29th ACM Symposium on Computational Geometry

June 17-20, 2013 - Rio de Janeiro, Brazil

Young Researchers Forum

Collection of Abstracts

YRF Program Committee

Jean-Daniel Boissonnat (INRIA Sophia Antipolis)
Erin Chambers (Saint Louis University)
Sándor Fekete (TU Braunschweig)
Danny Halperin (Tel Aviv University)
Michael Hoffmann (ETH Zürich)
Maarten Löffler (Utrecht University)
David Mount (University of Maryland)
Bettina Speckmann (TU Eindhoven, chair)

Tuesday, June 18th 2013

Session 1 (3:00pm-4:30pm)

Manifold Reconstruction Using the Relaxed Delaunay Triangulation
Renata Rego, Ramsay Dyer and Jean-Daniel Boissonnat

Longest-Edge Algorithms for Quality Delaunay Refinement
Carlos Bedregal and Maria-Cecilia Rivara

NTU-Neutralized Complexes for the Optimal Homologous Chain Problem
Bala Krishnamoorthy and Gavin Smith

On the Subdivision Containment Problem for Random 2-Complexes
Anna Gundert and Uli Wagner

Dimension Detection by Local Homology
Tamal Dey, Fengtao Fan and Yusu Wang

Computationally proving triangulated 4-manifolds to be diffeomorphic
Benjamin A. Burton and Jonathan Spreer

Session 2 (5:00pm-6:30pm)

Lower bounds for LSH and nearest neighbor on the Bregman divergences
Amirali Abdullah and Suresh Venkatasubramanian

Fault Tolerant Clustering Revisited
Nirman Kumar and Benjamin Raichel

Discrete Morse Gradient Fields and Stable Matchings
Joao Paixao and Thomas Lewiner

On the Most Likely Convex Hull of Uncertain Points in the Plane
Hakan Yildiz

An arc-based integer programming formulation for proportional symbol maps
Rafael G. Cano, Cid C. de Souza, Pedro J. de Rezende and Tallys Yunes

Algorithms for tracking movement and guarding regions using radio tomographic imaging
Samira Daruki, Peter Hillyard, Neal Patwari and Suresh Venkatasubramanian

Manifold Reconstruction Using the Relaxed Delaunay Triangulation

Renata L.M.E. Rêgo* Ramsay Dyer† Jean-Daniel Boissonnat†

Abstract

This work presents the first implementation of a manifold reconstruction which is based on the relaxed Delaunay triangulation. Its theoretical foundation stems from the idea of protection developed in [2]. We empirically study its validity and practical aspects regarding the sampling conditions including fallback solutions in case of violation. Experiments indicate that it is possible to reconstruct k-manifolds embedded in high-dimensional space, given that the input point cloud is dense enough.

1 Introduction

The manifold learning process aims at retrieving some information about an unknown manifold $\mathbb{M}$ from a given set of sample points L drawn from $\mathbb{M}$. In this context, subcomplexes of the Delaunay triangulation have been used to build a faithful approximation of $\mathbb{M}$ from L. Particularly, the restricted Delaunay complex $Del(L, \mathbb{M})$ (defined in [5]) can be shown to provide good topological and geometric approximation of planar curves or surfaces in $3D$-space, assuming a sufficiently dense sampling [1].

De Silva [3] introduced a relaxed version of the Delaunay complex, and Boissonnat et al. [2] showed that the relaxed Delaunay complex $Del^\rho(L, W)$ is equivalent to $Del(L, \mathbb{M})$, where W is a set of input points sampled from $\mathbb{M}$, L is a subset of W, and ρ is the relaxation parameter, defined in Section 2. The equivalence is valid under sampling conditions specified for L and W. The motivation for introducing $Del^\rho(L, W)$ is that for manifold reconstruction only a discrete set of samples is available that represents $\mathbb{M}$. Motivated by the results of [2] we developed a manifold reconstruction algorithm that takes as input a finite point set W sampled from a k-manifold $\mathbb{M}$ and outputs an approximation of $\mathbb{M}$.

2 Foundations

First we introduce some notions of [2], which are the foundation of our reconstruction algorithm. Essentially the notion of a Delaunay center, i.e. an equidistant point to all vertices of a simplex such that the corresponding circumball does not contain other points, is too strict to be applicable in manifold reconstruction given only a finite number of sample points. In this context the generalized ρ-Delaunay center, i.e. a point not too far away from a Delaunay center, as illustrated in Figure 1 (middle) and defined below is more useful.

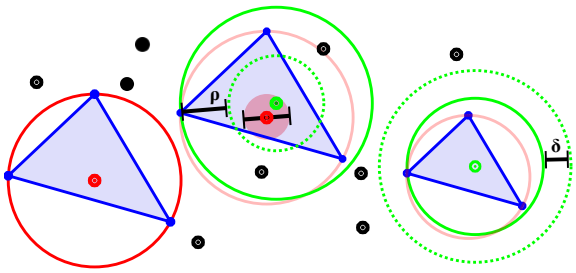

Figure 1: (left) The red Delaunay center is equidistant to the blue vertices of a simplex while all other points (black) are outside the red circumcircle. (middle) For the green ρ-Delaunay center all points are outside the smaller dashed green circle with radius $d_{max}(x, \sigma) - \rho$. Each point within the red disc of diameter ρ around a Delaunay center is a ρ-Delaunay center, however, note that the red circumcircle is not necessarily empty anymore. (right) Protection means that the black points are outside the dashed green circle and thus sufficiently far away for ensuring an empty red circumcircle.

Defnition 1 *Assume $\rho \geq 0$. A point $x \in X \subset \mathbb{R}^d$ is a **ρ-Delaunay center** for a simplex $\sigma \subset L$ if the minimal distance $d_{min}(x, L) = \min ||q - x||$, $q \in L$ to x is larger or equal to its ρ-shortened maximal distance to a simplex point $d_{max}(x, \sigma) = \max ||p - x||$, $p \in \sigma$:*

$$d_{min}(x, L) \geq d_{max}(x, \sigma) - \rho \qquad (1)$$

Note that for vanishing ρ we have a Delaunay center. A relaxed Delaunay complex then generalizes to

Defnition 2 *The **ρ-Delaunay complex** of L restricted to X is the abstract simplicial complex, $Del^\rho(L, X)$, on L defined by:*

$$\sigma \in Del^\rho(L, X) \Longleftrightarrow \sigma \text{ has a } \rho\text{-Delaunay center in } X.$$

Boissonnat et al. [2] proved that a ρ-Delaunay complex is identical to the restricted Delaunay complex, if all ρ-Delaunay centers are sufficiently *protected*. Notice that this powerful insight enables the generation of the true restricted Delaunay complex even if only a finite sampling of a surface is available. Protection means that

*Center of Informatics-UFPE and Federal Institute of Pernambuco (IFPE) rlmer@cin.ufpe.br

†Geometrica, Inria Sophia Antipolis, {ramsay.dyer, jean-daniel.boissonnat}@inria.fr

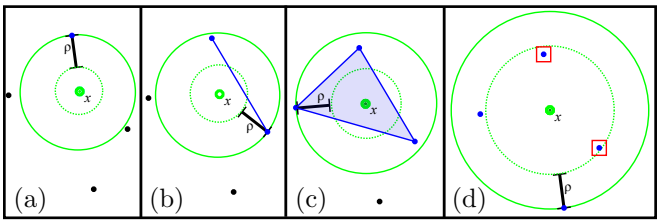

Figure 2: Stepwise simplex creation: (a) Initially a 0-simplex is created which is the closest point (blue) to x; (b,c) A 1-simplex (blue edge) and 2-simplex (blue triangle) are created adding the second and third closest vertex respectively; (d) No 3-simplex is created since adding the fourth closest point would violate Equation 1 due to the two points highlighted in red.

a δ-thickened version of the circumball is still empty as illustrated in Figure 1 (right) and formally defined as

Defnition 3 *A point $x \in X$ is a δ-protected center for $\sigma \subset L$ if:*

$$d_{min}(x, L \setminus \sigma) \geq d_{max}(x, \sigma) + \delta \qquad (2)$$

Our main objective is to turn these insights into a practical algorithm. This includes the determination of landmarks L and estimation of ρ as well as strategies for obtaining protected centers.

3 Proposed Algorithm

The main algorithm is straightforward, for each $x \in W$, the algorithm creates successively higher dimensional simplices, until the condition of Equation 1 is violated (cf. Figure 2). In the following we address some additional aspects that complement our algorithm:

Landmark Selection is done with a self-organizing learning based approach [4] choosing $L \subset W$.

Relaxation Parameter Estimation. A simple argument with covering disks is used to estimate ρ: Let ϵ be the sampling radius of W with respect to $\mathbb{M}$, the disks of radius ϵ centered on the input points $w \in W$ should cover the surface, i.e., if the surface area is A then $|W|\pi\epsilon^2 \geq A$, and likewise $|L|\pi\lambda^2 \geq A$, where λ is the sampling radius of L. Since the landmarks are well-spaced and the input points are clustered according to a uniform random distribution, it is reasonable to suppose that $|W|\pi\epsilon^2 \geq |L|\pi\lambda^2$. We know from [2] that if a simplex $\sigma \subset L$ has a Delaunay center in $\mathbb{M}$, then by triangle inequality it will have a ρ-relaxed Delaunay center in W provided $\rho \geq 2\epsilon$ and thus $\rho \geq 2\epsilon \geq 2\sqrt{\frac{|L|}{|W|}}\lambda$.

Protection. Our experiments indicate that W needs to be significantly denser than L. Such a requirement is explained as follows: The most protection that we could obtain is $\delta = C_1\lambda$, with a very small value for constant C_1. According to [2], we require $\delta \geq C_2\rho$, with a fairly big value of C_2. As a consequence we need the sampling radius ϵ of W to be much smaller than that of the vertices λ: $2\epsilon \leq \rho \leq \frac{C_1}{C_2}\lambda$. Therefore even for very dense point clouds, the relaxed Delaunay centers are sometimes not sufficiently protected. In practice we improve protection with a perturbation strategy which moves the vertex closest to the relaxed Delaunay center x of a high-dimensional simplex by ρ towards x.

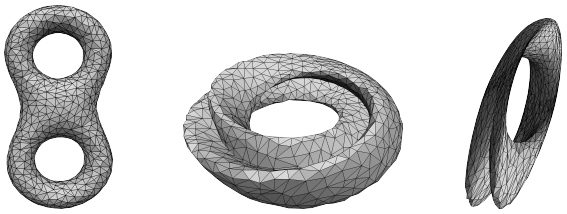

(a) (b) (c)

Figure 3: Reconstructions: (a) double torus; (b,c) Klein bottle projected on dimensions 1,2,3 and 1,4,5.

Results. Figure 3 shows the reconstruction of a double torus, and 3D projections of the reconstruction of a Klein bottle embedded in 5D. The proposed algorithm is able to produce a 2-manifold triangular mesh after 4 and 6 iterations of the vertex perturbation strategy for the Klein bottle and double torus respectively. In both cases $|L| = 1K$, while $|W| = 5M$ for the double torus, and $|W| = 30M$ for the Klein bottle. By reducing $|W|$ at some point we experienced fail cases, i.e. high dimensional simplices. We believe that the number of required samples is strongly related to the extrinsic curvature of a manifold within its ambient space, which will be analyzed more carefully in future work.

4 Conclusion

First experiments indicate that our implementation enables manifold reconstruction based on relaxed Delaunay complex [2]. Currently, the main drawback is that the density of the point cloud needs to be very high in relation to the desired vertex density, i.e., number of vertices in the mesh. For the future we hope that our practical algorithm will enable improvements of the theory as well as practical reconstruction.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22:481–504, 1998.

[2] J.-D. Boissonnat, R. D. Dyer, A. Ghosh, and S. Oudot. Equating the witness and restricted Delaunay complexes. Technical Report CGL-TR-24, 2011.

[3] V. De Silva. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*, 135:39–64, 2008.

[4] R. do Rego and A. Araújo. A surface reconstruction method based on self-organizing maps and intrinsic Delaunay triangulation. In *IJCNN*, pages 1–8, 2010.

[5] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *SOCG*, pages 285–292, 1994.

Longest-Edge Algorithms for Quality Delaunay Refinement

Carlos Bedregal* Maria-Cecilia Rivara*

Abstract

We study of the computational cost of the Lepp-Delaunay algorithm for the quality triangulation of PSLG geometries, studying their theoretical and geometric properties. Our goal is to demonstrate that this is a provably good algorithm, proving that satisfying size-optimal meshes are obtained.

1 Introduction

Chew [3] and Ruppert [5] developed Delaunay-based refinement algorithms for constrained and un-constrained Delaunay triangulations, offering theoretical guarantees on both the shape and size of the mesh produced for the quality triangulation problem of PSLG geometries.

Lepp-Bisection algorithm [4] was designed for the (local and iterative) refinement of quality triangulations for finite element applications. It was recently proved that this produces size-optimal refined triangulations [2, 1].

Lepp-Delaunay algorithm [4] is supported by the properties of both the longest-edge propagating path and the Delaunay algorithm. The longest edge propagation strategy allows a more natural *local* improvement of the target triangles, producing the most equilateral triangulation, with the advantage of being independent of the order of processing.

Our current study focuses on the Lepp-Delaunay algorithm for the quality triangulation of PSLG geometries. We aim to provide the proofs that this algorithm produces size-optimal triangulations.

1.1 Longest Edge Propagating Path and Terminal Triangles

When refining a non-quality triangle, the Lepp-Delaunay algorithm *propagates* the refinement to a sequence of neighbor triangles in order to maintain a good graded, conforming triangulation (where the intersection of neighbor triangles is either a common vertex or a common edge). Given a triangle t, a *longest edge propagating path*, $\texttt{Lepp}(t)$, is defined as the sequence of increasing neighbor triangles by the longest edge that starts with t and finishes either with two *terminal triangles* sharing a common – terminal – longest edge, or with a boundary terminal triangle. (See Fig. 1)

*Department of Computer Science, University of Chile, {cbedrega,mcrivara}@dcc.uchile.cl

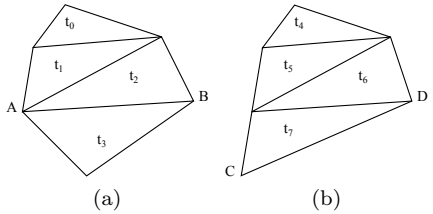

Figure 1: (a) AB is an interior terminal edge shared by terminal triangles $\{t_2, t_3\}$ of $\texttt{Lepp}(t_0) = \{t_0, t_1, t_2, t_3\}$. (b) CD is a boundary terminal edge with terminal triangle $\{t_7\}$ of $\texttt{Lepp}(t_4) = \{t_4, t_5, t_6, t_7\}$

2 Study on the Lepp-Delaunay Algorithm

We consider the quality, size-optimal, Delaunay triangulation of a PSLG geometry, also called Delaunay refinement. Given a PSLG-polygon P (a non-convex polygon including interior points and/or interior edges), construct a quality triangulation of P considering both the interior of P and its boundary. Algorithm 1 briefly describes the Lepp-Delaunay algorithm.

Algorithm 1 Lepp-Delaunay Algorithm

Input: PSLG polygon P and min angle parameter ϵ
  Construct τ, the Constrained Delaunay Triangulation of P
  Find $S \subset \tau$, set of bad angled triangles defined by ϵ
  for each t in S **do**
    while t remains in τ **do**
      Find $\texttt{Lepp}(t)$, terminal triangles t_1, t_2, terminal edge l, and p the midpoint of l
      Perform constrained Delaunay insertion of p
      Update S
    end while
  end for

Our analysis is based on the facts that (a) Lepp-Delaunay algorithm inserts new points after finding a pair of terminal triangles; (b) the new point inserted corresponds to the midpoint of the edge shared by these triangles. Since terminal triangles satisfy the Delaunay property, the following result holds [4].

Proposition 1: For any pair of non-boundary Delaunay terminal triangles, their biggest angle $\gamma \leq 120°$. (See Fig. 2)

This bound on the angles translates into a bound

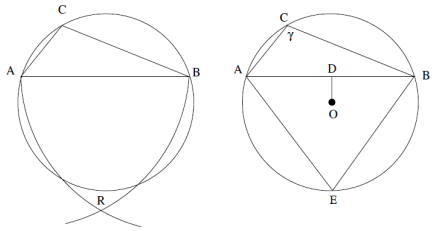

Figure 2: (left) Region R defines the area for point E of neighbor triangle ABE. (right) Region R becomes a point when $\gamma = 120°$, triangle ABE is equilateral.

on the distance from the new point to existing points, bounding also the length of new edges. Then, this distance can be bounded by a constant that depends on the size of existing edges (like the shortest edge or the radius of the circumcircle of the terminal triangle).

