

Improved Implementation of Point Location in General Two-Dimensional Subdivisions^{*}

Michael Hemmer, Michal Kleinbort, and Dan Halperin

Tel-Aviv University, Israel

Abstract. We present a major revamp of the point-location data structure for general two-dimensional subdivisions via randomized incremental construction, implemented in CGAL, the Computational Geometry Algorithms Library. We can now guarantee that the constructed directed acyclic graph \mathcal{G} is of linear size and provides logarithmic query time. Via the construction of the Voronoi diagram for a given point set S of size n , this also enables nearest-neighbor queries in guaranteed $O(\log n)$ time. Another major innovation is the support of general unbounded subdivisions as well as subdivisions of two-dimensional parametric surfaces such as spheres, tori, cylinders. The implementation is exact, complete, and general, i.e., it can also handle non-linear subdivisions. Like the previous version, the data structure supports modifications of the subdivision, such as insertions and deletions of edges, after the initial preprocessing. A major challenge is to retain the expected $O(n \log n)$ preprocessing time while providing the above (deterministic) space and query-time guarantees. We describe efficient preprocessing algorithms, which explicitly verify the length \mathcal{L} of the longest query path. However, instead of using \mathcal{L} , our implementation is based on the depth \mathcal{D} of \mathcal{G} . Although we prove that the worst case ratio of \mathcal{D} and \mathcal{L} is $\Theta(n/\log n)$, we conjecture, based on our experimental results, that this solution achieves expected $O(n \log n)$ preprocessing time.

1 Introduction

Birn et al. [1] presented a structure for planar nearest-neighbor queries, based on Delaunay triangulations, named Full Delaunay Hierarchies (FDH). The FDH is a very simple, and thus light, data structure that is also very easy to construct. It outperforms many other methods in several scenarios, but it does not have a worst-case optimal behavior. However, it is claimed [1] that methods that do have this behavior are too cumbersome to implement and thus not available. We got challenged by this claim.

^{*} This work has been supported in part by the 7th Framework Programme for Research of the European Commission, under FET-Open grant number 255827 (CGL—Computational Geometry Learning), by the Israel Science Foundation (grant no. 1102/11), and by the Hermann Minkowski–Minerva Center for Geometry at Tel Aviv University.

In this article we present an improved version of CGAL’s planar point location that implements the famous incremental construction (RIC) algorithm as introduced by Mulmuley [2] and Seidel [3]. The algorithm constructs a linear size data structure that guarantees a logarithmic query time. It enables nearest-neighbor queries in guaranteed $O(\log n)$ time via planar point location in the Voronoi Diagram of the input points. In Section 4 we compare our revised implementation for point location, applied to nearest neighbor search, against the FDH. Naturally, this is only a byproduct of our efforts as planar point location is a very fundamental problem in Computational Geometry. It has numerous applications in a variety of domains including computer graphics, motion planning, computer aided design (CAD) and geographic information systems (GIS).

Previous Work: Most solutions can only provide an expected query time of $O(\log n)$ but cannot guarantee it, in particular, those that only require $O(n)$ space. Some may be restricted to static scenes that do not change, while others can only support linear geometry.

Triangulation-based point location methods, such as the approaches by Kirkpatrick [4] and Devillers [5] combine a logarithmic hierarchy with some walk strategy. Both require only linear space and Kirkpatrick can even guarantee logarithmic query time. However, both are restricted to linear geometry, since they build on a triangulation of the actual input.

Many methods can be summarized under the model of the trapezoidal search graph as pointed out by Seidel and Adamy [6]. Conceptually, the initial subdivision is further subdivided into trapezoids by emitting vertical rays (in both directions) at every endpoint of the input, which is the fundamental search structure. In principal, all these solutions can be generalized to support input curves that are decomposable into a finite number of x -monotone pieces.

The *slabs method* of Dobkin and Lipton [7] is one of the earliest examples. Every endpoint induces a vertical wall giving rise to $2n + 1$ vertical slabs. A point location is performed by a binary search to locate the correct slab and another search within the slab in $O(\log n)$ time. Preparata [8] introduced a method that avoids the decomposition into $n + 1$ slabs reducing the required space from $O(n^2)$ to $O(n \log n)$. Sarnak and Tarjan [9] went back to the slabs of Dobkin and Lipton and added the idea of persistent data structures, which reduced the space consumption to $O(n)$. Another example for this model is the separating chains method of Lee and Preparata [10]. Combining it with fractional cascading, Edelsbrunner et al. [11], achieved $O(\log n)$ query time as well. For other methods and variants the reader is referred to a comprehensive overview given in [12].

