Computational Geometry: Young Researchers Forum (CG:YRF'17)

Book of Abstracts

This volume contains the abstracts of papers at "Computational Geometry: Young Researchers Forum" (CG:YRF), a satellite event of the 33rd Symposium on Computational Geometry, held in Brisbane, Australia on July 4-7, 2017.

The CG:YRF program committee consisted of the following people:

- Therese Biedl, University of Waterloo, Canada (chair)
- Arijit Bishnu, Indian Statistical Institute, India
- Stephane Durocher, University of Manitoba, Canada
- Seok-hee Hong, University of Sydney, Australia
- John Iacono, New York University, USA
- Clement Maria, University of Queensland, Australia
- David Mount, University of Maryland, USA
- Yoshio Okamoto, The University of Electro-Communications, Japan

There were 8 papers submitted to CG:YRF. Of these, 6 were accepted with minor revisions, and the remaining 2 were accepted after a major revision and a second review. One paper had to be withdrawn later due to difficulties in obtaining a visa in a timely manner.

Persistent Nerves Revisited

Nicholas J. Cavanna and Donald R. Sheehy

University of Connecticut

1 Introduction

A nerve is a simplicial complex derived from a cover of a topological space. Nerves appear all over computational topology and geometry, e.g. a Delauney triangulation is the nerve of a Voronoi diagram and the Čech complex is the nerve of a collection of metric balls. They are used to solve problems concerning surface reconstruction, homology inference, and homological sensor networks, among other areas.

If one has an open cover of a paracompact space in which all non-empty intersections of the cover elements are contractible, i.e. it is a *good cover*, then its nerve is homotopy equivalent to the covered space. This result is known as the Nerve Theorem. It has a natural extension to the setting of persistent homology called the Persistent Nerve Lemma (PNL), due to Chazal and Oudot [2]. The PNL implies that given a filtration of covers of a filtration of spaces such that at each time the cover is good, then the persistent homology of the space filtration is that of the nerve filtration. Good covers are not always an option, e.g. if a metric space is not convex then the metric balls of a finite point sample may cover the space, but they won't be a good cover without adding other conditions. The requirement of having a good cover in order to invoke the PNL is the motivation for our work – instead we assume the cover elements' homology is trivial when included into a later scale.

Recently, Botnan and Spreemann assumed an interleaving between two cover filtrations to prove a bound on the bottleneck distance between the persistence diagrams of their nerve filtrations [1]. Govc and Skraba [3] considered a simplicial filtration and a cover of the terminal complex, using it to construct a cover filtration by taking the intersection of the cover with the simplicial filtration at each scale. They assumed that the persistence modules of all the non-empty k-wise intersections of the filtration's cover elements were ε-interleaved with 0 and proved a tight bound on the bottleneck distance between the persistence diagrams of the filtered simplicial complex and the nerve filtration, linear with respect to dimension and ε.

We consider a more general cover assumption—that we have an arbitrary cover filtration consisting of simplicial complexes which collectively cover a filtered simplicial complex, rather than a cover filtration induced by intersection like Govc and Skraba. We also assume the cover filtration is what we define as *ε-good*. The major contributions are that there are interleavings between the filtered simplicial complex and the nerve filtration's homology groups, induced by chain maps, which imply a tight dimension-dependent bound on the bottleneck distance between the two persistence diagrams, linear with respect to dimension and ε. A noteworthy corollary of our result is the Persistent Nerve Lemma by considering a 0-good cover.

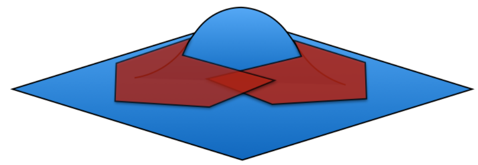

Figure 1: A bump in $\mathbb{R}^3$ and subsets with non-contractible intersection.

2 Background

Let $\mathcal{U} := \{U_1, \ldots, U_n\}$ be an arbitrary collection of *filtrations*, growing sequences of spaces, where $U_i := (U_i^\alpha)_{\alpha \geq 0}$, and each U_i^α is a simplicial complex. We call $\mathcal{U}$ a *cover filtration*. For each $\alpha \geq 0$, define $\mathcal{U}^\alpha := \{U_1^\alpha, \ldots, U_n^\alpha\}$ and $W^\alpha := \bigcup_{i\in[n]} U_i^\alpha$. For each non-empty $v \subseteq [n] = \{1, \ldots, n\}$, let $U_v^\alpha := \bigcap_{i\in v} U_i^\alpha$.

The *nerve* of the cover $\mathcal{U}^\alpha$ is defined as $\operatorname{Nrv} \mathcal{U}^\alpha := \{v \subseteq [n] \mid U_v^\alpha \neq \emptyset\}$. One can check this is a simplicial complex. The *nerve filtration* is defined as $\operatorname{Nrv} \mathcal{U} := (\operatorname{Nrv} \mathcal{U}^\alpha)_{\alpha\geq 0}$. When we consider the collection of spaces that each $\mathcal{U}^\alpha$ covers over all $\alpha \geq 0$, we get the *union filtration*, $\mathcal{W} := (W^\alpha)_{\alpha\geq 0}$. $\mathcal{U}^\alpha$ is a *good cover* of W^α if for all subsets $v \subseteq [n]$, we have U_v^α is empty or contractible. For filtrations, we say $\mathcal{U}$ is a *good cover* of $\mathcal{W}$ if $\mathcal{U}^\alpha$ is a good cover of W^α for all $\alpha \geq 0$. The Persistent Nerve Lemma implies that if $\mathcal{U}$ is a good cover of $\mathcal{W}$, then $\operatorname{Dgm}(\operatorname{Nrv} \mathcal{U}) = \operatorname{Dgm}(\mathcal{W})$, where $\operatorname{Dgm}(\cdot)$ is the *persistence diagram* over all dimen-

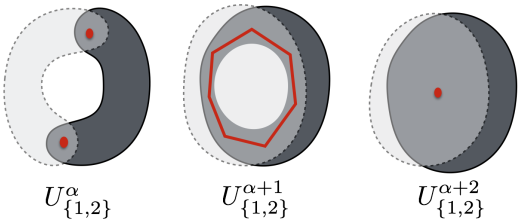

Figure 2: A cover filtration that is not good, but is 1-good.

sions of the input filtration– a multiset representing the "birth" and "deaths" of homological features as $\alpha \to \infty$. When we refer to the persistence diagram of just the k-dimensional homological features we will write $\mathrm{Dgm}_k(\cdot)$.

Given a cover filtration $\mathcal{U}$, we say it is *ε-good* if for all non-empty $v \subseteq [n]$, and for all $\alpha \geq 0$, $\tilde{H}_*(U_v^\alpha \hookrightarrow U_v^{\alpha+\varepsilon}) = 0$, so any nontrivial homology classes of U_v^α are trivial when mapped into $U_v^{\alpha+\varepsilon}$. Note that due to the definition of contractibility, $\mathcal{U}$ being a good cover of $\mathcal{W}$ implies that it is 0-good.

For each $\mathcal{U}^\alpha$ there is a corresponding commutative diagram $D\mathcal{U}^\alpha$, where the spaces are the non-empty sets U_v^α for non-empty $v \subseteq [n]$ and there is an inclusion map $U_v^\alpha \hookrightarrow U_{v'}^\alpha$ whenever $v' \subset v$. Let N^α be the barycentric subdivision of $\mathrm{Nrv}\,\mathcal{U}^\alpha$ which has simplices of the form $\sigma = v_0 \to \ldots \to v_k$, where $v_i \subset v_{i+1}$, and each v_i corresponds to a simplex of $\mathrm{Nrv}\,\mathcal{U}^\alpha$. This is an abstract simplicial complex and we denote the associated geometric filtration as $\mathcal{N} := (|N^\alpha|)_{\alpha \geq 0}$. We define the *homotopy colimit* of $D\mathcal{U}^\alpha$ as

$$\mathrm{hocolim}\; D\mathcal{U}^\alpha := \bigcup_{N^\alpha \ni \sigma = v_0 \to \ldots} U_{v_0}^\alpha \times |\sigma|,$$

where $|\cdot|$ is the geometric realization functor. This homotopy colimit is also known as the Meyer-Vietoris blowup complex [4]. It yields another filtration, $\mathcal{B} = (B^\alpha)_{\alpha \geq 0}$, where $B^\alpha = \mathrm{hocolim}\; D\mathcal{U}^\alpha$.

We note that there is a (pseudo)-metric between two persistence diagrams D and D' called the *bottleneck distance*, denoted $\mathrm{d}_B(D, D')$, which is the standard measure of the similarity of two persistence diagrams, and with that, the persistent homology of two filtrations.

3 Results

Theorem 1 *If $\mathcal{U} = \{U_1, \ldots, U_n\}$ is a set of simplicial filtrations that is an ε-good cover of the simplicial filtration $\mathcal{W} = \bigcup_{i=1}^n U_i$, then*

$$\mathrm{d}_B(Dgm_k(\mathcal{W}), Dgm_k(\mathrm{Nrv}\,\mathcal{U})) \leq \frac{(k+1)\varepsilon}{2}.$$

Furthermore, there is an upper-bound of $\frac{(D+1)\varepsilon}{2}$, where D is the dimension of the nerve filtration.

