Approximation Results for Kinetic Variants of TSP

Mikael Hammar[1] and Bengt J. Nilsson[2]

[1] Department of Computer Science, Lund University,
Box 118, S-221 00 Lund, Sweden
e-mail: mikael.hammar@cs.lth.se

[2] School of Technology and Society, Malmö University College,
Citadellsvägen 7, 205 06 Malmö, Sweden
e-mail: bengt.nilsson@ts.mah.se

Abstract. We study the approximation complexity of certain kinetic variants of the Traveling Salesman Problem in the plane where we consider instances in which each point moves with a fixed constant speed in a fixed direction. We prove the following results.

1. If the points all move with the same velocity, then there is a PTAS for the Kinetic TSP.
2. The Kinetic TSP cannot be approximated better than by a factor of two by a polynomial time algorithm unless P=NP, even if there are only two moving points in the instance.
3. The Kinetic TSP cannot be approximated better than by a factor of $2^{\Omega(\sqrt{n})}$ by a polynomial time algorithm unless P=NP, even if the maximum velocity is bounded. The n denotes the size of the input instance.

1 Introduction

Consider a cat in a field with an ample supply of mice. The cat's objective is to catch all the mice while exerting the minimum amount of energy. The cat therefore wishes to use the shortest possible path to chase the mice. A major difficulty is the fact that the mice are moving. This problem is an instance of the *Kinetic Traveling Salesman Problem*.

The *Traveling Salesman Problem*, TSP for short, is probably the best known intractable problem. It asks for the shortest closed tour that visits the nodes in a given weighted complete graph exactly once. This deceptively simple problem lies at the heart of combinatorial optimization and it has spawned a wealth of research in complexity theory and operations research; see Lawler et al. [9].

TSP optimization has been shown to be NP-hard, even for restricted instances. Specifically, if the graph is metric or even embeddable in the Euclidean plane, the problem is NP-hard [7, 12].

Even among NP-hard problems, the TSP is considered to be a computationally difficult problem. It cannot be approximated to within *any* constant factor by a polynomial time algorithm, unless P=NP [14]. For metric graphs the situation is better. Christofides [6] presents a polynomial time $\frac{3}{2}$-approximation algorithm, but better than a constant factor approximation algorithm is not possible in polynomial time unless P=NP [3].

Arora [1] and (independently) Mitchell [11], in their seminal papers, showed that there exist *polynomial time approximation schemes* for TSP when the graph is embedded in the Euclidean plane. Polynomial time approximation schemes, PTAS for short, are polynomial time algorithms that, for any $\epsilon > 0$, produce a $(1 + \epsilon)$-approximation to a given problem. The running time of a PTAS is polynomial in the input size, for any fixed ϵ. Recently, the running times of PTAS's for d-dimensional Euclidean TSP has been significantly improved [2, 13].

In the kinetic traveling salesman problem, we look at TSP for moving points in the Euclidean plane. We consider instances in which each point moves with a fixed velocity. This is a natural and both theoretically and practically important generalization of TSP (e.g. several scheduling problems can be reduced to solving variants of kinetic TSP).

Previous Work This research topic evolved out of a problem posed by Mordecai Golin in the early 1990s. Golins suggestion was the TSP for moving points on the line and it was solved using a $O(n^2)$ time dynamic programming algorithm in 1998 by Helvig, Robins and Zelikovsky [8]. They also gave a $2+\epsilon$-approximation algorithm for the Kinetic TSP if the number of points with non-zero speed is $O(\frac{\log n}{\log \log n})$, and they also considered different versions of the k-delivery problem. Previously, in 1996, Chalasani et al. gave a constant factor approximation algorithm for the Kinetic TSP when all points have the same velocity. Their research into this area was justified by its application in the framework of item collection on conveyor belts. They also considered different versions of the k-delivery problem but a major disadvantage of all their solutions is that their algorithms operate in the L_1-metric, which implies that they cannot benefit from the PTAS of Arora et al. that surfaced into the research community one year after these results were published.

The Kinetic TSP has also been addressed by researchers in the operations research field using linear programming relaxations [10], although their main approach is to discretizise the problem, resulting in a Time Dependent TSP rather than a TSP for moving points.

Results We prove the following results.

1. If the points all move with the same speed and in the same direction then there is a PTAS for the Kinetic TSP. This improves the result of Chalasani et al. [4, 5]. In fact it also generalizes their result, since they only consider moving points with positive initial coordinates and they assume that the starting point is at the origin.
2. The Kinetic TSP cannot be approximated to within a factor of two in polynomial time unless P=NP, even if there are only two moving points in the instance. This proves that Helvig et al.'s algorithm is in fact asymptotically optimal, disregarding the small ε factor, and that P might be NP.
3. The Kinetic TSP cannot be approximated to within a factor of $2^{\Omega(\sqrt{n})}$ in polynomial time unless P=NP, even if the maximum speed is bounded. The n denotes the size of the input instance.

The last result in particular is surprising in the light of existing polynomial time approximation schemes [1, 2, 11] for the static version of the problem.

In the next section, we state definitions and give preliminary results concerning Kinetic TSP. Specifically, we give an overview of the original reduction of Garey et al. [7] proving the NP-hardness of the Euclidean TSP, since it plays an important role in our later reductions. In Section 3, we prove the existence of a PTAS for the case when all points move with the same speed in the same direction. In Sections 4 and 5, we prove the stated inapproximability results and we conclude the presentation with a discussion of open problems.

2 Preliminaries and Notation

In the *Euclidean Traveling Salesman Problem* we are given a set of points $S = \{s_1, s_2, \ldots, s_n\}$ in the Euclidean plane. The objective is to compute the shortest tour that visits all points (the optimal TSP tour). The corresponding shortest path that visits all points of S, starting at s_1 and ending at s_n is called the optimal TSP path.

As in Chalasani et al. [5] we distinguish between *space-points* and *moving points*. A space-point is a point in a coordinate system, whereas a moving point is a point-object in space, the Euclidean plane in our case, that travels with a given velocity. The coordinates of a moving point s can be described by the function $s(t) = (x + tv\cos\alpha, y + tv\sin\alpha)$, where $v \geq 0$ is the point's speed and α is its direction. If $v = 0$ we say that the point is static. The *traveling salesman* is described by a special point that can move with variable speed and direction. The initial position s_0 of the salesman is assumed to be $(0, 0)$ and its maximal speed is assumed to be 1. The path taken by the salesman is denoted P and $P(t)$ is the position of the salesman at time t. If $P(t) = s(t)$ then we say that the salesman *visits* s at time t. P is called a *salesman path* of S if all points in S have been visited by the salesman. If the salesman also returns to its initial position, then we call the resulting tour a *salesman tour*.

We can now define the kinetic traveling salesman problem for moving points in the plane.

Definition 1. *A set of moving points* $S(t) = \{s_1(t), s_2(t), \dots, s_n(t)\}$ *in the plane with the Euclidean metric is given. A point* $s_i(t)$ *in* $S(t)$ *moves with the speed* $v_i < 1$. *Consider a traveling salesman as defined above. The objective of the* Kinetic *Traveling Salesman Problem (KTSP) is to compute a salesman tour, starting and ending at the initial point* $s_0 = (0, 0)$, *that minimizes the traveling time of the salesman. The* Translational *Traveling Salesman Problem (TTSP) is a restricted version of KTSP, where all points of* S *have the same speed and direction.*

A convenient way to characterize kinetic TSP instances is as complete directed graphs with time dependent edge weights. The weight of an edge is defined to be equal to the time a salesman need to traverse the edge starting at time t.

The following lemma provides an important fact concerning the speed of the salesman. It is a direct consequence of the fact that the salesman can travel faster than any other moving point. Note that the salesman needs to be faster than the points that are being visited since otherwise some "mice may never be caught".

Lemma 1. *An optimal salesman moves with maximal speed.*

From now on we assume that the salesman travels with maximal speed and since the maximal speed is 1, the distance traveled equals the traveling time.

