Smoothed Analysis of the 2-Opt Heuristic for the TSP under Gaussian Noise *

Marvin Künnemann¹, Bodo Manthey², and Rianne Veenstra²

¹Max Planck Institute for Informatics, Saarbrücken, Germany, and Saarbrücken Graduate School of Computer Science, marvin@mpi-inf.mpg.de

²University of Twente, Department of Applied Mathematics, Enschede, Netherlands, b.manthey@utwente.nl

August 2, 2023

The 2-opt heuristic is a very simple local search heuristic for the traveling salesperson problem. In practice it usually converges quickly to solutions within a few percentages of optimality. In contrast to this, its running-time is exponential and its approximation performance is poor in the worst case.

Englert, Röglin, and Vöcking (Algorithmica, 2014) provided a smoothed analysis in the so-called one-step model in order to explain the performance of 2-opt on *d*-dimensional Euclidean instances, both in terms of running-time and in terms of approximation ratio. However, translating their results to the classical model of smoothed analysis, where points are perturbed by Gaussian distributions with standard deviation σ , yields only weak bounds.

We prove bounds that are polynomial in n and $1/\sigma$ for the smoothed runningtime with Gaussian perturbations. In addition, our analysis for Euclidean distances is much simpler than the existing smoothed analysis.

Furthermore, we prove a smoothed approximation ratio of $O(\log(1/\sigma))$. This bound is almost tight, as we also provide a lower bound of $\Omega(\frac{\log n}{\log \log n})$ for $\sigma = O(1/\sqrt{n})$. Our main technical novelty here is that, different from existing smoothed analyses, we do not separately analyze objective values of the global and local optimum on all inputs (which only allows for a bound of $O(1/\sigma)$), but simultaneously bound them on the same input.

1 2-Opt and Smoothed Analysis

The traveling salesperson problem (TSP) is one of the classical combinatorial optimization problems. Euclidean TSP is the following variant: given points $X \subseteq [0, 1]^d$, find the shortest Hamiltonian cycle that visits all points in X (also called a *tour*). Even this restricted variant is

^{*}This paper is based on results presented at ISAAC 2013 [22] and ICALP 2015 [19].

NP-hard for $d \ge 2$ [24]. We consider Euclidean TSP with Manhattan and Euclidean distances as well as squared Euclidean distances to measure the distances between points. For the former two, there exist polynomial-time approximation schemes (PTAS) [2,23]. The latter, which has applications in power assignment problems for wireless networks [14], admits a PTAS for d = 2and is APX-hard for $d \ge 3$ [28].

As it is unlikely that there are efficient algorithms for solving Euclidean TSP optimally, heuristics have been developed in order to find near-optimal solutions quickly. One very simple and popular heuristic is 2-opt: starting from an initial tour, we iteratively replace two edges by two other edges to obtain a shorter tour until we have found a local optimum. Experiments indicate that 2-opt converges to near-optimal solutions quite quickly [15,16], but its worst-case performance is bad: the worst-case running-time is exponential even for d = 2 [11] and the approximation ratio is $\Omega(\log n/\log \log n)$ for Euclidean instances [6].

An alternative to worst-case analysis is average-case analysis, where the expected performance with respect to some probability distribution is measured. The average-case runningtime for Euclidean and random metric instances and the average-case approximation ratio for non-metric instances of 2-opt have been analyzed [4, 6, 10, 18]. However, while worst-case analysis is often too pessimistic because it is dominated by artificial instances that are rarely encountered in practice, average-case analysis is dominated by random instances, which have often very special properties with high probability that they do not share with typical instances.

In order to overcome the drawbacks of both worst-case and average-case analysis and to explain the performance of the simplex method, Spielman and Teng invented smoothed analysis [26], a hybrid of worst-case and average-case analysis: an adversary specifies an instance, and then this instance is slightly randomly perturbed. The smoothed performance is the expected performance, where the expected value is taken over the random perturbation. The underlying assumption is that real-world instances are often subjected to a small amount of random noise. This noise can stem from measurement or rounding errors, or it might be a realistic assumption that the instances are influenced by unknown circumstances, but we do not have any reason to believe that these are adversarial. Smoothed analysis often allows more realistic conclusions about the performance than worst-case or average-case analysis. Since its invention, it has been applied successfully to explain the performance of a variety of algorithms. We refer to two surveys for an overview of smoothed analysis in general [21, 27] and a more recent survey about smoothed analysis applied to local search algorithms [20].

1.1 Related Results

Running-time. Englert, Röglin, and Vöcking [11] provided a smoothed analysis of 2-opt in order to explain its performance. They used the one-step model: an adversary specifies n probability density functions $f_1, \ldots, f_n : [0,1]^d \to [0,\phi]$. Then the n points X_1, \ldots, X_n are drawn independently according to the densities f_1, \ldots, f_n , respectively. Here, ϕ is the perturbation parameter. If $\phi = 1$, then the only possibility is the uniform distribution on $[0,1]^d$, and we obtain an average-case analysis. The larger ϕ , the more powerful the adversary. Englert et al. [11] proved that the expected number of iterations of 2-opt is $O(n^4\phi)$ and $O(n^{4+\frac{1}{3}}\phi^{\frac{8}{3}}\log(n\phi))$ for Manhattan and Euclidean distances, respectively. These bounds can be improved slightly by choosing the initial tour with an insertion heuristic. However, if we transfer these bounds to the classical model of points perturbed by Gaussian distributions of standard deviation σ , we obtain bounds that are polynomial in n and $1/\sigma^d$ [11, Section 6], since the maximum density of a d-dimensional Gaussian with standard deviation σ is $\Theta(\sigma^{-d})$. While this is polynomial for any fixed d, it is unsatisfactory that the degree of the polynomial depends on d.

Approximation ratio. Much less is known about the smoothed approximation performance of algorithms. Karger and Onak have shown that multi-dimensional bin packing can be approximated arbitrarily well for smoothed instances [17] and there are frameworks to approximate Euclidean optimization problems such as TSP for smoothed instances [3, 7]. However, these approaches mostly consider algorithms tailored to solving smoothed instances.

With respect to concrete algorithms other than 2-opt, we are only aware of analyses of the jump and lex-jump heuristics for scheduling [5,12].

Englert et al. [11] proved a bound of $O(\phi^{1/d})$. Translated to Gaussians, this yields a bound of $O(1/\sigma)$ if we truncate the Gaussians such that all points lie in a hypercube of constant side length. This result, however, does not explain the approximation performance 2-opt, as the bound is still quite large, even for larger values of σ or smaller values of ϕ .

1.2 Our Contribution

In order to improve our understanding of the practical performance of 2-Opt, we provide an improved smoothed analysis of both its running-time and its approximation ratio. To do this, we use the classical smoothed analysis model: an adversary chooses n points from the *d*-dimensional unit hypercube $[0, 1]^d$, and then these points are independently randomly perturbed by Gaussian random variables of standard deviation σ .

Running-time. The bounds that we prove are polynomial in n and $1/\sigma$. Different to earlier results, the degree of the polynomial is independent of d. As distance measures, we consider Manhattan (Section 3.3), Euclidean (Section 3.5), and squared Euclidean distances (Section 3.4).

The analysis for Manhattan distances is essentially an adaptation of the existing analysis by Englert et al. [11, Section 4.1]. Note that our bound does not have any factor that is exponential in d.

Our analysis for Euclidean distances is considerably simpler than the one by Englert et al., which is rather technical and takes more than 25 pages [11, Section 4.2 and Appendix C].

The analysis for squared Euclidean distances is, to our knowledge, not preceded by a smoothed analysis in the one-step model. Because of the nice properties of squared Euclidean distances and Gaussian perturbations, this smoothed analysis is relatively compact and elegant.

Table 1 summarizes our bounds for the number of iterations.

Approximation ratio. As the earlier smoothed analysis by Englert et al. [11], we provide bounds on the quality of the worst local optimum. While this measure is rather unrealistic and pessimistic, it decouples the analysis from the seeding of the heuristic. Taking into account the seeding would probably severely complicate the analysis.

Our bound of $O(\log(1/\sigma))$ improves significantly upon the direct translation of the bound of Englert et al. [11] to Gaussian perturbations (see Section 4.2 for how to translate the bound to Gaussian perturbations without truncation). It smoothly interpolates between the averagecase constant approximation ratio and the worst-case bound of $O(\log n)$.

In order to obtain our improved bound for the smoothed approximation ratio, we take into account the origins of the points, i.e., their unperturbed positions. Although this information

	Manhattan	Euclidean	squared Euclidean
Englert et al. [11]	$2^{O(d)}n^4\phi$	$O_d\left(n^{4+\frac{1}{3}}\phi^{\frac{8}{3}}\log(n\phi)\right)$	_
general	$O\left(\frac{d^2n^4D_{\max}}{\sigma}\right)$	$Oig(rac{\sqrt{d}n^4D_{ ext{max}}^4}{\sigma^4}ig)$	$O(rac{\sqrt{d}n^4D_{ ext{max}}^2}{\sigma^2})$
$\sigma = O(1/\sqrt{n\log n})$	$Oig(rac{d^2n^4}{\sigma}ig)$	$O(\frac{\sqrt{d}n^4}{\sigma^4})$	$O(\frac{\sqrt{d}n^4}{\sigma^2})$
$\sigma = \Omega(1/\sqrt{n\log n})$	$O(d^2 n^5 \sqrt{\log n})$	$O(\sqrt{d}n^6\log^2 n)$	$O(\sqrt{d}n^5\log n)$
remarks	· · · · ·	only for $d \ge 4$	only for $d \ge 3$; a weaker bound
			holds for $d = 2$ (Theorem 3.14)

Table 1: Our bounds compared to the bounds obtained by Englert et al. [11] for the one-step model. The bounds can roughly be transferred to Gaussian noise by replacing ϕ with σ^{-d} . For convenience, we added our bounds for small and large values of σ : for $\sigma = O(1/\sqrt{n \log n})$, we have $D_{\max} = \Theta(1)$, for larger σ , we have $D_{\max} = \Theta(\sigma\sqrt{n \log n})$. The notation O_d means that terms depending on d are hidden in the O. The remarks are only for our bounds.

is not available to the algorithm, it can be exploited in the analysis. The smoothed analyses of approximation ratios so far [3,5,7,11,12,17] essentially ignored this information. While this simplifies the analysis, being oblivious to the unperturbed positions seems to be too pessimistic. In fact, we see that the bound of Englert et al. [11] cannot be improved beyond $O(1/\sigma)$ by ignoring the positions of the points (Section 4.2). The reason for this limitation is that the lower bound for the global optimum is obtained if all points have the same origin, which corresponds to an average-case rather than a smoothed analysis. On the other hand, the upper bound for the local optimum has to hold for all choices of the unperturbed points, most of which yield higher costs for the global optimum than the average-case analysis. Taking this into account carefully yields our bound of $O(\log(1/\sigma))$ (Section 4.3).

To complement our upper bound, we show that the lower bound of $\Omega(\log n/\log \log n)$ by Chandra et al. [6] remains true for $\sigma = O(1/\sqrt{n})$ (Section 4.4). This implies that a smoothed bound of $o(\log(1/\sigma)/\log\log(1/\sigma))$ is impossible, and, thus, our bound cannot be improved significantly.

2 2-Opt and Smoothing Model

2.1 2-Opt Heuristic for the TSP

Let $X = \subseteq \mathbb{R}^d$ be a set of *n* points. The goal of the TSP is to find a Hamiltonian cycle (also called a tour) *T* through *X* that has minimum length according to some distance measure. In this paper, we consider standard Euclidean distances for both approximation ratio and running-time as well as squared Euclidean distances and Manhattan distances for the running-time.

Given a tour T, a 2-change replaces two edges $\{X_1, X_2\}$ and $\{X_3, X_4\}$ of T by two new edges $\{X_1, X_3\}$ and $\{X_2, X_4\}$, provided that this yields again a tour (this is the case if X_1, X_2, X_3, X_4 appear in this order in the tour) and that this decreases the length of the tour, i.e., $d(X_1, X_2) + d(X_3, X_4) - d(X_1, X_3) - d(X_2, X_4) > 0$, where $d(a, b) = ||a - b||_2$ (Euclidean distances), $d(a, b) = ||a - b||_2$ (Buclidean distances), $d(a, b) = ||a - b||_2$ (Squared Euclidean distances). The 2-

opt heuristic iteratively improves an initial tour by applying 2-changes until it reaches a local optimum. A local optimum is called a 2-optimal tour.

2.2 Smoothing Model

Throughout the rest of this paper, let $\bar{X} = \{x_1, \ldots, x_n\} \in [0, 1]^d$ be a set of n points from the unit hypercube. In the smoothed analysis, these points are chosen by an adversary, and they serve as unperturbed *origins*. Let $Z_1, \ldots, Z_n \sim \mathcal{N}(0, \sigma^2)$ be n independent random variables with mean 0 and standard deviation σ . By slight abuse of notation, \mathcal{N} refers here to the multivariate normal distribution with covariance matrix diag (σ^2) . We obtain the *perturbed* point set $X = \{X_1, \ldots, X_n\} \subseteq \mathbb{R}^d$ by adding $X_i = x_i + Z_i$ for each $i \in [n] = \{1, \ldots, n\}$. We write $X \leftarrow \text{pert}_{\sigma}(\bar{X})$ to make explicit from which point set \bar{X} the points in X are obtained.

We assume that $\sigma \leq 1$ throughout the paper. This is justified by two reasons. First, small σ are the interesting case, i.e., when the order of magnitude of the perturbation is relatively small. Second, smoothed performance guarantees are monotonically decreasing in σ : if we have $\sigma > 1$, then this is equivalent to adversarial instances in $[0, 1/\sigma]^d$ that are perturbed with standard deviation 1. This in turn is dominated by adversarial instances in $[0, 1]^d$ that are perturbed with standard deviation 1, as $[0, 1/\sigma]^d \subseteq [0, 1]^d$. Thus, any upper bound for $\sigma = 1$ (be it for the number of iterations or the approximation ratio) holds also for larger σ .

Let us make a final remark about the smoothing model: while the algorithm itself, the 2-opt heuristic in our case, only sees X and does not know anything about the origins \bar{X} , we can of course exploit the positions of the unperturbed points in the analysis.

3 Smoothed Analysis of the Running-time

In this section, we make the dependence on all parameters (the number n of points, the dimension d, and the perturbation parameter σ) explicit. This means that the O or Ω do not hide any factors, not even factors depending on d, which is often considered as a constant and therefore ignored. (This is also in contrast to our analysis of the approximation ratio, where the hidden constant can indeed depend on d.)

3.1 Probability Theory for the Running-time

In order to get an upper bound for the length of the initial tour, we need an upper bound for the diameter of the point set X. Such an upper bound is also necessary for the analysis of 2-changes with Euclidean distances (Section 3.5). We choose D_{\max} such that $X \subseteq [-D_{\max}, D_{\max}]^d$ with a probability of at least 1 - 1/n!. For fixed d and $\sigma \leq 1$, we can choose $D_{\max} = \Theta(1 + \sigma\sqrt{n\log n})$ according to the following lemma. For $\sigma = O(1/\sqrt{n\log n})$, we have $D_{\max} = \Theta(1)$.

Lemma 3.1. Let $c \ge 2$ be a sufficiently large constant, and let $D_{\max} = c \cdot (\sigma \sqrt{n \log n} + 1)$. Then $\mathbb{P}(X \not\subseteq [-D_{\max}, D_{\max}]^d) \le 1/n!$.

Proof. We have $X \not\subseteq [-D_{\max}, D_{\max}]^d$ only if there is a point x_i and a coordinate of x_i that is perturbed by more than $D_{\max} - 1 \ge c\sigma \cdot \sqrt{n \log n}$. According to Durrett [9, Theorem 1.2.3], the probability that a 1-dimensional Gaussian of standard deviation σ is more than $c\sigma\sqrt{n \log n}$ away from its mean is bounded from above by $2 \cdot \frac{\exp(-c^2 n \log n/2)}{c\sqrt{2\pi n \log n}}$. Thus, the probability that $X \not\subseteq [-D_{\max}, D_{\max}]^d$ is bounded from above by $2dn \cdot \frac{\exp(-c^2 n \log n/2)}{c\sqrt{2\pi n \log n}}$. For sufficiently large c, this is at most 1/n!. Note that the constant c in Lemma 3.1 does not depend on the dimension d.

The following lemma is well known and follows from the fact that the density of a *d*-dimensional Gaussian with standard deviation σ is bounded from above by $(2\sigma)^{-d}$ and the volume of a *d*-dimensional ball of radius ε is bounded from above by $(2\varepsilon)^d$.

Lemma 3.2. Let $a \in \mathbb{R}^d$ be drawn according to a d-dimensional Gaussian distribution of standard deviation σ , and let $B = \{b \in \mathbb{R}^d \mid ||b - c||_2 \leq \varepsilon\}$ be a d-dimensional hyperball of radius ε centered at $c \in \mathbb{R}^d$. Then $\mathbb{P}(a \in B) \leq (\varepsilon/\sigma)^d$.

For $x, y \in \mathbb{R}^d$ with $x \neq y$, let $L(x, y) = \{\xi \cdot (y - x) + x \mid \xi \in \mathbb{R}\}$ denote the straight line through x and y.

Lemma 3.3. Let $a, b \in \mathbb{R}^d$ be arbitrary with $a \neq b$. Let $c \in \mathbb{R}^d$ be drawn according to a *d*-dimensional Gaussian distribution with standard deviation σ . Then the probability that *c* is ε -close to L(a,b), i.e., $\min_{c^* \in L(a,b)} \|c - c^*\|_2 \leq \varepsilon$, is bounded from above by $(\varepsilon/\sigma)^{d-1}$.

Proof. We divide drawing c into drawing a 1-dimensional Gaussian c^* in the direction of a - b and drawing a (d - 1)-dimensional Gaussian c' in the hyperplane orthogonal to a - b and containing c^* . Then the distance of c to L(a, b) is $||c - c^*||_2$. For every c^* , the point c is ε -close to L(a, b) only if c' falls into a (d - 1)-dimensional hyperball of radius ε around c^* in the (d - 1)-dimensional subspace orthogonal to a - b. Now the lemma follows by applying Lemma 3.2.

We need the following lemma in Section 3.5.

Lemma 3.4. Let $f : \mathbb{R} \to \mathbb{R}$ be a differentiable function. Let B be an upper bound for the absolute value of the derivative of f. Let c be distributed according to a Gaussian distribution with standard deviation σ . Let I be an interval of size ε , and let $f(I) = \{f(x) \mid x \in I\}$ be the image of I. Then $\mathbb{P}(c \in f(I)) = O(B\varepsilon/\sigma)$.

