Deterministic Rectangle Enclosure and Offline Dominance Reporting on the RAM

Peyman Afshani[1], Timothy M. Chan[2] ⋆, and Konstantinos Tsakalidis[3]

[1] MADALGO ⋆⋆, Department of Computer Science, Aarhus University, Denmark
peyman@madalgo.au.dk
[2] David R. Cheriton School of Computer Science, University of Waterloo, Canada
tmchan@uwaterloo.ca
[3] Computer Science and Engineering Department, CUHK, Hong Kong
tsakalid@cse.cuhk.edu.hk

Abstract. We revisit a classical problem in computational geometry that has been studied since the 1980s: in the *rectangle enclosure* problem we want to report all k enclosing pairs of n input rectangles in 2D. We present the first deterministic algorithm that takes $O(n \log n + k)$ worst-case time and $O(n)$ space in the word-RAM model. This improves previous deterministic algorithms with $O((n \log n + k) \log \log n)$ running time. We achieve the result by derandomizing the algorithm of Chan, Larsen and Pătraşcu [SoCG'11] that attains the same time complexity but in expectation.

The 2D rectangle enclosure problem is related to the *offline dominance range reporting* problem in 4D, and our result leads to the currently fastest deterministic algorithm for offline dominance reporting in any constant dimension $d \geq 4$.

A key tool behind Chan et al.'s previous randomized algorithm is *shallow cuttings for 3D dominance ranges*. Recently, Afshani and Tsakalidis [SODA'14] obtained a deterministic $O(n \log n)$-time algorithm to construct such cuttings. We start with an improved deterministic construction algorithm that runs in $O(n \log \log n)$ time in the word-RAM; this result is of independent interest. Many additional ideas are then incorporated, including a linear-time algorithm for *merging* shallow cuttings.

1 Introduction

We study the problem of *rectangle enclosure*: given a set of n axis-aligned rectangles on the plane, report all k pairs (r_1, r_2) of input rectangles where r_1 completely encloses r_2. This is a classic problem in the field of computational geometry [16] with applications to VLSI design, image processing, computer graphics and databases [17,13,10,5].

⋆ The work of this author was done during his visit to the Department of Computer Science and Engineering, Hong Kong University of Science and Technology.

⋆⋆ Center for Massive Data Algorithmics - a Center of the Danish National Research Foundation.

Previous Results. An early paper by Bentley and Wood [3] presented an $O(n \log n + k)$ worst-case time and linear-space algorithm for the related *rectangle intersection* problem (reporting all k pairs (r_1, r_2) where r_1 intersects r_2), raising the question whether the same bound could be achieved for rectangle enclosure. Vaishnavi and Wood [17] first addressed the question presenting an $O(n \log^2 n + k)$-time algorithm that uses $O(n \log^2 n)$ space. Lee and Preparata [13] improved the space bound to linear.

Further improvements were discovered in the 1990s. The linear-space algorithm of Gupta, Janardan, Smid and Dasgupta [10] and an alternative implementation by Lagogiannis, Makris and A. Tsakalidis [12] take $O((n \log n + k) \log \log n)$ worst-case time. Recently, Chan, Larsen and Pătraşcu [5] succeeded in improving the running time to the desired bound of $O(n \log n + k)$ using linear space. However their algorithm uses randomization and thus the time bound holds in expectation. All presented algorithms operate in the standard RAM model with word size $w \geq \log n$ and with input in rank-space.

It is well-known that the rectangle enclosure problem is reducible to the 4D version of the *offline dominance reporting* problem: given n input and query points in $\mathbb{R}^d$, report the input points that are dominated by each query point ($(p_1, \ldots, p_d)$ is dominated by $(q_1, \ldots, q_d)$ if $p_i < q_i$ for all i). For the reduction it suffices to map each input rectangle $[x_1, x_2] \times [y_1, y_2]$ to a 4D point $(x_1, y_1, -x_2, -y_2)$ and equate the query points with the input points.

