

Open Problems from CCCG 2009

Erik D. Demaine*

Joseph O'Rourke†

The following is a list of the problems presented on August 17, 2009 at the open-problem session of the *21st Canadian Conference on Computational Geometry* held in Vancouver, British Columbia, Canada.

Simply Developing Geodesic Polygons

Joseph O'Rourke

Smith College

orourke@cs.smith.edu

Let \mathcal{P} be a simple piecewise-geodesic curve on a sphere, i.e., a curve composed of a finite sequence of arcs of great circles that do not intersect except at shared endpoints. A *development* of \mathcal{P} is a polygonal chain in the plane whose segments have the same lengths as the arcs of \mathcal{P} , and whose angles correspond to the angles between the arcs on the sphere. A development can be viewed as created by rolling the sphere on the plane, starting at one end of \mathcal{P} , maintaining \mathcal{P} in contact with the plane at all times, and rolling without twisting or slippage until reaching the other endpoint.

The goal is to establish that certain classes of (closed) geodesic polygons always develop without intersection, regardless of the starting point for the development. The only class known to develop without intersection are geodesic convex polygons: those with vertex angles (to one consistent side) at most π [OS89].

During the presentation of the problem, the poser illustrated an example of a simple open curve that develops to a crossing chain, and Jack Snoeyink observed that closing it to a simple geodesic polygon \mathcal{P} shows that not all star-shaped geodesic polygons develop without intersection. The example is shown in Figure 1. Every arc from a to a point inside \mathcal{P} is nowhere exterior to \mathcal{P} , so indeed this is a star-shaped geodesic polygon.

Any natural class of nonconvex geodesic polygons that always develop without intersection would be of interest. More ambitious would be to *characterize* those geodesic polygons that develop




Figure 1: This polygon on the sphere is star-shaped from a . When cut open at d , it develops to a crossing chain in the plane.

without intersection. More ambitiously still, generalize to geodesic polygons on convex polyhedra, or on smooth convex surfaces.

References

- [OS89] Joseph O'Rourke and Catherine Schevon. On the development of closed convex curves on 3-polytopes. *J. Geom.*, 13:152–157, 1989.

Random Graph Spanners

Joseph O'Rourke

Smith College

orourke@cs.smith.edu

Select n points uniformly at random from the unit square, and then form a graph $G = G(n, p)$

*MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar Street, Cambridge, MA 02139, USA, edemaine@mit.edu

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

by connecting points according to the Erdős–Rényi model, adding an edge between each pair of points with probability p . This is not a “random geometric graph” in the usual sense of that term, because points are connected without regard to their geometric distance. Rather, it is a random graph built on geometric points.

1. It is known that the threshold for the creation of a “giant component” $C \subseteq G$ is $p_1 = 1/n$. Is C (almost surely) a spanner for the points it connects, for $p = p_1 + \varepsilon$? My conjecture: NO.
2. It is known that the threshold for complete connection of the point set is $p_2 = \ln n/n$. Is $G(n, p_2 + \varepsilon)$ (almost surely) a spanner? My conjecture: YES.

Stretch Embedding in Labeled Grids

Belén Palop

Universidad de Valladolid

b.palop@infor.uva.es

Find two labelings Λ_1 and Λ_2 of the n^2 vertices in the $n \times n$ grid graph so as to maximize the minimum “stretch” between Λ_1 and Λ_2 . Here the labelings must use the same n^2 labels, say $1, 2, \dots, n^2$. Define the *combined distance* between a pair of labels (λ_1, λ_2) as the sum of the ℓ_1 distances between the two vertices marked with those labels under Λ_1 , and between the two vertices marked with those labels under Λ_2 . The goal is to find two labelings that maximize the minimum combined distance over all labels λ .

Shortly after the conference, two groups independently established that the optimal stretch is $\Theta(\sqrt{n})$. See [JPT10] in this proceedings.

