International Journal of Computational Geometry & Applications, Vol. 2, No. 4 (1992) 417–435

COMPUTING SHORTEST TRANSVERSALS OF SETS

BINAY BHATTACHARYA
School of Computing Science, Simon Fraser University
Burnaby, British Columbia, Canada V5A 1S6

JUREK CZYZOWICZ
Département d'Informatique, Université du Québec à Hull
Hull, Québec, Canada J8X 3X7

PETER EGYED
GODFRIED TOUSSAINT
School of Computer Science, McGill University
Montreal, Quebec, Canada H3A 2A7

and

IVAN STOJMENOVIC
JORGE URRUTIA
Department of Computer Science, University of Ottawa
Ottawa, Ontario, Canada K1N 6N5

Received 29 July 1991
Revised 29 October 1992

ABSTRACT

Given a family of objects in the plane, the line transversal problem is to compute a line that intersects every member of the family. In this paper we examine a variation of the line transversal problem that involves computing a shortest line segment that intersects every member of the family. In particular, we give $O(n \log n)$ time algorithms for computing a shortest transversal of a family of n lines, a family of n line segments, and a family of convex polygons with a total of n vertices. In general, finding a line transversal for a family of n objects takes $\Omega(n \log n)$ time. This time bound holds for a family of n line segments as well as for a family of convex polygons with a total of n vertices. Hence, our shortest transversal algorithms for these families are optimal.

Keywords: Computational Geometry, shortest transversal, envelope, butterfly polygon.

1. Introduction

Consider a family of objects in the plane. The family is said to admit a *line transversal* if there exists a line that intersects every member of the family. Given a family of convex sets, mathematicians are interested in establishing necessary and sufficient conditions for the existence of line transversals and enumerating the essentially different line transversals.[1,2] Computer scientists, on the other hand, are interested in the so called *line transversal problem*, which is concerned with the computation of line transversals.

The line transversal problem has received considerable attention in the computing literature. O'Rourke presented an $O(n \log n)$ time algorithm for computing a line transversal of a family of n vertical line segments.[3] As observed by O'Rourke, the problem can be expressed as a 2 variable linear program with $2n$ constraints. Hence, with the aid of more recent results the problem can be solved in $O(n)$ time.[4,5] Edelsbrunner, Maurer, Preparata, Rosenberg, Welzl and Wood showed how to compute a representation of the line transversals of a family of n line segments in $O(n \log n)$ time.[6] By relating the problem of computing a representation of the line transversals of a family of convex polygons with a total of n vertices to the computation of the upper envelope of a set of line segments, Edelsbrunner, Guibas and Sharir were able to develop an $O(n\alpha(n) \log n)$ time algorithm for the problem, where $\alpha(n)$ is the inverse of Ackerman's function.[7] The running time of their algorithm was subsequently reduced to $O(n \log n)$ by a result of Hershberger which shows that the upper envelope of m line segments can be computed in $O(m \log m)$ time.[8] Atallah and Bajaj proposed algorithms for computing a representation of the line transversals of a family of n simple convex sets (sets with a constant size storage description which can be used to compute the intersection and common tangents of any pair of sets in constant time) whose running time complexity depends on the number of times the boundaries of any pair may intersect.[9] For example, for a family of n convex simple homothets their algorithm runs in $O(n \log n)$ time since the boundaries of any pair of sets may intersect at most twice.

Various specialized algorithms have also been proposed. Bajaj and Li described an $O(n \log n)$ time algorithm for computing a line transversal of a family of n equal radius circles.[10] Houle, Imai, Imai and Robert proposed an $O(n \log n)$ time algorithm for computing a line transversal of n circles.[11] Edelsbrunner gave two $O(n \log n)$ time algorithms, one for computing a line transversal of n translates and another for computing a line transversal of n homothets.[12]

The optimality of many line transversal algorithms was established by Avis, Robert and Wenger who showed that to compute a line transversal for a family of n line segments or a family of n circles takes $\Omega(n \log n)$ time.[13] In contrast, several families of objects have special properties which allow for faster line transversal algorithms. Avis and Doskas as well as Edelsbrunner proposed $O(kn)$ time algorithms that compute a line transversal of a family of convex polygons with a total of n vertices, where k is the number of different slopes over all edges of the polygons.[14,12] Egyed and Wenger presented an algorithm that computes a line transversal of a family of n pairwise disjoint convex translates in $O(n)$ time.[15] Furthermore, their algorithm runs in optimal $O(n \log n)$ time if the pairwise disjointness condition does not hold. Egyed and Wenger also proposed an $O(n)$ time algorithm that finds the directed line transversals that intersect an ordered family of n pairwise disjoint simple objects in order.[16]

We consider the following variation of the line transversal problem which, until recently, was completely unexplored. Given a family of objects in the plane, the *shortest transversal problem* is to compute a shortest line segment that intersects every member of the family. Bhattacharya and Toussaint proposed $O(n \log^2 n)$ time algorithms for computing the shortest transversal of a family of n lines and of a family of n line segments.[17] In this paper we extend their work by presenting $O(n \log n)$ time algorithms for the problems they considered, as well as an $O(n \log n)$ time algorithm for computing a shortest transversal of a family of convex polygons with a total of n vertices. The $\Omega(n \log n)$ time bound for the problem of computing a line transversal of a family of n line segments, implies that our algorithms for the line segment and convex polygon versions of the shortest transversal problem, are optimal.

The sequel of this paper is arranged as follows. In Section 2, butterfly polygons are introduced and some key theorems pertaining to butterfly polygons, which will find application in our solutions to the line segment and convex polygon versions of the shortest transversal problem, are presented. In Section 3 an $O(n \log n)$ time algorithm for computing a shortest transversal of a family of n lines is outlined. This algorithm is generalized in Section 4 where it is shown that a shortest transversal of a family of n line segments can be computed in $O(n \log n)$ time. In Section 5 an $O(n \log n)$ time algorithm for computing a shortest transversal of a family of convex polygons with a total of n vertices is developed. Lastly, in Section 6 we conclude with some final remarks and open problems.

2. Geometric Preliminaries

Consider a simple polygon P with exactly four convex vertices. The polygon P is a *butterfly* polygon if there exists an ordered labeling a, b, c, d of the convex vertices of P (see Figure 1) and if P is constrained as follows: the line segments $[a, c]$ and $[b, d]$ intersect; for each pair of opposite concave chains, their common separating tangents intersect each of the remaining two concave chains exactly once. An X-polygon is a butterfly polygon in which a and c as well as b and d are visible. Bhattacharya and Toussaint considered X-polygons in their paper on computing shortest transversals.[17] Our solutions to the line segment and polygon versions of the shortest transversal problem are based on the following fundamental geometric minimization problem: given a butterfly polygon P, compute a shortest line segment lying inside P that joins a pair of opposite concave chains of P. The following theorems are due to Bhattacharya, Egyed and Toussaint.[18,19]

Theorem 1 *Given a butterfly polygon P with n vertices, a shortest line segment that joins a pair of opposite concave chains of P and lies inside P can be computed in $O(\log^2 n)$ time.*

Theorem 2 *Given a butterfly polygon P with opposite pairs of concave chains L, R and T, B where $|L| + |R| = n$ and $|T| + |B| = m$, a shortest line segment that joins L to R and lies inside P can be computed in $O(m + \log n)$ time.*

3. Shortest Transversals of Lines

In the interest of readability, we begin by examining the simplest version of our

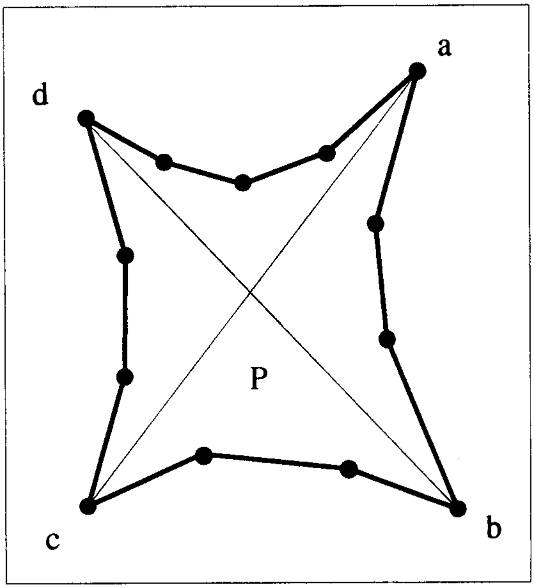

Fig. 1. A butterfly polygon.

problem, that of computing a shortest transversal of a family of lines. The techniques developed here, will, in the subsequent sections, be generalized and applied to the line segment and polygon versions of the problem. Let $\boldsymbol{H} = \{h_1, h_2, \ldots, h_n\}$ denote a family of lines in $\boldsymbol{R}^2$. For simplicity, we assume the lines of $\boldsymbol{H}$ are in general position, that is, no line is vertical, no two lines have the same slope and no three lines are concurrent. These assumptions can be removed without affecting the asymptotic complexity of our algorithm.

