

Unsplittable Euclidean Capacitated Vehicle Routing: A $(2 + \epsilon)$ -Approximation Algorithm

Fabrizio Grandoni ✉

IDSIA, USI-SUPSI, Lugano, Switzerland

Claire Mathieu ✉

CNRS, IRIF, Université de Paris, France

Hang Zhou ✉

École Polytechnique, Institut Polytechnique de Paris, France

Abstract

In the unsplittable capacitated vehicle routing problem, we are given a metric space with a vertex called depot and a set of vertices called terminals. Each terminal is associated with a positive demand between 0 and 1. The goal is to find a minimum length collection of tours starting and ending at the depot such that the demand of each terminal is covered by a single tour (i.e., the demand cannot be split), and the total demand of the terminals in each tour does not exceed the capacity of 1.

Our main result is a polynomial-time $(2 + \epsilon)$ -approximation algorithm for this problem in the two-dimensional Euclidean plane, i.e., for the special case where the terminals and the depot are associated with points in the Euclidean plane and their distances are defined accordingly. This improves on recent work by Blauth, Traub, and Vygen [IPCO'21] and Friggstad, Mousavi, Rahgoshay, and Salavatipour [IPCO'22].

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases capacitated vehicle routing, approximation algorithms, Euclidean plane

Digital Object Identifier 10.4230/LIPIcs.ITCS.2023.63

Funding This work was partially funded by the SNSF Grant 200021 200731 / 1, the Grant ANR-19-CE48-0016 from the French National Research Agency (ANR), and the Hi!PARIS Grant “Efficiency in Algorithms”.

Acknowledgements We thank Vincent Cohen-Addad for helpful preliminary discussions. We thank the anonymous reviewers for their valuable comments.

1 Introduction

In the *unsplittable capacitated vehicle routing problem* (*unsplittable CVRP*), we are given a metric space with a vertex O called *depot* and a set of n vertices V called *terminals*. Each terminal $v \in V$ is associated with a positive demand $demand(v) \in (0, 1]$. A feasible solution S is a collection of tours starting and ending at the depot O such that the demand of each terminal $v \in V$ is covered by a *single* tour in S (i.e., the demand of v cannot be split), and, for each tour $t \in S$, the total demand $demand(t)$ of the terminals covered by t does not exceed the *tour capacity* of 1. Our goal is to find a feasible solution S minimizing the total length of the tours $cost(S) := \sum_{t \in S} cost(t)$, where $cost(t)$ denotes the length of t , i.e., the overall weight of the edges of t . As an application, the reader might think about a set of identical vehicles, each with the same capacity, located in a depot; These vehicles have to deliver a set of items at different locations and return to the depot.

Originally introduced by Dantzig and Ramser in 1959 [9], the unsplittable CVRP is arguably one of the most basic problems in Operations Research, and it generalizes famous problems in a natural way. For example, if the sum of the demands is at most 1 (hence



© Fabrizio Grandoni, Claire Mathieu, and Hang Zhou;
licensed under Creative Commons License CC-BY 4.0

14th Innovations in Theoretical Computer Science Conference (ITCS 2023).

Editor: Yael Tauman Kalai; Article No. 63; pp. 63:1–63:13

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

one tour is sufficient), the problem is equivalent to Traveling Salesman Problem (TSP); if all the terminals are placed at the same location, the problem is equivalent to Bin Packing. Since Bin Packing is $(3/2)$ -hard to approximate (i.e., a better than $3/2$ polynomial-time approximation algorithm does not exist unless $P = NP$) [17], the unsplittable CVRP is also $(3/2)$ -hard to approximate.

On general metrics, the first constant-factor approximation algorithm (namely a 3.5 approximation) for the unsplittable CVRP is the *iterated tour partitioning (ITP)* algorithm, which was proposed and analyzed in the 1980s by Haimovich and Rinnooy Kan [11] and Altinkemer and Gavish [2]. The approximation ratio for the unsplittable CVRP was only recently slightly improved by Blauth, Traub, and Vygen [7], and then further improved to roughly 3.194 by Friggstad, Mousavi, Rahgoshay, and Salavatipour [10].

We study the unsplittable CVRP in the *two-dimensional Euclidean plane*, called the *unsplittable Euclidean CVRP*. Here the depot and terminals correspond to points in the Euclidean plane, and for any two vertices $u, v \in V \cup \{O\}$, the weight of the edges $\{u, v\}$ is given by the corresponding Euclidean distance. We remark that the $(3/2)$ -hardness of approximation mentioned before extends to this case (indeed, even to the one-dimensional Euclidean case). The approach in [10] directly implies a 2.694 approximation for the unsplittable Euclidean CVRP, which is also the best known approximation factor for this problem. More precisely, the algorithm in [10] has approximation ratio strictly below $1.694 + \alpha$, where α is the approximation ratio of a TSP algorithm in the considered metric. Since the Euclidean TSP admits a PTAS [4, 15], the algorithm in [10] leads to a 2.694 approximation.

Our main result is a $(2 + \epsilon)$ -approximation algorithm for the unsplittable Euclidean CVRP (Theorem 1). This matches the best known approximation factor for the special case when all the demands are equal [3].

► **Theorem 1.** *Let $\epsilon > 0$. There is a polynomial time $(2 + \epsilon)$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem.*

As a by-product, we design a *polynomial time approximation scheme (PTAS)* for the case when all demands are *big* (Theorem 3). This matches the best known approximation factor for the special case when all the demands are big and equal [11].

► **Definition 2.** *Let $\epsilon > 0$. We say that a terminal is big (w.r.t. ϵ) if its demand is at least ϵ , and it is small otherwise.*

► **Theorem 3.** *Let $\epsilon > 0$. Assuming that all terminals are big w.r.t. ϵ , there is a polynomial time $(1 + \epsilon)$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem.*

Note that the special case with big terminals only remains NP-hard by a reduction from 3-Partition.¹ See Table 1 for a comparison of our results with the closely related previous work.

