4.88pt@
International Journal of Computational Geometry & Applications

Minimum Polygon Transversals of Line Segments

David Rappaport*

Department of Computing and Information Science, Queen's University
Kingston, Ontario K7L 3N6 CANADA

ABSTRACT

Let S be used to denote a finite set of planar geometric objects. Define a *polygon transversal* of S as a closed simple polygon that simultaneously intersects every object in S, and a *minimum polygon transversal* of S as a polygon transversal of S with minimum perimeter. If S is a set of points then the minimum polygon transversal of S is the convex hull of S. However, when the objects in S have some dimension then the minimum polygon transversal and the convex hull may no longer coincide. We consider the case where S is a set of line segments. If the line segments are constrained to lie in a fixed number of orientations we show that a minimum polygon transversal can be found in $O(n \log n)$ time. More explicitely, if m denotes the number of line segment orientations, then the complexity of the algorithm is given by $O(3^m n + \log n)$. The general problem for line segments is not known to be polynomial nor is it known to be NP-hard.

Keywords: Computational geometry, transversal, simple polygon, convex hull.

1. Introduction

The problem of intersecting a collection of objects with a common line has received considerable attention in the area of discrete and computational geometry. Such a line is known as a *line transversal* in the mathematics literature, or a *line stabber* in the computer science literature.[20] Edelsbrunner et. al. show that a stabbing line for a set of n line segments can be determined in $O(n \log n)$ time.[10] O'Rourke shows that in $O(n)$ time suffices to find a line stabber if the line segments are parallel.[16] In[1,9,11] algorithms are given for stabbing collections of simple objects with a line. Efficient algorithms for stabbing lines, line segments and polyhedra with a line in three dimensions are also known.[3,17] Avis and Doskas present a general approach based on linear programming for stabbing d-dimensional polyhedra with a $d-1$ hyperplane.[2] Houle et. al. use a linear programming approach for stabbing hyperspheres with hyperplanes.[13] Optimizing the length of stabbers is discussed in[4] where efficient algorithms are given for computing a shortest line segment stabber of a set of planar line segments. In[5] the smallest radius disc that intersects a set of line segments in the plane is found in linear time. Jones and Ke find a *maximal stabbing* of planar line segments in which the maximum number of line segments is stabbed by a fixed length stabber.[14] Goodrich and Snoeyink present an algorithm that determines whether a set of parallel line segments can be stabbed by the boundary of a convex polygon[12,19]. Meijer and Rappaport allow the interior and the boundary of the polygon to stab the set S of parallel line segments, and find a stabbing polygon of smallest perimeter called a *minimum stabbing polygon* of S in $O(n \log n)$ time.[15] Czyzowicz et. al. present a linear time algorithm to find a minimum stabbing polygon for the case when the line segments are edges of a polygon.[7]

2. Preliminaries

Consider a set of line segments $S = \{s_1, s_2, \ldots 1, s_n\}$. A simple polygon $\mathcal{P}$ is a *polygon transversal* of S, if for all $s \in S$, $\mathcal{P} \cap s \neq \emptyset$. That is, every line segment in S has at least one point in the interior or on the boundary of $\mathcal{P}$. A simple polygon $\mathcal{P}$ is a *minimum polygon transversal* of S if $\mathcal{P}$ is a perimeter minimizing polygon transversal of S. The perimeter of $\mathcal{P}$ is computed as the sum of the Euclidean lengths of its edges. We assume that real arithmetic can be performed in constant time. Although this may be an unrealistic assumption, it allows us to focus on the combinatorial issues of the problem without getting involved in numerical intricacies. Observe that a minimum

Figure 1: A set of line segments and its critical sequence.

as a *critical segment*. An extreme line is *critical extreme* if it passes through at least two critical points of different critical segments. Let $\phi(L) \in [0..2\pi)$ denote the polar angle of a directed line L. We define the *critical sequence* of S as a sequence $\Xi(S) = (\xi_1, \xi_2, \ldots, \xi_m)$ so that for all i = 1, 2, ... m, ξ_i is a critical extreme line, and $\phi(\xi_i) \leq \phi(\xi_i + 1)$. In figure 1 a set of line segments is shown with its critical sequence illustrated with dashed lines.

The critical sequence of S plays an important role in developing an algorithm to determine a minimum polygon transversal of S. This structure can also be used to solve other problems. It is a previously unexplored property of line segments that is interesting in its own right. In section 3 we will consider the critical sequence of line segments in more detail.

