for moderately large values of k . In: Dong, Y., Du, D.-Z., Ibarra, O. (eds.) ISAAC 2009. LNCS, vol. 5878, pp. 994–1003. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10631-6_100
2. Altinkemer, K., Gavish, B.: Heuristics for unequal weight delivery problems with a fixed error guarantee. *Oper. Res. Lett.* 6(4), 149–158 (1987). <https://www.sciencedirect.com/science/article/pii/0167637787900125>
3. Arora, S.: Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM* 45(5), 753–782 (1998)
4. Asano, T., Katoh, N., Tamaki, H., Tokuyama, T.: Covering points in the plane by k -tours: Towards a polynomial time approximation scheme for general k . In: 29th ACM Symposium on the Theory of Computing (STOC), pp. 275–283 (1997)
5. Becker, A.: A tight $4/3$ approximation for capacitated vehicle routing in trees. In: 21st International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX), pp. 3:1–3:15 (2018)
6. Becker, A., Klein, P.N., Saulpic, D.: A quasi-polynomial-time approximation scheme for vehicle routing on planar and bounded-genus graphs. In: 25th Annual European Symposium on Algorithms (ESA), pp. 12:1–12:15 (2017)
7. Becker, A., Klein, P.N., Saulpic, D.: Polynomial-time approximation schemes for k -center, k -median, and capacitated vehicle routing in bounded highway dimension. In: 26th Annual European Symposium on Algorithms (ESA), pp. 8:1–8:15 (2018)
8. Becker, A., Klein, P.N., Schild, A.: A PTAS for bounded-capacity vehicle routing in planar graphs. In: 16th International Algorithms and Data Structures Symposium (WADS), pp. 99–111 (2019)
9. Becker, A., Paul, A.: A framework for vehicle routing approximation schemes in trees. In: Friggstad, Z., Sack, J.-R., Salavatipour, M.R. (eds.) WADS 2019. LNCS, vol. 11646, pp. 112–125. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-24766-9_9

10. Blauth, J., Traub, V., Vygen, J.: Improving the approximation ratio for capacitated vehicle routing. In: Singh, M., Williamson, D.P. (eds.) IPCO 2021. LNCS, vol. 12707, pp. 1–14. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-73879-2_1
11. Cohen-Addad, V., Filtser, A., Klein, P.N., Le, H.: On light spanners, low-treewidth embeddings and efficient traversing in minor-free graphs. In: 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS), pp. 589–600 (2020)
12. Dantzig, G.B., Ramser, J.H.: The truck dispatching problem. *Manag. Sci.* **6**(1), 80–91 (1959)
13. Das, A., Mathieu, C.: A quasipolynomial time approximation scheme for euclidean capacitated vehicle routing. *Algorithmica* **73**(1), 115–142 (2015). <https://doi.org/10.1007/s00453-014-9906-4>
14. Friggstad, Z., Mousavi, R., Rahgoshay, M., Salavatipour, M.R.: Improved approximations for cvrp with unsplittable demands (2021). arXiv preprint [arXiv:2111.08138](https://arxiv.org/abs/2111.08138)
15. Gabow, H.N.: An efficient implementation of edmonds’ algorithm for maximum matching on graphs. *J. ACM (JACM)* **23**(2), 221–234 (1976)
16. Golden, B.L., Wong, R.T.: Capacitated arc routing problems. *Networks* **11**(3), 305–315 (1981)
17. Haimovich, M., Kan, A.H.G.R.: Bounds and heuristics for capacitated routing problems. *Math. Oper. Res.* **10**(4), 527–542 (1985)
18. Hamaguchi, S., Katoh, N.: A capacitated vehicle routing problem on a tree. In: Chwa, K.-Y., Ibarra, O.H. (eds.) ISAAC 1998. LNCS, vol. 1533, pp. 399–407. Springer, Heidelberg (1998). https://doi.org/10.1007/3-540-49381-6_42
19. Jayaprakash, A., Salavatipour, M.R.: Approximation schemes for capacitated vehicle routing on graphs of bounded treewidth, bounded doubling, or highway dimension. In: Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms SODA (2022)
20. Khachay, M., Dubinin, R.: PTAS for the euclidean capacitated vehicle routing problem in R^d . In: Kochetov, Y., Khachay, M., Beresnev, V., Nurminski, E., Pardalos, P. (eds.) DOOR 2016. LNCS, vol. 9869, pp. 193–205. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44914-2_16
21. Khachay, M., Ogorodnikov, Y.: QPTAS for the CVRP with a moderate number of routes in a metric space of any fixed doubling dimension. In: Kotsireas, I.S., Pardalos, P.M. (eds.) LION 2020. LNCS, vol. 12096, pp. 27–32. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-53552-0_4
22. Khachay, M., Ogorodnikov, Y., Khachay, D.: An extension of the das and mathieu QPTAS to the case of polylog capacity constrained CVRP in metric spaces of a fixed doubling dimension. In: Kononov, A., Khachay, M., Kalyagin, V.A., Pardalos, P. (eds.) MOTOR 2020. LNCS, vol. 12095, pp. 49–68. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-49988-4_4
23. Labbé, M., Laporte, G., Mercure, H.: Capacitated vehicle routing on trees. *Oper. Res.* **39**(4), 616–622 (1991)
24. Mathiue, C., Zhou, H.: A PTAS for Capacitated Vehicle Routing on Trees (2020). CoRR [arXiv:2111.03735](https://arxiv.org/abs/2111.03735), <https://arxiv.org/abs/2111.03735>
25. Papadimitriou, C.H., Yannakakis, M.: The traveling salesman problem with distances one and two. *Math. Oper. Res.* **18**(1), 1–11 (1993)