¹ Recall that in the 3-Partition problem, we are given as input positive integers a_1, \dots, a_{3n}, b , such that $b/4 < a_i < b/2$ for all i , and such that $\sum_{i=1}^{3n} a_i = nb$. We must decide whether there exists a partition of $\{1, \dots, 3n\}$ into n sets T_j such that $\sum_{i \in T_j} a_i = b$ for all $j = 1, \dots, n$. The 3-Partition problem is NP-complete [17].

■ **Table 1** Approximation ratios for the Euclidean CVRP.

	equal demands	unequal demands
arbitrary demands	$2 + \epsilon$ [3]	2.694 [10] $2 + \epsilon$ [this work: Theorem 1] (3/2)-hardness [folklore]
big demands	$1 + \epsilon$ [11]	2.694 [10] $1 + \epsilon$ [this work: Theorem 3] NP-hard [folklore]

1.1 Related Work

1.1.1 Equal Demands

The special case when all the demands are equal is called the *unit demand* version of the Euclidean CVRP. Let the (equal) demand be $1/k$ where k is a positive integer. For general k , the best polynomial-time approximation factor is $2 + \epsilon$ [3]. For small values of k , PTASs are known. Haimovich and Rinnooy Kan [11] described a PTAS when k is constant. Asano et al. [5], extending techniques in [11], obtained a PTAS for $k = O(\log n / \log \log n)$. This was improved to $k \leq 2^{\log^{f(\epsilon)}(n)}$ for some function $f(\epsilon)$ depending on ϵ by Adamaszek, Czumaj, and Lingas [1]. For higher dimensional Euclidean metrics, Khachay and Dubinin [14] gave a PTAS for fixed dimension ℓ and $k = O(\log^{\frac{1}{\ell}}(n))$.

1.1.2 Relation with Heuristics in Practice

From a more practical perspective, we note that the standard setting in which the CVRP arises consists of real-life road networks. Such inputs are related to two theoretical models: Euclidean metrics (because distances between points are not exactly equal to but related to Euclidean distances) and planar graph metrics (because the road network is not exactly a planar graph but is related to planar embedded graphs), thus those are the two theoretical settings which are most relevant in practice. This paper deals with the Euclidean CVRP; we leave the complementary planar graph CVRP for future work.

Many algorithms use a 2-phase *route-first-cluster-second* method, the ITP algorithm being the best-known one. However, the reverse *cluster-first-route-second* method is more common in practice, e.g., used by the famous *OR-tools* developed by Google [16].² As described in [8], “*In a cluster-first-route-second method, customers are first grouped into clusters and the routes are then determined by suitably sequencing customers within each cluster. Different techniques have been proposed for the clustering phase, while the routing phase amounts to solving a TSP.*”

The *cluster-first-route-second* method, although widely used in practice, does not have theoretical guarantees in terms of approximation factor except in special cases. For example, Blauth, Traub, and Vygen [7] considered the worst-case instances for the ITP algorithm in metric spaces, and designed a *cluster-first-route-second* algorithm with a theoretical guarantee for those worst-case instances. In our work, we design another kind of *cluster-first-route-second* algorithm that achieves a theoretical guarantee for general instances in Euclidean metrics. Indeed, existing heuristics of that method construct clusters of size at most equal

² See Lines 79–88 of the code: https://github.com/google/or-tools/blob/stable/ortools/constraint_solver/routing.h.

to the tour capacity, so their routing creates an intra-cluster route for each cluster. Our main idea is to construct finer clusters over the small terminals (of total demand roughly ϵ). We also have an internal routing to connect the terminals within each cluster into an intra-cluster route. In addition, we have an external routing phase so that each route is obtained by combining intra-cluster routes and big terminals.

2 Overview of our Techniques

Our approach uses a new way to cluster small terminals. In addition to that, we combine several ideas from the literature in a non-trivial way, including techniques developed in the framework of Euclidean TSP and Bin Packing.

2.1 Big Terminals Only (Section 4)

Let us first consider the case of big terminals only (i.e., all the demands are at least some constant $\epsilon > 0$).

In the related Bin Packing problem, it is not hard to solve the problem exactly when (1) all items are big, i.e., their sizes are at least some constant $\epsilon > 0$, and (2) there is a constant number of distinct item sizes. Exploiting the adaptive rounding technique, Karmarkar and Karp [12] obtained a PTAS in the special case where all items are big (but the number of distinct sizes is arbitrary).

Our PTAS for the case of big terminals (Theorem 3) is inspired by that approach. We start by observing that the problem can be solved exactly when (1) all demands are big, (2) there is a constant number of distinct demands, and in addition, (3) the terminals are placed at a constant number of distinct locations. Similarly to Bin Packing, up to losing a $1 + \epsilon$ factor in the approximation, with adaptive rounding we get rid of the assumption (2). To get rid of the assumption (3), it is natural to partition the instance into subinstances of *bounded distance* (Definition 4 and Lemma 5) and then discretize the Euclidean plane, similarly to the approach from [1]. This leads to Theorem 3.

2.2 General Terminals (Section 5)

Now we consider the general case of big and small terminals together. To show Theorem 1, we distinguish the following two cases.

Case 1: When the optimum solution OPT has sufficiently many tours (Section 5.1).

