

Worst Case and Probabilistic Analysis of the 2-Opt Algorithm for the TSP*

Matthias Englert Heiko Röglin Berthold Vöcking

Department of Computer Science
RWTH Aachen, Germany
{englert,roeglin,voecking}@cs.rwth-aachen.de
January 25, 2008

Abstract

2-Opt is probably the most basic local search heuristic for the TSP. This heuristic achieves amazingly good results on “real world” Euclidean instances both with respect to running time and approximation ratio. There are numerous experimental studies on the performance of 2-Opt. However, the theoretical knowledge about this heuristic is still very limited. Not even its worst case running time on Euclidean instances was known so far. We clarify this issue by presenting, for every $p \in \mathbb{N}$, a family of L_p instances on which 2-Opt can take an exponential number of steps.

Previous probabilistic analyses were restricted to instances in which n points are placed uniformly at random in the unit square $[0, 1]^2$, where it was shown that the expected number of steps is bounded by $\tilde{O}(n^{10})$ for Euclidean instances. We consider a more advanced model of probabilistic instances in which the points can be placed according to general distributions on $[0, 1]^d$, for an arbitrary $d \geq 2$. In particular, we allow different distributions for different points. We study the expected number of local improvements in terms of the number n of points and the maximal density ϕ of the probability distributions. We show an upper bound on the expected length of any 2-Opt improvement path of $\tilde{O}(n^{4+1/3} \cdot \phi^{8/3})$. When starting with an initial tour computed by an insertion heuristic, the upper bound on the expected number of steps improves even to $\tilde{O}(n^{4+1/3-1/d} \cdot \phi^{8/3})$. If the distances are measured according to the Manhattan metric, then the expected number of steps is bounded by $\tilde{O}(n^{4-1/d} \cdot \phi)$. In addition, we prove an upper bound of $O(\sqrt[d]{\phi})$ on the expected approximation factor w. r. t. all L_p metrics.

Let us remark that our probabilistic analysis covers as special cases the uniform input model with $\phi = 1$ and a smoothed analysis with Gaussian perturbations of standard deviation σ with $\phi \sim 1/\sigma^d$. Besides random metric instances, we also consider an alternative random input model in which an adversary specifies a graph and distributions for the edge lengths in this graph. In this model, we achieve even better results on the expected number of local improvements of 2-Opt.

*This work was supported in part by the EU within the 6th Framework Programme under contract 001907 (DELIS) and by DFG grants VO 889/2 and WE 2842/1. An extended abstract appeared in Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA).

1 Introduction

In the *traveling salesperson problem (TSP)*, we are given a set of *vertices* and for each pair of distinct vertices a distance. The goal is to find a *tour* of minimum length that visits every vertex exactly once and returns to the initial vertex at the end. Despite many theoretical analyses and experimental evaluations of the TSP, there is still a considerable gap between the theoretical results and the experimental observations. One important special case is the *Euclidean TSP* in which the vertices are points in \mathbb{R}^d , for some $d \in \mathbb{N}$, and the distances are measured according to the Euclidean metric. This special case is known to be NP-hard in the strong sense [Pap77], but it admits a polynomial time approximation scheme (PTAS), shown independently in 1996 by Arora [Aro98] and Mitchell [Mit99]. These approximation schemes are based on dynamic programming. However, the most successful algorithms on practical instances rely on the principle of local search and very little is known about their complexity.

The *2-Opt* algorithm is probably the most basic local search heuristic for the TSP. 2-Opt starts with an arbitrary initial tour and incrementally improves this tour by making successive improvements that exchange two of the edges in the tour with two other edges. More precisely, in each *improving step* the 2-Opt algorithm selects two edges $\{u_1, u_2\}$ and $\{v_1, v_2\}$ from the tour such that u_1, u_2, v_1, v_2 are distinct and appear in this order in the tour, and it replaces these edges by the edges $\{u_1, v_1\}$ and $\{u_2, v_2\}$, provided that this change decreases the length of the tour. The algorithm terminates in a local optimum in which no further improving step is possible. We use the term *2-change* to denote a local improvement made by 2-Opt. This simple heuristic performs amazingly well on “real-life” Euclidean instances like, e. g., the ones in the well-known TSPLIB [Rei91]. Usually the 2-Opt heuristic needs a clearly subquadratic number of improving steps until it reaches a local optimum and the computed solution is within a few percentage points of the global optimum [JM97].

There are numerous experimental studies on the performance of 2-Opt. However, the theoretical knowledge about this heuristic is still very limited. Let us first discuss the number of local improvement steps made by 2-Opt before it finds a locally optimal solution. When talking about the number of local improvements, it is convenient to consider the *state graph*. The vertices in this graph correspond to the possible tours and an arc from a vertex v to a vertex u is contained if u is obtained from v by performing an improving 2-Opt step. On the positive side, van Leeuwen and Schoone consider a 2-Opt variant for the Euclidean plane in which only steps are allowed that remove a crossing from the tour. Such steps can introduce new crossings, but van Leeuwen and Schoone [vLS81] show that after $O(n^3)$ steps, 2-Opt has found a tour without any crossing. On the negative side, Lueker [Lue75] constructs TSP instances whose state graphs contain exponentially long paths. Hence, 2-Opt can take an exponential number of steps before it finds a locally optimal solution. This result is generalized to k -Opt, for arbitrary $k \geq 2$, by Chandra, Karloff, and Tovey [CKT99]. These negative results, however, use arbitrary graphs whose edge lengths do not satisfy the triangle inequality. Hence, they leave open the question about the worst case complexity of 2-Opt on metric TSP instances. In particular, Chandra, Karloff, and Tovey ask whether it is possible to construct Euclidean TSP instances on which 2-Opt can take an exponential number of

steps. We resolve this question by constructing such instances in the Euclidean plane. In chip design applications, often TSP instances arise in which the distances are measured according to the Manhattan metric. Also for this metric and for every other L_p metric, we construct instances with exponentially long paths in the 2-Opt state graph.

Theorem 1.1. *For every $p \in \mathbb{N} \cup \{\infty\}$ and $n \in \mathbb{N}$, there is a two-dimensional TSP instance with $16n$ vertices in which the distances are measured according to the L_p metric and whose state graph contains a path of length $2^{n+4} - 22$.*

For Euclidean instances in which n points are placed uniformly at random in the unit square, Kern [Ker89] shows that the length of the longest path in the state graph is bounded by $O(n^{16})$ with probability $1 - c/n$ for some constant c . Chandra, Karloff, and Tovey [CKT99] improve this result by bounding the expected length of the longest path in the state graph by $O(n^{10} \log n)$. That is, independent of the initial tour and the choice of the local improvements, the expected number of 2-changes is bounded by $O(n^{10} \log n)$. For instances in which n points are placed uniformly at random in the unit square and the distances are measured according to the Manhattan metric, Chandra, Karloff, and Tovey show that the expected length of the longest path in the state graph is bounded by $O(n^6 \log n)$.

We consider a more general probabilistic input model and improve the previously known bounds. The probabilistic model underlying our analysis allows that different vertices are placed according to different continuous probability distributions in the unit hypercube $[0, 1]^d$, for some constant *dimension* $d \geq 2$. The distribution of a vertex v_i is defined by a density function $f_i: [0, 1]^d \rightarrow [0, \phi]$ for some given $\phi \geq 1$. Our upper bounds depend on the number n of vertices and the upper bound ϕ on the density. We denote instances created by this input model as ϕ -perturbed Euclidean or Manhattan instances, depending on the underlying metric. The parameter ϕ can be seen as a parameter specifying how close the analysis is to a worst case analysis since the larger ϕ is, the better can worst case instances be approximated by the distributions. For $\phi = 1$ and $d = 2$, every point has a uniform distribution over the unit square, and hence the input model equals the uniform model analyzed before. Our results narrow the gap between the subquadratic number of improving steps observed in experiments [JM97] and the upper bounds from the probabilistic analysis. With slight modifications, this model also covers a smoothed analysis, in which first an adversary specifies the positions of the points and after that each position is slightly perturbed by adding a Gaussian random variable with small standard deviation σ . In this case, one has to set $\phi = 1/(\sqrt{2\pi}\sigma)^d$.

We also consider a model in which an arbitrary graph $G = (V, E)$ is given and for each edge $e \in E$, a probability distribution according to which the edge length $d(e)$ is chosen independently of the other edge lengths. Again, we restrict the choice of distributions to distributions that can be represented by density functions $f_e: [0, 1] \rightarrow [0, \phi]$ with maximal density at most ϕ for a given $\phi \geq 1$. We denote inputs created by this input model as ϕ -perturbed graphs. Observe that in this input model only the distances are perturbed whereas the graph structure is not changed by the randomization. This can be useful if one wants to explicitly prohibit certain edges. However, if the graph G is not complete, one has to initialize 2-Opt with a Hamiltonian cycle to start with.

We prove the following theorem about the expected length of the longest path in

the 2-Opt state graph for the three probabilistic input models discussed above. It is assumed that the dimension $d \geq 2$ is an arbitrary constant.

Theorem 1.2. *The expected length of the longest path in the 2-Opt state graph*

- a) *is $O(n^4 \cdot \phi)$ for ϕ -perturbed Manhattan instances with n points.*
- b) *is $O(n^{4+1/3} \cdot \log(n\phi) \cdot \phi^{8/3})$ for ϕ -perturbed Euclidean instances with n points.*
- c) *is $O(m \cdot n^{1+o(1)} \cdot \phi)$ for ϕ -perturbed graphs with n vertices and m edges.*

Usually, 2-Opt is initialized with a tour computed by some tour construction heuristic. One particular class are *insertion heuristics*, which insert the vertices one after another into the tour. We show that also from a theoretical point of view, using such an insertion heuristic yields a significant improvement for metric TSP instances because the initial tour 2-Opt starts with is much shorter than the longest possible tour. In the following theorem, we summarize our results on the expected number of local improvements.

Theorem 1.3. *The expected number of steps performed by 2-Opt*

- a) *is $O(n^{4-1/d} \cdot \log n \cdot \phi)$ on ϕ -perturbed Manhattan instances with n points when 2-Opt is initialized with a tour obtained by an arbitrary insertion heuristic.*
- b) *is $O(n^{4+1/3-1/d} \cdot \log^2(n\phi) \cdot \phi^{8/3})$ on ϕ -perturbed Euclidean instances with n points when 2-Opt is initialized with a tour obtained by an arbitrary insertion heuristic.*

In fact, our analysis shows not only that the expected number of local improvements is polynomially bounded but it also shows that the second moment and hence the variance is bounded polynomially for ϕ -perturbed Manhattan and graph instances. For the Euclidean metric, we cannot bound the variance but the 3/2-th moment polynomially.

Similar to the running time, the good approximation ratios obtained by 2-Opt on practical instances cannot be explained by a worst-case analysis. In fact, there are quite negative results on the worst-case behavior of 2-Opt. For example, Chandra, Karloff, and Tovey [CKT99] show that there are Euclidean instances in the plane for which 2-Opt has local optima whose costs are $\Omega\left(\frac{\log n}{\log \log n}\right)$ times larger than the optimal costs. However, the same authors also show that the expected approximation ratio of the worst local optimum for instances with n points drawn uniformly at random from the unit square is bounded from above by a constant. We generalize their result to our input model in which different points can have different distributions with bounded density ϕ and to all L_p metrics.

Theorem 1.4. *Let $p \in \mathbb{N} \cup \{\infty\}$. For ϕ -perturbed L_p instances, the expected approximation ratio of the worst tour that is locally optimal for 2-Opt is bounded by $O(\sqrt[d]{\phi})$.*

The remainder of the paper is organized as follows. We start by stating some basic definitions and notations in Section 2. In Section 3, we present the lower bounds. In Section 4, we analyze the expected number of local improvements and prove Theorems 1.2 and 1.3. Finally, in Sections 5 and 6, we prove Theorem 1.4 about the expected approximation factor and we discuss the relation between our analysis and a *smoothed analysis*.

2 Preliminaries

An *instance of the TSP* consists of a set $V = \{v_1, \dots, v_n\}$ of *vertices* (depending on the context, synonymously referred to as *points*) and a symmetric *distance function* $d: V \times V \rightarrow \mathbb{R}_{\geq 0}$ that associates with each pair $\{v_i, v_j\}$ of distinct vertices a distance $d(v_i, v_j) = d(v_j, v_i)$. The goal is to find a Hamiltonian cycle of minimum length. We also use the term *tour* to denote a Hamiltonian cycle. For a natural number $n \in \mathbb{N}$, we denote the set $\{1, \dots, n\}$ by $[n]$.

A pair (V, d) of a nonempty set V and a function $d: V \times V \rightarrow \mathbb{R}_{\geq 0}$ is called a *metric space* if for all $x, y, z \in V$ the following properties are satisfied:

- (a) $d(x, y) = 0$ if and only if $x = y$ (*reflexivity*),
- (b) $d(x, y) = d(y, x)$ (*symmetry*),
- (c) $d(x, z) \leq d(x, y) + d(y, z)$ (*triangle inequality*).

If (V, d) is a metric space, then d is called a *metric on V* . A TSP instance with vertices V and distance function d is called *metric TSP instance* if (V, d) is a metric space.

A well-known class of metrics on \mathbb{R}^d is the class of L_p *metrics*. For $p \in \mathbb{N}$, the distance $d_p(x, y)$ of two points $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^d$ with respect to the L_p metric is given by $d_p(x, y) = \sqrt[p]{|x_1 - y_1|^p + \dots + |x_d - y_d|^p}$. The L_1 metric is often called *Manhattan metric*, and the L_2 metric is well-known as *Euclidean metric*. For $p \rightarrow \infty$, the L_p metric converges to the L_∞ metric defined by the distance function $d_\infty(x, y) = \max\{|x_1 - y_1|, \dots, |x_d - y_d|\}$. A TSP instance (V, d) with $V \subseteq \mathbb{R}^d$ in which d equals d_p restricted to V is called an L_p *instance*. We also use the terms *Manhattan instance* and *Euclidean instance* to denote L_1 and L_2 instances, respectively. Furthermore, if p is clear from context, we write d instead of d_p .

A *tour construction heuristic* for the TSP incrementally constructs a tour and stops as soon as a valid tour is created. Usually, a tour constructed by such a heuristic is used as the initial solution 2-Opt starts with. A well-known class of tour construction heuristics for metric TSP instances are so-called *insertion heuristics*. These heuristics insert the vertices into the tour one after another, and every vertex is inserted between two consecutive vertices in the current tour where it fits best. To make this more precise, let T_i denote a subtour on a subset S_i of i vertices, and suppose $v \notin S_i$ is the next vertex to be inserted. If (x, y) denotes an edge in T_i that minimizes $d(x, v) + d(v, y) - d(x, y)$, then the new tour T_{i+1} is obtained from T_i by deleting the edge (x, y) and adding the edges (x, v) and (v, y) . Depending on the order in which the vertices are inserted into the tour, one distinguishes between several different insertion heuristics. Rosenkrantz et al. [RSI77] show an upper bound of $\lceil \log n \rceil + 1$ on the approximation factor of any insertion heuristic on metric TSP instances. Furthermore, they show that two variants which they call *nearest insertion* and *cheapest insertion* achieve an approximation ratio of 2 for metric TSP instances. The nearest insertion heuristic always inserts the vertex with the smallest distance to the current tour, and the cheapest insertion heuristic always inserts the vertex whose insertion leads to the cheapest tour T_{i+1} .

3 Exponential Lower Bounds

In this section, we answer Chandra, Karloff, and Tovey’s question [CKT99] whether it is possible to construct TSP instances in the Euclidean plane on which 2-Opt can take an exponential number of steps. We present, for every $p \in \mathbb{N} \cup \{\infty\}$, a family of two-dimensional L_p instances with exponentially long sequences of improving 2-changes. In Section 3.1, we present our construction for the Euclidean plane, and in Section 3.2 we extend this construction to general L_p metrics.

3.1 Exponential Lower Bound for the Euclidean Plane

In Lueker’s construction [Lue75] many of the 2-changes remove two edges that are far apart in the current tour in the sense that many vertices are visited between them, no matter of how the direction of the tour is chosen. Our construction differs significantly from the previous one as the 2-changes in our construction affect the tour only locally. The instances we construct are composed of gadgets of constant size. Each of these gadgets has a *zero state* and a *one state*, and there exists a sequence of improving 2-changes starting in the zero state and eventually leading to the one state. Let G_0, \dots, G_{n-1} denote these gadgets. If gadget G_i with $i > 0$ has reached state one, then it can be reset to its zero state by gadget G_{i-1} . The crucial property of our construction is that whenever a gadget G_{i-1} changes its state from zero to one, it resets gadget G_i twice. Hence, if in the initial tour, gadget G_0 is in its zero state and every other gadget is in state one, then for every i with $0 \leq i \leq n-1$, gadget G_i performs 2^i state changes from zero to one as, for $i > 0$, gadget G_i is reset 2^i times.

Every gadget is composed of 2 subgadgets, which we refer to as *blocks*. Each of these blocks consists of 4 vertices that are consecutively visited in the tour. For $i \in \{0, \dots, n-1\}$ and $j \in [2]$, let \mathcal{B}_1^i and \mathcal{B}_2^i denote the blocks of gadget G_i and let A_j^i, B_j^i, C_j^i , and D_j^i denote the four points \mathcal{B}_j^i consist of. If one ignores certain intermediate configurations that arise when one gadget resets another one, our construction ensures the following properties: The points A_j^i, B_j^i, C_j^i , and D_j^i are always consecutive in the tour, the edge between B_j^i and C_j^i is contained in every tour, and B_j^i and C_j^i are always the inner points of the block. That is, if one excludes the intermediate configurations, only the configurations $A_j^i B_j^i C_j^i D_j^i$ and $A_j^i C_j^i B_j^i D_j^i$ occur during the sequence of 2-changes. Observe that the change from one of these configurations to the other corresponds to a single 2-change in which the edges $A_j^i B_j^i$ and $C_j^i D_j^i$ are replaced by the edges $A_j^i C_j^i$ and $B_j^i D_j^i$, or vice versa. In the following, we assume that the sum $d(A_j^i, B_j^i) + d(C_j^i, D_j^i)$ is strictly smaller than the sum $d(A_j^i, C_j^i) + d(B_j^i, D_j^i)$, and we refer to the configuration $A_j^i B_j^i C_j^i D_j^i$ as the *short state* of the block and to the configuration $A_j^i C_j^i B_j^i D_j^i$ as the *long state*. Another property of our construction is that neither the order in which the blocks are visited nor the order of the gadgets is changed during the sequence of 2-changes. Again with the exception of the intermediate configurations, the order in which the blocks are visited is $\mathcal{B}_1^0 \mathcal{B}_2^0 \mathcal{B}_1^1 \mathcal{B}_2^1 \dots \mathcal{B}_1^{n-1} \mathcal{B}_2^{n-1}$ (see Figure 3.1).

Due to the aforementioned properties, we can describe every non-intermediate tour that occurs during the sequence of 2-changes completely by specifying for every block if it is in its short state or in its long state. In the following, we denote the state of

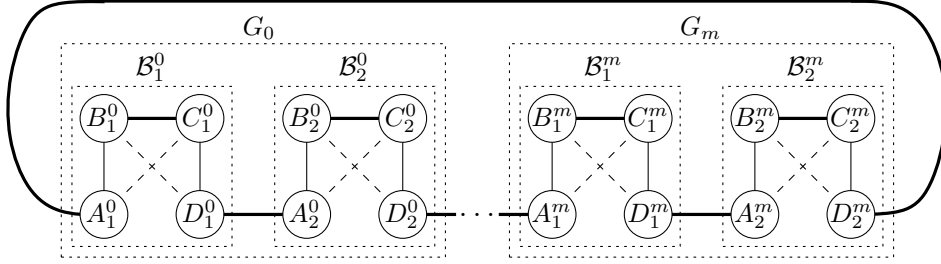


Figure 3.1: In the illustration, we use m to denote $n - 1$. Every tour that occurs in the sequence of 2-changes contains the thick edges. For each block, either both solid or both dashed edges are contained. In the former case the block is in its short state; in the latter case the block is in its long state.

a gadget G_i by a pair (x_1, x_2) with $x_j \in \{S, L\}$, meaning that block \mathcal{B}_j^i is in its short state if and only if $x_j = S$. Since every gadget consists of two blocks, there are four possible states for each gadget. However, only three of them appear in the sequence of 2-changes, namely (L, L) , (S, L) , and (S, S) . We call state (L, L) the *zero state* and state (S, S) the *one state*. In order to guarantee the existence of an exponentially long sequence of 2-changes, the gadgets we construct possess the following property.