Generally, we use OPT to denote a shortest tour or path, unless more specific notation is needed. Given a tour or a path P in the Euclidean plane, we let $\mathcal{C}(P)$ denote the length of P.

2.1 X3C Reduction of Garey et al.

Garey et al. [7] prove that the Euclidean traveling salesman problem is NP-hard by a reduction from the problem *exact cover by 3-sets*, X3C:

> Given a family $\mathcal{F} = \{F_1, F_2, \dots, F_r\}$ of 3-element subsets of a set U of $3k$ elements (represented by the integers $1, \dots, 3k$), does there exist a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ of pairwise disjoint sets such that $\bigcup_{F \in \mathcal{F}'} F = U$?

We say that if there exists such a subfamily for $\mathcal{F}$ then $\mathcal{F} \in X3C$, otherwise we say that $\mathcal{F} \notin X3C$.

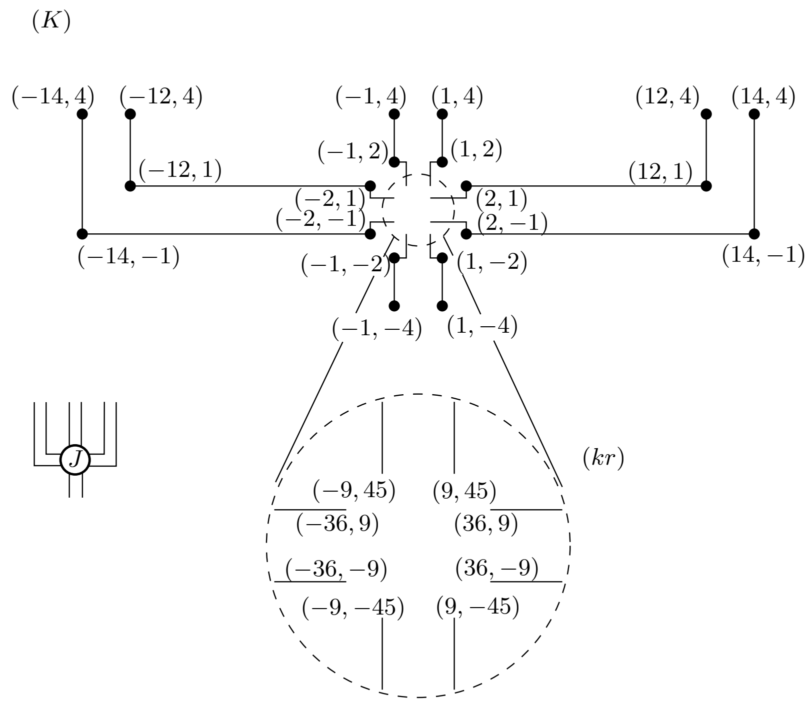

Fig. 1. TSP junctions.

We will describe the instance used in the reduction performed by Garey et al. This section follows the description of Garey et al. closely, using identical notation. For a complete presentation of the reduction we refer to the paper [7].

Given an instance of X3C, i.e. a family $\mathcal{F}$ of 3-elements, the objective is to construct a set of points S and a bound L^* such that an optimal TSP tour in S has length at most L^* if and only if $\mathcal{F}$ has an exact cover. S is constructed from $\mathcal{F}$ in stages, starting with two basic structures: *junctions* and *crossovers*. Let us start with the junctions. For each set $F \in \mathcal{F}$ we construct one junction in S, that is, a gadget designed to represent F in the TSP instance. Fig. 1 describes the geometry of a junction in detail. Note that throughout the reduction, a line segment represents the set of points with integer coordinates it contains. For simplicity, we assume that the junction gadget is centered around the origin $(0, 0)$. The region within the dotted circle is called the *active region*. The end point coordinates of the line segments in the junction's active region are dependent on the number of sets in $\mathcal{F}$ as well as the number of possible elements, since $r = |\mathcal{F}|$ and $3k = |U|$; see the active region in Fig. 1. Moreover, the line segment end point coordinates outside the active region are dependent on the value K,

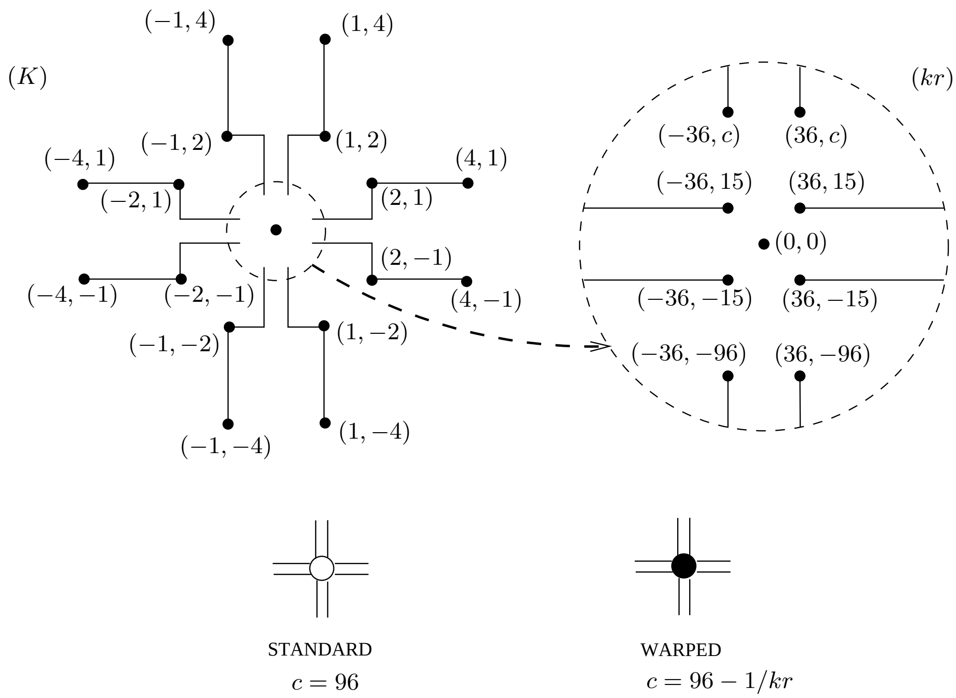

Fig. 2. TSP crossovers.

which is defined as

$$K = 108kr^2 + 1008k^2r^2 + 108k^2r,$$

see Fig. 1.

The second structure of the reduction is the crossover. There are two types of crossovers: the standard type and the warped type. The crossover structure is used to represent the elements of a set in the family. Fig. 2 describes the standard and the warped crossover in detail. Note that the origin itself is included in this gadget. The difference between the standard and the warped crossover lies in the coordinates of the topmost line segments' lower end points.

The crossovers are assembled into vertical sequences called crossover stacks; see Fig. 3. The two topmost end points of each crossover coincide with the two lowest points in the crossover above it. The topmost crossover is warped and the other crossovers in the sequence are of the standard type.

Each set $F_i = \{a_i, b_i, c_i\}$ in the family $\mathcal{F}$ will be represented in S by a set structure consisting of one junction and three crossover stacks of height a_i, b_i

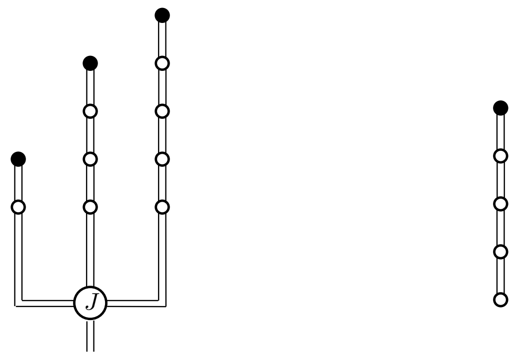

Set structure for $F = \{2, 4, 5\}$. Crossover stack of height 5.