Each line h_i is specified by the pair (m_i, c_i) such that $h_i = \{(x, y) \mid y = m_i x + c_i\}$. Let us label the lines in $\boldsymbol{H}$ so that $m_1 < m_2 < \cdots < m_n$. We say that two lines in $\boldsymbol{H}$ are *neighbors* if and only if their slopes are consecutive. Hence, for $i = 1, 2, \ldots, n$ the lines h_i and h_{i+1} (indices taken modulo n) are neighbors. In the following we assume all indices are modulo n. Each line h_i determines two half-planes $h_i^+ = \{(x, y) \mid y \geq m_i x + c_i\}$ and $h_i^- = \{(x, y) \mid y \leq m_i x + c_i\}$. We call each h_i^+ a *positive half-plane* and each h_i^- a *negative half-plane*. Given $t, t' \in \boldsymbol{R}$, let $[t, t']$ denote the range $\{s : t \leq s \leq t'\}$ if $t \leq t'$ and $\{s : s \geq t \cup s \leq t'\}$ if $t > t'$. Similarily, given $t, t' \in \boldsymbol{R}$, let (t, t') denote the range $\{s : t < s < t'\}$ if $t < t'$ and $\{s : s > t \cup s < t'\}$ if $t > t'$.

Our algorithm partitions the real line $\boldsymbol{R}$ corresponding to all possible slopes of lines in $\boldsymbol{R}^2$ into the ranges $r_1, r_2, \ldots, r_n$ and for each range computes a shortest transversal of $\boldsymbol{H}$ with slope in the range. Each range is determined by a pair of neighbors of $\boldsymbol{H}$. For $i = 1, 2, \ldots, n$ the range $r_i = [m_i, m_{i+1}]$. Given a pair h_i, h_{i+1} of neighbors, we call the regions that lie between h_i and h_{i+1}, in the counterclockwise sense, L_i and R_i. Thus, for $i = 1, 2, \ldots, n-1$ let $L_i = h_i^- \cap h_{i+1}^+$ and $R_i = h_i^+ \cap h_{i+1}^-$. As well, let $L_n = h_1^- \cap h_n^-$ and $R_n = h_1^+ \cap h_n^+$.

We now investigate some properties of transversals of families of lines. Having partitioned $\boldsymbol{R}$ into the ranges $r_1, r_2, \ldots, r_n$, we characterize the transversals of the particular ranges. We begin with a lemma that forms a basis for the solution of each shortest transversal problem considered in this paper. The simple proof is left to the reader.

Lemma 1 *Let l_1, l_2 denote a pair of lines in $\boldsymbol{R}^2$ where $l_i = \{(x,y) \mid y = m_i x + c_i\}$ and $m_1 < m_2$. For $i = 1,2$ let $h_i^+ = \{(x,y) \mid y \geq m_i x + c_i\}$ and $h_i^- = \{(x,y) \mid y \leq m_i x + c_i\}$. Consider the range $r = [m_1, m_2]$ and the regions $L = h_1^- \cap h_2^+$ and $R = h_1^+ \cap h_2^-$. A line segment s is a transversal of $\{l_1, l_2\}$ with slope$(s) \in r$ if and only if s joins L to R.*

Suppose s is a transversal of $\boldsymbol{H}$ and *slope*$(s) \in r_i$, it follows from Lemma 1 that each of L_i and R_i contains an endpoint of s. Hence, when considering a particular range r_i, we can confine our search for a shortest transversal of $\boldsymbol{H}$ to those line segments that have an endpoint in each of L_i and R_i.

The lines of $\boldsymbol{H}$ induce a planar map called the arrangement $\boldsymbol{A}$ of $\boldsymbol{H}$, whose vertices are the intersection points of lines in $\boldsymbol{H}$, edges are the maximal connected portions of lines in $\boldsymbol{H}$ not containing a vertex and faces are maximal connected portions of the plane not meeting any edge or vertex of $\boldsymbol{A}$. We shall call each region of the plane that is the union of a face f of $\boldsymbol{A}$ and the vertices and edges bounding f, a cell of $\boldsymbol{A}$. Each cell of $\boldsymbol{A}$ is the intersection of a family of n closed half-planes where for $i = 1, 2, \ldots, n$ either h_i^+ or h_i^- is a member of the family. For $i = 1, 2, \ldots, n$ let

$$H_i^L = \{h_1^-, h_2^-, \ldots, h_i^-, h_{i+1}^+, \ldots, h_n^+\}$$

$$H_i^R = \{h_1^+, h_2^+, \ldots, h_i^+, h_{i+1}^-, \ldots, h_n^-\}$$

It is easy to verify that $\boldsymbol{A}$ has $2n$ unbounded cells and that each such cell is determined by the half-planes of some H_i^L or H_i^R. For $i = 1, 2, \ldots, n$ let A_i^L and A_i^R denote the intersection of the half-planes in H_i^L and H_i^R, respectively. Observe that $A_i^L \subseteq L_i$ and $A_i^R \subseteq R_i$. We call each pair A_i^L, A_i^R of unbounded cells *antipodes* of $\boldsymbol{A}$.

For each range r_i, only a segment with an endpoint in each of L_i and R_i can have the property that it is a transversal of $\boldsymbol{H}$. However, not every such segment has this property. The following lemma characterizes those pairs of points, one from each of L_i and R_i, that define a transversal of $\boldsymbol{H}$.

Lemma 2 *For $i = 1, 2, \ldots, n$ a line segment s is a transversal of $\boldsymbol{H}$ with slope$(s) \in r_i$ if and only if s joins A_i^L to A_i^R.*

Proof. ($\Rightarrow$) Suppose that $s = [a,b]$ is a transversal of $\boldsymbol{H}$ and *slope*$(s) \in r_i$. From Lemma 1 we know that each of L_i and R_i contains an endpoint of s. Suppose that $a \in L_i$ and $b \in R_i$. Consider any line h_j where either $j < i$ or $j > i+1$. By definition $m_j \notin [m_i, m_{i+1}]$. Hence, if $j < i$ then $m_j < m_i$ and so $a \in h_j^-$ and $b \in h_j^+$. On the other hand, if $j > i+1$ then $m_j > m_{i+1}$ and so $a \in h_j^+$ and $b \in h_j^-$. Clearly, $a \in A_i^L$ and $b \in A_i^R$, establishing the necessity of the lemma. ($\Leftarrow$) Let $s = [a,b]$ and suppose that $a \in A_i^L$ and $b \in A_i^R$. Observe that s is a transversal of each line h_j since if $j \leq i$ then $a \in h_j^-$ and $b \in h_j^+$, and if $j \geq i+1$ then $a \in h_j^+$ and $b \in h_j^-$.

Hence s is a transversal of $\boldsymbol{H}$. Furthermore, since $a \in L_i$ and $b \in R_i$ we know from Lemma 1 that $slope(s) \in r_i$, establishing the sufficiency of the lemma. □

It follows from Lemma 2 that for a particular range r_i we need only determine the shortest line segment joining A_i^L to A_i^R. Note that since we assumed the lines of $\boldsymbol{H}$ are in general position, each pair A_i^L, A_i^R of antipodes is disjoint. This immediately suggests the following algorithm: for $i = 1, 2, \ldots, n$ determine the pair A_i^L, A_i^R and compute the line segment that realizes the shortest distance between the pair; choose the smallest of the line segments. Although this algorithm is correct, we show that it can be simplified.

For $i = 1, 2, \ldots, n$ let σA_i^L and σA_i^R denote the boundaries of A_i^L and A_i^R, respectively, and let E_i^L and E_i^R denote the portion of the boundaries of σA_i^L and σA_i^R, respectively, that include only bounded edges. The *envelope* $\boldsymbol{E}$ of $\boldsymbol{H}$ is defined as the polygon whose boundary consists of the bounded edges and vertices of all the unbounded cells of $\boldsymbol{A}$ (see Figure 2). Hence,

$$\boldsymbol{E} = \bigcup_{i=1}^{n} E_i^L \cup E_i^R$$

Consider the following observations concerning $\boldsymbol{E}$: E_i^L and E_i^R are the portions of $\boldsymbol{E}$ contained in L_i and R_i, respectively; the endpoints of E_i^L and E_i^R are the intersection points of $\boldsymbol{E}$ with h_i and h_{i+1}; if we traverse the boundary of $\boldsymbol{E}$ in the counterclockwise direction, beginning with an intersection point of h_1 and the boundary of $\boldsymbol{E}$, then we encounter the lines of $\boldsymbol{H}$ twice in order. Furthermore, Suri showed that $\boldsymbol{E}$ has at most $4n - 2$ vertices.[20]

We now present a lemma that characterizes, in terms of $\boldsymbol{E}$, the shortest transversals of $\boldsymbol{H}$ with slope in a particular range r_i (see Figure 3). Given a pair p, q of points, let $d(p, q)$ denote the euclidean distance between them.

