

ON THE SHORTEST SEPARATING CYCLE*

Adrian Dumitrescu[†]

December 15, 2018

Abstract

According to a result of Arkin et al. (2016), given n point pairs in the plane, there exists a simple polygonal cycle that separates the two points in each pair to different sides; moreover, a $O(\sqrt{n})$ -factor approximation with respect to the minimum length can be computed in polynomial time. Here the following results are obtained:

(I) We extend the problem to geometric hypergraphs and obtain the following characterization of feasibility. Given a geometric hypergraph on points in the plane with hyperedges of size at least 2, there exists a simple polygonal cycle that separates each hyperedge if and only if the hypergraph is 2-colorable.

(II) We extend the $O(\sqrt{n})$ -factor approximation in the length measure as follows: Given a geometric graph $G = (V, E)$, a separating cycle (if it exists) can be computed in $O(m + n \log n)$ time, where $|V| = n$, $|E| = m$. Moreover, a $O(\sqrt{n})$ -approximation of the shortest separating cycle can be found in polynomial time. Given a geometric graph $G = (V, E)$ in \mathbb{R}^3 , a separating polyhedron (if it exists) can be found in $O(m + n \log n)$ time, where $|V| = n$, $|E| = m$. Moreover, a $O(n^{2/3})$ -approximation of a separating polyhedron of minimum perimeter can be found in polynomial time.

(III) Given a set of n point pairs in convex position in the plane, we show that a $(1 + \varepsilon)$ -approximation of a shortest separating cycle can be computed in time $n^{O(\varepsilon^{-1/2})}$. In this regard, we prove a lemma on convex polygon approximation that is of independent interest.

Keywords: Minimum separating cycle, traveling salesman problem, geometric hypergraph, 2-colorability, convex body approximation.

1 Introduction

Given a set of n pairs of points in the plane with no common elements, $\{(p_i, q_i) \mid i = 1, \dots, n\}$, a *shortest separating cycle* is a plane cycle (a closed curve, a.k.a. *tour*) of minimum length that contains inside exactly one point from each of the n pairs. The problem SHORTEST SEPARATING CYCLE is that of finding such a cycle, given the input pairs. It was introduced by Arkin et al. [3] motivated by applications in data storage and retrieval in distributed sensor networks. The authors gave a $O(\sqrt{n})$ -factor approximation for the general case and better approximations for some special cases. On the other hand, using a reduction from VERTEX COVER, they showed that the problem is hard to approximate for a factor of 1.36 unless $P = NP$, and is hard to approximate for a factor of 2 assuming the Unique Games Conjecture; see, e.g., [24, Ch. 16] for technical background.

The assumption that no point appears more than once, i.e., $|\{p_1, \dots, p_n\} \cup \{q_1, \dots, q_n\}| = 2n$, is sometimes necessary for the existence of a separating cycle; i.e., there are instances of sets of

*A preliminary version of this paper appeared in the *Proceedings of the 29th Canadian Conference on Computational Geometry*, Ottawa, ON, Canada, July 2017.

[†]Department of Computer Science, University of Wisconsin–Milwaukee, USA. Email: dumitres@uwm.edu

pairs with common elements and no separating cycle; see for instance Fig. 2 (the edges in these graphs represent pairs of input points). For convenience, points on the boundary of the cycle are considered inside; it is easy to see that requiring points to lie strictly in the interior or also on the boundary are equivalent variants in regards to the existence of a separating cycle. Moreover, the equivalence is almost preserved in the length measure: given any positive $\varepsilon > 0$, and a separating cycle C for n pairs, enclosing $P = \{p_1, \dots, p_n\}$ (after relabeling each pair, if needed), with some of the points of P on its boundary, a separating cycle of length at most $(1 + \varepsilon) \text{len}(C)$ can be constructed, having all points of P in its interior.

In this paper we study the extension of the concept of separating cycle to arbitrary graphs and hypergraphs, and to higher dimensions; in the original version introduced by Arkin et al. [3], the input graph is a *matching*, i.e., it consists of n edges with no common endpoints; see Fig. 1 for an example. Two instances with 8 and respectively 3 point pairs that do not admit separating cycles are illustrated in Fig. 2; the common reason is that both graphs contain odd cycles and odd cycles do not admit separating cycles.



Figure 1: A shortest separating cycle for a matching.

We observe that for arbitrary input graphs, one cannot use the algorithm from [3]. That algorithm (in [3, Subsec. 3.5]) first computes a minimum-size square Q containing at least one point from each pair, and then computes a constant-factor approximation of a shortest cycle (tour) of the points contained in Q , in the form of a simple polygon. In the end, this tour is refined to a separating cycle of the given set of point pairs with only a small increase in length. Here we note that there exist instances, such as that in Fig. 2 (right), for which there is no separating cycle confined to Q ; moreover, the length of a shortest separating cycle can be arbitrarily larger than any function of $\text{diam}(Q)$ and n , and so a new approach is needed for the general version with arbitrary input graphs, or its extension to hypergraphs; i.e., the current $O(\sqrt{n})$ -factor approximation does not carry through to these settings.