2.1 Bounds Based on Geometrical Properties

Consider point D the midpoint of longest edge AB, $d(.)$ the distance from new point D to a previous point $p = \{A, B, C, E\}$ and r the circumradius of triangle ABC.

Proposition 2: For any Delaunay terminal triangle it holds that $d(p) \geq r/2$.

Consider the two terminal triangles ABC and ABE as illustrated in Fig. 2.

For the obtuse triangle ABC with $\gamma = 120°$ it holds that $OD = r/2$ and angle $DOB = 60$, so $d(A) = d(B) = r\sqrt{3}/2$. Distance $d(C)$ will depend on the position of point C, minimized when ABC is isosceles, and maximized when angle α tends to 0, so $r/2 \leq d(C) < r$. Distance $d(E) = 3r/2$, since AEB is equilateral.

For obtuse triangles with $90° < \gamma < 120°$, $d(A)$, $d(B)$ and $d(C)$ increases as point D approaches point O, so $r\sqrt{3}/2 < d(A) < r$ and $r/2 < d(C) < r$. On the other hand, $d(E)$ decreases the closer point E lies to the circumcircle, and increases the further it lies, therefore $r < d(E) < r\sqrt{3}$.

For rectangle triangles, $\gamma = 90°$, point D corresponds to the circumcenter of triangle ABC, so the distance to previous points A, B or C is r. For $d(E)$ the argument used obtuse with obtuse triangles still holds, so $r \leq d(E) \leq r\sqrt{3}$.

For acute triangles $d(A)$ decreases as point D moves away from O; the furthest place for inserting point D happens when triangle ABC is equilateral, so $r\sqrt{3}/2 \leq d(A) < r$. Distance $d(C)$ on the other hand increases as D moves away from the circumcenter of ABC, and is maximized for the equilateral triangle, therefore $r < d(C) \leq 3r/2$. The argument of obtuse triangles for $d(E)$ holds for acute triangles, with $d(E)$ minimized for the isosceles triangle AEB with $\gamma = 120$, and maximized for the equilateral triangle AEB (which is similar to ABC); therefore $r/2 \leq d(E) \leq r$.

Additionally, we found that the distance of the new point D to existing points can also be bounded by the shortest edge of the terminal triangles. This yields a similar result to the one described above: the new edge is at least half the length of the shortest existing edge. For lack of space this proof is not included here.

3 Current Work

With these results we aim to establish a bound on $d()$, as well to define a bound on the number of points inserted by the algorithm, using the concept of *local feature size* introduced by Ruppert [5]. This theoretical framework would help us to prove that the Lepp-Delaunay algorithm generates meshes that are both graded and size-optimal. A similar machinery has been used in the demonstrations of other Delaunay refinement algorithms [6].

Proving that the edge function is within a constant factor of the local feature size yields to an output mesh that is size-optimal within a constant factor of the minimum mesh meeting the same quality bounds. The constant is independent of the size of the geometry and depends on the requirements of the application. After the constant is determined, the rest of the demonstration of size optimality follows from the theoretical framework.

We are also developing a benchmark to experimentally evaluate the performance of the algorithm. We plan to measure the generation of new triangles during the refinement process and to compare these results with other known techniques.

References

[1] C. Bedregal and M.-C. Rivara. A study on size-optimal longest edge refinement algorithms. In *Proceedings of the 21st Int. Meshing Roundtable*, pages 121–136. Springer Berlin Heidelberg, 2012.

[2] C. Bedregal and M.-C. Rivara. Longest-edge algorithms for size-optimal refinement of triangulations. Technical report, Unpublished manuscript, 2013.

[3] L. P. Chew. Guaranteed-quality triangular meshes. Technical report, Department of Computer Science, Cornell University, 1989.

[4] M.-C. Rivara. New longest-edge algorithms for the refinement and/or improvement of unstructured triangulations. *Int. J. Mumer. Meth. Engrg.*, 40(18):3313–3324, 1997.

[5] J. Ruppert. A delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. of Algorithms*, 18(3):548 – 585, 1995.

[6] A. Üngör. Off-centers: A new type of steiner points for computing size-optimal quality-guaranteed delaunay triangulations. In *LATIN*, pages 152–161, 2004.

NTU-Neutralized Complexes for the Optimal Homologous Chain Problem

Bala Krishnamoorthy* Gavin W. Smith†

Abstract

Given a simplicial complex K with weights on its simplices and a chain on K, the Optimal Homologous Chain Problem (OHCP) is to find a chain with minimal weight that is homologous (over $\mathbb{Z}$) to the given chain. OHCP has been shown to be NP-complete, but if the boundary matrix of K is totally unimodular (TU), it becomes solvable in polynomial time when modeled as a linear program (LP). We define a condition on the simplicial complex called non-total-unimodularity neutralized, or *NTU neutralized*, which ensures that even when the boundary matrix is not TU, the OHCP LP must contain an integral optimal vertex for every input chain. Our complete paper is at http://arxiv.org/abs/1304.4985.

1 Introduction

Consider the LP model of OHCP [1] for a given p-chain $\mathbf{c}$ on a simplicial complex K with m p-simplices and n q-simplices, where $q = p+1$. Let $\mathbf{w}$ be a set of nonnegative weights on the p-simplices of K, and B represent the boundary matrix $[\partial_q]$ of K.

$$\begin{aligned} \min \quad & \begin{bmatrix} \mathbf{w}^T & \mathbf{w}^T & \mathbf{0}^T & \mathbf{0}^T \end{bmatrix} \mathbf{z} \\ \text{subject to} \quad & \begin{bmatrix} I & -I & -B & B \end{bmatrix} \mathbf{z} = \mathbf{c} \\ & \mathbf{z} \geq \mathbf{0} \end{aligned} \qquad (1)$$

There are two variables in $\mathbf{z}$ for each p- and q-simplex in K, one each for possible positive and negative coefficients of these simplices. We refer to the coordinates corresponding to p-simplices (q-simplices) as x-variables (y-variables). For two x-variables (y-variables) corresponding to the same simplex, we call their difference the p-coefficient (q-coefficient). We refer to (1) as the OHCP LP, and let P denote its feasible region, and A denote the equality constraint matrix.

It was shown [1] that A is totally unimodular (TU) iff B is so. If A is TU, then all vertices of the OHCP LP are integral for every input chain $\mathbf{c} \in \mathbb{Z}^m$, which in turn guarantees that solving the OHCP LP solves the associated OHCP instance. For $p = 1$, this result amounts to K not having any Möbius strips. We demonstrate that even when B is *not* TU, the OHCP LP could still solve the instance. We provide a characterization of all simplicial complexes where the OHCP LP can be used to solve every OHCP instance, even when B is not TU.

*Department of Mathematics, Washington State University, bkrishna@math.wsu.edu

†Department of Mathematics, Washington State University, gsmith@math.wsu.edu

1.1 An Example

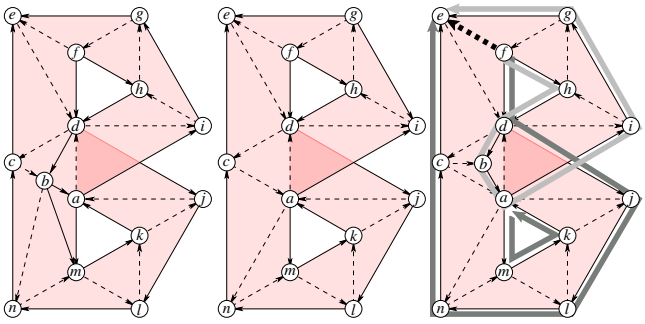

Consider the three different triangulations of a space. In the left and right figures, we have a Möbius strip self-intersecting at one (d) and two vertices (a, d), respectively. In the middle figure, the self intersection is along the *edge* ad, and so we do not have a Möbius strip. Hence A is TU for the middle figure, while it is not for the other two triangulations. Still, in the right triangulation, the OHCP LP has an integral optimal solution for every input chain. For example, consider the edge fe (thick, black, and dashed) with multiplier 1 as the input chain. Let the weights of dashed edges be 1, of db and ba be 0.1, and of all other edges be 0.05. In the right complex, there are two integral optimal solutions to the OHCP LP (shown with thick arrows in contrasting shades of gray). Any convex combination of these two chains also represents an optimal solution for the OHCP LP, including the solution consisting of all solid edges with multipliers ± 0.5, which is also a vertex of P. This behavior may be explained by the presence of a disk whose boundary is an odd number of dashed edges, e.g., triangle adc, which *neutralizes* the Möbius strip. If we alter the left triangulation by adding the triangle akm, then it also becomes NTU-neutralized; $adcbnmlkj$ is a disc whose boundary is 9 dashed edges.

1.2 Related Work

Dey, Hirani, and Krishnamoorthy [1] showed A is TU if and only if there are no pure subcomplexes $L_0 \subset L (\subseteq K)$ of dimension p and q where $\mathsf{H}_p(L, L_0)$ has torsion.

2 Solutions and Projections with Minimal Inputs

Consider the special case of (1) with elementary input chain $\mathbf{e}_i$ for some $i \leq m$. We refer to this instance of an OHCP LP as OHCP_i, and its feasible region as P_i.

Lemma 1 *For a given complex K, there is an OHCP LP with integral input chain $\mathbf{c}$ that has a non-integral basic solution if and only if there is some i such that $OHCP_i$ has a basic solution where all y-variables that are non-zero are non-integral.*

Call the space of x variables $\mathscr{X}$ $\left(= \mathbb{R}^{2m}\right)$. We discuss projections of $P, P_A, \text{Ker}(A)$, basic solutions, and vertices onto $\mathscr{X}$. For any of these objects Ω, let $\Omega|_{\mathscr{X}}$ be the projection of Ω onto $\mathscr{X}$. Note that because all coefficients of y-variables are 0 in the objective function of (1), the optimum value occurs at a vertex of $P|_{\mathscr{X}}$.

Corollary 2 *For a given complex K, there is an OHCP LP with integral input chain $\mathbf{c}$ and polyhedron P that has a non-integral vertex $\mathbf{z}$ where $\mathbf{z}|_{\mathscr{X}}$ is a vertex of $P|_{\mathscr{X}}$ if and only if there is some i such that $OHCP_i$ has a non-integral vertex $\mathbf{z}'$ with all non-zero y-variables non-integral, and where $\mathbf{z}'|_{\mathscr{X}}$ is a vertex of $P_i|_{\mathscr{X}}$.*

3 Column-Wise Containment in MNTU Matrices

We create a few definitions helpful for discussion. $\mathbf{z}$ is *concise* if for each pair of variables for the same simplex, one (or both) is 0. A set of vectors is *linearly concise* if any linear combination of the set is concise. For any OHCP LP, there is the unique feasible concise solution where all the y-variables are 0. We call this solution the *identity* solution, and denote it $\mathbf{z}^I$.

We examine vertices whose non-zero q-coefficients are column-wise contained in minimal non-TU (MNTU) submatrices of $[\partial_q]$. For any MNTU submatrix M, let Q_M represent the set of columns intersecting M, and let $\mathscr{M}$ be the set of elements of $\text{Ker}(A)$ whose non-zero q-coefficients are contained in Q_M. We call rows of $[\partial_q]$ that intersect M *interior* rows. For any concise $\mathbf{z} \in \mathbb{R}^{2(m+n)}$, we define $\mathbf{m}(\mathbf{z})$ to be a particular element of $\mathscr{M}$ with $\{\mathbf{z}, \mathbf{m}(\mathbf{z})\}$ linearly concise, and for each interior row i, the i^{th} p-coefficient of $\mathbf{z}$ and $\mathbf{m}(\mathbf{z})$ are equal.

Theorem 3 *For any MNTU submatrix M of $[\partial_q]$, and any interior row i, there is a unique vertex $\mathbf{z}^i$ of P_i whose non-zero q-coefficients are contained in Q_M and whose p-coefficients of interior rows are all 0. This vertex is $\mathbf{z}^I - \mathbf{m}(\mathbf{z}^I)$ where $\mathbf{z}^I$ is the identity solution to $OHCP_i$.*

4 NTU-Neutralized Complexes

Definition 1 *For any interior row i of a MNTU submatrix M of $[\partial_q]$, let $\mathbf{k}^i$ represent a concise integral element of $\text{Ker}(A)$ where the sum of p-coefficients of interior rows is odd, $\left(\mathbf{k}^i - \mathbf{m}(\mathbf{k}^i)\right)|_{\mathscr{X}} \neq \mathbf{0}$, and the absolute value of each p-coefficient of $\mathbf{k}^i - \mathbf{m}(\mathbf{k}^i)$ is not greater than its absolute value in $\mathbf{z}^i$. If each interior row i of M has such a $\mathbf{k}^i$, then M is* neutralized. *If all MNTU submatrices of $[\partial_q]$ are neutralized, then K is* NTU-neutralized in the q^{th} dimension.

Theorem 4 *For any MNTU submatrix M of $[\partial_q]$, the projection $\mathbf{z}^i|_{\mathscr{X}}$ for each interior row i is a convex combination of $\mathbf{z}^1|_{\mathscr{X}}$ and $\mathbf{z}^2|_{\mathscr{X}}$ where both $\mathbf{z}^1$ and $\mathbf{z}^2$ are integral elements of P_i if and only if M is neutralized.*

Theorem 5 *For a given complex K with boundary matrix $[\partial_q]$, the projection $\mathbf{z}|_{\mathscr{X}}$ of each non-integral vertex $\mathbf{z}$ of any OHCP LP over K with integral input p-chain $\mathbf{c}$ and polyhedron P is not a vertex of $P|_{\mathscr{X}}$ if and only if K is NTU-neutralized in the q^{th} dimension.*

In the right triangulation of the Example 1.1 for the OHCP with input chain fe, $\mathbf{z}^i|_{\mathscr{X}}$ is all solid edges, each with a coefficient of 0.5. $\mathbf{m}(\mathbf{z}^I)|_{\mathscr{X}}$ is the union of these solid edges, each with coefficient -0.5, together with edge fe with coefficient 1. The light and dark grey chains are projections of two integral vertices $\mathbf{z}^1|_{\mathscr{X}}$ and $\mathbf{z}^2|_{\mathscr{X}}$, respectively. All coefficients in both of these projections are 1. Triangle adc satisfies all criteria for $\mathbf{k}^i$. Then $\mathbf{z}^i|_{\mathscr{X}} = \frac{1}{2}(\mathbf{z}^1|_{\mathscr{X}} + \mathbf{z}^2|_{\mathscr{X}})$. Also, $\mathbf{m}(\mathbf{k}^i)|_{\mathscr{X}} = \mathbf{k}^i|_{\mathscr{X}} - \frac{1}{2}(\mathbf{z}^1|_{\mathscr{X}} - \mathbf{z}^2|_{\mathscr{X}})$, and $\mathbf{z}^1|_{\mathscr{X}} - (\mathbf{k}^i - \mathbf{m}(\mathbf{k}^i))|_{\mathscr{X}} = \mathbf{z}^2|_{\mathscr{X}} + (\mathbf{k}^i - \mathbf{m}(\mathbf{k}^i))|_{\mathscr{X}} = \mathbf{z}^i|_{\mathscr{X}}$.

5 Conclusion

Our main result implies that when K is NTU neutralized, if an optimal solution of the OHCP LP is non-integral then there must exist another integral optimal solution with the same total weight. If a standard LP algorithm finds the fractional optimal solution, we should be able to find an adjacent integral optimal solution using an approach similar to that of Güler et al. [2] for the same task in the context of interior point methods for linear programming. This approach should run in strongly polynomial time. An unresolved question of interest is the complexity of checking whether a complex is NTU neutralized.

References

[1] T. K. Dey, A. N. Hirani, and B. Krishnamoorthy. Optimal Homologous Cycles, Total Unimodularity, and Linear Programming. *SIAM Journal on Computing*, 40(4):1026–1040, 2011.

[2] O. Güler, D. den Hertog, C. Roos, T. Terlaky, and T. Tsuchiya. Degeneracy in interior point methods for linear programming: a survey. *Annals of Operations Research*, 46- 47(1):107138, March 1993.

