

External-Memory Computational Geometry

(Preliminary Version)

Michael T. Goodrich*

Dept. of Computer Science
The Johns Hopkins University
Baltimore, MD 21218

Jyh-Jong Tsay†

Dept. of Computer Science
and Information Engineering
National Chung Cheng University
Chiayi 621, Taiwan, ROC

Darren Erik Vengroff‡

Dept. of Computer Science
Brown University
Providence, RI 02912-1910

Jeffrey Scott Vitter§

Dept. of Computer Science
Duke University
Durham, NC 27708-0129

Abstract

In this paper we give new techniques for designing efficient algorithms for computational geometry problems that are too large to be solved in internal memory. We use these techniques to develop optimal and practical algorithms for a number of important large-scale problems. We discuss our algorithms primarily in the context of single processor/single disk machines, a domain in which they are not only the first known optimal results but also of tremendous practical value. Our methods also produce the first known optimal algorithms for a wide range of two-level and hierarchical multilevel memory models, including parallel models. The algorithms are optimal both in terms of I/O cost and internal computation.

1 Introduction

Input/Output (I/O) communication between fast internal memory and slower secondary storage is the bottleneck in many large-scale information-processing applications, and its relative significance is increasing as parallel computing gains popularity. In this paper we consider the important application area of computational geometry and develop several paradigms for optimal geometric computation using secondary storage.

*This research was supported in part by the National Science Foundation under Grants CCR-9003299 and IRI-9116843, and by the NSF and DARPA under Grant CCR-8908092, email: goodrich@cs.jhu.edu

†This research was supported by the National Science Council under grant NSC81-0408-E-194-508, Republic of China, email: jjt@cs.ccu.edu.tw

‡This research was supported in part by Army Research Office grant DAAL03-91-G-0035 and by National Science Foundation grant DMR-9217290. Author is currently visiting Duke University, email: dev@cs.duke.edu

§This research was supported in part by National Science Foundation grant CCR-9007851 and by Army Research Office grant DAAL03-91-G-0035, email: jsv@cs.duke.edu

Large-scale problems involving geometric data are ubiquitous in spatial databases [24,32,33], geographic information systems (GIS) [10,24,32], constraint logic programming [19,20], object oriented databases [39], statistics, virtual reality systems, and computer graphics [32]. As an example, NASA's soon-to-be petabyte-sized databases are expected to facilitate a variety of complex geometric queries [10]. Important operations on geometric data include range queries, constructing convex hulls, nearest neighbor calculations, finding intersections, and ray tracing, to name but a few.

1.1 Our I/O model

In I/O systems, data are usually transferred in units of *blocks*, which may consist of several kilobytes. This blocking takes advantage of the fact that the seek time is usually much longer than the time needed to transfer a record of data once the disk read/write head is in place. An increasingly popular way to get further speedup is to use many disk drives and/or many CPUs working in parallel [12,13,18,25,28,36]. We model such systems, examples of which are shown in Figure 1, using the following four parameters:

- M = # items that can fit in internal memory;
- B = # records per block;
- P = # CPUs (internal processors);
- D = # disk drives.

For the problems we consider, we define three general parameters:

- N = # items or updates in the problem instance;
- K = # query operations in the problem instance;
- T = # items in the solution to the problem.

We will assume that $M < N$, $1 \leq P \leq M/\log M$, and $1 \leq DB \leq M/2$. The measures of performance

Figure 1: (a) The parallel disk model. Each of the D disks can simultaneously transfer B records to and from internal memory in a single I/O. The internal memory can store $M \geq DB$ records. (b) Multiprocessor generalization of the I/O model in (a), in which each of the $P = D$ internal processors controls one disk and has an internal memory of size M/P . The P processors are connected by some topology such as a hypercube or an EREW PRAM and their memories collectively have size M .

that we would like to minimize simultaneously are the number of I/Os and the internal computation time. The relevant terms that enter the formulæ for the I/O bounds are often in units of blocks, such as N/B , M/B , and so on. For that reason we define the following shorthand notation:

$$\nu = \frac{N}{B}, \quad \mu = \frac{M}{B}, \quad \kappa = \frac{K}{B}, \quad \tau = \frac{T}{B}.$$

In order to get across our techniques in the minimum space, we illustrate our results in this paper for the special case of the I/O model in which $P = 1$ and $D = 1$. Even in this simplified model, our results are significant since $P = 1$ and $D = 1$ accurately model the vast majority of I/O systems currently installed and being produced, and no I/O optimal algorithms were previously known for the problems we discuss. Our results are optimal in the parallel I/O model for $P \geq 1$ and $D \geq 1$ and in the parallel hierarchy models [27,38]. In particular, in the parallel I/O model, using P processors reduces the internal computation time by a factor of P and using D disks reduces the number of I/O steps by a factor of D . This generalization is discussed in Section 7 and in greater detail in the full version of this paper.