An asymptotically optimal solution is the randomized incremental construction (RIC), which was introduced by Mulmuley [2] and Seidel [3]. In the static setting, it achieves $O(n \log n)$ preprocessing time, $O(\log n)$ query time and $O(n)$ space, all in expectancy. As pointed out in [13], the latter two can even be worst-case guaranteed. It is also claimed there that one can achieve these worst-case bounds in an expected preprocessing time of $O(n \log^2 n)$, but no concrete proof is given. The approach is able to handle dynamic scenes; that is, it is possible to add or delete edges later on. This method is discussed in more detail in Section 2.

Contribution: We present here a major revision of the trapezoidal-map random incremental construction algorithm for planar point location in CGAL. As the previous implementation, it provides a linear size data structure for non-linear subdivisions that can handle static as well as dynamic scenes. The new version is now able to guarantee $O(\log n)$ query time and $O(n)$ space. Following recent changes in the “2D Arrangements” package [14], the implementation now also supports unbounded subdivisions as well as ones that are embedded on two-dimensional parametric surfaces. After a review of the RIC in Section 2, we discuss, in Section 3, the difference between the length \mathcal{L} of the longest search path and the depth \mathcal{D} of the DAG. We prove that the worst-case ratio of \mathcal{D} and \mathcal{L} is $\Theta(n/\log n)$. Moreover, we describe two algorithms for the preprocessing stage that achieve guaranteed $O(n)$ size and $O(\log n)$ query time. Both are based on a verification of \mathcal{L} after the DAG has been constructed: An implemented one that runs in expected $O(n \log^2 n)$ time, and a more efficient one that runs in expected $O(n \log n)$ time. The latter is a very recent addition that was not included in the reviewed submission. However, the solution that is integrated into CGAL is based on a verification of \mathcal{D} . Based on our experimental results, we conjecture that it also achieves expected $O(n \log n)$ preprocessing time. Section 4 demonstrates the performance of the new implementation by comparing our point location in a Voronoi Diagram with the nearest neighbor implementation of the FDH and others. Section 5 presents more details on the new implementation. To the best of our knowledge, this is the only available implementation for guaranteed logarithmic query time point location in general two-dimensional subdivisions.

2 Review of the RIC for Point Location

We review here the random incremental construction (RIC) of an efficient point location structure, as introduced by [2, 3] and described in [13, 15]. For ease of reading we discuss the algorithm in case the input is in general position. Given an arrangement of n pairwise interior disjoint x -monotone curves, a random permutation of the curves is inserted incrementally, constructing the Trapezoidal Map, which is obtained by extending vertical walls from each endpoint upward and downward until an input curve is reached or the wall extends to infinity. During the incremental construction, an auxiliary search structure, a directed acyclic graph (DAG), is maintained. It has one root and many leaves, one for every trapezoid in the trapezoidal map. Every internal node is a binary decision node, representing either an endpoint p , deciding whether a query lies to the left or to the right of the vertical line through p , or a curve, deciding if a query is above or below it. When we reach a curve-node, we are guaranteed that the query point lies in the x -range of the curve. The trapezoids in the leaves are interconnected, such that each trapezoid knows its (at most) four neighboring trapezoids, two to the left and two to the right. In particular, there are no common x -coordinates for two distinct endpoints¹.

¹ In the general case all endpoints are lexicographically compared; first by the x -coordinate and then by the y -coordinate. This implies that two covertical points produce a virtual trapezoid, which has a zero width.

