

ROBUST PLANE SWEEP FOR INTERSECTING SEGMENTS*

JEAN-DANIEL BOISSONNAT[†] AND FRANCO P. PREPARATA[‡]

Abstract. In this paper, we reexamine in the framework of robust computation the Bentley–Ottmann algorithm for reporting intersecting pairs of segments in the plane. This algorithm has been reported as being very sensitive to numerical errors. Indeed, a simple analysis reveals that it involves predicates of degree 5, presumably never evaluated exactly in most implementations. Within the exact-computation paradigm we introduce two models of computation aimed at replacing the conventional model of real-number arithmetic. The first model (predicate arithmetic) assumes the exact evaluation of the signs of algebraic expressions of some degree, and the second model (exact arithmetic) assumes the exact computation of the value of such (bounded-degree) expressions. We identify the characteristic geometric property enabling the correct report of all intersections by plane sweeps. Verification of this property involves only predicates of (optimal) degree 2, but its straightforward implementation appears highly inefficient. We then present algorithmic variants that have low degree under these models and achieve the same performance as the original Bentley–Ottmann algorithm. The technique is applicable to a more general case of curved segments.

Key words. computational geometry, segment intersection, plane sweep, robust algorithms

AMS subject classifications. 68Q20, 68Q25, 68Q40, 68U05

PII. S0097539797329373

1. Introduction. As is well known, computational geometry has traditionally adopted the arithmetic model of exact computation over the real numbers. This model has been extremely productive in terms of algorithmic research, since it has permitted a vast community to focus on the elucidation of the combinatorial (topological) properties of geometric problems, thereby leading to sophisticated and efficient algorithms. Such an approach, however, has a substantial shortcoming, since all computer calculations have finite precision, a feature which affects not only the quality of the results but even the validity of specific algorithms. In other words, in this model, algorithm correctness does not automatically translate into program correctness. In fact, there are several reports of failures of implementations of theoretically correct algorithms (see, e.g., [For87, Hof89]). This state of affairs has engendered a vigorous debate within the research community, as is amply documented in the literature. Several proposals have been made to remedy this unsatisfactory situation. They can be split into two broad categories according to whether they perform exact computations (see, e.g., [BKM⁺95, FV93, Yap97, She96]) or approximate computations (see, e.g., [Mil88, HHK89, Mil89]).

This paper fine-tunes the exact-computation paradigm. The numerical computations of a geometric algorithm are basically of two types: tests (predicates) and constructions, each with clearly distinct roles. Tests are associated with branching

*Received by the editors October 20, 1997; accepted for publication (in revised form) September 10, 1998; published electronically March 15, 2000. This work was partially supported by ESPRIT LTR 21957 (CGAL) and by the U.S. Army Research Office under grant DAAH04-96-1-0013. Part of the work was done while the first author was visiting the Center for Geometric Computing at Brown University.

<http://www.siam.org/journals/sicomp/29-5/32937.html>

[†]INRIA, 2004 Route des Lucioles, BP 93, 06902 Sophia-Antipolis, France (Jean-Daniel.Boissonnat@sophia.inria.fr).

[‡]Center for Geometric Computing, Computer Science Department, Brown University, 115 Waterman Street, Providence, RI 02912-1910, (franco@cs.brown.edu).

decisions in the algorithm that determine the flow of control, whereas constructions are needed to produce the output data. While approximations in the execution of constructions are often acceptable, approximations in the execution of tests may produce incorrect branching, leading to the inconsistencies which are the object of the criticisms leveled against geometric algorithms. The exact-computation paradigm therefore requires that tests be executed with total accuracy. This will guarantee that the result of a geometric algorithm will be topologically correct albeit geometrically approximate. This also means that robustness is in principle achievable if one is willing to employ the required precision. The reported failures of structurally correct algorithms are entirely attributable to noncompliance with this criterion.

Therefore, geometric algorithms can also be characterized on the basis of the complexity of their predicates. The complexity of a predicate is expressed by the degree of a homogeneous polynomial embodying its evaluation. The degree of an algorithm is the maximum degree of its predicates, and an algorithm is robust if the adopted precision matches the degree requirements.

The “degree criterion” is a design principle aimed at developing low-degree algorithms. This approach involves reexamining under the degree criterion the rich body of geometric algorithms known today, possibly without negatively affecting traditional algorithmic efficiency. A previous paper [LPT96] considered as an illustration of this approach the issue of proximity queries in two and three dimensions. As an additional case of degree-driven algorithm design, in this paper we confront another class of important geometric problems, which have caused considerable difficulties in actual implementations: plane sweep problems for sets of segments. As we shall see, plane sweep applications involve a number of predicates of different degree and algorithmic power. Their analysis not only will lead to new and robust implementations (an outcome of substantial practical interest) but will elucidate on a theoretical level some deeper issues pertaining to the structure of several related problems and the mechanism of plane sweeps.

2. Three problems associated with intersecting segments. Given is a finite set \mathcal{S} of line segments in the plane. Each segment is defined by the coordinates of its two endpoints. We discuss the three following problems (see Figure 2.1):

- Pb1: Report the pairs of segments of \mathcal{S} that intersect.
- Pb2: Construct the arrangement \mathcal{A} of \mathcal{S} , i.e., the incidence structure of the graph obtained interpreting the union of the segments as a planar graph.
- Pb3: Construct the trapezoidal map \mathcal{T} of \mathcal{S} . \mathcal{T} is obtained by drawing two vertical line segments (*walls*), one above and one below each endpoint of the segments and each intersection point. The walls are extended either until they meet another segment of \mathcal{S} or to infinity.

Let S_1, \dots, S_n be the segments of \mathcal{S} , and let k be the number of intersecting pairs. We say that the segments are in *general position* if any two intersecting segments intersect in a single point, and all endpoints and intersection points are distinct.

The number of intersection points is no more than the number of intersecting pairs of segments, and both are equal if the segments are in general position. Therefore, the number of vertices of \mathcal{A} is at most k , the number of edges of \mathcal{A} is at most $n + 2k$, and the number of vertical walls in \mathcal{T} is at most $2(n + k)$, the bounds being tight when the segments are in general position. Thus the sizes of both \mathcal{A} and \mathcal{T} are $O(n + k)$. We didn’t consider here the two-dimensional faces of either \mathcal{A} or \mathcal{T} . Including them would not change the problems we address.




FIG. 2.1. A set of segments \mathcal{S} , the corresponding arrangement \mathcal{A} and its trapezoidal map \mathcal{T} .

3. Algebraic degree and arithmetic models. It is well known that the efficient algorithms that solve Pb1–Pb3 are very unstable when implemented using nonexact arithmetic, and several frustrating experiences have been reported [For85]. This motivates us to carefully analyze the predicates involved in those algorithms. We first introduce here some terminology borrowed from [LPT96]. See also [Bur96, For93]. We consider each input data (i.e., coordinates of an endpoint of some segment of \mathcal{S}) as a *variable*.

An *elementary predicate* is the sign $-$, 0 , or $+$ of a homogeneous multivariate polynomial whose arguments are a subset of the input variables. The degree of an elementary predicate is defined as the maximum degree of the irreducible factors (over the rationals) of the polynomials that occur in the predicate and that do not have a constant sign. A *predicate* is more generally a Boolean function of elementary predicates. Its degree is the maximum degree of its elementary predicates.

The *degree of an algorithm* A is defined as the maximum degree of its predicates.

The *degree of a problem* P is defined as the minimum degree of any algorithm that solves P .

In most problems in computational geometry, $d = O(1)$. However, as d affects the speed and/or robustness of an algorithm, it is important to measure d precisely.

In the rest of this paper we consider the degree as an additional measure of algorithmic complexity. Note that, qualitatively, degree and memory requirements are similar, since the arithmetic capabilities demanded by a given degree must be available, albeit they may be never resorted to in an actual run of the algorithm (since the input may be such that predicates may be evaluated reliably with lower precision).

We will consider two arithmetic models. In the first one, called the *predicate arithmetic of degree d* , the only numerical operations that are allowed are the evaluations of predicates of degree at most d . Algorithms of degree d can therefore be implemented exactly in the predicate arithmetic model of degree d . This model is motivated by recent results that show that evaluating the sign of a polynomial expression may be faster than computing its value (see [ABD⁺97, BY97, BEPP97, Cla92, She96]). This model, however, is very conservative since the nonavailability of the arithmetics required by a predicate is assimilated to an entirely random choice of the value of the predicate.

The second model, called the *exact arithmetic of degree d* , is more demanding. It assumes that values (and not just signs) of polynomials of degree at most d be represented and computed exactly (i.e., roughly as d -fold precision integers). However, higher-degree operations (e.g., a multiplication operation of whose factors is a d -fold precision integer) are appropriately rounded. Typical rounding is rounding to the nearest representable number, but less accurate rounding can also be adequate as will be demonstrated later.