1.2 Our results

In this paper we present a number of general techniques for designing external-memory algorithms for computational geometry problems. These techniques include the following:

- *distribution sweeping*: a generic method for externalizing plane-sweep algorithms;
- *persistent B-trees*: an off-line method for constructing an optimal-space persistent version of the B-tree data structure. For batched problems this gives a factor of B improvement over the generic persistence techniques of Driscoll *et al.* [11];
- *batch filtering*: a general method for performing K simultaneous external-memory searches in data structures that can be modeled as planar layered dags and in certain fractional cascaded data structures;
- *on-line filtering*: A technique based on the work of Tamassia and Vitter [35] that allows I/O optimal on-line queries in fractional cascaded data structures based on balanced binary trees.
- *external marriage-before-conquest*: an external-memory analog to the well-known technique of Kirkpatrick and Seidel [22] for performing output-sensitive hull constructions.

We apply these techniques to derive optimal external-memory algorithms for the following fundamental problems in computational geometry: computing the pairwise intersection of N orthogonal segments, answering K range queries on N points, constructing the 2-d and 3-d convex hull of N points, performing K point location queries in a planar subdivision of size N , finding all nearest neighbors for a set of N points in the plane, finding the pairwise intersections of N rectangles, computing the measure of the union of N rectangles, computing the visibility of N segments from a point, performing K ray-shooting queries in CSG models of size N , and several geometric dominance problems. Our results are summarized in the following theorem, individual parts of which are discussed in the remaining sections of the paper.

Theorem 1.1: *Each of the problems mentioned in the preceding paragraph can be solved in external memory using $O((\nu + \kappa) \log_{\mu} \nu + \tau)$ I/Os. If D disks are used in parallel, the number of I/Os required can be reduced by a factor of D .*

For problems in which there are no queries as part of the problem instance, we use $K = 0$ (and thus $\kappa = 0$); if the output (solution) size is fixed, we use $T = 1$ (and thus $\tau = 1/B = o(1)$).

2 Distribution sweeping

The well-known *plane sweep* paradigm [30] is a powerful approach for developing computational geometry algorithms that are efficient in terms of internal computation. In this section we develop a new plane sweep approach that for the first time achieves optimal I/O performance (and a subsequent improvement in practice) for a large number of large-scale off-line problems in computational geometry.

A typical internal memory plane sweep algorithm proceeds by examining the geometric objects that make up its input (e.g., points, line segments, vertices of polygons, etc.) in sequence along a given dimension. As the scan proceeds, a dynamic data structure, typically a search tree, is maintained. Each time a new input object is processed, either an update, a query, or both are performed in the dynamic data structure. If each object results in only a constant number of updates and/or queries, each of which can be performed in time logarithmic in the number of objects, then the total running time of the algorithm is $O((N + K) \log N + T)$. For problems to which this technique is typically applied, this performance is optimal. An example of this is the standard plane sweep algorithm for orthogonal segment intersection, where the dynamic data structure is an interval tree [30].

An obvious way of implementing algorithms of this type in secondary memory is to replace the dynamic search tree with a dynamic B-tree [6,9]. Unfortunately, this requires $\Theta((N + K) \log_\mu \nu) = \Theta(B(\nu + \kappa) \log_\mu \nu)$ I/O operations in the worst case, which is prohibitive. Previous work using lazy batched updates on the B-tree yielded algorithms with $O((\nu + \kappa) \log_2 \nu)$ I/Os [34].

Our new method uses an off-line top-down implementation of the sweep, which is based upon a novel application of the subdivision technique used in the “distribution sort” algorithms of [3,27,37]. The central idea is that we divide the input into $O(\mu)$ strips, each containing an equal number of input objects. We then scan down these strips simultaneously, looking for components of the solution involving interactions between objects among different strips. Once we have done this, we are left with the problem of determining the components of the solution involving interactions completely within the strips, which we can solve recursively. Because we sweep the strips and then distribute the data to recursive subproblems for solution, we refer to our technique as *distribution sweeping*.

In Sections 2.1 and 2.2 we discuss optimal distribution sweeping algorithms for orthogonal segment intersection and all nearest neighbors respectively. There

are many other problems that this technique can be applied to. These are mentioned in Section 2.3 and discussed at length in the full version of this paper.

2.1 Orthogonal segment intersection reporting

The orthogonal segment intersection problem is that of reporting all intersecting pairs from a set of N orthogonal line segments. This problem is important in graphics and VLSI design systems.

To solve this problem, we first use an optimal sorting algorithm, for example [27,37], to sort the endpoints of all segments into two lists, one sorted by x and the other by y . The list sorted by x is used to locate medians which we will use to split the input into $\lfloor \mu \rfloor$ vertical strips γ_i . The list sorted by y is used to perform the sweep, which moves from top to bottom. Points encountered during the sweep are processed as follows:

- If the top endpoint of a vertical segment is encountered, the segment is inserted into an *active list* A_i associated with the strip γ_i in which the segment lies, and later, when the bottom endpoint is encountered, the segment is deleted from A_i .
- If the endpoints of a horizontal segment R are encountered, we consider the strips that R passes completely through and report all the vertical segments in the active lists of those strips.