Fig. 3. Compound gadgets of the construction S.

and c_i respectively. These are joined by connecting the three topmost pairs of points in the junction to the lowest pair of points in the crossover stacks. In this way a tube system is built for each set structure. These tube systems are connected into a huge system as Fig. 6 suggests. The figure describes the final TSP construction S for the family $\mathcal{F} = \{\{1,2,3\}, \{2,4,5\}, \{1,2,6\}\}$.

Let T_0 denote the set of unit length edges between pairs of points in S. The tube system consists of edges from T_0, and at tube intersection points crossovers and junctions are located. The following lemma serves to justify our attention to the tube segments.

Lemma 2. *An optimal TSP tour OPT in S contains all edges of* T_0.

To complete a Hamiltonian cycle in this graph, following the line segments of the tubes, some of these tubes must be connected to each other. These connections are realized in the active regions of the gadgets. For junctions there are two types of connections: either they are connected or they are not; see Fig. 4. For crossovers (standard and warped), there are three possible connections: no connection, the upward connection and the downward connection; see Fig. 5. Note that the non-connection alternative is the cheapest one so an optimal TSP tour needs to minimize the number of other connections. We say that a gadget is active if it performs a connection.

Let $L^* = |T_0| + 72kr^2 + 312krq + 108k^2r - 6k$, where $q < 3kr$ is the number of crossovers in S. L^* is going to be the lower bound of a TSP tour in S if no exact cover exists. We state the lemma as follows:

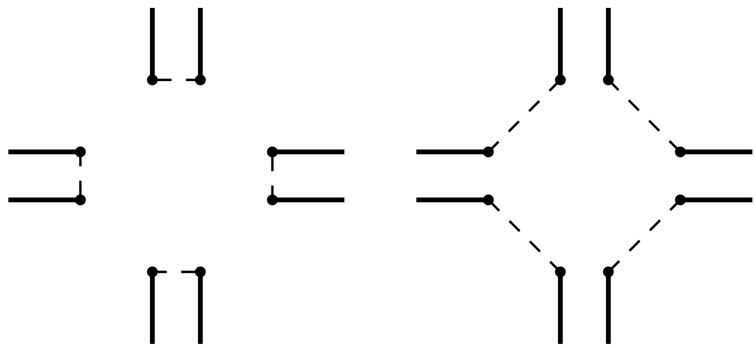

a) NO CONNECTION Length: $72kr$

b) CONNECTION Length: $180kr$

Fig. 4. Possible connections in active regions of junctions.

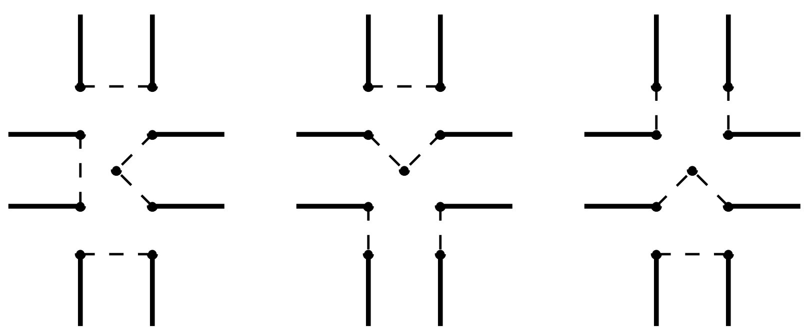

a) NO CONNECTION Length: $252kr$

b) DOWNWARD CONNECTION Length: $312kr$

c) UPWARD CONNECTION Length: $312kr(-2$ if WARPED)

Fig. 5. Possible connections in active regions of crossovers.

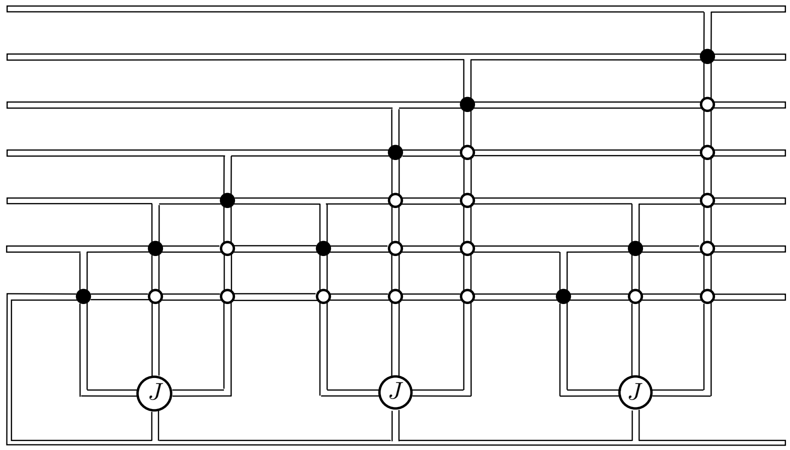

Fig. 6. Final S when $\mathcal{F} = \{\{1,2,3\},\{2,4,5\},\{1,2,6\}\}$.

Lemma 3. *There is an optimal TSP tour OPT of length less or equal to L^* if and only if there is an exact cover of the family $\mathcal{F}$.*

Garey et al. prove this theorem by showing that only connections corresponding to an exact cover can give a tour of length less or equal to L^*. This means that only the junctions corresponding to the sets involved in the exact cover are active, and in the same way, only the crossovers in the crossover stacks representing the elements in the exact cover are active. For the proof of this theorem we refer to Garey et al. [7].

We end this section with a lemma that appeared as the final remark of Garey et al. in their paper. This lemma is crucial for our reductions in later sections.

Lemma 4. *Let OPT denote the minimum length circuit of S, i.e. an optimal TSP tour. Then OPT has integral length.*

This follows from Lemma 2 and the way in which the active regions of the gadgets were constructed. Out of this fact next lemma follows directly.

Lemma 5. *There is no polynomial time approximation algorithm for TSP producing a tour APX such that $\mathcal{C}(APX) < \mathcal{C}(OPT) + 1$ unless $P = NP$.*

The lemma holds also for the optimal TSP path problem, since the reduction of Garey et al. also works for that problem. In this case, the reduction from X3C results in instances where we have opened the bottom tube. Observe that the length of the optimal TSP path for instances thus constructed is equal to the length of the corresponding optimal TSP tour in S. Instances produced by the

X3C reduction will from now on be called *GGJ-instances*. We refer to the original instances as type a whereas type b instances are those with an opened bottom tube.

3 A PTAS for TTSP

To find a PTAS for the Translational TSP we are going to find a bijective mapping between this problem and the static Euclidean TSP. Let us start this section by analyzing the translational TSP. At this point we are interested in finding the optimal translational TSP path. A set $S(t) = \{s_1(t), \ldots, s_n(t)\}$ of moving points is given together with a starting point $s_0 = (0, 0)$. All points move in the same direction α and with the same speed v. We can assume, w. l. o. g., that $\alpha = \frac{\pi}{2}$. Thus a point $s_i(t)$ is defined as

$$s_i(t) = (x + tv\cos\frac{\pi}{2}, y + tv\sin\frac{\pi}{2}) = (x, y + tv).$$

The traveling distance c_{ij} between two points $s_i(t)$ and $s_j(t)$ is the time needed by a salesman, moving with speed 1, to travel from s_i to s_j, i.e.

$$c_{ij} = \frac{v}{1-v^2}(y_j - y_i) + \sqrt{\frac{(x_j - x_i)^2}{1-v^2} + \frac{(y_j - y_i)^2}{(1-v^2)^2}}$$

Note that the traveling distance is independent of time and that the cost function c_{ij} is asymmetric. Consider the function

$$d(s_i, s_j) = \frac{c_{ij} + c_{ji}}{2} = \sqrt{\frac{(x_j - x_i)^2}{1-v^2} + \frac{(y_j - y_i)^2}{(1-v^2)^2}}$$

and the bijective mapping $f_v(s)$ from a moving point $s = (x, y)$ to a static point in the Euclidean plane:

$$f_v(s) = \left(\frac{x}{\sqrt{1-v^2}}, \frac{y}{1-v^2}\right)$$

Let $P = (s_{i_1}, \ldots, s_{i_n})$ be a path in the translational instance, then

$$\mathcal{C}(P) = \sum_{k=1}^{n-1} c_{i_k, i_{k+1}} = \frac{v(y_{i_n} - y_{i_1})}{1-v^2} + \sum_{k=1}^{n-1} d(s_{i_k}, s_{i_{k+1}}).$$

With this it is easy to prove the following result.

