

PTAS for k -Tour Cover Problem on the Plane for Moderately Large Values of k^*

Anna Adamaszek¹, Artur Czumaj¹, and Andrzej Lingas²

¹ Centre for Discrete Mathematics and its Applications (DIMAP) and Department of Computer Science, University of Warwick, UK

A.M.Adamaszek@warwick.ac.uk, A.Czumaj@warwick.ac.uk

² Department of Computer Science, Lund University, 22100 Lund, Sweden
Andrzej.Lingas@cs.lth.se

Abstract. Let P be a set of n points in the Euclidean plane and let O be the origin point in the plane. In the k -tour cover problem (called frequently the *capacitated vehicle routing problem*), the goal is to minimize the total length of tours that cover all points in P , such that each tour starts and ends in O and covers at most k points from P .

The k -tour cover problem is known to be \mathcal{NP} -hard. It is also known to admit constant factor approximation algorithms for all values of k and even a polynomial-time approximation scheme (PTAS) for small values of k , $k = \mathcal{O}(\log n / \log \log n)$.

In this paper, we significantly enlarge the set of values of k for which a PTAS is provable. We present a new PTAS for all values of $k \leq 2^{\log^\delta n}$, where $\delta = \delta(\varepsilon)$. The main technical result proved in the paper is a novel reduction of the k -tour cover problem with a set of n points to a small set of instances of the problem, each with $\mathcal{O}((k/\varepsilon)^{\mathcal{O}(1)})$ points.

1 Introduction

The k -tour cover problem (k -TC), is a very natural and well known generalization of the traveling salesperson problem (TSP) to include several tours [2,3,8,12]. Namely, we are given a set P of points (sites), a distinguished point O outside P , called the origin as well as a distance function defined on $P \cup \{O\}$. A tour is a cycle whose vertices are in $P \cup \{O\}$. The length of a tour is the sum of distances between the adjacent points on the tour. The objective is to find a set of tours, each including the origin and at most k points in P , which covers all points in P and achieves the minimum total length.

In Operations Research, the k -TC problem is well known as the *capacitated vehicle routing problem* [12]. The name comes from its standard application when the points in P represent customer locations, and the origin O stands for a depot. Then, a fleet of vehicles located at the depot must serve all the customers, so

* Research supported in part by the *Centre for Discrete Mathematics and its Applications (DIMAP)*, EPSRC award EP/D063191/1, and by VR grant 621-2005-408.

that each vehicle can serve at most k customers. The objective is to minimize the total distance traveled by the fleet. The k -TC problem (capacitated vehicle routing problem) is one of the central special cases of a more general vehicle routing problem, introduced by Dantzig and Ramser [5] fifty years ago, and studied very extensively in the literature ever since (cf. [9,12]).

The k -TC problem contains the TSP problem as a special case and it is known to be \mathcal{NP} -hard for all $k \geq 3$. For this reason, the research on k -TC has focused on heuristic algorithms and approximation algorithms. The most extensively studied variants of k -TC are the metric one, when the distance function is symmetric and satisfies the triangle inequality, and in particular the *two-dimensional Euclidean* one, when the points are placed in the plane and the distance is Euclidean.

The general metric case of k -TC for $k \geq 3$ has been shown to be APX-complete [2], i.e., complete for the class of optimization problems admitting constant factor approximations. However, the approximability status of the two-dimensional *Euclidean k -TC* problem, in particular, the problem of the existence of a PTAS, has not been completely settled yet. One of the first studies of two-dimensional Euclidean k -TC has been due to Haimovich and Rinnooy Kan [8], who presented several heuristics for the metric and Euclidean k -TC, including a PTAS for the two-dimensional Euclidean k -TC with $k < c \log \log n$, for some constant c [8, Section 6]. Asano et al. [3] substantially subsumed this result by designing a PTAS for $k = \mathcal{O}(\log n / \log \log n)$. They also observed that Arora's [1] or Mitchell's [10] PTAS for the two-dimensional Euclidean TSP implies a PTAS for the corresponding k -TC where $k = \Omega(n)$. There has not been any significant progress since the paper by Asano et al. [3] until very recently, when Das and Mathieu [6] showed a *quasi-polynomial* time approximation scheme (QPTAS) for the two-dimensional Euclidean k -TC for every k . Their algorithm combines the approach developed by Arora [1] for Euclidean TSP with some new ideas to deal with k -TC and gives a $(1 + \varepsilon)$ -approximation for the two-dimensional Euclidean k -TC in time $n^{\log^{\mathcal{O}(1/\varepsilon)} n}$ (this bound holds for any k).