Lemma 3 *For $i = 1, 2, \ldots, n$ a line segment s is a shortest transversal of $\boldsymbol{H}$ with $slope(s) \in r_i$ if and only if s is a shortest line segment joining E_i^L to E_i^R.*

Proof. Let $s = [a, b]$ denote a transversal of $\boldsymbol{H}$ with $slope(s) \in r_i$. From Lemma 2 we know that each of A_i^L and A_i^R contains an endpoint of s. Without loss of generality assume $a \in A_i^L$ and $b \in A_i^R$. Clearly, if $a \notin \sigma A_i^L$ or $b \notin \sigma A_i^R$ then s can be shortened until $a \in \sigma A_i^L$ and $b \in \sigma A_i^R$ while remaining a transversal of $\boldsymbol{H}$. Combining this with Lemma 2, we have that a line segment s is a shortest transversal of $\boldsymbol{H}$ with $slope(s) \in r_i$ if and only if s is a shortest line segment joining σA_i^L to σA_i^R. All that remains to be shown is, given any line segment $s = (a, b)$ with $a \in \sigma A_i^L$ and $b \in \sigma A_i^R$, if either $a \in \sigma A_i^L - E_i^L$ or $b \in \sigma A_i^R - E_i^R$ there exists another line segment $s' = (a', b')$ with $a' \in E_i^L$ and $b' \in E_i^R$ such that $d(a', b') < d(a, b)$. Let $E_i^L = (p_1, p_2, \ldots, p_u)$ and $E_i^R = (q_1, q_2, \ldots, q_v)$ where $p_1 \in h_i$ and $q_1 \in h_{i+1}$. Also, let e_1, e_2, e_3 and e_4 denote the open, unbounded edges of σA_i^L and σA_i^R bounded by p_1, q_1, p_u and q_v, respectively. Let $s = (a, b)$ be the line segment under consideration, then we need only examine the following three cases: i) $a \in e_1$ and $b \in e_4$; ii) $a \in e_1$ and $b \in e_2$; and iii) $a \in e_1$ and $b \in E_i^R$. Each of the cases not mentioned is equivalent to one of these three. For case i) we simply note that $d(p_1, q_v) < d(a, b)$. We reduce case

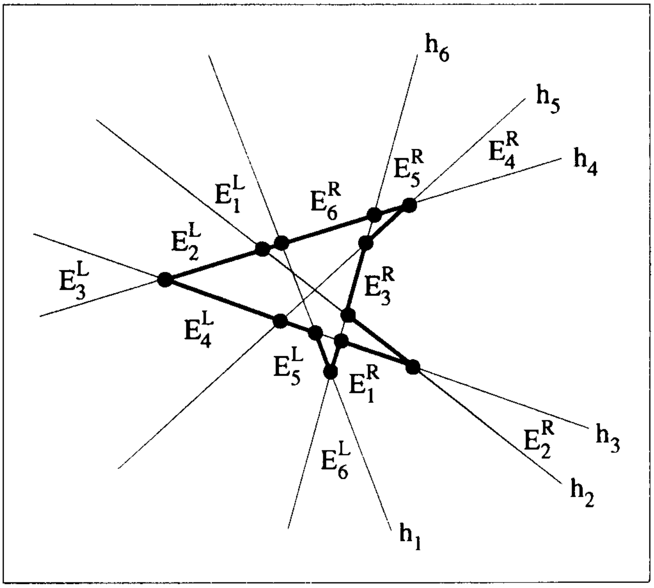

Fig. 2. The envelope of a set of lines.

ii) to case iii) by remarking that $d(a, q_1) < d(a, b)$. Lastly, for case iii) we note that $d(p_1, b) < d(a, b)$. □

It follows from Lemma 3 that for a particular range r_i we need only determine the shortest line segment joining E_i^L to E_i^R. This suggests the following algorithm, **Shortest_Transversal**, which takes as input the family $\boldsymbol{H}$ of lines and returns a shortest transversal of $\boldsymbol{H}$. Initialize the length of the current solution to be infinitely large. Then, for $i = 1, 2, \ldots, n$ compute E_i^L and E_i^R and determine a shortest line segment s that joins E_i^L to E_i^R. If the length of the current solution is larger than the length of s then set the current solution to s and set the length of the current solution to the length of s. Once this process is complete, return the current solution.

We now show that **Shortest_Transversal** returns a solution that is indeed a shortest transversal of $\boldsymbol{H}$. By Lemma 3, when computing a shortest transversal of $\boldsymbol{H}$ with slope in a particular range r_i, it suffices to compute a shortest line segment joining E_i^L to E_i^R. Also, since $r_1, r_2, \ldots, r_n$ partition $\boldsymbol{R}$, the slopes of all potential solutions are considered. Hence, **Shortest_Transversal** computes a shortest transversal of $\boldsymbol{H}$.

What remains to be shown is that there exists an implementation of **Shortest_Transversal** achieving the desired $O(n \log n)$ running time. We first consider the construction for $i = 1, 2, ..., n$ of E_i^L and E_i^R. Suri proposed an $O(n \log n)$ time divide-and-conquer algorithm for computing $\boldsymbol{E}$.[20] Edelsbrunner and Guibas also gave an $O(n \log n)$ time algorithm for computing $\boldsymbol{E}$ in the initialization phase of their algorithm for topologically sweeping an arrangement.[21] Using a balanced tree

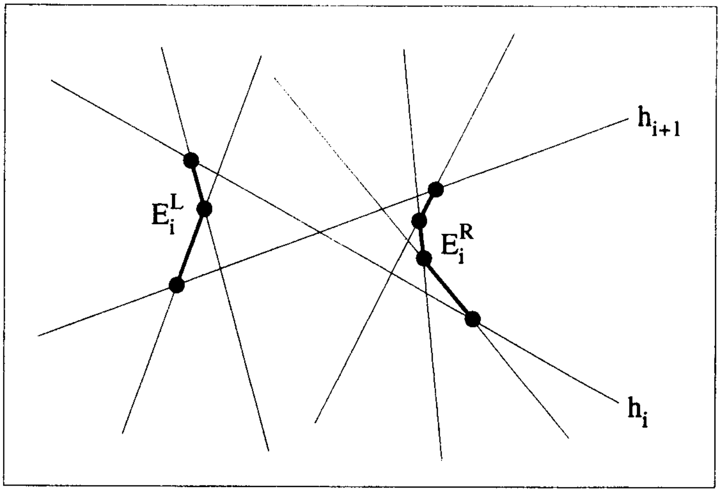

Fig. 3. Illustration of Lemma 3.

representation of a concatenable queue, Vegter obtained a simple $O(n \log n)$ time algorithm for computing $\boldsymbol{E}$.[22] Most recently, Keil described a very simple $O(n \log n)$ time algorithm for computing $\boldsymbol{E}$ that uses no data structure more complex than a stack.[23] Although all the above algorithms could be used to develop $O(n \log n)$ time solutions to the shortest transversal problem when the input consists of either a family of n lines or a family of n line segments, Keil's algorithm adapts most easily to the convex polygon version of the problem and also permits for an optimal $O(n \log n)$ time solution to this problem. Without going into details, we now give a brief description of Keil's algorithm.

The algorithm is iterative in that it constructs A_i^L and A_i^R from A_{i+1}^L and A_{i+1}^R. Note that, given A_i^L and A_i^R, it is straightforward to compute E_{i+1}^L and E_{i+1}^R. Begin by partitioning each of H_i^L and H_i^R into two sets corresponding to their positive and negative half-planes. For $i = 1, 2, \ldots, n$ let

$$H_i^{L+} = \{h_{i+1}^+, h_{i+2}^+, \ldots, h_n^+\} \qquad H_i^{L-} = \{h_1^-, h_2^-, \ldots, h_i^-\}$$

$$H_i^{R-} = \{h_{i+1}^-, h_{i+2}^-, \ldots, h_n^-\} \qquad H_i^{R+} = \{h_1^+, h_2^+, \ldots, h_i^+\}$$

Rather than directly maintain A_i^L and A_i^R from one iteration to the next, the algorithm separately maintains the components of A_i^L and A_i^R that correspond to their positive and negative half-planes. For $i = 1, 2, \ldots, n$ let A_i^{L+}, A_i^{L-}, A_i^{R+} and A_i^{R-} be the intersection of the half-planes in H_i^{L+}, H_i^{L-}, H_i^{R+} and H_i^{R-}, respectively. Hence, $A_i^L = A_i^{L+} \cap A_i^{L-}$ and $A_i^R = A_i^{R+} \cap A_i^{R-}$.