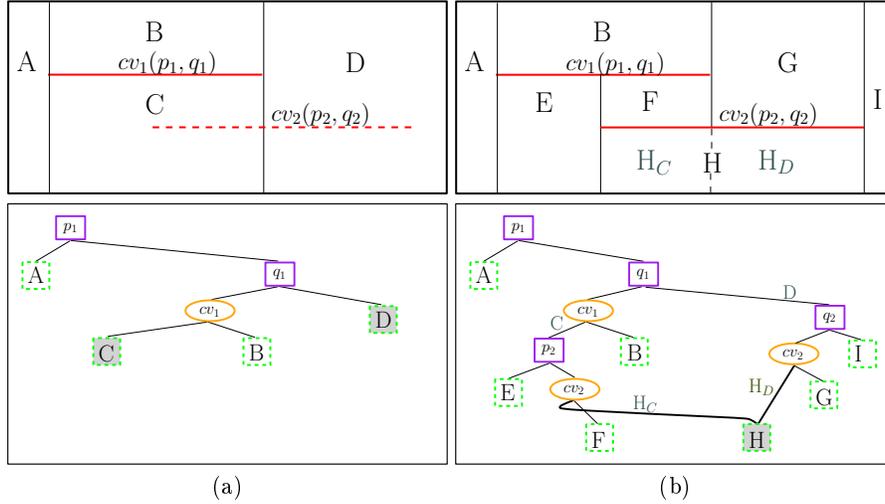


Fig. 1: Trapezoidal decomposition and the constructed DAG for two segments cv_1 and cv_2 : (a) before and (b) after the insertion of cv_2 . The insertion of cv_2 splits the trapezoids C, D into E, F, H_C and G, I, H_D , respectively. H_C and H_D are merged into H , as they share the same top (and bottom) curves.

When a new x -monotone curve is inserted, the trapezoid containing its left endpoint is located by a search from root to leaf. Then, using the connectivity mechanism described above, the trapezoids intersected by the curve are gradually revealed and updated. Merging new trapezoids, if needed, takes time that is linear in the number of intersected trapezoids. The merge makes the data structure become a DAG (as illustrated in Figure 1) with expected $O(n)$ size, instead of an $\Omega(n \log n)$ size binary tree [6]. For an unlucky insertion order the size of the resulting data structure may be quadratic, and the longest search path may be linear. However, due to the randomization one can expect $O(n)$ space, $O(\log n)$ query time, and $O(n \log n)$ preprocessing time.

3 On the Difference between Paths and Search Paths

As shown in [13], one can build a data structure, which guarantees $O(\log n)$ query time and $O(n)$ size, by monitoring the size and the length of the longest search path \mathcal{L} during the construction. The idea is that as soon as one of the values becomes too large, the structure is rebuilt using a different random insertion order. It is shown that only a small constant number of rebuilds is expected. However, in order to retain the expected construction time of $O(n \log n)$, both values must be efficiently accessible. While this is trivial for the size, it is not clear how to achieve this for \mathcal{L} . Hence, we resort to the depth \mathcal{D} of the DAG, which is an upper bound on \mathcal{L} as the set of all possible search paths is a subset of all paths in the DAG. Thus, the resulting data structure still guarantees a logarithmic query time.

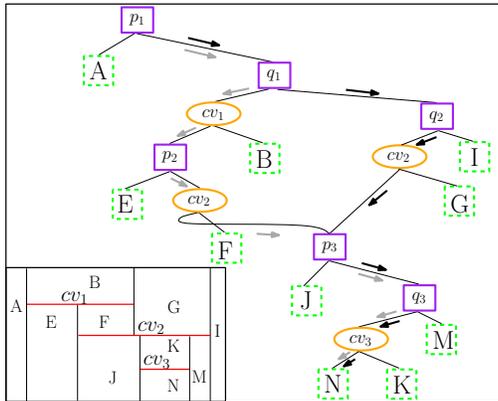
The depth \mathcal{D} can be made accessible in constant time by storing the depth of each leaf in the leaf itself, and maintaining the maximum depth in a separate variable. The cost of maintaining the depth can be charged to new nodes, since

existing nodes never change their depth value. This is not possible for \mathcal{L} while retaining linear space, since each leaf would have to store a non-constant number of values, i.e., one for each valid search path that reaches it. In fact the memory consumption would be equivalent to the data structure that one would obtain without merging trapezoids, namely the trapezoidal search tree, which for certain scenarios requires $\Omega(n \log n)$ memory as shown in [6]. In particular, it is necessary to merge as (also in practice) the sizes of the resulting search tree and the resulting DAG considerably differ.

In Section 3.1 we show that for a given DAG its depth \mathcal{D} can be linear while \mathcal{L} is still logarithmic, that is, such a DAG would trigger an unnecessary rebuild. It is thus questionable whether one can still expect a constant number of rebuilds when relying on \mathcal{D} . Our experiments in Subsection 3.3 show that in practice the two values hardly differ, which indicates that it is sufficient to rely on \mathcal{D} . However, a theoretical proof to consolidate this is still missing. Subsection 3.2 provides efficient preprocessing solutions for the static scenario (where all segments are given in advance). As such, we see it as a concretization of, and an improvement over, the claim mentioned in [13].