Figure 2: Left and center: instances with no separating cycle. Right: instance where the minimum axis-parallel square (or rectangle) that contains at least one point from each pair does not lead to a solution; a solution is indicated by the red cycle.

We first show that a planar geometric graph $G = (V, E)$ admits a separating cycle (for all its edge-pairs) if and only if it is bipartite. This result can be extended to hypergraphs in \mathbb{R}^d . Given a geometric hypergraph on points in \mathbb{R}^d with no singleton edges¹, there exists a simple polyhedron

¹A *singleton* edge is an edge with one vertex.

that separates each hyperedge if and only if the hypergraph is 2-colorable.

Definitions and notations. A *hypergraph* is a pair $H = (V, E)$, where V is finite set of *vertices*, and E is a family of subsets of V , called *edges*. H is said to be *2-colorable* if there is a 2-coloring of V such that no edge is monochromatic; see, e.g., [2, Ch. 1.3].

For a polygonal cycle C , let $\overset{\circ}{C}$ and \overline{C} denote the interior and exterior of C , respectively; and ∂C denote its boundary. Consider a geometric hypergraph $H = (V, E)$ on points in the plane with no singleton edges. A polygonal cycle C is said to be a *separating cycle* for H if (i) C is simple (i.e., with no self-intersections); and (ii) each edge of H has points inside C (in its interior or on its boundary) and points in the exterior of C ; that is, for each edge $A \in E$, both $A \cap (\overset{\circ}{C} \cup \partial C)$ and $A \cap \overline{C}$ are nonempty.

A simple polygonal cycle is said to have *zero area*, if $\text{Area}(C) \leq \varepsilon$, for a sufficiently small given $\varepsilon > 0$. Similarly, a polyhedron P is said to have *zero volume*, if $\text{Vol}(P) \leq \varepsilon$, for a sufficiently small given $\varepsilon > 0$.

For $r > 0$, let $B(r)$ denote the ball (i.e., disk in the plane) of radius r . For two convex bodies, A and B let $A + B$ denote their *Minkowski sum*, namely $A + B = \{a + b \mid a \in A, b \in B\}$.

Preliminaries and related work. Let S be a finite set of points in the plane. According to an old result of Few [11], the length of a minimum spanning path (resp., minimum spanning tree) of any n points in the unit square is at most $\sqrt{2n} + 7/4$ (resp., $\sqrt{n} + 7/4$). Both upper bounds are constructive; for example, the construction of a short spanning path works as follows. Lay out about \sqrt{n} equidistant horizontal lines, and then visit the points layer by layer, with the path alternating directions along the horizontal strips. In particular, the length of the minimum spanning tree of any n points in the unit square is bounded from above by the same expression. An upper bound with a slightly better multiplicative constant for a path was derived by Karloff [19]. Fejes Tóth [10] had observed earlier that for n points of a regular hexagonal lattice in the unit square, the length of the minimum spanning path is asymptotically equal to $(4/3)^{1/4}\sqrt{n}$, where $(4/3)^{1/4} = 1.0745\dots$. As such, the maximum length of the minimum spanning tree of any n points in the unit square is $\Theta(\sqrt{n})$, for a small constant (close to 1). The $O(\sqrt{n})$ upper bound also holds for points in a convex polygon of diameter $O(1)$, in particular for n points in a rectangle of diameter $O(1)$. In every dimension $d \geq 3$, Few showed that the maximum length of a shortest path (or tree) through n points in the unit cube is $\Theta(n^{1-1/d})$; the $O(n^{1-1/d})$ upper bound is again constructive and extends to rectangular boxes of diameter $O(1)$.

The topic of “separation” has appeared in multiple interpretations; here we only give a few examples: [1, 6, 7, 13, 15, 16, 17]. Some results on watchman tours relying on Few’s bounds can be found in [8]; others can be found in [4]. For instance, in the problem of finding a separating cycle for a given set of segment pairs, that we study here, it is clear that the edges of the cycle must hit all of the given segments. As such, this problem is related to the classic problem of hitting a set of segments by straight lines [16]. In a broader context, coloring of geometric hypergraphs has been studied, e.g., in [23].

2 Separating cycles for graphs and hypergraphs

By adapting results on hypergraph 2-colorability to a geometric setting, we obtain the following.

Theorem 1. *Let $H = (V, E)$ be a geometric hypergraph on points in the plane with no singleton edges. Then H admits a separating cycle if and only if H is 2-colorable.*

Proof. For the direct implication, assume that C is a separating cycle: then for each $A \in E$, both $A \cap \overset{\circ}{C}$ and $A \cap \overline{C}$ are nonempty. Color the points in the interior of C by red and those in its exterior by blue. As such, the hypergraph H is 2-colorable.

