



# Vector TSP: A Traveling Salesperson Problem with Racetrack-like acceleration constraints<sup>☆</sup>

Arnaud Casteigts<sup>a,b</sup>, Mathieu Raffinot<sup>b</sup>, Mikhail Raskin<sup>b</sup>, Jason Schoeters<sup>c,\*</sup>

<sup>a</sup> University of Geneva, Switzerland

<sup>b</sup> LaBRI, Université de Bordeaux, CNRS, Bordeaux INP, France

<sup>c</sup> University of Cambridge, United Kingdom

## ARTICLE INFO

### Article history:

Received 24 September 2024

Received in revised form 5 February 2025

Accepted 18 February 2025

Available online xxxx

We dedicate this paper to Ralf Klasing for his 60th birthday. Ralf had many contributions to the TSP, addressing both approximation and exact algorithms, and he also considered less common versions of the problem, such as the *Bamboo Garden Trimming Problem*, which certainly shares the feature of being as exotic as Vector TSP. Ralf, we hope you will enjoy, happy birthday!

### Keywords:

Traveling Salesperson Problem

Acceleration constraints

Motion planning

Path finding

Algorithms and heuristics

Complexity and hardness

## ABSTRACT

We study a new version of the Traveling Salesperson Problem, called VECTOR TSP, where the traveler is subject to discrete acceleration constraints, as defined in the paper-and-pencil game Racetrack (also known as Vector Racer). In this model, the degrees of freedom at a certain point in time depends on the current velocity, and the speed is not limited.

The paper introduces this problem and initiates its study, discussing also the main differences with existing versions of TSP. Not surprisingly, the problem turns out to be NP-hard. A key feature of VECTOR TSP is that it deals with acceleration in a discrete, combinatorial way, making the problem more amenable to algorithmic investigation. The problem involves two layers of trajectory planning: (1) the order in which cities are visited, and (2) the physical trajectory realizing such a visit, both interacting with each other. This interaction is formalized as an interactive protocol between a high-level tour algorithm and a trajectory oracle, the former calling the latter repeatedly. We present an exact implementation of the trajectory oracle, adapting the A\* algorithm for paths over multiple checkpoints whose ordering is given (this algorithm being possibly of independent interest). To motivate the problem further, we perform experiments showing that the naive approach consisting of solving the instance as an EUCLIDEAN TSP first, then optimizing the trajectory of the resulting tour, is typically suboptimal and outperformed by simple (but dedicated) heuristics.

© 2025 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

The problem of visiting a given set of locations and returning to the starting point, while minimizing the total cost, is known as the TRAVELING SALESPERSON PROBLEM (TSP, for short). The problem was independently formulated by Hamilton and Kirkman in the 1800s [4] and has been extensively studied since then. Many versions of this problem exist, motivated by applications in fields as varied as delivery planning [8], stock cutting [36], and DNA computation/reconstruction [25,27]. In the original version, an instance of the problem is equivalent to a graph whose vertices represent the *cities* (places to be visited) and whose weights on the edges represent the cost of moving from one city to another. The goal is to find

<sup>☆</sup> A preliminary version of this work was presented at ALGOSENSORS 2020. This work has been supported by ANR project TEMPOGRAL (ANR-22-CE48-0001).

\* Corresponding author.

E-mail addresses: [arnaud.casteigts@labri.fr](mailto:arnaud.casteigts@labri.fr) (A. Casteigts), [mathieu.raffinot@labri.fr](mailto:mathieu.raffinot@labri.fr) (M. Raffinot), [mikhail.raskin@labri.fr](mailto:mikhail.raskin@labri.fr) (M. Raskin), [js2807@cam.ac.uk](mailto:js2807@cam.ac.uk) (J. Schoeters).

the minimum cost tour (optimization version) or to decide whether a tour having at most some cost exists (decision version) subject to the constraint that every city is visited *exactly* once. Karp proved that the Hamiltonian Cycle problem is NP-hard, which implies that TSP is NP-hard [21]. On the positive side, while the trivial algorithm has a factorial running time (considering all the possible tours), Held and Karp presented in [18] a dynamic programming algorithm running in time  $O(n^2 2^n)$ , which as of today remains the fastest known deterministic algorithm for TSP (a faster randomized algorithm exists [5]). TSP was shown to be inapproximable to within a constant factor by Orponen and Manilla in 1990 [29]. Some generalizations of TSP include a version with unknown environment which is discovered while visiting the cities, known as the COVERING CANADIAN TRAVELER PROBLEM [17,24], and more recently, a version operating on a temporal graph where costs between cities are time-dependent, known as TIME DEPENDENT TSP [13]. Another problem closely related to the TSP is the BAMBOO GARDEN TRIMMING PROBLEM [16,22], where one aims to maintain the height of a forest of bamboos to a minimum.

In many cases, the problem can be restricted to a more tractable setting. In METRIC TSP, the costs must respect the triangle inequality, namely  $cost(u, v) \leq cost(u, w) + cost(w, v)$  for all  $u, v, w$ , or equivalently, the constraint of visiting a city exactly once is relaxed. METRIC TSP was shown to be approximable within a ratio of 1.5 by Christofides [9]. This ratio has been slightly improved recently by Karlin et al. in [20]. A lower bound of 1.0045 is known for the approximation ratio (i.e., no PTAS exists for METRIC TSP [31]). A particular case of METRIC TSP is the EUCLIDEAN TSP, where cities are points in the plane, and the weights are the Euclidean distance between them. This problem is still NP-hard (see Papadimitriou [30] and Garey et al. [15]), but admits a PTAS, discovered independently by Arora [2] and Mitchell [26].

One attempt to add physical constraints to the EUCLIDEAN TSP is DUBINS TSP [23], which encode inertia through bounding the curvature of a trajectory by some radius, thereby offering a geometric abstraction to the problem. However, the radius being a fixed value, the speed is considered constant. Savla et al. present multiple results on DUBINS TSP, such as an elegant algorithm in [34] which modifies every other segment of an EUCLIDEAN TSP solution to enforce the curvature constraint.

More flexible models for acceleration have been considered in the context of path planning problems, where one aims to find an optimal trajectory between *two* given locations (typically, with obstacles), while satisfying acceleration constraints. More generally, the literature on *kinodynamics* is pretty vast (see, e.g. [6,7,11]). The constraints are often formulated in terms of a number of dimensions, a bounded acceleration and a bounded speed. The positions are typically represented by real coordinates, using tools from control theory and analytic functions [32,33,37].

In a recreational column of the *Scientific American* in 1973 [14], Martin Gardner presented a paper-and-pencil game known as RACETRACK (not to be confused with a similarly titled TSP heuristic [38] that is not related to acceleration). The physical model is as follows. In each step, a vehicle moves according to a discrete-coordinate vector (initially the zero vector), with the constraint that the vector at step  $i + 1$  cannot differ from the vector at step  $i$  by more than one unit in each dimension. The game consists of finding the best trajectory (smallest number of vectors, regardless of distance) in a given race track. A nice feature of such models is the ability to think of the state of the vehicle at a given time as a point  $(x, y, dx, dy)$  in a *configuration space*. Note that this game is not a version of TSP, since only the starting and final positions are specified. In fact, an optimal trajectory in this game can be found by performing a simple breadth-first search in the configuration graph. These techniques, which we describe in more details in the present paper, were rediscovered many times, both in the RACETRACK context (see e.g. [3,12,28,35]) and in the kinodynamics literature (see e.g. [6,11])—they can now be considered as folklore. Note that a more efficient algorithm for RACETRACK instances of uniform width was proposed by Bekos et al. [3], as well as “local view” algorithms in which the track is discovered in an online fashion during the race. The theoretical complexity of the RACETRACK problem has also been studied. In particular, Holzer and McKenzie [19] proved that RACETRACK is NL-complete. They also proved that the reachability problem with a given time limit (referred to as single-player RACETRACK) is NL-complete, and that deciding the existence of a winning strategy in Gardner’s original game when the positions cannot be occupied by two players simultaneously is P-complete (under polylogarithmic time reductions).

### 1.1. Contributions

In this article, we combine TSP and RACETRACK into a single problem called VECTOR TSP, where a vehicle must visit an (initially unordered) set of coordinates (cities), using a trajectory that obeys racetrack-like constraints. The cost measure is the total number of vectors used, regardless of the distance traveled. Motivations for such a problem are varied. On the practical side, there is an increasing interest for using UAVs in delivery planning and patrolling. In the longer term, outer space travel will also be subject to similar isotropic constraints, where inertia matters more than speed. On the theoretical side, this problem is quite natural and the constraints are simplest possible in a discrete setting. Yet, it requires solutions that are significantly different from its standard versions, making this problem intriguing.

The paper starts by describing the RACETRACK physical model and some of its basic properties, as well as the concept of configuration graph and algorithmic tricks based on this graph. Then, we define the VECTOR TSP problem, together with a number of preliminary observations. In particular, we investigate the differences between VECTOR TSP and EUCLIDEAN TSP, showing that solutions visiting the cities according to an optimal EUCLIDEAN TSP tour are suboptimal for VECTOR TSP. This motivates the study of VECTOR TSP as a problem of its own.

Then, we prove that VECTOR TSP is NP-hard. This result is not surprising in itself, but the proof is far from trivial. It proceeds by defining another version of the problem called VECTOR TSP WITH STOPS, where every city must be visited

at zero speed. This problem is closer in spirit to EUCLIDEAN TSP, making it possible to adapt the arguments used by Papadimitriou's reduction from EXACT COVER to EUCLIDEAN TSP. Then, we show that VECTOR TSP WITH STOPS reduces to VECTOR TSP.

On the positive side, we present an algorithmic framework for tackling VECTOR TSP, based on an interactive scheme between a high-level algorithm and a trajectory oracle. The first is responsible for exploring the space of possible visit orders, while making queries to the second for assessing the quality of a particular visit order. We present algorithms for both. The high-level algorithm adapts known heuristics for EUCLIDEAN TSP, trying to gradually improve the solution through performing 2-permutations in the tour until a local optimum is found. As for the trajectory oracle, we present an original algorithm that generalizes  $A^*$  to the search of multi-point paths in the configuration space, with a cost estimation function that uses unidimensional projections of the distances.

