

# I/O-Efficient Map Overlay and Point Location in Low-Density Subdivisions\*

Mark de Berg<sup>†</sup>    Herman Haverkort<sup>†</sup>    Shripad Thite<sup>†</sup>    Laura Toma<sup>‡</sup>

## Abstract

We present improved and simplified I/O-efficient algorithms for two problems on planar low-density subdivisions, namely map overlay and point location. More precisely, we show how to preprocess a low-density subdivision with  $n$  edges in  $O(\text{sort}(n))$  I/O's into a compressed linear quadtrees such that one can:

- (i) compute the overlay of two such preprocessed subdivisions in  $O(\text{scan}(n))$  I/O's, where  $n$  is the total number of edges in the two subdivisions,
- (ii) answer a single point location query in  $O(\log_B n)$  I/O's and  $k$  batched point location queries in  $O(\text{scan}(n) + \text{sort}(k))$  I/O's.

For the special case where the subdivision is a fat triangulation, we show how to obtain the same bounds with an ordinary (uncompressed) quadtree, and we show how to make the structure fully dynamic using  $O(\log_B n)$  I/O's per update. Our algorithms and data structures improve on the previous best known bounds for general subdivisions both in the number of I/O's and storage usage, they are significantly simpler, and several of our algorithms are cache-oblivious.

## 1 Introduction

The traditional approach to algorithms design considers each atomic operation to take roughly the same amount of time. Unfortunately this simplifying assumption is invalid when the algorithm operates on data stored on disk: reading data from or writing data to disk can be a factor 100,000 or more slower than doing an operation on data that is already present in main memory. Thus, when the data is stored on disk it is usually much more important to minimize the number of disk accesses, rather than to minimize the CPU computation time.

This has led to the study of so-called I/O-efficient algorithms, also known as external-memory or out-of-core algorithms. The by now standard way of analyzing I/O-efficient algorithms is with the model introduced by Aggarwal and Vitter [1]. In this model, a computer has an internal memory of size  $M$  and an arbitrarily large disk. The data on disk is stored in blocks of size  $B$ , and whenever an algorithm wants to work on data not present in internal memory, the block(s) containing the data are read from disk. The I/O-complexity of an algorithm is the number of I/O's it performs—that is, the number of block transfers between the internal memory and the disk. In this model, scanning—reading a set of  $n$  consecutive items from disk—can be done in  $\text{scan}(n) = \lceil n/B \rceil$  I/O's, and sorting takes  $\text{sort}(n) = \Theta((n/B) \log_{M/B}(n/B))$  I/O's.

One of the main application areas for I/O-efficient algorithms has always been the area of geographic information systems (GIS), because GIS typically work with massive amounts of data and loading all of it into memory is often infeasible. In GIS the data for a particular geographic

---

\*MdB and ST were supported by the Netherlands' Organisation for Scientific Research (NWO)

<sup>†</sup>Department of Mathematics and Computer Science, Technische Universiteit Eindhoven, the Netherlands, mdberg@win.tue.nl, cs.herman@haverkort.net, sthite@win.tue.nl

<sup>‡</sup>Department of Computer Science, Bowdoin College, Brunswick, ME, USA, ltoma@bowdoin.edu

region is stored as a number of separate thematic layers. There can be a layer storing the road network, a layer storing the river network, a layer storing a subdivision of the region according to land usage or soil type, and so on. To combine the information from two such layers—for example to find the crossings between the river network and the road network—one has to compute the overlay of the layers. Even though the map-overlay problem is one of the most basic algorithmic problems in spatial databases and GIS, which are main application areas for I/O-efficient algorithms, the map-overlay problem has still not been solved satisfactorily in the I/O-model.

**Background.** The problem of map overlay can be formulated as a red-blue intersection problem: given a set of non-intersecting blue segments and a set of non-intersecting red segments in the plane, compute all intersections between the red and blue segments. Arge et al. [3] showed how to solve this problem in  $O(\text{sort}(n) + k/B)$  I/O's, where  $k$  is the number of intersections. Even though this is optimal in the worst case, it is not satisfactory for several reasons. First, as pointed out by the authors, their solution is complicated: the endpoint-dominance problem is solved to decide which segments are above the others; this information is then encoded in a planar  $(s, t)$ -graph which is sorted topologically according to the above-below order with the algorithm of Chiang et al. [8]; finally a distribution sweep is performed using a modified version of an external-memory segment tree. Hence, the practical value of the solution is unclear. Second, the algorithm uses  $\Theta(n \log_{M/B} n/B)$  storage (that is,  $\Theta(n/B \cdot \log_{M/B} n/B)$  disk blocks). A randomized solution for computing the intersections in a set of line segments is described by Crauser et al. [10]. They give the same expected I/O-bound of  $O(\text{sort}(n) + k/B)$  I/O's and linear space under some (realistic) assumptions for  $M, B, n$ . Whether this algorithm is practical is not clear either.

Although the I/O-complexity of the above algorithms is optimal for general sets of line segments, there are important special cases for which this is not clear. For example, the overlay of two triangulations that are suitably stored—in a doubly-connected edge list [14], say—can trivially be computed in  $O(n+k)$  time in internal memory, which raises the question whether the overlay of two suitably stored triangulations can be computed in  $O(\text{scan}(n+k))$  I/O's. In fact, in internal memory one can overlay two subdivisions in  $O(n+k)$  time when these subdivisions are connected [15]. This brings us to the topic of our paper: is it possible to do the overlay of two planar maps in  $O(\text{scan}(n+k))$  I/O's? And can it be done *cache-obliviously* [16] so that no parameter tuning is necessary and the I/O-behavior is good over all levels in a multilevel memory hierarchy?

Most research into I/O-efficient algorithms has focused on algorithms that are efficient for any worst-case input. However, worst-case inputs are often artificial constructions that do not usually occur in real life. In computational geometry this has led to the study of input models where inputs are assumed to have properties that make them resemble realistic inputs better [13]. The most common assumption in such models is that the objects are *fat*, that is, they are not arbitrarily long and narrow. Fatness has been studied extensively in computational geometry in the recent years, and solutions to many fundamental problems have been improved by exploiting fatness and related notions, see De Berg et al. [13, 12] and references therein.

In this paper, we consider two types of subdivisions: fat triangulations and low-density subdivisions. A  $\delta$ -*fat triangulation* is a triangulation in which every angle is bounded from below by a fixed positive constant  $\delta$ . A  $\lambda$ -*low-density subdivision* is a subdivision such that any disk  $D$  is intersected by at most  $\lambda$  edges whose length is at least the diameter of  $D$ , for some fixed constant  $\lambda$ . We believe these two types of subdivisions include most subdivisions encountered in practice, for reasonable values of  $\delta$  and  $\lambda$ .