References

- [JPT10] Minghui Jiang, Vincent Pilaud, and Pedro J. Tejada. On a dispersion problem in grid labeling. In *Proc. 22nd Canad. Conf. Comput. Geom.*, August 2010.

Crowd Simulation

Jack Snoeyink and Glenn Elliott

Univ. North Carolina

snoeyink@cs.unc.edu

Agent-based collision avoidance in crowd simulations leads to the following problem.

Suppose that a floor plan is given as a polygon P with holes, and we wish to simulate agents, represented as points $A = \{a_1, \dots, a_m\}$, traveling through this environment and avoiding collisions with what they can “see.” To this end we would

like to preprocess the floor plan to support “limited visibility queries”—reporting a region whose extent may depend on the complexity of an agent’s visibility polygon and number of visible agents. The idea is that if you are seeing a single agent approaching along a long, empty corridor, you will recognize this and make very early adjustments to avoid collisions, but if you are surrounded by many agents and/or visual clutter, then you will pay attention only to immediate neighbors.

One possible way to make this more precise would be to define the visual complexity within radius r for agent a_1 as the number of sides plus agents in the intersection of the ball $B_r(a_1)$ of radius r centered at a_1 with the visibility polygon of P from a_1 . In notation,

$$vc_1(r) = |B_r(a_1) \cap VP(a_1)| + |B_r(a_1) \cap VP(a_1) \cap A|.$$

If we decide that clutter is more than ξ things, we could then define the cluttered radius $r_1(\xi) = \operatorname{argmin}_r vc_1(r) > \xi$, and ask for a limited visibility query to return a description of size $O(\xi)$ that contains matches the visibility polygon $VP(a_1)$ inside $B_{r_1(\xi)}(a_1)$.

Note that this goes a step beyond output-sensitive algorithms—we are allowed to modify the output to keep it small, as long as we preserve the portion that is near the agent.

Compatible Tetrahedralizations

Stephen Kobourov

Univ. Arizona

kobourov@cs.arizona.edu

Given two polyhedra P_1 and P_2 whose surfaces are combinatorially equivalent (in particular, their 1-skeletons are isomorphic graphs), do they necessarily have a compatible tetrahedralization with Steiner points? A *compatible tetrahedralization with Steiner points* of two polyhedra is a tetrahedralization of the interior of each polyhedron, allowing additional (Steiner) vertices in addition to the boundary vertices given by the polyhedron, such that the two tetrahedralizations are combinatorially equivalent: there is a bijection between the vertices of the two tetrahedralizations that preserves triangles and tetrahedra in both directions.

The analogous problem for 2D polygons is known to be always possible [ASS93], requiring $\Theta(n^2)$ Steiner points in the worst case, where n is the number of vertices of each polygon.

In a more general version of the problem (which skirts the issue of how to define “polyhedron”), we are given two tetrahedralizations whose two-dimensional boundary surfaces are combinatorially

equivalent. Here we wish to find a tetrahedralization that refines each given tetrahedralization (i.e., includes all existing triangles), such that the two refined tetrahedralizations are combinatorially equivalent.

Compatible tetrahedralizations permit natural morphing between the geometries.

References

- [ASS93] Boris Aronov, Raimund Seidel, and Diane Souvaine. On compatible triangulations of simple polygons. *Computational Geometry: Theory and Applications*, 3(1):27–35, June 1993.

Landmark Navigation

Leo Guibas
Stanford University
guibas@cs.stanford.edu

The task is to navigate from point A to B , using distances to *landmark* points L_i .

For any point p , define $d_i(p) = d(p, L_i)$ and $d^2(p) = (1/n) \sum_{i=1}^n d_i^2(p)$. Define “coordinates” for A ,

$$c(A) = \langle d_1^2(A) - d^2(A), \dots, d_i^2(A) - d^2(A), \dots \rangle$$

and similarly for $c(B)$. Let

$$d^*(A, B) = \|c(A) - c(B)\|^2$$

With a sufficient number of landmarks, it is known that following gradient descent of $d^*(A, B)$ will get you from A to B [FGG⁺05].