This process finds all intersections except those between vertical segments and horizontal segments or portions thereof that do not completely span vertical strips. These are found when the problem is solved recursively on each strip.

In constructing the recursive subproblems, R is only distributed to the two strips containing its endpoints, thus at each level of recursion each segment is represented only twice. This is the key to preventing the total size of the subproblems at a given level of recursion from exceeding the input size by more than a constant factor. Once the number of points in a recursive subproblem falls below M , we simply solve the problem in main memory.

Insertions and vertical segments can be processed efficiently using blocks. With the exception of deleting segments from active lists, the total number of I/Os performed by this method is optimal $O(\nu \log_\mu \nu + \tau)$, where $\tau = T/B$ and T is the number of intersections reported. If “vigilant” deletion is used to delete each segment as soon as the sweep line reaches the bottom endpoint, a nonoptimal $O(N) = O(B\nu)$ term is added to the I/O bound. Instead we use the following “lazy”

deletion approach: For each strip, we maintain A_i as a stack. When a new segment is inserted, we simply add it to the stack. We keep all but the B most recently added elements of this stack in blocks of size B in external memory. When we are asked to output the active list, we scan the entire stack, outputting the segments still current and removing the segments whose deletion time has passed. A simple amortization argument shows that this method achieves the bound of Theorem 1.1.

2.2 All nearest neighbors

Given a set S of N points in the plane, the all nearest neighbors problem is that of finding, for each point $p \in S$, a nearest neighbor $NN(p)$, i.e. a point in $S \setminus \{p\}$ at least as close to p as any other. This problem has many applications in answering basic proximity questions on sets of objects. In this section we show how the all nearest neighbors problem can be solved by distribution sweeping.

To find the nearest neighbors of all points we will solve two subproblems, one to find the nearest neighbor above each point and the other to find the nearest neighbor below it. These solutions can then be compared to find the actual nearest neighbor of each point. Without loss of generality we will discuss only the method of finding the nearest neighbor below each point.

Just as we did for orthogonal segment intersection, we initially sort the input points using an optimal sorting algorithm [27,37] and then divide the set of points into vertical strips $\gamma_1, \gamma_2, \dots, \gamma_{\mu/5}$.

For each point p in strip γ_i we define $FN_b(p)$ to be the closest point to p that is below p and outside γ_i . We perform a downward sweep in which we produce, for each point p , either the identity and location of $FN_b(p)$ or a certificate $C_b(p)$ that $NN(p)$ lies in the same strip γ_i as p . The details of how these neighbors and certificates are produced are presented in the full version of this paper. An important point is that just because a scan finds $FN_b(p)$ does not mean that $FN_b(p)$ is $NN(p)$; a recursive step may find a point within γ_i that is closer. Thus $FN_b(p)$ is stored with each p in order that a recursive step can compare solutions it finds with the best solution from any higher level.

As we scan, we maintain an active set A_i for each strip γ_i . A_i is the set of points in γ_i for which we do not yet have either a certificate or a definite answer as to the identity of $FN_b(p)$. The following lemma, due to Atallah and Tsay [5] bounds the size of A_i .

Lemma 2.1 [5]: *At all times during the sweep, $|A_i| \leq 4$ for all γ_i .*

Although there can be no more than a constant number (4) of points in A_i at a time, these are not required to be the last four points in γ_i that the sweep line passed. Nevertheless, since there are only a constant number, we can keep the four or fewer blocks of output in which they appear in main memory until certificates or neighbors are found for them, at which time those blocks can be reinserted into the output stream.

Recursion continues until subproblems are small enough to fit in main memory, at which point they are solved directly. This gives us the optimal running time $O(\nu \log_\mu \nu)$.

2.3 Other applications of distribution sweeping

Though space precludes a full exposition, the distribution sweeping method can be used to solve a number of other off-line problems in computational geometry that are traditionally solved by plane sweep techniques. The resulting algorithms use an optimal $O((\nu + \kappa) \log_\mu \nu + \tau)$ I/Os. Problems in this category include batched range queries, computing the visibility from a point in the plane, finding pairwise rectangle intersections, computing the measure of a union of rectangles, and the 3-d maxima problem. These problems are discussed in greater detail in the full version of this paper.

3 Persistent B-trees

The B-tree data structure [6,9] is a fundamental structure for maintaining a dynamically-changing dictionary in external memory. In some cases, however, it may be advantageous to be able to access previous versions of the data structure. Being able to access such previous versions is known as *persistence*, and there exist very general techniques for making most data structures persistent [11]. Persistence can be implemented either in an on-line fashion (i.e., where the tree updates are coming on-line) or in an off-line fashion (i.e., where one is given the sequence of tree updates in advance).

For the on-line case, the method of Driscoll *et al.* Lspace [11] can be applied to hysterical B-trees as described by Maier and Salveter [26]. Since it is on-line, this structure requires $O(N \log_\mu \nu)$ I/Os to construct, which is optimal in an on-line setting. Unfortunately, this is a factor of B away from optimal for the sort of

batch geometric problems we would like to consider. For these we need an off-line strategy that requires only $O(\nu \log_\mu \nu)$ I/Os. In the following section we describe just such a method.