The data structures on which we will base our solutions are modifications of the so-called *linear quadtree*, which was introduced by Gargantini [17]. The linear quadtree is a quadtree variant where only the leaf regions are stored, and not the internal nodes. To facilitate a search in the quadtree, a linear order is defined on the leaves based on some space-filling curve; then a B-tree is constructed on the leaves using this ordering—see Section 2 for details. The idea of using linear quadtrees to

store planar subdivisions has been used by Hjaltason and Samet [19]. They present algorithms for constructing (or: bulk-loading, as it is often called in GIS) the quadtree, for insertions, and for bulk-insertions. Although their experiments indicate their method performs well in practice, it has several disadvantages. First, the I/O-complexity of their algorithms is analysed in terms of various parameters that depend on the data and the algorithm in a way that is not well-understood. In particular, the performance of their algorithms does not seem to be worst-case optimal. Second, the stopping rule for splitting quadtree cells is based on two user-defined parameters (the maximum depth and a so-called splitting threshold), and so the method is not fully automatic.

**Our results.** In this paper we show how to overcome these disadvantages for fat triangulations and low-density subdivisions and present improved and simplified algorithms for map overlay and point location in external memory. Our results are based on a quadtree which we define to ensure that (i) each leaf intersects only a constant number of edges of the subdivision, (ii) that we create only  $O(n)$  leaves, and (iii) that we can construct the leaves efficiently.

For fat triangulations our quadtree is defined by recursively splitting the unit square into quadrants until all edges that intersect a cell are incident to a common vertex. We prove that this stopping rule yields a quadtree of linear size. Nevertheless, due to the potentially large depth of the quadtree, it is still difficult or impossible to build the quadtree I/O-efficiently by distributing the edges from the root down into the quadtree while splitting nodes as needed. Fortunately our stopping rule makes a completely different approach possible: we give an algorithm that is simple and elegant—simpler than the algorithm of Hjaltason and Samet [19]—and uses only  $O(\text{sort}(n))$  I/O's.

For low-density subdivisions we continue splitting until each cell contains only a single bounding-box vertex of any edge. This stopping rule leads to cells with a constant number of edges, but the number of cells cannot be bounded. Therefore we combine the ideas of compressed quadtrees and linear quadtrees to get a *linear compressed quadtree*, rather than a regular quadtree. We show that with the stopping rule just defined, the compressed quadtree has linear size. We also give a construction algorithm that uses only  $O(\text{sort}(n))$  I/O's.

Once we have proved that these quadtrees have linear size and each leaf region intersects a constant number of edges, our other results come almost for free: overlaying two subdivisions boils down to a simple merge of the ordered lists of quadtree leaves taking  $O(\text{scan}(n))$  I/O's, point location can be done in  $O(\log_B n)$  I/O's by searching in the B-tree built on top of the list of quadtree leaves, and performing  $k$  batched point location queries can be done in  $O(\text{scan}(n) + \text{sort}(k))$  I/O's by sorting the points along the space-filling curve and merging the sorted list with the list of quadtree leaves. The results for map overlay apply to pairs of fat triangulations, low-density subdivisions, or low-density sets of line segments, as well as to mixed pairs of maps of these types. The structure for fat triangulations can be made fully dynamic at the cost of  $O(\log_B n)$  I/O's per update.

An optimal static structure for point location in general planar subdivisions was already given by Goodrich et al. [18] for the standard I/O-model and by Bender et al. [5] for the cache-oblivious model. Batched point location can be done with  $O(\text{sort}(n + k))$  I/O's in the I/O-model using the algorithm by Arge et al. [3]. The result on dynamization, however, is new as far as we know: the best known dynamic I/O-efficient point location structures use  $O(\log_B^2 n)$  I/O's per query [2] in the I/O-model. All our data structures and query algorithms are cache-oblivious. Our construction and update algorithms for triangulations can be made cache-oblivious, except that updates will then take  $O(\log_B n + \frac{1}{B} \log^2 n)$  I/O's. These results constitute the first results for cache-oblivious map overlay, batched point location and dynamic point location.

## 2 Fat triangulations

In this section we describe our solution for fat triangulations. A  $\delta$ -fat triangulation is a triangulation consisting of  $\delta$ -fat triangles, that is, triangles all of whose angles are at least  $\delta$  for some fixed constant

$\delta > 0$ . We assume that  $B = \Omega(1/\delta)$  and  $M = \Omega(1/\delta^3)$ . We assume that all triangulations are triangulations of the unit square  $[0, 1]^2$ . (Our algorithms and proofs extend to triangulations of convex regions—we leave the details for the full paper.) In this section we show the following:

**Theorem 2.1.** *Let  $\mathcal{F}$  be a  $\delta$ -fat triangulation with  $n$  edges. Knowing the memory size  $M$  and the block size  $B$ , we can construct, in  $O(\text{sort}(n/\delta^2))$  I/O's, a linear quadtree for  $\mathcal{F}$  with  $O(n/\delta^2)$  cells such that each cell intersects  $O(1/\delta)$  triangles and the total number of intersections between cells and triangles is  $O(n/\delta^2)$ . With this structure we can perform the following operations:*

- (i) **Map overlay:** *Given two  $\delta$ -fat triangulations with  $n$  triangles in total, each stored in such a linear quadtree, we can find all pairs of intersecting triangles in  $O(\text{scan}(n/\delta^2))$  I/O's.*
- (ii) **(Batched) point location:** *for any query point  $p$  we can find the triangle of  $\mathcal{F}$  that contains  $p$  in  $O(\log_B(n/\delta))$  I/O's, and for any set  $P$  of  $k$  query points we can find for each point  $p \in P$  the face of  $\mathcal{F}$  that contains  $p$  in  $O(\text{scan}(n/\delta^2) + \text{sort}(k))$  I/O's.*
- (iii) **Updates:** *Inserting a vertex, moving a vertex, deleting a vertex, and flipping an edge can all be done in  $O((\log_B n)/\delta^4)$  I/O's.*

*In the cache-oblivious model the same bounds hold, except that updates then take  $O((\log_B n + \frac{1}{B} \log^2 n)/\delta^4)$  I/O's.*

## 2.1 The quadtree subdivision for fat triangulations

A quadtree is a hierarchical subdivision of the unit square into quadrants, where the subdivision is defined by a criterion to decide what quadrants are subdivided further, and what quadrants are leaves of the hierarchy. A *canonical square* is any square that can be obtained by recursively splitting the unit square into quadrants. For a canonical square  $\sigma$ , let  $\text{mom}(\sigma)$  denote its parent, that is, the canonical square that contains  $\sigma$  and has twice its width. The leaves of the quadtree form the quadtree subdivision; that is, a *quadtree subdivision* for a set of objects in the unit square is a subdivision into disjoint canonical squares (*quadtree cells*), such that each cell obeys the stopping rule while its parent does not. The stopping rule we use is as follows:

*Stopping rule for fat triangulations:* Stop splitting when all edges intersecting the cell  $\sigma$  under consideration are incident to a common vertex.

