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

(Extended Abstract)

Matthias Englert Heiko Röglin Berthold Vöcking

Abstract

2-Opt is probably the most basic and widely used 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. In this paper, we clarify this issue by presenting a family of Euclidean 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 $[0,1]^2$. In particular, we allow different distributions for different points. We study the expected running time 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 $\tilde{O}(n^{3+5/6} \cdot \phi^{8/3})$. If the distances are measured according to the Manhattan metric, then the expected number of steps is bounded by $\tilde{O}(n^{3+1/2} \cdot \phi)$. In addition, we prove an upper bound of $O(\sqrt{\phi})$ on the expected approximation factor with respect to both of these 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 $\phi \sim 1/\sigma^2$. 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 running time of 2-Opt.

1 Introduction

In the *traveling salesperson problem (TSP)*, we are given a set $\{v_1, v_2, \ldots, v_n\}$ of *vertices* and for each pair $\{v_i, v_j\}$ of distinct vertices a distance $d(v_i, v_j)$. The goal is to find a tour of minimal length visiting each vertex exactly once and returning to the initial vertex at the end, that is, the goal is to compute a permutation π minimizing

$$\sum_{i=1}^{n-1} d(v_{\pi(i)}, v_{\pi(i+1)}) + d(v_{\pi(n)}, v_{\pi(1)}) \ .$$

Despite many theoretical analyses and experimental evaluations of the TSP, there is still a considerable gap between the theoretical results and the experimental observations. The Euclidean TSP, for example, is known to be NP-hard in the strong sense [Pap77]. In a breakthrough result, Arora has shown that the Euclidean TSP admits a polynomial time approximation scheme (PTAS) and, hence, can be approximated arbitrarily well in polynomial time [Aro98]. Arora's PTAS is 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 and widely used local search heuristic for the TSP. 2-Opt starts with an arbitrary initial tour and incrementally improves this tour by exchanging 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 the algorithm 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

*Department of Computer Science, RWTH Aachen, Germany. This work was supported by DFG grants VO 889/2 and WE 2842/1. {englert,roeglin,voecking}@cs.rwth-aachen.de. A full version appeared as technical report [ERV06].

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 lies 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 u to a vertex v is contained if v is obtained from u by performing one 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. Observe that such steps can introduce new crossings. However, van Leeuwen and Schoone show that after $O(n^3)$ steps, 2-Opt has found a tour without any crossing [vLS80]. On the negative side, Lueker constructs TSP instances whose state graphs contain exponentially long paths, that is, 2-Opt can take an exponential number of steps before it finds a locally optimal solution [Lue75]. This result is generalized to k-Opt for arbitrary $k \geq 2$ by Chandra, Karloff, and Tovey [CKT99]. These 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. In this paper, we settle this question. We construct Euclidean instances whose state graphs contain exponentially longs paths, that is, Euclidean instances on which 2-Opt can take an exponential number of steps before finding a locally optimal solution. In chip design applications, often TSP instances arise in which the distances are measured according to the Manhattan metric. Also for this metric, we construct instances with exponentially long paths in the 2-Opt state graph.

THEOREM 1.1. *a) For every $n \in \mathbb{N}$, there is a graph in the plane with Manhattan metric with $16n$ vertices whose corresponding state graph contains a path of length $2^{n+4} - 22$.*

b) For every $n \in \mathbb{N}$, there is a graph in the Euclidean plane with $8n$ vertices whose corresponding state graph contains a path of length $2^{n+3} - 14$.

For Euclidean instances in which n points are placed uniformly at random in the unit square, Kern 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 [Ker89]. Chandra, Karloff, and Tovey improve this result by bounding the expected length of the longest path in the state graph by $O(n^{10} \log n)$ [CKT99]. 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 in the plane using different continuous probability distributions. The distribution of vertex v_i is defined by a density function $f_i : [0,1]^2 \to [0, \phi]$ 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 worst case instances can be approximated by the distributions. For $\phi = 1$, 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 $\phi \sim 1/\sigma^2$.

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 which can be specified by density functions $f : [0, 1] \to [0, \phi]$ 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 touched by the randomization. This can be useful if one wants to explicitely 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.