Keil exploits the fact that the intersection of n positive or negative half-planes can be maintained in total $O(n)$ time when the half-planes are either inserted or deleted in sorted order based on the slopes of the lines determining the half-planes. Hence, over all iterations computing A_i^{L+}, A_i^{L-}, A_i^{R+} and A_i^{R-} takes $O(n)$ time. It is straightforward to compute A_i^L from A_i^{L+} and A_i^{L-} in $O(|A_i^L|)$ time, and then E_i^L

from A_i^L in constant time. The same is true for the construction of E_i^R. Recall that $\boldsymbol{E}$ has at most $4n-2$ vertices, and so, over all iteration this process takes $O(n)$ time. Hence, once the lines of $\boldsymbol{H}$ have been sorted in $O(n \log n)$ time, the envelope $\boldsymbol{E}$ of $\boldsymbol{H}$ can be computed in $O(n)$ time.

We still require a technique for computing the shortest line segment joining E_i^L to E_i^R. Edelsbrunner showed that the problem of computing the minimum distance between two non-intersecting convex polygons, with a total of m vertices, can be solved in $O(\log m)$ time.[24] Since the lines of $\boldsymbol{H}$ are in general position, E_i^L and E_i^R are disjoint. Hence, $O(\log n)$ time is sufficient to compute the shortest line segment joining E_i^L to E_i^R. Since $\boldsymbol{E}$ has at most $4n-2$ vertices,

$$\sum_{i=1}^{n} |E_i^L| + |E_i^R| \leq 6n - 2$$

Clearly, the running time of the algorithm is dominated by the time it takes to compute $\boldsymbol{E}$. Therefore, the algorithm runs in $O(n \log n)$ time. Recall that lines of $\boldsymbol{H}$ were labeled so that $m_1 < m_2 < \cdots < m_n$. Since this labeling process reduces to sorting, this requirement does not affect the running time of the algorithm. We have therefore established the following theorem.

Theorem 3 *Given the family $\boldsymbol{H}$ of n lines, a shortest transversal of $\boldsymbol{H}$ can be computed in $O(n \log n)$ time.*

Let $\boldsymbol{I}$ be the intersection points of the lines in $\boldsymbol{H}$. Since $\boldsymbol{E}$ has $O(n)$ vertices, we are not obliged to use an $O(\log n)$ time algorithm to compute the shortest distance between each E_i^L, E_i^R pair in order to obtain an $O(n \log n)$ time algorithm. Clearly, an $O(n)$ time algorithm would suffice. Thus we can also modify our algorithm to compute shortest constrained transversals. McKenna and Toussaint showed that the minimum vertex distance between two non-intersecting convex polygons can be computed in $O(n)$ time.[25] Therefore, we have the following theorem.

Theorem 4 *Given the family $\boldsymbol{H}$ of n lines, a shortest transversal of $\boldsymbol{H}$ with endpoints in $\boldsymbol{I}$ can be computed in $O(n \log n)$ time.*

4. Shortest Transversals of Line Segments

In this section we consider the problem of computing a shortest transversal of a family of line segments. Let $\boldsymbol{S} = \{s_1, s_2, ..., s_n\}$ denote a family of line segments in $\boldsymbol{R}^2$. In addition, let $\boldsymbol{H} = \{h_1, h_2, ..., h_n\}$ denote the corresponding family of lines where h_i is the line containing s_i. We assume that $\boldsymbol{S}$ is conditioned so that the lines of $\boldsymbol{H}$ are in general position. As before, this assumption can be removed without affecting the asymptotic complexity of our algorithm.

Let V denote the union of the endpoints of $s_1, s_2, ..., s_n$. Each line segment s_i is specified by its endpoints p_i^1 and p_i^2. For $i = 1, 2, \ldots, n$ let $V_i = \{p_i^1, p_i^2\}$. As in the previous section, each line h_i is specified by the pair (m_i, c_i) such that $h_i = \{(x, y) \mid y = m_i x + c_i\}$. Furthermore, the lines of $\boldsymbol{H}$, and accordingly the line segments of $\boldsymbol{S}$, are labeled so that $m_1 < m_2 < \cdots < m_n$. Now, for $i = 1, 2, \ldots, n$ define r_i, A_i^L, A_i^R, E_i^L and E_i^R as in the previous section.

Our algorithm computes for each of the ranges $r_1, r_2, \ldots, r_n$ a shortest transversal of $\boldsymbol{S}$ with slope in the range, provided such a transversal exists. Suppose s is a shortest transversal of $\boldsymbol{S}$ with $slope(s) \in r_i$, then since s is also a transversal of $\boldsymbol{H}$, s joins A_i^L to A_i^R. Moreover, by Lemma 2, s joins σA_i^L and σA_i^R. Although a line segment joining σA_i^L and σA_i^R is a transversal of $\boldsymbol{H}$, it is not necessarily a transversal of $\boldsymbol{S}$. Thus, our goal is to ascertain which pairs of points, one from each of σA_i^L and σA_i^R, define a transversal of $\boldsymbol{S}$. We begin by examining the problem of computing a representation of the line transversals of $\boldsymbol{S}$.

Some problems arise when vertical transversals are considered. Depending on the particular situation, a vertical line is said to have a slope of either $+\infty$ or $-\infty$. In order to simplify the presentation involving a potentially vertical line l and the *above/below* relationship, we adopt the following conventions: if $slope(l) = +\infty$ then above and below refer to left and right, respectively; if $slope(l) = -\infty$ then above and below refer to right and left, respectively.

Each line segment s_i has two tangents for a given slope u. Let $uptan(s_i, u)$ denote the *upper tangent* of s_i with slope u, and let $lowtan(s_i, u)$ denote the *lower tangent* of s_i with slope u. Note that $uptan(s_i, u) = lowtan(s_i, u)$ when $u = m_i$. For $i = 1, 2, \ldots, n$ and each slope u, we specify two subsets, $upper(V_i, u)$ and $lower(V_i, u)$ of V_i. Let $upper(V_i, u)$ and $lower(V_i, u)$ be the sets of $p \in V_i$ that lie on $uptan(s_i, u)$ and $lowtan(s_i, u)$, respectively. For each slope u we also specify two subsets of V. Let $upper(V, u)$ and $lower(V, u)$ denote the unions, for $i = 1, 2, \ldots, n$, of $upper(V_i, u)$ and $lower(V_i, u)$, respectively.

If A and B are two sets of points then a line l *separates* A and$B$ if A lies in one of the closed half-planes bounded by l while B lies in the other. A line l has slope u and is a transversal of $\boldsymbol{S}$ if and only if l separates $upper(V, u)$ and $lower(V, u)$. For $i = 1, 2, \ldots, n$ we consider any $u \in (m_i, m_{i+1})$ and let $V_i^U = upper(V, u)$ and $V_i^L = lower(V, u)$. The following lemma, which is a simple adaptation of ideas presented by Edelsbrunner,[12] characterizes the line transversals of $\boldsymbol{S}$ with slope in a particular range r_i.