As before, we assume bounded distance w.l.o.g. The main novelty of our approach is to *cluster* small terminals into groups with total demand roughly ϵ to reduce to the special case with big terminals only. The clustering process should put together small terminals that are near one another in the plane: to that end, we use a Voronoi cell decomposition of the plane, compute a traveling salesman tour on the small terminals in each cell, and then apply iterated tour partitioning with capacity ϵ to each such tour, with a virtual depot located at the center of the cell. For the analysis, thanks to the bounded distance and to the degree of discretization into cells, we prove (Lemma 17) that the cost of clustering small terminals is at most $(1 + \epsilon) \cdot \text{opt}$, where opt is the optimal cost. Besides, thanks to the “Assignment Lemma” from [6], we prove (Lemma 18) that an optimal solution covering the *clustered* terminals and the big terminals has cost at most $(1 + \epsilon) \cdot \text{opt}$. Since the clustered terminals behave the same as the big terminals, we apply our PTAS for big terminals only to obtain a solution of cost at most $(1 + 2\epsilon) \cdot \text{opt}$ covering the clustered terminals and the big terminals. Together with the cost of clustering small terminals, the overall cost of our solution to the initial instance is at most $(2 + O(\epsilon)) \cdot \text{opt}$.

Case 2: When the optimum solution OPT has a bounded number of tours (Section 5.2). A $2 + O(\epsilon)$ approximation is obtained by a simple adaptation of the techniques in [4, 5].

This completes the proof of Theorem 1.

2.3 Open Questions

To summarize, for the unsplittable Euclidean CVRP, in this paper we have designed algorithms that match the best known results for the special case in which all demands are equal. Can we go beyond that? Looking at Table 1, we see that for the unsplittable Euclidean CVRP, Theorem 1 is not the end of the story: a gap remains in the approximation factor, between $2 + \epsilon$ and $3/2$. A good starting point might be to focus on the special case when all demands are equal (first column of Table 1). It is an interesting open question whether Algorithm 2, the algorithm in Theorem 1, would yield a better-than-2 approximation for that case.

Note that the running time of our algorithms is not practical due to the exponential dependency on $1/\epsilon$. Improving the running time deserves further investigation.

3 Preliminaries

Whenever needed, we will assume w.l.o.g. that $\epsilon > 0$ is upper bounded by a sufficiently small constant. For a subset of terminals $U \subseteq V$, we let $\text{demand}(U) := \sum_{v \in U} \text{demand}(v)$. For any point $p \in \mathbb{R}^2$, we let $\text{dist}(p)$ denote the p -to- O distance in the Euclidean plane.

► **Definition 4** (bounded distance). *Let $C := (1/\epsilon)^{1/\epsilon}$. We say that a set of terminals V has bounded distance if $\frac{D_{\max}}{D_{\min}} \leq C$, where $D_{\max} := \max_{v \in V} \text{dist}(v)$ and $D_{\min} := \min_{v \in V} \text{dist}(v)$.*

► **Lemma 5** (adaptation from Theorem 3 of [1]). *There is a polynomial time algorithm to partition the set V into disjoint subsets $V_1, \dots, V_m \subseteq V$, for some $m \in \mathbb{N}$, such that for each $i \in [1, m]$, V_i has bounded distance, and for any $\rho \geq 1$, ρ -approximate solutions to the instances on V_i 's for all $i \in [1, m]$ yield a $(\rho + \epsilon)$ -approximate solution to the instance on V .*

Proof sketch. We first partition the set of points into subsets, such that each subset has bounded distance. Then we solve the instance on each subset independently. Finally, we return the union of the solutions to each instance. ◀

The following Assignment Lemma due to Becker and Paul [6] will be used in Lemma 18 to combine subtours from different tours, while ensuring that the resulting tours violate the capacity only slightly.

► **Lemma 6** (Assignment Lemma, Lemma 1 in [6]). *Let $G = (A \uplus B, E)$ be a bipartite graph with $E \subseteq A \times B$, where for each $b \in B$, $N(b) \neq \emptyset$ denotes the set of neighbors of b . Each edge $(a, b) \in E$ has a weight $w(a, b) \geq 0$ and each vertex $b \in B$ has a weight $w(b)$ satisfying $0 \leq w(b) \leq \sum_{a \in N(b)} w(a, b)$. Then there exists a function $f : B \rightarrow A$ such that each vertex $b \in B$ is assigned by f to a vertex $a \in N(b)$ and, for each vertex $a \in A$, we have*

$$\sum_{b \in B | f(b)=a} w(b) - \sum_{b \in B | (a,b) \in E} w(a, b) \leq \max_{b \in B} \{w(b)\}.$$

The following lemma, which will be used to prove Lemma 18, was initially due to Altinkemer and Gavish [2] and rephrased by Blauth, Traub, and Vygen [7].

63:6 Unsplittable Euclidean Capacitated Vehicle Routing

► **Lemma 7** ([2, 7]). *Given an instance of the unsplittable capacitated vehicle routing problem with terminals V , depot O , and a traveling salesman tour t_{TSP} on $V \cup \{O\}$, there exists a feasible solution of cost at most $\text{cost}(t_{\text{TSP}}) + \sum_{v \in V} 4 \cdot \text{dist}(v) \cdot \text{demand}(v)$.*

An elementary but critical step in our construction is to partition the plane into a collection of Voronoi cells as follows.

► **Definition 8.** *Let Z denote the set of points in \mathbb{R}^2 whose polar coordinates (r, Θ) satisfy*

$$r = \left(1 + \frac{\epsilon^2}{4} \cdot i\right) \cdot D_{\min}, \text{ where } i \text{ is an integer such that } 0 \leq i < K_1 := \frac{4C}{\epsilon^2}$$

and

$$\Theta = \frac{2\pi}{K_2} \cdot j, \text{ where } j \text{ is an integer such that } 0 \leq j < K_2 := \frac{8\pi C}{\epsilon^2}.$$

The points in Z are called centers. For each $z \in Z$, the cell $c(z)$ of z is the set of points $p \in \mathbb{R}^2$ such that $\text{dist}(p) \in [D_{\min}, D_{\max}]$ and p is closer to z than to any other center (breaking ties arbitrarily). The boundary of $c(z)$ is denoted by $\text{bd}(z)$.