Let P be a predicate (polynomial) of degree d . We ignore for simplicity the size of the coefficients of P , because they are typically small constants; if the input data are all b -bit integers, the size of each monomial in predicate P is upper bounded by 2^{bd} . Moreover, let v be the number of variables that occur in a predicate; for most geometric problems and, in particular, for those considered in this paper, v is a small number. Since the polynomial P is homogeneous, it may contain only highest-degree monomials, whose number is bounded by v^d . It follows that an algorithm of degree d requires precision $p \leq db + d \log v = db + O(1)$ in the exact arithmetic model of degree d .

4. The predicates and the degree of problems Pb1–Pb3. In this section, we analyze the degree of problems Pb1–Pb3 (in subsections 4.1–4.3) and of the standard algorithms for solving these problems (in subsection 4.4).

We use the following notations. The coordinates of point A_i are denoted x_i and y_i . $A_i <_x A_j$ means that the x -coordinate of point A_i is smaller than the x -coordinate of point A_j . The same is true for $<_y$. $[A_i A_j]$ denotes the line segment whose left and right endpoints are, respectively, A_i and A_j , while $(A_i A_j)$ denotes the line containing $[A_i A_j]$. $A_i <_y (A_j A_k)$ means that point A_i lies below line $(A_j A_k)$.

4.1. Predicates. Pb1 requires only that we check if two line segments intersect (Predicate 2' below).

Pb2 requires in addition the ability to sort intersection points along a line segment (Predicate 4 below).

Pb3 requires the ability to execute all the predicates listed below.

- Predicate 1: $A_0 <_x A_1$.
- Predicate 2: $A_0 <_y (A_1 A_2)$.
- Predicate 2': $[A_0 A_1] \cap [A_2 A_3] \neq \emptyset$.
- Predicate 3: $A_0 <_x [A_1 A_2] \cap [A_3 A_4]$.
- Predicate 4: $[A_0 A_1] \cap [A_2 A_3] <_x [A_0 A_1] \cap [A_4 A_5]$.
- Predicate 5: $[A_0 A_1] \cap [A_2 A_3] <_x [A_4 A_5] \cap [A_6 A_7]$.

Two other predicates appear in some algorithms that report segment intersections.

- Predicate 3': $(x = x_0) \cap [A_1 A_2] <_y (x = x_0) \cap [A_3 A_4]$.
- Predicate 4': $[A_0 A_1] \cap [A_2 A_3] <_y (A_4 A_5)$.

4.2. Algebraic degree of the predicates. We now analyze the algebraic degree of the predicates introduced above.

PROPOSITION 4.1. *The degree of Predicates i and i' ($i = 1, \dots, 5$) is i .*

Proof. We first provide explicit formulae for the predicates.

Evaluating Predicate 2 is equivalent to evaluating the sign of

$$\text{orient}(A_0 A_1 A_2) = \begin{vmatrix} x_0 & x_1 & x_2 \\ y_0 & y_1 & y_2 \\ 1 & 1 & 1 \end{vmatrix} = \begin{vmatrix} x_1 - x_0 & x_2 - x_0 \\ y_1 - y_0 & y_2 - y_0 \end{vmatrix}.$$

Predicate 2' can be implemented as follows for the case $A_0 <_x A_2$ (otherwise we exchange the roles of $[A_0 A_1]$ and $[A_2 A_3]$):

```

if  $A_1 <_x A_2$  then return false
if  $A_3 <_x A_1$ 
    if  $\text{orient}(A_0 A_1 A_2) \times \text{orient}(A_0 A_1 A_3) < 0$  then return true
    else return false
else
    if  $\text{orient}(A_0 A_1 A_2) \times \text{orient}(A_2 A_3 A_1) > 0$  then return true
    else return false
    
```

Therefore, in all cases, Predicate 2' reduces to Predicate 2.

The intersection point $I = [A_i A_j] \cap [A_k A_l]$ is given by

$$(4.1) \quad I = A_i + (A_j - A_i) \frac{N_I}{D_I}$$

with $N_I = \text{orient}(A_i A_k A_l)$ and

$$\begin{aligned} D_I &= \begin{vmatrix} x_j - x_i & x_k - x_l \\ y_j - y_i & y_k - y_l \end{vmatrix} \\ &= \text{orient}(A_i A_j A_k) - \text{orient}(A_i A_j A_l) \\ &\stackrel{\text{def}}{=} \text{Orient}(A_i A_j A_k A_l). \end{aligned}$$

Predicate 4' reduces to evaluating $\text{orient}(I, A_0, A_1)$, where I is the intersection of $[A_2 A_3]$ and $[A_4 A_5]$. It follows from (4.1) that this is equivalent to evaluating the sign of $\text{orient}(A_2 A_3 A_4 A_5)$ and of

$$\text{orient}(A_0 A_1 A_2) \times \text{Orient}(A_2 A_3 A_4 A_5) - \text{orient}(A_2 A_4 A_5) \times \text{Orient}(A_0 A_1 A_2 A_3).$$

Predicates 3–5: Explicit formulas for Predicates 3, 4, and 5 can be immediately deduced from the coordinates of the intersection points $I = [A_0 A_1] \cap [A_2 A_3]$ and $J = [A_4 A_5] \cap [A_6 A_7]$ which are given by (4.1). If $A_4 A_5 = A_0 A_1$, it is clear from (4.1) that $(x_1 - x_0)$ is a common factor of $x_I - x_0$ and $x_J - x_0$.

If $[A_1A_2]$ and $[A_3A_4]$ do not intersect, Predicate 3' reduces to Predicate 2. Otherwise, it reduces to Predicate 3.

The above discussion shows that the degree of predicates i and i' is *at most* i . To establish that it is *exactly* i , we have shown in the appendix that the polynomials of Predicates 2, 3, 4', and 5, as well as the factor other than $(x_1 - x_0)$ involved in Predicate 4, are irreducible over the rationals.

It follows that the proposition is proved for all predicates. \square

Recalling the requirements of the various problems in terms of predicates, we have the following proposition.

PROPOSITION 4.2. *The algebraic degrees of Pb1, Pb2, and Pb3 are, respectively, 2, 4, and 5.*

4.3. Implementation of Predicate 3 with exact arithmetic of degree 2.

As they will be useful in section 7, in this subsection we present two approximate implementations of Predicate 3 (of degree 3) under the exact arithmetic of degree 2.

From (4.1) we know that Predicate 3 can be written as

$$(4.2) \quad x_{01} D < x_{21} N,$$

where $x_{01} = x_0 - x_1$, $x_{21} = x_2 - x_1$, $D = |\text{Orient}(A_1A_2A_3A_4)|$, and $N = \text{sign}(D) \times \text{orient}(A_1A_3A_4)$.

We stipulate to employ floating-point arithmetic conforming to the IEEE 754 standard [Gol91]. In this standard, simple precision allows us to represent b -bit integers with $b = 24$ and double precision allows us to represent b' -bit integers with $b' = 2b+5 = 53$. The coordinates of the endpoints of the segments are represented in simple precision, and the computations are carried out in double precision. We denote \oplus , \otimes , and \circlearrowleft the rounded arithmetic operations $+$, \times , and $/$. In the IEEE 754 standard, all four arithmetic operations are exactly rounded; i.e., the computed result is the floating-point number that best approximates the exact result.

Since $\text{Orient}(A_1A_2A_3A_4)$ and $\text{orient}(A_1A_3A_4)$ are $(2b+3)$ -bit integers, the four terms x_{01} , x_{21} , N , and D in inequality (4.2) can be computed exactly in double precision, and the following monotonicity property is a direct consequence of exact rounding of arithmetic operations.

Monotonicity property 1: $x_{01} \otimes D < x_{21} \otimes N \implies x_{01} \times D < x_{21} \times N$.

This implies that the comparison between the two computed expressions $x_{01} \otimes D$ and $x_{21} \otimes N$ evaluates Predicate 3 except when these numbers are equal.

Since in most algorithms, an intersection point is compared with many endpoints, it is more efficient to compute and store the coordinates of each intersection point and to perform comparisons with the computed abscissae rather than to evaluate (4.2) repeatedly. We now illustrate an effective rounding procedure of the x -coordinates of intersection points which gives an alternative approximate implementation of Predicate 3.

LEMMA 4.3. *If the coordinates of the endpoints of the segments are simple precision integers, then the abscissa x_I of an intersection point can be rounded to one of its two nearest simple precision integers using only double precision floating-point arithmetic operations.*

Proof. We assume that the coordinates of the endpoints of the segments are represented as b -bit integers stored as simple precision floating-point numbers. The computations are carried out in double precision.

