Open Problems from CCCG 2011

Erik D. Demaine* Joseph O'Rourke†

The following is a description of the problems presented on August 10, 2011 at the open-problem session of the *23rd Canadian Conference on Computational Geometry* held in Toronto, Ontario, Canada.

Blocking visibility with cylinders
Joseph O'Rourke
Smith College
orourke@cs.smith.edu

Suppose you have a supply of infinite-length, opaque, unit-radius cylinders, and you would like to block all visibility from a point $p \in \mathbb{R}^3$ to infinity with as few cylinders as possible. (The cylinders are infinite length in both directions.) The cylinders may touch but not interpenetrate, and they should be disjoint from p, leaving a small ball around p empty. (Another variation would insist that cylinders be pairwise disjoint, i.e., not touching one another.)

A collection of parallel cylinders arranged to form a "fence" around p do not suffice, leaving two line-of-sight $\pm$ rays to infinity. Perhaps a grid of cylinders in the pattern illustrated in Figure 1 (left) suffice, but at least if there are not many cylinders, there is a view from an interior point to infinity (Figure 1, right).

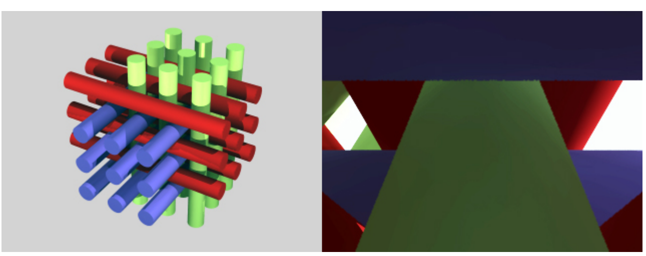

Figure 1: A grid of cross cylinders. A view from inside shows not all visiblity is blocked.

This question was originally posed on MathOverflow [OR11a], and several ideas contributed there suggest to start with the six cylinder arrangment in Figure 2 (left), supplemented by a circular "forest" to block the remaining lines of sight, three-quarters of which are illustrated in Figure 2 (right). The illustrated configuration needs 18 cylinders, but perhaps as few as 10 suffice for this plan?

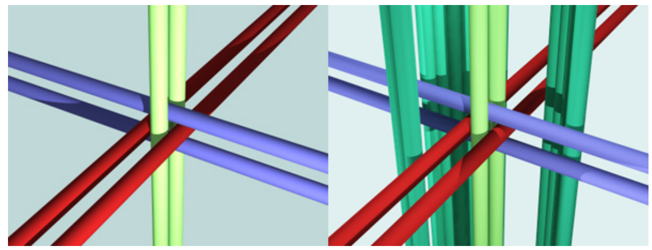

Figure 2: Six cylinders block all but some "diagonal" lines of sight. Erecting a vertical fence should then block all lines of sight.

What is the minimum number of infinite cylinders that can block visibility from a point?

References

[OR11a] J. O'Rourke. Blocking visibility with cylinders. http://mathoverflow.net/questions/69963/ 11 May 2011.

The Rain Hull and the Rain Ridge
Joseph O'Rourke
Smith College
orourke@cs.smith.edu

Rain falls steadily on an island, a 2-manifold M, which you may assume, as you prefer, is: (a) smooth, or (b) a PL-manifold, or perhaps (c) a triangulated irregular network (TIN). After a time, M is saturated, in the sense that every raindrop drains into the ocean rather than filling yet-unfilled crevices or basins. At this point, we have what I will dub the *rain hull* of M, $H_R(M)$, a unidirectional version of the the reflex-free hull defined by Jack Snoeyink at the 13th CCCG [ACCS04]

(1) How difficult is to compute the rain hull $H_R(M)$?

This question was originally posed on MathOverflow [OR11b] and a respondent there (Joel Hamkins) argued that at least it can be computed in polynomial time. Nonlocal effects such as that illustrated in Figure 3 must be accommodated.

