An elementary algorithm for reporting intersections of red/blue curve segments ☆

Jean-Daniel Boissonnat [a,*], Antoine Vigneron [b]

[a] *INRIA Sophia Antipolis, 2004 Route des Lucioles, BP 93, 06904 Sophia-Antipolis, France*
[b] *Hong Kong University of Science and Technology, China*

Abstract

Let E_r and E_b be two sets of x-monotone and non-intersecting curve segments, $E = E_r \cup E_b$ and $|E| = n$. We give a new sweep-line algorithm that reports the k intersecting pairs of segments of E. Our algorithm uses only three simple predicates that allow to decide if two segments intersect, if a point is left or right to another point, and if a point is above, below or on a segment. These three predicates seem to be the simplest predicates that lead to subquadratic algorithms. Our algorithm is almost optimal in this restricted model of computation. Its time complexity is $O(n \log n + k \log \log n)$ and it requires $O(n)$ space. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Computational geometry; Segment intersection; Plane sweep; Robust algorithms

1. Introduction

The usual model to analyze geometric algorithms is the Real RAM which is assumed to compute exactly with real numbers [15]. This model hides the fact that the arithmetic of real computers has a limited precision and ignores numerical and robustness issues. As a consequence a direct implementation of an algorithm that is correct under the Real RAM model does not necessarily translate into a robust and/or efficient program, and catastrophic behaviors are commonly observed.

A first approach to remedy this problem is to use exact arithmetic. In the context of geometric algorithms, much progress has been done in the recent past [14,16,18,19]. Another approach, to be followed here, has emerged recently. Decisions in geometric algorithms depend on geometric predicates which are usually algebraic expressions. For example, for a triple of points given by their Cartesian coordinates,

---

* Corresponding author.
 E-mail address: jean-daniel.boissonnat@sophia.inria.fr (J.-D. Boissonnat).

deciding what is the orientation of the triangle reduces to evaluating (the sign of) a multivariate polyno-mial of degree two. If an algorithm only uses predicates of degree 2 as a function of the input data, and if the input data are coded as simple fixed precision numbers, computations can be done exactly using the native double precision hardware of the computer. The degree of an algorithm is therefore related to the precision required to run an algorithm using exact arithmetic. This motivates the design of efficient algorithms of low degree. Reducing the degree of the algorithms will reduce the number and the com-plexity of the degenerate configurations, make the algorithms more elementary and more general, reduce the amount of numerical computations, which is usually quite a large fraction of the total execution time especially if multi-precision computing is invoked, and possibly also refrain the use of complicated data structures resulting in low space requirements. However, reducing the degree of an algorithm may in-crease its time complexity in the Real RAM model. Following Liotta et al. [13], we consider the degree of the predicates as an additional measure of the complexity of problems and algorithms, and intend to elucidate the relationship between time-complexity and degree of the predicates. Related research can be found in Knuth's seminal work [12] and in some recent papers [3,4,10].

In this paper, we consider the problem of reporting the k intersecting pairs among a set of n x-monotone curve segments. We address the red/blue case, where this set is partitioned into two subsets of non-intersecting segments. This problem can be solved in optimal $O(n \log n + k)$ time [1,6,9]. However, these algorithms use predicates of high degree, e.g., to compare the abscissae of two intersection points or to locate an intersection point with respect to a vertical slab. These predicates have a degree and an algebraic complexity that are usually higher than the intersection predicate: this is, in particular, the case for line segments and circle segments [4]. Our algorithm only uses the intersection predicate and two other simple predicates: the predicate that sorts two points by abscissae, and the predicate that says if a point is below, on, or above a segment. In particular, we do not compute the arrangement nor the trapezoidal map of the segments. Moreover, the predicates we use do not say anything about the number or the positions of the intersection points, and the time complexity of these algorithms depends only on the number of intersecting pairs of segments, not the number of intersection points (differently from the other non trivial algorithms [1,2,6,8,9]).

The time complexity of our algorithm is $O(n \log n + k \log \log n)$, which is close to optimal, and it uses optimal space $O(n)$. This result generalizes a similar result for the case of pseudo-segments, i.e., segments that intersect in at most one point [4].

Recently, T. Chan has independently obtained a different algorithm that uses a segment-tree [7]. Its time-complexity is $O(n \log n + k \log_{2+k/n} n)$ and it uses $O(n + k)$ space. We note that this time bound is better than ours for $k = \Omega(n^{1+\varepsilon})$ but worse for $k = O(n \text{ polylog } n)$.