The rounded value \tilde{x}_I of x_I is given by

$$\tilde{x}_I = \lfloor ((x_{21} \otimes N) \odot D) \rfloor \oplus x_1,$$

where $\lfloor X \rfloor$ denotes the integer nearest to X (with any tie-breaking rule). If $\varepsilon = 2^{-b'}$ is a strict bound to the modulus of the relative error of all arithmetic operations, $\tilde{X} = (x_{21} \otimes N) \odot D$ satisfies the following relations:

$$\frac{x_{21}N}{D}(1 - 2\varepsilon) \approx \frac{x_{21}N}{D}(1 - \varepsilon)^2 < \tilde{X} < \frac{x_{21}N}{D}(1 + \varepsilon)^2 \approx \frac{x_{21}N}{D}(1 + 2\varepsilon).$$

As $\frac{x_{21}N}{D} = x_I - x_1 \leq 2^{b+1}$, we obtain

$$|(x_{21} \otimes N) \odot D - x_{21}N/D| \lesssim 2^{b+1}22^{-b'} = 2^{-b-3} \ll 1.$$

We round \tilde{X} to the nearest integer $\lfloor \tilde{X} \rfloor$. Since $\lfloor \tilde{X} \rfloor$ and x_1 are $(b+1)$ -bit integers, there is no error in the addition. Therefore, \tilde{x}_I is a $(b+2)$ -bit integer and the absolute error on \tilde{x}_I is smaller than 1. \square

It follows that, under the hypothesis of the lemma, if E is an endpoint, I is an intersection point, and \tilde{I} is the corresponding rounded point, the following monotonicity property holds.

Monotonicity property 2:
$$\begin{aligned} \tilde{I} <_x E &\implies I <_x E, \\ E <_x \tilde{I} &\implies E <_x I. \end{aligned}$$

This implies that the comparison between the x -coordinates of \tilde{I} and E evaluates Predicate 3, except when the abscissae of \tilde{I} and E coincide.

Notice that the monotonicity property does not necessarily hold for two intersection points.

Remark 1. A result similar to Lemma 4.3 has been obtained by Priest [Pri92] for points with floating-point coordinates. More precisely, if the endpoints of the segments are represented as simple precision floating-point numbers, Priest [Pri92] has proposed a rather complicated algorithm that uses double precision floating-point arithmetic and rounds x_I to the nearest simple precision floating-point number. Since it applies to the integers as well, this stronger result also implies the monotonicity property.

4.4. Algebraic degree of the algorithms. The naive algorithm for detecting segment intersections (Pb1) evaluates $\Theta(n^2)$ Predicates 2' and thus is of degree 2, which is degree-optimal by the proposition above. Although the time-complexity of the naive algorithm is worst-case optimal, since $0 \leq k \leq \frac{1}{2}n(n-1)$, it is worth looking for an output sensitive algorithm whose complexity depends on both n and k . Chazelle and Edelsbrunner [CE92] have shown that $\Omega(n \log n + k)$ is a lower bound for Pb1 and therefore also for Pb2 and Pb3. A very recent algorithm of Balaban [Bal95] solves Pb1 optimally in $O(n \log n + k)$ time using $O(n)$ space. This algorithm does not solve Pb2 nor Pb3 and, since it uses Predicate 3', its degree is 3.

Pb2 can be solved by first solving Pb1 and subsequently sorting the reported intersection points along each segment. This can easily be done in $O((n+k) \log n)$ time by a simple algorithm of degree 4 using $O(n)$ space. A direct (and asymptotically more efficient) solution to Pb2 has been proposed by Chazelle and Edelsbrunner [CE92]. Its time complexity is $O(n \log n + k)$, and it uses $O(n+k)$ space. Their algorithm, which constructs the arrangement of the segments, is of degree 4.

A solution to Pb3 can be deduced from a solution to Pb2 in $O(n+k)$ time using a very complicated algorithm of Chazelle [Cha91]. A deterministic and simple algorithm due to Bentley and Ottmann [BO79] solves Pb3 in $O((n+k) \log n)$ time, which

is slightly suboptimal, using $O(n)$ space. This classical algorithm uses the sweep-line paradigm and evaluates $O((n+k)\log n)$ predicates of all types discussed above, and therefore has degree 5. Incremental randomized algorithms [CS89, BDS⁺92] construct the trapezoidal map of the segments and thus solve Pb3 and have degree 5. Their time complexity and space requirements are optimal (although only as expected performances).

In this paper, we revisit the Bentley–Ottmann algorithm and show that a variant of degree 3 (instead of 5) [Mye85, Sch91] can solve Pb1 with no sacrifice of performance (section 6.1). Although this algorithm is slightly suboptimal with respect to time complexity, it is much simpler than Balaban’s algorithm. We also present two variants of the sweep-line algorithm. The first one (section 6.2) uses only predicates of degree at most 2 and applies to the restricted but important special case where the segments belong to two subsets of nonintersecting segments. The second one (section 7) uses the exact arithmetic of degree 2. All these results are based on a (nonefficient) lazy sweep-line algorithm (to be presented in section 5) that solves Pb1 by evaluating predicates of degree at most 2.

Remark 2. When the segments are not in general position, the number s of intersection points can be less than the number k of intersecting pairs. In the extreme, $s = 1$ and $k = n(n-1)/2$. Some algorithms can be adapted so that their time complexities depend on s rather than k [BMS94]. However, a lower bound on the degree of such algorithms is 4 since they must be able to detect if two intersection points are identical, therefore to evaluate Predicate 4’.

5. A lazy sweep-line algorithm. Let \mathcal{S} be a set of n segments whose endpoints are E_1, \dots, E_{2n} . For a succinct review, the standard algorithm first sorts E_1, \dots, E_{2n} by increasing x -coordinates and stores the sorted points in a priority queue X , called the *event schedule*. Next, the algorithm begins sweeping the plane with a vertical line L and maintains a data structure Y that represents a subset of the segments of \mathcal{S} (those currently intersected by L , ordered according to the ordinates of their intersections with L). Intersections are detected in correspondence of adjacencies created in Y , either by insertion/deletion of segment endpoints or by order exchanges at intersections. An intersection, upon detection, is inserted into X according to its abscissa. Of course, a given intersection may be detected several times. Multiple detections can be resolved by performing a preliminary membership test for an intersection in X and omitting insertion if the intersection has been previously recorded. An intersection is reported when the sweep-line reaches its abscissa. We stipulate to use another policy to resolve multiple detections, namely to remove from X an intersection point I whose associated segments are no longer adjacent in Y . Event I will be reinserted in X when the segments become again adjacent in Y . This policy has also the advantage of reducing the storage requirement of Bentley–Ottmann’s algorithm to $O(n)$ [Bro81].

5.1. Description of the lazy algorithm. We now describe a modification of the sweep-line algorithm that does not need to process the intersection points by increasing x -coordinates.

First, the algorithm sorts the endpoints of the segments by increasing x -coordinates into an array X . Let E_1, \dots, E_{2n} be the sorted list of endpoints.

Then the algorithm starts processing *events* and maintains a dictionary Y that stores an ordered subset of the line segments. The events consist of the endpoints and of a subset of the pairs of segments that intersect and are adjacent in Y . Such pairs are called *processable* and will be precisely defined below. The algorithm processes the endpoints by increasing order of their x -coordinates, but, contrary to the standard

algorithm, the processable pairs are processed in any order. As a consequence, two intersection points or even an intersection point and an endpoint won't necessarily be processed in the order of their x -coordinates, and Y won't necessarily represent the ordered set of segments intersecting some vertical line L (as in the standard algorithm).

The events are processed in the same way as in the standard algorithm, i.e., processing an endpoint means either the insertion of the corresponding segment into Y or its deletion from it (as appropriate), and processing an intersection means its report and the exchange of the two segments involved. The only difference is that we maintain a set P of processable pairs: each time a pair of intersecting segments become adjacent in Y , we check whether the pair is processable and, in the affirmative, add it to P .

While there are processable pairs, the algorithm extracts any of them from P and processes it. When there are no more processable pairs, the algorithm proceeds to the next endpoint. When there are no more processable pairs and no more endpoints to be processed, the algorithm stops.