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.

One way of improving the approximation ratio and running time of 2-Opt is to use an *insertion heuristic* for computing the initial tour. We show that using such an insertion heuristic yields a significant improvement since the initial tour 2-Opt starts with is an $O(\log n)$-approximation of the optimal tour and hence 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 $O(n^{3+1/2} \cdot \log n \cdot \phi)$ on ϕ-perturbed Manhattan instances with n points when one starts with a tour obtained by an arbitrary insertion heuristic.

b) is $O(n^{3+5/6} \cdot \log(n\phi) \cdot \log n \cdot \phi^{8/3})$ on ϕ-perturbed Euclidean instances with n points when one starts with a tour obtained by an arbitrary insertion heuristic.

The bounds in the previous theorem can be improved by one $\log n$-factor if one does not use an arbitrary insertion heuristic but *nearest insertion* or *cheapest insertion* since these heuristics yield a 2-approximation of the optimal tour.

In fact, our analysis shows not only that the expected running time 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 show that there are Euclidean instances 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 [CKT99]. However, the same authors also show that the expected approximation ratio 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. For both Euclidean and Manhattan instances, we obtain the following theorem.

THEOREM 1.4. *For ϕ-perturbed Manhattan and Euclidean instances, the expected approximation ratio of the worst tour that is locally optimal for 2-Opt is bounded by $O(\sqrt{\phi})$.*

Let us remark that this result is merely of theoretical interest since a 2-approximation of the optimal tour is already achieved by computing the initial tour 2-Opt starts with by an appropriate insertion heuristic. However, this result shows that the approximation ratio is bounded by $O(\sqrt{\phi})$ no matter how the initial tour is chosen.

In this extended abstract, we focus on the results about the number of 2-changes for the Euclidean TSP, that is, we outline the proofs of Theorem 1.1 b), Theorem 1.2 b), and Theorem 1.3 b). The complete proofs of these and the other results can be found in the full version of this paper [ERV06].

2 Preliminaries

We begin by stating some basic definitions and notations. First of all, an instance of the TSP consists of a set $V = \{v_1, \ldots, v_n\}$ of *vertices* (depending on the context, synonymously referred to as *points*) and a symmetric *distance function* $d : V \times V \to \mathbb{R}_{\geq 0}$ which 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 tour of minimal length visiting each vertex exactly once and returning to the initial vertex at the end. Since we focus on Euclidean instances, we assume in the following that $V \subseteq \mathbb{R}^2$ and that the distance between two points $P_1 = (x_1, y_1) \in \mathbb{R}^2$ and $P_2 = (x_2, y_2) \in \mathbb{R}^2$ is given by $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$.

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. show an upper bound of $\lceil \log n \rceil + 1$ on the approximation factor of any insertion heuristic on metric TSP instances [RSI77]. 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 Euclidean Instances with Exponentially Long Sequences of Improving 2-Changes

In this section, we present a family of Euclidean instances with exponentially long sequences of improving 2-changes. In Lueker's construction for the general TSP many of the 2-changes remove two edges which 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. This is crucial to the construction and also to its generalization to k-changes. Our construction differs significantly from the previous ones as the 2-changes 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 $G_0, \ldots, G_{n-1}$ 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. 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 the j-th block of the i-th gadget consist of. If one ignores certain intermediate configurations that arise when one gadgets

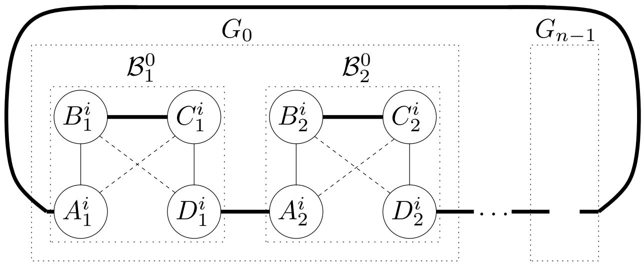

Figure 1: Every tour that occurs in the sequence of 2-changes contains the thick edges. For each block either both solid or both dashed edges are contained. In the former case the block is in its short state; in the latter case the block is in its long state.

resets another one, our construction ensures the following properties: The points 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 $d(A_j^i, B_j^i) + d(C_j^i, D_j^i)$ of the distances between A_j^i and B_j^i and between C_j^i and 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 $\mathcal{B}_1^0 \mathcal{B}_2^0 \mathcal{B}_1^1 \mathcal{B}_2^1 \ldots \mathcal{B}_1^{n-1} \mathcal{B}_2^{n-1}$. See Figure 1 for an illustration.