We now prove the converse implication. Let $V = R \cup B$ be a partition of the points into red and blue points, such that no edge in E is monochromatic. We construct a simple polygonal cycle containing only the red points in its interior. To this end, we first compute a minimum spanning tree T for the points in R ; T is non-crossing [22, Ch. 6], however there could be blue points contained in edges of T . Replace each such edge s with a two-segment polygonal path \tilde{s} connecting the same pair of points and lying very close to the original segment, and so that \tilde{s} is not incident to any other point.

The resulting tree, \tilde{T} is still non-crossing and spans all points in R . By doubling the edges of \tilde{T} and adding short connection edges, if needed, construct a simple polygonal cycle C of zero area that contains it and lies very close to it; as such, C contains all red points and none of the blue points, as required. \square

Since hypergraph 2-colorability is NP-complete [14], Theorem 1 yields the following.

Corollary 1. *Given a geometric hypergraph $H = (V, E)$ on points in the plane with no singleton edges, the problem of deciding whether H admits a separating cycle is NP-complete.*

We next present an approximation algorithm for computing a shortest separating cycle of a geometric graph. A key fact in our algorithm is the following observation.

Lemma 1. *Let G be connected bipartite graph. Then (apart from a color flip), G admits a unique 2-coloring.*

Proof. Recall that a graph is bipartite if and only if it contains no odd cycle [18, Ch. 3.3]. Consider an arbitrary vertex s and color it red. Then the color of any other vertex, say v , is uniquely determined by the parity of the length of the shortest path from s to v in G : red for even length and blue for odd length. Indeed, the vertices are colored alternately on any path, and since any cycle has even length, all lengths of paths from s to v have the same parity, as required. \square

Let $G = (V, E)$ be the input geometric graph no isolated vertices, where $|V| = n$, $|E| = m$. Let G_1, \dots, G_k denote the connected components of G , where $G_i = (V_i, E_i)$, for $i = 1, \dots, k$.

Theorem 2. (i) *Given a geometric graph $G = (V, E)$, a separating cycle (if it exists) can be computed in $O(m + n \log n)$ time, where $|V| = n$, $|E| = m$. (ii) *Further, a $O(\sqrt{n})$ -approximation of the shortest separating cycle can be found in polynomial time.**

Proof. (i) The graph is first tested for bipartiteness and the input instance is declared infeasible if the test fails (by Theorem 1). This test takes $O(m+n)$ time; see, e.g., [18, Ch. 3.3]. We subsequently assume that G is bipartite, with vertices colored by red and blue: $V = R \cup B$. Then the algorithm constructs a plane spanning tree T of the red points (for instance, a minimum spanning tree or a strictly monotone path), and outputs a simple cycle by doubling its edges and avoiding the blue points on its edges by bending those edges as indicated in the proof of Theorem 1.

To this end, the following parameters are computed: For each red point $r \in R$, $\delta_1(r)$ is the minimum distance to a blue point. For each edge e of T , $\delta_2(e) \geq 0$ is the minimum distance from e to a blue point ($\delta_2(e) = 0$ if e is incident to a blue point); and $\delta_3(e) > 0$ is the minimum nonzero distance from e to a blue point ($\delta_3(e) = \infty$ if no blue point is close to e , as explained below). The set of values $\delta_2(e), \delta_3(e)$, $e \in T$, are used for doubling T , and the set of values $\delta_1(r)$, $r \in R$, are

used to determine the separating cycle in the vicinity of red vertices; here we omit the details. The set of values $\delta_1(r)$, $r \in R$, and $\delta_2(e), \delta_3(e)$, $e \in T$, can be determined using point location for the blue points (as query points) in a planar triangulated subdivision containing the edges of T , all in $O(n \log n)$ time [5, Ch. 6]. The overall time complexity of the algorithm is $O(m + n \log n)$.

(ii) The algorithm above is modified as follows; the first step is the same bipartiteness test. The algorithm 2-colors the vertices in each connected component by red and blue: $V_i = R_i \cup B_i$, for $i = 1, \dots, k$. By Lemma 1, the 2-coloring of each component is unique (apart from a color flip). The initial coloring of a component may be subsequently subject to a color flip if the algorithm so later decides. Obviously, the coloring of each component is done independently of the others.

Then, the algorithm guesses the diameter of OPT , as determined by one of the $\binom{n}{2}$ pairs of points in V (by trying all such pairs). In each iteration, the algorithm may compute a separating cycle and record its length; the shortest cycle found in the process will be output by the algorithm; some iterations may be abandoned earlier, without the need for this calculation.