Offline dominance reporting is a fundamental problem in the area of orthogonal range searching; it has even found applications outside of computational geometry [**?**]. Chan, Larsen and Pătraşcu's result implies an algorithm with $O(n \log^{d-3} n + k)$ expected time for any constant dimension $d \geq 4$, where k is the total number of reported points. However their algorithm is randomized. The best deterministic algorithm known requires $O(n \log^{d-2} n + k)$ or $O((n \log^{d-3} n + k) \log \log n)$ time.

Our Contributions. We present the first deterministic algorithm for rectangle enclosure that takes $O(n \log n + k)$ worst-case time and $O(n)$ space in the standard word-RAM model. Our result thus improves over the previous deterministic algorithms of Gupta et al. [10] and Lagogiannis et al. [12] and removes randomization from the algorithm of Chan et al. [5].

Our approach also gives the currently fastest deterministic algorithm for the offline dominance reporting problem for any constant dimension $d \geq 4$, with worst-case running time $O(n \log^{d-3} n + k)$.

Our Approach. Our algorithm may be viewed as a derandomization of Chan, Larsen and Pătraşcu's [5], however significant new ideas are required.

In Chan et al.'s algorithm, randomization was used to construct combinatorial objects that have properties similar to those of *shallow cuttings for 3D dominance ranges*. Shallow cuttings were introduced by Matoušek [**?**], and a complicated randomized $O(n \log n)$-time algorithm was given by Ramos [**?**] for constructing shallow cuttings in the more general setting of 3D halfspace ranges.

In a recent SODA'14 paper, Afshani and K. Tsakalidis [2] presented a deterministic $O(n \log n)$-time algorithm for constructing shallow cuttings for 3D dominance ranges using linear space in the pointer machine model. In Section 2 we first improve their algorithm to run in $O(n \log\log n)$ worst-case time on the word-RAM. As an immediate consequence we obtain a deterministic algorithm for offline 3D dominance reporting that takes $O(n \log\log n + k)$ worst-case time and linear space in the word-RAM (Section 4); this result is new. Previously only $O(n \log n + k)$ worst-case time could be achieved using linear space [1,15,2]

Much further work is needed to derive our result on offline 4D dominance reporting and 2D rectangle enclosure. The crucial new ingredient is an algorithm that can *merge* two shallow cuttings for 3D dominance ranges in *linear* time; this result is obtained by modifying our shallow cutting construction algorithm in interesting ways and is described in Section 3. Then in Section 5 we use an intricate combination of Chan et al.'s approach with the deterministic shallow cutting construction and merging subroutines to achieve our final result. The combination requires a re-organization of the previous algorithm. In particular we isolate a subproblem we call *tree point location* for which we obtain a deterministic algorithm by incorporating planar separators [14] to ideas from Chan et al. This problem may be viewed as a new 2D variant of *fractional cascading* [7] and is thus of independent interest.

Notation and Definitions. Point p *dominates* point q if and only if each coordinate of p is greater than or equal to that of q. To avoid ambiguity for points on the plane we use the term "*covers*" (instead of "dominates"). Let P be a set of n points in $\mathbb{R}^d$. The *level* of any point $p \in \mathbb{R}^d$ (with respect to P) is the number of points in P that are dominated by p. The region of $\mathbb{R}^d$ dominated by p is called a *cell*. The conflict list of a cell is the subset of P that lies inside the cell.

A k-*shallow cutting* for dominance ranges on pointset P is a set of *vertices* (points) S such that (i) every vertex in S has level at most $c_{\max}k$ in P for a constant $c_{max} > 1$ and (ii) any point in $\mathbb{R}^d$ with level in P at most k is dominated by some vertex of S. Shallow cutting S is *optimal* when it contains at most $c_{\max}\frac{n}{k}$ vertices. A planar shallow cutting has the shape of an orthogonal *staircase* curve $c_1 d_1 c_2 \dots d_{t-1} c_t$ of alternating vertical line segments $c_i d_i = [c_i(x)] \times [c_i(y), d_i(y)]$ and horizontal line segments $d_i c_{i+1} = [d_i(x), c_{i+1}(x)] \times [d_i(y)]$ (Fig. 1a).

