

Open Problems from CCCG 2007

Erik D. Demaine*

Joseph O'Rourke†

The following is a list of the problems presented on August 20, 2007 at the open-problem session of the 19th Canadian Conference on Computational Geometry held in Ottawa, Ontario, Canada.

Rolling a Sphere Upside-Down

Joseph O'Rourke

Smith College

orourke@cs.smith.edu

Imagine rolling a unit sphere on a plane without slipping or twisting, so that the point of contact follows a closed curve \mathcal{C} . The south pole touches the plane at the start. What is the shortest length $L = |\mathcal{C}|$ that results in the north pole touching after one complete circuit of \mathcal{C} ? Figure 1 shows a curve achieving $L = 3\pi$. Also, \mathcal{C} cannot be shorter than the geodesic distance between the poles: $L \geq \pi$. Another \mathcal{C} that achieves $L = 3\pi$ is an equilateral triangle of side length π . It is established in [Joh07] that the poles cannot be interchanged by any \mathcal{C} that is a circle.




Figure 1: Rolling a sphere to interchange the north and south pole. Here $|\mathcal{C}| = 3\pi$.

*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

Update. At the conference, Jack Snoeyink found a path with length approximately 2.72π . Vishal Verma, a graduate student at UNC Chapel Hill, joined him [VS07] to improve this to $< 2.44\pi$ for a teardrop-shaped path depicted in Figure 2.




Figure 2: Path of plane/sphere contact in plane (left) and on sphere (right). Dashed curve at right is the path of the south pole as the sphere rolls.

Hammersley [Ham83] had posed a more general form of this problem in the literature on optimal control: for a unit sphere lying on the plane at (x_0, y_0) and having initial orientation C_0 , determine the shortest path for it to roll without twisting that brings it to point (x_1, y_1) with orientation C_1 . Although the solution to Hammersley's problem does not in general have a closed form, Arthurs and Walsh [AW86] have given an expression as a boundary-value problem with ten coupled partial differential equations that they use to derive information on the curvature of optimal paths. For the special case of a closed path above, their result implies that the curvature is proportional to the distance along the axis of symmetry, which establishes that the curve depicted in Figure 2 is optimal if self-intersections are ruled out.

References

- [AW86] A. M. Arthurs and G. R. Walsh. On Hammersley's minimum problem for a rolling sphere. *Mathematical Proceedings of Cambridge Philosophical Society*, 99:529–534, 1986.
- [Ham83] J. M. Hammersley. Oxford commemoration ball. In *Probability, Statistics*

and *Analysis*, Cambridge University Press, 1983, pp. 112–142.

- [Joh07] Brody D. Johnson. The nonholonomy of the rolling sphere. *American Mathematical Monthly*, 116(6):500–508, 2007.
- [VS07] Vishal Verma and Jack Snoeyink. Upend- ing a sphere. In *Abstracts from the 17th Annual Fall Workshop on Computational Geometry*, 2007. http://www.research.ibm.com/people/l/lenchner/fwcg2007/Proceedings/submission_35.pdf

Spanning Trees of the Graph of a Polyhedron

Alex Benton

Cambridge University

pb355@damtp.cam.ac.uk

What is the best possible bound on the number of spanning trees of the 1-skeleton of a polyhedron, i.e., a 3-connected planar graph?

Update. An upper bound on the number of spanning trees of a polyhedron graph is derived in [DO07, p. 431], based on a result of McKay [McK83]: the number is $O((16/3)^n/n)$, which is $2^{O(n)}$.

References

- [DO07] Erik D. Demaine and Joseph O’Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, July 2007. <http://www.gfalop.org>.
- [McK83] Brendan D. McKay. Spanning trees in regular graphs. *European Journal of Combinatorics*, 4(2):149–160, 1983.

Minimum Length Barrier to X-rays in a Square

Jonathan Lenchner

IBM T.J. Watson Research Center

lenchner@us.ibm.com

Given a unit square, construct a barrier of minimum length that intersects every line passing through a portion of the square. The barrier should consist of one or more piecewise-smooth curves. The barrier need not be connected and portions of the barrier may be located inside, outside, or on the boundary of the square.