We present experimental results based on this framework, giving empirical evidence that our algorithmic framework, as simple as it is, already outperforms solutions based on an optimal resolution of EUCLIDEAN TSP. The goal of these experiments is not so much to evaluate the algorithms themselves as to confirm that VECTOR TSP is an original problem that deserves further independent study.

## 1.2. Organization of the paper

In Section 2, we present the RACETRACK physical model and the VECTOR TSP problem, together with a number of algorithmic tricks and observations. In Section 3, we characterize a number of basic properties of VECTOR TSP and compare it with EUCLIDEAN TSP. We prove that VECTOR TSP is NP-hard in Section 4. The algorithmic framework is presented in Section 5, and experiments are reported in Section 5.3. Finally, we conclude with some remarks and future research directions in Section 6.

## 2. Model and definitions

In this section, we present the RACETRACK model and some of its algorithmic features. Then, we define VECTOR TSP and establish a few facts about it.

### 2.1. The racetrack model

The RACETRACK model [14] is an acceleration model where both time and space are discrete. The model specifies what movements a vehicle can make based on its current position and velocity. In the original formulation of [14], the space is two-dimensional, though the model generalizes naturally to higher dimensions.

#### Configuration

The state of the vehicle at a certain time is called a *configuration*. A configuration  $c$  is encoded by a position  $pos(c)$  and a velocity  $vel(c)$ , both in  $\mathbb{Z}^2$ , being possibly abbreviated as a 4-tuple  $(x, y, dx, dy)$ , where  $dx$  and  $dy$  correspond to the movements the vehicle just made in the last move (initially 0). Given a configuration  $c$ , the constraints consist of restricting the set of legal successors  $succ(c)$ , which are the configurations  $c'$  whose velocity differ from  $vel(c)$  by at most one unit in each dimensions, and whose positions correspond to moving the vehicle according to this chosen velocity. (Morally, a configuration encodes a vector/movement that *has just been* performed.) Precisely, the successors of a configuration  $c_i = (x_i, y_i, dx_i, dy_i)$  are all the configurations  $c_j = (x_j, y_j, dx_j, dy_j)$  such that  $|dx_j - dx_i| \leq 1$ ,  $|dy_j - dy_i| \leq 1$ ,  $x_j = x_i + dx_j$  and  $y_j = y_i + dy_j$ . These constraints are illustrated in Fig. 1. Some versions of the model consider stronger restrictions, by allowing a single dimension to change in one time step. We refer to both versions as the 9-successor model and the 5-successor model, respectively. These models generalize naturally to  $d$ -dimensional spaces by applying similar constraints for each dimension. For any fixed number of dimensions  $d$ , a key aspect is that the number of successors is *constant*.

#### Trajectory

A trajectory is a sequence of configurations  $c_1, c_2, \dots, c_k$ , with  $c_{i+1} \in succ(c_i)$  for all  $i < k$ . The cost of a trajectory is its length  $k$ . The fact that acceleration constraints are isotropic in the RACETRACK model imply that trajectories are *reversible*. The *reverse* of a configuration  $c = (x, y, dx, dy)$  is the configuration  $c^{-1} = (x + dx, y + dy, -dx, -dy)$ , i.e. the same movement backwards. (One can think of two superposed vectors in opposite directions.) Thus, if  $(c_1, c_2, \dots, c_k)$  is a trajectory, then  $(c_k^{-1}, \dots, c_2^{-1}, c_1^{-1})$  is also a trajectory.

#### Configuration space

The set of all configurations together with the successor relation define an infinite graph called the *configuration graph*. Let  $\mathcal{C}$  be the infinite set of all configurations, the configuration graph is the directed graph  $G = (V, E)$  where  $V = \mathcal{C}$  and  $E = \{(c_i, c_j) \subseteq \mathcal{C} \times \mathcal{C} : c_j \in succ(c_i)\}$ .

Basic path-finding algorithms using the configuration graph have been independently discovered many times, we consider them as folklore. Concretely, a trajectory in the original space is a path in the configuration graph. This correspondence allows for searching optimal trajectories using standard graph search like breadth-first search (BFS), in an appropriate subset of  $G$  (computed on the fly).

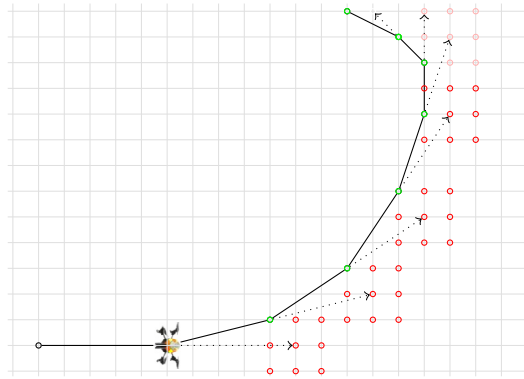


Fig. 1. Example of a trajectory in the RACETRACK model.

### 2.2. The VECTOR TSP problem

We introduce the VECTOR TSP problem, which consists of finding a minimum length trajectory that visits a given set of cities and returns to the starting point, starting and finishing at zero speed, and subject to RACETRACK-like physical constraints. Note that we allow to pass over a city without a visit. The order in which the cities are visited (i.e., the *tour*) is not specified in the input. The goal can thus be seen as determining a *tour* and a trajectory *realizing* this tour.

To define the problem more formally, let us define what *visiting* a city means. Here, we say that a city (or more generally, a point)  $p$  is *visited* by a configuration  $c$  if  $p$  lies on the segment between the point  $p_a = pos(c) - vel(c)$  and  $p_b = pos(c)$ . If several cities are visited by a same configuration, then these cities are considered visited in increasing order of distances from  $p_a$ .

#### VECTOR TSP

**Input:** A set of  $n$  cities (points)  $P \subseteq \mathbb{Z}^d$ , for some fixed  $d$ , a distinguished city  $p_0 \in P$ , and an integer  $k$  encoded in unary.

**Question:** Does there exist a trajectory  $\mathcal{T} = (c_1, \dots, c_k)$  of length at most  $k$  that visits all the cities in  $P$ , with  $pos(c_1) = pos(c_k) = p_0$  and  $vel(c_1) = vel(c_k) = \vec{0}$ .

Receiving parameter  $k$  in unary guarantees that the problem is not hard for artificial reasons. For example, it prevents a trivial instance from being hard just because the cities are exponentially far from each other (in the size of the input, see also [12,19]). The optimization version of VECTOR TSP is defined analogously, except that instead of providing  $k$  in unary, one can replace it with any number in unary that is a (polynomially related) lower bound on the optimal trajectory, such as the length of a trivial trajectory between the two more distant cities.

#### Stopping at the cities


Although the spirit of the model is that speed is unbounded, it could be useful in some scenarios to stop physically at the cities (e.g. to avoid blurred pictures, or to drop objects in a delivery scenario). We define the problem VECTOR TSP WITH STOPS similarly as VECTOR TSP, with the additional requirement that the configurations visiting a city must have zero speed. Technically, this version of the problem can be seen as somewhat intermediate between EUCLIDEAN TSP and VECTOR TSP. We do not investigate it *per se* in the present paper, but we use it as a corner stone in our proof of NP-hardness of VECTOR TSP in Section 4.

#### Tour/trajectory/solution

As already explained, in VECTOR TSP, a *tour* denotes a visit order of the cities (i.e. a permutation), and a *trajectory* denotes a sequence of racetrack configurations (that may realize such a tour). Given a tour  $\pi$ , we call  $racetrack(\pi)$  the set of optimal trajectories realizing the tour  $\pi$ . It is important to keep in mind that the optimality of trajectories in  $racetrack(\pi)$  is only relative to  $\pi$ , it does not imply that these trajectories are optimal solutions to the problem instance (if  $\pi$  is itself suboptimal). If no other tour than  $\pi$  exists that admit a shorter realizing trajectory, then, in this case only, a trajectory in  $racetrack(\pi)$  is an *optimal solution* for the VECTOR TSP problem.

### 3. Intermediate results and observations

In this section, we establish a number of preliminary results that highlight the differences between VECTOR TSP and EUCLIDEAN TSP. We also present key properties of the configuration graph, used in subsequent sections.



**Fig. 2.** On the left, a trajectory in  $\text{racetrack}(\pi)$ , where  $\pi$  is an optimal tour for EUCLIDEAN TSP; on the right, an optimal VECTOR TSP solution using less vectors along a different tour. The loop on  $p_0$  illustrates the final “move” required to finish at zero speed. The alternation of colors is used to help distinguish individual moves.

### 3.1. Differences between EUCLIDEAN TSP and VECTOR TSP

VECTOR TSP having essentially the same input as EUCLIDEAN TSP, namely, a set of coordinates, it is tempting to see it as just a small variant of that problem. Both problems are however truly different.

To start, observe that the initial city does matter in VECTOR TSP. This is easy to see on a set of aligned cities. If the vehicle starts at one of the endpoint cities, then it is able to visit all the cities and return after two phases of acceleration and deceleration, whereas if it starts in the middle, at least three such phases will be needed, resulting in a solution of higher cost despite the fact that the traveled distance is the same.

Beyond the starting point, we now show that optimal tours for EUCLIDEAN TSP may be suboptimal for VECTOR TSP.

**Lemma 1.** *There exists an instance  $\mathcal{I}$  of VECTOR TSP, such that the optimal tour  $\pi$  for the corresponding EUCLIDEAN TSP instance is suboptimal for  $\mathcal{I}$ .*

**Proof.** We provide a simple example in Fig. 2. On the left, a trajectory in  $\text{racetrack}(\pi)$ , where  $\pi$  is an optimal tour for EUCLIDEAN TSP, starting and ending at  $p_0$  (whence the final deceleration loop). Realizing this trajectory requires 22 moves, but an optimal trajectory for VECTOR TSP takes only 20 (right picture) and corresponds to a tour that would be suboptimal for EUCLIDEAN TSP.  $\square$

From the same example, observe that optimal VECTOR TSP solution may self-cross. In fact, the cost of traveling between two given cities in VECTOR TSP depends crucially on choices that are made earlier in the trajectory. This feature, together with the fact that traveling some distance costs essentially the square root of that distance, makes VECTOR TSP significantly different from EUCLIDEAN TSP and not reducible to it (at least, not straightforwardly), nor even to general TSP (whose costs between cities are fixed).