On the Subdivision Containment Problem for Random 2-Complexes

Anna Gundert* Uli Wagner†

Abstract

For random graphs, the *containment problem* considers the probability that a binomial random graph $G(n,p)$ contains a given graph as a substructure. When asking for a copy of a *subdivision* of the given graph, it is well-known that the (sharp) threshold is at $p = 1/n$.

We consider the analogous question for random 2-dimensional complexes $X^2(n,p)$. Improving previous results, we show that $p = \Theta(1/\sqrt{n})$ is the (coarse) threshold for containing a subdivision of any fixed complete 2-complex.

Introduction A basic problem in graph theory is to determine whether a given graph G, which may be thought of as "large", contains a fixed graph H as a substructure. The most straightforward form of containment is that G contains a copy of H as a *subgraph*. Another important variant is that G contains some *subdivision* of H as a subgraph; in this case, one also says that G contains H as a *topological minor*.

For random graphs, the *containment problem* considers the probability that a binomial random graph $G(n,p)$ contains a copy of a given graph H. For subgraph containment, it is well-known [4] that this probability has a (coarse) threshold of $\Theta(n^{-1/m(H)})$, where $m(H)$ is the density of the densest subgraph of H. The (sharp) threshold for containment of any fixed graph as a topological minor is $p = 1/n$ by a well-known result of Ajtai, Komlós and Szemerédi [1].

A random model $X^k(n,p)$ for simplicial complexes of arbitrary fixed dimension k which generalizes the random graph model $G(n,p)$ was introduced by Linial and Meshulam [9] and has since then been studied extensively, see, e.g., [10, 3, 2, 6, 7, 11]. Subgraph containment admits a direct generalization to higher dimensions: We can ask whether a given simplicial complex X contains a fixed complex F as a subcomplex. The proof methods for random graphs extend directly to random 2-complexes, and the threshold probability for $X^2(n,p)$ to contain a fixed complex F as a subcomplex is given by the density (in terms of triangles versus vertices) of the densest subcomplex of F, see [3, 6]. For the question of topological minors, a natural analogue is whether X contains a subdivision of F. Cohen, Costa, Farber and Kappeler [6] show that for any $\epsilon > 0$ and $p \geq n^{-1/2+\epsilon}$, the random complex $X^2(n,p)$ *asymptotically almost surely (a.a.s.)*, i.e., with probability tending to 1 as $n \to \infty$, contains a subdivision of any fixed F.

We improve their result by showing that the (coarse) threshold for containing a subdivision of any fixed complete 2-complex is at $p = \Theta(1/\sqrt{n})$.

Theorem 1

(i) For every $t \in \mathbb{N}$ there is a constant $c_t > 1$ such that $X^2(n,p)$ with $p = \sqrt{c_t/n}$ a.a.s. contains a subdivision of the complete 2-complex K_t^2 on t vertices.

(ii) For any $t \geq 7$ the complex $X^2(n,p)$ with $p \ll \sqrt{1/n}$ a.a.s. does not contain a subdivision of K_t^2.

Here we focus on part (i). The result in [6] is proved by reduction to the subcomplex containment problem, by showing that a given F can be subdivided to decrease its triangle density. We use a different approach, based on an idea going back at least to Brown, Erdős and Sós [5] and used also in [3]: Our proof is based on studying common links of pairs of vertices, which form random graphs of the type $G(n-2,p^2)$.

The somewhat technical proof of part (ii) is based on bounds on the number of triangulations of a fixed surface and will be presented in the full version.

Preliminaries A (finite) 2-*dimensional complex* (or 2-*complex*) is a set system of the form $X = V \cup E \cup T$, $E \subseteq \binom{V}{2}$, $T \subseteq \binom{V}{3}$, such that $(e \subset \tau, \tau \in T)$ implies $e \in E$. The random 2-complex $X^2(n,p)$ has vertex set $V = [n]$, a *complete* 1-*skeleton*, i.e., $E = \binom{[n]}{2}$, and every $\tau \in \binom{[n]}{3}$ is added to T independently with probability p, which may be constant or, more generally, a function $p(n)$ depending on n. The *complete* 2-*complex* K_t^2 has $V = [n]$, $E = \binom{[n]}{2}$ and $T = \binom{[n]}{3}$.

A *subdivision* of $X = (V,E,T)$ is a 2-complex $X' = (V',E',T')$ with $V \subseteq V'$ that is obtained by replacing the edges of X with internally-disjoint paths and the triangles of X with internally-disjoint triangulated disks such that for every triangle the subdivision of the triangle agrees with the subdivisions of its edges. The *link* $\mathrm{lk}_X(v)$ of a vertex $v \in V$ in X is the graph $(\{u \in V : \{u,v\} \in E\}, \{e : e \cup \{v\} \in T\})$.

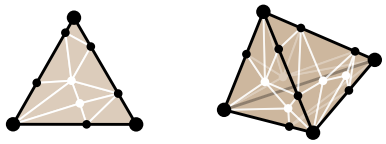

Figure 1: Subdivisions of K_3^2 (a *triangulated disk*) and of K_4^2. Vertices and edges internal to triangles are drawn in white.

*Institut für Theor. Informatik, ETH Zürich, Switzerland, anna.gundert@inf.ethz.ch. Supported by Swiss National Science Foundation (SNF Projects 200021-125309, 200020-138230).

†Institute of Science and Technology Austria, uli@ist.ac.at

The Main Idea For fixed $t > 0$, we aim to find a subdivision of K_t^2 in $X^2(n,p)$. We would be allowed to subdivide the edges of K_t^2, but there will be no need for this: we will find t vertices such that the triangles spanned by the edges between any three of them can be filled with disjoint triangulated disks.

Fix a partition of the vertex set $V = [n]$ into two sets U and W, each of size $\frac{n}{2}$. We will choose the t vertices of K_t^2 from U, whereas the internal vertices for fillings will come from W. To ensure disjointness of the fillings of different triangles, we partition W into $\binom{t}{3}$ sets W_σ, $\sigma \in \binom{[t]}{3}$, each of size $n/(2\binom{t}{3})$, and choose the internal vertices of the filling for each triangle $\sigma \in \binom{[t]}{3}$ from W_σ.

For two vertices a and b, denote by $G_{a,b}$ the graph $\mathrm{lk}_X(a) \cap \mathrm{lk}_X(b)$ which has vertex set $[n] \backslash \{a,b\}$ and edge set $\{\{v,w\} : \{a,v,w\}, \{b,v,w\} \in X\}$. Denote by $C_{a,b}^\sigma$ the largest connected component of $G_{a,b}[W_\sigma]$. If there are several components of maximum size, let $C_{a,b}^\sigma$ be the one containing the smallest vertex.

The basis of our proof is the following lemma:

Lemma 2 *Let X be a 2-complex with vertex set $[n]$ and complete 1-skeleton. Suppose there is a set $A \subset U$, $|A| = t$ such that for every $\{a,b\} \subset A$, $c \in A \backslash \{a,b\}$, and $\sigma \in \binom{[t]}{3}$ there are $v, w \in C_{a,b}^\sigma$ with $\{a,b,v\} \in X$ and $\{c,w\} \in G_{a,b}$. Then X contains a subdivision of K_t^2.*

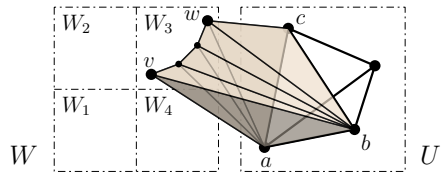

Figure 2: Part of a subdivision of a K_4^2.

Proof. For $\{a,b,c\} \subset A$ there exists a path in $G_{a,b}[W_\sigma]$ between v and w which together with the triangles $\{a,b,v\}$, $\{a,w,c\}$ and $\{b,w,c\}$ creates a disk filling the triangle created by the edges $\{a,b\}, \{b,c\}, \{a,c\}$. See Figure 2 for an illustration. By fixing a distinct $\sigma \in \binom{[t]}{3}$ for every $\{a,b,c\}$ we ensure disjoint fillings. □

For fixed $a, b \in U$ and $\sigma \in \binom{[t]}{3}$, call $u \in U \setminus \{a,b\}$ *connected to* $C_{a,b}^\sigma$ if $\{u,w\} \in G_{a,b}$ for some $w \in C_{a,b}^\sigma$ and let $N_{a,b}^\sigma = \{u \in U \setminus \{a,b\} : u \text{ connected to } C_{a,b}^\sigma\}$. Consider two families of complexes $X = \left([n], \binom{[n]}{2}, T\right)$:

- $\mathcal{A}_{a,b,\sigma} = \{X : \exists v \in C_{a,b}^\sigma \text{ with } \{a,b,v\} \in T\}$.
- For $\delta > 0$: $\mathcal{B}_{a,b,\sigma,\delta} = \{X : |N_{a,b}^\sigma| \geq (1-\delta)(|U|-2)\}$.

Lemma 3 *Let X be a 2-complex with vertex set $[n]$ and complete 1-skeleton. If there is a $\delta < 1/\binom{t}{3}^2$ such that $X \in \mathcal{A}_{a,b,\sigma} \cap \mathcal{B}_{a,b,\sigma,\delta}$ for all $a, b \in U$ and $\sigma \in \binom{[t]}{3}$, there exists $A \in \binom{U}{t}$ fulfilling the conditions of Lemma 2.*

Proof. $A \in \binom{U}{t}$ chosen uniformly at random fulfills the conditions of Lemma 2 with non-zero probability. □

Sketch of Proof for Theorem 1. For $t \in \mathbb{N}$ we show that there is $c_t > 1$ and $\delta < 1/\binom{t}{3}^2$ such that $X^2(n,p)$ for $p = \sqrt{c_t/n}$ a.a.s. fulfills the conditions of Lemma 3, i.e., $X^2(n,p) \in \bigcap_{a,b,\sigma} \mathcal{A}_{a,b,\sigma} \cap \mathcal{B}_{a,b,\sigma,\delta}$.

The probability of the events $\mathcal{A}_{a,b,\sigma}$ and $\mathcal{B}_{a,b,\sigma,\delta}$ depends on the size of $C_{a,b}^\sigma$, the largest connected component of the graph $G_{a,b}[W_\sigma]$, which is a random graph of type $G(|W_\sigma|, p^2)$. As $p^2 = c_t/n$, for c_t large enough such a graph *fails* to have a giant component of size linear in $|W_\sigma|$ with exponentially small probability, see, e.g., [8].

For fixed $a, b \in U$, $\sigma \in \binom{[t]}{3}$, both the number of $v \in C_{a,b}^\sigma$ with $\{a,b,v\} \in X$ as well as the number of $u \in N_{a,b}^\sigma$ are binomially distributed. Conditioning on the highly probable event that $C_{a,b}^\sigma$ is of linear size, we can then use Chernoff's inequality to show that $X^2(n,p) \notin \mathcal{A}_{a,b,\sigma} \cap \mathcal{B}_{a,b,\sigma,\delta}$ with probability exponentially small in n. As the number of choices of a, b and σ is polynomial in n, a union bound finishes the proof. □

References

[1] M. Ajtai, J. Komlós, and E. Szemerédi. Topological complete subgraphs in random graphs. *Studia Sci. Math. Hungar.*, 14:293–297, 1979.

[2] L. Aronshtam, N. Linial, T. Łuczak, and R. Meshulam. Collapsibility and vanishing of top homology in random simplicial complexes. *Discrete Comput. Geom.*, 49:317–334, 2013.

[3] E. Babson, C. Hoffman, and M. Kahle. The fundamental group of random 2-complexes. *J. Amer. Math. Soc.*, 24:1–28, 2011.

[4] B. Bollobás. Random graphs. In *Combinatorics*, number 52 in London Mathematical Society Lecture Note Series. Cambridge University Press, 1981.

[5] W. G. Brown, P. Erdős, and V. T. Sós. Some extremal problems on r-graphs. In *New directions in the theory of graphs (Proc. Third Ann Arbor Conf., 1971)*, pages 53–63. Academic Press, New York, 1973.

[6] D. Cohen, A. Costa, M. Farber, and T. Kappeler. Topology of random 2-complexes. *Discrete Comput. Geom.*, 47:117–149, 2012.

[7] D. Kozlov. The threshold function for vanishing of the top homology group of random d-complexes. *Proc. Amer. Math. Soc.*, 138:4517–4527, 2010.

[8] M. Krivelevich and B. Sudakov. The phase transition in random graphs: A simple proof. *Random Struct. Alg.*, 2012. (Published online before inclusion in an issue).

[9] N. Linial and R. Meshulam. Homological connectivity of random 2-complexes. *Combinatorica*, 26(4):475–487, 2006.

[10] R. Meshulam and N. Wallach. Homological connectivity of random k-dimensional complexes. *Random Struct. Alg.*, 34(3):408–417, 2009.

[11] U. Wagner. Minors in random and expanding hypergraphs. In *Proc. of the 27th annual SoCG*, pages 351–360, 2011.

Dimension Detection by Local Homology

Tamal K. Dey* Fengtao Fan† Yusu Wang‡

Abstract

In this paper, we propose a purely topological method to detect the dimension of a manifold from its sampled points. Our approach detects the dimension by determining the local homology of the manifold. We show that the local homology at a point z can be computed *exactly* for manifolds using Vietoris-Rips complexes whose vertices are confined within a local neighborhood of z.

1 Introduction

Learning about a manifold embedded in $\mathbb{R}^d$ from its point data is a key problem in various manifold learning applications. Most times, the intrinsic dimension of the manifold $\mathcal{M}$ is one of the simplest, yet still very important, quantities that one would like to infer from input data. Therefore, considerable research has focused on estimating the dimension for manifolds. Under the statistical setting, different approaches estimate the manifold dimension based on the growth rate of the volume (or some analog of it) of an intrinsic ball [6, 7]. They generally assume that the input points are sampled from some probabilistic distribution supported on the hidden manifold. In the computational geometry community, there are provable dimension detection algorithms which all require (ε, δ)-sampling condition that points are both ε-dense and δ-sparse. Cheng et al. [4] developed an improved algorithm from them to tolerate a small amount of Hausdorff noise (of the order ε^2 times the local feature size). More recently, Cheng et al. [3] proposed an algorithm to estimate the dimension by detecting the so-called *slivers*. This algorithm assumes that the input points are sampled from the hidden manifold using a Poisson process without noise.

In this paper we develop a dimension detection method based on the topological concept of local homology. The idea of using local homology to study stratified spaces from sampled points was first proposed by Bendich et al. [1] and further explored in [2]. Both of them used Delaunay triangulations in their algorithms. In a recent paper [10], Skraba and Wang proposed to approximate the multi-scale representations of local homology using families of Rips complexes.

*Dept. of Comp. Sci. Eng., The Ohio State University, Columbus, OH 43210, USA. Email: tamaldey@cse.ohio-state.edu

†Dept. of Comp. Sci. Eng., The Ohio State University, Columbus, OH 43210, USA. Email: fanf@cse.ohio-state.edu

‡Dept. of Comp. Sci. Eng., The Ohio State University, Columbus, OH 43210, USA. Email: yusu@cse.ohio-state.edu

2 Our results

Given a smooth m-dimensional manifold $\mathcal{M}$ embedded in $\mathbb{R}^d$, the local homology group $\mathsf{H}(\mathcal{M}, \mathcal{M}-z)$ at a point $z \in \mathcal{M}$ is isomorphic to the reduced homology group of an m-dimensional sphere, that is $\mathsf{H}(\mathcal{M}, \mathcal{M}-z) \cong \tilde{\mathsf{H}}(\mathbf{S}^m)$. Hence, given a set of noisy sampled points P of $\mathcal{M}$, we aim to detect the dimension of $\mathcal{M}$ by estimating $\mathsf{H}(\mathcal{M}, \mathcal{M}-z)$ from P. Specifically, we assume that P is an ε-sample of a manifold $\mathcal{M}$ with positive reach $\rho(\mathcal{M})$ in the sense that the Hausdorff distance between P and $\mathcal{M}$ is at most ε. All homology groups are assumed to be under $\mathbb{Z}_2$. The definitions of sampling related concepts such as reach, ε-sample and Hausdorff distance can be found in the book [5] whereas the topology related concepts such as homology, reduced homology and relative homology can be found in the book [9].

