Smoothed Analysis of Resource Augmentation

Jeff Erickson[1], **Ivor van der Hoog**[2], **and Tillmann Miltzow**[2]

1 University of Illinois
`jeffe@illinois.edu`
2 Utrecht University
`i.d.vanderhoog@uu.nl`
`t.mitzow@gmail.com`

Abstract

The predominant approach to find decent solutions for hard optimization problems is to compute an approximation. An alternative approach is resource augmentation (a form of problem relaxation), where you consider an optimal solution subject to slightly weaker problem constraints. This alternative approach has considerably less traction in theoretical computer science than approximation algorithms have. We study optimization problems with natural resource augmentations and show that the bit-precision of their optimal solution can be bounded using smoothed analysis of their augmentation. Our results imply that for realistic problem constraints, the optimal solution to a augmented version of a problem yields an optimal solution for the original problem. We hope our results help solidify the traction that resource augmentation has in theoretical computer science.

1 Introduction

This paper is an extended abstract from [14]. The RAM is a mathematical model of a computer which emulates how a computer can manipulate data. Within computational geometry, algorithms are often analyzed within the real RAM [15, 18, 23] where values with infinite precision can be stored and compared in constant space and time. By allowing these infinite precision computations, it becomes possible to verify geometric primitives in constant time, which simplifies the analysis of geometric algorithms. Mairson and Stolfi [19] point out that "without this assumption it is virtually impossible to prove the correctness of any geometric algorithms." The downside of the real RAM is that it neglects the bit-precision of the underlying algorithms. If an algorithm can be correctly executed with a limited bit-precision then the algorithm is called *robust*. Many classical examples in computational geometry are inherently nonrobust [23].

Often inputs which require excessive bit-precision are contrived and do not resemble *realistic* inputs. A natural way to theoretically model this is smoothed analysis, which interpolates *smoothly* between worst case analysis and average case analysis [28]. Practical inputs are constructed inherently with small amount of noise and random perturbation. This perturbation helps to show performance guarantees in terms of the input size and and the magnitude of the perturbation. By now smoothed analysis is well-established, for instance Spielman and Teng received the Gödel

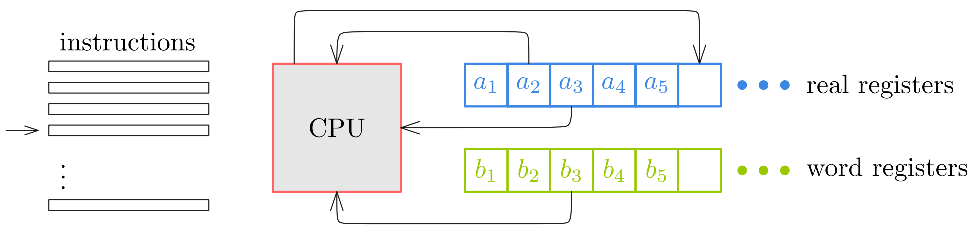

Figure 1. The dominant model in computational geometry is the real RAM. It consists of a central processing unit, which can operate on real and word registers in constant time, following a set of instructions.

36th European Workshop on Computational Geometry, Würzburg, Germany, March 16–18, 2020.
This is an extended abstract of a presentation given at EuroCG'20. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