To complete the description of the algorithm, we need to define the *processable* pairs. The definition rests on the notions of *active* and *prime* pairs to be given below. We need the following notations. We denote by $L(E_i)$ the vertical line passing through E_i . Slab (E_i, E_{i+1}) denotes the open vertical slab bounded by $L(E_i)$ and $L(E_{i+1})$, and $(E_i, E_{i+1}]$ denotes the semiclosed slab obtained by adjoining line $L(E_{i+1})$ to the open slab (E_i, E_{i+1}) . For two segments S_k and S_l , we denote by A_{kl} their rightmost left endpoint, by B_{kl} their leftmost right endpoint, and by I_{kl} their intersection point when they intersect (see Figure 5.1). In addition, W_{kl} denotes the set of segment endpoints that belong to the (closed) region bounded by the vertical lines $L(A_{kl})$ and $L(B_{kl})$ and by the two segments (a trapezoid if the two segments do not intersect, a double-wedge otherwise). We denote by E_{kl} the most recently processed element of W_{kl} and by F_{kl} the element of W_{kl} to be processed next. (Note that E_{kl} and F_{kl} are always defined, since they may respectively coincide with A_{kl} and B_{kl} .) Last, we define sets $W_{kl}^+(E_{kl})$ and $W_{kl}^-(E_{kl})$ as follows. If S_k and S_l do not intersect, $W_{kl}^+(E_{kl}) = \emptyset$ and $W_{kl}^-(E_{kl})$ consists of all points $E \in W_{kl}$, $E_{kl} \leq_x E$. Otherwise, an endpoint $E \in W_{kl}$ belongs to $W_{kl}^+(E_{kl})$ (respectively, to $W_{kl}^-(E_{kl})$) if $E_{kl} \leq_x E$ and if the slab $(E_{kl}, E]$ does (respectively, does not) contain I_{kl} .¹

DEFINITION 5.1. *Let (S_k, S_l) be a pair of segments, and assume without loss of generality that $S_k \cap L(E_{kl}) <_y S_l \cap L(E_{kl})$. The pair is said to be active if the following conditions are satisfied:*

- (1) S_k and S_l are adjacent in Y ,
- (2) $S_k < S_l$ in Y ,
- (3) $F_{kl} \in W_{kl}^+(E_{kl})$ (emptiness condition).

Observe that the emptiness condition implies that the segments intersect (since $W_{kl}^+(E_{kl}) = \emptyset$ if they do not). We now identify a subset of the active pairs, whose processing, as we shall see, has priority.

DEFINITION 5.2. *An active pair of segments (S_k, S_l) is said to be prime if the next endpoint to be processed belongs to W_{kl} (i.e., it coincides with $F_{kl} \in W_{kl}^+(E_{kl})$).*

It should be noted that deciding if a pair of intersecting segments is active or prime reduces to the evaluation of Predicates 2 only and the condition $F_{kl} \in W_{kl}^+(E_{kl})$ should

¹For line segments or even pseudosegments, $W_{kl}^-(E_{kl})$ and $W_{kl}^+(E_{kl})$ do not depend on E_{kl} . However, we keep E_{kl} as a parameter in preparation for the more general case of monotone arcs. See Remark 4 at the end of this section.




FIG. 5.1. For the definitions of $W_{kl}^-(E_{kl})$ and $W_{kl}^+(E_{kl})$.

not be construed as implying the comparison of the abscissae of I_{kl} and F_{kl} (Predicate 3). Indeed, if $S_k \cap L(A_{kl}) <_y S_l \cap L(A_{kl})$ (which can be decided by Predicate 2), then the emptiness condition corresponds to $F_{kl} <_y S_k$ and $S_l <_y F_{kl}$.

For reasons that will be clear below, the set of processable pairs is not specified in the finest detail in this section, since several different implementations are possible. We require only that, at any time, the two following assertions remain true.

- (1) All prime pairs are processable.
- (2) All processable pairs are active.

In other words, the next endpoint may be processed once there are no more prime pairs, without placing any deadline on the processing of the current active pairs as long as they are not prime.

In the rest of this section, we will simply assume that we have an oracle at our disposal that can decide if a given pair is processable. Clearly, for some instances of the lazy algorithm (e.g., when considering all active pairs as processable), the oracle can be implemented with Predicate 2 only. In such a case, the lazy algorithm involves only Predicates 1 and 2 and is of degree 2 by Proposition 4.1. We will prove its correctness in the next subsection.

The important issue of efficiently detecting the processable pairs will be considered in sections 6 and 7 where several oracles will be introduced. These instances of the lazy algorithm will still be correct but, in some cases, will have a degree higher than 2.

5.2. Correctness of the lazy algorithm. Let $Y^-(E_i)$ and $Y^+(E_i)$ be, respectively, snapshots of the data structure Y immediately before and after processing event E_i , $i = 1, \dots, 2n$. Observe that $Y^-(E_i)$ and $Y^+(E_i)$ differ only by the segment S that has E_i as one of its endpoints. Let $Y(E_i) = Y^-(E_i) \cup Y^+(E_i)$. The order relation in Y is denoted by $<$.

THEOREM 5.3. *If Predicates 1 and 2 are evaluated exactly, the described lazy sweep-line algorithm will detect all pairs of segments that intersect.*

Proof. The algorithm (correctly) sorts the endpoints E_1, \dots, E_{2n} of the segments by increasing x -coordinates into X . Consequently, the set of segments that intersect $L(E)$ and the set of segments in $Y(E)$ coincide for any endpoint E . The proof of the theorem is articulated now as two lemmas and their implications.

LEMMA 5.4. *Two segments have exchanged their positions in Y if and only if they intersect and if the pair has been processed.*

Proof. Let us consider two segments, say S_k and S_l , that do not intersect. Without loss of generality, let $S_k < S_l$ in $Y(A_{kl})$. Assume for a contradiction that $S_l < S_k$ in $Y^-(B_{kl})$. S_k and S_l cannot exchange their positions because they will never form an active pair. Therefore, $S_l < S_k$ in $Y^-(B_{kl})$ can happen only if there exists a segment S_m , $m \neq k, l$, that at some stage in the execution of the algorithm was present in Y together with S_k and S_l and caused one of the following two events to occur.

- (1) $S_m > S_l$ and the positions of S_m and S_k are exchanged in Y .
- (2) $S_m < S_k$ and the positions of S_m and S_l are exchanged in Y .

In both cases, the segments that exchange their positions are not consecutive in Y , violating condition 1 of Definition 5.1.

Therefore, two segments can exchange their positions in Y only if they intersect and this can happen only when their intersection is processed. Moreover, when the intersection has been processed, the segments are no longer active and cannot exchange their positions a second time. \square

We say that an endpoint E of S is *correctly placed* if and only if the subset of the segments that are below E (in the plane) coincides with the subset of the segments $< S$ in $Y(E)$, i.e.,

$$\forall S' \in Y(E), \quad S' < S \iff S' \cap L(E) <_y S \cap L(E).$$

Otherwise, E is said to be *misplaced*. (Note that $S \cap L(E)$ coincides with E .)

LEMMA 5.5. *If Predicates 1 and 2 are evaluated exactly, both endpoints of every segment are correctly placed.*

Proof. Assume, for a contradiction, that E of S is the *first* endpoint to be misplaced by the algorithm.

CLAIM 1. *E can be misplaced only if there exist at least two intersecting segments S_k and S_l in $Y^-(E)$ such that E belongs to W_{kl} .*

Proof. First recall that Predicate 2 is the only predicate involved in placing S in Y .

Consider first the case where E is the left endpoint of S . Let (S_k, S_l) be *any* pair of segments in $Y^-(E)$. If $E \notin W_{kl}$, then both S_l and S_k are either above or below E , and their relative order does not affect the placement of E . This establishes that E will be correctly placed in $Y^+(E)$ (i.e., the contrapositive of the necessary condition expressed by the claim).

Suppose now that E is a right endpoint. In this case we shall establish the lemma in its direct form. The left endpoint of S has been correctly placed since it was processed earlier and E is the first one to be misplaced. If E has been misplaced,

then, by Lemma 5.4, there exists a segment $S' \in Y^-(E)$ intersecting S to the left of E such that the relative positions of S and S' in $Y^+(A)$ and $Y^-(E)$ are the same (A is the rightmost left endpoint of S and S'), i.e., their order in Y has not been reversed by the algorithm. In this case, $E = B_{kl} \in W_{kl}$ for $S_k = S$ and $S_l = S'$, thus proving that there exists a pair S_l and S_k such that S_l and S_k intersect and $E \in W_{kl}$. \square

Let S_k and S_l be two segments of $Y^-(E)$ such that $E \in W_{kl}$. Assume without loss of generality that $S_k < S_l$ in $Y(E_{kl})$. Since E is the first endpoint to be misplaced, we have $S_k \cap L(E_{kl}) <_y S_l \cap L(E_{kl})$. For convenience, we will say that two segments S_p and S_q have been *exchanged between E' and E''* for two endpoints $E' <_x E''$ if $S_p < S_q$ in $Y^+(E')$ and $S_q < S_p$ in $Y^-(E'')$.

The case where $E \in W_{kl}^-(E_{kl})$ cannot cause any difficulty since S_k and S_l cannot be active between E_{kl} and E , and therefore, S_k and S_l cannot be exchanged between E_{kl} and E , which implies that E is correctly placed with respect to S_k and S_l .