2. The problem

A curve segment is *x-monotone* if it is the graph of a partially defined univariate continuous function (i.e., any vertical line intersects such a segment in at most one point). Let $E = E_r \cup E_b$ be a set of n x-monotone curve segments such that no two segments of E_r (respectively E_b) intersect. The problem is to report the k pairs of segments of E that intersect.

Let s and s' be two segments of E and let p and q be two endpoints of some segments of E, not necessarily of s or s'. $x(p)$ and $y(p)$ denote the coordinates of p, and $s(x)$ denotes the point of s whose abscissae is x (if such a point exists). We consider the following predicates:

Table 1

Predicate	Degree			
	line segment	half circle	circle segment 3 points	pol. of degree d
1	2	2	12	d
2	1	1	1	1
3	2	2	4	d
4, 5	3	4	12	d

Predicate 1. $s \cap s' \neq \emptyset$.
Predicate 2. $x(p) \leqslant x(q)$.
Predicate 3. $y(p) \leqslant y(s(x(p)))$.
Predicate 4. $y(s(x(p))) \leqslant y(s'(x(p)))$.
Predicate 5. $\exists x \in [x(p), x(q)]$ such that $s(x) = s'(x)$.

Predicate 1 is mandatory. Predicate 2 allows to sort the endpoints. Predicate 3 tells whether an endpoint lies above or below a segment. Predicate 4 provides the order of two segments along a vertical line passing through an endpoint. Predicate 5 checks whether two segments intersect within a vertical slab defined by two endpoints.

We do not specify a precise representation for the segments, which may depend on the application. For example, the function associated with a segment may be given explicitly together with the interval where it is defined, or the function may be given implicitly as the algebraic function of degree d that interpolates $d + 1$ given points (including the endpoints of the segment). Other representations are also possible. The degrees of the predicates clearly depend on the chosen representation. In Table 1 (see also [4,11]), the degrees of Predicates 2–5 are given for line segments represented by the coordinates of their endpoints, for half-circles defined by centers, radii plus a Boolean (to distinguish between the upper and the lower arc), for circle segments defined by three points (including the two endpoints of the segment), and for curves defined by a polynomial equation $y = f(x)$. Observe that the predicates are ordered by increasing degrees in all those cases. However, this order may be different for some other types of segments. For instance, for half-circles defined by three non x-extreme points, the degree of Predicate 2 is 20 while the degree of Predicate 3 is only 4.

Our algorithm only requires Predicates 1, 2 and 3 while other non-trivial algorithms [1,2,6,8,9] make use of our five predicates. Thus they have higher algebraic degree when we consider line segments, half circles or circles defined by three points.

3. Algorithm

In this section, e_b (respectively e_r) denotes a blue (respectively red) segment, i.e., $e_b \in E_b$ and $e_r \in E_r$. We will prove the following result:

Theorem 1. *The red–blue curve segment intersection problem can be solved in* $O(n \log n + k \log \log n)$ *time and* $O(n)$ *space using predicates 1, 2 and 3.*

Fig. 1. $U(e_r)$ and a good pair.

We present a new plane sweep algorithm for this problem. A vertical sweep line moves across the plane from left to right. When it reaches an endpoint—a sweep *event*, some data structures are updated and intersections can be reported. At any time, we only consider *active* segments, which are the segments that cross the sweep line.

Definition 1. A good pair (see Fig. 1) consists of a blue and a red intersecting segments, denoted, respectively, e_b and e_r, such that the left endpoint of e_b is above e_r and on the right of the left endpoint of e_r. A bad pair is an intersecting pair that is not a good pair.

In the following subsections, we describe an algorithm that reports the *good* pairs. The other intersections are reported by the same algorithm after exchanging the orientation of the *y*-axis or the colors of the segments. The algorithm to be described is therefore applied four times.

3.1. Data structures

The idea is to report the good pairs while pushing the blue segments downwards. Each blue segment is stored in a set $U(e_r)$ for some active red segment e_r. The set of segments in $U(e_r)$ will be maintained during the sweep. It consists of the blue segments that cannot be pushed down because they are blocked either by e_r or by the lowest segment of $U(e_r)$.