Lemma 2 *Let p be any vertex of $\mathcal{P}$, then p is incident to the boundary of some minimal stabbing half-plane of S.*

Proof. Let L be a support line of $\mathcal{P}$ incident to p, and assume, for the sake of contradiction, that p is not incident to the boundary of any minimal stabbing half-plane. Therefore there is a minimal stabbing half-plane of S, call it H, such that L is entirely interior to H or L is entirely exterior to H.

Suppose L is interior to H. Therefore, there exists a line segment $s' \in S$ so that $\mathcal{P}$ and s' lie on opposite sides of L. This implies that the intersection of s' and $\mathcal{P}$ is empty, and $\mathcal{P}$ is not a transversal of S.

Suppose on the other hand that L is exterior to H. If p is not adjacent to any segment in S then there exists a line L' contained in $H^{\ell}(L)$ so that $H^{\ell}(L') \cap \mathcal{P}$ is a polygon transversal of S smaller than $\mathcal{P}$. Otherwise if p is adjacent to a segment $s \in S$ we can find a line L' and do as above as long as s intersects the interior of $\mathcal{P}$. That s intersects the interior of $\mathcal{P}$ is guaranteed, otherwise s would lie on a support line of $\mathcal{P}$, the boundary of a minimum stabbing half-plane, contrary to our initial assumption.

Therefore, we have shown that if p is a vertex of a polygon transversal and p is not incident to a minimum stabbing half-plane, then that polygon transversal is not a minimum polygon transversal. $\square$

Lemma 3 *Let p be any vertex of $\mathcal{P}$, then p is incident to a critical segment.*

Proof. Assume first that p is not incident to any line segment of S. We can construct a small disc, c, centred at P such that the intersection of c and S is empty. Now construct a polygon $\mathcal{P}'$ by using the chord made by the intersection of c and $\mathcal{P}$. This new polygon is a polygon transversal and furthermore its perimeter is smaller than the perimeter of $\mathcal{P}$, thus the assumption that $\mathcal{P}'$ is a minimum polygon transversal is false.

Now assume that p is incident to one or more line segments in S and none of them are critical line segments. If p is incident to the boundary of a minimal stabbing half-plane then each of the segments intersect the interior of $\mathcal{P}$. As before we can construct a small disc c centred at p so that we can form a polygon transversal of S that is smaller than $\mathcal{P}$. Let S' denote the set of segments in S passing through p. The disc c is centred at p and small enough so that the intersection of $S - S'$ and c is empty, and every segment in S' intersects the boundary of c inside $\mathcal{P}$. A polygon $\mathcal{P}'$ constructed by using the chord made by the intersection of c and $\mathcal{P}$ has smaller perimeter and is a polygon transversal. This shows that $\mathcal{P}$ is not optimal. Therefore we can conclude that all vertices $p \in \mathcal{P}$ are incident to critical segments. □

A non-degenerate (a segment consisting of more than a single point) critical segment collinear to an extreme line, is denoted as a *rim segment*. If a vertex $p \in \mathcal{P}$ intersects a critical segment s so that a line through s supports $\mathcal{P}$, then we say that p is a *reflection vertex*. This implies that every reflection vertex of $\mathcal{P}$ must lie on a rim segment.

We have given a necessary condition for the occurrence of reflection vertices. However, as we shall see, it is not necessarily true that every rim segment intersects a reflection vertex of $\mathcal{P}$. There are some nice rim segments for which we can guarantee intersection with a reflection vertex. We can define a rim segment s as a *nice rim segment* if the directed extreme line supporting s, $L(s)$, contains the set of critical points of S on the boundary of or in its left half-plane, $H^{\ell}(L(s))$. Those rim segments in S that are not nice are denoted as *pesky rim segments*. Thus the rim segments are partitioned into two equivalence classes, and the adjectives nice and pesky used with a rim segment denotes membership in one class or the other. See figure 6 for an example of a set of all pesky line segments.

Lemma 4 *If s is a nice rim segment then a reflection vertex of $\mathcal{P}$ will lie on s.*

Proof. Observe that $H^{\ell}(L(s))$ contains all of the critical points of S in its interior or on its boundary and is convex. So all points in the interior of any rim segment of S must also be contained in $H^{\ell}(L(s))$. Every vertex of $\mathcal{P}$ is incident to a critical segment. Therefore every vertex of $\mathcal{P}$ must either be in the interior or on the boundary of $H^{\ell}(L(s))$. It follows that there must be a vertex of $\mathcal{P}$ incident to s, and that this vertex is a reflection vertex. □

In the next section we examine critical sequences in more detail. Afterwards we return to the problem of computing minimum polygon transversals.