An overview of the proof is as follows. As for all α, $B^\alpha \subseteq W^\alpha \times |N^\alpha|$, we have natural chain maps induced by projection $b^\alpha : C_*(B^\alpha) \to C_*(W^\alpha)$ and $p^\alpha : C_*(B^\alpha) \to C_*(|N^\alpha|)$, where b^α at the space level is a homotopy equivalence, so $\mathrm{Dgm}(\mathcal{W}) = \mathrm{Dgm}(\mathcal{B})$.

Since $|N^\alpha|$ is homeomorphic to $|\mathrm{Nrv}\,\mathcal{U}^\alpha|$, it follows that $\mathrm{Dgm}(\mathcal{N}) = \mathrm{Dgm}(|\mathrm{Nrv}\,\mathcal{U}|)$, and as simplicial homology is equivalent to singular homology, we have that $\mathrm{Dgm}(\mathcal{N}) = \mathrm{Dgm}(\mathrm{Nrv}\,\mathcal{U})$. Define $t := (k+1)\varepsilon$, where k is the maximal dimension of homology groups being considered. We create a chain map $q^\alpha : C_*(|N^\alpha|) \to C_*(W^{\alpha+t})$ such that $a^\alpha := q^\alpha \circ p^\alpha$ is chain homotopic to $i_{\mathcal{B}}^{\alpha,\alpha+t} \circ b^\alpha$, via chain homotopy $c^\alpha : C_*(|N^\alpha|) \to C_*(W^{\alpha+t})$. These maps can be viewed in diagram 1.

$$\begin{tikzcd}
C_k(W^\alpha) \arrow[r, "i_{\mathcal{W}}^{\alpha,\alpha+t}"] & C_k(W^{\alpha+t}) \\
C_k(B^\alpha) \arrow[u, "b^\alpha"] \arrow[d, "p^\alpha"'] \arrow[r, "i_{\mathcal{B}}^{\alpha,\alpha+t}"'] & C_k(B^{\alpha+t}) \arrow[u, "b^{\alpha+t}"'] \arrow[d, "p^{\alpha+t}"] \\
C_k(|N^\alpha|) \arrow[r, "i_{\mathcal{N}}^{\alpha,\alpha+t}"] \arrow[uur, "q^\alpha"] & C_k(|N^{\alpha+t}|)
\end{tikzcd} \tag{1}$$

We use q^α to define a chain map $\bar{q}^\alpha : C_*(|N^\alpha|) \to C_*(B^{\alpha+t})$, where $\bar{q}^\alpha(\sigma) := \sum_{i=0}^{k} q(\sigma_i) \otimes \bar{\sigma}_i$, with $\sigma_i := v_0 \to \ldots \to v_i$ and $\bar{\sigma}_i := v_i \to \ldots \to v_k$, such that $p^{\alpha+t} \circ \bar{q}^\alpha$ commutes with $i_{\mathcal{N}}^{\alpha,\alpha+t}$ and $\bar{q}^\alpha \circ p^\alpha$ is chain homotopic to $i_{\mathcal{B}}^{\alpha,\alpha+t}$, via chain homotopy $\bar{c}^\alpha$, defined analogously to $\bar{q}^\alpha$.

By applying the homology functor to the diagram at all α, the chain maps p^α and $\bar{q}^\alpha$ commute with all the inclusion homomorphisms, forming interleaving homomorphisms between $\mathcal{N}$ and $\mathcal{B}$ thus implying our result.

References

[1] Magnus Bakke Botnan and Gard Spreemann. Approximating persistent homology in Euclidean space through collapses. *Applicable Algebra in Engineering, Communication and Computing*, pages 1–29, 2015.

[2] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SCG '08, pages 232–241, New York, NY, USA, 2008. ACM.

[3] Dejan Govc and Primoz Skraba. An approximate nerve theorem. *arXiv*, https://arxiv.org/pdf/1608.06956v2.pdf, 2016.

[4] Afra Zomorodian and Gunnar Carlsson. Localized homology. *Computational Geometry: Theory and Applications*, 41(3):126–148, 2008.

Towards Canonical Ideal Triangulations of Convex Projective Surfaces via Area Estimates

Dominic Tate

Department of Mathematics and Statistics, The University of Sydney

1 Introduction

This paper outlines a method for calculating a canonical ideal triangulation of a member of a family of non-compact surfaces.[1] Canonical triangulations and methods for their construction such as [5] have applications in a diverse array of problems including the uniformisation of Riemann surfaces and graphical problems such as discretizations of texture analysis and shape classification.

The Hilbert metric and Hilbert area are notions of length and area which may be defined on convex, open, bounded domains in $\mathbb{R}^2$.[2] A surface of the kind described above, hereafter denoted S, may be imbued with a length and area locally isometric to the Hilbert length and area of some domain. An identification of this kind is called a convex projective structure on S. A convex projective structure on a surface is uniquely determined by the generalized shear parameters defined by Fock and Goncharov in [3].

Cooper and Adeboye in [1] use the generalized shear parameters to assign to each triangulation of S, a real number which bounds below the Hilbert area of the given surface. This provides a function on the set of all ideal triangulations which serves as a tool for determining if a particular triangulation is canonically determined by the geometry. The author shows that this assignment does not constitute a uniform bound on that area. An improvement on [1] for estimating the area of such surfaces is subsequently constructed by the author. This function has the benefit of avoiding cases appearing in [1] where each triangulation is assigned the same value. Moreover, this work serves as a means of estimating the area of the surface with respect to the Hilbert geometry. This area is particularly difficult to calculate with exactitude as the Hilbert area form depends intrinsically on the domain in which the area is calculated.

[1]The particular surfaces under consideration are non-compact surfaces of negative Euler characteristic, finite topological type and empty boundary. Definitions of these terms may be found, for example, in [4].

[2]The formal definition of this geometry and the subsequent geometric structures on S can be found in, for example, [3].

2 Canonical Triangulations and Edge Flips

It follows from the definition above that S has some number of punctures.[3] An ideal triangulation of S is the result of 'filling in' the aforementioned punctures, triangulating the resulting surface with vertices only at the filled in points, then removing the punctures once again. Hereafter the word triangulation will refer to ideal triangulations and the word 'ideal' will be omitted.

Fix a triangulation Π of S. Every edge E in Π is shared by two triangles which together form a quadrilateral Q of which E is one diagonal. An edge flip is the process of removing E and replacing it with the other diagonal of Q, thus forming a new triangulation as in Figure 1. It is well-known that any two triangulations of S differ by a finite sequence of these moves. Two triangulations will be considered equivalent if their edges differ by ambient isotopy of S. While any two triangulations of the once-punctured torus are combinatorially equivalent, there are countably many ambient isotopy classes of triangulations, see for example, [2].

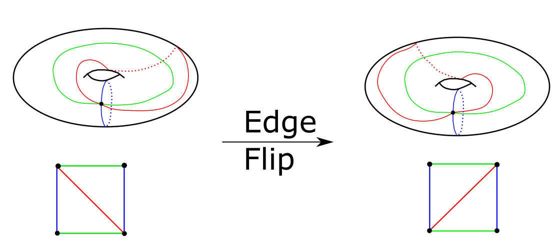

Figure 1: An edge flip on the once-punctured torus.

3 Different Geometric Structures on a Surface

Fock and Goncharov [3] devise a set of coordinates on the set of isomorphism classes of convex projective structures on a fixed surface S. These coordinates are well-suited to comparing triangulations of Ω and for use

[3]For the purposes of this section they may be thought of as marked points but it is pertinent to the geometry of S that the construction proceeds with punctures as opposed to marked points.

in estimating the Hilbert area of ideal triangles. One parameter, called a triple ratio, is assigned to each ideal triangle τ in Π and denoted $\mathrm{Tr}(\tau)$. Two parameters, called edge ratios, are assigned to each edge.

Theorem 1 *Cooper and Adeboye [1]*
Let S be a surface with convex projective structure. If Θ is the set of ideal triangles in Π

$$\phi(S) := \frac{1}{8}\sum_{T\in\Theta}\left(\pi^2 + \log^2\left(Tr(T)\right)\right) < Area(S)$$

If there exists a triangulation of S at which ϕ attains a global maximum or minimum then it may be used as a canonical triangulation for the purpose of classifying convex projective structures on S. Note that ϕ depends only on triple ratios, it will take the same value on different triangulations with the same set of triple ratios.

4 Area Estimates and Canonical Triangulations

The author presents two results demonstrating the use and limitations of ϕ as a tool for estimating area and determining a canonical triangulation. Thereafter, ϕ is explicitly calculated on the graph of triangulations of a torus with fixed convex projective structure. This informs the subsequent construction of an alternative to ϕ for estimating the Hilbert area of an ideal triangle in a convex domain.