The case where $E \in W_{kl}^+(E_{kl})$ is more difficult. E is not correctly placed only if S_k and S_l are not exchanged between E_{kl} and E , i.e., $S_k < S_l$ in both $Y^+(E_{kl})$ and $Y^-(E)$. We shall prove that this is not possible and therefore conclude that S is correctly placed into Y in this case as well.

Assume, for a contradiction, that S_k and S_l have not been exchanged between E_{kl} and E . As E belongs to $W_{kl}^+(E_{kl})$, S_k and S_l cannot be adjacent in $Y^-(E)$ since otherwise they would constitute a prime pair and they would have been exchanged. Let $(S_k = S_{k_0}, S_{k_1}, \dots, S_{k_r}, S_{k_{r+1}} = S_l)$ ($r \geq 1$) be the subsequence of segments of $Y^-(E)$ occurring between S_k and S_l with the further (legitimate) assumption that (S_k, S_l) is a pair of intersecting segments such that $E \in W_{kl}^+(E_{kl})$, for which r is *minimal* (i.e., for which the above subsequence is shortest). For an arbitrary $1 \leq i \leq r$, consider segment S_{k_i} and assume, without loss of generality, that $E <_y S_{k_i}$. As a direct consequence of the fact that $(S_k, S_{k_1}, \dots, S_{k_r}, S_l)$ is shortest, we observe that E cannot belong to $W_{k_i k_i}^+(E_{k_i k_i})$ nor to $W_{k_i l}^+(E_{k_i l})$. We distinguish two cases.

(i) $E <_y S_{k_i}$ (see Figure 5.2). Clearly, $E \in W_{k_i l}$, and, by the above observation, we must have $E \in W_{k_i l}^-(E_{k_i l})$. It follows that the relative y -orders of S_{k_i} and S_l along $L(E_{k_i l})$ and $L(E)$ are the same, hence $S_l \cap L(E_{k_i l}) <_y S_{k_i} \cap L(E_{k_i l})$. As E is the first endpoint to be misplaced, the order in $Y^+(E_{k_i l})$ agrees with the geometry; i.e., we have $S_l < S_{k_i}$ in $Y^+(E_{k_i l})$. Moreover, since the pair (S_{k_i}, S_l) is not active (between $E_{k_i l}$ and E , because $E \in W_{k_i l}^-(E_{k_i l})$) and therefore cannot be exchanged, the same inequality holds in $Y^-(E)$, which contradicts the definition of S_{k_i} .

(ii) $S_{k_i} <_y E$. This case is entirely symmetric to the previous one. It suffices to exchange the roles of S_k and S_l and to reverse the relations $<$ and $<_y$.

Since a contradiction has been reached in both cases, the lemma is proved. \square

We now complete the proof of the theorem. The previous lemma implies that the endpoints are correctly processed. Indeed let E_i be an endpoint. If E_i is a right endpoint, we simply remove the corresponding segment from Y and update the set of active segments. This can be done exactly since predicates of degree ≤ 2 are evaluated correctly. If E_i is a left endpoint, it is correctly placed in Y on the basis of the previous lemma.

The lemma also implies that all pairs that intersect have been processed. Indeed if S_p and S_q are two intersecting segments such that $S_p \cap L(A_{pq}) <_y S_q \cap L(A_{pq})$ and $S_q \cap L(B_{pq}) <_y S_p \cap L(B_{pq})$, the lemma shows that $S_p < S_q$ in $Y^+(A_{pq})$ and $S_p > S_q$ in $Y^-(B_{pq})$, which implies that the pair (S_p, S_q) has been processed (Lemma 5.4).

This concludes the proof of the theorem. \square

Remark 3. Handling the degenerate cases does not cause any difficulty, and the




FIG. 5.2. For the proof of correctness of the lazy algorithm.

previous algorithm will work with only minor changes. For the initial sorting of the endpoints, we can take any order relation compatible with the order of their x -coordinates, e.g., the lexicographic order.

Remark 4. Theorem 5.3 applies directly to pseudosegments, i.e., curved segments that intersect in at most one point. Lemmas 5.4 and 5.5 also extend to the case of monotone arcs that may intersect at more than one point. To be more precise, in Lemma 5.4, we have to replace “intersect” with “intersect an odd number of times”; Lemma 5.5 and its proof are unchanged provided that we define $W_{kl}^+(E_{kl})$ (respectively, $W_{kl}^-(E_{kl})$) as the subset of W_{kl} consisting of the endpoints $E, E_{kl} \leq_x E$, such that the slab $(E_{kl}, E]$ contains an odd number (respectively, none or an even number) of intersection points. As a consequence, the lazy algorithm (which still uses only Predicates 1, 2, and 2') will detect all pairs of arcs that intersect an odd number of times.

Remark 5. For line segments, observe that checking whether a pair of segments is active does not require knowing (and therefore maintaining) E_{kl} . In fact, we can replace condition 3 in the definition of an active pair by the following condition: $S_l <_y F_{kl} <_y S_k$ and $I_{kl} <_x F_{kl}$. If $E_{kl} <_x I_{kl}$, the two definitions are identical and if $I_{kl} <_x E_{kl}$, the pair is not active since, by Lemma 5.5, condition 2 of the definition won't be satisfied.

6. Efficient implementations of the lazy algorithm in the predicate arithmetic model. The difficulty of efficiently implementing the lazy sweep-line algorithm using only predicates of degree at most 2 (i.e., in the predicate arithmetic model of degree 2) is due to verification of the emptiness condition in Definition 5.1 and of the prime-pair condition expressed by Definition 5.2: both conditions require examination of all endpoints that have not been processed yet. One can easily check that various known implementations of the sweep achieve straightforward verification

of the emptiness condition by introducing algorithmic complications. The following subsection describes an efficient implementation of the lazy algorithm in the predicate arithmetic model of degree 3. The second subsection improves on this result in a special but important instance of Pb1, namely the case of two sets of nonintersecting segments. The algorithm presented there uses only predicates of degree at most 2.

6.1. Robustness of the standard sweep-line algorithm. We shall run our lazy algorithm under the predicate arithmetic model of degree 3. We then have the capability to correctly compare the abscissae of an intersection and of an endpoint. We refine the lazy algorithm in the following way. Let E_i be the last processed endpoint and let E_{i+1} be the endpoint to be processed next. An active pair (S_k, S_l) that occurs in Y between $Y(E_i)$ and $Y^-(E_{i+1})$ will be processed if and only if its intersection point I_{kl} lies to the right of E_i and not to the right of E_{i+1} . As the slab is free of endpoints in its interior, any pair of adjacent segments encountered in Y (between $Y(E_i)$ and $Y^-(E_{i+1})$) and that intersect within the slab is active. Moreover the intersection points of all prime pairs belong to the slab. It follows that this instance of the lazy algorithm need not explicitly check whether a pair is active or not and therefore is much more efficient than the lazy algorithm of section 5. This algorithm is basically what the original algorithm of Bentley–Ottmann² becomes when predicates of degree at most 3 are evaluated. (Recall that the standard algorithm requires the capability to correctly execute predicates of degree up to 5.)

We therefore conclude with the following theorem.

THEOREM 6.1. *If Predicates 1, 2, and 3 are evaluated exactly, the standard sweep-line algorithm will solve Pb1 in $O((n+k)\log n)$ time.*

It is now appropriate to briefly comment on the implementation details of the just described modified algorithm. Data structure Y is implemented as usual as a dictionary. Data structure X , however, is even simpler than in the standard algorithm (which uses a priority queue with dictionary access). Here X has a primary component realized as a static search tree on the abscissae of the endpoints E_1, \dots, E_{2n} . Leaf E_j points to a secondary data structure $\mathcal{L}(E_j)$ realized as a conventional linked list, containing (in an arbitrary order) adjacent intersecting pairs in slab $(E_j, E_{j+1}]$. Remember that, when E_j has been processed, all intersecting pairs of $\mathcal{L}(E_j)$ are active. Insertion into $\mathcal{L}(E_j)$ is performed at one of its ends, and so is access for reporting (when the plane sweep reaches slab $(E_j, E_{j+1}]$). Since access to events occurs in data structure X , each pair in X must have pointers to the two corresponding segments in Y , in order to enable the necessary updates. In addition, to effect constant-time removal of a pair (S_h, S_k) due to loss of adjacency, all that is needed is to make bidirectional the mentioned pointers. Notice that the elements of X correspond to pairs of adjacent segments in Y , so that at most two records in X are pointed to by any member of Y . We finally observe that a segment adjacency arising in Y during the execution of the algorithm must be tested for intersection; however, an intersecting pair of adjacent segments is eligible for insertion into X only as long as the plane sweep has not gone beyond the slab containing the intersection in question. As regards the running time, beside the initial sorting of the endpoints and the creation of the corresponding primary tree in time $O(n \log n)$, it is easily seen that each intersection uses $O(\log n)$ time (amortized), thereby achieving the performance of the standard algorithm.