Property 3.1. *If, for $i \in \{0, \dots, n - 2\}$, gadget G_i is in state (L, L) or (S, L) and gadget G_{i+1} is in state (S, S) , then there exists a sequence of seven consecutive 2-changes terminating with gadget G_i being in state (S, L) or (S, S) , respectively, and gadget G_{i+1} in state (L, L) . In this sequence only edges of and between the gadgets G_i and G_{i+1} are involved.*

If this property is satisfied and if in the initial tour gadget G_0 is in its zero state (L, L) and every other gadget is in its one state (S, S) , then there exists an exponentially long sequence of 2-changes in which gadget G_i changes 2^i times from state zero to state one, as the following lemma shows.

Lemma 3.2. *If, for $i \in \{0, \dots, n - 2\}$, gadget G_i is in the zero state (L, L) and all gadgets G_j with $j > i$ are in the one state (S, S) , then there exists a sequence of $2^{n+3-i} - 14$ consecutive 2-changes in which only edges of and between the gadgets G_j with $j \geq i$ are involved and that terminates in a state in which all gadgets G_j with $j \geq i$ are in the one state.*

Proof. We prove the lemma by induction on i . If gadget G_{n-1} is in state (L, L) , then it can change its state with two 2-changes to (S, S) without affecting the other gadgets. Hence, the lemma is true for $i = n - 1$. Now assume that the lemma is true for $i + 1$ and consider a state in which gadget G_i is in state (L, L) and all gadgets G_j with $j > i$ are in state (S, S) . Due to Property 3.1, there exists a sequence of seven consecutive 2-changes in which only edges of and between G_i and G_{i+1} are involved terminating with G_i being in state (S, L) and G_{i+1} being in state (L, L) . By the induction hypothesis there exists a sequence of $2^{n+2-i} - 14$ 2-changes after which all gadgets G_j with $j > i$ are in state (S, S) . Then, due to Property 3.1, there exists a sequence of seven consecutive 2-changes

in which only G_i changes its state from (S, L) to (S, S) while resetting gadget G_{i+1} again from (S, S) to (L, L) . Hence, we can apply the induction hypothesis again, yielding that after another $2^{n+2-i} - 14$ 2-changes all gadgets G_j with $j \geq i$ are in state (S, S) . This concludes the proof as the number of 2-changes performed is $14 + 2(2^{n+2-i} - 14) = 2^{n+3-i} - 14$. \square

In particular, this implies that, given Property 3.1, one can construct instances consisting of $2n$ gadgets, i. e., $16n$ points, whose state graphs contain paths of length $2^{2n+3} - 14 > 2^{2n+4} - 22$, as desired in Theorem 1.1.

3.1.1 Detailed description of the sequence of steps

Now we describe in detail how a sequence of 2-changes satisfying Property 3.1 can be constructed. First, we assume that gadget G_i is in state (S, L) and that gadget G_{i+1} is in state (S, S) . Under this assumption, there are three consecutive blocks, namely \mathcal{B}_2^i , \mathcal{B}_1^{i+1} , and \mathcal{B}_2^{i+1} , such that the leftmost one is in its long state, and the other blocks are in their short states. We need to find a sequence of 2-changes in which only edges of and between these three blocks are involved and after which the first block is in its short state and the other blocks are in their long states. Remember that when the edges $\{u_1, u_2\}$ and $\{v_1, v_2\}$ are removed from the tour and the vertices appear in the order u_1, u_2, v_1, v_2 in the current tour, then the edges $\{u_1, v_1\}$ and $\{u_2, v_2\}$ are added to the tour and the subtour between u_1 and v_2 is visited in reverse order. If, e. g., the current tour corresponds to the permutation $(1, 2, 3, 4, 5, 6, 7)$ and the edges $\{1, 2\}$ and $\{5, 6\}$ are removed, then the new tour is $(1, 5, 4, 3, 2, 6, 7)$. The following sequence of 2-changes has the desired properties. Brackets indicate the edges that are removed from the tour.

$$\begin{array}{l}
1) \quad [A_2^i \ C_2^i] \ B_2^i \ D_2^i \ \left| \ A_1^{i+1} \ B_1^{i+1} \ C_1^{i+1} \ D_1^{i+1} \ \left| \ A_2^{i+1} \ B_2^{i+1} \ [C_2^{i+1} \ D_2^{i+1}] \right. \\
2) \quad A_2^i \ C_2^{i+1} \ [B_2^{i+1} \ A_2^{i+1}] \ \left| \ D_1^{i+1} \ C_1^{i+1} \ B_1^{i+1} \ A_1^{i+1} \ \left| \ [D_2^i \ B_2^i] \ C_2^i \ D_2^{i+1} \right. \\
3) \quad A_2^i \ C_2^{i+1} \ [B_2^{i+1} \ D_2^i] \ \left| \ A_1^{i+1} \ B_1^{i+1} \ [C_1^{i+1} \ D_1^{i+1}] \ \left| \ A_2^{i+1} \ B_2^i \ C_2^i \ D_2^{i+1} \right. \\
4) \quad A_2^i \ C_2^{i+1} \ B_2^{i+1} \ C_1^{i+1} \ \left| \ [B_1^{i+1} \ A_1^{i+1}] \ D_2^i \ D_1^{i+1} \ \left| \ A_2^{i+1} \ B_2^i \ [C_2^i \ D_2^{i+1}] \right. \\
5) \quad [A_2^i \ C_2^{i+1}] \ B_2^{i+1} \ C_1^{i+1} \ \left| \ B_1^{i+1} \ C_2^i \ [B_2^i \ A_2^{i+1}] \ \left| \ D_1^{i+1} \ D_2^i \ A_1^{i+1} \ D_2^{i+1} \right. \\
6) \quad A_2^i \ B_2^i \ C_2^i \ B_1^{i+1} \ \left| \ [C_1^{i+1} \ B_2^{i+1}] \ C_2^{i+1} \ A_2^{i+1} \ \left| \ D_1^{i+1} \ D_2^i \ [A_1^{i+1} \ D_2^{i+1}] \right. \\
7) \quad A_2^i \ B_2^i \ [C_2^i \ B_1^{i+1}] \ \left| \ C_1^{i+1} \ A_1^{i+1} \ [D_2^i \ D_1^{i+1}] \ \left| \ A_2^{i+1} \ C_2^{i+1} \ B_2^{i+1} \ D_2^{i+1} \right. \\
\quad \quad A_2^i \ B_2^i \ C_2^i \ D_2^i \ \left| \ A_1^{i+1} \ C_1^{i+1} \ B_1^{i+1} \ D_1^{i+1} \ \left| \ A_2^{i+1} \ C_2^{i+1} \ B_2^{i+1} \ D_2^{i+1} \right.
\end{array}$$

If gadget G_i is in state (L, L) instead of state (S, L) , a sequence of steps that satisfies Property 3.1 can be constructed analogously. Additionally, one has to take into account that the three involved blocks \mathcal{B}_1^i , \mathcal{B}_1^{i+1} , and \mathcal{B}_2^{i+1} are not consecutive in the tour but that block \mathcal{B}_2^i lies between them. However, one can easily verify that this block is not affected by the sequence of 2-changes, as after the seven 2-changes have been performed, the block is in the same state and at the same position as before.

3.1.2 Embedding the construction into the Euclidean plane

The only missing step in the proof of Theorem 1.1 for the Euclidean plane is to find points such that all of the 2-changes that we described in the previous section are




Figure 3.2: This illustration shows the points of the gadgets G_{n-1} and G_{n-2} . One can see that G_{n-2} is a scaled, rotated, and translated copy of G_{n-1} .

improving. We specify the positions of the points of gadget G_{n-1} and give a rule how the points of gadget G_i can be derived when all points of gadget G_{i+1} have already been placed. In our construction it happens that different points have exactly the same coordinates. This is only for ease of notation; if one wants to obtain a TSP instance in which distinct points have distinct coordinates, one can slightly move these points without affecting the property that all 2-changes are improving. For $j \in [2]$, we choose $A_j^{n-1} = (0, 0)$, $B_j^{n-1} = (1, 0)$, $C_j^{n-1} = (-0.1, 1.4)$, and $D_j^{n-1} = (-1.1, 4.8)$. Then $A_j^{n-1}B_j^{n-1}C_j^{n-1}D_j^{n-1}$ is the short state because

$$d(A_j^{n-1}, C_j^{n-1}) + d(B_j^{n-1}, D_j^{n-1}) - (d(A_j^{n-1}, B_j^{n-1}) + d(C_j^{n-1}, D_j^{n-1})) > 2.09 .$$

We place the points of gadget G_i as follows (see Figure 3.2):

1. Start with the coordinates of the points of gadget G_{i+1} .
2. Rotate these points around the origin by $3\pi/2$.
3. Scale each coordinate with a factor of 3.
4. Translate the points by the vector $(-1.2, 0.1)$.

For $j \in [2]$, this yields $A_j^{n-2} = (-1.2, 0.1)$, $B_j^{n-2} = (-1.2, -2.9)$, $C_j^{n-2} = (3, 0.4)$, and $D_j^{n-2} = (13.2, 3.4)$.

From this construction it follows that each gadget is a scaled, rotated, and translated copy of gadget G_{n-1} . If one has a set of points in the Euclidean plane that admit certain improving 2-changes, then these 2-changes are still improving if one scales, rotates, and translates all points in the same manner. Hence, it suffices to show that the sequences in which gadget G_{n-2} resets gadget G_{n-1} from (S, S) to (L, L) are improving. There are two such sequences; in the first one, gadget G_{n-2} changes its state from (L, L) to (S, L) , in the second one, gadget G_{n-2} changes its state from (S, L) to (S, S) . Since the coordinates of the points in both blocks of gadget G_{n-2} are the same, the inequalities

for both sequences are also identical. The improvements made by the steps in both sequences are bounded from below by 0.03, 0.91, 0.06, 0.05, 0.43, 0.06, and 0.53. This concludes the proof of Theorem 1.1 for the Euclidean plane as it shows that all 2-changes in Lemma 3.2 are improving.

3.2 Exponential Lower Bound for L_p Metrics

We were not able to find a set of points in the plane such that all 2-changes in Lemma 3.2 are improving with respect to the Manhattan metric. Therefore, we modify the construction of the gadgets and the sequence of 2-changes. Our construction for the Manhattan metric is based on the construction for the Euclidean plane, but it does not possess the property that every gadget resets its neighboring gadget twice. This property is only true for half of the gadgets. To be more precise, we construct two different types of gadgets which we call *reset gadgets* and *propagation gadgets*. Reset gadgets perform the same sequence of 2-changes as the gadgets that we constructed for the Euclidean plane. Propagation gadgets also have the same structure as the gadgets for the Euclidean plane, but when such a gadget changes its state from (L, L) to (S, S) , it resets its neighboring gadget only once. Due to this relaxed requirement it is possible to find points in the Manhattan plane whose distances satisfy all necessary inequalities. Instead of n gadgets, our construction consists of $2n$ gadgets, namely n propagation gadgets G_0^P, \dots, G_{n-1}^P and n reset gadgets G_0^R, \dots, G_{n-1}^R . The order in which these gadgets appear in the tour is $G_0^P G_0^R G_1^P G_1^R \dots G_{n-1}^P G_{n-1}^R$.

As before, every gadget consists of two blocks and the order in which the blocks and the gadgets are visited does not change during the sequence of 2-changes. Consider a reset gadget G_i^R and its neighboring propagation gadget G_{i+1}^P . Then Property 3.1 is still satisfied. That is, if G_i^R is in state (L, L) or (S, L) and G_{i+1}^P is in state (S, S) , then there exists a sequence of seven consecutive 2-changes resetting gadget G_{i+1}^P to state (L, L) and leaving gadget G_i^R in state (S, L) or (S, S) , respectively. The situation is different for a propagation gadget G_i^P and its neighboring reset gadget G_i^R . In this case, if G_i^P is in state (L, L) , it first changes its state with a single 2-change to (S, L) . After that, gadget G_i^P changes its state to (S, S) while resetting gadget G_i^R from state (S, S) to state (L, L) by a sequence of seven consecutive 2-changes. In both cases, the sequences of 2-changes in which one block changes from its long to its short state while resetting two blocks of the neighboring gadget from their short to their long states are chosen analogously to the ones for the Euclidean plane described in Section 3.1.1.

In the initial tour, only gadget G_0^P is in state (L, L) and every other gadget is in state (S, S) . With similar arguments as for the Euclidean plane, we can show that gadget G_i^R is reset from its one state (S, S) to its zero state (L, L) 2^i times and that the total number of steps is $2^{n+4} - 22$.

3.2.1 Embedding the construction into the Manhattan plane

Similar to the construction in the Euclidean plane, the points in both blocks of a reset gadget G_i^R have the same coordinates. Also in this case one can slightly move all the points without affecting the inequalities if one wants distinct coordinates for

distinct points. Again, we choose points for the gadgets G_{n-1}^P and G_{n-1}^R and describe how the points of the gadgets G_i^P and G_i^R can be chosen when the points of the gadgets G_{i+1}^P and G_{i+1}^R are already chosen. For $j \in [2]$, we choose $A_{R,j}^{n-1} = (0, 1)$, $B_{R,j}^{n-1} = (0, 0)$, $C_{R,j}^{n-1} = (-0.7, 0.1)$, and $D_{R,j}^{n-1} = (-1.2, 0.08)$. Furthermore, we choose $A_{P,1}^{n-1} = (-2, 1.8)$, $B_{P,1}^{n-1} = (-3.3, 2.8)$, $C_{P,1}^{n-1} = (-1.3, 1.4)$, $D_{P,1}^{n-1} = (1.5, 0.9)$, $A_{P,2}^{n-1} = (-0.7, 1.6)$, $B_{P,2}^{n-1} = (-1.5, 1.2)$, $C_{P,2}^{n-1} = (1.9, -1.5)$, and $D_{P,2}^{n-1} = (-0.8, -1.1)$.

Before we describe how the points of the other gadgets are chosen, we first show that the 2-changes within and between the gadgets G_{n-1}^P and G_{n-1}^R are improving. For $j \in [2]$, $A_{R,j}^{n-1}B_{R,j}^{n-1}C_{R,j}^{n-1}D_{R,j}^{n-1}$ is the short state because

$$d(A_{R,j}^{n-1}, C_{R,j}^{n-1}) + d(B_{R,j}^{n-1}, D_{R,j}^{n-1}) - (d(A_{R,j}^{n-1}, B_{R,j}^{n-1}) + d(C_{R,j}^{n-1}, D_{R,j}^{n-1})) = 1.36 .$$

Also the 2-change in which G_{n-1}^P changes its state from (L, L) to (S, L) is improving because

$$d(A_{P,1}^{n-1}, C_{P,1}^{n-1}) + d(B_{P,1}^{n-1}, D_{P,1}^{n-1}) - (d(A_{P,1}^{n-1}, B_{P,1}^{n-1}) + d(C_{P,1}^{n-1}, D_{P,1}^{n-1})) = 2.2 .$$

The improvements made by the 2-changes in the sequence in which G_{n-1}^P changes its state from (S, L) to (S, S) while resetting G_{n-1}^R are 0.04, 0.4, 0.04, 0.16, 0.4, 0.04, and 0.6.

Again, our construction possesses the property that each pair of gadgets G_i^P and G_i^R is a scaled and translated version of the pair G_{n-1}^P and G_{n-1}^R . Since we have relaxed the requirements for the gadgets, we do not even need rotations here. We place the points of G_i^P and G_i^R as follows:

1. Start with the coordinates specified for the points of gadgets G_{i+1}^P and G_{i+1}^R .
2. Scale each coordinate with a factor of 7.7.
3. Translate the points by the vector $(1.93, 0.3)$.

For $j \in [2]$, this yields $A_{R,j}^{n-2} = (1.93, 8)$, $B_{R,j}^{n-2} = (1.93, 0.3)$, $C_{R,j}^{n-2} = (-3.46, 1.07)$, and $D_{R,j}^{n-2} = (-7.31, 0.916)$. Similar to our construction for the Euclidean plane, it suffices to show that the sequences in which gadget G_{n-2}^R resets gadget G_{n-1}^P from (S, S) to (L, L) are improving. As the coordinates of the points in the two blocks of gadget G_{n-2}^R are the same, the inequalities for both sequences are also identical. The improvements made by the steps in both sequences are 1.06, 1.032, 0.168, 1.14, 0.06, 0.4, and 0.012. This concludes the proof of Theorem 1.1 for the Manhattan metric as it shows that all 2-changes are improving.

Let us remark that this also implies Theorem 1.1 for the L_∞ metric because distances w. r. t. the L_∞ metric coincide with distances w. r. t. the Manhattan metric if one rotates all points by $\pi/4$ around the origin and scales every coordinate with $1/\sqrt{2}$.

3.2.2 Embedding the construction into general L_p metrics

It is also possible to embed our construction into the L_p metric for $p \geq 3$. For $j \in [2]$, we choose $A_{R,j}^{n-1} = (0, 1)$, $B_{R,j}^{n-1} = (0, 0)$, $C_{R,j}^{n-1} = (3.5, 3.7)$, and $D_{R,j}^{n-1} = (7.8, -3.2)$. Moreover, we choose $A_{P,1}^{n-1} = (-2.5, -2.4)$, $B_{P,1}^{n-1} = (-4.7, -7.3)$, $C_{P,1}^{n-1} = (-8.6, -4.6)$,

$D_{P,1}^{n-1} = (3.7, 9.8)$, $A_{P,2}^{n-1} = (3.2, 2)$, $B_{P,2}^{n-1} = (7.2, 7.2)$, $C_{P,2}^{n-1} = (-6.5, -1.6)$, and $D_{P,2}^{n-1} = (-1.5, -7.1)$. We place the points of G_i^P and G_i^R as follows:

1. Start with the coordinates specified for the points of gadgets G_{i+1}^P and G_{i+1}^R .
2. Rotate these points around the origin by π .
3. Scale each coordinate with a factor of 7.8.
4. Translate the points by the vector $(7.2, 5.3)$.

It can be calculated that the distances of these points when measured according to the L_p metric for any $p \geq 3$ satisfy all necessary inequalities.

4 Expected Number of 2-Changes

We analyze the expected number of 2-changes on random d -dimensional Manhattan and Euclidean instances, for an arbitrary constant dimension $d \geq 2$, and on general TSP instances. The previous results on the expected number of 2-changes due to Kern [Ker89] and Chandra, Karloff, and Tovey [CKT99] are based on the analysis of the improvement made by the smallest improving 2-change. If the smallest improvement is not too small, then the number of improvements cannot be large. In our analyses for the Manhattan and the Euclidean metric, we consider not only a single step but certain pairs of steps. We show that the smallest improvement made by any such pair is typically much larger than the improvement made by a single step, which yields our improved bounds. Our approach is not restricted to pairs of steps. One could also consider sequences of steps of length k for any small enough k . In fact, for general ϕ -perturbed graphs with m edges, we consider sequences of length $\sqrt{\log m}$. The reason why we can analyze longer sequences for general graphs is that these inputs possess more randomness than ϕ -perturbed Manhattan and Euclidean instances because every edge length is a random variable that is independent of the other edge lengths. Hence, the analysis for general ϕ -perturbed graphs demonstrates the limits of our approach under optimal conditions. For Manhattan and Euclidean instances, the gain of considering longer sequences is small due to the dependencies between the edge lengths.

4.1 Manhattan Instances

In this section, we analyze the expected number of 2-changes on ϕ -perturbed Manhattan instances. First we prove a weaker bound than the one in Theorem 1.2. The proof of this weaker bound illustrates our approach and reveals the problems one has to tackle in order to improve the upper bounds. It is solely based on an analysis of the smallest improvement made by any of the possible 2-Opt steps. If with high probability every 2-Opt step decreases the tour length by a polynomially large amount, then with high probability only polynomially many 2-Opt steps are possible before a local optimum is reached.