This problem appeared in the July 2007 edition of IBM’s online puzzle column *Ponder This*¹. According to the column, the problem has also appeared in [Jon64] and in an internal publication of the Lawrence Livermore National Laboratory.

¹See <http://domino.research.ibm.com/Comm/wwwr-ponder.nsf/challenges/July2007.html>.

The obvious barrier is the entire perimeter of the square, with cost 4. However, we can also use just three sides of the square, at cost 3. Even better, we can use the two diagonals, at cost $2\sqrt{2} \simeq 2.828$. Better still is to use two adjacent edges of the square and half of the opposite diagonal, at cost $2 + \sqrt{2}/2 \simeq 2.707$, as shown in Figure 3.




Figure 3: A barrier of length $2 + \sqrt{2}/2 \simeq 2.707$.

As a final observation, we can do better than using the two adjacent edges in the previous barrier, by instead using a three-segment Steiner tree as in Figure 4, at cost $\sqrt{2} + \sqrt{6}/2 \simeq 2.639$. This barrier is conjectured, but not known, to be optimal.




Figure 4: Place the bottom left corner of the square at $(0, 0)$. Then a barrier consisting of the diagonal segment $[(1/2, 1/2), (1, 1)]$ together with three segments formed by joining the corners $(0, 1), (0, 0), (1, 0)$ to a point at $(\frac{1}{2} - \frac{\sqrt{3}}{6}, \frac{1}{2} - \frac{\sqrt{3}}{6})$ yields a barrier of length $\sqrt{2} + \frac{\sqrt{6}}{2} \simeq 2.639$.

At the conference, Otfried Cheong suggested that a lower bound of 2 could be derived from the Cauchy-Crofton formula, and this was later verified to be the case.

A clever and even more elementary proof that 2 is a lower bound was found by Ozgur Ozkan, a student of John Iacono’s at Polytechnic University. Ozgur’s argument runs as follows:

By a limiting argument, we may assume that the (approximately) optimal solution is piecewise linear. Thus let $S = \{x_1, x_2, \dots\}$ be the set of line segments making up a barrier. Let the two diagonals of the square be denoted by d_1 and d_2 . In order to block just the rays which are perpendicular to each of these two diagonals, the projection of S onto d_1 and d_2 must cover d_1 and d_2 . Thus, if θ_i is the angle x_i makes with d_1 , then $|x_i| \cos \theta_i$ is the length of x_i ’s projection onto d_1 , and $|x_i| \sin \theta_i$

is the length of x_i 's projection onto d_2 . Therefore

$$\begin{aligned} 2\sqrt{2} &\leq \sum_{x_i \in S} |x_i|(\cos \theta_i + \sin \theta_i) \\ &\leq \sum_{x_i \in S} |x_i|\sqrt{2}, \quad \text{so} \\ 2 &\leq \sum_{x_i \in S} |x_i|. \end{aligned}$$

References

- [Jon64] R. E. D. Jones. Opaque sets of degree α . *American Mathematical Monthly*, 71:535–537, May 1964.

Doubly Orthogonal Point Set?

Therese Biedl

University of Waterloo

biedl@softbase.uwaterloo.ca

Is there a point set in 2D that is the vertex set of an orthogonal polygon such that a rotation of the point set by a nonmultiple of 90° is also the vertex set of an orthogonal polygon?

Precisely, an orthogonal polygon is a (simple) polygon whose edges are all either horizontal or vertical. A vertex is a point on the polygon where two edges of different slopes meet. (Put differently, we do not allow “extra” vertices along the edges.) The open problem is motivated by the question of whether it is possible to reconstruct an orthogonal polygon when given only a set of points that is supposed to be its vertex set. (Think of the popular children’s game connect-the-dots, except that the numbers on the dots are illegible.)