In this paper we focus on the two-dimensional Euclidean variant of k -TC. (To simplify the notation, we shall further refer to this variant as to k -TC).

Our **main result** is a new PTAS for k -TC for all values of $k \leq 2^{\log^\delta n}$, where $\delta = \delta(\varepsilon)$. This significantly enlarges the set of values of k for which a PTAS is known. Our PTAS relies on a novel reduction of an instance of k -TC with a set of n points to an instance or a small number of independent instances of the problem with a small number of points. Our first reduction takes any instance of k -TC on n points and reduces it to an instance with $\mathcal{O}((k/\varepsilon)^{\mathcal{O}(1)} \log^2(n/\varepsilon))$ points. Then we present a refinement, where the instance of k -TC is reduced to a small set of instances of k -TC, each with $\mathcal{O}((k/\varepsilon)^{\mathcal{O}(1)})$ points. These results, when combined with the recent QPTAS due to Das and Mathieu [6], give the aforementioned PTAS for k -TC for all values $k \leq 2^{\log^\delta n}$, where $\delta = \delta(\varepsilon)$.

For simplicity of the presentation, we will present $(1 + \mathcal{O}(\varepsilon))$ -approximation algorithms; reduction to $(1 + \varepsilon)$ -approximation is straightforward.

2 Preliminaries

We assume a fixed origin in the plane and denote it by O . For a tour \mathcal{T} , its (Euclidean) length is denoted by $|\mathcal{T}|$. For a set U of tours, we set $|U|$ to $\sum_{\mathcal{T} \in U} |\mathcal{T}|$.

For a set P of points in the plane, we denote by $TSP(P)$ the minimum length of a TSP-tour through P and by $OPT(P)$ the minimum length of a solution to k -TC (i.e., the minimum length of a set of tours, each through the origin and containing at most k points of P , which covers all points in P). When P is clear from the context, we shall simply use the notation OPT .

For a point $p \in P$, we denote by $r(p)$ the distance of p from the origin O .

The following simple lower bound plays a very important role in the previous approaches to k -TC, see [3, Proposition 2] and [8, Lemma 1].

Fact 1. $OPT(P) \geq \frac{2}{k} \sum_{p \in P} r(p)$.

Following [3], we shall term $\frac{2}{k} \sum_{p \in P} r(p)$ as the *radial cost* of P , and denote by $rad(P)$. Among other things, Haimovich and Kan considered the so called *iterated tour partitioning heuristic* for k -TC in [8]. The heuristic starts from constructing a TSP-tour T through P . Then, it considers all k -tour covers resulting from partitioning T into paths visiting exactly k points (assuming that n is divisible by k), and connecting the endpoints of the paths with O . The heuristic outputs the shortest among these solutions.

Fact 2. [3] *If the iterated tour partitioning heuristic uses a TSP tour U , then it returns a k -tour cover of total length not exceeding $(1 - \frac{1}{k}) \cdot |U| + rad(P)$.*

Note that given a TSP tour, the iterated tour partitioning heuristic can be implemented in time $\mathcal{O}(k\frac{n}{k} + n)$ by repeatedly updating the previous partition and k -tour cover to the next one in time $\mathcal{O}(\frac{n}{k})$. Using the minimum spanning tree heuristic for TSP we can find a 2-approximation of the TSP in time $\mathcal{O}(n \log n)$. Hence, we obtain the following.

Corollary 1. *If the iterated tour partitioning heuristic uses the minimum spanning tree heuristic for TSP then it returns a $(3 - \frac{2}{k})$ -approximation of an optimal k -tour cover of an n -point set and it can be implemented in time $\mathcal{O}(n \log n)$.*

3 PTAS for Moderate Values of k

In this section we present a reduction that takes as an input any instance of the k -tour cover problem on a set of n points in the Euclidean plane and reduces it to an instance of the problem with $\mathcal{O}((k/\varepsilon)^{\mathcal{O}(1)} \log^2(n/\varepsilon))$ points. Then, we apply this reduction to obtain a PTAS for the k -tour cover problem for all $k \leq 2^{\log^\delta n}$, where δ is some positive constant, $\delta = \delta(\varepsilon)$.

