Open Problems from CCCG 2015

Joseph O'Rourke*

The following is a description of the problems presented on August 11, 2015 at the open-problem session of the *27th Canadian Conference on Computational Geometry* held in Kingston, Ontario, Canada.

Largest cell in an arrangement
Joseph O'Rourke
Smith College
orourke@cs.smith.edu

Is it possible to find the largest-area bounded cell in an arrangement of n lines in $\mathbb{R}^2$ (Figure 1) in subquadratic time? Sariel Har-Peled observed that

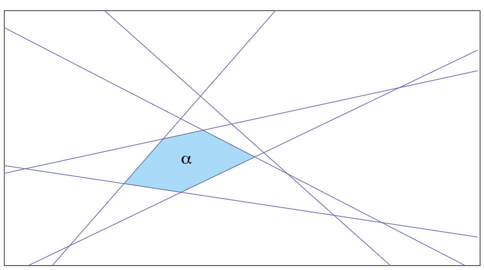

Figure 1: An arrangement and its largest cell area α.

if the largest cell's area α is much greater than its expected area $1/n^2$, then random sampling permits achieving expected subquadratic time.

I conjecture $\Omega(n^2)$ is a lower bound.

References

[OR15a] J. O'Rourke. Largest cell in an arrangement. http://cstheory.stackexchange.com/q/30802/337 13 March 2015.

Avoiding points on a sphere
Joseph O'Rourke
Smith College
orourke@cs.smith.edu

Let S be a unit-radius sphere in $\mathbb{R}^3$, and let P be a finite set of points contained in a hemisphere of S, viewed as a rigid pattern. Let R be any region/subset of S, not necessarily connected. Say that P *fits* in R if there is some placement of P such that each of its points is strictly interior to R. Say that R *avoids* P is no placement fits in R.

What is the largest area R that avoids a given set P?

An example is shown in Figure 2, with $|P| = 5$. The geodesic convex hull H avoids P, because the four points on the boundary of H are not strictly inside H. Similarly, the minimum enclosing disk D also avoids P for the same reason. But D is not the largest avoiding region for this particular P.

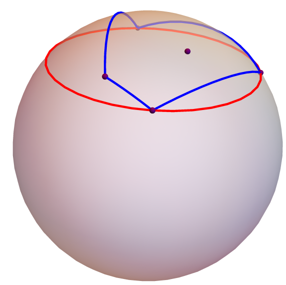

Figure 2: A 5-point set, its geodesic hull, and its minimum enclosing disk.

Updates. (1) Paz Carmi observed that if the diameter d of D is small relative to π (the length of the great circle arc between the poles), two copies of D centered at each pole avoid P, i.e., if $d < \pi/2$ those two pole-copies will be separated by a band of width $\geq d$. Even smaller d permit more copies of D.

(2) Let the surface area of the sphere be A. Alexandru Damain proved that, for an n-point set P, the area of an avoiding set cannot be larger than $(\frac{n-1}{n})A$. Suppose to the contrary that there is a region R whose measure A' exceeds $(\frac{n-1}{n})A$ and avoids P, with $|P| = n$. Rotate the set of points $P = \{p_1, \ldots, p_n\}$ randomly. Associate a random indicator variable X_i with each point p_i, with $X_i = 1$ if p_i lands in R after the random rotation, and $X_i = 0$ otherwise. We have that $\mathrm{E}[X_i] = \Pr[X_i = 1] > \frac{n-1}{n}$, because $A'/A > \frac{n-1}{n}$. Let X be the total number of vertices of P that lie

*Department of Computer Science, Smith College, Northampton, MA 01063, USA. orourke@cs.smith.edu

in R. Then by linearity of expectation,

$$\mathrm{E}[X] = \mathrm{E}[\sum_{i=1}^{n} X_i] = \sum_{i=1}^{n} \mathrm{E}[X_i] > n \cdot \frac{n-1}{n} = n-1 \ .$$

Because the expected number of points covered by a random rotation of P is greater than $n-1$, there must be at least one rotation of the points for which all n are covered. Thus P fits in R, and R does not avoid P.