Proof. Since the derivative of f is bounded by B, the set f(I) is contained in some interval of length $B\varepsilon$. The lemma follows since the density of c is bounded from above by $O(1/\sigma)$. \Box

The chi distribution [13, Section 8] is the distribution of the Euclidean length of a *d*dimensional Gaussian random vector of standard deviation σ and mean 0. In the following, we denote its density function by χ_d . It is given by

$$\chi_{d,\sigma}(x) = \frac{2^{1-\frac{d}{2}} \cdot \left(\frac{x}{\sigma}\right)^{d-1} \cdot \exp\left(-(x/\sigma)^2/2\right)}{\sigma \cdot \Gamma(d/2)},\tag{1}$$

where Γ denotes the gamma function. We need the following lemma several times.

Lemma 3.5. Assume that $c \in \mathbb{N}$ is a fixed constant and $d \in \mathbb{N}$ is arbitrary with d > c. Then we have

$$\int_0^\infty \chi_{d,\sigma}(x) x^{-c} \, \mathrm{d}x = \frac{2^{-c/2} \Gamma\left(\frac{d-c}{2}\right)}{\sigma^c \cdot \Gamma\left(\frac{d}{2}\right)} = \Theta\left(\frac{1}{d^{c/2} \cdot \sigma^c}\right)$$

Proof. The first equality follows by integration. For the second inequality, we observe $2^{-c/2}$ is a fixed constant (which also never depends on d when we apply the lemma) and that

$$\Gamma(x) = \sqrt{2\pi} x^{x-1/2} e^{-x+\mu(x)}$$

for some function μ with $\mu(x) \in \left[0, \frac{1}{12x}\right]$ according to Stirling's formula [1, 6.1.37]. We have $\frac{d-c}{2} \geq \frac{1}{2}$ as d > c and both are integers. Then

$$\frac{\Gamma(\frac{d-c}{2})}{\Gamma(\frac{d}{2})} = \frac{\sqrt{2\pi} \cdot \left(\frac{d-c}{2}\right)^{\frac{d-c-1}{2}} \cdot \exp\left(-\frac{d-c}{2} + \mu\left(\frac{d-c}{2}\right)\right)}{\sqrt{2\pi} \cdot \left(\frac{d}{2}\right)^{\frac{d-1}{2}} \cdot \exp\left(-\frac{d}{2} + \mu\left(\frac{d-c}{2}\right)\right)} \\
= \frac{\left(\frac{d-c}{2}\right)^{\frac{d-c-1}{2}}}{\left(\frac{d}{2}\right)^{\frac{d-1}{2}}} \cdot \underbrace{\exp\left(\frac{c}{2} + \mu\left(\frac{d-c}{2}\right) - \mu\left(\frac{d}{2}\right)\right)}_{=\Theta(1)} \\
= \underbrace{\left(\frac{d-c}{d}\right)^{\frac{d-1}{2}}}_{=A} \cdot \underbrace{\left(\frac{d-c}{2}\right)^{-\frac{c}{2}}}_{=B} \cdot \Theta(1).$$

Here, the third equality follows from two facts: first, c is a fixed constant, thus $\exp(c/2) = \Theta(1)$. Second, $\frac{d-c}{2}, \frac{d}{2} \geq \frac{1}{2}$. Thus, $\mu(\frac{d-c}{2})$ and $\mu(\frac{d}{2})$ lie between 0 and a constant. Hence, the exponential term is $\Theta(1)$.

Analyzing A and B remains to be done: We have $B \leq (d/2)^{-c/2}$, thus $B = O(d^{-c/2})$. If $d \leq 2c$, then B is bounded from below by a constant and so is $d^{-c/2}$. If $d \geq 2c$, then $B \geq (d/4)^{-c/2} = \Omega(d^{-c/2})$. Hence, $B = \Theta(d^{-c/2})$.

 $B \ge (d/4)^{-c/2} = \Omega(d^{-c/2}).$ Hence, $B = \Theta(d^{-c/2}).$ We have $A = \left(1 - \frac{c}{d}\right)^{\frac{d-1}{2}} \le \exp\left(-\frac{(d-1)\cdot c}{2d}\right) = O(1).$ Distinguishing the cases $d \le 2c$ and d > 2c in the same way as for B yields $A = \Omega(1)$. Thus, $A = \Theta(1).$

The analysis with Euclidean and squared Euclidean distances depends on the distribution of the distance between two points perturbed by Gaussians, where a larger distance between the two points is better for the analysis. The following two lemmas show that, given that larger distance is better, we can replace the distribution of the distance by the corresponding chi distribution. Since we do not know the original positions of the points involved, this allows us to replace unknown distributions by the chi distribution.

Lemma 3.6. Assume that a is drawn according to a d-dimensional Gaussian distribution with standard deviation σ and mean 0. Assume that b is drawn according to a d-dimensional Gaussian distribution with standard deviation σ and mean μ . Then $\|b\|_2$ stochastically dominates $\|a\|_2$, i.e., $\mathbb{P}(\|b\|_2 \leq t) \leq \mathbb{P}(\|a\|_2 \leq t)$ for all $t \in \mathbb{R}$.

Proof. For d = 1, we have the following:

$$\mathbb{P}(\|b\|_{2} \leq t) = \mathbb{P}(b \in [-t,t]) = \mathbb{P}(a \in [-t-\mu,t-\mu])$$
$$= \mathbb{P}(a \in [-t,t]) + \underbrace{\mathbb{P}(a \in [-t-\mu,-t]) - \mathbb{P}(a \in [t-\mu,t])}_{\leq 0} \leq \mathbb{P}(\|a\|_{2} \leq t)$$

Now we prove the lemma for larger d. Since Gaussian distributions are rotation symmetric, we can assume that $\mu = (\delta, 0, ..., 0)$ for some $\delta \ge 0$.

We observe that $||b||_2$ dominates $||a||_2$ if and only if $||b||_2^2$ dominates $||a||_2^2$. Let $b' = a + \mu$. It suffices to prove the lemma for this choice of b', as b' follows the same distribution as b. Fixing a_2, \ldots, a_d fixes also a'_2, \ldots, a'_d . Then $||b'||_2^2$ dominates $||a||_2^2$ if $|a_1 + \delta|$ dominates $||a_1|$. This is true because the lemma holds for d = 1.

Lemma 3.7. Let b be as in Lemma 3.6, and let $h : [0, \infty] \to [0, \infty)$ be a monotonically decreasing function. Let g be the density function of ||b||. Then

$$\int_0^\infty g(x)h(x)\,\mathrm{d} x \le \int_0^\infty \chi_{d,\sigma}(x)h(x)\,\mathrm{d} x,$$

provided that both integrals exist.

Proof. Let a denote the d-dimensional Gaussian random variable of standard deviation σ and mean 0. Then $||a||_2$ has density $\chi_{d,\sigma}$. By Lemma 3.6, $||a||_2$ is dominated by $||b||_2$. This implies that $h(||a||_2)$ dominates $h(||b||_2)$ since h is monotonically decreasing. The lemma follows by observing that the two integrals are the two expected values of $h(||a||_2)$ and $h(||b||_2)$.

For Euclidean and squared Euclidean distances, it turns out to be useful to study $\Delta_{a,b}(c) = d(c,a) - d(c,b)$ for points $a, b, c \in X$. By abusing notation, we sometimes write $\Delta_{i,j}(k)$ instead of $\Delta_{X_i,X_j}(X_k)$ for short. A 2-change that replaces $\{X_1, X_2\}$ and $\{X_3, X_4\}$ by $\{X_1, X_3\}$ and $\{X_2, X_4\}$ improves the tour length by $\Delta_{1,4}(2) - \Delta_{1,4}(3) = \Delta_{2,3}(1) - \Delta_{2,3}(4)$.

3.2 2-Opt State Graph and Linked 2-Changes

The number of iterations that 2-opt needs depends of course heavily on the initial tour and on which 2-change is chosen in each iteration. We do not make any assumptions about the initial tour and about which 2-change is chosen. Following Englert et al. [11], we consider the 2-opt state graph: we have a node for every tour and a directed edge from tour T to tour T'if T' can be obtained by one 2-change. The 2-opt state graph is a directed acyclic graph, and the length of the longest path in the 2-opt state graph is an upper bound for the number of successful iterations that 2-opt needs.

In order to improve the bounds, we also consider pairs of linked 2-changes [11]. Two 2changes form a pair of linked 2-changes if there is one edge added in one 2-change and removed in the other 2-change. Formally, one 2-change replaces $\{X_1, X_2\}$ and $\{X_3, X_4\}$ by $\{X_1, X_3\}$ and $\{X_2, X_4\}$ and the other 2-change replaces $\{X_1, X_3\}$ and $\{X_5, X_6\}$ by $\{X_1, X_5\}$ and $\{X_2, X_6\}$. The edge $\{X_1, X_3\}$ is the one that appears and disappears again (or the other way round). It can happen that $\{X_2, X_4\}$ and $\{X_5, X_6\}$ intersect. Englert et al. [11] called a pair of linked 2-changes a type *i* pair if $|\{X_2, X_4\} \cap \{X_5, X_6\}| = i$. As type 2 pairs, which involve only four nodes, are difficult to analyze because of dependencies, we ignore them. Fortunately, the following lemma states that we will find enough disjoint pairs of linked 2-changes of type 0 and 1 in any sufficiently long sequence of 2-changes.

Lemma 3.8 (Englert et al. [11, Lemma 9 of corrected version]). Every sequence of t consecutive 2-changes contains at least t/7 - 3n/28 disjoint pairs of linked 2-changes of type 0 or type 1.

Following Englert et al. [11, Figure 8], we subdivide type 1 pairs into type 1a and type 1b depending on how $\{X_2, X_4\}$ and $\{X_5, X_6\}$ intersect. One of the 2-changes replaces $\{X_1, X_2\}$ and $\{X_3, X_4\}$ by $\{X_1, X_3\}$ and $\{X_2, X_4\}$. Then other 2-change, i.e., the one that removes the edge $\{X_1, X_3\}$ shared by the linked pair, determines its type:

Type 0: $\{X_1, X_3\}$ and $\{X_5, X_6\}$ are replaced by $\{X_1, X_5\}$ and $\{X_3, X_6\}$.

Type 1a: $\{X_1, X_3\}$ and $\{X_2, X_5\}$ are replaced by $\{X_1, X_5\}$ and $\{X_2, X_3\}$.

Type 1b: $\{X_1, X_3\}$ and $\{X_2, X_5\}$ are replaced by $\{X_1, X_2\}$ and $\{X_3, X_5\}$.

The main idea in the proofs by Englert et al. [11] and also in our proofs is to bound the minimal improvement of any 2-change or the minimal improvement of any pair of linked 2-changes. We denote the smallest improvement of any 2-change by Δ_{\min} and the smallest improvement of any pair of linked 2-changes of type 0, 1a, or 1b by $\Delta_{\min}^{\text{link}}$. It will be clear from the context which distance measure is used for Δ_{\min} and $\Delta_{\min}^{\text{link}}$.

Suppose that the initial tour has a length of at most L, then 2-opt cannot run for more than L/Δ_{\min} iterations and not for more than $\Theta(L/\Delta_{\min}^{\text{link}})$ iterations, provided that $L/\Delta_{\min}^{\text{link}} = \Omega(n^2)$ because of Lemma 3.8.

The following lemma formalizes this and shows how to bound the expected number of iterations using a tail bound for Δ_{\min} or Δ_{\min}^{link} .

Lemma 3.9. Suppose that, with a probability of at least 1 - 1/n!, any tour has a length of at most L. Let $\gamma > 1$. Then

- (1) If $\mathbb{P}(\Delta_{\min} \leq \varepsilon) = O(P\varepsilon)$, then the expected length of the longest path in the 2-opt state graph is bounded from above by $O(PLn \log n)$.
- (2) If $\mathbb{P}(\Delta_{\min} \leq \varepsilon) = O(P\varepsilon^{\gamma})$, then the expected length of the longest path in the 2-opt state graph is bounded from above by $O(P^{1/\gamma}L)$.
- (3) The same bounds as ((1)) and ((2)) hold if we replace Δ_{\min} by $\Delta_{\min}^{\text{link}}$, provided that $PL = \Omega(n^2)$ for Case (1) and $P^{1/\gamma}L = \Omega(n^2)$ for Case (2).

Proof. If the length of the longest tour is longer than L, then we use the trivial upper bound of n!. This contributes only O(1) to the expected value.

Consider the first statement. Let T be the longest path in the 2-opt state graph. If $T \ge t$, then $\Delta_{\min} \le L/t$. Plugging this in and observing that n! is an upper bound for T yields

$$\mathbb{E}(T) = \sum_{t=1}^{n!} \mathbb{P}(T \ge t) \le \sum_{t=1}^{n!} O(PL/t) = O(\log(n!) \cdot PL) = O(PLn \log n)$$

Now consider the second statement, and let T be as above. Let $K = O(L \cdot P^{1/\gamma})$. Then

$$\mathbb{E}(T) = \sum_{t=1}^{n!} \mathbb{P}(T \ge t) \le \sum_{t=1}^{n!} \min\left\{1, O\left(P \cdot (L/t)^{\gamma}\right)\right\}$$
$$= K + PL^{\gamma} \cdot \sum_{t \ge K} O(t^{-\gamma}) = K + PL^{\gamma} \cdot O(K^{1-\gamma}) = O(K).$$

Finally, we consider the third statement. The statement follows from the observation that the maximal number of disjoint pairs of linked 2-changes and the length of the longest path in the 2-opt state graph are asymptotically equal if they are of length at least $\Omega(n^2)$ (Lemma 3.8) and the probability statements become nontrivial only for $t = \Omega(PL)$ in the first and $t = \Omega(P^{1/\gamma}L)$ in the second case.

3.3 Manhattan Distances

The essence of our analysis for Manhattan distances is a straightforward adaptation of the analysis in the one-step model. The extra factor of D_{max} comes from the bound of the initial tour, and the extra factor of d^2 stems from stating the dependence on d explicitly and getting rid of the exponential dependence on d [11, Proofs of Theorem 7 and Lemma 10].

Lemma 3.10. $\mathbb{P}(\Delta_{\min}^{\text{link}} \leq \varepsilon) = O(d^2 n^6 \varepsilon^2 / \sigma^2).$

Proof. We consider a pair of linked 2-changes as described in Section 3.2. The improvement of the first 2-change is

$$\Gamma_1 = \sum_{i=1}^d |x_{1i} - x_{2i}| + |x_{3i} - x_{4i}| - |x_{1i} - x_{3i}| - |x_{2i} - x_{4i}|,$$

where x_{ji} is the *i*-th coordinate of $X_j \in X$. The improvement of the second 2-change is

$$\Gamma_2 = \sum_{i=1}^d |x_{1i} - x_{3i}| + |x_{5i} - x_{6i}| - |x_{1i} - x_{5i}| - |x_{3i} - x_{6i}|.$$

Note that we can have a type 1 pair, i.e., two of the points X_2, X_4, X_5, X_6 can be identical.

Each ordering of the x_{ji} gives rise to a linear combination for Γ_1 and Γ_2 . We have $(6!)^d$ such orderings. If we examine the case distinctions by Englert et al. [11, Lemmas 11, 12, 13] closely, we see that any pair of linear combinations is either impossible (it uses a different ordering of the variables for Γ_1 and Γ_2 or one of Γ_1 and Γ_2 is non-positive, thus the corresponding 2-change is in fact not a 2-change) or we have one variable x_{ji} that has a non-zero coefficient in Γ_1 and a coefficient of 0 in Γ_2 and another variable $x_{j'i'}$ that has a non-zero coefficient in Γ_2 and a coefficient of 0 in Γ_1 . The absolute values of the non-zero coefficients of x_{ji} and $x_{j'i'}$ is 2. Now Γ_1 falls into $(0, \varepsilon]$ only if x_{ji} falls into an interval of length $\varepsilon/2$. This happens with a probability of at most $O(\varepsilon/\sigma)$. By independence, the same holds for Γ_2 and $x_{j'i'}$.

However, we would incur an extra factor of $(6!)^d$ in this way, and we would like to remove all exponential dependence of d. In order to do this, we assume that we know i and i' already. This comes at the expense of a factor of $O(d^2)$ for taking a union bound over the choices of iand i'. We let an adversary fix values for all $x_{j\tilde{i}}$ with $\tilde{i} \neq i, i'$. Since we know i and i', we are left with at most $(6!)^2 = O(1)$ possible linear combinations.

Finally, the lemma follows by taking a union bound over all $O(n^6)$ possible pairs of linked 2-changes.

Theorem 3.11. The expected length of the longest path in the 2-opt state graph corresponding to d-dimensional instances with Manhattan distances is at most $O(d^2n^4D_{\text{max}}/\sigma)$.

Proof. The initial tour has a length of at most $O(ndD_{\text{max}})$ with a probability of at least 1-1/n! by Lemma 3.1. We apply Lemma 3.9 for linked 2-changes using Lemma 3.10 and $\gamma = 2$.

3.4 Squared Euclidean Distances

3.4.1 Preparation

In this section, we have $\Delta_{a,b}(c) = \|c - a\|_2^2 - \|c - b\|_2^2$ for $a, b, c \in \mathbb{R}^d$.

Assume that we have a 2-change that replaces $\{X_1, X_2\}$ and $\{X_3, X_4\}$ by $\{X_1, X_3\}$ and $\{X_2, X_4\}$. The improvement caused by this 2-change is $\Delta_{2,3}(1) - \Delta_{2,3}(4) = \Delta_{1,4}(2) - \Delta_{1,4}(3)$. Given the positions of the four nodes except for a single $i \in \{1, 2, 3, 4\}$, such a 2-change yields a small improvement only if the corresponding $\Delta_{.,.}(i)$ falls into some interval of size ε . The following lemma gives an upper bound for the probability that this happens.

