

# Chapter 9

## Taking a Walk in a Planar Arrangement

### 9.1 Introduction

Let  $\hat{S}$  be a set of  $n$   $x$ -monotone arcs in the plane, each pair of which intersect in at most  $t$  points. Computing the whole (or parts of the) arrangement  $\mathcal{A}(\hat{S})$ , induced by the arcs of  $\hat{S}$ , is one of the fundamental problems in computational geometry, and has received a lot of attention in recent years [SA95]. One of the basic techniques used for such constructions is based on *randomized incremental* construction of the *vertical decomposition* of the arrangement (see Chapter 8 and [BY98]).

If we are interested in only computing parts of the arrangement (e.g., a single face or a zone), the randomized incremental technique can still be used, but it requires non-trivial modifications [CEG<sup>+</sup>93, BDS95]. Intuitively, the added complexity is caused by the need to “trim” parts of the plane as the algorithm advances, so that it will not waste energy on regions which are no longer relevant. In fact, this requirement implies that such an algorithm has to know in advance what are the regions we are interested in at any stage during the randomized incremental construction.

A variation of this theme, with which the existing algorithms cannot cope efficiently, is the following *online* scenario: We start from a point  $p = p(0) \in \mathbb{R}^2$ , and we find the face  $f$  of  $\mathcal{A}(\hat{S})$  that contains  $p(0)$ . Now the point  $p$  starts moving and traces a connected curve  $\{p(t)\}_{t \geq 0}$ . As our walk continues, we wish to keep track of the face of  $\mathcal{A}(\hat{S})$  that contains the current point  $p(t)$ . The collection of these faces constitutes the *zone* of the curve  $p(t)$ . However, the function  $p(t)$  is not assumed to be known in advance, and it may change when we cross into a new face or abruptly change direction in the middle of a face (see [BDH99] for an application where such a scenario arises). The only work we are aware of that can deal with this problem efficiently is due to Overmars and van Leeuwen [OvL81], and it only applies to the case of lines (and, with some simple modifications, to certain restricted cases involving line segments as well).<sup>1</sup> It can compute such a walk in (deterministic)  $O((n+m) \log^2 n)$  time, inside an arrangement of  $n$  lines, where  $m$  is the number of intersections of the walk with the lines of  $\hat{S}$ . This is done by maintaining dynamically the intersection of half-planes that corresponds to the current face.

In this chapter, we propose a new randomized algorithm that computes the zone of the walk in a general arrangement of arcs, as above, in  $O(\lambda_{t+2}(n+m) \log n)$  expected time, where  $\lambda_{t+2}(n+m)$  is the maximum length of a Davenport-Schinzel sequence of order  $t+2$  having  $n+m$  symbols [SA95]. The new algorithm can be interpreted as a third “online” alternative to the algorithms of [CEG<sup>+</sup>93, BDS95]. The algorithm is rather simple and appears to be practical. As a matter of fact, we had implemented and experimented with a variant of the algorithm [AHH<sup>+</sup>99].

---

<sup>1</sup>Very recently, an improvement has been given by Chan [Cha99a]; it can perform an update in  $O(\log^{1+\epsilon} n)$  amortized time, and answer queries (of the kind used in this application) in  $O(\log n)$  time per query.

As an application of the new algorithm, we present an algorithm for computing a level in an arrangement of arcs. It computes a single level in  $O(\lambda_{t+2}(n+m) \log n)$  expected time, where  $m$  is the complexity of the level. We also show how to adapt the main algorithm to obtain a point-location algorithm that locates  $m$  points in an arrangement of  $n$  arcs, as above, in expected time  $O(\lambda_{t+2}(n+m+w) \log n)$ , where  $w$  is the minimum number of intersections between a spanning tree connecting those query points and the given arcs.

Both results improve by almost a logarithmic factor over the best previous result of [OvL81], for the case of lines (and for certain cases involving line segments).<sup>2</sup> For the case of general arcs, we are not aware of any similarly efficient previous result.

The chapter is organized as follows. In Section 9.2 we describe the algorithm. In Section 9.3 we analyze its performance. In Section 9.4 we mention a few applications of the algorithm, including the construction of a single level, and multiple point-location. Concluding remarks are given in Section 9.5.

## 9.2 The Algorithm

In this section, we present the algorithm for performing an online walk inside a planar arrangement.

**Randomized Incremental Construction of the Zone Using an Oracle.** Given a set  $\hat{S}$  of  $n$   $x$ -monotone arcs in the plane, so that any pair of arcs of  $\hat{S}$  intersect at most  $t$  times (for some fixed constant  $t$ ), let  $\mathcal{A}(\hat{S})$  denote the arrangement of  $\hat{S}$ ; namely, the partition of the plane into faces, edges, and vertices as induced by the arcs of  $\hat{S}$  (see [SA95] for details). We assume that  $\hat{S}$  is in general position, meaning that no three arcs of  $\hat{S}$  have a common point, and that the  $x$ -coordinates of the intersections and endpoints of the arcs of  $\hat{S}$  are pairwise distinct. As in Chapter 8, the *vertical decomposition* of  $\mathcal{A}(\hat{S})$ , denoted by  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$ , is the partition of the plane into vertical pseudo-trapezoids, obtained by erecting two vertical segments up and down from each vertex of  $\mathcal{A}(\hat{S})$  (i.e., each point of intersection between a pair of arcs and each endpoint of an arc), and by extending each of them until it either reaches an arc of  $\hat{S}$ , or otherwise all the way to infinity. See, e.g., [BY98, SA95] and Chapter 8 for more details concerning vertical decompositions. To simplify (though slightly abuse) the notation, we refer to the cells of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  as *trapezoids*.

We repeat some notations from Chapter 8: A *selection*  $R$  of  $\hat{S}$  is an ordered sequence of distinct elements of  $\hat{S}$ . By a slight abuse of notation, we also denote by  $R$  the unordered set of its elements. Let  $\sigma(\hat{S})$  denote the set of all selections of  $\hat{S}$ . For a permutation  $S$  of  $\hat{S}$ , let  $S_i$  denote the subsequence consisting of the first  $i$  elements of  $S$ , for  $i = 0, \dots, n$ .

Computing the decomposed arrangement  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  can be done in a randomized incremental manner, as described in Section 8.2. Let  $\gamma$  be the curve traced by the walk. For a selection  $R \in \sigma(\hat{S})$ , let  $\mathcal{D}_\gamma(R)$  (resp.  $\mathcal{Z}_\gamma(R)$ ) denote the *district* (resp. *zone*) of  $\gamma$  in  $\mathcal{A}(R)$ ; these are, respectively, the set of all trapezoids of  $\mathcal{A}_{\mathcal{VD}}(R)$  and the set of all faces of  $\mathcal{A}(R)$  that have a nonempty intersection with  $\gamma$ . Let  $\mathcal{A}_{\gamma, \mathcal{VD}}(R)$  denote the set of all trapezoids in  $\mathcal{A}_{\mathcal{VD}}(R)$  that cover  $\mathcal{Z}_\gamma(R)$ . Our goal is to compute  $\mathcal{A}_{\gamma, \mathcal{VD}}(\hat{S})$ . (Alternatively, we may be interested only in computing the district  $\mathcal{D}_\gamma(\hat{S})$  of  $\gamma$ .)

We assume for the moment that we are supplied with an oracle  $\mathcal{O}(S_i, \gamma, \Delta)$ , that can decide in constant time whether a given vertical trapezoid  $\Delta$  is in  $\mathcal{A}_{\gamma, \mathcal{VD}}(S_i)$ . Equipped with this oracle, computing  $\mathcal{A}_{\gamma, \mathcal{VD}}(S)$  is fairly easy, using a variant of the randomized incremental construction, outlined above. The algorithm, called `CompZoneWithOracle`, is depicted in Figure 9.1. We present this algorithm at a

<sup>2</sup>Our results are also asymptotically faster and much simpler to implement than what is yielded by the recent results of Chan [Cha99a, Cha99b].

conceptual level only, because this is not the algorithm that we shall actually use. It is given to help us to describe and analyze the actual online algorithm that we shall present later.

Note that the set of trapezoids  $\mathcal{C}_i$  maintained by the algorithm in the  $i$ -th iteration is a superset of  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(S_i)$ . There might be trapezoids in  $\mathcal{C}_i$  that are no longer in  $\mathcal{Z}_\gamma(S_i)$  (typically these are trapezoids that are separated from  $\mathcal{Z}_\gamma(S_i)$  by an arc that does not cross their interior and is intersected after they have been created). However, this implies that any such trapezoid will be eliminated the first time an arc that crosses it will be handled, or, if no such arc exists, at the final clean-up step of the algorithm. Moreover, the algorithm `CompZoneWithOracle` can be augmented to compute a *history* DAG (as in [SA95]), whose nodes are the trapezoids created by the algorithm and where each trapezoid destroyed during the execution of the algorithm points to the trapezoids that were created from it. Let  $\mathcal{HT}_\gamma(S_i)$  denote this structure after the  $i$ -th iteration of the algorithm. Note that the out-degree of each node of  $\mathcal{HT}_\gamma$  is bounded by a constant that depends on  $t$ .

**Definitions.** A trapezoid created by the `SplitGeom` operation of `CompZoneWithOracle` is called a *transient* trapezoid if it is later merged (in the same iteration) to form a larger trapezoid. A trapezoid generated by `CompZoneWithOracle` is *final* if it is not transient. The *rank*  $\text{rank}(\Delta)$  of a trapezoid  $\Delta$  is the maximum of the indices  $i, j$  of the arcs containing the bottom and top edges of  $\Delta$  in the permutation  $S$ . We denote by  $D(\Delta)$  the *defining set of a final trapezoid*  $\Delta$ ; this is the minimal set  $D$  such that  $\Delta \in \mathcal{A}_{\mathcal{V}\mathcal{D}}(D)$ . It is easy to verify that  $|D(\Delta)| \leq 4$ . We can also define  $D(\Delta)$  for a transient trapezoid  $\Delta$ , to be the minimal set  $D$  such that  $\Delta$  can be transient during an incremental construction of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(D)$ . Here it is easy to verify that  $|D(\Delta)| \leq 6$ . The *index*  $\text{index}(\Delta)$  of a trapezoid  $\Delta$  is the minimum  $i$  such that  $D(\Delta) \subseteq S_i$ . For a trapezoid  $\Delta$ , we denote by  $\text{cl}(\Delta)$  the *conflict list* of  $\Delta$ ; that is, the set of arcs of  $\hat{S}$  that intersect  $\Delta$  in its interior. Let  $\text{next}(\Delta)$  denote the first element of  $\text{cl}(\Delta)$ , according to the ordering of  $S$ .

For a trapezoid  $\Delta$  generated by `CompZoneWithOracle`, which was not merged into a larger trapezoid, we denote by  $\text{father}(\Delta)$  the trapezoid that  $\Delta$  was generated from. A vertical side of a trapezoid  $\Delta$  is called a *splitter*. A splitter  $\nu$  is *transient* if it is not incident to the intersection point (or endpoint) that induced the vertical edge that contains  $\nu$  (this means that the two trapezoids adjacent to  $\nu$  are transient, and will be merged into a larger final trapezoid). See Figure 9.2 for an illustration of some of these definitions. It is easy to verify that a trapezoid  $\Delta$  is transient if and only if at least one of its bounding splitters is transient. Thus, one can decide whether a trapezoid is transient, by inspecting its splitters, in constant time.