The above/below relationship within E_b (respectively E_r) is a partial order, which is a total order when restricted to active segments. This relation can be extended to all the pairs (e, e') of non-intersecting segments, in which case *e is below e'* is denoted by $e < e'$, and $e \leqslant e'$ denotes $e < e'$ or $e = e'$. We will maintain the ordered list of the active red segments in a balanced search tree just like in Chan's algorithm [6] for red–blue segments. In order to deal with boundary cases, we will introduce an additional red segment e_∞ that is minimum with respect to our order and will always be considered active.

For each blue segment e_b, we associate a red segment $l(e_b)$. We will require that the three following properties remain satisfied during the course of the algorithm:

Property 1. *As the plane is swept, $l(e_b)$ can only decrease with respect to the vertical order of red segments.*

Property 2. *The good pairs consisting of a blue segment e_b and the red segments lying above $l(e_b)$ have been reported previously* (*see Fig. 3*).

Property 3. *A blue segment e_b can only intersect $e_r \leqslant l(e_b)$ to the right of the sweep line* (*see Fig. 1*).

Property 1 will be used in the analysis of our algorithm to show that an intersecting pair is reported at most once. It formalizes the idea that we "push" the blue segments downwards as the plane is swept. The next two properties will help to prove the correctness of the algorithm since they imply that, when the sweep line reaches the right endpoint of a segment, all the good pairs involving that segment have been reported.

For each active red segment e_r, we denote by $U(e_r)$ the set of all active blue segments e_b such that $l(e_b) = e_r$. $U(e_r)$ is stored in a mergeable heap data structure, that is a data structure that allows union, minimum extraction and deletion (e.g., a binomial heap). The underlying order is the vertical order of active red segments along the sweep line. Note that we shall not maintain $l(e_b)$ explicitly, as it is not clear whether we can do it within the same time bounds. Its value can be retrieved from the heaps $U(e_r)$. For any red segment e_r, the set $U(e_r)$ satisfies the following property:

Property 4. *If $U(e_r) \neq \emptyset$, then its minimum does not intersect e_r and is above e_r* (*see Fig. 1*).

However, the sets $U(e_r)$ do not have any simple monotonicity property. For example, if e_r is below e_r', then $U(e_r)$ may be entirely above $U(e_r')$ or they may be interlaced (see Fig. 2).

Properties 1–4 will be the invariants of our algorithm. First we note that they are maintained between two events of the sweep. Indeed, no change in our data structure is performed between two events of the sweep, so Properties 1, 2 and 4 are not affected. For all $e_b \in U(e_r)$, Property 4 ensures that the intersection of e_b with the sweep line remains above e_r until the next event is reached, so Property 3 is preserved as well.

3.2. Handling the events

We distinguish four kinds of events depending whether the sweep line reaches a left or a right endpoint of a segment, and whether the segment is red or blue. Each event corresponds to the insertion, the deletion of a blue or a red segment. We will now explain in detail how these events are handled while maintaining our invariants.

3.2.1. Inserting a red segment
When the sweep line reaches a left endpoint of a red segment e_r, e_r is just inserted in the list of active red segments and the heap $U(e_r)$ is initialized as an empty set. Our invariants are obviously maintained. Note that Property 2 holds because it only deals with *good* pairs. This is where this notion is crucial.


Fig. 2. $U(e_r)$ is above $U(e_r')$.

Fig. 3. Insertion of e_b.

3.2.2. Inserting a blue segment

When the sweep line reaches the left endpoint of a blue segment e_b, we first determine the red segment e_r that lies just below the endpoint. If e_b does not intersect e_r, then it is inserted in $U(e_r)$. Otherwise, we report this intersecting pair and repeat the same process with the active red segment that is below e_r until we reach a red segment that does not intersect e_b, which will eventually happen since $e_\infty < e_b$.

This procedure can be written in the following way, where $e_r - 1$ denotes the active red segment that lies immediately below e_r (see Fig. 3).

Insert(e_b, e_r)
while e_b intersects e_r
    report the intersection
    $e_r \leftarrow e_r - 1$
Insert e_b in $U(e_r)$

The intersections between e_b and the red segments that are below the left endpoint of e_b and above $l(e_b)$ have been reported, therefore Property 2 holds for e_b. We note that here e_b is above $l(e_b)$ which shows that Properties 3 and 4 are maintained.