Our construction uses a series of transformations that eliminate most of the input points and reduce the input problem instance to one significantly smaller.




Fig. 1. The structure of circles, rays, and locations. The point labeled O is the origin. Other fat dots represent the points from P . In the right picture each point has been moved to its nearest location.

3.1 Removing Close Points

Let L be the maximum distance from a point in P to the origin O , that is, $L = \max\{p \in P : r(p)\}$. Since $\text{OPT} \geq 2L$, we can ignore any point that is at a distance at most $L\varepsilon/n$ from the origin: covering all such points with 1-tours will give us additional cost not greater than $n \cdot 2\frac{L\varepsilon}{n} \leq \varepsilon \cdot \text{OPT}$. Therefore, from now on, we will consider only the points p with $r(p) \geq L\varepsilon/n$.

3.2 Circles, Rays, and Locations

Let us create *circles* around the origin, the i -th circle with a radius

$$c_i = \frac{L\varepsilon}{n} \cdot \left(1 + \frac{\varepsilon}{k}\right)^i, \quad \text{for } 0 \leq i \leq \left\lceil \log_{(1+\varepsilon/k)} \frac{n}{\varepsilon} \right\rceil.$$

Let us draw *rays* from the origin with the angle between any pair of neighboring rays equal to $2\pi/s$ (that is, partition the space into s sectors) with $s = \lceil \frac{2\pi k}{\varepsilon} \rceil$.

Define a *location* to be any point on the plane that is the intersection of a circle and a ray. Since

$$\log_{(1+\varepsilon/k)} \frac{n}{\varepsilon} = \frac{\log \frac{n}{\varepsilon}}{\log(1 + \varepsilon/k)} = \Theta\left(\frac{k}{\varepsilon} \cdot \log \frac{n}{\varepsilon}\right),$$

there are $\Theta\left(\frac{k}{\varepsilon} \log(n/\varepsilon)\right)$ circles and $\Theta\left(\frac{k}{\varepsilon}\right)$ rays. Therefore we obtain:

Claim 1. *The total number of locations T satisfies $T = \Theta(k^2\varepsilon^{-2} \log(n/\varepsilon))$.*

Now, we modify P by moving each point from P to its nearest location.

Claim 2. *The operation of moving each point to its nearest location can change the cost of a k -tour cover by at most $\varepsilon \cdot \text{OPT}$.*

Proof. Let p be a point in P . Suppose that p lies between the circles with radius c_i and c_{i+1} (the distance between p and the origin is in the interval $[L\varepsilon/n, L]$, so we know such circles exist). The distance between these circles equals $c_{i+1} - c_i = \frac{\varepsilon}{k} \cdot c_i$. The distance between two consecutive locations at the i -th circle is less than $2\pi c_i/s \leq \frac{\varepsilon}{k} \cdot c_i$. Therefore the distance between p and its nearest location is at most $\sqrt{2} \cdot (\frac{1}{2} \cdot \frac{\varepsilon}{k} c_i) < \frac{\varepsilon}{k} \cdot c_i \leq \frac{\varepsilon}{k} \cdot r(p)$.

If we move a point $p \in P$ by a distance at most $\frac{\varepsilon}{k} \cdot r(p)$, the cost of a tour can change by at most $2\frac{\varepsilon}{k} \cdot r(p)$. If we add up the changes of the cost generated by moving all points in P , then this total change is upper bounded by $\sum_{p \in P} 2\frac{\varepsilon}{k} \cdot r(p)$. Next, we use Fact 1 to conclude that the total cost of moving all the points is at most $\varepsilon \cdot \text{OPT}$. □

From a k -tour cover U' for a modified instance of the problem (where all points have been moved to their nearest locations) we can easily get a k -tour cover U for the original version of the problem such that $|U| \leq |U'| + \varepsilon \cdot \text{OPT}$. So a PTAS for the modified version yields a PTAS for the original version. In the rest of this paper we will consider the modified version of the problem.

3.3 Trivial and Nontrivial Tours

We say that a tour *visits* a location if it contains at least one point from that location. (If an edge passes through a location, but the tour does not contain any point from that location, then the tour does not visit that location.) We call a tour *trivial* if it visits only a single location in P ; it is *nontrivial* otherwise.