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 a gadget G_i by a vector (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 gad-

0)	[A_2^i	C_2^i]	B_2^i	D_2^i	A_1^{i+1}	B_1^{i+1}	C_1^{i+1}	D_1^{i+1}	A_2^{i+1}	B_2^{i+1}	[C_2^{i+1}	D_2^{i+1}]
1)	A_2^i	C_2^{i+1}	[B_2^{i+1}	A_2^{i+1}]	D_1^{i+1}	C_1^{i+1}	B_1^{i+1}	A_1^{i+1}	[D_2^i	B_2^i]	C_2^i	D_2^{i+1}
2)	A_2^i	C_2^{i+1}	[B_2^{i+1}	D_2^i]	A_1^{i+1}	B_1^{i+1}	[C_1^{i+1}	D_1^{i+1}]	A_2^{i+1}	B_2^i	C_2^i	D_2^{i+1}
3)	A_2^i	C_2^{i+1}	B_2^{i+1}	C_1^{i+1}	[B_1^{i+1}	A_1^{i+1}]	D_2^i	D_1^{i+1}	A_2^{i+1}	B_2^i	[C_2^i	D_2^{i+1}]
4)	[A_2^i	C_2^{i+1}]	B_2^{i+1}	C_1^{i+1}	B_1^{i+1}	C_2^i	[B_2^i	A_2^{i+1}]	D_1^{i+1}	D_2^i	A_1^{i+1}	D_2^{i+1}
5)	A_2^i	B_2^i	C_2^i	B_1^{i+1}	[C_1^{i+1}	B_2^{i+1}]	C_2^{i+1}	A_2^{i+1}	D_1^{i+1}	D_2^i	[A_1^{i+1}	D_2^{i+1}]
6)	A_2^i	B_2^i	[C_2^i	B_1^{i+1}]	C_1^{i+1}	A_1^{i+1}	[D_2^i	D_1^{i+1}]	A_2^{i+1}	C_2^{i+1}	B_2^{i+1}	D_2^{i+1}
7)	A_2^i	B_2^i	C_2^i	D_2^i	A_1^{i+1}	C_1^{i+1}	B_1^{i+1}	D_1^{i+1}	A_2^{i+1}	C_2^{i+1}	B_2^{i+1}	D_2^{i+1}

Figure 2: A sequence of seven 2-changes in which the first block changes from its long state to its short state while resetting the two other blocks from their short to their long states. Brackets indicate the edges that are removed from the tour.

gets we construct possess the following properties:

1. If gadget G_i with $0 \le i \le n-2$ is in state (L, L) and gadget G_{i+1} is in state (S, S), then there exists a sequence of 7 consecutive 2-changes involving only edges of and between the gadgets G_i and G_{i+1} terminating with gadget G_i being in state (S, L) and gadget G_{i+1} in state (L, L).
2. If gadget G_i with $0 \le i \le n-2$ is in state (S, L) and gadget G_{i+1} is in state (S, S), then there exists a sequence of 7 consecutive 2-changes involving only edges of and between the gadgets G_i and G_{i+1} terminating with gadget G_i being in state (S, S) and gadget G_{i+1} in state (L, L).

If these properties are 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. In order to see this, we prove the following lemma.