Lemma 3.12. Let $a, b \in \mathbb{R}^d$, $a \neq b$, and let c be drawn according to a Gaussian distribution with standard deviation σ . Let $I \subseteq \mathbb{R}$ be an interval of length ε . Then

$$\mathbb{P}(\Delta_{a,b}(c) \in I) \le \frac{\varepsilon}{4\sigma \cdot \|a - b\|_2}.$$

Proof. Since Gaussian distributions are rotationally symmetric, we can assume without loss of generality that $a = (0, \ldots, 0)$ and $b = (\delta, 0, \ldots, 0)$ with $\delta = ||a - b||_2$. Let $c = (c_1, \ldots, c_d)$. Then $\Delta_{a,b}(c) = c_1^2 - (c_1 - \delta)^2 = 2c_1\delta + \delta^2$. Thus, $\Delta_{a,b}(c) \in I$ if and only if c_1 falls into an interval of length $\frac{\varepsilon}{2\delta}$. Since c_1 is a 1-dimensional Gaussian random variable with a standard deviation of σ , the probability for this is bounded from above by $\frac{\varepsilon}{4\delta\sigma}$ since the maximum density of a 1-dimensional Gaussian of standard deviation σ is bounded from above by $\frac{1}{2\sigma}$.

3.4.2 Single 2-Changes

In this section, we prove a simple bound for the expected number of iterations of 2-opt with squared Euclidean distances. This bounds holds for all $d \ge 2$. In the next section, we improve this bound for the case $d \ge 3$ using pairs of linked 2-changes.

Lemma 3.13. For $d \ge 2$, we have $\mathbb{P}(\Delta_{\min} \in (0, \varepsilon]) = O(\frac{n^4 \varepsilon}{\sigma^2 \sqrt{d}})$.

Proof. Consider a 2-change where $\{X_1, X_2\}$ and $\{X_3, X_4\}$ are replaced by $\{X_1, X_3\}$ and $\{X_2, X_4\}$. Its improvement is given by $\Delta_{2,3}(1) - \Delta_{2,3}(4)$. We let an adversary fix X_3 . Then we draw X_2 . This fixes the distance $\delta = ||X_2 - X_3||_2$. Now we draw X_4 . This fixes $\Delta_{2,3}(4)$. The 2-change yields an improvement of at most ε only if $\Delta_{2,3}(1)$ falls into an interval of size at most ε . According to Lemma 3.12, the probability that this happens is at most $\frac{\varepsilon}{4\delta\sigma}$.

Now let g be the probability density of $\delta = ||X_2 - X_3||$. Then the probability that the 2-change yields an improvement of at most δ is bounded from above by

$$\int_0^\infty g(\delta) \cdot \frac{\varepsilon}{4\sigma\delta} \,\mathrm{d}\delta \le \int_0^\infty \chi_{d,\sigma}(\delta) \cdot \frac{\varepsilon}{4\sigma\delta} \,\mathrm{d}\delta = O\left(\frac{\varepsilon}{\sigma^2\sqrt{d}}\right).$$

The first step is due to Lemma 3.7. The second step is due to Lemma 3.5 using c = 1 and $d \ge 2$. The lemma follows by a union bound over the $O(n^4)$ possible 2-changes.

Theorem 3.14. For all $d \geq 2$, the expected length of the longest path in the 2-opt state graph corresponding to d-dimensional instances with squared Euclidean distances is at most $O(\sqrt{d}D_{\max}^2 n^6 \log(n)/\sigma^2)$.

Proof. With a probability of at least 1 - 1/n!, the instance is contained in a hypercube of sidelength D_{max} . Thus, the longest edge has a length of at most $\sqrt{d}D_{\text{max}}$. Therefore, the initial tour has a length of at most ndD_{max}^2 . We combine this with Lemmas 3.9 and 3.13 to complete the proof.

3.4.3 Pairs of Linked 2-Changes

We can obtain a better bound than in the previous section by analyzing pairs of linked 2changes. With the following three lemmas, we analyze the probability that pairs of linked 2-changes of type 0, 1a, or 1b yield an improvement of at most ε .

Lemma 3.15. For $d \ge 2$, the probability that there exists a pair of type 0 of linked 2-changes that yields an improvement of at most ε is bounded from above by $O(\frac{n^6 \varepsilon^2}{\sigma^4 d})$.

Proof. Consider a fixed pair of type 0 of linked 2-changes involving the six points X_1, \ldots, X_6 as described in Section 3.2. We show that the probability that it yields an improvement of at most ε is at most $O(\varepsilon \sigma^{-2}/\sqrt{d})$. A union bound over the $O(n^6)$ possibilities of pairs of type 0 yields the lemma.

The basic idea is that we restrict ourselves to analyzing $\Delta_{1,4}(3)$ and $\Delta_{1,6}(5)$ only in order to bound the probability that we have a small improvement. In this way, we use the principle of deferred decision to show that we can analyze the improvements of the two 2-changes as if they were independent:

- 1. We let an adversary fix X_1 arbitrarily.
- 2. We draw X_4 , which determines the distance $||X_1 X_4||$.
- 3. We draw X_2 . This fixes the position of the "bad" interval for $\Delta_{1,4}(3)$. Its size is already fixed since we know the positions of X_1 and X_4 . The position of X_3 is still random.
- 4. We draw X_3 . The probability that X_3 assumes a position such that the first 2-change yields an improvement of at most ε is thus at most $\frac{\varepsilon}{4\sigma \cdot \|X_1 X_4\|}$.
- 5. We draw X_6 . This determines the distance $||X_1 X_6||$.
- 6. We draw X_5 . The probability that X_5 assumes a position such that the second 2-change yields an improvement of at most ε is thus at most $\frac{\varepsilon}{4\sigma \cdot \|X_1 X_6\|}$.

Let g be the probability density function of the distance between X_1 and X_4 , and let g' be the probability density function of the distance between X_1 and X_6 . By independence of the points, the probability that both 2-changes of the pair yield an improvement of at most ε is bounded from above by

$$\int_{\delta=0}^{\infty} g(\delta) \cdot \frac{\varepsilon}{4\sigma\delta} \,\mathrm{d}\delta \cdot \int_{\delta=0}^{\infty} g'(\delta) \cdot \frac{\varepsilon}{4\sigma\delta} \,\mathrm{d}\delta.$$

We observe that $\frac{\varepsilon}{4\sigma\delta}$ is monotonically decreasing in δ . Thus, by Lemma 3.7, we can replace g and g' by the density $\chi_{d,\sigma}$ of the chi distribution to get the following upper bound for the probability that a pair of type 0 yields an improvement of at most ε :

$$\left(\int_0^\infty \chi_{d,\sigma}(\delta) \cdot \frac{\varepsilon}{4\delta\sigma} \,\mathrm{d}\delta\right)^2 = O\left(\frac{\varepsilon^2}{\sigma^4 d}\right).$$

Here, we use Lemma 3.5 with c = 1, which is allowed since $d \ge 2$.

Lemma 3.16. For $d \ge 2$, the probability that there exists a pair of type 1a of linked 2-changes that yields an improvement of at most ε is bounded from above by $O(\frac{n^5 \varepsilon^2}{\sigma^4 d})$.

Proof. We can analyze pairs of type 1a in the same way as type 0 pairs in Lemma 3.15. To do this, we analyze $\Delta_{2,3}(4)$ and $\Delta_{1,2}(5)$:

- 1. We let an adversary fix the position of X_2 .
- 2. We draw X_3 . This fixes $||X_2 X_3||$.

- 3. We draw X_1 . This fixes $||X_1 X_2||$. In addition, this fixes the positions of the intervals into which $\Delta_{2,3}(4)$ and $\Delta_{1,2}(5)$ must fall if the first or second 2-change yield an improvement of at most ε .
- 4. We draw X_4 .
- 5. We draw X_5 .

The remainder of the proof is identical to the proof of Lemma 3.15, except that we have to take a union bound only over $O(n^5)$ possible choices.

Lemma 3.17. For $d \ge 3$, the probability that there exists a pair of type 1b of linked 2-changes that yields an improvement of at most ε is bounded from above by $O(\frac{n^5 \varepsilon^2}{\sigma^4 d})$.

Proof. Again, we proceed similarly to Lemma 3.15. We analyze a fixed pair of type 1b, where $\{X_1, X_2\}$ and $\{X_3, X_4\}$ are replaced by $\{X_1, X_3\}$ and $\{X_2, X_4\}$ in one step and $\{X_1, X_3\}$ and $\{X_2, X_5\}$ are replaced by $\{X_1, X_2\}$ and $\{X_3, X_5\}$, and apply a union bound over the $O(n^5)$ possible type 1a pairs. We analyze the probability that $\Delta_{2,3}(4)$ or $\Delta_{2,3}(5)$ assume a bad value.

We draw the points in the following order:

- 1. We fix X_2 .
- 2. We draw X_3 . This fixes the distance $||X_2 X_3||_2$, which is crucial for both 2-changes.
- 3. We draw X_1 .
- 4. We draw X_4 . The probability that the first 2-change yields an improvement of at most ε is at most $\frac{\varepsilon}{4\sigma \cdot ||x_2 x_3||}$.
- 5. We draw X_5 . The probability that the second 2-change yields an improvement of at most ε is at most $\frac{\varepsilon}{4\sigma \cdot ||x_2 x_3||}$.

The main difference to Lemma 3.15 is that the sizes of the bad intervals are not independent. However, once the size of the bad intervals is fixed, we can analyze the probabilities that $\Delta_{2,3}(4)$ or $\Delta_{2,3}(5)$ fall into their bad intervals as independent. Given that $||X_2 - X_3|| = \delta$ is fixed, the probability that the first and the second 2-change yield an improvement of at most ε is bounded from above by $\frac{\varepsilon^2}{16\delta^2\sigma^2}$. Since this is decreasing in δ , we can replace the distribution of δ by the chi distribution to obtain an upper bound according to Lemma 3.7. Thus, using Lemma 3.5 with c = 2 and $d \geq 3$, we obtain the following upper bound for the probability that a pair of type 1b yields an improvement of at most ε :

$$\int_{\delta=0}^{\infty} \chi_{d,\sigma}(\delta) \cdot \frac{\varepsilon^2}{16\delta^2 \sigma^2} \,\mathrm{d}\delta = O\left(\frac{\varepsilon^2}{d\sigma^4}\right).$$

With the three lemmas above, we can obtain a bound on the expected number of iterations of 2-opt for TSP with squared Euclidean distances.

Theorem 3.18. For $d \ge 3$, the expected length of the longest path in the 2-opt state graph corresponding to d-dimensional instances with squared Euclidean distances is at most $O(\frac{n^4\sqrt{d}D_{\max}^2}{\sigma^2})$.

Proof. The probability that any pair of linked 2-changes of type 0, 1a, or 1b yields an improvement of at most ε is bounded from above by $O(\frac{\varepsilon^2 n^6}{\sigma^4 d})$. We apply Lemma 3.9 with $\gamma = 2$ and observe that the initial tour has a length of at most $O(ndD_{\text{max}}^2)$ with a probability of at least 1 - 1/n!.

$$b = (0, \frac{\delta}{2}) \bullet z = (x, y)$$

$$a = (0, -\frac{\delta}{2}) \bullet \sqrt{(x + \frac{\delta}{2})^2 + y^2}$$

Figure 1: The situation for Lemma 3.19.

3.5 Euclidean Distances

3.5.1 Differences of Euclidean Distances

In this section, we have $\Delta_{a,b}(z) = ||z-a||_2 - ||z-b||_2$ for $a, b, z \in \mathbb{R}^d$. Analyzing $||z-a||_2 - ||z-b||_2$ turns out to be more difficult than analyzing $||z-a||_2 - ||z-b||_2$ in the previous section. In particular the case when $||z-a||_2 - ||z-b||_2$ is close to its maximal value of $||a-b||_2$ requires special attention. Intuitively, this is for the following reason: if $\Delta_{a,b}(z) \approx ||a-b||_2$, then z is close L(a,b). Assume that $z \in L(a,b)$ for the moment. Then either z is between a and b, which is fine. Or z is not between a and b. Then moving z in the direction of L(a,b) does not change $\Delta_{a,b}(z)$ at all.

We observe that $\eta = \Delta_{a,b}(z)$ behaves essentially 2-dimensionally: it depends only on the distance of z from L(a, b) (this is x in the following lemma) and on the position of the projection z onto L(a, b) (this is y in the following lemma). It also depends on the distance $||a-b||_2$ between a and b (this is δ in the following lemma, and we had this dependency also in the previous section about squared Euclidean distances). The following lemma makes the connection between x and y explicit for a given η . Figure 1 depicts the situation described in the lemma.

Lemma 3.19. Let $z = (x, y) \in \mathbb{R}^2$, $x \ge 0$, $y \ge 0$. Let $a = (0, -\delta/2)$ and $b = (0, \delta/2)$ be two points at a distance of δ . Let $\eta = ||z - a||_2 - ||z - b||_2$. Then we have

$$y^{2} = \frac{\eta^{2}\delta^{2} + 4\eta^{2}x^{2} - \eta^{4}}{4\delta^{2} - 4\eta^{2}} = \frac{\eta^{2}}{4} + \frac{\eta^{2}x^{2}}{\delta^{2} - \eta^{2}}$$
(2)

for $0 \leq \eta < \delta$ and

$$x^{2} = \frac{y^{2} \cdot \left(4\delta^{2} - 4\eta^{2}\right) + \eta^{4} - \eta^{2}\delta^{2}}{4\eta^{2}} = \frac{y^{2} \cdot \left(\delta^{2} - \eta^{2}\right)}{\eta^{2}} - \frac{\delta^{2} - \eta^{2}}{4}.$$
(3)

for $\delta \geq \eta > 0$. Furthermore, $\eta > \delta$ is impossible.

Proof. The last statement follows from the triangle inequality. We have $\eta = \sqrt{(y+\delta/2)^2 + x^2} - \sqrt{(y-\delta/2)^2 + x^2}$. Rearranging terms and squaring implies

$$\begin{split} \eta^2 + (y - \delta/2)^2 + x^2 + 2\eta \sqrt{(y - \delta/2)^2 + x^2} &= (y + \delta/2)^2 + x^2 \\ \Leftrightarrow 2\eta \sqrt{(y - \delta/2)^2 + x^2} &= 2y\delta - \eta^2. \end{split}$$

Squaring again yields

$$4\eta^{2} \cdot \left((y - \delta/2)^{2} + x^{2} \right) = 4y^{2}\delta^{2} - 4y\delta\eta^{2} + \eta^{4}$$

$$\Leftrightarrow 4\eta^{2} \cdot \left(y^{2} - \delta y + \delta^{2}/4 + x^{2} \right) = 4y^{2}\delta^{2} - 4y\delta\eta^{2} + \eta^{4}$$

$$\Leftrightarrow 4\eta^{2}y^{2} - 4\eta^{2}\delta y + \eta^{2}\delta^{2} + 4\eta^{2}x^{2} = 4y^{2}\delta^{2} - 4y\delta\eta^{2} + \eta^{4}$$

$$\Leftrightarrow 4\eta^{2}y^{2} + \eta^{2}\delta^{2} + 4\eta^{2}x^{2} = 4y^{2}\delta^{2} + \eta^{4}.$$

By rearranging terms again, we obtain

$$y^{2} \cdot \left(4\delta^{2} - 4\eta^{2}\right) = \eta^{2}\delta^{2} + 4\eta^{2}x^{2} - \eta^{4}$$

Using the assumption $\eta < \delta$ or $\eta > 0$ implies the two claims.

As said before, the difficult case in analyzing $\Delta_{a,b}(c) = \eta$ is when $\eta \approx ||a - b||_2$. In terms of the previous lemma, this can only happen if x is small, i.e., if c is close to L(a, b), but not between a and b. The following lemmas makes a quantitative statement about this connection.

Lemma 3.20. Let $a, b, z \in [-D_{\max}, D_{\max}]^d$. Assume that $||a-b|| = \delta$ and that z has a distance of x from L(a, b). Then

$$|\Delta_{a,b}(z)| \le \delta - \frac{x^2 \delta}{32dD_{\max}^2}.$$
(4)

Proof. Let y be the distance of z from $m = \frac{a+b}{2}$, and let $\eta = \Delta_{a,b}(z)$. Then, according to (3), we have

$$x^{2} = \frac{y^{2} \cdot (\delta^{2} - \eta^{2})}{\eta^{2}} - \frac{\delta^{2} - \eta^{2}}{4}$$

We have $\delta \leq \eta$. This and the upper bound $y \leq 2\sqrt{d}D_{\text{max}}$ yields the following weaker bound:

$$x^2 \le \frac{4dD_{\max}^2 \cdot \left(\delta^2 - \eta^2\right)}{\eta^2}.$$
(5)

We distinguish two cases. The first case is that $\eta < \delta/2$. In this case, it suffices to show that $\delta/2 \leq \delta - \frac{x^2\delta}{32dD_{\max}^2}$ in order to prove (4). Since $|x| \leq 2\sqrt{d}D_{\max}$, this holds because $\delta/2 \leq \delta - \delta/2$. The second case is that $\eta \geq \delta/2$. We have

$$\delta^2 - \eta^2 = (\delta + \eta) \cdot (\delta - \eta) \le 2\delta \cdot (\delta - \eta).$$

Replacing $\delta^2 - \eta^2$ by $2\delta \cdot (\delta - \eta)$ in the numerator and η^2 by $\delta^2/4$ in the denominator of (5), we obtain

$$x^{2} \leq \frac{4dD_{\max}^{2} \cdot \left(\delta^{2} - \eta^{2}\right)}{\eta^{2}} \leq \frac{32dD_{\max}^{2} \cdot \left(\eta - \delta\right)}{\delta}$$

Rearranging terms completes the proof.

In order to be able to apply Lemma 3.4, we need the following upper bound on the derivative of y with respect to η , given that x is fixed.