*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

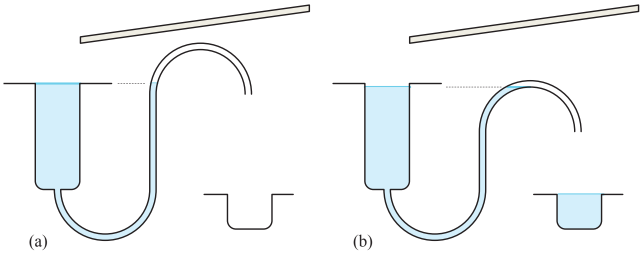

Figure 3: (a) The connecting tube rises too high to fill the other, protected basin. (b) A lower tube does overflow into the other catch-basin.

Let us assume we have $\overline{M} = H_R(M)$ computed or given. A raindrop falling on $p \in \overline{M}$ might follow a unique *trickle path* (that is the technical term: e.g., see [dBHT11]) to the ocean, or the drop may randomly 'fracture' to follow distinct paths to the ocean. Define the *rain ridge* (my terminology) $R(\overline{M})$ to be the complement of the points of $\overline{M}$ that have a unique trickle path.

So points on the rain ridge are akin to points on a *cut locus*, in that they have two or more distinct paths to $\partial\overline{M}$. They are, in a sense, *continental-divide* points [Hay09].

(2) What can be said about the structure of the rain ridge $R(\overline{M})$? And how quickly can it be computed?

Unlike the cut locus or "ridge tree," the rain ridge is not always a tree. All the points in a filled basin are in the rain ridge, for when a raindrop lands in a filled basin, it is natural to assume it "spreads out" and spills in equal portions over every boundary point of the basin. But surely there are substantive properties to investigate. Surely the rain ridge $R(\overline{M})$ cannot be an arbitrary subset of $\overline{M}$?

(3) Can an extended metric be assigned to $\overline{M}$ so that its geodesics are its trickle paths?

An *extended metric* is one that permits $d(x, y) = \infty$ (e.g., for points not on the same trickle path). What I am hoping for here is a way to view the rain ridge as a cut locus of $\partial\overline{M}$, and then apply a century of knowledge on the cut locus to the rain ridge.

References

[ACCS04] H.K. Ahn, S.W. Cheng, O. Cheong, and J. Snoeyink. The reflex-free hull. *International Journal of Computational Geometry and Applications*, 14(6):453–474, 2004.

[dBHT11] M. de Berg, H. Haverkort, and C. Tsirogiannis. Implicit flow routing on terrains with applications to surface networks and drainage structures. In *Proc. 22nd ACMSIAM Symp. on Discrete Algorithms (SODA)*, pages 285–296, 2011.

[Hay09] B. Hayes. Dividing the Continent. In *Group Theory in the Bedroom, and Other Mathematical Diversions.* Hill and Wang, 2009. pp. 107–123.

[OR11b] J. O'Rourke. The rain hull and the rain ridge. http://mathoverflow.net/questions/69963/ 10 July 2011.

Long Alternating Paths

Jorge Urrutia
Universidad Nacional Autónoma de México
urrutia@matem.unam.mx

Let P_{kn} be a point set with kn points in general position. A k-coloring of P_{kn} is a partitioning of P_{kn} into k disjoint subsets $S_1, \ldots, S_k$, each with n elements. The sets $S_1, \ldots, S_k$, are called the chromatic classes of P_{kn}.

An *alternating* path Π of P_{kn} is a simple polygonal path connecting a subset of the points of P_{kn} such that there are no monochromatic edges in the path.

Conjecture Any 3-colored point set P_{3n} contains an alternating path with at least $2n$ elements.

We have been unable to prove that P_{3n} always contains an alternating path with $\frac{3}{2}n$ points; this seems to be a challenging weaker open problem. For 3-colored point sets P_{3n} in convex position, it is known there always exists a path that covers $2n$ points, and that this bound is tight [MSU]. Tight bounds for 2-colored point sets are not known for point sets in convex, or in general position [AGHNP].

References

[AGHNP] M. Abellanas, J. Garcia, G. Hernandez, M. Noy, and P. Ramos. Bipartite embeddings of trees in the plane. Discrete Appl. Math. 93 (1999), 141148.

[AU] J. Akiyama, and J. Urrutia. Simple alternating path problem. *Discrete Math*, (1990) 84: 101-103.