3.2.3. Removing a red segment

Now suppose the sweep line reaches the right endpoint of a red segment e_r. We denote by m the minimum of $U(e_r)$. If m intersects $e_r - 1$, then report the intersection, try again with $e_r - 2, \ldots$, until a segment $e_r - j$ is found that does not intersect m. Then m is moved from $U(e_r)$ to $U(e_r - j)$, and we start again with the new minimum of $U(e_r)$. Otherwise, if m does not intersect e_r, we simply merge $U(e_r)$ and $U(e_r - 1)$ thus maintaining Property 4.

The procedure is as follows (**Insert** is the procedure we described in the previous section).

while $m = \min(U(e_r))$ intersects $e_r - 1$
    report the intersection

extract m from $U(e_r)$
   Insert$(m, e_r - 2)$
$U(e_r - 1) = U(e_r - 1) \cup U(e_r)$

Bad pairs may be reported during this operation: we just discard them by comparing the coordinates of the left endpoints.

Two kinds of blue segments are dealt with at this step, the ones that will be merged into $U(e_r - 1)$ and the others. As already mentioned, Property 4 is maintained for the segments of the first category. It is easy to see that the other invariants are maintained as well. A blue segment of the second category is handled by a call to **Insert**$(m, e_r - 2)$, which implies that Properties 1 and 2 are maintained. By Property 4, the intersection of the blue segment m with the sweep line is above e_r. Segment m will, therefore, be above the new segment $l(m)$, and, since m and $l(m)$ do not intersect, Property 3 and Property 4 hold.

3.2.4. Removing a blue segment

Suppose now that the sweep line reaches the right endpoint of a blue segment e_b and let $e_r = l(e_b)$. As we said before, $l(e_b)$ can be obtained from the mergeable heap data structure that stores the sets $U(\cdot)$.

If e_b is not the minimum of $U(e_r)$, we just remove it. Otherwise, we still remove e_b from $U(e_r)$, but then we need to check whether the new minimum m of $U(e_r)$ intersects e_r. If it does, we push it downwards by running the procedure **Insert**$(m, e_r - 1)$, then repeat the whole process with the new minimum until it does not intersect e_r, so that Property 4 remains true. It follows that the other invariants are maintained as well.

extract e_b from $U(e_r)$
while $m = \min(U(e_r))$ intersects e_r
   report this intersection
   extract m from $U(e_r)$
   Insert$(m, e_r - 1)$

Once again, this procedure may report bad pairs but this can be fixed by a simple test.

3.3. Proof of correctness

Let us consider a good pair $(e_r, e_b) \in E_r \times E_b$. We assume that the sweep line has just reached the right endpoint of e_r or e_b, whichever comes first. By Property 3 we know that $l(e_b) < e_r$. Then Property 2 shows that our intersection has been reported previously.

3.4. Analysis

Maintaining the ordered red segments list takes $O(n \log n)$ time. Localizing or removing an endpoint takes $O(\log n)$ time. The other parts of the algorithm take one mergeable heap operation for each event or reported intersection.

We first observe that an intersection cannot be reported twice. Let e_r and e_b be two intersecting segments. After reporting their intersection once, $l(e_b) < e_r$. Moreover, Property 1 states that $l(e_b)$ can only go down the list of red active segments, and we never test the intersection between e_b and a segment

that lies above $l(e_\mathrm{b})$. Altogether, that means that our algorithm runs in $\mathrm{O}(n \log n + kh)$ time if h is the time required to perform a heap operation.

Plainly, we can implement the sets $U(e_\mathrm{r})$ with binomial heaps [17] so that $h = \mathrm{O}(\log n)$. We can do better if we first pre-sort the blue segments, and then implement our algorithm using range-restricted mergeable heaps [5]. This leads to $h = \mathrm{O}(\log \log n)$ and an overall running time $\mathrm{O}(n \log n + k \log \log n)$.

3.5. Degenerate cases

Since this algorithm is elementary, we only need to consider a few degenerate cases which turn out to be very easy to handle. Two kinds of degeneracy may occur: either two endpoints have the same abscissae or an endpoint lies on a curve.

The first case can be solved by extending the partial order on endpoints abscissae to any total order. The order should be the same for each one of the four plane sweeps we perform, for otherwise some intersecting pairs of segments would never be good.

If a point lies on a segment, we just consider that it is above the segment during two sweeps and below the segment during the two sweeps with the y-axis reversed.

Our algorithm can be easily generalized to the case where segments of the same color are allowed to share endpoints, whose main application is map overlay. We only need to consider the segments without their endpoints when we define their vertical order.

