Open Problems from CCCG 2003

Erik D. Demaine*   Joseph O'Rourke†

The following is a list of the problems presented on August 11, 2003 at the open-problem session of the 14th Canadian Conference on Computational Geometry held in Halifax, Nova Scotia, Canada.

Boxed problem numbers indicate appearance in The Open Problem Project (TOPP); see `http://www.cs.smith.edu/~orourke/TOPP/`.

Flipping to the 3D Delaunay Triangulation
Jonathan Shewchuk
Univ. California Berkeley
jrs@cs.berkeley.edu

Is there a local operator for modifying triangulations (tetrahedralizations) of a 3D point set along with a global objective function such that repeatedly applying local operations that improve the objective function ultimately leads to a unique global optimum equal to the 3D Delaunay triangulation? In 2D, edge flips are well-known to suffice, with the objective of lexically maximizing the vector of angles. Bistellar flips generalize to higher dimensions, but the resulting flip graph of configurations is not known to be connected in 3D (and is known to be disconnected in 5D). Thus the problem asks for an algorithm in the spirit of flipping, based purely on local transformations, that can convert any 3D triangulation to the Delaunay triangulation, and do so quickly if the triangulation is "almost" Delaunay. Ideally a solution could also be generalized to any dimension.

Update: Settled positively, in some sense, by Jonathan Shewchuk. The algorithm locally transforms the "star" around each vertex in the triangulation, allowing inconsistencies between the stars of different vertices. (A manuscript is in progress, available on request from the author.) It remains open whether the problem can be solved by local operators that preserve a triangulation at all times.

*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`. Supported by NSF Distinguished Teaching Scholars award DUE-0123154.

Variable-Speed Linear Morphing
Anna Lubiw
Univ. Waterloo
alubiw@uwaterloo.edu

This problem concerns "morphing" or "reconfiguring" graphs in the plane from one embedding to another while avoiding edge crossings throughout the motion. Precisely, an *embedding* of a graph $G = (V, E)$ is a function $f : V \to \mathbb{R}^2$. A *proper embedding* is an embedding with the property that no two edges e_1, e_2 intersect except at common endpoints: $f(e_1) \cap f(e_2) = f(e_1 \cap e_2)$ where $f(e_i)$ is the line segment between the f-mapped endpoints of e_i. A *morph* between two embeddings f_1 and f_2 of a common graph G is a continuum of embeddings of G that start at f_1 and end at f_2: $m : [0, 1] \to \{\text{embeddings of } G\}$ with $m(0) = f_1$ and $m(1) = f_2$. In a *proper morph* m, each embedding $m(t)$ must be proper, for $t \in [0, 1]$.

The *constant-speed linear morph* moves each vertex at a constant speed from the initial position to the target position: $m(t)(v) = (1-t)f_1(v) + tf_2(v)$. A *variable-speed linear morph* is a more general family of morphs where each vertex v monotonically traces out the line segment from the initial position $f_1(v)$ to the target position $f_2(v)$, but the speed may vary over time.

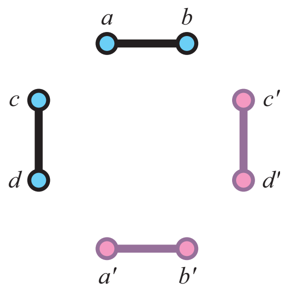

(a) Linearly morphable with variable speeds but not constant speed.

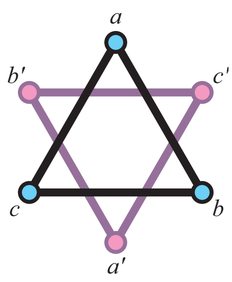

(b) Not linearly morphable even with variable speeds.

Figure 1: Examples of linear morphing.

Figure 1(a) shows an example of initial and target configurations that have proper variable-speed linear morphs, but for which the constant-speed linear morph is not proper. Figure 1(b) shows an example where no proper variable-speed linear morph is possible. What is the complexity of deciding whether a proper variable-speed linear morph exists for a given initial and target configuration of a given graph? What if the vertices must move one at a time, each going from its initial position to its target position?