[MSU] C. Merino, G. Salazar and J. Urrutia. On the length of the longest alternating path for multicoloured point sets in convex position. *Discrete Mathematics*, Vol. 360, no. 15, pp. 1791-1797, 2006.

[PKT] J. Pach, J. Kynčl, and G. Tóth. Long alternating paths in bicolored point sets. *Discrete Mathematics*, 308 (2008), 4315–4322.

Monochromatic Empty Triangles
Jorge Urrutia
Universidad Nacional Autónoma de México
urrutia@matem.unam.mx

Let P_n be a set of n points in general position on the plane, each of which is colored red or blue. A triangle with vertices in P is called *empty* if it contains no point of P in its interior, it is called *monochromatic* if all of its vertices are red, or all are blue.

Conjecture Any bicolored point set P_n contains $\Omega(n^2)$ monochromatic empty triangles.

A liner bound was established in [DHKS]. It was improved to $cn^{\frac{5}{4}}$ in [AFHU], and to $cn^{\frac{4}{3}}$ in [PT].

References

[AFHU] O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, and J. Urrutia, Monochromatic empty triangles. *Computational Geometry,* Vol. 42, Issue 9, November 2009, Pages 934–938.

[DHKS] O. Devillers, F. Hurtado, Gy. Károlyi, and C. Seara, Chromatic variants of the Erdős-Szekeres theorem on points in convex position, *Computational Geometry,* Vol. 26, Issue 3, 2003, pp. 193–208.

[PT] J. Pach, and G. Tóth. Monochromatic empty triangles in two-colored point sets. Geometry, Games, Graphs and Education: the Joe Malkevitch Festschrift (S. Garfunkel, R. Nath, eds.), COMAP, Bedford, MA, 2008, 195-198. Also in: Discrete Applied Mathematics, submitted.

Shortest Periodic Light Ray
Boaz Ben-Moshe
Ariel University Center of Samaria
benmo@g.ariel.ac.il

Given a simple polygon, find the shortest periodic path of a light ray reflecting from the polygon edges as perfect mirrors. This problem is solved for *rational triangles*, those whose angles are rational multiples of π, but seems to be open for arbitrary triangles.

The Geometry of Golf
Alejandro López-Ortiz
University of Waterloo
alopez-o@uwaterloo.ca

After repeated unsuccessful attempts to get the ball in the hole from a particular point in the green, a golfer walks away in frustration and declares: That shot is impossible!

A mathematician happens to be standing nearby and says outloud: Hmmm, is it true that one can always putt a golf ball into the hole on this or any other arbitrary green?

A computer scientist overhears the mathematician and thinks: for given a green and ball location can I use my smartphone to determine if the shot is possible and if so in what direction and speed should I hit the ball?

More formally, the mathematician's question becomes: does every smooth two-dimensional manifold under a gravitational potential field is "connected" in the sense that a point particle at an arbitrary point on it can be made to roll into any other point on the manifold given a proper nudge (initial velocity vector) in the right direction.

The answer for this question is NO. A simple counterexample folds the surface into caves, but this is not strictly necessary: there are C^∞ manifolds which are described by the plot of a function $f : R \times R \to R$ and yet do not allow rolling motions into the hole. To see this consider a green with a mountain ridge between the ball and the hole, as depicted in Figure 4. If the slope on the hole side of

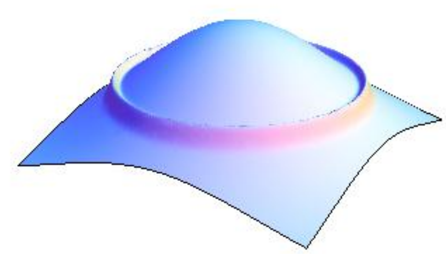

Figure 4: The hole is at the top center. This particular example is due to Jaap Eldering at http://mathoverflow.net/questions/84033/.

the ridge is sufficiently steep the ball can be made to become airborne and overfly the hole; see Figure 5.