The online algorithm constructs portions of  $\mathcal{HT}_\gamma(S)$  incrementally, as they are needed. This is done by performing a sequence of point locations in the arrangement, where each such query constructs the final trapezoid of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(\hat{S})$  that contains the query point, plus all ancestor trapezoids that lie on paths of  $\mathcal{HT}_\gamma(S)$  from the root to that trapezoid. Informally, instead of building  $\mathcal{HT}_\gamma(S)$  layer-by-layer, as is done by `CompZoneWithOracle`, the online algorithm constructs the DAG in an ‘orthogonal’ DFS manner. The intuition behind the design and efficiency of the algorithm is that the expected number of nodes in  $\mathcal{HT}_\gamma(S)$  is only  $O(\lambda_{t+2}(n+m))$ , whereas the overall expected running time of `CompZoneWithOracle` is  $O(\lambda_{t+2}(n+m) \log n)$  (this is the expected overall size of the conflict lists of the computed nodes, which the algorithm needs to construct and manipulate). Both bounds are easy consequences of known results, and will be discussed in Section 9.3.2. Thus, in our quest to usurp the oracle, we can afford to pay an extra  $O(\log n)$  time during the search for and construction of each node of  $\mathcal{HT}_\gamma(S)$ . This indeed will be the cost of a point-location query (ignoring the time required for constructing any new node of  $\mathcal{HT}_\gamma(S)$  that the query has to pass through). Informally, the online algorithm performs essentially the same operations as the preceding algorithm, except that it executes them in a different order, and,

```

ALGORITHM  CompZoneWithOracle( $\hat{S}, \gamma, \mathcal{O}$ )
  Input: A set  $\hat{S}$  of  $n$  arcs, a curve  $\gamma$ , an oracle  $\mathcal{O}$ 
  Output:  $\mathcal{A}_{\gamma, \mathcal{VD}}(\hat{S})$ 
begin
  Choose a random permutation  $S = \langle s_1, s_2, \dots, s_n \rangle$  of  $\hat{S}$ .
   $\mathcal{C}_0 \leftarrow \{\mathbb{R}^2\}$ 
  for  $i$  from 1 to  $n$  do
     $\mathcal{D}_i \leftarrow \left\{ \Delta \mid \Delta \in \mathcal{C}_{i-1}, \text{int } \Delta \cap s_i \neq \emptyset \right\}$ 
     $Temp \leftarrow \emptyset$ 
    for each  $\Delta \in \mathcal{D}_i$  do
       $Temp \leftarrow Temp \cup \text{SplitGeom}(\Delta, s_i)$ ,
      where  $\text{SplitGeom}(\Delta, s)$  is the operation of splitting a vertical
      trapezoid  $\Delta$  crossed by an arc  $s$  into a constant number of
      vertical trapezoids, as in [BCKO08] and Chapter 8, such that the
      new trapezoids cover  $\Delta$ , and they do not intersect  $s$  in their interior.
    end for
    Merge all the adjacent trapezoids of  $Temp$  that have the same top
    and bottom arcs. Let  $Temp_1$  be the resulting set of trapezoids.
    Let  $Temp_2$  be the set of all trapezoids of  $Temp_1$  that are in  $\mathcal{A}_{\gamma, \mathcal{VD}}(S_i)$ .
    Compute this set using  $|Temp_1|$  calls to  $\mathcal{O}$ .
     $\mathcal{C}_i \leftarrow (\mathcal{C}_{i-1} \setminus \mathcal{D}_i) \cup Temp_2$ 
  end for
  Remove from  $\mathcal{C}_n$  all trapezoids not belonging to  $\mathcal{A}_{\gamma, \mathcal{VD}}(\hat{S})$ , by checking
  with  $\mathcal{O}$  each trapezoid of  $\mathcal{C}_n$ .
  return  $\mathcal{C}_n$ 
end CompZoneWithOracle

```

Figure 9.1: A randomized incremental algorithm for constructing the zone of a walk in an arrangement of arcs, using an oracle

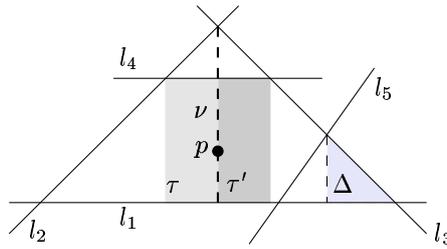


Figure 9.2: Illustration of the definitions:  $\nu$  is a transient splitter, and thus  $\tau, \tau'$  are both transient. We have  $\text{rank}(\tau) = \text{index}(\tau) = 4$ ,  $\text{rank}(\Delta) = 3$ , and  $\text{index}(\Delta) = 5$ , where  $S = \langle l_1, l_2, l_3, l_4, l_5 \rangle$ .

in addition, it may revisit again and again portions of the DAG that have already been constructed as it searches down the DAG while performing point locations. However, since this extra cost is only logarithmic, it does not increase the asymptotic complexity of the algorithm.

**An Online Algorithm for Constructing the Zone.** Let us assume that the random permutation  $S$  of  $\hat{S}$  has been fixed in advance. Note that  $S$  predetermines  $\mathcal{HT}_\gamma = \mathcal{HT}_\gamma(S)$ . The key observation in the online algorithm is that in order to construct a specific leaf of  $\mathcal{HT}_\gamma(S)$  we do not have to maintain the entire DAG, and it suffices to compute only the parts of the DAG that lie on paths connecting the leaf with the root of  $\mathcal{HT}_\gamma$  (there might be several such paths, since our structure is a DAG, and not a tree).

To facilitate this computation, we maintain a *partial history* DAG, denoted by  $T$ . The nodes of  $T$  are of two types: (i) *final nodes*: those are nodes whose corresponding trapezoids appear in  $\mathcal{HT}_\gamma = \mathcal{HT}_\gamma(S)$ , and (ii) *transient nodes*: these are some of the leaves of  $T$ , whose corresponding trapezoids are transient. A transient node can be easily detected since its trapezoid is transient. A final node is simply a node which is not transient. In particular, all the internal nodes of  $T$  are copies of identical nodes of  $\mathcal{HT}_\gamma$  (whose corresponding trapezoids are final), while some of the leaves of  $T$  might be transient. Intuitively,  $T$  stores the portion of  $\mathcal{HT}_\gamma$  that we have computed explicitly so far. The transient leaves of  $T$  delimit portions of  $\mathcal{HT}_\gamma$  that have not been expanded yet. Inside each node of  $T$ , we also maintain the conflict list of the corresponding trapezoid  $\Delta$  and its first element  $\text{next}(\Delta)$ .

Suppose we wish to compute a leaf of  $\mathcal{HT}_\gamma$  that contains a given point  $p$ . We denote this operation by  $\text{PointLocate}(p)$ . We first locate the leaf of  $T$  that contains  $p$ . This is done by traversing a path in  $T$ , starting from the root of  $T$ , and going downward in each step into the child of the current trapezoid that contains  $p$  (each such step requires  $O(1)$  time, because, as already noted, the out-degree of any node of  $\mathcal{HT}_\gamma$ , and thus of  $T$ , is bounded by a constant that depends on  $t$ ). (A technical issue that we face is that usually  $p$  lies on some trapezoid boundary, so we need additional local information to determine which child to descend to at each of the above steps — see below for more details) At the end we either reach a final leaf (with an empty conflict list) which is the required leaf of  $\mathcal{HT}_\gamma$ , or we encounter a leaf  $v$  of  $T$ . In the latter case, we need to expand  $T$  further below  $v$ . If  $v$  represents a transient trapezoid, then the first step is to replace  $v$  by the corresponding node  $v^*$  of  $\mathcal{HT}_\gamma$ , obtained by merging the transient trapezoid of  $v$  with adjacent transient trapezoids, with identical top and bottom arcs, to form the final trapezoid associated with  $v^*$ . If  $v$  is a final node we expand it by splitting it with the first arc that crosses its trapezoid, using steps (iv) and (v) below.

Assume for the moment that we are supplied (for the case where  $v$  is transient) with a method (to be described shortly) to generate all those adjacent transient trapezoids, whose union forms the final trapezoid that is stored at  $v^*$  in  $\mathcal{HT}_\gamma$ . Then we do the following: (i) Merge all those transient trapezoids into a new (final) trapezoid  $\Delta$ ; (ii) Compute the conflict list  $\text{cl}(\Delta)$  from the conflict lists of the transient

trapezoids; (iii) Compute the first element  $s_\Delta = \text{next}(\Delta)$  in  $\text{cl}(\Delta)$  according to the permutation  $S$ ; (iv) Compute all the transient or final trapezoids generated from  $\Delta$  by splitting it by  $s_\Delta$  (this generates  $O(1)$  new trapezoids); and (v) Extract from  $\text{cl}(\Delta)$  the conflict list  $\text{cl}(\Delta')$  of each new trapezoid  $\Delta'$ , and compute  $\text{next}(\Delta')$  as well.

Overall, this requires  $O(k + l)$  time, where  $k$  is the number of transient trapezoids that are merged, and  $l$  is the total length of the conflict lists of these transient trapezoids. This is trivial to show for steps (i), (iii), (iv) and (v). To perform in step (ii) the merging of the conflict lists in linear time, one may use a global bit-vector structure. Namely, we initialize before the execution of the algorithm a bit-vector  $b$  of size  $n$  to be everywhere zero. To merge several conflict lists  $L_1, \dots, L_k$ , we scan each list in turn, and for each of its elements  $s_j$ , we first test whether  $b_j = 0$ ; if so, we add  $s_j$  to the output conflict list, and change  $b_j$  to 1. After creating the output conflict list in this manner, we scan the output list, turning off all the bits that got turned on.

In this manner, we have upgraded a transient leaf  $v$  of  $T$  into a final node  $v^*$ . We denote this operation by **Expand**( $v$ ). We can now continue going down in  $T$ , passing to the child of  $\Delta$  that contains  $p$  and repeating recursively the above procedure at that child, until constructing and reaching the desired leaf of  $\mathcal{HT}_\gamma$  that contains  $p$  (namely, until we reach a node that is final and had empty conflict list).

To complete the presentation of this point location mechanism, we describe **Expand**( $v$ ), the procedure that computes the ‘sibling’ transient trapezoids that are adjacent to the transient trapezoid of  $v$ .

Let  $\tau$  be the transient trapezoid. Then either the top arc or the bottom arc of  $\tau$  are the cause of the splitting that generated  $\tau$ . In particular,  $\text{next}(\tau_f)$  is either the top or bottom arc of  $\tau$ , where  $\tau_f = \text{father}(\tau)$  denotes the trapezoid that  $\tau$  was generated (split) from. This also implies that  $\text{rank}(\tau) = \text{index}(\tau)$ . Since  $\tau$  is transient, one of the splitters of  $\tau$  must be transient. Let  $\nu$  denote such a transient splitter, and let us assume that  $\nu$  is the right edge of  $\tau$ . Note also that  $\tau_f$  must be a final trapezoid, so in particular  $\nu$  was generated from a final splitter (by the insertion of an arc that separated  $\nu$  from the vertex that induced the bigger final splitter).