For $n = 2$, let P be two antipodal points. Then for R a hemisphere of S, P cannot fit in R, as both points cannot be strictly interior to R. So R avoids P. But increasing R slightly allows P to fit in R. So here, the $(\frac{n-1}{n})A = (\frac{1}{2})A$ bound is tight.

References

[OR15b] J. O'Rourke. Regions on a sphere that avoid a fixed point set. http://mathoverflow.net/q/208230/6094 2 June 2015.

The point-set knot number

Joseph O'Rourke
Smith College
orourke@cs.smith.edu

For a knot K, define $z(K)$, the *point-set knot number*, to be the smallest number n such that every general-position point set S of n points may be connected to a simple (i.e., non-self-intersecting) polygonal knot P equivalent to K. S must be in general position in the sense that no 3 points are collinear, and no 4 points coplanar. The vertices of the polygon P must be exactly the points in S. Let $z(K) = \infty$ if there is no such n.

For example, $z(K_{0_1}) = 3$, for K_{0_1} the unknot. Let $p_1, p_2, \ldots, p_n$ be the points of S sorted top to bottom, and let $P = (p_1, p_2, \ldots, p_n, p_1)$, connecting the points in vertical order and closing with the segment $p_n p_1$. It is clear that no self-intersections can occur with the vertically sorted connections $(p_1, p_2, \ldots, p_n)$. So only the last segment $s = p_n p_1$ needs to be checked. If s passes through an intermediate vertex, then 3 points of S are collinear. If s passes through an edge interior point, then 4 points of S are coplanar, as shown in Figure 3. Thus P is simple. It should also be clear that P is an unknot.

Obviously $z(K)$ is at least the stick number of K. But is $z(K)$ finite for every K? In particular, what is $z(K_{3_1})$, for the trefoil knot K_{3_1}?

Update. Günter Ziegler informed me that my question was answered positively by Negami [Neg91]. "His first main observation was that for any sufficiently large n, by a Ramsey Theory argument, every set of n points in general position in $\mathbb{R}^3$ contains n points on an order-n curve (without loss of generality: the moment curve);. The second part of the argument shows that every knot occurs on a sufficiently large set of points on the moment curve. So $z(K)$ is finite for all knots K.

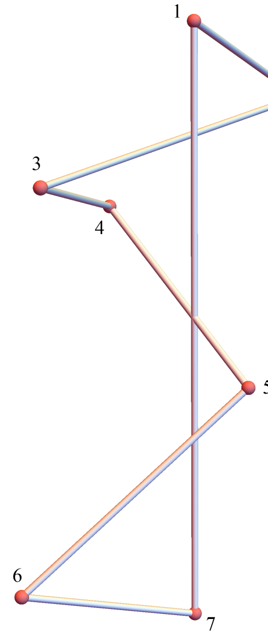

Figure 3: Edge $(7, 1)$ crosses $(4, 5)$.

An answer for the trefoil knot was given by Alfonsin [Alf99]: Any set of 7 points in general position contains a trefoil or its mirror image."

References

[Alf99] Jorge Alfonsin Spacial graphs and oriented matroids: the trefoil. *Discrete Comput. Geom.* 22 (1999), 149-158.

[Neg91] Seiya Negami. Ramsey theorems for knots, links and spacial graphs. *Trans. Amer. Math. Soc.* (2)324 (1991), 527–541.

[OR15c] J. O'Rourke. Can every large point set be connected to a given knot?. http://mathoverflow.net/q/204436/6094 1 May 2015.

Compact source unfoldings of convex polyhedra

Anna Lubiw
Univ. Waterloo
alubiw@uwaterloo.ca

Given a convex polyhedron, find an unfolding that has the minimum enclosing circle. One contender might be the *center source unfolding* which we define to be the source unfolding from the point c that is the *center* of the polyhedron's surface, i.e., c is the point that minimizes $\max d(p, c)$ as p ranges over all points on the surface of the polyhedron, and distance d is measured on the surface of the polyhedron. For background on source unfolding, see [DO07]. The radius of the minimum enclosing circle of the center source unfolding is the radius R of the polyhedron's surface, i.e., the value $\max d(p, c)$.