We first observe that the local homology $\mathsf{H}(\mathcal{M}, \mathcal{M}-z)$ can be inferred by the image $\mathrm{Im}(i_*)$ of the map $i_* : \mathsf{H}(\mathcal{M}, \mathcal{M}-\mathring{D}_2) \to \mathsf{H}(\mathcal{M}, \mathcal{M}-\mathring{D}_1)$ which is induced from the inclusion map $i : (\mathcal{M}, \mathcal{M}-\mathring{D}_2) \hookrightarrow (\mathcal{M}, \mathcal{M}-\mathring{D}_1)$, where D_1 and D_2 ($D_1 \subset D_2 \subset \mathcal{M}$) are two closed topological balls containing z. However, in general, we only have an ε-sample of the hidden manifold. Interestingly, under appropriate sampling conditions, the manifold $\mathcal{M}$ is a deformation retract of $\mathbb{X}_\alpha = \cup_{q\in P} B_\alpha(q)$ which is the union of Euclidean balls centered at samples with radius α. This deformation retraction and our observation enable us to prove that the local homology of $\mathcal{M}$ can be inferred from the relative homology groups in $\mathbb{X}_\alpha$. Let $\bar{p} \in \mathcal{M}$ denote the closest point to $p \in P$. Denote $B_r(z)$ the Euclidean ball centered at z with radius r and $\mathring{B}_r(z)$ the interior of $B_r(z)$. Let $\mathbb{B}^{\alpha,\beta} = \mathbb{X}_\alpha - \mathring{B}_\beta(p) \cap \mathbb{X}_\alpha$ where $\mathbb{X}_\alpha = \cup_{q\in P} B_\alpha(q)$. We show the following.

Proposition 1 *Let $0 < \varepsilon < \frac{\rho(\mathcal{M})}{22}$, and $\theta_1 \le \alpha \le \alpha' \le \theta_2'$, where $\theta_1 = \frac{(\varepsilon+\rho(\mathcal{M})) - \sqrt{\varepsilon^2+\rho(\mathcal{M})^2 - 6\varepsilon\rho(\mathcal{M})}}{2}$ and $\theta_2' = \frac{\rho(\mathcal{M})-13\varepsilon}{4}$. Set $\delta = \alpha + 3\varepsilon$ and $\delta' = \alpha' + 3\varepsilon$. For $\varepsilon < \lambda' < \rho(\mathcal{M}) - 4\delta'$ and $\lambda \ge \lambda' + 2(\alpha' - \alpha)$, we have,*

$$\mathrm{Im}(\mathsf{H}(\mathbb{X}_\alpha, \mathbb{B}^{\alpha,\lambda+3\delta}) \to \mathsf{H}(\mathbb{X}_{\alpha'}, \mathbb{B}^{\alpha',\lambda'+\delta'})) \cong \mathsf{H}(\mathcal{M}, \mathcal{M}-\bar{p}).$$

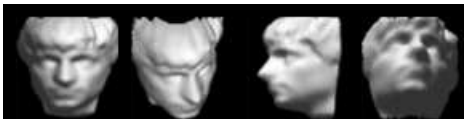

(a) **Head**

(b) **D1**

(c) **D0**

Figure 1: Image data : **Head**, **D1** and **D0**

	Shift	**Head**	**D1**	**D0**
Ours	2	3	4	3
SLIVER	3	4	3	2
MLE	4.27	4.31	11.47	14.86
MA	3.35	4.47	10.77	13.93
PN	3.62	3.98	6.22	8.86
LPCA	3	3	5	8.86
ISOMAP	2	3	5	[3, 6]

Table 1: Comparison results

In fact, knowing the local neighborhood of the point p is enough for extracting the local homology at $\bar{p}$. Specifically, denote $\mathbb{X}_{\alpha,r} = \mathbb{X}_\alpha \cap B_r(p)$ and $\mathbb{X}^\beta_{\alpha,r} = \mathbb{X}_{\alpha,r} \cap B^\beta(p)$, where $B^\beta(p) = \mathbb{R}^d \setminus \mathring{B}_\beta(p)$. We prove that $\mathsf{H}(\mathcal{M}, \mathcal{M} - \bar{p})$ can be obtained from the relative homology groups defined over the neighborhood of p such as $\mathsf{H}(\mathbb{X}_{\alpha,r}, \mathbb{X}^\beta_{\alpha,r})$ by applying the Excision theorem ([9]). We then push our results further by relating Rips complexes to the nerves of $\mathbb{X}_{\alpha,r}$ and $\mathbb{X}^\beta_{\alpha,r}$. The Rips complex $\mathcal{R}^r(P)$ on the point set P contains all simplices $\sigma = [p_0 p_1 \dots p_k]$ $(p_i \in P)$ with the pairwise Euclidean distance $d(p_i, p_j) < r$ for all $0 \le i, j \le k$. Denote $P_{\alpha,r} = \{p_i \in P \mid B_\alpha(p_i) \cap B_r(p) \neq \emptyset\}$ and $P^\beta_{\alpha,r} = \{P_{\alpha,r} \cap B^\beta(p) \neq \emptyset\}$. We show that the local homology of $\mathcal{M}$ at $\bar{p}$ can be computed from the Rips complexes built on $P_{\alpha,r}$'s and $P^\beta_{\alpha,r}$'s.

Theorem 2 *Let* $0 < \varepsilon < \frac{\rho(\mathcal{M})}{58}$ *and* $\theta_1 \le \alpha \le \frac{\rho(\mathcal{M}) - 13\varepsilon}{22}$. *Let* $\varepsilon < \eta_1, \eta_2 < \rho(\mathcal{M})$ *such that* $\eta_1 \ge 9\alpha + 4\varepsilon$ *and* $\eta_2 \ge \eta_1 + 12\alpha + 6\varepsilon$. *The inclusion*

$$j_\alpha : (\mathcal{R}^{2\alpha}(P_{\alpha,r}), \mathcal{R}^{2\alpha}(P^{\eta_2}_{\alpha,r})) \hookrightarrow (\mathcal{R}^{6\alpha}(P_{3\alpha,r}), \mathcal{R}^{6\alpha}(P^{\eta_1}_{3\alpha,r}))$$

satisfies $\mathrm{Im}(j_{\alpha *}) \cong \mathsf{H}(\mathcal{M}, \mathcal{M} - \bar{p})$ *for any* $r \ge \eta_1 + \eta_2$.

Our main theorem provides the promised topological method for dimension detection. The implementation of our algorithm is straightforward from the theorem.

3 Experimental results

Table 1 presents the comparison results on the real data. The real data contains images of a 2D translation of a smaller image within a black image (**Shift**) (see [3]), a rotating head (**Head**, Fig. 1(a)), handwritten 1's (**D1**, Fig. 1(b)) and 0's (**D0**, Fig. 1(c)) from MNIST database. Our method is compared with the dimension detection method via slivers (SLIVER) [3], the maximum likelihood estimation (MLE) [8], the manifold adaptive method (MA) [6], the packing number method (PN) [7], the local PCA (LPCA) [4], and the isomap method (ISOMAP) [11]. Notice that although **Shift** is uniform and noise free, only ISOMAP and ours get the correct dimension. The dimension of **Head** is considered to be around 3 or 4 in the literature. Ours falls into this range. Although the ground truth dimensions for **D1** and **D0** are unknown, ours along with SLIVER, PN, LPCA and ISOMAP report dimensions in range $[3, 7]$ for **D1** and in range $[2, 9]$ for **D0**.

References

[1] P. Bendich, D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and D. Morozov. Inferring local homology from sampled stratified spaces. In *Proc. 48th Ann. IEEE Sympos. Foundat. Comp. Sci.*, pages 536–546, 2007.

[2] P. Bendich, B. Wang, and S. Mukherjee. Local homology transfer and stratification learning. In *Proc. 23rd Ann. ACM-SIAM Sympos. Discrete Alg.*, pages 1355–1370, 2012.

[3] S.-W. Cheng and M.-K. Chiu. Dimension detection via slivers. In *Proc. 20th Ann. ACM-SIAM Sympos. Discrete Alg.*, pages 1001–1010, 2009.

[4] S.-W. Cheng, Y. Wang, and Z. Wu. Provable dimension detection using principal component analysis. In *Proc. 21st Ann. Sympos. Comput. Geom.*, pages 208–217, 2005.

[5] T. K. Dey. *Curve and surface reconstruction: Algorithms with mathematical analysis*. Cambridge University Press, New York, 2006.

[6] A. M. Farahmand, C. Szepesvári, and J.-Y. Audibert. Manifold-adaptive dimension estimation. In *Proc. 24th Conf. Machine Learning (ICML).*, pages 265–272, 2007.

[7] B. Kégl. Intrinsic dimension estimation using packing numbers. In *Neural Infor. Proc. Sys. Foundation (NIPS)*, pages 681–688, 2002.

[8] E. Levina and P. J. Bickel. Maximum likelihood estimation of intrinsic dimension. In *Advances in Neural Information Processing Systems 17*, pages 777–784, 2005.

[9] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley Publishing Company, Menlo Park, 1984.

[10] P. Skraba and B. Wang. Sampling for local homology with vietoris-rips complexes, 2012. CoRR, abs/1206.0834.

[11] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323, 2000.

Computationally proving triangulated 4-manifolds to be diffeomorphic

Benjamin A. Burton Jonathan Spreer*

Abstract

We present new computational methods for proving diffeomorphy of triangulated 4-manifolds, including algorithms and topological software that can for the first time effectively handle the complexities that arise in dimension four and be used for large scale experiments.

1 Introduction

In dimensions ≤ 3, every topological manifold has a unique smooth structure up to diffeomorphism. In dimensions ≥ 4 this is no longer true: there are pairs of 4-manifolds which are *homeomorphic* (they represent the same topological manifold) but not *diffeomorphic* (they represent two distinct smooth manifolds) [9]. Finding such pairs is important; indeed, the only outstanding variant of the Poincaré conjecture asks whether one can find two non-diffeomorphic smooth 4-spheres [9].

We move this problem to the *piecewise linear* setting, which is better suited for computation. Here manifolds are given as *triangulations* (decompositions into simplices). Piecewise linear manifolds are in 1-to-1 correspondence with smooth manifolds for dimensions ≤ 6, and so results translate between both settings; we use both languages interchangably in this paper.

Despite this equivalence, work on non-diffeomorphic pairs is done exclusively in the smooth setting. The only example of two 4-manifold *triangulations* that are homeomorphic but not diffeomorphic follows a well-established result for the smooth setting [1]. One of the few candidates from the PL-world is a pair of triangulations of the *K3 surface* (one of the four fundamental building blocks of simply connected 4-manifolds): the 16-vertex $(K3)_{16}$ of Casella and Kühnel [6], and the 17-vertex $(K3)_{17}$ of Spreer and Kühnel [12]. The smooth type of $(K3)_{17}$ is canonical, but the smooth type of $(K3)_{16}$ remains unknown. It is conjectured [12]:

Conjecture 1 *$(K3)_{16}$ and $(K3)_{17}$ are diffeomorphic.*

In the computational setting, proving that $(K3)_{16}$ and $(K3)_{17}$ are *homeomorphic* is easy, using the software `simpcomp` [7] in conjunction with Freedman's celebrated classification of simply connected 4-manifolds [8]. Proving they are *diffeomorphic* is much harder. Our approach is based on a theorem of Pachner [11], which states that two triangulated manifolds are diffeomorphic if and only if they are related by a sequence of *bistellar moves* (local modifications).

Bistellar moves offer significant challenges in dimensions ≥ 4: "effective" sequences of moves can be extremely difficult to find. Indeed, the *number* of moves required to connect two diffeomorphic triangulations of size n must have no computable upper bound [10].

Here we describe work in progress towards resolving Conjecture 1, including effective heuristics and fast algorithms for manipulating 4-manifold triangulations with bistellar flips, and a tight lower bound on the size of a 1-vertex triangulation of the $K3$ surface with over a million distinct realisations. This work has wider relevance, and forms the beginning of a larger project to explicitly construct and study "exotic" triangulated 4-manifolds.

2 Minimal triangulations of the K3 surface

For computation, we want to triangulate manifolds using few top-dimensional simplices. We therefore work with *generalised triangulations*, which are collections of abstract simplices whose facets are identified in pairs—these can allow far fewer simplices than the more rigid simplicial complexes. We also favour triangulations with just one vertex, which in lower dimensions offer significant advantages for both theory and computation.

Proving minimality is extremely difficult in three dimensions. In four dimensions, we solve this completely for the $K3$ surface in the one-vertex setting:

Proposition 1 *For any triangulation of the $K3$ surface we have $f_4 \geq 146 - 6f_0$, where f_4 denotes the number of 4-dimensional simplices and f_0 denotes the number of vertices. In the case where $f_0 = 1$, this bound is tight.*

We prove $f_4 \geq 146 - 6f_0$ by combining (i) the fact that the $K3$ surface is simply connected and has Euler characteristic $\chi(K3) = f_0 - f_1 + f_2 - f_3 + f_4 = 24$, with (ii) the *Dehn-Sommerville equations* $2f_1 - 3f_2 + 4f_3 - 5f_4 = 0$ and $2f_3 - 5f_4 = 0$. Here each f_i denotes the number of i-faces of the triangulation.

We prove this bound is tight by reducing both $(K3)_{16}$ and $(K3)_{17}$ using bistellar moves to one-vertex triangulations with $(f_0, f_1, f_2, f_3, f_4) = (1, 1, 234, 350, 140)$. In three dimensions, such a "simplification" of triangulations is fast and effective [4], but in four dimensions it is far more difficult and requires the interaction of many different tools and heuristics.

*School of Mathematics and Physics, The University of Queensland, AUS. bab@maths.uq.edu.au, j.spreer@uq.edu.au

Our approach incorporates: (i) classical greedy techniques, which reduce a triangulation as far as possible using local moves; (ii) "composite" moves that collapse edges and reduce triangulations near low-degree edges and triangles; (iii) simulated annealing [2], where we apply the inverses of reducing moves to escape local minima; (iv) breadth-first searching through the Pachner graph (or "flip graph") [4].

We note that the *interaction* between these techniques is crucial: each technique failed to reduce $(K3)_{16}$ and $(K3)_{17}$ on its own. Again we contrast this with three dimensions, where these techniques are found to be highly effective even in isolation.

All computations were performed using the new 4-manifold toolkit in the software package `Regina` [5].

3 Connecting triangulations $(K3)_{16}$ and $(K3)_{17}$

To prove Conjecture 1 we must find a sequence of local modifications connecting $(K3)_{16}$ and $(K3)_{17}$. In three dimensions, the following approach is often successful: (i) simplify both triangulations as far as possible, and then (ii) repeatedly apply random local modifications that preserve the number of simplices until both triangulations are identical. This often succeeds (provided both triangulations represent the same manifold) because many manifolds appear to have only few distinct minimal triangulations.

In contrast, for the $K3$ surface the number of minimal triangulations appears to be *much* larger, and so the classical approach above does not work. Instead we use a more sophisticated method to ensure that (i) every minimal triangulation is visited only once and (ii) we can detect if a longer "detour" through larger triangulations is required. Specifically, we run a dual-source breadth-first search through the *Pachner graph*, whose nodes represent triangulations of the $K3$ surface and whose arcs represent bistellar flips that preserve the number of simplices. The two sources are our minimal one-vertex triangulations of $(K3)_{16}$ and $(K3)_{17}$.

Each time we perform a local move, we must test whether the resulting triangulation has been seen before (up to combinatorial isomorphism). For this we compute the *isomorphism signature* of the triangulation [3], a polynomial-time computable hash that uniquely identifies the isomorphism type. This reduces the comparison to a fast lookup, and the overall algorithm runs in time $O(T \log T \cdot n^2 \log n)$, where T is the (large) number of triangulations, and n is the (small) number of simplices in each. The search parallelises well, since the bottlenecks are the hashing and performing local moves.

Thus far, the algorithm has detected 1 738 260 distinct minimal one-vertex triangulations of the $K3$ surface. The search is ongoing, and has neither exhausted the list of minimal triangulations nor connected $(K3)_{16}$ with $(K3)_{17}$. This enormous number of minimal triangulations is both interesting and surprising, offering a stark contrast to observations from dimension three.

4 Conclusion and future research

Proving Conjecture 1 would eliminate an important candidate for a pair of homeomorphic but non-diffeomorphic simply connected 4-manifolds. Moreover, as noted earlier, this work has a wider appeal: it shows for the first time how difficult problems of diffeomorphism and "exotic structures" in 4-manifold topology can be realistically tackled using computational tools.

Future developments will include: multiple-vertex triangulations containing fewer 4-simplices; a richer set of local modifications; and distributed algorithms for use on high-performance computing facilities.

