BRICS

Basic Research in Computer Science

BRICS RS-96-12 Arge et al.: External-Memory Algorithms for Processing Line Segments in GIS

External-Memory Algorithms for Processing Line Segments in Geographic Information Systems

Lars Arge
Darren E. Vengroff
Jeffrey S. Vitter

BRICS Report Series **RS-96-12**

ISSN 0909-0878 **May 1996**

Copyright © 1996, BRICS, Department of Computer Science
University of Aarhus. All rights reserved.

Reproduction of all or part of this work is permitted for educational or research use on condition that this copyright notice is included in any copy.

See back inner page for a list of recent publications in the BRICS Report Series. Copies may be obtained by contacting:

BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK - 8000 Aarhus C
Denmark

Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk

BRICS publications are in general accessible through WWW and anonymous FTP:

`http://www.brics.dk/`
`ftp ftp.brics.dk (cd pub/BRICS)`

External-Memory Algorithms for Processing Line Segments in Geographic Information Systems*

Lars Arge[†] Darren E. Vengroff[‡] Jeffrey S. Vitter[§]

May 1996

Abstract

In the design of algorithms for large-scale applications it is essential to consider the problem of minimizing I/O communication. Geographical information systems (GIS) are good examples of such large-scale applications as they frequently handle huge amounts of spatial data. In this paper we develop efficient new external-memory algorithms for a number of important problems involving line segments in the plane, including trapezoid decomposition, batched planar point location, triangulation, red-blue line segment intersection reporting, and general line segment intersection reporting. In GIS systems, the first three problems are useful for rendering and modeling, and the latter two are frequently used for overlaying maps and extracting information from them.

*An extended abstract version of this paper was presented at the Third Annual European Symposium on Algorithms (ESA '95).

[†]BRICS, Department of Computer Science, University of Aarhus, 8000 Aarhus C., Denmark. BRICS is an acronym for Basic Research in Computer Science, a Center of the Danish National Research Foundation. Supported in part by the ESPRIT II Basic Research Actions Program of the EC under contract No. 7141 (Project ALCOM II). This work was done while a Visiting Scholar at Duke University. Email: large@brics.dk.

[‡]Department of Computer Science, Brown University, Providence, RI 02912, USA. Supported in part by the U.S. Army Research Office under grant DAAH04–93–G–0076 and by the National Science Foundation under grant DMR–9217290. This work was done while a Visiting Scholar at Duke University. Email: dev@cs.brown.edu.

[§]Department of Computer Science, Duke University, Durham, NC 27708-0129, USA. Supported in part by the National Science Foundation under grant CCR–9007851 and by the U.S. Army Research Office under grant DAAH04–93–G–0076. Email: jsv@cs.duke.edu.

1 Introduction

The Input/Output communication between fast internal memory and slower external storage is the bottleneck in many large-scale applications. The significance of this bottleneck is increasing as internal computation gets faster, and especially as parallel computing gains popularity [27]. Currently, technological advances are increasing CPU speeds at an annual rate of 40–60% while disk transfer rates are only increasing by 7–10% annually [29]. Internal memory sizes are also increasing, but not nearly fast enough to meet the needs of important large-scale applications, and thus it is essential to consider the problem of minimizing I/O communication.

Geographical information systems (GIS) are a rich source of important problems that require good use of external-memory techniques. GIS systems are used for scientific applications such as environmental impact, wildlife repopulation, epidemiology analysis, and earthquake studies and for commercial applications such as market analysis, facility location, distribution planning, and mineral exploration [20]. In support of these applications, GIS systems store, manipulate, and search through enormous amounts of spatial data [16, 21, 30, 32]. NASA's EOS project GIS system [16], for example, is expected to manipulate petabytes (thousands of terabytes, or millions of gigabytes) of data!

Typical subproblems that need to be solved in GIS systems include point location, triangulating maps, generating contours from triangulated elevation data, and producing map overlays, all of which require manipulation of line segments. As an illustration, the computation of new scenes or maps from existing information—also called map overlaying—is an important GIS operation. Some existing software packages are completely based on this operation [32]. Given two thematic maps (piecewise linear maps with, e.g., indications of lakes, roads, pollution level), the problem is to compute a new map in which the thematic attributes of each location is a function of the thematic attributes of the corresponding locations in the two input maps. For example, the input maps could be a map of land utilization (farmland, forest, residential, lake), and a map of pollution levels. The map overlay operation could then be used to produce a new map of agricultural land where the degree of pollution is above a certain level. One of the main problems in map overlaying is "line-breaking," which can be abstracted as the red-blue line segment intersection problem.

In this paper, we present efficient external-memory algorithms for large-scale geometric problems involving collections of line segments in the plane, with applications to GIS systems. In particular, we address region decomposition problems such as trapezoid decomposition and triangulation, and line segment intersection problems such as the red-blue segment intersection problem and more general formulations.

1.1 The I/O Model of Computation

The primary feature of disks that we model is their extremely long access time relative to that of solid state random-access memory. In order to amortize this access time over a large amount of data, typical disks read or write large blocks of contiguous data at once. Our problems are modeled by the following parameters:

$$\begin{aligned} N &= \#\text{ of items in the problem instance;} \\ M &= \#\text{ of items that can fit into internal memory;} \\ B &= \#\text{ of items per disk block,} \end{aligned}$$

where $M < N$ and $1 \leq B \leq M/2$. Depending on the size of the data items, typical values for workstations and file servers in production today are on the order of $M = 10^6$ or 10^7 and $B = 10^3$. Large-scale problem instances can be in the range $N = 10^{10}$ to $N = 10^{12}$.

In order to study the performance of external-memory algorithms, we use the standard notion of I/O complexity [1, 34]. We define an *input/output operation* (or simply *I/O* for short) to be the process of reading or writing a block of data to or from the disk. The I/O complexity of an algorithm is simply the number of I/Os it performs. For example, reading all of the input data requires N/B I/Os. We will use the term *scanning* to describe the fundamental primitive of reading (or writing) all items in a set stored contiguously on external storage by reading (or writing) the blocks of the set in a sequential manner.

For the problems we consider we define two additional parameters:

$$\begin{aligned} K &= \#\text{ of queries in the problem instance;} \\ T &= \#\text{ of items in the problem solution.} \end{aligned}$$

Since each I/O can transmit B items simultaneously, it is convenient to introduce the following notation:

$$n = \frac{N}{B}, \qquad k = \frac{K}{B}, \qquad t = \frac{T}{B}, \qquad m = \frac{M}{B}.$$

We will say that an algorithm uses a linear number of I/O operations if it uses at most $O(n)$ I/Os to solve a problem of size N.

An increasingly popular approach to further increase the throughput of I/O systems is to use a number of disks in parallel. The number D of disks range up to 10^2 in current disk arrays. One method of using D disks in parallel is *disk striping* [34], in which the heads of the disks are moved synchronously, so that in a single I/O operation each disk reads or writes a block in the same location as each of the others. In terms of performance, disk striping has the effect of using a single large disk with block size $B' = DB$. Even though disk striping does not in theory achieve asymptotic optimality [34] when D is very large, it is often the method of choice in practice for using parallel disks, especially when D is moderately sized [33].

1.2 Previous Results in I/O-Efficient Computation

Early work on I/O algorithms concentrated on algorithms for sorting and permutation related problems [1, 15, 23, 24, 34]. External sorting requires $\Theta(n \log_m n)$ I/Os,[1] which is the external-memory equivalent of the well-known $\Theta(N \log N)$ time bound for sorting in internal memory. Work has also been done on matrix algebra and related problems arising in scientific computation [1, 34, 33]. More recently, researchers have designed external-memory algorithms for a number of problems in different areas, such as in computational geometry [19] and graph theoretic computation [5, 14]. In [6] a general connection between the comparison-complexity and the I/O complexity of a given problem is shown, and in [4] alternative solutions for some of the problems in [14] and [19] are derived by developing and using dynamic external-memory data structures.

1.3 Our Results

In this paper, we combine and modify in novel ways several of the previously known techniques for designing efficient algorithms for external memory. In particular we use the *distribution sweeping* and *batch filtering* paradigms of [19] and the *buffer tree* data structure of [4]. In addition we also develop a powerful new technique that can be regarded as a practical external-memory version of batched fractional cascading on an external-memory version of

[1] We define for convenience $\log_m n = \max\{1, (\log n)/\log m\}$.

a segment tree. This enables us to improve on existing external-memory algorithms as well as to develop new algorithms and thus partially answer some open problems posed in [19].

In Section 2 we introduce the *endpoint dominance problem*, which is a subproblem of *trapezoid decomposition*. We introduce an $O(n \log_m n)$-I/O algorithm to solve the endpoint dominance problem, and we use it to develop an algorithm with the same asymptotic I/O complexity for trapezoid decomposition, *planar point location*, *triangulation of simple polygons* and for the *segment sorting* problem. In Section 3 we give external-memory algorithms for line segment intersection problems. First we show how our segment sorting algorithm can be used to develop an $O(n \log_m n + t)$-I/O algorithm for *red-blue line segment intersection*, and then we discuss an $O((n+t) \log_m n)$-I/O algorithm for the *general segment intersection* problem.