Lemma 3.1. *If gadget G_i with $0 \le i \le n-1$ 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 $2^{n+3-i} - 14$ 2-changes in which only edges of and between the gadgets G_j with $j \ge i$ are involved and that terminates in a state in which all gadgets G_j with $j \ge i$ 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 the first property, there exists a sequence of 7 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 the second property, there exists a sequence of 7 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 $j \ge i$ are in state (S, S). This concludes the proof as the number of 2-changes performed is $14 + 2(2^{n+2-i} - 14) = 2^{n+3-i} - 14$. □

We still need to show how the aforementioned properties are achieved, that is, how the sequence of 2-changes exactly looks like. We first present a sequence of 2-changes that satisfies the second property. Observe that the initial situation in the second property is as follows: There are three consecutive blocks, namely $\mathcal{B}_2^i$, $\mathcal{B}_1^{i+1}$, and $\mathcal{B}_2^{i+1}$, the leftmost one is in its long state, and the other blocks are in their short states. We need to present 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 sequence of seven 2-changes shown in Figure 2 satisfies the second property.

A sequence of steps that satisfies the first property 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, that is, after the seven 2-changes have been performed, the block is in the same state and at the same position as before.

It remains to find a set of points in the plane such that all 2-changes in the construed sequence of 2-changes decrease the length of the tour. The remaining details about how such points in the Euclidean plane can be chosen can be found in the full version of this paper [ERV06].

4 The Expected Number of 2-Changes

In this section, we analyze the expected length of the longest path in the state graph of ϕ-perturbed Euclidean instances. In order to prove an upper bound on this length, one could analyze the smallest improvement $\Delta_{\min}$ made by any of the 2-changes. If $\Delta_{\min}$ is bounded from below by some term δ, then there cannot be a sequence of more than $\sqrt{2}n/\delta$ consecutive improving 2-changes since the initial tour has length at most $\sqrt{2}n$. We like to mention that already the analysis of $\Delta_{\min}$ that we present in the following yields a bound of $O(n^7 \cdot \log^2 n \cdot \phi^3)$ on the expected number of 2-changes which improves the previously known bound.

Intuitively, this approach is too pessimistic since most of the steps performed by 2-Opt yield a larger improvement than $\Delta_{\min}$. In particular, two consecutive steps yield an improvement of at least $\Delta_{\min}$ plus the improvement $\Delta'_{\min}$ 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 in fact much larger since 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 Euclidean instances claimed in Theorem 1.2.

4.1 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 4.1. *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 $t/3 - n(n-1)/12$.*

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 linked pairs of 2-changes is created. However, these pairs are not necessarily disjoint. Hence, after the list has been created, pairs have to be removed from the list until there are no non-disjoint pairs left. 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. In this case, the pair (S_i, S_j) is added to the constructed list $\mathcal{L}$. Analogously, 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. In this case, the pair $(S_i, S_{j'})$ is added to the list $\mathcal{L}$.

After the sequence has been processed completely, each pair in $\mathcal{L}$ is linked by an edge but we still have to identify a subset $\mathcal{L}'$ of $\mathcal{L}$ consisting only of pairwise disjoint pairs. This subset is constructed in a greedy fashion. We process the list $\mathcal{L}$ step by step, starting with an empty list $\mathcal{L}'$. For each pair in $\mathcal{L}$, we check whether it is disjoint from all pairs which have already been inserted into $\mathcal{L}'$ or not. In the former case, the current pair is inserted into $\mathcal{L}'$. This way, we obtain a list $\mathcal{L}'$ of disjoint pairs such that each pair is linked by an edge. The number of pairs in $\mathcal{L}$ is at least $2t - n(n-1)/2$ since each of the t steps gives rise to 2 different pairs, unless an edge is added to the tour which is never removed again. Each 2-change occurs in at most 4 different pairs in $\mathcal{L}$, hence, each pair in $\mathcal{L}$ is non-disjoint from at most 6 other pairs in $\mathcal{L}$. This implies that $\mathcal{L}$ contains at most 6 times as many pairs as $\mathcal{L}'$. □

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

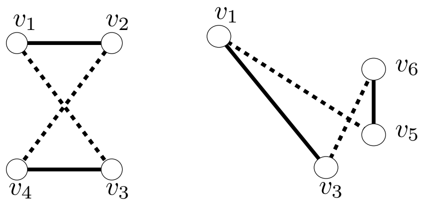

Figure 3: A pair of type 1: First the edges $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are exchanged with the edges $\{v_1, v_3\}$ and $\{v_2, v_4\}$, then $\{v_1, v_3\}$ and $\{v_5, v_6\}$ are exchanged with $\{v_1, v_5\}$ and $\{v_3, v_6\}$.