We compute the transient trapezoid  $\tau'$  that lies to the right of  $\tau$  and has the same top and bottom arcs, by taking the midpoint  $p$  of  $\nu$ , and by performing a point-location query of  $p$  in  $T$  using (recursively) the same mechanism described above. During this point-location process, we always go down into the trapezoid  $\Delta$  that contains  $p$  in its interior or on its left edge; see below for details. We stop as soon as we encounter a transient trapezoid  $\tau'$  that has a left edge identical to  $\nu$ . This happens when  $\tau$  and  $\tau'$  have the same top and bottom edges; namely, we stop when  $\text{rank}(\tau) = \text{rank}(\tau')$ . (Intuitively, if the trapezoid  $\tau'$  has rank smaller than  $\text{rank}(\tau)$ , then either it fully contains  $\nu$  in its interior, or its left edge is longer than (and contains)  $\nu$ ; the first time when both  $\tau$  and  $\tau'$  have the same connecting edge is when their top and bottom edges are identical, namely, when  $\text{rank}(\tau) = \text{rank}(\tau')$ .) See below for more details in the proof of correctness of the algorithm. We continue this process of collecting adjacent transient trapezoids using point-location queries on midpoints of transient splitters, until the two extreme splitters (the left splitter of the leftmost trapezoid in the sequence and the right splitter of the rightmost trapezoid) are non-transient. We take the union of those trapezoids to be the new expanded trapezoid. See [Figure 9.2](#) for a scenario where such a merging occurs. A more detailed illustration is given in [Appendix 9.B](#) below.

Of course, during this point-location process, we might be forced into going into parts of  $\mathcal{HT}_\gamma$  that are rather remote from the point  $p$ . In such a case, we will compute those parts in an online manner, by performing **PointLocate** and **Expand** calls on the relevant trapezoids that we encounter while going down  $T$ . Thus, the point-location process is recursive, and might be quite substantial. Nevertheless, as will be argued below, the overall cost of the **PointLocate** and **Expand** operations is not excessive, so these operations are efficient in an amortized sense.

Let  $\gamma$  be the curve of the online walk whose zone we wish to compute. We consider  $\gamma$  to be a directed curve, supplied to us by the user through a function **EscapePoint** $_\gamma$ . The function **EscapePoint** $_\gamma(p, \Delta)$

receives as input a point  $p \in \gamma$ , and a trapezoid  $\Delta$  that contains  $p$ , and outputs the next intersection point of  $\gamma$  with  $\partial\Delta$  following  $p$ . If  $\gamma$  ends before we reach  $\partial\Delta$ , the function returns `nil`. We assume (although this is not crucial for the algorithm) that  $\gamma$  does not intersect itself.

Let  $G_S$  denote the adjacency graph of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$ . This is a graph having a vertex for each trapezoid in  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$ , such that an edge connects two vertices if their corresponding trapezoids share a common vertical side. Under general position assumptions, each vertex in  $G_S$  has degree at most 4. It is easy to verify that a connected component of  $G_S$  corresponds to a face of  $\mathcal{A}(\hat{S})$ . Given a final leaf-trapezoid  $\Delta$ , we can compute the face of  $\mathcal{A}(\hat{S})$  that contains  $\Delta$  by performing a DFS in  $G_S$ . This is done by performing point-location queries on appropriate points on the splitters of  $\Delta$ , in a manner similar to that used in the `Expand` operation. This yields all the neighbors of  $\Delta$  in  $G_S$ , and we continue in this manner until the entire connected component of  $G_S$  containing  $\Delta$  is constructed.

Thus, given a walk  $\gamma$ , we can compute its zone by the algorithm `CompZoneOnline` depicted in [Figure 9.3](#). See [Appendix 9.A](#) for the pseudo-code of the main subroutines used by `CompZoneOnline`, and [Appendix 9.B](#) for an illustration of the execution of `CompZoneOnline`.

As will be shown in [Section 9.3.1](#), by the time the algorithm terminates, the final parts of  $T$  are contained in  $\mathcal{HT}_\gamma$ . (A proper inclusion might arise; see [Remark 9.3.13](#).) In analyzing the performance of the algorithm, we first bound the overall expected time required to compute  $\mathcal{HT}_\gamma$ , which can be done by bounding the expected running time of `CompZoneWithOracle` (in an appropriate model of computation). Next, we will bound the additional time spent by the algorithm in traversing between adjacent trapezoids (i.e., the additional time spent in performing the point-location queries).

**Remark 9.2.1** By skipping the expansion of the face that contains the current point  $p$  in `CompZoneOnline`, we get a more efficient algorithm that only computes the district  $\mathcal{D}$  of the walk, that is, the collection of trapezoids in  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that  $\gamma$  crosses. There are cases where this will be sufficient; see [Section 9.4](#) (e.g., in the adaptation of the algorithm for computing a level in an arrangement).

## 9.3 Analysis of `CompZoneOnline`

### 9.3.1 Correctness

In this subsection, we establish the correctness of `CompZoneOnline`. Before starting, we note that the correctness of `CompZoneWithOracle` is easier to establish, and follows routinely from standard considerations. We therefore omit any further details concerning this issue.

The main technical issues that arise in the proof of correctness have to do with the potentially complex patterns of exploring the DAG  $T$  that can arise during the recursive execution of the `PointLocate` and `Expand` routines. The first main step in the proof is to show that the `Expand` routine always terminates properly, with a transient trapezoid that is compatible with the input one. Once this is shown, the rest of the proof is a routine, though somewhat involved, task, which employs induction on the structure of  $T$  and on the sequence of steps executed by the algorithm.

**Lemma 9.3.1** *During the execution of `CompZoneOnline`, the union of trapezoids of the leaves of  $T$  form a pairwise disjoint covering of the plane by vertical trapezoids.*

*Proof:* By induction on the steps of `CompZoneOnline`, noting that this is true initially, and that each step that modifies  $T$  either merges leaf trapezoids into a larger leaf trapezoid or splits a leaf trapezoid into subtrapezoids. ■

```

ALGORITHM  CompZoneOnline( $\hat{S}$ ,  $p$ , EscapePoint $_{\gamma}$ )
  Input: A set  $\hat{S}$  of  $n$  arcs, a starting point  $p$  of the walk,
           and a function EscapePoint $_{\gamma}$  that represents the walk
  Output: The decomposed zone  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(\hat{S})$  of  $\gamma$  in  $\mathcal{A}(\hat{S})$ 
begin
  Choose a random permutation  $S = \langle s_1, s_2, \dots, s_n \rangle$  of  $\hat{S}$ .
   $T \leftarrow \{(\mathbb{R}^2, S)\}$  - a partial history DAG with a root corresponding to
    the whole plane; the conflict list of the root is the whole  $S$ .
   $v \leftarrow \text{PointLocate}(p, \cdot)$ ,
    where PointLocate( $p, \cdot$ ) is the leaf of  $\mathcal{HT}_{\gamma}$  whose associated trapezoid
    contains  $p$ . (The procedure has an additional flag parameter
    that we ignore here; it is used in cases where  $p$  lies on trapezoid
    boundaries; see Appendix 9.A and below.)
  /* All the paths in  $\mathcal{HT}_{\gamma}$  from  $v$  to the root now exist in  $T$ . */
   $\mathcal{D} \leftarrow \{\Delta_v\}$  ( $\Delta_v$  is the trapezoid stored at  $v$ .)
  while ( $p \neq \text{nil}$ ) do
     $p \leftarrow \text{EscapePoint}_{\gamma}(p, \Delta_v)$ 
     $w \leftarrow \text{PointLocate}(p, +)$ ,
      where  $(p, +)$  denotes a point  $p^+$  on  $\gamma$  just past  $p$ , and  $w$  is the next leaf
      of  $\mathcal{HT}_{\gamma}$ , such that  $p^+ \in \Delta_w$ . This is done by performing a point-location
      query in  $T$ , as described in the text, and expanding  $T$  accordingly.
     $v \leftarrow w$ 
     $\mathcal{D} \leftarrow \mathcal{D} \cup \{\Delta_v\}$ 
  end while
  if only the district of  $\gamma$  needs to be computed then
    return  $\mathcal{D}$  (the district of  $\gamma$  in  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(\hat{S})$ )
   $\mathcal{Z} \leftarrow \emptyset$ 
  for each  $\Delta \in \mathcal{D}$ 
    Compute the face  $F$  of  $\Delta_v$  in  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(S)$  (if it had not yet been computed).
     $\mathcal{Z} \leftarrow \mathcal{Z} \cup F$ 
  end for
  return  $\mathcal{Z}$ .
end CompZoneOnline

```

Figure 9.3: Algorithm for constructing the zone of a walk in an arrangement of arcs in an online manner. See [Appendix 9.A](#) for the pseudo-code of the main subroutines used by CompZoneOnline, and [Appendix 9.B](#) for an illustration of the execution of CompZoneOnline.

**Corollary 9.3.2** *Each conflict list, as computed for a (transient or final) trapezoid  $\Delta$  by the procedure CompZoneOnline, is the list of all arcs of  $S$  that cross (the interior of)  $\Delta$ .*

*Proof:* By induction on the steps of CompZoneOnline. Observe that the region(s) that  $\Delta$  was generated from cover  $\Delta$ , and thus the union of their conflict lists must contain, by the induction hypothesis, the correct conflict list of  $\Delta$ , which is thus correctly constructed by the appropriate Expand or PointLocate step. ■

**Corollary 9.3.3** *For a trapezoid  $\Delta$  created by CompZoneOnline, all the arcs of  $D(\Delta)$  appear in  $S$  before all the arcs of  $\text{cl}(\Delta)$ .*

Consider an Expand operation that is triggered by point location at the midpoint  $p$  of a transient splitter  $\nu$  that bounds a transient trapezoid  $\tau$ , where  $\tau$  has already been generated in  $T$  and we wish to find the transient trapezoid  $\tau'$  that shares  $\nu$  with  $\tau$  as a common splitter. (It is easily verified that  $\tau'$  is uniquely defined, given the permutation  $S$ .) Let  $i$  denote the rank of  $\tau$ . As noted,  $i$  is also the index of  $\tau$ . Clearly  $s_i$  must be either the top or the bottom arc of  $\tau$ . Assume, without loss of generality, that  $\nu$  is the right splitter of  $\tau$  and that  $s_i$  is the top edge of  $\tau$ . Let  $s_j$  be the bottom arc of  $\tau$ , with  $j < i$ .