Our results are summarized in Table 1. For all but the batched planar point location problem, no algorithms specifically designed for external memory were previously known. The batched planar point location algorithm that was previously known [19] only works when the planar subdivision is monotone, and the problems of triangulating a simple polygon and reporting intersections between other than orthogonal line segments are stated as open problems in [19].

For the sake of contrast, our results are also compared with modified internal-memory algorithms for the same problems. In most cases, these modified algorithms are plane-sweep algorithms modified to use B-tree-based dynamic data structures rather than binary tree-based dynamic data structures, following the example of a class of algorithms studied experimentally in [13]. Such modifications lead to algorithms using $O(N \log_B n)$ I/Os. For two of the algorithms the known optimal internal-memory algorithms [9, 10] are not plane-sweep algorithms and can therefore not be modified in this manner. It is difficult to analyze precisely how those algorithms perform in an I/O environment; however it is easy to realize that they use at least $\Omega(N)$ I/Os. The I/O bounds for algorithms based on B-trees have a logarithm of base B in the denominator rather than a logarithm of base m. But the most important difference between such algorithms and our results is the fact that the updates to the dynamic data structures are handled on an individual basis, which leads to an extra multiplicative factor of B in the I/O bound, which is very significant in practice.

As mentioned, the red-blue line segment intersection problem is of special interest because it is an abstraction of the important map-overlay problem,

Problem	I/O bound of new result	Result using mod. internal memory alg.
Endpoint dominance	$O(n \log_m n)$	$O(N \log_B n)$
Trapezoid decomposition	$O(n \log_m n)$	$O(N \log_B n)$
Batched planar point location	$O((n+k) \log_m n)$	
Triangulation	$O(n \log_m n)$	$\Omega(N)$
Segment sorting	$O(n \log_m n)$	$O(N \log_B n)$
Red-blue line segment intersection	$O(n \log_m n + t)$	$O(N \log_B n + t)$
Line segment intersection	$O((n+t) \log_m n)$	$\Omega(N)$

Figure 1: Summary of results.

which is the core of several vector-based GISs [2, 3, 26]. Although a time-optimal internal-memory algorithm for the general intersection problem exists [10], a number of simpler solutions have been presented for the red-blue problem [8, 11, 22, 26]. Two of these algorithms [11, 26] are not plane-sweep algorithms, but both sort segments of the same color in a preprocessing step with a plane-sweep algorithm. The authors of [26] claim that their algorithm will perform well with inadequate internal memory owing to the fact that data are mostly referenced sequentially. A closer look at the main algorithm reveals that it can be modified to use $O(n \log_2 n)$ I/Os in the I/O model, which is only a factor of $\log m$ from optimal. Unfortunately, the modified algorithm still needs $O(N \log_B n)$ I/Os to sort the segments.

In this paper we focus our attention on the single disk model. As described in Section 1.1, striping can be used to implement our algorithms on parallel disk systems with $D > 1$. Additionally, techniques from [23] and [25] can be used to extend many of our results to parallel disk systems. In the conference version of this paper we conjectured that all our results could be improved by the optimal factor of D on parallel disk systems with D disks, but it is still an open problem whether the required merges can be done efficiently enough to allow this.

2 The Endpoint Dominance Problem

In this section we consider the *endpoint dominance problem* (EPD) defined as follows: Given N non-intersecting line segments in the plane, find the segment directly above each endpoint of each segment.

EPD is a powerful tool for solving other important problems as we will illustrate in Section 2.1. As mentioned in the introduction a number of techniques for designing efficient I/O-efficient algorithms have been developed in recent years, including distribution sweeping, batch filtering [19] and buffer trees [4]. However, we do not seem to be able to efficiently solve EPD using these techniques directly. Section 2.2 briefly review some of the techniques and during that process we try to illustrate why they are inadequate for solving EPD. Fortunately, as we will demonstrate in Section 2.3, we are able to combine the existing techniques with several new ideas in order to develop an I/O-efficient algorithm for the problem, and thus for a number of other important problems.

2.1 Using EPD to solve other Problems

In this section we with three lemmas illustrate how an I/O-efficient solution to EPD can be used in the construction of I/O-efficient solutions to other problems.

Lemma 1 *If EPD can be solved in $O(n \log_m n)$ I/Os, then the trapezoid decomposition of N non-intersecting segments can be computed in $O(n \log_m n)$ I/Os.*

Proof: We solve two instances of EPD, one to find the segments directly above each segment endpoint and one (with all y coordinates negated) to find the segment directly below each endpoint—see Figure 2 for an example of this on a simple polygon. We then compute the locations of all $O(N)$ vertical trapezoid edges. This is done by scanning the output of the two EPD instances in $O(n)$ I/Os. To explicitly construct the trapezoids, we sort all trapezoid vertical segments by the IDs of the input segments they lie on, breaking ties by x coordinate. This takes $O(n \log_m n)$ I/Os. Finally, we scan this sorted list, in which we find the two vertical edges of each trapezoid in adjacent positions. The total amount of I/O used is thus $O(n \log_m n)$. □

Lemma 2 *If EPD can be solved in $O(n \log_m n)$ I/Os, then a simple polygon with N vertices can be triangulated in $O(n \log_m n)$ I/O operations.*

Proof: After computing the trapezoid decomposition of a simple polygon, the polygon can be triangulated in $O(n)$ I/Os using a slightly modified version of an algorithm from [18].

□

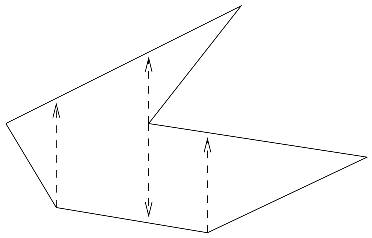

Figure 2: Using EPD to compute the trapezoid decomposition of a simple polygon.

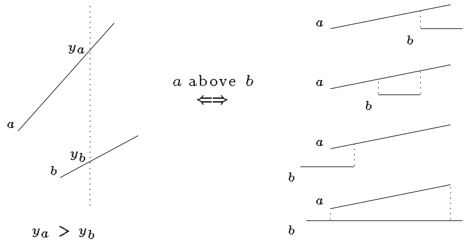

Figure 3: Comparing segments. Two segments can be related in four different ways.

We define a segment $\overline{AB}$ in the plane to be above another segment $\overline{CD}$ if we can intersect both $\overline{AB}$ and $\overline{CD}$ with the same vertical line l, such that the intersection between l and $\overline{AB}$ is above the intersection between l and $\overline{CD}$. Note that two segments are in comparable if they cannot be intersected with the same vertical line. Figure 3 demonstrates that if two segments are comparable then it is enough to consider vertical lines through the four endpoints to obtain their relation. The problem of sorting N non-intersecting segments in the plane is to extending the partial order defined in the above way to a total order. This problem will become important in the solution to the red-blue line segment intersection problem in section 3.1.

Lemma 3 *If EPD can be solved in $O(n \log_m n)$ I/Os, then a total ordering of N non-intersecting segments can be found in $O(n \log_m n)$ I/Os.*

Proof: We first solve EPD on the input segments augmented with the segment S_∞ with endpoints $(-\infty, \infty)$ and (∞, ∞). The existence of S_∞ ensures that all input segment endpoints are dominated by some segment. We define an aboveness relation $\searrow$ on elements of a non-intersecting set of segments S such that $\overline{AB} \searrow \overline{CD}$ *if and only if* either $(C, \overline{AB})$ or $(D, \overline{AB})$ is in the solution to EPD on S. Here $(A, \overline{BC})$ denotes that $\overline{BC}$ is the segment immediately above A. Similarly, we solve EPD with negated y coordinates and a special segment $S_{-\infty}$ to establish a belowness relation $\nearrow$. As discussed sorting the segments corresponds to extending the partial order defined by $\searrow$ and $\nearrow$ to a total order.

In order to obtain a total order we define a directed graph $G = (V, E)$ whose nodes consist of the input segments and the two extra segments S_∞ and $S_{-\infty}$. The edges correspond to elements of the relations $\searrow$ and $\nearrow$. For each pair of segments $\overline{AB}$ and $\overline{CD}$, there is an edge from $\overline{AB}$ to $\overline{CD}$

iff $\overline{CD} \searrow \overline{AB}$ or $\overline{AB} \nearrow \overline{CD}$. To sort the segments we simply have to topologically sort G. As G is a planar s,t-graph of size $O(N)$ this can be done in $O(n \log_m n)$ I/Os using an algorithm of [14]. $\square$

2.2 Buffer Trees and Distribution Sweeping

In internal memory EPD can be solved optimally with a simple plane-sweep algorithm; We sweep the plane from left to right with a vertical line, inserting a segment in a search tree when its left endpoint is reached and removing it again when the right endpoint is reached. For every endpoint we encounter we also do a search in the tree to identify the segment immediately above the point.

