

42 PARALLEL ALGORITHMS IN GEOMETRY

Michael T. Goodrich

INTRODUCTION

The goal of parallel algorithm design is to develop parallel computational methods that run very fast with as few processors as possible, and there is an extensive literature of such algorithms for computational geometry problems. There are several different parallel computing models, and in order to maintain a focus in this chapter, we will describe results in the Parallel Random Access Machine (PRAM) model, which is a synchronous parallel machine model in which processors share a common memory address space (and all inter-processor communication takes place through this shared memory). Although it does not capture all aspects of parallel computing, it does model the essential properties of parallelism. Moreover, it is a widely-accepted model of parallel computation, and all other reasonable models of parallel computation can easily simulate a PRAM.

Interestingly, parallel algorithms can have a direct impact on efficient sequential algorithms, using a technique called *parametric search*. This technique, which is discussed in Chapter 43, involves the use of a parallel algorithm to direct searches in a parameterized geometric space so as to find a critical location (e.g., where an important parameter changes sign or achieves a maximum or minimum value).

The PRAM model is subdivided into submodels based on how one wishes to handle concurrent memory access to the same location. The Exclusive-Read, Exclusive-Write (EREW) variant does not allow for concurrent access. The Concurrent-Read, Exclusive-Write (CREW) variant permits concurrent memory reads, but memory writes must be exclusive. Finally, the Concurrent-Read, Concurrent-Write (CRCW) variant allows for both concurrent memory reading and writing, with concurrent writes being resolved by some simple rule, such as having an arbitrary member of a collection of conflicting writes succeed. One can also define randomized versions of each of these models (e.g., an rCRCW PRAM), where in addition to the usual arithmetic and comparison operations, each processor can generate a random number from 1 to n in one step.

Early work in parallel computational geometry, in the way we define it here, began with the work of Chow [Cho80], who designed several parallel algorithms with polylogarithmic running times using a linear number of processors. Subsequent to this work, several researchers initiated a systematic study of work-efficient parallel algorithms for geometric problems, including Aggarwal *et al.* [ACG⁺88], Akl [Akl82, Akl84, Akl85], Amato and Preparata [AP92, AP95], Atallah and Goodrich [AG86, Goo87], and Reif and Sen [RS92a, Sen89].

In Section 42.1 we give a brief discussion of general techniques for parallel geometric algorithm design. We then partition the research in parallel computational geometry into problems dealing with convexity (Section 42.2), arrangements and decompositions (Section 42.3), proximity (Section 42.4), geometric searching (Section 42.5), and visibility, envelopes, and geometric optimization (Section 42.6).

42.1 SOME PARALLEL TECHNIQUES

The design of efficient parallel algorithms for computational geometry problems often depends upon the use of powerful general parallel techniques (e.g., see [AL93, J92, KR90, Rei93]). We review some of these techniques below.

PARALLEL DIVIDE-AND-CONQUER

Possibly the most general technique is parallel divide-and-conquer. In applying this technique one divides a problem into two or more subproblems, solves the subproblems recursively in parallel, and then merges the subproblem solutions to solve the entire problem. As an example application of this technique, consider the problem of constructing the upper convex hull of a S set of n points in the plane presorted by x -coordinates. Divide the list S into $\lceil \sqrt{n} \rceil$ contiguous sublists of size $\lfloor \sqrt{n} \rfloor$ each and recursively construct the upper convex hull of the points in each list. Assign a processor to each pair of sublists and compute the common upper tangent line for the two upper convex hulls for these two lists, which can be done in $O(\log n)$ time using a well-known “binary search” computation [Ede87, O’R98, PS85]. By maximum computations on the left and right common tangents, respectively, for each subproblem S_i , one can determine which vertices on the upper convex hull of S_i belong to the upper convex hull of S . Compressing all the vertices identified to be on the upper convex hull of S constructs an array representation of this hull, completing the construction.

The running time of this method is characterized by the recurrence relation $T(n) \leq T(\sqrt{n}) + O(\log n)$, which implies that $T(n)$ is $O(\log n)$. It is important to note that the coefficient for the $T(\sqrt{n})$ term is 1 even though we had $\lceil \sqrt{n} \rceil$ subproblems, for all these subproblems were processed simultaneously in parallel. The number of processors needed for this computation can be characterized by the recurrence relation $P(n) \leq \max\{\lceil \sqrt{n} \rceil P(\sqrt{n}), n\}$, which implies that $P(n)$ is $O(n)$. Thus, the *work* needed for this computation is $O(n \log n)$, which is not quite optimal. Still, this method can be adapted to result in an optimal work bound [BSVar, Che95, GG97].