2 Construction of 3D Dominance Shallow Cuttings

Theorem 1. *An optimal k-shallow cutting for 3D dominance ranges on n input points and for any integer k can be constructed deterministically in $O(n \log\log n)$ worst-case time and $O(n)$ space.*

Algorithm Sketch. Following [2] we sort the points and sweep a plane parallel to the xy-plane considering the points in non-decreasing z-coordinate. This reduces the problem to the problem of maintaining a 2D shallow cutting under only

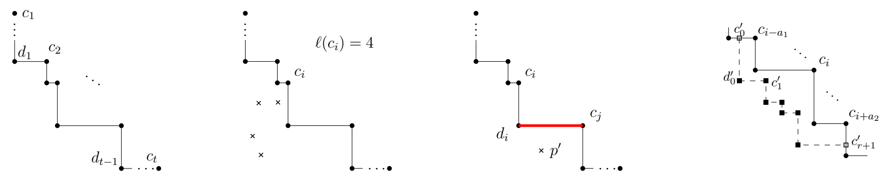

(a) (b) (c) (d)

insertions, specifically the planar shallow cutting of the xy-projection of the points of P that lie *below* the sweep plane. When the point with the next highest z-coordinate is considered (i.e., the next insertion), we update the corresponding 2D shallow cutting. The idea is that we do not need to change the planar shallow cutting for most insertions. However each insertion can increase the level of the shallow cutting corners. Thus, once in a while, the planar shallow cutting needs to be fixed. This is done by removing some consecutive parts of it and then adding new staircase "patches" that are covered by the parts just removed (details will follow). Every time a planar shallow cutting cell is removed, a 3D shallow cutting cell is created using the z-coordinate of the sweep plane (i.e., when a planar staircase corner with coordinates (x, y) is removed, a 3D vertex (x, y, z) is created where z is the coordinate of the sweep plane). Finally when the sweep terminates, the remaining planar shallow cutting cells are turned into 3D shallow cutting cells using $+\infty$ as the z-coordinate. It is easily verified that the produced 3D shallow cutting of P is correct and that its size is equal to the number of planar shallow cutting corners removed throughout the algorithm plus the number of planar shallow cutting corners that remain when the sweep terminates.

Remark. The sketched algorithm is a variant of Afshani and Tsakalidis' [2, Section 3] with the significant difference that it is insertion-only (sweeping upwards) as opposed to their deletion-only algorithm (sweeping downwards). Their algorithm has the advantage that it can compute $O(\log n)$ different shallow cuttings in total $O(n \log n)$ time. However a crucial ingredient of that algorithm is an auxiliary data structure ([2, Lemma 2]) that needs to be updated at every sweep point. Unfortunately, we cannot see a way to update the auxiliary data structure in $O(\log \log n)$ time in the word RAM model. Fortunately, as we shall see later, we can achieve the desired $O(n \log \log n)$ running time without any auxiliary data structures, by just changing the direction of the sweep.

The invariant. The planar shallow cutting is maintained in the form of a staircase $S = c_1 d_1 \dots c_t$, composed of inner corners $d_1, \dots, d_{t-1}$ and outer corners $c_1, \dots, c_t$. The outer corners c_1 and c_t are (conceptually) at infinity, i.e. $c_1(y) = +\infty$ and $c_t(x) = -\infty$ (Fig.1a). We maintain the invariant that the inner corners dominate at least k and the outer corners at most $10k$ points.