3.1 Off-line persistence

In the off-line case we can build a persistent tree by the distribution sweep method. We slightly modify our application of distribution sweeping for this construction, however, in that we follow the recursive calls on the sequences of suboperations by a non-recursive “merge” step.

We begin by applying using the techniques of [3,37] to divide the set \mathcal{X} of elements mentioned in σ into s groups of size roughly N/s each, where $s = \lceil \sqrt{\mu} \rceil$. This, of course, divides σ into s subsequences, one for each group. We then recursively construct a persistent data structure for each subsequence. Each such recursive call returns a list of “roots” of s -way trees, each of which is marked with a time stamp that represents the index in σ when this root was created. We mark every s th element in each list as a “bridge” element and we merge these bridge elements into a single list \mathcal{Y} . We store pointers from each element $y \in \mathcal{Y}$ to all its bridge predecessors in the recursively-constructed lists. The list \mathcal{Y} , together with these pointers, defines the roots of the persistent structure. Since we only choose every s th element from each list as a bridge, it is easy to see that total space needed is $O(\nu)$ blocks, and the depth of the resulting (layered dag) persistent structure is $O(\log_\mu \nu)$.

A search in the past, say at time position i , begins by locating the root active for time i and searching down in the structure from there, always searching in nodes whose time stamp is the largest value $\leq i$. Performing only one such search would not be an efficient strategy, however, unless $s = \sqrt{\mu}$ is $O(B)$. Nevertheless, as we show in the next section, this is a very efficient data structure (e.g., for point location) if it is searched using the batched filtering technique.

4 Batch filtering

In this section we demonstrate how, for many query problems in computational geometry, we can represent a data structure of size N in ν disk blocks in such a way that K constant-sized output queries of the data structure can be answered in $O((\nu + \kappa) \log_\mu \nu)$ I/O operations. Because we represent the data structure as a dag through which the K queries filter down from source to sinks, we call this technique *batch filtering*.

Given a data structure that supports queries, we can often model the processing of a query as the

traversal of a decision dag isomorphic to the data structure. We begin at a source node in the dag, and at each node we visit, we make a decision based on the outcome of comparisons between the query value and some number d of values stored at the node. We then make a decision as to which of the node’s $O(d)$ children to visit next. This process continues until we reach a sink in the dag, at which point we report the outcome of the query.

By restricting the class of such dags we are willing to consider, we are able to prove the following lemma, which will serve as a building block for optimal algorithms to solve a number of important geometric problems.

Lemma 4.1: *Let $G = (V, E)$ be a planar layered decision dag with a single source such that the maximum out degree of any node is μ . Let the graph be represented in ν blocks, with the nodes ordered by level and the nodes within a level ordered from left to right. Let $N = |V|$ and let h be the height of G . We can filter K queries through G in $O(\nu + h\kappa)$ I/O operations.*

Proof Sketch: We traverse the levels one by one, sending all K inputs to the i th level before any are sent to the $(i + 1)$ st. We do this by maintaining two FIFO queues, one for the current level and one for the next level. Each queue is a left-to-right list of edges between its level and the next one. If less than B queries traverse an edge then they are explicitly stored in the queue. If B or more traverse the edge, then the queue contains a pointer to a linked list of blocks storing them. Since the graph is planar, there exists an efficiently blocked method of producing one queue from the previous queue. \square

Luckily, the restrictions imposed on the type of decision dags we can handle with batch filtering is not too severe. In particular, many computations use decision trees, which clearly constitute a special case of the lemma. Often these trees are binary, but we can divide a binary tree into layers of height $O(\log \mu)$ and then store each node on a layer boundary along with all its descendants in the layer below it as a single node with branching factor μ . This allows us to reduce h by a factor of $O(\log \mu)$ yet still meet the conditions of the lemma. We will see this approach used in solving subproblems of the 3-d convex hull problem in Section 6.3.

Another way of using batch filtering, which we discuss in Section 4.1, is by structuring more complicated decision dags as recursive constructions in order to get around the planarity restrictions of the lemma.

4.1 Application: multiple-point planar point location

Planar point location is one of the fundamental problems of computational geometry. In the version of the problem considered here, we are given a monotone planar decomposition having N vertices, and a series of K query points. For each query point, we are to return the identifier of the region in which it lies. In main memory, this problem can be solved in optimal time $O((N+K) \log N)$ using fractional cascading [7,8]; $O(N \log N)$ is spent on preprocessing and $O(K \log N)$ is needed to perform the queries.

We can apply the technique of Lemma 4.1 to the main tree, but the bridge pointers connecting the catalogs make the dag non-planar. To get around this, we note that as queries traverse the edges between nodes in the main tree, they are ordered by the catalog values they query. This ordering is established at the root of the data structure, where a μ -ary tree is used to locate the queries in the first catalog. By relying on this ordering, we can efficiently process the queries that arrive at each node of the main tree. The overall complexity of this technique is thereby maintained at $O((\nu + \kappa) \log_\mu \nu)$. Details of the construction appear in the full version of the paper.