BUILD-AND-SEARCH

Another important technique in parallel computational geometry is the build-and-search technique. It is a paradigm that often yields efficient parallel adaptations of sequential algorithms designed using the powerful plane sweeping technique. In the build-and-search technique, the solution to a problem is partitioned into a *build* phase, where one constructs in parallel a data structure built from the geometric data present in the problem, and a *search* phase, where one searches this data structure in parallel to solve the problem at hand. An example of an application of this technique is for the trapezoidal decomposition problem: given a collection of nonintersecting line segments in the plane, determine the first segments intersected by vertical rays emanating from each segment endpoint (cf. Figure 40.?). The existing efficient parallel algorithm for this problem is based upon first building in parallel a data structure on the input set of segments that allows for such vertical

ray-shooting queries to be answered in $O(\log n)$ time by a single processor, and then querying this structure for each segment endpoint in parallel. This results in a parallel algorithm with an efficient $O(n \log n)$ work bound and fast $O(\log n)$ query time.

42.2 CONVEXITY

Results on the problem of constructing the convex hull of n points in \mathbb{R}^d are summarized in Table 42.2.1, for various fixed values of d , and, in the case of $d = 2$, under assumptions about whether the input is presorted. We restrict our attention to parallel algorithms with efficient work bounds, where we use the term *work* of an algorithm here to refer to the product of its running time and the number of processors used by the algorithm. A parallel algorithm has an *optimal* work bound if the work used asymptotically matches the sequential lower bound for the problem. In the table, h denotes the size of the hull, and c is some fixed constant. Also, we use (throughout this chapter) $\bar{O}(f(n))$ to denote an asymptotic bound that holds with high probability.

TABLE 42.2.1 Parallel convex hull algorithms.

PROBLEM	MODEL	TIME	WORK	REF
2Dim presorted	rand-CRCW	$\bar{O}(\log^* n)$	$\bar{O}(n)$	[GG91]
2Dim presorted	CRCW	$O(\log \log n)$	$O(n)$	[BSVar]
2Dim presorted	EREW	$O(\log n)$	$O(n)$	[Che95]
2Dim polygon	EREW	$O(\log n)$	$O(n)$	[Che95]
2Dim	rand-CRCW	$\bar{O}(\log n)$	$\bar{O}(n \log h)$	[GG91]
2Dim	EREW	$O(\log n)$	$O(n \log n)$	[MS88]
2Dim	EREW	$O(\log^2 n)$	$O(n \log h)$	[GG91]
3D	rand-CRCW	$\bar{O}(\log n)$	$\bar{O}(n \log n)$	[RS92b]
3D	CREW	$O(\log n)$	$O(n^{1+1/c})$	[AP93]
3D	EREW	$O(\log^2 n)$	$O(n \log n)$	[AGR94]
3D	EREW	$O(\log^3 n)$	$O(n \log h)$	[AGR94]
Fixed $d \geq 4$	rand-EREW	$\bar{O}(\log^2 n)$	$\bar{O}(n^{\lfloor d/2 \rfloor})$	[AGR94]
Even $d \geq 4$	EREW	$O(\log^2 n)$	$O(n^{\lfloor d/2 \rfloor})$	[AGR94]
Odd $d > 4$	EREW	$O(\log^2 n)$	$O(n^{\lfloor d/2 \rfloor} \log^c n)$	[AGR94]

We discuss a few of these algorithms to illustrate their flavor.

2-DIMENSIONAL CONVEX HULLS

The two-dimensional convex hull algorithm of Miller and Stout [MS88] is based upon a parallel divide-and-conquer scheme where one presorts the input and then divides it into many subproblems ($O(n^{1/4})$ in their case), solves each subproblem independently in parallel, and then merges all the subproblem solutions together

in $O(\log n)$ parallel time. Of course, the difficult step is the merge of all the sub-problems, with the principal difficulty being the computation of common tangents between hulls. The total running time is characterized by the recurrence

$$T(n) \leq T(n^{1/4}) + O(\log n),$$

which solves to $T(n) = O(\log n)$.

3-DIMENSIONAL CONVEX HULLS

All of the 3D convex hull algorithms listed in Table 42.2.1 are also based upon this many-way, divide-and-conquer paradigm, except that there is no notion of presorting in three dimensions, so the subdivision step also becomes nontrivial. Reif and Sen [RS92b] use a random sample to perform the division, and the methods of Amato, Goodrich, and Ramos [AGR94] derandomize this approach. Amato and Preparata [AP93] use parallel separating planes, an approach extended to higher dimension in [AGR94].

LINEAR PROGRAMMING

A problem strongly related to convex hull construction, which has also been addressed in a parallel setting, is d -dimensional linear programming, for fixed dimensions d (see Chapter 45). Of course, one could solve this problem by transforming it to its dual problem, constructing a convex hull in this dual space, and then evaluating each vertex in the simplex that is dual to this convex hull. This would be quite inefficient, however, for $d \geq 4$. The best parallel bounds for this problem are listed in Table 42.2.2. See Section 45.6 for a detailed discussion.

TABLE 42.2.2 Fixed d -dimensional parallel linear programming.