► **Fact 9.** *The number of centers in Z is at most $K := \frac{32\pi \cdot C^2}{\epsilon^4}$. For any center $z \in Z$ and any point $v \in c(z)$, the z -to- v distance is at most $\frac{\epsilon^2 \cdot D_{\min}}{2}$. Furthermore, $\text{bd}(z) \leq \epsilon^2 \cdot D_{\min}$.*

4 Algorithm for Big Terminals Only (Proof of Theorem 3)

Theorem 3 follows directly from Lemma 5 and the following Lemma 10.

► **Lemma 10.** *Let $\epsilon > 0$. Algorithm 1 is a polynomial time $(1 + \epsilon)$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem under the assumption that all terminals are big and the set of terminals has bounded distance.*

■ **Algorithm 1** Algorithm for big terminals only. Let I be the input instance. Let $\beta = \epsilon^2/(4C)$.

-
- 1: Move each big terminal in I to the closest center
 - 2: **for** each center $z \in Z$ that contains at least $1/\beta$ terminals **do**
 - 3: Let $Q(z)$ denote the set of terminals at z
 - 4: Sort the terminals in $Q(z)$ in non-decreasing order of their demands
 - 5: Partition $Q(z)$ into $1/\beta$ groups s.t., each group consists of $\beta \cdot |Q(z)|$ consecutive terminals in the above ordering³
 - 6: Round the demand of each terminal $v \in Q(z)$ to the maximum demand in the group of v
 - 7: Let I' be the resulting instance of the unsplittable Euclidean CVRP
 - 8: Compute an optimal solution S' to I' ▷ Lemma 12
 - 9: **for** each tour $t' \in S'$ and for each terminal v covered by t' **do**
 - 10: Let $z \in Z$ denote the center where v is located
 - 11: Add to t' two copies of the edge between z and the corresponding terminal of v in I
 - 12: **return** the resulting solution S
-

In the rest of the section, we prove Lemma 10.

³ When $\beta \cdot |Q(z)|$ is not an integer, each of the first groups has one less element than each of the remaining groups.

► **Lemma 11.** *Let $\text{opt}(I')$ denote the cost of an optimal solution to the instance I' . We have $\text{opt}(I') \leq (1 + \epsilon) \cdot \text{opt}$.*

Proof. First, we analyze the cost increase due to moving terminals to the closest centers. Consider any tour t in OPT. Since all terminals are big, t contains at most $1/\epsilon$ terminals. By Fact 9, the cost increase of moving each terminal to the closest center is at most $\epsilon^2 \cdot D_{\min}$. Thus the overall cost of moving the terminals on t to the closest centers is at most $\epsilon \cdot D_{\min} \leq \epsilon \cdot \text{cost}(t)/2$. Summing over all tours in OPT, the overall cost increase of moving terminals to the closest centers is at most $(\epsilon/2) \cdot \text{opt}$.

Next, we analyze the extra cost due to the adaptive rounding. The analysis is similar to the one for adaptive rounding for bin packing (Lemma 3.11 in [17]), except that, for each center z , each terminal in the group with the largest demands at z is connected to the depot by adding a separate tour. Consider any center z . If $|Q(z)| < 1/\beta$, then there is no extra cost due to the adaptive rounding; If $|Q(z)| \geq 1/\beta$, the number of terminals at z in the group with the largest demands is at most $\lceil \beta \cdot |Q(z)| \rceil \leq \beta \cdot |Q(z)| + 1 \leq 2\beta \cdot |Q(z)|$. So the extra cost due to the adaptive rounding is at most $2\beta \cdot |Q(z)| \cdot 2D_{\max}$. Summing over all centers z , we have the overall extra cost is at most

$$2\beta \cdot n \cdot 2D_{\max} = (\epsilon^2/(2C)) \cdot n \cdot 2D_{\max} \leq (\epsilon^2/2) \cdot n \cdot 2D_{\min} \leq (\epsilon/2) \cdot \text{opt},$$

using the fact that $\text{opt} \geq \epsilon \cdot n \cdot 2D_{\min}$ (since OPT has at least $\epsilon \cdot n$ tours of length at least $2D_{\min}$ each).

The claim follows by summing the extra costs in both phases. ◀

► **Lemma 12.** *The optimal solution S' for I' can be computed in polynomial time.*

Proof. There are only $O_\epsilon(1)$ centers, so there are only $O_\epsilon(1)$ positions for the terminals in I' . Thanks to the adaptive rounding of the demands, each center has $O_\epsilon(1)$ distinct demands in I' . Hence I' can be described concisely by $O_\epsilon(1)$ pairs of “position” and “demand” with an integer for each pair giving the number of terminals located at that position and with that demand. Moreover, since the demands are big, at most $1/\epsilon$ terminals are covered by each tour, hence there are only $O_\epsilon(1)$ possible distinct types of tours. Thus the instance I' can be solved optimally in polynomial time by exhaustive search. ◀

► **Lemma 13.** $\text{cost}(S) \leq (1 + \epsilon/2) \cdot \text{cost}(S')$.

Proof. Consider a given tour $t' \in S'$, and let t be the corresponding tour in S . Tour t' covers at most $1/\epsilon$ terminals, and the increase of the length of t w.r.t. t' due to each such terminal v is at most $\epsilon^2 \cdot D_{\min} \leq \epsilon^2 \cdot \text{cost}(t')/2$ by Fact 9. The claim follows. ◀

Lemma 10 follows from Lemmas 11–13.