In [4] a number of external-memory data structures called buffer trees are developed for use in plane-sweep algorithms. Buffer trees are data structures that can support the processing of a batch of N updates and K queries on an initially empty dynamic data structure of elements *from a totally ordered set* in $O((n+k)\log_m n + t)$ I/Os. They can be used to implement plane-sweep algorithms in which the entire sequence of updates and queries is known in advance. The queries that such plane-sweep algorithms ask of their dynamic data structures need not be answered in any particular order; the only requirement on the queries is that they must all eventually be answered. Such problems are known as *batch dynamic problems* [17]. The plane-sweep algorithm for EPD sketched above can be stated as a batched dynamic problem. However, the requirement that the element stored in the buffer tree is taken from a totally ordered set is not fulfilled in the algorithm, as we do not know any total order of the segments. Actually, as demonstrated in Lemma 3, finding such an ordering is an important application of EPD. Therefore, we cannot use the buffer tree as the tree structure in the plane-sweep algorithm to get an I/O-efficient algorithm. For the other problems we are considering in this paper, the known internal-memory plane-sweep solutions cannot be stated as batched dynamic algorithms (since the updates depend on the queries) or else the elements involved are not totally ordered.

In [19] a powerful external memory version of the plane-sweep paradigm called distribution sweeping is introduced. Unfortunately, direct application of distribution sweeping appears insufficient to solve EPD. In order to illustrate why distribution sweeping is inadequate for the task at hand, let us briefly review how it works. We divide the plane into m vertical *slabs*, each of which contains $\Theta(n/m)$ input objects, for example points or line segment endpoints. We then sweep down vertically over all of the slabs to locate components of the solution that involve interaction of objects in different slabs or objects (such as line segments) that completely span one or more slabs. The choice of m slabs is to ensure that one block of data from each slab fits in main memory. To find components of the solution involving interaction between objects residing in the same slab, we recursively solve the problem in each slab. The recursion stops after $O(\log_m n/m) = O(\log_m n)$ levels when the subproblems are small enough to fit in internal memory. In order to get an $O(n \log_m n)$ algorithm one therefore need to be able to do one sweep in $O(n)$ I/Os. Normally this is accomplished by preprocessing the objects by using an optimal algorithm to sort them by y-coordinate. This e.g. allows one to avoid having to perform a sort before each recursive application of the technique, because as the objects are distributed to recursive subproblems their y ordering is retained. The reason that distribution sweeping fails for EPD is that there is no necessary relationship between the y ordering of endpoints of segments and their endpoint dominance relationship. In order to use distribution sweeping to get an optimal algorithm for EPD we instead need to sort the segments in a preprocessing step which leaves us with the same problem we encountered in trying to use buffer trees for EPD.

As know techniques fails to solve EPD optimally we are led instead to other approaches as discussed in the next section.

2.3 External-Memory Segment Trees

The *segment tree* [7, 28] is a well-known dynamic data structure used to store a set of segments in one dimension, such that given a query point all segments containing the point can be found efficiently. Such queries are called *stabbing queries*. An external-memory segment tree based on the approach in [4] is shown in Figure 4. The tree is perfectly balanced over the endpoints of the segments it represents and has branching factor $\sqrt{m/4}$. Each leaf represents $M/2$ consecutive segment endpoints. The first level of the tree partitions the data into $\sqrt{m/4}$ intervals σ_i—for illustrative reasons we call them slabs—separated by dotted lines on Figure 4. Multislabs are defined as contiguous ranges of slabs, such as for example $[\sigma_1, \sigma_4]$. There are $m/8 - \sqrt{m}/4$ multislabs. The key point is that the number of multislabs is a quadratic function of the branching factor. The reason why we choose the branching factor to be $\Theta(\sqrt{m})$ rather than $\Theta(m)$ is so that we have room in internal memory for a constant number of blocks for each of the $\Theta(m)$ multislabs. The smaller

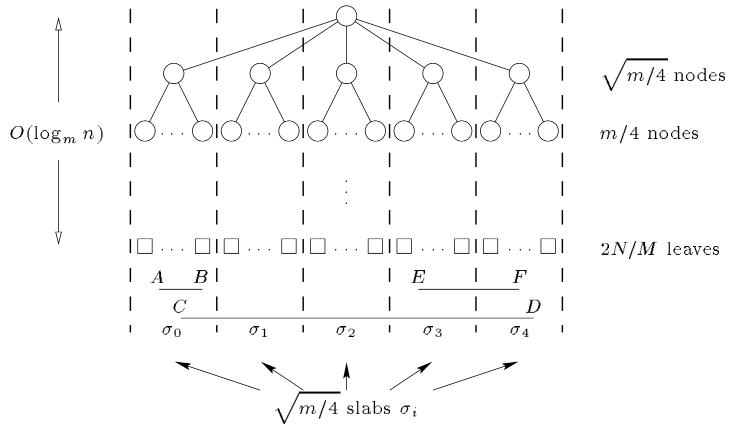

Figure 4: An external-memory segment tree based on a buffer tree over a set of N segments, three of which, $\overline{AB}$, $\overline{CD}$, and $\overline{EF}$, are shown.

branching factor at most about doubles the height of the tree.

Segments such as $\overline{CD}$ that completely span one or more slabs are called *long segments*. A copy of each long segment is stored in the largest multislab it spans. Thus, $\overline{CD}$ is stored in $[\sigma_1, \sigma_3]$. All segments that are not long are called *short segments* and are not stored in any multislab. Instead, they are passed down to lower levels of the tree where they may span recursively defined slabs and be stored. $\overline{AB}$ and $\overline{EF}$ are examples of short segments. The portions of long segments that do not completely span slabs are treated as small segments. There are at most two such synthetically generated short segments for each long segment and total space utilization is thus $O(n \log_m n)$ blocks.

To answer a stabbing query, we simply proceed down a path in the tree searching for the query value. At each node we encounter, we report all the long segments associated with each of the multislabs that span the query value.

Because of the size of the nodes and auxiliary multislab data, the buffer tree approach is inefficient for answering single queries. In batch dynamic environments, however, it can be used to develop optimal algorithms. In [4], techniques are developed for using external-memory segment trees in a batch dynamic environment such that inserting N segments in the tree and performing K queries requires $O((n + k) \log_m n + t)$ I/Os.

It is possible to come close to solving EPD by first constructing an external-memory segment tree over the projections of the segments onto the x-axis and then performing stabbing queries at the x coordinates of the endpoints of the segments. However, what we want is the single segment directly above each query point in the y dimension, as opposed to all segments it stabs. Fortunately, we are able to modify the external segment tree in order to efficiently answer a batch of this type of queries. The modification requires two significant improvements over existing techniques. First, as discussed in Section 2.3.1, we need to strengthen the definition of the structure, and the tree construction techniques of [4] must be modified in order to guarantee optimal performance when the structure is built. Second, as discussed in Section 2.3.2 the batched query algorithm must be augmented using techniques similar to fractional cascading [12].

2.3.1 Constructing Extended External Segment Trees

We will construct what we call an *extended external segment tree* using an approach based on distribution sweeping. When we are building an external segment tree on *non-intersecting* segments in the plane we can compare all segments in the same multislab just by comparing the order of their endpoints on one of the boundaries. An extended external segment tree is just an external segment tree as described in the last section built on non-intersecting segments, where the segments in each of the multislabs are sorted. Before discussing how to construct an extended external segment tree I/O-efficiently we will show a crucial property, namely that the segments stored in the multislab lists of a node in such a structure can be sorted efficiently. We will use this extensively in the rest of the paper. When we talk about sorting segments in the multislab lists of a node we imagine that they are "cut" to the slab boundaries, that is, that we have removed the part of the segments that are stored recursively further down the structure. Note that this might result in another total order on the segments than if we considered the whole segment.

Lemma 4 *The set of N segment stored in the multislab lists of an internal node of an extended external segment tree can be sorted in $O(n)$ I/O operations.*

Proof: We claim that we can construct a sorted list of the segments by repeatedly looking at the top segment in each of the multislabs, and selecting one of them to go to the sorted list.

To prove the claim, assume for the sake of contradiction that there exists a top segment s in one of the multislab lists which is above the top segment in all the other multislab lists it is comparable with, but which must be below a segment t in a total order. If this is the case there exist a series of segment $s_1, s_2 \ldots, s_i$ such that t is above s_1 which is above s_2 and so on ending with s_i being above s. But if s_i is above s then so is the top segment in the multislab list containing s_i contradicting the fact that s is above the top segment in all multislab lists it is comparable with.