Consider the iteration in which the guess is correct, with pair $a, b \in V$; we may assume for concreteness that ab is a horizontal segment of unit length; refer to Fig 3. As such, we have that $\text{len}(\text{OPT}) \geq 2|ab| = 2$. In this iteration, the algorithm computes a separating cycle whose length is bounded from above by $O(\sqrt{n})$. First, the algorithm computes a rectangle Q of unit width and height $\sqrt{3}$ centered at the midpoint of ab . By the diameter assumption, OPT is contained in Q . In the next step the algorithm computes a separating cycle C containing only red points in Q in its interior (however, the initial coloring of some of the components may be flipped, as needed). Note that any separating cycle must contain for each component either all red points or all blue points but not a mix of two colors. By Lemma 1, the coloring of each component is unique (modulo a color flip) and so for each component at least one of its color classes is entirely contained in Q . As such, all points in V not contained in Q can be discarded from further consideration.




Figure 3: Left: input bipartite graph. Center: a separating cycle can be computed from the MST of the red points (after color flips). Right: a shortest separating cycle.

Each of the components G_i , $i = 1, \dots, k$ is checked against this containment condition: if a component is found where neither of its two color classes lies in Q , the algorithm abandons this iteration (and assumed diameter pair, $ab = \text{diam}(\text{OPT})$). For each component G_i : (i) if $R_i \subset Q$, then the coloring of this component remains unchanged, regardless of whether $B_i \subset Q$ or $B_i \not\subset Q$. (ii) if $R_i \not\subset Q$ and $B_i \subset Q$, then the coloring of this component is flipped: $R_i \leftrightarrow B_i$, so that $R_i \subset Q$ after the color flip.

Once the recoloring of components is complete, the algorithm computes a minimum spanning

tree T of the red points in Q . Its length is bounded from above by the length of the spanning tree computed by Few’s algorithm. Since the number of red points does not exceed n , we have $\text{len}(T) = O(\sqrt{n})$. Finally, T is converted into a separating cycle C by a factor of at most $2 + \varepsilon$ increase in length, for any given $\varepsilon > 0$, as in the proof of part (i). Recalling that $\text{len}(\text{OPT}) \geq 2$, it follows that C is a $O(\sqrt{n})$ -factor approximation of a shortest separating cycle. \square

An instance on which our algorithm—as well as that of Arkin et al. [3]—performs poorly appears in Fig. 4.




Figure 4: An instance for which the $O(\sqrt{n})$ -approximation ratio is tight; the left and respectively the right segment endpoints are each enclosed in a unit square and form the vertices of a $\sqrt{n} \times \sqrt{n}$ uniform grid. Left: a shortest separating cycle. Right: a minimum spanning tree of the red points (as basis for the separating cycle constructed by the algorithm).

3 Separating cycles for matchings in convex position in the plane

A matching of n point pairs is said to be in *convex position* if the $2n$ points are in convex position. In this section we develop a polynomial time approximation scheme (PTAS) for this setting; given n point pairs in convex position and $\varepsilon > 0$, the algorithm computes a $(1 + \varepsilon)$ -approximation of a shortest separating cycle. Denoting an optimal solution by OPT , note that OPT is a convex polygon with n vertices. Moreover, observe that a shortest separating cycle is a shortest TSP tour for the set of neighborhood pairs $\{p_i, q_i\}$, $i = 1, \dots, n$; see [20, Sec. 7.4] for an overview of the traveling salesman problem (TSP).

Theorem 3. *Given a set of n point pairs in convex position, a $(1 + \varepsilon)$ -approximation of a shortest separating cycle can be computed in time $n^{O(\varepsilon^{-1/2})}$.*

We need the following technical lemma for convex polygon approximation. Our lemma is clearly of independent interest; while it answers a basic question, we could not find such a result in the literature; there exist however related results, see, e.g., [12].

Lemma 2. *Given a convex polygon P and any $\varepsilon > 0$, there exists a subpolygon $Q \subset P$ with $O(1/\sqrt{\varepsilon})$ vertices, such that $P \subset \text{conv}(Q) + B(\varepsilon \cdot \text{diam}(P))$. Apart from the constant factor, this bound cannot be improved.*

Proof. Observe that the number of vertices of Q does not depend on the number of vertices of P , it only depends on ε . We will assume without loss of generality that $\text{diam}(P) = 1$. Let $P = p_1, p_2, \dots, p_m$ be the vertices of P labeled clockwise.

First construct a subpolygon $R \subset P$ iteratively. Set $i \leftarrow 1$ and include p_i into R . Scan P clockwise until we find the first vertex, p_j such that at least one vertex among p_{i+1}, \dots, p_{j-1} is at distance at least ε from the chord $p_i p_j$. Include p_j into R . Set $i \leftarrow j$ and continue in the same

manner, scan P clockwise until we find the first vertex, p_j such that at least one vertex among p_{i+1}, \dots, p_{j-1} is at distance at least ε from the chord $p_i p_j$, and so on. Suppose that the scanning ends after r phases; then R has either r or $r + 1$ vertices; see Fig. 5 (left). Note that the last side of R can be unconstrained, i.e., with no guarantee of a vertex of P at distance at least ε from it.