Lemma 4 *For $i = 1, 2, \ldots, n$ a line l is transversal of $\boldsymbol{S}$ with $slope(l) \in r_i$ if and only if l separates V_i^U and V_i^L.*

Proof. ($\Rightarrow$) If l has $slope(l) \in r_i$ and l is a transversal of S then by definition l separates V_i^U and V_i^L. ($\Leftarrow$) If l separates V_i^U and V_i^L then l is a transversal of S. Suppose $slope(l) \notin r_i$ then there exists some k, $k \neq i$, such that $slope(l) \in r_k$. Clearly, l separates V_k^U and V_k^L. Assume, without loss of generality, that $k < i$. Clearly l cannot contain s_i since this would imply that $slope(l) \in r_i$. Let l^+ and l^- denote the positive and negative half-planes determined by l. Also, let p and q be the endpoints of s_i. Assume $p \in V_i^U$ and $q \in V_i^L$ then $q \in V_k^U$ and $p \in V_k^L$. This implies that $p \in l^+$, $q \in l^-$, $p \in l^-$ and $q \in l^+$. However, since l does contain s_i this is impossible, contradicting the assumption that $slope(l) \notin r_i$. Therefore $slope(l) \in r_i$. □

For $i = 1, 2, \ldots, n$ let B_i^U denote the intersection of all the closed half-planes that contain V_i^U and whose bounding lines separate V_i^U and V_i^L. Similarly, for

$i = 1, 2, \ldots, n$ let B_i^L denote the intersection of all the closed half-planes that contain V_i^L and whose bounding lines separate V_i^U and V_i^L. Observe that if V_i^U and V_i^L are not separable then $B_i^U = B_i^L = \emptyset$. For $i = 1, 2, \ldots, n$ let σB_i^U and σB_i^L be the boundaries of B_i^U and B_i^L, respectively. Note that if $B_i^U = B_i^L = \emptyset$ then $\sigma B_i^U = \sigma B_i^L = \emptyset$. As an immediate consequence of Lemma 4 we have the following lemma which characterizes the line transversals of S with slope in a particular range r_i in terms of σB_i^U and σB_i^L (see Figure 4).

Lemma 5 *For $i = 1, 2, \ldots, n$ a line l is transversal of S with $slope(l) \in r_i$ if and only if l separates σB_i^U and σB_i^L.*

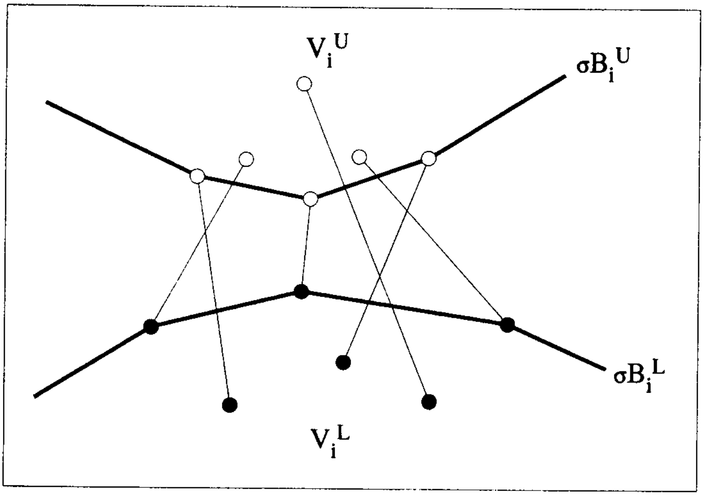

Fig. 4. Illustration of Lemma 5.

Hence, for $i = 1, 2, \ldots, n$ a representation of the line transversals of S with slope in a particular range r_i is equivalent to a representation of σB_i^U and σB_i^L. Edelsbrunner et al. presented an $O(n \log n)$ time, duality based algorithm for computing this representation.[6] Furthermore, they indirectly showed that

$$\sum_{i=1}^{n} |\sigma B_i^U| + |\sigma B_i^L| \leq 10n + 4$$

We now return to the problem of computing a shortest transversal of S. To begin, we present a lemma that characterizes the shortest transversals of S with slope in a particular range r_i (see Figure 5).

Lemma 6 *For $i = 1, 2, \ldots, n$ a line segment s is a shortest transversal of S with $slope(s) \in r_i$ if and only if s is shortest among those line segments that join σA_i^L to σA_i^R and separate σB_i^U and σB_i^L.*

Proof. We first show that a line segment s is a transversal of S with $slope(s) \in r_i$ if and only if s joins A_i^L to A_i^R and separates σB_i^U and σB_i^L. ($\Rightarrow$) Suppose s is a transversal of S with $slope(s) \in r_i$. Since s is also a transversal of H, by Lemma 2, s joins A_i^L to A_i^R. Let l be the line containing s. Since l is a transversal of S and

$slope(l) \in r_i$, by Lemma 5, l separates σB_i^U and σB_i^L, establishing necessity. ($\Leftarrow$) Now suppose that s joins A_i^L to A_i^R and s separates σB_i^U and σB_i^L. Since s joins A_i^L to A_i^R, by Lemma 2, $slope(s) \in r_i$ and s is a transversal of $\boldsymbol{H}$. Suppose s is not a transversal of $\boldsymbol{S}$. Since s is a transversal of $\boldsymbol{H}$, some s_j lies either above or below s such that l_j intersects s. Let p denote the intersection point of h_j and s. Clearly the points p_j^1, p_j^2 and p are collinear, however, p does not lie between p_j^1 and p_j^2 on l_j. Furthermore, each of B_i^U and B_i^L contains an endpoint of s_j. Hence, p does not lie between σB_i^U and σB_i^L, which is a contradiction. Therefore, s is a transversal of $\boldsymbol{S}$, establishing sufficiency. To complete the proof, note that since A_i^L and A_i^R are convex and disjoint, every shortest transversal of $\boldsymbol{S}$ with slope in the range r_i joins σA_i^L and σA_i^R. □

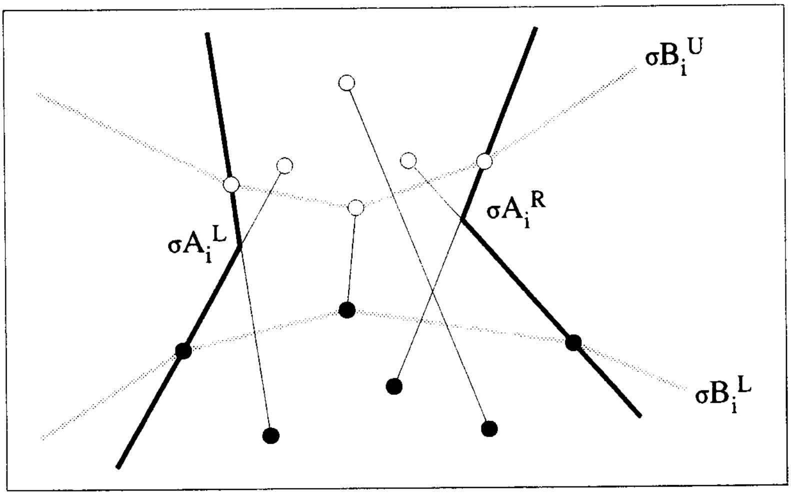

Fig. 5. Illustration of Lemma 6.

For $i = 1, 2, \ldots, n$ let G_i denote the region bounded by σA_i^L and σA_i^R, let H_i denote the region bounded by σB_i^U and σB_i^L, and let $Q_i = H_i \cap G_i$. The following lemma characterizes the regions $Q_1, Q_2, \ldots, Q_n$ (see Figure 6).

Lemma 7 *For $i = 1, 2, \ldots, n$ the region Q_i is a butterfly polygon.*

Proof. Consider some region Q_i such that $H_i \neq \emptyset$. Let r denote the subrange of r_i for which $\boldsymbol{S}$ admits line transversals. Every line contained in $\boldsymbol{H}_i$ has slope in the range r. Let $r' = [m_{i+1}, m_i]$. Every line contained in G_i has slope in the range r'. Observe that r and r' may only overlap at m_i or m_{i+1}. For any slope m, σA_i^L or σA_i^R admits a tangent with slope m if and only if $m \in r$. Similarly, for any slope m, σB_i^U or σB_i^L admits a tangent with slope m if and only if $m \in r'$. Consequently, each of σA_i^L and σA_i^R intersects each of σB_i^U and σB_i^L exactly once. Clearly, Q_i is a simple polygon containing only four convex vertices labeled q_1, q_2, q_3 and q_4, which are the intersections of σA_i^L and σB_i^U, σB_i^U and σA_i^R, σA_i^R and σB_i^L, and σB_i^L and σA_i^L, respectively. Furthermore, the line segments $[q_1, q_3]$ and $[q_2, q_4]$ intersect. Therefore, the region Q_i is a butterfly polygon. □

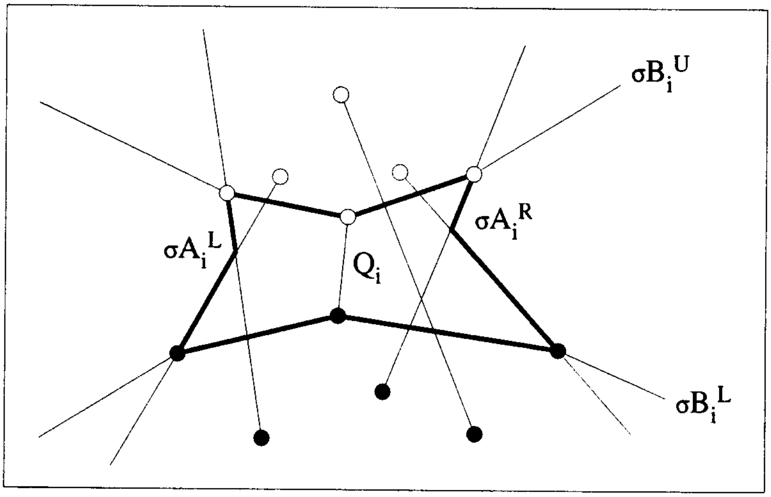

Fig. 6. Illustration of Lemma 7.