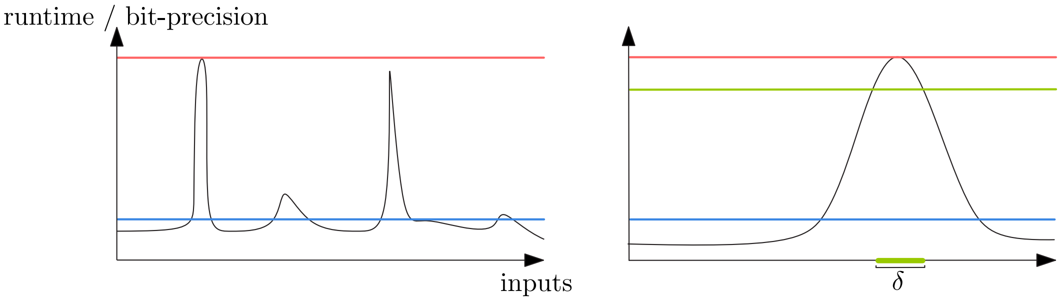

Figure 2. The x-axis symbolizes all inputs. The red line indicates the worse case running time or required bit-precision. The blue line indicates the average however, typical instances are not always average. Smoothed analysis considers the average of inputs near some worst instance (shown in green).

Prize for it. However, within computational geometry its application is limited to smoothed analysis of the bit-precision of the art gallery problem [10] and order type realisability [29], and smoothed analysis of the runtime of k-means clustering [3, 20], Euclidean TSP [12, 21], and partitioning algorithms for Euclidean functionals [5].

In this paper, we introduce a framework applicable to a wide class of real RAM optimization problems and show that under smoothed analysis of their resource augmentation, the optimal solution to these problems can be computed with logarithmic bit-precision. This is an extended abstract of Section 3 of [14].

Smoothed analysis. In *smoothed analysis*, the performance of an algorithm is studied for worst case input which is randomly perturbed by a magnitude of δ. Intuitively, smoothed analysis interpolates between average case and worst case analysis (Figure 2). The smaller δ, the closer we are to true worst case input. Correspondingly larger δ is closer to the average case analysis.

Formally, for smoothed analysis we fix some $\delta \in [0, 1]$, which describes the *magnitude of perturbation*. In this paper, we consider an array $I = (a, b) \in \mathbb{R}^n \times \mathbb{Z}^m$ of n real numbers and m integers as the input of the optimization problem (for an extensive overview of the real RAM model that takes both real and integer input refer to [14], A.1). We assume that each real number is perturbed independently and that the integers stay as they are. We denote by $(\Omega_\delta, \mu_\delta)$ the probability space where each $x \in \Omega_\delta$ defines for each instance I a new 'perturbed' instance $I_x = (a + x, b)$. We denote by $\mathscr{C}(I_x)$ the cost of instance I_x (note that traditionally, smoothed analysis is applied to algorithms where the cost of an instance is the runtime required by that algorithm on the instance. In this paper, the cost is the required number of bits to represent the optimal solution of the instance). The smoothed expected cost of instance I equals:

$$\mathscr{C}_\delta(I) = \mathop{\mathbb{E}}_{x \in \Omega_\delta} \mathscr{C}(I_x) = \int_{\Omega_\delta} \mathscr{C}(I_x) \mu_\delta(x) \, \mathrm{d}x.$$

If we denote by Γ_n the set of all instances of size n, then the smoothed complexity equals:

$$\mathscr{C}_{\text{smooth}}(n, \delta) = \max_{I \in \Gamma_n} \mathop{\mathbb{E}}_{x \in \Omega_\delta} [\mathscr{C}(I_x)].$$

Intuitively smoothed analysis shows that not only do the majority of instances behave nicely, but actually in every neighborhood (bounded by the maximal perturbation δ) the majority of instances behave nicely. The smoothed complexity is measured in terms of n and δ. If the expected complexity is small in terms of $1/\delta$ then we have a theoretical verification of the hypothesis that worst case examples are well-spread. Following [8, 28] we perceive an algorithm to have polynomial cost in practice, if the expected cost of the algorithm is polynomial in n and in $1/\delta$.

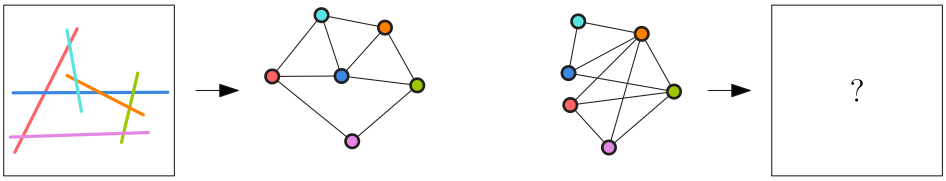

Figure 3. Left: given a set of segments S, they define a segment intersection graph G_S. Right: given a graph G, is there a set of segments S' such that $G_{S'} = G$?