5 Algorithm for General Terminals (Proof of Theorem 1)

Theorem 1 follows directly from Lemma 5 and the following Theorem 14.

► **Theorem 14.** *Let $\epsilon > 0$. Algorithm 2 is a polynomial time $(2 + O(\epsilon))$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem assuming that the set of terminals has bounded distance.*

■ **Algorithm 2** Algorithm for general terminals. Here $\Gamma = K \cdot C/\epsilon = 32\pi C^3/\epsilon^5$.

```

1:  $Y \leftarrow \sum_{v \in V} \text{demand}(v)$ 
2: if  $Y \geq \Gamma$  then
3:    $S \leftarrow$  solution computed by Algorithm 3
4: else
5:    $S \leftarrow$  solution computed by Algorithm 4
return  $S$ 

```

In the rest of this section we prove Theorem 14.

Define $\Gamma = K \cdot C/\epsilon = 32\pi C^3/\epsilon^5$. We compute the overall demand $Y := \sum_{v \in V} \text{demand}(v)$. We assume that $Y \geq 1$.⁴ It is easy to see that Y is a 2-approximation on the number of tours in OPT, i.e., the number of tours in OPT is in $[Y, 2Y]$.⁵ If $Y \geq \Gamma$, then we apply Algorithm 3 (see Section 5.1), otherwise we apply Algorithm 4 (see Section 5.2).

By Theorem 15 (see Section 5.1), the solution computed by Algorithm 3 has cost at most $(2 + O(\epsilon)) \cdot \text{opt}$ when the number of tours in OPT is at least Γ . By Theorem 19 (see Section 5.2), the solution computed by Algorithm 4 has cost at most $(2 + O(\epsilon)) \cdot \text{opt}$ when the number of tours in OPT is at most $2 \cdot \Gamma$. Thus the cost of the solution S is at most $(2 + O(\epsilon)) \cdot \text{opt}$.

5.1 Case 1: OPT Has Sufficiently Many Tours

In this subsection, we prove the following theorem.

► **Theorem 15.** *Let $\epsilon > 0$. Algorithm 3 is a polynomial time $(2 + O(\epsilon))$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem under the assumptions that the set of terminals has bounded distance and the number of tours in OPT is at least Γ .*

We consider a Voronoi-cell decomposition of the plane with respect to the centers Z , where $c(z)$ denotes the cell of z (breaking ties arbitrarily). Let $Q(z)$ be the small terminals contained in $c(z)$. For each cell $c(z)$, we cluster the small terminals $Q(z)$ inside that cell such that each resulting cluster has total demand roughly ϵ . For each cluster of small terminals, we replace them by a single *clustered* terminal at the center z of that cell, whose demand is the total demand of those small terminals. This results in a new instance I' consisting uniquely of big terminals, which can be solved near-optimally by Lemma 10.

How do we achieve a good clustering of the small terminals in each cell $c(z)$? We first compute a traveling salesman tour on the small terminals in that cell, and then greedily partition that tour into segments of consecutive terminals.

► **Fact 16.** *Assume that the number of tours in OPT is at least Γ . We have $K \cdot 2D_{\max} \leq \epsilon \cdot \text{opt}$.*

Proof. $K \cdot 2D_{\max} = \frac{\Gamma \cdot \epsilon}{C} \cdot 2D_{\max} \leq \Gamma \cdot \epsilon \cdot 2D_{\min} \leq \epsilon \cdot \text{opt}$ since OPT uses at least Γ tours of length at least $2D_{\min}$ each. ◀

The following lemma shows that the clustering of small terminals is not too expensive.

► **Lemma 17.** *Let W denote the total cost of the segments and their connections to the centers. Then $W \leq (1 + O(\epsilon))\text{opt}$.*

⁴ This assumption is without loss of generality, since otherwise the problem is equivalent to the Euclidean TSP and thus admits a PTAS [4, 15].

⁵ Indeed, if there are two tours in OPT whose total demand is at most 1, we can combine them into a single tour.

■ **Algorithm 3** Algorithm when OPT has sufficiently many tours. Let I be the input instance.

```

1: for each center  $z \in Z$  do
2:   Let  $Q(z)$  denote the small terminals in the cell  $c(z)$  of  $z$ 
3:   Compute a  $(1 + \epsilon)$ -approximate traveling salesman tour  $t(z)$  on  $Q(z)$  ▷ [4, 15]
4:   Partition  $t(z)$  into segments of consecutive terminals, such that the total demand of
   the terminals in each segment is in  $[\epsilon, 2\epsilon]$ , except possibly the last segment with smaller
   total demand
5:   for each segment of  $t(z)$  do
6:     Let  $d$  be the total demand of the segment
7:     if  $d < \epsilon$  then  $d \leftarrow \epsilon$ 
8:     Create a clustered terminal located at  $z$  and of demand  $d$ 
9: Let  $I'$  be the instance of the unsplittable Euclidean CVRP induced by big and clustered
   terminals
10: Compute an optimal solution  $S'$  to  $I'$  ▷ Lemma 10.
11: for each tour  $t' \in S'$  and for each clustered terminal  $v$  covered by  $t'$  do
12:   Let  $z \in Z$  denote the center where  $v$  is located
13:   Let  $x$  denote the segment of  $t(z)$  corresponding to  $v$ 
14:   Add to  $t'$  the segment  $x$ , plus the edges between the endpoints of  $x$  and the center  $z$ 
15: return the resulting solution  $S$ 