**Lemma 9.3.4** *During the execution of this Expand step, ignoring recursive calls, all the trapezoids that are either visited or generated fully contain  $\nu$  either in their interior or within their left edge. Consequently, for any such trapezoid, except for the last one, either its conflict list contains  $s_i$  and  $s_j$ , or it contains  $s_i$  and  $s_j$  is the bottom arc of the trapezoid. Moreover, if this sequence of trapezoids includes a trapezoid that does contain  $\nu$  on its left edge then all subsequent trapezoids in the sequence have this property.*

*Proof:* This is shown by induction on the sequence of steps of this (nonrecursive portion of the) execution of Expand. For this proof, we assume that during this execution, any recursive call to Expand terminates correctly, with a transient trapezoid that is adjacent to the trapezoid that initiated the recursive call and has the same top and bottom arcs. If this is not the case, we assume that the algorithm aborts at that point, and from that point on there is nothing to prove. (We thus modify the algorithm for the purpose of the proof, but a subsequent argument (in Lemma 9.3.5 below) will show that the algorithm never aborts.)

The first node that the Expand procedure visits is the root of  $T$ , and the claims clearly hold in this case. Assume they hold for all trapezoids up to and including a trapezoid  $\Delta$ . Suppose first that  $\Delta$  is final, and let  $s_k = \text{next}(\Delta)$ . By induction hypothesis, the conflict list of  $\Delta$  contains  $s_i$ , so we must have  $k \leq i$ . Then Corollary 9.3.3 implies that  $s_k$  does not cross the relative interior of  $\nu$ . This implies that the subtrapezoid  $\Delta'$  of  $\Delta$  that is split from it by  $s_k$  and contains (a point slightly to the right of) the midpoint of  $\nu$  must fully contain  $\nu$  in its closure. If  $\Delta$  contained  $\nu$  on its left side then clearly this also holds for  $\Delta'$ .

If  $\Delta$  is transient then there are two subcases: If the top and bottom edges of  $\Delta$  are, respectively,  $s_i$  and  $s_j$ , then the Expand procedure terminates and returns  $\Delta$ ; the claims clearly hold in this case. Otherwise, the Expand procedure executes a recursive call with the midpoint of some (transient) splitter of  $\Delta$ . As argued above, we may assume that this recursive call returns a transient trapezoid compatible with  $\Delta$ , in the above sense. We repeat this step, if needed, until we obtain a sequence of compatible transient trapezoids, including  $\Delta$ , which cannot be expanded any further. We then merge all these trapezoids into a trapezoid  $\Delta'$ , which clearly must be final. It is obvious, by construction, that  $\Delta'$  satisfies the first assertion of the lemma.

To prove the second assertion (for transient trapezoids), assume that  $\Delta$  contains  $\nu$  on its left edge  $\nu_0$ . It suffices to show that  $\nu_0$  is not transient (which will imply that  $\Delta$  is not merged with other trapezoids

on its left, so that this left splitter will be contained in  $\nu_0$ ), so assume to the contrary that it is transient. As already noted, a transient splitter like  $\nu_0$  is generated when we insert an arc  $s_\ell$  that delimits  $\nu_0$  and separates it from the vertex  $w$  that induced it. Once this occurs,  $\nu_0$  will trigger a call to **Expand** which, if properly terminated, will erase  $\nu_0$  as it merges  $\Delta$  with adjacent transient trapezoids. As argued above, we have  $\ell < i$  (because  $s_i \in \text{cl}(\Delta)$ ).

The same argument also implies that the trapezoid  $\tau_f$  from which  $\tau$  has been split (by the insertion of  $s_i$ ) has a non-transient right splitter  $\nu^*$  that contains  $\nu$  and extends all the way to the vertex  $w$ . Thus, by **Corollary 9.3.2**, the conflict list of  $\tau_f$  contains  $s_\ell$ , so  $\text{next}(\tau_f)$  cannot be  $s_i$ , contrary to assumption.

This induction step completes the proof of the lemma.  $\blacksquare$

As a consequence, it is easily verified by induction that, during the execution of this **Expand** step, including all recursive calls, no arc  $s_k$  with  $k > i$  is processed.

We now show that each call to **Expand** terminates correctly.

**Lemma 9.3.5** *Each point-location query at the midpoint of a transient splitter generates a “compatible” transient trapezoid; that is, a transient trapezoid adjacent to the current transient trapezoid, that has the same top and bottom arcs.*

*Proof:* Suppose the lemma is false, and consider the first call where this happens, where we order the calls in the order of returns from **Expand** (i.e. in postorder on the recursion forest). Let  $\Delta$  be the transient trapezoid that initiated this call, and assume, without loss of generality, that the call started at the right splitter  $\nu$  of  $\Delta$ , and that the top and bottom edges of  $\Delta$  are, respectively,  $s_i$  and  $s_j$ , with  $i > j$ . Arguing as in the proof of **Lemma 9.3.4**, we have that during the execution of the nonrecursive portion of this call, the procedure visits or generates a sequence  $\Sigma$  of trapezoids, each of which contains  $\nu$ , and, by the induction hypothesis, all recursive calls that it executes terminate properly.

Let  $s_k, s_\ell$ , with  $k < \ell$ , be the two arcs that intersect at the vertex  $w$  that induced a splitter that contains  $\nu$  (the case where  $w$  is an endpoint of an arc  $s_k$  is handled similarly). Clearly, we must have  $k < \ell < i$ . Moreover,  $w$  must lie above  $\nu$ , or else the insertion of  $s_j$  would have disconnected  $\nu$  from  $w$ ; since  $s_i$  has not yet been inserted at that stage, it easily follows that  $\nu$  could not have been formed at all. It is easily seen that any trapezoid  $\tau \in \Sigma$  that contains  $\nu$  in its interior must either contain  $s_k$  and  $s_\ell$  in its conflict list or be bounded by  $s_k$  and contain  $s_\ell$  in its conflict list. Hence, as can be easily verified,  $s_k$  and  $s_\ell$  will eventually be processed in splitting operations during this execution, and will consequently generate trapezoids in  $\Sigma$  that contain  $\nu$  on their left edge. Moreover, either  $s_i$  and  $s_j$  appear in the conflict list of any such trapezoid, or  $s_i$  appears in the conflict list and  $s_j$  bounds the trapezoid on the bottom. Eventually,  $s_i$  will thus be inserted, and then the resulting trapezoid will be compatible with  $\Delta$  and the procedure will terminate correctly, contrary to assumption.  $\blacksquare$

**Lemma 9.3.6** *For any final trapezoid  $\Delta$  created by the **Expand** procedure, during the execution of **CompZoneOnline**,  $\Delta$  is a trapezoid of  $\mathcal{A}_{\mathcal{VD}}(S_i)$ , where  $i = \text{index}(\Delta)$ .*

*Proof:* By induction on the *depth* of the nodes in  $T$ , where the depth of a node is defined to be the length of the longest path from the root of  $T$  to this node.

The only node of depth 0 is the root of  $T$ , which is being computed during the initialization of the algorithm, and is also the only trapezoid in  $\mathcal{A}_{\mathcal{VD}}(S_0)$ .

Assume that the induction hypothesis holds for all trapezoids of depth  $< k$ , and let  $\Delta$  be a final trapezoid of depth  $k$  in  $T$ . There are two cases to consider: (a)  $\Delta$  has been generated by splitting a final trapezoid  $\tau$  by inserting some arc  $s_i$ . (b)  $\Delta$  has been obtained by merging several transient trapezoids.

In case (a), by induction hypothesis,  $\tau$  is a trapezoid of  $\mathcal{A}_{\mathcal{VD}}(S_j)$ , where  $j = \text{index}(\tau)$ . By **Corollary 9.3.2**, the conflict list of  $\tau$  is computed correctly, so no arc  $s_\ell$ , with  $j < \ell < i$ , crosses  $\tau$ . Hence  $\tau$  is

also a trapezoid of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(S_{i-1})$ , and, by construction, any final trapezoid obtained by splitting  $\tau$  with  $s_i$  is a trapezoid of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(S_i)$ . Since  $i = \text{index}(\Delta)$ , this completes the induction step in this case.

In case (b), let  $\tau_1, \dots, \tau_m$  be the transient trapezoids whose merging forms  $\Delta$ . By construction, all of them have the same top arc, say  $s_i$ , and the same bottom arc, say  $s_j$ . Suppose, without loss of generality, that  $i > j$ . Since these trapezoids are transient, we have, as argued above,  $\text{index}(\tau_\ell) = \text{rank}(\tau_\ell) = i$  for each  $\ell = 1, \dots, m$ . In particular, this also holds for the leftmost and rightmost trapezoids in this sequence, which is easily seen to imply that all the arcs in  $D(\Delta)$  belong to  $S_i$ . Finally, since the conflict lists of the  $\tau_\ell$ 's are computed correctly, and the conflict list of  $\Delta$  is the union of these lists, it follows that no arc in that list belongs to  $S_i$ . In other words,  $\Delta$  is a final trapezoid defined by arcs of  $S_i$  and no arc of  $S_i$  crosses its interior. This readily implies that  $\Delta$  is a trapezoid of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(S_i)$ , and this completes the induction step in this case, and thus completes the proof of the lemma. ■

**Lemma 9.3.7** *All the (final) non-leaf nodes computed by CompZoneOnline appear in  $\mathcal{HT}_\gamma$ .*

*Proof:* What we really need to show is that each non-leaf trapezoid  $\Delta$  in  $T$  belongs to  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(S_{i-1})$ , where  $s_i$  is the arc that has split  $\Delta$ , thus making it a non-leaf. Indeed, any such trapezoid belongs to  $\mathcal{HT}_\gamma$ , by construction and by correctness of CompZoneWithOracle.

The proof proceeds by induction on the sequence of trapezoid-splitting steps taken by CompZoneOnline. That is, a final non-leaf trapezoid  $\Delta$  will be considered when it is split by an arc (thus becoming a non-leaf). The claim clearly holds initially for the whole plane, stored at the root of  $T$  (and of  $\mathcal{HT}_\gamma$ ). Let  $\Delta$  be a non-leaf final trapezoid generated in  $T$  and then split by an arc  $s_i$  by CompZoneOnline, and suppose that all previously-split non-leaf trapezoids in  $T$  appear in  $\mathcal{HT}_\gamma$ . The trapezoid  $\Delta$  has been split as part of a point-location query with some point  $p$ . Suppose first that  $p \in \gamma$  or  $p$  lies on a splitter of a final trapezoid of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(\hat{S})$  (the later case occurs when we expand the district of  $\gamma$  into its zone). Since  $\Delta$  is split at that point, it follows by construction that  $p \in \Delta$  and therefore  $\Delta$  belongs to  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(S_{i-1})$ .