3. Computing Critical Sequences

Recall we use $\Xi(S)$ to denote the critical sequence of S. Let $|\Xi(S)|$ denote the cardinality of $\Xi(S)$. We prove that $|\Xi(S)|$ is in $O(n)$. The proof follows the results of[8] very closely. To obtain a linear upper bound on the size of $\Xi(S)$ we cast $\Xi(S)$ into a transformed domain. Let the dual transformation $\mathcal{D}$ map a point $p = (\pi_1, \pi_2)$ to the non-vertical line $\mathcal{D}(p) : y = \pi_1 x + \pi_2$, and map a non-vertical line $L : y = \lambda_1 x + \lambda_2$ to the point $\mathcal{D}(L) = (\lambda_1, \lambda_2)$. This (or similar variants) is a commonly used transformation in the computational geometry literature, see[8]. To represent a line segment s with endpoints p and q under the dual transformation apply $\mathcal{D}$ to p and to q. Thus $\mathcal{D}(p)$ and $\mathcal{D}(q)$ are two non vertical lines that define a double wedge. Observe that for every point r in the interior of s, $\mathcal{D}(r)$ is a line contained in the interior of the double wedge. Every line passing through s maps to a point contained in the interior or boundary of the double wedge. We can use $\mathcal{D}(s)$ to denote the closed double wedge as obtained above. What is a critical extreme line under the dual transformation? Given two line segments s and t we use $L(s,t)$ to denote a directed line that is simultaneously tangent to s and to t and keeps both s and t to the right. Consider a critical extreme line $L(s,t)$, $\mathcal{D}(L(s,t))$ is a point that intersects the boundaries of both $\mathcal{D}(s)$ and $\mathcal{D}(t)$. If we consider only those critical extreme lines where left and above coincide then, every line parallel and to the left (above) of $L(s,t)$ does not intersect s or t and for every line segment $u \in S$ there is a line L_u parallel to the left (above) of $L(s,t)$ that intersects u. In the dual plane these lines correspond to points on a vertical line passing through $\mathcal{D}(L(s,t))$ above $\mathcal{D}(L(s,t))$. We can conclude that critical extreme lines that intersect S in the closed lower half-plane correspond to vertices in the lower envelope of the the upper rays of $\mathcal{D}(S)$. Similarly, the critical extreme lines that intersect S in the closed upper half-plane correspond to vertices in the upper envelope of the the lower rays of $\mathcal{D}(S)$. To obtain an upper bound on the number of critical extreme lines, effectively the size of $\Xi(S)$, we only have to determine the number of vertices in these lower and upper envelopes.

Lemma 5 *There are at most $3n-1$ vertices in the lower envelope of the upper rays of $\mathcal{D}(s)$.*

Proof. We will give a counting argument on the number of edges in the envelope that are contained in upper right rays. Observe, that the envelope consists of at most $n+1$ bays bounded to the left and to the right by wedge centres. Each bay is a convex chain of edges, and the edges of the envelope in a single bay and on upper right rays appear in decreasing order by slope. We can say that the last upper right edge of a bay is the edge with smallest slope in that bay. Therefore, there are at most n last edges contained in upper right rays. (The leftmost bay cannot contain any upper right edges.) Also observe that no upper right ray can contain more that one non-last edge. To elucidate this point consider a non-last edge e and let e^+ denote its successor edge. Note that $r(e)$ lies above $r(e^+)$, the upper right rays containing e and (e^+) respectively, everywhere to the right of the intersection of e and e^+. Thus $r(e)$ cannot appear anywhere on the envelope to the right of the intersection of (e) and (e^+). Observe too that the upper right ray that contains the rightmost last edge contains no non-last edges at all. Thus there are at most $n-1$ non-last upper right ray edges in the envelope. A similar symmetric argument can be used for upper left rays. So we can establish a bound of $4n-2$ edges in the envelope. We are interested in the vertices of the envelope, in particular the vertices that are not centers of wedges. Since there are n wedge centres there can be no more than $3n-1$ vertices on the envelope that correspond to critical extreme lines. □

Using the cardinality of $\Xi(S)$ we can describe an efficient divide and conquer algorithm to compute $\Xi(S)$. The details follow.