MODEL	TIME	WORK	REF
Rand-CRCW	$\bar{O}(1)$	$\bar{O}(n)$	[AM90]
CRCW	$O((\log \log n)^{d-1})$	$O(n)$	[GR97]
EREW	$O(\log n (\log \log n)^{d-1})$	$O(n)$	[Goo96]

OPEN PROBLEMS

There are a number of interesting open problems regarding convexity:

1. Can d -dimensional linear programming be solved (deterministically) in $O(\log n)$ time using $O(n)$ work in the CREW PRAM model?
2. Is there an efficient output-sensitive parallel convex hull algorithm for $d \geq 4$?
3. Is there an optimal-work $O(\log^2 n)$ -time CREW PRAM convex hull algorithm for odd dimensions greater than 4?

42.3 ARRANGEMENTS AND DECOMPOSITIONS

Another important class of geometric problems that has been addressed in the parallel setting are arrangement and decomposition problems, which deal with ways of partitioning space. We review the best parallel bounds for such problems in Table 42.3.1.

GLOSSARY

Arrangement: The partition of space determined by the intersections of a collection of geometric objects, such as lines, line segments, or (in higher dimensions) hyperplanes. In this chapter, algorithms for constructing arrangements produce the *incidence graph*, which stores all adjacency information between the various primitive topological entities determined by the partition, such as intersection points, edges, faces, etc. See Section 24.3.1.

Red-blue arrangement: An arrangement defined by two sets of objects A and B such that the objects in A (resp. B) are nonintersecting.

Axis-parallel: All segments/lines are parallel to one of the coordinate axes.

Polygon triangulation: A decomposition of the interior of a polygon into triangles by adding diagonals between vertices. See Section 26.2.

Trapezoidal decomposition: A decomposition of the plane into trapezoids (and possibly triangles) by adding appropriate vertical line segments incident to vertices. See Section 34.3.

Star-shaped polygon: A (simple) polygon that is completely visible from a single point. A polygon with nonempty kernel. See Section 26.1.

$1/r$ -cutting: A partition of \mathbb{R}^d into $O(r^d)$ simplices such that each simplex intersects at most n/r hyperplanes. See Sections 36.2 and 40.1.

TABLE 42.3.1 Parallel arrangement and decomposition algorithms.

PROBLEM	MODEL	TIME	WORK	REF
d -dim hyperplane arr	EREW	$O(\log n)$	$O(n^d)$	[AGR94]
2D seg arr	rand-CRCW	$\bar{O}(\log n)$	$\bar{O}(n \log n + k)$	[CCT92a, CCT92b]
2D axis-par seg arr	CREW	$O(\log n)$	$O(n \log n + k)$	[Goo91]
2D red-blue seg arr	CREW	$O(\log n)$	$O(n \log n + k)$	[GSG92, GSG93, Rüb92]
2D seg arr	EREW	$O(\log^2 n)$	$O(n \log n + k)$	[AGR95]
Polygon triangulation	CRCW	$O(\log n)$	$O(n)$	[Goo95]
Polygon triangulation	CREW	$O(\log n)$	$O(n \log n)$	[Goo89, Yap88]
2D nonint seg trap decomp	CREW	$O(\log n)$	$O(n \log n)$	[ACG89]
2D quadtree decomp	EREW	$O(\log n)$	$O(n \log n + k)$	[BET99]

We sketch the one randomized algorithm in Table 42.3.1 to illustrate how randomization and parallel computation can be mixed. Let S be a set of segments in

the plane with k intersecting pairs. The goal is to construct $\mathring{A}(S)$, the arrangement induced by S . First, an estimate \hat{k} for k is obtained from a random sample. Then a random subset $R \subset S$ of a size r dependent on \hat{k} is selected. $\mathring{A}(R)$ is constructed using a suboptimal parallel algorithm, and processed (in parallel) for point location. Next the segments intersecting each cell of $\mathring{A}(R)$ are found using a parallel point-location algorithm, together with some ad hoc techniques. Visibility information among the segments meeting each cell is computed using another suboptimal parallel algorithm. Finally, the resulting cells are merged in parallel. Because various key parameters in the suboptimal algorithms are kept small by the sampling, optimal expected work is achieved.

All of the algorithms for computing segment arrangements are *output-sensitive*, in that their work bounds depend upon both the input size and the output size. In these cases we must slightly extend our computational model to allow for the machine to request additional processors if necessary. In all these algorithms, this request may originate only from a single “master” processor, however, so this modification is not that different from our assumption that the number of processors assigned to a problem can be a function of the input size. Of course, to solve a problem on a real parallel computer, one would simulate one of these efficient parallel algorithms to achieve an optimal speed-up over what would be possible using a sequential method.

A related class of intersection-related problems is the class of problems dealing with methods for detecting intersections. Testing if a collection of objects has at least one intersection is frequently easier than finding all such intersections, and Table 42.3.2 reviews such results in the parallel domain.

GLOSSARY

Star-shaped polygon: A (simple) polygon that is completely visible from a single point; a polygon with non-empty kernel. See Chapter 26.

TABLE 42.3.2 Parallel intersection detection algorithms.