The computer scientist's questions, posed formally, become: first, given a description of a green (perhaps discretized as a TIN) give an efficient algorithm that determines if the hole can always be

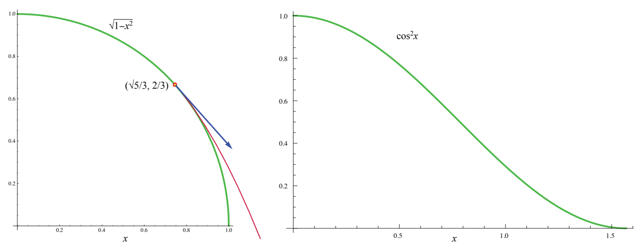

Figure 5: (a) Not gently falling. (b) Gently falling. Figure from [OR11c].

reached from all points, and second, given a ball position in said green can we compute the direction and speed of the putting action that will roll the ball into the hole?

Several variations of the question

Putting. Under what conditions can a given ball on a C^∞ manifold, with a quadratic gravitational field, reach the hole?

Golf green design. Under what conditions can every putt on a C^∞ manifold, with a quadratic gravitational field, reach the hole?

Hole location design. Given a C^∞ manifold, which points on it are reachable from all others and hence would be reasonable choices for location of the hole?

Chipping. What if you can chip, i.e., loft the ball?

Driving. Under what conditions is it possible to achieve a "hole-in-one" from the driving tee if we consider obstacles such as trees?

Sand Save. Under what shape conditions can you chip out of a sand trap and always move closer to the hole.

To understand the physics of the problem we study the 2D setting of a ball rolling down a curve.

First consider the instantaneous version of the problem: given a ball on the curve and moving at a certain speed will it become airborne at this instant?

We consider first the case where the particle was at rest. In this setting two forces are acting on the ball, namely gravity and surface resistance. Gravity is a vertical vector pointing downwards with magnitude $9.8m/s^2$. The surface resistance is a vector perpendicular to the tangent to the curve.

Observe that the magnitude of the resistance vector is exactly equal to to the projection of the gravitational vector on the normal direction to the tangent to the curve at all times. This is easier to see when the particle is at rest: if the forces on the direction of the surface were not perfectly canceled with that of gravity then the particle would either burrow into the surface until equilibrium is achieved (think of a really heavy ball making a dent) or would magically start hovering over it, both of which do not happen with a rolling golf ball.

Hence the only movement possible for a particle at rest is in the direction of the tangent to the surface and the surface resistance must perfectly cancel any force in any other direction.

The ball will then move along the direction of the tangent at a speed which is given by the addition of the resistance vector to that of the force of gravity. Let $\Gamma(t) = (x(t), y(t))$ denote the trajectory of the particle parameterized by the time t. Then the speed vector $v(t)$ is given by $d\Gamma/dt$, and the instantaneous change in speed is given by the differential equation $v'(t) = ||\Gamma'(t)||^{-1}(\Gamma'(t)\cdot(0,-g))+(0,-g)$.

The particle becomes airborne if the speed vector ever lies above the tangent to the curve which would result in a ski-like take off along the direction of the speed vector.

If the particle is already in motion then the same equations apply and the only change is in the initial condition $v(t)$ which for the particle at rest case was $v(t_0) = \vec{0}$ and now becomes $v(t_0) = (v_x(t_0), v_y(t_0))$.

We can test for the airborne state if we recall that the cross product of two vectors $\vec{a} = (a_1, a_2), = \vec{b} = (b_1, b_2)$ in the plane is the vector $(0, 0, a_1b_2 - a_2b_1)$, where the last coordinate is positive if and only if $\vec{a}$ is below $\vec{b}$. Substituting the speed vector and the tangent vector above we get that the ball remains on the surface iff

$$v_x(t)a_y(t) - v_y(t)a_x(t) \leq 0$$

We can now use this equation together with an iterative differential equation solver to numerically test this property along the entire trajectory of a putting path.

Update. In [OR11c] differential equations are derived for what can be considered a one-dimensional version of the putting problem.

References

[OR11c] J. O'Rourke. Gently falling functions. MathOverflow. http://mathoverflow.net/questions/68114/ 18 June 2011.