Note that the stopping rule includes the case where  $\sigma$  is not intersected by any edges. To prove that the stopping rule gives a linear number of cells, we use the following lemma, which is a special case of known results on fatness.

**Lemma 2.2.** *Let  $0 < \beta \leq 1$  be a constant, and let  $R$  be a region in the plane. Then  $R$  is intersected by  $O(1/\beta^2)$  quadtree cells  $\sigma$  such that  $\text{diam}(\sigma) \geq \beta \cdot \text{diam}(R)$ .*

Now we are ready to prove bounds on the complexity of the quadtree subdivision.

**Lemma 2.3.** *Let  $\mathcal{F}$  be a  $\delta$ -fat triangulation of the unit square with  $n$  edges. Then the stopping rule defined above leads to a quadtree subdivision with  $O(n/\delta^2)$  cells, such that each cell is intersected by at most  $2\pi/\delta$  triangles, and the total number of triangle-cell intersections is  $O(n/\delta^2)$ .*

*Proof:* Since the degree of each vertex in  $\mathcal{F}$  is at most  $2\pi/\delta$ , the stopping rule implies that each cell is intersected by at most that many triangles. Below we bound the number of triangle-cell intersections, from which the bound on the number of cells follows immediately.

For a vertex  $v$  of  $\mathcal{F}$ , let  $d_v$  be the closest distance to an edge opposite to  $v$ , that is, an edge of a triangle incident to  $v$  that is not incident to  $v$  itself. Let  $D_v$  be the disk centered at  $v$  and of radius


$r_v := d_v/2$ . Using the fact that the minimum angle in any triangle is  $\delta$ , we can show that for any edge  $e$  incident to  $v$  we have  $r_v \geq (\sin \delta)|e|/2^{\pi/\delta+1}$ , where  $|e|$  is the length of  $e$ .

Let  $\sigma$  be a square in the quadtree subdivision. We will charge each intersection of a cell-triangle pair  $(\sigma, \Delta)$  either to a vertex or an edge of  $\mathcal{F}$ , as follows.

If  $\sigma$  intersects some disk  $D_v$ , we charge  $(\sigma, \Delta)$  to  $v$ . (If it intersects more disks, we take an arbitrary one and charge  $(\sigma, \Delta)$  to the vertex defining that disk.) Since  $\text{mom}(\sigma)$  cannot be completely contained in the union of the triangles incident to  $v$ —otherwise  $\text{mom}(\sigma)$  would not have been split by our stopping rule—we have  $\text{diam}(\text{mom}(\sigma)) \geq r_v$ . Hence,  $\text{diam}(\sigma) \geq r_v/2$ . From Lemma 2.2 we now know that  $v$  is charged by intersections with  $O(1)$  cells. Since each cell intersects at most  $2\pi/\delta$  triangles, the total number of cell-triangle intersections charged in this way is  $O(n/\delta)$ .

Now assume  $\sigma$  does not intersect any disk  $D_v$ .

If  $\sigma$  does not intersect any edge, we charge  $(\sigma, \Delta)$  to the edge  $e$  of  $\Delta$  that is closest to  $\sigma$ , breaking ties arbitrarily. If  $\sigma$  intersects one or more edges, we charge  $(\sigma, \Delta)$  to an edge  $e$  of  $\Delta$  that intersects  $\sigma$ —thus charging at most two such pairs  $(\sigma, \Delta)$  to  $e$  for any fixed  $\sigma$ . From at least one of the two triangles incident to  $e$  the square  $\text{mom}(\sigma)$  intersects an edge different from  $e$ . Let  $v_1, v_2$  be the vertices of  $e$ , and let  $v_3$  be the third vertex of this triangle  $\Delta'$ .



Consider the Voronoi diagram of the edges of  $\Delta'$  inside  $\Delta'$ . Because  $\sigma$  is charged to  $e$ , it must intersect the Voronoi cell  $\mathcal{V}(e)$  of  $e$ .

We partition  $\mathcal{V}(e)$  into a number of regions, and assign  $(\sigma, \Delta)$  to one of these regions, as follows. Extend the bisector at  $v_3$  (which separates  $\mathcal{V}(v_1v_3)$  from  $\mathcal{V}(v_2v_3)$ ) into  $\mathcal{V}(e)$ , thus splitting  $\mathcal{V}(e)$  into two regions. At least one of these regions intersects  $\sigma$ —in the figure this region is shown in grey. We partition this region,  $R$ , into subregions  $R_1, R_2, \dots, R_m$  as follows. Let  $p_0$  be the Voronoi vertex inside  $\Delta'$ . Let  $p_1$  be the midpoint of  $v_1p_0$ . Region  $R_1$  is obtained by drawing a segment from  $p_1$  orthogonal to  $v_1p_0$  into  $R$ . Region  $R_2$  is obtained by splitting the remainder of  $R$  by a segment from the midpoint of  $v_1p_1$  and orthogonal to  $v_1p_0$ , etc. This process is repeated until the remaining region is completely contained in  $D_v$ —see the figure.

Using a packing argument and the fact that all angles in  $\Delta'$  are at least  $\delta$ , we can now show that there are  $O(1/\delta)$  regions per edge  $e$ , that each region gets assigned  $O(1/\delta)$  different cells, and that for each such cell, at most two cell-triangle intersections are charged to  $e$ . Thus there are  $O(n/\delta^2)$  cell-triangle intersections in total. ■

## 2.2 Storing the quadtree subdivision and the triangulation

We will store the quadtree subdivision defined above as a so-called *linear quadtree* [17]. To this end, we define an ordering on the leaf cells of the quadtree subdivision. The ordering is based on a space-filling curve defined recursively by the order in which it visits the quadrants of a canonical square. We will use the *z-order space-filling curve* for this, which visits the quadrants in the order bottom left, top left, bottom right, top right, and within each quadrant, the z-order curve visits its subquadrants recursively in the same order. Since the intersection of every canonical square with this curve is a contiguous section of the curve, this yields a well-defined ordering of the leaf cells of the quadtree subdivision. We call the resulting order the z-order.