The author has shown that for S a surface with convex projective structure,

$$\mathrm{Area}(S) - \phi(S)$$

is not uniformly bounded. This makes essential use of the fact that hyperbolic geometry is a special case of Hilbert geometry, as shown in [3]. In this case, all ideal triangles are isometric and consequently have area π and triple ratio equal to 1. Taking a sequence of surfaces with hyperbolic structure and increasing number of triangles in their triangulation provides a counterexample to uniform boundedness. It is shown in [3] that the triple ratios resulting from an edge flip are rational polynomials in the pre-flip triple ratios and edge ratios. It follows by induction on edge flips that the infimum of ϕ is not always achieved on some triangulation. These facts limit the efficacy of efforts to estimate area or yield a canonical triangulation from ϕ.

An example in which ϕ is more informative about the underlying geometry of the given surface is shown in Figure 2. Each vertex in the graph represents an ideal triangulation of the once-punctured torus. Two vertices are joined by an edge if and only if their respective triangulations differ by a single edge flip. In the image, vertices are labelled with the value of the largest triple ratio of the two in its corresponding triangulation. Either triple ratio uniquely determines ϕ because direct calculation shows that the product of the two triple ratios in a finite-area structure is 1. A priori there is nothing special about the vertex labelled 2, about which the labelled graph has rotational symmetry. This vertex represents a triangulation for which ϕ is a better area approximation than any other.

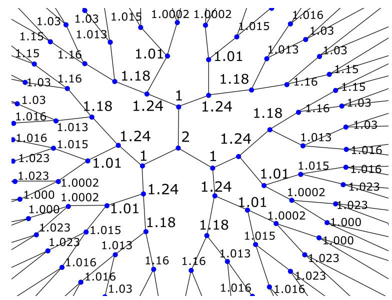

Figure 2: A reparameterisation of ϕ on the flip graph of a once-punctured torus.

If S was chosen to have a hyperbolic geometry, all possible triangles have triple ratio equal to 1, so ϕ will take a constant value, regardless of changes in triangulation. It does not, in this case, distinguish a particular triangulation. An alternative function to ϕ has been constructed by the author which takes the edge ratios as input, in order to make such a dinstinction. An analysis of the edge-flip functions given in [3] shows that unlike triple ratios, the set of edge ratios is never constant under edge-flips so this function is better able to distinguish between differen triangulations. Having constructed a more discerning function on the flip graph, these two functions can be used as tools for computing effective estimates for the Hilbert area of an ideal triangle in a bounded, convex domain in $\mathbb{R}^2$.

References

[1] Cooper, D. and Adeboye, I, *The Area of Convex Projective Surfaces and Fock-Goncharov Coordinates*, arXiv pre-print:1506.08245 [math.GT] (20 May, 2016), 1-10.

[2] Farb, Benson and Margalit, Dan, *A Primer on Mapping Class Groups*, Princeton Mathematical Series, Vol. 49, Princeton University Press, Princeton, NJ (2012).

[3] Fock, V. and Goncharov, A. B., *Moduli Spaces of Convex Projective Structures on Surfaces*, Advances in Mathematics 208, 1 (2007), 249-273.

[4] Munkres, James R., *Topology*, Prentice Hall Incorporated (2000).

[5] Epstein, D.B.A. and Penner, R.C., *Euclidean Decompositions of Noncompact Hyperbolic Manifolds*, J. Differential Geometry 27, 1 (1988), 67-80.

Length spectrum estimates for hyperbolic 3-manifolds

Robert C. Haraway III, Neil Hoffman, Matthias Görner, and Maria Trnkova

Many 3-manifolds are *hyperbolic*, admitting a complete metric locally isometric to hyperbolic space. Weeks and others have developed and implemented procedures to estimate this metric given a triangulation of a 3-manifold, and to estimate properties of this metric ([3], [5]). However, these procedures rely on floating-point calculations, and do not provide bounds on the error of their estimates. Recent progress has already been made in providing such error bounds on estimates of the actual hyperbolic metric and of some basic properties of the metric, like volume ([4]). The current project builds on this work, using interval arithmetic to provide rigorous estimates of the length spectrum.

1 Basic facts and definitions

A 3-triangulation, or, for the purposes of this paper, simply a *triangulation*, is a face-pairing T of finitely many tetrahedra. Identifying the sides of these simplices according to the face-pairing yields a topological space, which we also call T. In this abstract we will assume every point in T admits a neighborhood homeomorphic to a ball. That is, we assume T is a *3-manifold*. Furthermore, we will assume T is orientable.

Suppose that the tetrahedra of T are realized as geodesic tetrahedra of hyperbolic space H^3, and suppose that each face-pairing map is an isometry of H^3 taking one face to another. Then the metrics on the tetrahedra descend to a metric on T. If the dihedral angles around every edge class add up to 2π, then the resulting metric restricted to $T \setminus V(T)$ is locally isometric to H^3—it is a *hyperbolic* metric. For each $v \in V(T)$, there is an additional "completion" condition c_v involving only the tetrahedra incident to v, analogous to the 2π requirement on dihedral angles. The metric on $T \setminus V(T)$ completes to a hyperbolic metric on T if and only if for all $v \in V(T)$, c_v holds true. A 3-manifold T with a complete, finite-volume hyperbolic metric is a *hyperbolic 3-manifold*. The multiset of lengths of simple closed geodesics in T with length at most L is the *length spectrum* $\mathcal{L}(T)|_L$ of T *up to* L.

A hyperbolic 3-manifold T is determined by the isometry classes of its tetrahedra. One can specify the isometry class of a tetrahedron by a few parameters. For instance, compact tetrahedra are determined by their edge lengths. Conversely, the condition that edges of given length fit together into a hyperbolic tetrahedron is some analytic inequality on these lengths. Furthermore, the 2π-restrictions on dihedral angles around edge classes and the completion conditions are likewise analytic equations on the edge lengths. A hyperbolic structure on T is a solution to this system of analytic conditions, or, in interval arithmetic, a convergent sequence of interval approximations to such a solution.

We can identify every hyperbolic 3-manifold T as the quotient of H^3 by the action of a discrete group of isometries Γ. For $g \in \Gamma$ and $U \subset H^3$ we write $gU = \{g(u) : u \in U\}$; we also write $\Gamma U = \bigcup_{g\in\Gamma} gU$. Because T is a 3-manifold, every nontrivial $g \in \Gamma$ is nontrivial translation along and possibly trivial rotation around some geodesic of H^3, called its *axis* $\alpha(g)$. The length of the translation is called $\ell(g)$. If S is a subset of Γ, we write $\ell(S) = \{\ell(g) : g \in \Gamma\}$, a multiset. If γ is a simple closed geodesic in T and $\tilde{\gamma}$ lifts γ, there is $g \in \Gamma$ such that $\tilde{\gamma} = \alpha(g)$ and $\ell(g) = \ell(\gamma)$.

A *fundamental domain* for Γ acting on H^3 is a polyhedron $D \subset H^3$ such that (cf. [2], p. 259)

- $\bigcup_{g\in\Gamma} gD = H^3$;
- for all $g, g' \in \Gamma$, $gD = g'D$ implies $g = g'$; and
- for all $p \in H^3$, there is $\epsilon > 0$ such that the ball $B_\epsilon(p)$ intersects gD for only finitely many $g \in \Gamma$.

We say a *skeleton* of T is a finite 2-complex in T which intersects every simple closed geodesic. Suppose $\pi : H^3 \to T$ is the quotient map, and suppose D is a fundamental domain of Γ. If K is a skeleton, then we call $\pi^{-1}(K) \cap D$ a *sinew* for Γ in D. For instance, if the interior of D is a ball, then ∂D descends to a skeleton in T, so $\partial D = \pi^{-1}(\pi(\partial D)) \cap D$ is a sinew for Γ in D.

For every skeleton K, every simple closed geodesic γ of T intersects K. Therefore, every lift $\tilde{\gamma}$ of γ intersects $\pi^{-1}(K)$ at some point p. Translating p to D via $g \in \Gamma$, we see that $g\tilde{\gamma}$ is a lift of γ intersecting $\pi^{-1}(K) \cap D$. Thus, *for any sinew X of Γ in D, every simple closed geodesic in T admits a lift intersecting X.*

2 Tiling procedure for length spectrum superset

We recently developed the following procedure inspired by the work of Hodgson and Weeks in [3]. Let $L > 0$, and let T be a hyperbolic 3-manifold as above.

- Let T, Γ, D, and X be as above. Initialize a polyhedron $H := D$.

- Let S be $\{g \in \Gamma \mid gD \subset H \wedge g \neq id\}$. Let B be ∂H. The following loop maintains these definitions.
- While there is a face F of B with $d(X, F) \leq L$:
  - Let g be the element of Γ taking D to the "other side" of F.
  - Set $S, H, B := S \cup \{g\}, H \cup gD, \partial(H \cup gD)$.