This problem has been well-studied. O’Rourke [O’R88] showed that a simple scanning algorithm can recover the orthogonal polygon in $O(n \log n)$ time. Rappaport [Rap86] showed that, if extra vertices are allowed, then the problem becomes NP-hard.

We recently started studying the problem where the given point set is allowed to be rotated before reconstructing the orthogonal point set ([Gen07]; see also [BG07]). Clearly only $O(n)$ rotations could possibly yield an orthogonal polygon, because at least four edges of the polygon have to be on an edge of the convex hull. Therefore, an $O(n^2 \log n)$ algorithm for this problem is trivial. We managed to improve the time complexity to $O(n \log n)$ for orthogonally convex polygons. The crucial insight was that, for orthogonally convex polygons, there can be only one rotation that could possibly work; and the proof of this insight yielded an algorithm to find this rotation efficiently.

This raises the natural question of what happens with orthogonal polygons that are not orthogonally convex. Our proof very clearly fails for such polygons, hence the open question: is there only one rotation that works? Or could there be two different rotations for which the set of points is the vertex set of an orthogonal polygon? The only negative example that we could find consists of the points of a regular octagon, which can be interpreted as the vertex sets of two rectangles in two different rotations. But this is neither a single polygon nor simple.

We would also be interested in whether one could find the rotation efficiently, if there is only one.

Update: Maarten Löffler and Elena Mumford [LM07] resolved this question in the negative. In fact, they consider a more general problem: given a set of points in \mathbb{R}^d , is it the vertex set of some connected rectilinear graph (not necessarily a polygon)? They prove that any point set in \mathbb{R}^d has at most one orientation where it is a vertex set of a connected rectilinear graph. For contrast, Figure 5 shows an example of two rectilinear graphs on the same point set, but note that G' is not connected (similar to the regular-octagon example above). For the special case where the points are in the plane and have rational coordinates, Fekete and Woeginger [FW97, Theorem 4.7] already proved that at most one orientation is possible.




Figure 5: Two rectilinear graphs in the plane with the same vertex set, but different slopes. Note that G' is not connected.

The details of the argument in the planar case are mostly algebraic, but the main idea is to determine the “greatest common divisor” in some sense between the coordinates of the vertices of an axis-aligned connected rectilinear graph, and argue that, if there is a connected rectilinear graph in another orientation, we can use that to travel to a point that is not a multiple of this divisor away from the others, which leads to a contradiction. The result can be extended to arbitrary dimensions by simply projecting it down to a suitable plane.

To find the right orientation of the graph (in the planar case), we can improve the $O(n^2 \log n)$ result to $O(n^2)$ by taking the dual of the problem. A relatively simple algorithm can sweep the arrangement

of lines and identify potential good orientations, of which there are at most $O(n)$. We maintain the sorted order of the other lines, which allows us to test an orientation in linear time.

If the goal is to find a orientation that allows a *planar* connected rectilinear graph (a more general case than simple polygons), then the problem becomes NP-hard. Because we know that there is at most one such orientation, we can use the algorithm described above to find it. Then we can compute the maximal rectilinear graph with this slope. However, now we need to decide whether this graph has a planar connected subgraph. Jansen and Woeginger [JW93] proved this problem NP-complete.

The main remaining open question is whether the (unique) rectilinear orientation can be computed in less than quadratic time. Also, algorithmic extensions to higher dimensions are open.