As the number of multislab lists is $O(m)$ there is room for a block from each of them in internal memory. Thus the sorted list can be constructed in $O(n)$ I/Os by performing a standard external-memory merge of $O(m)$ sorted lists into a single sorted list. $\square$

In order to construct an extended external segment tree on N segments, we first use an optimal sorting algorithm to create a list of all the endpoints of the segments sorted by x-coordinate. This list is used during the whole algorithm to find the medians we use to split the interval associated with a given node into $\sqrt{m/4}$ vertical slabs. We now construct the $O(m)$ sorted multislab lists associated with the root in the following way: First we scan through the segments and distribute the long segments to the appropriate multislab list. This can be done in $O(n)$ I/Os because we have enough internal memory to hold a block of segments for each multislab list. Then we sort each of these lists individually with an optimal sorting algorithm. Finally, we recursively construct an extended external segment tree for each of the slabs. The process continues until the number of endpoints in the subproblems falls below $M/2$.

Unfortunately, this simple algorithm requires $O(n \log_m^2 n)$ I/Os, because we use $O(n \log_m n)$ I/Os to sort the multislab lists on each level of the recursion. To avoid this problem, we modify our algorithm to construct the multislab lists of a node not only from a list of segments but also from two other *sorted* lists of segments. One sorted list consists of segments that have one endpoint in the x range covered by the node under construction and one to the left thereof. The other sorted list is similar but contains segments entering the range from the right. Both lists are sorted by the y coordinate at which the segments enter the range of the node being constructed. In the construction of the structure the two sorted lists will contain segments which was already stored further up the tree. We begin to build a node just as we did before, by scanning through the unsorted list of segments, distributing the long segments to the appropriate multislab lists, and then sorting each multislab list. Next, we scan through the two sorted lists and distribute the long segments to the appropriate multislab lists. Segments will be taken from these lists in sorted order, and can thus be merged into the previously sorted multislab lists at no additional asymptotic cost. This completes the construction of the sorted multislab lists, and now we simply have to produce the input for the algorithm at each of the $\sqrt{m/4}$ children of the current node. The $\sqrt{m/4}$ unsorted lists are created by scanning through the list of segments as before, distributing the segments with both endpoints in the same slab to the list associated with the slab in question. The $2\sqrt{m/4}$ sorted lists of boundary crossing segments are constructed from the sorted multislab lists generated at the current level; First we use a linear number of I/Os sort the segments (Lemma 4) and then the $2\sqrt{m/4}$ lists can be constructed by scanning through the sorted list of segments, distributing the boundary crossing segments to the appropriate of $2\sqrt{m/4}$ lists. These lists will automatically be sorted.

In the above process all the distribution steps can be done in a linear number of I/Os, because the number of lists we distribute into always is $O(m)$, which means that we have enough internal memory to hold a block of segments for each output list. Thus, each level of recursion uses $O(n)$ I/Os plus the number of I/Os used on sorting. The following lemma then follows from the fact that each segment only ones is contained in a list that is sorted:

Lemma 5 *An extended external segment tree on N non-intersecting segments in the plane can be constructed in $O(n \log_m n)$ I/O operations.*

2.3.2 Filtering Queries Through an Extended Tree

Having constructed an extended external segment tree, we can now use it to find the segments directly above each of a series of K query points. In solving EPD, we have $K = 2N$, and the query points are the endpoints of the original segments. To find the segment directly above a query point p, we examine each node on the path from the root of the tree to the leaf containing p's x coordinate. At each such node, we find the segment directly above p by examining the sorted segment list associated with each multislab containing p. This segment can then be compared to the segment that is closest to the query point p so far, based on segments seen further up the tree, to see if it is the new globally closest segment. All K queries can be processed through the tree at once using a technique similar to batch filtering [19], in which *all* queries are pushed through a given level of the tree before moving on to the next level.

Unfortunately, the simple approach outlined in the preceding paragraph is not efficient. There are two problems that have to be dealt with. First, we must be able to look for a query point in many of the multislabs lists corresponding to a given node simultaneously. Second, searching for the position of a point in the sorted list associated with a particular multislab may require many I/Os, but as we are looking for an $O(n \log_m n)$ solution we are only allowed to use a linear number of I/Os to find the positions off *all* the query points. To solve the first problem, we will take advantage of the internal memory that is available to us. The second problem is solved with a notion similar to fractional cascading [11, 12, 31]. The idea behind fractional cascading on internal-memory segment trees is that instead of searching for the same element in a number of sorted lists of different nodes, we augment the list at a node with sample elements from lists at the node's children. We then build bridges between the augmented list and corresponding elements in the augments lists of the node's children. These bridges obviate the need for full searches in the lists at the children. We take a similar approach for our external-memory problem, except that we send sample elements from parents to children. Furthermore, we do not use explicit bridges.

As a first step towards a solution based on fractional cascading, we preprocess the extended external segment tree in the following way (corresponding to "building bridges"): For each internal node, starting with the root, we produce a set of sample segments. For each of the $\sqrt{m/4}$ slabs (*not* multislabs) we produce a list of samples of the segments in the multislab lists that span it. The sample list for a slab consists of every $(2\sqrt{m/4}\,)$th segment in the sorted list of segments that spans it, and we "cut" the segments to the slab boundaries. All the samples are produced by scanning through the sorted list of all segments in the node produced as in Lemma 4, distributing the relevant segments to the relevant sample lists. This can be done efficiently simply by maintaining $\sqrt{m/4}$ counters during the scan, counting how many segments so far have been seen spanning a given slab. For every slab we then augment the multislab lists of the corresponding child by merging the sampled list with the multislab list of the child that contains segments spanning the whole x-interval. This merging happens before we proceed to preprocessing the next level of the tree. At the lowest level of internal nodes, the sampled segments are passed down to the leaves.

We now prove a crucial lemma about the I/O complexity of the preprocessing steps and the space of the resulting data structure:

Lemma 6 *The preprocessing described above uses $O(n \log_m n)$ I/Os. After the preprocessing there are still $O(N)$ segments stored in the multi-lists on each level of the structure. Furthermore, each leaf contain less than M segments.*

Proof: Before any samples are passed down the tree, we have at most $2N$ segments represented at each level of the tree. Let N_i be the number of long segments, both original segments and segments sent down from the previous level, among all the nodes at level i of the tree after the preprocessing step. At the root, we have $N_0 \leq 2N$. We send at most $N_i/(2\sqrt{m/4}\,) \cdot \sqrt{m/4} = N_i/2$ segments down from level i to level $i+1$. Thus, $N_{i+1} \leq 2N + N_i/2$. By induction on i, we can show that for all i, $N_i = (4 - (1/2)^{i-1})\,N = O(N)$. From Lemma 4 and the fact that the number of multislab lists is $O(m)$—and thus that we can do a distribution or a merge step in a single pass of the data—it follows that each segment on a given level is read and written a constant number of times during the preprocessing phase. The number of I/Os used at level i of the tree is thus $O(n_i)$, where $n_i = N_i/B$. Since there are $O(\log_m n)$ levels, we in total use $O(n \log_m n)$ I/Os.

Before preprocessing, the number of segments stored in a node is less than the number of endpoints in the leaves below the node. To be precise, a leaf contains less than $M/2$ segments and a node i levels up the tree from a leaf contains less than $M/2 \cdot (\sqrt{m/4})^i$ segments. After preprocessing, the number of segments N_l in a leaf at level l in the tree must be $N_l \leq M/2 + \frac{N_{l-1}}{2\sqrt{m/4}}$, where N_{l-1} is the maximal number of segments in a node at level $l-1$; this is because at most every $(2\sqrt{m/4})$th of these segments are sent down to the leaf. Thus,

$$N_l \leq M/2 + \frac{M/2 \cdot \sqrt{m/4} + N_{l-2}/2\sqrt{m/4}}{2\sqrt{m/4}} \leq M/2 + M/4 + \frac{N_{l-2}}{(2\sqrt{m/4})^2}$$

and so on, which means that $N_l < M$. □

Having preprocessed the tree, we are now ready to filter the K query points through it. We assume without loss of generality that $K = O(N)$. If $K = \Omega(N)$ we break the queries into K/N groups of $K' = N$ queries and process each group individually. For EPD, we have $K = 2N$, so this grouping is not necessary. But as we will see later, grouping reduces the overall complexity of processing a batch of queries when K is very large. Since our fractional cascading construction is done backwards (sampled segments sent downwards), we filter queries from the leaves to the root rather than from the root to the leaves. To start off, we sort the K query points by their x coordinates in $O(k \log_m k)$ I/Os. We then scan the sorted list of query points to determine which leaf a given query belongs to. This produces an unsorted list of queries for each leaf as indicated on Figure 5a). Next we iterate through the leaves, and for each leaf find all dominating segments of the queries assigned to the leaf that are among the segments in the leaf. This is done by loading the entire set of segments stored at that leaf (which fits in memory according to Lemma 6), and then use an internal-memory algorithm to find the dominating segment for each query. As the total size of the data in all the leaves is $O(N)$, the total I/O complexity of the process is $O(k+n)$. In order to prepare for the general step of moving queries up the tree, we sort the queries that went into each leaf based on the order of the segments that we found to be directly above them, ending up in a situation as indicated in Figure 5b). This takes $O(k \log_m k)$ I/Os.