The z-order curve not only orders the leaf cells of the quadtree subdivision, but it also provides an ordering for any two points in the unit square—namely the z-order of any two disjoint canonical squares containing the points. (We assume that canonical squares are closed at the bottom and the left side, and open at the top and the right side.) The z-order of two points can be determined as follows. For a point  $p = (p_x, p_y)$  in  $[0, 1)^2$ , define its z-index  $Z(p)$  to be the value in the range  $[0, 1)$  obtained by interleaving the bits of the fractional parts of  $p_x$  and  $p_y$ , starting with the first

bit of  $p_x$ . The value  $Z(p)$  is sometimes called the *Morton block index* of  $p$ . The z-order of two points is now the same as the order of their z-indices [19]. The z-indices of all points in a canonical square  $\sigma$  form a subinterval  $[z_1, z_2)$  of  $[0, 1)$ , where  $z_1$  is the z-index of the bottom left corner of  $\sigma$ . Note that any subdivision of the unit square  $[0, 1)^2$  into  $k$  leaf cells of a quadtree corresponds directly to a subdivision of the unit interval  $[0, 1)$  of z-indices into  $k$  subintervals.

A simple (but novel) way of storing a triangulation in a linear quadtree is now obtained by storing all cell-triangle intersections in a B-tree [9]: each cell-triangle intersection  $(\sigma, \Delta)$  of a cell  $\sigma$  corresponding to the z-index interval  $[z_1, z_2)$  is represented by storing triangle  $\Delta$  with key  $z_1$ . With this way of storing the linear quadtree, the leaf cells of the quadtree are stored implicitly: each pair of consecutive different keys  $z_1$  and  $z_2$  constitutes the z-index interval of a quadtree leaf cell. To obtain a cache-oblivious solution, we can use a cache-oblivious B-tree [5, 6, 7].

In the remainder we will sometimes need to compute or compare z-indices. Whether this takes constant time depends on the operations allowed by the model of computation. In any case, since we care mainly about I/O-efficiency, such computations do not effect the analysis of our algorithms.

### 2.3 Building the quadtree I/O-efficiently

The natural algorithm to build a quadtree would take a set of triangles and a canonical square (initially all triangles and the unit square) as input, check if the condition of the stopping rule is satisfied, and if not, distribute the triangles among the four children and subdivide the children recursively. Unfortunately, this algorithm takes  $O(n^2)$  time and  $O(n^2/B)$  I/O's, as the quadtree may have height  $O(n)$ . Below we describe a faster algorithm that computes the leaf cells that result with our stopping rule directly, using local computations instead of a top-down approach.

For any vertex  $v$  of the given triangulation  $\mathcal{F}$ , let  $star(v)$  be the *star* of  $v$  in  $\mathcal{F}$ ; namely, it is the set of triangles of  $\mathcal{F}$  that have  $v$  as a vertex. Recall that a canonical square is any square that can be obtained by recursively subdividing the unit square into quadrants. For a set  $S$  of triangles inside the unit square, we say that a canonical square  $\sigma$  is *active* in  $S$  if it lies completely inside  $S$  and all edges from  $S$  that intersect  $\sigma$  are incident to a common vertex, while  $mom(\sigma)$  intersects multiple edges of  $S$  that are not all incident to a common vertex. Thus the cells of the quadtree subdivision we wish to compute for  $\mathcal{F}$  are exactly the active canonical squares in  $\mathcal{F}$ .

**Lemma 2.4.** *Let  $\Delta = (u, v, w)$  be a triangle of  $\mathcal{F}$  and  $\sigma$  a canonical square that intersects  $\Delta$ . Then  $\sigma$  is active in  $\mathcal{F}$  if and only if  $\sigma$  is active in  $star(u)$ ,  $star(v)$  or  $star(w)$ .*

*Proof:* Consider the canonical cell  $\tau = mom(\sigma)$ . Let  $\gamma_u$  denote the boundary of  $star(u)$ . The curve  $\gamma_u$  is a simple closed polygonal curve.

If  $\sigma$  is active in  $star(u)$  then there exist three edges of  $star(u)$  that intersect  $\tau$  and are not all incident to the same vertex. Since  $\sigma$  is completely contained inside  $star(u)$ , it is separated from any edges of  $\mathcal{F}$  outside  $star(u)$  by  $\gamma_u$ . Hence  $\sigma$  is active in  $\mathcal{F}$ . The same argument implies that if  $\sigma$  is active in  $star(v)$  or in  $star(w)$ , then it is active in  $\mathcal{F}$ .

As for the other direction, suppose  $\sigma$  is active in  $\mathcal{F}$  and, without loss of generality, assume that all edges that intersect  $\sigma$  are incident to  $u$ . We distinguish three cases:  $\tau$  intersects at least two, one, or zero edges incident to  $u$ .

If  $\tau$  intersects at least two edges incident to  $u$ , then  $\tau$  must also intersect a third edge that is not incident to  $u$ ; in particular it intersects an edge of  $\gamma_u$ . The square  $\tau$  thus intersects three edges of  $star(u)$  that are not all incident to a common vertex, and its child  $\sigma$  is active in  $star(u)$ .

If  $\tau$  intersects only one edge incident to  $u$ , then, without loss of generality, let  $(u, v)$  be this edge. If  $\tau$  intersects another edge incident to  $v$ , we apply the argument for the first case to  $v$  and conclude that  $\sigma$  is active in  $star(v)$ . Otherwise  $\tau$  is fully contained inside the union of the two triangles incident to  $(u, v)$ . But then  $\sigma$  cannot be active for  $\mathcal{F}$ . This is a contradiction.

Finally, if  $\tau$  does not intersect any edges incident to  $u$ , then  $\sigma$  does not intersect any edges at all and lies completely inside  $\Delta$ . Furthermore,  $\tau$  intersects  $(v, w)$ . We now apply the arguments given above with  $u$  and  $v$  changing roles, and find that  $\sigma$  is active in  $star(v)$  or  $star(w)$ . ■

On the basis of the above lemma, we can construct the linear quadtree as follows:

1. Compute an adjacency list for each vertex.
2. Scan the adjacency lists for all vertices: for each vertex  $u$  load its adjacency list in memory and compute the active cells of  $star(u)$ , with for each cell  $\sigma$  the triangles that intersect  $\sigma$ . Output each triangle with the key  $z_1$  of the Z-index interval  $[z_1, z_2)$  that corresponds to  $\sigma$ .
3. Sort the triangles by key, removing duplicates.
4. Finally, build a (cache-oblivious) B-tree on the list of triangles with their keys.

**Lemma 2.5.** *The quadtree described above for a  $\delta$ -fat triangulation with  $n$  edges can be constructed with  $O(\text{sort}(n/\delta^2))$  I/O's. (Proof straightforward from the previous Lemmas)*

**Eliminating superfluous cells.** The I/O-complexity of the construction and the storage requirements in practice can be reduced with an easy optimization: we merge all active cells that lie properly inside triangles with their successors or predecessors in the Z-order. In fact, in step 2 of the algorithm, we will not even output such cells. Instead we only output triangle-key pairs for triangle-cell intersections such that an edge of the triangle intersects the cell. We sort these triangle-key pairs, and then scan them. Whenever two consecutive triangles have different keys  $z_1$  and  $z_2$ , we identify the most significant bit that differs between them. Let  $z$  be the lowest Z-index in  $[z_1, z_2]$  for which this bit has value 1. For each triangle stored with key  $z_2$ , we now replace its key by  $z$ . Thus all cells with Z-indices in  $[z_1, z)$  are merged with each other, and all cells with  $z$ -values in  $[z, z_2]$  are merged with each other.