Computing a shortest line segment that joins σA_i^L to σA_i^R and separates σB_i^U and σB_i^L is equivalent to computing a shortest line segment that lies entirely inside Q_i and joins the two opposite concave chains of Q_i that correspond to σA_i^L and σA_i^R. This suggests the following variation of **Shortest_Transversal** that allows for computing a shortest transversal of the family $\boldsymbol{S}$ of line segments. Initialize the length of the current solution to be infinitely large. Then, for $i = 1, 2, \ldots, n$ determine Q_i. If $Q_i \neq \emptyset$ then compute a shortest line segment s lying entirely inside Q_i that joins the opposite concave chains of Q_i that correspond to σA_i^L and σA_i^R. If the length of the current solution is larger than the length of s then set the current solution to s and set the length of the current solution to the length of s. Once this process terminates return the current solution if the length of the current solution is not infinitely large, and otherwise return $\emptyset$.

The correctness of the algorithm follows from Lemma 6 and Lemma 7 and that all possible solutions are considered since $r_1, r_2, \ldots, r_n$ partition $\boldsymbol{R}$. Therefore, the algorithm returns a shortest transversal of $\boldsymbol{S}$, provided one exists.

It remains to be shown that there exists an implementation of the algorithm that runs in $O(n \log n)$ time. As shown in Section 3, it is possible to compute all the σA_i^L, σA_i^R pairs in $O(n \log n)$ time. Furthermore, as already indicated, it is possible to compute $\sigma B_1^U, \sigma B_2^U, \ldots, \sigma B_n^U$ and $\sigma B_1^L, \sigma B_2^L, \ldots, \sigma B_n^L$ in $O(n \log n)$ time. Since σA_i^L and σA_i^R intersect the unbounded edges of σB_i^U and σB_i^L, we can construct Q_i in $O(\log(|\sigma A_i^L| + |\sigma A_i^R|))$ time. Recall that Suri showed that $\boldsymbol{E}$ has at most $4n - 2$ vertices, and so

$$\sum_{i=1}^{n} |\sigma A_i^L| + |\sigma A_i^R| \leq 10n - 2$$

Therefore, once the regions $G_1, G_2, \ldots, G_n$ and $H_1, H_2, \ldots, H_n$ have been determined in $O(n \log n)$ time, we can construct $Q_1, Q_2, \ldots, Q_n$ in $O(n)$ time. Recall that Edelsbrunner et al. indirectly showed that

$$\sum_{i=1}^{n} |\sigma B_i^U| + |\sigma B_i^L| \leq 10n + 4$$

Hence,

$$\sum_{i=1}^{n} |Q_i| \leq 20n + 2$$

Applying Theorem 1 to a polygon Q_i, the shortest line segment joining a pair of opposite concave chains of Q_i can be computed in $O(\log^2 |Q_i|)$ time. Hence, computing the required shortest line segments inside $Q_1, Q_2, \ldots, Q_n$ requires only $O(n)$ time. Clearly, the running time of the algorithm is $O(n \log n)$, which is optimal. Furthermore, note that the assumed labeling of the lines of $\boldsymbol{H}$ does not affect the asymptotic running time of the algorithm. We have therefore established the following theorem.

Theorem 5 *Given the family $\boldsymbol{S}$ of line segments, a shortest transversal of $\boldsymbol{S}$ can be computed in $\Theta(n \log n)$ time.*

5. Shortest Transversals of Convex Polygons

We now investigate the final problem, that of computing a shortest transversal of a family of convex polygons. Let $\boldsymbol{P} = \{P_1, P_2, ..., P_m\}$ be a family of simple convex polygons in $\boldsymbol{R}^2$. In addition, let $\boldsymbol{S} = \{s_1, s_2, ..., s_n\}$ denote the family of line segments corresponding to the edges of the polygons of $\boldsymbol{P}$, and let $\boldsymbol{H} = \{h_1, h_2, ..., h_n\}$ denote the family of lines where h_i is the line containing s_i. We assume that $\boldsymbol{P}$ is conditioned so that the lines of $\boldsymbol{H}$ are in general position and $\boldsymbol{P}$ admits no point transversal. As before, this assumption can be handled without affecting the asymptotic running time of our algorithm.

As in the previous section, let V be the union of the endpoints of $s_1, s_2, ..., s_n$. Each line segment s_i is specified by its endpoints p_i^1 and p_i^2. Furthermore, for $i = 1, 2, \ldots, n$ let $h_i = \{(x, y) \mid y = m_i x + c_i\}$ and assume the lines of $\boldsymbol{H}$, and accordingly the line segments of $\boldsymbol{S}$, are labeled so that $m_1 < m_2 < \cdots < m_n$. For $j = 1, 2, \ldots, m$ let V_j be the vertices of the polygon P_j.

For $i = 1, 2, \ldots, n$ define r_i as in the previous section. For each of the ranges r_i, our algorithm computes a shortest transversal of $\boldsymbol{S}$ with slope in the range, provided such a transversal exists. For each of the problems previously considered, the envelope $\boldsymbol{E}$ of $\boldsymbol{H}$ captured all the necessary distance information. Unfortunately, this no longer holds true when convex polygons are considered. For convex polygons we must instead maintain two envelopes, one for each of the left and right sides, that change only slightly as we proceed from one range to another. We begin by examining the problem of computing a representation of the line transversals of $\boldsymbol{P}$.

Each polygon P_j has two tangents for each slope u. Let $uptan(P_j, u)$ denote the upper tangent of P_j with slope u, and let $lowtan(P_j, u)$ denote the lower tangent of P_j with slope u. For $j = 1, 2, \ldots, m$ and each slope u, we specify two subsets, $upper(V_j, u)$ and $lower(V_j, u)$ of V_j. Let $upper(V_j, u)$ be the set of $p \in V_j$ that lie on $uptan(P_j, u)$. Similarly, let $lower(V_j, u)$ be the set of $p \in V_j$ that lie on $lowtan(P_j, u)$. For each slope u we also specify two subsets of V. Let $upper(V, u)$ denote the union of $upper(V_j, u)$ for $j = 1, 2, \ldots, m$. Similarly, let $lower(V, u)$ denote the union of $lower(V_j, u)$ for $j = 1, 2, \ldots, m$.

A line l with slope u is a transversal of $\boldsymbol{P}$ only if l separates $upper(V, u)$ and $lower(V, u)$. For $i = 1, 2, \ldots, n$ we consider any $u \in (m_i, m_{i+1})$ and let $V_i^U = upper(V, u)$ and $V_i^L = lower(V, u)$. The following lemma partially characterizes the line transversals of $\boldsymbol{P}$ with slope in a particular range r_i. The proof mimics that of Lemma 4.

Lemma 8 *For $i = 1, 2, \ldots, n$ if a line l is a transversal of $\boldsymbol{P}$ with $slope(l) \in r_i$ then l separates V_i^U and V_i^L.*

For $i = 1, 2, \ldots, n$ let B_i^U denote the intersection of all the closed half-planes with slope in the range r_i that contain V_i^U and whose bounding lines separate V_i^U and V_i^L. Similarly, for $i = 1, 2, \ldots, n$ let B_i^L denote the intersection of all the closed half-planes with slope in the range r_i that contain V_i^L and whose bounding lines separate V_i^U and V_i^L. Observe that if V_i^U and V_i^L are not separable then $B_i^U = B_i^L = \emptyset$. For $i = 1, 2, \ldots, n$ let σB_i^U and σB_i^L be the boundaries of B_i^U and B_i^L, respectively. Note that if $B_i^U = B_i^L = \emptyset$ then $\sigma B_i^U = \sigma B_i^L = \emptyset$. As an immediate consequence of Lemma 8 we have the following characterization of the line transversals of $\boldsymbol{P}$ with slope in a particular range r_i in terms of σB_i^U and σB_i^L.