Secondary question: Is there an efficient algorithm to find point c? The algorithm to find the diameter D of a convex polyhedron [AAOS97] is surely relevant. Note that $D/2 \leq R \leq D$. Tight examples for the two extremes are a cigar-shaped ellipsoid and a sphere, respectively.

For a unit cube, $R = D = 2$, and a center source unfolding is shown in Figure 4. This is probably the unfolding of the cube that has the minimum enclosing circle. The same unfolding yields the minimum sized square needed to wrap a unit cube [P14].

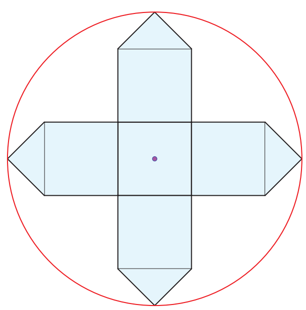

Figure 4: Source unfolding of a cube from the center of a face.

References

[AAOS97] Pankaj K. Agarwal, Boris Aronov, Joseph O'Rourke and Catherine A. Schevon. Star unfolding of a polytope with applications. *SIAM Journal on Computing*, 26:1689–1713, 1997.

[DO07] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, New York, 2007.

[P14] Qinxuan Pan. Uniqueness of Optimal Cube Wrapping. *Proc. 26th Canadian Conference on Computational Geometry*, Halifax, 2014.

Core set for median

Michael Kerber
Max Planck Institute for Informatics
mkerber@mpi-inf.mpg.de

Let P be a set of n points in $\mathbb{R}^d$. The *median* opt for P is the point in $\mathbb{R}^d$ that minimizes the sum of the distances to P:

$$\text{opt} := \text{argmin}_{q \in \mathbb{R}^d} \sum_{p \in P} d(q, p)$$

where $d(\cdot, \cdot)$ is the Euclidean distance function. We call a point q' an *ϵ-approximate median*, if its sum of distances to P is at most $(1+\epsilon)$ the optimal sum:

$$\sum_{p \in P} d(q', p) \leq (1 + \epsilon) \sum_{p \in P} d(\text{opt}, p)$$

We call $E \subseteq P$ an *ϵ-coreset*, if the affine subspace spanned by the points in E contains an ϵ-approximate median. It is known that coresets of constants size exists: Precisely, any point set P has an ϵ-coreset of size $O(1/\epsilon \log 1/\epsilon)$, regardless of the number of points or the ambient dimension. See [KR15] for the details, which are based on a result by Shyamalkumar and Varadarajan [SV12].

Question: Does every point set permit an ϵ-coresets of size $O(1/\epsilon)$?

This is the best bound one can hope for, because points sets exist for which any coreset must be of size $\Omega(1/\epsilon)$. Moreover, it is known that corests of size $O(1/\epsilon)$ exist for the related problems of approximating the *mean* of the point set (minimizing the sum of squared distances) and the *center* of the point set (where the sum of distances is replaced by the maximal distance).

References

[KR15] Michael Kerber and Sharath Raghvendra Approximating and Streaming Algorithms for Projective Clustering via Random Projections arxiv:1407.2063 May 2015

[SV12] N. Shyamalkumar and K. Varadarajan Efficient Subspace Approximation Algorithms Discrete & Computational Geometry 47, pp.44–63 2012

A puzzle with convex sets

Michael Kerber
Max Planck Institute for Informatics
mkerber@mpi-inf.mpg.de

This puzzle was posed to me some weeks ago:

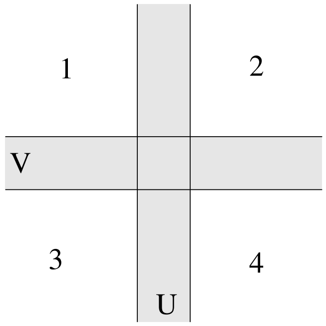

Figure 5: An inifinite vertical stripe U and an infinite horizontal stripe V split the plane into 4 regions.

Find two convex sets U and V in the plane such that $\mathbb{R}^2 \setminus U \cup V$ splits into 5 components.

It should be emphasized that no further assumptions beyond convexity on U and V are made (U and V need not to be finite, closed, or open). Four components are easy to achieve as shown in Figure 5. During the workshop, the puzzle was solved by Don Sheehy and Nicholas Cavanna. In order to not spoil the nice solution, the explanation is not given here.