Since the interval  $[z_1, z)$  of the Z-order curve covers a connected area in the plane, all cells in the range  $[z_1, z)$  that do not intersect any edge must be completely contained in a triangle that already intersects the cell that starts with Z-index  $z_1$ . Similarly,  $[z, z_2]$  and the cell that starts with  $z_2$  together cover a connected area in the plane, so the triangles that intersect it must have been stored with key  $z_2$  already. Hence no more triangles need to be stored as a result of merging cells.

**Updates.** We support the following operations: inserting a vertex, moving a vertex, deleting a vertex, and flipping an edge. By Lemma 2.4, all leaf cells that intersect a triangle  $\Delta = (u, v, w)$  can be computed from the local quadtrees of  $star(u)$ ,  $star(v)$  and  $star(w)$ . Since the size of  $star(u)$ ,  $star(v)$  and  $star(w)$  is  $O(1/\delta)$ , by Lemma 2.3 the total number of cells that intersect  $\Delta$  is  $O(1/\delta^3)$ . Since each of the supported operations changes only  $O(1/\delta)$  triangles, we can compute the structure of the quadtree locally in the area of the update, and determine what the changes entail for the data stored on the disk. All changes can thus be made in  $O((\log_B n)/\delta^4)$  I/O's when a normal B-tree is used, and in  $O((\log_B n + \frac{1}{B} \log^2 n)/\delta^4)$  I/O's when a cache-oblivious B-tree is used [5, 6, 7].

## 2.4 Overlaying maps and point location

**Lemma 2.6.** *The linear quadtree for  $\delta$ -fat triangulations as described above supports map overlay in  $O(\text{scan}(n/\delta^2))$  I/O's, and point location in  $O(\log_B(n/\delta))$  I/O's, where  $n$  is the number of points in the triangulation. Batched point location for  $k$  points takes  $O(\text{scan}(n/\delta^2) + \text{sort}(k))$  I/O's.*

*Proof:* Each triangulation's quadtree, or rather, subdivision of the Z-order curve, is stored on disk as a sorted list of Z-indices with triangles. To overlay the two triangulations, we will scan the two quadtrees simultaneously in Z-order, at any point keeping in memory the triangles stored with the last key read from the first list and those stored with the last key read from the second list.

Starting from the beginning of the lists, we repeat the following until both lists have been read completely: we read the next key from the list with the smallest unread key, we load all triangles stored with that key into memory, and we compute the intersections with the triangles in memory that were read from the other list. The input has size  $O(n/\delta^2)$ . The output consists of  $O(n/\delta)$  intersections since a  $\delta$ -fat triangulation has density  $O(1/\delta)$  [13], which implies the claim.

To perform point location with a point  $p$ , we compute the z-index  $Z(p)$  of  $p$  and search the B-tree for the triangles with the highest keys less than or equal to  $Z(p)$ . To do batched point location, we sort the set  $P$  of query points by z-index, and scan the leaves of the B-tree and  $P$  in parallel (similar to the overlay operation as described above). ■

### 3 Low-density subdivisions

In this section we describe our solution for storing planar low-density subdivisions. For a planar object  $o$ , let  $\text{size}(o)$  denote the radius of the smallest enclosing disk of  $o$ . The *density* of a set  $S$  of objects in the plane is the smallest number  $\lambda$  such that the following holds: any disk  $D$  is intersected by at most  $\lambda$  objects  $o \in S$  such that  $\text{diam}(o) \geq \text{diam}(D)$  [13]. We say that a planar subdivision  $\mathcal{F}$  has density  $\lambda$  if its edge set has density  $\lambda$ . In other words, any disk  $D$  is intersected by at most  $\lambda$  edges whose length is at least the diameter of  $D$ . We assume that  $B = \Omega(\lambda)$ , and that the input lies in the unit square  $[0, 1]^2$ . In this section we will prove the following result.

**Theorem 3.1.** *Let  $\mathcal{F}$  be a subdivision or a set of non-intersecting line segments of density  $\lambda$  with  $n$  edges. Knowing the memory size  $M$  and the block size  $B$ , we can construct, in  $O(\text{sort}(\lambda n))$  I/O's, a linear compressed quadtree for  $\mathcal{F}$  with  $O(n)$  cells that each intersect  $O(\lambda)$  edges. With this structure we can perform the following operations:*

- (i) **Map overlay:** *If we have two subdivisions (or sets of non-intersecting line segments) of density  $\lambda$  with  $n$  edges in total, both stored in such a linear compressed quadtree, then we can find all pairs of intersecting edges in  $O(\text{scan}(\lambda n))$  I/O's.*
- (ii) **(Batched) point location:** *for any query point  $p$  we can find the face of  $\mathcal{F}$  that contains  $p$  in  $O(\log_B n)$  I/O's, and for any set  $P$  of  $k$  query points we can find for each point  $p \in P$  the face of  $\mathcal{F}$  that contains  $p$  in  $O(\text{scan}(\lambda n) + \text{sort}(k))$  I/O's.*

*The data structure, the overlay algorithm and the query algorithms are cache-oblivious.*

Any set of disjoint  $\delta$ -fat triangles in the plane has density  $O(1/\delta)$  [13]. Thus the results of this section can be used for  $\delta$ -fat triangulations. However, the solution from the previous section is simpler and dynamic.

Below we explain our data structure, and how to construct it. The query algorithms are the same as described in the previous section.

#### 3.1 The compressed quadtree subdivision for low-density maps

Let  $\mathcal{F}$  be a subdivision of the unit square with  $n$  edges and of density  $\lambda$ . In general it is impossible to construct a standard quadtree on  $\mathcal{F}$  consisting of a linear number of cells that are each intersected by a constant number of edges. Indeed, in a general subdivision of the unit square there can be many vertices arbitrarily close together, even if the subdivision has constant density. To overcome this problem we shall use a so-called *compressed quadtree*.

Let  $\mathcal{G}$  be the set of vertices of the axis-parallel bounding boxes of the edges of  $\mathcal{F}$ . This set has a nice property:

**Lemma 3.2** ([11]). *Any square  $\sigma$  that does not contain any bounding-box vertex of an object in a set  $S$  with density  $\lambda$ , intersects  $O(\lambda)$  objects from  $S$ .*

We now construct a quadtree for  $\mathcal{F}$  with the following stopping rule.