Lemma 9 *For $i = 1, 2, \ldots, n$ a line l is a transversal of $\boldsymbol{P}$ with $slope(l) \in r_i$ if and only if l separates σB_i^U and σB_i^L.*

Hence, for $i = 1, 2, \ldots, n$ a representation of the line transversals of $\boldsymbol{P}$ with slope in the range r_i is equivalent to a representation of σB_i^U and σB_i^L. A result of Edelsbrunner, Guibas and Sharir combined with an improvement due to Hershberger contain all the elements of a duality based, $O(n \log n)$ time algorithm for computing this representation.[7,8]

We now return to the problem of computing a shortest transversal of $\boldsymbol{P}$. As in Section 3, let h_i^+ and h_i^- denote the positive and negative half-planes determined by h_i. Each convex polygon P_j can be expressed as the intersection of some set of half-planes determined by the lines in $\boldsymbol{H}$. Let

$$H^+ = \{h_k \mid p_k^1, p_k^2 \in lowtan(P_j, m_k), 1 \leq j \leq m\}$$

$$H^- = \{h_k \mid p_k^1, p_k^2 \in uptan(P_j, m_k), 1 \leq j \leq m\}$$

Each convex polygon P_j can be expressed as the intersection of some subset of the positive half-planes determined by the lines in H^+, and some subset of the negative half-planes determined by the lines in H^-. For $i = 1, 2, \ldots, n-1$ let

$$H_i^{L-} = \{h_k^- \mid h_k \in H^-, k \geq i+1\} \qquad H_i^{L+} = \{h_k^+ \mid h_k \in H^+, k \leq i\}$$

$$H_i^{R+} = \{h_k^+ \mid h_k \in H^+, k \geq i+1\} \qquad H_i^{R-} = \{h_k^- \mid h_k \in H^-, k \leq i\}$$

As well, let H_n^{L+} and H_n^{R-} be the positive and negative half-planes determined by H^+ and H^-, respectively, and let $H_n^{L-} = H_n^{R+} = \emptyset$. Also, for $i = 1, 2, \ldots, n$ let

H_i^L be the union of H_i^{L+} and H_i^{L-}, let H_i^R be the union of H_i^{R+} and H_i^{R-}, let A_i^L and A_i^R denote the intersection of the half-planes in H_i^L and H_i^R, respectively, and let σA_i^L and σA_i^R be the boundaries of A_i^L and A_i^R, respectively. Finally, for $j = 1, 2, \ldots, m$ let H_{ij}^L and H_{ij}^R be the half-planes of H_i^L and H_i^R, respectively, contributed by P_j, and let A_{ij}^L and A_{ij}^R be the intersection of the half-planes of H_{ij}^L and H_{ij}^R, respectively. The following lemma characterizes the shortest transversals of $\boldsymbol{P}$.

Lemma 10 *For $i = 1, 2, \ldots, n$ a line segment s is a shortest transversal of $\boldsymbol{P}$ with slope$(s) \in r_i$ if and only if it is shortest among those line segments that join σA_i^L to σA_i^R and whose containing line separates σB_i^U and σB_i^L.*

Proof. We first show that a line segment s is a transversal of $\boldsymbol{P}$ with *slope*$(s) \in r_i$ if and only if it joins A_i^L to A_i^R such that its containing line separates σB_i^U and σB_i^L. ($\Rightarrow$) Suppose is a transversal of $\boldsymbol{P}$ with *slope*$(s) \in r_i$. Consider some P_j and observe that since s is a transversal of P_j there exists some $p \in s$ such that $p \in P_j$. However, P_j is the intersection of A_{ij}^L and A_{ij}^R which implies that $p \in A_{ij}^L$ and $p \in A_{ij}^R$. Furthermore, since *slope*$(s) \in r_i$, each of A_{ij}^L and A_{ij}^R contains an endpoint of s. Clearly, each of A_i^L and A_i^R contains an endpoint of s and so s joins A_i^L to A_i^R. Let l be the line containing s. Since l is a transversal of $\boldsymbol{P}$ and *slope*$(s) \in r_i$, by Lemma 9, l must separate σB_i^U and σB_i^L, establishing necessity. ($\Leftarrow$) Now suppose that s joins A_i^L to A_i^R such that its containing line separates σB_i^U and σB_i^L. Let l be the line containing s. Since l separates σB_i^U and σB_i^L, by Lemma 9, l is a transversal of $\boldsymbol{P}$ and *slope*$(l) \in r_i$. Observe that each of A_i^L and A_i^R contains an endpoint of s, and so, for $j = 1, 2, \ldots, m$ each of A_{ij}^L and A_{ij}^R contains an endpoint of s. Consider some P_j. Recall that *slope*$(l) \in r_i$ and l is a transversal of $\boldsymbol{P}$. Therefore, since s joins A_{ij}^L to A_{ij}^R, there exists some point $p \in s$ such that $p \in A_{ij}^L$ and $p \in A_{ij}^R$. However, P_j is the intersection A_{ij}^L and A_{ij}^R. Therefore, $p \in P_j$ and s is a transversal of P_j. Clearly, s is a transversal of $\boldsymbol{P}$, establishing sufficiency. To complete the proof, observe that since A_i^L and A_i^R are convex and disjoint, every shortest transversal of $\boldsymbol{P}$ with slope in the range r_i joins σA_i^L to σA_i^R. □

For $i = 1, 2, \ldots, n$ let G_i be the region bounded by σA_i^L and σA_i^R, let H_i be the region bounded by σB_i^U and σB_i^L, and let $Q_i = H_i \cap G_i$. The following lemma characterizes the regions $Q_1, Q_2, \ldots, Q_n$. The proof closely follows that of Lemma 7.

Lemma 11 *For $i = 1, 2, \ldots, n$ the region Q_i is a butterfly polygon.*

Computing a shortest line segment that joins σA_i^L to σA_i^R such that the line containing the line segment separates σB_i^U and σB_i^L, is equivalent to computing a shortest line segment lying inside Q_i with slope in the range r_i. Some line segments that join σA_i^L to σA_i^R and separate σB_i^U and σB_i^L do not have slope in the range r_i. The following lemma characterizes these line segments.

Lemma 12 *For $i = 1, 2, \ldots, n$ every line segment s with slope$(s) \notin r_i$ that joins σA_i^L to σA_i^R and separates σB_i^U and σB_i^L is a transversal of $\boldsymbol{P}$.*

Proof. Consider some region Q_i such that $Q_i \neq \emptyset$. Let $s = [a, b]$ where $a \in \sigma A_i^L$ and $b \in \sigma A_i^R$. Without loss of generality, assume that *slope*$(s) > m_{i+1}$. Let l_a and

l_b be the lines through a and b, respectively, with slope m_{i+1}. Clearly, l_a and l_b are line transversals of $\boldsymbol{P}$. Let a' and b' denote the intersections of l_a with σA_i^R and l_b with σA_i^L, respectively. Also, let $s_a = [a, a']$ and $s_b = [b, b']$. The points a, a', b, b' form a quadrilateral with parallel sides s_a and s_b. Consider any polygon P_j. Since s_a and s_b are transversals of P_j, there exists $p \in P_j$ and $q \in P_j$ such that $p \in s_a$ and $q \in s_b$. By convexity the line segment $[p, q]$ is contained in P_j. Furthermore, s and $[p, q]$ intersect. Therefore s intersects P_j and so s is a transversal of $\boldsymbol{P}$. □

Consider the following variation of **Shortest_Transversal** for computing a shortest transversal of the family $\boldsymbol{P}$ of convex polygons. Initialize the length of the current solution to be infinitely large. Then, for $i = 1, 2, \ldots, n$ determine Q_i. If $Q_i \neq \emptyset$ then compute a shortest line segment s lying inside Q_i that joins the concave chains of Q_i that correspond to σA_i^L and σA_i^R. If *slope*$(s) \notin r_i$ then disregard s. On the other hand, if *slope*$(s) \in r_i$, then if the length of the current solution is larger than the length of s then set the current solution to s and set the length of the current solution to the length of s. Once this process terminates return the current solution if the length of the current solution is not infinitely large, and otherwise return $\emptyset$.

The correctness of the algorithm follows from Lemma 10, Lemma 11, Lemma 12 and the fact that all possible solutions are considered since $r_1, r_2, \ldots, r_n$ partition $\boldsymbol{R}$. Therefore, the algorithm returns a shortest transversal of $\boldsymbol{P}$, provided one exists.