Lemma 1 *If the procedure terminates, then $\ell(S)$ contains $\mathcal{L}(T)|_L$, and $\min(\ell(S))$ is the systole length.*

Proof. If the procedure terminates, $d(X, F) > L$ for all faces F of B. So $d(X, B) > L$. Suppose γ is a simple closed geodesic in T. Some lift $\tilde{\gamma}$ intersects X in a point q. Let $g \in \Gamma$ be such that $\tilde{\gamma} = \alpha(g)$ and $\ell(g) = \ell(\gamma)$. Then $\ell(\gamma) = d(q, g(q)) \geq d(X, gX)$. If $g \notin S$, then by definition of S, $gD \not\subseteq H$; in particular, $gX \not\subseteq H$. Thus, $d(X, gX) \geq d(X, B)$. Therefore, if $g \notin S$, then $\ell(\gamma) > L$. So if $\ell(\gamma) \leq L$, then $g \in S$. Now, each g corresponds to exactly one γ. So $\ell(S)$ contains $\mathcal{L}(T)|_L$.

Suppose $h \in \Gamma$ has smallest length among nontrivial elements. Then by [1], $\alpha(h)$ descends to a shortest simple closed geodesic η, and $\ell(h) = \ell(\eta)$. There is some lift $\tilde{\eta}$ intersecting X; let $h' \in \Gamma$ such that $\alpha(h') = \tilde{\eta}$ and $\ell(h') = \ell(\eta)$. Then by the above argument, $h' \in S$. So $\min(\ell(S)) = \ell(\eta)$, the systole length. □

Refining S until $\ell(S) = \mathcal{L}(T)|_L$ is work in progress. The above procedure is the first to provide interval bounds on systole length without an appeal to exact arithmetic, as far as the authors know. We should now address some potential concerns about implementing the above procedure in interval arithmetic.

3 Interval arithmetic

First, the initial definitions. Computing generators of Γ from the metric on T is straightforward, and generalizes easily to interval arithmetic. Less straightforward are choices of D and X. Our procedure allows any such choices as input. This contrasts with the approach of Hodgson and Weeks in [3], which begins with a calculation of a *Dirichlet domain* D, the calculation of which has, historically, been expensive. (New methods in numerical analysis for solving overdetermined systems may reduce this cost.) Our method doesn't require this calculation. Instead, one can build up a suitable D by repeatedly gluing on tetrahedra in the face-pairing using face-pairing maps. Defining a sinew X is also straightforward, using a "butterfly" in each tetrahedron, and ensuring that the butterflies glue up in T (see Fig. 1).

Second, the test $d(X, F) \leq L$ in fact goes as follows. We have some approximation $[d_l, d_r]$ of $d(X, F)$ with $d_l, d_r \in \mathbf{Q}$. We actually test $d_l \leq L$ (assuming $L \in \mathbf{Q}$). For termination, we argue that eventually, either for all faces F, $d_l > L$, or else for some face F, $[d_l, d_r]$ is too wide. In the latter case, one picks a better approximation of the structure on T and begins again. Showing this process terminates is work in progress.

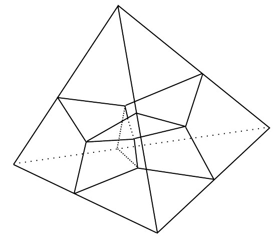

Figure 1: A butterfly in a tetrahedron. We can pick butterflies consistently by picking points on edges and faces which agree in T, then coning off these.

Finally, in the assignment $dH := \partial(H \cup g.D)$, we have to distinguish which faces of the old H and of $g.D$ become faces of the new ∂H, and which get *identified*, and thereby become "interior" faces of the new H. This appears to be an *equality* test. But we can use a trick to circumvent this apparent problem. Suppose a face of $g.D$ is the same as a face F of the old ∂H. Note that F is a face of $h.D$ for some $h \in S$. Then $h^{-1}.(g.D)$ is adjacent to D along $h^{-1}.F$. Let f be the face-pairing of D associated to $h^{-1}.F$. Then $f^{-1}.(h^{-1}.(g.D)) = D$. So $g^{-1}hf$ is the identity, essentially by the definition of fundamental domain.

Now pick a tetrahedron in D of largest inradius R, and let p be its incenter. Then every nontrivial element of Γ translates p at least distance $2R$ away. On the other hand, the identity translates p not at all. So we test for $d(p, (g^{-1}hf).p) < 2R$ and $d(p, (g^{-1}hf).p) > \epsilon$ for some small $\epsilon > 0$. If both tests fail, as before, we better approximate T and begin again. Showing this terminates likewise is work in progress.

References

[1] C. Adams, J. Hass, and P. Scott, "Simple closed geodesics in hyperbolic 3-manifolds," Bull. London Math. Soc. **31** (1999), no. 1, 81–86.

[2] F. Bonahon, *Low-dimensional geometry: from Euclidean surfaces to hyperbolic knots*, American Mathematical Society, 2009.

[3] C. D. Hodgson and J. R. Weeks, "Symmetries, isometries, and length spectra of closed hyperbolic 3-manifolds," Exp. Math. **3** (1994), no. 4, pp. 261–274.

[4] N. Hoffman, K. Ichihara, M. Kashiwagi, H. Masai, S. Oishi, and A. Takayasu, "Verified computations for hyperbolic 3-manifolds," Exp. Math. **25** (2016), no. 1, pp. 66–78.

[5] J. R. Weeks, "Computation of hyperbolic structures in knot theory," *Handbook of knot theory*, Elsevier, 2005, pp. 461–480.

An Optimal Lower Bound for the Hilbert-type Planar Universal Traveling Salesman Problem

Patrick Eades and Julián Mestre

1 Introduction

The *traveling salesman problem* (TSP) is one of the most studied problems in theoretical computer science. For a given set of locations and pairwise distances it asks to find the shortest path such that all locations are visited. In the case when the locations are points in $\mathbb{R}^2$ with their Euclidean distances this problem is known as the *Planar TSP*.

The *universal traveling salesman* (UTSP) is a heuristic approach to the TSP in which all possible locations are ordered ahead of time (generally according to some space-filling curve). Then when presented with a subset of the possible locations which need service, the UTSP visits them in the induced order. The goal of the UTSP is to find an ordering such that the *competitive ratio*, that is, the ratio of induced tour to optimal tour, is minimized over all subsets of the possible locations.

It is known that in the planar case there exists an ordering of the unit square such that the competitive ratio is bounded by $O(\log n)$. We show this ratio is optimal for a large category of orderings of the unit square (Hilbert-type orderings) by proving the existence of a subset with competitive ratio $\Omega(\log n)$.

Work is ongoing to determine whether all non-Hilbert type orderings admit a subset with competitive ratio also at least $\Omega(\log n)$.

2 Problem Statement

Definition 1 *Suppose that $\preceq$ is a total ordering of $[0,1]^2$ and $S \subset [0,1]^2$ is a set of n points, indexed such that $s_1 \preceq \cdots \preceq s_n$. Let $\mathcal{S}_n$ be the symmetric group of order $n!$ and let $s_{n+1} := s_1$. We define:*

$$\textsc{Tsp}(S) = \min_{\sigma \in \mathcal{S}_n} \sum_{i=1}^{n} d(\sigma(s_i), \sigma(s_{i+1})),$$

$$\textsc{Utsp}_{\preceq}(S) = \sum_{i=1}^{n} d(s_i, s_{i+1}),$$

$$\rho_{\preceq}(S) = \frac{\textsc{Utsp}_{\preceq}(S)}{\textsc{Tsp}(S)}.$$

We call $\rho_{\preceq}(S)$ the competitive ratio *of $\preceq$ with respect to S.*

Let us introduce a category of orderings of the unit square, called *Hilbert-type* by analogy to the Hilbert space-filling curve, which induces the most well-known ordering of this type. These orderings are worth examining because they have a high degree of locality, and so form the basis for most UTSP heuristics.

Definition 2 *A* Hilbert-type *ordering of the unit square divides $[0,1]^2$ into k^2 equal squares, and assigns a total ordering to the squares. This ordering induces a partial ordering on $[0,1]^2$, defined for pairs of points that lie in different squares. Points lying in the same square are ordered by recursively dividing that square into k^2 equal squares and applying the same method.*

A Hilbert-type ordering is uniquely determined by k and the ordering used.

The main result of our work is a lower bound on the competitive ratio for Hilbert-type orderings. In particular, for any Hilbert-type ordering we prove the existence of a set with constant optimal cost but logarithmic cost under the UTSP heuristic.

Theorem 1 *For any Hilbert-type ordering $\preceq$ of the unit square and any sufficiently large n there exists a set $S \subset [0,1]^2$ of size n such that $\rho_{\preceq}(S) = \Omega(\log n)$.*

3 Prior Work

The UTSP heuristic was formally introduced by Platzman and Bartholdi in 1989 in response to a need for faster performing TSP algorithms [5]. They used an ordering induced by the Sierpinski space-filling curve to achieve an approximation ratio of $O(\log n)$ which was shown to be tight (for the Sierpinski ordering) by Bertsimas and Grigni [1].

The first general lower bound in the planar case was provided by Hajiaghayi et al. [4] who used a grid-graph to show a lower bound of $\Omega\left(\sqrt[6]{\frac{\log n}{\log\log n}}\right)$. They further conjecture their method could be improved to provide a tight lower bound of $\Omega(\log n)$

Gorodezky et al. [3] provide a lower bound of $\Omega(\log n)$ in general metric spaces. However, their worst case metric space is a Ramanujan graph which is not the setting of most applications, so interest remains in the problem when restricted to the Euclidean plane.