Lemma 3.21. For $x, y \ge 0$, let $y = \sqrt{\frac{\eta^2}{4} + \frac{\eta^2 x^2}{\delta^2 - \eta^2}}$ with $0 < \eta < \delta$. Assume further that $\eta \le \delta - \frac{x^2 \delta}{8 d D_{\max}^2}$ and that $x \le 2\sqrt{d} D_{\max}$. Then the derivative of y with respect to η is bounded by

$$\frac{1}{2} + \frac{32^{3/2} D_{\max}^3 d^{3/2}}{\delta x^2}$$

Proof. The derivative of y with respect to η is given by

$$y' = \frac{\mathrm{d}y}{\mathrm{d}\eta} = \frac{\delta^4 - 2\delta^2\eta^2 + \eta^4 + 4\delta^2x^2}{2 \cdot (\delta^2 - \eta^2)^{3/2} \cdot \sqrt{\delta^2 - \eta^2 + 4x^2}}$$

= $\frac{(\delta^2 - \eta^2)^2 + 4\delta^2x^2}{2 \cdot (\delta^2 - \eta^2)^{3/2} \cdot \sqrt{\delta^2 - \eta^2 + 4x^2}}$
= $\frac{(\delta^2 - \eta^2)^2}{2 \cdot (\delta^2 - \eta^2)^{3/2} \cdot \sqrt{\delta^2 - \eta^2 + 4x^2}} + \frac{4\delta^2x^2}{2 \cdot (\delta^2 - \eta^2)^{3/2} \cdot \sqrt{\delta^2 - \eta^2 + 4x^2}}$
 $\leq \frac{1}{2} + \frac{2\delta^2x^2}{(\delta^2 - \eta^2)^{3/2} \cdot \sqrt{\delta^2 - \eta^2 + 4x^2}}.$

We observe that $y' \ge 0$ for all x and allowed choices of η and δ . For the second term, we have

$$\frac{2\delta^2 x^2}{(\delta^2 - \eta^2)^{3/2} \cdot \sqrt{\delta^2 - \eta^2 + 4x^2}} \le \frac{2\delta^2 x^2}{(\delta^2 - \eta^2)^{3/2} \cdot \sqrt{4x^2}} = \frac{\delta^2 x}{(\delta^2 - \eta^2)^{3/2}}$$

By assumption, we have $\delta - \eta \ge \frac{x^2 \delta}{32 dD_{\max}^2}$ and $\eta \le \delta$. Thus, we have

$$\frac{\delta^2 x}{(\delta^2 - \eta^2)^{3/2}} = \frac{\delta^2 x}{\left((\delta + \eta) \cdot (\delta - \eta)\right)^{3/2}} \le \frac{\delta^2 x}{\delta^{3/2} \cdot \left(\frac{x^2 \delta}{32dD_{\max}^2}\right)^{3/2}} = \frac{32^{3/2} D_{\max}^3 d^{3/2}}{\delta x^2}.$$

Using Lemmas 3.21 and 3.4, we can bound the probability that $\Delta_{a,b}(z)$ assumes a value in an interval of size ε .

Lemma 3.22. Let $d \ge 4$. Let $a, b \in [-D_{\max}, D_{\max}]^d$ be arbitrary, $a \ne b$, and let z be drawn according to a Gaussian distribution with standard deviation σ . Let $\delta = ||a - b||_2$. Let I be an interval of length ε . Then

$$\mathbb{P}(\Delta_{a,b}(z) \in I) \le O\left(\frac{\varepsilon D_{\max}^3 \sqrt{d}}{\sigma^3 \delta}\right) + \mathbb{P}(z \notin [-D_{\max}, D_{\max}]^d).$$

Proof. We assume throughout this proof that $z \in [-D_{\max}, D_{\max}]^d$. The case that this is not satisfied is taken care of by the second term in the upper bound for the probability in the statement of the lemma.

Let x denote the distance of z to L(a, b), and let y denote the position of the projection of z onto L(a, b). First, let us assume that x is fixed. Then, by Lemmas 3.21 and 3.4, the probability that $\Delta_{a,b}(z) \in I$ is bounded from above by

$$O\left(\left(1+\frac{D_{\max}^3 d^{3/2}}{\delta x^2}\right)\cdot\frac{\varepsilon}{\sigma}\right).$$

Here, the requirements of Lemma 3.21 are satisfied because of Lemma 3.20, or we have $z \notin [-D_{\max}, D_{\max}]^d$.

We observe that this probability is decreasing in x. Thus, in order to get an upper bound for the probability with random x, we can use the (d-1)-dimensional chi distribution for xaccording to Lemma 3.7. We obtain

$$\int_{x=0}^{\infty} \chi_{d-1,\sigma}(x) \cdot O\left(\left(1 + \frac{D_{\max}^{3} d^{3/2}}{\delta x^{2}}\right) \cdot \frac{\varepsilon}{\sigma}\right) dx$$
$$= O\left(\frac{\varepsilon}{\sigma}\right) + \int_{x=0}^{\infty} \chi_{d-1,\sigma}(x) \cdot O\left(\frac{D_{\max}^{3} d^{3/2}\varepsilon}{\sigma \delta x^{2}}\right) dx = O\left(\frac{\varepsilon}{\sigma} + \frac{D_{\max}^{3} \sqrt{d\varepsilon}}{\sigma^{3}\delta}\right).$$

by Lemma 3.5 using c = 2 and $d - 1 \ge 3$. Since $\delta \le 2\sqrt{d}D_{\text{max}}$, the lemma follows.

3.5.2 Analysis of Pairs of 2-Changes

We immediately go to pairs of linked 2-changes, as these yield the better bounds.

Lemma 3.23. For $d \ge 4$, the probability that a pair of linked 2-changes of type 0 yields an improvement of at most ε or some point lies outside $[-D_{\max}, D_{\max}]^d$ is bounded from above by

$$O\left(\frac{n^6\varepsilon^2 D_{\max}^6}{\sigma^8}\right).$$

Proof. We proceed similarly as in the proof of Lemma 3.15 for type 0 pairs for squared Euclidean distances. We draw the points of a fixed pair of linked 2-changes as in the proof of Lemma 3.15.

In the same way as in the proof of Lemma 3.15, using Lemma 3.22 instead of Lemma 3.12, we obtain that the probability that one fixed of the two 2-changes yields an improved of at most ε is bounded from above by

$$\int_{\delta=0}^{\infty} \chi_{d,\sigma}(\delta) \cdot O\left(\frac{\varepsilon D_{\max}^3 \sqrt{d}}{\sigma^3 \delta}\right) \, \mathrm{d}\delta = O\left(\frac{\varepsilon D_{\max}^3}{\sigma^4}\right).$$

Here, we applied Lemma 3.5 with c = 1.

Again in the same way as in the proof of Lemma 3.15, we can analyze both 2-changes of the type 0 pair is if they are independent. Finally, the lemma follows by a union bound over the $O(n^6)$ possibilities for a type 0 pair.

Lemma 3.24. For $d \ge 4$, the probability that a pair of linked 2-changes of type 1a yields an improvement of at most ε or some point lies outside $[-D_{\max}, D_{\max}]^d$ is bounded from above by

$$O\left(\frac{n^5\varepsilon^2 D_{\max}^6}{\sigma^8}\right).$$

Proof. The lemma can be proved in the same way as Lemma 3.16 with differences analogous to the proof of Lemma 3.23. \Box

Lemma 3.25. For $d \ge 4$, the probability that a pair of linked 2-changes of type 1b yields an improvement of at most ε or some point lies outside $[-D_{\max}, D_{\max}]^d$ is bounded from above by

$$O\left(\frac{n^5\varepsilon^2 D_{\max}^6}{\sigma^8}\right).$$

Proof. Similar to the proof of Lemma 3.17 and using Lemma 3.22, the probability that the two 2-changes of the pair both yield an improvement of at most ε is bounded from above by

$$\int_{\delta=0}^{\infty} \chi_{d,\sigma}(\delta) \cdot O\left(\frac{\varepsilon D_{\max}^3 \sqrt{d}}{\sigma^3 \delta}\right)^2 d\delta = \int_{\delta=0}^{\infty} \chi_{d,\sigma}(\delta) \cdot O\left(\frac{\varepsilon^2 D_{\max}^6 d}{\sigma^6 \delta^2}\right) d\delta.$$

Now the lemma follows by applying Lemma 3.5 with c = 2.

Theorem 3.26. For $d \ge 4$, the expected length of the longest path in the 2-opt state graph corresponding to d-dimensional instances with Euclidean distances is at most $O(\frac{\sqrt{dn^4}D_{\max}^4}{\sigma^4})$.

Proof. We have $\mathbb{P}(\Delta_{\min}^{\ln k} \leq \varepsilon) = O(\frac{n^6 \varepsilon^2 D_{\max}^6}{\sigma^8})$ by Lemmas 3.23, 3.24, and 3.25. If all points are in $[-D_{\max}, D_{\max}]^d$, then the longest edge has a length of $O(\sqrt{d}D_{\max})$. Thus, the initial tour has a length of at most $O(n\sqrt{d}D_{\max})$. Plugging this into Lemma 3.9 yields the result. \Box

4 Smoothed Analysis of the Approximation Ratio

4.1 Technical Preparation

The following standard lemma provides a convenient way to bound the deviation of a perturbed point from its mean in the two-step model.

Lemma 4.1 (Chi-square bound [26, Cor. 2.19]). Let x be a Gaussian random vector in \mathbb{R}^d of standard deviation σ centered at the origin. Then, for $t \geq 3$, we have $\mathbb{P}(||x|| \geq \sigma 3\sqrt{d \ln t}) \leq t^{-2.9d}$.

To give large-deviation bounds on sums of independent variables with bounded support, we will make use of a standard Chernoff-Hoeffding bound.

Lemma 4.2 (Chernoff-Hoeffding Bound [8, Exercise 1.1]). Let $X := \sum_{i=1}^{n} X_i$, where $X_i, i = 1, \ldots, n$ are independently distributed in [0,1], and $\mu_L \leq \mathbb{E}(X) \leq \mu_H$. Then, for $0 \leq \varepsilon \leq 1$, we have

$$\mathbb{P}(X > (1 + \varepsilon) \cdot \mu_H) \le \exp(-(\varepsilon^2/3) \cdot \mu_H),$$

$$\mathbb{P}(X < (1 - \varepsilon) \cdot \mu_L) \le \exp(-(\varepsilon^2/2) \cdot \mu_L).$$

Throughout this paper, we assume that the dimension $d \ge 2$ is a fixed constant. Given a sequence of points $X_1, \ldots, X_n \in \mathbb{R}^d$, we call a collection $T \subseteq [n] \times [n]$ of edges a *tour*, if T is connected and every $i \in [n] = \{1, \ldots, n\}$ has in- and outdegree exactly one in T. Note that we consider directed tours, which is useful in the analysis in this chapter, but our distances are always symmetric.

Given any collection of edges S, its length is denoted by $L(S) = \sum_{(u,v)\in S} d(u,v)$, where d(u,v) denotes the Euclidean distance $||X_u - X_v||$ between points X_u and X_v .

We call a collection $T \subseteq [n]^2$ a partial 2-optimal tour if T is a subset of a tour and $d(u, v) + d(w, z) \leq d(u, w) + d(v, z)$ holds for all edges $(u, v), (w, z) \in T$. Our main interests are the traveling salesperson functional $\mathsf{TSP}(X) := \min_{\mathsf{tour } T} L(T)$ as well as the functional $\mathsf{2OPT}(X) := \max_{2-\mathsf{optimal tour } T} L(T)$ that maps the point set X to the length of the longest 2-optimal tour through X.

We note that the results in Section 4.2 hold for metrics induced by arbitrary norms in \mathbb{R}^d (Lemma 4.4 and 4.5) or typical ℓ_p norms (Lemma 4.6 and 4.7), not only for the Euclidean metric. We conjecture that also the upper bound in Section 4.3 holds for more general metrics, while the lower bound in Section 4.4 is probably specific for the Euclidean metric. Still, we think that the construction can be adapted to work for most natural metrics.

For obtaining lower bounds on the length of optimal tours, we consider the boundary functional $\mathsf{TSP}_B(X)$ that attains the length of the shortest tour through all points in X that is allowed to traverse the boundary of $[0, 1]^d$ at zero cost. For a proof of the following lemma, we refer to the monograph by Yukich [29].

Lemma 4.3 (Boundary Functional [29, Lemma 3.7]). There is a constant C > 0 such that for all sets $X \subseteq [0,1]^d$ of n points, we have $\mathsf{TSP}_B(X) \ge \mathsf{TSP}(X) - Cn^{\frac{d-2}{d-1}}$.

4.2 Length of 2-Optimal Tours under Perturbations

In this section, we provide an upper bound for the length of any 2-optimal tour and a lower bound for the length of any global optimum. These two results yield an upper bound of $O(1/\sigma)$ for the approximation ratio.

Chandra et al. [6] proved a bound on the worst-case length of 2-optimal tours that, in fact, already holds for the more general notion of *partial* 2-optimal tours. For an intuition why this is true, let us point out that their proof strategy is to argue that not too many long arcs in a tour may have similar directions due to the 2-optimality of the edges, while short edges do not contribute much to the length. The claim then follows from a packing argument. It is straight-forward to verify that it is never required that the collection of edges is closed or connected.

Lemma 4.4 (Length of partial 2-optimal tours [6, Theorem 5.1], paraphrased). There exists a constant c_d such that for every sequence X of n points in $[0,1]^d$, any partial 2-optimal tour has length less than $c_d \cdot n^{1-1/d}$.

While this bound directly applies to any perturbed instance under the one-step model, Gaussian perturbations fail to satisfy the premise of bounded support in $[0, 1]^d$. However, Gaussian tails are sufficiently light to enable us to translate the result to the two-step model by carefully taking care of outliers.

Lemma 4.5. There exists a constant b_d such that for any $\sigma \leq 1$ the following statement holds. For any \bar{X} , the probability that any partial 2-optimal tour on $X \leftarrow \text{pert}_{\sigma}(\bar{X})$ has length greater than $b_d \cdot n^{1-1/d}$, i.e., $2\mathsf{OPT}(X) > b_d \cdot n^{1-1/d}$, is bounded by $\exp(-\Omega(\sqrt{n}))$. Furthermore,

$$\mathbb{E}_{X \leftarrow \text{pert}_{\sigma}(\bar{X})} (2\mathsf{OPT}(X)) \le b_d \cdot n^{1-1/d}.$$

Proof. By translation, assume without loss of generality that the input points are contained in $[-1/2, 1/2]^d$. We define cubes $C_1 = \ell_1 [-1, 1]^d$, $C_2 = \ell_2 [-1, 1]^d$, ... with $\ell_k := 6\sqrt{kd \ln 3}$. The side length of cube C_k is $2\ell_k$. We consider the partitioning of \mathbb{R}^d into the regions C_1 and $C_i \setminus C_{i-1}$ for $i \ge 2$. For some cube C and any tour T, let $E_C(T)$ denote the edges in T that are completely contained in C. For any tour T, the sequence E_1, E_2, \ldots defined by $E_1 := E_{C_1}(T)$ and $E_k := E_{C_k}(T) \setminus E_{C_{k-1}}(T)$, for $k \ge 2$, partitions the edges of T. Thus, $L(T) = \sum_{k=1}^{\infty} L(E_k)$.

For any outcome of the perturbed points, let T be the longest 2-optimal tour. Then, each E_k is a partial 2-optimal tour in C_k . Let n_k be the (random) number of points in $\mathbb{R}^d \setminus C_{k-1}$,

which is an upper bound on the number of points in $C_k \setminus C_{k-1}$. At most $3n_k$ vertices are incident to the edges E_k , since each such edge is incident to at least one endpoint in $C_k \setminus C_{k-1}$ and every point has degree 2 in T. Since $C_k = \ell_k [-1, 1]^d$ is a translated unit cube scaled by $2\ell_k$, Lemma 4.4 yields $L(E_k) \leq c_d \cdot (2\ell_k)(3n_k)^{1-1/d}$.

Observe that X_i is not contained in C_k only if its origin has been perturbed by noise of length at least $\ell_k/2$. Thus, let $Z \sim \mathcal{N}(0, \sigma^2)$ and note that $\sigma \leq 1$ implies that $\ell_k/2 \geq 3\sqrt{dk \ln 3\sigma}$. Hence, for each point X_i , Lemma 4.1 yields

$$\mathbb{P}(X_i \notin C_k) \le \mathbb{P}\left(\|Z\| \ge \frac{\ell_k}{2} \right) \le 3^{-(2.9d)k}.$$

By linearity of expectation, we conclude that $\mathbb{E}(n_k) \leq n 3^{-(2.9d)(k-1)}$ for $k \geq 1$. This yields

$$\mathbb{E}(L(T)) = \sum_{k=1}^{\infty} \mathbb{E}(L(E_k)) \le \sum_{k=1}^{\infty} c_d \cdot (2\ell_k) (3 \mathbb{E}(n_k))^{1-1/d}$$
$$\le c_d \cdot 12\sqrt{d\ln 3} \cdot (3n)^{1-1/d} \left(\sum_{k=1}^{\infty} \sqrt{k} 3^{-2.9(d-1)(k-1)}\right) = O(n^{1-1/d}),$$

where we used Jensen's inequality for the first inequality.

To derive tail bounds for the length of any 2-optimal tour, let $N_k := n3^{-2.9d(k-1)}$ be the upper bound on $\mathbb{E}(n_k)$ derived above. By the Chernoff bound (Lemma 4.2), we have

$$\mathbb{P}(n_k \ge 2N_k) \le \exp(-N_k/3).$$

This guarantee is only strong as long as N_k is sufficiently large. Hence, we regard this guarantee only for $1 \le k \le k_1$, where k_1 is chosen such that $\sqrt{n} \le N_{k_1} \le 3^{2.9d} \sqrt{n}$. Assume that $n_k \le 2N_k$ for all $1 \le k \le k_1$. Then, analogously to the above calculation, the contribution of E_1, \ldots, E_{k_1} is bounded by

$$\sum_{k=1}^{k_1} L(E_k) \le \sum_{k=1}^{k_1} c_d \cdot (2\ell_k) (6N_k)^{1-1/d}$$
$$\le c_d \cdot 12\sqrt{d\ln 3} \cdot (6n)^{1-1/d} \left(\sum_{k=1}^{\infty} \sqrt{k} 3^{-2.9(d-1)(k-1)}\right) = O(n^{1-1/d}).$$

Let p_1 denote the probability that some $1 \le k \le k_1$ fails to satisfy $n_k \le 2N_k$. Then,

$$p_1 \le \sum_{k=1}^{k_1} \mathbb{P}(n_k > 2N_k) \le k_1 \exp(-N_{k_1}/3) = \exp(-\Omega(\sqrt{n})).$$

Let us continue assuming that all $1 \le k \le k_1$ satisfy $n_k \le 2N_k$. Since in particular $n_{k_1} \le 2N_{k_1}$, at most $n_{k_1} \le 2 \cdot 3^{2.9d} \sqrt{n}$ vertices remain outside C_{k_1-1} . Let $k_2 := \lceil \sqrt{n} \rceil$. By a union bound,

$$p_2 := \mathbb{P}(\exists j : X_j \notin C_{k_2}) \le n3^{-2.9d(k_2-1)} = \exp(-\Omega(\sqrt{n})).$$

Assume that the corresponding event holds (i.e., $X \subseteq C_{k_2}$), then the remaining points outside C_{k_1-1} (and hence, outside C_{k_1}) are contained in C_{k_2} . We conclude that, with probability at

least $1 - (p_1 + p_2) = 1 - \exp(-\Omega(\sqrt{n}))$, we have

$$\sum_{k=k_1+1}^{\infty} L(E_k) = \sum_{k=k_1+1}^{k_2} L(E_k) \le c_d \cdot (2\ell_{k_2})(6N_{k_1})^{1-1/d}$$
$$= O(\sqrt{k_2}(N_{k_1})^{1-1/d}) = O(n^{\frac{1}{4} + \frac{1}{2}\left(1 - \frac{1}{d}\right)}) = O(n^{1-1/d})$$

This finishes the claim, since we have shown that with probability $1 - \exp(-\Omega(\sqrt{n}))$, both the contribution of E_1, \ldots, E_{k_1} and E_{k_1+1}, \ldots is bounded by $O(n^{1-1/d})$.