Theorem 4.1. *Starting with an arbitrary tour, the expected number of steps performed by 2-Opt on ϕ -perturbed Manhattan instances with n vertices is $O(n^6 \cdot \log n \cdot \phi)$.*

Proof. In order to prove the desired bound on the expected convergence time, we only need two simple observations. First, the initial tour can have length at most dn as the number of edges is n and every edge has length at most d . And second, every 2-Opt step decreases the length of the tour by a polynomially large amount with high probability. The latter can be shown by a union bound over all possible 2-Opt steps. Consider a fixed 2-Opt step S , let e_1 and e_2 denote the edges removed from the tour in step S , and let e_3 and e_4 denote the edges added to the tour. Then the improvement $\Delta(S)$ of step S can be written as

$$\Delta(S) = \mathbf{d}(e_1) + \mathbf{d}(e_2) - \mathbf{d}(e_3) - \mathbf{d}(e_4) . \quad (4.1)$$

Without loss of generality let $e_1 = (v_1, v_2)$ be the edge between the vertices v_1 and v_2 , and let $e_2 = (v_3, v_4)$, $e_3 = (v_1, v_3)$, and $e_4 = (v_2, v_4)$. Furthermore, for $i \in \{1, \dots, 4\}$, let $x^i \in \mathbb{R}^d$ denote the coordinates of vertex v_i . Then the improvement $\Delta(S)$ of step S can be written as

$$\Delta(S) = \sum_{i=1}^d (|x_i^1 - x_i^2| + |x_i^3 - x_i^4| - |x_i^1 - x_i^3| - |x_i^2 - x_i^4|) .$$

Depending on the order of the coordinates, $\Delta(S)$ can be written as linear combination of the coordinates. If, e. g., for all $i \in [d]$, $x_i^1 \geq x_i^2 \geq x_i^3 \geq x_i^4$, then the improvement $\Delta(S)$ can be written as $\sum_{i=1}^d (-2x_i^2 + 2x_i^3)$. There are $(4!)^d$ such orders and each one gives rise to a linear combination of the x_i^j 's with integer coefficients. For each of these linear combinations, the probability that it takes a value in the interval $(0, \varepsilon]$ is bounded from above by $\varepsilon\phi$, following, e. g., from Lemma A.1. Since $\Delta(S)$ can only take a value in the interval $(0, \varepsilon]$ if one of the linear combinations takes a value in this interval, the probability of the event $\Delta(S) \in (0, \varepsilon]$ can be upper bounded by $(4!)^d \varepsilon\phi$.

Let Δ_{\min} denote the improvement of the smallest improving 2-Opt step S , i. e., $\Delta_{\min} = \min\{\Delta(S) \mid \Delta(S) > 0\}$. We can estimate Δ_{\min} by a union bound, yielding

$$\Pr[\Delta_{\min} \leq \varepsilon] \leq (4!)^d \varepsilon n^4 \phi$$

as there are at most n^4 different 2-Opt steps. Let T denote the random variable describing the number of 2-Opt steps before a local optimum is reached. Observe that T can only exceed a given number t if the smallest improvement Δ_{\min} is less than dn/t , and hence

$$\Pr[T \geq t] \leq \Pr\left[\Delta_{\min} \leq \frac{dn}{t}\right] \leq \frac{d(4!)^d n^5 \phi}{t} .$$

Since there are at most $(n!)$ different TSP tours and none of these tours can appear twice during the local search, T is always bounded by $(n!)$. Altogether, we can bound the expected value of T by

$$\mathbf{E}[T] = \sum_{t=1}^{n!} \Pr[T \geq t] \leq \sum_{t=1}^{n!} \frac{d(4!)^d n^5 \phi}{t} .$$

Since we assumed the dimension d to be a constant, bounding the n -th harmonic number by $\ln(n) + 1$ and using $\ln(n!) = O(n \log n)$ yields

$$\mathbf{E}[T] = O(n^6 \cdot \log n \cdot \phi) . \quad \square$$

The bound in Theorem 4.1 is only based on the smallest improvement Δ_{\min} made by any of the 2-Opt steps. Intuitively, this is too pessimistic since most of the steps performed by 2-Opt yield a larger improvement than Δ_{\min} . In particular, two consecutive steps yield an improvement of at least Δ_{\min} plus the improvement Δ'_{\min} of the second smallest step. This observation alone, however, does not suffice to improve the bound substantially. Instead, we regroup the 2-changes to pairs such that each pair of 2-changes is *linked* by an edge, i. e., one edge added to the tour in the first 2-change is removed from the tour in the second 2-change, and we analyze the smallest improvement made by any pair of linked 2-Opt steps. Obviously, this improvement is at least $\Delta_{\min} + \Delta'_{\min}$ but one can hope that it is much larger because it is unlikely that the 2-change that yields the smallest improvement and the 2-change that yields the second smallest improvement form a pair of linked steps. We show that this is indeed the case and use this result to prove the bound on the expected length of the longest path in the state graph of 2-Opt on ϕ -perturbed Manhattan instances claimed in Theorem 1.2.

4.1.1 Construction of pairs of linked 2-changes

Consider an arbitrary sequence of consecutive 2-changes of length t . The following lemma guarantees that the number of disjoint linked pairs of 2-changes in every such sequence increases linearly with the length t .

Lemma 4.2. *In every sequence of t consecutive 2-changes, the number of disjoint pairs of 2-changes that are linked by an edge, i. e., pairs such that there exists an edge added to the tour in the first 2-change of the pair and removed from the tour in the second 2-change of the pair, is at least $t/3 - n(n-1)/4$.*

Proof. Let S_1, \dots, S_t denote an arbitrary sequence of consecutive 2-changes. The sequence is processed step by step and a list \mathcal{L} of disjoint linked pairs of 2-changes is created. Assume that the 2-changes S_1, \dots, S_{i-1} have already been processed and that now 2-change S_i has to be processed. Assume further that in step S_i the edges e_1 and e_2 are exchanged with the edges e_3 and e_4 . Let j denote the smallest index with $j > i$ such that edge e_3 is removed from the tour in step S_j if such a step exists, and let j' denote the smallest index with $j' > i$ such that edge e_4 is removed from the tour in step $S_{j'}$ if such a step exists. If the index j is defined, the pair (S_i, S_j) is added to the constructed list \mathcal{L} . If the index j is not defined but the index j' is defined, the pair $(S_i, S_{j'})$ is added to the constructed list \mathcal{L} . After that, both steps S_j and $S_{j'}$ (if defined) are removed from the sequence of 2-changes, that is, they are not processed in the following in order to guarantee disjointness of the pairs in \mathcal{L} .

If one 2-change is processed, it excludes at most two other 2-changes from being processed. Hence, the number of pairs added to \mathcal{L} is at least $t/3 - n(n-1)/4$ because there can be at most $\lfloor n(n-1)/4 \rfloor$ steps S_i for which neither j nor j' is defined. \square

Consider a fixed pair of 2-changes linked by an edge. Without loss of generality assume that in the first step the edges $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are exchanged with the edges $\{v_1, v_3\}$ and $\{v_2, v_4\}$, for distinct vertices v_1, \dots, v_4 . Also without loss of generality assume that in the second step the edges $\{v_1, v_3\}$ and $\{v_5, v_6\}$ are exchanged with the edges $\{v_1, v_5\}$ and $\{v_3, v_6\}$. However, note that the vertices v_5 and v_6 are not necessarily




Figure 4.1: A pair of type 1.




Figure 4.2: Pairs of type 2.

distinct from the vertices v_2 and v_4 . We distinguish between three different types of pairs.

1. $|\{v_2, v_4\} \cap \{v_5, v_6\}| = 0$. This case is illustrated in Figure 4.1.
2. $|\{v_2, v_4\} \cap \{v_5, v_6\}| = 1$. We can assume w.l.o.g. that $v_2 \in \{v_5, v_6\}$. We have to distinguish between two subcases: a) The edges $\{v_1, v_5\}$ and $\{v_2, v_3\}$ are added to the tour in the second step. b) The edges $\{v_1, v_2\}$ and $\{v_3, v_5\}$ are added to the tour in the second step. These cases are illustrated in Figure 4.2.
3. $|\{v_2, v_4\} \cap \{v_5, v_6\}| = 2$. The case $v_2 = v_5$ and $v_4 = v_6$ cannot appear as it would imply that the tour is not changed by performing the considered pair of steps. Hence, for pairs of this type, we must have $v_2 = v_6$ and $v_4 = v_5$.

When distances are measured according to the Euclidean metric, pairs of type 3 result in vast dependencies and hence the probability that there exists a pair of this type in which both steps are improvements by at most ε w.r.t. the Euclidean metric cannot be bounded appropriately. In order to reduce the number of cases we have to consider and in order to prepare the analysis of ϕ -perturbed Euclidean instances, we exclude pairs of type 3 from our probabilistic analysis by leaving out all pairs of type 3 when constructing the list \mathcal{L} in the proof of Lemma 4.2.

We only need to show that there are always enough pairs of type 1 or 2. Consider two steps S_i and S_j with $i < j$ that form a pair of type 3. Assume that in step S_i the edges $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are replaced by the edges $\{v_1, v_3\}$ and $\{v_2, v_4\}$, and that in step S_j these edges are replaced by the edges $\{v_1, v_4\}$ and $\{v_2, v_3\}$. Now consider the next step S_l with $l > j$ in which the edge $\{v_1, v_4\}$ is removed from the tour if such a step exists and the next step $S_{l'}$ with $l' > j$ in which the edge $\{v_2, v_3\}$ is removed from

the tour if such a step exists. Observe that neither (S_j, S_l) nor $(S_j, S_{l'})$ can be a pair of type 3 because otherwise the improvement of one of the steps S_i, S_j , and S_l , or $S_{l'}$, respectively, must be negative. In particular, we must have $l \neq l'$.

If we encounter a pair (S_i, S_j) of type 3 in the construction of the list \mathcal{L} , we mark step S_i as being processed without adding a pair of 2-changes to \mathcal{L} and without removing S_j from the sequence of steps to be processed. Let x denote the number of pairs of type 3 that we encounter during the construction of the list \mathcal{L} . Our argument above shows that the number of pairs of type 1 or 2 that are added to \mathcal{L} is at least $x - n(n-1)/4$. This implies $t \geq 2x - n(n-1)/4$ and $x \leq t/2 + n(n-1)/8$. Hence, the number of relevant steps reduces from t to $t' = t - x \geq t/2 - n(n-1)/8$. Using this estimate in Lemma 4.2 yields the following lemma.

Lemma 4.3. *In every sequence of t consecutive 2-changes the number of disjoint pairs of 2-changes of type 1 or 2 is at least $t/6 - 7n(n-1)/24$. \square*

4.1.2 Analysis of pairs of linked 2-changes

The following lemma gives a bound on the probability that there exists a pair of type 1 or 2 in which both steps are small improvements.

Lemma 4.4. *In a ϕ -perturbed Manhattan instance with n vertices, the probability that there exists a pair of type 1 or type 2 in which both 2-changes are improvements by at most ε is bounded by $O(n^6 \cdot \varepsilon^2 \cdot \phi^2)$.*

Proof. First, we consider pairs of type 1. We assume that in the first step the edges $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are replaced by the edges $\{v_1, v_3\}$ and $\{v_2, v_4\}$ and that in the second step the edges $\{v_1, v_3\}$ and $\{v_5, v_6\}$ are replaced by the edges $\{v_1, v_5\}$ and $\{v_3, v_6\}$. For $i \in [6]$, let $x^i \in \mathbb{R}^d$ denote the coordinates of vertex v_i . Furthermore, let Δ_1 denote the (possibly negative) improvement of the first step and let Δ_2 denote the (possibly negative) improvement of the second step. The random variables Δ_1 and Δ_2 can be written as

$$\Delta_1 = \sum_{i=1}^d |x_i^1 - x_i^2| + |x_i^3 - x_i^4| - |x_i^1 - x_i^3| - |x_i^2 - x_i^4|$$

and

$$\Delta_2 = \sum_{i=1}^d |x_i^1 - x_i^3| + |x_i^5 - x_i^6| - |x_i^1 - x_i^5| - |x_i^3 - x_i^6| .$$

For any fixed order of the coordinates, Δ_1 and Δ_2 can be expressed as linear combinations of the coordinates with integer coefficients. For $i \in [d]$, let σ_i denote an order of the coordinates x_i^1, \dots, x_i^6 , let $\sigma = (\sigma_1, \dots, \sigma_d)$, and let Δ_1^σ and Δ_2^σ denote the corresponding linear combinations. We denote by \mathcal{A} the event that both Δ_1 and Δ_2 take values in the interval $(0, \varepsilon]$, and we denote by \mathcal{A}^σ the event that both linear combinations Δ_1^σ and Δ_2^σ take values in the interval $(0, \varepsilon]$. Obviously \mathcal{A} can only occur if for at least one σ , the event \mathcal{A}^σ occurs. Hence, we obtain

$$\Pr[\mathcal{A}] \leq \sum_{\sigma} \Pr[\mathcal{A}^\sigma] .$$

Since there are $(6!)^d$ different orders σ , which is constant for constant dimension d , it suffices to show that for every tuple of orders σ , the probability of the event \mathcal{A}^σ is bounded from above by $O(\varepsilon^2\phi^2)$. Then a union bound over all possible pairs of linked 2-changes of type 1 yields the lemma for pairs of type 1.

We divide the set of possible pairs of linear combinations $(\Delta_1^\sigma, \Delta_2^\sigma)$ into three classes. We say that a pair of linear combinations belongs to class A if at least one of the linear combinations equals 0, we say that it belongs to class B if $\Delta_1^\sigma = -\Delta_2^\sigma$, and we say that it belongs to class C if Δ_1^σ and Δ_2^σ are linearly independent. For tuple of orders σ that yield pairs from class A or B, the event \mathcal{A}^σ can never occur because in both cases the value of at least one linear combination is at most 0. For tuples σ that yield pairs from class C, we can apply Lemma A.1 from Appendix A, which shows that the probability of the event \mathcal{A}^σ is bounded from above by $(\varepsilon\phi)^2$. Hence, we only need to show that every pair $(\Delta_1^\sigma, \Delta_2^\sigma)$ of linear combinations belongs either to class A, B, or C.

Consider a fixed tuple of orders $\sigma = (\sigma_1, \dots, \sigma_d)$. We split Δ_1^σ and Δ_2^σ into d parts that correspond to the d dimensions. To be precise, for $j \in [2]$, we write $\Delta_j^\sigma = \sum_{i \in [d]} X_i^{\sigma_i, j}$, where $X_i^{\sigma_i, j}$ is a linear combination of the variables x_i^1, \dots, x_i^6 . For $i \in [d]$, we show that the pair of linear combinations $(X_i^{\sigma_i, 1}, X_i^{\sigma_i, 2})$ belongs either to class A, B, or C. This directly implies that also $(\Delta_1^\sigma, \Delta_2^\sigma)$ must belong to one of these classes.

Assume that the pair of linear combinations $(X_i^{\sigma_i, 1}, X_i^{\sigma_i, 2})$ is linearly dependent for the fixed order σ_i . Observe that this can only happen if $X_i^{\sigma_i, 1}$ does not contain x_i^2 and x_i^4 and if $X_i^{\sigma_i, 2}$ does not contain x_i^5 and x_i^6 . The former can only happen if either $x_i^3 \geq x_i^4$, $x_i^2 \geq x_i^4$, and $x_i^2 \geq x_i^1$ or if $x_i^3 \leq x_i^4$, $x_i^2 \leq x_i^4$, and $x_i^2 \leq x_i^1$. The latter can only happen if either $x_i^5 \geq x_i^6$, $x_i^3 \geq x_i^6$, and $x_i^5 \geq x_i^1$ or if $x_i^5 \leq x_i^6$, $x_i^3 \leq x_i^6$, and $x_i^5 \leq x_i^1$.

If one chooses the order such that x_i^2 , x_i^4 , x_i^5 , and x_i^6 cancel out and such that $x_i^1 \geq x_i^3$, one can verify by a case distinction that $X_i^{\sigma_i, 1} \in \{0, -2x_i^1 + 2x_i^3\}$ and $X_i^{\sigma_i, 2} \in \{0, 2x_i^1 - 2x_i^3\}$. Hence, in this case the resulting pair of linear combinations belongs either to class A or B. Analogously, if one chooses the order such that x_i^2 , x_i^4 , x_i^5 , and x_i^6 cancel out and such that $x_i^3 \geq x_i^1$, we have $X_i^{\sigma_i, 1} \in \{0, 2x_i^1 - 2x_i^3\}$ and $X_i^{\sigma_i, 2} \in \{0, -2x_i^1 + 2x_i^3\}$. Hence, also in this case, the pair of resulting linear combinations belongs either to class A or B.

With similar arguments we prove the lemma for pairs of type 2. We first prove the lemma for pairs of type 2 a). Using the same notations as for pairs of type 1, we can write the improvement Δ_2 as

$$\Delta_2 = \sum_{i \in [d]} |x_i^1 - x_i^3| + |x_i^2 - x_i^5| - |x_i^1 - x_i^5| - |x_i^2 - x_i^3|.$$

Again we show that, for every $i \in [d]$ and every order σ_i , the pair of linear combinations $(X_i^{\sigma_i, 1}, X_i^{\sigma_i, 2})$ belongs either to class A, B, or C. Assume that the pair is linearly dependent for the fixed order σ_i . Observe that this can only happen if $X_i^{\sigma_i, 1}$ does not contain x_i^4 and if $X_i^{\sigma_i, 2}$ does not contain x_i^5 . The former can only happen if either $x_i^3 \geq x_i^4$ and $x_i^2 \geq x_i^4$ or if $x_i^3 \leq x_i^4$ and $x_i^2 \leq x_i^4$. The latter can only happen if either $x_i^2 \geq x_i^5$ and $x_i^1 \geq x_i^5$ or if $x_i^2 \leq x_i^5$ and $x_i^1 \leq x_i^5$.

If one chooses the order such that x_i^4 and x_i^5 cancel out and such that $x_i^1 \geq x_i^3$, one can verify by a case distinction that $X_i^{\sigma_i, 1} \in \{0, -2x_i^1 + 2x_i^2, -2x_i^1 + 2x_i^3, -2x_i^1 + 2x_i^3\}$

and $X_i^{\sigma_i,2} \in \{0, 2x_i^1 - 2x_i^2, 2x_i^1 - 2x_i^3, 2x_i^2 - 2x_i^3\}$. Hence, under the assumption that the linear combinations $X_i^{\sigma_i,1}$ and $X_i^{\sigma_i,2}$ are linearly dependent, the pair of resulting linear combinations in the case $x_i^1 \geq x_i^3$ belongs either to class A or B. Analogously, if one chooses the order such that x_i^4 and x_i^5 cancel out and such that $x_i^3 \geq x_i^1$, we have $X_i^{\sigma_i,1} \in \{0, 2x_i^1 - 2x_i^2, 2x_i^1 - 2x_i^3, 2x_i^2 - 2x_i^3\}$ and $X_i^{\sigma_i,2} \in \{0, -2x_i^1 + 2x_i^2, -2x_i^1 + 2x_i^3, -2x_i^2 + 2x_i^3\}$. Hence, also in this case, the pair of resulting linear combinations belongs either to class A or B.

It remains to consider pairs of type 2 b). For these pairs, we can write Δ_2 as

$$\Delta_2 = \sum_{i=1}^d |x_i^1 - x_i^3| + |x_i^2 - x_i^5| - |x_i^1 - x_i^2| - |x_i^3 - x_i^5| .$$

Assume that the pair of linear combinations $(X_i^{\sigma_i,1}, X_i^{\sigma_i,2})$ is linearly dependent for the fixed order σ_i . Observe that this can only happen if $X_i^{\sigma_i,1}$ does not contain x_i^4 and if $X_i^{\sigma_i,2}$ does not contain x_i^5 . As we have already seen for pairs of type 2 a), the former can only happen if either $x_i^3 \geq x_i^4$ and $x_i^2 \geq x_i^4$ or if $x_i^3 \leq x_i^4$ and $x_i^2 \leq x_i^4$. The latter can only happen if either $x_i^2 \geq x_i^5$ and $x_i^3 \geq x_i^5$ or if $x_i^2 \leq x_i^5$ and $x_i^3 \leq x_i^5$.