Algorithm CRITSEQ

1. If $|S| \geq 3$ then split S into two roughly equal subsets A and B, such that $|A| - |B| \leq 1$.
2. Recursively compute $\Xi(A) = \alpha_1, \alpha_2, \ldots, \alpha_{|A|}$ and $\Xi(B) = \beta_1, \beta_2, \ldots, \beta_{|B|}$.
3. Merge $\Xi(A)$ and $\Xi(B)$ to obtain $\Xi(S)$.

It remains to show how to merge $\Xi(A)$ and $\Xi(B)$. We introduce some notation to keep track of segments that are tangent to critical extreme lines. For a critical extreme line ξ we use $s(\xi)$ and $t(\xi)$ to denote segments in S, such that, ξ is tangent to both $s(\xi)$ and $t(\xi)$, ξ is directed from $s(\xi)$ to $t(\xi)$, $s(\xi)$ and $t(\xi)$ are both in the right half-plane of ξ, and if ξ is tangent to more than two segments in S, then $s(\xi)$ and $t(\xi)$ are extreme on ξ. On the other hand, if a and b are distinct line segments then we use $L(a,b)$ to denote the line tangent to a and to b, directed from a to b, such that a and b are contained in the closed right half plane of $L(a,b)$. Thus $\xi = L(s(\xi), t(\xi))$.

Algorithm Merge

1. Let M $= \Xi(A) +_M \Xi(B)$, where $+_M$ denotes merging the sequences $\Xi(A)$ and $\Xi(B)$ into one sequence where the polar angles are sorted.

2. For each $\alpha \in \Xi(A)$ there exists β^- and β^+ such that the angles $\phi(\beta^-) \leq \phi(\alpha) \leq \phi(\beta^+)$ and β^- and β^+ are adjacent in $\Xi(B)$. We denote by $b \in B$ the segment $b = t(\beta^-) = s(\beta^+)$. Let L_b denote the directed line with polar angle $\phi(\alpha)$ that is tangent to b.

   if $H^\ell(\alpha)$ is a subset of $H^\ell(L_b)$ then delete α from M.

3. For each $\beta \in \Xi(B)$ we define a L_a using a symmetric definition to the one in step 2.

   if $H^\ell(\beta)$ is a subset of $H^\ell(L_a)$ then delete β from M.

4. Consider α^- and α^+ adjacent in $\Xi(A)$. Let $a = t(\alpha^-) = s(\alpha^+)$, and let $b = t(\beta^-) = s(\beta^+)$ as above. Then $\theta(a) = [\phi(\alpha^-)\ \ldots\ \phi(\alpha^+)]$ (respectively $\theta(b) = [\phi(\beta^-)\ \ldots\ \phi(\beta^+)]$. For every pair $a \in A$ and $b \in B$ such that $\theta(a)\ \cap\ \theta(b) \neq \emptyset$ let $\phi_{lo} = \max(\phi(\alpha^-,\ \beta^-)$ and let $\phi_{hi} = \min(\phi(\alpha^+,\ \beta^+)$.(Note: The usual precautions must be taken for angle wraparound.)

   if $\phi_{lo} \leq \phi(L(a,b)) \leq \phi_{hi}$ then insert $L(a,b)$ into M.

5. Set $\Xi(S)$ to M.

Theorem 1 *Algorithm CRITSEQ determines $\Xi(S)$ in $O(n \log n)$ time and $O(n)$ space.*

Proof. The complexity of algorithm CRITSEQ is characterized by the recurrence $T(n) \leq T(n/2)$ + cost of merging $\Xi(A)$ and $\Xi(B)$. Thus we must show that the cost of the merge algorithm is in $O(n)$. Steps 1 and 5 are obviously in $O(n)$. It is straightforward to show that steps 2-4 in algorithm merge are in $O(n)$ after performing $\Xi(A) +_M \Xi(B)$.