5 On-line filtering

In the Section 4 we discussed batch filtering, a technique which allows the efficient processing of a batch of queries. In some applications, this approach is not suitable because the queries arrive one at a time and must be processed individually. If we use batch filtering with a batch of size $K = 1$ for each query, this would require $O(\nu + \log_\mu \nu)$ I/Os per query. More desirable is a method for preprocessing the data structure so that individual queries can be answered with an optimal $O(\log_B \nu)$ I/Os. In this section we briefly describe how this can be done with a modified version of a parallel fractional cascading technique of Tamassia and Vitter [35].

The method of Tamassia and Vitter [35] works with data structures whose underlying graphs are balanced binary trees. Preprocessing takes $O(N)$ work. Once this is done, individual queries can be answered on a p processor CREW PRAM in $O(\log_p N)$ time. The access patterns of the processors during the search are such that data is accessed in groups that can be conveniently put into blocks for I/O purposes. This technique, details of which are presented in the full version of this paper, allows us to effectively replace a single step of p processors by one I/O, thereby giving us a

search complexity of $O(\log_B \nu)$ I/Os.

6 Convex hull algorithms

The convex hull problem is that of computing the smallest convex polytope completely enclosing a set of N points in d -dimensional space. This problem has important applications ranging from statistics to graphics to metallurgy. In this section we will examine the problem in external memory for two and three dimensions. The three-dimensional case is particularly interesting because of the number of two-dimensional geometric structures closely related to it, such as Voronoi diagrams and Delaunay triangulations. In fact, by well-known reductions [17], our 3-d convex hull algorithm immediately gives external-memory algorithms for planar Voronoi diagrams and Delaunay triangulations with the same I/O performance.

In main memory the lower bound for computing the convex hull of N points in dimension $d = 2$ and $d = 3$ is $\Omega(N \log N)$ [30]. In secondary memory, this bound becomes $\Omega(\nu \log_\mu \nu)$. In this section we give optimal algorithms that match this lower bound. For the two-dimensional case we show how to beat this lower bound for the case when the output size T is much smaller than N (in the extreme case, $T = O(1)$). We develop an output-sensitive algorithm based upon an external-memory version of the marriage-before-conquest paradigm of Kirkpatrick and Seidel [22].

Our 3-d convex hull is somewhat esoteric, so we also describe a simplified version that, although not optimal asymptotically, is simpler to implement and is faster for the vast majority of practical cases.

6.1 A worst-case optimal two-dimensional convex hull algorithm

For the two-dimensional case, a number of main memory algorithms are known that operate in optimal time $O(N \log N)$ [30]. A simple way to solve the problem optimally in external memory is to modify one of the main memory approaches, namely Graham's scan [16]. Graham's scan requires that we sort the points, which can be done in $O(\nu \log_\mu \nu)$ I/O operations, and then scan linearly through them, at times backtracking, but only over each input point at most once. Clearly this scanning stage can be accomplished in $O(\nu)$ I/O operations, so the overall complexity of the algorithm is $O(\nu \log_\mu \nu)$.

6.2 An output-sensitive two-dimensional convex hull algorithm

If the output size T is significantly smaller than N (for example, T can be $O(1)$) then we can do better than the Graham scan approach. In this section we show how to construct a two-dimensional convex hull using a number of I/Os that is output-size sensitive in a stronger sense than any of the algorithms discussed thus far. Note that when $T = o(N)$, we actually do better than Theorem 1.1 indicates. Our results are optimal, as stated in the following theorem.

Theorem 6.1: *The convex hull of a set S of N points in the plane having T extreme points can be computed in $O(\nu \log_\mu \tau)$ I/Os, which is optimal.*

We omit details in this preliminary version, but the main idea of our method is as follows: First, we observe that we may restrict our attention to the upper hull (i.e., edges with normals with positive y -components) without loss of generality. We use the techniques of [3,37] to divide the set of input points into $s = \lceil \sqrt{\mu} \rceil$ buckets divided by vertical lines. We then use an external-memory implementation of a method of Goodrich [15] for combining prune-and-search bridge finding [22] with the Graham scan technique [16] to find all the upper hull edges intersecting our given vertical lines. Our implementation uses $O(\nu)$ I/Os. Given these hull edges we may then recurse on any buckets that are not completely spanned by the hull edges we just discovered. Our analysis is based on the fact that at least one of the following two conditions holds in any such divide step:

1. Half of the s buckets are completely covered and eliminated from further consideration, reducing the number of points to consider by a constant fraction.
2. We discover $s/2$ hull edges that do not completely cover a bucket.

This implies that the total number of I/Os is $O(\nu \log_\mu \tau)$, which is optimal for any value of T .

6.3 Three-dimensional convex hulls

Even in main memory, space sweeping algorithms fail to solve the 3-d convex hull problem, and we must resort to more advanced divide and conquer approaches [29]. One idea is to use a plane to partition the points into equally sized sets, recursively construct the convex hull for each set, and then merge the recursive solutions together in linear time. Unfortunately, we know no way of implementing an algorithm of this type in secondary memory; the problem

Halfspace Intersection

Input: A set S of N halfspaces in 3-d space.