Pointed Spanning Trees in Triangulations

Oswin Aichholzer
Graz Univ. Technology
oaich@ist.tugraz.at

TOPP #50

Does every triangulation of a set of points in the plane (in general position) contain a pointed spanning tree as a subgraph? A vertex is pointed if one of its incident faces has an angle larger than π at this vertex. A spanning tree is pointed if all of its vertices are pointed. The poser conjectures the answer to be YES. This is obviously true for any triangulation that contains a Hamiltonian path or a pointed pseudotriangulation as a subgraph. However, there exist triangulations not containing either of these structures. (See, e.g., [O'R02] for a discussion of pseudotriangulations.)

Update: Settled negatively by Aichholzer et al. [AHK04] with a 124-point counterexample.

References

[AHK04] O. Aichholzer, C. Huemer, and H. Krasser. Triangulations without pointed spanning trees, January 2004. In *Abstracts 20th European Workshop on Comput. Geom.*, Seville, 2004. `http://www.us.es/ewcg04/Articulos/krasser.ps`

[O'R02] J. O'Rourke. Computational geometry column 43. *Internat. J. Comput. Geom. Appl.*, 12(3):263–265, 2002. Also in SIGACT News, 33(1):58-60 (2002), Issue 122.

Monotonicity of Pointedness

Bettina Speckmann
TU Eindhoven
speckman@win.tue.nl

Let S be a set of points in the plane whose interior points (points not on the convex hull) are marked as either "pointed" or "nonpointed". See Figure 2.

Consider the number of possible pseudotriangulations of S respecting these marks; see the previous

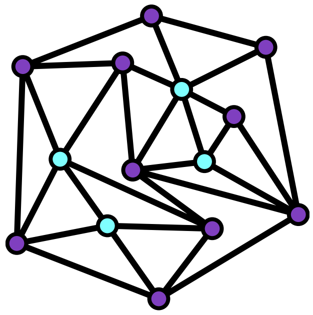

Figure 2: A pseudotriangulation with pointed points marked (dark).

problem for the definition of pointedness. Suppose some vertex $p \in S$ is marked as nonpointed.

Conjecture: When changing the mark of p to pointed, the number of possible pseudotriangulations of S respecting the marks of S increases.

The conjecture implies a partial order on the number of pseudotriangulations. If true, it would prove that there are always at least as many pointed pseudotriangulations as triangulations; this problem was posed by Jack Snoeyink during the CCCG 2001 open problem session, and appears as TOPP Problem 40.

The conjecture is not true if instead of changing the mark of one specific vertex we require only that the number of pointed vertices increases by one.

Moving a Ladder in a Forest

David Kirkpatrick
Univ. British Columbia
kirk@cs.ubc.ca

Given a collection of point obstacles in $\mathbb{R}^2$ ("trees") and two placements S and T of a unit-length line segment (the "ladder"), what is the complexity of finding the shortest rigid motion of the line segment from S to T that avoids all of the point obstacles? Here the *length* of a motion is measured as the arc length of the trace of a specified fixed point of the line segment (e.g., the midpoint). If the obstacles are polygonal, the problem is NP-hard [AKY03]. See [AKY96] [AKY02] for earlier work on special instances of the problem.

References

[AKY96] Tetsuo Asano, David Kirkpatrick, and Chee K. Yap. d_1-optimal motion for a rod. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 252–263, 1996.

[AKY02] Tetsuo Asano, David Kirkpatrick, and Chee K. Yap. Pseudo approximation algorithms, with applications to optimal motion planning. In *Proc. 18th Annu. ACM Sympos. Comput. Geom.*, pages 170–178, 2002.

[AKY03] Tetsuo Asano, David Kirkpatrick, and Chee Yap. Minimizing the trace length of a rod endpoint in the presence of polygonal obstacles is NP-hard. In *Proc. 15th Canad. Conf. Comput. Geom.*, Halifax, 2003, pp. 10–13. `http://torch.cs.dal.ca/~cccg/papers/54.pdf`

Flipping between Convex Decompositions
Ferran Hurtado
Univ. Politecnica de Catalunya
hurtado@ma2.upc.es

Fix a point set S, no three points of which are collinear, and fix a number k. Consider all convex decompositions with vertex set equal to S and with exactly k convex regions. Is it possible to reconfigure any such convex decomposition into any other by a sequence of $O(1)$-edge flips? A *j-edge flip* consists of replacing j edges with j different edges in such a way that the convex decomposition remains a convex decomposition with exactly k convex regions. Figure 3 shows examples of one-edge flips.