If one chooses the order such that x_i^4 and x_i^5 cancel out and such that $x_i^1 \geq x_i^3$, one can verify by a case distinction that $X_i^{\sigma_i,1} \in \{0, -2x_i^1 + 2x_i^2, -2x_i^1 + 2x_i^3, -2x_i^2 + 2x_i^3\}$ and $X_i^{\sigma_i,2} \in \{0, 2x_i^1 - 2x_i^2, 2x_i^1 - 2x_i^3, 2x_i^2 - 2x_i^3\}$. Hence, under the assumption that the linear combinations $X_i^{\sigma_i,1}$ and $X_i^{\sigma_i,2}$ are linearly dependent, the pair of resulting linear combinations belongs either to class A or B. If one chooses the order such that x_i^4 and x_i^5 cancel out and such that $x_i^3 \geq x_i^1$, one can verify by a case distinction that $X_i^{\sigma_i,1} \in \{0, 2x_i^1 - 2x_i^2, 2x_i^1 - 2x_i^3, 2x_i^2 - 2x_i^3\}$ and $X_i^{\sigma_i,2} \in \{0, -2x_i^1 + 2x_i^2, -2x_i^1 + 2x_i^3, -2x_i^2 + 2x_i^3\}$. Hence, also in this case, the pair of resulting linear combinations belongs either to class A or B. \square

4.1.3 Expected number of 2-changes

Based on Lemmas 4.3 and 4.4, we are now able to prove part a) of Theorem 1.2.

Proof of Theorem 1.2 a). Let T denote the random variable that describes the length of the longest path in the state graph. If $T \geq t$, then there must exist a sequence S_1, \dots, S_t of t consecutive 2-changes in the state graph. We start by identifying a set of linked pairs of type 1 and 2 in this sequence. Due to Lemma 4.3, we know that we can find at least $t/6 - 7n(n-1)/24$ such pairs. Let Δ_{\min}^* denote the smallest improvement made by any pair of improving 2-Opt steps of type 1 or 2. For $t > 2n^2$, we have $t/6 - 7n(n-1)/24 > t/48$ and hence due to Lemma 4.4,

$$\Pr [T \geq t] \leq \Pr \left[\Delta_{\min}^* \leq \frac{48dn}{t} \right] = O \left(\min \left\{ \frac{n^8 \phi^2}{t^2}, 1 \right\} \right) .$$

Since T cannot exceed $(n!)$, this implies the following bound on the expected number of 2-changes:

$$\mathbf{E}[T] \leq 2n^2 + \sum_{t=2n^2+1}^{n!} O \left(\min \left\{ \frac{n^8 \phi^2}{t^2}, 1 \right\} \right) = O(n^4 \cdot \phi) .$$

This concludes the proof of part a) of the theorem. \square

Chandra, Karloff, and Tovey [CKT99] show that for every metric that is induced by a norm on \mathbb{R}^d , and for any set of n points in the unit hypercube $[0, 1]^d$, the optimal tour visiting all n points has length $O(n^{(d-1)/d})$. Furthermore, every insertion heuristic finds an $O(\log n)$ -approximation [RSI77]. Hence, if one starts with a solution computed by an insertion heuristic, the initial tour has length $O(n^{(d-1)/d} \cdot \log n)$. Using this observation yields part a) of Theorem 1.3.

Proof of Theorem 1.3 b). Since the initial tour has length $O(n^{(d-1)/d} \cdot \log n)$, we obtain for an appropriate constant c and $t > 2n^2$,

$$\Pr[T \geq t] \leq \Pr\left[\Delta_{\min}^* \leq \frac{c \cdot n^{(d-1)/d} \cdot \log n}{t}\right] = O\left(\min\left\{\frac{n^{8-2/d} \cdot \log^2 n \cdot \phi^2}{t^2}, 1\right\}\right).$$

This yields

$$\mathbf{E}[T] \leq 2n^2 + \sum_{t=2n^2+1}^{n!} O\left(\min\left\{\frac{n^{8-2/d} \cdot \log^2 n \cdot \phi^2}{t^2}, 1\right\}\right) = O(n^{4-1/d} \cdot \log n \cdot \phi). \quad \square$$

4.2 Euclidean Instances

In this section, we analyze the expected number of 2-changes on ϕ -perturbed Euclidean instances. The analysis is similar to the analysis of Manhattan instances in the previous section, only Lemma 4.4 needs to be replaced by its equivalent version for the L_2 metric.

Lemma 4.5. *For ϕ -perturbed L_2 instances, the probability that there exists a pair of type 1 or type 2 in which both 2-changes are improvements by at most $\varepsilon \leq 1/2$ is bounded by $O(n^6 \cdot \phi^5 \cdot \varepsilon^2 \cdot \log^2(1/\varepsilon)) + O(n^5 \cdot \phi^4 \cdot \varepsilon^{3/2} \cdot \log(1/\varepsilon))$.*

The bound that this lemma provides is slightly weaker than its L_1 counterpart, and hence also the bound on the expected running time is slightly worse for L_2 instances. The crucial step to prove Lemma 4.5 is to gain a better understanding of the random variable that describes the improvement of a single fixed 2-change. In the next section, we analyze this random variable under several conditions, e. g., under the condition that the length of one of the involved edges is fixed. With the help of these results, pairs of linked 2-changes can easily be analyzed. Let us mention that our analysis of a single 2-change yields a bound of $O(n^7 \cdot \log^2(n) \cdot \phi^3)$ for the expected number of 2-changes. For Euclidean instances in which all points are distributed uniformly at random over the unit square, this bound already improves the best previously known bound of $O(n^{10} \cdot \log n)$.

4.2.1 Analysis of a single 2-change

We analyze a 2-change in which the edges $\{O, Q_1\}$ and $\{P, Q_2\}$ are exchanged with the edges $\{O, Q_2\}$ and $\{P, Q_1\}$ for some vertices O, P, Q_1 , and Q_2 . In the input model we consider, each of these points has a probability distribution over the unit hypercube according to which it is chosen. In this section, we consider a simplified




Figure 4.3: The random variable Z is defined as $r - \mathbf{d}(P, Q)$.

random experiment in which O is chosen to be the origin and P , Q_1 , and Q_2 are chosen independently and uniformly at random from a d -dimensional hyperball with radius \sqrt{d} centered at the origin. In the next section, we argue that the analysis of this simplified random experiment helps to analyze the actual random experiment that occurs in the probabilistic input model.

Due to the rotational symmetry of the simplified model, we assume without loss of generality that P lies at position $(0, T)$ for some $T \geq 0$. For $i \in [2]$, Let Z_i denote the difference $\mathbf{d}(O, Q_i) - \mathbf{d}(P, Q_i)$. Then the improvement Δ of the 2-change can be expressed as $Z_1 - Z_2$. The random variables Z_1 and Z_2 are identically distributed, and they are independent if T is fixed. We denote by $f_{Z|T=\tau, R=r}$ the density of Z_1 and Z_2 under the conditions that $\mathbf{d}(O, Q_1) = r$ and $\mathbf{d}(O, Q_2) = r$, respectively, and $T = \tau$.

Lemma 4.6. For $\tau, r \in (0, \sqrt{d}]$, and $z \in (-\tau, \min\{\tau, 2r - \tau\})$,

$$f_{Z|T=\tau, R=r}(z) \leq \begin{cases} \sqrt{\frac{2}{\tau^2 - z^2}} & \text{if } r \geq \tau, \\ \sqrt{\frac{2}{(\tau+z)(2r-\tau-z)}} & \text{if } r \leq \tau. \end{cases}$$

For $z \notin [-\tau, \min\{\tau, 2r - \tau\}]$, the density $f_{Z|T=\tau, R=r}(z)$ is 0.

Proof. We denote by Z the random variable $\mathbf{d}(O, Q) - \mathbf{d}(P, Q)$, where Q is a point chosen uniformly at random from a d -dimensional hyperball with radius \sqrt{d} centered at the origin. In the following, we assume that the plane spanned by the points O , P , and Q is fixed arbitrarily, and we consider the random experiment conditioned on the event that Q lies in this plane. To make the calculations simpler, we use polar coordinates to describe the location of Q . Since the radius $\mathbf{d}(O, Q) = r$ is given, the point Q is completely determined by the angle α between the y -axis and the line between O and Q (see Figure 4.3). Hence, the random variable Z can be written as

$$Z = r - \sqrt{r^2 + \tau^2 - 2r\tau \cdot \cos \alpha} .$$

It is easy to see that Z can only take values in the interval $[-\tau, \min\{\tau, 2r - \tau\}]$, and hence the density $f_{Z|T=\tau, R=r}(z)$ is 0 outside this interval.

Since Q is chosen uniformly at random, the angle α is chosen uniformly at random from the interval $[0, 2\pi)$. For symmetry reasons, we can assume that α is chosen uniformly from the interval $[0, \pi)$. When α is restricted to the interval $[0, \pi)$, then there

exists a unique inverse function mapping Z to α , namely

$$\alpha(z) = \arccos\left(\frac{\tau^2 + 2zr - z^2}{2r\tau}\right) .$$

The density $f_{Z|T=\tau, R=r}$ can be expressed as

$$f_{Z|T=\tau, R=r}(z) = f_\alpha(\alpha(z)) \cdot \left| \frac{d}{dz} \alpha(z) \right| = -\frac{1}{\pi} \cdot \frac{d}{dz} \alpha(z) ,$$

where f_α denotes the density of α , i. e., the uniform density over $[0, \pi)$. For $|x| < 1$, the derivative of the arc cosine is

$$(\arccos(x))' = -\frac{1}{\sqrt{1-x^2}} .$$

Hence, the derivative of $\alpha(z)$ equals

$$\frac{r-z}{r\tau} \cdot \frac{-1}{\sqrt{1 - \frac{(\tau^2 + 2zr - z^2)^2}{4r^2\tau^2}}} = \frac{2(z-r)}{\sqrt{4r^2\tau^2 - 4r^2z^2 - 4r\tau^2z + 4rz^3 - \tau^4 + 2\tau^2z^2 - z^4}} .$$

In order to prove the lemma, we distinguish between the cases $r \geq \tau$ and $r \leq \tau$.

First case: $r \geq \tau$.

In this case, it suffices to show

$$4r^2\tau^2 - 4r^2z^2 - 4r\tau^2z + 4rz^3 - \tau^4 + 2\tau^2z^2 - z^4 \geq 2(z-r)^2(\tau^2 - z^2) ,$$

which is implied by

$$\begin{aligned} & 4r^2\tau^2 - 4r^2z^2 - 4r\tau^2z + 4rz^3 - \tau^4 + 2\tau^2z^2 - z^4 - 2(z-r)^2(\tau^2 - z^2) \\ &= 2r^2(\tau^2 - z^2) - \tau^4 + z^4 \geq 2\tau^2(\tau^2 - z^2) - \tau^4 + z^4 = (\tau^2 - z^2)^2 \geq 0 . \end{aligned}$$

Second case: $r \leq \tau$.

In this case, it suffices to show

$$4r^2\tau^2 - 4r^2z^2 - 4r\tau^2z + 4rz^3 - \tau^4 + 2\tau^2z^2 - z^4 \geq 2(z-r)^2(\tau+z)(2r-\tau-z) ,$$

which is implied by

$$\begin{aligned} & 4r^2\tau^2 - 4r^2z^2 - 4r\tau^2z + 4rz^3 - \tau^4 + 2\tau^2z^2 - z^4 - 2(z-r)^2(\tau+z)(2r-\tau-z) \geq 0 \\ & \iff (-2r+z+\tau)(\tau+z)(z^2+2\tau z-2rz+2r^2-\tau^2-2\tau r) \geq 0 \\ & \iff z^2+2\tau z-2rz+2r^2-\tau^2-2\tau r \leq 0 , \end{aligned}$$

which follows from

$$\begin{aligned} & z^2 + 2\tau z - 2rz + 2r^2 - \tau^2 - 2\tau r \\ &= z^2 + 2z(\tau - r) + 2r^2 - \tau^2 - 2\tau r \\ &\leq (2r - \tau)^2 + 2(2r - \tau)(\tau - r) + 2r^2 - \tau^2 - 2\tau r \\ &= 2(r^2 - \tau^2) \leq 0 . \end{aligned}$$

□

Based on Lemma 4.6, the density of the random variable $\Delta = Z_1 - Z_2$ under the conditions $R_1 := d(O, Q_1) = r_1$, $R_2 := d(O, Q_2) = r_2$, and $T := d(O, P) = \tau$ can be computed as the convolution of the density $f_{Z|T=\tau, R=r}$ with itself.

Lemma 4.7. *Let $\tau, r_1, r_2 \in (0, \sqrt{d}]$, and let Z_1 and Z_2 be independent random variables drawn according to the densities $f_{Z|T=\tau, R=r_1}$ and $f_{Z|T=\tau, R=r_2}$, respectively. For $\delta \in (0, 1/2]$ and a sufficiently large constant κ , the density $f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta)$ of the random variable $\Delta = Z_1 - Z_2$ is bounded from above by*

$$\begin{cases} \frac{\kappa}{\tau} \cdot \ln \delta^{-1} & \text{if } \tau \leq r_1, \tau \leq r_2, \\ \frac{\kappa}{\sqrt{r_1 r_2}} \cdot (\ln \delta^{-1} + \ln |2(r_1 - r_2) - \delta|^{-1} + \kappa) & \text{if } r_1 \leq \tau, r_2 \leq \tau, \\ \frac{\kappa}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} & \text{if } r_1 \leq \tau \leq r_2, \\ \frac{\kappa}{\sqrt{\tau r_2}} \cdot (\ln \delta^{-1} + \ln |2(\tau - r_2) - \delta|^{-1} + \kappa) & \text{if } r_2 \leq \tau \leq r_1. \end{cases}$$

The simple but somewhat tedious calculation that yields Lemma 4.7 is deferred to Appendix B.1. In order to prove Lemma 4.5, we need bounds on the densities of the random variables Δ , Z_1 , and Z_2 under certain conditions. We summarize these bounds in the following lemma.

Lemma 4.8. *Let $\tau, r \in (0, \sqrt{d}]$, $\delta \in (0, 1/2]$, and let κ denote a sufficiently large constant.*

a) *For $i \in [2]$, the density of Δ under the condition $R_i = r$ is bounded by*

$$f_{\Delta|R_i=r}(\delta) \leq \frac{\kappa}{\sqrt{r}} \cdot \ln \delta^{-1} .$$

b) *The density of Δ under the condition $T = \tau$ is bounded by*

$$f_{\Delta|T=\tau}(\delta) \leq \frac{\kappa}{\tau} \cdot \ln \delta^{-1} .$$

c) *The density of Δ is bounded by*

$$f_{\Delta}(\delta) \leq \kappa \cdot \ln \delta^{-1} .$$

d) *For $i \in [2]$, the density of Z_i under the condition $T = \tau$ is bounded by*

$$f_{Z_i|T=\tau}(z) \leq \frac{\kappa}{\sqrt{\tau^2 - z^2}}$$

if $|z| < \tau$. Since Z_i takes only values in the interval $[-\tau, \tau]$, the conditional density $f_{Z_i|T=\tau}(z)$ is 0 for $z \notin [-\tau, \tau]$.

Lemma 4.8 follows from Lemmas 4.6 and 4.7 by integrating over all values of the unconditioned distances. The proof can be found in Appendix B.2.

4.2.2 Simplified random experiments

In the previous section we did not analyze the random experiment that really takes place. Instead of choosing the points according to the given density functions, we simplified their distributions by placing point O in the origin and by giving the other points P , Q_1 , and Q_2 uniform distributions centered around the origin. In our input model, however, each of these points is described by a density function over the unit hypercube. We consider the probability of the event $\Delta \in [0, \varepsilon]$ in both the original input model as well as in the simplified random experiment. In the following, we denote this event by \mathcal{E} . We claim that the simplified random experiment that we analyze is only slightly dominated by the original random experiment, in the sense that the probability of the event \mathcal{E} in the simplified random experiment is smaller by at most some factor depending on ϕ .

In order to compare the probabilities in the original and in the simplified random experiment, consider the original experiment and assume that the point O lies at position $x \in [0, 1]^d$. Then one can identify a region $\mathcal{R}_x \subseteq \mathbb{R}^{3d}$ with the property that the event \mathcal{E} occurs if and only if the random vector (P, Q_1, Q_2) lies in \mathcal{R}_x . No matter of how the position x of O is chosen, this region always has the same shape, only its position is shifted. Let $\mathcal{V} = \sup_{x \in [0, 1]^d} \text{Vol}(\mathcal{R}_x \cap [0, 1]^{3d})$. Then the probability of \mathcal{E} can be bounded from above by $\phi^3 \cdot \mathcal{V}$ in the original random experiment. Since Δ is invariant under translating O , P , Q_1 , and Q_2 by the same vector, we obtain

$$\begin{aligned} \text{Vol}\left(\mathcal{R}_x \cap [0, 1]^{3d}\right) &= \text{Vol}\left(\mathcal{R}_{0^d} \cap ([-x_1, 1-x_1] \times \cdots \times [-x_d, 1-x_d])^3\right) \\ &\leq \text{Vol}\left(\mathcal{R}_{0^d} \cap [-1, 1]^{3d}\right) . \end{aligned}$$

Hence, $\mathcal{V} \leq \mathcal{V}' := \text{Vol}(\mathcal{R}_{0^d} \cap [-1, 1]^{3d})$. Since the hyperball centered around the origin with radius \sqrt{d} contains the hypercube $[-1, 1]^d$ completely, the probability of \mathcal{E} in the simplified random experiment can be bounded from below by $\mathcal{V}'/V_d(\sqrt{d})^3$, where $V_d(\sqrt{d})$ denotes the volume of a d -dimensional hyperball with radius \sqrt{d} . Since this hyperball is contained in a hypercube with side length $2\sqrt{d}$, its volume can be bounded from above by $(4d)^{d/2}$. Hence, the probability of \mathcal{E} in the simplified random experiment is smaller by at most a factor of $\phi^3(4d)^{3d/2}$ compared to the original random experiment.

Taking into account this factor and using Lemma 4.8 c) and a union bound over all possible 2-changes yields the following lemma about the improvement of a single 2-change.

Lemma 4.9. *The probability that there exists an improving 2-change whose improvement is at most $\varepsilon \leq 1/2$ is bounded from above by $O(n^4 \cdot \phi^3 \cdot \varepsilon \cdot \log(1/\varepsilon))$.*

Using similar arguments as in the proof of Theorem 4.1 yields the following upper bound on the expected number of 2-changes.

Theorem 4.10. *Starting with an arbitrary tour, the expected number of steps performed by 2-Opt on ϕ -perturbed Euclidean instances is $O(n^7 \cdot \log^2(n) \cdot \phi^3)$.*

Pairs of type 1. In order to improve upon Theorem 4.10, we consider pairs of linked 2-changes as in the analysis of ϕ -perturbed Manhattan instances. Since our analysis of

pairs of linked 2-changes is based on the analysis of a single 2-change that we presented in the previous section, we also have to consider simplified random experiments when analyzing pairs of 2-changes. For a fixed pair of type 1, we assume that point v_3 is chosen to be the origin and the other points v_1, v_2, v_4, v_5 , and v_6 are chosen uniformly at random from a hyperball with radius \sqrt{d} centered at v_3 . Let \mathcal{E} denote the event that both Δ_1 and Δ_2 lie in the interval $[0, \varepsilon]$, for some given ε . With the same arguments as above, one can see that the probability of \mathcal{E} in the simplified random experiment is smaller compared to the original experiment by at most a factor of $((4d)^{d/2}\phi)^5$.

Pairs of type 2. For a fixed pair of type 2, we consider the simplified random experiment in which v_2 is placed in the origin and the other points v_1, v_3, v_4 , and v_5 are chosen uniformly at random from a hyperball with radius \sqrt{d} centered at v_2 . In this case, the probability in the simplified random experiment is smaller by at most a factor of $((4d)^{d/2}\phi)^4$.