It remains to argue the correctness of the algorithm. The correctness of CRITSEQ relies on the correctness of algorithm merge. Thus consider a $\xi \in \Xi(S)$. ξ is either:

1. an $\alpha \in \Xi(A)$

2. a $\beta \in \Xi(B)$

   or a critical extreme line that *bridges* $\Xi(A)$ and $\Xi(B)$, that is,

3. ξ is of the form $s(\xi) \in A$ and $t(\xi) \in B$

4. ξ is of the form $s(\xi) \in B$ and $t(\xi) \in A$

An α is in $\Xi(S)$ if and only if α is critical extreme in A and the critical extreme line in B with $\phi(\alpha)$ is contained $H^\ell(\alpha)$. This is the test that is performed in step 2. of algorithm merge. Case 2. above is symmetric to case 1. and is handled by step 3. of algorithm merge.

On the other hand consider the case where ξ is a bridge (as in 3. and 4 above). Thus ξ must be critical extreme in both A and B. Using the notation of algorithm merge it suffices to show that $\xi = L(a,b)$ is critical extreme in A and B. All lines tangent to a in the range of angles $\theta(a)$ are critical extreme. Similarly lines tangent to b in the range $\theta(b)$ are critical extreme. Thus if $\phi(L(a,b))$ is in the range of both $\theta(a)$ and $\theta(b)$ then $L(a,b)$ is critical extreme in both A and B, and otherwise $L(a,b)$ is not critical extreme in both A and B. A symmetric argument can be used for case 4.

Thus we have argued that Merge correctly merges $\Xi(A)$ and $\Xi(B)$ to obtain $\Xi(S)$ in $O(n \log n)$ time and $O(n)$ space. □

4. Minimum Polygon Transversals of Line Segments

We begin this discussion by assuming that we have a set of line segments S, restricted so that every rim segment in S is a nice rim segment. Note this guarantees that no two rim segments intersect. However, we make no such restrictions on non-rim segments. We will use the term *nice*

Figure 2: A set of nice segments and its critical sequence.

segments to describe such a set of line segments. See figure 2. In the sequel we remove the niceness condition and show how the general problem can be solved.

4.1. The Nice Case

Critical extreme lines of S are lines of support of $\mathcal{P}$ and thus appear in angular order in a counter-clockwise traversal of $\mathcal{P}$. Thus consecutive reflection vertices of $\mathcal{P}$ are incident to consecutive rim segments. To make this notion more formal we introduce some more notation. Consider the critical extreme line $\xi_i = (L(s(\xi_i), t(\xi_i))$. We denote the endpoint of $s(\xi_i)$ that is incident to ξ_i as $p(s(\xi_i))$. So if $s(\xi_i)$ and $s(\xi_j)$ are rim segments so that for all $k, i < k < j$, $s(\xi_k)$ is not a rim segment, then we say that $s(\xi_i)$ and $s(\xi_j)$ are consecutive rim segments. We can denote the part of the boundary of $\mathcal{P}$ that lies between the reflection vertices incident to $s(\xi_i)$ and $s(\xi_j)$ as $\mathcal{P}_{ij}$. We denote a channel between $s(\xi_i)$ and $s(\xi_j)$ as C_{ij}. C_{ij} is a simple polygon constructed by taking the difference between two convex polygons. That is, take the convex hull of the critical points swept by the critical sequence from $s(\xi_i)$ to $s(\xi_j)$, that is, $T = \{p(t(\xi_{i-1})), p(s(\xi_i)), p(t(\xi_i)), p(s(\xi_{i+1})), \ \ldots, \ p(s(\xi_{j-1})), p(t(\xi_{j-1})), p(s(\xi_j))\}$ and remove from it the convex hull of the points $T - \{p(t(\xi_{i-1})), p(s(\xi_j))\}$. We denote the channel boundary derived from the first hull the *outer channel boundary* and the part derived from the second hull the *inner channel boundary*. It is possible that the inner and outer channel boundaries overlap. See figure 3.

Lemma 6 *Given a set of nice rim segments, with $s(\xi_i)$ and $s(\xi_j)$, two consecutive rim segments, the polygonal path $\mathcal{P}_{ij}$ is constrained to lie within the channel C_{ij}.*

Proof. The path $\mathcal{P}_{ij}$ must keep critical points of non rim segments in the interior or on the boundary of $\mathcal{P}$. Therefore, the inner channel boundary must lie to the left of $\mathcal{P}_{i,j}$. The outer channel boundary is the outer limit of a minimal length path. For suppose $\mathcal{P}_{i,j}$ leaves the channel. It must cross the outer boundary at least twice. But any such path can be made shorter by clipping it to the outer boundary. Therefore, we conclude that $\mathcal{P}_{ij}$ is constrained to lie within the channel C_{ij}. □