Later on, in this paper, we give evidence that this divergence between both problems is not occasional. For sufficiently large random instances, optimal EUCLIDEAN TSP tours become *typically* suboptimal for VECTOR TSP.

### 3.2. The configuration graph is (in effect) bounded

In theory, the positions and velocities are unbounded, which makes the configuration graph infinite. Nevertheless, we show below that the *useful part* of the configuration graph is bounded. Moreover, an optimal trajectory can always be found within a subgraph of the configuration graph whose size is polynomial in the size of the input.

**Lemma 2 (Length of an Optimal Solution).** *Let  $\mathcal{I}$  be a VECTOR TSP instance based on a set  $P$  of  $n$  cities, let  $L$  be the maximum distance in a dimension between two cities of  $P$  (euclidean max norm). An optimal solution contains at most  $O(nL)$  configurations.*

**Proof.** A naive (suboptimal) solution consists of visiting the cities in an arbitrary order at *unit speed* (i.e. the max norm of the speed is one). Such a trajectory contains  $O(nL)$  configurations, which is thus an upper bounds on the size of an optimal solution.  $\square$

Lemma 2 implies that the useful part of the configuration graph can be restricted to a subgraph of *polynomial* size in the size of the input.




Fig. 3. Computation of an optimal (one dimensional) trajectory for  $k = 12$ .

**Lemma 3** (Bounds on the Configuration Graph). Let  $\mathcal{I}$  be a ( $d$ -dimensional) VECTOR TSP instance based on a set  $P$  of  $n$  cities and a starting city  $p_0 \in P$ . Let  $L$  be the maximum distance between two cities of  $P$  in a dimension (euclidean max norm). An optimal trajectory for  $\mathcal{I}$  can be found in a subgraph of  $G$  that contains at most  $O((nL)^{3d})$  configurations.

**Proof.** For simplicity, consider a renormalized instance in which the coordinate of  $p_0$  is 0 in each dimension. By Lemma 2, an optimal solution has at most  $O(nL)$  configurations, thus the maximum absolute value of a position in any dimension is at most  $O((nL)^2)$ , which is attained, for example, if one accelerates all the time in this dimension. Similarly, the maximum absolute value of the velocity in any dimension is  $O(nL)$ , which implies at most  $O((nL)^3)$  values for the combined position and velocity in each dimension, and thus at most  $O((nL)^{3d})$  values in  $d$  dimensions.  $\square$

**Corollary 1.** In two dimensions, an optimal trajectory can be found in a subgraph of the configuration graph containing  $O((nL)^6)$  vertices and edges.

### 3.3. One-dimensional cost computation

Some of the computation in the subsequent sections rely on computing the cost of trajectories in one dimension. Here, we first prove a direct formula for expressing the cost of moving from one position to another in one dimension, starting and finishing at zero speed. Then, we explain how more general one-dimensional computation reduces to it.

**Lemma 4.** In one dimension, the cost of covering  $k$  space units, starting and finishing at zero speed is  $\lceil 2\sqrt{k} \rceil$ .

**Proof.** Starting at zero speed, a vehicle accelerating over  $s$  steps covers a distance of  $1 + 2 + \dots + s = s(s + 1)/2$  units. Symmetrically, a vehicle decelerating down to zero speed over  $s$  steps also covers  $s(s + 1)/2$  units. Thus, if  $k = s(s + 1)$  for some integer  $s$ , a strategy is to accelerate for  $s$  steps and then decelerate for  $s$  steps, which maximizes the average speed and is thus optimal (see the top scenario in Fig. 3). Together with the final loop for reaching zero speed, this corresponds to a total cost of  $2s + 1$ . And since  $s^2 < k < (s + \frac{1}{2})^2$ , this corresponds to a cost of  $\lceil 2\sqrt{k} \rceil$ .

If  $k$  cannot be expressed as the product of two consecutive integers, then let  $s$  be the largest integer such that  $s(s + 1) < k$ . There are two cases:

1. If  $s(s + 1) < k \leq (s + 1)^2$ , then  $s^2 + s + 1 \leq k \leq s^2 + 2s + 1$  (since  $k$  is an integer). In this case, there is a non-empty gap in the middle and a trajectory of cost  $2s + 1$  (including the final loop) cannot exist. However, we will show that a trajectory of cost  $2s + 2$  always exist in this case. Let  $k' \leq s + 1$  be the size of the gap ( $k' = 1$  in Fig. 3, middle scenario). One can complete the trajectory by adding a single additional move of length  $k'$ . If  $k' \geq s - 1$ , this move can be inserted in the middle. If  $k' < s - 1$ , then a move of length  $k'$  already exists in the deceleration part. It suffices to repeat this move and shift the previous ones accordingly (dashed green vector in Fig. 3, bottom scenario) in order to obtain a trajectory of cost  $2s + 2$ . Since  $(s + 0.5)^2 < k \leq (s + 1)^2$ , this cost amounts again to  $\lceil 2\sqrt{k} \rceil$ .
2. If  $(s + 1)^2 < k < (s + 1)(s + 2)$ , then  $s^2 + 2s + 2 \leq k \leq s^2 + 3s + 1$  (since  $k$  is an integer). In this case, there is a gap of  $k'$  units, with  $s + 2 \leq k' \leq 2s + 1$ . No trajectory with cost  $2s + 2$  exists, since  $k'$  exceeds the size allowed for an additional vector. However, one can complete the trajectory by splitting  $k'$  in a sum of two numbers  $\leq s + 1$  and insert two additional such moves in the decelerating segment. Since  $(s + 1)^2 < k \leq (s + 1.5)^2$ , this amounts to a cost of  $2s + 3 = \lceil 2\sqrt{k} \rceil$ .  $\square$

Observe, from the previous lemma, that an optimal trajectory between two points in one dimension, starting and finishing at zero speed, can always be realized by first accelerating as much as possible without overpassing the middle point, then filling the potential gaps upon deceleration. Another case of interest, used in Section 5.2, is when the vehicle (with position  $x$  and speed  $dx$ ) is moving towards point  $x'$ , and it is able to stop before bypassing  $x'$ . The strategy in this



case is to simulate that the vehicle has accelerated continuously  $dx$  time steps from a virtual starting point located at  $x'' = x - dx(dx + 1)/2$  and to compute the cost from  $x''$  to  $x'$  starting and finishing at zero speed. By the above remark, the resulting cost (minus the  $dx$  virtual steps) must correspond to the optimal cost towards  $x'$  from the current position and speed.

Also remark that for a pair of cities, placed in any number of dimensions, the cost when starting and stopping at zero speed, is  $\lceil 2\sqrt{K} \rceil$  where  $K$  is the largest distance among all dimensions.

#### 4. NP-hardness of VECTOR TSP

Although the result itself is not surprising, proving that VECTOR TSP is NP-hard does not appear to be straightforward. In this section, we do so in two steps. First, we prove that VECTOR TSP WITH STOPS is NP-hard (Theorem 2), by adapting some ideas from Papadimitriou [30] in his proof that EUCLIDEAN TSP is NP-hard, via a reduction from EXACT COVER. The main conceptual difficulty in this adaptation is to account for the fact that acceleration allows one to travel quadratic distance in linear time, and counter rounding effects by a scaling argument in continuous space (Theorem 1). Then, we show that VECTOR TSP WITH STOPS reduces to VECTOR TSP (Theorem 3).

##### 4.1. NP-hardness of VECTOR TSP WITH STOPS

We choose to do most of the reasoning in this section on a slightly modified version of VECTOR TSP WITH STOPS, referred to as CONTINUOUS VECTOR TSP WITH STOPS.

###### CONTINUOUS VECTOR TSP WITH STOPS

**Input:** A set of  $n$  cities (points)  $P \subseteq \mathbb{Z}^d$ , with intercity cost  $c(p_i, p_j) = 2\sqrt{K}$  where  $K$  is the largest distance between the cities among all dimensions, and an integer  $k$ .

**Question:** Does there exist a trajectory  $\mathcal{T} = (c_1, \dots, c_k)$  of cost at most  $k$  that visits all the cities in  $P$ , with  $pos(c_1) = pos(c_k) = p_0$  and  $vel(c_1) = vel(c_k) = \vec{0}$ .

CONTINUOUS VECTOR TSP WITH STOPS allows us to avoid some technical details of VECTOR TSP WITH STOPS, such as the rounding effects on costs between cities. These effects are taken care of later through a scaling argument which transfers NP-hardness results from CONTINUOUS VECTOR TSP WITH STOPS to VECTOR TSP WITH STOPS. This version is closer in spirit to the EUCLIDEAN TSP, allowing us to adapt Papadimitriou’s NP-hardness proof for EUCLIDEAN TSP [30]. Note that we can talk about paths for CONTINUOUS VECTOR TSP WITH STOPS instead of trajectories (as there is just the notion of a cost between cities, not of rounds/vectors). In the case of VECTOR TSP WITH STOPS, either paths or trajectories can be considered.

The reduction is from EXACT COVER, whose definition is as follows. Let  $\mathcal{U}$  be a set of  $m$  elements (the *universe*), EXACT COVER takes as input a set  $\mathcal{F} = \{F_i\}$  of  $n$  subsets of  $\mathcal{U}$ , and asks if there exists  $\mathcal{F}' \subseteq \mathcal{F}$  such that all sets in  $\mathcal{F}'$  are *disjoint* and  $\mathcal{F}'$  covers all the elements of  $\mathcal{U}$ . For example, if  $\mathcal{U} = \{1, 2, 3\}$  and  $\mathcal{F} = \{\{1, 2\}, \{3\}, \{2, 3\}\}$ , then  $\mathcal{F}' = \{\{1, 2\}, \{3\}\}$  is a valid solution, but  $\{\{1, 2\}, \{2, 3\}\}$  is not.