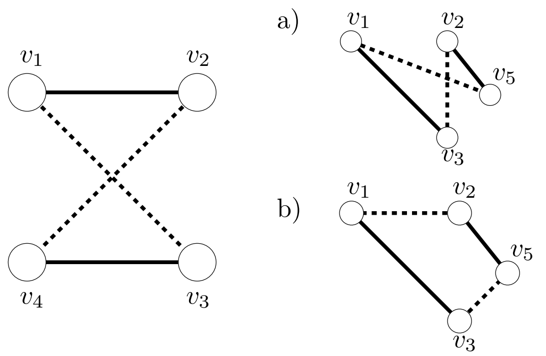

Figure 4: A pair of type 2: First the edges $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are exchanged with the edges $\{v_1, v_3\}$ and $\{v_2, v_4\}$, then $\{v_1, v_3\}$ and $\{v_2, v_5\}$ are exchanged with either $\{v_1, v_5\}$ and $\{v_2, v_3\}$ or $\{v_1, v_2\}$ and $\{v_3, v_5\}$.

1. $|\{v_2, v_4\} \cap \{v_5, v_6\}| = 0$. We can assume w. l. o. g. that in the second step the edges $\{v_1, v_5\}$ and $\{v_3, v_6\}$ are added to the tour. See Figure 3.
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.
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 $v_4 = v_5$.

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 cannot be bounded appropriately. Hence 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 4.1. The following lemma, whose proof can be found in the full version of this paper [ERV06], shows that there are always enough pairs of type 1 or 2.

Lemma 4.2. *In every sequence of t consecutive 2-changes the number of disjoint pairs of 2-changes of type 1 or 2 is at least $t/6 - 5n(n-1)/48$.*

4.2 Analysis of Pairs of Linked 2-Changes

We prove the following lemmas about pairs of type 1 and 2 for ϕ-perturbed Euclidean instances.

Lemma 4.3. *The probability that there exists a pair of type 1 in which both 2-changes are improvements by at most ε is bounded by $O(n^6 \cdot \varepsilon^2 \cdot (\log^2(1/\varepsilon) + 1) \cdot \phi^5)$.*

Lemma 4.4. *The probability that there exists a pair of type 2 in which both 2-changes are improvements by at most ε is bounded by $O(n^5 \cdot \varepsilon^{3/2} \cdot (\log(1/\varepsilon) + 1) \cdot \phi^4)$.*

Before we can prove Lemmas 4.3 and 4.4, we have to understand the random variable that describes the improvement of a single 2-change. In this section, we analyze this variable under certain conditions. If, e. g., we would like to analyze a pair of linked 2-changes that share an edge e, it is helpful to know the densities of the random variables that describe the improvements of the first and the second 2-change under the condition that the length of e is given.

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 considered input model, each of these points has a probability distribution over the unit square according to which it is chosen. We consider a simplified random experiment in which O is chosen to be the origin and P, Q_1, and Q_2 are chosen independently and uniformly from the interior of a circle with radius $\sqrt{2}$ around the origin. Due to the rotational symmetry of this model, we assume further that that P lies at position $(0, T)$, where T denotes the distance between O and P. In the next section, we explain how the analysis of this simple random experiment helps us to analyze the actual random experiment that occurs in the probabilistic input model. Let Z_1 denote the difference $d(O, Q_1) - d(P, Q_1)$ and let Z_2 denote the difference $d(O, Q_2) - d(P, Q_2)$. Then the improvement Δ of the 2-Opt step can be expressed as $Z_1 - Z_2$. In the full version of this paper [ERV06], we prove the following lemmas on the conditional densities of the random variables Δ and Z_i. In the following, let $R_1 := d(O, Q_1)$ and $R_2 := d(O, Q_2)$.