4. Conclusion

Given a set of general x-monotone segments, we have presented an algorithm to report the pairs of red–blue segments that intersect. Our algorithm uses only three simple predicates that allow to decide if two segments intersect, if a point is left or right to another point, and if a point is above, below or on a segment. These three predicates seem to be the simplest predicates that lead to subquadratic algorithms. Our algorithm is almost optimal in this restricted model of computation.

Interestingly, the time complexity of our algorithm depends on the number of pairs of intersecting segments, not on the number of intersection points. In particular, our algorithm works even if the segments intersect infinitely many times.

We conclude with some open problems. First, can we remove the $\log \log n$ factor in our time complexity results?

The $\Omega(n\sqrt{k} + n \log n)$ lower bound holds for general curve segments. It has been possible to do better for line segments [4]. Can we also do better for other special curve segments such as circle segments?

In this paper, we have restricted our attention to x-monotone segments. This may be a restriction when the points with a vertical tangent are difficult to compute. This is in particular the case for circles defined by three points where the predicate that compares x-extreme points has degree 20 while the intersection predicate has degree 12 only.

Acknowledgements

References

[1] I.J. Balaban, An optimal algorithm for finding segment intersections, in: Proc. 11th ACM Sympos. Comput. Geom., 1995, pp. 211–219.

[2] J.L. Bentley, T.A. Ottmann, Algorithms for reporting and counting geometric intersections, IEEE Trans. Comput. C-28 (9) (1979) 643–647.

[3] J.-D. Boissonnat, F.P. Preparata, Robust plane sweep for intersecting segments, SIAM J. Comput. 29 (5) (2000) 1401–1421.

[4] J.-D. Boissonnat, J. Snoeyink, Efficient algorithms for line and curve segment intersection using restricted predicates, Computational Geometry 16 (1) (2000) 35–52.

[5] J.D. Bright, Range-restricted mergeable priority queues, Inform. Process. Lett. 47 (3) (1993) 159–164.

[6] T.M. Chan, A simple trapezoid sweep algorithm for reporting red/blue segment intersections, in: Proc. 6th Canad. Conf. Comput. Geom., 1994, pp. 263–268.

[7] T.M. Chan, Reporting curve segment intersections using restricted predicates, Computational Geometry 16 (4) (2000) 245–256.

[8] B. Chazelle, H. Edelsbrunner, An optimal algorithm for intersecting line segments in the plane, J. ACM 39 (1) (1992) 1–54.

[9] K.L. Clarkson, P.W. Shor, Applications of random sampling in computational geometry, II, Discrete Comput. Geom. 4 (1989) 387–421.

[10] F. d'Amore, P.G. Franciosa, G. Liotta, A robust region approach to the computation of geometric graphs, in: G. Bilardi, G.F. Italiano, A. Pietracaprina, G. Pucci (Eds.), Algorithms—ESA'98, Lecture Notes in Computer Science, Vol. 1461, Springer, 1998, pp. 175–186.

[11] O. Devillers, A. Fronville, B. Mourrain, M. Teillaud, Exact predicates for circle arcs arrangements, in: Proc. 16th Annu. ACM Sympos. Comput. Geom., 2000, pp. 139–147.

[12] D.E. Knuth, Axioms and Hulls, Lecture Notes in Computer Sciences, Vol. 606, Springer, Heidelberg, 1992.

[13] G. Liotta, F.P. Preparata, R. Tamassia, Robust proximity queries: An illustration of degree-driven algorithm design, SIAM J. Comput. 28 (3) (1998) 864–889.

[14] S. Pion, Interval arithmetic: An efficient implementation and an application to computational geometry, in: Workshop on Applications of Interval Analysis to Systems and Control, 1999, pp. 99–110.

[15] F.P. Preparata, M.I. Shamos, Computational Geometry: An Introduction, 3rd Edition, Springer, 1990.

[16] S. Schirra, Robustness and precision issues in geometric computation, in: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geometry, Elsevier Science/North-Holland, Amsterdam, 2000, Chapter 14, pp. 597–632.

[17] J. Vuillemin, A data structure for manipulating priority queues, Comm. ACM 21 (4) (1978) 309–315.

[18] C. Yap, Towards exact geometric computation, Computational Geometry 7 (1) (1997) 3–23.

[19] C.K. Yap, Robust geometric computation, in: J.E. Goodman, J. O'Rourke (Eds.), Handbook of Discrete and Computational Geometry, CRC Press LLC, Boca Raton, FL, 1997, Chapter 35, pp. 653–668.