Spielman and Teng explain smoothed analysis by applying it to the simplex algorithm, which was known for a particularly good performance in practice that was seemingly impossible to verify theoretically [16]. Since the introduction of smoothed analysis, it has been applied to numerous problems. Most relevant for us is the recent smoothed analysis of the art gallery problem [10] and of order types [29]. Both papers deal with the required bit-precision needed in computations under slight perturbations. In the worst case, both problems need an exponential bit-precision, as both problems are complete for the existential theory of the reals.

The Existential Theory of the Reals. The required precision of an algorithm plays an important role if we want to show that a problem lies in the class NP. It is often easy to describe a potential witness to an NP-hard problem, but the bit-precision of the witness is unknown. A concrete example is the recognition of segment intersection graphs (Figure 3): given a graph, can we represent it as the intersection graph of segments? The canonical witness is the set of segments, but the required bit-precision is unclear. Matoušek [22] comments on this as follows:

Serious people seriously conjectured that the number of digits can be polynomially bounded—but it cannot.

Indeed, there are examples which require an exponential number of bits in any numerical representation. This *exponential bit-precision phenomenon* occurs not only for segment intersection graphs, but also for many other natural algorithmic problems [1,2,4,6,7,9,11,13,24–27]. It turns out that all of those algorithmic problems do not accidentally require exponential bit-precision, but are closely linked, as they are all complete for a certain complexity class called $\exists\mathbb{R}$. Thus either all of those problems belong to NP, or none of them do. Using our results on smoothed analysis, we show that for many $\exists\mathbb{R}$-hard optimization problems the exponential bit-precision phenomenon only occurs for near-degenerate input.

The complexity class $\exists\mathbb{R}$ can be defined as the set of decision problems that are polynomial-time equivalent to deciding if a formula of the *Existential Theory of the Reals* (ETR) is true or not. An ETR formula has the form:

$$\Psi = \exists x_1, \ldots, x_n \quad \Phi(x_1, \ldots, x_n),$$

where Φ is a well-formed sentence over the alphabet

$$\Sigma = \{0, 1, x_1, \ldots, +, \cdot, =, \leq, <, \wedge, \vee, \neg\}.$$

More specifically, Φ is quantifier-free and $x_1, \ldots, x_n$ are all variables of Φ. We say Ψ is true if and only if there are real numbers $x_1, \ldots, x_n \in \mathbb{R}$ such that $\Phi(x_1, \ldots, x_n)$ is true.

2 Results of Smoothed Analysis of Resource Augmentation

Under the resource augmentation of an algorithmic problem, you try to find an optimal solution to a problem formulation with weaker problem constraints. Resource augmentation does not compromise on optimality: the aim is to find an optimal solution to the newly augmented problem.

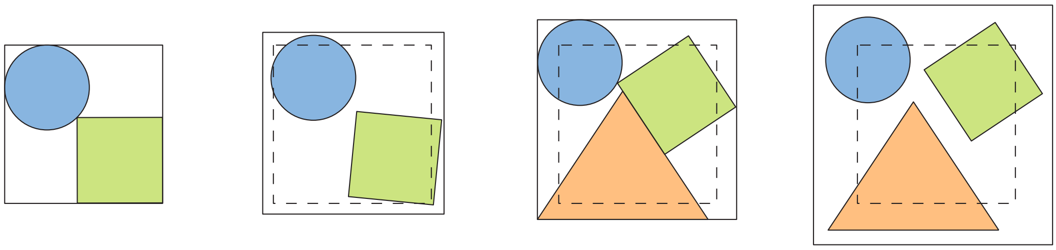

Figure 4. We augment the container from left to right. This extra space can lead to a better solution. If the optimal solution *value* does not change, the extra space allows for a solution with low bit-precision.