Lemma 4.5. *Let τ be an arbitrary distance with $0 \leq \tau \leq \sqrt{2}$. For a sufficiently large constant κ and for*

$i \in \{1,2\}$, *the conditional density* $f_{Z_i|T=\tau}$ *of* Z_i *under the condition* $T = \tau$ *can be bounded by*

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

if $|z| \le \tau$. *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 4.6. *Let* r *be an arbitrary distance with* $0 \le r \le \sqrt{2}$. *For a sufficiently large constant* κ *and for* $i \in \{1,2\}$, *the conditional density* $f_{\Delta|R_i=r}(\delta)$ *of* Δ *for* $\delta \ge 0$ *under the condition* $d(O, Q_i) = r$ *can be bounded by*

$$f_{\Delta|R_i=r}(\delta) \le \frac{\kappa}{\sqrt{r}} \cdot \left(\ln\left(\frac{1}{\delta}\right) + 1 \right) .$$

LEMMA 4.7. *Let* τ *be an arbitrary distance with* $0 \le \tau \le \sqrt{2}$. *For a sufficiently large constant* κ, *the conditional density* $f_{\Delta|T=\tau}(\delta)$ *of* Δ *for* $\delta \ge 0$ *under the condition* $T = \tau$ *can be bounded by*

$$f_{\Delta|T=\tau}(\delta) \le \frac{\kappa}{\tau} \cdot \left(\ln\left(\frac{1}{\delta}\right) + 1 \right) .$$

4.3 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 square. 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 has position $(x, y) \in [0,1]^2$. Then one can identify a region $\mathcal{R}_{(x,y)} \subseteq \mathbb{R}^6$ with the property that the event $\mathcal{E}$ occurs if and only if the random vector (P, Q_1, Q_2) lies in $\mathcal{R}_{(x,y)}$. No matter of how the position (x, y) of O is chosen, this region always has the same shape, only its position is shifted. Let $\mathcal{V} = \sup_{(x,y)\in[0,1]^2} \mathrm{vol}(\mathcal{R}_{(x,y)} \cap [0,1]^6)$. Then the probability of $\mathcal{E}$ can be bounded from above by $\phi^3 \cdot \mathcal{V}$ in the original random experiment. One can easily see that

$$|\mathcal{R}_{(x,y)} \cap [0,1]^6| \le |\mathcal{R}_{(0,0)} \cap [-1,1]^6| .$$

Hence for $\mathcal{V}' = \mathrm{vol}(\mathcal{R}_{(0,0)} \cap [-1,1]^6)$ we have $\mathcal{V} \le \mathcal{V}'$. Observe that the probability of $\mathcal{E}$ in the simplified random experiment can be bounded from below by $(1/(2\pi))^3 \cdot \mathcal{V}'$ since the circle centered around the origin with radius $\sqrt{2}$ contains the square $[-1,1]^2$ completely. Hence, the probability of $\mathcal{E}$ in the simplified random experiment is smaller by at most a factor of $(2\pi\phi)^3$ compared to the original random experiment.

Pairs of Type 1. 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 from a circle with radius $\sqrt{2}$ centered at v_3. Let $\mathcal{E}$ denote the event that both the improvement Δ_1 of the first step and the improvement Δ_2 of the second step 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 $(2\pi\phi)^5$.

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 from a circle with radius $\sqrt{2}$ centered at v_2. In this case the probability in the simplified random experiment is smaller by at most a factor of $(2\pi\phi)^4$.

4.4 Analysis of Pairs of Linked 2-Changes

Finally, we can prove Lemmas 4.3 and 4.4.

Proof of Lemma 4.3. 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 circle with radius $\sqrt{2}$ 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 $i \in \{1,2\}$, 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 4.6 and obtain, for a sufficiently large constant κ,

$$\begin{aligned} &\mathbf{Pr}\left[\Delta_i \in [0,\varepsilon] \mid d(v_1, v_3) = r\right] \\ &= \int_0^{\varepsilon} f_{\Delta_i|d(v_1,v_3)=r}(\delta)\, d\delta \le \frac{\kappa \cdot \varepsilon}{\sqrt{r}} \cdot \left(\ln\left(\frac{1}{\varepsilon}\right) + 1 \right) .\end{aligned}$$