```

Proof. Consider a center $z \in Z$. Let $t^*(z)$ denote an optimal traveling salesman tour on $Q(z)$. Let $\text{opt}(z)$ denote the cost of the part of OPT that is inside the cell $c(z)$ of z . Recall that $\text{bd}(z)$ is the boundary of $c(z)$. We have

$$\text{cost}(t^*(z)) \leq \text{opt}(z) + \frac{3}{2} \cdot \text{bd}(z) \leq \text{opt}(z) + \frac{3}{2} \cdot \epsilon^2 \cdot D_{\min},$$

where the first inequality is due to Karp (Theorem 3 in [13]) and the second inequality is by Fact 9. Since the traveling salesman tour $t(z)$ computed in Algorithm 3 is a $(1 + \epsilon)$ -approximation, we have

$$\text{cost}(t(z)) \leq (1 + \epsilon) \cdot \text{cost}(t^*(z)) \leq (1 + \epsilon) \cdot \left(\text{opt}(z) + \frac{3}{2} \cdot \epsilon^2 \cdot D_{\min} \right).$$

Next, we analyze the costs of the connections to the center z . Observe that for each segment of $t(z)$, the total cost to connect the two endpoints of the segment to z is at most $\epsilon^2 \cdot D_{\min}$ by Fact 9. Since the demand of each segment is at least ϵ excluding possibly one segment, the number of segments is at most $\frac{\text{demand}(Q(z))}{\epsilon} + 1$. Thus the cost of the connections to the center z is at most

$$\left(\frac{\text{demand}(Q(z))}{\epsilon} + 1 \right) \cdot \epsilon^2 \cdot D_{\min}.$$

Let $W(z)$ denote the total cost of the segments in $c(z)$ and their connections to z . We have

$$W(z) \leq (1 + \epsilon) \cdot \left(\text{opt}(z) + \frac{3}{2} \cdot \epsilon^2 \cdot D_{\min} \right) + \left(\frac{\text{demand}(Q(z))}{\epsilon} + 1 \right) \cdot \epsilon^2 \cdot D_{\min}.$$

Summing over all centers $z \in Z$ and since the number of centers is at most K , we have

$$W = \sum_z W(z) \leq (1 + \epsilon) \cdot \text{opt} + \frac{\text{demand}(V)}{\epsilon} \cdot \epsilon^2 \cdot D_{\min} + 3K\epsilon^2 \cdot D_{\min}.$$

63:10 Unsplittable Euclidean Capacitated Vehicle Routing

Observe that $\text{opt} \geq \text{demand}(V) \cdot 2D_{\min}$, since OPT must contain at least $\text{demand}(V)$ tours of length at least $2D_{\min}$ each. Then

$$\frac{\text{demand}(V)}{\epsilon} \cdot \epsilon^2 \cdot D_{\min} \leq \frac{\epsilon \cdot \text{opt}}{2}.$$

By Fact 16,

$$3K\epsilon^2 \cdot D_{\min} \leq 3K \cdot D_{\max} \leq \frac{3\epsilon \cdot \text{opt}}{2}.$$

Therefore, $W \leq (1 + 3\epsilon) \cdot \text{opt}$. ◀

► **Lemma 18.** *Let $\text{opt}(I')$ denote the cost of an optimal solution to the instance I' . We have $\text{opt}(I') \leq (1 + O(\epsilon)) \cdot \text{opt}$.*

Proof. First, we construct a bi-criteria solution \tilde{S} to the instance I' where each tour has total demand at most $1 + 2\epsilon$. Then we transform \tilde{S} into a solution S' to the instance I' without capacity violation, and we show that the cost of S' is at most $(1 + O(\epsilon)) \cdot \text{opt}$.

We say that a clustered terminal is *good* if its corresponding segment has total demand at least ϵ , and *bad* otherwise⁶. We construct the solution \tilde{S} in two phases.

1. Define a bipartite graph $G = (A \uplus B, E)$ as follows. Let A denote the set of tours in an optimal solution to the instance I . Let B denote the set of good clustered terminals. There is an edge in G between a tour $a \in A$ and a clustered terminal $b \in B$ if and only if a covers a small terminal in the segment corresponding to b . We define $w(b)$ as the demand of the segment b . We also define $w(a, b)$ as the part of $w(b)$ which is covered by a . Note that $w(b) = \sum_{a \in N(b)} w(a, b)$ and $N(b) \neq \emptyset$ for each $b \in B$. In addition, $w(b) \leq 2\epsilon$ by the construction of segments in Algorithm 3. We apply the Assignment Lemma (Lemma 6) to obtain a function $f : B \rightarrow A$ such that each clustered terminal $b \in B$ is assigned to exactly one tour $a \in A$ with $(a, b) \in E$ and for each tour $a \in A$, the total demand of the big terminals on a and the clustered terminals that are assigned to a by f is at most $1 + 2\epsilon$.
2. For each bad clustered terminal, we connect it to the depot by a separate tour.

Now we analyze the cost of \tilde{S} . Consider a tour $a \in A$ in the first phase. Recall that each clustered terminal $b \in B$ is good, so has demand at least ϵ . Thus there are at most $1/\epsilon$ clustered terminals $b \in B$ with $f(b) = a$. For each $b \in B$ such that $f(b) = a$, by definition, a contains a small terminal v in the segment corresponding to b . The cost to connect b to a is at most twice the v -to- b distance, which is at most $\epsilon^2 \cdot D_{\min}$ by Fact 9. So the cost to connect b to a over all $b \in B$ such that $f(b) = a$ is at most $\epsilon \cdot D_{\min}$, which is at most $\epsilon \cdot \text{cost}(a)$. Summing over all tours $a \in A$, the overall cost of the connection is at most $\epsilon \cdot \text{opt}$. In the second phase, at most one separate tour is created for each cell. Since the number of cells is at most K , the total connection in this phase has cost at most $K \cdot 2D_{\max}$, which is at most $\epsilon \cdot \text{opt}$ by Fact 16. Therefore, $\text{cost}(\tilde{S}) \leq (1 + 2\epsilon) \cdot \text{opt}$.