References

[1] B. Benedetti and F. H. Lutz. Random Discrete Morse Theory and a New Library of Triangulations. `arXiv:1303.6422 [cs.CG]`, 2013, p. 33.

[2] A. Björner and F. H. Lutz. Simplicial manifolds, bistellar flips and a 16-vertex triangulation of the Poincaré homology 3-sphere. *Experiment. Math.*, 9(2):275–289, 2000.

[3] B. A. Burton. The Pachner graph and the simplification of 3-sphere triangulations. *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry ACM*, pp. 153 – 162, 2011.

[4] B. A. Burton. Computational topology with Regina: Algorithms, heuristics and implementations. `arXiv:1208.2504`, 2012. To appear in Geometry & Topology Down Under, Amer. Math. Soc.

[5] B. A. Burton, R. Budney, W. Pettersson, et al. Regina: normal surface and 3-manifold topology software. `http://regina.sourceforge.net/`, 1999–2012.

[6] M. Casella and W. Kühnel. A triangulated $K3$ surface with the minimum number of vertices. *Topology*, 40(4):753–772, 2001.

[7] F. Effenberger and J. Spreer. simpcomp - a GAP toolbox for simplicial complexes. *ACM Communications in Computer Algebra*, 44(4):186 – 189, 2010.

[8] M. Freedman. The topology of four-dimensional manifolds. *J. Differential Geom.*, 17:357–453, 1982.

[9] R. E. Gompf and A. I. Stipsicz. *4-manifolds and Kirby calculus*, volume 20 of *Graduate Studies in Mathematics*. AMS, Providence, RI, 1999.

[10] A. Mijatović. Simplifying triangulations of S^3. *Pacific J. Math.*, 208(2):291–324, 2003.

[11] U. Pachner. Konstruktionsmethoden und das kombinatorische Homöomorphieproblem für Triangulierungen kompakter semilinearer Mannigfaltigkeiten. *Abh. Math. Sem. Uni. Hamburg*, 57:69–86, 1987.

[12] J. Spreer and W. Kühnel. Combinatorial properties of the K3 surface: Simplicial blowups and slicings. *Experiment. Math.*, 20(2):201–216, 2011.

Lower bounds for approximate nearest neighbor on the Bregman divergences.

Amirali Abdullah
University of Utah

Suresh Venkatasubramanian
University of Utah

Abstract

Bregman divergences are important distance measures that are used extensively in data-driven applications such as computer vision, text mining, and speech processing, and are a key focus of interest in machine learning. As an example, the Kullback-Liebler divergence is a natural distance measure between probability distributions widely used in statistics and image processing. The Itakura-Saito is also commonly used in signal analysis and distance between speech samples. Answering *approximate nearest neighbor* (ANN) queries under these measures is very important in these applications and has been the subject of extensive study. In the context of Euclidean geometry, there have been lower bounds demonstrated for the high-dimensional setting, and also on Locality Sensitive Hashing (LSH) techniques. We aim to extend these tools to a broad subclass of Bregman divergences to obtain at least matching lower bounds, and in some cases stronger results.

1 Introduction

Bregman Divergences. Let $\phi : M \subset \mathbb{R}^d \to \mathbb{R}$ be a *strictly convex* function that is differentiable in the relative interior of M. The *Bregman divergence* D_ϕ is defined as

$$D_\phi(x,y) = \phi(x) - \phi(y) - \langle \nabla\phi(y), x-y \rangle.$$

More intuitively, as we can see in Figure 1, D_ϕ is defined by a lifting map upto a convex function ϕ where $D_\phi(p,q)$ is the vertical distance of $\phi(p)$ from the estimate obtained by linearization of ϕ at q. This is just the Euclidean distance when ϕ is the paraboloid, and thus we can view D_ϕ as in some sense a natural generalization of the lifting map commonly used in Euclidean discrete geometry for analysis of Voronoi diagrams or convex hull algorithms.

An important subclass of Bregman divergences are the *uniform* Bregman divergences. Suppose ϕ has domain $M = \prod_{i=1}^d M_i$ and can be written as $\phi(x) = \sum_{i=1}^d \phi_{\mathbb{R}}(x_i)$, where $\phi_{\mathbb{R}} : M_i \subset \mathbb{R} \to \mathbb{R}$ is also strictly convex and differentiable in relint(S_i) and each . Then $D_\phi(x,y) = \sum_{i=1}^d D_{\phi_{\mathbb{R}}}(x_i,y_i)$ is a *uniform* Bregman divergence. Many commonly used Bregman divergences are uniform, including the Kullback-Leibler, Itakura-Saito, Exponential distance and Bit entropy.

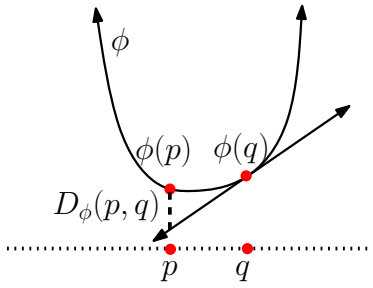

Figure 1: Illustrating the lifting map

1.1 Background and Overview of our work

The standard results for lower bounds on LSH (at least the earlier work by Panigrahy *et al.* [2] and Motwani *et al.* [1]) leverage heavily the combinatorial properties of the Hamming cube under the l_1 distance. Very roughly, they show that for any hashing function or partition of the cube into cells of tables, the data structure is highly sensitive to noise perturbation. If we obtain a query point by perturbation of our dataset and then random selection, then the query point has low probability of falling in the same hash bucket as it's nearest neighbor(or even approximate nearest neighbor). This means that any lookup algorithm which uses tables must touch a large number of cells (in expectation) to successfully retrieve an approximate nearest neighbour.

Our first simpler result gives a distance-preserving geometric isomorphism of the Hamming cube to a constructed set of points under any *uniform* Bregman divergence which we call a *pseudo-Hamming-cube*. The existence of such a set is not immediately obvious as the Hamming cube is inherently symmetric and distances are defined under l_1, while Bregman divergences D_ϕ are not symmetric, but we find a way to circumvent this difficulty.

Our second line of attack and subject of ongoing work, is to directly analyze a combinatorial object we introduce and stipulate as the *Bregman* cube. Here our distances are no longer symmetric as $D_\phi(a,b) \neq D_\phi(b,a)$. Indeed, we aim to exploit this asymmetry (which we capture by a parameter μ) to yield us stronger lower bounds than the Euclidean case for l_1. This requires defining an asymmetric noise perturbation according to the distribution of unequal distances across edges of our Bregman cube.

This is an abstract of a presentation given at CG:YRF 2013. 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 Reduction from the Hamming cube to an isomorphic set of Bregman points

First we define precisely the Hamming cube. The *Hamming cube* $(\{0,1\}^d, ||\cdot||_1)$ is the set of 2^d vertices of a unit length hypercube in $\mathbb{R}^d$ under the l_1 distance. We define now a combinatorial structure isomorphic to the Hamming cube, which we term the *pseudo-Hamming-cube*. Given a $D_\phi : \mathbb{R}^{2d} \times \mathbb{R}^{2d} \to \mathbb{R}$, a pseudo-Hamming-cube $C_\phi \subset \mathbb{R}^{2d}$ is a set of 2^d points such that there exists a bijection $f_\phi : \{0,1\}^d \to C_\phi$ and a fixed constant $c_0 \in \mathbb{R}$ so that $D_\phi(f_\phi(x), f_\phi(y)) = c_0||x-y||_1, \forall x,y \in \{0,1\}^d$.
We note that that c_0 can be set to 1 by scaling ϕ appropriately. We now show the existence of a *pseudo-Hamming-cube* for any uniform D_ϕ.

Lemma 1 *For any uniform $D_\phi : \mathbb{R}^{2d} \times \mathbb{R}^{2d} \to \mathbb{R}$, there exists a pseudo-cube $C_\phi \subset \mathbb{R}^{2d}$ and a suitable constant $c_0 \in \mathbb{R}$.*

Proof. Recall that for uniform D_ϕ, we have that $\phi(x) = \sum_{i=1}^d \phi_{\mathbb{R}}(x_i)$. Pick arbitrary a,b in the domain of $\phi_{\mathbb{R}}$ and hence obtain the two ordered pairs (a,b) and (b,a) in $\mathbb{R}^2$. Now stipulate $C_\phi = \{x_1 \times x_2 \ldots x_d \text{ s.t } x_i \in \{(a,b),(b,a)\}\} \subset \mathbb{R}^{2d}$. Note that $|C_\phi = 2^d|$. We define the isomorphism $h_\phi : \{0,1\}^d \to C_\phi$ by using the helper function $\bar{h} : \{0,1\} \to \{(a,b),(b,a)\}$ where $\bar{h}(0) = (a,b)$ and $\bar{h}(1) = (b,a)$. Now we state for $x \in \{0,1\}^d$, $h_\phi(x) = \bar{h}(x_1) \times \bar{h}(x_2) \times \ldots \bar{h}(x_d)$. The insight is that for any two $x,y \in C_\phi$, the ith component of D_ϕ on $\bar{h}(x_i)$ and $\bar{h}(y_i)$ is symmetrized, as

$$\begin{aligned} D_{\phi_{\mathbb{R}} \times \phi_{\mathbb{R}}}((a,b),(b,a)) &= D_{\phi_{\mathbb{R}} \times \phi_{\mathbb{R}}}((b,a),(a,b)) \\ &= D_{\phi_{\mathbb{R}}}(b,a) + D_{\phi_{\mathbb{R}}}(a,b). \end{aligned}$$

Direct computation shows that $\forall x,y \in \{0,1\}^d$, we have:

$$D_\phi(h_\phi(x), h_\phi(y)) = \left(D_{\phi_{\mathbb{R}}}(b,a) + D_{\phi_{\mathbb{R}}}(a,b)\right) ||x-y||_1.$$

This completes our proof, with $c_0 = D_{\phi_{\mathbb{R}}}(b,a) + D_{\phi_{\mathbb{R}}}(a,b)$. □

Now that we have defined a distance preserving isomorphism between the l_1 Hamming cube in $\mathbb{R}^d$ with our D_ϕ pseudo-Hamming-cube in $\mathbb{R}^d$, it is not hard to show that LSH and ANN lower bounds for the l_1 Hamming cube now transfer over to D_ϕ. In particular, the results from [2] in combination with our isomorphism imply:

Theorem 2 *For asymmetric and uniform Bregman divergences in $\mathbb{R}^{2d}$ (where d is $\Omega(\log n)$), any t-probe randomized data structure that returns an approximate nearest neighbor within a c multiplicative factor and with constant success probability must use space $n^{1+\Omega(\frac{1}{ct})}$.*

3 Asymmetric Bregman cube analysis

Our second line of attack and the heart of this paper is obtaining improved bounds over the l_1 analysis by employing the asymmetry in our favor. We first introduce a new structure, *the asymmetric Bregman* cube $C_\phi \subset \mathbb{R}^d$ or simply Bregman cube. We let D_ϕ be a uniform Bregman divergence, and pick $a,b \in \mathbb{R}$ such that $D_{\phi_{\mathbb{R}}}(a,b) = 1$. Assume that $D_{\phi_{\mathbb{R}}}(b,a) = \mu$, where $\mu > 1$ is a parameter dependent on $\phi_{\mathbb{R}}$ and the choice of a and b. Then C_ϕ is defined as $(\{a,b\}^d, D_\phi)$, i.e, the 2^d vertices of a cube along with the distance D_ϕ. Intuitively, this is combinatorially equivalent to a regular Hamming cube but now where distances of 1 and μ are associated with flipping a bit from 0 to 1 and 1 to 0 respectively.

To handle perturbation in this asymmetric setting, we define a random variable $\nu_{p_1,p_2} : \{0,1\}^d \to \{0,1\}^d$, such that $\nu_{p_1,p_2}(x)$ is obtained by by flipping each 0 bit of x to 1 with probability p_1 and each 1 bit of x to 0 with probability p_2. We can extend this to a perturbation operator R_{p_1,p_2} on boolean functions such that $R_{p_1,p_2} f(x) = E_{\nu_{p_1,p_2}}[f(\nu_{p_1,p_2}(x))]$.

R_{p_1,p_2} is a generalization of the regular symmetric noise operator which is defined as $T_p = R_{p,p}$ and is well studied in Fourier analysis. The key to many of the lower bounds obtained for the l_1 cube revolve around the fact that $\nu_{p,p}$ applied on any point in the Hamming cube displaces it a distance of pd in expectation, and that T_p scatters/smooths out subsets of the cube. This property is known as *hypercontractivity* and can be seen more precisely as a contraction under a certain norm. Let 1_S be the characteristic function of a subset of the Hamming cube, $1_S : \{0,1\}^d \to \{0,1\}$. Then it is well known that:

$$||T_p 1_S||_2 \leq ||1_S||_{1+(1-2p)^2} \tag{1}$$

Matters are different for an asymmetric cube C_ϕ with 0 to 1 bit-distance of 1 and 1 to 0 bit-distance of μ. $\nu_{p,p}$ no longer would displace every point on our cube by an equal amount in expectation, due to the asymmetry. We can prove however, that $\nu_{p_1,\frac{p_1}{\mu}}$ *does* have this property and displaces each point by $p_1 d$ distance in expectation. Problematically though, useful bounds for hypercontractivity on all subsets do not hold universally for R_{p_1,p_2} due to certain worst case scenarios, and we are forced to analyze hypercontractivity on partitions of the cube under conditions specific to our problem. This is the central technical challenge to our work going forward.

References

[1] Motwani, R., Naor, A., and Panigrahy, R. Lower bounds on locality sensitive hashing. *SIAM Journal on Discrete Mathematics 21*, 4 (2007), 930–935.

[2] Panigrahy, R., Talwar, K., and Wieder, U. A geometric approach to lower bounds for approximate near-neighbor search and partial match. In *Foundations of Computer Science, 2008. FOCS '08. IEEE 49th Annual IEEE Symposium on* (oct. 2008), pp. 414 –423.

Fault-Tolerant Clustering Revisited*

Nirman Kumar† Benjamin Raichel‡

1 Introduction

Two of the most common clustering problems are k-center and k-median clustering. In both these problems, the goal is to find the minimum cost partition of a given point set P in some metric space M, into k clusters. Each cluster is defined by a point in the set of cluster ***centers***, $C \subseteq \mathsf{P}$, where $|C| = k$. In k-center clustering, the cost is the maximum distance of a point to its assigned cluster center, and in k-median clustering, the cost is the sum of distances of points to their assigned cluster center. In both cases, given a set of cluster centers C, the optimal strategy is to assign a point to its closest center in C. In the fault-tolerant versions of these problems, we assign each point to its lth closest center in C, where $1 \leq l \leq k$ is a given integer. Intuitively, each center is prone to failure, and the cost of clustering is the worst-case cost, when up to $l-1$ failures among the closest centers for any point can be tolerated. The problem is to compute the optimal set of centers C, which minimizes the cost. The fault-tolerant k-center problem was first studied by Krumke [4], who gave a 4-approximation algorithm for this problem. Chaudhuri *et al.* provided a 2-approximation algorithm for this problem [1], which is the best possible under standard complexity theoretic assumptions. In both these papers, the version considered, differs slightly from ours in that one only considers points which are not centers when computing the point that has the furthest distance to its lth closest center. Khuller *et al.* [3] later considered both versions of the k-center problem. Their first version is the same as ours, i.e. the cost is the maximum distance of any point (including centers) to its lth nearest center. They gave a 2-approximation when $l < 4$ and a 3-approximation otherwise. Their second version is the same as that of Krumke [4]. For this version, they provided a 2-approximation algorithm matching the result of Chaudhuri *et al.* [1].

For k-median clustering, a fault-tolerant version has been considered by Swamy and Shmoys [6]. However, our version is different from theirs.

Our Contribution. Our main contribution is in showing that there is a natural reduction of the fault-tolerant versions of k-center and k-median to the ordinary (non-fault-tolerant) versions. Though the approximation ratios of our algorithms are not optimal, the authors feel that the algorithms more than make up for this in their conceptual simplicity and practicality from an implementation standpoint. For both these problems, we can use any constant factor approximation algorithm for the ordinary version of the respective problem. However, using the algorithm by Gonzalez[2] for k-center gives a smaller approximation factor, and this result is presented in full in this writeup. Additionally, to the best of our knowledge, the fault-tolerant variant of k-median that we investigate, has not been considered before.

