Geometric kth Shortest Paths: the Applet

John Hershberger
Mentor Graphics Corporation
john_hershberger@mentor.com

Valentin Polishchuk
Helsinki Institute for IT, CS Dept
University of Helsinki
valentin.polishchuk@helsinki.fi

Bettina Speckmann*
Dep. Mathematics and Computer Science
TU Eindhoven
b.speckmann@tue.nl

Topi Talvitie
Helsinki Institute for IT, CS Dept
University of Helsinki
topi.talvitie@helsinki.fi

Introduction. Computing shortest paths in a polygonal domain is a classic problem in computational geometry. Efficient algorithms for computing such paths use the *continuous Dijkstra* paradigm [2], which not only allows one to find the shortest path between two points but also computes the "shortest path map" from a given source—a structure enabling efficient queries of shortest paths to points in the domain.

The applet. This note accompanies an applet that illustrates how the continuous Dijkstra approach can be "taken to the next level" and used to compute the "kth shortest path map'.' We briefly outline the relevant notions below; for formal definitions and proofs please refer to [1].

Shortest path map. Let P be a polygonal domain with holes and let s be a point in P (indicated by a green dot); all paths are assumed to start from s. Two paths to a point $q \in P$ are *homotopically different* (or have different *homotopy types*) if one cannot be continuously deformed to the other without intersecting holes. The *1st shortest path* (or the *1-path*) to q is just the shortest path from s to q. Say that q is on a *1-wall* if there exist two (necessarily homotopically different) 1-paths to it. The *homotopic shortest path map* of P (or *1-map*) cuts P along 1-walls (the map is simply connected—there is a unique 1-path to every point in its interior). The figure shows a 1-map; 1-walls are red.

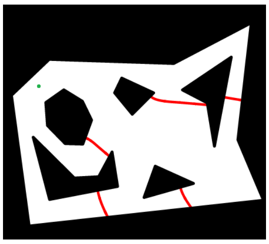

***k*-paths.** The *2nd shortest path* (or the *2-path*) is the shortest path homotopically different from the 1-path. The *kth shortest path* (the *k-path*) is defined recursively. Fig. 1 shows k-paths for $k = 1, \ldots, 5$ (the 5-path is nonsimple—it is equal to the 4-path plus the loop around the hole).

*Bettina Speckmann was supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.208.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

SoCG'14, June 8–11, 2014, Kyoto, Japan.
ACM 978-1-4503-2594-3/14/06.
http://dx.doi.org/10.1145/2582112.2595650

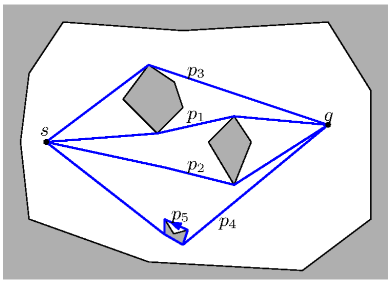

Figure 1: k-paths for $k = 1, \ldots, 5$.

Continuous Dijkstra. Imagine that a wave starts propagating from s at unit speed. At any time t the wavefront consists of the points at geodesic distance t from s. The wavefront is composed of circular-arc wavelets. When wavelets collide their intersection point traces a 1-wall. The wavelets do not propagate past the walls; i.e., the part of any wavelet on the other side of the wall dies. The figure shows a snapshot of the propagation; the area claimed by the wave is gray.

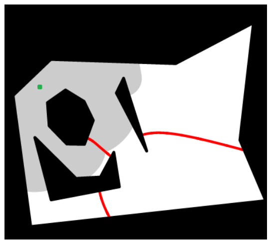

2-garage and partial wavelet resurrection. We allow wavelets to propagate beyond the walls; however, the propagation continues not in the domain P but on the "next floor." We define "parking garages" as follows: The 1-garage is just P. Recall that the 1-map is P sliced along 1-walls. The 1-map will be the *1st floor* (or 1-floor) of the *2-garage*. Take a copy of the 1-floor and put it on top of itself; the copy will be the *2-floor*—the last floor of the 2-garage. The 1- and 2-floors are glued along 1-walls; each side of any wall on the 1-floor is glued to the *opposite* side of the wall on the 2-floor. This way, when two wavelets meet at a 1-wall on the 1-floor, each continues propagating on the 2-floor (where the wavelets actually diverge, since they propagate into different sides of the wall). When wavelets collide on the 2-floor, their meeting point traces a *2-wall*. The figure shows wavelets on the 2-floor in the instance from the previous figure. The 1-walls are now blue, and the 2-walls are red.

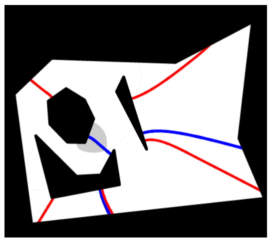

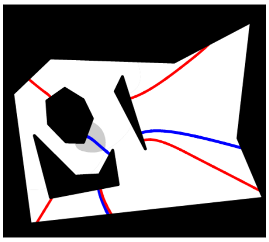

***k*-garage.** The k-garage for arbitrary k is defined recursively: Take a copy of P, slice it along $(k-1)$-walls, put the sliced copy on top of the $(k-1)$-garage, and glue the copy to the $(k-1)$-floor along $(k-1)$-walls, identifying the opposite sides of every wall. The copy is the k-floor—the top floor of the k-garage. When two wavelets meet at a $(k-1)$-wall on the $(k-1)$-floor, both wavelets continue propagating on the k-floor. When wavelets collide on the k-floor, their meeting point traces a k-wall. The next figure shows a snapshot of wave propagation on three floors. On a k-floor, the k-walls are red and the $(k-1)$-walls are blue:

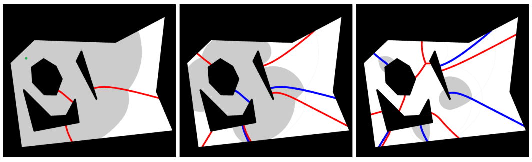

***k*-floor as *k*th homotopic shortest path map.** Note that k-walls and $(k-1)$-walls are comprised of the points that have two homotopically different k-paths. That is, if we cut P into cells along the walls, then k-paths to any point within one cell "have the same homotopy type" (homotopy types are originally defined only for paths with the same endpoint, but in [1] we extend the definition to compare homotopy types also for paths ending at different points). In this sense the cells define the *homotopic kth shortest path map* (or k-SPM), a generalization of the 1-SPM from 1-paths to k-paths for arbitrary $k > 1$. The figure below shows k-paths (in green) to a point, for $k = 1, 2, 3$:

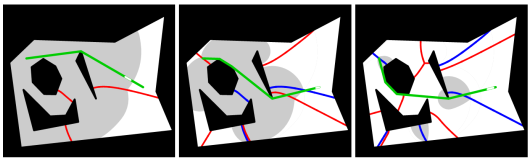

The next figure shows the k-paths to a point on the other side of the 1-wall; it can be seen that the 1-path and 2-path have "changed" their homotopy types:

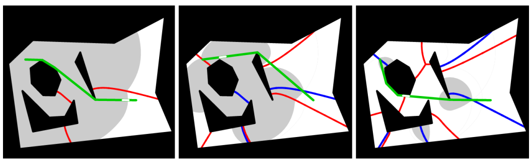

The next figure shows the k-paths to a point on the other side of the 2-wall; it can be seen that the 2-path and 3-path have "changed" their homotopy types:

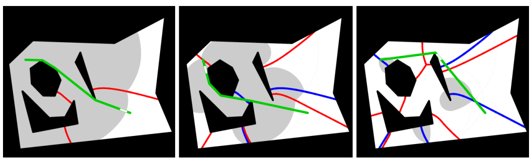

***k*-path as 1-path in the garage.** By construction, no wavelet collision happens in the garage until the top floor; the wavelets collide and trace walls only at the k-floor (these walls are 1-walls for the garage and k-walls for P). Hence, the shortest path map (1-map) in the garage is obtained simply by slicing the top floor along the k-walls. The most interesting structural property of the k-paths and the garage is the following: For any point q in P, let q' be the copy of q on the k-floor and let p' be the shortest path from q' to s in the garage; then the k-path to q in P is obtained by projecting p' onto the base sheet, P. That is, the path from q' to s starts on the k-floor, goes to a $(k-1)$-wall, uses it to get down to the $(k-1)$-floor, then uses the $(k-1)$-floor to reach a $(k-2)$-wall, uses it as a ramp down to the $(k-2)$-floor, and so on, until crossing a 1-wall to reach the 1-floor, on which the path goes to s. For example, the first canvas in the figure below shows the 3-floor and a 3-path; the path crosses a 2-wall (blue) on the 3-floor. The subpath after the crossing point is a 2-path, and the next canvas shows the 2-floor and the subpath; the subpath crosses a 1-wall (blue) on the 2-floor to reach the 1-floor. The rest of the path (shown in the last canvas) is just the shortest path (1-path) to s.

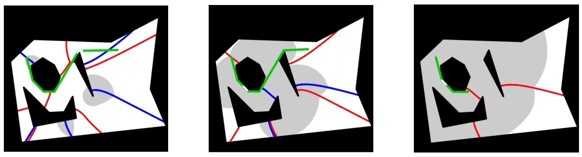

The applet: details. Using these definitions, we can state more precisely what the applet does: it shows how the continuous Dijkstra wavefront propagates in floors of the garage. Several canvases are shown; each is a floor. The 1st canvas shows wave propagation on the 1st floor of the garage (equivalently, wave propagation in P), the 2nd on the 2nd floor, etc. On the j-floor, for any j, the j-walls are red and the $(j-1)$-walls are blue. Hovering the mouse over a point shows the k-paths to it (green); a white worm moves along the path to signify how the path loops around holes. A slider is included to go back and forth in the wave propagation.

The applet can be found at `http://www.cs.helsinki.fi/group/compgeom/ksp/applet/`. The user can edit the domain, and when done, press *Propagate!* to start wave propagation. There is also a demo mode.

In the applet, the k-paths to the vertices of the domain are computed using the "simple visibility-based algorithm" in [1]. By casting beams toward the possible continuation paths from every vertex, the k-paths are computed for each pixel in a grid. The k-walls can be drawn approximately between pixels where k- and $(k+1)$-paths exchange identities. The wave propagation is shown by drawing the continuation beams of limited radius from every vertex and using the nonzero-filling rule of HTML5 canvas to draw the areas that have winding number of at least k on the kth level canvas.

REFERENCES

[1] S. Eriksson-Bique, J. Hershberger, V. Polishchuk, B. Speckmann, S. Suri, T. Talvitie, K. Verbeek, and H. Yıldız. Geometric kth shortest paths, 2013. Available at `http://www.cs.helsinki.fi/group/compgeom/ksp/socg2014_submission_32.pdf`.

[2] J. Hershberger and S. Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput.*, 28(6):2215–2256, 1999.