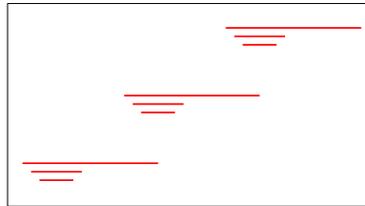
3.1 Worst Case Ratio of Depth and Longest Search Path

The figure to the right shows the DAG of Figure 1 after inserting a third segment. There are two paths that reach the trapezoid N (black and gray arrows). However, the gray path is not a valid search path, since all points in N are to the right of q_1 ; that is, such a search would never visit the left child of q_1 . It does, however, determine the depth of N , since it is the longer path of the two. In the sequel we use this observation to



construct an example that shows that the ratio between \mathcal{D} and \mathcal{L} can be as large as $\Omega(n/\log n)$. Moreover, we will show that this bound is tight.

We start by constructing a simple-to-demonstrate lower bound that achieves $\Omega(\sqrt{n})$ ratio between \mathcal{D} and \mathcal{L} . Assuming that $n = k^2 \in \mathbb{N}$, the construction consists of k blocks, each containing k horizontal segments. The blocks are arranged as depicted in the figure to the right. Segments are inserted from



top to bottom. A block starts with a large segment at the top, which we call the *cover segment*, while the other segments successively shrink in size. Now the next block is placed to the left and below the previous block. Only the cover segment of this block extends below the previous block, which causes a merge as illustrated in Figure 2. All $k = \sqrt{n}$ blocks are placed in this fashion. This

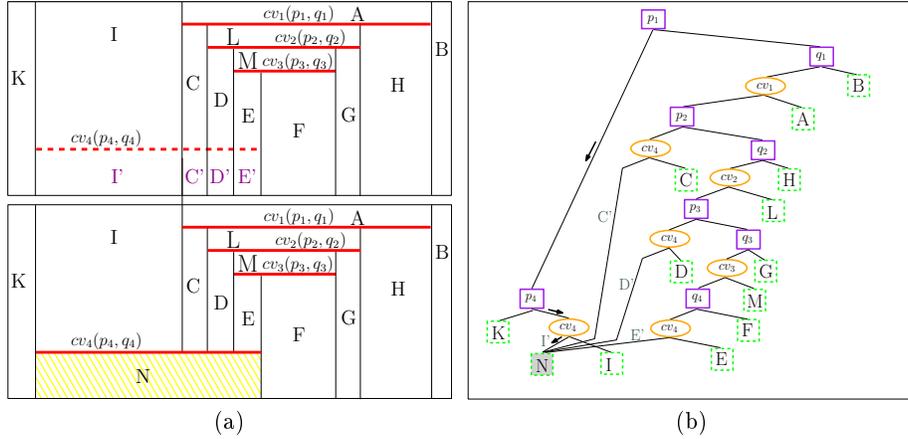
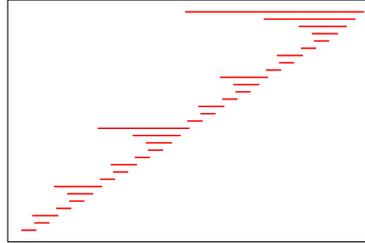


Fig. 2: (a) The trapezoidal-map after inserting cv_4 . The map is displayed before and after the merge of I' , C' , D' , and E' into N , in the top and bottom illustrations, respectively. A query path to the region of I' in N will take 3 steps, while the depth of N in this example is 11.

construction ensures that each newly inserted segment intersects the trapezoid with the largest depth, which increases \mathcal{D} . The largest depth of $\Omega(n)$ is finally achieved in the trapezoid below the lowest segment. However, the actual search path into this trapezoid has only $O(\sqrt{n})$ length, since for each previous block it only passes through one node in order to skip it and $O(\sqrt{n})$ in the last block.

The following construction, which uses a recursive scheme, establishes the lower bound $\Omega(n/\log n)$ for \mathcal{D}/\mathcal{L} . Blocks are constructed and arranged in a similar fashion as in the previous construction. However, this time we have $\log n$ blocks, where block i contains $n/2^i$ segments. Within each block we then apply the same scheme recursively as depicted in the figure to the right. Again segments are inserted top to bottom such that the depth of $\Omega(n)$ is achieved in the trapezoid below the lowest segment. The fact that the lengths of all search paths are logarithmic can be proven by the following argument. By induction we assume that the longest path within a block of size $n/2^i$ is some constant times $(\log_2 n - i)$. Obviously this is true for a block containing only one segment. Now, in order to reach block i with $n/2^i$ segments, we require $i - 1$ comparisons to skip the $i - 1$ st preceding blocks. Thus in total the search path is of logarithmic length.