A recent paper by Christodoulou and Sgouritsa in SODA 2017 [2] provides an ordering of the $m \times m$ grid graph with competitive ratio $O(\frac{\log n}{\log \log n})$, disproving a conjecture of Bertsimas and Grigni [1]. However, their construction only works on discrete grids and cannot be generalized to the unit square.

The *generalized Lebesgue orderings* used by Christodoulou and Sgouritsa, when naturally extended from the grid graph to the unit square are of Hilbert type. Thus, by our analysis they exhibit a competitive ratio of $\Omega(\log n)$ on the unit square and are not an improvement over Platzman and Bartholdi's Sierpinski approach.

4 Outline of Method

Intuitively our method for a $k \times k$ Hilbert-type ordering $\preceq$ is to draw a random line $\mathcal{L}$ through the unit square and let S_ℓ be k^ℓ points spaced uniformly along $\mathcal{L}$. Our use of a random line is inspired by [4] as a way to generalize the style of argument used by [1].

Clearly the optimal tour visits each point in the order they appear down the line, so $\text{TSP}(S_\ell) = \Theta(1)$. Then it remains to show that $\text{UTSP}_{\preceq}(S_\ell) = \Theta(\ell)$, which implies that $\rho_{\preceq}(S_\ell) = \Theta(\log |S_\ell|)$.

Lemma 1 *For a uniformly random line $\mathcal{L}$ through $[0,1]^2$, partitioned into k^2 equal squares, the probability of $\mathcal{L}$ intersecting a given square is $\frac{1}{k}$.*

Lemma 2 *For a uniformly random line $\mathcal{L}$ through a square Q of side length $\frac{1}{k}$, the expected length of the line segment $\mathcal{L} \cap Q$ is $\frac{1}{k}$.*

Let $a_\ell = \text{UTSP}_{\preceq}(S_\ell)$. We can decompose a_ℓ into a sum of smaller problems as follows: let $Q_1, \ldots, Q_{k^2}$ be the squares in the $k \times k$ division of $[0,1]^2$. Then a_ℓ is equal to the cost within each Q_i intersecting $\mathcal{L}$ plus the cost of jumping between them in order.

Applying Lemmas 1 and 2 the expected cost within each Q_i is $\frac{1}{k^2} a_{\ell-1}$, since an expected $k^{\ell-1}$ points lie in each square.

Definition 3 *Let $\mathcal{L}$ be a line through $[0,1]^2$ partitioned into k^2 equal squares with a total ordering $\preceq$. Let Q_1, Q_2 and Q_3 be three squares which intersect $\mathcal{L}$, labeled so that $Q_2 \cap \mathcal{L}$ lies between $Q_1 \cap \mathcal{L}$ and $Q_3 \cap \mathcal{L}$ along $\mathcal{L}$. If $Q_1, Q_3 \preceq Q_2$ or $Q_2 \preceq Q_1, Q_3$ we call $Q_2 \cap \mathcal{L}$ a* non order-respecting segment *of $\mathcal{L}$.*

Lemma 3 *For a random line $\mathcal{L}$ through $[0,1]^2$ partitioned into k^2 equal, ordered squares, the expectation of the sum of the non order-respecting segments of the line is constant (see Figure 1).*

Thus, we have that $\mathbb{E}(a_\ell) = \mathbb{E}(a_{\ell-1}) + C$ and so $\mathbb{E}(a_\ell) = \Theta(\ell)$ as required. By taking a line with cost at least as large as the expectation we have the S we needed to prove Theorem 1.

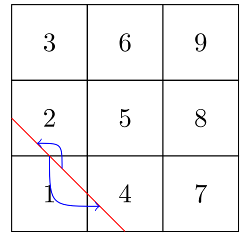

Figure 1: An non-order respecting line: the UTSP tour must first visit square 1, then 2 and then 4, resulting in a constant-size jump over square 1

5 Ongoing Work

Intuitively any non-Hilbert-type ordering will have very bad locality properties. We aim to use this intuition to construct a point set S with logarithmic competitive ratio.

References

[1] D. Bertsimas and M. Grigni. Worst-case examples for the spacefilling curve heuristic for the euclidean traveling salesman problem. *Operations Research Letters*, 8(5):241 – 244, 1989.

[2] G. Christodoulou and A. Sgouritsa. An improved upper bound for the universal tsp on the grid. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1006–1021. SIAM, 2017.

[3] I. Gorodezky, R. D. Kleinberg, D. B. Shmoys, and G. Spencer. Improved lower bounds for the universal and a priori tsp. In *Proceedings of the 13th International Conference on Approximation, and 14 the International Conference on Randomization, and Combinatorial Optimization: Algorithms and Techniques*, APPROX/RANDOM'10, pages 178–191, Berlin, Heidelberg, 2010. Springer-Verlag.

[4] M. T. Hajiaghayi, R. Kleinberg, and T. Leighton. Improved lower and upper bounds for universal tsp in planar metrics. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, pages 649–658, Philadelphia, PA, USA, 2006. Society for Industrial and Applied Mathematics.

[5] L. K. Platzman and J. J. Bartholdi III. Spacefilling curves and the planar travelling salesman problem. *Journal of the ACM (JACM)*, 36(4):719–737, 1989.

Voronoi Diagrams with Rotational Distance Costs *

Mark de Berg[1], Joachim Gudmundsson[2], Herman Haverkort[1], and Michael Horton[2,3]

[1]Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands
[2]School of Information Technologies, The University of Sydney, Australia
[3]Data61, CSIRO, Australia

1 Introduction

We analyse the problem of constructing Voronoi diagrams where the distance metric contains a term that represents a rotation cost.

The motivation for investigating the problem is the work by Taki et al. [1] who proposed the concept of the *dominant region* for the analysis of football (soccer) matches. Informally, the idea is to subdivide the football pitch into regions such that a particular player can reach the points within the region before all other players. The dominant region diagram is thus a Voronoi diagram where the distance function considers properties such as the *a priori* direction the player is facing, the player's velocity, and physiological factors such as the acceleration and turning rate of the player. There have been several efforts to construct algorithms for this problem [2, 4], but the algorithms have been approximations or heuristics. To date, no formal analysis of this problem has been carried out, nor has an exact algorithm been presented.

We aim to make progress in this direction by analysing Voronoi diagrams where the distance metric contains a term for cost of rotation. Under the model, each site is a tuple containing a location and a direction. The distance function from the site s to a point p in the plane contains two terms: a linear term for the cost of rotation from the site's initial direction to the direction towards the point p; and a term for the Euclidean distance from s to p. Each site can be visualised as a robot that can rotate about its axis when stationary and move only in a straight line in the direction it is facing, i.e. the distance to a point is the sum of the cost of rotation to face the point and the cost of moving to the point.

This abstract presents the research undertaken and results to date. The problem is formally defined in Section 2. Preliminary analysis of the problem is presented in Section 3 showing that the Voronoi regions can be disconnected, even when the number of sites $n = 2$, and the complexity of the Voronoi diagram can be at least quadratic in n. Section 4 lists the open problems.

*We thank Ali Mehrabi for useful discussions during the initial stages of this research.

2 Preliminaries

The input is a set S of n sites, and each site $s \in S$ is a triple (s_x, s_y, s_α), where $(s_x, s_y) \in \mathbb{R}^2$ specifies the *location* of s, and $s_\alpha \in [-\pi, \pi)$ is the *direction*, that is a counterclockwise angle relative to the positive x-axis. Where appropriate and unambiguous, s is treated as a point in $\mathbb{R}^2$, for example $|\overline{sp}|$ denotes the Euclidean distance between site s and a point $p \in \mathbb{R}^2$.

Let $C \geq 0$ be a fixed constant. The distance function $d_C \colon S \times \mathbb{R}^2 \to \mathbb{R}_{\geq 0}$ accepts as arguments a site s and a point p, and determines the distance from s to p, i.e. $d_C(s,p) = \theta + C|\overline{sp}|$, where θ is the absolute difference between s_α and the angle of $\overline{sp}$, relative to the positive x-axis, i.e. $\theta = \arccos(p' \cdot (\cos s_\alpha, \sin s_\alpha)/||p'||)$, where $p' = (p_x - s_x, p_y - s_y)$.

Now that we have the distance function d_C, we can define the Voronoi diagrams and related concepts in the usual way: the *bisector* of two sites $s, u \in S$ is defined as $b(s,u) = \{z \in \mathbb{R}^2 \mid d(s,z) = d(u,z)\}$; the *Voronoi region* of a site $s \in S$ is defined as $VR(s,S) = \{z \in \mathbb{R}^2 \mid d(s,z) < d(s',z) \,\forall\, s' \in S - s\}$; and the *Voronoi diagram* of S is defined as

$$V(S) = \mathbb{R}^2 \setminus \bigcup_{s \in S} VR(s,S).$$

Voronoi regions may be disconnected, and thus each Voronoi region $VR(s,S)$ can have one or more *faces*.