Otherwise,  $p$  is the midpoint of some transient splitter  $\nu$ , and the point location of  $p$  is part of some Expand operation. Again,  $p$  (and in fact the whole segment  $\nu$ ) belongs to  $\Delta$ . Let  $\tau$  be the transient trapezoid bounded by  $\nu$  that has triggered that Expand operation, and let  $\tau_0$  be the first transient trapezoid (in the execution order of CompZoneOnline) in the sequence of compatible trapezoids that includes  $\tau$ . Let  $\tau_f$  be the father of  $\tau_0$ , which is a final (non-leaf) trapezoid from which  $\tau_0$  has been split by an arc  $s_j$ . By induction hypothesis,  $\tau_f$  belongs to  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(S_{j-1})$ . Hence  $\tau_0$  is fully contained (as a set) in the zone of  $\gamma$  in  $\mathcal{A}(S_{j-1})$ . By the preceding analysis, we know that (a)  $s_j$  is the top or bottom arc of  $\tau_0$  and of  $\tau$ , and (b) during the whole Expand operation that started at  $\tau_0$  no arc beyond  $s_j$  is inserted. It follows that  $p$  belongs to the zone of  $\gamma$  in  $\mathcal{A}(S_{j-1})$ , and that  $s_i$ , the arc that has split  $\Delta$ , must satisfy  $i \leq j$ . Hence  $\Delta$  belongs to the zone of  $\gamma$  in  $\mathcal{A}(S_{i-1})$ , which establishes the induction step and thus completes the proof of the lemma. ■

**Lemma 9.3.8** *All the trapezoids of  $\mathcal{A}_{\gamma, \mathcal{V}\mathcal{D}}(\hat{S})$  are computed by CompZoneOnline.*

*Proof:* Consider first the sequence of trapezoids of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(\hat{S})$  that constitute the district of  $\gamma$  in the full arrangement, in the order in which they are traversed by  $\gamma$ . The proof first proceeds by induction on this sequence, arguing that each of these trapezoids is produced by CompZoneOnline. As implied by the preceding analysis, each point-location query with a point  $p$  returns the trapezoid of  $\mathcal{A}_{\mathcal{V}\mathcal{D}}(\hat{S})$  that contains  $p$ . Hence, if the  $k$ -th trapezoid in the above sequence has been generated, the *EscapePoint* $_\gamma$  function produces a point that lies in the next trapezoid, which will therefore also be produced by CompZoneOnline. A similar argument holds for the subsequent stage of the algorithm that expands the district of  $\gamma$  into its zone. ■

### 9.3.2 Running Time

In this subsection, we first analyze the performance of `CompZoneWithOracle`, and then use this analysis to bound the expected running time of `CompZoneOnline`. We assume that `CompZoneWithOracle` maintains for each trapezoid  $\Delta$  its conflict list  $\text{cl}(\Delta)$  that stores the set of arcs that cross it. We also assume that each conflict list  $\text{cl}(\Delta)$  stores its minimal element  $\text{next}(\Delta)$  in the ordering of  $S$ , and that each yet uninserted arc  $s$  maintains a list of all current trapezoids  $\Delta$  for which  $\text{next}(\Delta) = s$ . Then it is easy to see that the running time of the algorithm is proportional to the overall size of all the conflict lists that it generates. We also assume that a call to the Oracle  $\mathcal{O}$  takes  $O(1)$  time.

**Lemma 9.3.9** *The algorithm `CompZoneWithOracle` computes the zone of  $\gamma$  in  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  in  $O(\lambda_{t+2}(n+m)\log n)$  expected time, and the expected number of trapezoids that it generates is  $O(\lambda_{t+2}(n+m))$ .*

*Proof:* The proof is a straightforward adaptation of the proof of [CEG<sup>+</sup>93].<sup>3</sup> Specifically, we first make  $m$  cuts at the points where  $\gamma$  crosses the arcs of  $\hat{S}$ , thereby obtaining a collection of  $m+n$  subarcs, so that  $\mathcal{Z}_\gamma(\hat{S})$  becomes a single face in the new arrangement. We now insert the original arcs of  $\hat{S}$  one by one in the random order  $S$ . It is easily checked that the expected number of subarcs of the  $r$  random arcs have been inserted is  $O(r + \frac{m}{n}r)$ . Plugging this into the analysis in [CEG<sup>+</sup>93] and omitting the further easy details, the bounds follow easily. ■

**Lemma 9.3.10** *The expected number of transient trapezoids generated by `CompZoneOnline` is  $O(\lambda_{t+2}(n+m))$ , and the expected total size of their conflict lists is  $O(\lambda_{t+2}(n+m)\log n)$ .*

*Proof:* Each final trapezoid generated by `CompZoneOnline` might be split into  $O(1)$  transient trapezoids. Each final (non-leaf) trapezoid computed by `CompZoneOnline` is also computed by `CompZoneWithOracle`, as follows from Lemma 9.3.7. By Lemma 9.3.9, the expected number of such trapezoids is  $O(\lambda_{t+2}(n+m))$ .

The second part of the lemma follows by a similar argument. ■

**Definition 9.3.11** *A curve  $\gamma$  is locally  $x$ -monotone in  $\mathcal{A}(\hat{S})$ , if it can be decomposed inside each face of  $\mathcal{A}(\hat{S})$  into a constant number of  $x$ -monotone curves.*

**Theorem 9.3.12** *The algorithm `CompZoneOnline` computes the zone of  $\gamma$  in  $\mathcal{A}(\hat{S})$  in  $O(\lambda_{t+2}(n+m)\log n)$  expected time, provided that  $\gamma$  is a locally  $x$ -monotone curve in  $\mathcal{A}(\hat{S})$ .*

*Proof:* The time spent by `CompZoneOnline` is bounded by the time required to construct the history DAG, by the time spent in maintaining the conflict lists of the trapezoids, and by the time spent on performing point-location queries, as we move from one trapezoid to another in  $\mathcal{A}_{\gamma, \mathcal{VD}}(S)$ .

By Lemma 9.3.9 and Lemma 9.3.10, the expected time spent on maintaining the conflict lists of the trapezoids computed by the algorithm is  $O(\lambda_{t+2}(n+m)\log n)$ , since the total time spent on handling the conflict lists is proportional to their total length. By Lemma 9.3.10, the expected total size of those conflict lists is  $O(\lambda_{t+2}(n+m)\log n)$ .

Moreover, the depth of the history DAG constructed by the algorithm is  $O(\log n)$  with a probability polynomially close to 1 [Mul94]. Thus, the expected time spent directly on performing a single point-location query (that is, the number of trapezoids that contain the query point and are visited or generated during this point location step) is  $O(\log n)$ . The curve  $\gamma$  is locally  $x$ -monotone, which implies that it intersects each splitter of any trapezoid of any  $\mathcal{A}_{\gamma, \mathcal{VD}}(\hat{S})$  at most  $O(1)$  times. Thus, the expected number

<sup>3</sup>The algorithm of [CEG<sup>+</sup>93] has some additional overhead that is not required in `CompZoneWithOracle`

of point-location queries performed by the algorithm is proportional to the expected number of transient and final trapezoids created, plus  $O(m)$ . By [Lemma 9.3.10](#), we have that the expected running time is

$$O\left(\lambda_{t+2}(n+m)\log n + (\lambda_{t+2}(n+m) + m)\log n\right) = O(\lambda_{t+2}(n+m)\log n).$$

■

**Remark 9.3.13** Note that `CompZoneWithOracle` computes the zone of  $\gamma$  in  $\mathcal{A}_{\mathcal{VD}}(S_i)$ , for each  $i = 1, \dots, n$ . As a consequence, it might compute a trapezoid  $\Delta \in \mathcal{A}_{\gamma, \mathcal{VD}}(S_i)$  that does not intersect the zone of  $\gamma$  in  $\mathcal{A}_{\gamma, \mathcal{VD}}(S)$ . In particular, such a trapezoid  $\Delta$  will not be computed by `CompZoneOnline`. This is a slackness in our analysis that we currently do not know whether it can be exploited to further improve the analysis of the algorithm (we suspect that it cannot improve the above asymptotic bound on the running time).

**Remark 9.3.14** The only classical result of this type that we are aware of, is due to Overmars and van Leeuwen [[OvL81](#)]. It maintains dynamically the intersection of  $n$  halfplanes in (deterministic)  $O(\log^2 n)$  time for each insertion or deletion operation. This procedure can be used to perform walks inside line arrangements in (deterministic)  $O((n+m)\log^2 n)$  time, where  $m$  is the number of intersections of the walk with the lines. Our algorithm is somewhat simpler and is faster than the algorithm of [[OvL81](#)] by nearly a logarithmic factor, and, most importantly, applies to more general arrangements.<sup>4</sup>

The technique of [[OvL81](#)] can also be used to perform *x-monotone* online walks in arrangement of segments. We simply regard each segment as the full line that contains it, but make it active only when the walk passes vertically above or below the segment. This means that, in addition to the usual updates that the algorithm performs, we also insert (resp. delete) segments when the walk becomes co-vertical with their left (resp. right) endpoint. (If the walk is not *x-monotone*, this may slow down the algorithm considerably.)

As for general arcs (or non-monotone walks in arrangements of segments), we are not aware of any result of this type in the literature. Of course, if the curve  $\gamma$  is known in advance (and is simple, in the sense that one can compute quickly its intersections with any arc of  $\hat{S}$ ), we can compute the single face containing  $\gamma$  in an appropriately modified arrangement (as in the proof of the general planar Zone Theorem [[SA95](#), Theorem 5.11]; see also the proof of [Lemma 9.3.9](#) using the algorithms of [[BDS95](#), [CEG+93](#)]). These algorithms (especially the first one) are slightly simpler than the algorithm of [Theorem 9.3.12](#), although they have the same expected performance. However, these algorithms are useless for online walks, and they are probably slower than our algorithm in practice, as they either maintain additional complicated data-structures [[CEG+93](#)], or perform additional redundant computation of regions that lie outside the zone of  $\gamma$  [[BDS95](#)].

Recently, several algorithms for performing online walks were implemented [[AHH+99](#)], including a variant of the algorithm presented here, which exhibited satisfactory performance in practice.

## 9.4 Applications

In this section we present several applications of the algorithm `CompZoneOnline`.

<sup>4</sup>Very recently, Chan [[Cha99a](#)] improved the result of [[OvL81](#)], providing a data structure for maintaining intersection of halfplanes in  $O(\log^{1+\varepsilon} n)$  amortized time for each update. His data structure is considerably more complicated than ours and than that of Overmars and van Leeuwen, and currently seems to be only of theoretical significance. Moreover, our algorithm is still faster than Chan's (by a factor of  $O(\log^\varepsilon n)$ ).

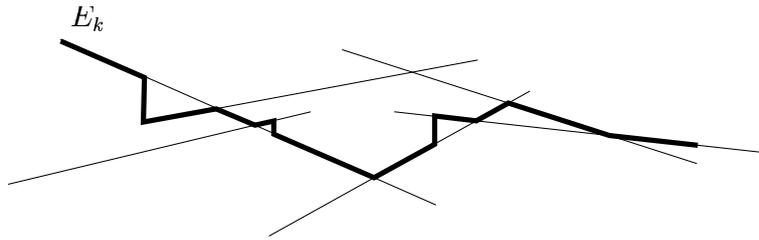


Figure 9.4: The first level in an arrangement of segments (the vertical edges show the jump discontinuities of the level, but are not part of the level).

### 9.4.1 Computing a Level in an Arrangement of Arcs

In this subsection we show how to modify the algorithm of the previous section to compute a level in an arrangement of  $x$ -monotone arcs. See also [Section 8.3](#) for a discussion of levels in line arrangements.