*Stopping rule for low-density subdivisions:* Stop splitting when the cell  $\sigma$  under consideration contains at most one point from  $\mathcal{G}$ .

Consider the quadtree that we get from this stopping rule. Its cells intersect  $O(\lambda)$  edges, but the number of cells cannot be bounded. Hence, we compress the quadtree [20], by building it as follows. We recursively subdivide each canonical square  $\sigma$  that contains more than one point from  $\mathcal{G}$  into *five* regions. Let  $\sigma'$  be the smallest canonical square that contains all points of  $\sigma \cap \mathcal{G}$ . The first region is the *donut*  $\sigma \setminus \sigma'$ . The remaining four regions are the four quadrants of  $\sigma'$ . Note that the first region does not contain any points of  $\mathcal{G}$ , so it is never subdivided further. When  $\sigma' = \sigma$ , the first region is skipped; when  $\sigma'$  is smaller than  $\sigma$ , we call  $\sigma \setminus \sigma'$  a *proper donut*.

**Lemma 3.3.** *Let  $\mathcal{F}$  be a subdivision of the unit square with  $n$  edges and of density  $\lambda$ . Then a compressed quadtree subdivision based on the stopping rule defined above has  $O(n)$  cells, and each cell is intersected by at most  $O(\lambda)$  edges.*

*Proof:* Note that each internal node of the compressed quadtree has points from  $\mathcal{G}$  in at least two of its subtrees, hence the total number of nodes is  $O(|\mathcal{G}|) = O(n)$ .

If a cell  $\sigma$  is a square, it contains at most one guarding point of  $\mathcal{G}$ . Hence  $\sigma$  can be covered by four (not necessarily disjoint) squares that lie completely in  $\sigma$  and do not have points from  $\mathcal{G}$  in their interior. By Lemma 3.2, we conclude that  $\sigma$  is intersected by at most  $O(\lambda)$  edges.

If  $\sigma$  is a proper donut, it is empty of guard points. Moreover, it can be covered by six squares that do not contain any points of  $\mathcal{G}$  in their interior, and hence only  $O(\lambda)$  edges intersect it. ■

### 3.2 Storing the compressed quadtree subdivision and the low-density map

We store the cell-edge intersections of the compressed quadtree subdivision in a list sorted by the z-order of the cells, indexed by a (cache-oblivious) B-tree. The only difference with the previous section is that we now have to deal with donuts as well as square cells. Recall that a canonical square (a square that can be obtained from the unit square by a recursive partitioning into quadrants) corresponds to an interval on the z-order curve. For a donut this is not true. However, a donut corresponds to at most two such intervals, because a donut is the set-theoretic difference of two canonical squares. Thus the solution of the previous section can be applied if we represent each donut by two intervals  $[z_1, z_2)$  and  $[z_3, z_4)$ ; edges intersecting the first part of the donut are stored with key  $z_1$  and edges intersecting the second part are stored with key  $z_3$ . As described in Section 2.3, we merge cells that do not intersect any edge with their immediate successors or predecessors in the z-order. We call the resulting structure—the B-tree on the cell-edge intersections whose keys imply a compressed quadtree subdivision—a *linear compressed quadtree*. Map overlay and point location are done in exactly the same way as with the linear quadtree described earlier.

### 3.3 Building the compressed quadtree I/O-efficiently

We construct the leaves of the compressed quadtree, or rather, the corresponding subdivision of the z-order curve, as follows. We sort  $\mathcal{G}$  into z-order, and scan the sorted points. For each pair of consecutive points, say  $u$  and  $v$ , we construct their lowest common ancestor  $lca(u, v)$  by examining the longest common prefix of the bit strings representing  $z(u)$  and  $z(v)$ . We output the five z-indices that bound and separate the z-order intervals of the four children of  $lca(u, v)$ . To complete the subdivision of the z-order curve, we sort the output into a list  $\mathcal{L}$  by z-order, removing duplicates.

**Lemma 3.4.** *The above algorithm generates a subdivision of the z-order curve that corresponds to a compressed quadtree on  $\mathcal{G}$  in  $O(\text{sort}(n))$  I/O's.*

*Proof:* The definition of the compressed quadtree implies that a canonical square  $\sigma$  is subdivided into its four quadrants if and only if it contains at least two points that lie in different quadrants of  $\sigma$ . If such a pair exists, at least one such pair  $(u, v)$  appears as a pair of consecutive points in the z-order, since the z-order curve visits all points in  $\sigma$  consecutively. The above algorithm generates the boundaries of the four quadrants of  $\sigma$  when the pair  $(u, v)$  is encountered. Furthermore, whenever the above algorithm generates the boundaries of the four quadrants of a cell  $\sigma$ , there is a pair  $(u, v)$  that has  $\sigma$  as their lowest common ancestor, and therefore lie in different quadrants of  $\sigma$ .

It remains to check that the boundaries of donut cells are generated correctly. The z-indices that define the outer square of a donut  $\sigma \setminus \tau$  appear when a pair of points is processed that consists of a point in the hole  $\sigma'$  and a point in a sibling of  $\sigma$ . The z-indices that define the hole  $\sigma'$  appear when a pair of points is processed that lie in different quadrants of the inner square. ■

Having constructed the compressed quadtree subdivision, we now distribute the edges in  $\mathcal{F}$  to the faces of the quadtree subdivision, or rather, to the corresponding sections of the z-order curve.

To do so, we first build a B-tree on the subdivision of the z-order curve as computed above. We then load the (roughly)  $M/B$  nodes of the B-tree, that reside  $\log_B(M/B)$  levels below the root, in memory. Note that each of these nodes covers a certain section of the z-order curve, and together they form a subdivision of the z-index interval covered by the root. We assign an output stream to each of these nodes, and reserve a buffer of one block in memory for each of them. We now read the edges from the input one by one, and distribute each edge to the output streams of the nodes whose section of the z-order curve is intersected by the edge. Note that each edge may be copied to several streams. Once all edges have been read, we distribute the edges in each node's stream recursively into the subtree rooted at that node.

After distributing all edges recursively to the leaves, we collect all edge-cell intersections, ordered by the z-indices of the cells, and put a new B-tree on top of them. Each cell  $\sigma$  without any intersecting edges is merged with the cells that precede or follow it in the z-order, up to a cell that stores an edge of the face of  $\mathcal{F}$  that contains  $\sigma$  (see Section 2.3 for an explanation).