Output: The set of all halfspaces $h_i \in S$ whose bounding planes lie on the boundary of the intersection $\bigcap_{h_j \in S_0} h_j$

1. For $j = 1$ to $\Theta(\log_\mu \nu)$, take a random sample S_j of S , where $|S_j| = N^\epsilon$ for a constant $0 < \epsilon < 1$.
2. Recursively solve the halfspace intersection problem on each sample S_j , giving a set of solutions I_j .
3. Use polling ([31]) to estimate the size of the partition of $S - S_j$ that each sample solution I_j will induce. Let S_r be the sample whose solution I_r generates the smallest such partition.
4. For each cone C_i of I_r , compute R_i , the set of halfspaces in $S - S_r$ whose boundaries intersect C_i .
5. Eliminate redundant planes from each R_i , yielding R_i^* .
6. Recursively solve the halfspace intersection problem on each set R_i^* .

Algorithm 6.1: An algorithm for computing the 3-d convex hull of a set of points.

is that we cannot adequately anticipate all possible paths through the recursive subsolutions that might be traversed during the merging phase. Another obstacle is that we need to be able to merge $O(\mu^\epsilon)$ recursive solutions in linear time, rather than just two. If we use any fewer, then the depth of the recursion will not be small enough to give us an optimal algorithm.

In order to get around the problems associated with a merging approach, we use a novel formulation of the distribution method. We consider the dual of the convex hull problem, namely that of computing the intersection of N half spaces all of which contain the origin [30]. Once we are dealing with the dual problem, we can use a distribution based approach along the lines of that proposed by Reif and Sen for computing 3-d convex hulls in parallel [31].

Let S be a set of N halfspaces all of which contain the origin. Let the boundary of halfspace $h_i \in S$ be denoted P_i . Suppose we have a subset $S_0 \subset S$ such that $|S_0| = N^\epsilon$. Let $I_0 = \bigcap_{h_j \in S_0} h_j$. A face of I_0 might have up to N^ϵ edges. We can reduce this complexity by triangulating each face, which can be done by sorting the vertices of I_0 along a vector not perpendicular to any face and then sweeping a plane along this sorted order. By Euler's law the size of the resulting set of faces is at most $O(N^\epsilon)$. We can now decompose I_0 into $O(N^\epsilon)$ cones C_i , each of which has one of these faces as a base and the origin as an apex. An obvious way of distributing the halfspaces into subproblems is to create a subproblem for each cone C_i consisting of finding the intersection of all halfspaces

Figure 3: Parallel multilevel memory hierarchies. The H hierarchies (of any of the types listed in Figure 2) have their base levels connected by H interconnected processors.

$h_j \in S \setminus S_0$ whose bounding planes P_j intersect C_i . Unfortunately, a given P_j may intersect many cones, so it is not clear that we can continue to work through the $O(\log \log N)$ required levels of recursion without causing a very large blow up in the total size of the subproblems. Luckily, using a form of random sampling called polling and eliminating redundant planes from within a cone prior to recursion [31], we can with high probability get around this problem. (In this discus-

sion, the phrase “with high probability” means with probability $1 - N^{-\alpha}$, for some constant α .)

Algorithm 6.1 is the resulting distribution algorithm for computing the intersection of all $h_i \in S$. Step 1 can be completed with $O(\nu \log_\mu \nu)$ I/Os by making a linear pass through S for each sample, as suggested by Knuth [23]. Step 2 consists of recursive calls that will be considered later. In Step 3 we decompose each S_j into cones using a plane sweep. This takes $O((|S_j|/B) \log_\mu (|S_j|/B))$ I/Os. We then take a random sample from $S - S_j$ for each S_j . This takes $O(\nu \log_\mu \nu)$ I/Os. Finally, we solve a tree structured point location problem on all elements of the sample. This is done by batch filtering as described in Section 4. The number of I/O operations needed by Step 4 is $O(\frac{r}{B} \log_\mu \frac{r}{B})$, where $r = \sum_i |R_i|$. In Step 5, redundant planes are eliminated using a variant of the 3-d maxima algorithm from Section 2 and a 2-d convex hull algorithm. Both require $O(\frac{r}{B} \log_\mu \frac{r}{B})$ I/O operations. Finally, Step 6 recursively solves the subproblems.

By methods analogous to the approach of Reif and Sen [31] for the parallel case, we can develop the following recurrence for the running time of our algorithm:

$$T(N) = O(\nu \log_\mu \nu) + T(N^\epsilon) \log_\mu \nu + \sum_i T(|R_i|).$$

The first term on the right-hand side is the I/O cost for sampling and partitioning, the second term is the I/O cost for sorting the samples, and the last term is for the recursive calls. In the recurrence the $|S_i|$ terms are actually random variables. It suffices to use Karp’s method for solving probabilistic recurrence relations

[21] to get the optimal solution $T(n) = O(\nu \log_\mu \nu)$ with high probability.