Figure 5: Convex polygon approximation; here $m = 21$, $r = 5$. $R = \{p_1, p_4, p_8, p_{13}, p_{17}\}$ is drawn in blue lines. $Q = \{p_1, p_3, p_4, p_7, p_8, p_{12}, p_{13}, p_{16}, p_{17}, p_{21}\}$ is drawn in magenta lines. The angles $\alpha(\sigma)$ are shaded (left side of figure): $\alpha(p_1 p_4) = \angle p_2 p_1 p_4$, $\alpha(p_4 p_6) = \angle p_5 p_4 p_8$, and so on.

Each side of R is a chord or side of P . A side of R is said to be *short* if its length is at most $\sqrt{\varepsilon}$ and *long* otherwise. Write $r = r_1 + r_2$, where r_1 and r_2 are the number of short and long sides of R , respectively.

Since $\text{diam}(P) = 1$, we have $\text{per}(P) \leq \pi$ by the classic isoperimetric inequality. Since R is a subpolygon of P , we have $\text{per}(R) \leq \text{per}(P)$ by the triangle inequality; thus $\text{per}(R) \leq \pi$. This further implies that $r_2 \leq \text{per}(R)/\sqrt{\varepsilon} \leq \pi/\sqrt{\varepsilon}$.

When scanning P clockwise, let σ be a side of R that is a nontrivial chord of P ; let $\alpha(\sigma)$ denote the angle made by the chord with the first clockwise edge of P on the boundary of P . By convexity, the angles $\alpha(\sigma)$ corresponding to all nontrivial chords of P that are edges in R are pairwise non-overlapping if their apices are placed at a common point; see Fig. 5 (left). As such,

$$\sum_{\sigma \in R} \alpha(\sigma) \leq 2\pi. \quad (1)$$

Recall that $\tan x \leq 2x$, for $0 \leq x \leq \pi/3$. Consider any short side $\sigma \in R$ that is a chord of P ; we thus have

$$2\alpha(\sigma) \geq \tan \alpha(\sigma) \geq \frac{\varepsilon}{|\sigma|} \geq \frac{\varepsilon}{\sqrt{\varepsilon}} = \sqrt{\varepsilon}, \quad \text{or } \alpha(\sigma) \geq \pi/3. \quad (2)$$

It follows from (1) and (2) that

$$r_1 \leq \frac{2\pi}{\sqrt{\varepsilon}/2} + \frac{2\pi}{\pi/3} = \frac{4\pi}{\sqrt{\varepsilon}} + 6.$$

Consequently,

$$r = r_1 + r_2 \leq \frac{4\pi}{\sqrt{\varepsilon}} + 6 + \frac{\pi}{\sqrt{\varepsilon}} = \frac{5\pi}{\sqrt{\varepsilon}} + 6.$$

To obtain Q we further subdivide each polygonal arc of P spanned by edges of R (with the possible exception of the last arc) as follows: the arc $p_i p_j$ is subdivided at p_{j-1} , that is, this vertex is included along with p_i and p_j into Q . Observe that Q has at most

$$2(r + 1) \leq \left\lceil \frac{10\pi}{\sqrt{\varepsilon}} + 14 \right\rceil = O(\varepsilon^{-1/2})$$

vertices. (In particular, Q has at most $\frac{32}{\sqrt{\varepsilon}}$ vertices, if $\varepsilon > 0$ is sufficiently small.) In addition, by construction, Q 's enlargement contains P : recall that vertices in R are iteratively chosen so that the corresponding arcs of P are minimally breaking the inclusion property stated in the lemma, and so the subdivision of each arc described above will satisfy this property. In particular, each vertex of P is at distance at most ε from the corresponding side of Q .

The case of a regular polygon shows that the bound is tight. Let P be a regular m -gon inscribed in a circle of unit radius, and $Q \subset P$ be a subpolygon, such that every vertex of P is at distance at most ε from the corresponding side of Q . Let 2β be the center angle spanned by the longest edge of Q ; for simplicity, assume that the number of vertices of P on the corresponding polygonal arc is odd (the other case is similar). The distance condition requires $1 - \cos \beta \approx \beta^2/2 = \varepsilon$, which solves to $\beta = \sqrt{2\varepsilon}$. This implies that the number of vertices of Q is $\Omega(\beta^{-1}) = \Omega(\varepsilon^{-1/2})$, as required. \square

Remark. If $P \subset \text{conv}(Q) + B(\varepsilon \cdot \text{diam}(P))$, i.e., Q is a suitable approximation as required in Lemma 2, then $\text{diam}(Q) \geq \text{diam}(P)/(1 + 2\varepsilon)$, and consequently,

$$P \subset \text{conv}(Q) + B(\varepsilon(1 + 2\varepsilon) \cdot \text{diam}(Q)).$$

Indeed, $\varepsilon \cdot \text{diam}(P) \leq \varepsilon(1 + 2\varepsilon) \cdot \text{diam}(Q)$, from which the inclusion follows.