²With the policy concerning multiple detections of intersections that is stipulated at the beginning of section 5.




FIG. 6.1. Assume that the computed x -coordinate of the intersection I_{12} of S_1 and S_2 is (erroneously) found to be smaller than the x -coordinate of the left endpoint of S_3 . Y is implemented as a balanced binary search tree. The key associated to a node is the $\lceil \frac{n}{2} \rceil$ th element of the corresponding subtree. If all predicates of degree at most 2 are evaluated exactly, S_3 is correctly inserted below S_2 , and S_4 is correctly inserted above S_1 and S_2 . The different states of Y are (S_1) , (S_1, S_2) , (S_2, S_1) , (S_3, S_2, S_1) , and (S_3, S_2, S_1, S_4) . Since segments S_3 and S_4 are never adjacent, their intersection I_{34} will not be detected. Observe that the missed intersection point can be arbitrarily far from the intersection point involved in the wrong decision.

Finally, we note that if only predicates of degree ≤ 2 are evaluated correctly, the algorithm of Bentley–Ottmann may fail to report the set of intersecting pairs of segments. See Figure 6.1 for an example.

Remark 6. The fact that the sweep-line algorithm does not need to sort intersection points had already been observed by several authors including Myers [Mye85], Schorn [Sch91], and Hobby [Hob93]. Myers does not use it for solving robustness problems but for developing an algorithm with an expected running time of $O(n \log n + k)$. Schorn uses this fact to decrease the precision required by the sweep-line algorithm from fivefold to threefold; i.e., Schorn’s algorithm uses exact arithmetic of degree 3. Using Theorem 5.3, we will show in section 7 that double precision suffices.

6.2. Reporting intersections between two sets of nonintersecting line segments. In this subsection, we consider two sets of line segments in the plane, \mathcal{S}_b (the blue set) and \mathcal{S}_r (the red set), where no two segments in \mathcal{S}_b (similarly, in \mathcal{S}_r) intersect. Such a problem arises in many applications, including the union of two polygons and the merge of two planar maps. We denote by n_b and n_r the cardinalities of \mathcal{S}_b and \mathcal{S}_r , respectively, and let $n = n_b + n_r$.

Mairson and Stolfi [MS88] have proposed an algorithm that works for arcs of curve as well as for line segments. Its time complexity is $O(n \log n + k)$, which is optimal, and requires $O(n + k)$ space ($O(n)$ in case of line segments). The same asymptotic time-bound has been obtained by Chazelle et al. [CEGS94] and by Chazelle and Edelsbrunner [CE92]. The latter algorithm is not restricted to two sets of nonintersecting line segments. Other algorithms have been proposed by Nievergelt and Preparata [NP82] and by Guibas and Seidel [GS87] in the case where the segments of \mathcal{S}_b (and \mathcal{S}_r) are




FIG. 6.2. Notations for the case of two sets of nonintersecting line segments.

the edges of a subdivision with convex faces. With the exception of the algorithm of Chazelle et al. [CEGS94], all these algorithms construct the resulting arrangement and therefore have degree 4. The algorithm of Chazelle et al. requires comparing the ordinates of the intersections of two segments with a vertical line passing through an endpoint. Therefore it is of degree 3.

We propose here an algorithm that computes all the intersections but not the arrangement. This algorithm uses only predicates of degree ≤ 2 and has time complexity $O((n+k)\log n)$.

Segments are assumed to be nonvertical since intersections with vertical segments can be easily handled with predicates of degree ≤ 2 . We say that a point E_i is *vertically visible* from a segment $S_b \in \mathcal{S}_b$ if the vertical line segment joining E_i with S_b does not intersect any other segment in \mathcal{S}_b (the same notion is applicable to \mathcal{S}_r). For two intersecting segments $S_b \in \mathcal{S}_b$ and $S_r \in \mathcal{S}_r$, let L be a vertical line to the right of A_{br} such that no other segment intersects L between S_b and S_r (i.e., S_b and S_r are adjacent). We let T_{br} denote the wedge defined by S_b and S_r in the slab between L and $L(I_{br})$ (see Figure 6.2). For a point set \mathcal{F} , we let $CH^+(\mathcal{F})$ and $CH^-(\mathcal{F})$ denote its upper and lower convex hulls, respectively.

Our algorithm is based on the following observation.

LEMMA 6.2. *T_{br} contains blue endpoints if and only if it contains a blue endpoint that is vertically visible from S_b . Similarly, T_{br} contains red endpoints if and only if it contains a red endpoint vertically visible from S_r .*

Proof. The sufficient condition is trivial, so we prove only necessity. Assume without loss of generality that $S_r \cap L(A_{br}) <_y S_b \cap L(A_{br})$. Let \mathcal{E} be the subset of the blue endpoints that belong to T_{br} . Clearly, all vertices of $CH^+(\mathcal{E})$ are vertically visible from S_b . \square

Our algorithm has two phases. The second one is the lazy algorithm of section

5. The first one can be considered as a preprocessing step that will help to efficiently find active pairs of segments.

More specifically, our objective is to develop a quick test of the emptiness condition based on the previous lemma. The preprocessing phase is aimed at identifying the candidate endpoints for their potential belonging to wedges formed by intersecting adjacent pairs. Referring to \mathcal{S}_b (and analogously for \mathcal{S}_r), we first sweep the segments of \mathcal{S}_b and construct for each blue segment S_b the lists \mathcal{L}_b^- and \mathcal{L}_b^+ of blue endpoints that are vertically visible from S_b and lie respectively below and above S_b . The sweep takes time $O(n \log n)$ and the constructed lists are sorted by increasing abscissa. Since there is no intersection point, only predicates of degree ≤ 2 are used. The total size of the lists $\mathcal{L}_b^-, \mathcal{L}_b^+, \mathcal{L}_r^-,$ and \mathcal{L}_r^+ is $O(n)$.

As mentioned above, the crucial point is to decide whether or not the wedge T_{br} of a pair of intersecting segments S_b and S_r , adjacent in Y , contains endpoints of other segments. Again, we assume that $S_r < S_b$ in Y . If such endpoints exist, then T_{br} contains either a blue vertex of $CH^+(\mathcal{L}_b^- \cap L^+)$ or a red vertex of $CH^-(\mathcal{L}_b^+ \cap L^+)$. (L^+ is the half-plane to the right of line L .)

We will show below that, using predicates of degree ≤ 2 , the lists can be preprocessed in time $O(n \log n)$ and that deciding whether or not T_{br} contains endpoints can be done in time $O(\log n)$, using only predicates of degree at most 2.

Assuming for the moment that this primitive is available, we can execute the plane sweep algorithm described earlier. Specifically, we sweep \mathcal{S}_b and \mathcal{S}_r simultaneously, using the lazy sweep-line algorithm of section 5.³ Each time we detect a pair of (adjacent) intersecting segments S_b and S_r , we can decide in time $O(\log n)$ whether they are “active” or “not active,” using only predicates of degree ≤ 2 .

We sum up the results of this section in the following theorem.

THEOREM 6.3. *Given n line segments in the plane belonging to two sets \mathcal{S}_b and \mathcal{S}_r , where no two segments in \mathcal{S}_b (analogously, in \mathcal{S}_r) intersect, there exists an algorithm of optimal degree 2 that reports all intersecting pairs in $O((n + k) \log n)$ time using $O(n)$ storage.*

We now return to the implementation of the primitive described above. Suppose that, for some segment S_i ($i = b$ or r) we have constructed the upper hull $CH^+(\mathcal{L}_i^- \cap L^+)$. Then we can detect in $O(\log n)$ time if an element of a list, say \mathcal{L}_b^- , lies above some segment S_r . More specifically, we first identify among the edges of $CH^+(\mathcal{L}_b^- \cap L^+)$ the two consecutive edges whose slopes are, respectively, smaller and greater than the slope of S_r . This requires only the evaluation of $O(\log |\mathcal{L}_b^-|)$ predicates of degree 2. It then remains to decide whether the common endpoint E of the two reported edges lies above or below the line containing S_r . This can be answered by evaluating the orientation predicate $\text{orient}(E, A_r, B_r)$.

The crucial requirement of the adopted data structure is the ability to efficiently maintain $CH^+(\mathcal{L}_i^- \cap L^+)$. To this purpose, we propose the following solution.