Lemma 6. *Let S_T be an instance of the translational TSP with specified starting* and endpoint, *and let S_E be the corresponding Euclidean instance after the transformation using f_v. A salesman path in S_T is optimal if and only if the corresponding Euclidean Hamiltonian path in S_E is optimal.*

Proof. Let $P = (s_{i_1}, \dots, s_{i_n})$ denote a salesman path in S_T and let P' denote the corresponding Hamiltonian path in S_E. Observe that $\mathcal{C}(P') = \sum_{j=1}^{n-1} d(s_{i_j}, s_{i_{j+1}})$. Therefore, $\mathcal{C}(P) = \mathcal{C}(P') + \frac{v(y_{i_n} - y_{i_1})}{1-v^2}$. Note that the last term is fixed since both the starting point and the endpoint were given. It follows that P is optimal if and only if P' is optimal. □

Consider an arbitrary instance of the translational TSP with specified starting point s_0 and ending point s_n. We can transform this instance into an instance of the Euclidean TSP using the bijective mapping f_v. Given this new instance, we compute a Hamiltonian path APX_E such that $\mathcal{C}(APX_E) \leq (1+\epsilon)\mathcal{C}(OPT_E)$, using a modified PTAS for the Euclidean TSP [1, 2, 11]. Here OPT_E denotes an optimal TSP path in the Euclidean instance. Let APX_T and OPT_T denote the corresponding paths in the translational instance. Observe that OPT_T is an optimal salesman path by Lemma 6. We have

$$\begin{aligned}\mathcal{C}(APX_T) &= \frac{v(y_n - y_1)}{1-v^2} + \mathcal{C}(APX_E) \\ &\leq \frac{v(y_n - y_1)}{1-v^2} + (1+\epsilon)\mathcal{C}(OPT_E) \\ &= \mathcal{C}(OPT_T) + \epsilon\mathcal{C}(OPT_E) \\ &\leq (1 + \frac{\epsilon}{1-v})\mathcal{C}(OPT_T).\end{aligned}$$

The last inequality holds since OPT_T is at least as long as the shortest path between s_1 and s_n, i.e. $\mathcal{C}(OPT_T) \geq \frac{|y_n - y_1|}{1+v}$. Hence,

$$\begin{aligned}\mathcal{C}(OPT_E) = \mathcal{C}(OPT_T) - \frac{v(y_n - y_1)}{1-v^2} &\leq \mathcal{C}(OPT_T) + \frac{v|y_n - y_1|}{(1-v)(1+v)} \\ \leq \mathcal{C}(OPT_T) + \frac{v\mathcal{C}(OPT_T)}{1-v} &= \frac{\mathcal{C}(OPT_T)}{1-v}.\end{aligned}$$

It follows that APX_T is a $(1 + \frac{\epsilon}{1-v})$-approximation of the optimal path. Thus, we have a PTAS for the translational TSP path problem. Note that the approximation factor does depend on the maximal speed. Thus, if we want to compute a constant factor approximation for an instance in which the points are moving with the speed $1 - \frac{1}{n}$ then our algorithm might need exponential time.

With this approximation scheme we can now give a PTAS for the translational TSP. We use the same approach as Chalasani et al. [4, 5]. The difficulty of the translational TSP is that the initial point s_0, is not moving, which creates an asymmetry that we must be able to handle. Assume that $OPT = (s_0, s_1(t), \dots, s_n(t), s_0)$ is an optimal salesman tour. It follows easily that the optimal salesman path starting from s_0 and ending at $s_n(t)$ is a part of that tour. For each possible such ending point $s_i(t)$, $i \geq 1$, we compute a $(1 + \frac{\epsilon}{2(1-v)})$-approximate salesman path. This gives us n paths that can easily be turned into salesman tours. The algorithm returns the shortest of these tours.

Theorem 1. *The algorithm described above is a PTAS for the translational TSP.*

Proof. Let OPT_H denote the length of the optimal salesman path, starting from s_0 and ending at $s_i(t)$. The length of the optimal salesman tour OPT is $\mathcal{C}(OPT) = OPT_H + l$, assuming s_i is the last moving point visited by the tour OPT and l is the length of the last segment in the tour (between the points $s_i(t)$ and s_0). There is a tour APX among the tours our algorithm computes, that also has $s_i(t)$ as the last unvisited point. The length of APX is

$$\mathcal{C}(APX) = (1 + \frac{\epsilon}{2(1-v)})OPT_H + l',$$

where l' is the length between $s_i(t)$ and s_0. A salesman that travels along APX visits $s_i(t)$ at time $t = (1 + \frac{\epsilon}{2(1-v)})OPT_H$, whereas the optimal salesman visits $s_i(t)$ at time $t = OPT_H$. The time difference is $\frac{\epsilon}{2(1-v)}OPT_H$ and in that time period, $s_i(t)$ moves the length $\frac{v\epsilon}{2(1-v)}OPT_H$. The triangle inequality assures us that l' does not exceed $l + \frac{v\epsilon}{2(1-v)}OPT_H$. The cost of our approximate tour is thus

$$\begin{aligned}\mathcal{C}(APX) &\leq (1 + \frac{\epsilon}{2(1-v)})OPT_H + l + \frac{v\epsilon}{2(1-v)}OPT_H \\ &= (1 + \frac{(1+v)\epsilon}{2(1-v)})OPT_H + l \\ &< (1 + \frac{\epsilon}{1-v})OPT_H + l.\end{aligned}$$

The last inequality holds since the speed does not exceed 1. The ratio between the costs of APX and OPT becomes

$$\frac{\mathcal{C}(APX)}{\mathcal{C}(OPT)} < \frac{(1 + \frac{\epsilon}{1-v})OPT_H + l}{OPT_H + l} \leq (1 + \frac{\epsilon}{1-v}).$$

The proof is completed, since the cost of the returned salesman tour does not exceed that of APX. □

Observe that our technique differs from that of Chalasani et al. [5] only in the bijective mapping f_v with which the transformation was performed. Thanks to this mapping we were able to generalize their result. If the speeds are bounded by a constant $c < 1$ then the algorithm is a true PTAS.

4 Two Reductions for KTSP

The kinetic traveling salesman problem is in general not as easily approximated as the translational TSP. Helvig, Robins and Zelikovsky [8] give a $2 +$

ϵ-approximation algorithm for kinetic TSP-instances with n points of which $O(\frac{\log n}{\log \log n})$ are moving. The simple construction described in Figure 7 shows that their algorithm is close to optimal. Note that all GGJ-instances used in this section are stationary.

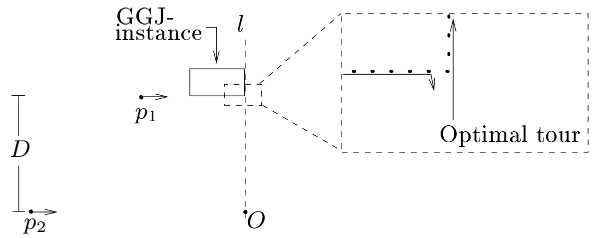

Fig. 7. Kinetic TSP-instance with two moving points.

The figure contains a GGJ-instance of type a, at distance D from the origin and two moving points p_1 and p_2, both with speed $v < 1$. We assume that the salesman starts at the origin at time $t = 0$. The point p_1 moves horizontally and lies at the same altitude as the bottom line of the GGJ-instance and intersects the line l at time D. Point p_2 moves parallel to p_1 and reaches the origin at time $2D + L^*$. Recall that a GGJ-instance is constructed so that if $\mathcal{F} \in X3C$, then an optimal salesman tour is at most L^* long. Otherwise the length is at least $L^* + 1$.