Theorem 1. *The $\Omega(n/\log n)$ worst-case lower bound on \mathcal{D}/\mathcal{L} is tight.*

Proof. Obviously, \mathcal{D} of $O(n)$ is the maximal achievable depth, since by construction each segment can only appear once along *any* path in the DAG. It remains to show that for any scenario with n segments there is no DAG for which \mathcal{L} is

smaller than $\Omega(\log n)$. Since there are n segments, there are at least n different trapezoids having these segments as their top boundary. Let T be a decision tree of the optimal search structure. Each path in the decision tree corresponds to a valid search path in the DAG and vice versa. The depth of T must be larger than $\log_2 n$, since it is only a binary tree. We conclude that the worst case ratio of \mathcal{D} and \mathcal{L} is $\Theta(n/\log n)$. \square

3.2 Efficient Solutions for Static Subdivisions

We first describe an algorithm for static scenes that runs in expected $O(n \log^2 n)$ time, constructing a DAG of linear size in which \mathcal{L} is $O(\log n)$. The result is based on the following lemma. The proof is given in the full version of this paper [16].

Lemma 1. *Let S be a planar subdivision induced by n pairwise interior disjoint x -monotone curves. The expected size of the trapezoidal search tree \mathcal{T} , which is constructed as the RIC above but without merges, is $O(n \log n)$.*

The following algorithm `compute_max_search_path_length` computes \mathcal{L} in expected $O(n \log^2 n)$ time. Starting at the root it descends towards the leaves in a recursive fashion. Taking the history of the current path into account, each recursion call maintains the interval of x values that are still possible. Thus, if an x -coordinate of a point node is not contained in the interval the recursion does not need to split. This means that the algorithm essentially mimics \mathcal{T} (as it would have been constructed), since the recursion only follows possible search paths. By Lemma 1 the expected number of leaves of \mathcal{T} , and thus of search paths, is $O(n \log n)$. Since the expected length of a query is $O(\log n)$ this algorithm takes expected $O(n \log^2 n)$ time.

Definition 1. $f(n)$ denotes the time it takes to verify that, in a linear size DAG constructed over a planar subdivision of n x -monotone curves, \mathcal{L} is bounded by $c \log n$ for a constant c .

Theorem 2. *Let S be a planar subdivision with n x -monotone curves. A point location data structure for S , which has $O(n)$ size and $O(\log n)$ query time in the worst case, can be built in $O(n \log n + f(n))$ expected time, where $f(n)$ is as defined above.*

Proof. The construction of a DAG with some random insertion order takes expected $O(n \log n)$ time. The linear size can be verified trivially on the fly. After the construction the algorithm `compute_max_search_path_length` is used to verify that \mathcal{L} is logarithmic. The verification of the size and \mathcal{L} may trigger rebuilds with a new random insertion order. However, according to [13], one can expect only a constant number of rebuilds. Thus, the overall expected runtime remains expected $O(n \log n + f(n))$. \square

The verification process described above takes expected $O(n \log^2 n)$ time. However, one can do better as we briefly sketch next. Let T be the collection of

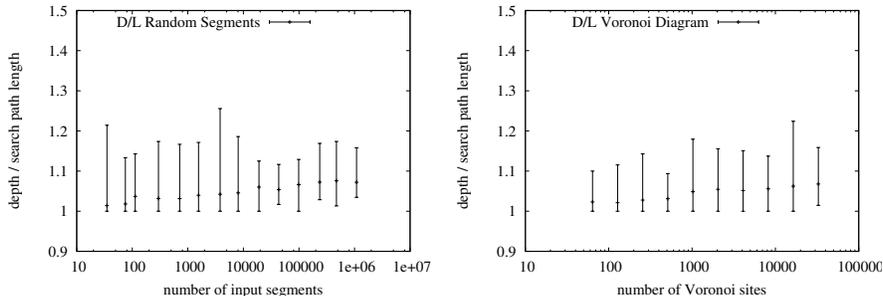


Fig. 3: Ratio of \mathcal{D} and \mathcal{L} in two scenarios: random segments (left), Voronoi diagram of random points (right). Plots show average value with error bars.

all the trapezoids created during the construction of the DAG, including intermediate trapezoids that are later killed by the insertion of later segments. Let $\mathcal{A}(T)$ denote the arrangement of the trapezoids. The *depth* of a point p in the arrangement is defined as the number of trapezoids in T that cover p . The key to the improved algorithm is the following observation by Har-Peled.