Question 1: Can two sets split the plane into 6 (or more) component?

Update. One day after I posed the problem, Sang Woo Bae contacted me with a proof outline. It is based on the same idea that Don, Nicolas and myself also had in mind: Show that at most one component can be bounded (otherwise, at least one of sets is not convex), and show that at most 4 components can be unbounded. It seems that the answer to Question 1 is therefore negative.

Question 2: In how many components can two convex sets split $\mathbb{R}^d$?

As Don pointed out, $2d + 1$ sets are possible by extending the 5-split example in $\mathbb{R}^2$. Is this the best possible?

Realizing trees with farthest-point Voronoi diagrams

Therese Biedl
Univ. Waterloo
biedl@uwaterloo.ca

Let T be a tree. Is there a set of points P such that T corresponds to the farthest-point Voronoi diagram of P?

Tree T could be given in three different ways:

1. T could be a *geometric tree*, i.e., be given with coordinates for the nodes of the tree and with the edges drawn straight-line between them.
2. T could be an *ordered tree*, i.e., it comes with a fixed order of arcs around each node, and in the farthest-point Voronoi diagram these orders must be respected.
3. T could be an *abstract tree*, i.e., with nodes and arcs but no further information.

Related work: The above question is completely answered for Voronoi diagrams (see [H92, BHH13] for setting (1) and [LM03] for settings (2,3)) and for Straight Skeletons (see [BHH13] for setting (1) and [ACD+12] for settings (2,3)). In a nutshell, in setting (2) (and therefore also (3)) *any* tree can be represented, and in setting (1) there exists a polytime algorithm to test whether a given tree can be represented.

Progress at CCCG: Multiple CCCG participants discovered that the situation is the same for farthest-point Voronoi diagrams: Any ordered tree can be represented (even by points in convex position), and for a given geometric tree we can test in polynomial time, using linear programming, whether it can be represented. We are now in the process of working out the details and writing up the results.

References

[ACD+12] Oswin Aichholzer, Howard Cheng, Satyan L. Devadoss, Thomas Hackl, Stefan Huber, Brian Li, and Andrej Risteski. What makes a tree a straight skeleton? In *Proc. 24th Canadian Conference on Computational Geometry*, 2012, Charlottetown, pages 253–258, 2012.

[BHH13] T. Biedl, M. Held, and S. Huber. Recognizing Straight Skeletons and Voronoi Diagrams and Reconstructing Their Input. In *10th Internat. Symposium on Voronoi Diagrams in Science and Engineering*, pages 37–46. IEEE, 2013.

[H92] D. Hartvigsen. Recognizing Voronoi Diagrams with Linear Programing. *ORSA J. Computing* 4(4):369–374, 1992.

[LM03] Giuseppe Liotta and Henk Meijer. Voronoi drawings of trees. *Comput. Geom.*, 24(3):147–178, 2003.

Flipping open problems
Prosenjit Bose
Carleton University
jit@scs.carleton.ca

In Jit's Ferran Hurtado Memorial lecture, he posed more than 18 open problems. Here we mention a theme that ran throughout his presentation: finding flip-sequences sensitive to the difference between the start and end triangulations. See [BH09] for more information.

1. Given two triangulations, is it possible to determine the minimum number of flips to convert one into the other?
2. Can we find an approximation to the minimum number of flips or a sequence that is sensitive to the minimum number somehow? E.g., if k is the minimum number of flips, can we find a sequence of $f(k) = 2^k$ flips?
3. Can one compute a set of simultaneous flips that converts one triangulation into another that is sensitive to the minimum number of simultaneous flips required?

Update. Just after the conference [F15]: "the diameter of the flip graph is at least $\frac{7n}{3} + \Theta(1)$, improving upon the previous $2n + \Theta(1)$ lower bound."

References

[BH09] Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Computational Geometry.* Volume 42, Issue 1, 2009, pages 60–80.

[F15] Fabrizio Frati. A Lower Bound on the Diameter of the Flip Graph. arXiv:1508.03473 August 2015.