Open Problems from CCCG 2006

Erik D. Demaine* Joseph O'Rourke†

The following is a list of the problems presented on August 14, 2006 at the open-problem session of the 18th Canadian Conference on Computational Geometry held in Kingston, Ontario, Canada.

Edge Reconnections in Unit Chains
Anna Lubiw
University of Waterloo
alubiw@uwaterloo.ca

Is it possible to reconfigure between any two configurations of a unit-length 3D polygonal chain, or more generally unit-length 3D polygonal trees, by "edge reconnections"?

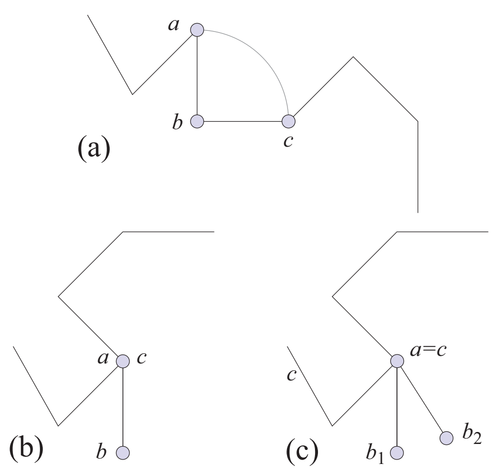

Figure 1: An edge reconnection.

An *edge reconnection* is an operation on a path of three consecutive vertices (a, b, c) of a chain or tree linkage; refer to Figure 1. If it is possible to fold the linkage, preserving edge lengths and avoiding self-intersection, to bring the two incident edges ab and bc into coincidence, then the edge reconnection breaks the connection between the edges at b, splits b into two vertices b_1 and b_2, and fuses $a{=}c$ together. This operation changes a chain into a tree. If we apply this operation to a tree, there will generally be attachments at b; allow them to attach to either b_1 or b_2. Is it possible to reconfigure between any two unit-length 3D trees of the same number of links using edge reconnections?

Update: At the conference, a related move was defined: a *tetrahedral swap* alters consecutive vertices (a, b, c, d) of a chain to either (a, c, b, d) or (a, d, b, c), whichever of the two yields a chain. It was established (by Erik Demaine, Anna Lubiw, and Joseph O'Rourke) that, if tetrahedral swaps are permitted without regard to self-intersection, then they suffice to unlock any chain.

Edge Swaps in Planar Matchings
Ferran Hurtado (via Henk Meijer)
Univ. Politécnica de Catalunya
Ferran.Hurtado@upc.edu

Is the space of (noncrossing) plane matchings on a set of n points in general position connected by "two-edge swaps"? A *two-edge swap* replaces two edges ab and cd of a plane matching M with a different pair of edges on the same vertices, say ac and bd, to form another plane matching M'. (The swap is valid only if the resulting matching is noncrossing.) Notice that ac and bd together with their replacement edges form a length-4 *alternating cycle* that does not cross any segments in M. Is it possible to reach any plane matching from any other on the same point set by a sequence of such two-edge swaps? In other words, if we form the graph where vertices correspond to plane matchings and edges correspond to two-edge swaps, is this graph connected?

This is a reposing of one question from a series of related questions posed at CCCG 2003 [DO04]; see also [MR04]. It is known that k-edge swaps suffice to connect the space of plane matchings, but only if no bound is imposed on k [HHNR05]. The same work also proved that a two-edge swap always exists, i.e., the two-edge-swap graph on plane matchings cannot have an isolated node. (The proof, however, was not included in the journal version of the paper.) For $n = 2m$ points in convex position, the two-edge-swap graph is well understood: its vertex connectivity is $m - 1$, it contains no Hamiltonian path for m odd and greater than 3, and it contains a Hamiltonian cycle for even values of m greater than 2 [HHN02].

*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

References

[DO04] Erik D. Demaine and Joseph O'Rourke. Open problems from CCCG 2003. In *Proc. 16th Canad. Conf. Comput. Geom.*, pages 209–211, 2004.

[HHN02] Carmen Hernando, Ferran Hurtado and Marc Noy. Graphs of non-crossing perfect matchings. *Graphs and Combinatorics* 18(3):517–532, 2002.

[HHNR05] Michael H. Houle, Ferran Hurtado, Marc Noy, Eduardo Rivera-Campo. Graphs of triangulations and perfect matchings. *Graphs and Combinatorics*, 21(3):325–331, 2005.