2 Problem Definition and Algorithms

We are given a set of n points $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$ in a metric space M. Let $\mathsf{d}(\mathsf{p}, \mathsf{p}')$ denote the distance between the points p and p'. Let $\mathsf{ball}(\mathsf{p}, x)$ denote the closed ball of radius x with center p. For a given point $\mathsf{p} \in \mathsf{P}$, a subset $S \subseteq \mathsf{P}$ and an integer $1 \leq i \leq |S|$, let $\mathsf{d}_i(\mathsf{p}, S)$ denote the minimum x such that $|\mathsf{ball}(\mathsf{p}, x) \cap S| \geq i$. We abbreviate $\mathsf{d}_1(\mathsf{p}, S)$ to $\mathsf{d}(\mathsf{p}, S)$. Let $\mathsf{nn}_i(\mathsf{p}, S)$ denote the ith nearest neighbor of p in S, i.e. the point in S such that $\mathsf{d}(\mathsf{p}, \mathsf{nn}_i(\mathsf{p}, S)) = \mathsf{d}_i(\mathsf{p}, S)$. Let $\mathsf{NN}_i(\mathsf{p}, S) = \cup_{j=1}^{i} \{\mathsf{nn}_j(\mathsf{p}, S)\}$ be the set of i nearest neighbors of p in S. For both the problems, in the first step we run an approximation algorithm for the non-fault-tolerant version of the problem for $m = \lfloor k/l \rfloor$ centers, and in the second step additional centers are added. In our analysis below for fault-tolerant k-center, we use the well known 2-approximation algorithm for m-center by Gonzalez [1] in the first step, which we denote by **kCenter**(P, m). Let **kMedian**(P, m) denote any constant factor approximation algorithm for m-median.

We first run the algorithm **kCenter**(P, m) (resp. **kMedian**(P, m)). Let $Q \subseteq \mathsf{P}$ denote the set of m centers output and let $Q = \{\mathsf{q}_1, \ldots, \mathsf{q}_m\}$. Let $C = \bigcup_{i=1}^{m} \mathsf{NN}_l(\mathsf{q}_i, \mathsf{P})$, i.e. we take the l nearest neighbors in P of each point q_i, for $i = 1, \ldots, m$. We only use this set C in the analysis. If however C has less than k points we can throw in $k - |C|$ additional points chosen arbitrarily from $\mathsf{P} \setminus C$, since adding additional centers can only decrease the cost of our solution. Let **fCenter**(l, k) and **fMedian**(l, k) denote these algorithms for the fault-tolerant k-center and k-median problems, respectively.

*See [5] for the full version.

†University of Illinois; nkumar5@illinois.edu; http://www.cs.uiuc.edu/~nkumar5/.

‡University of Illinois; raichel2@illinois.edu; http://www.cs.uiuc.edu/~raichel2/.

[1] For k iterations, add the point furthest from the current set of centers, to the set of centers, where the first center is arbitrary.

3 Results

3.1 Fault-tolerant k center

For an instance of the problem on a point set P, let $C^* = \{w_1, w_2, .., w_k\}$ be an optimal set of centers, and let r_{opt} be its cost (i.e. $r_{\mathsf{opt}} = \max_{\mathsf{p}\in\mathsf{P}} \mathsf{d}_l(\mathsf{p}, C^*)$). Similarly, let $C = \{c_1, \ldots, c_k\}$ be the set of centers returned by **fCenter**(l, k), and let r_{alg} be its cost. We assume $l|k$, i.e. $m = k/l$. Algorithm **fCenter**(l, k) first calls **kCenter**(P, m). Let $Q = \{q_1, \ldots, q_m\}$ be the returned set of centers, computed in that order by Gonzalez's algorithm. Let $r_i = \mathsf{d}(\mathsf{q}_i, Q_{i-1})$ for $2 \le i \le m$, where $Q_{i-1} = \{q_1, \ldots, q_{i-1}\}$. (The $m = 1$ case is easier.)

The following is easy to see, and is used in the correctness proof of Gonzalez's algorithm.

Lemma 1 *For $i \neq j$, $\mathsf{d}(\mathsf{q}_i, \mathsf{q}_j) \ge r_m$.*

Lemma 2 *For any q_i, $\mathsf{NN}_l(\mathsf{q}_i, C^*) \subseteq \mathsf{ball}(\mathsf{q}_i, r_{\mathsf{opt}})$ and $\mathsf{NN}_l(\mathsf{q}_i, \mathsf{P}) \subseteq \mathsf{ball}(\mathsf{q}_i, r_{\mathsf{opt}})$.*

Proof. The first claim follows since $\mathsf{q}_i \in \mathsf{P}$ and so $\mathsf{d}_l(\mathsf{q}_i, C^*) \le r_{\mathsf{opt}}$. As $C^* \subseteq \mathsf{P}$, the second claim follows. □

Lemma 3 *We have that, $r_{\mathsf{alg}} \le r_m + r_{\mathsf{opt}}$.*

Proof. As in Gonzalez's algorithm, we have $r_m = \max_{\mathsf{p}\in\mathsf{P}} \mathsf{d}(\mathsf{p}, Q_{m-1})$, and so $\mathsf{d}(\mathsf{p}, Q) \le r_m$ for any $\mathsf{p} \in \mathsf{P}$. Consider any point $\mathsf{p} \in \mathsf{P}$ and let $\mathsf{q} = \mathsf{nn}_1(\mathsf{p}, Q)$. By how **fCenter**$(l, k)$ is defined, $\mathsf{NN}_l(\mathsf{q}, \mathsf{P}) \subseteq C$. Now, $\mathsf{d}(\mathsf{p}, \mathsf{v}) \le \mathsf{d}(\mathsf{p}, \mathsf{q}) + \mathsf{d}(\mathsf{q}, \mathsf{nn}_l(\mathsf{q}, \mathsf{P}))$ for any $\mathsf{v} \in \mathsf{NN}_l(\mathsf{q}, \mathsf{P})$, by the triangle inequality. Now, $\mathsf{d}(\mathsf{p}, \mathsf{q}) + \mathsf{d}(\mathsf{q}, \mathsf{nn}_l(\mathsf{q}, \mathsf{P})) \le r_m + r_{\mathsf{opt}}$, by Lemma 2. As such, $|\mathsf{ball}(\mathsf{p}, r_m + r_{\mathsf{opt}}) \cap C| \ge l$, and so $r_{\mathsf{alg}} \le r_m + r_{\mathsf{opt}}$. □

The proof of the following lemma is easy and is present in the full version [5].

Lemma 4 *If $r_{\mathsf{alg}} > 3r_{\mathsf{opt}}$, then for any $1 \le i \neq j \le m$, $\mathsf{ball}(\mathsf{q}_i, r_{\mathsf{opt}})$ and $\mathsf{ball}(\mathsf{q}_j, r_{\mathsf{opt}})$ are disjoint and each contains at least l centers from C^*.*

Lemma 5 *We have that, $r_{\mathsf{alg}} \le 3r_{\mathsf{opt}}$.*

Proof. Suppose otherwise that $r_{\mathsf{alg}} > 3r_{\mathsf{opt}}$. By Lemma 4, for $i = 1, \ldots, m$, $|\mathsf{ball}(\mathsf{q}_i, r_{\mathsf{opt}}) \cap C^*| \ge l$, and for $1 \le i < j \le m$, $\mathsf{ball}(\mathsf{q}_i, r_{\mathsf{opt}}) \cap \mathsf{ball}(\mathsf{q}_j, r_{\mathsf{opt}}) = \emptyset$. Assign all points in $C^* \cap \mathsf{ball}(\mathsf{q}_i, r_{\mathsf{opt}})$ to q_i. Notice, q_i is the unique point from Q within distance r_{opt} for any point assigned to it. Now $|Q| = m = k/l$, and each point in Q gets at least l points of C^* assigned to it uniquely. As such, there are at least $ml = k$ points of C^* assigned to some point of Q. Since $|C^*| = k$, it follows that each center in C^* gets assigned to a center in Q within distance r_{opt}. For $\mathsf{p} \in \mathsf{P}$, let v be its closest center in C^*. Let q be v's center from Q in distance $\le r_{\mathsf{opt}}$. We have $\mathsf{d}(\mathsf{p}, \mathsf{q}) \le \mathsf{d}(\mathsf{p}, \mathsf{v}) + \mathsf{d}(\mathsf{v}, \mathsf{q}) \le r_{\mathsf{opt}} + r_{\mathsf{opt}} = 2r_{\mathsf{opt}}$, by the triangle inequality. Moreover, since for $1 \le j \le l$, $\mathsf{d}(\mathsf{q}, \mathsf{nn}_j(\mathsf{q}, \mathsf{P})) \le r_{\mathsf{opt}}$, we have, $\mathsf{d}(\mathsf{p}, \mathsf{nn}_j(\mathsf{q}, \mathsf{P})) \le \mathsf{d}(\mathsf{p}, \mathsf{q}) + \mathsf{d}(\mathsf{q}, \mathsf{nn}_j(\mathsf{q}, \mathsf{P})) \le 2r_{\mathsf{opt}} + r_{\mathsf{opt}} = 3r_{\mathsf{opt}}$. Since for $\mathsf{q} \in Q$, $\mathsf{NN}_l(\mathsf{q}, \mathsf{P}) \subseteq C$, this implies that for any point $\mathsf{p} \in \mathsf{P}$, we have $|\mathsf{ball}(\mathsf{p}, 3r_{\mathsf{opt}}) \cap C| \ge l$ and so $r_{\mathsf{alg}} \le 3r_{\mathsf{opt}}$, a contradiction. □

Theorem 6 *For a given point set P in a metric space M, **fCenter**(l, k) achieves a 3-approximation to the optimal solution of the fault-tolerant k-center problem when $l|k$, and a 4-approximation otherwise.*

Proof. The $l|k$ case follows from Lemma 5. The proof for the other case is similar, and can be found in [5]. □

3.2 Fault-tolerant k median

Theorem 7 *For a given point set P in a metric space M, the algorithm **fMedian**(l, k) achieves a $(1 + 4c)$-approximation to the optimal solution of the fault-tolerant k-median problem, where c is the approximation guarantee of the subroutine **kMedian**.*

In the full version [5] we show that this theorem is implied by the following lemma, which is used to bound the cost of the optimum non-fault-tolerant m-median solution in terms of the optimum fault-tolerant k-median solution. See [5] for the proof of the lemma.

Lemma 8 *For any point set P in a metric space M and an integer $1 \le l \le |\mathsf{P}|$, there exists $Q \subseteq \mathsf{P}$ such that (1) $|Q| \le |\mathsf{P}|/l$, and (2) $\forall \mathsf{p} \in \mathsf{P}$, $\mathsf{d}_1(\mathsf{p}, Q) \le 2\mathsf{d}_l(\mathsf{p}, \mathsf{P})$.*

This lemma can also be used to get a result for fault-tolerant k-center similar to Theorem 7. This gives a weaker approximation guarantee than that of Theorem 6, but has the advantage, that any constant factor approximation algorithm can be used for **kCenter**(P, m).

References

[1] S. Chaudhuri, N. Garg, and R. Ravi. The p-neighbor k-center problem. *Inf. Proc. Lett.*, 65(3):131–134, 1998.

[2] T. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoret. Comput. Sci.*, 38:293–306, 1985.

[3] S. Khuller, R. Pless, and Y. J. Sussmann. Fault tolerant k-center problems. *Theor. Comput. Sci.*, 242(1-2):237–245, 2000.

[4] S. O. Krumke. On a generalization of the p-center problem. *Inf. Proc. Lett.*, 56:67–71, 1995.

[5] N. Kumar and B. Raichel. Fault tolerant clustering revisited. `http://cs.uiuc.edu/~raichel2/lnnkc.pdf`, 2013.

[6] C. Swamy and D. B. Shmoys. Fault-tolerant facility location. In *Proc. 14th ACM-SIAM Sympos. Discrete Algorithms*, pages 735–736, 2003.

Discrete Morse Gradient Fields and Stable Matchings

João Paixão* Thomas Lewiner*

Abstract

Forman's discrete Morse theory for cell complexes essentially relies on a discrete analogue of gradient fields, defined as an acyclic matching between incident cells. One can easily extract from this object several pieces of topological and geometrical information of a cell complex such as homology or the Morse-Smale decomposition. Most constructions of this discrete gradient are based on some variant of greedy pairing of adjacent cells, similar to a steepest descent algorithm. In this work, an equivalent formulation in terms of stable matching is proposed to study the geometric accuracy of the construction. It further simplifies the proofs of previous results regarding the behavior of the gradient field and the position of critical points.

1 Introduction

Forman introduced a combinatorial Morse theory [3] to study the topology of discrete objects such as cell and simplicial complexes. In this theory, the main objects from the smooth theory such as a Morse function, the corresponding gradient field, its critical points, and its flow have a discrete analogue. In particular, the discrete gradient is essentially an alternating cycle-free matching on the Hasse diagram of the complex. The Morse-Smale complex can be seamlessly obtained from the discrete gradient field, by purely combinatorial algorithms [7, 2, 5] avoiding any numerical integration or differentiation. The main issue is how faithful the discrete gradient is to the geometry of a function sampled at the vertices of the complex. This geometric accuracy can be formulated as the position of the critical elements (i.e. the cells of the Morse-Smale decomposition) with respect to Banchoff's critical points, obtained by a simplex-wise interpolation of the function [1].

There are numerous constructions proposed for a geometric faithful discrete gradient field. Lewiner et al. [7, 2] propose a modified greedy weighted matching algorithm with weights based on the difference of the function values at incident cells. The algorithm checks if the matching would have an alternating cycle at each step before adding a new edge to the matching. The geometric accuracy of the algorithm was proven robust if the triangulation is subdivided *twice* [6]. The results rely on a complicated proof technique, making it difficult to generalize and refine the results. Gyulassy et al. [5] suggested that a priority-queue based algorithm avoids checking for cycles. Robins et al.[8] developed an algorithm which uses homotopy expansions to build the gradient on the lower star of each vertex, and the authors were able to prove a one–to–one correspondence between the critical cells and Banchoff's critical points.

In this work, we propose a construction of a discrete gradient field for triangulated surfaces using a stable matching algorithm [4]. This construction is equivalent to the greedy one [7] for the same weights. The main advantage is that stable matchings are characterized by a local stability condition, which simplifies the proofs for the geometric accuracy, and gives a better description of the constructed gradient.

2 Stable matching

Given a triangulated surface K, one defines its Hasse diagram H to be a directed graph in which the set of nodes of H is the set of simplexes of K and there is an arc (τ, σ) if and only if simplex τ is contained in simplex σ and $dim(\tau)+1 = dim(\sigma)$. A scalar function $f : K_0 \to \mathbb{R}$ sampled on the vertices of a triangulated surface, can be extended to all the simplexes $\sigma \in K$ as the mean of the values of f on the vertices of simplex σ. For the stable matching, the arcs of the Hasse diagram of K are sort by the weight function $\mathsf{W}(\tau, \sigma) = f(\sigma) - f(\tau)$.

Given a matching M on H, an arc $(\tau, \sigma) \notin M$ is unstable if there exists $\rho, \rho' \in H$ matched with τ and σ respectively such that:

$$\mathsf{W}(\tau, \sigma) > \mathsf{W}(\tau, \rho) \quad \text{and} \quad \mathsf{W}(\tau, \sigma) > \mathsf{W}(\rho', \sigma) \ .$$

If there is no unstable arc, then M is called stable matching [4]. A stable matching always exists since H is a bipartite graph. In this work, the weights are assumed to be distinct, therefore the stable matching is *unique* and can be found by the classical Gale–Shapley algorithm [4].

A matching M on H defines a discrete vector field $\mathcal{V}$, where $\mathcal{V}$-paths are alternating paths in M. A discrete vector field is a gradient [3] if M is alternating cycle-free. One of the advantages of using stable matchings is that they are characterized by the local stability condition which is used to prove the results of the next section.

*Department of Mathematics, Pontifícia Universidade Católica, jpaixao@mat.puc-rio.br, thomas@lewiner.org. This work is partially financed by CNPq, FAPERJ, PUC-Rio, and CAPES.

3 Results

Gradient paths In Morse theory, the gradient field is $-\nabla f$, and f is thus expected to decrease along the gradient path. Let $\mathcal{V}$ be the discrete vector field on a triangulated surface K, obtained from stable matching M, ordered by a scalar function f, as in the previous section. A $\mathcal{V}$–path is *decreasing* if the value of f at the first vertex of the path is greater than its value at the last vertex.