Because of the large speed of the moving points, it follows that the first point to be visited is p_1. Furthermore, the point p_2 must be visited before it reaches the origin. Otherwise, the tour has unbounded length.

If $\mathcal{F} \in X3C$, then an optimal salesman can catch p_1 and continue along the GGJ-instance's optimal salesman tour. After all points in the GGJ-instance have been visited, the salesman still has time to catch p_2 before it reaches the origin.

If $\mathcal{F} \notin X3C$, then an optimal salesman does not have time to catch all the points in the GGJ-instance before he has to turn back and catch p_2. This means that he must traverse the distance D at least four times. From this we get:

Theorem 2. *It is NP-Hard to get an approximation ratio less than* 2 *for the kinetic TSP, even if there are only two moving points in the instance.*

In Theorem 2 we need the speed v to be arbitrarily close to 1. With more moving points we can restrict the speed to $v = 1/2$ and still get a lower bound of 2 on the approximation ratio.

Theorem 3. *It is NP-Hard to get an approximation ratio less than* 2 *for the kinetic TSP, even if the speed of the moving points is restricted to* $v = 1/2$.

Proof. To prove the theorem we use two pairs of sets with $k = m$ points in each set evenly distributed on a line, and a GGJ-instance of type b with m static points. Each pair contains one set with static points and one with moving points. The static points lie on a vertical line of length $D = m^2$, whereas the moving points lie on a line with the same height but with the length $\frac{\sqrt{5}}{2}D$. The slope of the second line is -2 in pair 1 and 2 in pair 2. The moving points perform a horizontal motion from left to right with the speed $v = 1/2$.

Given these building blocks, we construct an instance of the Kinetic TSP as described in Figure 8. The salesman begins at the origin at time $t = 0$. At the same time, the lowest moving point of pair 1 intersects the origin and at time $D + L^*$, the topmost moving point in pair 2 collides with the corresponding static point in that pair. The pairs are constructed so that the salesman can move (with speed 1) along the vertical line of the static points, visiting pairs of points from the two sets at their collision point if and only if the salesman visits the first two colliding points at the time of impact. It thus takes time D to visit all points in a pair, assuming the conditions above, since the length of the vertical line is D.

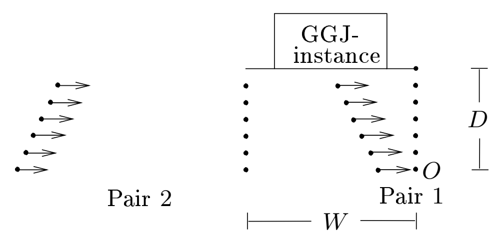

Fig. 8. KTSP-instance with moving points of speed 1/2.

If there exists an exact cover for the X3C-instance, then an optimal salesman has time to visit all points in both pair 1 and the GGJ-instance before the first two points in pair 2 collide. The time that an optimal salesman needs in order to visit all points in the whole instance is thus at most $2D + L^* + W$, where $W \leq L^*$ is the length of the bottom line segment of the GGJ-instance.

If there is no exact cover for the X3C-instance, then an optimal salesman does not have enough time to visit all points in the GGJ-instance before the first two points in pair 2 collide. Thus, he is left with two options. Either he visits all points in the GGJ-instance before visiting the points of pair 2 or he saves some points in the GGJ-instance in order to fetch the other points optimally.

If the salesman chooses to save some points in the static instance, then he must go back again after visiting the moving points. The total traveling time is thus at least $4D + L^*$.

If he decides to visit all points in the static GGJ-instance at once, then he misses the first collision in pair 2 by at least one time unit. The salesman ends up chasing each point separately. If he visits the points set-wise, he ends up close to the GGJ-instance and at least a distance D from the origin. Such an approach takes thus at least time $4D + L^*$. To avoid this he might zig-zag between the two sets in pair 2. For each pair of points that the salesman visits, the gain is at most $3D/k$. But since the points in the second set move with speed $v = 1/2$, the cost of moving between the two sets back and forth doubles each time. The cost of this approach is thus $4D + L^* - 3Dz/k + 2^{z-1}$, where z is the number of zig-zag movements. Minimizing this expression, using the knowledge that $L^* \geq m$, $D = m^2$ and $k = m$, proves that the length of the tour is at least $4m^2 - cm \log m$, for some constant c.

Since the salesman moves with speed 1, length equals time. The length of an optimal salesman tour is therefore at least $4m^2 - cm \log m$ if $\mathcal{F} \notin X3C$ and at most $2(m^2 + L^*)$ if $\mathcal{F} \in X3C$. The limit of the ratio gives a lower bound on the approximation ratio and concludes the proof of Theorem 3:

$$\lim_{m \to \infty} \frac{4m^2 - cm \log m}{2(m^2 + L^*)} = 2.$$

□

5 An Exponential Lower Bound for the general KTSP

In this section, we present a gap producing reduction from X3C to the kinetic TSP. The instance used in the reduction is a uniformly expanding KTSP-instance. A uniformly expanding instance contains moving points of the form $s_i(t) = (v_i t \cos \alpha, v_i t \sin \alpha)$. This implies that at time $t = 0$ all points are located at the origin, and that the relative distances within the instance do not change over time. We assume that the salesman begins his pursuit at time $t_0 > 0$ (if he starts at time $t = 0$, he visits all points at once without moving). Observe that the length of a salesman tour P in a uniformly expanding TSP-instance is direction independent. That is, if we reverse the order in which the points are visited, the length of the new tour is equal to the length of P. As gadgets in the KTSP-instance we use a special kind of GGJ-instances suitable for translation. We therefore begin with a description of these, and continue with a detailed description of the KTSP-instance.

Consider a static type b GGJ-instance. From this instance we want to create a translational instance with speed v whose optimal tour has the same length as the static instance. This is achieved by applying the inverse of the function f_v described in Section 3 to all points in the static instance. Let us call such instances *translational inverse GGJ-instances*. Given such an instance we create a uniformly expanding GGJ-instance in which the relative distances match the relative distances in the translational instance. We call these *small expanding GGJ-instances*, or simply GGJ-instances.

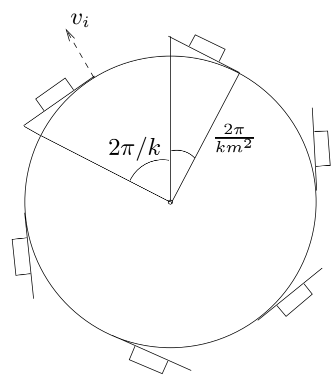

Fig. 9. One circle of the inapproximable KTSP-instance.

The KTSP-instance consists of l concentric circles $C_1, \ldots, C_l$, with radii $r_1, r_2, \ldots, r_l$, centered around the origin. Each circle consists of $k = \lceil m^2 l \rceil$ identical small expanding GGJ-instances. The GGJ-instances are placed on the circles as shown in Figure 9.

Each GGJ-instance contains m points. We let $l = m^a$, for some $a > 1$. The total number of points in the instance, denoted n, is therefore $mkl = m^{2a+3}$. The rightmost point of a GGJ-instance is placed on the circle and the bottom line of the GGJ-instance follows the circle's tangent line at that point. Each circle C_i is expanding with a speed v_i, i.e. the rightmost point of each GGJ-instance on C_i has a speed v_i, directed away from the center and the radius of C_i is $r_i = v_i t$.