Assume now that P is the *unknown* convex polygon OPT. Set $k = \lceil \frac{10\pi}{\sqrt{\varepsilon}} + 14 \rceil$. By convexity, P contains exactly one point from each pair; as such, $|P| = n$. The algorithm finds a subpolygon $Q \subset P$ satisfying the property in Lemma 2 by generating all subpolygons Q with at most k vertices. It does so by generating all $\binom{n}{\leq k}$ subsets of at most k segments; for a subset of $i \leq k$ segments, it goes through all 2^i possible endpoint selections (one point from each pair). The reason is that the shortest separating cycle of some $i \leq k$ endpoints (one point from each pair) may be an infeasible candidate for Q ; see Fig. 6, but generating shortest separating cycles for each choice of $i \leq k$ endpoints will yield a feasible candidate, as required by Lemma 2.




Figure 6: Left: a set of seven pairs with points in convex position. Right: the shortest separating cycle for the pairs $\{p_i, q_i\}$, $i = 1, 2, 3$, $Q = \{q_1, q_2, q_3\}$ (drawn in red) is an infeasible candidate, i.e., $\text{conv}(Q) + B(\varepsilon \cdot \text{diam}(\text{OPT}))$ does not contain the optimal convex polygon $P = \text{OPT} = \{p_1, p_6, q_4, p_3, q_5, p_7, p_2\}$. On the other hand, $Q = \{p_1, p_3, p_2\}$ (drawn in blue) is a valid candidate.

Algorithm A1.

STEP 1: Set $k = \left\lceil \frac{10\pi}{\sqrt{\varepsilon}} + 14 \right\rceil$. Generate all $\sum_{i=1}^k \binom{n}{i}$ subsets of at most k segments.

STEP 2: For a subset of size i , go through all 2^i endpoint selections (one point from each pair).

STEP 3: For a subset Q of $i \leq k$ points as described above, if $\text{conv}(Q) + B(\varepsilon(1 + 2\varepsilon) \cdot \text{diam}(Q))$ contains at least one point from each of the n input pairs, keep one such point from each pair (if both points of a pair are enclosed, choose one arbitrarily); then compute the perimeter of the convex polygon made by these n points, i.e., the length of the separating cycle. Otherwise skip this subset.

STEP 4: Return the cycle of minimum length from among those computed in STEP 3.

The running time of Algorithm A1 is determined by the number of candidates examined, namely

$$\sum_{i=1}^k \binom{n}{i} 2^i \leq n^k = n^{O(\varepsilon^{-1/2})}.$$

The algorithm correctness follows from Lemma 2; indeed, $P = \text{OPT}$ is contained in

$$\text{conv}(Q) + B(\varepsilon \cdot \text{diam}(P)) \subset \text{conv}(Q) + B(\varepsilon(1 + 2\varepsilon) \cdot \text{diam}(Q)),$$

namely the enlargement of one of the candidates Q that are generated. The length of the separating cycle that is returned is bounded from above by the perimeter of the enlargement. We employ the following standard fact.

Lemma 3. *Let Q be a planar convex body. Then $\text{per}(\text{conv}(Q) + B(r)) = \text{per}(Q) + 2\pi r$.*

Proof. Let $w(\alpha)$ denote the width of Q in direction α , i.e., the minimum width of a strip of parallel lines enclosing Q , whose lines are orthogonal to direction α . According to Cauchy's surface area formula [21, pp. 283–284], for any planar convex body Q , we have

$$\int_0^\pi w(\alpha) \, d\alpha = \text{per}(Q). \tag{3}$$

Observe that the width of $\text{conv}(Q) + B(r)$ in direction α is $w(\alpha) + 2r$, for any $\alpha \in [0, \pi)$. Using the stated formula we have

$$\text{per}(\text{conv}(Q) + B(r)) = \int_0^\pi (w(\alpha) + 2r) \, d\alpha = \int_0^\pi w(\alpha) \, d\alpha + 2\pi r = \text{per}(Q) + 2\pi r,$$

as required. □

By Lemma 3, the length of the separating cycle C returned by Algorithm A1 is bounded from above (for ε sufficiently small) by

$$\begin{aligned} \text{len}(C) &\leq \text{per}(\text{conv}(Q) + B(\varepsilon(1 + 2\varepsilon) \cdot \text{diam}(Q))) \\ &= \text{per}(Q) + 2\pi\varepsilon(1 + \varepsilon) \cdot \text{diam}(Q) \\ &\leq \text{per}(P) + 2\pi\varepsilon(1 + \varepsilon) \cdot \text{diam}(P) \\ &\leq \text{per}(P) + 2\pi\varepsilon(1 + \varepsilon) \cdot \text{per}(P)/2 \\ &= (1 + \pi\varepsilon(1 + \varepsilon)) \text{per}(P) = (1 + \pi\varepsilon(1 + \varepsilon)) \text{len}(\text{OPT}) \\ &\leq (1 + 4\varepsilon) \text{len}(\text{OPT}). \end{aligned}$$

The required approximation follows by rescaling ε , and this completes the proof of Theorem 3.