Theorem 1. *There is an optimal solution in which there are at most T non-trivial tours.*

Proof. We say that a set of tours t_1, t_2, \dots, t_m ($m \geq 2$) forms a *cycle* if there is a set of locations $\ell_1, \ell_2, \dots, \ell_m, \ell_{m+1} = \ell_1$ such that each tour t_i visits locations ℓ_i and ℓ_{i+1} . Note that the origin is not considered as a location.

To prove our theorem we will need the following:

Lemma 1. *There is an optimal solution in which there are no cycles.*

Proof. Let U be such an optimal solution which minimizes the sum over all its nontrivial tours of the number of locations visited by that tour.

Let us suppose that U has a cycle, and let t_1, t_2, \dots, t_m be a minimal cycle (m is minimal). Let $\ell_1, \ell_2, \dots, \ell_m$ be the locations in which the consecutive tours meet. From the minimality of the cycle we know that both tours and locations are pairwise distinct.

Let $v(t, \ell)$ denote the number of points from a location ℓ visited by a tour t . Let $\min = \min_{i \in \{1, \dots, m\}} \{v(t_i, \ell_i)\}$. Now we are ready to swap points between the tours: the i -th tour, instead of visiting $v(t_i, \ell_i)$ points in the location ℓ_i and $v(t_i, \ell_{i+1})$ points in the location ℓ_{i+1} will now visit $(v(t_i, \ell_i) - \min)$ points in ℓ_i and $(v(t_i, \ell_{i+1}) + \min)$ points in ℓ_{i+1} . Here ℓ_{m+1} denotes ℓ_1 .

Observe that the modification does not change the number of points visited by each tour. It also does not increase the length of any tour. Therefore, we

obtain another optimal solution, in which the sum over all nontrivial tours of the number of locations visited by that tour is smaller than in U (we managed to remove one location from each tour t_i for which $v(t_i, \ell_i) = \min$). This is a contradiction with the minimality of that sum in U .

Therefore the optimal solution U has no cycles. □

Consider an optimal solution without cycles. Note that the lack of 2-cycles means that no two tours visit the same pair of locations. To each nontrivial tour we can assign a pair of distinct locations visited by this tour. The chosen pairs are in one-to-one correspondence with the nontrivial tours and they induce an acyclic undirected graph on the locations.

Hence, we can have at most $T - 1$ nontrivial tours in an acyclic solution, so using Lemma 1 we have proved the theorem. □

3.4 Reduction to an Instance of k -TC with $(k \log n/\varepsilon)^{O(1)}$ Points

Observe that Theorem 1 implies that there is an optimal solution in which at most Tk points are covered by nontrivial tours. Therefore it is enough to consider only solutions which fulfill that property.

If the number of points in a location ℓ is greater than Tk , some of the points will have to be covered by trivial tours. We may assume, without loss of generality, that among all trivial tours visiting a given location there is at most one that visits less than k points. Moreover, if at least one point from some location is visited by a nontrivial tour, we can assume that all trivial tours visiting that location contain exactly k elements. Therefore, for each location ℓ containing c_ℓ points, we only have to consider at most $\min\{c_\ell, c_\ell - k \cdot \lceil \frac{c_\ell - Tk}{k} \rceil\} \leq Tk$ points for nontrivial tours. After finding a $(1 + \varepsilon)$ -approximation for such reduced case, we will add trivial tours covering all remaining points. That will give us $(1 + \varepsilon)$ -approximation for the original problem.

Corollary 2. *One can reduce the k -TC problem on n points to one on at most T^2k points.*

3.5 PTAS for k -TC with $k \leq 2^{\log^\delta n}$

We use Corollary 2 to reduce any instance of k -TC with the input set of n points P to an instance of k -TC with $N = T^2k = \Theta(k^5 \varepsilon^{-4} \log^2(n/\varepsilon))$ input points. For such input instance, we apply the quasi-polynomial time approximation scheme for k -TC due to Das and Mathieu [6]. The obtained algorithm returns a $(1 + \varepsilon)$ -approximation in time $N^{\log^{O(1/\varepsilon)} N}$. This gives polynomial time for all $k \leq 2^{\log^\delta n}$ for some constant $\delta = \delta(\varepsilon) > 0$. Hence, we have the following main theorem.