References

- [BG07] Therese Biedl and Burkay Genç. Reconstructing orthogonal polyhedra from putative vertex sets. Technical Report CS-2007-28, School of Computer Science, University of Waterloo, August 2007. <http://www.cs.uwaterloo.ca/research/tr/2007/CS-2007-28.pdf>
- [FW97] Sándor P. Fekete and Gerhard J. Woeginger. Angle-restricted tours in the plane. *Computational Geometry: Theory and Applications*, 8(4):195–218, 1997.
- [Gen07] Burkay Genç. Reconstruction of Orthogonal Polyhedra. Ph.D. Thesis, School of Computer Science, University of Waterloo, December 2007.
- [JW93] Klaus Jansen and Gerhard J. Woeginger. The complexity of detecting crossing-free configurations in the plane. *BIT*, 33(4):580–595, 1993.
- [LM07] Maarten Löffler and Elena Mumford. Rotating rectilinear graphs. In *Abstracts from the 17th Annual Fall Workshop on Computational Geometry*, 2007. http://www.research.ibm.com/people/l/lenchner/fwcg2007/Proceedings/submission_23.pdf
- [O’R88] Joseph O’Rourke. Uniqueness of orthogonal connect-the-dots. In Godfried T. Toussaint, editor, *Computational Morphology*, pages 97–104. North-Holland, Amsterdam, Netherlands, 1988.
- [Rap86] David Rappaport. On the complexity of computing orthogonal polygons from a set

of points. Technical Report TR-SOCS-86.9, McGill University, 1986.

Pushing Cubes Around

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

Dumitrescu and Pach [DP04] have showed that any configuration of unit squares on the integer lattice can be reconfigured to any other such shape via two types of moves, while remaining connected throughout. “Connected” here means 4-connected, i.e., edge-edge connected. I ask the same question for a configuration of unit cubes, where connectivity must be via face-face connections. Such objects, formed by face-to-face gluing of unit cubes, are called *polycubes*. The two moves permitted are exactly the Dumitrescu-Pach moves, shown in Figure 6, except now available parallel to any of the three coordinate planes. Perhaps applying the Dumitrescu-Pach algorithm to each xy -layer of cubes will help, but even if this maintains connectivity, in general cubes must be transferred between layers.




Figure 6: Cube moves based on Dumitrescu-Pach square moves.

Update. Zachary Abel and Scott D. Kominers [AK08] have recently announced a solution to this problem in the affirmative and have given a generalization to configurations of hypercubes of any dimension. Their method uses iterative relocation of modules from a configuration V to form a canonical chain at a distinguished module of V . They efficiently locate a module on the boundary of V which can either be relocated to the canonical chain immediately or can be relocated after recursive modification of the interior V .

Specifically, their methods prove the existence of a module x on V ’s outer boundary such that if $V \setminus \{x\}$ is not connected then $V \setminus \{x\}$ consists of exactly two components, one of which is disjoint from the outer boundary of V . Furthermore, this module x may be located quickly. Indeed, the solution yields an algorithm which requires at most $O(n^2)$ calculation time and at most $O(n^2)$ moves; this is asymptotically optimal.

References

- [AK08] Zachary Abel and Scott D. Kominers. Pushing hypercubes around. Preprint, [arXiv:0802.3414](https://arxiv.org/abs/0802.3414).
- [DP04] Adrian Dumitrescu and Janos Pach. Pushing squares around. In *Proceedings of the 20th Annual ACM Symposium on Computational Geometry*, pages 116–123, 2004.

Surface Flips

Joseph O’Rourke

Smith College

orourke@cs.smith.edu

Although this problem can be stated with more generality, I choose to pose it only for a polycube object P as defined in the previous problem. For any portion S of the surface P that constitutes a topological disk, and whose boundary is a cycle C of edges of P all lying in one plane Π , we define a *surface flip* as reflecting S through Π , as long as this operation maintains weak simplicity of the resulting surface P' . See Figure 7.




Figure 7: Surface flip about red 4-cycle.

Note that, whereas the previous problem preserved the volume, this move preserves the surface area. Also, the combinatorial structure of the surface does not change under these surface flips.

Characterize the class of polycube shapes that are connected by surface flips. In particular, are all shapes with the same combinatorial type connected?

Polycube Pops

Joseph O’Rourke

Smith College

orourke@cs.smith.edu

Again this problem is restricted to polycubes, and the moves preserve the surface area, but now they

will alter the combinatorial structure. My goal is to define a set of moves that will be able to convexify a polycube, or convert to some other canonical form, in a manner that will serve in some sense as a generalization of the vertex pop moves for unit orthogonal polygons (i.e., *polyominoes*) explored in [ABB⁺07].