The data structure associated with a list \mathcal{L}_i^- ($i = b$ or r) represents the upper convex hull $CH^+(\mathcal{L}_i^-)$ of \mathcal{L}_i^- . (Similarly, the data structure associated with a list \mathcal{L}_i^+ represents the lower convex hull $CH^-(\mathcal{L}_i^+)$ of \mathcal{L}_i^+ .) This implies that a binary search on the convex hull slopes uniquely identifies the test vertex. Since the elements of each list are already sorted by increasing x -coordinates, the data structures can be constructed in time proportional to their sizes and therefore in $O(n)$ time in total. It can be easily checked that only orientation predicates (of degree 2) are involved in this process. To guarantee the availability of $CH^+(\mathcal{L}_i^- \cap L^+)$, we have to ensure

³We can adopt the policy of processing all active pairs before the next endpoint.

that our data structure can efficiently handle the deletion of elements. As elements are deleted in order of increasing abscissa, this can be done in amortized $O(\log |S|)$ time per deletion [HS90, HS96]. It follows that preprocessing all lists takes $O(n)$ time, uses $O(n)$ space, and requires only the evaluation of predicates of degree ≤ 2 .

7. An efficient implementation of the lazy algorithm under the exact arithmetic model of degree 2. We shall run the lazy algorithm of section 6.1 under the exact arithmetic model of degree 2, i.e., Predicates 1 and 2 are evaluated exactly but Predicate 3 is implemented with exact arithmetic of degree 2 as explained in section 4.3. Several intersection points may now be found to have the same abscissa as an endpoint. We refine the lazy algorithm in the following way. Let E_i be the last processed endpoint and let E_{i+1} be the endpoint with an abscissa strictly greater than the abscissa of E_i to be processed next. An active pair (S_k, S_l) will be processed if and only if its intersection point is found to lie to the right of E_i and not to the right of E_{i+1} .

We claim that this policy leads to efficient verification of the emptiness condition. Indeed, the intersections of all prime pairs belong to $(E_i, E_{i+1}]$, because $E_{i+1} \in W_{kl}^+(E_{kl}) \implies I_{kl} \leq_x E_{i+1}$, and, by the monotonicity property, both implementations of Predicate 3 of section 4.3 will report that $I_{kl} \leq_x E_{i+1}$.

The crucial observation that drastically reduces the time complexity is the following. A pair of adjacent segments (S_k, S_l) encountered in Y between $Y(E_i)$ and $Y(E_{i+1}^-)$ whose intersection point is reported to lie in slab $(E_i, E_{i+1}]$ is active if and only if $E_{i+1} \notin W_{kl}^-(E_{kl})$. Indeed, since I_{kl} is found to be $<_x E_{i+2}$, the monotonicity property implies that $I_{kl} <_x E_{i+2}$. Therefore, when checking if a pair is active, it is sufficient to consider just the next endpoint, not all of them.

Theorem 5.3 therefore applies. If no two endpoints have the same x -coordinate, the algorithm is the same as the algorithm in section 6.1 apart from the implementation of Predicate 3. Otherwise, we construct X on the distinct abscissae of the endpoints and store all endpoints with identical x -coordinates in a secondary search structure with endpoints sorted by y -coordinates. This secondary structure will allow us to determine if a pair is active in logarithmic time by binary search. We conclude with the following theorem.

THEOREM 7.1. *Under the exact arithmetic model of degree 2, the instance of the lazy algorithm described above solves Pb1 in $O((n+k)\log n)$ time.*

8. Conclusion. Further pursuing our investigations in the context of the exact-computation paradigm, in this paper we have illustrated that important problems on segment sets (such as intersection report, arrangement, and trapezoidal map), which are viewed as equivalent under the Real-RAM model of computation, are distinct if their algebraic degree is taken into account. This sheds new light on robustness issues which are intimately connected with the notion of algebraic degree and illustrates the richness of this new direction of research.

For example, we have shown that the well-known plane sweep algorithm of Bentley–Ottmann uses more machinery than strictly necessary and can be appropriately modified to report segment intersections with arithmetic capabilities very close to optimal and no sacrifice in performance.

Another result of our work is that exact solutions of some problems can be obtained even if approximate (or even random) evaluations of some predicates are performed. More specifically, using less powerful arithmetic than demanded by the application, we have been able to compute the vertices of an arrangement of line segments

by constructing an arrangement which may be different from the actual one (and may not even correspond to any set of straight line segments) but still has the same vertex set.

Our work shows that the sweep-line algorithm is more robust than usually believed, proposes practical improvements leading to robust implementations, and provides a better understanding of the sweeping line paradigm. The key to our technique is to relax the horizontal ordering of the sweep. This is one step further after similar attempts, aimed though at different purposes [Mye85, MS88, EG89].

A host of interesting open questions remain. One such question is to devise an output-sensitive algorithm for reporting segment intersections with optimal time complexity and with optimal algebraic degree (that is, 2). It would also be interesting to examine the plane sweep paradigm in general. For example, with regard to the construction of Voronoi diagrams in the plane, one should elucidate the reasons for the apparent gap between the algebraic degrees of Fortune’s plane sweep solution and of the (optimal) divide-and-conquer and incremental algorithms.

Appendix. In connection with reducibility of polynomials over a domain of rationality, we choose the rationals as the latter. “Reducible” means “reducible over the rationals.”

We first recall that a general determinant is irreducible if its entries are regarded as independent variables [Boc07]. This suffices to prove directly that the degree of Predicate 2 is 2 and will be crucial for completing the proof of irreducibility of the polynomials pertaining to the other predicates. Use will be made of the following theorem.

Let $p(x_1, \dots, x_n)$ be a multivariate homogeneous polynomial, and let I be a subset of $\{1, \dots, n\}$, such that for any $j \in I$, x_j is a variable of degree 1 in p . For $i \in \{1, \dots, n\}$ p can be expressed as

$$p = p_i x_i + p_{i0},$$

where p_i and p_{i0} are polynomials in all variables except x_i . Two polynomials are not considered distinct if they differ just by a multiplicative constant. We have the following.

THEOREM 8.1. *Let $p(x_1, \dots, x_n)$ be a multivariate homogeneous polynomial with degree at least 3 and $|I| \geq 2$. If for some $i, j \in I$, coefficients p_i and p_j are distinct and irreducible, then p is irreducible.*

Proof. Assume, for a contradiction, that p is reducible. This means that p can be expressed as $p = \phi\psi$, where ϕ and ψ are multivariate homogeneous polynomials over the same variables satisfying $1 \leq \text{degree}(\phi), \text{degree}(\psi) < \text{degree}(p)$. Each degree-1 variable obviously appears in exactly one of the factors. With complete generality, assume that x_i appears in ϕ . We distinguish two cases.

- (1) $\text{degree}(\phi) > 1$. Expanding ϕ around x_i we obtain $\phi = \phi_i x_i + \phi_{i0}$, so that

$$p = \phi\psi = (\phi_i x_i + \phi_{i0})\psi = \phi_i \psi x_i + \phi_{i0} \psi.$$

This means that $p_i = \phi_i \psi$, with $\text{degree}(\phi_i), \text{degree}(\psi) \geq 1$, i.e., p_i is reducible, a contradiction.

- (2) $\text{degree}(\phi) = 1$. Variable x_j appears either in ϕ or in ψ . In the first case, p_i and p_j are both proportional to ψ , i.e., they may differ only by a multiplicative constant and are not distinct, a contradiction.

In the second case, by an argument analogous to that of Case 1, we reach the conclusion that $p_j = \phi\psi_j$ and that $\text{degree}(\psi_j) = \text{degree}(p) - \text{degree}(\phi) - 1 \geq$

$3 - 1 - 1 = 1$. Since $p_j = \phi\psi_j$ and $\text{degree}(\phi), \text{degree}(\psi_j) \geq 1$, p_j is reducible, another contradiction.

This completes the proof of the theorem. \square

The calculation of coefficients p_i, p_j corresponds to taking formal derivatives of p with respect to x_i, x_j . Notice that a derivation reduces the degree by 1. The reader may verify, by explicit and not very enlightening calculations, that the following schedules of derivations lead to irreducible 2×2 determinants. The indices conform with the notation of section 4.1.

Predicate 3: Apply Theorem 8.1 to the pair (x_1, x_2) .

Predicate 4: Apply Theorem 8.1 first to (x_2, x_3) and then to (y_6, y_7) .

Predicate 4': Apply Theorem 8.1 first to (x_0, x_1) and then to (x_2, x_3)

Predicate 5: Apply Theorem 8.1 first to (x_1, x_2) , next to (x_7, x_8) , and finally to (y_5, y_6) .