**High-level view of the reduction:** Given an instance  $\mathcal{I}$  of EXACT COVER, we construct an instance  $\mathcal{I}'$  of CONTINUOUS VECTOR TSP WITH STOPS such that  $\mathcal{I}$  admits a solution if and only if there is a trajectory visiting all the cities of  $\mathcal{I}'$  using at most a prescribed number of vectors.  $\mathcal{I}'$  is composed of several types of gadgets, representing respectively the subsets  $F_i \in \mathcal{F}$  and the elements of  $\mathcal{U}$  (with some repetition). For each  $F_i$ , a subset gadget  $C_i$  is created which consists of a number of cities placed horizontally (wavy horizontal segments in Fig. 4). For now, we state that each gadget can be traversed optimally in exactly two possible ways (without considering direction), which ultimately corresponds to including (mode 1) or excluding (mode 2) subset  $F_i$  in the EXACT COVER solution. The  $C_i$ ’s are located one below another, starting with  $C_1$  at the top. Between every two consecutive gadgets  $C_i$  and  $C_{i+1}$ , copies of *element* gadgets are placed for each element in  $\mathcal{U}$ , thus the element gadgets  $H_{i,j}$  are indexed by both  $1 \leq i \leq n - 1$  and  $1 \leq j \leq m$  (see again Fig. 4). The element gadgets are also made of a number of cities, whose particular organization is described later on. Finally, every subset gadget  $C_i$  above or below an element gadget representing element  $j$  is modified in a way that represents whether  $F_i$  contains element  $j$  or not.

Intuitively, a tour visiting all the cities must choose between inclusion or exclusion of each  $F_i$  (i.e., mode 1 or 2 for each  $C_i$ ). An element  $j \in \mathcal{U}$  is considered as covered by a subset  $F_i$  if  $C_i$  does *not* visit any of the adjacent gadgets representing  $j$ . Each element gadget  $H_{i,j}$  must be visited either from above (from  $C_i$ ) or from below (from  $C_{i+1}$ ). The number of subset gadgets is  $n$ , the number of element gadgets for each element is  $n - 1$  (one between every two consecutive subset gadgets), and the construction guarantees that at most one element gadget for each element  $j \in \mathcal{U}$  is visited from a subset gadget  $C_i$  (or the tour is suboptimal). These three properties collectively imply that for each element  $j \in \mathcal{U}$ , there is exactly one subset gadget  $C_i$  that does not visit any of the element gadgets representing  $j$ .

In summary, the tour proceeds from the top left corner through the  $C_i$ ’s (in order), visiting all the  $H_{i,j}$ . So long as a  $C_i$  visits a  $H_{i,j}$  (thus, from above), this means that element  $j$  has not yet been covered in the EXACT COVER solution. Element  $j$  is covered by subset  $F_i$  in the EXACT COVER solution if  $C_i$  is the subset gadget that does *not* visit the corresponding  $H_{i,j}$ , after which all the  $H_{k \geq i,j}$  will necessarily be visited from below by the corresponding  $C_{k+1}$ , ensuring the element  $j$  will not be covered more than once.

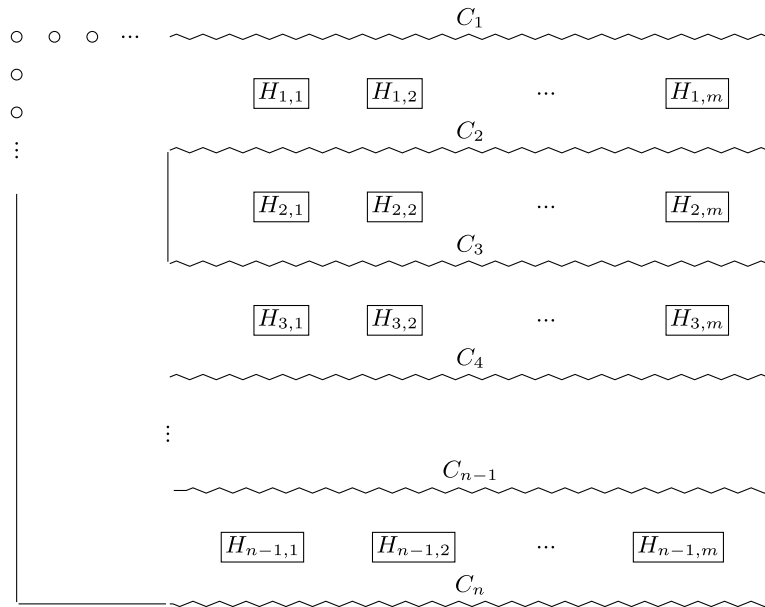


Fig. 4. Papadimitriou's high-level construction.




Fig. 5. 1-chain structure which turns 90° twice. The distance between consecutive cities is 1. The optimal visit order is shown in (b). We abbreviate a 1-chain schematically as shown in (c).

**Construction 1:** Let us now present the details of the gadgets and how to assemble them together, so as to transform any given EXACT COVER instance into a CONTINUOUS VECTORTSP WITH STOPS instance.<sup>1</sup> Consider the 1-chain structure presented in Fig. 5. This structure is composed of cities positioned on a line, at distance 1 from one another. 1-chains can bend at 90 degrees angles. Next, consider the structure in Fig. 6, referred to as a 2-chain. The distance between the leftmost (or rightmost) city and its nearby cities is  $\frac{\sqrt{2}}{2}$ . The closest distance between vertical cities is 1, and 2 between horizontal cities.


Finally, consider the A, B, and H structures, presented in Figs. 7 to 9 respectively. Structures A and B will replace parts of 2-chains as shown, and consist of cities positioned on a grid with distance 1 between nearest cities, except for the cities closest to the middle, which are horizontally distanced by 2 and 6 respectively. Structure H consists of a symmetric set of cities which will be positioned far away from all other gadgets.

Construct the structure shown in Fig. 10, where  $n$  is the number of subsets given in the corresponding EXACT COVER instance, and  $m$  the number of elements in the universe.


Then, for every 2-chain  $C_i$ , remove the cities positioned directly above or below an H. In other words, if the H structure is contained in a bounding box of coordinates  $(x_1, y_1)$  to  $(x_2, y_2)$ , the cities of  $C_i$  between  $x$ -coordinates  $x_1$  and  $x_2$  are concerned. Insert in this emptied space one of two structures, depending on the elements in  $F_i$ ,  $C_i$ 's corresponding subset. If the subset contains the element corresponding to the above (or below) H, then replace by structure A (see Fig. 8), otherwise replace by structure B (see Fig. 9). Finally, rotate the whole city set by 45°. As the city coordinates should be

<sup>1</sup> The cost  $k$  will be missing from the instance, but we will specify this cost in Theorem 1.

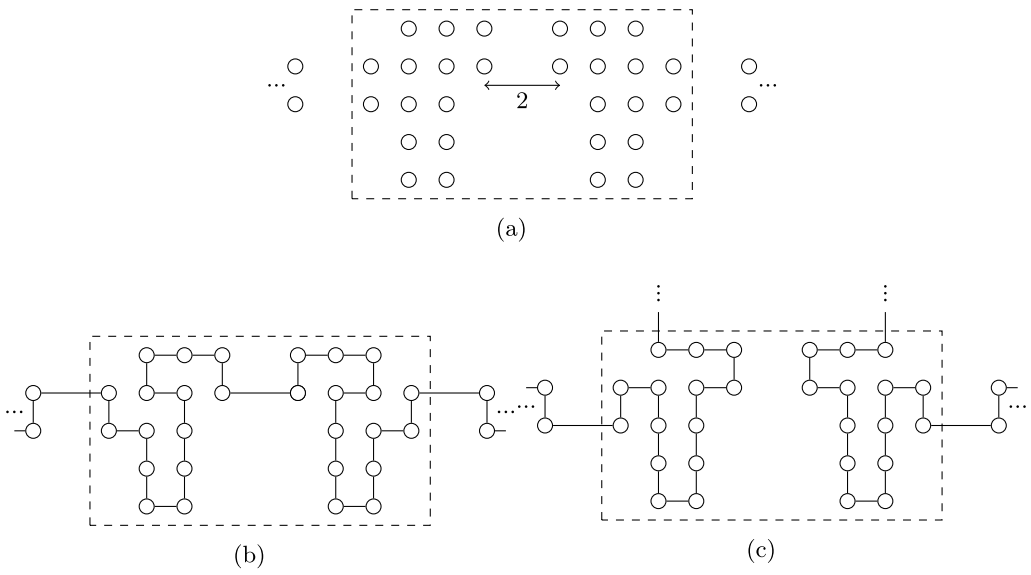




**Fig. 6.** 2-chain structure (a). A 2-chain has precisely two optimal 1-trajectories, mode 1 shown in (b) and mode 2 shown in (c). We abbreviate a 2-chain schematically as shown in (d).



**Fig. 7.** Structure  $H$ . An optimal 1-trajectory in  $H$  is shown in (b). We abbreviate an  $H$  structure schematically as shown in (c).



**Fig. 8.** Structure  $A$  with parts of a 2-chain on the sides (see (a)). Visiting structure  $A$  in mode 1 makes it suboptimal to visit an  $H$  structure above or below (see (b)). Visiting structure  $A$  in mode 2 however, makes it less costly to visit an  $H$  structure above or below (see for example (c) for the former).

integer, after the rotation we scale everything by  $\sqrt{2}$ . (The rotation and the scaling are the only changes we add w.r.t. Papadimitriou’s original construction.) This concludes Construction 1.