Theorem 2. *There is a PTAS for the k -TC problem provided that $k \leq 2^{\log^\delta n}$ for some positive constant $\delta = \delta(\varepsilon)$.*

4 Refinement: Reduction to $(k/\varepsilon)^{\mathcal{O}(1)}$ Points

In the preceding section, we have demonstrated that the problem of close approximation of the k -TC problem on the input set of n points in the plane reduces to that for a multi-point-set of size polynomial in k/ε and polylogarithmic in n in the relevant locations. In this section, we shall eliminate the polylogarithmic dependency of n in the reduction. This will have only a relatively small effect on the asymptotics for the size of the largest k in terms of n for which we can attain a PTAS and we will obtain a PTAS for all $k \leq 2^{\log^{\delta'} n}$, where comparing to the bound in Theorem 2, we will have $\delta' > \delta$. However, for small values of k this will lead to a faster PTAS. Hopefully, because it removes completely the dependency on n from the size of the reduced instance, it also might be a step towards a PTAS for arbitrary values of k .

Our approach resembles Baker’s method [4] of closely approximating several hard problems on planar graphs. It relies on the following separation lemma.

Lemma 2. *Let P be a set of points situated in the locations and let $\varepsilon > 0$. There is a clustering of the circles into rings of $\lceil \log_{1+\frac{\varepsilon}{k}}(6/\varepsilon) \rceil$ consecutive circles and there are positive integers $a = \mathcal{O}(\varepsilon^{-1})$ and $b \in \{1, \dots, a\}$ such that if we mark each $(b+ja)$ -th ring then any k -tour cover U of P can be transformed to a k -tour cover U' of the points in the unmarked rings such that*

1. *no tour in U' visits two points in P separated by a marked ring, and*
2. *$|U'| \leq (1 + \frac{\varepsilon}{2})|U|$.*

Furthermore, the points in the marked rings can be covered with k -tours of total length at most $\frac{\varepsilon}{2}|U|$ produced by the iterated tour partitioning heuristic from [8] (cf. Section 2).

Proof. Let t denote a tour obtained by removing its edges incident to O . Suppose that t crosses one of the marked rings. Let i be the number of the most inner circle of the ring. Denote the circle by C_i . It follows by straightforward calculation and the definition of the circles that each minimal fragment of t crossing the aforementioned ring is at least $\frac{2}{\varepsilon}$ times longer than the doubled radius of C_i . We can appropriately split the tour t along C_i into smaller ones by connecting pairs of crossing points on C_i with O or just with themselves, see Figure 2.

The total length of the smaller tours is longer than $|t|$ by at most $\frac{\varepsilon}{2}$ of the total length of the aforementioned fragments of t .

We may assume, without loss of generality, that the aforementioned marked ring is the outermost among those crossed by t . We can iterate the elimination of the crossings of the smaller resulting tours but for their edges incident to O with more inner marked rings. Note that then other disjoint fragments of t will be charged with the increase of the length of the union of the resulting smaller tours. Finally, by applying short-cutting, we can drop the points in the marked rings from the resulting tours.

We conclude that we can transform U into a k -tour cover U' of the points in P in the unmarked rings such that no tour in U' crosses any marked ring (but for its edges incident to O) and $|U'| \leq (1 + \frac{\varepsilon}{2})|U|$.




Fig. 2. Splitting t into smaller tours. The grey area is the marked ring. In the left picture dotted lines represent the lines which will be added to our solution. The right picture shows two separate tours obtained from the original tour (one is marked with a dashed line, and the other with a solid one), before the short-cutting.

It remains to show that we can set a and $b \in \{1, \dots, a\}$ such that one can easily cover the points in P contained in the marked rings with k -tours of total length not exceeding $\frac{\varepsilon|U|}{2}$.

Let R_j denote the set of points from P lying in the j -th ring. Set a to $\lceil \frac{24}{\varepsilon} \rceil$. For each $b \in \{1, \dots, a\}$, let P_b be the set of points in P in the marked rings, $P_b = \sum_{j \equiv b \pmod a} R_j$. We shall show that there is some $b \in \{1, \dots, a\}$ such that by applying the k -TC heuristic given in Corollary 1 for P_b , we can cover P_b with k -tours of length at most $\frac{\varepsilon|U|}{2}$. For this purpose, we shall observe that $\sum_j TSP(R_j) \leq 3 \cdot TSP(P)$.