[MR04] Henk Meijer and David Rappaport. Simultaneous edge flips for convex subdivisions. In *Proc. 16th Canad. Conf. Comput. Geom.*, pages 57–59, 2004.

Shortest Aspect Tour

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

How quickly can one compute an exterior tour of a polyhedron that "views it from all different angles"? More precisely, the *aspect graph* of a polyhedron P represents the cells of the arrangement of planes containing the faces of P: from each cell, one sees a different "aspect" of P. The aspect graph has $\Theta(n^6)$ nodes [PD90]. The goal is to find a closed path π in $\mathbb{R}^3$ of minimum possible length that visits every aspect, i.e., π meets every cell of the arrangement, and does not penetrate the interior of the polyhedron P.

The same problem may be posed in the plane, and with or without the constraint that the path be a closed tour.

Update: The problem is related to the *external watchman route* problem, where the goal is to find a tour exterior to the polyhedron such that every point on the polyhedron's surface is visible from some point on the tour. This problem was first posed and solved in [NG94]; details even for planar convex polygons remain unresolved [AW06]. An aspect tour is a more stringent requirement, so in general is longer than the optimal external watchman route.

References

[AW06] Rafa Absar and Sue Whitesides. On computing shortest external watchman routes for convex polygons. In *Proc. 18th Canad. Conf. Comput. Geom.*, pages 59–62, 2006.

[NG94] S. Ntafos and L. Gewali. External watchman routes. *Visual Comput.*, 10:474–483, 1994.

[PD90] H. Plantinga and C. R. Dyer. Visibility, occlusion, and the aspect graph. *Internat. J. Comput. Vision*, 5(2):137–160, 1990.

The SS-Divide

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

Fix a nonvertex point p on the surface of a convex polyhedron P of n vertices. Draw shortest paths to all vertices, breaking ties arbitrarily. Connect the vertices by geodesics in the angular order of their shortest paths from p. See Figure 2. This geodesic

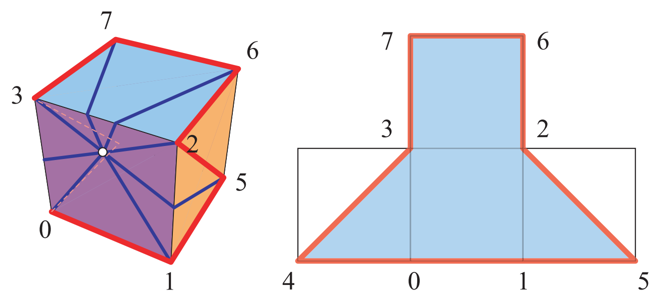

Figure 2: The ss-divide for a cube with respect to a point $\frac{3}{4}$ up the middle of the front face.

polygon is called the *ss-divide* in [DO07], because it divides the source unfolding from the star unfolding with respect to p. It partitions the surface into two halves, neither of which contains any vertices, and both of which unfold without overlap.

1. Under what conditions are both unfolded polygons convex?
2. Under what conditions are both unfolded polygons congruent? (They are in the example in Figure 2.)
3. It is established in [AAOS97] that there are $\Theta(n^4)$ distinct permutations of the vertices realized by ss-divides for different points p. Can the shortest ss-divide for a fixed polyhedron be computed in $o(n^4)$ time?
4. Given a set of permutations of vertices corresponding to all the ss-divides for some unknown polyhedron, how difficult is it to construct some polyhedron P that realizes those permutations?

References

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

[DO07] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, July 2007. http://www.gfalop.org.

Divide-and-Conquer: The Game

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

What is the outcome of the following two-player combinatorial game? The board is the 3D integer lattice, and a board position is a set S of unit-cube lattice cells marked as *solid*. For example, these cells might form an orthogonal polyhedron. The two-player game is impartial: both players can make the same moves from the same board position. In a move, a player chooses an integral coordinate plane that has at least one cube of S to either side, and removes from S all the cubes in one half-space bounded by the plane. The player who reduces S to the empty set (i.e., who cannot choose a plane with at least one cube on either side) loses. For which S can the first player force a win?

The problem should first be analyzed in 2D, as it is always an option for either player to reduce the problem to 2D at any step. The game generalizes to arbitrary dimensions.