Recall that $\mathcal{P}$ encounters extreme lines in polar order. Motivated by this fact, we construct a reflection polygon by reflecting channels about common edges. Consider two channels C_{ij} and C_{jk}, sharing a common rim segment $s(\xi_j)$. C_{jk} is concatenated to C_{ij} by reflecting C_{jk} about $s(\xi_j)$, and translating and rotating it until the two copies of $s(\xi_j)$ are aligned. This is repeated for every pair of adjacent channels. The final outcome is a polygon, $\mathcal{Q}$, we will denote the *reflection polygon* for S. Observe that the reflection polygon thus constructed may or may not be simple, but this is of no consequence. See figure 4. Note, that the boundary of $\mathcal{Q}$ contains two images of rim segments.

Figure 3: The channels obtained from the nice rim segments of figure 2.

Figure 4: The reflection polygon obtained from the nice rim segments of figure 2.

In fact these are both copies of the same rim segment. We will denote the left copy by ℓ_Q and the right copy by r_Q.

Lemma 7 *A shortest path in $\mathcal{Q}$ from a point q on ℓ_Q to its reflected image q' on r_Q maps to a minimum polygon transversal of S.*

Proof. There exists a 1-1 correspondence between every path in $\mathcal{Q}$ from q a point on ℓ_Q to its reflected image q' on r_Q and polygonal chains with common endpoints that are constrained to lie in the channels between consecutive rim segments. Thus using lemma 4.1 we can conclude that the shortest path obtained as above corresponds to a minimum polygon transversal of S. □

We can state an algorithm to obtain a minimum polygon transversal of a set of nice line segments.

Algorithm MPT

1. Obtain the critical sequence $\Xi(S)$.
2. Use $\Xi(S)$ to construct the reflection polygon $\mathcal{Q}$.
3. Solve for a shortest path inside $\mathcal{Q}$ that begins at a point q on ℓ_Q and ends at a point q' on r_Q, such that q' is a reflected image of q.

To determine the computational complexity of Algorithm MPT we examine each of the steps individually. Using the results of section 3, step 1 is in $O(n \log n)$. Step 2 can be done in $O(n \log n)$ time using any $O(n \log n)$ convex hull algorithm[18,8]. In fact the polar order of $\Xi(S)$ can be explolited so that the convex hulls could be found in linear time . To solve step 3 we can use a method that was used in[6] or[7] to find a shortest path in $\mathcal{Q}$ from the edge ℓ_Q to r_Q in $O(n)$ time.

Theorem 2 *A minimum polygon transversal of a set of nice segments can be obtained in $O(n \log n)$ time.*

4.2. The Pesky Case

The approach used in finding a minimum polygon transversal of a set of line segments with pesky segments, is to convert the input into numerous different sets of nice segments, and use the algorithm MPT developed in the previous section on each of the individual cases.

Consider a set of line segments, S, and a polygon $\mathcal{P}$ which is a minimum polygon transversal. For each pesky line segment s in S there are several ways in which s can interact with $\mathcal{P}$. That is, one or the other or both of the endpoints of s can be stabbed by $\mathcal{P}$, or $\mathcal{P}$ stabs an interior point of s and lies in $H^\ell(L(s))$. Thus consider a pesky rim segment s with endpoints $p(s(\xi_i))$ and $p(s(\xi_{i+1}))$, we label these endpoints as $p^-(s)$ and $p^+(s)$ respectively. We can assign a symbol from $\{-,+,0\}$ to each pesky line segment $s \in S$ denoting how s interacts with $\mathcal{P}$, where the symbols are assigned as follows:

if ($p^-(s)$ is stabbed by $\mathcal{P}$ **) then** symbol$(s) \leftarrow -$

  else if ($p^+(s)$ is stabbed by $\mathcal{P}$ **) then** symbol$(s) \leftarrow +$

    else symbol$(s) \leftarrow 0$

The precedence used above is deliberate, resulting in a unique characterization of each polygon transversal of S with a string from $\{-,+,0\}^*$. We use $\rho(\mathcal{P})$ to denote the string, the *route plan* for polygon $\mathcal{P}$. Note that the unique first symbol of $\rho(\mathcal{P})$ is established using the $\Xi(S)$ order. Using these route plans we can partition the set of all polygon transversals of S into equivalence classes. Thus, $[\mathcal{P}] = \{P : P \text{ is a polygon transversal of } S \text{ and } \rho(P) = \rho(\mathcal{P})\}$. For a set S of k pesky rim segments this amounts to at most $O(3^k)$ different equivalence classes.