We complement the bound above by a lower bound on tour lengths of perturbed inputs, making use of the following result by Englert et al. [11] for the one-step model.

Lemma 4.6 (Englert et al. [11, Proof of Theorem 1.4]). Let X_1, \ldots, X_n be a ϕ -perturbed instance. Then with probability $1 - \exp(-\Omega(n))$, any tour on X_1, \ldots, X_n has length at least $\Omega(n^{1-1/d}/\sqrt[d]{\phi})$.

It also follows from their results that this bound translates to the two-step model consistently with the intuitive correspondence of $\phi \sim \sigma^{-d}$ between the one-step and the two-step model.

Lemma 4.7. Let X_1, \ldots, X_n be an instance of points in the unit cube perturbed by Gaussians of standard deviation $\sigma \leq 1$. Then with probability $1 - \exp(-\Omega(n))$ any tour on X_1, \ldots, X_n has length at least $\Omega(\sigma n^{1-1/d})$.

Proof. We summarize the arguments of Englert et al. [11, Section 6] first, who considered truncated Gaussian perturbations: Here, we condition the Gaussian perturbation Z_i for each input point X_i to be contained in $A := [-\alpha, \alpha]^d$ for some $\alpha \ge 1$. Conditioned on this event, the resulting input instance is contained in the cube $C := [-\alpha, 1 + \alpha]^d$. By straight-forward calculations, the conditional distribution of each point in C has maximum density bounded by $O(\alpha^d/\sigma^d)$. Moreover, the probability that the condition fails for a single point is bounded by $\mathbb{P}(Z_i \notin A) \le d\sigma \exp(-\alpha^2/(2\sigma^2))$ for all *i*. Thus, by choosing $\alpha \ge 1$ sufficiently large, each point has at least constant probability to satisfy the condition $Z_i \in A$.

Given any instance (with Gaussian perturbations which are not truncated), first reveal the (random) subinstance of those points for which the condition $Z_i \in A$ is satisfied and let n' be the number of such points. By the Chernoff bound (Lemma 4.2), and $\mathbb{P}(Z_i \in A) = \Omega(1)$, we have $n' \geq c \cdot n$ for some c > 0 with probability at least $1 - \exp(-\Omega(n))$. If this event occurs, we obtain a random instance of $n' \geq cn$ points and maximum density $\phi = O(\alpha^d/\sigma^d)$. Hence an application of Lemma 4.6 yields that, for some constant c' > 0, the probability that a tour of length less than $c' \cdot (n')^{1-1/d} / \sqrt[d]{\phi} = O((\sigma/\alpha)n^{1-1/d}) = O(\sigma n^{1-1/d})$ exists is at most $\exp(-\Omega(n)) + \exp(-\Omega(n')) = \exp(-\Omega(n))$.

Note that Lemmas 4.5 and 4.7 almost immediately yield the following bound on the approximation performance for the two-step model. (The large-deviation bound is immediate. For the expected approximation ratio, we make use of the worst-case bound of $O(\log n)$, given in Lemma 4.9 below.)

Observation 4.8. Let X_1, \ldots, X_n be an instance of points in the unit cube perturbed by Gaussians of standard deviation $\sigma \leq 1$. Then the approximation performance of 2-Opt is bounded by $O(1/\sigma)$ in expectation and with probability $1 - \exp(-\Omega(\sqrt{n}))$.

We remark that this bound is best possible for an analysis of perturbed instances that separately bounds the lengths of any 2-optimal tour from above and gives a lower bound on any optimal tour. To see this, we argue that Lemma 4.6, Lemma 4.4 (even under ϕ -perturbed input), Lemma 4.7 and Lemma 4.5 cannot be improved in general. This is straight-forward for Lemma 4.6, since *n* points distributed uniformly at random in a cube of volume $1/\phi$ always have, by scaling and Lemma 4.4, a tour of length $O(n^{1-1/d}/\sqrt[d]{\phi})$. Hence, the lower bound on optimal tours on perturbed instances is tight. To see that the upper bound on any 2-optimal tour is tight, take *n* uniformly distributed points that have, by Lemma 4.6, an optimal tour of length $\Omega(n^{1-1/d})$ with high probability and thus also in expectation.

Naturally, this transfers to the case of Gaussian perturbations, albeit more technical to verify: If we place n identical points in $[0, 1]^d$, say at the origin, and perturb them with Gaussians of standard deviation σ , then we may without loss of generality scale the unit cube to $[0, 1/\sigma]^d$ and perturb the points with standard deviation 1 instead. By Lemma 4.5, any 2-optimal tour and, thus, any optimal tour on these points has a length of $O(n^{1-1/d})$ on the scaled instance, since the origins are still contained in the unit cube. Thus, the optimal tour on the original instance has a length of at most $O(\sigma \cdot n^{1-1/d})$ in expectation and with high probability.

We only sketch that 2-optimal tours can have a length of at least $\Omega(n^{1-1/d})$: We distribute the *n* (unperturbed) points into $1/\sigma^d$ groups of $\sigma^d n$ points each, and we partition the cube $[0,1]^d$ into $1/\sigma^d$ subcubes of equal side length. Let c > 0 be a constant such that with high probability, at least $c\sigma^d n$ points of a group remain in their subcube after perturbation. We call these points successful. Since successful points are identically distributed, conditioned on falling into a compact set, the shortest tour through these (at least) $c\sigma^d n$ points has a length of at least $\sigma \cdot c'(\sigma^d n)^{1-1/d} = c'\sigma^d n^{1-1/d}$ for some other constant c' > 0 [29]. (This is just a scaled version of perturbing and truncating a Gaussian of standard deviation 1 to a unit hypercube, which would result in a tour length of $m^{1-1/d}$ for m points.) By closeness of the tour on all points to the boundary functional and geometric superadditivity of the boundary functional (see Yukich [29] for details), it follows that the optimal tour on all successful points is at least $\Omega((1/\sigma^d) \cdot \sigma^d n^{1-1/d}) = \Omega(n^{1-1/d})$.

4.3 Upper Bound on the Approximation Performance

In this section, we establish an upper bound on the approximation performance of 2-Opt under Gaussian perturbations. We achieve a bound of $O(\log 1/\sigma)$. Due to the lower bound presented in Section 4.4, improving the smoothed approximation ratio to $o(\log(1/\sigma)/\log\log(1/\sigma))$ is impossible. Thus, our bound is almost tight.

As noted in the previous section, to beat $O(1/\sigma)$ it is essential to exploit the structure of the unperturbed input. This will be achieved by classifying edges of a tour into *long* and *short* edges and bounding the length of long edges by a (worst-case) global argument and short edges locally against the partial optimal tour on subinstances (by a reduction to an (almost-)average case). The local arguments for short edges will exploit how many unperturbed origins lie in the vicinity of a given region.

The global argument bounding long edges follows from the worst-case $O(\log n)$ bound on the worst-case approximation performance [6] that we rephrase here for our purposes.

Lemma 4.9 ([6, Proof of Theorem 4.3]). Let T be a 2-optimal tour and OPT denote the length of the optimal traveling salesperson tour T_{OPT} . Let T_i contain the set of all edges in T whose length is in $[\text{OPT}/2^i, \text{OPT}/2^{i-1}]$. Then $L(T_i) = O(\text{OPT})$. In particular, it follows that $L(T) = O(\log n) \cdot \text{OPT}$.

In the proof of our bound of $O(\log 1/\sigma)$, the above lemma accounts for all edges of length $[\Omega(\sigma), O(1)]$. A central idea to bound all shorter edges is to apply the one-step model result to small parts of the input space. In particular, we will condition sets of points to be perturbed into cubes of side length σ . The following technical lemma helps to capture what values of ϕ suffice to express the conditional density function of these points depending on the distance of their unperturbed origins to the cube. This allows for appealing to the one-step model result of Lemma 4.6.

Lemma 4.10. Let $c \in [0, \sigma]^d$ and $k = (k_1, \ldots, k_d) \in \mathbb{N}_0^d$. Let Y be the random variable $X \sim \mathcal{N}(c, \sigma^2)$ conditioned on $X \in Q := [k_1\sigma, (k_1+1)\sigma] \times \cdots \times [k_d\sigma, (k_d+1)\sigma]$ and f_Y be the corresponding probability density function. Then f_Y is bounded from above by $\exp(||k||_1 + (3/2)d)\sigma^{-d}$.

Proof. Let $f_X(x) = \frac{1}{(2\pi)^{d/2}\sigma^d} \cdot \exp(-\frac{\|x-c\|^2}{2\sigma^2})$ be the probability density function of X. Let $q := \operatorname{argmin}_{z \in Q} \|z - c\|$ be the point in Q that is closest to c. Then, since $f_X(x)$ is rotationally invariant around c and decreasing in $\|x - c\|$, the density $f_X(x)$ inside Q is maximized at x = q. Likewise, $q' := \operatorname{argmax}_{z \in Q} \|z - c\|$ minimizes the density inside Q. Since Q is a $(\sigma \times \cdots \times \sigma)$ -cube in $\mathbb{R}_{\geq 0}$, $\|q' - c\| \leq \|(q + \sigma \mathbf{1}) - c\|$, where $\mathbf{1} = (1, \ldots, 1) \in \mathbb{R}^d$ denotes the all-ones vector. Given $g(q) := \frac{f_X(q)}{f_X(q+\sigma \mathbf{1})}$, we can thus bound the conditional probability density function f_Y for $x \in Q$ by

$$f_Y(x) = \frac{f_X(x)}{\int_Q f_X(y)dy} \le \frac{f_X(x)}{f_X(q')\operatorname{vol}(Q)} \le \frac{f_X(q)}{f_X(q+\sigma\mathbf{1})} \cdot \frac{1}{\sigma^d} = \frac{g(q)}{\sigma^d}.$$

It remains to bound, for $x \in Q$,

$$g(x) = \frac{f(x)}{f(x+\sigma\mathbf{1})} = \exp\left(\frac{\|(x-c)+\sigma\mathbf{1}\|^2 - \|x-c\|^2}{2\sigma^2}\right) = \exp\left(\frac{\|x-c\|_1}{\sigma} + \frac{d}{2}\right).$$

Since for all $x \in Q$, $||x - c||_1 \le (||k||_1 + d)\sigma$, we can bound $g(x) \le \exp(||k||_1 + (3/2)d)$, yielding the claim.

The main result of this section is the following theorem, which will be proved in the remainder of the section.

Theorem 4.11. Let $X = (X_1, \ldots, X_n)$ be an instance of points in $[0, 1]^d$ perturbed by Gaussians of standard deviation $\sigma \leq 1$. With probability $1 - \exp(-\Omega(n^{1/2-\varepsilon}))$ for any constant $\varepsilon > 0$, we have $2\mathsf{OPT}(X) = O(\log(1/\sigma)) \cdot \mathsf{TSP}(X)$. Furthermore,

$$\mathbb{E}\left(\frac{2\mathsf{OPT}(X)}{\mathsf{TSP}(X)}\right) = O(\log(1/\sigma)).$$

Since the approximation performance of 2-Opt is bounded by $O(\log n)$ in the worst-case, we may assume that $1/\sigma = O(n^{\varepsilon})$ for all constant $\varepsilon > 0$, since otherwise our smoothed result is superseded by Lemma 4.9. Furthermore, we may also assume that $1/\sigma = \omega(1)$, since otherwise Observation 4.8 already yields the result. In what follows, let T_{OPT} and T be any optimal and longest 2-optimal, respectively, traveling salesperson tour on X_1, \ldots, X_n . Furthermore, we let $\text{OPT} = L(T_{\text{OPT}})$ denote the length of the shortest traveling salesperson tour.

4.3.1 Outliers and Long Edges

We will first show that the contribution of almost all points outside $[0,1]^d$ is bounded by $O(\sigma n^{1-1/d})$ with high probability and in expectation, similar to Lemma 4.5. For this, we define growing cubes $A_i := [-a_i, 1 + a_i]^d$, where we set $a_i := 3\sigma\sqrt{di\ln(3/\sigma)}$ for $i \ge 1$ and $A_0 = [0,1]^d$. Let n_i be the number of points not contained in A_{i-1} . For every point X_j , Lemma 4.1 with $t := (3/\sigma)^i$ bounds $\mathbb{P}(X_j \notin A_{i-1}) \le (\sigma/3)^{2.9d(i-1)}$ (note that we have chosen the a_i such that $t \ge 3$). Thus, $\mathbb{E}(n_i) \le n(\sigma/3)^{2.9d(i-1)}$. We define E_i as the set of edges of the longest 2-optimal tour T contained in A_i with at least one endpoint in $A_i \setminus A_{i-1}$. We first bound the contribution of the E_i with $i \ge 2$.

Lemma 4.12. With probability $1 - \exp(-\Omega(n^{1/2-\varepsilon}))$ for any constant $\varepsilon > 0$, we have

$$\sum_{i=2}^{\infty} L(E_i) = O\left(\sigma n^{1-1/d}\right).$$

In addition, we have $\mathbb{E}\left(\sum_{i=2}^{\infty} L(E_i)\right) = O\left(\sigma n^{1-1/d}\right)$.

Proof. The proof is analogous to the proof of Lemma 4.5. Linearity of expectation, Lemma 4.4, and Jensen's inequality yield

$$\begin{split} \sum_{i=2}^{\infty} \mathbb{E}(L(E_i)) &\leq \sum_{i=2}^{\infty} c_d \cdot (3 \mathbb{E}(n_i))^{1-1/d} (1+2a_i) \\ &\leq \sum_{i=2}^{\infty} 3c_d \cdot n^{1-1/d} \left(\frac{\sigma}{3}\right)^{2.9(d-1)(i-1)} \left(1+6\sigma\sqrt{i\ln(3/\sigma)}\right) \\ &\leq 3c_d \cdot n^{1-1/d} \left(\frac{\sigma}{3}\right)^{2.9(d-1)} \left(1+6\sigma\sqrt{\ln(3/\sigma)}\right) \left(\sum_{i=0}^{\infty} \sqrt{i+2} \left(\frac{\sigma}{3}\right)^{2.9(d-1)i}\right). \end{split}$$

By observing that $\sum_{i=0}^{\infty} \sqrt{i+2} (\sigma/3)^{2.9(d-1)i}$ is bounded by a constant, we conclude that $\sum_{i=2}^{\infty} \mathbb{E}(L(E_i))$ is bounded by $O(\sigma n^{1-1/d})$.

Let $N_k := n(\sigma/3)^{2.9d(k-1)}$ be the upper bound on $\mathbb{E}(n_k)$ derived above. By the Chernoff bounds (Lemma 4.2), we have

$$\mathbb{P}(n_k \ge 2N_k) \le \exp(-N_k/3).$$

Choose k_1 such that $(\sigma/3)^{2.9d} \sigma \sqrt{n} \leq N_{k_1} \leq \sigma \sqrt{n}$. Thus, $k_1 = O(\log n)$. Assume that $n_k \leq 2N_k$ for all $1 \leq k \leq k_1$. Then, analogously to the above calculation, the contribution of E_2, \ldots, E_{k_1} is bounded by

$$\sum_{k=2}^{k_1} L(E_k) \le \sum_{k=2}^{k_1} c_d \cdot (1+2a_k) (6N_k)^{1-1/d}$$
$$\le c_d \cdot \left(1+6\sigma\sqrt{d\ln(3/\sigma)}\right) \cdot (6n)^{1-1/d} (\sigma/3)^{2.9(d-1)}$$
$$\cdot \left(\sum_{k=0}^{\infty} \sqrt{k+2}(\sigma/3)^{2.9(d-1)k}\right) = O(\sigma n^{1-1/d}).$$

Note that the probability that some $1 \le k \le k_1$ fails to satisfy $n_k \le 2N_k$ is bounded by

$$\sum_{k=1}^{k_1} \mathbb{P}(n_k > 2N_k) \le k_1 \exp(-N_{k_1}/3) = \exp(-\Omega(n^{1/2-\varepsilon})),$$

for any constant $\varepsilon > 0$. Since $n_{k_1} \leq 2N_{k_1}$, at most $n_{k_1} \leq 2\sigma\sqrt{n}$ vertices remain outside A_{k_1-1} . Let $k_2 := \lceil \sigma\sqrt{n} \rceil$. By a union bound, for any constant $\varepsilon > 0$,

$$\mathbb{P}(\exists j : X_j \notin A_{k_2}) \le n(\sigma/3)^{2.9d(k_2-1)} = \exp(-\Omega(n^{1/2-\varepsilon})).$$

Assume that we have the – very likely – event that all points are in A_{k_2} , then the remaining points outside A_{k_1-1} are contained in A_{k_2} . We conclude that

$$\sum_{k=k_1}^{\infty} L(E_k) = \sum_{k=k_1}^{k_2} L(E_k)$$

$$\leq c_d \cdot (1+2a_{k_2})(6N_{k_1})^{1-1/d}$$

$$= O(\sqrt{k_2}(N_{k_1})^{1-1/d})$$

$$= O(\sigma^{3/2-1/d}n^{\frac{1}{4}+\frac{1}{2}\left(1-\frac{1}{d}\right)}) = O(\sigma n^{1-1/d}).$$

In the remainder of the proof, we bound the total length of edges inside A_1 . Define $C := A_1$ and note that all edges in C have bounded length $\sqrt{d}(1 + 2a_1) = O(1)$. We let T_i contain the set of all those edges within C (in the longest 2-optimal tour T) whose lengths are in $[OPT/2^i, OPT/2^{i-1}]$. Let k_1 be such that $\sqrt{d}(1+2a_1) \in [OPT/2^{k_1}, OPT/2^{k_1-1}]$. Then $L(T_k) =$ 0 for all $k < k_1$, since no longer edges exist. Let k_2 be such that $\sigma \in [OPT/2^{k_2}, OPT/2^{k_2-1}]$. Then $\sum_{k=k_1}^{k_2} L(T_k) = O((k_2 - k_1) \cdot OPT) = O(\log(1/\sigma)OPT)$ by Lemma 4.9. This argument bounds the contribution of *long edges*, i.e., edges longer than σ , in the worst case, after observing the perturbation of the input points. It remains to bound the length of short edges in C, which we do in the next section.