Using smoothed analysis, we argue that studying slight augmentations of algorithmic problems is justifiable for practical applications of the algorithm as we show that the problem conditions that make the problem hard are brittle.

An example of resource augmentation exists for the geometric packing problem (Figure 4) where an algorithm needs to pack a set of convex objects into a unit-size square container. To pack the optimal number of objects into this container is $\exists\mathbb{R}$-complete [2] and therefore a word RAM algorithm cannot hope to correctly find an optimal solution with limited time or memory. A resource augmentation algorithm looks to find a way to pack as many objects into a container C' which is larger by a factor $(1+\alpha)$ (α being the augmentation parameter). We apply smoothed analysis to resource augmentation problems, where we study these problems under a slight perturbation of such an augmentation parameter. We prove in [14] that the resource augmentation problems that we study have an optimal solution with expected logarithmic bit-precision.

► **Theorem 2.1.** *Let P be a resource augmentation problem that is monotonous, moderate and smoothable. Under perturbations of the augmentation of magnitude δ, the problem P has an optimal solution with an expected bit-precision of $O(\log(n/\delta))$.*

In the proof of this theorem (see full version) we argue about the solution space of the problem P and we define three natural properties of this solution space. The *monotonous property* demands that as we augment P more and more, the solution space only gains more candidate solutions. The *moderate property* demands that as we continuously augment the problem, we do not encounter more than a polynomial number of new optimums. In many hard optimization problems, the optimum is a value between 1 and n and the moderate property is then immediately implied. The *smoothable property* is the least intuitive of the three, it demands if you augment a problem P by ε, then it contains a solution which is optimal for the *original* problem and has a bit-precision of $O(\log(n/\varepsilon))$. It might appear as though the third property immediately implies the theorem, yet recall that we look for an optimal solution for the newly augmented problem. The other two properties, together with common bounds in probability theory, bound the expected bit-precision of an optimal solution to the perturbed problem.

Implications of Theorem 2.1. To illustrate the applicability of our findings, we give 3 corollaries:

The art gallery problem has been shown to be $\exists\mathbb{R}$-complete [1] which (assuming $\exists\mathbb{R} \neq NP$) prohibits a compact representation for all art gallery solutions. Yet our corollary states that under realistic conditions, the solution to the art gallery problem can be represented using logarithmic bit-precision. This result was already shown in [10], however with Theorem 2.1 this result can be re-proven by showing that the art gallery problem, with a resource augmentation of edge inflation is in fact monotonous, moderate and smoothable.

Recently Kostitsyna et al. showed that an optimal solution to the minimum-link path in a

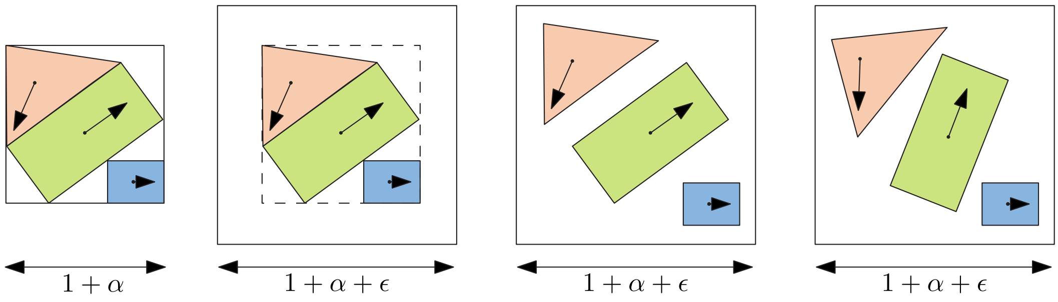

Figure 5. We increase a container of size $(1+\alpha)$ to size $(1+\alpha+\varepsilon)$. This extra space allows us to take the original solution, and space each object by a distance of $O(\varepsilon/n)$, which in turn allows us to find a more favourable rotation and / or translation for the object.

simple polygon has linear bit-precision in the worst case [17]. Just as the art gallery problem, this problem can be augmented by inflating the edges of the simple polygon. With a similar analysis, it then swiftly follows that the problem of computing the minimum-link path in a simple polygon is monotonous, moderate and smoothable.