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

$$
\begin{aligned}
& \mathbf{Pr}\left[\Delta_1, \Delta_2 \in [0,\varepsilon] \mid d(v_1,v_3) = r\right] \\
\leq\ & \frac{\kappa^2}{r} \cdot \varepsilon^2 \cdot \left(\ln\left(\frac{1}{\varepsilon}\right) + 1\right)^2 \\
\leq\ & \frac{\kappa'}{r} \cdot \varepsilon^2 \cdot \left(\ln^2\left(\frac{1}{\varepsilon}\right) + 1\right) \qquad (4.1)
\end{aligned}
$$

for a sufficiently large constant κ'. For $r \in [0, \sqrt{2}]$, the density $f_{d(v_1,v_3)}$ of the random variable $d(v_1, v_3)$ in the simplified random experiment is $f_{d(v_1,v_3)}(r) = r$. Combining this observation with the bound given in (4.1) yields

$$
\begin{aligned}
& \mathbf{Pr}\left[\Delta_1, \Delta_2 \in [0,\varepsilon]\right] \\
=\ & \int_0^{\sqrt{2}} r \cdot \mathbf{Pr}\left[\Delta_1, \Delta_2 \in [0,\varepsilon] \mid d(v_1,v_3) = r\right] \, dr \\
\leq\ & \sqrt{2}\kappa' \cdot \varepsilon^2 \cdot \left(\ln^2\left(\frac{1}{\varepsilon}\right) + 1\right) .
\end{aligned}
$$

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

The proof of Lemma 4.4 uses similar arguments.

4.5 The Expected Number of 2-Changes

Based on Lemmas 4.2, 4.3, and 4.4, we are now able to prove part b) of Theorem 1.2.

Proof of Theorem 1.2 b). 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 $S_1, \ldots, S_t$ 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 4.2, we know that we can find at least $t/6 - 5n(n-1)/48$ such pairs. For $i \in \{1, 2\}$, let $\Delta_{\min}^{(i)}$ denote the smallest improvement made by any pair of improving 2-Opt steps of type i. For $t > n^2$, we have $t/6 - 5n(n-1)/48 > t/16$ and hence due to Lemmas 4.3 and 4.4,

$$
\begin{aligned}
& \mathbf{Pr}\left[T \geq t\right] \\
\leq\ & \mathbf{Pr}\left[\Delta_{\min}^{(1)} \leq \frac{16\sqrt{2}n}{t}\right] + \mathbf{Pr}\left[\Delta_{\min}^{(2)} \leq \frac{16\sqrt{2}n}{t}\right] \\
=\ & O\left(\min\left\{\frac{n^8(\log^2(t/n)+1)}{t^2}\phi^5, 1\right\}\right) \\
& + O\left(\min\left\{\frac{n^{13/2}(\log(t/n)+1)}{t^{3/2}}\phi^4, 1\right\}\right) .
\end{aligned}
$$

This implies the following bound on the expected length of the longest path in the state graph

$$
\begin{aligned}
\mathbf{E}[T] &= n^2 + \sum_{t=1}^{n!} O\left(\min\left\{\frac{n^8 \log^2 t}{t^2}\phi^5, 1\right\}\right) \\
& \quad + O\left(\min\left\{\frac{n^{13/2}\log t}{t^{3/2}}\phi^4, 1\right\}\right) .
\end{aligned}
$$

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

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

This concludes the proof of part b) of the theorem. □

It is well-known that for an arbitrary set of n points in the unit square and for an arbitrary metric on $\mathbb{R}^2$ the optimal tour visiting all n points has length $O(\sqrt{n})$ (see, e. g., [CKT99]). Furthermore, every insertion heuristic finds an $O(\log n)$-approximation [RSI77]. Hence, if one starts with a solution calculated by an insertion heuristic, the initial tour has length $O(\sqrt{n} \cdot \log n)$. Using this observation yields part b) of Theorem 1.3.

5 Smoothed Analysis

Smoothed Analysis was introduced by Spielman and Teng as a hybrid of worst case and average case analysis [ST04]. 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 an average case analysis alone as it allows an adversary to specify a worst case 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.