Theorem 1 *f is globally decreasing on each $\mathcal{V}$–path [6]. In addition $\mathcal{V}$ is a valid discrete gradient vector field, i.e. $\mathcal{V}$ has no closed $\mathcal{V}$–path.*

The proof idea is the following: if there exists a non-decreasing $\mathcal{V}$-path $\blacktriangleleft\ \rho\ \sigma\ \tau\ \rho'\ \blacktriangleright$ of length 2, then it is easy to show that (τ, σ) is unstable. Since the paths are non-decreasing, there is no alternating cycle in M.

Relationship to greedy construction When the edges weights are distinct, it is known that the greedy weighted matching is equivalent to the stable matching. Moreover, Theorem 1 ensures that the stability of the matching automatically avoids alternating cycle. Therefore the stable matching construction is equivalent to the greedy construction [7, 2]. The stable matching formulation leads to the following results regarding the geometric accuracy of the greedy construction.

Critical points The following results are concerned with the relation between Forman's critical vertices, edges, and triangles [3] (i.e. unmatched nodes of H) to the Banchoff minima, saddles, and maxima [1] respectively. They are proven solely with the stability condition. Due to space constraints, only the result for critical vertices is presented.

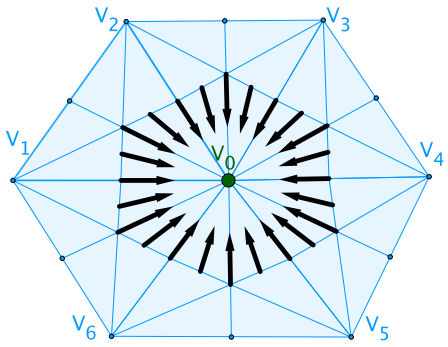

Figure 1: Critical vertex V_0 (unmatched green node) surrounded by the arcs (black arrows) in the stable matching on the barycentric subdivision (blue) of K. By Theorem 3, V_0 is a critical vertex if and only if $V_i > V_0$ for $1 \leq i \leq 6$.

Theorem 2 *A vertex v is a critical vertex of $\mathcal{V}$ if and only if v is minimum over its 2-neighborhood.*

Barycentric subdivision On a subdivision of the triangulated surface, our results about the relation of Forman's theory and Banchoff's theory are more precise. Let K' be the *first* barycentric subdivision of K, with scalar function f' linearly interpolated from f. Let $\mathcal{V}'$ be the discrete gradient on K', obtained from the stable matching ordered by f'.

Theorem 3 *A vertex v is a Banchoff minimum of K' if and only v is critical vertex of $\mathcal{V}'$. If a vertex v is a Banchoff saddle of K', then in $\mathcal{V}'$, there exists a critical edge in the lower star of v. If a vertex v is a Banchoff maximum, then in $\mathcal{V}'$, there exists a critical triangle in the lower star of v.*

This result improves upon previous results on the geometric accuracy of the greedy construction [6] since it requires one less barycentric subdivision and it gives a more precise location of the critical points. Moreover the proof is considerably simpler, which might yield a generalization to higher dimensions.

References

[1] T. Banchoff. Critical points and curvature for embedded polyhedra. *Journal of Differential Geometry*, 1:257–268, 1967.

[2] F. Cazals, F. Chazal, and T. Lewiner. Molecular shape analysis based upon Morse–Smale complex and the Connolly function. In *Symposium on Computational Geometry*, pages 351–360. ACM, 2003.

[3] R. Forman. A discrete Morse theory for cell complexes. In S. T. Yau, editor, *Geometry, Topology and Physics for Raoul Bott*, pages 112–115. International Press, 1995.

[4] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, pages 9–15, 1962.

[5] A. G. Gyulassy, P.-T. Bremer, B. Hamann, and V. Pascucci. A practical approach to Morse-smale Complex computation: Scalability and generality. *Transactions on Visualization and Computer Graphics*, 14(6):1619–1626, 2008.

[6] T. Lewiner. Critical sets in discrete Morse theories: relating Forman and piecewise-linear approaches. *Computer Aided Geometric Design*, 2013. to appear.

[7] T. Lewiner, H. Lopes, and G. Tavares. Applications of Forman's discrete Morse theory to topology visualization and mesh compression. *Transactions on Visualization and Computer Graphics*, 10(5):499–508, 2004.

[8] V. Robins, P. J. Wood, and A. P. Sheppard. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1646–1658, 2011.

On the Most Likely Convex Hull of Uncertain Points in the Plane

Hakan Yıldız *

Abstract

Consider a set of points in the plane, where the existence each point is uncertain and only known probabilistically. We study the problem of finding the *most likely convex hull*, i.e, the polygon with the highest probability of being the convex hull of the points. We present an $O(n^3)$ dynamic programming algorithm to solve the problem.

1 Introduction

Given a set of points $\mathcal{P}$ in the plane, the *convex hull* of $\mathcal{P}$ is a minimal convex polygon that contains all points in $\mathcal{P}$. Computing convex hulls is a well-studied problem and worst-case optimal algorithms exist for both its planar and high-dimensional forms [3].

In this work, we revisit the convex hull problem in an uncertainty setting where the existence of the input points is uncertain and only known probabilistically. In particular, we are given a point set $\mathcal{P}$ in the plane such that each point of $\mathcal{P}$ is known to exist with an independent probability. Under these conditions, we want to compute a convex polygon C such that the probability that C is the convex hull of the outcome of a probabilistic experiment is maximum. For simplicity, we refer to this polygon as the *most likely convex hull* of $\mathcal{P}$.

Our contribution is a dynamic programming algorithm to compute the most likely convex hull in $O(n^3)$ time and $O(n^2)$ space. We note that our method resembles some similar approaches previously proposed for various convex subset problems [2, 1].

2 Definitions

Let $\mathcal{P}$ be an uncertain set of n points. We represent $\mathcal{P}$ as a set of pairs $\{(s_1, \pi_1), \ldots, (s_n, \pi_n)\}$ where each s_i is a point in the plane and each π_i is a real number in the range $(0, 1]$. The point s_i is the *site* where the ith point of $\mathcal{P}$ may exist, and the number π_i is the associated probability of existence. In a probabilistic experiment, the ith point exists at s_i with probability π_i and is absent otherwise.

Let S be the set of all sites for the points in $\mathcal{P}$, i.e., $S = \{s_1, \ldots, s_n\}$. A subset $A \subseteq S$ is the outcome of a probabilistic experiment on $\mathcal{P}$ with probability:

$$\prod_{i \mid s_i \in A} \pi_i \times \prod_{i \mid s_i \notin A} \overline{\pi_i}$$

where $\overline{\pi_i}$ is the complementary probability $(1 - \pi_i)$. For ease of reference, we denote this probability by $\pi(A)$. Consider a convex polygon C. We define the *likeliness* of C, denoted $\mathcal{L}(C)$, to be the probability that C is the convex hull of the outcome of a probabilistic experiment on $\mathcal{P}$. We can write $\mathcal{L}(C)$ as

$$\mathcal{L}(C) = \Pr\big[\, \mathcal{CH}(A) = C \,\big] = \sum_{\substack{A \subseteq S \\ \mathcal{CH}(A) = C}} \pi(A)$$

where $\mathcal{CH}(A)$ is the convex hull of A. The *most likely convex hull* of $\mathcal{P}$ is the polygon C that maximizes $\mathcal{L}(C)$.

3 Algorithm

We now explain our algorithm to find the most likely convex hull. We begin with an observation. Let C be a convex polygon and V be the set of its vertices such that $V \subseteq S$. Let $S_{out} \subseteq S$ be the set of sites outside C. Observe that C is the convex hull of an experiment outcome $A \subseteq S$ if and only if A contains all points in V and no points in S_{out}. This implies the following relation, whose formal proof we omit from this abstract:

Lemma 1

$$\begin{aligned} \mathcal{L}(C) &= \Pr\big[\, V \subseteq A \ \wedge \ S_{out} \cap A = \emptyset \,\big] \\ &= \prod_{i \mid s_i \in V} \pi_i \times \prod_{i \mid s_i \in S_{out}} \overline{\pi_i} \end{aligned}$$

For each site $s_i \subseteq S$, we compute the most likely convex hull C such that s_i is the lowest vertex (along y-axis) of C. The hull with the maximum likeliness among these is the most likely convex hull of $\mathcal{P}$. We compute the most likely hull for a particular lowest vertex s_i as follows. Let s_j and s_k be two sites in S above s_i such that s_k is radially to the left of s_j with respect to s_i. Let G_j^k denote open region bounded by the segment $s_j s_k$ and the rays $\overrightarrow{s_i s_j}$ and $\overrightarrow{s_i s_k}$. (See Figure 1a.) We define the *contribution* of the directed edge $\vec{s_j s_k}$, denoted $\mathcal{C}(\vec{s_j s_k})$, as π_j times the product of the complementary probabilities of all sites in G_j^k. That is,

$$\mathcal{C}(\vec{s_j s_k}) = \pi_j \times \prod_{u \mid s_u \in G_j^k} \overline{\pi_u}$$

*Department of Computer Science, University of California Santa Barbara, hakan@cs.ucsb.edu

For a site s_j above s_i, let G_i^j to be the open region bounded by the downward ray extending from s_i and the ray $\overrightarrow{s_i s_j}$. We define G_j^i be the complementary region of G_i^j. (See Figure 1b.) We define the contribution of the directed edge $s_i\vec{s}_j$, denoted $\mathcal{C}(s_i\vec{s}_j)$, as π_i times the product of the complementary probabilities of all sites in G_i^j. Similarly, we define the contribution of the reverse edge $s_j\vec{s}_i$, denoted $\mathcal{C}(s_j\vec{s}_i)$, to be π_j times the complementary probabilities of all sites in G_j^i. The following lemma shows how edge contributions relate to the likeliness of a convex polygon.

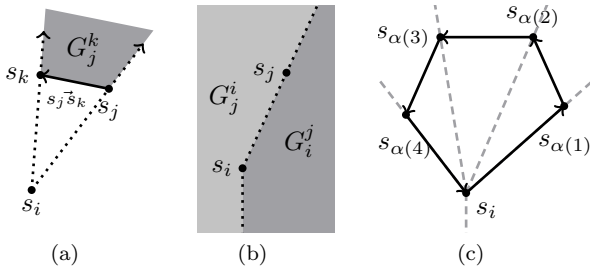

Figure 1

Lemma 2 *Let C be a convex polygon with vertices $s_i, s_{\alpha(1)}, \ldots, s_{\alpha(m)}$ in counter-clockwise order such that s_i is the lowest vertex of C. Then,*

$$\mathcal{L}(C) = \mathcal{C}(s_i\vec{s_{\alpha(1)}}) \times \mathcal{C}(s_{\alpha(1)}\vec{s}_{\alpha(2)}) \times \cdots \times \mathcal{C}(s_{\alpha(m-1)}\vec{s}_{\alpha(m)}) \times \mathcal{C}(s_{\alpha(m)}\vec{s}_i)$$

Proof. One can partition the space outside C into the regions $G_i^{\alpha(1)}, G_{\alpha(1)}^{\alpha(2)}, \ldots, G_{\alpha(m-1)}^{\alpha(m)}, G_{\alpha(m)}^i$ by drawing a downward ray from s_i and drawing rays $\overrightarrow{s_i s_{\alpha(j)}}$ for each $1 \le j \le m$. (See Figure 1c for an example.) Then, by Lemma 1, it is easy to see that the $\mathcal{L}(C)$ is the product of the contributions of the edges of C. □

We note that the contribution of each edge can be computed in constant time after an $O(n^2)$-time preprocessing. The idea in brief is to utilize a modified version of a triangle query structure by Eppstein et al. [4]. We omit the details of this structure from this manuscript.

Our dynamic programming strategy is as follows. For each directed edge $s_j\vec{s}_k$, we compute the convex chain starting at s_i and ending with edge $s_j\vec{s}_k$, such that the product of contributions of the edges participating in the chain is maximum. Let $\mathcal{T}(s_j\vec{s}_k)$ denote the product of contributions for this chain (and also the chain itself as an abuse of notation). By Lemma 2, the chain $\mathcal{T}(s_j\vec{s}_i)$ with the maximum value among all sites s_j is the most likely convex hull whose lowest vertex is s_i. The following recurrence is easy to observe and is the core of our dynamic programming algorithm:

$$\mathcal{T}(s_j\vec{s}_k) = \mathcal{C}(s_j\vec{s}_k) \times \max_{\substack{s \in S \\ s \text{ is radially right to } s_j \text{ w.r.t. } s_i \\ s \to s_j \to s_k \text{ is a left turn}}} \mathcal{T}(s\vec{s}_j)$$

We utilize the above recurrence to compute all $\mathcal{T}(\cdot)$ values as follows. We begin by setting $\mathcal{T}(s_i\vec{s}_j)$ to $\mathcal{C}(s_i\vec{s}_j)$ for all sites s_j. Then, we process all sites above s_i in right-to-left radial order (with respect to s_i). When we process a site s_j, we compute $\mathcal{T}(s_j\vec{s}_k)$ values for all sites s_k radially to the left of s_j. This can be done in $O(n)$ time as follows. Let L_j be the set of sites in S that are radially left to s_j plus s_i. Similarly, let R_j be the set of sites radially right to s_j plus s_i. Let $s_{\beta(1)}, \ldots, s_{\beta(m)}$ be the sites in R_j in clockwise order around s_j. For each site $s_{\beta(u)}$ in R_j, we define s_u^* be the site s among the sequence $s_{\beta(u)}, s_{\beta(u+1)}, \ldots, s_{\beta(m)}$ that maximizes $\mathcal{T}(s\vec{s}_j)$. The site s_u^* can be computed for all sites $s_{\beta(u)}$ in linear time by sweeping the sites in R_j in counter-clockwise order. (We compute this counter-clockwise order in a preprocessing step.)

For each site $s_{\beta(u)}$ in R_j, we set the value $\mathcal{T}(s_j\vec{s}_k)$ to $\mathcal{C}(s_j\vec{s}_k) \times \mathcal{T}(s_u^*\vec{s}_j)$ for all sites s_k in L_j inside the wedge bounded by the lines $\overleftrightarrow{s_{\beta(u-1)}s_j}$ and $\overleftrightarrow{s_{\beta(u)}s_j}$. Notice that the sites in this wedge are the sites that form a left turn when connected to $s_{\beta(u)}, s_{\beta(u+1)}, \ldots, s_{\beta(m)}$ through s_j (which is the condition in the recurrence relation). Finally, we note that by considering the sites $s_{\beta(u)}$ in radial order around s_j, we can locate each site in the wedge of interest in constant time.

The processing of a single point s_j takes $O(n)$ time, and thus we can find the most likely hull where s_i is the lowest vertex in $O(n^2)$ time. It follows that the most likely hull of $\mathcal{P}$ can be found in $O(n^3)$ time. The space cost is $O(n^2)$, dominated by the storage of the $\mathcal{T}(\cdot)$ values.

Theorem 3 *The most likely convex hull of n uncertain points in the plane can be computed in $O(n^3)$ time and $O(n^2)$ space.*

References

[1] C. Bautista-Santiago, J. M. Díaz-Báñez, D. Lara, P. Pérez-Lantero, J. Urrutia, and I. Ventura. Computing optimal islands. *Operations Research Letters*, 39(4):246–251, 2011.

[2] V. Chvátal and G. Klincsek. Finding largest convex subsets. *Congresus Numeratium*, 29:453–460, 1980.

[3] M. De Berg, O. Cheong, M. Van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 2008.

[4] D. Eppstein, M. Overmars, G. Rote, and G. Woeginger. Finding minimum area k-gons. *Discrete & Computational Geometry*, 7(1):45–58, 1992.

An arc-based integer programming formulation for proportional symbol maps

Rafael G. Cano* Cid C. de Souza* Pedro J. de Rezende* Tallys Yunes†

Abstract

Proportional symbol maps are a cartographic tool that employs scaled symbols to represent data associated with specific locations. The symbols we consider are opaque disks, which may be partially covered by other overlapping disks. We address the problem of creating a drawing that corresponds to stacking the disks, in sequence, in such a way as to maximize the total visible length of their boundaries. We propose a novel integer programming formulation and test it in on real-world instances with a branch-and-cut algorithm. When compared with state-of-the-art models from the literature, our model significantly reduces computation times.