3 Analysis

The case $n = 2$. Given two sites s and u, the bisector $b(s,u)$ can be a set of closed or infinite curves, and may have multiple disconnected components.

We will need the following lemma.

Lemma 1 *$\mathcal{B}$ is a set of disjoint 2D regions in $\mathbb{R}^2$ such that any line ℓ crosses the boundaries of $\mathcal{B}$ at most three times. Then $\mathcal{B}$ has at most one bounded face and at most three unbounded faces.*

We define *disjoint* in this context to be that for any $B', B'' \in \mathcal{B}$, the closures of B' and B'' only intersect at a

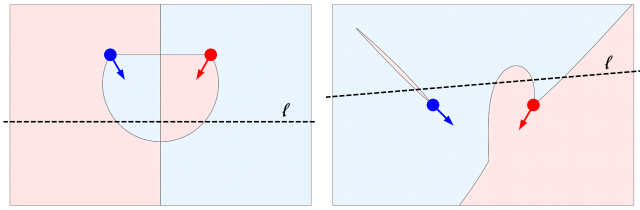

(a) $C = 0$: *angle-only* case (b) $C > 0$: *general* case

Figure 1: Example showing line dashed line ℓ intersecting the bisector $b(s, u)$ at (a) three points for $C = 0$ and (b) five points for $C > 0$.

finite number of points. The condition that the number of intersection points between ℓ and the boundaries of $\mathcal{B}$ is at most three excludes instances where a boundary $\delta(B)$, $B \in \mathcal{B}$ and ℓ intersect in a line-segment $\overline{pq}$. However, such cases can be dealt with using a perturbation argument.

Consider the number of intersections between an arbitrary line ℓ and $b(s, u)$ in Voronoi diagrams of two sites s and u. We analyse two distance functions: where $C = 0$, i.e. the distance from site s to point p is the rotation cost only—the *angle-only* case; and where $C > 0$—the *general* case. In the angle-only case, we can prove that the number of intersections is no greater than three, see Fig. 1(a). Using this fact and Lemma 1 leads to the following result.

Corollary 2 *When $C = 0$, that is, in the angle-only case, the Voronoi diagram of two sites in the d_C-distance has at most four faces.*

However, in the case of the general distance function, there are examples where line ℓ intersects the bisector line in five points, and therefore a bound cannot be constructed using Lemma 1, see Fig. 1(b). A bound for this case remains an open question.

The case $n > 2$. In the general case with more than two sites, we show a lower bound of $\Omega(n^2)$ on the complexity of the diagram, see Fig. 2. Furthermore, it is possible to construct cases where the number of faces of a single Voronoi region is $\Omega(n^2)$, giving the following theorem.

Theorem 3 *The number of faces of the Voronoi diagram of n sites based on the d_C-distance is $\Omega(n^2)$ in the worst case.*

Thus far, we have been unable to find an example where the complexity of the Voronoi diagram is cubic. An upper bound on the complexity of $V(S)$ when $n > 2$ is an open question, although we conjecture that it is $O(n^2)$.

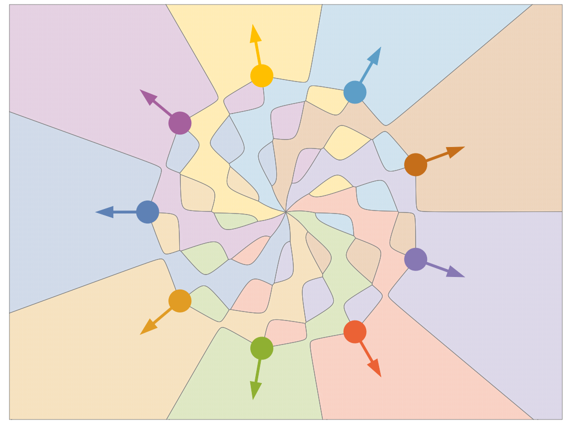

Figure 2: Example of Voronoi diagram of quadratic complexity with respect to the number of sites.

4 Open Problems

The main open problems are to determine the complexity of the Voronoi diagram when $n = 2$ using the general distance function; and the complexity when $n > 2$ using both distance functions.

It also remains to determine an algorithm to efficiently compute the Voronoi diagram in the d_C-metric. (Note: the framework for *abstract Voronoi diagrams* [3] cannot be directly used, since this requires the bisector $b(s, u)$ to be connected.)

References

[1] T. Taki, J. Hasegawa, and T. Fukumura. "Development of motion analysis system for quantitative evaluation of teamwork in soccer games". In: *Proc. Int. Conf. on Image Processing*. IEEE, Sept. 1996, pp. 815–818. DOI: 10.1109/ICIP.1996.560865.

[2] T. Taki and J. Hasegawa. "Visualization of Dominant Region in Team Games and Its Application to Teamwork Analysis". In: *Proc. IEEE Computer Graphics Int. Conf.* IEEE, June 2000, pp. 227–235. DOI: 10.1109/CGI.2000.852338.

[3] R. Klein, E. Langetepe, and Z. Nilforoushan. "Abstract Voronoi diagrams revisited". In: *Computational Geometry* 42.9 (2009), pp. 885–902. DOI: 10.1016/j.comgeo.2009.03.002.

[4] R. Nakanishi et al. "An Approximate Computation of the Dominant Region Diagram for the Real-Time Analysis of Group Behaviors". In: *Proc. 13th Annual RoboCup Int. Symp.* Springer, June 2009, pp. 228–239. DOI: 10.1007/978-3-642-11876-0_20.

Unwinding Annular Curves and Electrically Reducing Planar Networks

Hsien-Chih Chang* Jeff Erickson*

1 Introduction

Any continuous deformation of a closed curve on any surface can be decomposed into a finite sequence of *homotopy moves*, consisting of the following three operations and their inverses. A 1→0 move removes an empty loop; a 2→0 move removes an empty bigon, and a 3→3 move flips an empty triangle. A classical argument of Steinitz [6] implies that any planar curve with n vertices can be simplified using $O(n^2)$ homotopy moves; using a later result of Hass and Scott [4] one can extend this upper bound to contractible curves on arbitrary surfaces. We prove that Hass and Scott's quadratic upper bound is tight.

Theorem 1 *Simplifying a contractible curve in the annulus (or any surface that has the annulus as its covering space) requires $\Omega(n^2)$ homotopy moves in the worst case.*

This improves our previous $\Omega(n^{3/2})$ lower bound, which follows from an analysis of curves in the plane, and generalizes our previous $\Omega(n^2)$ lower bound for simplifying *non-contractible* curves on higher-genus surfaces [1].

Our second result concerns the reduction of plane graphs via *electrical transformations*: leaf reductions, loop reductions, series-parallel reductions, and $\Delta\leftrightarrow Y$ transformations. We distinguish between two types of electrical transformations in plane graphs: A loop reduction, parallel reduction, or $\Delta\to Y$ transformation is *facial* if the edges deleted by the operation bound a face in G, and *non-facial* otherwise. Dual pairs of facial electrical transformations correspond to local transformations in the medial graph of G, which we call *medial electrical moves*.

Theorem 2 *Reducing an n-vertex plane graph with two terminals as much as possible requires $\Omega(n^2)$ facial electrical transformations in the worst case.*

The proof uses our quadratic homotopy lower bound for annular curves. This result matches the upper bound implied by Hass and Scott [4], and it strengthens and generalizes our earlier $\Omega(n^{3/2})$ lower bound for reducing plane graphs to a single vertex [1].

2 Unwinding Annular Curves

To simplify our analysis of annular curves, it is convenient to work in the punctured plane $\mathbb{R}^2 \setminus \{o\}$, where o is an arbitrary point called the *origin*. In any homotopy in the punctured plane, homotopy moves that contract either the face containing the origin or the outer face of the curve are forbidden. It is precisely these forbidden homotopy moves that make the quadratic lower bound possible; if we only forbid homotopy moves on the outer face, then any curve can be simplified using at most $O(n^{3/2})$ moves [1].

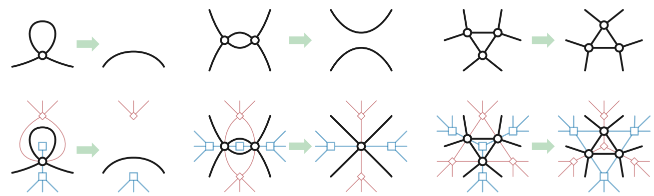

Figure 1. Top row: Homotopy moves 1→0, 2→0, and 3→3. Bottom row: Medial electrical moves 1→0, 2→1, and 3→3.

Let γ be an arbitrary oriented closed curve in the punctured plane, and let p be any point outside the image of γ. The *winding number* $wind(\gamma, p)$ is the number of times γ crosses a generic ray ρ based at p from right to left, minus the number of times γ crosses ρ from left to right. For any vertex x of γ, the winding number $wind(\gamma, x)$ is defined as the average of the winding numbers around the four faces incident to x.