We let $v_1 = \frac{1}{2e}$ and $v_i = (1 + \frac{m^2}{k})v_{i-1}$. This implies that the largest circle's speed, v_l, is bounded by

$$v_l = (1 + \frac{m^2}{k})^l v_1 = (1 + \frac{1}{l})^l v_1 < e v_1 = \frac{1}{2}.$$

Note also that $r_1 \leq r_2 \leq \ldots \leq r_l$. Observe that the GGJ-instances are uniformly expanding, i.e. the distance between two points s_i and s_j is $d(s_i, s_j) = d_{ij} t$, where d_{ij} is constant. The expansion rate is such that the distance between the leftmost and the rightmost point of a GGJ-instance is less than $t v_i \tan \frac{2\pi}{m^2 k}$. We

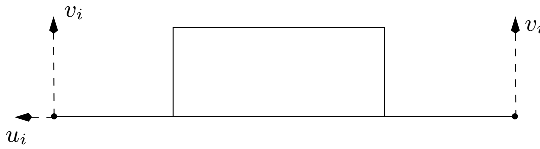

Fig. 10. A small expanding GGJ-instance.

let $u_i = v_i \tan \frac{2\pi}{m^2 k}$ define the expansion rate; see Figure 10. This implies that the angle between the leftmost point L, the origin, and the rightmost point R in a GGJ-instance is $\angle LOR = \frac{2\pi}{m^2 k}$ and the angle between two consecutive GGJ-instances is $\frac{2\pi}{k}$; see Fig. 9. The fastest point of a GGJ-instance is the leftmost point and this point has the speed

$$\sqrt{v_i^2 + (v_i \tan \frac{2\pi}{m^2 k})^2} \leq v_i(1 + \frac{2\pi}{m^2 k}).$$

Thus, no point belonging to circle C_i has a speed exceeding $v_i(1 + \frac{2\pi}{m^2 k})$, and this implies that the maximal speed is bounded. From now on we simply denote our newly constructed instance the *KTSP-instance*.

We will show that the construction above guarantees that an optimal salesman visits the GGJ-instances circlewise, one at a time. Furthermore, the time it takes to go from any moving point s_i to any other moving point s_j in the KTSP-instance depends linearly on the starting time. Assuming that a salesman starts from s_i at time t it would take him time $\kappa_{ij} t$ to go to s_j, where we consider κ_{ij} to be constant since it does not depend on t. As a consequence, traveling time is multiplicative in the edge weight constants κ_{ij}. Thus, if the salesman begins his pursuit at time t_0 and finishes the first GGJ-instance at time Kt_0, then it will take him at least time $CK^{kl}t_0$ to visit all points in the KTSP-instance, where C denotes a constant and K is independent of t. Note that a non-optimal salesman must do worse on each GGJ-instance and since the error is also multiplicative, the inapproximability ratio is going to become large. The rest of the presentation is devoted to proving the following theorem.

Theorem 4. *For any $\gamma > 0$, there exists no polynomial-time algorithm that achieves an approximation ratio of $2^{\Omega(n^{1/2-\gamma})}$ for the kinetic traveling salesman problem, unless P=NP.*

Consider a salesman moving between the points in the KTSP-instance. Let P be the path taken by the salesman. We define $\mathcal{C}_P(t)$ to be the length of this path (which is synonymous with the time it takes for the salesman to traverse path P). We let $\mathcal{T}_P(t)$ denote the time when the salesman arrives at the end of the path P, given the starting time t, i.e., $\mathcal{T}_P(t) = t + \mathcal{C}_P(t)$. A simple proof of induction yields that

$$\mathcal{T}_P(t) = t \prod_{[s_i,s_j]\in P} (1+\kappa_{ij}) = tK_P$$

and therefore,

$$\mathcal{C}_P(t) = t \prod_{[s_i,s_j]\in P} (1+\kappa_{ij}) - t,$$

where $t\kappa_{ij}$ is the time it takes for the salesman to travel between s_i and s_j, starting from s_i at time t. Let us examine a small expanding GGJ-instance produced by an X3C-instance. We assume that the instance is located somewhere on circle C_i. The optimal salesman path for the GGJ-instance starting at the rightmost point and ending at the leftmost point is denoted *opt* if $\mathcal{F} \in X3C$ and *opt*$'$ if $\mathcal{F} \notin X3C$. We would like to find an upper bound on *opt* and a lower bound on *opt*$'$. To do this, we need some definitions. Consider a translational inverse GGJ-instance of a satisfying X3C instance with the same size as the small expanding GGJ-instance at time t. Let all points in the translational instance move with the same velocity as the rightmost point in the expanding GGJ-instance at time t. We define $D(t)$ as the length of the optimal path, starting at the rightmost point and ending at the leftmost point. Let $D'(t)$ denote the length of the optimal path in a similar instance constructed from an X3C-instance with no exact cover. Clearly,

$$D(t) \le \mathcal{C}_{opt}(t) \le D(t+\mathcal{C}_{opt}(t)), \text{ and} \tag{1}$$

$$D'(t) \le \mathcal{C}_{opt'}(t) \le D'(t+\mathcal{C}_{opt'}(t)) \tag{2}$$

Furthermore, let $f(t)$ denote the distance between the two closest points in the expanding instance at time t and let t_i be the point in time such that $f(t_i) = 1 - v_i^2$. Note that t_i is the time when the expanding GGJ-instances reach the size of the original GGJ-instances. That is, if we apply the function f_v on an expanding instance at time t_i, then the static instance that results will have the same size as the original GGJ-instances of Section 2. Thus, $D(t_i) \le L^*$ and $D'(t_i) \ge L^* + 1$. Observe also that t_i depends on the circle where the instance is located. $D(t)$, $D'(t)$ and $f(t)$ are linear mappings, which implies that

$$\frac{D(t)}{D(\tilde{t})} = \frac{D'(t)}{D'(\tilde{t})} = \frac{f(t)}{f(\tilde{t})} = \frac{t}{\tilde{t}}.$$

Lemma 7. *The following bounds hold:*

$$c_1 km^2 \le t_i \le c_2 km^3 \text{ (for some constants } c_1 \text{ and } c_2\text{)}, \tag{3}$$

$$D'(t) > D(t) + f(t) \tag{4}$$

$$\mathcal{C}_{opt}(t) \leq \frac{D(t_i)}{t_i - D(t_i)} t, \tag{5}$$

$$\mathcal{C}_{opt'}(t) \geq \frac{D(t_i) + 1 - v_i^2}{t_i} t. \tag{6}$$

Proof. (3) Let v_f denote the expansion rate of the distance between the two closest points in the instance. Then $f(t) = v_f t$ and by definition, $1 - v_i^2 = f(t_i) = v_f t_i$. The length of the bottom line is between 1 and m at time t_i. Furthermore, the expansion rate of the bottom line is u_i. Given this we bound v_f to:

$$u_i \geq v_f \geq \frac{u_i}{m}.$$

Now, $t_i = \frac{1-v_i^2}{v_f}$, so

$$\frac{1 - v_i^2}{u_i} \leq t_i \leq \frac{m(1 - v_i^2)}{u_i}.$$

Recall that $u_i = v_i \tan \frac{2\pi}{m^2 k}$ and so by Taylor expansion, we get that

$$\frac{2\pi + \frac{1}{3}}{m^2 k} \geq \tan \frac{2\pi}{m^2 k} \geq \frac{2\pi}{m^2 k}.$$

Using this we can bound t_i as follows:

$$\frac{m^2 k(1 - v_i^2)}{(2\pi + \frac{1}{3})v_i} \leq t_i \leq \frac{m^3 k(1 - v_i^2)}{2\pi v_i}.$$

Inserting the bounds of v_i yields that $c_1 k m^2 \leq t_i \leq c_2 k m^3$, for some constants c_1 and c_2.

(4) It follows from the linearity of the functions:

$$D'(t) = D'(t_i) \cdot \frac{t}{t_i} > (D(t_i) + f(t_i)) \cdot \frac{t}{t_i} = \frac{D(t_i)D(t)}{D(t_i)} + \frac{f(t_i)f(t)}{f(t_i)} = D(t) + f(t)$$

(5) We start with the inequality

$$\mathcal{C}_{opt}(t) \leq D(t + \mathcal{C}_{opt}(t))$$

taken from (1). Using the linearity of $D(t)$ we get that

$$D(t + \mathcal{C}_{opt}(t)) = D(t_i)(t + \mathcal{C}_{opt}(t))/t_i.$$

This implies that

$$\mathcal{C}_{opt}(t) \leq \frac{D(t_i)}{t_i - D(t_i)} t,$$

since $t_i \geq c_1 k m^2 > D(t_i)$.