Next, we construct the solution S' from the solution \tilde{S} . Consider any tour t in \tilde{S} . If t exceeds the capacity, then we remove from t clustered terminals until t is within the capacity. Let t' be the resulting tour. We include t' into S' . Let R be the set of removed clustered terminals from the tours in \tilde{S} . We create additional depot tours to cover the terminals in R as follows. First, we connect those terminals to the depot by a traveling salesman tour t_{TSP}

⁶ We remark that here we do not consider the rounding up to ϵ of the quantity d in the algorithm.

visiting R , such that the terminals in R located at the same center are visited consecutively by t_{TSP} . Next, we apply Lemma 7 on t_{TSP} to obtain a set of tours within the tour capacity and covering the terminals in R . We include those tours into S' .

From the construction and Lemma 7,

$$\text{cost}(S') \leq \text{cost}(\tilde{S}) + \text{cost}(t_{\text{TSP}}) + \sum_{v \in R} 4 \cdot \text{dist}(v) \cdot \text{demand}(v).$$

Since there are at most K centers, all clustered terminals are located at centers, and using Fact 16,

$$\text{cost}(t_{\text{TSP}}) \leq K \cdot 2D_{\max} \leq \epsilon \cdot \text{opt}.$$

At the same time, we have

$$\begin{aligned} \sum_{v \in R} 4 \cdot \text{dist}(v) \cdot \text{demand}(v) &= \sum_{t \in \tilde{S}} \sum_{v \in R \text{ and } v \in t} 4 \cdot \text{dist}(v) \cdot \text{demand}(v) \\ &\leq \sum_{t \in \tilde{S}} \sum_{v \in R \text{ and } v \in t} 2 \cdot \text{cost}(t) \cdot \text{demand}(v) \leq \sum_{t \in \tilde{S}} 4\epsilon \cdot 2 \cdot \text{cost}(t) = 8\epsilon \cdot \text{cost}(\tilde{S}). \end{aligned}$$

In the first inequality above we used the fact that $2 \cdot \text{dist}(v) \leq \text{cost}(t)$. In the last inequality above we used the fact that $\text{demand}(t) \leq 1 + 2\epsilon$ and $\text{demand}(t') > 1 - 2\epsilon$, where t' is the tour obtained from t by removing a minimal subset of clustered terminals (covered by t) in order to enforce $\text{demand}(t') \leq 1$. Thus

$$\sum_{v \in R \text{ and } v \in t} \text{demand}(v) = \text{demand}(t) - \text{demand}(t') \leq 4\epsilon.$$

Altogether, for $\epsilon > 0$ small enough,

$$\text{cost}(S') \leq \epsilon \cdot \text{opt} + (1 + 8\epsilon)\text{cost}(\tilde{S}) < (1 + 12\epsilon) \cdot \text{opt}. \quad \blacktriangleleft$$

The cost of output solution in Algorithm 3 is at most $W + (1 + \epsilon) \cdot \text{opt}(I')$ by Lemma 10. This is at most $(2 + O(\epsilon)) \cdot \text{opt}$ by Lemmas 17 and 18. Hence Algorithm 3 has approximation ratio $2 + O(\epsilon)$.

5.2 Case 2: OPT Has a Bounded Number of Tours

In this subsection, we prove the following theorem.

► **Theorem 19.** *Let $\epsilon > 0$. Algorithm 4 is a polynomial time $(2 + O(\epsilon))$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem under the assumption that the number of tours in OPT is at most $2 \cdot \Gamma$.*

► **Lemma 20.** *Let $\epsilon > 0$. There is a polynomial time $(1 + \epsilon)$ -approximation algorithm for the unsplittable Euclidean capacitated vehicle routing problem under the assumptions that the number of tours in OPT is at most $2 \cdot \Gamma$ and each demand is an integer multiple of $\frac{1}{2n}$.*

Proof sketch of Lemma 20. The proof is a straightforward adaptation of the algorithm of Asano et al. [5] for the unit-demand case (i.e., for the case when all demands are equal) of the Euclidean CVRP, which in turn is a simple adaptation of Arora's algorithm [4] for the Euclidean TSP.

■ **Algorithm 4** Algorithm when OPT has a bounded number of tours.

-
- 1: Round each demand down to the next smaller integer multiple of $\frac{1}{2n}$, obtaining an instance \hat{I} .
 - 2: Compute a $(1 + \epsilon)$ -approximate solution \hat{S} to \hat{I} ▷ Lemma 20
 - 3: $S \leftarrow \emptyset$
 - 4: **for** each tour t in \hat{S} **do**
 - 5: Make two copies t_1 and t_2 of t
 - 6: Let v_1, \dots, v_ℓ be the terminals covered by t in non-increasing order of their unrounded demands
 - 7: Let i be the largest integer in $[1, \ell]$ such that $\sum_{j=1}^i \text{demand}(v_j) \leq 1$
 - 8: Let t_1 cover the demands of v_1, \dots, v_i
 - 9: Let t_2 cover the demands of v_{i+1}, \dots, v_ℓ
 - 10: $S \leftarrow S \cup \{t_1, t_2\}$
 - 11: **return** S
-

Recall that Arora's algorithm defines a randomized hierarchical quadtree decomposition, such that a near-optimal solution intersects the boundary of each square only $O_\epsilon(1)$ times and those crossings happen at one of a small set of prespecified points, called *portals*, and then uses a polynomial time dynamic program to find the best solution with this structure.