Observation 1. *The length of a path in the DAG for a query point p is at most three times the depth of p in $\mathcal{A}(T)$.*

It follows that we need to verify that the maximum depth of a point in $\mathcal{A}(T)$ is some constant $c_1 \log n$. We do this as follows. We first transform the collection T into a collection R of axis-parallel rectangles such that the maximum depth in $\mathcal{A}(R)$ is the same as the maximum depth in $\mathcal{A}(T)$. This can be done in $O(n \log n)$ time since the input segments in S are interior pairwise disjoint, and we rely on the separation property of the segments in S stemming from [17]. We then apply an algorithm by Alt and Scharf [18], which detects the maximum depth in an arrangement of n rectangles in $O(n \log n)$ time. Notice that we only apply this verification algorithm on DAGs of linear size. Putting everything together we obtain:

Theorem 3. *Let S be a planar subdivision with n x -monotone curves. A point location data structure for S , which has $O(n)$ size and $O(\log n)$ query time in the worst case, can be built in $O(n \log n)$ expected time.*²

3.3 Experimental Results

Since \mathcal{D} is an upper bound on \mathcal{L} and since \mathcal{D} is accessible in constant time our implementation explores an alternative that monitors \mathcal{D} instead of \mathcal{L} . Though this may cause some additional rebuilds, the experiments in this section give strong

² Theorem 3 has been added to the paper after the ESA review process. The main result of this section in the original submission was Theorem 2 with $f(n) = O(n \log^2 n)$.

evidence that one can still expect $O(n \log n)$ preprocessing time. We compared \mathcal{D} and \mathcal{L} in two different scenarios: random non-intersecting line segments and Voronoi diagram for random sites.³ Each scenario was tested with an increasing number of subdivision edges, with several runs for each input. Figure 3 displays the average \mathcal{D}/\mathcal{L} ratio, and also the minimal and maximal ones. Obviously, the average ratio is close to 1 and never exceeded a value of 1.3.

These experimental results indicate that replacing the test for the length of the longest path \mathcal{L} by the depth \mathcal{D} of the DAG in the randomized incremental construction essentially does not harm the runtime. However, the following conjecture remains to be proven.

Conjecture 1. There exists a constant $c > 0$ such that the runtime of the randomized incremental algorithm, modified such that it rebuilds in case the depth \mathcal{D} of the DAG becomes larger than $c \log n$, is expected $O(n \log n)$, i.e., the number of expected rebuilds is still constant.

4 Nearest Neighbor Search in Guaranteed $O(\log n)$ Time

As stated in the Introduction, we were challenged by the claim of Birn et al. [1] that guaranteed logarithmic nearest-neighbor search can be achieved via efficient point location on top of the Voronoi Diagram of the input points, but that this approach “*does not seem to be used in practice*”. With this section we would like to emphasize that such an approach is available and that it should be considered for use in practice. Using the RIC planar point location, the main advantage would be that query times are stable and independent of the actual scenario.

4.1 Nearest Neighbor Search via Voronoi Diagram

Given a set P of n points, which we wish to preprocess for efficient point location queries, we first create a Delaunay triangulation (DT) which takes $O(n \log n)$ expected time. The Voronoi diagram (VD) is then obtained by dualizing. Using a sweep, the arrangement representing the VD, which has at most $3n - 6$ edges, can be constructed in $O(n \log n)$ time. However, taking advantage of the spatial coherence of the edges, we use a more efficient method that directly inserts VD edges while crawling over the DT. The resulting arrangement is then further processed by our RIC implementation. If Conjecture 1 is true then this takes expected $O(n \log n)$ time. Alternatively, it would have been possible to implement the solution presented in Subsection 3.2 (for which we can prove expected $O(n \log n)$ preprocessing time).

4.2 Nearest Neighbor Search via Full Delaunay Hierarchy

The full Delaunay hierarchy (FDH) presented in [1] is based on the fact that one can find the nearest neighbor by performing a greedy walk on the edges of the Delaunay triangulation (DT). The difference is that the FDH keeps all edges

³ Additional experimental results that include also the scenarios constructed in Section 3.1 can be found at <http://acg.cs.tau.ac.il/projects/lppl>.