The proof for $\exists\mathbb{R}$-completeness of the packing problems is in preparation [2] and just as for the art gallery problem this implies that the optimal solution to a packing problem cannot always be compactly represented. The packing problem has a natural resource augmentation, where one simply increases the size of the container. In the full version we show that the packing problem with container augmentation is monotonous, moderate and smoothable in the following way: if a container of size 1 can fit a collection I of items then a container of size $(1+\alpha)$ can also fit I and possibly more, thereby the monotonous property is trivial. If the input is a set of n elements that need to be packed in a container, then an optimal solution can pack at most k elements with $k \in [n]$. Therefore as we increase the container size continuously, there can be at most $O(n)$ new optimal solutions which implies the moderate property. The monotonous property is the hardest to show (Figure 5). In a solution to the packing problem, every object is rotated and translated and especially describing the rotation of an object is hard if you have to use limited bit-precision. In the full version we consider an optimal solution to a given container size, and show that if that container size increases by a value ε, then all the convex objects in the container can freely move and rotate a distance of $O(\varepsilon/n)$. This in turn, allows us to describe the translation and rotation of each object with a bit-precision of $O(\log(n/\varepsilon))$. Note that computing an embedding of an object with such a translation and rotation, might require more bit-precision.

► **Corollary 2.2.** *Under perturbations of the augmentation of magnitude δ, the following problems have an optimal solution with an expected bit-precision of $O(\log(n/\delta))$.*
- *the art gallery problem under perturbation of edge inflation [10].*
- *packing polygonal objects into a square container under perturbation of the container width.*
- *computing the minimum-link path in a simple polygon under perturbation of edge inflation.*

Limitations. We hope that Corollary 2.2 provides a convincing argument that our framework applies to a wide set of algorithmic problems that have a natural resource augmentation. Yet, our result is not wihout limiations: given an algorithmic problem, it is not clear *a priori* whether there is a way to augment resources such that it is both mathematically sound, satisfying, as well as practically plausible. For example, if we search for the smallest square container that fits a given set of items, the number of changes in the optimum is unbounded thus the moderate property does not hold.

Acknowledgments. The third author acknowledges the generous support of the NWO Veni grant EAGER.The second author acknowledges the support of the NWO grant 614.001.504.

References

1 Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$-complete. In *STOC*, pages 65–73, 2018.

2 Mikkel Abrahamsen, Tillmann Miltzow, and Nadja Seiferth. A framework for $\exists\mathbb{R}$-completeness of two-dimensional packing problems. *in Preparation*, 2020.

3 David Arthur, Bodo Manthey, and Heiko Röglin. k-means has polynomial smoothed complexity. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 405–414. IEEE, 2009.

4 Daniel Bienstock. Some provably hard crossing number problems. *Discrete & Computational Geometry*, 6(3):443–459, 1991.

5 Markus Bläser, Bodo Manthey, and Raghavendra Rao. Smoothed analysis of partitioning algorithms for euclidean functionals. *Algorithmica*, 66(2):397–418, 2013.

6 Jean Cardinal, Stefan Felsner, Tillmann Miltzow, Casey Tompkins, and Birgit Vogtenhuber. Intersection graphs of rays and grounded segments. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 153–166. Springer, 2017.

7 Jean Cardinal and Udo Hoffmann. Recognition and complexity of point visibility graphs. *Discrete & Computational Geometry*, 57(1):164–178, 2017.