Given a set of line segments we know that a minimum polygon transversal must follow one of the possible route plans. We show that given a set of pesky segments S, and a route plan $\rho(\mathcal{P})$ we can transform S into a set of line segments S^ρ such that S^ρ is nice, and $\mathcal{P}$ is a minimum polygon transversal for S^ρ.

The transformation of S works as follows:

procedure Modify($\rho(\mathcal{P})$**)**

for every pesky rim segment $s \in S$

  if (symbol$(s) = -$ **) then** replace s by its endpoint $p^-(s)$

    else if (symbol$(s) = +$ **) then** replace s by its endpoint $p^+(s)$

      else clip the rim segments consecutive to s to the interior of $H^\ell(L(s))$

end for

We give an example of a set of line segments in figure 5, with a polygon transversal and the modified set of segments after applying procedure Modify.

Lemma 8 *The set of segments S^ρ as obtained by procedure Modify is nice, and $\mathcal{P}$ is a minimum polygon transversal of S^ρ.*

Proof. Clearly the modified pesky segments labeled $-$ or $+$ are nice, since as points they are no longer rim segments. For those pesky segments s labeled 0 observe that the rim segments consecutive to s are clipped to the interior of $H^\ell(L(s))$. Furthermore, since all rim segments are interior to the left half plane of their consecutive rim segments, by transitivity all rim segments are inside $H^\ell(L(s))$. Therefore S^ρ is nice.

To see that $\mathcal{P}$ stabs S^ρ observe that all transformed segments remain stabbed by $\mathcal{P}$. The segments replace by their endpoints are clearly stabbed by $\mathcal{P}$. Those segments that are clipped are rim segments and so $\mathcal{P}$ must lie within the intersection of the left half planes bounding the rim segments. So all original segments in S are stabbed by $\mathcal{P}$ in a point that is also within the intersection of the left half planes bounding the rim segments. This is true of all clipped rim segments, so we conclude that $\mathcal{P}$ does stab S^ρ. □

Figure 5: By our labelling scheme we get symbol(a) = 0, symbol(b) = $-$, symbol(c) = $+$, symbol(d) = 0, and symbol(e) = $-$. On the right we illustrate the modifications used to convert the original set into a nice set of line segments.

Lemma 9 *Given a set S with k pesky rim segments, there exists a string from $\{+,-,0\}^k = \rho(\mathcal{P})$, and a minimum polygon transversal of S^ρ is a minimum polygon transversal of S.*

Proof. Let $\mathcal{P}^\rho$ denote a minimum polygon transversal of S^ρ. Clearly, $\mathcal{P}^\rho$ is also a polygon transversal of S. By the previous lemma we saw that $\mathcal{P}$ a minimum polygon transversal of S is also a polygon transversal of S^ρ. Therefore, the two polygon transversals must be of the same perimeter, and $\mathcal{P}^\rho$ is therefore a minimum polygon transversal of S. □

Our approach should now be clear. Given a set of line segments, we consider all possible route plans. For each route plan we apply procedure Modify to obtain a nice set of segments. We then apply algorithm MPT of the previous section to each such set of nice rim segments. The global minimum polygon transversal is chosen as the best of these solutions.

Theorem 3 *Given a set of n line segments, k of them pesky, a minimum polygon transversal of S can be found in $O(3^k n + n\log n)$ time, and $O(n)$ space.*

Proof. The result is straightforward. The only comment is that the critical sequence does not have to be recomputed from scratch with every set of modified segments. Rather, the endpoints of critical extreme lines for the modified segments are either the replaced endpoints, or the endpoints of clipped line segments. It is a routine matter to accomplish these computations in linear time for each modified set of segments. Thus the complexity of algorithm MPT for each modified set of segments is $O(n)$. The result follows immediately. □

We can restate the result above based on a different, possibly more natural parameter of a set of line segments.

Theorem 4 *Given a set of n line segments each line segment lying in one of m orientations, a minimum polygon transversal of S can be found in $O(3^m n + n\log n)$ time, and $O(n)$ space.*

Proof. There are at most m rim segments for any set of line segments with at most m orientations. Thus, at most m segments can be pesky. The result follows immediately. □

This suggests that if the set of line segments lie in a fixed number of orientations, (like the lines directly obtainable in TEX) then the approach above is polynomial.