4 Concluding remarks

Remark 1. If the input is a set of pairs so that the corresponding graph is bipartite, it admits a separating cycle by Theorem 1. (If the corresponding graph is not bipartite, no separating cycle exists.) Similarly, if the input is a 2-colorable hypergraph, it admits a separating cycle. For illustration, we recall some common instances of 2-colorable hypergraphs. A hypergraph $H = (V, E)$ is called k -uniform if all $A \in E$ have $|A| = k$. A random 2-coloring argument gives that any k -uniform hypergraph with fewer than 2^{k-1} edges is 2-colorable [2, Ch. 1.3]; as such, by Theorem 1, it admits a separating cycle. Slightly better bounds have been recently obtained; see [2, Ch. 3.5]. Similarly, let $H = (V, E)$ be a hypergraph in which every edge has size at least k and assume that every edge $A \in E$ intersects at most Δ other edges, i.e., the maximum degree in H is at most Δ . If $e(\Delta + 1) \leq 2^{k-1}$ (here $e = \sum_{i=0}^{\infty} 1/i!$ is the base of the natural logarithm), then by the Lovász Local Lemma, H can be 2-colored [2, Ch. 5.2] and so by Theorem 1, it admits a separating cycle; moreover, if a 2-coloring is given, it can be used to obtain a separating cycle. While testing for 2-colorability can be computationally expensive in a general setting (recall that hypergraph 2-colorability is NP-complete [14]), it can be always achieved in exponential time.

Remark 2. Theorem 2 generalizes to 3-dimensional polyhedra. A polyhedron in 3-space is a simply connected solid bounded by piecewise linear 2-dimensional manifolds. The *perimeter* $\text{per}(P)$ of a polyhedron P is the total length of the edges of P (as in [8]).

For part (i), a method similar to that used in the planar case can be used to construct a separating polyhedron in \mathbb{R}^3 (or \mathbb{R}^d). However, since computing minimum spanning trees in \mathbb{R}^3 is more expensive [9, Ch. 9], we employ a slightly different approach (in particular, this approach is also applicable to the planar case). We may assume a coordinate system so that no pair of points have the same x -coordinate. First, the points in V are colored by red or blue as a result of the bipartiteness test, in $O(m + n)$ time. The algorithm then computes a (spanning tree of the red points in the form of a) x -monotone polygonal path P spanning the red points; this step takes $O(n \log n)$ time. From P , it then obtains a x -monotone polygonal path \tilde{P} spanning the red points and not incident to any blue point ($P = \tilde{P}$ if no blue points are incident to edges of P); \tilde{P} is constructed in $O(n \log n)$ time.

To this end, P and all blue points are projected onto the xoy plane. Let $\sigma(\cdot)$ denote the projection function. Note that $\sigma(P)$ is x -monotone and that the projection $\sigma(b)$ of a blue point b can be incident to at most one edge of $\sigma(P)$; given b , such an edge can be determined in $O(\log n)$ time by binary search. Checking the projection points $\sigma(b)$ against corresponding edges of $\sigma(P)$ allows for testing whether the original edges of P are incident to the respective blue points. Further, this test allows replacing each such edge s with a two-segment polygonal path \tilde{s} connecting the same pair of points and lying very close to the original segment, and so that \tilde{s} is not incident to any other point; see Fig 7. Such a replacement can be executed in $O(1)$ time per edge. Finally the algorithm computes a polyhedron of *zero volume* that contains \tilde{P} ; as such, the polyhedron contains all red points but no blue points; this step takes $O(n \log n)$ time.




Figure 7: Replacing two edges of the x -monotone path spanning the red points.

For part (ii), instead of a rectangle based on segment ab as an assumed diameter pair, the algorithm works with a rectangular box where ab is parallel to a side of the box and is incident to its center. The upper bound on the perimeter of the separating polyhedron follows from Few's bound mentioned in the preliminaries: it is roughly three times the length of a shortest path (or tree) spanning the red points.