PROBLEM	MODEL	TIME	WORK	REF
2 convex polygons	CREW	$O(1)$	$O(n^{1/c})$	[DK89a]
2 star-shaped polygons	CREW	$O(\log n)$	$O(n)$	[GM91]
2 convex polyhedra	CREW	$O(\log n)$	$O(n)$	[DK89a]

Given a collection of n hyperplanes in \mathbb{R}^d , another important decomposition problem is the construction of a $(1/r)$ -cutting. Here an EREW algorithm running in $O(\log n \log r)$ time using $O(nr^{d-1})$ work has been obtained [Goo93].

OPEN PROBLEMS

1. Is there an optimal-work $O(\log n)$ -time polygon triangulation algorithm that doesn't use concurrent writes?
2. Can a line segment arrangement be constructed in $O(\log n)$ time using $O(n \log n + k)$ work in the CREW PRAM model?

42.4 PROXIMITY

An important property of Euclidean space is that it is a metric space, and distance plays an important role in many computational geometry applications. For example, computing a closest pair of points can be used in collision detection, as can the more general problem of computing the nearest neighbor of each point in a set S , a problem we will call the *all-nearest neighbors (ANN)* problem. Perhaps the most fundamental problem in this domain is the subdivision of space into regions where each region $V(s)$ is defined by a *site* s in a set S of geometric objects such that each point in $V(s)$ is closer to s than to any other object in S . This subdivision is the *Voronoi diagram* (Chapter 23); its graph-theoretic dual, which is also an important geometric structure, is the *Delaunay triangulation* (Section 25.1). For a set of points S in \mathbb{R}^d , there is a simple “lifting” transformation that takes each point $(x_1, x_2, \dots, x_d) \in S$ to the point $(x_1, x_2, \dots, x_d, x_1^2 + x_2^2 + \dots + x_d^2)$, forming a set of points S' in \mathbb{R}^{d+1} (Section 23.1). Each simplex on the convex hull of S' with a negative $(d+1)$ -st component in its normal vector projects back to a simplex of the Delaunay triangulation in \mathbb{R}^d . Thus, any $(d+1)$ -dimensional convex hull algorithm immediately implies a d -dimensional Voronoi diagram (VD) algorithm. Table 42.4.1 summarizes the bounds of efficient parallel algorithms for constructing Voronoi diagrams in this way, as well as methods that are designed particularly for Voronoi diagram construction or other specific proximity problems. (In the table, the underlying objects are points unless stated otherwise.)

GLOSSARY

Convex position: A set of points that are all on the boundary of their convex hull.

Voronoi diagram for line segments: A Voronoi diagram that is defined by a set of nonintersecting line segments, with distance from a point p to a segment s being defined as the distance from p to a closest point on s . See Section 23.3.

OPEN PROBLEMS

1. Can a 2-dimensional Voronoi diagram be constructed in $O(\log n)$ time using $O(n \log n)$ work under either the CREW or EREW PRAM models?

TABLE 42.4.1 Parallel proximity algorithms.

PROBLEM	MODEL	TIME	WORK	REF
2-dim ANN in convex pos	EREW	$O(\log n)$	$O(n)$	[CG92]
2-dim ANN	EREW	$O(\log n)$	$O(n \log n)$	[CG92]
d -dim ANN	CREW	$O(\log n)$	$O(n \log n)$	[Cal93]
2-dim VD in L_1 metric	CREW	$O(\log n)$	$O(n \log n)$	[WC90]
2-dim VD	rand-CRCW	$\bar{O}(\log n)$	$\bar{O}(n \log n)$	[RS92b]
2-dim VD	CRCW	$O(\log n \log \log n)$	$O(n \log n \log \log n)$	[CGÓ90]
2-dim VD	EREW	$O(\log^2 n)$	$O(n \log n)$	[AGR94]
2-dim VD for segments	CREW	$O(\log^2 n)$	$O(n \log^2 n)$	[GÓY93]
3-dim VD	EREW	$O(\log^2 n)$	$O(n^2)$	[AGR94]

2. Is there an efficient output-sensitive parallel algorithm for constructing 3D Voronoi diagrams?

42.5 GEOMETRIC SEARCHING

Given a subdivision of space by a collection S of geometric objects, such as line segments, the point location problem is to build a data structure for this set that can quickly answer *vertical ray-shooting queries*, where one is given a point p and asked to report the first object in S hit by a vertical ray from p . We summarize efficient parallel algorithms for planar point location in Table 42.5.1. The time and work bounds listed, as well as the computational model, are for building the data structure to achieve an $O(\log n)$ query time. We do not list the space bounds for any of these methods in the table since, in every case, they are equal to the preprocessing work bounds.

GLOSSARY

- Arbitrary planar subdivision:** A subdivision of the plane (not necessarily connected), defined by a set of line segments that intersect only at their endpoints.
- Monotone subdivision:** A connected subdivision of the plane in which each face is intersected by a vertical line in a single segment.
- Triangulated subdivision:** A connected subdivision of the plane into triangles whose corners are vertices of the subdivision (see Chapter 25).
- Shortest path in a polygon:** The shortest path between two points that does not go outside of the polygon (see Section 26.4).
- Ray-shooting query:** A query whose answer is the first object hit by a ray oriented in a specified direction from a specified point.