Suppose for the sake of this observation that the tour t considered in the first part of the proof is an n -tour, i.e., an optimal TSP tour of $P \cup \{O\}$. Apply almost the same transformation to the tour t as before with the exception that instead of connecting the outer cut part by two rays to O , we connect the cutting points directly. By the triangle inequality, the total length of the so modified TSP tour t is at most $(1 + \frac{\varepsilon}{2}) \cdot TSP(P)$. The modified TSP tour t can be easily reduced to the non-necessarily optimal TSP tours of the unmarked regions by short-cutting. Assuming first for a moment that the unmarked rings are the even ones, and then conversely, that the unmarked rings are the odd ones, and that $\varepsilon < \frac{1}{2}$, we conclude that $\sum_j TSP(R_j) \leq 3 \cdot TSP(P)$.

Using Fact 2 we get that

$$\begin{aligned} \sum_{b \in \{1, \dots, a\}} \text{OPT}(P_b) &\leq \sum_{b \in \{1, \dots, a\}} \sum_{j \equiv b \pmod a} \text{OPT}(R_j) = \sum_j \text{OPT}(R_j) \\ &\leq \sum_j (\text{rad}(R_j) + TSP(R_j)) \leq \text{rad}(P) + 3 \cdot TSP(P) \leq 4|U|. \end{aligned}$$

There must be some $b \in \{1, \dots, a\}$ such that $\text{OPT}(P_b) \leq \frac{4}{a}|U| \leq \frac{\varepsilon|U|}{6}$. Thus, if we apply the 3-approximation algorithm for the k -tour cover of P_b , which

is a composition of the iterated tour partitioning heuristic with the minimum spanning tree heuristic for TSP, we obtain a k -tour cover of P_b of length at most $\frac{\varepsilon|U|}{2}$. \square

Theorem 3. *The k -TC problem on a set P of n points on the plane can be reduced to a collection of $\mathcal{O}(\varepsilon^{-1} \log(n/\varepsilon)/\log(1/\varepsilon))$ disjoint k -tour cover problems, each on $\mathcal{O}(k^5\varepsilon^{-6} \log^2(1/\varepsilon))$ -point set and each having the maximum distance to the origin at most $(1/\varepsilon)^{\mathcal{O}(1/\varepsilon)}$ larger than the minimum one, such that $(1 + \varepsilon)$ -approximate solutions to each of the latter problems yield a $(1 + \mathcal{O}(\varepsilon))$ -approximation to the original k -tour cover problem. The reduction can be done in time $\mathcal{O}(n \log n)$ for a fixed ε .*

Proof. Move the points to the locations and compute the sets R_j of input points lying in the rings for a fixed ε . This all can be easily done in time $\mathcal{O}(n \log n)$ by using standard data structures for point location [11].

Next, compute the value a (the distance between marked rings) and for each $b \in \{1, \dots, a\}$, compute a 3-approximate k -tour cover of the set P_b of points contained in the marked rings. All the a computations take $\mathcal{O}(an \log n) = \mathcal{O}(n \log n)$ time by Corollary 1.

Fix b to that minimizing the length of the aforementioned tour. It follows from Lemma 2 that the produced cover of P_b has length at most $\frac{\varepsilon}{2}\text{OPT}$. Now we will have to compute approximate solutions for each maximal sequence of consecutive not marked rings. Let us denote the number of such sequences by q . We can easily compute that $q = \mathcal{O}(\varepsilon^{-1} \log \frac{n}{\varepsilon} / \log \frac{1}{\varepsilon})$. For $i = 1, \dots, q$, let I_i denote the set of points contained in such i -th sequence. Note that these point sets can be also easily computed in time $\mathcal{O}(n \log n)$.

It follows from Lemma 2 that if we compute separately $(1 + \varepsilon)$ -approximation of the optimal cover with k -tours for each set I_i , then the union of these coverings will have length at most $(1 + \mathcal{O}(\varepsilon))\text{OPT}$.

Note that for a given i , the number of locations in I_i is $\mathcal{O}(a \cdot \frac{k}{\varepsilon} \cdot \log_{(1+\frac{\varepsilon}{k})} \frac{1}{\varepsilon}) = \mathcal{O}(k^2\varepsilon^{-3} \log \frac{1}{\varepsilon})$. Hence, by the discussion in Section 3, we can account to the intended $(1 + \varepsilon)$ -approximation of $\text{OPT}(I_i)$ the trivial tours decreasing the point-multiplicity in each location to $\mathcal{O}(k^3\varepsilon^{-3} \log \frac{1}{\varepsilon})$. Thus, for each I_i we can reduce the problem to one with $\mathcal{O}(k^5\varepsilon^{-6} (\log \frac{1}{\varepsilon})^2)$ points.