**Definition 9.4.1** *Let  $\hat{S}$  be, as above, a set of  $n$   $x$ -monotone arcs in the plane, any pair of which intersect at most  $t$  times (for some fixed constant  $t$ ). We also assume, as above, that  $\hat{S}$  is in general position. The level of a point in the plane is the number of arcs of  $\hat{S}$  lying strictly below it. Consider the closure  $E_l$  of the set of all points on the arcs of  $\hat{S}$  having level  $l$  (for  $0 \leq l < n$ ).  $E_l$  is an  $x$ -monotone (not necessarily connected) curve (which is polygonal in the case of lines or segments), which is called the level  $l$  of the arrangement  $\mathcal{A}(\hat{S})$ . At  $x$ -coordinates where a vertical line intersects less than  $l + 1$  arcs of  $S$ , we consider  $E_l$  to be undefined.*

Levels are a fundamental structure in computational and combinatorial geometry, and have been subject to intensive research in recent years (see [[AACCS98](#), [Dey98](#), [TT97](#), [TT98](#)]). Tight bounds on the complexity of a single level, even for arrangements of lines, proved to be surprisingly hard to obtain. Currently, the best known upper bound for the case of lines is  $O(n(l + 1)^{1/3})$  [[Dey98](#)], while the lower bound is  $\Omega(n \log(l + 1))$  [[Ede87](#)].<sup>5</sup> See also [[AACCS98](#), [TT98](#)] for weaker upper bounds for other classes of arcs.

First, note that if  $\hat{S}$  is a set of lines, then, once we know the leftmost ray that belongs to  $E_l$ , the level  $l$  is locally defined: as we move from left to right along  $E_l$ , each time we encounter an intersection point (a vertex of  $\mathcal{A}(\hat{S})$ ) we have to change the line that we traverse. (This is also depicted in [Figure 9.4](#).) In particular, we can compute the level  $E_l$  in  $O(\lambda_3(n + |E_l|) \log n)$  time, using `CompZoneOnline`. The same procedure can be used to compute a level in an arrangement of more general arcs. The only non-local behavior we have to watch for are jump discontinuities of the level caused when an endpoint of an arc appears below the current level, or when the current level reaches an endpoint of an arc (see [Figure 9.4](#)). See below for details concerning the handling of those jumps.

In the following, let  $0 \leq l < n$  be a prescribed parameter. Let  $E_l$  denote the level  $l$  in the arrangement  $\mathcal{A}(\hat{S})$ .

The following adaptation of `CompZoneOnline` to our setting is rather straightforward, but we include it for the sake of completeness. We sort the endpoints of the arcs of  $\hat{S}$  by their  $x$ -coordinates. Each time our walk reaches the  $x$ -coordinate of the next endpoint, we update  $E_l$  by jumping up or down to the next arc, if needed. This additional work requires  $O(n \log n)$  time.

If the level reaches the  $x$ -coordinate  $x_0$  of a right endpoint of an arc, past which there are fewer than  $l + 1$  arcs intersecting a vertical line, then the level lies on the highest arc just to the left of  $x_0$  and it ceases to be defined just to the right of  $x_0$ . In this case, our walk climbs to the line  $y = +\infty$

<sup>5</sup>Recently, a slightly larger lower bound has been announced by G. Tóth [[Tót01](#)].

and moves along the line (effectively tracing a sequence of topmost trapezoids of  $\mathcal{A}(\hat{S})$ ) until it reaches the  $x$ -coordinate of a left endpoint of an arc, following which we might have again  $l + 1$  arcs crossing a vertical line. If so the walk then descends on the topmost arc and continues to trace the level  $l$ . For simplicity, we continue the discussion of the algorithm assuming that  $E_l$  is everywhere defined.

During the walk, we maintain the invariant that the top edge of the current trapezoid is part of  $E_l$ . To compute the first trapezoid in the walk, we compute the intersection of level  $l$  with the  $y$ -axis (this can be done by sorting the arcs according to their intersections with the  $y$ -axis). Let  $p_0$  be this starting point. We perform a point-location query with  $p_0$  in our virtual history DAG to compute the starting trapezoid  $\Delta_0$  (containing  $p_0$  on its top arc).

Now, by walking to the right of  $\Delta_0$  we can compute the part of  $E_l$  lying to the right of the  $y$ -axis. Let  $\Delta$  be the current trapezoid maintained by the algorithm, such that its top edge is a part of  $E_l$ . Let  $p(\Delta)$  denote the top right vertex of  $\Delta$ . By performing point-location queries in our partial history DAG  $T$ , we can compute all the trapezoids of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that contain  $p(\Delta)$  (by our general position assumption, the number of such trapezoids is at most 6; this number materializes when  $p(\Delta)$  lies in the intersection of two arcs). By inspecting this set of trapezoids, one can decide where  $E_l$  continues to the right of  $\Delta$ , and determine the next trapezoid having  $E_l$  as its roof. The algorithm sets  $\Delta$  to be this trapezoid.

If the algorithm reaches an  $x$ -coordinate of an endpoint of an arc, we have to update  $E_l$  by jumping up (if this is the right endpoint of an arc and it lies on or below the level) or down (if it is a left endpoint and lies below the level); namely, we set  $\Delta$  to be the trapezoid lying above (or below) the current  $\Delta$ .

The algorithm continues in this manner, until reaching the rightmost edge of  $E_l$ . The algorithm then performs a symmetric walk to the left of the  $y$ -axis to compute the other portion of the level.

Let `CompLevel` denote this modified algorithm. We summarize our result:

**Theorem 9.4.2** *The algorithm `CompLevel` computes the level  $l$  in  $\mathcal{A}(\hat{S})$  in  $O(\lambda_{t+2}(n + |E_l|) \log n)$  expected time.*

**Remark 9.4.3** Since `CompLevel` is online, we can use it to compute the first  $m'$  edges or vertices of  $E_l$ , in expected  $O(\lambda_{t+2}(n + m') \log n)$  time.

**Remark 9.4.4** A straightforward extension of `CompLevel` allows us to compute any connected path within the union of  $\hat{S}$  (i.e., we restrict our “walk” to the arcs of  $\hat{S}$ ) in an on-line manner, in randomized expected time  $O(\lambda_{t+2}(m + n) \log n)$ , where  $m$  is the number of vertices of the path. As above, the extended version can also handle vertical jumps between adjacent arcs during the walk.

**Remark 9.4.5** For the case of lines, one can use the algorithm of [OvL81] to construct a level  $E_l$  in  $O(n \log n + |E_l| \log^2 n)$  deterministic time, as described, e.g., in [EW86]. The same technique, with a simple modification, also works for the case of line segments, with the same complexity bounds. Our algorithm is faster in these cases by nearly a logarithmic factor.

As already mentioned, very recently, Chan [Cha99a] presented a faster algorithm for the dynamic maintenance of the intersection of halfplanes, requiring  $O(\log^{1+\epsilon} n)$  amortized time for each operation. Thus, one can compute the level  $l$  in  $O(n \log n + |E_l| \log^{1+\epsilon} n)$  deterministic time. Chan [Cha99b] also showed that by using the algorithm of [ABMS98] one can compute the level  $l$  in  $O(n + |E_l| \alpha(n)^2 \log n)$  randomized expected time. We note, however, that our algorithm is still faster and simpler than those two algorithms.

## 9.4.2 Other Applications

In this subsection, we provide some additional applications of `CompZoneOnline` and `CompLevel`. The first application has been presented in [Chapter 8](#): Applying the algorithm of [Theorem 9.4.2](#) and [Remark 9.4.3](#) to a set  $L$  of  $n$  lines in the plane, we showed that one can compute a  $(1/r)$ -cutting of  $\mathcal{A}(L)$ , having at most  $(1 + \varepsilon)(8r^2 - 2r + 4)$  vertical trapezoids, in randomized expected time  $O\left(\left(1 + \frac{1}{\varepsilon}\right)nr\alpha(n)\log n\right)$ , for any  $\varepsilon > 0$ .

Once we have computed the level  $l$  (in an arrangement of general arcs), we can clip the arcs to their portions below the level. Using those clipped arcs as input, we can compute the portion of the arrangement below the level  $l$  (i.e., the first  $l$  levels of  $\mathcal{A}(\hat{S})$ ) in  $O((m + n)\log n + r)$  time, where  $m = |E_l|$  is the complexity of the level  $l$ , and  $r$  is the complexity of the first  $l$  levels of  $\mathcal{A}(\hat{S})$ , using, e.g., the algorithm described in [\[Mul94\]](#). This improves over the previous result of [\[ERK96\]](#) that computes this portion of the arrangement of lines in  $O(n\log n + nl)$  time. (Note that this running time is *not* output sensitive: It is easy to come up with examples where the complexity of the first  $l$  levels is only  $O(l^2)$ .)

A byproduct of the technique of `CompZoneOnline` is the ability to perform point-location queries in an arrangement using the partial history DAG mechanism:

**Definition 9.4.6** *For a point set  $P$ , and a set of arcs  $\hat{S}$ , let  $\mathcal{M}(P, \hat{S})$  denote a connected polygonal set, such that: (i)  $P \subseteq \mathcal{M}(P, \hat{S})$ , and (ii) the number of intersections between  $\mathcal{M}(P, \hat{S})$  and the arcs of  $\hat{S}$  is minimum. Let  $w(P, \hat{S})$  denote the number of such intersections.*

*The set  $\mathcal{M}(P, \hat{S})$  can be interpreted as the minimum spanning Steiner tree of  $P$ , under the metric that measures distances by the number of intersections with the arcs of  $\hat{S}$ .*

**Theorem 9.4.7** *Given a set  $\hat{S}$  of  $n$  arcs in the plane, one can answer point-location queries for a set  $P$  of  $m$  points, where the points are given in an online manner, such that the overall expected time to answer those queries is  $O(\lambda_{t+2}(n + m + w(P, \hat{S}))\log n)$  time. Here the output of a query is the trapezoid of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that contains the query point.*

*Proof:* We precompute a random permutation  $S$  of  $\hat{S}$ , and answer the point-location queries, by computing the relevant parts of the history DAG of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$ , as in `CompZoneOnline`. (Note that here there is no “continuity”: consecutive point locations may explore far away portions of the plane. Nevertheless, this does not effect the performance of the algorithm directly.)

By the time the algorithm terminates, the partial history DAG  $T$  is contained in  $\mathcal{A}_{\gamma, \mathcal{VD}}(S)$ , where  $\gamma = \mathcal{M}(P, \hat{S})$ . (To see this, imagine that we run `CompZoneOnline` to compute the district of  $\gamma$ . It is easily seen that the resulting  $T$  contains the tree generated by our point location queries.) The expected total weight of the trapezoids of  $\mathcal{HT}_{\gamma}(S)$  is thus  $O(\lambda_{t+2}(n + m + w(P, \hat{S}))\log(n))$ , by [Lemma 9.3.9](#). This bounds the overall expected query time. ■

**Remark 9.4.8** The result of [Theorem 9.4.7](#) is somewhat disappointing, since, in the worst case,  $w(P, \hat{S}) = \Theta(nm^{1/2})$  [\[Aga91b\]](#), while for the case of lines,  $m$  faces can be computed in, roughly,  $O(n^{2/3}m^{2/3})$  time [\[AMS98\]](#) (a task that clearly subsumes our point location queries). Currently, for the case of general arcs, no bound better than  $O(m^{1/2}\lambda_{t+2}(n))$  is known for the complexity of  $m$  faces in an arrangement of  $n$  arcs (see [\[EGP<sup>+</sup>92, Har99c\]](#)). We leave the problem of improving the performance of the algorithm of [Theorem 9.4.7](#), either for the general case or just for the case of lines, as an open problem for further research.