Details. We now discuss the details of the algorithm. Let $p = (x, y, z)$ be the next point swept by the sweep plane. Remember that we need to insert the point

$p' = (x, y)$ into the dynamic planar shallow cutting. To do that we maintain the following structures: first, a dynamic van Emde Boas tree on the x-coordinates of the staircase, and second, for every point corner c of the staircase, we keep track of the number of points it covers, $\ell(c)$ (Fig. **??**), as well as a linked list containing them, $L(c)$ (i.e., the conflict list and its size). Using these we can now insert the point p'. The dynamic van Emde Boas tree enables us to find the inner corner d_i immediately to the left of the point p', helping us decide whether p' lies above or below the staircase (Fig. 1d). In the former case, we are done. In the latter case, for every outer corner c_j that covers p', we increase $\ell(c_j)$ by one and then append p' to $L(c_j)$. If for all such corners we still have $\ell(c_j) \leq 10k$, then the invariant is maintained and thus we are done. However, it is possible that for some corners this invariant is violated. Below we describe how to "patch" such violated invariants.

Complexity. Sorting by z-coordinate takes $O(n \log\log n)$ time in total [11]. Finding d_i takes $O(\log\log n)$ time [8] and thus $O(n \log\log n)$ time in total. It turns out the rest of the algorithm consumes linear time. If there are $m(p')$ corners that cover p', then updating their relevant information takes $O(m(p'))$ time. Note that p' is now added to the conflict lists of $m(p')$ corners. Notice that since the size of each conflict list $O(k)$, the total running time of this step is $O(Tk)$ where T is size of the shallow cutting. If we can prove that $T = O(n/k)$, then this running time is linear. Now we describe how to maintain the invariant which also guarantees the upper bound on T.

Patching. Let c_i be the leftmost outer corner whose invariant is violated. Let a_1 and a_2 be the largest integers such that all the outer corners $c_{i-a_1}, d_{i-a_1+1}, \ldots, c_{i+a_2}$ all have levels greater than $3k$. To patch the staircase, we begin by finding a new outer corner c'_0 at the same y-coordinate as c_{i-a_1}, such that it covers $3k$ points, as depicted in Fig. **??**. c'_0 can be found in $O(k)$ time using a linear time selection algorithm on the conflict list of c_{i-a_1}. Next, we find the inner corner d'_0 directly below c'_0 that covers k points. Now, we alternate between finding new outer and inner corners: at the j-th step, we find the outer corner c'_j that dominates $2k$ points, and then the inner corner d'_j that dominates k points. As before, using the right conflict list, each of these corners can be found in $O(k)$ time. The patching is terminated at a point c'_{r+1} with level $3k$ that lies below the outer corner c_{i+a_2} (Fig. **??**). Finally, the new outer and inner corners (from c'_0 to c'_{r+1}) are incorporated into the staircase, the old corners $(c_{i-a_1}, \ldots, c_{i+a_2})$ are removed, and the van Emde Boas tree is also updated to reflect the changes in the staircase.

Analysis of Patching. If it easy to see that the overall cost of patching is $O(Tk)$ since each new inner or outer corner can be found in $O(k)$ time. Moreover, the facts that the levels of the removed corners differ from the levels of the created ones by at least k (except for only c'_0 and c'_{r+1}) and that at least 5 corners are created for every patch, suffices to claim that $T \leq C_0 \frac{n}{k}$ for a positive constant

C_0. We set $c_{\max} := \max\{C_0, 10\}$. A detailed proof on the symmetric approach is found in [2, Section 3].

3 Merging Two 3D Dominance Shallow Cuttings

We begin by a naive merging algorithm.