TABLE 42.5.1 Parallel geometric searching algorithms.

QUERY PROBLEM	MODEL	TIME	WORK	REF
Point loc in arb subdivision	CREW	$O(\log n)$	$O(n \log n)$	[ACG89]
Point loc in monotone subdivision	EREW	$O(\log n)$	$O(n)$	[TV91]
Point loc in triangulated subdivision	CREW	$O(\log n)$	$O(n)$	[CZ90]
Point loc in d -dim hyp arr	EREW	$O(\log n)$	$O(n^d)$	[AGR94]
Shortest path in triangulated polygon	CREW	$O(\log n)$	$O(n)$	[GSG92]
Ray shooting in triangulated polygon	CREW	$O(\log n)$	$O(n)$	[HS93]
Line & convex polyhedra intersection	CREW	$O(\log n)$	$O(n)$	[DK89b, CZ90]

OPEN PROBLEMS

1. Is there an efficient data structure that allows n simultaneous point locations to be performed in $O(\log n)$ time using $O(n)$ processors in the EREW PRAM model?
2. Is there an efficient data structure for 3-dimensional point location in convex subdivisions that can be constructed in $O(n \log n)$ work and at most $O(\log^2 n)$ time and which allows for a query time that is at most $O(\log^2 n)$?

42.6 VISIBILITY, ENVELOPES, AND OPTIMIZATION

We summarize efficient parallel methods for various visibility and lower envelope problems for a simple polygon in Table 42.6.1. In the table, m denotes the number of edges in a visibility graph. For definitions see Chapter 28.

TABLE 42.6.1 Parallel visibility algorithms for a simple polygon.

PROBLEM	MODEL	TIME	WORK	REF
Kernel	EREW	$O(\log n)$	$O(n)$	[Che95]
Vis from a point	EREW	$O(\log n)$	$O(n)$	[ACW91]
Vis from an edge	CRCW	$O(\log n)$	$O(n)$	[Her92]
Vis from an edge	CREW	$O(\log n)$	$O(n \log n)$	[GSG92, GSG93]
Vis graph	CREW	$O(\log n)$	$O(n \log^2 n + m)$	[GSG92, GSG93]

We sketch the algorithm for computing the point visibility polygon [ACW91], which is notable for two reasons: first, it is employed as a subprogram in many other algorithms; and second, it requires much more intricate processing and analysis than the relatively simple optimal sequential algorithm (Section 25.3). The parallel algorithm is recursive, partitioning the boundary into $n^{1/4}$ subchains, and computing *visibility chains* from the source point of visibility x . Each of these chains

is star-shaped with respect to x , i.e., effectively “monotone” (see Section 26.1). This monotonicity property is, however, insufficient to intersect the visibility chains quickly enough in the merge step to obtain optimal bounds. Rather, the fact that the chains are subchains of the boundary of a simple polygon must be exploited to achieve logarithmic-time computation of the intersection of two chains. This then leads to the optimal bounds quoted in Table 42.6.1.

The bounds of efficient parallel methods for visibility problems on general sets of segments and curves in the plane are summarized in Table 42.6.2.

GLOSSARY

Lower envelope: The function $F(x)$ defined as the point-wise minimum of a collection of functions $\{f_1, f_2, \dots, f_n\}$: $F(x) = \min_i f_i(x)$ (see Section 21.2).

k -intersecting curves: A set of curves every two of which intersect at most k times (where they cross).

$\lambda_s(n)$: The maximum length of a Davenport-Schinzel sequence [SA95, AS00] of order s on n symbols. If s is a constant, $\lambda_s(n)$ is $o(n \log^* n)$. See Section 40.4.

TABLE 42.6.2 General parallel visibility and enveloping algorithms.

PROBLEM	MODEL	TIME	WORK	REF
Lower env for segments	EREW	$O(\log^2 n)$	$O(n \log n)$	[Her89]
Lower env for k -int curves	EREW	$O(\log^2 n)$	$O(\lambda_{k+2}(n) \log n)$	[BM87]

Finally, we summarize some efficient parallel algorithms for solving several geometric optimization problems in Table 42.6.3.

GLOSSARY

Largest-area empty rectangle: For a collection S of n points in the plane, the largest-area rectangle that does not contain any point of S in its interior.

All-farthest neighbors problem in a simple polygon: Determine for each vertex p of a simple polygon the vertex q such that the shortest path from p is longest.

Closest visible-pair between polygons: A closest pair of mutually-visible vertices between two nonintersecting simple polygons in the plane.

Minimum circular-arc cover: For a collection of n arcs of a given circle C , a minimum-cardinality subset that covers C .

Optimal-area inscribed/circumscribed triangle: For a convex polygon P , the largest-area triangle inscribed in P , or, respectively, the smallest-area triangle circumscribing P .

Min-link path in a polygon: A piecewise-linear path of fewest “links” inside a simple polygon between two given points p and q ; see Sections 23.4 and 24.3.

TABLE 42.6.3 Parallel geometric optimization algorithms.

