Programming Assignment (optional)

due: January 7th, 2016

Please submit the solution to this assignment directly to Efi Fogel (`efifogel@gmail.com`). Further details will be given in the course website.

Exercise: Given a simple polygon P with n vertices and a set G of points contained in the (closed) polygon P, we locate a camera at every point $g \in G$. A camera at point g sees a point p if the segment $\overline{gp}$ is contained in the (closed) polygon P.

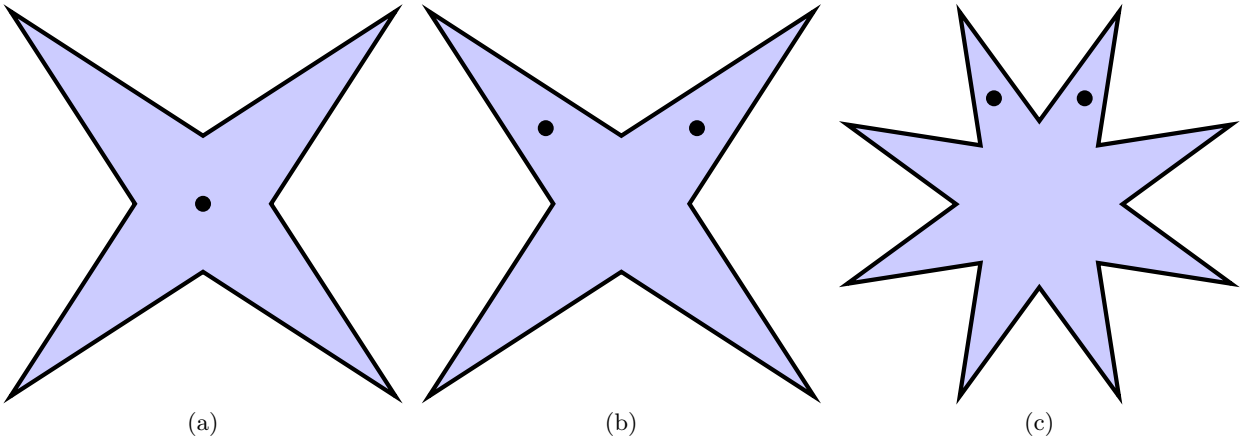

Figure 1: Examples of simple polygons and camera placements. (a) The single camera covers the polygon. (b) The two cameras cover the polygon. (c) The two cameras do not cover the polygon.

1. Design and implement an algorithm that decides whether G covers P; that is, whether for every point $p \in P$ there is at least one camera that sees p. If the answer is negative, namely G does not cover P, then the algorithm returns a point in the polygon that is not visible from any camera.

   In your implementation you may first assume general position, e.g., you may assume that none of the camera points coincide with any one of the polygon vertices and no three points are collinear. Next, improve your implementation to lift the general position restriction.

2. What are the time and space complexities of your algorithm?

Your implementation will be benchmarked for time efficiency and the results will be published in the course website.

References

[1] Efi Fogel, Ron Wein, and Dan Halperin. *CGAL Arrangements and Their Applications, A Step-by-Step Guide*. Springer, 2012.