4.2.3 Analysis of pairs of linked 2-changes

Finally, we can prove Lemma 4.5.

Proof of Lemma 4.5. We start by considering pairs of type 1. We consider the simplified random experiment in which v_3 is chosen to be the origin and the other points are drawn uniformly at random from a hyperball with radius \sqrt{d} centered at v_3 . If the position of the point v_1 is fixed, then the events $\Delta_1 \in [0, \varepsilon]$ and $\Delta_2 \in [0, \varepsilon]$ are independent as only the vertices v_1 and v_3 appear in both the first and the second step. In fact, because the densities of the points v_2, v_4, v_5 , and v_6 are rotationally symmetric, the concrete position of v_1 is not important in our simplified random experiment anymore, but only the distance R between v_1 and v_3 is of interest.

For $i \in [2]$, we determine the conditional probability of the event $\Delta_i \in [0, \varepsilon]$ under the condition that the distance $d(v_1, v_3)$ is fixed with the help of Lemma 4.8 a) and obtain

$$\begin{aligned} \Pr[\Delta_i \in [0, \varepsilon] \mid d(v_1, v_3) = r] &= \int_0^\varepsilon f_{\Delta \mid R_i=r}(\delta) d\delta \\ &\leq \int_0^\varepsilon \frac{\kappa}{\sqrt{r}} \ln \delta^{-1} d\delta \\ &\leq \frac{3\kappa}{\sqrt{r}} \cdot \varepsilon \cdot \ln(1/\varepsilon) . \end{aligned} \quad (4.2)$$

Since for fixed distance $d(v_1, v_3)$ the random variables Δ_1 and Δ_2 are independent, we obtain

$$\Pr[\Delta_1, \Delta_2 \in [0, \varepsilon] \mid d(v_1, v_3) = r] \leq \frac{9\kappa^2}{r} \cdot \varepsilon^2 \cdot \ln^2(1/\varepsilon) . \quad (4.3)$$

For $r \in [0, \sqrt{d}]$, the density $f_{d(v_1, v_3)}$ of the random variable $d(v_1, v_3)$ in the simplified random experiment is $r^{d-1}/d^{d/2-1}$. Combining this observation with the bound given in (4.3) yields

$$\Pr[\Delta_1, \Delta_2 \in [0, \varepsilon]] \leq \int_0^{\sqrt{d}} \frac{9\kappa^2 \cdot r^{d-2}}{d^{d/2-1}} \cdot \varepsilon^2 \cdot \ln^2(1/\varepsilon) dr = O(\varepsilon^2 \cdot \ln^2(1/\varepsilon)) .$$

There are $O(n^6)$ different pairs of type 1, hence a union bound over all of them concludes the first part of the proof when taking into account the factor $((4d)^{d/2}\phi)^5$ that results from considering the simplified random experiment.

It remains to consider pairs of type 2. We consider the simplified random experiment in which v_2 is chosen to be the origin and the other points are drawn uniformly at random from a hyperball with radius \sqrt{d} centered at v_2 . In contrast to pairs of type 1, pairs of type 2 exhibit larger dependencies as only 5 different vertices are involved in these pairs. Fix one pair of type 2. The two 2-changes share the whole triangle consisting of v_1, v_2 , and v_3 . In the second step, there is only one new vertex, namely v_5 . Hence, there is not enough randomness contained in a pair of type 2 such that Δ_1 and Δ_2 are nearly independent as for pairs of type 1.

We start by considering pairs of type 2 a). First, we analyze the probability that Δ_1 lies in the interval $[0, \varepsilon]$. After that, we analyze the probability that Δ_2 lies in the interval $[0, \varepsilon]$ under the condition that the points v_1, v_2, v_3 , and v_4 have already been chosen. In the analysis of the second step we cannot make use of the fact that the distances $d(v_1, v_3)$ and $d(v_2, v_3)$ are random variables anymore since we exploited their randomness already in the analysis of the first step. The only distances whose randomness we can exploit are the distances $d(v_1, v_5)$ and $d(v_2, v_5)$. We pessimistically assume that the distances $d(v_1, v_3)$ and $d(v_2, v_3)$ have been chosen by an adversary. This means the adversary can determine an interval of length ε in which the random variable $d(v_2, v_5) - d(v_1, v_5)$ must lie in order for Δ_2 to lie in $[0, \varepsilon]$.

Analogously to (4.2), the probability of the event $\Delta_1 \in [0, \varepsilon]$ under the condition $d(v_1, v_2) = r$ can be bounded by

$$\Pr[\Delta_1 \in [0, \varepsilon] \mid d(v_1, v_2) = r] \leq \frac{3\kappa}{\sqrt{r}} \cdot \varepsilon \cdot \ln(1/\varepsilon) . \quad (4.4)$$

Due to Lemma 4.8 d), the conditional density of the random variable $Z = d(v_2, v_5) - d(v_1, v_5)$ under the condition $d(v_1, v_2) = r$ can be bounded by

$$f_{Z \mid d(v_1, v_2) = r}(z) \leq \frac{\kappa}{\sqrt{r^2 - z^2}}$$

for $|z| \leq r$.

The intervals the adversary can specify which have the highest probability of Z falling into them are $[-r, -r + \varepsilon]$ and $[r - \varepsilon, r]$. Hence, the conditional probability of the event $\Delta_2 \in [0, \varepsilon]$ under the condition $d(v_1, v_2) = r$ and for fixed points v_3 and v_4 is bounded from above by

$$\int_{\max\{r-\varepsilon, -r\}}^r \frac{\kappa}{\sqrt{r^2 - z^2}} dz \leq \frac{\kappa}{\sqrt{r}} \cdot \int_{\max\{r-\varepsilon, -r\}}^r \frac{1}{\sqrt{r - |z|}} dz \leq \frac{\kappa' \sqrt{\varepsilon}}{\sqrt{r}}$$

for a sufficiently large constant κ' . Combining this inequality with (4.4) yields

$$\Pr[\Delta_1, \Delta_2 \in [0, \varepsilon] \mid d(v_1, v_2) = r] \leq \frac{3\kappa\kappa'}{r} \varepsilon^{3/2} \cdot \ln(1/\varepsilon) .$$

In order to get rid of the condition $d(v_1, v_2) = r$, we integrate over all possible values the random variable $d(v_1, v_2)$ can take, yielding

$$\Pr[\Delta_1, \Delta_2 \in [0, \varepsilon]] \leq \int_0^{\sqrt{d}} \frac{3\kappa\kappa' \cdot r^{d-2}}{d^{d/2} - 1} \cdot \varepsilon^{3/2} \cdot \ln(1/\varepsilon) dr = O\left(\varepsilon^{3/2} \cdot \ln(1/\varepsilon)\right) .$$

Applying a union bound over all $O(n^5)$ possible pairs of type 2 a) concludes the proof when one takes into account the factor $((4d)^{d/2}\phi)^4$ due to considering the simplified random experiment.

For pairs of type 2 b), the situation looks somewhat similar. We analyze the first step and in the second step, we can only exploit the randomness of the distances $d(v_2, v_5)$ and $d(v_3, v_5)$. Due to Lemma 4.8 b) and similar to (4.2), the probability of the event $\Delta_1 \in [0, \varepsilon]$ under the condition $d(v_2, v_3) = \tau$ can be bounded by

$$\Pr [\Delta_1 \in [0, \varepsilon] \mid d(v_2, v_3) = \tau] \leq \frac{3\kappa}{\tau} \cdot \varepsilon \cdot \ln(1/\varepsilon) . \quad (4.5)$$

The remaining analysis of pairs of type 2 b) can be carried out completely analogously to the analysis of pairs of type 2 a). \square

4.2.4 The expected number of 2-changes

Based on Lemmas 4.3 and 4.5, we are now able to prove part b) of Theorem 1.2.

Proof of Theorem 1.2 b). We use the same notations as in the proof of part a) of the theorem. For $t > 2n^2$, we have $t/6 - 7n(n-1)/24 > t/48$ and hence due to Lemma 4.5,

$$\begin{aligned} \Pr [T \geq t] &\leq \Pr \left[\Delta_{\min}^* \leq \frac{48\sqrt{dn}}{t} \right] \\ &= O \left(\min \left\{ \frac{n^8 \cdot \log^2(t) \cdot \phi^5}{t^2}, 1 \right\} \right) + O \left(\min \left\{ \frac{n^{13/2} \cdot \log(t) \cdot \phi^4}{t^{3/2}}, 1 \right\} \right) . \end{aligned}$$

This implies that the expected length of the longest path in the state graph is bounded from above by

$$2n^2 + \sum_{t=2n^2+1}^{n!} \left(O \left(\min \left\{ \frac{n^8 \cdot \log^2(t) \cdot \phi^5}{t^2}, 1 \right\} \right) + O \left(\min \left\{ \frac{n^{13/2} \cdot \log(t) \cdot \phi^4}{t^{3/2}}, 1 \right\} \right) \right) .$$

Splitting the sums at $t = n^4 \cdot \log(n\phi) \cdot \phi^{5/2}$ and $t = n^{13/3} \cdot \log^{2/3}(n\phi) \cdot \phi^{8/3}$, respectively, yields

$$\mathbf{E} [T] = O \left(n^4 \cdot \log(n\phi) \cdot \phi^{5/2} \right) + O \left(n^{13/3} \cdot \log^{2/3}(n\phi) \cdot \phi^{8/3} \right) .$$

This concludes the proof of part b) of the theorem. \square

Using the same observations as in the proof of Theorem 1.3 a) also yields part b).

Proof of Theorem 1.3 b). Estimating the length of the initial tour by $O(n^{(d-1)/d} \cdot \log n)$ instead of $O(n)$ improves the upper bound on the expected number of 2-changes by a factor of $\Theta(n^{1/d}/\log n)$ compared to Theorem 1.2 b). This observation yields the bound claimed in Theorem 1.3 b). \square

4.3 General Graphs

In this section, we analyze the expected number of 2-changes on ϕ -perturbed graphs. Observe that ϕ -perturbed graphs contain more randomness than ϕ -perturbed Manhattan or Euclidean instances because each edge length is a random variable that is independent of the other edge lengths. It is easy to obtain a polynomial bound on the expected number of local improvements by just estimating the smallest improvement made by any of the 2-changes. For Manhattan and Euclidean instances we improved this simple bound by considering pairs of linked 2-changes. For ϕ -perturbed graphs we pursue the same approach but due to the larger amount of randomness, we are now able to consider not only pairs of linked steps but longer sequences of linked steps.

We know that every sequence of steps that contains k distinct 2-changes shortens the tour by at least $\Delta^{(k)} = \Delta_{\min}^{(1)} + \dots + \Delta_{\min}^{(k)}$, where $\Delta_{\min}^{(i)}$ denotes the i -th smallest improvement. This observation alone, however, does not suffice to improve the simple bound substantially. Instead we show that one can identify in every long enough sequence of consecutive 2-changes, subsequences that are *linked*, where a sequence S_1, \dots, S_k of 2-changes is called *linked* if for every $i \in [k - 1]$, there exists an edge that is added to the tour in step S_i and removed from the tour in step S_{i+1} . We analyze the smallest improvement of a linked sequence that consists of k distinct 2-Opt steps. Obviously, this improvement must be at least $\Delta^{(k)}$ as in the worst-case, the linked sequence consists of the k smallest improvements. Intuitively, one can hope that it is much larger than $\Delta^{(k)}$ because it is unlikely that the k smallest improvements form a sequence of linked steps. We show that this is indeed the case and use this result to prove the desired upper bound on the expected number of 2-changes.

We introduce the notion of *witness sequences*, i. e., linked sequences of 2-changes that satisfy some additional technical properties. We show that the smallest total improvement made by a witness sequence yields an upper bound on the running time. That is, whenever the 2-Opt heuristic needs many local improvement steps to find a locally optimal solution, there must be a witness sequence whose total improvement is small. Furthermore, our probabilistic analysis reveals that it is unlikely that there exists a witness sequence whose total improvement is small. Together, these results yield the desired bound on the expected number of 2-changes.

4.3.1 Definition of witness sequences

In this section, we give a formal definition of the notion of a *k-witness sequence*. As mentioned above, a witness sequence S_1, \dots, S_k has to be linked, i. e., for $i \in [k - 1]$, there must exist an edge that is added to the tour in step S_i and removed from the tour in step S_{i+1} . Let m denote the number of edges in the graph. Then there are at most $4^{k-1} \cdot m^{k+1}$ such linked sequences as there are at most m^2 different choices for S_1 , and once S_i is fixed, there are at most $4m$ different choices for S_{i+1} . For a fixed 2-change, the probability that it is an improvement by at most ε is bounded by $\varepsilon\phi$. We would like to show an upper bound of $(\varepsilon\phi)^k$ on the probability that each step in the witness sequence S_1, \dots, S_k is an improvement by at most ε . For general linked sequences, this is not true as the steps can be dependent in various ways. Hence, we need to introduce further restrictions on witness sequences.

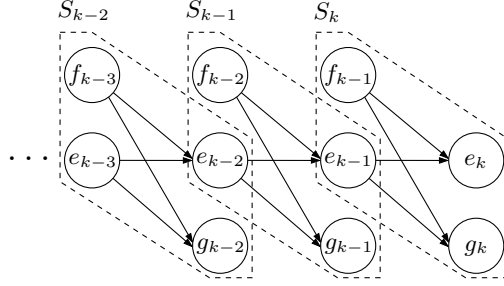


Figure 4.4: Illustration of the notations used in Definitions 4.11, 4.12, and 4.13. Every node in the DAG corresponds to an edge involved in one of the 2-changes. An arc from a node u to a node v indicates that in one of the 2-changes, the edge corresponding to node u is removed from the tour and the edge corresponding to node v is added to the tour. Hence, every arc is associated with one 2-change.

In the following definitions, we assume that a linked sequence S_1, \dots, S_k of 2-changes is given. For $i \in [k]$, in step S_i the edges e_{i-1} and f_{i-1} are removed from the tour and the edges e_i and g_i are added to the tour, i. e., for $i \in [k-1]$, e_i denotes an edge added to the tour in step S_i and removed from the tour in step S_{i+1} . These definitions are illustrated in Figure 4.4.

Definition 4.11 (witness sequences of type 1). *If for every $i \in [k]$, the edge e_i does not occur in any step S_j with $j < i$, then S_1, \dots, S_k is called a k -witness sequence of type 1.*

Intuitively, witness sequences of type 1 possess enough randomness as every step introduces an edge that has not been seen before. Based on this observation, we prove in Lemma 4.14 the desired bound of $(\varepsilon\phi)^k$ on the probability that every step is an improvement by at most ε for these sequences.

Definition 4.12 (witness sequences of type 2). *Assume that for every $i \in [k]$, the edge e_i does not occur in any step S_j with $j < i$. If both endpoints of f_{k-1} occur in steps S_j with $j < k$, then S_1, \dots, S_k is called a k -witness sequence of type 2.*

Also for witness sequences of type 2, we obtain the desired bound of $(\varepsilon\phi)^k$ on the probability that every step is an improvement by at most ε . Due to the additional restriction on f_{k-1} , there are less than $4^{k-1}m^{k+1}$ witness sequences of type 2. As the two endpoints of f_{k-1} must be chosen among those vertices that occur in steps S_j with $j < k$, there are only $O(k^2)$ choices for the last step S_k . This implies that the number of k -witness sequences of type 2 can be upper bounded by $O(4^k k^2 m^k)$.

Definition 4.13 (witness sequences of type 3). *Assume that for every $i \in [k-1]$, the edge e_i does not occur in any step S_j with $j < i$. If the edges e_k and g_k occur in steps S_j with $j < k$ and if f_{k-1} does not occur in any step S_j with $j < k$, then S_1, \dots, S_k is called a k -witness sequence of type 3.*

Also witness sequences of type 3 possess enough randomness to bound the probability that every step is an improvement by at most ε by $(\varepsilon\phi)^k$ as also the last step introduces a new edge, namely f_{k-1} .

4.3.2 Improvement made by witness sequences

In this section, we analyze the probability that there exists a k -witness sequence in which every step is an improvement by at most ε .

Lemma 4.14. *The probability that there exists a k -witness sequence in which every step is an improvement by at most ε*

- a) *is bounded from above by $4^{k-1}m^{k+1}(\varepsilon\phi)^k$ for k -witness sequences of type 1.*
- b) *is bounded from above by $k^24^{k-1}m^k(\varepsilon\phi)^k$ for k -witness sequences of type 2.*
- c) *is bounded from above by $k^24^k m^k(\varepsilon\phi)^k$ for k -witness sequences of type 3.*

Proof. We use a union bound to estimate the probability that there exists a witness sequence in which every step is a small improvement.

a) We consider k -witness sequences of type 1 first. As already mentioned in the previous section, the number of such sequences is bounded by $4^{k-1}m^{k+1}$ as there are at most m^2 choices for the first step S_1 , and once S_i is fixed, there are at most $4m$ choices for step S_{i+1} . The number $4m$ follows since if S_i is fixed, there are two choices for the edge added to the tour in step S_i and removed from the tour in step S_{i+1} , there are at most m choices for the other edge removed in step S_{i+1} , and once these edges are determined, there are two possible 2-Opt steps in which these edges are removed from the tour.

Now fix an arbitrary k -witness sequence S_1, \dots, S_k of type 1. We use the same notations as in Figure 4.4 to denote the edges involved in this sequence. In the first step, the edges e_0 and f_0 are replaced by the edges e_1 and g_1 . We assume that the lengths of the edges e_0 , f_0 , and g_1 are determined by an adversary. The improvement of step S_1 can be expressed as a simple linear combination of the lengths of the involved edges. Hence, for fixed lengths of e_0 , f_0 , and g_1 , the event that S_1 is an improvement by at most ε corresponds to the event that the length $d(e_1)$ of e_1 lies in some fixed interval of length ε . Since the density of $d(e_1)$ is bounded by ϕ , the probability that $d(e_1)$ takes a value in the given interval is bounded by $\varepsilon\phi$. Now consider a step S_i and assume that arbitrary lengths for the edges e_j and f_j with $j < i$ and for g_j with $j \leq i$ are chosen. Since the edge e_i is not involved in any step S_j with $j < i$, its length is not determined. Hence, analogously to the first step, the probability that step S_i is an improvement by at most ε is bounded by $\varepsilon\phi$ independent of the improvements of the steps S_j with $j < i$. Applying this argument to every step S_i yields the desired bound of $(\varepsilon\phi)^k$.

b) In witness sequences of type 2, there are at most m^2 choices for step S_1 . Analogously to witness sequences of type 1, the number of possible choices for S_i with $1 < i < k-1$ is at most $4m$. The number of different vertices involved in steps S_j with $j < k$ is at most $4 + 2(k-2) = 2k$ as the first step introduces four new vertices and every other step at most 2. Since the endpoints of the edge f_{k-1} must be chosen among those vertices that have been involved in the steps S_j with $j < k$, there are at most $4k^2$ possible choices for step S_{k-1} . This implies that the number of different k -witness sequences of type 2 is bounded by $4 \cdot k^2 4^{k-2} m^k = k^2 4^{k-1} m^k$.

For a fixed witness sequence of type 2, applying the same arguments as for witness sequences of type 1, yields a probability of at most $(\varepsilon\phi)^k$ that every step is an improvement by at most ε .

c) The number of different edges involved in the steps S_i with $i < k$ is at most $4 + 3(k - 2) < 3k$. Hence, the number of k -witness sequences of type 3 is bounded by $9k^2 4^{k-2} m^k < k^2 4^k m^k$. Furthermore, similar to witness sequences of type 1, we can bound the probability that a fixed k -witness sequence of type 3 consists only of improvements by at most ε by $(\varepsilon\phi)^k$ since the last step introduces an edge which does not occur in the steps S_i with $i < k$, namely f_{k-1} . \square

Definition 4.15. *In the following, we use the term k -witness sequence to denote a k -witness sequence of type 1 or an i -witness sequence of type 2 or 3 with $i \leq k$. We call a k -witness sequence improving if every 2-change in the sequence is an improvement. Moreover, by $\Delta_{\text{ws}}^{(k)}$ we denote the smallest total improvement made by any improving k -witness sequence.*

Due to Lemma 4.14, it is unlikely that there exists an improving witness sequence whose total improvement is small.