Update: It was pointed out (by Erik Demaine) that this game is similar to the classic combinatorial game of *Chomp* [BCH03]. The initial configuration in Chomp is usually a solid box in the integer lattice. (Chomp is often played in the plane, but it has been analyzed in all dimensions.) One corner cube of the box, with minimum coordinates in all dimensions, is special and considered *poisoned*. In a move, a player chooses a solid cube c and removes all cubes whose coordinates are all at least the corresponding coordinates of c. This operation corresponds to removing all cubes in a nonempty octant (or quarter-plane in 2D) that extends to $+\infty$ in all coordinate axes. The player to eat the poisoned corner, and thus remove all remaining cubes, loses. Chomp has been studied extensively; one simple result is that the box initial configuration is always a first-player win. Of course, these results have no direct impact on Divide-and-Conquer.

On the other hand, the Sprague-Grundy theory of impartial games gives a polynomial-time algorithm for determining the winner from a position of Divide-and-Conquer. Specifically, a board position fitting in a $x \times y \times z$ box has at most $x^2y^2z^2$ future board positions; each can be associated with a "nimber" describing the game-theoretic value of the position; and the nimber of a position can be computed from the nimber of all its possible moves using the "mex rule". See [Dem01]. Is there a simpler characterization of the winner?

References

[BCH03] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning Ways for Your Mathematical Plays*. 2nd edition, A K Peters, 2003. Chomp is described in volume 3 on page 632.

[Dem01] Erik D. Demaine. Playing games with algorithms: Algorithmic combinatorial game theory. In *Proc. 26th Sympos. Math. Found. in Comput. Sci.*, LNCS 2136, August 2001, pages 18–32. arXiv:cs.CC/0106009

Blokus

Erik Demaine
MIT
edemaine@mit.edu

What is the computational complexity of the popular board game, Blokus[1]? Blokus is made by Sekkoia Sarl, based in France. In its original form, four players take turns placing colored polyominoes onto an initially empty 20×20 board. Each player starts with exactly one copy of each n-omino (edge-to-edge joining of n unit squares) for all $n \in \{1, 2, 3, 4, 5\}$. The first move of a player can place any polyomino in any position that overlaps the designated square for that player. Every subsequent move must place a polyomino (1) corner-adjacent to a piece of the same player, (2) not edge-adjacent to a piece of that player, and (3) disjoint from other players' pieces. A player can pass if and only if no piece can be placed. The game ends when no player can place any more pieces. A player's score is the total number of unit squares placed, plus small bonuses (+5 if the last piece placed is the monomino, +15 if all pieces are placed). The player with the maximum score wins.

This game can be generalized and simplified in many directions, making it suitable for complexity analysis. First, the number of players can be reduced to two (and indeed real-life Blokus has

[1] http://www.blokus.com

such variations). Second, the board size and piece-set size must be generalized to be part of the input. Third, the pieces can be chosen to be any desired set of polyominoes (or other polyforms such as polyiamonds). For example, each player may simply start with exactly n dominoes.

All such versions of Blokus are certainly in PSPACE, because the number of moves in polynomial in the input size. Are they PSPACE-complete?

Weighted-Region Shortest Path
Sada Nayanappa
University of Denver
snarayan@cs.du.edu

Is there a combinatorial solution to the weighted-region shortest-path problem when there are just two convex weighted regions? More precisely, suppose that the plane is divided into two regions according to a convex polygon P. The interior of P has *weight* $w > 1$, while the exterior of P has weight 1. Given two points s and t exterior to P, the goal is to find the shortest weighted path from s to t. Can an exact solution be computed in polynomial time? The problem already seems difficult when P is a triangle.

The general weighted-region shortest-path problem was considered and solved, up to a desired error factor of $1+\varepsilon$, in [MP91]. What is desired here is a combinatorial algorithm with an exact solution for this simple special case.

References

[MP91] Joseph S. B. Mitchell and Christos H. Papadimitriou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *J. ACM*, 38:18–73, 1991.

Optimal Perfect Matching
Jeff Phillips
Duke University
jeffp@cs.duke.edu

Given two point sets A and B in the plane, each of n points, how efficiently can we find a perfect matching between A and B and a rigid motion of A that together minimize the sum of the squared distances between matched points? We can factor out translation by aligning the centroids of A and B; then the rigid motion only needs to rotate.

Two related problems are easy to solve. If no rigid motions (rotations) are permitted, the problem reduces to minimum-weight bipartite matching in $K_{n,n}$, which the Hungarian Method solves in $O(n^3)$ time. On the other hand, if the matching is known, finding the rotation is easy. The Iterated Closest Point (ICP) algorithm is a heuristic for this problem, but it does not guarantee an optimal matching.