The question is this: what happens if there are obstacles? In a polygon with holes, one can get stuck. The poser conjectures that, in a simple polygon, with the vertices as landmarks, and distance defined by the geodesic distance within the polygon, gradient descent works: one never gets stuck. There is experimental support for this conjecture. Perhaps using art-gallery guards as the landmarks would also suffice.

Update. Jack Snoeyink found a counterexample, drawn in Figure 2. The steepest descent path from A to B , with the vertices as landmarks, will (non-obviously) stop somewhere in the interior of the polygon.

References

- [FGG⁺05] Q. Fang, J. Gao, L. Guibas, V. de Silva, and L. Zhang. GLIDER: Gradient landmark-based distributed routing for sensor networks. In *Proc. IEEE Conference on Computer Communications (INFOCOM)*, 2005.




Figure 2: Note the unequal vertical and horizontal scales.

Nearest Neighbors for Well-Spaced Points

Don Sheehy
Carnegie Mellon University
dsheehy@cs.cmu.edu

Given n well-spaced points in \mathbb{R}^d , can they be preprocessed to achieve $2^{O(d)} \log n$ -time (exact) nearest-neighbor queries using only $2^{O(d)}n$ space? Points are *well-spaced* if the ratio of the radii of the balls enclosing and enclosed in every point’s Voronoi cell is $O(1)$.

Update. It is known that the the Voronoi diagram of well-spaced points has size $2^{O(d^2)}n$ in the worst case, counting faces of all dimensions. The large constant is the crux of the problem. Two results are known that solve the problem using similar, quadratic exponential constants in d , [HT08] and [Gra10]. The dependence on the dimension for the constants is implicit in the result in [HT08] (although not worked out in that paper). Still it may be possible to improve this constant to $2^{O(d)}$, especially given that the number of Voronoi facets is only $2^{O(d)}n$. Even $2^{O(d \log d)}$ would be interesting.

References

- [HT08] Benoît Hudson Duru Türkoğlu. An efficient query structure for mesh refinement. In *20th Canad. Conf. Comput. Geom.*, pages 115–118, 2008.
- [Gra10] Chris Gray. Nearest-neighbor queries with well-spaced points. In *Proc. EuroCG*, 2010.

Hausdorff Core of a Polygon**Robert Fraser****University of Waterloo****r3fraser@cs.uwaterloo.ca**

Given a simple polygon P , find a convex polygon contained in P that minimizes the maximum Hausdorff distance between the two polygons (in both directions). This problem is related to the “potato-peeling” problem, which instead asks for a contained convex polygon of maximum area. The latter problem can be solved by a dynamic-programming algorithm in $O(n^7)$ time. Is there a similar dynamic-programming algorithm for exact construction of this “Hausdorff core”? An approximation algorithm appeared in [DDF⁺09].

Joseph O’Rourke suggested a variant which removes the requirement that the convex polygon be enclosed in P , which the poser thought might be easier.

Update. The problem has been solved in a special case: see [FN10] in this proceedings.

References

- [DDF⁺09] Reza Dorriviv, Stephane Durocher, Arash Farzan, Robert Fraser, Alejandro López-Ortiz, J. Ian Munro, Alejandro Salinger, and Matthew Skala. Finding a Hausdorff core of a polygon: On convex polygon containment with bounded Hausdorff distance. In *WADS ’09: Proc. 11th Internat. Symp. Algorithms Data Structures*, pages 218–229, Berlin, Heidelberg, 2009. Springer-Verlag.
- [FN10] Robert Fraser and Patrick K. Nicholson. Hausdorff core of a one reflex vertex polygon. In *Proc. 22nd Canad. Conf. Comput. Geom.*, August 2010.