The algorithm of [Theorem 9.4.7](#) is simple and online, and it has the additional favorable property of being adaptive. Namely, if  $w(P, \hat{S})$  is small (i.e., the query points are “close together”) the overall query

time improves. Furthermore, if there are a lot of queries close together, the first query will be slow, and the later ones will be faster (since the later queries use parts of paths that already exist in the partial history DAG).

## 9.5 Conclusions

In this chapter we have presented a new randomized algorithm for computing a zone in a planar arrangement, in an online fashion. This algorithm is the first efficient algorithm for the case of planar arcs, it performs faster (by nearly a logarithmic factor) than the algorithm of [OvL81] for the case of lines and for the case of an  $x$ -monotone walk in an arrangement of segments, and it is considerably simpler. (It is also faster and much simpler than the recent algorithm of [Cha99a].) We also presented an efficient randomized algorithm for computing a level in an arrangement of arcs in the plane, whose expected running time is faster than any previous algorithm for this problem.

The main result of this chapter relies on the application of point-location queries to compute the relevant parts of an “off-line” structure (i.e., the history DAG). The author believes that this technique should have additional applications. In particular, this approach might be useful also for algorithms in higher dimensions. We leave this as an open question for further research.

## 9.A Appendix - Pseudo-code for Subroutines of CompZoneOnline

```

ALGORITHM Split( $v$ )
  Input: A final node  $v$  in the partial history DAG  $T$ 
  begin
     $s \leftarrow \text{next}(\Delta_v)$ , where  $\Delta_v$  is the trapezoid associated with the node  $v$ 
     $L \leftarrow \text{SplitGeom}(\Delta_v, s)$ ,
      where  $\text{SplitGeom}(\Delta_v, s)$ , as above, returns the collection of trapezoids
      that cover  $\Delta_v$ , so that  $s$  does not intersect any of them in its interior.
    for each  $\tau \in L$  do
      Create a new node  $w$  and attach it as a child of  $v$  in  $T$ .
      Set  $\Delta_w$  to  $\tau$ 
      Compute  $\text{cl}(\Delta_w)$  from  $\text{cl}(\Delta_v)$ 
      Compute  $\text{next}(\Delta_w)$ , the first element of  $\text{cl}(\Delta_w)$ 
    end for
  end Split

```

Figure 9.5: Splitting a final node in  $T$  and creating its children

```

ALGORITHM PointLocateLeftCompatible( $v, p, r$ )
  Input:  $v$  - current node of  $T$ 
          $p$  - query point
          $r$  - target rank of output trapezoid
  Output: A transient trapezoid of rank  $r$  having  $p$  on its left splitter
  begin
    if  $\text{rank}(v) = r$  then
      return  $v$ 
    if  $\text{isTransient}(v)$  then
       $v \leftarrow \text{Expand}(v)$ 
    if  $\text{isLeaf}(v)$  then
      Split( $v$ )
    Let  $w$  be the child of  $v$ , so that  $\Delta_w$  contains  $p$  either in its interior
      or on its left splitter
    return PointLocateLeftCompatible( $w, p, r$ )
  end PointLocateLeftCompatible

```

Figure 9.6: Computing a transient trapezoid in  $T$  that is left “compatible” with an input transient trapezoid, by carrying out a point-location query in  $T$ . The algorithm also uses a symmetric routine  $\text{PointLocateRightCompatible}$ , whose code is omitted.

```

ALGORITHM Expand( $v$ )
  Input:  $v$  - current transient leaf node of  $T$ 
  Output: A final node of  $T$  whose trapezoid contains  $\Delta_v$ 
begin
  if  $isFinal(\Delta_v)$  then
    return  $v$ 
   $L \leftarrow \{v\}$ 

  /* collect the sequence of transient trapezoids adjacent to each other
  to the right of  $\Delta_v$  */
   $temp \leftarrow v$ 
  while  $isTransientRightSplitter(\Delta_{temp})$  do
     $temp \leftarrow \text{PointLocateLeftCompatible}( \text{root}(T),$ 
       $\text{midPointRightSplitter}( \Delta_{temp}, \text{rank}(\Delta_{temp}) )$ 
    )
     $L \leftarrow L \cup \{temp\}$ 
  end while

  /* Similarly collect the sequence of transient trapezoids to the left of  $\Delta_v$  */
   $temp \leftarrow v$ 
  while  $isTransientLeftSplitter(\Delta_{temp})$  do
     $temp \leftarrow \text{PointLocateRightCompatible}( \text{root}(T),$ 
       $\text{midPointLeftSplitter}( \Delta_{temp}, \text{rank}(\Delta_{temp}) )$ 
    )
     $L \leftarrow L \cup \{temp\}$ 
  end while

   $\Delta \leftarrow \bigcup_{u \in L} \Delta_u$ 
  Compute  $cl(\Delta)$  and  $next(\Delta)$  from the conflict lists of the
  nodes of  $L$ , using the global bit-vector technique.
  Add a new leaf node  $x$  to the partial history DAG  $T$ , mark  $x$  as final,
  and replace all nodes of  $L$  in  $T$  by  $x$ .
  Set  $\Delta_x$  to  $\Delta$ 

  return  $x$ 
end Expand

```

Figure 9.7: Expanding a transient leaf trapezoid of  $T$  to a final trapezoid containing it

```

ALGORITHM PointLocate(  $p$ ,  $flag$  )
  Input:  $p$  - a query point
            $flag$  - since  $p$  usually lies on the boundary of a trapezoid,  $flag$ 
                 indicates the side of the splitter or arc that contains  $p$ 
                 where the point location should take place
  Output: The trapezoid of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that contains  $p$  (in its interior
                 or on the appropriate edge dictated by  $flag$ )

  begin
     $v \leftarrow root(T)$ 

    while (  $cl(\Delta_v) \neq \emptyset$  ) do
      Expand( $v$ )
      Split( $v$ )
       $v \leftarrow$  child of  $v$  whose trapezoid contains  $p$ ; if  $p$  lies on
                    the boundary of several children trapezoids, choose the one
                    that is compatible with  $flag$ 
    end while

    return  $\Delta_v$ 
  end PointLocate

```

Figure 9.8: The function that performs a point-location query; that is, it computes the necessary parts of the partial history DAG  $T$ , and returns the trapezoid of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that contains a query point.

## 9.B Appendix - Taking a Walk in Ten Easy Figures

In this appendix we illustrate, step by step, the action of processing a single point-location query by `CompZoneOnline`.

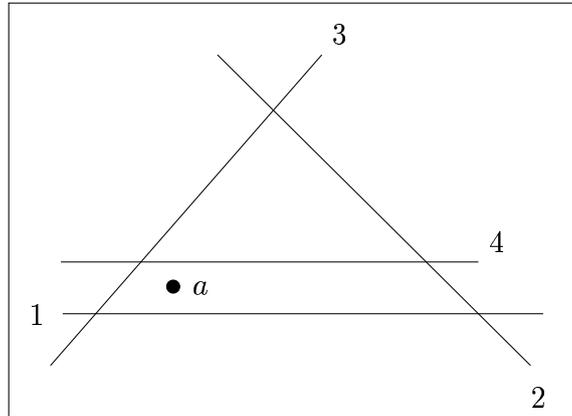


Figure 9.9: The input for `CompZoneOnline` is the set of segments  $\hat{S} = \{1, 2, 3, 4\}$ . We assume that the algorithm uses the permutation  $S = (1, 2, 3, 4)$ . We illustrate how `CompZoneOnline` carries out a point-location query to compute the trapezoid of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that contains the point  $a$ . We assume that this is the first query processed by the algorithm.



Figure 9.10: `CompZoneOnline` starts with the trapezoid  $\Delta_0 = \mathbb{R}^2$  that corresponds to the root of the partial DAG  $T$ .

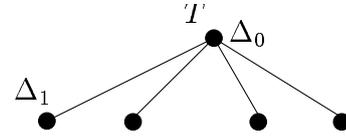


Figure 9.11: `CompZoneOnline` splits  $\Delta_0$  by the segment  $\text{next}(\Delta_0) = 1$ , and goes down in the DAG into the new node that corresponds to the trapezoid  $\Delta_1$ . All children of  $\Delta_0$  (including  $\Delta_1$ ) are final.

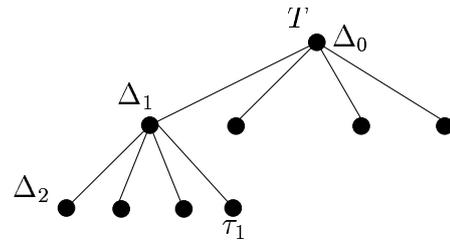
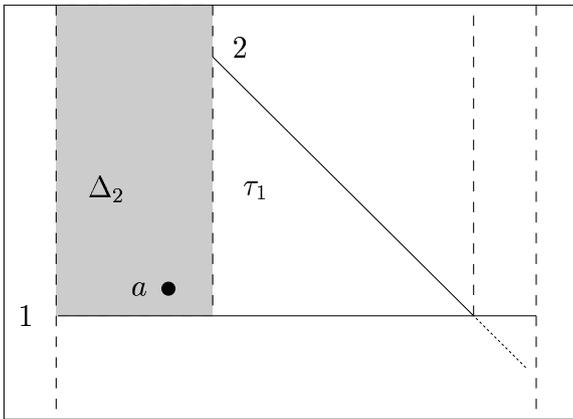


Figure 9.12: Since  $\Delta_1$  is final, it is split by  $\text{next}(\Delta_1) = 2$  into four final subtrapezoids, and we go down to the child  $\Delta_2$  that contains  $a$ .

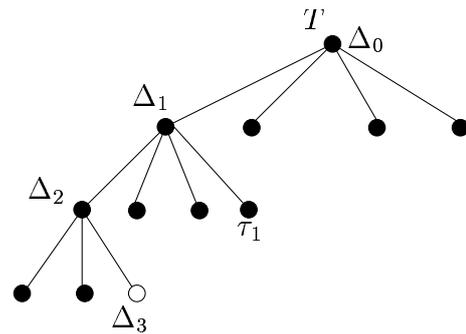
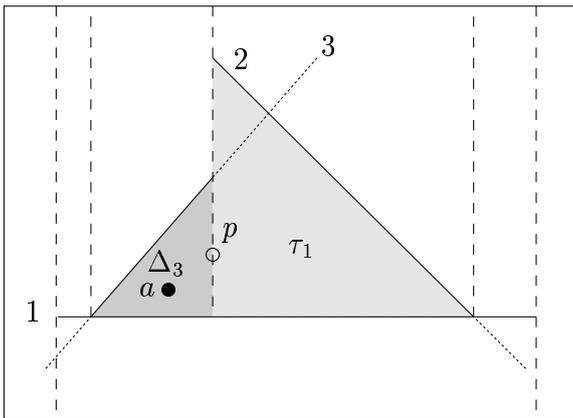


Figure 9.13: `CompZoneOnline` splits  $\Delta_2$  by  $\text{next}(\Delta_2) = 3$  into three subtrapezoids. Two of them are final, and the third one,  $\Delta_3$ , that contains  $a$  is transient (its right splitter is transient). We thus execute `Expand`( $\Delta_3$ ), which performs a point-location query (with the midpoint  $p$  of the right splitter). The point location goes down in the partial DAG, through  $\Delta_0$  and  $\Delta_1$ , and reaches the (final) leaf that stores  $\tau_1$ . Since  $\text{rank}(\Delta_3) = 3$  and  $\text{rank}(\tau_1) = 2$ , those two trapezoids are not compatible (this holds also because  $\tau_1$  is final whereas  $\Delta_3$  is not), and the algorithm continues by further splitting  $\tau_1$  by  $\text{next}(\tau_1) = 3$ .