(6) By the linearity of $D(t)$ and $f(t)$ we have that

$$\mathcal{C}_{opt'}(t) \geq D'(t) > D(t) + f(t) = \frac{D(t_i) + 1 - v_i^2}{t_i} \cdot t.$$

□

With these bounds we can prove a lower bound on the approximation ratio for the small expanding GGJ-instances. To simplify the analysis later on, we compute the approximation ratio in terms of the arrival time. We assume that the salesman visits his first point of the instance at time t.

Lemma 8. *It is NP-hard to find a salesman path apx, for an expanding GGJ-instance, with*

$$\mathcal{T}_{apx}(t) \leq (1 + \frac{1}{ckm^3})\mathcal{T}_{opt}(t).$$

Proof. The following ratio measures the gap between the arrival times of *apx* and *opt*. According to Lemma 7

$$\begin{aligned} \frac{\mathcal{T}_{apx}(t)}{\mathcal{T}_{opt}(t)} &\geq \frac{\mathcal{T}_{opt'}(t)}{\mathcal{T}_{opt}(t)} = \frac{\mathcal{C}_{opt'}(t) + t}{\mathcal{C}_{opt}(t) + t} \\ &\geq \frac{\frac{D(t_i)+1-v_i^2}{t_i} t + t}{\frac{D(t_i)}{t_i - D(t_i)} t + t} \geq 1 + \frac{1 - v_i^2}{t_i} - \frac{2D(t_i)^2}{t_i^2} \end{aligned}$$

$D(t_i) < 12m$ and $c_1 km^2 \leq t_i \leq c_2 km^3$. Expressed in terms of k and m, the quotient is at least

$$\frac{\mathcal{T}_{apx}(t)}{\mathcal{T}_{opt}(t)} > 1 + \frac{1 - v_i^2}{c_2 km^3} - \frac{2m^2}{(c_1 km^2)^2} > 1 + \frac{1}{ckm^3}$$

for sufficiently large k, and $c = 2c_2$. □

Let us return to the KTSP-instance that contains the small expanding GGJ-instances. We assume that the salesman starts at the origin at time $t_0 > 0$. We will prove that the KTSP-instance is inapproximable using the inapproximability result in Lemma 8 for the small expanding GGJ-instances.

Consider a salesman tour P for the KTSP-instance. A subpath of P starting at a point s_i and ending at the point s_j is called an *edge* of P if no other points are visited along the subpath. Edges that connect two different GGJ-instances are called *leaps*. A leap that exits the jth GGJ-instance on circle C_i is denoted J_{ij}. The shortest leap that connects circle C_i with circle C_{i+1} is the leap that connects the leftmost point of the topmost GGJ-instance on circle C_i with the rightmost point of the corresponding GGJ-instance on circle C_{i+1}. This leap will be denoted E_i (exterior leap exiting circle C_i); $\mathcal{T}_{E_i}(t) = K_{E_i} t$, for some K_{E_i}. We define E_l as the leap between the furthest point on circle C_l and the origin. The shortest leaps that connect two GGJ-instances on circle C_i are the leaps that connect the leftmost point of a GGJ-instance on circle C_i with the rightmost point of the next instance to the left. These leaps are identical in lengths, thus we can denote these simply as I_i (interior leaps performed within circle C_i); $\mathcal{T}_{I_i}(t) = K_{I_i} t$ for some constant K_{I_i}. Let opt_i denote the optimal salesman path in a small expanding GGJ-instance on circle C_i; $\mathcal{T}_{opt_i}(t) = K_{opt_i} t$ for some constant K_{opt_i}.

Lemma 9. *The length of an optimal path opt_i in an expanding GGJ-instance and the length of the leaps I_i and E_i are:*

$$\mathcal{C}_{opt_i}(t_i) = \Theta(m), \tag{7}$$

$$\mathcal{C}_{I_i}(t_i) = \Omega(m^2), \text{ and } \mathcal{C}_{I_i}(t_i) = O(m^3), \tag{8}$$

$$\mathcal{C}_{E_i}(t_i) = \Omega(m^4). \tag{9}$$

Proof. (7) $D(t_i) < D'(t_i) < 2m$, and since $t_i \geq c_1 k m^2$, we have that:

$$\mathcal{C}_{opt}(t_i) < \mathcal{C}_{opt'}(t_i) \leq \frac{D'(t_i)t_i}{t_i - D'(t_i)} = D'(t_i) + \frac{D'(t_i)^2}{t_i - D'(t_i)} = O(m).$$

Now, $D(t_i) \geq m$ since there are m points in the instance and the minimum distance between pairs of points is 1 at time t_i. It follows that $\mathcal{C}_{opt}(t_i)$ and therefore also $\mathcal{C}_{opt'}(t_i)$ is $\Theta(m)$. We conclude that $\mathcal{C}_{opt_i}(t_i) = \Theta(m)$.

(8) The distance, and thus the length of the traveling path, to the closest GGJ-instance on the same circle is shown to be $\Theta(t_i/k)$ in Figure 11, and $c_1 k m^2 \leq t_i \leq c_2 k m^3$.

(9) The length of any leap J, connecting circle C_i with another circle at time t_i is at least the minimum of

$$\min_{1 \leq j < i} \left\{ \frac{v_i - (1 + 2\pi/m^2 k)v_j}{1 + v_j} t_i \right\}$$

and

$$\min_{i < j \leq n} \left\{ \frac{v_j - (1 + 2\pi/m^2 k)v_i}{1 - v_j} t_i \right\}.$$

Inserting the bounds of t_i and v_i and observing that $v_j = (1 - \frac{m^2}{k})^{j-i} v_i$ yields that the length of any such leap is $\Omega(m^4)$.

Corollary 1. *For sufficiently large m*

$$K_{E_i} > K_{I_i} > K_{opt_i}.$$

Proof. It follows from Lemma 9 and the definition of $\mathcal{T}_P(t)$ and $\mathcal{C}_P(t)$.

Using this corollary we give a tight bound on the cost of an optimal salesman tour.

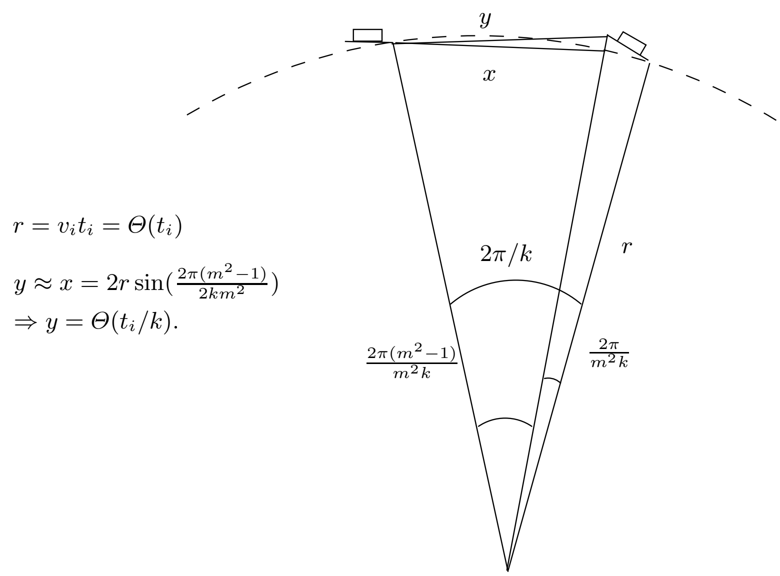

Fig. 11. Shortest distance between two instances at time t_i.