Each filtering step of the algorithm begins with a set of queries at a given level, partitioned by the nodes at that level and ordered within the nodes by the order of the segments found to be directly above them on the level. This is exactly what the output of the leaf processing was. The filtering step should produce a similar configuration on the next level up the tree. For one node this is indicated on Figure 5c). Remember that throughout the

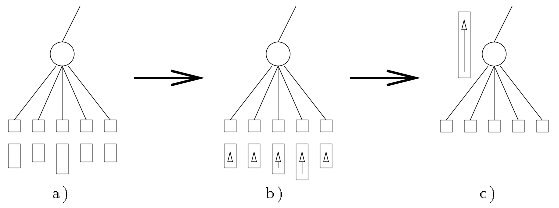

Figure 5: Filtering queries through the structure. An arrow in a list indicate that it is sorted.

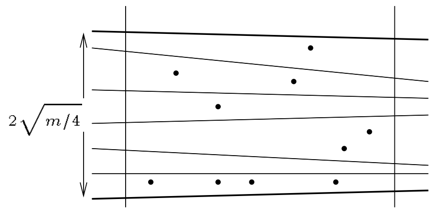

Figure 6: All queries between sampled segments (indicated by fat lines) must appear together in the list of queries for the slab.

algorithm we also keep track of the segment found to be closest to a given query point so far, such that when the root is reached we have found the dominating segment off all query points.

To perform one filtering step on a node we merge the list of queries associated with its children (slabs) and the node's multislab lists. The key property that allows us to find the dominating segments among the segments stored in the node in an I/O-efficient manner, and sort the queries accordingly, is that the list of queries associated with a son of the node cannot be to unsorted relative to their dominating segment in the node. This is indicated in Figure 6.

In order to produce for each slab a list of the queries in the slab, sorted according to dominating segment in the node, we again produce and scan through a sorted list of segments in the multislab list of the node, just like when we generated the samples that were passed down the tree in the the preprocessing phase. This time, however, instead of generating samples to pass down the tree, we insert a given segment in a list for each slab it spans. Thus if a segment completely spans four slabs it is inserted in four lists. If, during the scan, we encounter a segment which was sampled in slab s in the sampling phase then we stop the scan and process the queries in the list of queries for s between the sampled segment just encountered and the last sampled segment. As previously discussed these queries appear together in the sorted (according to dominating segment on the last level) list of queries for s. When this is done we clear the list of segments spanning s and continue the scan. The scan continues until all multislab segments have been processed. The crucial property is now that during the scan we can hold all the relevant segments in main memory because at no time during the scan do we store more than $2\sqrt{m/4}$ segments for each slab, that is, $2\sqrt{m/4} \cdot \sqrt{m/4} = m/2$ segments in total. Thus we can perform the scan, not counting the I/Os used to process the queries, in a linear number of I/Os.

To process the queries in a slab between two sampled segments we maintain $2\sqrt{m/4}$ output blocks, each of which corresponds to a segment between the two sampled segments. The block for a segment is for queries with the segment as dominating segment among the segments in the multislab list. As we read queries from the output of the child, we place them in the appropriate output block for the slab. If these output blocks become full, we write them back to disk. Once all queries between the two sampled segments have been processed, we concatenate the outputs associated with each of the segments between the samples. This results in a list of queries sorted according to dominating segment in the node, and this list is appended to an output list for the slab. All of the above is done in a number of I/Os linear in the number of queries processed.

When we finish the above process, we merge the sorted output query lists of all the slabs to produce the output of the current node in a linear number of I/Os.

As discussed above, once this process has reached the root, we have the correct answers to all queries. The total I/O complexity of the algorithm is given by the following theorem.

Theorem 1 *An extended external segment tree on N non-intersecting segments in the plane can be constructed, and K query points can be filtered through the structure in order to find the dominating segments for all these points, in $O((n+k)\log_m n)$ I/O operations.*

Proof: According to Lemma 5 and 6 construction and preprocessing together require $O(n \log_m n)$ I/Os.

Assuming $K \leq N$, sorting the K queries takes $O(n \log_m n)$ I/Os. Filtering the queries up one level in the tree takes $O(n)$ I/Os for the outer scan and $O(k)$ I/Os to process the queries. This occurs through $O(\log_m n)$ levels, giving an overall I/O complexity of $O(n \log_m n)$.

When $K > N$, we can break the problem into $K/N = k/n$ sets of N queries. Each set of queries can be answered as shown above in $O(n \log_m n)$ I/Os, giving a total I/O complexity of $O(k \log_m n)$. □

Theorem 1 immediately gives us the following bound for EPD, for which $K = 2N$.

Corollary 1 *The endpoint dominance problem can be solved in $O(n \log_m n)$ I/O operations.*

We then immediately get the following from Lemma 1, 2 and 3.

Corollary 2 *The trapezoid decomposition and the total order of N non-intersecting segments in the plane, as well as the triangulation of a simple polygon, can all be computed in $O(n \log_m n)$ I/O operations.*

It remains open whether a simple polygon can be triangulated in $O(n)$ I/Os when the input vertices are given by their order on the boundary of the polygon, which would match the linear internal-memory bound [9].

As a final direct application of our algorithm for EPD we consider the *multi-point planar point location problem.* This is the problem of reporting the location of K query points in a planar subdivision defined by N line segments. In [19] an $O((n+k)\log_m n)$-I/O algorithm for this problem is given for *monotone* subdivisions of the plane. Using Theorem 1 we can immediately extended the result to arbitrary planar subdivisions.

Lemma 7 *The multi-point planar point location problem can be solved using $O((n+k)\log_m n)$ I/O operations.*

3 Line Segment Intersection

In this section we design algorithms for line segment intersection reporting problems. In Section 3.1 we develop an I/O-efficient algorithm for the red-blue line segment intersection problem and in Section 3.2 we develop an algorithm for the general line segment intersection problem.

3.1 Red-Blue Line Segment Intersection

Using our ability to sort segments as described in Section 2, we can now overcome the problems in solving the red-blue line segment intersection problem with distribution sweeping. Given input sets S_r of non-intersecting red segments and S_b of non-intersecting blue segments, we construct two intermediate sets

$$\begin{aligned} T_r &= S_r \cup \bigcup_{(p,q)\in S_b} \{(p,p),(q,q)\} \\ T_b &= S_b \cup \bigcup_{(p,q)\in S_r} \{(p,p),(q,q)\} \end{aligned}$$

Each new set is the union of the input segments of one color and the endpoints of the segments of the other color (or rather zero length segments located at the endpoints). Both T_r and T_b are of size $O(|S_r|+|S_b|)=O(N)$. We sort both T_r and T_b using the algorithm from the previous section, and from now on assume they are sorted. This preprocessing sort takes $O(n\log_m n)$ I/Os.

We now locate intersections between the red and blue segments with a variant of distribution sweeping with a branching factor of $\sqrt{m}$. As discussed in Section 2.2, the structure of distribution sweeping is that we divide the plane into $\sqrt{m}$ slabs, not unlike the way the plane was divided into slabs to build an external segments tree in Section 2.3. We define *long segments* as those crossing one or more slabs and *short segments* as those completely contained in a slab. Furthermore, we shorten the long segments by "cutting" them at the right boundary of the slab that contain their left endpoint, and at the left boundary of the slab containing their right endpoint. This may produce up to two new short segments for each long segment, and below we show how to update T_r and T_b accordingly in $O(n)$ I/Os. We also show how to report all T_i intersections between the long segments of one color and the long and short segments of the other color in $O(n+t_i)$ I/Os. Next, we use one scan to partition the sets T_r and T_b into $\sqrt{m}$ parts, one for each slab, and we recursively solve the problem on the short segments contained in each slab to locate their intersections. Each original segment is represented at most twice at each level of recursion, thus the total problem size at each level of recursion remains $O(N)$ segments. Recursion continues through $O(\log_m n)$ levels until the subproblems are of size $O(M)$ and thus can be solved in internal memory. This gives us the following result:

Theorem 2 *The red-blue line segment intersection problem on N segments can be solved in $O(n \log_m n + t)$ I/O operations.*

Now, we simply have to fill in the details of how we process the segments on one level of the recursion. First, we consider how to insert the new points and segments generated when we cut a long segment at the slab boundaries into the sorted orders T_r and T_b. Consider a cut of a long red segment s into three parts. Changing T_r accordingly is easy, as we just need to insert the two new segments just before or after s in the total order. In order to insert all new red endpoints generated by cutting long red segments (which all lie on a slab boundary) in T_b, we first scan through T_r generating the points and distributing them to $\sqrt{m}$ lists, one for each boundary. The lists will automatically be sorted and therefore it is easy to merge them into T_r in a simple merge step. Altogether we update T_r and T_b in a $O(n)$ I/Os.