There are five moves, two primary moves, and three moves concerned with collocated faces, which we will call “fences” (the analog of “pins” in 2D). They are named as follows:

1. Corner pop.
2. U-pop.
3. Fence walk.
4. Fence pop.
5. Fence corner pop.

The two primary moves are illustrated in Figure 8. The first move, a *corner pop*, is the one most closely inspired by a vertex pop: three faces incident to the corner of a cube are replaced by the three other faces of that cube. The second move, a *U-pop*, also replaces three faces of a cube by three others. Note that both moves preserve surface area.




Figure 8: Two “pop” moves: (a,b) Corner pop; (c,d) U-pop.

To complete the definition of these moves, we need to specify how the adjacent faces are connected. This is shown in Figure 9, which makes it clear that in an abstract network of quadrilaterals forming the surface, the two moves replace the “wirings” internal to a 6- or 8-cycle; all exterior connections remain in place.

To have hope of reaching an interesting canonical form, it is important to permit the surface to become *weakly simple*, with perhaps several faces lying back-to-back on top of one another. We call




Figure 9: Combinatorial diagrams corresponding to the 3D moves in Fig. 8. (a,b) Corner pop; (c,d) U-pop.

two collocated faces a *fence*. It seems that three fence moves are needed. The *fence walk*, illustrated geometrically in Figure 10 and combinatorially in Figure 11(a,b), can be viewed as replacing three faces of a cube with three others, but this time two of the three faces are collocated.




Figure 10: Fence walk.

A *fence pop* takes a connected tree of fences all perpendicular to the same plane Π , and reflects them through Π . This move involves no combinatorial change. Finally, a variation on this is a *fence corner pop*, which does involve a combinatorial change, shown in Figure 11(c,d).




Figure 11: (a,b) Fence walk. (c,d) Fence corner pop.

Figure 12 shows an example of two shapes that can be connected by these polycube pop moves:

two corner pops, one per dent, two fence corner pops, and finally a U-pop. Note that the dents in (a) of the figure consist of 4 faces each, and are replaced by 2 faces. The extra 4 faces gained are then enough to build the highest cube in (b).




Figure 12: Object (a) can be converted to (b) via polycube pops.

There are (at least) two questions here: (1) Which class of shapes can be convexified (converted to an orthogonally convex polyhedron, e.g., Figure 12(b)) via the above five moves? (2) If not all shapes, then is there a natural set of moves that does suffice to connect all polycubes of the same surface area?

References

[ABB⁺07] Greg Aloupis, Brad Ballinger, Prosenjit Bose, Mirela Damian, Erik D. Demaine, Martin L. Demaine, Robin Flatland, Ferran Hurtado, Stefan Langerman, Joseph O'Rourke, Perouz Taslakian, and Godfried Toussaint. Vertex pops and popturns. In *Proceedings of the 19th Canadian Conference on Computational Geometry*, pages 137–140, 2007.

Dice Rolling in a Rectangle

Erik Demaine
MIT
edemaine@mit.edu

When a rectangle R is fully labeled, and so there are no free nor blocked cells inside R , what is the complexity of deciding whether a cube can roll over R compatibly with the labels? For definitions, see [BBD⁺07]. That paper conjectures this restricted decision problem is solvable in polynomial time.

Two variants were suggested at the presentation: (1) What if R is labeled only with a subset of the six die labels? (2) What if the sides of R are glued to form a torus (left glued to right, top glued to bottom)?

References

- [BBD⁺07] Kevin Buchin, Maike Buchin, Erik D. Demaine, Martin L. Demaine, Dania El-Khechen, Sándor Fekete, Christian Knauer, André Schulz, and Perouz Taslakian. On rolling cube puzzles. In *Proceedings of the 19th Canadian Conference on Computational Geometry*, pages 141–148, 2007.