Smoothing γ at a vertex x replaces a small neighborhood of x with two disjoint simple paths. There are two possible smoothings at each vertex, one of which splits γ into two subcurves; let γ_x^+ and γ_x^- denote the subcurves locally to the left of x and locally to the right of x, respectively. We define the *type* of any vertex x as $type(\gamma, x) := wind(\gamma_x^+, o)$. A vertex x is *irrelevant* if either $type(\gamma, x) = 0$ or $type(\gamma, x) = wind(\gamma, o)$ and *relevant* otherwise. Two vertices x and y have *complementary* types if $type(\gamma, x) + type(\gamma, y) = wind(\gamma, o)$.

Case analysis shows that homotopy moves modify the types and winding numbers of vertices as follows: (a) Each 1↔0 move creates or destroys one irrelevant vertex. (b) Each 2↔0 move creates or destroys two vertices with complementary types and identical winding numbers. (c) Each 3↔3 move changes the winding numbers of the three vertices, each by exactly 1. (d) Otherwise, homotopy moves do not change the type or winding number of any vertex.

To prove Theorem 1, consider any contractible γ and any homotopy that contracts γ to a point, and let x be a relevant vertex at any stage of the homotopy. We can follow x through the homotopy to a 2→0 move that deletes x and a complementary vertex; symmetrically, we can follow

*Department of Computer Science, University of Illinois at Urbana-Champaign. Supported in part by NSF grant CCF-1408763.

x backward through the homotopy either to a vertex of γ or to a 0→2 move that creates x and a complementary vertex. Proceeding inductively, we obtain a matching between complementary relevant vertices of γ. Each matched pair is connected by a path through the homotopy that alternately moves forward and backward through the homotopy, switching direction and type at 2↔0 moves.

During the homotopy, the winding number of a vertex changes precisely when it participates in a 3→3 move. Thus, the path between two matched vertices x and y must pass through at least $\left|wind(\gamma,x)-wind(\gamma,y)\right|$ 3→3 moves, and the entire homotopy must contain at least $\sum_{x\sim y}\left|wind(\gamma,x)-wind(\gamma,y)\right|/3$ 3→3 moves, where the sum is over all matched pairs of vertices of γ.

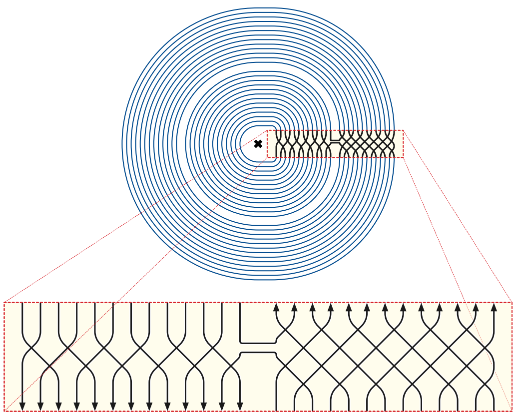

Figure 2. Our bad example curve X_{13} in the punctured plane.

For any relatively prime integers p and q, the *flat torus knot* $T(p,q)$ winds $|p|$ times around the origin and oscillates $|q|$ times between two concentric circles. For any odd integer p, let X_p denote the connected sum of $T(-p,1)$ and $T(p,2)$, where the former curve is scaled to lie inside the innermost face of the latter. This curve is contractible and has $3(p-1)$ vertices. Analysis of the types and winding numbers of the vertices of X_p implies that any homotopy that contracts X_p contains at least $p(p-1)/6$ 3→3 moves.

3 Planar Graphs with Two Terminals

Most applications of electrical transformations designate two vertices as *terminals*. In this context, leaf reductions, series reductions, and $Y\to\Delta$ transformations at terminals are not proper electrical transformations. Every 2-terminal plane graph can be reduced by facial electrical transformations to a unique graph, which we call a *bullseye*, but not necessarily to a single edge. (The medial graph of a bullseye is the curve $T(p,1)$, for some even integer p.)

A *multicurve* is an immersion of one or more closed curves. For a plane graph G with two terminals, the medial graph of G is the image of a multicurve embedded in the annulus. Arguments of Truemper [7] and Noble and Welsh [5], described in detail in our earlier paper [1], imply that reducing a plane curve γ using medial electrical moves requires at least as many steps as reducing γ using homotopy moves. To prove Theorem 2, we extend these arguments to annular curves. Specifically, we show that the number of medial electrical moves required to reduce an annular curve is at least the number of homotopy moves required to reduce the same curve; the quadratic lower bound now follows directly from Theorem 1.

Two key ingredients of our proof may be of independent interest. First, we show that a multicurve γ on *any* surface can be further reduced using medial electrical moves if and only if γ can be further reduced using homotopy moves. We prove this fact using a classical result [3] that any multicurve can be simplified as much as possible via homotopy moves without ever increasing the number of vertices.

Let γ be an arbitrary connected multicurve in the punctured plane. Let $X(\gamma)$ denote the minimum number of medial electrical moves required to reduce γ. Let $depth(\gamma)$ denote the minimum number of times a generic ray based at the origin crosses γ (not considering the directions of crossings). Our second key observation, which relies on the first, is that the inequality

$$X(\overline{\gamma}) < X(\gamma) + \left(depth(\gamma) - depth(\overline{\gamma})\right)/2$$

holds for every connected proper smoothing $\overline{\gamma}$ of γ. This generalizes the simpler inequality $X(\overline{\gamma}) < X(\gamma)$ for connected multicurves in the plane [1].

As a final remark, by including one more operation called the *terminal leaf reduction* in addition to facial electrical transformations, any 2-terminal plane graph can be reduced to a single edge; indeed, existing electrical reduction algorithms for plane graphs rely exclusively on these operations [2, 7]. Unfortunately, our lower bound does *not* apply when terminal leaf reductions are allowed.

References

[1] Hsien-Chih Chang and Jeff Erickson. Untangling planar curves. *Proc. 32nd Int. Symp. Comput. Geom.*, 29:1–29:15, 2016. Leibniz International Proceedings in Informatics 51. Full version to appear in *Disc. Comput. Geom.*

[2] Thomas A. Feo and J. Scott Provan. Delta-wye transformations and the efficient reduction of two-terminal planar graphs. *Oper. Res.* 41(3):572–582, 1993.

[3] Maurits de Graaf and Alexander Schrijver. Making curves minimally crossing by Reidemeister moves. *J. Comb. Theory Ser. B* 70(1):134–156, 1997.

[4] Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel J. Math.* 51:90–120, 1985.

[5] Steven D. Noble and Dominic J. A. Welsh. Knot graphs. *J. Graph Theory* 34(1):100–111, 2000.

[6] Ernst Steinitz. Polyeder und Raumeinteilungen. *Enzyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen* III.AB(12):1–139, 1916.

[7] Klaus Truemper. On the delta-wye reduction for planar graphs. *J. Graph Theory* 13(2):141–148, 1989.

Conditional Nontrivial Lower Bounds for 3SUM and Friends

Jean Cardinal* John Iacono† Stefan Langerman‡ Aurélien Ooms§

Abstract

3SUM-hardness is a key tool when it comes to identifying the complexity of geometric problems. Moreover, the more general k-SUM problem can be cast as a point location problem. We show that, if $(k-1)$-SUM can be solved in $\tilde{O}(n^{\frac{k}{3}-\delta})$ time, then k-SUM can be solved in $\tilde{O}(n^{\frac{k}{2}-\frac{3\delta k/2-3\delta}{4k/3-2-\delta}})$ time in the linear decision tree and algebraic computation tree models. This result holds for all integers $k \geqslant 3$. In particular, a corollary of our results is that a $\omega(n^{2-\varepsilon})$ lower bound on the depth of 8-linear decision trees for sorting $X+Y$ would imply a nontrivial lower bound of $\omega(n^{4/3-\varepsilon})$ on the depth of 4-linear decision trees for 3SUM.

1 Introduction

Most geometric problems suffer from missing tight complexity lower bounds. A remedy to this annoyance is *conditional complexity*: conjecture some key problem is hard, then show that this conjecture implies hardness for other problems. The 3SUM problem is one such key problem. For instance, Gajentaan and Overmars [5] show that, the problem of computing the separator of an input set of line segments (see Figure 1), the problem of computing the minimum area triangle spawned by input points, and many other geometric problems, are all as hard as 3SUM. The 3SUM problem is defined as the k-SUM problem when $k = 3$.

Problem (k-SUM) Given a set of n real numbers $S = \{ s_1 < s_2 < \cdots < s_n \}$, decide whether there exists $i = (i_1, i_2, \ldots, i_k) \in [n]^k$ such that $\sum_{j=1}^{k} s_{i_j} = 0$.

Until spring 2014, the conjecture had been that 3SUM requires quadratic time. Had this conjecture been true, it would have implied a quadratic lower bound on the problems mentioned above. However, Grønlund and Pettie [6] refuted this conjecture. In their breakthrough paper, they gave the first subquadratic time algorithms for 3SUM in both uniform and nonuniform models of computation: a real-RAM algorithm in time $O(n^2(\log\log n/\log n)^{2/3})$ and a 4-linear decision tree of depth $O(n^{\frac{3}{2}}\sqrt{\log n})$. The respective complexities of those algorithms have since been slightly improved. The outdated 3SUM conjecture has been updated to take into account this important progress:

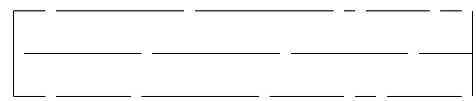

Figure 1: Separating line segments is 3SUM-hard.