Next, we consider how intersections involving long segments are found. We divide the algorithm into two parts; reporting intersections between long and short segments of different colors and between long segments of different colors.

Because T_r and T_b are sorted, we can locate interactions between long and short segments using the distribution-sweeping algorithm used to solve the orthogonal segment intersection problem in [19]. We use the algorithm

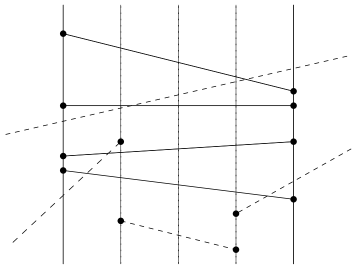

Figure 7: Long blue segments (dashed lines) can interact with multislab in three ways.

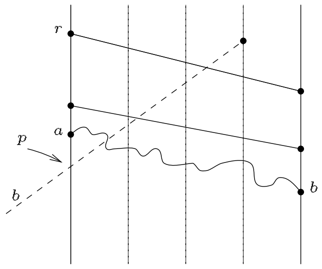

Figure 8: Proof of Lemma 8. The segment between a and b must intersect b.

twice and treat long segments of one color as horizontal segments and short segments of the other color as vertical segments. We sketch the algorithm for long red and blue short segments (details can be found in [19]); We sweep from top to bottom by scanning through the sorted list of red segments and blue endpoints T_r. When a top endpoint of a small blue segment is encountered, we insert the segment in an *active list* (a stack where we keep the last block in internal memory) associated with the slab containing the segment. When a long red segment is encountered we then scan through all the active lists associated with the slabs it completely spans. During this scan we know that every small blue segment in the list either is intersected by the red segment or will not be intersected by any of the following red segments (because we process the segments in sorted order), and can therefore be removed from the list. A simple amortization argument then shows that we use $O(n + t_i)$ I/Os to do this part of the algorithm.

Next we turn to the problem of reporting intersections between long segments of different colors. We define a multislab as in Section 2.3.1 to be a slab defined by two of the $\sqrt{m}$ boundaries. In order to report the intersections we scan through T_r and distribute the long red segments into the $O(m)$ multislabs. Next, we scan through the blue set T_b, and for each blue segment we report the intersections with the relevant long red segments. This is the same as reporting intersections with the appropriate red segments in each of the multislab lists. Now consider Figure 7. A long blue segments can "interact" with a multislab in three different ways. It can have one endpoint in the multislab, it can cross the multislab completely, or it can be totally contained in the multislab. First, let us concentrate on reporting intersections with red segments in multislabs for which the blue segment intersects the left boundary. Consider a blue segment b and a multislab m containing its right endpoint, and define y_p to be the y coordinate of a point p. We have the following:

Lemma 8 *If a blue segment b intersects the left boundary of a multislab at point p then all blue segments processed after b will intersect the same boundary at a point q below p. Let r be the left endpoint of a red segment in the multislab list. If $y_r \geq y_p$ and b intersects the red segment, then b intersects all red segments in the multislab list with left endpoints in the y-range $[y_p, y_r]$. The case $y_r \leq y_p$ is symmetric.*

Proof: The first part follows immediately from the fact that we process the segments in sorted order. Figure 8 demonstrates that the second part holds. □

Using this lemma we can now complete the design of the algorithm for our problem using a merging scheme. As discussed above, we process the blue segment in T_b one at a time and report intersections with red segments in multislabs list where the blue segment intersect the left boundary. For each such multislab list we do the following: We scan *forward* from the current position in the list until we find the first red segment s_r whose left endpoint lies below the intersection between the blue segment and the multislab boundary. Then we scan backward or forward as necessary in the multislab list in order to report intersections. Lemma 8 shows that the algorithm reports all intersections because all intersected segments lies consecutively above or belove s_r. Furthermore, it shows that we can use blocks efficiently such that we in total only scan through each multislabs list once without reporting intersections. Thus, our algorithm uses a total of $O(n + t_i)$ I/Os.

This takes care of the cases where the blue segment completely spans a multislab or where it has its right, and only the right, endpoint in the multislab. The case where the blue segment only has its left endpoint in the multislab can be handled analogously. The remaining case can be handled with the same algorithm, just by distributing the blue segments instead of the red segments, and then processing one long blue segment at a time. To summarize, we have shown how to perform one step of the distribution sweeping algorithm in $O(n + t_i)$ I/Os, and thus proven Theorem 2.

3.2 General Line Segment Intersection

The general line segment intersection problem cannot be solved by distribution sweeping as in the red-blue case, because the $\nearrow$ and $\searrow$ (Lemma 3) relations for sets of intersecting segments are not acyclic, and thus the preprocessing phase to sort the segments cannot be used to establish an ordering for distribution sweeping. However, as we show below, extended external segment trees can be used to establish enough order on the segments to make distribution sweeping possible. The general idea in our algorithm is to build an extended external segment tree on all the segments, and during this process to eliminate any inconsistencies that arise because of intersecting segments *on the fly*. This leads to a solution for the general problem that integrates all the elements of the red-blue algorithm into one algorithm. In this algorithm, intersections are reported both during the construction of an extended external segment tree and during the filtering of endpoints through the structure.

In order to develop the algorithm we need an external-memory priority queue [4]. Given m_p blocks of internal memory, N insert and delete-min operations can be performed on such a structure in $O(n \log_{m_p} n)$ I/Os. If we chose m_p to be m^c for some constant c $(0 < c < 1)$, we can perform the N operations using $O(n \log_m n)$ I/Os. In the construction of an extended external segment tree for general line segment intersection, we use two priority queues for each multislab. In order to have enough memory to do this, we reduce the fan-out of the extended segment tree from $\sqrt{m/4}$ to $(m/4)^{1/4}$. This does not change the asymptotic height of the tree, but it means that each node will have less than $\sqrt{m}/4$ multislabs. We chose m_p to be $\sqrt{m}$. Thus, with two priority queues per multislab, each node of the external segment tree still requires less than $m/2$ blocks of internal memory. Exactly what goes into the priority queues and how they are used will become clear as we describe the algorithm.

3.2.1 Constructing the Extended External Segment Tree

In the construction of an extended external segment tree in Section 2.3.1 we used the fact that the segments did not intersect in order to establish an ordering on them. The main idea in our algorithm is a mechanism for breaking *long segments* into smaller pieces every time we discover an intersection during construction of the multislab lists of a node. In doing so we manage to construct an extended segment tree with no intersections between long segments stored in the multislab lists of the same node.

In order to construct the extended external segment tree on the N (now possibly intersecting) segments, we as in Section 2.3.1 first in $O(n \log_m n)$ I/Os create a sorted list of all the endpoints of the segments. The list is sorted by x coordinate, and used during the whole algorithm to find the medians we use to split the interval associated with a node into $(m/4)^{1/4}$ vertical slabs. Recall that in Section 2.3.1 one node in the tree was constructed from three lists, one sorted list of segments for each of the two extreme boundaries and one unsorted list of segments. In order to create a node we start as in the non-intersecting case by scanning through the unsorted list of segments, distributing the long segments to the appropriate multislab lists. Next, we sort the multislab lists individually according to the left (or right) segment endpoint. Finally, we scan through the two sorted lists and distribute the segments from these lists. The corresponding multislab lists will automatically be sorted according to the endpoint on one of the boundaries.

Now we want to remove inconsistencies by removing intersections between long segments stored in the multislab lists. We start by removing intersections between segments stored in the same list. To do so we initialize two external priority queues for each of the multislabs, one for each boundary. Segments in these queues are sorted according to the order of the their endpoint on the boundary in question, and the queues are structured such that a delete-min operation returns the topmost segment. We process each of the multislab lists individually as follows: We scan through the list and check if any two consecutive segments intersect. Every time we detect an intersection we report it, remove one of the segment from the list, and break it at the intersection point as indicated on Figure 9. This creates two new segments.

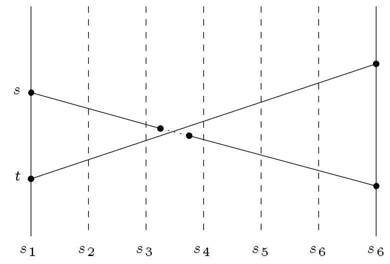

Figure 9: Breaking a segment.

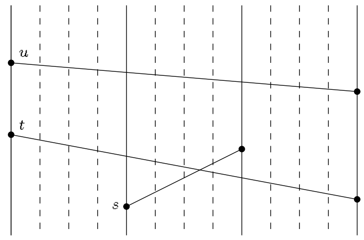

Figure 10: Proof of lemma 9.

If either one of them are long we insert it in both the priority queues corresponding to the appropriate multislab list. Any small segments that are created are inserted into a special list of segments which is distributed to the children of the current node along with normal small segments. The left part of s on Figure 9 between s_1 and s_3 is for example inserted in the queues for multislab $[s_1, s_3]$, and the part to the right of s_3 is inserted in the special list. It should be clear that after processing a multislab list in this way the remaining segments are non-intersecting (because every consecutive pair of segments are non-intersecting), and it will thus be consistently sorted. As we only scan through a multislab list ones the whole process can be done in a linear number of I/Os in the number of segments processed, plus the I/Os used to manipulate the priority queues.