Lemma 1. *Consider two pointsets P_1 and P_2 that contain n_1 and n_2 points respectively. For $i = 1, 2$, assume we are given a k_i-shallow cutting C_i on P_i of size m_i such that the conflict list of every cell contains at most $\beta_i k_i$ points of P_i. A $(\min\{k_1, k_2\})$-shallow cutting C on the union pointset $P = P_1 \cup P_2$ can be built in $O((m_1 + m_2)\log\log(m_1 + m_2))$ time, such that C contains $O(m_1 + m_2)$ cells and the conflict list of every cell in C contains at most $\beta_1 k_1 + \beta_2 k_2$ input points.*

Proof. Let R be the subset of $\mathbb{R}^3$ that is dominated by at least one vertex in C_1 and at least one vertex in C_2. It is easily seen that R is orthogonally convex and in fact any orthogonal ray to $y = -\infty$ or $x = -\infty$ crosses the boundary of R at most one. Thus, the complexity of the boundary of R is $O(|C_1| + |C_2|) = O(m_1 + m_2)$. Furthermore, the boundary of R can be computed with a straightforward sweep plane algorithm in $O((m_1 + m_2)\log\log(m_1 + m_2))$ time by employing a van Embe Boas tree as the search structure [8]. The shallow cutting C is defined by the vertices of the boundary of R. It is clear that every such vertex dominates either k_1 points from P_1 or k_2 points from P_2 and thus it dominates at least $\min\{k_1, k_2\}$ points of P. Similarly, each vertex on R can dominate at most $\beta_1 k_1$ points of P_1 and $\beta_2 k_2$ points of P_2. □

While the above merging algorithm is quite fast, it worsens the constants behind the parameters of the shallow cutting and thus it cannot be applied more than a constant number of times. In the next theorem, we show how such a "bad" shallow cutting can be refined into an optimal one. For this purpose we also use the following lemma.

Lemma 2. *[4, Theorem 4.3] Online planar orthogonal point location queries on a planar orthogonal decomposition of size n regions are supported in $O(\log\log n)$ worst-case time and $O(n)$ space.*

Theorem 2. *Let P be a set of n points in $\mathbb{R}^3$ with presorted z-coordinates and let C be a k-shallow cutting on P of size $\alpha n/k$, where the conflict list of every cell has size at most βk, for arbitrary constants $\alpha, \beta > 0$. C can be refined into an optimal k'-shallow cutting C' on P in $O(n + \frac{n}{k}\log\log n)$ time, such that $k' = \frac{k}{c_{\max}}$ for a universal constant $c_{\max}$ that does not depend on α and β. C' contains at most $c_{\max}\frac{n}{k'}$ cells and the conflict list of every cell in C' contains at most $c_{\max}k'$ points of P.*

Proof. We build C' using the plane sweep algorithm from the previous section. To review, the algorithm maintains a planar shallow cutting in the form of a staircase S'. To process the next point p, a predecessor query is used to find one corner of the staircase that covers the projection p' of p. As we noted during the proof of Theorem 1, other than this predecessor search, the rest of the algorithm runs in linear time. To remove this bottleneck, we use C to augment each point of P with a correct pointer to a staircase corner that covers it, thus negating the need for the predecessor search.

Hence we project the cutting C on the xy-plane and obtain an orthogonal planar decomposition of disjoint polygonal *regions* in order to support online planar orthogonal point location queries, i.e. report the region that any given query point lies in. Thus, Lemma 2 enables us to perform the following operation in $O(\log\log n)$ time: given a point q in the xy-plane, find the shallow cutting vertex $C(q)$ in C with the largest z-coordinate whose projection covers q.

Observe that the level of every vertex in C' is at most $c_{\max}k' = k$ which implies every vertex of C' is contained in at least one cell of C. We thus maintain one additional invariant in our sweep. Consider a staircase corner $v \in S'$, the shallow cutting vertex $C(v) \in C$ and its conflict list $\ell(C(v))$. We maintain the invariant that if the projection p' of a point $p \in \ell(C(v))$ lies below the staircase S', then p is assigned a pointer to a staircase corner that covers p'. This invariant removes the need for the predecessor search during the sweep.