PROBLEM	MODEL	TIME	WORK	REF
Largest-area empty rectangle	CREW	$O(\log^2 n)$	$O(n \log^3 n)$	[AKPS90]
All-farthest neighbors in polygon	CREW	$O(\log^2 n)$	$O(n \log^2 n)$	[Guh92]
Closest visible-pair btw polygons	CREW	$O(\log n)$	$O(n \log n)$	[HCL92]
Min circular-arc cover	EREW	$O(\log n)$	$O(n \log n)$	[AC89]
Opt-area inscr/circum triangle	CRCW	$O(\log \log n)$	$O(n)$	[CM92]
Opt-area inscr/circum triangle	CREW	$O(\log n)$	$O(n)$	[CM92]
Min-link path in a polygon	CREW	$O(\log n \log \log n)$	$O(n \log n \log \log n)$	[CGM ⁺ 90]

OPEN PROBLEMS

1. Can the visibility graph of a set of n nonintersecting line segments be constructed using $O(n \log n + m)$ work in time at most $O(\log^2 n)$ in the CREW model, where m is the size of the graph?
2. Can the visibility graph of a triangulated polygon be computed in $O(\log n)$ time using $O(n + m)$ work in the CREW model?

42.7 SOURCES AND RELATED MATERIAL

FURTHER READING

Our presentation has been results-oriented and has not provide much problem intuition or algorithmic techniques. There are several excellent surveys available in the literature [Ata92, AC94, AC00, AG93, RS93, RS00] that are more techniques-oriented. Another good location for related material is the book by Akl and Lyons [AL93].

RELATED CHAPTERS

Chapter 22: Convex hull computations
 Chapter 23: Voronoi diagrams and Delaunay triangulations
 Chapter 24: Arrangements
 Chapter 26: Polygons
 Chapter 34: Point location
 Chapter 38: Geometric intersection
 Chapter 40: Randomization and derandomization
 Chapter 45: Linear programming in low dimensions

REFERENCES

- [AS00] P.K. Agarwal and M. Sharir. Davenport-Schinzel sequences and their geometric applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 1–47. Elsevier North-Holland, Amsterdam, 2000.
- [ACG⁺88] A. Aggarwal, B. Chazelle, L.J. Guibas, C. Ó’Dúnlaing, and C. Yap. Parallel computational geometry. *Algorithmica*, 3:293–327, 1988.
- [AKPS90] A. Aggarwal, D. Kravets, J.K. Park, and S. Sen. Parallel searching in generalized Monge arrays with applications. In *Proc. 2nd Annu. ACM Sympos. Parallel Algorithms Architect.*, pages 259–268, 1990.
- [Akl82] S.G. Akl. A constant-time parallel algorithm for computing convex hulls. *BIT*, 22:130–134, 1982.
- [Akl84] S.G. Akl. Optimal algorithms for computing convex hulls and for sorting. *Computing*, 33:1–11, 1984.
- [Akl85] S.G. Akl. Optimal parallel algorithms for selection, sorting and computing convex hulls. In G.T. Toussaint, editor, *Computational Geometry*, pages 1–22. North-Holland, Amsterdam, Netherlands, 1985.
- [AL93] S.G. Akl and K.A. Lyons. *Parallel Computational Geometry*. Prentice Hall, 1993.
- [AM90] N. Alon and N. Megiddo. Parallel linear programming in fixed dimension almost surely in constant time. In *Proc. 31st Annu. IEEE Sympos. Found. Comput. Sci.*, pages 574–582, 1990.
- [AGR94] N.M. Amato, M.T. Goodrich, and E.A. Ramos. Parallel algorithms for higher-dimensional convex hulls. In *Proc. 35th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 683–694, 1994.
- [AGR95] N.M. Amato, M.T. Goodrich, and E.A. Ramos. Computing faces in segment and simplex arrangements. In *Proc. 27th Annu. ACM Sympos. Theory Comput.*, pages 672–682, 1995.
- [AP92] N.M. Amato and F.P. Preparata. The parallel 3D convex hull problem revisited. *Internat. J. Comput. Geom. Appl.*, 2:163–173, 1992.
- [AP93] N.M. Amato and F.P. Preparata. An NC¹ parallel 3D convex hull algorithm. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 289–297, 1993.
- [AP95] N.M. Amato and F.P. Preparata. A time-optimal parallel algorithm for three-dimensional convex hulls. *Algorithmica*, 14:169–182, 1995.
- [Ata92] M.J. Atallah. Parallel techniques for computational geometry. *Proc. IEEE*, 80:1435–1448, 1992.
- [AC89] M.J. Atallah and D.Z. Chen. An optimal parallel algorithm for the minimum circle-cover problem. *Inform. Process. Lett.*, 34:159–165, 1989.
- [AC94] M.J. Atallah and D.Z. Chen. Parallel computational geometry. In A.Y. Zomaya, editor, *Parallel Computations: Paradigms and Applications*. World Scientific, 1994.
- [ACW91] M.J. Atallah, D.Z. Chen, and H. Wagener. Optimal parallel algorithm for visibility of a simple polygon from a point. *J. Assoc. Comput. Mach.*, 38:516–553, 1991.
- [ACG89] M.J. Atallah, R. Cole, and M.T. Goodrich. Cascading divide-and-conquer: A technique for designing parallel algorithms. *SIAM J. Comput.*, 18:499–532, 1989.
- [AG86] M.J. Atallah and M.T. Goodrich. Efficient parallel solutions to some geometric problems. *J. Parallel Distrib. Comput.*, 3:492–507, 1986.