We suggest the following perturbation model for Euclidean instances of the TSP. First an adversary chooses a set of n points in the unit square. Then the coordinates of these points are perturbed by adding independent random variables to them. The random variables we add are basically Gaussian random variables with standard deviation $\sigma \leq 1$. If, however, one of the added random variables has an absolute value larger than some given $\alpha \geq 1$, then we draw another Gaussian random variable with standard deviation σ until the absolute value is bounded by α. Let X denote one such random variable and let Y denote a Gaussian random variable with standard deviation σ and density function

f_Y. Then the density f_X of X can be bounded by

$$f_X(x) \leq \frac{\sup_{y\in\mathbb{R}} f_Y(y)}{\mathbf{Pr}\left[|Y| \leq \alpha\right]} \leq \frac{1/(\sigma\sqrt{2\pi})}{1 - \sigma/\sqrt{2\pi}\cdot\exp(-\alpha^2/(2\sigma^2))} .$$

Observe that after the perturbation all points lie in the square $[-\alpha, 1+\alpha]^2$. Hence, in order to apply Theorems 1.2, 1.3, and 1.4, we first have to scale and shift the instance such that every point lies in the unit square. This can increase the density f_X of X by at most a factor of $(2\alpha+1)^2$. Thus with

$$(5.2) \qquad \phi = \frac{(2\alpha+1)^2}{(\sigma\sqrt{2\pi} - \sigma^2\exp(-\alpha^2/(2\sigma^2)))^2} = O\left(\frac{\alpha^2}{\sigma^2}\right) .$$

we can apply the aforementioned theorems.

Finally, let us remark that if the standard deviation is small enough, then it is not necessary to redraw the Gaussian random variables until they lie in the interval $[-\alpha, \alpha]$. For $\sigma \leq \min\{\alpha/\sqrt{4n\ln n}, 1\}$, the probability that one of the Gaussian random variables has an absolute value larger than $\alpha \geq 1$ can be bounded by

$$\frac{2n}{\sigma\sqrt{2\pi}} \cdot \int_{x=\alpha}^{\infty} \exp\left(-x^2/(2\sigma^2)\right)\, dx \leq \frac{2n\sigma}{\sqrt{2\pi}} \cdot \exp\left(-\alpha^2/(2\sigma^2)\right) \leq n^{-2n} .$$

In this case, even if one does not redraw the random variables outside $[-\alpha, \alpha]$, the Theorems 1.2, 1.3, and 1.4 can be applied with the corresponding ϕ given in (5.2). To see this, one must only 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^{-2n} 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 $O(\text{poly}(n, 1/\sigma))$.

6 Open Questions

We constructed Euclidean and Manhattan instances which possess exponentially long sequences of improving 2-changes. We leave open the question whether instances exist whose state graphs contain nodes from which every path leading to a sink has exponential length, that is, instances on which every 2-Opt variant has to make an exponential number of steps. Another open problem is to further narrow the gap between the upper bound on the expected number of 2-changes on uniform and ϕ-perturbed instances and the experimentally observed number of steps.

As we have already mentioned, one easily obtains a 2-approximation for metric TSP instances by using an appropriate insertion heuristic and, hence, our results about the expected approximation factor and the results in [CKT99] are merely of theoretical interest. In experimental studies it has been observed that 2-Opt usually yields an approximation which lies within a few percentage points of the global optimum on Euclidean instances [JM97]. Hence, an interesting open question is to show that the expected approximation factor of 2-Opt on uniform Euclidean instances is a small constant.

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.

[ERV06] Matthias Englert, Heiko Röglin, and Berthold Vöcking. Worst case and probabilistic analysis of the 2-Opt algorithm for the TSP. Technical Report TR06-092, ECCC, 2006.

[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.

[Lue75] George S. Lueker. Unpublished manuscript, 1975. Princeton University.

[Pap77] Christos H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. *Theoretical Computer Science*, 4(3):237–244, 1977.

[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.

[vLS80] Jan van Leeuwen and Anneke A. Schoon. Untangling a traveling salesman tour in the plane. Technical report, University of Utrecht, 1980. Report No. RUU-CS-80-11.