1 Introduction

Proportional symbol maps are a cartographic tool to visualize data associated with specific locations (e.g., earthquake magnitudes and city populations). Symbols whose area is proportional to the numerical values they represent are placed at the locations where those values were collected. Although symbols can be of any geometric shape, opaque disks are the most frequently used and, for this reason, the focus of our study. Portions of a disk may not be visible when overlapping occurs and, when large portions of a disk are covered, it becomes difficult to deduce its size. Therefore, the way in which the disks are drawn affects the quality of a symbol map.

Let S be a set of n disks and $\mathcal{A}$ be an *arrangement*, which is the subdivision defined by the boundaries of the disks in S (see Fig. 1). We denote the set of arcs of $\mathcal{A}$ by R. A *drawing* of S is a subset of the arcs and vertices of $\mathcal{A}$ that is drawn on top of the filled interiors of the disks in S. We focus on *stacking drawings*, i.e., a drawing that corresponds to the disks being stacked up, in sequence, starting from the one at the bottom of the stack. According to Cabello et al. [1], two metrics can be considered to determine the quality of a drawing: the minimum visible boundary length of any disk and the total visible boundary length over all disks. The *Max-Min* and *Max-Total* problems consist in maximizing the former and the latter values, respectively.

*Institute of Computing, University of Campinas, Brazil, {rgcano,cid,rezende}@ic.unicamp.br. Supported by CNPq, FAPESP and FAEPEX/UNICAMP.

†School of Business Administration, University of Miami, USA, tallys@miami.edu.

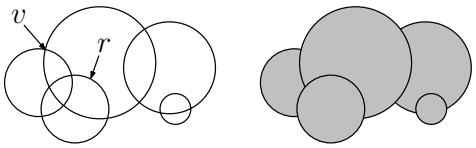

Figure 1: An arrangement $\mathcal{A}$ with vertex v and arc r (left), and a stacking drawing of the disks in $\mathcal{A}$ (right).

Cabello et al. [1] describe a greedy algorithm to solve the Max-Min problem in $O(n^2 \log n)$ time. Kunigami et al. [3] propose an integer linear programming (ILP) model to solve the Max-Total problem. Their model is based on two sets of binary variables: an arc variable x_r for each $r \in R$ (to indicate whether r is visible in the solution) and an ordering variable w_{ij} for each pair of disks $i, j \in S$ (to indicate the relative order between i and j). The computational complexity of the Max-Total problem for stacking drawings remains open.

We propose a novel ILP model for the Max-Total problem for stacking drawings (Sect. 2). Our formulations are in terms of arc variables only, thus reducing the dimension of the resulting polyhedra and the execution times. We prove that this model is a projection of the one described in [3] and also show how to make it stronger by lifting some of its inequalities (Sect. 3). We implement and test a branch-and-cut algorithm on real-world benchmark instances. A comparison with the formulations from [3] shows that our algorithm significantly improves computation times (Sect. 4).

2 Integer Linear Programming Model

Given an arc $r \in R$, we denote by ℓ_r the length of r, by d_r the disk whose boundary contains r and by S_r the set of disks that contain r in their interior. For each $r \in R$, we define a binary variable x_r that is equal to 1 if r is visible in the drawing, and equal to 0 otherwise. The objective is then to maximize $\sum_{r \in R} \ell_r x_r$.

Note that when an arc r is visible in a drawing, it induces an order between d_r and each disk in S_r (d_r must be drawn above each $i \in S_r$). We define an *induced order graph* $G_O = (V, E)$ as a directed multigraph with a vertex $v_i \in V$ for every disk $i \in S$, and a directed edge $e_{rj} = (v_{d_r}, v_j)$ for each arc $r \in R$ and disk $j \in S_r$. An example is shown in Fig. 2. Let $\mathcal{C}$ be the set of all cycles in G_O. Given a cycle $C \in \mathcal{C}$, let R_C be the set of arcs

that give rise to the edges of C, i.e., $R_C = \{r : e_{rj} \in C, \text{ for some } j \in S\}$. For each cycle $C \in \mathcal{C}$, we cannot have all arcs from R_C visible in a solution because they induce a cyclic order among the corresponding disks. Thus, the following constraints must be satisfied:

$$\sum_{r \in R_C} x_r \leq |C| - 1, \qquad \forall\, C \in \mathcal{C}. \tag{1}$$

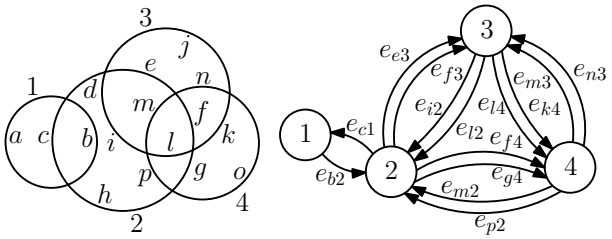

Figure 2: An arrangement with four disks (left) and its graph G_O (right).

3 Polyhedral Study

We refer to the model presented in Section 2 as *arc model* (AM) and to the one proposed in [3] as *graph orientation model* (GOM). Proofs are omitted due to space limitations.

Theorem 1 *The polyhedron defined by the linear relaxation of the AM model is the projection onto the x-space of the one associated with the GOM model.*

Let P be the polyhedron defined by the convex hull of all integer feasible solutions of the AM model. It follows from Theorem 1 and from the results presented in [3] that the dimension of P is $|R|$. Furthermore, given an arc $r \in R$, the inequality $x_r \geq 0$ always defines a facet of P, but $x_r \leq 1$ is facet-defining if and only if $S_r = \emptyset$.

Generally, inequalities (1) do not define facets of P, but they can be strengthened by lifting procedures, as shown in Fig. 3 (left). Arcs a, b and c form a cycle that yields the inequality $x_a + x_b + x_c \leq 2$. However, if arc d is visible, b and c must be covered by disk d_d. Thus, the inequality $x_a + x_b + x_c + x_d \leq 2$ is valid. Whereas the original AM model gives the same dual bounds as GOM (due to Theorem 1), the formulation obtained by lifting inequalities (1) (denoted by AM++) is stronger.

It should be noted that, even after lifting, the inequalities obtained might still not be facet-defining, as shown in Fig. 3 (right). Arcs $r_1, \ldots, r_4$ induce a 4-cycle in G_O and yield the inequality $\sum_{k=1}^{4} x_{r_k} \leq 3$, which can be lifted by adding $\sum_{k=1}^{4} x_{s_k}$. The resulting constraint does not define a facet of P because it is the sum of the valid inequalities $\sum_{k=1}^{4} x_{r_k} \leq 1$ and $\sum_{k=1}^{4} x_{s_k} \leq 2$. So far, it is unclear which conditions should be met by a cycle so that it leads to a strong lifted inequality.

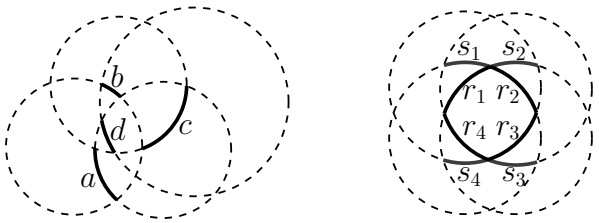

Figure 3: An arrangement that allows for lifting (left) and a lifted inequality that is not facet-defining (right).

4 Computational Results

We assess the effectiveness of our algorithm using a set of 28 instances[1] generated from data on the population of cities from several countries. Inequalities (1) (without lifting) are separated in polynomial time with a routine based on a procedure described by Grötschel et al. [2]. The program was coded in C++ and compiled with gcc 4.4.3. We used CGAL 3.5.1 to build the arrangements and CPLEX 12.4 to solve the integer programs. The experiments were run on an Intel Xeon X3430, 2.40GHz CPU with 8GB RAM.

Table 1 summarizes the results for the hardest instances, showing the number of disks and arcs in each one and the execution times, in seconds. Column *Speedup* shows the ratio between the execution times of GOM and AM++. Model AM++ performs better for all instances, especially for the hardest ones (reported here), for which it presents an average speedup of 23.6. The remaining instances are solved in less than two minutes. Still, AM++ achieves a speedup of 2.2.

Instance	$\lvert S\rvert$	$\lvert R\rvert$	GOM	AM	AM++	Speedup
Egypt	98	2033	99	19	8	12.4
France	135	3230	5039	435	62	81.3
Germany	150	2072	124	15	8	15.5
Greece	102	3482	5755	495	78	73.8
Japan	150	3544	249	17	13	19.2
Portugal	150	5070	304	39	30	10.1
USA (West)	87	3717	709	124	39	18.2

Table 1: Results for the seven hardest instances.

References

[1] S. Cabello, H. Haverkort, M. van Kreveld, and B. Speckmann. Algorithmic aspects of proportional symbol maps. *Algorithmica*, 58(3):543–565, 2010.

[2] M. Grötschel, M. Jünger, and G. Reinelt. On the acyclic subgraph polytope. *Math. Programming*, 33:28–42, 1985.

[3] G. Kunigami, C. C. de Souza, P. J. de Rezende, and T. Yunes. Generating optimal drawings of physically realizable symbol maps with integer programming. *The Visual Computer*, 28:1015–1026, 2012.

[1] `www.ic.unicamp.br/~cid/Problem-instances/Symbol-Maps`

Algorithms for tracking movement and guarding regions using radio tomographic imaging

Samira Daruki* Peter Hillyard† Neal Patwari‡ Suresh Venkatasubramanian§

Abstract

Radio Tomographic Imaging (RTI) is an emerging technology that locates moving objects in areas surrounded by simple and inexpensive radios. RTI is useful in emergencies, rescue operations, and security breaches, since the objects being tracked need not carry an electronic device. Tracking humans moving through a building, for example, could help firefighters save lives by locating victims quickly. RTI works by placing small inexpensive radios in a region of interest. The radios can send and receive wireless signals, and form a network of links that cover the region of interest. When a person walks through the region, they interfere with the links, creating a "shadow" of broken links that can be used to infer presence and track individuals.

This yields the following problem: given a collection of radios and a set of "visible" links, infer the trajectory of a person moving through the region. This inference must be robust under link errors and occlusion, as well as be performed in real time. In addition, there is an associated planning problem of where to place the radios in order to make the tracking algorithm as effective as possible.

In this paper, we present geometric algorithms for these questions. The key technical developments include a generalization of stabbing line and transversal problems.

1 Problem Definitions

We are given a collection of "sensors" represented as points in plane, in which pairs of sensors have a link between them. Someone moves in a straight line through the room. We are given a set of sensors which fire and we have to identify a candidate trajectory for the object's path.

Defnition 1 *Given a set of n sensors $P = \{p_1, ..., p_n\}$ in the region, the signature for input line ℓ is defined as the set of all the line segments $\overline{ij}$ connecting sensors p_i and p_j which ℓ hits. Clearly, the maximum length of signature for a line is $\binom{n}{2}$.*

Problem 1 *Given a region R in the plane, the set of n locations for sensors $\{p_1, ..., p_n\}$ and input signature $\sigma(\ell)$, we desire to find all lines ℓ^* with the same signature as input path ℓ.*

In fact, this is "reverse" line stabbing. In the traditional problem, you are given segments and a line and want to find which segments are hit. In the reverse version, given a collection of segments that are hit, we need to find the region of feasibility. To solve this, we use *duality*. The dual point to the line $\ell : y = ax + b$ is the point $\ell^* = (a, -b)$. Similarly, for a point $p = (c, d)$ its dual line is $p^* : y = cx - d$. Based on this transformation, the sensors described as points in primal space will be mapped to lines and line segments connecting sensors in primal space will be shown as double wedges in dual space. It yields a dual arrangement $\mathcal{A}(P)$ of complexity $O(n^2)$ in dual space. The first observation we make connects primal and dual signatures.

Observation 1 *Points inside some face of arrangement in dual space correspond to some set of lines in primal space with the same signature.*

This implies that we can compactly describe all valid lines by returning a list of dual cells.

2 Problem Hierarchy

Here we describe different versions of the problem and we show two different approaches for attacking the problem using the underlying graph and geometric structures. We try to give a brief sketch of the algorithms and basic ideas.

2.1 Case 0: Stabbing Line Segments

This is the basic case of our problem in which a set of m segments is given as input and the goal is to find all the possible lines which hit all of these segments. This can be done by finding the intersection of multiple double wedges corresponding to the link which are hit. There has been some previous work on this which solves the

*Department of Computer Science, University of Utah, daruki@cs.utah.edu

†Department of Electrical Engineering, University of Utah, peterhillyard@gmail.com

‡Department of Electrical Engineering, University of Utah, npatwari@ece.utah.edu

§Department of Computer Science, University of Utah, suresh@cs.utah.edu

problem in $m \log m$ time and linear space using some nice characterization of double wedges in dual space [2]. They exploit a key combinatorial property of the feasible space (which they call the stabbing region) in the dual that allows them to design a divide-and-conquer strategy that can merge two partial stabbing regions in linear time.

2.2 Case 1: Full Information, All Links

In this version, we assume all the links between pairs of nodes are "active". The input is the set of segments which are hit (called "blue") and the remaining links are called "red". The goal is to find all the lines which hit the blue links and avoid red ones. In this setting we have the following nice property:

Lemma 1 *If f, f' are two* bounded *cells of $\mathcal{A}(P)$, then $\sigma(f) \neq \sigma(f')$. If f is unbounded, then at most one other face has the same signature as f.*

Lemma 2 *(Nested signatures property) There are no two lines ℓ, ℓ' such that $\sigma^+(\ell) \subset \sigma^+(\ell')$.*

We have three different approaches to attack this problem:

Approach I Based on Lemma 2, since in the all links case the nested signatures property does *not* hold, this problem can be reduced to problem in [2] with just the input as blue links, and can be solved by the same divide and conquer approach.

Approach II The set of blue segments induce a bipartite subgraph on the network and results in a unique partitioning on the nodes. The resulting candidate region, can be interpreted as the intersection of upper and lower envelopes defined by the dual lines of the partitioned nodes.

Approach III The trivial solution is to enumerate and store all the signatures of cells in the dual arrangement and find the desired region based on the query signature, But it needs too much storage (n^2 cells, each with their own signature of length n^2). We can reduce the complexity by using hashing: Identify a canonical point $p(c)$ in each cell of the dual arrangement. There are only n^2 strings(one for each region). Hash each string into a table of size n^2 and store the canonical point associated with the region. When a query comes in, hash its signature and return the $O(1)$ sized list of potential cells. For each canonical (dual) point, check if it matches the signature. By Lemma 1 and 2 since all regions have a distinct signature (and no signatures are contained in others), we can return the desired bounded region.

2.3 Case 2: Full Information, Not All Links

In this case, we have the same blue-red setup as before, but not all sensor pairs are active (because of occlusion or other problems). Furthermore, Lemma 1 is no longer true: two different cells can have signatures that nest, and multiple cells can have the same signature. Furthermore, we do *not* have the nice characterization of dual structure as before to use the divide and conquer approach in [2].

Approach I If the graph is connected, the blue segments induce the bipartite characterization on the underlying network and approach will be the same as before in Case 1. If not, then blue segments will result in multiple partitionings on the network with size exponential in the number of components. Here the main challenge is that to design an algorithm that runs in linear space and polynomial time to compute the desired answer.

Approach II We use a generalized hashing strategy. The main problem is how to store compact IDs for cells so that we can identify matches. This can be done in the preprocessing step: Identify all the cells that have the same signatures. Then choose a single canonical point from each, hash the canonical points, also store the number of segments in the signature. When a query arrives, we probe the hash table as before, and eliminate all table entries that do not match the signature size. For the rest, we check if the point signature matches the query signature, and return that point if it does. Finally, we use the returned points to identify the group of faces to be returned.

2.4 Case 3: Online Tracking

All above approaches work when we are tracking offline, but what about if we want to track the object in real time ? The challenge here is that the query signature just determines the blue segments in each step and does *not* give complete information about the red set. We need to take more effort to compute and update the red set incrementally in each step using simplex range searching [1] on segments. More precisely, this case is a new version of the problem including blue-red-gray segments. Gray segments are the ones which are not hit till this step, but it might be hit in future. These segments cause incomplete information in the query signature and we cannot use the same hashing approach as we used for offline case anymore. We need to use a different strategy and design an efficient algorithm to do the online updates.

References

[1] B. Chazelle, M. Sharir, and E. Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica*, 8(1-6):407–429, 1992.

[2] H. Edelsbrunner, H. A. Maurer, F. P. Preparata, A. L. Rosenberg, E. Welzl, and D. Wood. Stabbing line segments. *BIT Numerical Mathematics*, 22(3):274–281, 1982.