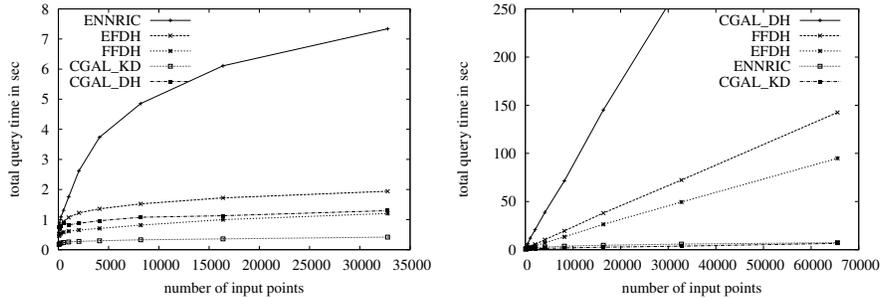


Fig. 4: Performance of 500k nearest-neighbor queries for different methods on two scenarios: (left) random points; (right) circle with outliers.

that appear during the randomized construction [19] of the DT in a flattened n -level hierarchy structure, where level i contains the DT of the first i points. Thus, a walk that starts at the first point is accelerated due to long edges that appeared at an early stage of the construction process while the DT was still sparse. The FDH is a very light, easy to implement, and fast data structure with expected $O(n \log n)$ construction time that achieves an expected $O(\log n)$ query path length. However, a query may take $O(n)$ time since the degree of nodes can be linear. For the experiments we used two exact variants: a basic exact version (EFDH) and a (usually faster) version (FFDH) that first performs a walk using inexact floating point arithmetic and then continues with an exact walk.

4.3 Experimental Results

We compared our implementation for nearest-neighbor search using the RIC point location on the Voronoi-diagram (ENNRIC) to the following exact methods: EFDH, FFDH, CGAL’s Delaunay hierarchy (CGAL_DH) [5], and CGAL’s kd-tree (CGAL_KD).⁴

All experiments have been executed on a Intel(R) Core(TM) i5 CPU M 450 with 2.40GHz, 512 kB cache and 4GB RAM memory, running Ubuntu 10.10. Programs were compiled using g++ version 4.4.5 optimized with `-O3` and `-DNDEBUG`. The left plot of Figure 4 displays the total query time in a random scenario, in which both input points and query points are randomly chosen within the unit square. Clearly, all methods have logarithmic query time, however due to larger constants ENNRIC is slower. The other plot presents a combined scenario of $(n - \lfloor \log n \rfloor)$ equally spaced input points on the unit circle and $\lfloor \log n \rfloor$ random outliers. The queries are random points in the same region. In this experiment the CGAL_KD and ENNRIC are significantly faster and maintain a stable query time. A similar scenario that was tested contains equally spaced input points on a circle and a point in the center with random query points inside the circle. The differences there are even more significant than in the previous scenario. As for the preprocessing time in all tested scenarios, obviously ENNRIC cannot compete with the fast construction time of the other methods.

⁴ Due to similar performance we elided the kd-tree implementation in ANN [20].

5 CGAL’s New RIC Point Location

With this article we announce our revamp of CGAL’s implementation of planar point location via the randomized incremental construction of the trapezoidal map, which is going to be available in the upcoming CGAL release 4.1.

Like the previous implementation by Oren Nechushtan [21], it is part of the “2D Arrangements” package [22] of CGAL. It allows both insertions and deletions of edges. The implementation is exact and covers all degenerate cases. Following the *generic-programming paradigm* [23] it can be easily applied to linear geometry but also to non-linear geometry such as algebraic curves or Bézier curves. The main new feature, and this is what triggered this major revision, is the support for unbounded curves, as it was introduced for the “2D Arrangements” package in [14], enabling point location on two-dimensional parametric surfaces (e.g., spheres, tori, etc.) as well.

In addition we did a major overhaul of the code basis. In particular, we maintain the depth \mathcal{D} of the DAG as described in Section 3 such that \mathcal{D} is accessible in constant time. Thus we can now guarantee logarithmic query time after every operation. Moreover, the data structure now operates directly on the entities of the arrangement. In particular, it avoids copying of geometric data which can significantly reduce the amount of additional memory that is used by the search structure. This is important, since due to the generic nature of the code it is not clear whether the geometric types (user provided) are referenced.

To the best of our knowledge, this is the only available implementation of a point location method with a guaranteed logarithmic query time that can handle two-dimensional subdivisions to this generality. Furthermore, it is the fastest available point location method, in terms of query time, for CGAL arrangements.⁵

6 Open Problem

Prove Conjecture 1, that is, prove that it is possible to rely on the depth \mathcal{D} of the DAG and still expect only a constant number of rebuilds. This solution would not require any changes to the current implementation.