Corollary 4.16. *For $k \in \mathbb{N}$ and $0 < \varepsilon \leq (4m^{(k-1)/(k-2)}\phi)^{-1}$,*

$$\Pr \left[\Delta_{\text{ws}}^{(k)} \leq \varepsilon \right] \leq 64k^3 (m\varepsilon\phi)^2 .$$

Proof. Due to Lemma 4.14 and the fact that witness sequences of type 2 or 3 must consist of at least two steps, applying a union bound yields the following bound on the probability that there exists an improving k -witness sequence whose total improvement is at most ε :

$$\begin{aligned} \Pr \left[\Delta_{\text{ws}}^{(k)} \leq \varepsilon \right] &\leq 4^{k-1} m^{k+1} (\varepsilon\phi)^k + \sum_{i=2}^k i^2 4^{i-1} m^i (\varepsilon\phi)^i + \sum_{i=2}^k i^2 4^i m^i (\varepsilon\phi)^i \\ &\leq 4^{k-1} m^{k+1} (\varepsilon\phi)^k + 2k^2 \sum_{i=2}^k (4m\varepsilon\phi)^i . \end{aligned}$$

Since $4m\varepsilon\phi < 1$, we can bound the sum by

$$\Pr \left[\Delta_{\text{ws}}^{(k)} \leq \varepsilon \right] \leq 4^{k-1} m^{k+1} (\varepsilon\phi)^k + 2k^3 (4m\varepsilon\phi)^2 ,$$

which implies the corollary because for $\varepsilon \leq (4m^{(k-1)/(k-2)}\phi)^{-1}$, the second term in the sum is at least as large as first one. \square

4.3.3 Finding witness sequences

In the previous section, we have shown an upper bound on the probability that there exists an improving k -witness sequence whose total improvement is small. In this section, we show that in every long enough sequence of consecutive 2-changes, one can identify a certain number of disjoint k -witness sequences. This way, we obtain a lower bound on the improvement made by any long enough sequence of consecutive 2-changes in terms of $\Delta_{\text{ws}}^{(k)}$.

Lemma 4.17. *Let $k \in \mathbb{N}$, and let S_1, \dots, S_t denote a sequence of consecutive 2-changes performed by the 2-Opt heuristic with $t > n4^{k+1}$. The sequence S_1, \dots, S_t shortens the tour by at least $t/4^{k+3} \cdot \Delta_{\text{ws}}^{(k)}$.*

Basically, we have to show that one can find $t/4^{k+3}$ disjoint k -witness sequences in the given sequence of consecutive 2-changes. Therefore, we first introduce a so-called *witness DAG* (directed acyclic graph) which represents the sequence S_1, \dots, S_t of 2-changes. In order to not confuse the constructed witness DAG W with the input graph G , we use the terms *nodes* and *arcs* when referring to the DAG W and the terms *vertices* and *edges* when referring to G . Nodes of W correspond to edges of G combined with a time stamp. The construction is started by adding the edges of the initial tour as nodes into W . These nodes get the time stamps $1, \dots, n$ in an arbitrary order. Then the sequence S_1, \dots, S_t is processed step by step. Assume that the steps S_1, \dots, S_{i-1} have already been processed and that step S_i is to be processed next. Furthermore, assume that in step S_i the edges e_{i-1} and f_{i-1} are exchanged with the edges e_i and g_i . Since the edges e_{i-1} and f_{i-1} are contained in the tour after the steps S_1, \dots, S_{i-1} , there are nodes in W corresponding to these edges. Let u_1 and u_2 denote the nodes with the most recent time stamps corresponding to e_{i-1} and f_{i-1} , respectively. We create two new nodes u_3 and u_4 corresponding to the edges e_i and g_i , each with time stamp $n+i$. Finally, four new arcs are added to W , namely the arcs (u_1, u_3) , (u_1, u_4) , (u_2, u_3) , and (u_2, u_4) . We refer to these four arcs as *twin arcs*. Observe that each node in W has indegree and outdegree at most 2. We call the resulting DAG W a *t -witness DAG*.

By the *height* of a node u , we denote the length of a shortest path from u to a leaf of W . After the witness DAG has been completely constructed, we associate with each node u with height at least k a sub-DAG of W . The sub-DAG W_u associated with such a node u is the induced sub-DAG of those nodes of W that can be reached from u by traversing at most k arcs. The following two lemmas imply Lemma 4.17.

Lemma 4.18. *For every sub-DAG W_u , the 2-changes represented by the arcs in W_u and their twin arcs yield a total improvement of at least $\Delta_{\text{ws}}^{(k)}$.*

Lemma 4.19. *For $t > n4^{k+1}$, every t -witness DAG contains at least $t/4^{k+2}$ nodes u whose corresponding sub-DAGs W_u are pairwise disjoint.*

Proof of Lemma 4.18. Assume that a sub-DAG W_u with root u is given. Since node u has height k , one can identify 2^{k-1} distinct sequences of linked 2-changes of length k in the sub-DAG W_u . In the following, we show that at least one of these sequences is a k -witness sequence or a sequence whose total improvement is as large as the total improvement of one of the k -witness sequences. We give a recursive algorithm *Sequ* that constructs such a sequence step by step. It is initialized with the sequence which consists only of the first step S_1 that is represented by the two outgoing arcs of the root u and their twin arcs.

Assume that *Sequ* is called with a sequence of steps S_1, \dots, S_i that has been constructed so far. Given this sequence, it has to decide if the sequence is continued with a step S_{i+1} such that S_i and S_{i+1} are linked or if the construction is stopped since a k -witness sequence is found. In Figure 4.5, we summarize the notations that we use in the following. In step S_j for $j \leq i-1$ and $j = i+1$, the edges e_{j-1} and f_{j-1} are




Figure 4.5: Construction of a path in the witness DAG: The path has been constructed up to step S_i and now it has to be decided whether to continue it along e_i or e'_i .

exchanged with the edges e_j and g_j . In step S_i , the edges e_{i-1} and f_{i-1} are exchanged with the edges e_i and e'_i , and in step S'_{i+1} , the edges e'_i and f'_i are exchanged with the edges e'_{i+1} and g'_{i+1} . We denote by E_i all edges that are involved in steps S_j with $j \leq i$. Similarly, by E_{i-1} we denote all edges that are involved in steps S_j with $j \leq i-1$.

Our construction ensures that whenever `Sequ` is called with a sequence S_1, \dots, S_i as input, then at least one of the edges that is added to the tour in step S_i is not contained in E_{i-1} . In the following, assume without loss of generality that $e_i \notin E_{i-1}$. When we call the algorithm recursively with the sequence S_1, \dots, S_{i+1} or with the sequence $S_1, \dots, S_i, S'_{i+1}$, then either the recursive call never gives back a return value since a witness sequence is found in the recursive call, which immediately stops the construction, or a 2-change S is returned. Whenever a 2-change S is returned, the meaning is as follows: There exists a sequence of linked 2-changes in the sub-DAG W_u starting with S_{i+1} or S'_{i+1} , respectively, whose net effect equals the 2-change S . That is, after all steps in the sequence have been performed, the same two edges as in S are removed from the tour, the same two edges are added to the tour, and all other edges either stay in or out of the tour. In this case, we can virtually replace step S_{i+1} or S'_{i+1} , respectively, by the new step S .

When `Sequ` is called with the sequence S_1, \dots, S_i , then it first identifies the steps S_{i+1} and S'_{i+1} based on the last step S_i . If $i = k$, then S_1, \dots, S_i is a k -witness sequence of type 1, and `Sequ` stops. Otherwise, the following steps are performed, where we assume that whenever `Sequ` has identified a witness sequence, it immediately stops the construction.

1. **TYPE 2 SEQUENCE:** If $f_{i-1} \in E_{i-1}$, then S_1, \dots, S_i is a witness sequence of type 2 because we assumed $e_i \notin E_{i-1}$.
2. **CONTINUATION:** If $e_{i+1} \notin E_i$ or $g_{i+1} \notin E_i$, then call `Sequ` recursively with `Sequ`(S_1, \dots, S_{i+1}). If $e'_i \notin E_{i-1}$ and ($e'_{i+1} \notin E_i$ or $g'_{i+1} \notin E_i$), then call `Sequ` recursively with `Sequ`($S_1, \dots, S_i, S'_{i+1}$).

If in one of the recursive calls a step S is returned, which happens only in case 3 (c), then replace the corresponding step S_{i+1} or S'_{i+1} virtually by the returned step. That is, in the following steps of the algorithm, assume that S_{i+1} or S'_{i+1}

equals step S . The algorithm ensures that the edges that are added to the tour in the new step S are always chosen from the set E_i .

3. NO CONTINUATION I: $e'_i \in E_{i-1}$ and $e_{i+1}, g_{i+1} \in E_i$ and $f_{i-1} \notin E_{i-1}$
 - (a) If $f_i \notin E_i$, then S_1, \dots, S_{i+1} is a witness sequence of type 3.
 - (b) If $e_{i+1}, g_{i+1} \in E_{i-1}$, then S_1, \dots, S_i is a witness sequence of type 2 since one endpoint of f_{i-1} equals one endpoint of e'_i and the other one equals one endpoint of either e_{i+1} or g_{i+1} .
 - (c) If $f_i \in E_i$ and ($e_{i+1} \in E_i \setminus E_{i-1}$ or $g_{i+1} \in E_i \setminus E_{i-1}$), then one can assume w.l.o.g. that $g_{i+1} = f_{i-1}$ and $e_{i+1} \in E_{i-1}$ since $e_{i+1} \neq e'_i$ and $g_{i+1} \neq e'_i$ (e_{i+1} and g_{i+1} share one endpoint with e_i ; e'_i does not share an endpoint with e_i .) In this case, return the step $S = (e_{i-1}, f_i) \rightarrow (e_{i+1}, e'_i)$. Observe that $e_{i+1}, e'_i \in E_{i-1}$, as desired.
4. NO CONTINUATION II: $e'_i \notin E_{i-1}$ and $e_{i+1}, g_{i+1}, e'_{i+1}, g'_{i+1} \in E_i$ and $f_{i-1} \notin E_{i-1}$
 - (a) If $e_{i+1}, g_{i+1}, e'_{i+1}, g'_{i+1} \in E_{i-1}$, then S_1, \dots, S_i is a witness sequence of type 2.
 - (b) If $f_i \notin E_i$, then S_1, \dots, S_{i+1} is a witness sequence of type 3.
 - (c) If $f'_i \notin E_i$, then $S_1, \dots, S_i, S'_{i+1}$ is a witness sequence of type 3.
 - (d) If $f_i, f'_i \in E_i$ and ($e_{i+1} \in E_i \setminus E_{i-1}$ or $g_{i+1} \in E_i \setminus E_{i-1}$) and ($e'_{i+1} \in E_i \setminus E_{i-1}$ or $g'_{i+1} \in E_i \setminus E_{i-1}$), then as in case 3 (c), assume w.l.o.g. $g_{i+1} = g'_{i+1} = f_{i-1}$ and $e_{i+1}, e'_{i+1} \in E_{i-1}$. In this case, it must be $f_i \neq e'_i$ and $f'_i \neq e_i$ as otherwise step S_i would be reversed in step S_{i+1} or S'_{i+1} . Hence, $f_i, f'_i \in E_{i-1}$, and S_1, \dots, S_i is a witness sequence of type 2 since one endpoint of f_{i-1} equals one endpoint of f_i and the other endpoint equals one endpoint of f'_i .
 - (e) If $|\{e_{i+1}, e'_{i+1}, g_{i+1}, g'_{i+1}\} \cap (E_i \setminus E_{i-1})| = 1$, assume w.l.o.g. $e_{i+1}, g_{i+1}, e'_{i+1} \in E_{i-1}$ and $g'_{i+1} = f_{i-1}$. As in the previous case, it must $f'_i \in E_{i-1}$. For the step $S = (e_{i-1}, f'_i) \rightarrow (e_i, e'_{i+1})$, the sequence S_1, \dots, S_{i-1}, S is a witness sequence of type 2 as $f'_i \in E_{i-1}$. Observe that the original sequence S_1, \dots, S_i together with the step S'_{i+1} yields the same net effect and hence the same improvement as the sequence S_1, \dots, S_{i-1}, S .

Observe that basically **Sequ** just constructs a path through the DAG starting at node u . When a path corresponding to the sequence S_1, \dots, S_i of 2-changes has been constructed, **Sequ** decides to either stop the construction since a witness sequence has been found, or, if possible, to continue the path with an arc corresponding to a step S_{i+1} or S'_{i+1} . In some situations, it can happen that **Sequ** has not found a witness sequence yet but cannot continue the construction. In such cases, step S_i is pruned and **Sequ** reconsiders the path S_1, \dots, S_{i-1} . Based on the pruned step S_i , it can then either decide that a witness sequence has been found, that also S_{i-1} has to be pruned, or it can decide to continue the path with S'_i instead of S_i .

This concludes the proof as the presented algorithm always identifies a k -witness sequence whose total improvement is at most as large as the improvement made by the steps in the sub-DAG W_u . \square

Proof of Lemma 4.19. A t -witness DAG W consists of $n+2t$ nodes and n of these nodes are leaves. Since the indegree and the outdegree of every node is bounded by 2, there are at most $n2^k$ nodes in W whose height is less than k . Hence, there are at least $n+2t-n2^k \geq t$ nodes in W with an associated sub-DAG. We construct a set of disjoint sub-DAGs in a greedy fashion: We take an arbitrary sub-DAG W_u and add it to the set of disjoint sub-DAGs that we construct. After that, we remove all nodes, arcs, and twin arcs of W_u from the DAG W . We repeat these steps until no sub-DAG W_u is left in W .

In order to see that the constructed set consists of at least $t/4^{k+2}$ disjoint sub-DAGs, observe that each sub-DAG consists of at most $2^{k+1} - 1$ nodes as its height is k . Hence, it contains at most $2^k - 1$ pairs of twin arcs, and there are at most $2^{k+2} - 4$ arcs that belong to the sub-DAG or that have a twin arc belonging to the sub-DAG. Furthermore, observe that each of these arcs can be contained in at most $2^k - 1$ sub-DAGs. Hence, every sub-DAG W_u can only be non-disjoint from at most $2^{2k+2} = 4^{k+1}$ other sub-DAGs. Thus, the number of disjoint sub-DAGs must be at least $\lfloor t/4^{k+1} \rfloor > t/4^{k+2}$, where the last inequality follows because we assumed $t > n4^{k+1}$. \square

4.3.4 The expected number of 2-changes

Now we can prove Theorem 1.2.

Proof of Theorem 1.2 c). We combine Corollary 4.16 and Lemma 4.17 to obtain an upper bound on the probability that the length T of the longest path in the state graph exceeds t . For $t > n4^{k+1}$, the tour is shortened by the sequence of 2-changes by at least $t/4^{k+2} \cdot \Delta_{\text{ws}}^{(k)}$. Hence, for $t > n4^{k+1}$,

$$\Pr [T \geq t] \leq \Pr \left[\frac{t}{4^{k+2}} \cdot \Delta_{\text{ws}}^{(k)} \leq n \right] = \Pr \left[\Delta_{\text{ws}}^{(k)} \leq \frac{n \cdot 4^{k+2}}{t} \right].$$

Combining this inequality with Corollary 4.16 yields for $t \geq 4^{k+3} \cdot n \cdot \phi \cdot m^{(k-1)/(k-2)}$,

$$\Pr [T \geq t] \leq 64k^3 \left(\frac{4^{k+2} \cdot n \cdot m \cdot \phi}{t} \right)^2.$$

Hence, we can bound the expected number of 2-changes by

$$\mathbf{E} [T] \leq 4^{k+3} \cdot n \cdot \phi \cdot m^{(k-1)/(k-2)} + \sum_{t=1}^{n!} \min \left\{ 64k^3 \left(\frac{4^{k+2} \cdot n \cdot m \cdot \phi}{t} \right)^2, 1 \right\}.$$

Splitting the sum at $t = k^{3/2} \cdot 4^{k+2} \cdot n \cdot m \cdot \phi$ yields

$$\mathbf{E} [T] = O \left(k^{3/2} \cdot 4^k \cdot n \cdot m^{(k-1)/(k-2)} \cdot \phi \right).$$

Setting $k = \sqrt{\log m}$ yields the theorem. \square

5 Expected Approximation Ratio

In this section, we consider the expected approximation ratio of the solution found by 2-Opt on ϕ -perturbed L_p instances. Chandra, Karloff, and Tovey [CKT99] show that if one has a set of n points in the unit hypercube $[0, 1]^d$ and the distances are measured according to a metric that is induced by a norm, then every locally optimal solution has length at most $c \cdot n^{(d-1)/d}$ for an appropriate constant c depending on the dimension d and the metric. Hence, it follows for every L_p metric that 2-Opt yields a tour of length at most $O(n^{(d-1)/d})$ on ϕ -perturbed L_p instances. This implies that, in order to bound the expected approximation ratio of 2-Opt on these instances, we just need to upper bound the expected value of $1/\text{Opt}$, where Opt denotes the length of the shortest tour.

Lemma 5.1. *Let $p \in \mathbb{N} \cup \{\infty\}$. For ϕ -perturbed L_p instances with n points, it holds*

$$\mathbf{E} \left[\frac{1}{\text{Opt}} \right] = O \left(\frac{\sqrt[d]{\phi}}{n^{(d-1)/d}} \right) .$$

Proof. Let $v_1, \dots, v_n \in \mathbb{R}^d$ denote the points of the ϕ -perturbed instance. We partition the unit hypercube into $k = \lceil n\phi \rceil$ smaller hypercubes with volume $1/k$ each and analyze how many of these smaller hypercubes contain at least one of the points. Assume that X of these hypercubes contain a point, then the optimal tour must have length at least $X/(3^d \sqrt[d]{k})$. In order to see this, we construct a set $P \subseteq \{v_1, \dots, v_n\}$ of points as follows: Consider the points v_1, \dots, v_n one after another, and insert a point v_i into P if P does not contain a point in the same hypercube as v_i or in one of its $3^d - 1$ neighboring hypercubes yet. Due to the triangle inequality, the optimal tour on P is at most as long as the optimal tour on v_1, \dots, v_n . Furthermore, P contains at least $X/3^d$ points and every edge between two points from P has length at least $1/\sqrt[d]{k}$ since P does not contain two points in the same or in two neighboring hypercubes.