4.3.2 Short Edges

To account for the length of the remaining edges, we take a different route than for the long edges: Call an edge that is shorter than σ a short edge and partition the bounding box $C = [-a_1, 1+a_1]^d$ into a grid of $(\sigma \times \cdots \times \sigma)$ -cubes C_1, \ldots, C_M with $M = \Theta((\sigma/(1+a_1))^{-d}) = \Theta(\sigma^{-d})$, which we call cells. All edges in T_k for $k \ge k_2$, i.e., short edges, are completely contained in a single cell or run from some cell C_i to one of its $3^d - 1$ neighboring cells. For a given tour T, let $E_{C_i}(T)$ denote the short edges of T for which at least one of the endpoints lies in C_i .

We aim to relate the length of the edges $E_{C_i}(T)$ for the longest 2-optimal tour T to the length of the edges $E_{C_i}(T_{\text{OPT}})$ of the optimal tour T_{OPT} . This local approach is justified by the following property.

Lemma 4.13. For any tour T', the contribution $L(E_{C_i}(T'))$ of cell C_i is lower bounded by $\mathsf{TSP}(X \cap C_i) - O(\sigma | X \cap C_i |^{\frac{d-2}{d-1}}).$

Proof. Consider all edges S in T' that have at least one endpoint in C_i . Replacing those edges $(u, v) \in S$ with $u \in C_i$ and $v \notin C_i$ by the shortest edge connecting u to the boundary

of C_i does not increase the total edge length by triangle inequality. If C_i were the unit cube, $L(E_{C_i}(T'))$ would thus be lower bounded by the boundary functional $\mathsf{TSP}^B(X \cap C_i)$. Instead, we scale the instance $X \cap C_i$ by $1/\sigma$ to obtain an instance X' in the unit cube, satisfying $\mathsf{TSP}(X \cap C_i) = \sigma \mathsf{TSP}(X')$ and, as argued above, $L(E_{C_i}(T')) \geq \sigma \mathsf{TSP}^B(X')$. Thus an application of Lemma 4.3 yields

$$L(E_{C_i}(T')) \ge \sigma \left(\mathsf{TSP}(X') - O\left(|X'|^{\frac{d-2}{d-1}}\right) \right) = \mathsf{TSP}(X \cap C_i) - O\left(\sigma \cdot |X \cap C_i|^{\frac{d-2}{d-1}}\right).$$

Intuitively, a cell C_i is of one of two kinds: either few points are expected to be perturbed into it and hence it cannot contribute much to the length of any 2-optimal tour (a sparse cell), or many unperturbed origins are close to the cell (a heavy cell). In the latter case, either the conditional densities of points perturbed into C_i are small, hence any optimal tour inside C_i has a large value by Lemma 4.6, or we find another cell close to C_i that has a very large contribution to the length of any tour.

To formalize this intuition, fix a cell C_i and let n_i be the expected number of points X_j with $X_j \in C_i$. Assume for convenience that a_1/σ and $(1 + a_1)/\sigma$ are integer. We describe the position of a cube C_i canonically by indices $pos(C_i) \in \{-\frac{a_i}{\sigma}, \ldots, \frac{1+a_i}{\sigma}\}^d$. For two cells C_i and C_j , we define their distance as $dist(C_i, C_j) = \|pos(C_i) - pos(C_j)\|_1$. For $k \ge 0$, let D_k denote all cells of distance k to C_i and let $n(D_k)$ denote the cardinality of unperturbed origins located in a cell in D_k . We call a perturbed point $X_\ell \in C_i$ with unperturbed origin $x_\ell \in C_j$, for some $C_j \in D_k$, a k-successful point. Let S_k denote the set of all k-successful points. Then $n_i = \sum_{k=0}^{\infty} \mathbb{E}(|S_k|)$.

Our first technical lemma shows that any cell C_i , having (in expectation) a large number μ of points perturbed into it from cells of distance at most K, contributes at least $\sigma \mu^{1-1/d} \exp(-O(K+1))$ to the length of the optimal tour.

Lemma 4.14. Let $K \ge 0$ and define $S_{\le K} := S_0 \cup \cdots \cup S_K$ as the set of k-successful points for $k \le K$. Let $\mu := \mathbb{E}(|S_{\le K}|)$. If $K = o(\log \mu)$, then with probability $1 - \exp(-\Omega(\mu))$, we have

$$L(E_{C_i}(T_{\text{OPT}})) \ge \frac{\sigma \mu^{1-1/d}}{\exp(O(K+1))}$$

Proof. Note that by Lemma 4.13, $L(E_{C_i}(T_{\text{OPT}})) \geq \mathsf{TSP}(S_{\leq K}) - O(\sigma \cdot |S_{\leq K}|^{\frac{d-2}{d-1}})$. Fix any realization of $S_{\leq K}$, i.e., choice of unperturbed origins inside some cell in D_0, \ldots, D_K whose perturbed points fall into C_i . We can simulate the distribution of $\mathsf{TSP}(S_{\leq K})$ (under this realization of $S_{\leq K}$) by appealing to the one-step model. Note that each point in $S_{\leq K}$ is distributed as a Gaussian conditioned on containment in cell C_j . By rotational invariance of the Gaussian distribution, Lemma 4.10 is applicable and bounds the conditional density function of each point in $S_{\leq K}$ by $\exp(K + (3/2)d)\sigma^{-d}$. By scaling, we obtain an instance in the unit cube with $N := |S_{\leq K}|$ points distributed according to density functions of maximum density $\exp(K + (3/2)d)$. Hence, by Lemma 4.6 we obtain that any tour has length $\Omega(N^{1-1/d}/\exp(K/d + 3/2))$ on the scaled instance with probability $1 - \exp(-\Omega(N))$. Scaling back to C_i , we obtain $\mathsf{TSP}(S_{\leq K}) \geq \Omega(\sigma N^{1-1/d}/\exp(K/d + 3/2))$. Since by Chernoff bounds (Lemma 4.2), $|S_{\leq K}| = \Omega(\mu)$ with probability $1 - \exp(-\Omega(\mu))$, we finally obtain, using Lemma 4.13,

$$L(E_{C_i}(T_{\text{OPT}})) \ge \Omega\left(\frac{\sigma\mu^{1-\frac{1}{d}}}{\exp(\frac{K}{d} + \frac{3}{2})}\right) - O(\sigma \cdot \mu^{\frac{d-2}{d-1}}) \ge \frac{\sigma\mu^{1-1/d}}{\exp(O(K+1))},$$

with probability $1 - \exp(-\Omega(\mu))$, where we used that $K = o(\log \mu)$.

The following simple technical lemma shows that with constant probability, a point is perturbed into the cell it originates in.

Lemma 4.15. Let $c \in Q_0 := [0, \sigma]^d$ and $Z \sim \mathcal{N}(c, \sigma^2)$. Then $\mathbb{P}(Z \in Q_0) \ge \frac{1}{(2\pi)^{d/2}} \exp(-\frac{d}{2})$.

Proof. Let $f(x) = \frac{1}{(2\pi)^{d/2}\sigma^d} \cdot \exp(-\frac{\|x-c\|^2}{2\sigma^2})$ be the probability density function of Z. For all $x \in Q_0$, we have $\|x-c\| \leq \sqrt{d\sigma}$ and hence $f(x) \geq \frac{1}{(2\pi)^{d/2}\sigma^d} \cdot \exp(-\frac{d}{2}) =: f_{\min}$. This yields

$$\mathbb{P}(Z \in Q_0) = \int_{Q_0} f(x) dx \ge \sigma^d f_{\min} = \frac{1}{(2\pi)^{d/2}} \exp(-d/2).$$

We are set-up to formally show the classification of heavy cells. Recall that $M = \Theta(\sigma^{-d})$ denotes the number of cells C_i .

Lemma 4.16. Let $\alpha := M^{\frac{d}{d-1}}$, $k_1 := \gamma \log \log(1/\sigma)$ and $k_2 := (1/\gamma')\sqrt{\log 1/\sigma}$ for sufficiently small constants γ, γ' . Then we can classify each cell C_i with $n_i \ge \frac{n}{\alpha}$ into one of the following two types.

(T1) With probability $1 - \exp(-\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$, we have

$$L(E_{C_i}(T)) \le O(\log 1/\sigma) L(E_{C_i}(T_{\text{OPT}})).$$

(T2) There is some $C_j \in D_{k_1} \cup \cdots \cup D_{k_2}$ such that for any $f(1/\sigma) = \text{polylog}(1/\sigma)$, we have

$$L(E_{C_i}(T)) \le \frac{L(E_{C_j}(T_{\text{OPT}}))}{f(1/\sigma)},$$

with probability $1 - \exp(-\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$.

Proof. We start with some intuition. By Lemma 4.4, we can bound $L(E_{C_i}(T)) = O(\sigma n_i^{1-1/d})$. If we have $\mathbb{E}(|S_{\leq k_1}|) = \Omega(n_i)$, then Lemma 4.14 already proves C_i to have type (T1). Otherwise, by tail bounds for the Gaussian distribution, we argue that some cell C_j in distance at most k_2 contains at least $n_i \exp(\Omega((\log \log 1/\sigma)^2))$ unperturbed origins. These are sufficiently many to let C_j contribute $f(1/\sigma)\sigma n_i^{1-1/d}$, for any $f(1/\sigma) = \operatorname{polylog}(1/\sigma)$, to the optimal tour length.

To make the intuition formal, note that all edges in $E_{C_i}(T)$ are contained in a cube of side length 3σ around C_i . By Chernoff bounds (Lemma 4.2), at most $2n_i$ points are contained in C_i with probability $1 - \exp(-\Omega(n_i))$. Hence, Lemma 4.4 bounds

$$L(E_{C_i}(T)) \le 3\sigma c_d (6n_i)^{1-1/d},$$
(6)

with probability $1 - \exp(-\Omega(n_i))$.

Case 1: $\mathbb{E}(S_{\leq k_1}) > n_i/2$. In this case, we may appeal to Lemma 4.14 (since $k_1 = o(\log n_i)$) and obtain

$$L(E_{C_i}(T_{\text{OPT}})) \ge \frac{\sigma(|S_{\le k_1}|)^{1-1/d}}{\exp(O(k_1))} = \Omega\left(\frac{\sigma n_i^{1-1/d}}{\log(1/\sigma)}\right),\tag{7}$$

with probability $1 - \exp(-\Omega(n_i))$, since $k_1 = \gamma \log \log 1/\sigma$ and γ can be chosen sufficiently small. By a union bound, (6) and (7) hold with probability $1 - \exp(-\Omega(n_i)) = 1 - \exp(-\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$, proving that C_i has type (T1). **Case 2:** $\mathbb{E}(S_{\leq k_1}) \leq n_i/2$. Every point in C_i has an ℓ_1 -distance of at least $\sigma(\operatorname{dist}(C_i, C_j) - d)$ to every point in C_j . Thus, by Lemma 4.1, we have

$$\mathbb{E}(|S_k|) \le n(D_k) \mathbb{P}\left(\|Z\| \ge \frac{k-d}{\sqrt{d}}\sigma\right) \le n(D_k) \exp\left(-0.32\frac{(k-d)^2}{d}\right),\tag{8}$$

for sufficiently large k. Since $\alpha = \text{poly}(1/\sigma)$, we can choose a sufficiently small constant γ' such that $k_2 = (1/\gamma')\sqrt{\log 1/\sigma}$ satisfies $\exp(-0.32(k_2 - d)^2/d) \leq 1/(4\alpha)$. From $\sum_{k=0}^{\infty} n(D_k) = n$, we conclude

$$\sum_{k=k_2+1}^{\infty} \mathbb{E}(|S_k|) \le \sum_{k=k_2+1}^{\infty} n(D_k) \exp(-0.32(k-d)^2/d) \le \frac{n}{4\alpha} \le \frac{n_i}{4}$$

Hence, we have

$$\sum_{k=k_1+1}^{k_2} \mathbb{E}(|S_k|) = n_i - \mathbb{E}(|S_{\leq k_1}|) - \sum_{k=k_2+1}^{\infty} \mathbb{E}(|S_k|) \ge \frac{n_i}{4}.$$

By (8), it follows that

$$N := \sum_{k=k_1+1}^{k_2} n(D_k) \ge \exp\left(0.32\frac{(k_1 - d)^2}{d}\right) \sum_{k=k_1+1}^{k_2} \mathbb{E}(|S_k|) = n_i \exp(\Omega((\log \log 1/\sigma)^2))$$

unperturbed origins are situated in cells in distance $k_1 < k \leq k_2$ from C_i . Note that there are at most $\sum_{k=k_1+1}^{k_2} |D_k| = O(k_2^d) = \text{polylog}(1/\sigma)$ such cells and $\exp(\Omega((\log \log 1/\sigma)^2)) = \omega(\log^c(1/\sigma))$ for any $c \in \mathbb{N}$. By pigeon hole principle, there is a cell $C_j \in D_{k_1} \cup \cdots \cup D_{k_2}$ with $\Omega(N/k_2^d) = n_i \exp(\Omega((\log \log 1/\sigma)^2))$ many unperturbed origins.

Let S'_0 be the 0-successful points for cell C_j , i.e., the points with origin in C_j that are perturbed into C_j . By Lemma 4.15, each unperturbed origin $x_\ell \in C_j$ has constant probability to be perturbed into C_j , i.e., $\mathbb{P}(X_\ell \in C_j) = \Omega(1)$. Hence, $\mathbb{E}(|S'_0|) = n_i \exp(\Omega((\log \log 1/\sigma)^2))$. Thus, Lemma 4.14 bounds

$$L(E_{C_j}(T_{\text{OPT}})) \ge \frac{\sigma(\mathbb{E}(|S'_0|))^{1-\frac{1}{d}}}{\exp(O(1))} = \sigma n_i^{1-\frac{1}{d}} \exp(\Omega((\log \log 1/\sigma)^2)),$$
(9)

with probability $1 - \exp(-\Omega(\mathbb{E}(|S'_0|))) = 1 - \exp(-\Omega(n_i))$. Since (6) and (9) hold simultaneously with probability $1 - \exp(-\Omega(n_i)) = 1 - \exp(-\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$, this proves that C_i has type (T2).

4.3.3 Total Length of 2-Optimal Tours

With the analyses of the previous subsections, we can finally bound the total length of 2optimal tours. To bound the total length of short edges, consider first sparse cells C_i , i.e., cells containing $n_i \leq n/\alpha$ perturbed points in expectation (recall that $\alpha = M^{\frac{d}{d-1}}$, where $M = \Theta(\sigma^{-d})$ is the number of cells). For each such cell, the Chernoff bound (Lemma 4.2) yields that with probability $1 - \exp(-\Omega(n/\alpha))$, at most $2n/\alpha$ points are contained in C_i , since each point is perturbed independently. By a union bound, no sparse cell contains more than $2n/\alpha$ points with probability at least $1 - M \exp(-\Omega(n/\alpha)) = 1 - \exp(-\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$. In this event, Lemma 4.4 allows for bounding the contribution of sparse cells by

$$\sum_{n_i \le n/\alpha} L(E_{C_i}(T)) \le M(3\sigma)c_d \left(\frac{6n}{\alpha}\right)^{1-\frac{1}{d}} = O\left(\frac{M\sigma n^{1-\frac{1}{d}}}{\alpha^{1-\frac{1}{d}}}\right) = O(\sigma n^{1-\frac{1}{d}}).$$
(10)

For bounding the length in the remaining cells (the heavy cells), let $\mathcal{T}_1 := \{i \mid C_i \text{ has type (T1)}\}$ and $\mathcal{T}_2 := \{i \mid C_i \text{ has type (T2)}\}$. We observe the following: with probability at least $1 - M \exp(-\Omega(n^{1-\varepsilon})) = 1 - \exp(-\Omega(n^{1-\varepsilon}))$, all type-(T1) cells C_i satisfy $L(E_{C_i}(T)) = O(\log 1/\sigma)L(E_{C_i}(T_{OPT}))$. Thus,

i:

$$\sum_{i \in \mathcal{T}_1} L(E_{C_i}(T)) \le O(\log 1/\sigma) \cdot \left(\sum_{i \in \mathcal{T}_1} L(E_{C_i}(T_{\text{OPT}})) \right) \le O(\log 1/\sigma) \cdot \text{OPT},$$
(11)

where the last inequality follows from $\sum_{i=1}^{M} L_{C_i}(T_{\text{OPT}}) \leq 2 \cdot \text{OPT}$, which holds since every edge in OPT (inside C) is counted at most twice on the left-hand side.

Let $A : \mathcal{T}_2 \to \{1, \ldots, M\}$ be any function that assigns to each cell C_i of type-(T2) a corresponding cell $C_{A(i)} \in D_{k_1} \cup \cdots \cup D_{k_2}$ satisfying the condition (T2). We say that C_i charges $C_{A(i)}$. We can choose any $f(1/\sigma) = \text{polylog}(1/\sigma)$ and have with probability at least $1 - M \exp(-\Omega(n^{1-\varepsilon})) = 1 - \exp(-\Omega(n^{1-\varepsilon}))$ that $L(E_{C_i}(T)) \leq \frac{L(E_{C_{A(i)}}(T_{\text{OPT}}))}{f(1/\sigma)}$ for all $i \in \mathcal{T}_2$. Assume that this event occurs. Since every cell C_i can only be charged by cells in distance $k_1 < k \leq k_2$, each cell can only be charged $\sum_{k=k_1+1}^{k_2} |D_k| = O(k_2^d)$ times. Hence,

$$\sum_{i \in \mathcal{T}_2} L(E_{C_{A(i)}}(T_{\text{OPT}})) \le O(k_2^d) \sum_{i=1}^M L(E_{C_i}(T_{\text{OPT}})) = O(k_2^d) \text{OPT}.$$

Since $k_2^d = \text{polylog}(1/\sigma)$, choosing $f(1/\sigma) = \text{polylog}(1/\sigma)$ sufficiently large yields

$$\sum_{i \in \mathcal{T}_2} L(E_{C_i}(T)) \le \sum_{i \in \mathcal{T}_2} \frac{L(E_{C_{A(i)}}(T_{\text{OPT}}))}{f(1/\sigma)} \le \frac{O(k_2^d)\text{OPT}}{f(1/\sigma)} = O(\text{OPT}).$$
 (12)

Proof of Theorem 4.11. By a union bound, we can bound by $1 - \exp(-\Omega(n^{1/2-\varepsilon}))$, for any constant $\varepsilon > 0$, the probability that (i) OPT = $\Omega(\sigma n^{1-1/d})$ (by Lemma 4.7), (ii) all edges outside C contribute $O(\sigma n^{1-1/d}) = O(\text{OPT})$ (by Lemma 4.12), (iii) all sparse cells contribute $O(\sigma n^{1-1/d}) = O(\text{OPT})$ (by (10)), (iv) the type-(T1) cells C_i induce a cost of $O(\log 1/\sigma)\text{OPT}$ (by (11)), and (v) the type-(T2) cells induce a cost of O(OPT) (by (12)). Since the remaining edges are long edges and contribute only $O(\log(1/\sigma) \cdot \text{OPT})$, we obtain that every 2-optimal tour has a length of at most $O(\log 1/\sigma)\text{OPT}$ with probability $1 - \exp(-\Omega(n^{1/2-\varepsilon}))$.