Theorem 4. (i) *Given a geometric graph $G = (V, E)$ in \mathbb{R}^3 , a separating polyhedron (if it exists) can be found in $O(m + n \log n)$ time, where $|V| = n$, $|E| = m$. (ii) *Further, a $O(n^{2/3})$ -approximation of a separating polyhedron of minimum perimeter can be found in polynomial time.**

We offered a characterization of geometric hypergraphs that admit separating cycles and gave several approximation algorithms. We conclude with the following questions regarding the shortest separating cycle in the plane.

1. Can the $O(\sqrt{n})$ approximation factor for the general version of the problem be improved?
2. Can sharper results be obtained for plane (noncrossing) geometric graphs? For the case of a plane matching?
3. What is the computational complexity of the problem for matchings in convex position? Does the problem admit a polynomial-time algorithm?

Acknowledgment. The author is grateful to an anonymous reviewer for his careful reading of the manuscript and pertinent remarks.

References

- [1] N. Alon, Z. Füredi, and M. Katchalski, Separating pairs of points by standard boxes, *European Journal of Combinatorics* **6(3)** (1985), 205–210.
- [2] N. Alon and J. Spencer, *The Probabilistic Method*, 4th edition, Wiley, Tel Aviv and New York, 2015.
- [3] E. Arkin, J. Gao, A. Hesterberg, J. S. B. Mitchell, and J. Zeng, The shortest separating cycle problem, *Proc. 14th International Workshop on Approximation and Online Algorithms (WAOA 2016)*, vol. 10138 of Lecture Notes in Computer Science, 2016, pp. 1–13.
- [4] E. M. Arkin, J. S. B. Mitchell and C. D. Piatko, Minimum-link watchman tours, *Information Processing Letters* **86** (2003), 203–207.
- [5] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd edition, Springer, 2008.
- [6] S. Cabello and P. Giannopoulos, The complexity of separating points in the plane, *Algorithmica* **74(2)** (2016), 643–663.
- [7] G. Călinescu, A. Dumitrescu, H. Karloff, and P.-J. Wan, Separating points by axis-parallel lines, *International Journal of Computational Geometry & Applications* **15(6)** (2005), 575–590.
- [8] A. Dumitrescu and C. D. Tóth, Watchman tours for polygons with holes, *Computational Geometry: Theory and Applications*, **45** (2012), 326–333.

- [9] D. Eppstein, Spanning trees and spanners, in J.-R. Sack and J. Urrutia (editors), *Handbook of Computational Geometry*, pages 425–461, Elsevier Science, Amsterdam, 2000.
- [10] L. Fejes Tóth, Über einen geometrischen Satz, *Mathematische Zeitschrift* **46** (1940), 83–85.
- [11] L. Few, The shortest path and shortest road through n points, *Mathematika* **2** (1955), 141–144.
- [12] R. Fleischer, K. Mehlhorn, G. Rote, E. Welzl, and C.-K. Yap, Simultaneous inner and outer approximation of shapes, *Algorithmica* **8(5-6)** (1992), 365–389.
- [13] R. Freimer, J.S.B. Mitchell, and C. Piatko, On the complexity of shattering using arrangements. Extended abstract in *Proc. of the 2nd Canadian Conference on Computational Geometry* (CCCG 1990), 1990, pp. 218–222. Technical Report 91-1197, Dept. of Comp. Sci., Cornell University, April, 1991.
- [14] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman and Co., New York, 1979.
- [15] D. R. Gaur, T. Ibaraki, and R. Krishnamurti, Constant ratio approximation algorithm for the rectangle stabbing problem and the rectilinear partitioning problem, *Journal of Algorithms* **43** (2002), 138–152.
- [16] R. Hassin and N. Megiddo, Approximation algorithms for hitting objects with straight lines, *Discrete Applied Mathematics* **30(1)** (1991), 29–42.
- [17] F. Hurtado, M. Noy, P. A. Ramos, and C. Seara, Separating objects in the plane with wedges and strips, *Discrete Applied Mathematics* **109** (2001), 109–138.
- [18] D. Jungnickel, *Graphs, Networks and Algorithms*, Springer Verlag, Berlin, 1999.
- [19] H. J. Karloff, How long can a Euclidean traveling salesman tour be? *SIAM Journal on Discrete Mathematics* **2** (1989), 91–99.
- [20] J. S. B. Mitchell, Geometric shortest paths and network optimization, in J.-R. Sack and J. Urrutia (editors), *Handbook of Computational Geometry*, pages 633–701, Elsevier Science, Amsterdam, 2000.
- [21] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, John Wiley, New York, 1995.
- [22] F. P. Preparata and M. I. Shamos, *Computational Geometry*, Springer-Verlag, New York, 1985.
- [23] S. Smorodinsky, On the chromatic number of geometric hypergraphs, *SIAM Journal on Discrete Mathematics* **21(3)** (2007), 676–687.
- [24] D. Williamson and D. Shmoys, *The Design of Approximation Algorithms*, Cambridge University Press, 2011.