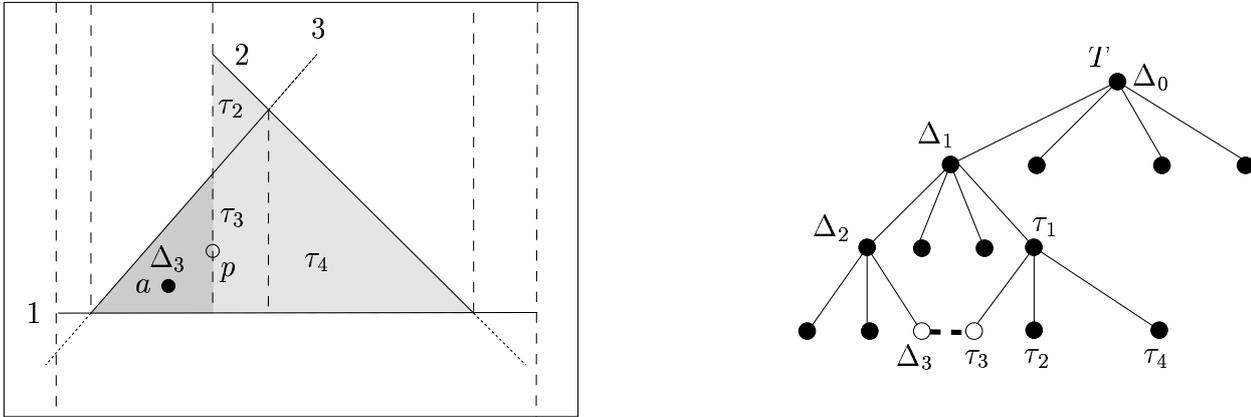


Figure 9.14: The splitting of  $\tau_1$  creates three children  $\tau_2, \tau_3, \tau_4$ , of which  $\tau_2$  and  $\tau_4$  are final, whereas  $\tau_3$  is transient. `CompZoneOnline` goes down to the newly created  $\tau_3$ , which contains  $p$  on its left splitter. Since  $\tau_3$  is transient and  $\text{rank}(\tau_3) = \text{rank}(\Delta_3)$ , it is compatible with  $\Delta_3$ . Since the right splitter of  $\tau_3$  is final, and so is the (empty) left splitter of  $\Delta_3$  the execution of `Expand`( $\Delta_3$ ) terminates.

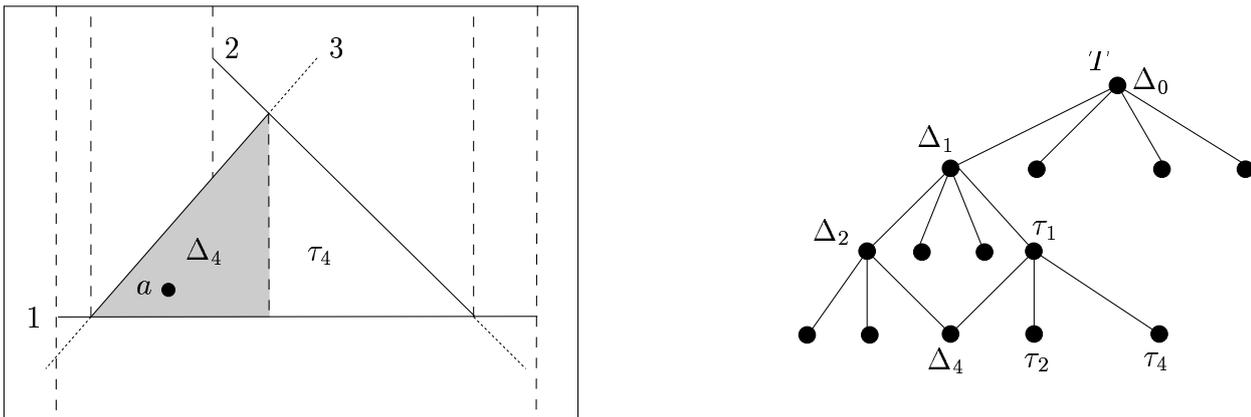


Figure 9.15: The trapezoids  $\Delta_3$  and  $\tau_3$ , from the previous figure, are merged by `CompZoneOnline` to form the final trapezoid  $\Delta_4$ . Since  $\text{cl}(\Delta_4)$  is not empty, we split it further by  $\text{next}(\Delta_4) = 4$ .

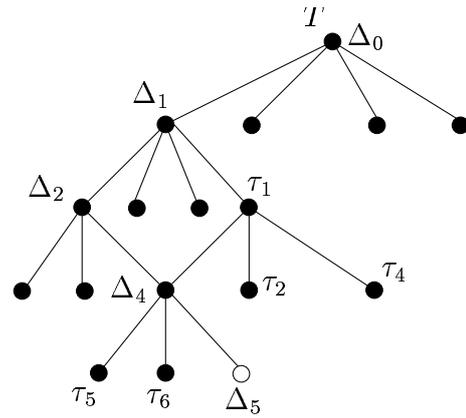
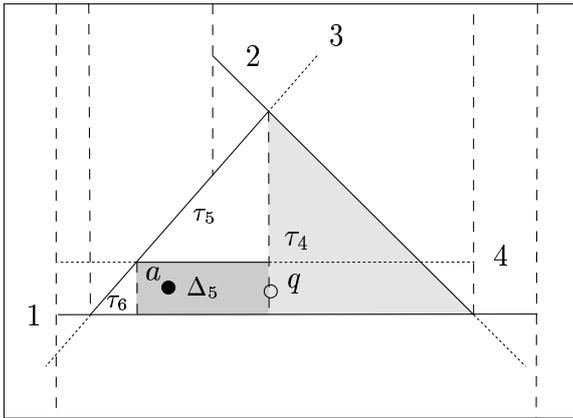


Figure 9.16: The splitting of  $\Delta_4$  creates three children  $\Delta_5, \tau_5, \tau_6$ , of which  $\tau_5$  and  $\tau_6$  are final whereas  $\Delta_5$  is transient and contains  $a$ . `CompZoneOnline` goes into  $\Delta_5$ , and since it is transient, `CompZoneOnline` calls `Expand( $\Delta_5$ )`. A point-location query is performed, at the midpoint  $q$  of the right splitter of  $\Delta_5$ . This query traverses in  $T$  the path  $(\Delta_0, \Delta_1, \tau_1, \tau_4)$ . The (final) trapezoid  $\tau_4$  is not compatible with  $\Delta_5$  so it is further split by the segment  $\text{next}(\tau_4) = 4$ .

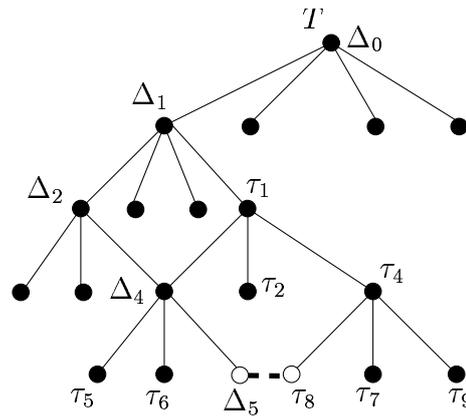
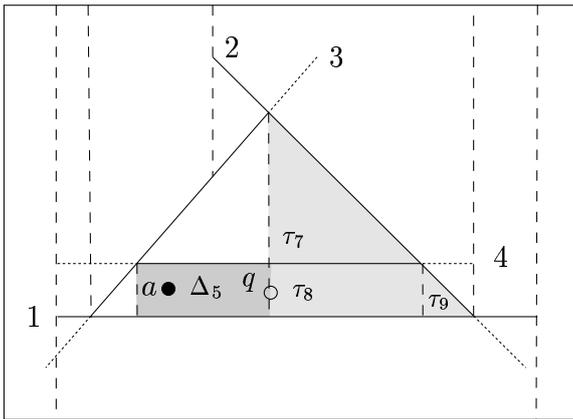


Figure 9.17:  $\tau_4$  is split into  $\tau_7, \tau_8, \tau_9$ , of which only  $\tau_8$  is transient. `CompZoneOnline` goes into  $\tau_8$ , which contains  $q$  on its left splitter. This trapezoid is compatible with  $\Delta_5$ , and since both the left splitter of  $\Delta_5$  and the right splitter of  $\tau_8$  are final, `Expand( $\Delta_5$ )` terminates and merges both trapezoids.

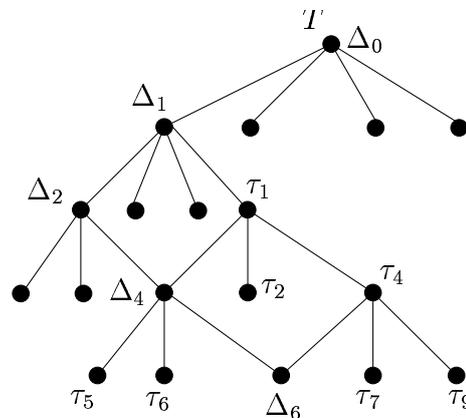
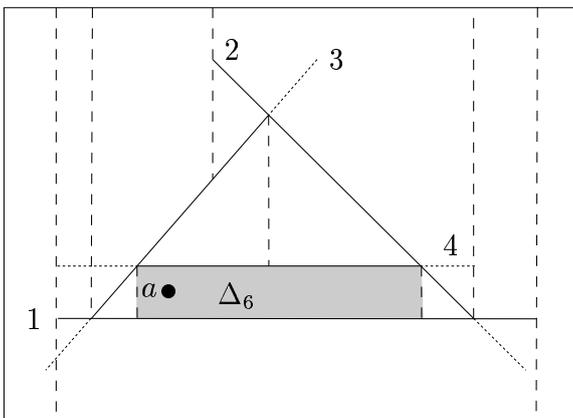


Figure 9.18: Voila! The newly formed final trapezoid  $\Delta_6$  contains our query point  $a$ , and its conflict list is empty. Thus,  $\Delta_6$  is the trapezoid of  $\mathcal{A}_{\mathcal{VD}}(\hat{S})$  that contains  $a$ , and is output as such by `CompZoneOnline`.