Each I_i consists of $\mathcal{O}(\varepsilon^{-1})$ consecutive rings and for a point in a ring the maximum distance to the origin is at most $\mathcal{O}(\varepsilon^{-1})$ times larger than the minimum one. Hence, for a point in I_i the maximum distance to the origin is at most $(1/\varepsilon)^{\mathcal{O}(1/\varepsilon)}$ times larger than the minimum one.

The appropriate q sets of points can be computed in time $\mathcal{O}(n \log n)$ and they specify the problems to which we approximately reduce the original k -tour cover problem. \square

5 Final Remarks

The central open question left is whether there is a PTAS for the k -TC problem for all values of k . While we have enlarged the set of values of k for which a

PTAS exists, we still do not know how to reach polynomial values for k , even $k = n^{0.001}$. In particular, a PTAS k -TC for $k = \Theta(\sqrt{n})$ is elusive. For arbitrary values of k , the best currently known result is either a quasi-polynomial time approximation scheme by Das and Mathieu [6] that runs in time $n^{\log^{\mathcal{O}(1/\varepsilon)} n}$, or the polynomial-time constant-factor approximation algorithm due to Haimovich and Rinnooy Kan [8]. Similarly as in [3], we believe that the case $k = \Theta(\sqrt{n})$ is the hardcore of the difficulty in obtaining a PTAS for all values of k .

Following [8], let us observe that if we divide the range of k into a logarithmic number of intervals of the form $[\varepsilon^{-2i}, \varepsilon^{-2(i+1)})$, then for k in at most one of the intervals none of the inequalities $TSP(P) \leq \varepsilon \cdot rad(P)$, $rad(P) \leq \varepsilon \cdot TSP(P)$ hold. Note that if any of the inequalities holds then by plugging any PTAS for TSP in the iterated tour partitioning heuristic, we obtain an $(1 + \mathcal{O}(\varepsilon))$ -approximation of k -TC. Thus, the aforementioned heuristic is a PTAS for a substantial range of k depending on P : for every set of points P there is k_0 such that there is a polynomial-time $(1 + \mathcal{O}(\varepsilon))$ -approximation algorithm for k -TC for every k with $k_0/\varepsilon < k < \varepsilon k_0$. Despite this observation and despite recent progress in [3,6], the problem of designing a PTAS for *all* k remains open: we believe that our paper sheds the light on this problem and is a step towards a PTAS for arbitrary k .

References

1. Arora, S.: Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM* 45(5), 753–782 (1998)
2. Asano, T., Katoh, N., Tamaki, H., Tokuyama, T.: Covering points in the plane by k -tours: a polynomial time approximation scheme for fixed k . IBM Tokyo Research Laboratory Research Report RT0162 (1996)
3. 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: *Proc. 29th Annual ACM Symposium on Theory of Computing*, pp. 275–283 (1997)
4. Baker, B.S.: Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM* 41(1), 153–180 (1994)
5. Dantzig, G.B., Ramser, R.H.: The truck dispatching problem. *Management Science* 6(1), 80–91 (1959)
6. Das, A., Mathieu, C.: A quasi-polynomial time approximation scheme for Euclidean capacitated vehicle routing. In: *Proc. 21st Annual ACM-SIAM Symposium on Discrete Algorithms* (2010)
7. Garey, M.R., Johnson, D.S.: *Computers and Intractability*. In: *A Guide to the Theory of NP-completeness*, W.H. Freeman and Company, New York (1979)
8. Haimovich, M., Rinnooy Kan, A.H.G.: Bounds and heuristics for capacitated routing problems. *Mathematics of Operation Research* 10(4), 527–542 (1985)
9. Laporte, G.: The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research* 59(3), 345–358 (1992)
10. Mitchell, J.S.B.: 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)
11. Preparata, F., Shamos, M.: *Computational Geometry – an Introduction*. Springer, New York (1985)
12. Toth, P., Vigo, D.: *The Vehicle Routing Problem*. SIAM, Philadelphia (2001)