Since a 2-optimal tour always constitutes a $O(\log n)$ -approximation to the optimal tour length by Lemma 4.9, we also obtain that the expected cost of the worst 2-optimal tour is bounded by

$$O(\log 1/\sigma) \cdot \operatorname{OPT} + \exp(-\Omega(n^{1/2-\varepsilon})) \cdot O(\log n) \cdot \operatorname{OPT} = O(\log 1/\sigma) \cdot \operatorname{OPT}.$$



Figure 2: Parts V_1 and V_3 of the lower bound instance. Each point is contained in a corresponding small container (depicted as brown circle) with high probability. The black lines indicate the constructed 2-optimal tour, which runs analogously on V_2 .

4.4 Lower Bound on the Approximation Ratio

We complement our upper bound on the approximation performance by the following lower bound: for $\sigma = O(1/\sqrt{n})$, the worst-case lower bound is robust against perturbations. For this, we face the technical difficulty that in general, a single outlier might destroy the 2-optimality of a desired long tour, potentially cascading into a series of 2-Opt iterations that result in a substantially different or even optimal tour.

Theorem 4.17. Let $\sigma = O(1/\sqrt{n})$. For infinitely many n, there is an instance X of points in \mathbb{R}^2 perturbed by normally distributed noise of standard deviation σ such that with probability $1 - O(n^{-s})$ for any constant s > 0, we have $2\mathsf{OPT}(X) = \Omega(\log n/\log \log n) \cdot \mathsf{TSP}(X)$. This also yields

$$\mathbb{E}\left(\frac{2\mathsf{OPT}(X)}{\mathsf{TSP}(X)}\right) = \Omega\left(\frac{\log n}{\log\log n}\right).$$

We remark that our result transfers naturally to the one-step model with $\phi = \Omega(n)$ and interestingly, holds with probability 1 over such random perturbations.

Proof of Theorem 4.17. We alter the construction of Chandra et al. [6] to strengthen it against Gaussian perturbations with standard deviation $\sigma = O(1/\sqrt{n})$ (see Figure 2). Let $p \geq 3$ be an odd integer and $P := 3p^{2p}$. The original instance of [6] is a subset of the $(P \times P)$ grid, which we embed into $[0,1]^2$ by scaling by 1/P, and consists of three parts V_1 , V_2 and V_3 . The vertices in V_1 are partitioned into the layers L_0, \ldots, L_p . Layer *i* consists of $p^{2i} + 1$ equidistant vertices, each of which has a vertical distance of $c_i = p^{2p-2i-1}/P$ to the point above it in Layer i + 1 and a horizontal distance of $a_i = p^{2p-2i}/P$ to the nearest neighbor(s) in the same layer. The set V_2 is a copy of V_1 shifted to the right by a distance of 2/3. The remaining part V_3 consists of a copy of Layer p of V_1 shifted to the right by 1/3 to connect V_1 and V_2 by a path of points. We regard L_i as the set of Layer-*i* points in $V_1 \cup V_2 \cup V_3$.

As in the original construction, we will construct an instance of $n = \Theta(p^{2p})$ points, which implies $p = \Theta(\log n / \log \log n)$. Let $0 \le t \le p$ be the largest odd integer such that $p^{2t+1} \le (3\sigma)^{-1}$. In our construction, we drop all Layers $t + 1, \ldots, p$ in both V_1 and V_2 , as well as Layer p in V_3 . Instead, we connect V_1 and V_2 already in Layer t by an altered copy of Layer t of V_1 shifted to the right by 1/3. Let C be an arbitrary point of our construction, for convenience we will use the central point of Layer t in V_3 . We introduce $p^{2p} - 1$ additional copies of this point C. These surplus points serve as a "padding" of the instance to ensure $n = \Theta(p^{2p})$. Note that the resulting instance has t + 1 layers L_0, \ldots, L_t . We choose t such that the magnitude of perturbation is negligible compared to the pairwise distances of all non-padding points. Furthermore, the restriction on σ ensures that incorporating the padding points increases the optimal tour length only by a constant.

Lemma 4.18. With probability $1 - O(n^{-s})$ for any constant s > 0, the optimal tour has length O(1).

Proof. Let n be the number of points in the constructed instance. Note that $\overline{X} = (x_1, \ldots, x_n)$ consists of (i) a subset $\overline{X}^{\text{orig}}$ of the instance of Chandra et al. [6], plus (ii) an additional copy \overline{X}^t of Layer t and (iii) the padding points $\overline{X}^{\text{pad}}$ in V_3 . Denote the number of points in $\overline{X}^{\text{orig}} \cup \overline{X}^t$ by n'. We have

$$n' = p^{2t} + 2\left(\sum_{i=0}^{t} p^{2i} + 1\right) \le p^{2t} + 2(1 - p^{-2})^{-1}p^{2t} + 2t = O(\sigma^{-1}/p),$$

by choice of t. Hence $n = (p^{2p} - 1) + n' = \Theta(p^{2p})$. It is easy to see [6] that the original instance of Chandra et al. has a minimum spanning tree of length $\mathsf{MST}(\overline{X}^{\mathrm{orig}}) \leq 9p^{2p}/P$. (This is achieved by the spanning tree that includes, for each Layer-*i* vertex with $0 \leq i < p$, the vertical edge to the point above it, and each edge between consecutive points on Layer p.) Clearly,

$$\mathsf{MST}(\overline{X}^{\mathrm{orig}} \cup \overline{X}^t) \le \mathsf{MST}(\overline{X}^{\mathrm{orig}}) + \mathsf{MST}(\overline{X}^t) + 2a_t \le 9p^{2p}/P + p^{2p}/P + 2a_t = O(1).$$

Consider the perturbed instance $X \leftarrow \text{pert}_{\sigma}(\overline{X})$. Note that for every constant s > 0, we have $p\sigma \geq 3\sigma\sqrt{d(s+1)\ln n}$ for sufficiently large n. Thus for each $1 \leq i \leq n$, the Gaussian noise $Z_i \sim \mathcal{N}(0,\sigma)$ satisfies $||Z_i|| \leq p\sigma$ with probability at least $1 - O(n^{-s+1})$ by Lemma 4.1. By a union bound, we have $\sum_{i=1}^{n'} ||Z_i|| \leq O(n'p\sigma) = O(1)$ with probability at least $1 - O(n^{-s})$. In this case, by the triangle inequality, the fact that $\mathsf{TSP}(Y) \leq 2 \cdot \mathsf{MST}(Y)$ for all point sets Y and since only a constant number of edges connects the three parts, we obtain

$$\begin{aligned} \mathsf{TSP}(X_1, \dots, X_n) &\leq 2 \cdot \mathsf{MST}(\overline{X}^{\mathrm{orig}} \cup \overline{X}^t) + 2\left(\sum_{i=1}^{n'} \|Z_i\|\right) + \mathsf{TSP}(X^{\mathrm{pad}}) + O(1) \\ &\leq \mathsf{TSP}(X^{\mathrm{pad}}) + O(1). \end{aligned}$$

Note that we may translate and scale $\overline{X}^{\text{pad}}$ to be contained in $[0, \sigma]^d$, by which $\mathsf{TSP}(X^{\text{pad}})$ may be regarded as the optimal tour length on an instance of $p^{2p} = \Theta(n)$ points in $[0, 1]^d$ perturbed by Gaussians with standard deviation 1. By Lemma 4.5, any 2-optimal tour and hence also the optimal tour on the scaled instance has length $O(\sqrt{n})$ with probability $1 - \exp(-\Omega(\sqrt{n}))$. Scaling back to the original instance, we obtain $\mathsf{TSP}(X^{\text{pad}}) = O(\sqrt{n}\sigma) = O(1)$ with probability $1 - \exp(-\Omega(\sqrt{n}))$. This yields the result by a union bound.

We find a long 2-optimal tour on all non-padding points analogously to the original construction by taking a shortcut of the original 2-optimal tour, which connects V_1 and V_2 already in Layer t (see Figure 2).

Consider the padding points, which are yet to be connected. Let C_{ℓ} denote the nearest point in Layer t of V_3 that is to the left of C. Symmetrically, C_r is the nearest point to the

right of C. Let T^p be any 2-optimal path from C_{ℓ} to C_r that passes through all the padding points (including C). We replace the edges (C_{ℓ}, C) and (C, C_r) by the path T^p , completing the construction of our tour T.

Lemma 4.19. Let s > 0 be arbitrary. With probability $1 - O(n^{-s})$, T is 2-optimal and has a length of $\Omega(\log n / \log \log n)$.

Note that given Lemma 4.19, Theorem 4.17 follows directly using Lemma 4.18. The (rather technical) proof of Lemma 4.19 hence concludes our lower bound.

Probability of 2-optimality. To account for the perturbation in the analysis, we define a safe region for every point. More formally, let x_j be any unperturbed origin. We define its container B_j as the circle centered at x_j with radius $\beta := a_t/8 = p^{2p-2t}/(8P) \ge \sigma p/8$. Very likely, all perturbed points lie in their containers.

Lemma 4.20. For sufficiently large p, the tour T constructed as described in Section 4.4 is 2-optimal, provided that all points X_i lie in their corresponding containers B_i .

We first show that this lemma implies Lemma 4.19.

Proof of Lemma 4.19. Let $Z \sim \mathcal{N}(0, \sigma^2)$, and let s > 0 be arbitrary. By $\beta \geq \sigma p/8 = \Omega(\sigma \log(n)/\log \log n) = \omega(\sigma \sqrt{\log n})$, we have $\beta \geq 3\sigma \sqrt{d(s+1)\ln n}$ for sufficiently large n. By definition of the containers, Lemma 4.1 yields that for any point X_j and sufficiently large n,

$$\mathbb{P}(X_j \notin B_j) \le \mathbb{P}(\|Z\| \ge \beta) \le \mathbb{P}(\|Z\| \ge \sigma 3\sqrt{d(s+1)\ln n}) \le n^{-(s+1)}.$$

By a union bound, we conclude that with probability $1 - n^{-s}$, all points are contained in their corresponding containers and hence, by the previous lemma, T is 2-optimal.

Recall that t is the largest odd integer satisfying $p^{2t+1} \leq (3\sigma)^{-1}$. Since $\sigma^{-1} = \Omega(\sqrt{n})$, this implies $t \geq \frac{p-1}{2} - 1$. Observe that T visits $t = \Omega(p)$ many layers and crosses a horizontal distance of 2/3 in each of them. Hence, it has a length of at least $\Omega(p) = \Omega(\log n/\log \log n)$.

In the remainder of this section, we prove Lemma 4.20, i.e., show that the constructed tour is 2-optimal, provided all points stay inside their respective containers. Clearly, it suffices to show for any pair of edges (u, v) and (w, z) in the tour, the corresponding 2-change, i.e., replacing these edges by (u, w) and (v, z) does not reduce the tour length, i.e., $d(u, w) + d(v, z) \ge d(u, v) + d(w, z)$. We first state the technical lemmas capturing the ideas behind the construction.

The first lemma treats pairs of horizontal edges and establishes how large their vertical distance must be in order to make swapping these edges increase the length of the tour. It is a generalization of a similar lemma of Chandra et al. [6] to a perturbation setting, in which points are placed arbitrarily into small containers.

Note that in what follows, for a point $p \in \mathbb{R}^2$, we let p_x denote its x-coordinate and p_y its y-coordinate. Furthermore, for any points $p, q \in \mathbb{R}^2$, we let $d_x(p,q) := |p_x - q_x|$ and $d_y(p,q) := |p_y - q_y|$ denote their horizontal and vertical distance, respectively.

Lemma 4.21. Let \overline{pq} and \overline{rs} be horizontal line segments in the Euclidean plane with $p_x < q_x$ and $r_x < s_x$. Let B_p , B_q , B_r and B_s be circles of radius β with centers p, q, r and s, respectively. If $d(r,s) \ge d(p,q) + 4\beta$ and the vertical distance $v := d_y(p,r) = d_y(q,s)$ between \overline{pq} and \overline{rs} is at least

$$\sqrt{d(p,q)d(r,s)} + 4\beta d(r,s) + 2\beta,$$

then, for all $\tilde{p} \in B_p, \tilde{q} \in B_q, \tilde{r} \in B_r, \tilde{s} \in B_s$, we have

$$d(\tilde{p}, \tilde{r}) + d(\tilde{q}, \tilde{s}) \ge d(\tilde{p}, \tilde{q}) + d(\tilde{r}, \tilde{s}).$$

Proof. Note that $d_y(\tilde{p}, \tilde{r}), d_y(\tilde{q}, \tilde{s}) \geq v - 2\beta$. Furthermore, we have that

$$d(r,s) - 2\beta \leq \tilde{s}_x - \tilde{r}_x = (\tilde{p}_x - \tilde{r}_x) + (\tilde{q}_x - \tilde{p}_x) + (\tilde{s}_x - \tilde{q}_x)$$
$$\leq (\tilde{p}_x - \tilde{r}_x) + d(p,q) + 2\beta + (\tilde{s}_x - \tilde{q}_x),$$

and hence

$$(\tilde{p}_x - \tilde{r}_x) + (\tilde{s}_x - \tilde{q}_x) \ge d(r, s) - d(p, q) - 4\beta,$$

where the right-hand side expression is at least 0, since $d(r,s) \ge d(p,q) + 4\beta$ by assumption. Let $L := \tilde{p}_x - \tilde{r}_x$ and $R := \tilde{s}_x - \tilde{q}_x$, then it is straight-forward to verify that the expression

$$d(\tilde{p}, \tilde{r}) + d(\tilde{q}, \tilde{s}) \geq \sqrt{(v - 2\beta)^2 + L^2} + \sqrt{(v - 2\beta)^2 + R^2},$$
 (13)

subject to $L + R \ge d(r, s) - d(p, q) - 4\beta$ is minimized when $L = R = \frac{d(r, s) - d(p, q)}{2} - 2\beta$. Hence, we can bound (13) by

$$\begin{split} d(\tilde{p},\tilde{r}) + d(\tilde{q},\tilde{s}) \\ &\geq 2\sqrt{(v-2\beta)^2 + \left(\frac{d(r,s) - d(p,q)}{2} - 2\beta\right)^2} \\ &\geq 2\sqrt{d(p,q)d(r,s) + 4\beta d(r,s) + \left(\frac{d(r,s) - d(p,q)}{2} - 2\beta\right)^2} \\ &= 2\sqrt{\left(\frac{d(r,s) - d(p,q)}{2}\right)^2 + d(p,q)d(r,s) + 4\beta d(r,s) - 2\beta(d(r,s) - d(p,q)) + (2\beta)^2} \\ &= 2\sqrt{\left(\frac{d(r,s) + d(p,q)}{2}\right)^2 + 2\beta(d(r,s) + d(p,q)) + (2\beta)^2} \\ &= 2\left(\frac{d(r,s) + d(p,q)}{2} + 2\beta\right) = d(r,s) + d(p,q) + 4\beta \ge d(\tilde{p},\tilde{q}) + d(\tilde{r},\tilde{s}), \end{split}$$

where the third line follows from our assumption on v.

The following very basic lemma shows that a sequence of edges that share roughly the same direction will always be 2-optimal.

Lemma 4.22. Let p_1, p_2, p_3 and p_4 be a sequence of points in $[0, 1]^2$ such that all connecting segments $p_{i+1} - p_i$ fulfill $|(p_{i+1} - p_i)_y| \le (p_{i+1} - p_i)_x$. Then,

$$d(p_1, p_3) + d(p_2, p_4) \ge d(p_1, p_2) + d(p_3, p_4).$$

Proof. For any point p, let C_p denote the cone $C_p := \{q \mid |(q-p)_y| \leq (q-p)_x\}$. Let $\Delta := p_2 - p_1$, then by assumption, we have $p_2 \in C_{p_1}$ and thus $|\Delta_y| \leq \Delta_x$. Let us assume that $0 \leq \Delta_y \leq \Delta_x$ (the other case is symmetric). Since by assumption, $p_3 \in C_{p_2}$, we have for $\Delta' := p_3 - p_1$ that $\Delta'_x = \Delta_x + \delta_x$ and $\Delta'_y = \Delta'_y + \delta_y$ for some $\delta_x > 0$ and δ_y with $|\delta_y| < \delta_x$. If $\delta_x \geq \Delta_y$, the claim

is immediate from $d(p_1, p_3) \ge \Delta_x + \delta_x \ge \Delta_x + \Delta_y \ge d(p_1, p_2)$. Otherwise, for $\delta_x < \Delta_y$, we obtain

$$d(p_1, p_3) = \sqrt{(\Delta_x + \delta_x)^2 + (\Delta_y + \delta_y)^2}$$

$$\geq \sqrt{(\Delta_x + \delta_x)^2 + (\Delta_y - \delta_x)^2}$$

$$\geq \sqrt{\Delta_x^2 + \Delta_y^2 + 2\delta_x(\Delta_x - \Delta_y)} \geq \sqrt{\Delta_x^2 + \Delta_y^2} = d(p_1, p_2).$$

By an analogous computation, $d(p_2, p_4) \ge d(p_3, p_4)$ follows and hence the claim.

We can now prove Lemma 4.20. Assume that all points are contained in their respective containers. We call an edge between X_i and X_j horizontal (or vertical) if the edge between x_i and x_j is horizontal (or vertical) and neither x_i nor x_j belong to the set of padding points. In what follows, we will first consider horizontal-horizontal, horizontal-vertical and vertical-vertical edge pairs and then turn to pairs of edges for which at least one edge is adjacent to some padding point. Recall that β is chosen such as to satisfy $a_t = 8\beta$.