Acknowledgments. We are indebted to H. Brönnimann for having pointed out an error in a previous version of this paper and to O. Devillers for discussions that lead to Lemma 4.3. S. Pion, M. Teillaud, and M. Yvinec are also gratefully acknowledged for their comments on this work. Special thanks are due to two anonymous referees, whose insightful comments have significantly contributed to the quality of the paper.

REFERENCES

- [ABD⁺97] F. AVNAIM, J.-D. BOISSONNAT, O. DEVILLERS, F. PREPARATA, AND M. YVINEC, *Evaluating signs of determinants using single-precision arithmetic*, *Algorithmica*, 17 (1997), pp. 111–132.
- [Bal95] I. J. BALABAN, *An optimal algorithm for finding segment intersections*, in Proceedings of the 11th Annual ACM Symposium on Computational Geometry, Vancouver, Canada, 1995, pp. 211–219.
- [BO79] J. L. BENTLEY AND T. A. OTTMANN, *Algorithms for reporting and counting geometric intersections*, *IEEE Trans. Comput.*, C-28 (1979), pp. 643–647.
- [Boc07] M. BOCHER, *Introduction to Higher Algebra*, Macmillan, New York, 1907.
- [BDS⁺92] J.-D. BOISSONNAT, O. DEVILLERS, R. SCHOTT, M. TEILLAUD, AND M. YVINEC, *Applications of random sampling to on-line algorithms in computational geometry*, *Discrete Comput. Geom.*, 8 (1992), pp. 51–71.
- [BEPP97] H. BRÖNNIMANN, I. EMIRIS, V. PAN, AND S. PION, *Computing exact geometric predicates using modular arithmetic with single precision*, in Proceedings of the 13th Annual ACM Symposium on Computational Geometry, Nice, France, 1997, pp. 174–182.
- [BY97] H. BRÖNNIMANN AND M. YVINEC, *Efficient exact evaluation of signs of determinants*, in Proceedings of the 13th Annual ACM Symposium on Computational Geometry, Nice, France, 1997, pp. 166–173.
- [Bro81] K. Q. BROWN, *Comments on "Algorithms for reporting and counting geometric intersections"*, *IEEE Trans. Comput.*, C-30 (1981), pp. 147–148.
- [Bur96] C. BURNIKEL, *Exact Computation of Voronoi Diagrams and Line Segment Intersections*, Ph.D. thesis, Universität des Saarlandes, Saarbrücken, Germany, 1996.
- [BKM⁺95] C. BURNIKEL, J. KÖNNEMANN, K. MEHLHORN, S. NÄHER, S. SCHIRRA, AND C. UHRIG, *Exact geometric computation in LEDA*, in Proceedings of the 11th Annual ACM Symposium on Computational Geometry, Vancouver, Canada, 1995, pp. C18–C19.
- [BMS94] C. BURNIKEL, K. MEHLHORN, AND S. SCHIRRA, *On degeneracy in geometric computations*, in Proceedings of the 5th ACM-SIAM Symposium on Discrete Algorithms, Arlington, VA, 1994, pp. 16–23.
- [Cha91] B. CHAZELLE, *Triangulating a simple polygon in linear time*, *Discrete Comput. Geom.*, 6 (1991), pp. 485–524.
- [CE92] B. CHAZELLE AND H. EDELSBRUNNER, *An optimal algorithm for intersecting line segments in the plane*, *J. ACM*, 39 (1992), pp. 1–54.
- [CEGS94] B. CHAZELLE, H. EDELSBRUNNER, L. J. GUIBAS, AND M. SHARIR, *Algorithms for bichromatic line segment problems and polyhedral terrains*, *Algorithmica*, 11 (1994), pp. 116–132.

- [Cla92] K. L. CLARKSON, *Safe and effective determinant evaluation*, in Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science, IEEE Computer Society, Los Alamitos, CA, 1992, pp. 387–395.
- [CS89] K. L. CLARKSON AND P. W. SHOR, *Applications of random sampling in computational geometry*, II. Discrete Comput. Geom., 4 (1989), pp. 387–421.
- [EG89] H. EDELSBRUNNER AND L. J. GUIBAS, *Topologically sweeping an arrangement*, J. Comput. System Sci., 38 (1989), pp. 165–194. Corrigendum in 42 (1991), pp. 249–251.
- [For85] A. R. FORREST, *Computational geometry in practice*, in Fundamental Algorithms for Computer Graphics, NATO Adv. Sci. Inst. Ser. F Comput. Systems Sci. 17, R. A. Earnshaw, ed., Springer-Verlag, Berlin, 1985, pp. 707–724.
- [For87] A. R. FORREST, *Computational geometry and software engineering: Towards a geometric computing environment*, in Techniques for Computer Graphics, D. F. Rogers and R. A. Earnshaw, eds., Springer-Verlag, Berlin, 1987, pp. 23–37.
- [For93] S. FORTUNE, *Progress in computational geometry*, in Directions in Computational Geometry, R. Martin, ed., Information Geometers, Winchester, UK, 1993, pp. 81–128.
- [FV93] S. FORTUNE AND C. J. VAN WYK, *Efficient exact arithmetic for computational geometry*, in Proceedings of the 9th Annual ACM Symposium on Computational Geometry, San Diego, CA, 1993, pp. 163–172.
- [Gol91] D. GOLDBERG, *What every computer scientist should know about floating-point arithmetic*, ACM Comput. Surv., 23 (1991), pp. 5–48.
- [GS87] L. J. GUIBAS AND R. SEIDEL, *Computing convolutions by reciprocal search*, Discrete Comput. Geom., 2 (1987), pp. 175–193.
- [HS90] J. HERSHBERGER AND S. SURI, *Applications of a semi-dynamic convex hull algorithm*, In Proceedings of the 2nd Scandinavian Workshop on Algorithm Theory (Bergen, 1990), Lecture Notes Comput. Sci. 447, Springer-Verlag, Berlin, 1990, pp. 380–392.
- [HS96] J. HERSHBERGER AND S. SURI, *Off-line maintenance of planar configurations*, J. Algorithms, 21 (1996), pp. 453–475.
- [Hob93] J. HOBBY, *Practical Segment Intersection with Finite Precision Output*, Technical Report 93/2-27, Bell Laboratories, Murray Hill, NJ, 1993.
- [Hof89] C. M. HOFFMANN, *The problems of accuracy and robustness in geometric computation*, IEEE Computer, 22 (1989), pp. 31–41.
- [HHK89] C. M. HOFFMANN, J. E. HOPCROFT, AND M. T. KARASICK, *Robust set operations on polyhedral solids*, IEEE Comput. Graph. Appl., 9 (1989), pp. 50–59.
- [LPT96] G. LIOTTA, F. P. PREPARATA, AND R. TAMASSIA, *Robust proximity queries: An illustration of degree-driven algorithm design*, SIAM J. Comput., 28 (1999), pp. 864–889.
- [MS88] H. G. MAIRSON AND J. STOLFI, *Reporting and counting intersections between two sets of line segments*, in Theoretical Foundations of Computer Graphics and CAD, NATO Adv. Sci. Inst. Ser. F. Comput. Systems Sci. 40, R. A. Earnshaw, ed., Springer-Verlag, Berlin, 1988, pp. 307–325.
- [Mil88] V. J. MILENKOVIC, *Verifiable implementations of geometric algorithms using finite precision arithmetic*, Artificial Intelligence, 37 (1988), pp. 377–401.
- [Mil89] V. MILENKOVIC, *Double precision geometry: A general technique for calculating line and segment intersections using rounded arithmetic*, in Proceedings of the 30th Annual IEEE Symposium on Foundations Computer Science, IEEE Computer Society, Los Alamitos, CA, 1989, pp. 500–505.
- [Mye85] E. W. MYERS, *An $O(E \log E + I)$ expected time algorithm for the planar segment intersection problem*, SIAM J. Comput., 14 (1985), pp. 625–637.
- [NP82] J. NIEVERGELT AND F. P. PREPARATA, *Plane-sweep algorithms for intersecting geometric figures*, Commun. ACM, 25 (1982), pp. 739–747.
- [Pri92] D. PRIEST, *On Properties of Floating Point Arithmetics: Numerical Stability and the Cost of Accurate Computations*, Ph.D. thesis, Dept. of Mathematics, Univ. of California at Berkeley, 1992.
- [Sch91] P. SCHORN, *Robust Algorithms in a Program Library for Geometric Computation*, volume 32 of Informatik-Dissertationen ETH Zürich, Verlag der Fachvereine, Zürich, Switzerland, 1991.
- [She96] J. R. SHEWCHUK, *Robust adaptive floating-point geometric predicates*, in Proceedings of the 12th Annual ACM Symposium on Computational Geometry, Philadelphia, PA, 1996, pp. 141–150.
- [Yap97] C. YAP, *Towards exact geometric computation*, Comput. Geom., 7 (1997), pp. 3–23.