The following definitions are from Papadimitriou [30]. A subset  $Q$  of the set of cities is an  $a$ -component if for all  $q \in Q$  we have  $\min(\text{cost}(q, p) : p \notin Q) \geq a$  and  $\max(\text{cost}(q, p) : p \in Q) < a$ . A  $k$ -path for a set of cities is a set of  $k$ , not necessarily closed paths visiting all cities. A solution for a CONTINUOUS VECTOR TSP WITH STOPS instance is thus a closed

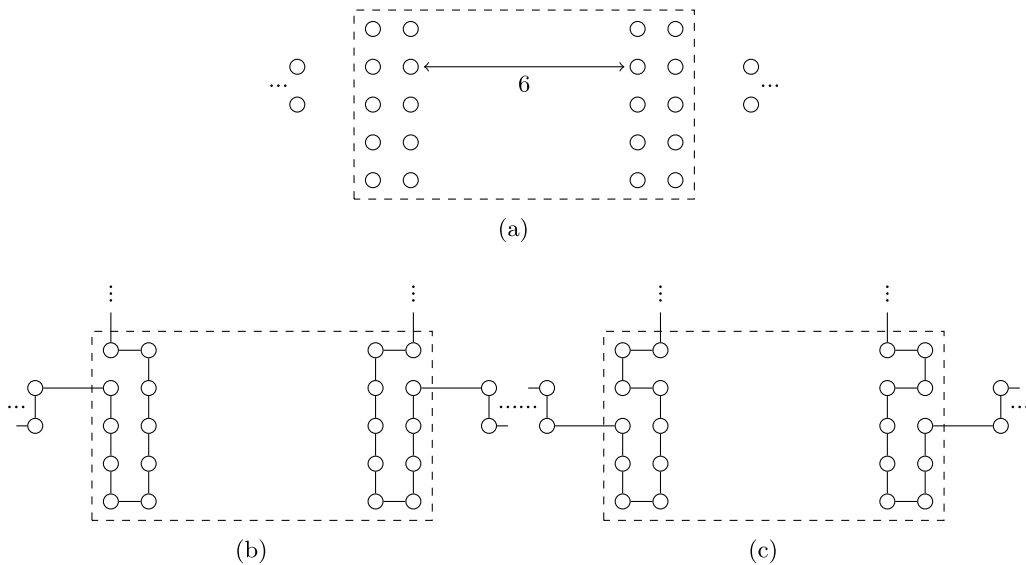


Fig. 9. Structure B (see (a)). Visiting structure B in any mode allows it to optimally visit an H structure above or below (see (b) and (c)).




Fig. 10. Construction 1 from EXACT COVER to CONTINUOUS VECTORTSP WITH STOPS (before adding structures A and B and rotating the city set).

(or cyclic) 1-path. A subset of cities is *a*-compact if, for all positive integers *k*, an optimal *k*-path has cost less than the cost of an optimal (*k* + 1)-path plus *a*. Note that *a*-components are trivially *a*-compact.

**Lemma 5** (Papadimitriou [30]). Suppose we have *N* *a*-components  $P_1, \dots, P_N \subseteq P$ , such that the cost to connect any two components is at least  $2a$ , and  $P_0 = P \setminus \bigcup_{i=1}^N P_i$  is *a*-compact. Suppose that any optimal 1-path of this instance does not directly connect any two *a*-components (but rather indirectly connects them through  $P_0$ ). Let  $K_1, \dots, K_N$  be the costs of the optimal 1-paths of  $P_1, \dots, P_N$  and  $K_0$  the cost of the optimal (*N* + 1)-path of  $P_0$ . If there is a 1-path *S* of *P* consisting of the union of an optimal (*N* + 1)-path of  $P_0$ , *N* optimal 1-paths of  $P_1, \dots, P_N$  and  $2N$  paths of cost *a* connecting *a*-components to  $P_0$ , then *S* is optimal. If no such 1-path exists, the optimal 1-path of *P* has a cost greater than  $K = K_0 + K_1 + \dots + K_N + 2Na$ .

Papadimitriou proposed the following optimal *k*-paths in his construction:

- For 1-chains, the 1-path shown in Fig. 5;
- For 2-chains, the path zigzagging one way (mode 1) or the other (mode 2) as shown in Fig. 6;



**Fig. 11.** Example of an undesired but optimal 1-path for a 2-chain before rotation (a). Papadimitrou's proposed optimal 1-path (in mode 2) after rotation (b).

- For *A* structures, a 1-path for mode 1, and a 2-path for mode 2, such as those shown in Fig. 8. Note that no 2-path exists for mode 1 such that the endpoints would align with vertices of a structure *H* directly above or below (as they do for the 2-path in mode 2);
- For *B* structures, independent of the mode, a 2-path aligning with vertices of a structure *H* above or below such as shown in Fig. 9;
- For *H* structures, the 1-path shown in Fig. 7. Among all 1-paths for *H* having as endpoints two of the cities *a, a', b, b', c, c', d, d'*, there are 4 optimal 1-paths, namely those with endpoints (*a, a'*), (*b, b'*), (*c, c'*), (*d, d'*).

It can be shown that these *k*-paths are optimal for CONTINUOUS VECTORTSP WITH STOPS as well, although due to the nature of CONTINUOUS VECTORTSP WITH STOPS, other paths would have been optimal as well before rotating the city set, for example for 2-chains, the path shown in Fig. 11 is optimal, but it completely disobeys any "mode". After rotation, only the described *k*-paths are optimal (see again Fig. 11).

These optimal *k*-paths in our CONTINUOUS VECTORTSP WITH STOPS construction have a cost of:

- $2(n - 1)$  for a 1-chain of *n* cities, or  $2\ell$  for a 1-chain of length  $\ell$ . Let us denote the latter as  $c_{1,\ell}$ ;
- $6 + (2 + 2\sqrt{2})\frac{n-4}{2}$  for a 2-chain of *n* cities, or  $6 + (2 + 2\sqrt{2})\frac{\ell-1}{2}$  for a 2-chain of length  $\ell$ . Let us denote the latter as  $c_{2,\ell}$ ;
- $52 + 2\sqrt{2}$  for a structure *A* in mode 1 (resp. 52 in mode 2), or  $42 - 6\sqrt{2}$  (resp.  $42 - 8\sqrt{2}$ ) more than the part of the 2-chain tour it replaces. Let us denote the latter as  $c_{A,1}$  (resp.  $c_{A,2}$ );
- 36 for a structure *B* independent of the mode, or  $26 - 8\sqrt{2}$  more than the part of the 2-chain tour it replaces. Let us denote the latter as  $c_B$ ;
- $36 + 8\sqrt{2} + 2\sqrt{6}$  for a structure *H*. Let us denote this cost as  $c_H$ .

We are now ready to prove the main theorem based on the presented construction and lemmas.

**Theorem 1.** CONTINUOUS VECTORTSP WITH STOPS is NP-hard.

**Proof.** We reduce from EXACT COVER with Construction 1. We show that this construction fits the hypothesis of Lemma 5. Observe that connecting a *k*-path from some 2-chain  $C_i$  (with structures *A* and/or *B*) with a *k*-path from some  $H_{i,j}$  (or  $H_{i-1,j}$ ) has optimal cost  $4\sqrt{5}$ . This optimal cost is only attained by connecting some extremity city of structure *H* (*a, a', b, b', c, c', d, d'* in Fig. 7) to the closest city of 2-chain *C* positioned at distance 20 with a 45° angle. The optimal cost to connect two *k*-paths between two  $H_{i,j}$ , is more than  $4\sqrt{5}$  (it is at least  $4\sqrt{11}$  to connect  $H_{i,j}$  and  $H_{i+1,j}$  and at least 20 to connect  $H_{i,j}$  and  $H_{i,j+1}$ ) vectors. It should be clear that an optimal 1-path must have *X* and *Y* as endpoints (see again Fig. 10). This construction meets the hypotheses of Lemma 5 with  $a = 4\sqrt{5}$ ,  $N = m(n - 1)$ ,  $K_1 = \dots = K_N = c_H$  and  $K_0 = 2c_{1,1000} + 2(n - 1)c_{1,100} + (n - 1)c_{1,51} + nc_{2,8m+100(m-1)} + mc_{A,1} + (p - m)c_{A,2} + (nm - p)c_B$ , where *p* is the sum of cardinalities of all given subsets of the EXACT COVER instance, which in turn describes the ratio of *A* structures vs *B* structures, of which the costs are distinct. Note that some distances in the construction (such as 22 and 21) were chosen very carefully (so the cost between *A* (or *B*) structures and *H* structures is exactly *a*, as the former are larger by 2 units at the bottom, and 1 unit at the top w.r.t. 2-chains). We examine when this structure has an optimal 1-path *S*, as described in the lemma. *S* traverses all 1-chains in the obvious way, and all 2-chains in one of the two modes. Since its portion on  $P_0$  has to be optimal, *S* must visit a structure *H* from any structure *B* encountered, and it must return to the symmetric city of *B*, since its portion on *H* must be optimal, too. If *S* encounters a structure *A* and the corresponding chain is traversed in mode 2, *S* will also visit a structure *H*. However, if the corresponding 2-chain is traversed in mode 1, *S* will traverse *A* without visiting any structure *H*, since trajectories connecting  $P_0$  and any *H* structure must be of cost *a*. Moreover this must happen exactly once for each column of the structure, since there are  $n - 1$  copies of *H* and *n* structures *A* or *B* in each column. (In particular, this implies that an *H* structure not visited from the current 2-chain must be visited from the following 2-chain.) Hence, if we consider the fact that  $C_j$  is traversed in mode 1 (resp. mode 2) to mean that the corresponding subset is (resp. is not) contained in the EXACT COVER solution, we see that the existence of a 1-path *S*, as

described in Lemma 5, implies the EXACT COVER instance admits a solution. Conversely, if the EXACT COVER instance admits a solution, we assign, as above, modes to the chains according to whether or not the corresponding subset is included in the solution. It is then possible to exhibit a 1-path  $S$  meeting the requirements of Lemma 5. Hence the structure at hand has a 1-path of cost no more than  $K = K_0 + m(n - 1)(2a + c_H)$  if and only if the given instance of EXACT COVER is positive. Finally, to obtain a valid CONTINUOUS VECTOR TSP WITH STOPS tour, connect both endpoints  $X$  and  $Y$  in Fig. 10 with a 1-chain, and increase  $K$  accordingly.  $\square$

**Lemma 6.** *The difference in cost between positive CONTINUOUS VECTOR TSP WITH STOPS instances and negative ones (of a same size) in Construction 1 is at least  $\sqrt{42} - 2\sqrt{10} > 0.15$ .*

**Proof.** In Papadimitriou's construction, it is stated that the minimum precision necessary is  $\sqrt{401} - 20 > 0.02$ , which corresponds to distinguishing between optimally connecting a 1-path from a structure  $A$  or  $B$  to a 1-path from a structure  $H$ , or connecting it suboptimally (all other suboptimal connections/paths can be recognized with lower precision). This implies that a gap between positive and negative instances must be at least this large.

In Construction 1, this distinction remains the smallest one (i.e. the one requiring the highest precision), and the gap evaluates to  $\sqrt{42} - 2\sqrt{10} > 0.15$ .  $\square$

We remark that scaling up the construction by some factor  $r^2$  multiplies the cost of a solution tour, as well as the gap, by  $r$ .