8 Daniel Dadush and Sophie Huiberts. A friendly smoothed analysis of the simplex method. In *STOC*, pages 390–403. ACM, 2018.

9 Michael G. Dobbins, Linda Kleist, Tillmann Miltzow, and Paweł Rzążewski. $\forall\exists\mathbb{R}$-completeness and area-universality. *ArXiv 1712.05142*, 2017.

10 Michael Gene Dobbins, Andreas Holmsen, and Tillmann Miltzow. Smoothed analysis of the art gallery problem. *CoRR*, abs/1811.01177, 2018. arXiv:1811.01177.

11 Michael Gene Dobbins, Andreas Holmsen, and Tillmann Miltzow. A universality theorem for nested polytopes. *arXiv*, 1908.02213, 2019.

12 Matthias Englert, Heiko Röglin, and Berthold Vöcking. Worst case and probabilistic analysis of the 2-opt algorithm for the tsp. In *SODA*, pages 1295–1304, 2007.

13 Jeff Erickson. Optimal curve straightening is $\exists\mathbb{R}$-complete. *arXiv:1908.09400*, 2019.

14 Jeff Erickson, Ivor van der Hoog, and Tillmann Miltzow. A framework for robust realistic geometric computations. *CoRR*, abs/1912.02278, 2019. arXiv:1912.02278.

15 Steven Fortune and Christopher Van Wyk. Efficient exact arithmetic for computational geometry. In *Proceedings of the ninth annual symposium on Computational geometry*, pages 163–172. ACM, 1993.

16 Victor Klee and George Minty. How good is the simplex algorithm. Technical report, Washington Univ. Seattle Dept. of Mathematics, 1970.

17 Irina Kostitsyna, Maarten Löffler, Valentin Polishchuk, and Frank Staals. On the complexity of minimum-link path problems. *Journal of Computational Geometry*, 8(2):80–108, 2017.

18 Chen Li, Sylvain Pion, and Chee-Keng Yap. Recent progress in exact geometric computation. *The Journal of Logic and Algebraic Programming*, 64(1):85–111, 2005.

19 Harry Mairson and Jorge Stolfi. Reporting and counting intersections between two sets of line segments. In *Theoretical Foundations of Computer Graphics and CAD*, pages 307–325. Springer, 1988.

20 Bodo Manthey and Heiko Röglin. Improved smoothed analysis of the k-means method. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 461–470. Society for Industrial and Applied Mathematics, 2009.

21 Bodo Manthey and Rianne Veenstra. Smoothed analysis of the 2-opt heuristic for the tsp: Polynomial bounds for gaussian noise. In *International Symposium on Algorithms and Computation*, pages 579–589. Springer, 2013.

22 Jiří Matoušek. Intersection graphs of segments and $\exists\mathbb{R}$. 2014. ArXiv 1406.2636.

23 David Salesin, Jorge Stolfi, and Leonidas Guibas. Epsilon geometry: building robust algorithms from imprecise computations. In *Proceedings of the fifth annual symposium on Computational geometry*, pages 208–217. ACM, 1989.

24 Marcus Schaefer. Complexity of some geometric and topological problems. In *Proceedings of the 17th International Symposium on Graph Drawing (GD 2009)*, Lecture Notes in Computer Science (LNCS), pages 334–344. Springer, 2009.

25 Marcus Schaefer. Realizability of graphs and linkages. In János Pach, editor, *Thirty Essays on Geometric Graph Theory*, chapter 23, pages 461–482. Springer-Verlag New York, 2013.

26 Marcus Schaefer and Daniel Štefankovič. Fixed points, Nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, 60(2):172–193, 2017.

27 Yaroslav Shitov. A universality theorem for nonnegative matrix factorizations. 2016, journal=ArXiv 1606.09068.

28 Daniel Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.

29 Ivor van der Hoog, Tillmann Miltzow, and Martijn van Schaik. Smoothed analysis of order types. *arXiv:1907.04645*, 2019.