- [AG93] M.J. Atallah and M.T. Goodrich. Deterministic parallel computational geometry. In J.H. Reif, editor, *Synthesis of Parallel Algorithms*, pages 497–536. Morgan Kaufmann, 1993.
- [AC00] M.J. Atallah and D.Z. Chen. Deterministic parallel computational geometry. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 155–200. Elsevier North-Holland, Amsterdam, 2000.
- [BSVar] O. Berkman, B. Schieber, and U. Vishkin. A fast parallel algorithm for finding the convex hull of a sorted point set. *Internat. Journal on Computational Geometry & Applications*, to appear.
- [BET99] M. Bern, D. Eppstein, and S.-H. Teng. Parallel construction of quadtrees and quality triangulations. *Internat. J. Comput. Geom. Appl.*, 9:517–532, 1999.
- [BM87] L. Boxer and R. Miller. Parallel dynamic computational geometry. Report 87-11, Dept. Comput. Sci., SUNY-Buffalo, 1987.
- [Cal93] P.B. Callahan. Optimal parallel all-nearest-neighbors using the well-seated pair decomposition. In *Proc. 34th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 332–340, 1993.
- [CM92] S. Chandran and D.M. Mount. A parallel algorithm for enclosed and enclosing triangles. *Internat. J. Comput. Geom. Appl.*, 2:191–214, 1992.
- [CGM⁺90] V. Chandru, S.K. Ghosh, A. Maheshwari, V.T. Rajan, and S. Saluja. NC-algorithms for minimum link path and related problems. Technical Report CS-90/3, TATA inst., Bombay, India, 1990.
- [Che95] D. Chen. Efficient geometric algorithms on the EREW PRAM. *IEEE Trans. Parallel Distrib. Syst.*, 6:41–47, 1995.
- [Cho80] A.L. Chow. *Parallel algorithms for geometric problems*. Ph.D. thesis, Dept. Comput. Sci., Univ. Illinois, Urbana, IL, 1980.
- [CCT92a] K.L. Clarkson, R. Cole, and R.E. Tarjan. Erratum: Randomized parallel algorithms for trapezoidal diagrams. *Internat. J. Comput. Geom. Appl.*, 2:341–343, 1992.
- [CCT92b] K.L. Clarkson, R. Cole, and R.E. Tarjan. Randomized parallel algorithms for trapezoidal diagrams. *Internat. J. Comput. Geom. Appl.*, 2:117–133, 1992.
- [CG92] R. Cole and M.T. Goodrich. Optimal parallel algorithms for polygon and point-set problems. *Algorithmica*, 7:3–23, 1992.
- [CGÓ90] R. Cole, M.T. Goodrich, and C. Ó’Dúnlaing. Merging free trees in parallel for efficient Voronoi diagram construction. In *Proc. 17th Internat. Colloq. Automata Lang. Program.*, volume 443 of *Lecture Notes Comput. Sci.*, pages 432–445. Springer-Verlag, 1990.
- [CZ90] R. Cole and O. Zajicek. An optimal parallel algorithm for building a data structure for planar point location. *J. Parallel Distrib. Comput.*, 8:280–285, 1990.
- [DK89a] N. Dadoun and D.G. Kirkpatrick. Cooperative subdivision search algorithms with applications. In *Proc. 27th Allerton Conf. Commun. Control Comput.*, pages 538–547, 1989.
- [DK89b] N. Dadoun and D.G. Kirkpatrick. Parallel construction of subdivision hierarchies. *J. Comput. Syst. Sci.*, 39:153–165, 1989.
- [Ede87] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monogr. Theoret. Comput. Sci.* Springer-Verlag, Heidelberg, 1987.
- [GM91] S.K. Ghosh and A. Maheshwari. An optimal parallel algorithm for determining the intersection type of two star-shaped polygons. In *Proc. 3rd Canad. Conf. Comput. Geom.*, pages 2–6, 1991.