**Lemma 7.** *If an optimal CONTINUOUS VECTOR TSP WITH STOPS tour has cost  $K$  on an instance of  $n$  cities, then an optimal VECTOR TSP WITH STOPS trajectory has a cost of at most  $K + n$ .*

**Proof.** The cost between any pair of cities  $p_i, p_j$  in VECTOR TSP WITH STOPS is  $c(p_i, p_j) = \lceil 2\sqrt{K} \rceil$  where  $K$  is the largest distance between the cities among all dimensions, and for CONTINUOUS VECTOR TSP WITH STOPS this cost is  $2\sqrt{K}$ . Note that the difference between such costs is at most 1. An optimal tour in CONTINUOUS VECTOR TSP WITH STOPS is exactly  $n$  intercity costs (in some order), meaning the difference between optimal trajectories is at most  $n$ .  $\square$

**Theorem 2.** VECTOR TSP WITH STOPS is NP-hard.

**Proof.** Take Construction 1 and scale it up by a sufficiently large factor, say  $R = (100n)^2$ . The gap from Lemma 6 is now at least  $13n$ . Since CONTINUOUS VECTOR TSP WITH STOPS is NP-hard (Theorem 1), and since the cost of an optimal VECTOR TSP WITH STOPS trajectory is at most that of an optimal CONTINUOUS VECTOR TSP WITH STOPS tour plus  $n$  (Lemma 7), which fits within the gap of size at least  $13n$ , VECTOR TSP WITH STOPS is also NP-hard.  $\square$

Note that a gap of any (constant) size is attainable for VECTOR TSP WITH STOPS, by changing  $R$  in the proof of Theorem 2 accordingly. This will be useful for proving that VECTOR TSP is NP-hard in Section 4.2.

#### 4.2. NP-hardness of VECTOR TSP


Through a scaling argument and use of the gap from Section 4.1, we show that the NP-hardness from VECTOR TSP WITH STOPS (Theorem 2) transfers to VECTOR TSP.

**Construction 2 and Construction 3:** Consider Construction 1. Define grid factor  $G = 100n^2$ , and the scaling factor  $S = \frac{100}{0.13}G^2n^3$ . As established in Section 4.1, the gap between positive VECTOR TSP WITH STOPS instances and negative ones can be increased by scaling up the construction sufficiently; scale until the gap is at least  $(nG)^2$ . This concludes Construction 2. Scale up again by  $S^2$  and replace each city with a  $G \times G$  square of cities (so containing  $(G + 1)^2$  cities) with the left top corner city in the original city location. This concludes Construction 3.

We define each square of size  $5G \times 5G$  centered on such a square of cities as a **region**, the square itself as the **inner region** and the region without the inner region as the (empty) **outer region**. (See also Fig. 12.) A **region visit** is a segment of trajectory intersecting the inner region; it is counted from the entry into the corresponding region until the trajectory leaves the region (a single region visit can leave and re-enter the inner region). A region visit is **fast** if a speed component (in some dimension) is always larger than zero during the entire region visit. Otherwise the region visit is called **slow**. (See again Fig. 12.) Note that a slow visit might never come to a complete stop, as speed components may turn to zero at different points in time.

**Lemma 8.** *Any slow region visit can be locally converted to a region visit that stops within the inner region, then visits all the cities, stopping at each one. The conversion does not change the trajectory outside the region visit and costs at most  $2G^2 + 15G$  additional vectors.*

**Proof.** A slow region visit means that all speed components must be zero at some point in time (although not necessarily at the same time). Take the first time it comes to a stop in some dimension. Keep this dimension to a stop (instead of whatever acceleration it had during the slow visit) until the other dimension comes to a stop (as it does in the original



(a) Fast visit: at least one speed component (the one along the  $x$  coordinate in this example) remains non-zero. (b) Slow visit: both speed components are zero at some time during the visit (although not at the same time in this example).

**Fig. 12.** Illustrative examples of a region and visits of a region. The larger gray dashed square indicates the outer region, while the other square shows the inner region containing the cities. Note that due to scaling, the outer region is much larger in actuality than depicted. The alternating red and blue lines are successive segments of trajectories. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

slow visit). Now the vehicle is at a stop somewhere in the region. Using at most  $5G$  vectors, the vehicle can travel to the inner region and come at a stop there at a corner city. Visiting the cities then requires at most  $2G^2$  vectors. Finally, we describe how to make the vehicle join back up with the original slow region visit at the latest configuration where a speed component is zero, say *w.l.o.g.* configuration  $c = (x, y, dx, 0)$ . Move the vehicle to  $(x - \frac{dx(\lfloor dx \rfloor + 1)}{2}, y, 0, 0)$  using at most  $5G$  vectors since the position is in the region as in the original slow visit, the vehicle was able to accelerate from zero speed to  $dx$  speed at position  $(x, y)$  while staying inside the region. From that configuration it can reach  $c$  using at most  $5G$  vectors, linking it back up to the original slow visit.  $\square$

**Lemma 9.** A fast region visit can visit at most  $G$  cities.

**Proof.** *W.l.o.g.* suppose the speed component which is always larger than zero to be  $dx$ . The position coordinate  $x$  is thus monotonically increasing or monotonically decreasing during the fast region visit, meaning it can visit at most one city per coordinate on the  $x$  axis, of which there are at most  $G$  in a region.  $\square$

**Corollary 2.** If a trajectory visits all cities in a region, it must do so either with at least one slow region visit, or with at least  $G$  fast region visits.

**Lemma 10.** A segment of a trajectory containing five fast region visits to the same region of cost  $K$ , can be replaced by a slow region visit of cost at most  $K + 13G$ .

**Proof.** If a trajectory leaves the region and then returns, the velocity along at least one dimension has changed sign (passing through the zero in the process). Moreover, if the sign of the velocity along some other dimension has not changed, then this speed must be below  $G$ , as leaving and returning to the inner region takes more than one unit of time. Therefore between each two region visits for each dimension there is a moment with a relatively low speed along this dimension (not necessarily the same moment).

For each dimension we independently modify the trajectory as follows. After the low speed moment between the first two visits, try to get within the inner region with zero speed (both – along this coordinate) as fast as possible. This cannot take longer than until the second low-speed moment plus time  $2G$  in case there was no full stop in between. We do the same transformation in reversed time from the fourth low-speed moment towards the third one. Note that without changing the trajectory outside of the segment, these two modifications will replace at most  $6G$  vectors into the third visit. (This in principle can lead to an overlap.) Now we shift the second replaced segment and everything afterwards by  $6G$ , by simply adding null vectors in-between the first replaced segment and the second one.

Observe that thanks to this idle time in all dimensions, we must obtain a point in time where the modified trajectory is within the inner region with zero speed (in all dimensions), which is thus a slow region visit.  $\square$

Note that the new slow region visit may not visit other regions/cities potentially visited by the original segment.

**Lemma 11.** *Given a solution of cost  $K$  for VECTOR TSP in Construction 3, one can construct a solution of cost at most  $K + 4nG^2$  for VECTOR TSP WITH STOPS in Construction 3.*

**Proof.** Consider all the regions and all the region visits. Let us call a region processed whenever it has been modified so that the cities are visited with stops. All the slow region visits can be processed at the cost of  $3G^2$  per region. Afterwards, segments with five fast region visits to the same unprocessed region, no overlap to any already chosen segments, and the minimum number of visits to other unprocessed regions, are modified to a slow region visit at the cost of  $G^2$ , which in turn can again be processed at the cost of  $3G^2$  extra vectors. Note that if some region remains unprocessed, then we have processed fewer than  $n$  regions, and each processed region can spend no more than 4 visits to each of the other regions. Thus each unprocessed region has  $G - 4n$  usable fast region visits, split into at most  $n$  groups by the previously chosen slow region visits and segments of five fast region visits. Since  $\frac{G-4n}{n} > 5$ , we can always continue to process cities until no unprocessed city remains.  $\square$

We are now ready to prove the main result.

**Theorem 3.** *VECTOR TSP is NP-hard.*

**Proof.** We reduce from VECTOR TSP WITH STOPS, which we have proven NP-hard in Theorem 2. More precisely, we have proven it has an NP-hard gap, which in Construction 2 we have widened through a scaling argument to be at least  $(nG)^2$ . It suffices to show that Construction 3 has a gap corresponding to this one. Consider a solution for VECTOR TSP WITH STOPS in Construction 2 with cost  $K$ . It can be scaled to a solution for VECTOR TSP in Construction 3 with cost at most  $KS + n + 3nG^2$ . Consider a solution for VECTOR TSP in Construction 3 with cost at most  $KS + 4nG^2$ . We can modify it so it visits cities at a stop for a total cost of at most  $KS + 8nG^2$ , by forgetting the points except the original city in the corner. Then, scale it down to a solution of cost at most  $K + \frac{8nG^2}{S} + nG^2$ . This is a solution for VECTOR TSP WITH STOPS in Construction 2 and observe that  $\frac{8nG^2}{S} + nG^2 < (nG)^2$ . Thus we have constructed a gap for VECTOR TSP in Construction 3 to which the gap for VECTOR TSP WITH STOPS in Construction 2 reduces.  $\square$

## 5. Algorithmic aspects and experimental investigation

In this section, we investigate some algorithmic approaches for tackling VECTOR TSP. The proposed framework is based on an interaction between a high-level part that compute the tour (i.e. visit order), and a trajectory algorithm that evaluates the cost of  $\text{racetrack}(\pi)$ , i.e., the cost of an optimal trajectory realizing this tour. Our solution to the latter problem is perhaps the main contribution of this section, it relies on a multi-point adaptation of the A\* algorithm, which may be of independent interest. Experiments are conducted for quantifying how often a trajectory based on an optimal tour for EUCLIDEAN TSP tend to be *suboptimal* for VECTOR TSP. The main goal of these experiment is to legitimate the study of VECTOR TSP as an independent problem.