Hence, it remains to analyze the random variable X . For each hypercube i with $1 \leq i \leq k$, we define a random variables X_i which takes value 0 if hypercube i is empty and value 1 if hypercube i contains at least one point. The density functions that specify the locations of the points induce for each pair of hypercube i and point j a probability p_i^j such that point j falls into hypercube i with probability p_i^j . Hence, one can think of throwing n balls into k bins in a setting where each ball has its own probability distribution over the bins. Due to the bounded density, we have $p_i^j \leq \phi/k$. For each hypercube i , let M_i denote the probability mass associated with the hypercube, that is

$$M_i = \sum_{j=1}^n p_i^j \leq \frac{n\phi}{k} .$$

We can write the expected value of the random variable X_i as

$$\mathbf{E}[X_i] = \Pr[X_i = 1] = 1 - \prod_{j=1}^n (1 - p_i^j) \geq 1 - \left(1 - \frac{M_i}{n}\right)^n$$

as, under the constraint $\sum_j (1 - p_i^j) = n - M_i$, the term $\prod_j (1 - p_i^j)$ is maximized if all

p_i^j are equal. Due to linearity of expectation, the expected value of X is

$$\mathbf{E}[X] \geq \sum_{i=1}^k 1 - \left(1 - \frac{M_i}{n}\right)^n = k - \sum_{i=1}^k \left(1 - \frac{M_i}{n}\right)^n .$$

Observe that $\sum_i M_i = n$. Thus, the sum $\sum_i (1 - M_i/n)$ becomes maximal if the M_i 's are chosen as unbalanced as possible. Hence, we assume that $\lceil k/\phi \rceil$ of the M_i 's take their maximal value of $n\phi/k$ and the other M_i 's are zero. This yields, for sufficiently large n ,

$$\begin{aligned} \mathbf{E}[X] &\geq k - \left(\left\lceil \frac{k}{\phi} \right\rceil \left(1 - \frac{\phi}{k}\right)^n + \left(k - \left\lceil \frac{k}{\phi} \right\rceil\right) \right) \\ &\geq \frac{k}{\phi} - \frac{2k}{\phi} \left(1 - \frac{\phi}{k}\right)^n \\ &\geq \frac{k}{\phi} \left(1 - 2 \left(1 - \frac{1}{n+1}\right)^n\right) \geq \frac{n}{5} . \end{aligned}$$

Hence, we obtain the following bound on the expected length of the optimal tour

$$\mathbf{E}[\text{Opt}] \geq \frac{\mathbf{E}[X]}{3^d \sqrt[d]{k}} \geq \frac{n}{5 \cdot 3^d \sqrt[d]{k}} \geq \frac{n^{(d-1)/d}}{5 \cdot 3^d \sqrt[d]{\phi + 1}} .$$

We still need to determine the expected value of the random variable $1/\text{Opt}$. Therefore, we first show that X is sharply concentrated around its mean value. The random variable X is the sum of n 0-1-random variables. If these random variables were independent, we could simply use a Chernoff bound to bound the probability that X takes a value that is smaller than its mean value. The X_i 's are negatively associated, in the sense that whenever we already know that some of the X_i 's are zero, then the probability of the event that another X_i also takes the value zero becomes smaller. Hence, intuitively, the dependencies can only help to bound the probability that X takes a value smaller than its mean value. Dubhashi and Ranjan [DR98] formalize this intuition by introducing the notion of negative dependence and by showing that in the case of negative dependent random variables, one can still apply a Chernoff bound. This yields

$$\Pr \left[X \leq \frac{n}{10} \right] \leq e^{-n/40} .$$

Thus, as $1/X \leq 1$ with certainty, for sufficiently large n ,

$$\mathbf{E} \left[\frac{1}{X} \right] \leq (1 - e^{-n/40}) \cdot \frac{10}{n} + e^{-n/40} \leq \frac{11}{n}$$

Altogether, this implies

$$\mathbf{E} \left[\frac{1}{\text{Opt}} \right] \leq \mathbf{E} \left[\frac{3^d \cdot \sqrt[d]{\lceil n\phi \rceil}}{X} \right] = O \left(\frac{\sqrt[d]{\phi}}{n^{(d-1)/d}} \right) . \quad \square$$

If one combines Lemma 5.1 with the result of Chandra, Karloff, and Tovey that every locally optimal solution has length $O(n^{(d-1)/d})$, one obtains Theorem 1.4.

6 Smoothed Analysis

Smoothed Analysis was introduced by Spielman and Teng [ST04] as a hybrid of worst case and average case analysis. The semi-random input model in a smoothed analysis is designed to capture the behavior of algorithms on typical inputs better than a worst case or average case analysis alone as it allows an adversary to specify an arbitrary input which is randomly perturbed afterwards. In Spielman and Teng's analysis of the Simplex algorithm the adversary specifies an arbitrary linear program which is perturbed by adding independent Gaussian random variables to each number in the linear program. Our probabilistic analysis of Manhattan and Euclidean instances can also be seen as a smoothed analysis in which an adversary can choose the distributions for the points over the unit hypercube. The adversary is restricted to distributions that can be represented by densities that are bounded by ϕ . Our model cannot handle Gaussian perturbations directly because the support of Gaussian random variables is not bounded.

Assume that every point v_1, \dots, v_n is described by a density whose support is restricted to the hypercube $[-\alpha, \alpha]^d$, for some $\alpha \geq 1$. Then after appropriate scaling and translating, we can assume that all supports are restricted to the unit hypercube $[0, 1]^d$. Thereby, the maximal density ϕ increases by at most a factor of $(2\alpha)^d$. Hence, after appropriate scaling and translating, Theorems 1.2, 1.3, and 1.4 can still be applied if one takes into account the increased densities.

One possibility to cope with Gaussian perturbations is to consider *truncated Gaussian perturbations*. In such a perturbation model, the coordinates of each point are initially chosen from $[0, 1]^d$ and then perturbed by adding Gaussian random variables with some standard deviation σ to them that are conditioned to lie in $[-\alpha, \alpha]$ for some $\alpha \geq 1$. The maximal density of such truncated Gaussian random variables for $\sigma \leq 1$ is bounded from above by

$$\frac{1/(\sigma\sqrt{2\pi})}{1 - \sigma \cdot \exp(-\alpha^2/(2\sigma^2))} .$$

After such a truncated perturbation, all points lie in the hypercube $[-\alpha, 1 + \alpha]^d$. Hence, one can apply Theorems 1.2, 1.3, and 1.4 with

$$\phi = \frac{(2\alpha + 1)^d}{(\sigma\sqrt{2\pi} - \sigma^2\sqrt{2\pi}\exp(-\alpha^2/(2\sigma^2)))^d} = O\left(\frac{\alpha^d}{\sigma^d}\right) .$$

It is not necessary to truncate the Gaussian random variables if the standard deviation is small enough. For $\sigma \leq \min\{\alpha/\sqrt{2(n+1)\ln n + 2\ln d}, 1\}$, the probability that one of the Gaussian random variables has an absolute value larger than $\alpha \geq 1$ is bounded from above by n^{-n} . In this case, even if one does not truncate the random variables, Theorems 1.2, 1.3, and 1.4 can be applied with $\phi = O(\alpha^d/\sigma^d)$. To see this, it suffices to observe that the worst-case bound for the number of 2-changes is $(n!)$ and the worst-case approximation ratio is $O(\log n)$ [CKT99]. Multiplying these values with the failure probability of n^{-n} constitutes less than 1 to the expected values. In particular, this implies that the expected length of the longest path in the state graph is bounded by $O(\text{poly}(n, 1/\sigma))$.

7 Conclusions and Open Problems

We have shown several new results on the running time and the approximation ratio of the 2-Opt heuristic. However, there is still a variety of open problems regarding this algorithm. Our lower bounds only show that there exist families of instances on which 2-Opt takes an exponential number of steps if it uses a particular pivot rule. It would be interesting to analyze the diameter of the state graph and to either present instances on which every pivot rule needs an exponential number of steps or to prove that there is always an improvement sequence of polynomial length to a locally optimal solution. Also the worst number of local improvements for some natural pivot rules like, e.g., the one that always makes the largest possible improvement or the one that always chooses a random improving 2-change, is not known yet. Furthermore, the complexity of computing locally optimal solutions is open. The only result in this regard is due to Krentel [Kre89] who shows that it is PLS-complete to compute a local optimum for the metric TSP for k -Opt for some constant k . It is not known whether his construction can be embedded into the Euclidean metric and whether it is PLS-complete to compute locally optimal solutions for 2-Opt. Fischer and Torenvliet [FT95] show, however, that for the general TSP, it is PSPACE-hard to compute a local optimum for 2-Opt that is reachable from a given initial tour.

The obvious open question concerning the probabilistic analysis is how the gap between experiments and theory can be narrowed further. In order to tackle this question, new methods seem to be necessary. Our approach, which is solely based on analyzing the smallest improvement made by a sequence of linked 2-changes, seems to yield too pessimistic bounds. Another interesting area to explore is the expected approximation ratio of 2-Opt. In experiments, approximation ratios close to 1 are observed. For instances that are chosen uniformly at random, the bound on the expected approximation ratio is a constant but unfortunately a large one. It seems to be a very challenging problem to improve this constant to a value that matches the experimental results.

Besides 2-Opt, there are also other local search algorithms that are successful for the traveling salesperson problem. In particular, the Lin-Kernighan heuristic [LK73] is one of the most successful local search algorithm for the symmetric TSP. It is a variant of k -Opt in which k is not fixed and it can roughly be described as follows: Each local modification starts by removing one edge $\{a, b\}$ from the current tour, which results in a Hamiltonian path with the two endpoints a and b . Then an edge $\{b, c\}$ is added, which forms a cycle; there is a unique edge $\{c, d\}$ incident to c whose removal breaks the cycle, producing a new Hamiltonian path with endpoints a and d . This operation is called a rotation. Now either a new Hamiltonian cycle can be obtained by adding the edge $\{a, d\}$ to the tour or another rotation can be performed. There are a lot of different variants and heuristic improvements of this basic scheme, but little is known theoretically. Papadimitriou [Pap92] shows for a variant of the Lin-Kernighan heuristic that computing a local optimum is PLS-complete, which is a sharp contrast to the experimental results. Since the Lin-Kernighan heuristic is widely used in practice, a theoretical explanation for its good behavior in practice is of great interest. Our analysis of 2-Opt relies crucially on the fact that there is only a polynomial number of different 2-changes. For the Lin-Kernighan heuristic, however, the number of different local

improvements is exponential. Hence, it is an interesting question whether nonetheless the smallest possible improvement is polynomially large or whether different methods yield a polynomial upper bound on the expected running time of the Lin-Kernighan heuristic.

References

- [Aro98] Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.
- [CKT99] Barun Chandra, Howard J. Karloff, and Craig A. Tovey. New results on the old k-Opt algorithm for the traveling salesman problem. *SIAM Journal on Computing*, 28(6):1998–2029, 1999.
- [DR98] Devdatt P. Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *Random Structures and Algorithms*, 13(2):99–124, 1998.
- [FT95] Sophie Fischer and Leen Torenvliet. The malleability of TSP_{2opt} . In *Proc. of the 21st Int. Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 152–166, 1995.
- [JM97] David S. Johnson and Lyle A. McGeoch. The traveling salesman problem: A case study in local optimization. In E. H. L. Aarts and J. K. Lenstra, editors, *Local Search in Combinatorial Optimization*. John Wiley and Sons, 1997.
- [Ker89] Walter Kern. A probabilistic analysis of the switching algorithm for the Euclidean TSP. *Mathematical Programming*, 44(2):213–219, 1989.
- [Kre89] Mark W. Krentel. Structure in locally optimal solutions. In *Proc. of the 30th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 216–221, 1989.
- [LK73] Shen Lin and Brian W. Kernighan. An effective heuristic for the traveling salesman problem. *Operations Research*, 21:489–516, 1973.
- [Lue75] George S. Lueker. Unpublished manuscript, 1975. Princeton University.
- [Mit99] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999.
- [Pap77] Christos H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. *Theoretical Computer Science*, 4(3):237–244, 1977.
- [Pap92] Christos H. Papadimitriou. The complexity of the lin-kernighan heuristic for the traveling salesman problem. *SIAM Journal on Computing*, 21(3):450–465, 1992.

- [Rei91] Gerhard Reinelt. TSPLIB – A traveling salesman problem library. *ORSA Journal on Computing*, 3(4):376–384, 1991.
- [RSI77] Daniel J. Rosenkrantz, Richard Edwin Stearns, and Philip M. Lewis II. An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing*, 6(3):563–581, 1977.
- [ST04] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.
- [vLS81] Jan van Leeuwen and Anneke A. Schoon. Untangling a traveling salesman tour in the plane. In *Proc. of the 7th Int. Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 87–98, 1981.

A Some Probability Theory

Lemma A.1. *Let X^1, \dots, X^n be independent d -dimensional random row vectors, and, for $i \in [n]$ and some $\phi \geq 1$, let $f_i: [0, 1]^d \rightarrow [0, \phi]$ denote the joint density of the entries of X^i . Furthermore, let $\lambda^1, \dots, \lambda^k \in \mathbb{Z}^{dn}$ be linearly independent row vectors. For $i \in [n]$ and a fixed $\varepsilon \geq 0$, we denote by \mathcal{A}_i the event that $\lambda^i \cdot X$ takes a value in the interval $[0, \varepsilon]$, where X denotes the vector $X = (X^1, \dots, X^n)^T$. Under these assumptions,*

$$\Pr \left[\bigcap_{i=1}^k \mathcal{A}_i \right] \leq (\varepsilon \phi)^k .$$

Proof. The main tool for proving the lemma is a change of variables. Instead of using the canonical basis of the dn -dimensional vector space \mathbb{R}^{dn} , we use the given linear combinations as basis vectors. To be more precise, the basis \mathcal{B} that we use consists of two parts: it contains the vectors $\lambda^1, \dots, \lambda^k$ and it is completed by some vectors from the canonical basis $\{e^1, \dots, e^{dn}\}$, where e^i denotes the i -th canonical row vector, i. e., $e^i_i = 1$ and $e^i_j = 0$ for $j \neq i$. That is, the basis \mathcal{B} can be written as $\{\lambda^1, \dots, \lambda^k, e^{\pi(1)}, \dots, e^{\pi(dn-k)}\}$, for some injective function $\pi: [dn - k] \rightarrow [dn]$.

Let $\Phi: \mathbb{R}^{dn} \rightarrow \mathbb{R}^{dn}$ be defined by $\Phi(x) = Ax$, where A denotes the $(dn) \times (dn)$ -matrix

$$\begin{pmatrix} \lambda^1 \\ \vdots \\ \lambda^k \\ e^{\pi(1)} \\ \vdots \\ e^{\pi(dn-k)} \end{pmatrix} .$$

Since \mathcal{B} is a basis of \mathbb{R}^{dn} , the function Φ is a bijection. We define $Y = (Y_1, \dots, Y_{dn})^T$ as $Y = \Phi(X)$, and for $i \in [n]$, we denote by Y^i the vector $(Y_{d(i-1)+1}, \dots, Y_{di})$. Let $f: \mathbb{R}^{dn} \rightarrow \mathbb{R}$ denote the joint density of the entries of the random vectors X^1, \dots, X^n , and let $g: \mathbb{R}^{dn} \rightarrow \mathbb{R}$ denote the joint density of the entries of the random vectors Y^1, \dots, Y^n . Due to the independence of the random vectors X^1, \dots, X^n , we have $f(x_1, \dots, x_{dn}) = f_1(x_1, \dots, x_d) \cdots f_n(x_{d(n-1)+1}, \dots, x_{dn})$. We can express the joint density g as

$$g(y_1, \dots, y_{dn}) = |\det_{\partial} \Phi^{-1}(y_1, \dots, y_{dn})| \cdot f(\Phi^{-1}(y_1, \dots, y_{dn})) ,$$

where \det_{∂} denotes the determinant of the Jacobian matrix of Φ^{-1} .

The matrix A is invertible as \mathcal{B} is a basis of \mathbb{R}^{dn} . Hence, for $y \in \mathbb{R}^{dn}$, $\Phi^{-1}(y) = A^{-1}y$ and the Jacobian matrix of Φ^{-1} equals A^{-1} . Thus, $\det_{\partial} \Phi^{-1} = \det A^{-1} = (\det A)^{-1}$. Since all entries of A are integers, also its determinant must be an integer, and since it has rank dn , we know that $\det A \neq 0$. Hence, $|\det A| \geq 1$ and $|\det A^{-1}| \leq 1$. For $y \in \mathbb{R}^{dn}$, we decompose $\Phi^{-1}(y) \in \mathbb{R}^{dn}$ into n subvectors with d entries each, i. e., $\Phi^{-1}(y) = (\Phi_1^{-1}(y), \dots, \Phi_n^{-1}(y))$ with $\Phi_i^{-1}(y) \in \mathbb{R}^d$ for $i \in [n]$. This yields

$$g(y) = |\det A^{-1}| \cdot f(\Phi^{-1}(y)) \leq f_1(\Phi_1^{-1}(y)) \cdots f_n(\Phi_n^{-1}(y)) .$$

The probability we want to estimate can be written as

$$\Pr \left[\bigcap_{i=1}^k \mathcal{A}_i \right] = \int_{y_1=0}^{\varepsilon} \cdots \int_{y_k=0}^{\varepsilon} \int_{y_{k+1}=-\infty}^{\infty} \cdots \int_{y_{dn}=-\infty}^{\infty} g(y_1, \dots, y_{dn}) dy_{dn} \cdots dy_1 . \quad (\text{A.1})$$

Since all entries of the vectors X^1, \dots, X^n take only values in the interval $[0, 1]$ and since for $i \in \{k+1, \dots, dn\}$, the random variable Y_i coincides with one of these entries, (A.1) simplifies to

$$\Pr \left[\bigcap_{i=1}^k \mathcal{A}_i \right] = \int_{y_1=0}^{\varepsilon} \cdots \int_{y_k=0}^{\varepsilon} \int_{y_{k+1}=0}^1 \cdots \int_{y_{dn}=0}^1 g(y_1, \dots, y_{dn}) dy_{dn} \cdots dy_1 . \quad (\text{A.2})$$

Without loss of generality, we assume that $\{i \mid \exists j \in [dn - k]: \pi(j) = i\} \subseteq [dk]$, i. e., only vectors e^i from the canonical basis with $i \leq dk$ are replaced by the vectors $\lambda^1, \dots, \lambda^k$ in the basis \mathcal{B} . Furthermore, we can assume w. l. o. g. $\pi(i) = i$, for $i > dk$. Under these assumptions, the density g can be upper bounded as follows:

$$g(y_1, \dots, y_{dn}) \leq \phi^k \cdot f_{k+1}(y_{dk+1}, \dots, y_{d(k+1)}) \cdots f_n(y_{d(n-1)+1}, \dots, y_{dn}) . \quad (\text{A.3})$$

Putting together (A.2) and (A.3) yields

$$\begin{aligned} \Pr \left[\bigcap_{i=1}^k \mathcal{A}_i \right] &\leq (\varepsilon\phi)^k \cdot \int_{y_{dk+1}=0}^1 \cdots \int_{y_{d(k+1)}=0}^1 f_{k+1}(y_{dk+1}, \dots, y_{d(k+1)}) \\ &\quad \cdots \int_{y_{d(n-1)+1}=0}^1 \int_{y_{dn}=0}^1 f_n(y_{d(n-1)+1}, \dots, y_{dn}) dy_{dn} \cdots dy_{dk+1} \\ &= (\varepsilon\phi)^k , \end{aligned}$$

where the last equation follows because f_{k+1}, \dots, f_n are density functions. \square

B Proofs of some Lemmas from Section 4.2

B.1 Proof of Lemma 4.7

Let $a, c \in (0, C]$ for some $C > 0$. In the following proof, we use the following two identities:

$$\int_0^c \frac{1}{\sqrt{z(c-z)}} dz = \pi$$

and

$$\begin{aligned} \int_0^a \frac{1}{\sqrt{z(z+c)}} dz &= \ln \left(\frac{c}{2} + a + \sqrt{a(a+c)} \right) + \ln \left(\frac{2}{c} \right) \\ &\leq \ln \left(\frac{3}{2}c + 2a \right) + \ln \left(\frac{2}{c} \right) \leq \ln(4C) + \ln \left(\frac{2}{c} \right) . \end{aligned}$$

Proof of Lemma 4.7. The conditional density of Δ can be calculated as convolution of the conditional densities of Z_1 and Z_2 as follows:

$$f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) = \int_{-\infty}^{\infty} f_{Z|T=\tau, R=r_1}(z) \cdot f_{Z|T=\tau, R=r_2}(z - \delta) dz .$$

In order to estimate this integral, we distinguish between several cases. In the following, let κ denote a sufficiently large constant.

First case: $\tau \leq r_1$ and $\tau \leq r_2$.