-
- [GG97] M. Ghose and M. Goodrich. Fast randomized parallel methods for planar convex hull construction. *Comput. Geom. Theory Appl.*, 7:219–236, 1997.
- [GG91] M. Ghose and M.T. Goodrich. In-place techniques for parallel convex hull algorithms. In *Proc. 3rd Annu. ACM Sympos. Parallel Algorithms Architect.*, pages 192–203, 1991.
- [Goo87] M.T. Goodrich. *Efficient parallel techniques for computational geometry*. Ph.D. thesis, Dept. Comput. Sci., Purdue Univ., West Lafayette, IN, 1987.
- [Goo89] M.T. Goodrich. Triangulating a polygon in parallel. *J. Algorithms*, 10:327–351, 1989.
- [Goo91] M.T. Goodrich. Intersecting line segments in parallel with an output-sensitive number of processors. *SIAM J. Comput.*, 20:737–755, 1991.
- [Goo93] M.T. Goodrich. Geometric partitioning made easier, even in parallel. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 73–82, 1993.
- [Goo95] M.T. Goodrich. Planar separators and parallel polygon triangulation. *J. Comput. Syst. Sci.*, 51:374–389, 1995.
- [GÓY93] M.T. Goodrich, C. Ó’Dúnlaing, and C. Yap. Constructing the Voronoi diagram of a set of line segments in parallel. *Algorithmica*, 9:128–141, 1993.
- [GR97] M.T. Goodrich and E.A. Ramos. Bounded-independence derandomization of geometric partitioning with applications to parallel fixed-dimensional linear programming. *Discrete Comput. Geom.*, 18:397–420, 1997.
- [GSG92] M.T. Goodrich, S. Shauck, and S. Guha. Parallel methods for visibility and shortest path problems in simple polygons. *Algorithmica*, 8:461–486, 1992.
- [GSG93] M.T. Goodrich, S. Shauck, and S. Guha. Addendum to “parallel methods for visibility and shortest path problems in simple polygons”. *Algorithmica*, 9:515–516, 1993.
- [Goo96] M.T. Goodrich. Fixed-dimensional parallel linear programming via relative epsilon-approximations. In *Proc. 7th ACM-SIAM Sympos. Discrete Algorithms*, pages 132–141, 1996.
- [Guh92] S. Guha. Parallel computation of internal and external farthest neighbours in simple polygons. *Internat. J. Comput. Geom. Appl.*, 2:175–190, 1992.
- [Her89] J. Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Inform. Process. Lett.*, 33:169–174, 1989.
- [Her92] J. Hershberger. Optimal parallel algorithms for triangulated simple polygons. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 33–42, 1992.
- [HS93] J. Hershberger and S. Suri. A pedestrian approach to ray shooting: Shoot a ray, take a walk. In *Proc. 4th ACM-SIAM Sympos. Discrete Algorithms*, pages 54–63, January 1993.
- [HCL92] F.R. Hsu, R.C. Chang, and R.C.T. Lee. Parallel algorithms for computing the closest visible vertex pair between two polygons. *Internat. J. Comput. Geom. Appl.*, 2:135–162, 1992.
- [J92] J. JáJá. *An Introduction to Parallel Algorithms*. Addison-Wesley, Reading, Mass., 1992.
- [KR90] R.M. Karp and V. Ramachandran. Parallel algorithms for shared memory machines. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 869–941. Elsevier/The MIT Press, Amsterdam, 1990.
- [MS88] R. Miller and Q.F. Stout. Efficient parallel convex hull algorithms. *IEEE Trans. Comput.*, C-37:1605–1618, 1988.
- [O’R98] J. O’Rourke. *Computational Geometry in C*. Cambridge University Press, 2nd edition, 1998.

- [PS85] F.P. Preparata and M.I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, 1985.
- [RS93] S. Rajasekaran and S. Sen. Random sampling techniques and parallel algorithms design. In J.H. Reif, editor, *Synthesis of Parallel Algorithms*, pages 411–452. Morgan Kaufmann, 1993.
- [Rei93] J.H. Reif. *Synthesis of Parallel Algorithms*. Morgan Kaufmann, Inc., San Mateo, CA, 1993.
- [RS92a] J.H. Reif and S. Sen. Optimal parallel randomized algorithms for three-dimensional convex hulls and related problems. *SIAM J. Comput.*, 21:466–485, 1992.
- [RS92b] J.H. Reif and S. Sen. Optimal parallel randomized algorithms for three-dimensional convex hulls and related problems. *SIAM J. Comput.*, 21:466–485, 1992.
- [RS00] J.H. Reif and S. Sen. Parallel computational geometry: An approach using randomization. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 765–828. Elsevier North-Holland, Amsterdam, 2000.
- [Rüb92] C. Rüb. Computing intersections and arrangements for red-blue curve segments in parallel. In *Proc. 4th Canad. Conf. Comput. Geom.*, pages 115–120, 1992.
- [Sen89] S. Sen. *Random Sampling Techniques for Efficient Parallel Algorithms in Computational Geometry*. Ph.D. thesis, Dept. of CS, Duke Univ., 1989.
- [SA95] M. Sharir and P.K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
- [TV91] R. Tamassia and J.S. Vitter. Parallel transitive closure and point location in planar structures. *SIAM J. Comput.*, 20:708–725, 1991.
- [WC90] Y.C. Wee and S. Chaiken. An optimal parallel L_1 -metric Voronoi diagram algorithm. In *Proc. 2nd Canad. Conf. Comput. Geom.*, pages 60–65, 1990.
- [Yap88] C.K. Yap. Parallel triangulation of a polygon in two calls to the trapezoidal map. *Algorithmica*, 3:279–288, 1988.