**Lemma 3.5.** *The compressed quadtree as defined above for a subdivision of density  $\lambda$  with  $n$  edges can be constructed with  $O(\text{sort}(\lambda n))$  I/O's. (Proof straightforward from the above)*

## 4 Conclusions

We described how one can efficiently store and overlay planar maps in the I/O-model of computation. Our algorithms work for planar maps that are fat triangulations or have low density. The solution for triangulations is based on quadtrees, is considerably simpler than previous solutions, and supports even dynamic maintenance of the planar maps under updates. The second construction, for low density planar maps, is based on compressed quadtrees and is somewhat more complicated; however our analysis gives a better dependency on the parameter that describes the input. Unfortunately it is not clear if the construction algorithm can be made cache-oblivious and if the structure can be made to support updates. Both constructions use linear space, improving on the previous space bound of  $O(n \log_{M/B} n)$  of Arge et al. [3]. Which of our two structures would give the most compact data structure for triangulations in practice remains to be seen. The first approach's dependency on the fatness may be better than it seems (perhaps an analysis in terms of *average* fatness is possible), while the second approach may introduce many guards (a triangulation of  $n$  vertices has roughly  $3n$  edges and thus roughly  $6n$  extra bounding box vertices as guards).

Our data structures can be used for range searching queries. In general this would not be very efficient, but we believe it is possible to achieve good bounds for *approximate range searching* [4]. However, the data structure for low-density subdivisions as presented in this paper does not give good bounds immediately, and needs to be subjected to some post-processing for this purpose. We are currently working out the details.

**Acknowledgement** The authors thank Sariel Har-Peled for his extensive contribution.

## References

- [1] A. Aggarwal and J. S. Vitter. The input/output complexity of sorting and related problems. *Commun. ACM*, 31:1116–1127, 1988.
- [2] L. Arge and J. Vahrenhold. I/O-efficient dynamic planar point location. In *Proc. 16th Annu. ACM Symp. Comput. Geom.*, pages 191–200, 2000.
- [3] L. Arge, D. E. Vengroff, and J. S. Vitter. External-memory algorithms for processing line segments in geographic information systems. In *Proc. 3th Annu. Eur. Symp. Algorithms*, pages 295–310, 1995.
- [4] S. Arya and D. Mount. Approximate range searching. *Comput. Geom. Theory Appl.*, 17:135–152, 2000.
- [5] M. Bender, R. Cole, and R. Raman. Exponential structures for efficient cache-oblivious algorithms. In *Proc. 29th Internat. Colloq. Automata Lang. Prog.*, volume 2380 of *LNCS*, pages 195–207, Berlin, 2002.
- [6] M. A. Bender, E. D. Demaine, and M. Farach-Colton. Cache-oblivious B-trees. In *Proc. 41th Annu. IEEE Symp. Found. Comput. Sci.*, pages 339–409, 2000.
- [7] G. S. Brodal, R. Fagerberg, and R. Jacob. Cache oblivious search trees via binary trees of small height. In *Proc. 13th ACM-SIAM Symp. Discrete Algorithms*, pages 39–48, 2002.
- [8] Y.-J. Chiang, M. T. Goodrich, E. F. Grove, R. Tamassia, D. E. Vengroff, and J. S. Vitter. External-memory graph algorithms. In *Proc. 6th ACM-SIAM Symp. Discrete Algorithms*, pages 139–149, 1995.
- [9] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press / McGraw-Hill, Cambridge, Mass., 2001.
- [10] A. Crauser, P. Ferragina, K. Mehlhorn, U. Meyer, and E. Ramos. Randomized external-memory algorithms for some geometric problems. *Comput. Geom. Theory Appl.*, 11(3):305–337, June 2001.
- [11] M. de Berg. Linear size binary space partitions for uncluttered scenes. *Algorithmica*, 28:353–366, 2000.
- [12] M. de Berg. Improved bounds on the union complexity of fat objects. In *Proc. 25th Conf. Found. Soft. Tech. Theoret. Comput. Sci.*, pages 116–127, 2005.
- [13] M. de Berg, M. J. Katz, A. v. Stappen, and J. Vleugels. Realistic input models for geometric algorithms. *Algorithmica*, 34:81–97, 2002.
- [14] M. de Berg, M. van Kreveld, M. H. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2nd edition, 2000.
- [15] U. Finke and K. Hinrichs. Overlaying simply connected planar subdivisions in linear time. In *Proc. 11th Annu. ACM Symp. Comput. Geom.*, pages 119–126, 1995.
- [16] M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran. Cache-oblivious algorithms. In *Proc. 40th Annu. IEEE Symp. Found. Comput. Sci.*, pages 285–298, 1999.
- [17] I. Gargantini. An effective way to represent quadtrees. *Commun. ACM*, 25(12):905–910, 1982.
- [18] M. T. Goodrich, J.-J. Tsay, D. E. Vengroff, and J. S. Vitter. External-memory computational geometry. In *Proc. 34th Annu. IEEE Symp. Found. Comput. Sci.*, pages 714–723, 1993.
- [19] G. R. Hjaltason and H. Samet. Speeding up construction of pmr quadtree-based spatial indexes. *VLDB Journal*, 11:190–137, 2002.
- [20] H. Samet. *Spatial Data Structures: Quadtrees, Octrees, and Other Hierarchical Methods*. Addison-Wesley, Reading, MA, 1989.

## A The full proof of Lemma 2.3

*Proof:* Since the degree of each vertex in  $\mathcal{F}$  is at most  $2\pi/\delta$ , the stopping rule immediately implies that each cell is intersected by at most that many triangles. In the remainder of the proof we bound the number of triangle-cell intersections, from which the bound on the number of cells follows immediately.

For a vertex  $v$  of  $\mathcal{F}$ , let  $d_v$  be the closest distance to an edge opposite to  $v$ , that is, an edge of a triangle incident to  $v$  that is not incident to  $v$  itself. Let  $D_v$  be the disk centered at  $v$  and of radius  $r_v := d_v/2$ . Since the minimum angle in any triangle is  $\delta$ , we know that  $d_v$  is at least  $s_v \sin \delta$  where  $s_v$  is the length of the shortest edge incident to  $v$ .

We now bound the length ratio of the edges around  $v$ . Let  $\phi_1, \phi_2, \dots, \phi_k$  be the angles between the edges around  $v$  when going around  $v$  from the shortest edge to the longest edge incident to  $v$ —in clockwise or counterclockwise order such that  $\sum_{i=1}^k \phi_i \leq \pi$ . Using the sine rule and the fact that the minimum angle in any triangle is  $\delta$ , we find that the ratio between lengths of the edges that make angle  $\phi_i$  is at most:

$$\max_{\chi \in [\delta, \pi - \delta - \phi_i]} \frac{\sin(\pi - \phi_i - \chi)}{\sin \chi} = \max_{\chi \in [\delta, \pi - \delta - \phi_i]} \sin \phi_i / \tan \chi + \cos \phi_i = \sin \phi_i / \tan \delta + \cos \phi_i.$$

The maximum ratio of the length of any two edges around  $v$  is now:

$$\max_{\phi_1, \dots, \phi_k} \prod_{i=1}^k (\sin \phi_j / \tan \delta + \cos \phi_j) = \max_{\phi \geq \delta} (\sin \phi / \tan \delta + \cos \phi)^{\lfloor \pi / \phi \rfloor} \leq (2 \cos \delta)^{\pi / \delta} < 2^{\pi / \delta}.$$

Hence, for any edge  $e$  incident to  $v$  we have

$$r_v = \frac{d_v}{2} \geq \frac{s_v \sin \delta}{2} \geq \frac{\sin \delta}{2^{\pi / \delta + 1}} \cdot |e|,$$

where  $|e|$  denotes the length of  $e$ .