A further observation sheds more positive results. Suppose that the number of rim segments in the input is much smaller that the size of the input. It suffices for there to be $O(\log n)$ rim segments for a set of n segments, and the algorithm above is again polynomial in the size of the input.

Figure 6: Every line segment in this set of segments is a pesky rim segment.

On the other hand we can in general obtain a set of line segments all of which are pesky, as is shown in figure 6.

5. Discussion

We have shown that for certain special cases we can compute a minimum polygon transversal of a set of line segments. That is, Given a set of line segments can a minimum polygon transversal be found in polynomial time? Or is the problem NP-hard?

Acknowledgements

I would like to thank two anonymous referees for their helpful comments.

References

1. M. Atallah and C. Bajaj, "Efficient algorithms for common transversals", *Information Processing Letters* **25** (1987) 87–91.
2. D. Avis and M. Doskas, "Algorithms for high dimensional stabbing problems", *Discrete and Applied Mathematics* **27** (1990) 39–48.
3. D. Avis and R. Wenger, "Polyhedral line transversals in space", *Discrete and Computational Geometry* **3** (1988) 257–265.
4. B. K. Bhattacharya, J. Czyzowicz, P. Egyed, I Stojmenovic, G. Toussaint, and J. Urrutia, "Computing shortest transversals of sets", *Proc. 7th ACM Symp. on Comput. Geometry*, North Conway (1991) 71–80.
5. B.K. Bhattacharya, S. Jadhav, A. Mukhopadhayay, and J.-M. Robert, "Optimal algorithms for some smallest intersection radius problems", *Proc. 7th ACM Symp. on Comput. Geometry*, North Conway (1991) 81–86.
6. R.E. Burkard, G. Rote, E.Y. Tao, and Z.L. Yu, "Shortest polygonal paths in space", *Computing* **45** (1990) 51–68.
7. J. Czyzowicz, P. Egyed, H. Everett, D. Rappaport, T. Shermer, D. Souvaine, G. Toussaint, and J.

Urrutia, "The aquarium keeper's problem", *Proc. 2nd ACM SIAM Symp. on Discrete Algorithms* (1991) 459–464.

8. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, (Springer-Verlag, 1987).

9. H. Edelsbrunner, "Finding transversals for sets of simple geometric figures", *Theoretical Computer Science* **35** (1985) 31–45.

10. H. Edelsbrunner, H.A. Maurer, F.P. Preparata, A.L. Rosenberg, E. Welzl, and D. Wood, "Stabbing line segments", *BIT* **22** (1982) 274–281.

11. P. Egyed and R. Wenger, "Stabbing pairwise disjoint translates in linear time", *Proc. 5th ACM Symp. on Comput. Geometry* (1989) 364–369.

12. M.T. Goodrich and J.S. Snoeyink, "Stabbing parallel segments with a convex polygon", *Comput. Vision Graph. Image Process.* bf 49 (1990) 152–170.

13. M. Houle, H. Imai, K. Imai, and J-M Robert, "Weighted Orthogonal Linear L_∞ approximation and applications", *Proc. of the Workshop on Algorithms and Data Structures* (Lecture Notes in Computer Science, No. 382, Springer Verlag 1989) 183–191.

14. B. Jones and Y. Ke, "Maximal segment stabbing of segments in the plane", *Proc. 2nd Canadian Conference in Computational Geometry* (Ottawa, 1990) 85–88.

15. H. Meijer and D. Rappaport, "Minimum polygon covers of parallel line segments", *Department of Computing and Information Science Technical Report 90-279* (Queen's University, 1990).

16. J. O'Rourke, "An on-line algorithm for fitting straight lines between data ranges", *Comm. ACM* **24** (1981) 574–478.

17. M. Pellegrini and P. Shor, "Finding stabbing line in 3-dimensional space", *Discrete and Comput. Geom.* **8** (1992) 191–208.

18. F. P. Preparata and M. I. Shamos, *Computational Geometry*, (Springer Verlag, New York, 1985).

19. A. Tamir, *Problem 4-2,* (Problems presented at the fourth NYU Computational Geometry Day, 1987).

20. R. Wenger, *Stabbing and Separation*, (Ph.D. Thesis, School of Computer Science, McGill University, 1988).