Unfortunately, we still have inconsistencies in the node because segments in different multislab lists can intersect each other. Furthermore, the newly produced long segments in the priority queues can intersect each other as well as segments in the multislab lists. In order to remove the remaining intersections we need the following lemma.

Lemma 9 *If the minimal (top) segments of all the priority queues and the top segments of all the multislab lists are all non-intersecting, then the top-most of them is not intersected by any long segment in the queues or lists.*

Proof: First, consider the top segment in the two priority queues corresponding to the two boundaries of a single multislab. If these two segments do not intersect, then they must indeed be the same segment. Furthermore, no other segment in these queues can intersect this top segment. Now consider the top segment in the multislab list of the same multislab. As the two segments are non-intersecting one of them must be completely above the other. This segment is not intersected by any segment corresponding to the same multislab. Now consider this top segment in all the multislabs. Pick one of the top-most of these non intersecting segments and call it s. Now consider Figure 10. Assume that s is intersected by another segment t in one of the queues or multislab lists. By this assumption t is not the top segment in its multislab. Call the top segment in this multislab u. Because u does not intersect either t or s, and as it is on top of t, it also has to be on top of s. This contradicts the assumption that s is above all the top segments. □

Our algorithm for finding and removing intersections now proceeds as follows. We repeatedly look at the top segment in each of the priority queues and multislab lists. If any of these segments intersect, we report the intersection and break one of the segments as before. If none of the top segments intersect we know from Lemma 9 that the topmost segment is not intersected at all. This segment can then be removed and stored in a list that eventually becomes the final multislab list for the slab in question. When we have processed all segments in this way, we end up with $O(m)$ sorted multislab list of non-intersecting segments. We have enough internal memory to buffer a block from each of the lists involved in the process, so we only need a number of I/Os linear in the number of segments processed (original and newly produced ones), plus the number of I/Os used to manipulate the priority queues.

Finally, as in Section 2.3.1, we produce the input to the next level of recursion by distributing the relevant segments (remembering to include the newly produced small segments) to the relevant children. As before, this is done in a number of I/Os linear in the number of segments processed. We stop the recursion when the number of *original* endpoints in the subproblems fall below $M/4$.

If the total number of intersections discovered in the construction process is T then the number of new segments produced is $O(T)$, and thus the number of segments stored on each level of the structure is bounded by $O(N+T)$. As in Section 2.3.1 we can argue that each segment is only contained in one list being sorted and thus we use a total of $O((n+t)\log_m(n+t)) = O((n+t)\log_m n)$ I/Os to sort the segments. In constructing each node we only use a linear number of I/Os, plus the number of I/Os used on priority queue operations. Since the number of priority queue operations is $O(T)$, the total number of of I/Os we use to construct the whole structure is bounded by $O((n+t)\log_m n)$.

3.2.2 Filtering Queries Through the Structure

We have now constructed an extended external segment tree on the N segments, and in the process of doing so we have reported some of the intersections between them. The intersections that we still have to report must be between segments stored in different nodes. In fact intersections involving segments stored in a node v can only be with segments stored in nodes below v or in nodes on the path from v to the root. Therefore we will report all intersections if, for all nodes v, we report intersections between segments stored at v and segments stored in nodes below v. But in v segments stored in nodes below v must be similar to the segments we called small in the red-blue line segment algorithm. Thus, if in each node v we had a list of endpoints of segments stored in nodes below v, sorted according to the long segment in v immediately on top of them, we could report the remaining intersections with the algorithm that was used in section 3.1.

In order to report the remaining intersections we therefore preprocess the structure and filter the endpoints of the $O(N+T)$ segments stored in the structure through the structure as we did in section 2.3.2. At each node the filtering process constructs a list of endpoints below the node sorted according to dominating segment among the segments stored in the node. At each node we can then scan this list to collect the relevant endpoints and then report intersections with the algorithm used in Section 3.1. For all nodes on one level of the structure the cost is linear in the number of segments and endpoints processed, that is, $O(n+t)$ I/Os, plus a term linear in the number of new intersections reported.

Recall that the preprocessing of the structure in Section 2.3.2 consisted of a sampling of every $(2\sqrt{m/4})$th segment of every slab in a node, which was then augmented to the segments stored in the son corresponding to the slab. The process was done from the root towards the leaves. We will do the same preprocessing here, except that because we decreased the fanout to $(m/4)^{1/4}$ we only sample every $(2(m/4)^{1/4})$th segment in a slab. However, as we are building the structure on intersecting segments we should be careful not to introduce intersections between segments stored in the multislab lists of a node when augmenting the lists with sampled segments. Therefore we do the preprocessing while we are building the structure. Thus, in the construction process described in the previous section, after constructing the sorted multislab lists of a node, we sample every $(2(m/4)^{1/4})$th segment in each slab precisely as in Section 2.3.2. We then send these segments down to the next level together with the other "normal" segments that need to be recursively stored further down the tree. However, we want to make sure that the sampled segments are not broken, but stored on the next level of the structure. Otherwise we cannot I/O-efficiently filter the query points through the structure, as the sampled segments are stored on the next level to make sure that the points are not to unsorted relative to the segments stored in a node. Therefore we give the sampled segments a special mark and make sure that we only break unmarked segments. We can do so because two marked segments can never intersect, otherwise they would have been broken on the previous level.

By the same argument used in Section 2.3.2 to prove Lemma 6 we can prove that the augmentation of sampled segments does not asymptotically increase the number of segments stored on each level of the structure. Also all the sampling and augmentation work can be done in a linear number of I/Os on each level of the structure. This means that the number of I/Os used to construct the structure is kept at $O((n+t)\log_m n)$, even when the preprocessing is done as an integrated part of it.

After the construction and preprocessing we are ready to filter the $O(N+T)$ endpoints through the $O(\log_m n)$ levels of the structure. Recall by referring to Figure 5 that in order to do the filtering we first sort the points by x coordinate and distribute them among the leaves. Then for each leaf in turn we find the dominating segments of the points assigned to the leaf and sort the points accordingly. Finally, the points are repeatedly filtered one level up until they reach the root.

The sorting of the points by x coordinate can be done in $O((n+t)\log_m(n+t)) = O((n+t)\log_m n)$ I/Os. Also each of the filtering steps can be done in a linear number of I/Os by the same argument as in Section 2.3.2 and the previous discussion. However, our structure lacks one important feature which we used in Section 2.3.2 to find dominating segments in the leaves. As in Section 2.3.2 we can argue that a leaf represents less than $M/4$ endpoints of the *original* segments, but as new segments and thus endpoints are introduced during the construction of the structure we cannot guarantee that the number of segments stored in a leaf is less than $M/2$. Therefore, we cannot find the dominating segments by loading all segments stored in a leaf into internal memory and using an internal memory algorithm. Also, the segments stored in a leaf may intersect each other and we need to find and reports such intersections. However, assuming that we can report such intersections and produce the sorted list of endpoints for each leaf, the rest of the algorithm runs in $O((n+t)\log_m n + t')$ I/Os, where T' is the number of intersections found during the filtering of the endpoints through the structure. If $T_1 = T + T'$ is the total number of intersections reported then this number is clearly $O((n+t_1)\log_m n)$.

In order to overcome the problems with leaves containing more than M segments we do the following: We collect all the segments stored in such leaves. The number of collected segments must be less than $2T_1$ (actually less than $2T$), since the involved leaves contain more than M segments but less than $M/2$ of the original N segments. The same holds for the number of endpoints assigned to the leaves. We then recursively build an external extended segment tree on these segments and filter the relevant endpoints through the structure in order to report intersections between the segments and produce a list of the points sorted according to dominating segment. If we do not count the I/Os used to process the leaves in this tree this costs us $O((t_1 + t_2)\log_m n)$ I/Os. Here T_2 is the number of intersections reported. Now we again need to collect the less than $2T_2$ segments in the leaves of the new tree containing more than M segments and recursively solve the problem for those segments. The process stops when all leaves contain less than M segments, and the total number of I/Os used on all the structure is then $O(n\log_m n + 2\sum_i t_i \log_m n) = O((n + t_t)\log_m n)$, where T_t is the total number of intersections reported. This completes our algorithm for the general segment intersection problem, giving us the following theorem:

Theorem 3 *All T intersections between N line segments in the plane can be reported in $O((n+t)\log_m n)$ I/O operations.*

4 Conclusions

In this paper, we have presented efficient external-memory algorithms for large-scale geometric problems involving collections of line segments in the plane, with applications to GIS systems. We have obtained these algorithms by combining buffer trees and distribution sweeping with a powerful new variant of fractional cascading designed for external memory.