By looking at the algorithm in Section 2, it is clear that the invariant can only be violated when a new staircase corner v is created on S' (during the patching phase). To fix the invariant, by Lemma 2 we can find the vertex $C(v)$ and its conflict list $\ell(C(v))$. We go through each point in $\ell(C(v))$ and if its projection is covered by v, then we assign it a pointer to v. This takes $O(k)$ time which is proportional to the time required for creating the staircase corner v. Thus, this incurs only a $O(\log\log n)$ additive time per shallow cutting vertex. $\square$

Corollary 1. *Given two pointsets $A, B \in \mathbb{R}^3$ presorted by z-coordinate and their respective k-shallow cuttings, for an integer $k = \Omega(\log\log(|A|+|B|))$, an optimal $\frac{k}{c_{\max}}$-shallow cutting on the union pointset $A \cup B$ can be constructed deterministically in $O(|A|+|B|)$ worst-case time.*

4 Offline 3D Dominance Reporting

Theorem 3. *Offline 3D dominance reporting on n input points, n query points and k reported points can be solved deterministically in $O(n\log\log n + k)$ worst-case time and $O(n)$ space.*

Proof. We follow the approach of Afshani [1] for online 3D dominance reporting queries in internal memory. We presort all coordinates [11] and construct a single $(\log n)$-shallow cutting for 3D dominance ranges by using the algorithm of Theorem 1 in $O(n\log\log n)$ worst-case time. We *assign* every query point to a cell of the cutting whose vertex dominates it by online planar point location queries on a planar orthogonal decomposition obtained from the projection of the cutting. By Lemma 2 this takes $O(n\log\log n)$ worst case time [4]. We resolve all the assigned queries by solving independently for each of the $O(\frac{n}{\log n})$

cells an offline 3D dominance reporting *subproblem* on the conflict list of the cell. The online 3D dominance reporting algorithm of [15] reports all (at most k) output conflict list points in $O(\frac{n}{\log n} \log n \log\log n + k) = O(n \log\log n + k)$ total worst-case time. For the remaining *unresolved* queries we solve a single offline 3D dominance reporting problem on all input points. Each unresolved query reports $\Omega(\log n)$ points and thus the more expensive $O(n \log n + k)$-time algorithm of [15] takes $O(k)$ total time to resolve them. □

Corollary 2. *Offline 3D dominance reporting on* n *input points,* n *query points and* k *reported points can be solved deterministically in* $O(n+k)$ *worst-case time and* $O(n)$ *space, when* $n < \bar{w}^{O(1)}$ *and a global look-up table has been constructed in* $2^{O(\bar{w})}$ *worst case time for a parameter* $\bar{w} < w$.

Proof. We modify the proof of Theorem 3. Particularly to construct the $(\log n)$-shallow cutting we replace the van Emde Boas tree [8] with atomic heaps [9]. This increases the $O(\log\log n)$ searching cost to $O(\log_{\bar{w}} n)$, for a total of $O(n \log_{\bar{w}} n) = O(n)$ worst-case construction time. To assign the queries we sweep a line parallel to the y-axis maintaining a dynamic search tree implemented as an atomic heap for a total of $O(n \log_{\bar{w}} n) = O(n)$ worst-case time. The conflict lists of each cell have size $O(\log n) = O(\log(\bar{w}^c)) = o(\bar{w})$. Thus by using a look-up table, each assigned query can be resolved in $O(1)$ worst-case time for a total of $O(\frac{n}{\log n}) = O(n)$ worst-case time. The cost of the remaining unresolved queries is charged to the output by using the algorithm of [15]. □

5 Offline 4D Dominance Reporting

A preliminary $O(n \log n \log\log n + k)$ worst-case time and linear-space algorithm for the rectangle enclosure problem is implied by Theorem 3. We obtain a faster deterministic algorithm for rectangle enclosure and thus also for offline dominance reporting as corollaries of the following theorem.

Theorem 4. *Offline 4D dominance reporting problem on* n *input points,* n *query points and* k *reported points can be solved deterministically in* $O(n \log n + k)$ *worst-case time and* $O(n)$ *space.*

Algorithm. We follow the approach of Chan, Larsen and Pătraşcu [5, Section 4.3]. We build a complete binary range tree $\mathcal{T}$ over the fourth coordinate of the input points and we associate every query point with the left siblings (if they exist) of the nodes on the path from the root of $\mathcal{T}$ to the leaf that contains its successor input point. To obtain the total output it suffices to solve in every node of $\mathcal{T}$ an offline 3D dominance reporting problem between the 3D projections of its input pointset and its associated query pointset.