### 5.1. Computing the tours (FlipTour)

A classical heuristic in the TSP literature is the Flip algorithm [10] (also called 2-opt although it is not an approximation algorithm). In each step, every possible *flip*, i.e. inversion of a subtour, of the current tour  $\pi$  is evaluated. If the resulting tour  $\pi'$  improves upon  $\pi$ , then it is selected and the algorithm tries to improve  $\pi'$  in turn. Eventually, the algorithm finds a tour that is a local optimum with respect to the flip operation. The high level part of our algorithm, called FlipTour, implements a similar approach for exploring the tours for VECTOR TSP. The main difference between Flip and FlipTour is that the cost of a tour is not evaluated in terms of distance, but in terms of the cost of an optimal trajectory realizing it (obtained by calling the Multi Point A\* algorithm presented in Section 5.2).

Let  $\text{arbitrary\_tour}()$  be a function that returns a trivial initial tour of length at most  $O(nL)$ , starting and terminating at the desired city. Such a tour can be obtained easily by using the strategy described in the proof of Lemma 2. Let  $\text{racetrack}()$  be the function that returns the cost of an optimal trajectory for the given tour, and let  $\text{flip}()$  be the function that inverts the desired segment of the underlying tour. Then FlipTour is described by Algorithm 1.

**Theorem 4.** *The above algorithm always reaches a local optimum after calling the oracle at most  $O(n^3L)$  many times, where  $n$  is the number of cities and  $L$  is the maximum distance (in a dimension) between two cities.*

**Proof.** The initial tour corresponds to a trajectory of length at most  $O(nL)$  (Lemma 2), thus  $C_{opt}$  is initialized with a value of at most  $O(nL)$  in Line 2. Now, the main loop iterates only when a shorter trajectory is found, which can occur at most as many times as the length of the initial trajectory. Finally, in each iteration, at most  $O(n^2)$  modified tours are evaluated by the oracle.  $\square$



**Algorithm 1** : FlipTour.Input: a set  $P$  of cities and a starting city  $p_0 \in P$ .

Output: a tour whose optimal trajectory cannot be improved by inverting a subtour.

---

```

1:  $\pi_{opt} \leftarrow \text{arbitrary\_tour}(P, p_0)$ 
2:  $C_{opt} \leftarrow \text{oracle}(\pi_{opt})$ 
3:  $improved \leftarrow \text{true}$ 
4: while  $improved$  do
5:    $improved \leftarrow \text{false}$ 
6:   for each city  $i \in P \setminus \{p_0\}$  do
7:     for each city  $j \in P \setminus \{p_0, i\}$  do
8:        $\pi_{test} \leftarrow \text{flip}(\pi_{opt}, i, j)$ 
9:        $C_{test} \leftarrow \text{oracle}(\pi_{test})$ 
10:      if  $C_{test} < C_{opt}$  then
11:         $\pi_{opt} \leftarrow \pi_{test}$ 
12:         $C_{opt} \leftarrow C_{test}$ 
13:         $improved \leftarrow \text{true}$ 
14:      break
15:   if  $improved$  then
16:     break
17: return  $\pi_{opt}$ 

```

---

## 5.2. Optimal trajectory for a given tour (Multi Point A\*)

Here, we discuss the problem of computing an optimal trajectory that realizes a given tour, i.e., how to compute  $\text{racetrack}(\pi)$  for a given  $\pi$ . A previous work of interest is Bekos et al. [3], which addresses the problem of computing an optimal RACETRACK trajectory in the so-called “Indianapolis” track, where the track has a fixed width and right-angle turns. The Indianapolis setting makes it possible to decompose the computation by using dynamic programming between the (constant sized) corner areas of the track. In contrast, we consider an open space with no simple way to decompose the computation. Our strategy adapts the  $A^*$  algorithm to searching a multi-point optimal path in the configuration graph. Namely, given a sequence of points  $\pi = (p_1, p_2, \dots, p_n)$ , compute (the cost of) an optimal trajectory visiting the points in order, starting at  $p_1$  and ending at  $p_n$  at zero speed. (In the special case of VECTOR TSP,  $p_1$  and  $p_n$  coincide.)

Hence, the main difficulty is that the optimal trajectory for a tour in open space does not reduce to gluing together subtrajectories for given parts of the tour. Our contribution is to design an estimation function that guides  $A^*$  through these constraints and finds the actual optimum efficiently.

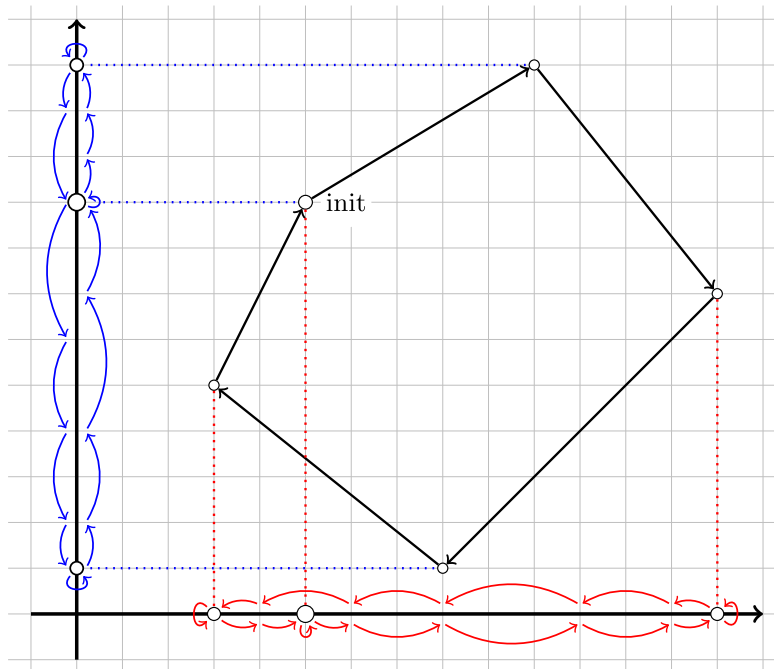
*Overview of the estimation function*

The general principle of  $A^*$  is to explore the search space by generating a set of successors of the current node (in our case, a configuration) whose remaining costs towards destination are *estimated* by a custom function. The most promising nodes are explored first (typically, based on a priority queue). Then,  $A^*$  is guaranteed to find the optimal solution (i.e., path in the search space) as long as the provided function does not overestimate the cost of a node. Furthermore, as is usual for  $A^*$ , the running time of the algorithm is primarily determined by the accuracy of the estimation.

Our estimation function relies on the following key idea. Given a tour  $\pi$  (in  $d$ -dimensions), a trajectory realizing this tour must have at least the cost of a trajectory realizing the projection of this tour separately in each dimension. Thus, one can use the maximum over these one-dimensional costs as an estimation for  $\pi$ . Consider for example the 2-dimensional tour  $\pi = ((5, 10), (10, 12), (14, 7), (8, 1), (3, 5), (5, 10))$  illustrated on Fig. 13. A vehicle realizing this tour must pass through the  $x$ -coordinates  $(5, 10, 14, 8, 3, 5)$  in this order. A similar statement holds for the  $y$ -coordinates. The sequence can be simplified further by retaining only the *inflexion points*, namely the coordinates where the vehicle begins, ends, or switches direction at zero speed (in the considered dimension). Here, the inflexion points are  $(5, 14, 3, 5)$ , their visit can be realized by the trajectory depicted in red along the  $x$ -axis. Again, a similar statement holds for the  $y$ -coordinates (in blue along the  $y$ -axis). Thus, the main work is to estimate the cost of such one-dimensional trajectories.

The fact that the speed is zero at the inflexion points makes it possible to decompose the computation as a sum of independent costs – one for each consecutive pair (whose calculation is described in Section 3.3). This does not imply that the  $d$ -dimensional trajectory computed by  $A^*$  will necessarily reach zero speed in the considered dimension, as other constraints may arise from combining the dimensions, which is fine because we are only lower bounding the costs here.

The main difficulty is that, as  $A^*$  progresses, what needs to be estimated is a *suffix* of the initial tour rather than the entire tour. Also, the current speed of the vehicle may be different from zero. Let  $(x, dx)$  be the projection of the current position and velocity of the vehicle in the considered dimension, and let  $x'$  be the position of the next inflexion point in the projection of the tour suffix. Either the vehicle is towards  $x'$  (that is,  $x' - x$  and  $dx$  have the same sign), or it is moving in the opposite direction. In the first case, the vehicle must stop first, then come back to  $x'$ . Stopping takes  $dx$  time units



**Fig. 13.** Lower bounding the cost of a tour by its one-dimensional projections. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and  $dx(dx + 1)/2$  space units, then the calculation reduces to computing the cost between two points with starting and finishing speed zero, as above. In the second case, either the vehicle is able to stop without bypassing  $x'$ , or it is not. If it is not, we replace  $x'$  in the projected suffix with the earliest position where the vehicle can stop, which is  $x + dx(dx + 1)/2$ , at a cost of  $dx$ . The last case, where the vehicle is going towards  $x'$  and is able to stop by  $x'$  is dealt with in Section 3.3, showing that this case can again be reduced to a case with zero initial speed.

In summary, the estimation of a suffix of the tour in a certain configuration reduces to computing, in each dimension separately (1) the cost from the current position and velocity to the next inflexion point (in the considered dimension), possibly changing the position of this point as explained above, and (2) the cost between all consecutive pairs of inflexion points in the projected suffix (in the considered dimension), each time starting and finishing at zero speed.

### 5.3. EUCLIDEAN TSP tours are typically not optimal

In this section, we present experimental evidence that optimal tours for EUCLIDEAN TSP are *typically* not optimal for VECTOR TSP. The experiments are conducted using the algorithmic components discussed in Sections 5.1 and 5.2. The instances are generated by placing cities uniformly at random within a two-dimensional square area. For each instance, we first use an external EUCLIDEAN TSP solver (namely, Concorde [1]) to obtain an optimal EUCLIDEAN TSP tour  $\pi$ , then we compute an optimal racetrack trajectory for this tour,  $\text{racetrack}(\pi)$ , using the Multi Point A\* algorithm presented above. Then, in order to see if a better solution than  $\text{racetrack}(\pi)$  exists for this instance, we take  $\pi$  as a starting point and try to improve this tour using the FlipTour heuristic (also presented above), with Multi Point A\* for computing the cost of the flipped tours.