The distribution approach used here is different from those of the distribution sort algorithms for the various I/O and memory hierarchy models [3,27,37,38] and the distribution sweeping algorithms discussed in Section 2, but it has the same asymptotic I/O complexity.

If desired, the randomization in our algorithm can be removed by an external memory implementation of the technique in [14]. Details are omitted for brevity.

In the full version of this paper we demonstrate how, for problems of any reasonable practical size, we can improve upon this algorithm by using samples of size μ instead of N^ε . The result is an algorithm that has asymptotic I/O performance of $O(\nu \log_\mu^2 \nu)$, but is far simpler to implement than Algorithm 6.1 and will generally perform better in practice. The main reason for the increase in performance is that to do polling efficiently the algorithm requires $\varepsilon < 1/8$ (see [31]) and thus in most practical situations $\mu < N^\varepsilon$.

7 Parallel and multi-level extensions

Up to this point our discussion has centered on the special case where $D = 1$ and $P = 1$. As has been mentioned, even in this restricted case the results presented here have much practical importance. More significantly, the paradigms described in this paper continue to work even when parallelism is added and D and P increase. Furthermore, they can be made to work optimally on hierarchical models having more than two levels; these include the well known HMM [1], BT [2], and UMH [4] (pictured in Figure 2), and their parallelizations [27,38] (pictured in Figure 3).

Details of the algorithms for these models are discussed in the full version of this paper. To a large extent they are based on modified versions of two of the main paradigms discussed above, namely distribution sweeping and batch filtering. We can also rely on the many-way divide-and-conquer approach of Atallah and Tsay [5], which can be extended to the I/O model. To implement distribution sweeping in these models we take advantage of the practical and optimal deterministic distribution techniques recently developed by Nodine and Vitter [27] for sorting. To implement batch filtering, we can use disk striping [28].

8 Conclusion

We have given a number of paradigms for external-memory computational geometry that yield the first

known I/O optimal algorithms for several interesting large-scale problems in computational geometry. Because they are simple and practical on most common systems ($P = 1, D = 1$) as well as on the parallel I/O systems likely to replace them in the not too distant future ($P \geq 1, D \geq 1$), the methods are likely to gain widespread use.

Nevertheless, there are a number of interesting problems that remain open:

- Is there a data structure for 2-d on-line range queries that achieves $O(\log_B \nu + \tau)$ I/Os for updates and range queries using $O(\nu)$ blocks of space? (The off-line version of the problem is solved optimally in this paper.) Updates and three-sided range queries can be handled by metablock trees [20] in $O(\log_B \nu + \log B + \tau)$ I/Os using $O(\nu)$ space. Two-sided range queries anchored on the diagonal can be done in $O(\log_B \nu + \tau)$ I/Os per query and $O(\log_B \nu + (\log_B \nu)^2/B)$ I/Os per (semidynamic) insertion [20].
- Can an N -vertex polygon be triangulated using $O(N/B)$ I/Os? Under what conditions?
- Can we find all intersecting pairs of N non-orthogonal segments using $O(\nu \log_\mu \nu + \tau)$ I/Os?

Acknowledgments

Figures 1, 2, and 3 were borrowed from [27] with our appreciation.

References

- [1] A. Aggarwal, B. Alpern, A. K. Chandra, and M. Snir, "A Model for Hierarchical Memory," *Proc. 19th ACM STOC* (1987), 305–314.
- [2] A. Aggarwal, A. Chandra, and M. Snir, "Hierarchical Memory with Block Transfer," *Proc. 28th IEEE FOCS* (1987), 204–216.
- [3] A. Aggarwal and J. S. Vitter, "The Input/Output Complexity of Sorting and Related Problems," *Comm. ACM* 31 (1988), 1116–1127.
- [4] B. Alpern, L. Carter, E. Feig, and T. Selker, "The Uniform Memory Hierarchy Model of Computation," *Proc. 31st IEEE FOCS* (1990), To appear in *Algorithmica*.
- [5] M. J. Atallah and J. -J. Tsay, "On the Parallel-Decomposability of Geometric Problems," *Algorithmica* 8 (1992), 209–231.
- [6] R. Bayer and E. McCreight, "Organization of Large Ordered Indexes," *Acta Inform.* 1 (1972), 173–189.
- [7] B. Chazelle and L. J. Guibas, "Fractional Cascading: I. A Data Structuring Technique," *Algorithmica* 1 (1986), 133–162.