Since Z_i takes only values in the interval $[-\tau, \tau]$, we can assume $0 < \delta \leq \min\{1/2, 2\tau\}$ and

$$f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) = \int_{-\tau+\delta}^{\tau} f_{Z|T=\tau, R=r_1}(z) \cdot f_{Z|T=\tau, R=r_2}(z - \delta) dz .$$

Due to Lemma 4.6, we can estimate the densities of Z_1 and Z_2 by

$$f_{Z|T=\tau, R=r_i}(z) \leq \sqrt{\frac{2}{\tau^2 - z^2}} \leq \sqrt{\frac{2}{\tau(\tau - |z|)}} \leq \sqrt{\frac{2}{\tau}} \left(\frac{1}{\sqrt{\tau - z}} + \frac{1}{\sqrt{\tau + z}} \right) .$$

For $\delta \in (0, \min\{1/2, 2\tau\}]$, we obtain the following upper bound on the density of Δ :

$$\begin{aligned} & f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) \\ & \leq \frac{2}{\tau} \int_{-\tau+\delta}^{\tau} \left(\frac{1}{\sqrt{\tau - z}} + \frac{1}{\sqrt{\tau + z}} \right) \left(\frac{1}{\sqrt{\tau - z + \delta}} + \frac{1}{\sqrt{\tau + z - \delta}} \right) dz \\ & = \frac{2}{\tau} \left(\int_0^{2\tau-\delta} \frac{1}{\sqrt{z(z+\delta)}} dz + \int_{\delta}^{2\tau} \frac{1}{\sqrt{z(2\tau+\delta-z)}} dz \right. \\ & \quad \left. + \int_0^{2\tau-\delta} \frac{1}{\sqrt{z(2\tau-\delta-z)}} dz + \int_0^{2\tau-\delta} \frac{1}{\sqrt{z(z+\delta)}} dz \right) \\ & \leq \frac{2}{\tau} \left(2\pi + 2 \ln(8\sqrt{\delta}) + 2 \ln(2\delta^{-1}) \right) \leq \frac{\kappa}{\tau} \cdot \ln \delta^{-1} . \end{aligned}$$

Second case: $r_1 \leq \tau$ and $r_2 \leq \tau$.

Since Z_i takes only values in the interval $[-\tau, 2r_i - \tau]$, we can assume $0 < \delta \leq \min\{1/2, 2r_1\}$ and

$$f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) = \int_{-\tau+\delta}^{\min\{2r_1-\tau, 2r_2-\tau+\delta\}} f_{Z|T=\tau, R=r_1}(z) \cdot f_{Z|T=\tau, R=r_2}(z - \delta) dz .$$

Due to Lemma 4.6, we can estimate the densities of Z_1 and Z_2 by

$$\begin{aligned} f_{Z|T=\tau, R_i=r_i}(z) & \leq \sqrt{\frac{2}{(\tau+z)(2r_i-\tau-z)}} \leq \begin{cases} \sqrt{\frac{2}{r_i(\tau+z)}} & \text{if } z \leq r_i - \tau \\ \sqrt{\frac{2}{r_i(2r_i-\tau-z)}} & \text{if } z \geq r_i - \tau \end{cases} \\ & \leq \sqrt{\frac{2}{r_i}} \left(\frac{1}{\sqrt{\tau+z}} + \frac{1}{\sqrt{2r_i-\tau-z}} \right) . \end{aligned}$$

Case 2.1: $\delta \in (\max\{0, 2(r_1 - r_2)\}, 2r_1]$.

We obtain the following upper bound on the density of Δ :

$$\begin{aligned}
& f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) \\
& \leq \frac{2}{\sqrt{r_1 r_2}} \int_{-\tau+\delta}^{2r_1-\tau} \left(\frac{1}{\sqrt{\tau+z}} + \frac{1}{\sqrt{2r_1-\tau-z}} \right) \left(\frac{1}{\sqrt{\tau+z-\delta}} + \frac{1}{\sqrt{2r_2-\tau-z+\delta}} \right) dz \\
& = \frac{2}{\sqrt{r_1 r_2}} \left(\int_0^{2r_1-\delta} \frac{1}{\sqrt{z(z+\delta)}} dz + \int_0^{2r_1-\delta} \frac{1}{\sqrt{z(2r_1-\delta-z)}} dz \right. \\
& \quad \left. + \int_{\delta}^{2r_1} \frac{1}{\sqrt{z(2r_2+\delta-z)}} dz + \int_0^{2r_1-\delta} \frac{1}{\sqrt{z(2(r_2-r_1)+\delta+z)}} dz \right) \\
& \leq \frac{2}{\sqrt{r_1 r_2}} \left(2\pi + 2 \ln(8\sqrt{d}) + \ln(2\delta^{-1}) + \ln(2(2(r_2-r_1)+\delta)^{-1}) \right) \\
& \leq \frac{\kappa}{\sqrt{r_1 r_2}} (\ln \delta^{-1} + \ln((2(r_2-r_1)+\delta)^{-1}) + \kappa) .
\end{aligned}$$

Case 2.2: $\delta \in (0, \max\{0, 2(r_1 - r_2)\}]$.

We obtain the following upper bound on the density of Δ :

$$\begin{aligned}
& f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) \\
& \leq \frac{2}{\sqrt{r_1 r_2}} \int_{-\tau+\delta}^{2r_2-\tau+\delta} \left(\frac{1}{\sqrt{\tau+z}} + \frac{1}{\sqrt{2r_1-\tau-z}} \right) \left(\frac{1}{\sqrt{\tau+z-\delta}} + \frac{1}{\sqrt{2r_2-\tau-z+\delta}} \right) dz \\
& = \frac{2}{\sqrt{r_1 r_2}} \left(\int_0^{2r_2} \frac{1}{\sqrt{z(z+\delta)}} dz + \int_0^{2r_2} \frac{1}{\sqrt{z(2r_1-\delta-z)}} dz \right. \\
& \quad \left. + \int_0^{2r_2} \frac{1}{\sqrt{z(2r_2+\delta-z)}} dz + \int_0^{2r_2} \frac{1}{\sqrt{z(2(r_1-r_2)-\delta+z)}} dz \right) \\
& \leq \frac{2}{\sqrt{r_1 r_2}} \left(2\pi + 2 \ln(8\sqrt{d}) + \ln(2\delta^{-1}) + \ln(2(2(r_1-r_2)-\delta)^{-1}) \right) \\
& \leq \frac{\kappa}{\sqrt{r_1 r_2}} (\ln \delta^{-1} + \ln((2(r_1-r_2)-\delta)^{-1}) + \kappa) .
\end{aligned}$$

Third case: $r_1 \leq \tau \leq r_2$.

Since Z_1 takes only values in the interval $[-\tau, 2r_1 - \tau]$ and Z_2 takes only values in the interval $[-\tau, \tau]$, we can assume $0 < \delta \leq \min\{1/2, 2r_1\}$ and

$$f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) = \int_{-\tau+\delta}^{2r_1-\tau} f_{Z|T=\tau, R=r_1}(z) \cdot f_{Z|T=\tau, R=r_2}(z-\delta) dz .$$

For $\delta \in (0, \min\{1/2, 2r_1\}]$, we obtain the following upper bound on the density of Δ :

$$\begin{aligned}
& f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) \\
& \leq \frac{2}{\sqrt{\tau r_1}} \int_{-\tau+\delta}^{2r_1-\tau} \left(\frac{1}{\sqrt{\tau+z}} + \frac{1}{\sqrt{2r_1-\tau-z}} \right) \left(\frac{1}{\sqrt{\tau-z+\delta}} + \frac{1}{\sqrt{\tau+z-\delta}} \right) dz \\
& = \frac{2}{\sqrt{\tau r_1}} \left(\int_{\delta}^{2r_1} \frac{1}{\sqrt{z(2\tau+\delta-z)}} dz + \int_0^{2r_1-\delta} \frac{1}{\sqrt{z(2(\tau-r_1)+\delta+z)}} dz \right. \\
& \quad \left. + \int_0^{2r_1-\delta} \frac{1}{\sqrt{z(z+\delta)}} dz + \int_0^{2r_1-\delta} \frac{1}{\sqrt{z(2r_1-\delta-z)}} dz \right) \\
& \leq \frac{2}{\sqrt{\tau r_1}} \left(2\pi + 2 \ln(8\sqrt{d}) + \ln(2\delta^{-1}) + \ln(2(2(\tau-r_1)+\delta)^{-1}) \right) \\
& \leq \frac{\kappa}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} .
\end{aligned}$$

Fourth case: $r_2 \leq \tau \leq r_1$.

Since Z_1 takes only values in the interval $[-\tau, \tau]$ and Z_2 takes only values in the interval $[-\tau, 2r_2 - \tau]$, we can assume $0 < \delta \leq \min\{1/2, 2\tau\}$ and

$$f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) = \int_{-\tau+\delta}^{\min\{2r_2-\tau+\delta, \tau\}} f_{Z|T=\tau, R=r_1}(z) \cdot f_{Z|T=\tau, R=r_2}(z-\delta) dz .$$

Case 4.1: $\delta \in (0, 2(\tau - r_2)]$.

We obtain the following upper bound on the density of Δ :

$$\begin{aligned}
& f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) \\
& \leq \frac{2}{\sqrt{\tau r_2}} \int_{-\tau+\delta}^{2r_2-\tau+\delta} \left(\frac{1}{\sqrt{\tau-z}} + \frac{1}{\sqrt{\tau+z}} \right) \left(\frac{1}{\sqrt{\tau+z-\delta}} + \frac{1}{\sqrt{2r_2-\tau-z+\delta}} \right) dz \\
& = \frac{2}{\sqrt{\tau r_2}} \left(\int_0^{2r_2} \frac{1}{\sqrt{z(z+\delta)}} dz + \int_0^{2r_2} \frac{1}{\sqrt{z(2\tau-\delta-z)}} dz \right. \\
& \quad \left. + \int_0^{2r_2} \frac{1}{\sqrt{z(2r_2+\delta-z)}} dz + \int_0^{2r_2} \frac{1}{\sqrt{z(2(\tau-r_2)-\delta+z)}} dz \right) \\
& \leq \frac{2}{\sqrt{\tau r_2}} \left(2\pi + 2 \ln(8\sqrt{d}) + \ln(2\delta^{-1}) + \ln(2(2(\tau-r_2)-\delta)^{-1}) \right) \\
& \leq \frac{\kappa}{\sqrt{\tau r_2}} \left(\ln \delta^{-1} + \ln((2(\tau-r_2)-\delta)^{-1}) + \kappa \right) .
\end{aligned}$$

Case 4.2: $\delta \in (2(\tau - r_2), 2\tau]$.

We obtain the following upper bound on the density of Δ :

$$\begin{aligned}
& f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) \\
& \leq \frac{2}{\sqrt{\tau r_2}} \int_{-\tau+\delta}^{\tau} \left(\frac{1}{\sqrt{\tau-z}} + \frac{1}{\sqrt{\tau+z}} \right) \left(\frac{1}{\sqrt{\tau+z-\delta}} + \frac{1}{\sqrt{2r_2-\tau-z+\delta}} \right) dz \\
& = \frac{2}{\sqrt{\tau r_2}} \left(\int_0^{2\tau-\delta} \frac{1}{\sqrt{z(2\tau-\delta-z)}} dz + \int_0^{2\tau-\delta} \frac{1}{\sqrt{z(z+\delta)}} dz \right. \\
& \quad \left. + \int_0^{2\tau-\delta} \frac{1}{\sqrt{z(2(r_2-\tau)+\delta+z)}} dz + \int_{\delta}^{2\tau} \frac{1}{\sqrt{z(2r_2+\delta-z)}} dz \right) \\
& \leq \frac{2}{\sqrt{\tau r_2}} \left(2\pi + 2 \ln(8\sqrt{d}) + \ln(2\delta^{-1}) + \ln(2(2(r_2-\tau)+\delta)^{-1}) \right) \\
& \leq \frac{\kappa}{\sqrt{\tau r_2}} \left(\ln \delta^{-1} + \ln((2(r_2-\tau)+\delta)^{-1}) + \kappa \right) .
\end{aligned}$$

Altogether, this yields the lemma. \square

B.2 Proof of Lemma 4.8

First, we derive the following lemma, which gives bounds on the conditional density of the random variable Δ when only one of the radii R_1 and R_2 is given.

Lemma B.1. *Let $r_1, r_2, \tau \in (0, \sqrt{d})$ and $\delta \in (0, 1/2]$. In the following, let κ denote a sufficiently large constant.*

a) *The density of Δ under the conditions $T = \tau$ and $R_1 = r_1$ is bounded by*

$$f_{\Delta|T=\tau, R_1=r_1}(\delta) \leq \begin{cases} \frac{\kappa}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} & \text{if } r_1 \leq \tau, \\ \frac{\kappa}{\tau} \cdot \ln \delta^{-1} & \text{if } r_1 \geq \tau. \end{cases}$$

b) *The density of Δ under the conditions $T = \tau$ and $R_2 = r_2$ is bounded by*

$$f_{\Delta|T=\tau, R_2=r_2}(\delta) \leq \begin{cases} \frac{\kappa}{\sqrt{\tau r_2}} \cdot (\ln \delta^{-1} + \ln |2(\tau - r_2) - \delta|^{-1} + \kappa) & \text{if } r_2 \leq \tau, \\ \frac{\kappa}{\tau} \cdot \ln \delta^{-1} & \text{if } r_2 \geq \tau. \end{cases}$$

Proof. a) We can write the density of Δ under the conditions $T = \tau$ and $R_1 = r_1$ as

$$\begin{aligned}
f_{\Delta|T=\tau, R_1=r_1}(\delta) &= \int_0^{\sqrt{d}} f_{R_2}(r_2) \cdot f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) dr_2 \\
&= \int_0^{\sqrt{d}} \frac{r_2^{d-1}}{d^{d/2-1}} \cdot f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) dr_2 ,
\end{aligned}$$

where f_{R_2} denotes the density of the length $R_2 = \mathbf{d}(O, Q_2)$. We use Lemma 4.7 to bound this integral. For $r_1 \leq \tau$ and sufficiently large constants κ' and κ'' , we obtain

$$\begin{aligned} f_{\Delta|T=\tau, R_1=r_1}(\delta) &\leq \int_0^\tau \frac{r_2^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{r_1 r_2}} (\ln \delta^{-1} + \ln |2(r_1 - r_2) - \delta|^{-1} + \kappa) dr_2 \\ &\quad + \int_\tau^{\sqrt{d}} \frac{r_2^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} dr_2 \\ &\leq \frac{\kappa'}{\sqrt{r_1}} \cdot \ln \delta^{-1} + \frac{\kappa'}{\sqrt{r_1}} \int_0^{\sqrt{d}} \ln |2(r_1 - r_2) - \delta|^{-1} dr_2 + \frac{\kappa}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} \\ &\leq \frac{\kappa''}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} . \end{aligned}$$

For $\tau \leq r_1$ and a sufficiently large constant κ' , we obtain analogously

$$\begin{aligned} f_{\Delta|T=\tau, R_1=r_1}(\delta) &\leq \int_0^\tau \frac{r_2^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r_2}} (\ln \delta^{-1} + \ln |2(\tau - r_2) - \delta|^{-1} + \kappa) dr_2 \\ &\quad + \int_\tau^{\sqrt{d}} \frac{r_2^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\tau} \cdot \ln \delta^{-1} dr_2 \\ &\leq \frac{\kappa'}{\tau} \cdot \ln \delta^{-1} . \end{aligned}$$

b) We can write the density of Δ under the conditions $T = \tau$ and $R_2 = r_2$ as

$$f_{\Delta|T=\tau, R_2=r_2}(\delta) = \int_0^{\sqrt{d}} \frac{r_1^{d-1}}{d^{d/2-1}} \cdot f_{\Delta|T=\tau, R_1=r_1, R_2=r_2}(\delta) dr_1 .$$

For $r_2 \leq \tau$ and sufficiently large constants κ' and κ'' , we obtain

$$\begin{aligned} f_{\Delta|T=\tau, R_2=r_2}(\delta) &\leq \int_0^\tau \frac{r_1^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{r_1 r_2}} (\ln \delta^{-1} + \ln |2(r_1 - r_2) - \delta|^{-1} + \kappa) dr_1 \\ &\quad + \int_\tau^{\sqrt{d}} \frac{r_1^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r_2}} (\ln \delta^{-1} + \ln |2(\tau - r_2) - \delta|^{-1} + \kappa) dr_1 \\ &\leq \frac{\kappa'}{\sqrt{r_2}} (\ln \delta^{-1} + \kappa) + \frac{\kappa'}{\sqrt{r_2}} \int_0^{\sqrt{d}} \ln |2(r_1 - r_2) - \delta|^{-1} dr_1 \\ &\quad + \frac{\kappa}{\sqrt{\tau r_2}} (\ln \delta^{-1} + \ln |2(\tau - r_2) - \delta|^{-1} + \kappa) \\ &\leq \frac{\kappa''}{\sqrt{\tau r_2}} (\ln \delta^{-1} + \ln |2(\tau - r_2) - \delta|^{-1} + \kappa) . \end{aligned}$$

For $\tau \leq r_2$ and a sufficiently large constant κ' , we obtain

$$\begin{aligned} f_{\Delta|T=\tau, R_2=r_2}(\delta) &\leq \int_0^\tau \frac{r_1^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r_1}} \cdot \ln \delta^{-1} dr_1 + \int_\tau^{\sqrt{d}} \frac{r_1^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\tau} \cdot \ln \delta^{-1} dr_1 \\ &\leq \frac{\kappa'}{\tau} \cdot \ln \delta^{-1} . \end{aligned} \quad \square$$

Now we are ready to prove Lemma 4.8.

Proof of Lemma 4.8. a) In order to prove part a), we integrate $f_{\Delta|T=\tau, R_1=r}(\delta)$ over all values τ that T can take:

$$\begin{aligned} f_{\Delta|R_1=r}(\delta) &= \int_0^{\sqrt{d}} \frac{\tau^{d-1}}{d^{d/2-1}} \cdot f_{\Delta|T=\tau, R_1=r}(\delta) d\tau \\ &\leq \int_0^r \frac{\tau^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\tau} \cdot \ln \delta^{-1} d\tau + \int_r^{\sqrt{d}} \frac{\tau^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r}} \cdot \ln \delta^{-1} d\tau \\ &\leq \frac{\kappa'}{\sqrt{r}} \cdot \ln \delta^{-1} . \end{aligned}$$

Furthermore, we integrate $f_{\Delta|T=\tau, R_2=r}(\delta)$ over all values τ that T can take:

$$\begin{aligned} f_{\Delta|R_2=r}(\delta) &= \int_0^{\sqrt{d}} \frac{\tau^{d-1}}{d^{d/2-1}} \cdot f_{\Delta|T=\tau, R_2=r}(\delta) d\tau \\ &\leq \int_0^r \frac{\tau^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\tau} \cdot \ln \delta^{-1} d\tau \\ &\quad + \int_r^{\sqrt{d}} \frac{\tau^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r}} (\ln \delta^{-1} + \ln |2(\tau - r) - \delta|^{-1} + \kappa) d\tau \\ &\leq \frac{\kappa'}{\sqrt{r}} \cdot \ln \delta^{-1} . \end{aligned}$$

b) For a sufficiently large constant κ' ,

$$\begin{aligned} f_{\Delta|T=\tau}(\delta) &= \int_0^{\sqrt{d}} \frac{r^{d-1}}{d^{d/2-1}} \cdot f_{\Delta|T=\tau, R_1=r}(\delta) dr \\ &\leq \int_0^{\tau} \frac{r^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\sqrt{\tau r}} \cdot \ln \delta^{-1} dr + \int_{\tau}^{\sqrt{d}} \frac{r^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\tau} \cdot \ln \delta^{-1} dr \\ &\leq \frac{\kappa'}{\tau} \cdot \ln \delta^{-1} . \end{aligned}$$

c) For a sufficiently large constant κ' ,

$$f_{\Delta}(\delta) = \int_0^{\sqrt{d}} \frac{\tau^{d-1}}{d^{d/2-1}} \cdot \frac{\kappa}{\tau} \cdot \ln \delta^{-1} d\tau \leq \kappa' \cdot \ln \delta^{-1} .$$

d) For sufficiently large constants κ' and κ'' ,

$$\begin{aligned} f_{Z|T=\tau}(z) &= \int_{r=0}^{\tau} \frac{r^{d-1}}{d^{d/2-1}} \sqrt{\frac{2}{(\tau+z)(2r-\tau-z)}} dr + \int_{r=\tau}^{\sqrt{d}} \frac{r^{d-1}}{d^{d/2-1}} \sqrt{\frac{2}{\tau^2-z^2}} dr \\ &\leq \sqrt{\frac{2}{\tau+z}} d^{1/2} \int_{r=0}^{\tau} \sqrt{\frac{1}{2r-\tau-r}} dr + \sqrt{\frac{2}{\tau^2-z^2}} \\ &\leq \frac{\kappa'}{\sqrt{\tau+z}} + \sqrt{\frac{2}{\tau^2-z^2}} \leq \frac{\kappa''}{\sqrt{\tau^2-z^2}} . \end{aligned}$$

□