Conjecture 1 *In the real-RAM model, 3SUM cannot be solved in $O(n^{2-\delta})$, for any $\delta > 0$.*

Another key problem is sorting the n^2 pairwise sums of elements of two sets of n real numbers. This problem is commonly refered to as sorting $X+Y$. The best known algorithms for sorting $X+Y$ all use a quadratic number of sum comparisons (4-linear queries) [4, 8, 9].

Our paper focuses on the (nonuniform) decision tree complexity of k-SUM. By leveraging Grønlund and Pettie's decision tree, we show how the complexity of the k-SUM problem relates to the complexity of the $(k-1)$-SUM problem. Our result implies that, unless sorting $X+Y$ can be solved in $O(n^{2-\varepsilon})$ 8-linear queries, 3SUM cannot be solved in $O(n^{4/3-\varepsilon})$ 4-linear queries.

Recent exposition on the n-linear decision tree complexity of k-SUM, casting k-SUM as a point location problem, suggested that the linear decision tree complexity of 3SUM should lie close to linear [2, 3]. This conjecture is indeed true; Kane, Lovett, and Moran [7] just showed that k-SUM can be solved using $O(n\log^2 n)$ $2k$-linear queries. This brings the state of the art close to the information-theoretic lower bound of $\Omega(n\log n)$.

This recent breakthrough limits the applications of our conditional results. Since sorting $X+Y$ can now be done in $O(n\log^2 n)$ 8-linear queries, our technique does not allow to unconditionally lower bound the depth of 4- or 5-linear decision trees for 3SUM ($\Omega(n^2)$ and $O(n\log^2 n)$ being respectively the best known lower and upper bound for 3- and 6-linear decision trees).

*Supported by the "Action de Recherche Concertée" (ARC) COPHYMA, convention number 4.110.H.000023. Université libre de Bruxelles (ULB), Brussels, Belgium, jcardin@ulb.ac.be.

†Research partially completed while on on sabbatical at the Algorithms Research Group of the Département d'Informatique at the Université Libre de Bruxelles with support from a Fulbright Research Fellowship, the Fonds de la Recherche Scientifique — FNRS, and NSF grants CNS-1229185, CCF-1319648, and CCF-1533564. New York University, New York, United States of America, cgyrf17@johniacono.com.

‡Directeur de recherches du F.R.S.-FNRS. Université libre de Bruxelles (ULB), Brussels, Belgium, slanger@ulb.ac.be.

§Supported by the Fund for Research Training in Industry and Agriculture (FRIA). Université libre de Bruxelles (ULB), Brussels, Belgium, aureooms@ulb.ac.be.

This is an abstract of a presentation given at CG:YRF 2017. 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 in a conference with formal proceedings and/or in a journal.

2 Results

Idea We proceed as in Grønlund and Pettie [6]. They reduce solving a 3SUM instance to a preprocessing phase on a large 2SUM instance (sorting) and a query phase that quickly searches through the sorted lists obtained via the preprocessing phase. Similarly, we reduce a k-SUM instance to a preprocessing phase followed by a query phase. The preprocessing phase involves solving a large $(k-1)$-SUM instance. The query phase quickly searches through the information gathered during the preprocessing phase.

Partitioning Let $S = \{ s_1 < s_2 < \cdots < s_n \}$. For some g to be chosen later, partition the interval $[s_1, s_n]$ into n/g blocks $S_1^*, S_2^* \ldots, S_{n/g}^*$ such that each block contains g numbers in S. For the sake of simplicity, and without loss of generality, we assume that g divides n. To each of the $(n/g)^{k-1}$ tuples of blocks $(S_{i_1}^*, S_{i_2}^*, \ldots, S_{i_{k-1}}^*)$ corresponds a cell $\times_{j=1}^{k-1} S_{i_j}^*$ in the $(k-1)$-dimensional grid generated by our partition of S. The two following lemmas are consequences of our partitioning scheme:

Lemma 1 *In $\mathbb{R}^{k-1}$ with variables $x_1, x_2, \ldots, x_{k-1}$, for a fixed $w \in \mathbb{R}$, the hyperplane of equation $-w = \sum_{i=1}^{k-1} x_i$ intersects $O((n/g)^{k-2})$ cells. Moreover, those cells can be found in $O((n/g)^{k-2})$ time.*

Proof. The hyperplane $-w = \sum_{i=1}^{k-1} x_i$ draws a $(k-1)$-dimensional staircase in the grid. We can output all the cells on this staircase in time proportional to their number (and k, which is constant). ▫

Lemma 2 *If the set S can be preprocessed in $T_g(n)$ time so that, for any given cell $\times_{j=1}^{k-1} S_{i_j}^*$ and any given $w \in S$, testing whether $-w \in +_{j=1}^{k-1} S_{i_j} = \{ \sum_{i=1}^{k-1} z_i : z \in \times_{j=1}^{k-1} S_{i_j} \}$ can be done in $O(\log g)$ time, then k-SUM can be solved in $T_g(n) + O(\frac{n^{k-1}}{g^{k-2}} \log g)$ time.*

Proof. Preprocess the input in $T_g(n)$ time. For each input number w, for each cell $\times_{j=1}^{k-1} S_{i_j}^*$ intersected by the hyperplane $-w = \sum_{i=1}^{k-1} x_i$, test whether $-w \in +_{j=1}^{k-1} S_{i_j}$ in $O(\log g)$ time. Use Lemma 1 to bound the number of intersected cells and the time it requires to find them. ▫

Lemma 3 *Assuming $(k-1)$-SUM on N numbers can be solved in $f(N) = O(poly(N))$ time, the set S can be preprocessed as in Lemma 2 using $T_g(n) = O(f(ng))$ time, in the linear decision tree and algebraic computation tree models.*

Proof. We sort all sets $+_{j=1}^{k-1} S_{i_j}$ by solving a $(k-1)$-SUM instance of size $O(ng)$. Sorting the set $+_{j=1}^{k-1} S_{i_j}$ amounts to answering the linear query $\sum_{i=1}^{k-1} z_i \leqslant \sum_{i=1}^{k-1} z'_i$ for all $z, z' \in \times_{j=1}^{k-1} S_{i_j}$. Those queries are equivalent to $\sum_{i=1}^{k-1} (z_i - z'_i) \leqslant 0$ which, by [1], can be answered by solving a $(k-1)$-SUM instance on the numbers $(z_i - z'_i)$. Because there are $O(n/g)$ blocks S_i^*, and because there are g^2 tuples (z_i, z'_i) in each block, there are $O(ng)$ such numbers. ▫

Theorem 4 *If $(k-1)$-SUM can be solved in $O(n^{\frac{k}{3}-\delta})$ time, with $k \geqslant 3$ and $0 \leqslant \delta \leqslant \frac{k}{3} - 1$, then k-SUM can be solved in $\tilde{O}(n^{\frac{k}{2} - \frac{3\delta k/2 - 3\delta}{4k/3 - 2 - \delta}})$ time in the linear decision tree and algebraic computation tree models.*

Proof. Assume we can solve $(k-1)$-SUM on N numbers in $\tilde{O}(N^{k/3-\delta})$ time, then the preprocessing phase can be done in $\tilde{O}((ng)^{k/3-\delta})$ time by Lemma 3. Balancing the sum in Lemma 2 gives an optimal solution $g = \tilde{\Theta}(n^{\frac{1}{2} + \frac{3\delta/2}{4k/3-2-\delta}})$. ▫

Remark Note that a more careful analysis allows to transfer polylog shavings from $(k-1)$-SUM to k-SUM.

References

[1] Michael Ben-Or. Lower bounds for algebraic computation trees (preliminary report). In *STOC*, pages 80–86. ACM, 1983.

[2] Jean Cardinal, John Iacono, and Aurélien Ooms. Solving k-SUM using few linear queries. In *ESA*, volume 57 of *LIPIcs*, pages 25:1–25:17, 2016.

[3] Esther Ezra and Micha Sharir. A nearly quadratic bound for the decision tree complexity of k-SUM. In *SoCG*, LIPIcs, 2017. To appear.

[4] Michael L. Fredman. How good is the information theory bound in sorting? *Theor. Comput. Sci.*, 1(4):355–361, 1976.

[5] Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995.

[6] Allan Grønlund and Seth Pettie. Threesomes, degenerates, and love triangles. In *FOCS 2014*, pages 621–630. IEEE, 2014.

[7] Daniel M. Kane, Sharar Lovett, and Shay Moran. Near-optimal linear decision trees for k-SUM and related problems. 2017. arXiv:1705.01720 [cs.CG].

[8] Jean-Luc Lambert. Sorting the sums $(x_i + y_j)$ in $O(n^2)$ comparisons. In *STACS 90*, pages 195–206. Springer, 1990.

[9] William L. Steiger and Ileana Streinu. A pseudo-algorithmic separation of lines from pseudo-lines. *Information Processing Letters*, 53(5):295–299, 1995.