Room Reconstruction from Point/Normal Data

Jack Snoeyink

Univ. North Carolina Chapel Hill

snoeyink@cs.unc.edu

Many sensors, such as the Delta-sphere constructed at UNC Chapel Hill, can send out beams and get range information for scattered points in a room. Suppose that your sensor gives not only a point, but also the normal to a plane when it hits a wall (perhaps by gathering individual point returns and using consensus to reduce error in wall positioning). You'd like to determine whether you have seen all the walls in your room. That is, from just the walls you have seen, can you construct a room that explains all observations? These questions are interesting even in the orthogonal case, where the walls must lie parallel to the coordinate planes.

This problem is relatively easy with a single sensor of known position in 2D: sort the sensed points radially and use the corresponding lines to construct a polygon, then check whether it is star-shaped with the sensor in the kernel. What about 3D, where the walls can have more complex shapes?

There are many variants on this problem: whether there are one or more sensors, whether you know each sensor(s) position, whether sensors return a few or all points (visibility polygon), whether the room must be orthogonal or simply connected, and whether the room reconstructed from the walls is unique.

Update: Biedl and Snoeyink have since been able to show that it is NP-hard to determine whether there is a unique reconstruction of an orthogonal polygon from a collection of 2D visibility polygons representing the information from several sensors.

Wireless Reflections

Boaz Ben-Moshe

College of Judea and Samaria

benmo@yosh.ac.il

Given a transmitter and receiver in an environment with barriers (walls, etc.), and given an integer

$c > 1$, find all the paths that go from transmitter to receiver by c billiard reflections (and perhaps also go through walls via transmission). In general, c will be small, say, $c < 100$. The problem arises in MIMO (Multiple-Input, Multiple-Output) communications.

Update: Ben-Moshe et al. have in some sense solved their problem and are implementing an algorithm as part of the Israeli Short Range Communication Consortium, <http://www.isrc.org.il/index.asp>.

Polygon that Projects as Chain

Prosenjit Bose

Carleton Univ.

jit@scs.carleton.ca

Is there a (closed) simple polygon in 3-space that projects to an open polygonal chain in three orthogonal directions? It is known that there is an open polygonal chain in 3-space that projects to a (closed) simple polygon in three orthogonal directions.

This problem was posed by Jack Snoeyink at an earlier CCCG, reporting a problem originally posed by Claire Kenyon.

Stretch Factor for Points in Convex Position

Prosenjit Bose

Carleton Univ.

jit@scs.carleton.ca

The *stretch factor* for a geometric graph G is the maximum, over all vertices u and v in G , of the ratio of the length of the shortest path from u to v in G to the Euclidean distance between them, $|u - v|$. If G has a stretch factor of t , it is called a t -spanner. Chew conjectured that the Delaunay triangulation is a t -spanner [Che89] for some constant t . Dobkin et al. [DFS90] established this for $t = \pi(1 + \sqrt{5})/2 \approx 5.08$. The value of t was improved to $2\pi/(3 \cos(\pi/6)) \approx 2.42$ by Keil and Gutwin [KG92], and further strengthened in [BM04]. Chew showed that t is $\pi/2 \approx 1.57$ for points on a circle, providing a lower bound. “It is widely believed that, for every set of points in \mathbb{R}^2 , the Delaunay triangulation is a $(\pi/2)$ -spanner” [NS07, p. 470].

This suggests the following special case: for points S in convex position (i.e., every point is on the hull of S), is the Delaunay triangulation of S a $(\pi/2)$ -spanner?

References

- [BM04] P. Bose and P. Morin. Online routing in triangulations. *SIAM Journal on Computing*, 33(4):937–951, 2004.

- [Che89] L. Paul Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences*, 39(2):205–219, 1989.
- [DFS90] David P. Dobkin, Steven J. Friedman, and Kenneth J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry*, 5(4):399–407, 1990.
- [KG92] J. Mark Keil and Carl A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7(1):13–28, 1992.
- [NS07] Giri Narasimhan and Michiel Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.