Acknowledgement: The authors thank Sarel Har-Peled for sharing Observation 1, which is essential to the expected $O(n \log n)$ time algorithm for producing a worst-case linear-size and logarithmic-time point-location data structure.

References

1. Birn, M., Holtgrewe, M., Sanders, P., Singler, J.: Simple and fast nearest neighbor search. In: Workshop on Algorithm Engineering and Experiments. (2010) 43–54
2. Mulmuley, K.: A fast planar partition algorithm, i. J. Symb. Comput. **10**(3/4) (1990) 253–280

⁵ See <http://acg.cs.tau.ac.il/projects/lppl> for a comparison to CGAL Landmarks point location [24].

3. Seidel, R.: A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *J. Comput. Geom.* **1** (1991) 51–64
4. Kirkpatrick, D.G.: Optimal search in planar subdivisions. *SIAM J. Comput.* **12**(1) (1983) 28–35
5. Devillers, O.: The Delaunay hierarchy. *Int. J. Found. Comput. Sci.* **13**(2) (2002) 163–180
6. Seidel, R., Adamy, U.: On the exact worst case query complexity of planar point location. *J. Algorithms* **37**(1) (2000) 189–217
7. Dobkin, D.P., Lipton, R.J.: Multidimensional searching problems. *SIAM J. Comput.* **5**(2) (1976) 181–186
8. Preparata, F.P.: A new approach to planar point location. *SIAM J. Comput.* **10**(3) (1981) 473–482
9. Sarnak, N., Tarjan, R.E.: Planar point location using persistent search trees. *Commun. ACM* **29**(7) (1986) 669–679
10. Lee, D.T., Preparata, F.P.: Location of a point in a planar subdivision and its applications. In: *ACM Symposium on Theory of Computing (STOC)*. STOC '76, New York, NY, USA, ACM (1976) 231–235
11. Edelsbrunner, H., Guibas, L.J., Stolfi, J.: Optimal point location in a monotone subdivision. *SIAM J. Comput.* **15**(2) (1986) 317–340
12. Snoeyink, J.: Point location. In Goodman, J.E., O'Rourke, J., eds.: *Handbook of Discrete and Computational Geometry*. CRC Press LLC, Boca Raton, FL (2004) 767–785
13. de Berg, M., van Kreveld, M., Overmars, M., Schwarzkopf, O.: *Computational Geometry: Algorithms and Applications*. Third edn. Springer-Verlag (2008)
14. Berberich, E., Fogel, E., Halperin, D., Melhorn, K., Wein, R.: Arrangements on parametric surfaces I: General framework and infrastructure. *Mathematics in Computer Science* **4** (2010) 67–91
15. Mulmuley, K.: *Computational geometry - an introduction through randomized algorithms*. Prentice Hall (1994)
16. Hemmer, M., Kleinbort, M., Halperin, D.: Improved implementation of point location in general two-dimensional subdivisions. *CoRR* **abs/1205.5434** (2012)
17. Guibas, L.J., Yao, F.F.: On translating a set of rectangles. In: *STOC*. (1980) 154–160
18. Alt, H., Scharf, L.: Computing the depth of an arrangement of axis-aligned rectangles in parallel. In: *Proceedings of the 26th European Workshop on Computational Geometry (EuroCG)*, Dortmund, Germany (March 2010) 33–36
19. Amenta, N., Choi, S., Rote, G.: Incremental constructions con brio. In: *Symposium on Computational Geometry*. (2003) 211–219
20. Mount, D.M., Arya, S.: Ann: A library for approximate nearest neighbor searching <http://www.cs.umd.edu/~mount/ANN/>.
21. Flato, E., Halperin, D., Hanniel, I., Nechushtan, O., Ezra, E.: The design and implementation of planar maps in CGAL. *ACM Journal of Experimental Algorithmics* **5** (2000) 13
22. Wein, R., Berberich, E., Fogel, E., Halperin, D., Hemmer, M., Salzman, O., Zuckerman, B.: 2D arrangements. In: *CGAL User and Reference Manual*. 4.0 edn. CGAL Editorial Board (2012)
23. Austern, M.H.: *Generic Programming and the STL*. Addison-Wesley (1999)
24. Haran, I., Halperin, D.: An experimental study of point location in planar arrangements in CGAL. *ACM Journal of Experimental Algorithmics* **13** (2008)