Given the goal of the experiments, the only metric that we consider is how often the optimal tour for EUCLIDEAN TSP turns out to be suboptimal for VECTOR TSP. In other words, how often at least one flip leads to an improvement of the resulting trajectory. Two series of experiments were conducted in two-dimensional space. The results are averaged over 100 iterations. In the first experiment (Fig. 14, left), the space is limited to a  $100 \times 100$  subset of  $\mathbb{Z}^2$ , in which  $n$  cities are distributed uniformly at random, for  $n$  ranging over 4, 6, 8, 10, 12. In order to eliminate the potential effects of a restricted space, the second experiments (Fig. 14, right) adapts the same scenario on a growing space, so that the density of cities is constant. The results suggest clearly that EUCLIDEAN TSP tours tend to be suboptimal as the number of cities grows, as well as when the space is less restricted, with nearly 60% suboptimal tours for instances of 12 cities. Recall that the FlipTour algorithm is only a heuristic that stops when a local optimum is reached, whereas the reference EUCLIDEAN TSP tour is a global optimum. As a result, our experiments can *under-estimate* the discrepancy between VECTOR TSP and EUCLIDEAN TSP (which only makes our point stronger).

To conclude this section, Fig. 15 shows a larger instance of 20 cities and two corresponding trajectories; on the left side, an optimal trajectory realizing an optimal EUCLIDEAN TSP tour  $\pi$  returned by the Concorde solver (with a cost of




Fig. 14. Probability that an EUCLIDEAN TSP solution is improved (by a flip).




Fig. 15. Example of an instance where the optimal EUCLIDEAN TSP tour can be improved.

128); on the right side, an optimal trajectory realizing a different tour  $\pi'$  obtained after two flips of  $\pi$  (with a cost of 120).

## 6. Conclusion and future work

In this paper, we have presented and studied a new version of the TSP called VECTOR TSP, where the movements of the vehicle must obey initial constraints according to the Racetrack model. After proving a number of basic properties about this problem and about the Racetrack model, we showed that VECTOR TSP is NP-hard. While not really surprising, the proof of this result turned out to be highly non-trivial. More direct transformation from existing problems may exist, we leave this as an open question.

We used several of the preliminary observations for devising and motivating an algorithmic framework based on two interacting layers, namely a high level tour computation algorithm, and an exact guided pathfinding algorithm for trajectory evaluation. As simple as it is, we showed through experiments that this framework outperforms optimal solutions based on the EUCLIDEAN TSP, thereby motivating further study of the problem.

One question of interest is how far an optimal trajectory of an optimal EUCLIDEAN TSP tour could be from the optimal VECTOR TSP solution. Our work also leaves a number of algorithmic questions open. For example, the high-level part of the algorithm relies on a standard TSP heuristics. Would there be dedicated heuristics for this component in VECTOR TSP? On the other hand, our Multi Point A\* algorithm for trajectory computation could be of independent interest and it would be interesting to find other problems beyond TSP where such a component is useful. Indeed, the Racetrack model is appealing for its simplicity and amenability to algorithmic investigation. Several geometric, robotic, and mobility-related problems could certainly be revisited from a Racetrack perspective, generating new insights and potential for useful applications.

## Data availability

No data was used for the research described in the article.

## References

- [1] D. Applegate, R. Bixby, V. Chvatal, W. Cook, Concorde TSP solver, 2006, <https://www.math.uwaterloo.ca/tsp/concorde.html>.
- [2] S. Arora, Polynomial time approximation schemes for Euclidean TSP and other geometric problems, in: Proceedings of 37th Conference on Foundations of Computer Science, IEEE, 1996, pp. 2–11.

- [3] M.A. Bekos, T. Bruckdorfer, H. Förster, M. Kaufmann, S. Poschenrieder, T. Stüber, Algorithms and insights for racetrack, *Theoret. Comput. Sci.* 748 (2018) 2–16.
- [4] N. Biggs, E.K. Lloyd, R.J. Wilson, *Graph Theory, 1736–1936*, Oxford University Press, 1986.
- [5] A. Björklund, Determinant sums for undirected hamiltonicity, *SIAM J. Comput.* 43 (1) (2014) 280–299.
- [6] J. Canny, B. Donald, J. Reif, P. Xavier, On the complexity of kinodynamic planning, in: [Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science, 1988, pp. 306–316, <http://dx.doi.org/10.1109/SFCS.1988.21947>.
- [7] J. Canny, A. Rege, J. Reif, An exact algorithm for kinodynamic planning in the plane, *Discrete Comput. Geom.* 6 (3) (1991) 461–484.
- [8] P. Chalasani, R. Motwani, Approximating capacitated routing and delivery problems, *SIAM J. Comput.* 28 (6) (1999) 2133–2149.
- [9] N. Christofides, Worst-Case Analysis of a New Heuristic for the Travelling Salesman Problem, *Tech. Rep.*, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.
- [10] G.A. Croes, A method for solving traveling-salesman problems, *Oper. Res.* 6 (6) (1958) 791–812.
- [11] B. Donald, P. Xavier, J. Canny, J. Reif, Kinodynamic motion planning, *J. ACM* 40 (5) (1993) 1048–1066.
- [12] J. Erickson, Ernie's 3D pancakes : "How hard is optimal racing?", 2009, <http://3dpancakes.typepad.com/ernie/2009/06/how-hard-is-optimal-racing.html>.
- [13] R. Fontaine, J. Dibangoye, C. Solnon, Exact and Anytime Approach for Solving the Time Dependent Traveling Salesman Problem with Time Windows, *Tech. Rep.*, Université de Lyon; INSA Lyon; INRIA, 2022, <https://hal.science/hal-03865036>.
- [14] M. Gardner, Sim, chomp and race track—new games for intellect (and not for lady luck), *Sci. Am.* 228 (1) (1973) 108–115.
- [15] M.R. Garey, R.L. Graham, D.S. Johnson, Some NP-complete geometric problems, in: *Proceedings of the Eighth Annual ACM Symposium on Theory of Computing*, ACM, 1976, pp. 10–22.
- [16] L. Gasieniec, R. Klasing, C. Levcopoulos, A. Lingas, J. Min, T. Radzik, Bamboo garden trimming problem (perpetual maintenance of machines with different attendance urgency factors), in: *International Conference on Current Trends in Theory and Practice of Informatics*, Springer, 2017, pp. 229–240.
- [17] N. Hahn, M. Xefteris, The covering canadian traveller problem revisited, 2023, arXiv preprint arXiv:2304.14319.
- [18] M. Held, R.M. Karp, A dynamic programming approach to sequencing problems, *J. Soc. Ind. Appl. Math.* 10 (1) (1962) 196–210.
- [19] M. Holzer, P. McKenzie, The computational complexity of RACETRACK, in: *International Conference on Fun with Algorithms*, Springer, 2010, pp. 260–271.
- [20] A.R. Karlin, N. Klein, S.O. Gharan, A (slightly) improved approximation algorithm for metric TSP, in: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021, pp. 32–45.
- [21] R.M. Karp, Reducibility among combinatorial problems, in: *Complexity of Computer Computations*, Springer, 1972, pp. 85–103.
- [22] J. Kuszmaul, Bamboo trimming revisited: Simple algorithms can do well too, in: *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures*, 2022, pp. 411–417.
- [23] J. Le Ny, E. Frazzoli, E. Feron, The curvature-constrained traveling salesman problem for high point densities, in: *46th IEEE Conference on Decision and Control*, 2007, pp. 5985–5990.
- [24] C.S. Liao, Y. Huang, The covering Canadian traveller problem, *Theoret. Comput. Sci.* 530 (2014) 80–88.
- [25] G.M. Mallén-Fullerton, G. Fernandez-Anaya, DNA fragment assembly using optimization, in: *15th IEEE Congress on Evolutionary Computation*, 2013, pp. 1570–1577.
- [26] J.S. Mitchell, Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems, *SIAM J. Comput.* 28 (4) (1999) 1298–1309.
- [27] A. Narayanan, S. Zorbalas, DNA algorithms for computing shortest paths, *Proc. Genet. Program.* 718 (1998) 723.
- [28] R. Olsson, A. Tarandi, A genetic algorithm in the game racetrack, 2011.
- [29] P. Orponen, H. Mannila, On Approximation Preserving Reductions: complete Problems and Robust Measures (Revised Version), Department of Computer Science, University of Helsinki, 1990.
- [30] C.H. Papadimitriou, The euclidean travelling salesman problem is NP-complete, *Theoret. Comput. Sci.* 4 (3) (1977) 237–244.
- [31] C. Papadimitriou, S. Vempala, On the approximability of the traveling salesman problem, in: *Conference Proceedings of the Annual ACM Symposium on Theory of Computing*, vol. 26, 2006, pp. 101–120, <http://dx.doi.org/10.1007/s00493-006-0008-z>.
- [32] I.M. Ross, R.J. Proulx, M. Karpenko, An optimal control theory for the traveling salesman problem and its variants, 2020, <http://dx.doi.org/10.48550/ARXIV.2005.03186>, <https://arxiv.org/abs/2005.03186>.
- [33] K. Savla, F. Bullo, E. Frazzoli, Traveling salesperson problems for a double integrator, *IEEE Trans. Autom. Control* 54 (4) (2009) 788–793.
- [34] K. Savla, E. Frazzoli, F. Bullo, On the point-to-point and traveling salesperson problems for Dubins' vehicle, in: *Proceedings of the American Control Conference*, vol. 2, 2005, p. 786.
- [35] J. Schmid, VectorRace - finding the fastest path through a two-dimensional track, 2005, URL: <https://schmid.dk/articles/vectorRace.pdf>.
- [36] S. Sonawane, P. Patil, R. Bharsakade, P. Gaigole, Optimizing tool path sequence of plasma cutting machine using TSP approach, in: *E3S Web of Conferences*, vol. 184, EDP Sciences, 2020, p. 01037.
- [37] S. Yakowitz, B. Rutherford, Computational aspects of discrete-time optimal control, *Appl. Math. Comput.* 15 (1) (1984) 29–45.
- [38] Y. Yuan, Y. Peng, RaceTrack: An approximation algorithm for the mobile sink routing problem, in: I. Nikolaidis, K. Wu (Eds.), *Ad-Hoc, Mobile and Wireless Networks*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 135–148.