All that remains, is to show that the algorithm runs in $O(n \log n)$ time. To begin, recall that $\sigma B_1^U, \sigma B_2^U, \ldots, \sigma B_n^U$ and $\sigma B_1^L, \sigma B_2^L, \ldots, \sigma B_n^L$ can be computed in $O(n \log n)$ time. We now consider the problem of computing $\sigma A_1^L, \sigma A_2^L, \ldots, \sigma A_n^L$ and $\sigma A_1^R, \sigma A_2^R, \ldots, \sigma A_n^R$. Consider how the pair A_i^L, A_i^R differs from the pair A_{i+1}^L, A_{i+1}^R. If $h_{i+1} \in H^+$ then $h_{i+1}^+ \in H_{i+1}^{L+}$ and $h_{i+1}^+ \notin H_{i+1}^{R+}$, but, $h_{i+1}^+ \notin H_i^{L+}$ and $h_{i+1}^+ \in H_i^{R+}$. Similarly. If $h_{i+1} \in H^-$ then $h_{i+1}^- \notin H_{i+1}^{L-}$ and $h_{i+1}^- \in H_{i+1}^{R-}$, but, $h_{i+1}^- \in H_i^{L-}$ and $h_{i+1}^- \notin H_i^{R-}$. Hence, if $h_{i+1} \in H^+$ then to transform H_{i+1}^L and H_{i+1}^R into H_i^L and H_i^R, we must delete h_{i+1}^+ from H_{i+1}^{L+} and insert h_{i+1}^+ into H_{i+1}^{R+}. Similarly, if $h_{i+1} \in H^-$ then to transform H_{i+1}^L and H_{i+1}^R into H_i^L and H_i^R, we must delete h_{i+1}^- from H_{i+1}^{R-} and insert h_{i+1}^- into H_{i+1}^{L-}.

Observe that each update operation either involves inserting into a set a half-plane whose slope is smaller than any other half-plane in the set, or, deleting from a set a half-plane whose slope is greater than any other half-plane in the set. Recall from the description of Keil's algorithm in Section 3 that the intersection of n positive or negative half-planes can be maintained in $O(n)$ total time when the half-planes are either inserted or deleted in sorted order based on the slopes of the lines determining the half-planes. Hence $O(n)$ time is sufficient to compute A_i^{L+}, A_i^{L-}, A_i^{R+} and A_i^{R-} for $i = 1, 2, \ldots, n$. Furthermore, given A_i^{L+}, A_i^{L-}, A_i^{R+} and A_i^{R-} it is possible to construct A_i^L and A_i^R in $O(\log n)$ time.[26] Clearly, $A_1^L, A_2^L, \ldots, A_n^L$ and $A_1^R, A_2^R, \ldots, A_n^R$ can be computed in $O(n \log n)$ time.

Since σA_i^L and σA_i^R intersect the unbounded edges of σB_i^U and σB_i^L, it is possible to construct Q_i in $O(\log n)$ time. Therefore, once the regions $G_1, G_2, \ldots, G_n$ and $H_1, H_2, \ldots, H_n$ have been determined in $O(n \log n)$ time, we can construct $Q_1, Q_2, \ldots, Q_n$ in $O(n \log n)$ time. Applying Theorem 2 to a polygon Q_i, a required shortest line inside Q_i can be computed in $O(|\sigma B_i^U| + |\sigma B_i^L| + \log n)$ time. However, Edelsbrunner, Guibas and Sharir showed that

$$\sum_{i=1}^{n} |\sigma B_i^U| + |\sigma B_i^L| \leq kn\alpha(n)$$

for some constant k.[7] Hence, computing the required shortest line segments inside $Q_1, Q_2, \ldots, Q_n$ requires $O(n\alpha(n) + n\log n)$ time, which is $O(n\log n)$. Clearly, the running time of the algorithm is $O(n\log n)$. Furthermore, note that the assumed labeling of the lines of $\boldsymbol{H}$ does not affect the asymptotic running time of the algorithm. We have therefore established the following theorem.

Theorem 6 *Given the family* $\boldsymbol{P}$ *of convex polygons, a shortest transversal of* $\boldsymbol{P}$ *can be computed in* $\Theta(n\log n)$ *time.*

6. Concluding Remarks

In this paper we presented $O(n\log n)$ time algorithms for computing a shortest transversal of a family of n lines, a family of n line segments, and a family of convex polygons with a total of n vertices. For the cases of line segments and convex polygons, the algorithms are optimal. Our solutions to the line segment and polygon versions of the shortest transversal problem are based on computing shortest distances inside butterfly polygons.

One open problem is to establish a non trivial lower bound for the problem of computing a shortest transversal of a family of lines. Other problems that remain open involve the development of efficient shortest transversal algorithms for families of objects not considered in this paper.

References

1. L. Danzer, B. Grünbaum and V. Klee, "Helly's theorem and its relatives", in *Convexity, Proc. Symposia in Pure Mathematics,* (American Mathematical Society, 1963).
2. J. E. Goodman, R. Pollack and R. Wenger, "Geometric transversal theory", in *Recent Trends in Discrete and Computational Geometry,* ed. J. Pach, to appear.
3. J. O'Rourke, "An on-line algorithm for fitting straight lines between data ranges", *Comm. ACM* **24** (1981) 574–578.
4. N. Megiddo, "Linear-time algorithms for linear programming in R^3 and related problems", *SIAM J. Comput.* **12** (1983) 759–776.
5. M. E. Dyer, "Linear time algorithms for two- and three-variable linear programs", *SIAM J. Comput.* **13** (1984) 31–45.
6. H. Edelsbrunner, H. A. Maurer, F. P. Preparata, A. L. Rosenberg, E. Welzl and D. Wood, "Stabbing line segments", *BIT* **22** (1982) 274–281.
7. H. Edelsbrunner, L. J. Guibas and M. Sharir, "The upper envelope of piecewise linear functions: algorithms and applications", *Discr. Comput. Geom.* **4** (1989) 311–336.
8. J. Hershberger, "Finding the upper envelope of n line segments in $O(nlogn)$ time", *Inform. Process. Lett.* **33** (1989) 169–174.

9. M. Atallah and C. Bajaj, "Efficient algorithms for common transversals", *Inform. Process. Lett.* **25** (1987) 87–91.

10. C. Bajaj and M. Li, "On the duality of intersection and closest points", *Proc. 21st Annual Allerton Conference on Communication and Control* (1983) pp. 459–460.

11. M. E. Houle, H. Imai, K. Imai and J.-M. Robert, "Weighted orthogonal linear L_∞-approximation and applications", *Proc. Workshop on Algorithms and Data Structures '89* (1989) pp. 183–191.

12. H. Edelsbrunner, "Finding transversals for sets of simple geometric figures", *Theo. Comp. Sc.* **35** (1985) 55–69.

13. D. Avis, J.-M. Robert and R. Wenger, "Lower bounds for line stabbing", *Inform. Process. Lett.* **33** (1989) 59–62.

14. D. Avis and M. Doskas, "Algorithms for high dimensional stabbing problems", *Discr. Appl. Math.* **27** (1990) 39–48.

15. P. Egyed and R. Wenger, "Stabbing pairwise disjoint translates in linear time", *Proc. 5th Annual ACM Symposium on Computational Geometry,* (1989) pp. 364–369.

16. P. Egyed and R. Wenger, "Ordered stabbing of pairwise disjoint convex sets in linear time", *Discr. Appl. Math.* **31** (1991) 133–140.

17. B. K. Bhattacharya and G. T. Toussaint, "Computing shortest transversals", *Computing* **46** (1991) 93–119.

18. B. K. Bhattacharya, P. Egyed and G. T. Toussaint, "Computing the wingspan of a butterfly", manuscript in preparation.

19. B. K. Bhattacharya, P. Egyed and G. T. Toussaint, "Computing the wingspan of a butterfly", *Proc. 3rd Canadian Conference on Computational Geometry,* (1991) pp. 88–91.

20. S. Suri, "Computing the envelope of a set of lines" (1985), extended abstract.

21. H. Edelsbrunner and L. J. Guibas, "Topologically sweeping an arrangement", *Proc. 18th Annual ACM Symposium on Theory on Computing,* (1986) pp. 389–403.

22. G. Vegter, "Computing the bounded region determined by finitely many lines in the plane", Technical Report CS 8703, University of Groningen, 1987.

23. M. Keil, "A simple algorithm for determining the envelope of a set of lines", *Inform. Process. Lett.* **39** (1991) 121–124.

24. H. Edelsbrunner, "Computing the extreme distances between two convex polygons", *J. Alg.* **6** (1985) 213–224.

25. M. McKenna and G. T. Toussaint, "Finding the minimum vertex distance between two disjoint convex polygons in linear time", *Comput. Math. Appl.* **11** (1985) 1227–1242.

26. M. Keil and G. T. Toussaint, "Detecting and computing intersections of convex chains", *Proc. 3rd Canadian Conference on Computational Geometry* (1991) pp. 7–10.