Lemma 10. *An optimal salesman tour, OPT, starting at the origin at time t_0 is finished at time*

$$\mathcal{T}_{OPT}(t_0) = K_{E_0} \prod_{i=1}^{l} \left((K_{opt_i} K_{I_i})^k \frac{K_{E_i}}{K_{I_i}} \right) t_0,$$

where E_0 is the leap between the origin and the first circle.

Proof. First of all let us only consider tours that visits the GGJ-instances counterclockwise one at a time. The length of the shortest such tour is equal to the given bound. To prove that this is also a lower bound we do as follows.

Let P be any salesman tour of the KTSP-instance. Let $K_{tsp_{ij}}$ denote the total cost spent in the jth GGJ-instance on circle C_i. Because of the salesman tour property we know that there must be at least one leap exiting each GGJ-instance and a leap J_0' from the origin to the first GGJ-instance in the tour. Thus we can characterize the finishing time of P as

$$\mathcal{T}_P(t_0) = K_{J_0'} \left(\prod_{i=1}^{l} \prod_{j=1}^{k} K_{tsp_{ij}} K_{J_{ij}} \left(\prod_{h=1}^{H_{ij}} K_{J_h} \right) \right) t_0,$$

where $K_{J_{ij}}$ describes the cost of exiting the jth GGJ-instance on circle C_i, and H_{ij} is the number of extra leaps exiting this GGJ-instance. The cost of these extra leaps are described by K_{J_h}.

Consider the total cost associated to the jth GGJ-instance on circle C_i. It follows from Corollary 1 that if $H_{ij} > 0$ then $K_{tsp_{ij}} K_{J_{ij}} \left(\prod_{h=1}^{H_{ij}} K_{J_h} \right) > K_{tsp_{ij}} K_{J_{ij}} K_{opt_i}$. If $H_{ij} = 0$ then $K_{tsp_{ij}} > K_{opt_i}$, since opt_i is the optimal path inside a GGJ-instance on circle C_i. In any case $K_{tsp_{ij}} K_{J_{ij}} \left(\prod_{h=1}^{H_{ij}} K_{J_h} \right) > K_{opt_i} K_{J_{ij}}$.

$K_{J_{ij}} > K_{I_i}$, since I_i is the shortest path between two GGJ-instances. Thus, $K_{opt_i} K_{J_{ij}} > K_{opt_i} K_{I_i}$, and since at least one leap is exiting each circle in P, a lower bound on $\mathcal{T}(t_0)$ is

$$\mathcal{T}_P(t_0) = K_{J'_0} \left(\prod_{i=1}^{l} \left(\prod_{j=1}^{k} K_{opt_i} K_{I_i} \right) \frac{K_{J'_i}}{K_{I_i}} \right) t_0,$$

where J'_i is a leap exiting circle C_i. This actually proves that the GGJ-instances must be taken circlewise. The question that remains is in what order the circles are taken.

Note that the cost of the shortest path between the origin and circle C_l is exactly $K_{E_0} K_{E_1} \cdots K_{E_{l-1}}$. Furthermore, the shortest path between the furthest point on circle C_l and the origin has the cost K_{E_l}. Now, any salesman path needs to visit these points. It turns out that

$$K_{E_0} K_{E_1} \cdots K_{E_l} \geq K_{J'_0} K_{J'_1} \cdots K_{J'_l}.$$

We get the following lower bound on any salesman path P:

$$\mathcal{T}_P(t_0) > K_{E_0} \prod_{i=1}^{l} \left((K_{opt_i} K_{I_i})^k \frac{K_{E_i}}{K_{I_i}} \right) t_0.$$

□

Corollary 2. *There exists no polynomial-time algorithm producing a tour APX for our KTSP-instance such that, unless P=NP,*

$$\mathcal{T}_{APX}(t_0) < K_{E_0} \prod_{i=1}^{l} \left(((1 + \frac{1}{ckm^3}) K_{opt_i} K_{I_i})^k \frac{K_{E_i}}{K_{I_i}} \right) t_0.$$

Proof. It follows by Lemma 8 and the proof of Lemma 10. □

We can now calculate a lower bound on the approximation ratio, assuming that $P \neq NP$.

Proof. (Theorem 4) We are actually interested in the ratio $\frac{\mathcal{C}_{APX}(t_0)}{\mathcal{C}_{OPT}(t_0)}$ but since $\frac{\mathcal{C}_{APX}(t_0)}{\mathcal{C}_{OPT}(t_0)} > \frac{\mathcal{T}_{APX}(t_0)}{\mathcal{T}_{OPT}(t_0)}$ we might as well use the latter ratio. From Lemma 10 and the resulting corollary we have that

$$\frac{\mathcal{T}_{APX}(t_0)}{\mathcal{T}_{OPT}(t_0)} \geq \frac{K_{E_0} \prod_{i=1}^{l} \left(((1 + \frac{1}{ckm^3}) K_{opt_i} K_{I_i})^k \frac{K_{E_i}}{K_{I_i}} \right) t_0}{K_{E_0} \prod_{i=1}^{l} \left((K_{opt_i} K_{I_i})^k \frac{K_{E_i}}{K_{I_i}} \right) t_0}$$

$$= (1 + \frac{1}{ckm^3})^{kl} \in 2^{\Omega(n^{1/2-\gamma})},$$

for $\gamma = 3/a$, and a can be chosen arbitrarily large. Remember that $n = mkl = m^{2a+3}$. □

6 Conclusions

We have investigated kinetic variants of TSP. Our major result proves an exponential lower bound on the approximation factor for such problems unless P=NP even when the velocities are bounded. Even so, we feel that the bound is coarse and can probably be improved. Also, the question of good upper bounds on the approximation ratio comes to mind.

The use of the mapping f_v described in Section 3 is actually a generic method that can be used to solve a large class of tour problems when instances perform constant translational movement in time.

The translational TSP PTAS that we present can be used in an $O(1)$-approximation algorithm for KTSP, assuming that the number of different velocities of the instances is bounded by a constant and that the maximum velocity of the instances considered is bounded.

References

1. S. Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *Proc. 37th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2–11, 1996.
2. S. Arora. Nearly linear time approximation schemes for Euclidean TSP and other geometric problems. In *Proc. 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 554–563, 1997.
3. S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and intractability of approximation problems. In *Proc. 33rd FOCS*, pages 13–22, 1992.
4. P. Chalasani and R. Motwani. Approximating capacitated routing and delivery problems. Technical Report STAN-CS-TN-95-24, Department of Computer Science, Stanford University, 1995.
5. P. Chalasani, R. Motwani, and A. Rao. Approximation algorithms for robot grasp and delivery. In *Proceedings of the 2nd International Workshop on Algorithmic Foundations of Robotics*, pages 347–362, 1996.
6. N. Christofides. Worst-case analysis of a new heuristic for the traveling salesman problem. Technical Report 388, Grad School of Industrial Administration, CMU, 1976.
7. M. R. Garey, R. L. Graham, and D. S. Johnson. Some NP-complete geometric problems. In *Proc. 8th Annual ACM Symposium on Theory of Computing*, pages 10–21, 1976.

8. C. S. Helvig, G. Robins, and A. Zelikovsky. Moving-target TSP and related problems. In *Proc. 6th Annual European Symposium on Algorithms (ESA)*, 1998.
9. E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys. *The TRAVELING SALESMAN PROBLEM, A Guided Tour of Combinatorial Optimization*. John Wiley & Sons Ltd., 1985.
10. C. Malandraki and M. S. Daskin. Time dependent vehicle routing problems: Formulations, properties and heuristic algorithms. *Transportation Science*, 26:185–200, 1992.
11. J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: Part II – a simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. To appear in SIAM J. Computing.
12. C. H. Papadimitriou. Euclidean TSP is NP-complete. *TCS*, 4:237–244, 1977.
13. S. B. Rao and W. D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *STOC*, pages 540–550, 1998.
14. S. Sahni and T. Gonzales. P-complete approximation problems. *JACM*, 23:555–565, 1976.