The following two important problems, which are related to those we have discussed in this paper, remain open:

- If given the vertices of a polygon in the order they appear around its perimeter, can we triangulate the polygon in $O(n)$ I/O operations?
- Can we solve the general line segment intersection reporting problem in the optimal $O(n\log_m n + t)$ I/O operations?

References

[1] A. Aggarwal and J. S. Vitter. The Input/Output complexity of sorting and related problems. *Communications of the ACM*, 31(9):1116–1127, 1988.

[2] D. S. Andrews, J. Snoeyink, J. Boritz, T. Chan, G. Denham, J. Harrison, and C. Zhu. Further comparisons of algorithms for geometric intersection problems. In *Proc. 6th Int'l. Symp. on Spatial Data Handling*, 1994.

[3] ARC/INFO. *Understanding GIS—the ARC/INFO method.* ARC/INFO, 1993. Rev. 6 for workstations.

[4] L. Arge. The buffer tree: A new technique for optimal I/O-algorithms. In *Proc. Workshop on Algorithms and Data Structures, LNCS 955*, pages 334–345, 1995.

[5] L. Arge. The I/O-complexity of ordered binary-decision diagram manipulation. In *Proc. Int. Symp. on Algorithms and Computation, LNCS 1004*, pages 82–91, 1995.

[6] L. Arge, M. Knudsen, and K. Larsen. A general lower bound on the I/O-complexity of comparison-based algorithms. In *Proc. Workshop on Algorithms and Data Structures, LNCS 709*, pages 83–94, 1993.

[7] J. L. Bentley. Algorithms for klee's rectangle problems. Dept. of Computer Science, Carnegie Mellon Univ., unpublished notes, 1977.

[8] T. M. Chan. A simple trapezoid sweep algorithm for reporting red/blue segment intersections. In *Proc. of 6th Canadian Conference on Computational Geometry*, 1994.

[9] B. Chazelle. Triangulating a simple polygon in linear time. In *Proc. IEEE Symp. on Foundations of Comp. Sci.*, 1990.

[10] B. Chazelle and H. Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *Journal of the ACM*, 39:1–54, 1992.

[11] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. Algorithms for bichromatic line-segment problems and polyhedral terrains. *Algorithmica*, 11:116–132, 1994.

[12] B. Chazelle and L. J. Guibas. Fractional cascading: I. A data structuring technique. *Algorithmica*, 1:133–162, 1986.

[13] Y.-J. Chiang. Experiments on the practical I/O efficiency of geometric algorithms: Distribution sweep vs. plane sweep. In *Proc. Workshop on Algorithms and Data Structures, LNCS 955*, pages 346–357, 1995.

[14] Y.-J. Chiang, M. T. Goodrich, E. F. Grove, R. Tamassia, D. E. Vengroff, and J. S. Vitter. External-memory graph algorithms. In *Proc. ACM-SIAM Symp. on Discrete Algorithms*, pages 139–149, 1995.

[15] T. H. Cormen, T. Sundquist, and L. F. Wisniewski. Asymptotically tight bounds for performing BMMC permutations on parallel disk systems. Technical Report PCS-TR94-223, Dartmouth College Dept. of Computer Science, July 1994.

[16] R. F. Cromp. An intellegent information fusion system for handling the archiving and querying of terabyte-sized spatial databases. In *S. R. Tate ed., Report on the Workshop on Data and Image Compression Needs and Uses in the Scientific Community, CESDIS Technical Report Series, TR–93–99*, pages 75–84, 1993.

[17] H. Edelsbrunner and M. H. Overmars. Batched dynamic solutions to decomposable searching problems. *Journal of Algorithms*, 6:515–542, 1985.

[18] A. Fournier and D. Y. Montuno. Triangulating simple polygons and equivalent problems. *ACM Trans. on Graphics*, 3(2):153–174, 1984.

[19] M. T. Goodrich, J.-J. Tsay, D. E. Vengroff, and J. S. Vitter. External-memory computational geometry. In *Proc. IEEE Symp. on Foundations of Comp. Sci.*, pages 714–723, 1993.

[20] L. M. Haas and W. F. Cody. Exploiting extensible dbms in integrated geographic information systems. In *Proc. of Advances in Spatial Databases, LNCS 525*, 1991.

[21] R. Laurini and A. D. Thompson. *Fundamentals of Spatial Information Systems*. A.P.I.C. Series, Academic Press, New York, NY, 1992.

[22] H. G. Mairson and J. Stolfi. Reporting and counting intersections between two sets of line segments. In *R. Earnshaw (ed.), Theoretical Foundation of Computer Graphics and CAD, NATO ASI Series, Vol. F40*, pages 307–326, 1988.

[23] M. H. Nodine and J. S. Vitter. Large-scale sorting in parallel memories. In *Proc. ACM Symp. on Parallel Algorithms and Architectures*, pages 29–39, 1991.

[24] M. H. Nodine and J. S. Vitter. Deterministic distribution sort in shared and distributed memory multiprocessors. In *Proc. ACM Symp. on Parallel Algorithms and Architectures*, pages 120–129, 1993.

[25] M. H. Nodine and J. S. Vitter. Paradigms for optimal sorting with multiple disks. In *Proc. of the 26th Hawaii Int. Conf. on Systems Sciences*, 1993.

[26] L. Palazzi and J. Snoeyink. Counting and reporting red/blue segment intersections. In *Proc. Workshop on Algorithms and Data Structures, LNCS 709*, pages 530–540, 1993.

[27] Y. N. Patt. The I/O subsystem — a candidate for improvement. *Guest Editor's Introduction in IEEE Computer*, 27(3):15–16, 1994.

[28] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, 1985.

[29] C. Ruemmler and J. Wilkes. An introduction to disk drive modeling. *IEEE Computer*, 27(3):17–28, 1994.

[30] H. Samet. *Applications of Spatial Data Structures: Computer Graphics, Image Processing, and GIS*. Addison Wesley, MA, 1989.

[31] V. K. Vaishnavi and D. Wood. Rectilinear line segment intersection, layered segment trees, and dynamization. *Journal of Algorithms*, 3:160–176, 1982.

[32] M. J. van Kreveld. Geographic information systems. Utrecht University, INF/DOC–95–01, 1995.

[33] D. E. Vengroff and J. S. Vitter. I/O-efficient scientific computation using TPIE. In *Proc. IEEE Symp. on Parallel and Distributed Computing*, 1995. Appears also as Duke University Dept. of Computer Science technical report CS-1995-18.

[34] J. S. Vitter and E. A. M. Shriver. Algorithms for parallel memory, I: Two-level memories. *Algorithmica*, 12(2–3):110–147, 1994.

Recent Publications in the BRICS Report Series

RS-96-12 **Lars Arge, Darren E. Vengroff, and Jeffrey S. Vitter.** ***External-Memory Algorithms for Processing Line Segments in Geographic Information Systems*****. May 1996. 34 pp. An shorter version of this paper was presented at the *Third Annual European Symposium on Algorithms*, ESA '95.**

RS-96-11 **Devdatt Dubhashi, David A. Grable, and Alessandro Panconesi.** ***Near-Optimal, Distributed Edge Colouring via the Nibble Method*****. May 1996. 17 pp. Invited to be published in in a special issue of *Theoretical Computer Science* devoted to the proceedings of ESA '95.**

RS-96-10 **Torben Braüner and Valeria de Paiva.** ***Cut-Elimination for Full Intuitionistic Linear Logic*****. April 1996. 27 pp. Also available as Technical Report 395, Computer Laboratory, University of Cambridge.**

RS-96-9 **Thore Husfeldt, Theis Rauhe, and Søren Skyum.** ***Lower Bounds for Dynamic Transitive Closure, Planar Point Location, and Parentheses Matching*****. April 1996. 11 pp. To appear in *Algorithm Theory: 5th Scandinavian Workshop*, SWAT '96 Proceedings, LNCS, 1996.**

RS-96-8 **Martin Hansen, Hans Hüttel, and Josva Kleist.** ***Bisimulations for Asynchronous Mobile Processes*****. April 1996. 18 pp. Appears in *Tbilisi Symposium on Language, Logic, and Computation*, 1995.**

RS-96-7 **Ivan Damgård and Ronald Cramer.** ***Linear Zero-Knowledegde - A Note on Efficient Zero-Knowledge Proofs and Arguments*****. April 1996. 17 pp.**

RS-96-6 **Mayer Goldberg.** ***An Adequate Left-Associated Binary Numeral System in the λ-Calculus (Revised Version)*****. March 1996. 19 pp. Accepted for *Information Processing Letters*. This report is a revision of the BRICS Report RS-95-38.**

RS-96-5 **Mayer Goldberg.** ***Gödelisation in the λ-Calculus (Extended Version)*****. March 1996. 10 pp.**

RS-96-4 **Jørgen H. Andersen, Ed Harcourt, and K. V. S. Prasad.** ***A Machine Verified Distributed Sorting Algorithm*****. February 1996. 21 pp. Abstract appeared in *7th Nordic Workshop***