- [8] B. Chazelle and L. J. Guibas, "Fractional Cascading: II. Applications," *Algorithmica* 1 (1986), 163–191.
- [9] D. Comer, "The Ubiquitous B-tree," *Comput. Surveys* 11 (1979), 121–137.
- [10] R. F. Crompt, "An Intelligent Information Fusion System for Handling the Archiving and Querying of Terabyte-sized Spatial Databases," in *Report on the Workshop on Data and Image Compression Needs and Uses in the Scientific Community*, S. R. Tate, ed., CESDIS Technical Report Series #TR-93-99, CESDIS, 1993, 75–84.
- [11] J. R. Driscoll, N. Sarnak, D. D. Sleator, and R. E. Tarjan, "Making Data Structures Persistent," *J. Comput. System Sci.* 38 (1989), 86–124.
- [12] G. Gibson, L. Hellerstein, R. M. Karp, R. H. Katz, and D. A. Patterson, "Coding Techniques for Handling Failures in Large Disk Arrays," U. C. Berkeley, UCB/CSD 88/477, December 1988.
- [13] D. Gifford and A. Spector, "The TWA Reservation System," *Comm. ACM* 27 (July 1984), 650–665.
- [14] M. H. Goodrich, "Geometric Partitioning Made Easier, Even in Parallel," *Proc. 9th ACM Comp. Geo.* (1993).
- [15] M. T. Goodrich, "Constructing the Convex Hull of a Partially Sorted Set of Points," *Computational Geometry: Theory and Applications* 2 (1993), 267–278.
- [16] R. L. Graham, "An Efficient Algorithm for Determining the Convex Hull of a Finite Planar Set," *Inform. Process. Lett.* 1 (1972), 132–133.
- [17] L. J. Guibas and J. Stolfi, "Primitives for the Manipulation of General Subdivisions and the Computation of Voronoi Diagrams," *ACM Trans. Graphics* 4 (1985), 74–123.
- [18] W. Jilke, "Disk Array Mass Storage Systems: The New Opportunity," Amperif Corporation, September 1986.
- [19] P. C. Kanellakis, G. M. Kuper, and P. Z. Revesz, "Constraint Query Languages," *Proc. 9th ACM PODS* (1990), 299–313.
- [20] P. C. Kanellakis, S. Ramaswamy, D. E. Vengroff, and J. S. Vitter, "Indexing for Data Models with Constraints and Classes," *Proc. 12th ACM PODS* (1993), 233–243.
- [21] R. M. Karp, "Probabilistic Recurrence Relations," *Proc. 23rd ACM STOC* (1991), 190–197.
- [22] D. G. Kirkpatrick and R. Seidel, "The Ultimate Planar Convex Hull Algorithm?," *SIAM J. Comput.* 15 (1986), 287–299.
- [23] D. E. Knuth, *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*, Addison Wesley, Reading, MA, 1973.
- [24] R. Laurini and D. Thompson, *Fundamentals of Spatial Information Systems*, A.P.I.C. Series, Academic Press, New York, NY, 1992.
- [25] N. B. Maginnis, "Store More, Spend Less: Mid-Range Options Around," *Computerworld* (November 16, 1987), 71–82.
- [26] D. Maier and S. C. Salveter, "Hysterical B-trees," *Inform. Process. Lett.* 12 (1981), 199–202.
- [27] M. H. Nodine and J. S. Vitter, "Deterministic Distribution Sort in Shared and Distributed Memory Multiprocessors," *Proc. 5th ACM SPAA* (1993), 120–129.
- [28] D. A. Patterson, G. Gibson, and R. H. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)," *Proc. 1988 ACM SIGMOD* (1988), 109–116.
- [29] F. P. Preparata and S. J. Hong, "Convex Hulls of Finite Sets of Points in Two and Three Dimensions," *Comm. ACM* 20 (1977), 87–93.
- [30] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, New York–Heidelberg–Berlin, 1985.
- [31] J. R. Reif and S. Sen, "Optimal Parallel Randomized Algorithms for Three-Dimensional Convex Hulls and Related Problems," *SIAM J. Comput.* 21 (1992), 466–485.
- [32] H. Samet, *Applications of Spatial Data Structures: Computer Graphics, Image Processing, and GIS*, Addison Wesley, Reading, MA, 1989.
- [33] H. Samet, *The Design and Analysis of Spatial Data Structures*, Addison Wesley, Reading, MA, 1989.
- [34] R. Tamassia and J. S. Vitter, mentioned in invited paper by J. S. Vitter, "Efficient Memory Access in Large-Scale Computation," *STACS '93*.
- [35] R. Tamassia and J. S. Vitter, "Optimal Cooperative Search in Fractional Cascaded Data Structures," *Proc. 2nd ACM SPAA* (1990), 307–316.
- [36] University of California at Berkeley, "Massive Information Storage, Management, and Use (NSF Institutional Infrastructure Proposal)," Technical Report No. UCB/CSD 89/493, January 1989.
- [37] J. S. Vitter and E. A. M. Shriver, "Algorithms for Parallel Memory I: Two-Level Memories," to appear in a special issue of *Algorithmica* on Large-Scale Memories, summary appears in "Optimal Disk I/O with Parallel Block Transfer" *Proc. 22nd ACM STOC* (1990), 159–169.
- [38] J. S. Vitter and E. A. M. Shriver, "Algorithms for Parallel Memory II: Hierarchical Multilevel Memories," to appear in a special issue of *Algorithmica* on Large-Scale Memories, summary appears in "Optimal Disk I/O with Parallel Block Transfer" *Proc. 22nd ACM STOC* (1990), 159–169.
- [39] S. B. Zdonik and D. Maier, eds., *Readings in Object-Oriented Database Systems*, Morgan Kaufman, 1990.