Let  $\sigma$  be a square in the quadtree subdivision. We will charge each intersection of a cell-triangle pair  $(\sigma, \Delta)$  either to a vertex or an edge of  $\mathcal{F}$ , as follows.

If  $\sigma$  intersects some disk  $D_v$ , we charge  $(\sigma, \Delta)$  to  $v$ . (If it intersects more disks, we take an arbitrary one and charge  $(\sigma, \Delta)$  to the vertex defining that disk.) Since  $\text{mom}(\sigma)$  cannot be completely contained in the union of the triangles incident to  $v$ —otherwise  $\text{mom}(\sigma)$  would not have been split by our stopping rule—we have  $\text{diam}(\text{mom}(\sigma)) \geq r_v$ . Hence,  $\text{diam}(\sigma) \geq r_v/2$ . From Lemma 2.2 we now know that  $v$  is charged by intersections with  $O(1)$  cells in this manner. Since each cell intersects at most  $2\pi/\delta$  triangles, the total number of cell-triangle intersections charged in this way is  $O(n/\delta)$ .

Now assume  $\sigma$  does not intersect any disk  $D_v$ . If  $\sigma$  does not intersect any edge, we charge  $(\sigma, \Delta)$  to the edge  $e$  of  $\Delta$  that is closest to  $\sigma$ , breaking ties arbitrarily. If  $\sigma$  intersects one or more edges, we charge  $(\sigma, \Delta)$  to an edge  $e$  of  $\Delta$  that intersects  $\sigma$ —thus charging at most two such pairs  $(\sigma, \Delta)$  to  $e$  for any fixed  $\sigma$ . From at least one of the two triangles incident to  $e$ , the square  $\text{mom}(\sigma)$  must intersect an edge different from  $e$ . Let  $v_1, v_2$  be the two vertices of  $e$ , and let  $v_3$  be the third vertex of this triangle  $\Delta'$ . Consider the Voronoi diagram of the edges of  $\Delta'$ , inside  $\Delta'$ . Because  $\sigma$  is charged to  $e$ , it must intersect the Voronoi cell  $\mathcal{V}(e)$  of  $e$ .

We partition  $\mathcal{V}(e)$  into a number of regions, and assign  $(\sigma, \Delta)$  to one of these regions, as follows. Extend the bisector at  $v_3$  (which separates  $\mathcal{V}(v_1 v_3)$  from  $\mathcal{V}(v_2 v_3)$ ) into  $\mathcal{V}(e)$ , thus splitting  $\mathcal{V}(e)$  into two regions. At least one of these regions intersects  $\sigma$ —in Fig. 1 this region is shown in grey. We partition this region,  $R$ , into subregions  $R_1, R_2, \dots, R_m$  as follows. Let  $p_0$  be the Voronoi vertex inside  $\Delta'$ . Let  $p_1$  be the midpoint of  $v_1 p_0$ . Region  $R_1$  is obtained by drawing a segment from  $p_1$  orthogonal to  $v_1 p_0$  into  $R$ . Region  $R_2$  is obtained by splitting the remainder of  $R$  by a segment from the midpoint of  $v_1 p_1$  and orthogonal to  $v_1 p_0$ , etc. This process is repeated until the remaining region is completely contained in  $D_v$ —see Fig. 1.




Figure 1: Illustration for the proof of Lemma 2.3.

Because both  $e$  and  $v_1p_0$  have length at most  $2r_v/(\sin(\delta)/2^{\pi/\delta+1})$ , the number of regions we get is

$$O\left(\log\left(\frac{2^{\pi/\delta+1}}{\sin\delta}\right)\right) = O\left(\frac{1}{\delta}\right).$$

We assign  $(\sigma, \Delta)$  to the largest subregion  $R_i$  (the one with smallest index) that it intersects. In the remainder of this proof, let us define the length of  $v_1p_0$  as the unit distance 1. Consider the region  $R_i$ . The point  $p_i$  has distance  $(1/2)^i$  to  $v_i$  and the edge  $p_i p_{i-1}$  has length  $(1/2)^i$ . Let  $\delta_1$  be the angle  $\angle p_0 v_1 v_2$ , which is equal to  $\angle v_3 v_1 p_0$ . The distance of any point in  $R_i$  to any edge of  $\Delta'$  is at least  $(1/2)^i \sin \delta_1$ , which is thus a lower bound on  $\text{diam}(\text{mom}(\sigma))$ . Hence,

$$\text{diam}(\sigma) \geq (1/2)^{i+1} \sin(\delta_1)$$

Now consider the region  $R_i^*$  obtained by taking the Minkowski sum of  $R_i$  with a disk of radius  $(1/2)^{i+1} \sin(\delta_1)$ . We have:

$$\text{area}(R_i^*) < (1/2)^{2i} (1 + \sin \delta_1) \cdot (2 \tan \delta_1 + \sin \delta_1).$$

Since  $\sigma$  intersects  $R_i$  and has width at least  $(1/2)^{i+1} \sin(\delta_1)/\sqrt{2}$ , while any square of that width that intersects  $R_i$  lies completely inside  $R_i^*$ , we have:

$$\text{area}(\sigma \cap R_i^*) \geq (1/2)^{2i+3} \sin^2 \delta_1.$$

Note that  $\delta/2 \leq \delta_1 \leq \pi/2 - \delta$ , and thus  $\sin \delta_1 > \sin \delta_1 \cos \delta_1 \geq \frac{1}{2} \sin \delta$  and  $\cos \delta_1 \geq \sin \delta$ . Hence,  $R_i$  gets assigned at most  $\text{area}(R_i^*)/\text{area}(\sigma \cap R_i^*) = O(1/\sin \delta) = O(1/\delta)$  cells. Since there are  $O(1/\delta)$  regions  $R_i$ ,  $e$  gets assigned  $O(1/\delta^2)$  different cells, and for each cell, at most two cell-triangle intersections are charged to  $e$ . Thus there are  $O(n/\delta^2)$  cell-triangle intersections in total.  $\blacksquare$