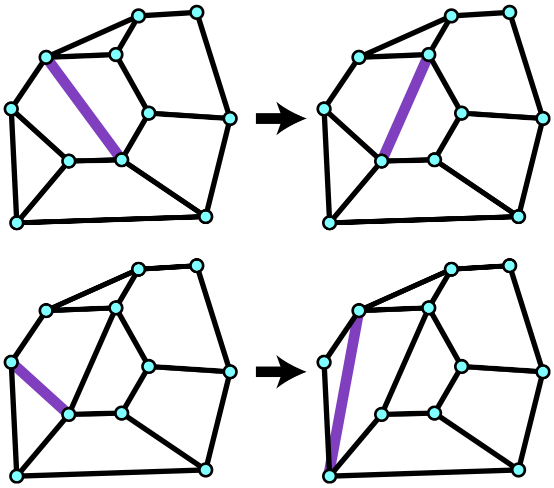

Figure 3: Examples of one-edge flips in convex decompositions with equal numbers of regions. The second flip cannot be made before the first flip.

Henk Meijer has shown that one-edge flips do not suffice; see Figure 4.

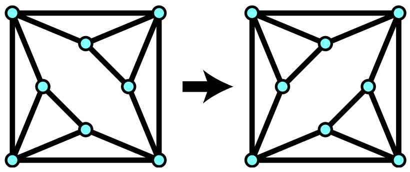

Figure 4: Two convex decompositions unreachable from each other by one-edge flips. Indeed, neither configuration has any allowable one-edge flip.

Update: Settled negatively by Henk Meijer and David Rappaport, as reported in this proceedings [MR04].

References

[MR04] Henk Meijer and David Rappaport. Simultaneous edge flips for convex subdivisions. In *Proc. 16th Canad. Conf. Comput. Geom.*, Montreal, 2004.

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

For a convex polyhedron of n vertices and F faces, what is the fewest number of nets (simple, nonoverlapping polygons) into which it may be cut along edges? Bounds both in terms of n and in terms of F would be of interest. An upper bound of F is trivial, but no general bound cF for $c < 1$ is known, despite the possibility that the answer is simply 1, independent of n or F. Two special cases have been established via matchings:

1. For triangulated (*simplicial*) polyhedra, $F/2$ pieces can be achieved by a perfect matching of pairs of triangles to quadrilaterals.
2. For *simple polyhedra*, those for which every vertex has degree 3, we can establish an upper bound of $2/3(F-2)$ pieces, again by matching to quadrilaterals.

Update: Shortly after the conference, Michael Spriggs established the first bound of cF for $c < 1$, proving $c = 2/3$. This bound was subsequently improved by Vida Dujmović, Pat Morin, and David Wood to $c = 1/2$. A bound sublinear in F or n remains elusive.

Painting a Polyhedron
Ferran Hurtado
Univ. Politecnica de Catalunya
hurtado@ma2.upc.es

Is it possible to "paint" any triangulated (simplicial) convex polyhedron? A *painting* of a polyhedron is an ordering of the facets $f_1, f_2, \ldots, f_n$ such that we can paint the faces red in that order and satisfy two properties at any time $1 \le t < n$:

1. the painted region $f_1 \cup f_2 \cup \cdots \cup f_t$ at time t is simply connected, as is the unpainted region $f_{t+1} \cup f_{t+2} \cup \cdots \cup f_n$; and

2. the painted-unpainted interface, i.e., the boundary of $f_1 \cup f_2 \cup \cdots \cup f_t$, consists of $O(\sqrt{n})$ edges.

The first constraint is equivalent to requiring the ordering of faces to be a *shelling order* (see, e.g., [Kal04]). Intuitively, the second constraint requires the wet edge of the painted surface to be small so as to allow a seamless paint job before that edge dries.

Update: Settled negatively by Richard Nowakowski and Norbert Zeh.

References

[Kal04] Gil Kalai. Polytope skeletons and paths. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry (Second Edition)*, chapter 20, pages 455–476. CRC Press LLC, Boca Raton, FL, 2004.