Horizontal-horizontal edge pair. Let (X_i, X_{i+1}) and (X_j, X_{j+1}) be two horizontal edges. Horizontal edges (X_i, X_{i+1}) with $x_i, x_{i+1} \in L_k$ appear only if $k \leq t$. We distinguish the following cases.

1. $x_i, x_{i+1}, x_j, x_{j+1} \in L_k$: Both edges are in the same layer. Note that no 2-change swaps neighboring edges. Assume without loss of generality that $(x_i)_x < (x_{i+1})_x < (x_j)_x < (x_{j+1})_x$ (the other case is symmetric). Since $a_k \ge a_t = 8\beta$, we have that

$$d_y(X_i, X_{i+1}) \le 2\beta \le a_k - 2\beta \le d_x(X_i, X_{i+1}).$$

Similarly, $d_y(X_j, X_{j+1}) \leq d_x(X_j, X_{j+1})$ and $d_y(X_{i+1}, X_j) \leq d_x(X_{i+1}, X_j)$. This shows that Lemma 4.22 is applicable to $X_i, X_{i+1}, X_j, X_{j+1}$, which yields that no 2-change can be profitable.

- 2. $x_i, x_{i+1} \in L_k$, and $x_j, x_{j+1} \in L_{k+1}$. By construction of T, the edges have opposite direction. Assume that $(x_i)_x < (x_{i+1})_x$ and hence $(x_j)_x > (x_{j+1})_x$ (the other case is symmetric). By construction $(x_{i+1})_x (x_i)_x \ge a_k$. We have that $(X_{i+1})_x (X_i)_x \ge a_k 2\beta \ge a_t 2\beta = 6\beta > 0$. The same reasoning shows that $(X_j)_x > (X_{j+1})_x$. Similarly, one can show that $p_y > q_y$ for all $p \in \{X_j, X_{j+1}\}$ and $q \in \{X_i, X_{i+1}\}$. Hence the 2-change to (X_i, X_j) and (X_{i+1}, X_{j+1}) has a crossing, which by triangle inequality cannot be profitable.
- 3. $x_i, x_{i+1} \in L_k$, and $x_j, x_{j+1} \in L_{k+\ell}$ with $\ell \ge 2$ and $k+\ell \le t$. Either both edges have opposite directions, then the previous argument shows that a 2-change is not profitable. Otherwise, note that the first requirement of Lemma 4.21, $a_k \ge a_{k+\ell} + 4\beta$, is fulfilled.

Also note that $\beta = \frac{a_t}{8} \leq \frac{a_k}{8p^{2\ell}}$, since $k \leq t - \ell$. We have

$$\begin{split} \sqrt{d(x_i, x_{i+1})d(x_j, x_{j+1}) + 4\beta d(x_i, x_{i+1}) + 2\beta} &= \sqrt{a_k a_{k+\ell} + 4\beta a_k} + 2\beta \\ &\leq \sqrt{\frac{a_k^2}{p^{2\ell}} + \frac{a_k^2}{2p^{2\ell}}} + \frac{a_k}{4p^{2\ell}} \\ &\leq \sqrt{\frac{3}{2}} \cdot \frac{a_k}{p^\ell} + \frac{a_k}{4p^{2\ell}} \\ &\leq \left(\sqrt{\frac{3}{2}} \cdot \frac{1}{p^{\ell-1}} + \frac{1}{4p^{2\ell-1}}\right) \frac{a_k}{p} \\ &\leq c_k \leq \sum_{m=0}^{\ell-1} c_{k+m} = d_y(x_i, x_j), \end{split}$$

since for sufficiently large p, we have $\sqrt{3/2}/p^{\ell-1} + 1/(4p^{2\ell-1}) \leq 1$. Consequently, Lemma 4.21 applies and shows that the 2-change does not yield an improvement.

Horizontal-vertical edge pair. Let (X_i, X_{i+1}) be a vertical edge and (X_j, X_{j+1}) be a horizontal edge. We assume that the vertical edge is in V_1 , since the case $x_i, x_{i+1} \in V_2$ is symmetric. Exactly one of the following cases occurs.

1. $x_i \in L_k, x_{i+1} \in L_{k+1}$ and $x_j, x_{j+1} \in L_{k'}$ with $k' \in \{k, k+1\}$. The horizontal edge is in the same layer as one of the end points of the vertical edge. Clearly, $d(X_i, X_{i+1}) \leq c_k + 2\beta$ and $d(X_j, X_{j+1}) \leq a_{k'} + 2\beta$. Since a 2-change cannot swap neighboring edges, at least one horizontal segment lies between both edges. By construction of the tour, one of the edges $\{x_i, x_j\}$ and $\{x_{i+1}, x_{j+1}\}$ crosses a vertical distance of at least c_k and the other a horizontal distance of at least $2a_{k'}$. Hence

$$d(X_i, X_j) + d(X_{i+1}, X_{j+1}) \ge 2a_{k'} + c_k - 4\beta \ge a_{k'} + c_k + 4\beta,$$

since $a_{k'} \ge a_t = 8\beta$.

2. $x_i \in L_k, x_{i+1} \in L_{k+1}$ and $x_j, x_{j+1} \in L_{k'}$ with $k' \notin \{k, k+1\}$. As in the previous case, $d(X_i, X_{i+1}) \leq c_k + 2\beta$ and $d(X_j, X_{j+1}) \leq a_{k'} + 2\beta$. Consider first the case that k' < k, then by construction of the tour, one of the edges $\{x_i, x_j\}$ and $\{x_{i+1}, x_{j+1}\}$ crosses a horizontal distance of at least $a_{k'}$ and the other edge crosses a vertical distance of at least $c_{k'}$, yielding

$$d(X_i, X_j) + d(X_{i+1}, X_{j+1}) \ge a_{k'} + c_{k'} - 4\beta \ge a_{k'} + c_k + 4\beta,$$

since $c_{k'} \ge c_{k-1} \ge c_k + 8\beta$. Otherwise, if k' > k+1, the edges x_i, x_j crosses a vertical distance of at least $c_{k+1} + c_k$ and hence

$$d(X_i, X_j) + d(X_{i+1}, X_{j+1}) \ge c_{k+1} + c_k - 2\beta \ge a_{k'} + c_k + 4\beta,$$

since $c_{k+1} \ge a_{k+2} + 6\beta \ge a_{k'} + 6\beta$. Thus in both cases, a 2-change is not profitable.

Vertical-vertical edge pair. Let (X_i, X_{i+1}) and (X_j, X_{j+1}) be vertical edges.

- 1. $x_i \in L_k, x_{i+1} \in L_{k+1}$ and $x_j \in L_{k'}, x_{j+1} \in L_{k'+1}$ with $(x_i)_x = (x_j)_x$, i.e., the vertical edges are above each other. By swapping the x- and y-axis in Lemma 4.22, we can show that a 2-change is not profitable, since it is easy to see that $|(p-q)_x| \leq (p-q)_y$ for all consecutive pairs (p,q) in $(X_i, X_{i+1}, X_j, X_{j+1})$.
- 2. $x_i \in L_k, x_{i+1} \in L_{k+1} \text{ and } x_j \in L_{k'}, x_{j+1} \in L_{k'+1} \text{ with } (x_i)_x \neq (x_{i'})_x$. Clearly, $d(X_i, X_j) \ge a_0 2\beta$ and $d(X_{i+1}, X_{j+1}) \ge a_0 2\beta$, while $d(X_i, X_{i+1}) \le c_k + 2\beta \le c_0 + 2\beta$ and $d(X_j, X_{j+1}) \le c_{k'} + 2\beta \le c_0 + 2\beta$. Hence a 2-change is not profitable, since $a_0 \ge 8\beta + c_0$.

Padding points. Since we assumed for convenience that the padding points are placed at the central vertex C of Layer t in V_3 , only the edges with at least one endpoint in V_3 are relevant candidates for the treatment of padding points. This is because all other edges have both endpoints at a distance of 1/6 to the padding points, which can never be accounted for by its edge length, since all edges except in Layer 0 are much shorter than 1/3. Separately, the Layer-0 edges can be handled easily as well: an edge $\{X_i, X_{i'}\}$ with $x_i = x_{i'} \in \overline{X}^{\text{pad}}$ is a horizontal edge, hence the pair $(X_i, X_{i'})$ and a Layer-0 edge trigger the corresponding case of horizontal-horizontal edge pairs with even smaller edge length of the edge $(X_i, X_{i'})$ in Layer t.

It remains to handle the following cases, where we regard C as a padding point, i.e., $C \in \overline{X}^{\text{pad}}$, not as a Layer-t point.

1. $x_i, x_{i'} \in \overline{X}^{\text{pad}}$, and $x_j, x_{j'} \in L_t$. Clearly, $d(X_j, X_{j'}) \leq a_t + 2\beta$ and $d(X_j, X_{j'}) \leq 2\beta$. Furthermore, at least one of $\{x_j, x_{j'}\}$ has a horizontal distance of at least $2a_t$ to $x_i = x_{i'}$. Hence,

$$d(X_i, X_j) + d(X_{i'}, X_{j'}) \ge 2a_t - 2\beta \ge a_t + 4\beta \ge d(X_j, X_{j'}) + d(X_i, X_{i'})$$

- 2. $x_i \in \{C_\ell, C_r\}, x_{i'} \in \overline{X}^{\text{pad}} \text{ and } x_j, x_{j'} \in L_t$. These edge pairs are exactly as regular pairs of Layer-*t* edges and the corresponding case of horizontal-horizontal edge pairs applies.
- 3. $x_i, x_{i'}, x_j, x_{j'} \in \overline{X}^{\text{pad}} \cup \{C_r, C_\ell\}$. All such edges are 2-optimal by construction, since a 2-optimal path from C_ℓ to C_r passing by all padding points was used.

This concludes the case analysis and thus the proof of Lemma 4.20.

5 Concluding Remarks

Running-time. Our approach for Euclidean distances does not work for d = 2 and d = 3. However, we can use the bound of Englert et al. [11] for Euclidean distances, which yields a bound polynomial in n and $1/\sigma$ for $d \in \{2, 3\}$.

In the same way as Englert et al. [11], we can slightly improve the smoothed number of iterations by using an insertion heuristic to choose the initial tour. We save a factor of $n^{1/d}$ for Manhattan and Euclidean distances and a factor of $n^{2/d}$ for squared Euclidean distances. The reason is that there always exist tours of length $O(D_{\max}n^{1-\frac{1}{d}})$ for n points in $[-D_{\max}, D_{\max}]^d$ for Euclidean and Manhattan distances and of length $O(D_{\max}^2n^{1-\frac{2}{d}})$ for squared Euclidean Euclidean distances and of length $O(D_{\max}^2n^{1-\frac{2}{d}})$ for squared Euclidean distances for $d \geq 2$ [29] (the constants in these upper bounds depend on d). Taking into

account also that, because Gaussians have light tails, only few points are far away from the hypercube $[0, 1]^d$ after perturbation, one might get an even better bound. However, we did not take these improvements into account in our analysis to keep the paper concise.

Of course, even our improved bounds do not fully explain the linear number of iterations observed in experiments. However, we believe that new approaches, beyond analyzing the smallest improvement, are needed in order to further improve the smoothed bounds on the running-time.

Approximation ratio. We have proved an upper bound of $O(\log 1/\sigma)$ for the smoothed approximation ratio of 2-Opt. Furthermore, we have proved that the lower bound of Chandra et al. [6] remains robust even for $\sigma = O(1/\sqrt{n})$. We leave as an open problem to generalize our upper bounds to the one-step model to improve the current bound of $O(\sqrt[4]{\phi})$ [11], but conjecture that this might be difficult, because of the special structure that Gaussian distributions provide.

While our bound significantly improves the previously known bound for the smoothed approximation ratio of 2-Opt, we readily admit that it still does not explain the performance observed in practice. A possible explanation is that when the initial tour is not picked by an adversary or the nearest neighbor heuristic, but using a construction heuristic such as the spanning tree heuristic or an insertion heuristic, an approximation factor of 2 is guaranteed even before 2-Opt has begun to improve the tour [25]. We chose to compare the worst local optimum to the global optimum in order, as this is arguably the simplest of all technically difficult possibilities.

However, a smoothed analysis of the approximation ratio of 2-Opt initialized with a good heuristic might be difficult: even in the average case, it is only known that the length of an optimal TSP is concentrated around $\gamma_d \cdot n^{\frac{d-1}{d}}$ for some constant $\gamma_d > 0$. But the precise value of γ_d is unknown [29]. Since experiments suggest that 2-Opt even with good initialization does not achieve an approximation ratio of 1 + o(1) [15, 16], one has to deal with the precise constants, which seems challenging.

Finally, we conjecture that many examples for showing lower bounds for the approximation ratio of concrete algorithms for Euclidean optimization such as the TSP remain stable under perturbation for $\sigma = O(1/\sqrt{n})$. The question remains whether such small values of σ , although they often suffice to prove polynomial smoothed running time, are essential to explain practical approximation ratios or if already slower decreasing σ provide a sufficient explanation.

References

- Milton Abramowitz and Irene A. Stegun, editors. Pocketbook of Mathematical Functions. Harri Deutsch, 1984.
- [2] Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.
- [3] Markus Bläser, Bodo Manthey, and B. V. Raghavendra Rao. Smoothed analysis of partitioning algorithms for Euclidean functionals. *Algorithmica*, 66(2):397–418, 2013.
- [4] Karl Bringmann, Christian Engels, Bodo Manthey, and B. V. Raghavendra Rao. Random shortest paths: Non-euclidean instances for metric optimization problems. In Krishnendu

Chatterjee and Jiří Sgall, editors, Proc. of the 38th Int. Symp. on Mathematical Foundations of Computer Science (MFCS), volume 8087 of Lecture Notes in Computer Science, pages 219–230. Springer, 2013.

- [5] Tobias Brunsch, Heiko Röglin, Cyriel Rutten, and Tjark Vredeveld. Smoothed performance guarantees for local search. *Mathematical Programming*, 146(1–2):185–218, 2014.
- [6] Barun Chandra, Howard Karloff, and Craig Tovey. New results on the old k-opt algorithm for the traveling salesman problem. *SIAM Journal on Computing*, 28(6):1998–2029, 1999.
- [7] Radu Curticapean and Marvin Künnemann. A quantization framework for smoothed analysis of Euclidean optimization problems. *Algorithmica*, 73(3):1–28, 2015.
- [8] Devdatt P. Dubhashi and Alessandro Panconesi. Concentration of Measure for the Analysis of Randomized Algorithms. Cambridge University Press, 2009.
- [9] Richard Durrett. Probability: Theory and Examples. Cambridge University Press, 2013.
- [10] Christian Engels and Bodo Manthey. Average-case approximation ratio of the 2-opt algorithm for the TSP. Operations Research Letters, 37(2):83–84, 2009.
- [11] Matthias Englert, Heiko Röglin, and Berthold Vöcking. Worst case and probabilistic analysis of the 2-Opt algorithm for the TSP. *Algorithmica*, 68(1):190–264, 2014. Corrected versions of Lemma 8 and 9 can be found at https://arxiv.org/abs/2302.06889.
- [12] Michael Etscheid. Performance guarantees for scheduling algorithms under perturbed machine speeds. Discrete Applied Mathematics, 195:84–100, 2015.
- [13] Merran Evans, Nicholas Hastings, and Brian Peacock. Statistical Distributions. John Wiley & Sons, 3rd edition, 2000.
- [14] Stefan Funke, Sören Laue, Zvi Lotker, and Rouven Naujoks. Power assignment problems in wireless communication: Covering points by disks, reaching few receivers quickly, and energy-efficient travelling salesman tours. Ad Hoc Networks, 9(6):1028–1035, 2011.
- [15] David S. Johnson and Lyle A. McGeoch. The traveling salesman problem: A case study. In Emile Aarts and Jan Karel Lenstra, editors, *Local Search in Combinatorial Optimization*, chapter 8. John Wiley & Sons, 1997.
- [16] David S. Johnson and Lyle A. McGeoch. Experimental analysis of heuristics for the STSP. In Gregory Gutin and Abraham P. Punnen, editors, *The Traveling Salesman Problem and its Variations*, chapter 9. Kluwer Academic Publishers, 2002.
- [17] David Karger and Krzysztof Onak. Polynomial approximation schemes for smoothed and random instances of multidimensional packing problems. In Proc. of the 18th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA), pages 1207–1216. SIAM, 2007.
- [18] Walter Kern. A probabilistic analysis of the switching algorithm for the TSP. Mathematical Programming, 44(2):213–219, 1989.

- [19] Marvin Künnemann and Bodo Manthey. Towards understanding the smoothed approximation ratio of the 2-opt heuristic. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, Proc. of the 42nd Int. Coll. on Automata, Languages and Programming (ICALP), volume 9134 of Lecture Notes in Computer Science, pages 859–871. Springer, 2015.
- [20] Bodo Manthey. Smoothed analysis of local search. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 285–308. Cambridge University Press, 2020.
- [21] Bodo Manthey and Heiko Röglin. Smoothed analysis: Analysis of algorithms beyond worst case. it – Information Technology, 53(6):280–286, 2011.
- [22] Bodo Manthey and Rianne Veenstra. Smoothed analysis of the 2-Opt heuristic for the TSP: Polynomial bounds for Gaussian noise. In Leizhen Cai, Siu-Wing Cheng, and Tak-Wah Lam, editors, Proc. of the 24th Ann. Int. Symp. on Algorithms and Computation (ISAAC), volume 8283 of Lecture Notes in Computer Science, pages 579–589. Springer, 2013.
- [23] 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.
- [24] Christos H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. Theoretical Computer Science, 4(3):237–244, 1977.
- [25] Daniel J. Rosenkrantz, Richard E. 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.
- [26] 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.
- [27] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: An attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):76–84, 2009.
- [28] Fred van Nijnatten, René Sitters, Gerhard J. Woeginger, Alexander Wolff, and Mark de Berg. The traveling salesman problem under squared Euclidean distances. In Jean-Yves Marion and Thomas Schwentick, editors, Proc. of the 27th Int. Symp. on Theoretical Aspects of Computer Science (STACS), volume 5 of LIPIcs, pages 239–250. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2010.
- [29] Joseph E. Yukich. Probability Theory of Classical Euclidean Optimization Problems, volume 1675 of Lecture Notes in Mathematics. Springer, 1998.