Polygon Triangulation Without Large Angles

Alexander Rand
University of Texas at Austin
arand@ices.utexas.edu

Let P be a generic convex polygon with vertices $V_1, V_2, \ldots, V_n$ (and define $V_0 := V_n$ and $V_{n+1} := V_1$ for simplicity). For $\gamma < \pi$, we will say that P belongs to the set S_γ if for any $i \notin \{j, j+1\}$ then $\angle V_j V_i V_{j+1} < \gamma$, i.e., no vertex forms a large angle with any opposite side of the polygon. See Figure 6.

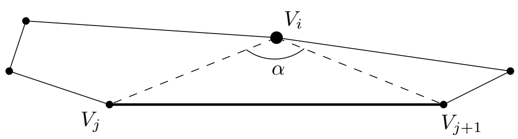

Figure 6: If $\angle V_j V_i V_{j+1}$ (denoted by α in the figure) is large, then no triangulation exists without a large angle. If this angle is bounded for all pairs of vertices and opposite edges, we expect some acceptable triangulation can be formed.

> **Open Problem** For $\gamma < \pi$, give an algorithm that, for any convex polygon in S_γ, adds some vertices to the *interior* of the polygon and produces a triangulation with no angles larger than $\theta(\gamma) < \pi$.

- Most related problems/algorithms in the literature (e.g., [BMR95, MPS07]) involve inserting vertices on the boundary of the polygon, which we have disallowed.
- The restriction to the set S_γ is essential: an obtuse triangle with largest angle very near π cannot be triangulated satisfying our requirements.
- The specific relationship between θ and γ can be selected at the discretion of the solver. The best solution hopefully has a form $\pi - \theta(\gamma) = \Omega(\pi - \gamma)$.
- The number of points added by the algorithm is unimportant for the original motivation of the problem, but it makes sense to ask what is the fewest number of vertices which can be added (which examples suggest is small; see conjecture below). Simple approaches often involve $O(n)$ vertices. Adding a very large number of vertices is not particularly helpful as any vertex placed too close to a boundary edge produces a large angle in the triangulation.
- At least one vertex can be necessary. Any triangulation of the regular n-gon without any additional vertices will produce at least one triangle of three consecutive vertices and, for large n, this has a large angle. Adding a single vertex at the center of the polygon gives an acute triangulation. See Figure 7.
- Moreover, it appears that at least two vertices must be inserted in some cases. See Figure 8.

> **Conjecture** There exists an algorithm and some function $\theta(\gamma)$ which solve the problem above using at most two additional vertices.

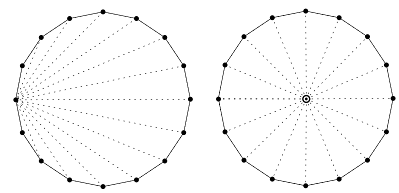

Figure 7: For a regular n-gon with n large, any triangulation without additional vertices includes a triangle of three consecutive vertices and thus a large angle. Adding a single vertex in the center yields an acute triangulation.

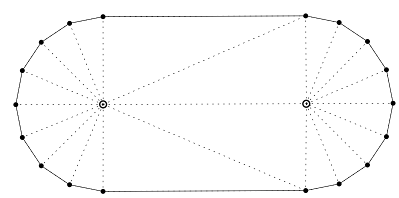

Figure 8: An example which requires two additional vertices (or at least appears to). This can be extended to any extreme γ or θ thresholds by adding more vertices to the semi-circles and making the full example wider in the horizontal direction.

Update. This problem is a special case of one less formally stated in the conclusion of [BDE92]. Specifically, Bern et al. ask for an algorithm and class of polygons yielding quality triangulations (i.e. without large angles) using only interior Steiner points.

References

[BDE92] M. Bern, D. Dobkin, and D. Eppstein. Triangulating polygons without large angles. *Proceedings of the 8th Annual Symposium on Computational Geometry*, pages 222–231, 1992.

[BMR95] M. Bern, S. Michell, and J. Ruppert. Linear-size nonobtuse triangulation of polygons. *Discrete Computational Geometry*, 14(1):411–428, 1995.

[MPS07] G. Miller, T. Phillips, and D. Sheehy. Size competitive meshing without large angles. In *34th International Colloquium on Automata, Languages and Programming*, pages 655–666, 2007.