In [5] it was observed that, when the number of tours in OPT is $O_\epsilon(1)$, there is a near-optimal solution in which the overall number of subtours passing through each square (via portals) is $O_\epsilon(1)$. Furthermore, for all equal demands, one can guess the number of terminals covered by each such subtour within a polynomial number of options. This leads to a polynomial number of configurations of subtours inside each square, which ensures the polynomial running time of a natural dynamic program.

Similarly to [5], when the demands are integer multiples of $\frac{1}{2n}$, one can guess the total demand covered by each subtour in the above sense. The details are left to the reader. ◀

Proof of Theorem 19. From Lemma 20, we obtain in polynomial time a set of tours \hat{S} covering the rounded demands of all terminals such that the total cost of the tours in \hat{S} is at most $(1 + \epsilon) \cdot \text{opt}$.

First, we show that the solution S constructed from \hat{S} is feasible for the unsplittable Euclidean CVRP. Consider two tours t_1 and t_2 in S corresponding to a tour $t \in \hat{S}$. Let v_1, \dots, v_ℓ be the terminals covered by t in non-increasing order of their unrounded demands. From the construction, t_1 is within the tour capacity. It suffices to show that t_2 is within the tour capacity. If $\sum_{j=1}^\ell \text{demand}(v_j) \leq 1$, then t_2 covers no demand and is trivially within the tour capacity. Next consider the case when $\sum_{j=1}^\ell \text{demand}(v_j) > 1$. Since the rounding creates an extra demand of at most $\frac{1}{2n}$ at each terminal and the number of terminals ℓ is at most n , we have $\sum_{j=1}^\ell \text{demand}(v_j) \leq 1 + \frac{\ell}{2n} \leq \frac{3}{2}$. To show that t_2 is within the tour capacity, it suffices to show that the overall demand on t_1 is at least $\frac{1}{2}$. There are two cases.

Case 1: $\text{demand}(v_1) \geq \frac{1}{2}$. Then $\text{demand}(t_1) \geq \text{demand}(v_1) \geq \frac{1}{2}$.

Case 2: $\text{demand}(v_1) < \frac{1}{2}$. By the choice of i , we have $\sum_{j=1}^{i+1} \text{demand}(v_j) > 1$. Since $\text{demand}(v_{i+1}) \leq \text{demand}(v_1) < \frac{1}{2}$, we have $\text{demand}(t_1) = \sum_{j=1}^i \text{demand}(v_j) > \frac{1}{2}$.

So both t_1 and t_2 are within the tour capacity. Thus the solution S is feasible.

For each tour $t \in \hat{S}$, each of the two corresponding tours t_1 and t_2 has cost at most the cost of t . Thus $\text{cost}(S) \leq 2 \cdot \text{cost}(\hat{S}) \leq 2(1 + \epsilon) \cdot \text{opt}$.

Hence Algorithm 4 returns a feasible solution to the unsplittable Euclidean CVRP with cost at most $2(1 + \epsilon) \cdot \text{opt}$. ◀

References

- 1 Anna Adamaszek, Artur Czumaj, and Andrzej Lingas. PTAS for k -tour cover problem on the plane for moderately large values of k . *International Journal of Foundations of Computer Science*, 21(06):893–904, 2010.
- 2 Kemal Altinkemer and Bezalel Gavish. Heuristics for unequal weight delivery problems with a fixed error guarantee. *Operations Research Letters*, 6(4):149–158, 1987.
- 3 Kemal Altinkemer and Bezalel Gavish. Heuristics for delivery problems with constant error guarantees. *Transportation Science*, 24(4):294–297, 1990.
- 4 Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM (JACM)*, 45(5):753–782, 1998.
- 5 Tetsuo Asano, Naoki Katoh, Hisao Tamaki, and Takeshi Tokuyama. Covering points in the plane by k -tours: towards a polynomial time approximation scheme for general k . In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 275–283, 1997.
- 6 Amariah Becker and Alice Paul. A framework for vehicle routing approximation schemes in trees. In *Workshop on Algorithms and Data Structures*, pages 112–125. Springer, 2019.
- 7 Jannis Blauth, Vera Traub, and Jens Vygen. Improving the approximation ratio for capacitated vehicle routing. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 1–14. Springer, 2021.
- 8 Jean-François Cordeau, Gilbert Laporte, Martin W.P. Savelsbergh, and Daniele Vigo. Vehicle routing. *Handbooks in operations research and management science*, 14:367–428, 2007.
- 9 George B. Dantzig and John H. Ramser. The truck dispatching problem. *Management Science*, 6(1):80–91, 1959.
- 10 Zachary Friggstad, Ramin Mousavi, Mirmahdi Rahgoshay, and Mohammad R. Salavatipour. Improved approximations for capacitated vehicle routing with unsplittable client demands. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 251–261. Springer, 2022.
- 11 Mordecai Haimovich and Alexander H. G. Rinnooy Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10(4):527–542, 1985.
- 12 Narendra Karmarkar and Richard M. Karp. An efficient approximation scheme for the one-dimensional bin-packing problem. In *23rd Annual Symposium on Foundations of Computer Science*, pages 312–320. IEEE, 1982.
- 13 Richard M. Karp. Probabilistic analysis of partitioning algorithms for the traveling-salesman problem in the plane. *Mathematics of operations research*, 2(3):209–224, 1977.
- 14 Michael Khachay and Roman Dubinin. PTAS for the Euclidean capacitated vehicle routing problem in \mathbb{R}^d . In *International Conference on Discrete Optimization and Operations Research*, pages 193–205. Springer, 2016.
- 15 Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k -MST, and related problems. *SIAM Journal on computing*, 28(4):1298–1309, 1999.
- 16 Laurent Perron and Vincent Furnon. OR-Tools, Google, version 9.3, 2022. <https://developers.google.com/optimization/>.
- 17 David P. Williamson and David B. Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.