For this reason first each node is *equipped* with an optimal $O(K)$-shallow cutting for 3D dominance ranges of its input points for a fixed parameter K. Then every associated query point is *assigned* to a cell of the equipped cutting whose vertex dominates it. All assigned queries are resolved by solving independently for each cell in $\mathcal{T}$ an offline 3D dominance reporting *subproblem* on the cell's conflict list. This leaves $O(\frac{k}{K})$ queries *unresolved*, i.e. assigned to a node where no vertex of the equipped cells dominates the query point, since each unresolved query reports $\Omega(K)$ points. We also decrease the number of input points to $O(\frac{k}{K})$ in the same way by repeating the above with the roles of the input and query points reversed. Finally we solve a single offline 4D dominance reporting problem on all the remaining input points and unresolved queries.

Complexity. Particularly we equip a node at level i of $\mathcal{T}$ with an optimal K_i-shallow cutting on the input points of its subtree such that $K_i = O(K)$. Thus $O(\frac{n}{K_i})$ subproblems are defined at every level of $\mathcal{T}$. Hence the total time cost of the algorithm on n input and query points and k reported points is $T_{4D}(n, k) \leq$

$$2 \cdot \left(T_E(n) + T_A(n) + \sum_{i=1}^{\log n} \sum_{j=1}^{O(\frac{n}{K_i})} T_{3D}(O(K_i), k_{ij}) \right) + T_{4D}(O(\frac{k}{K}), k) + O(k)$$

This is $O(n \log n + k)$ when (i) the time cost $T_E(n)$ of equipping the nodes with shallow cuttings and $T_A(n)$ of assigning the queries to the equipped cells is $O(n \log n)$, (ii) an offline 3D dominance reporting subproblem takes time linear to $O(K_i)$ input and k_{ij} output points and (iii) an offline 4D dominance reporting problem on $O(k/K)$ input and query points takes time linear to the total output size. To achieve (iii) we use the more expensive algorithm of [13] since $O(\frac{k}{K} \log^2 n + k) = O(k)$ when $K := O(\log^{2+\epsilon} n)$ for a constant $\epsilon > 0$. To achieve (ii) we use the algorithm of Corollary 3 at every equipped cell in $\mathcal{T}$ in combination with a single global look-up table that takes $2^{O(\bar{w})} = O(n)$ time to construct. The subproblem at the j-th equipped cell of a node at level i in $\mathcal{T}$ takes $O(K_i + k_{ij})$ time since $K_i = O(\log^{2+\epsilon} n) = \bar{w}^{O(1)}$. All subproblems at level i report $k_i = \sum_{j=1}^{O(\frac{n}{K_i})} k_{ij} = O(k)$ points in $O(\frac{n}{K_i} K_i + k_i) = O(n + k)$ time, for a total of $O(n \log n + k)$ time.

Equipping The Nodes With Optimal Shallow Cuttings. To achieve (i) first we show how to equip the nodes of $\mathcal{T}$ with optimal $O(K)$-shallow cuttings in $O(n \log n)$ time. We construct optimal K-shallow cuttings for the nodes at every level of $\mathcal{T}$ that is a multiple of $\delta \log \log n$, for a constant $\delta > 0$, using the algorithm of Theorem 1. In total this takes $O(n \log \log n \cdot \frac{\log n}{\delta \log \log n}) = O(n \log n)$ time. To equip the remaining nodes we consider all $O(\frac{n}{\log n})$ disjoint subtrees $\mathcal{T}' \in \mathcal{T}$ with $\log n$ leaves each where an optimal K-shallow cutting has already been constructed. We equip the nodes at every at level j of $\mathcal{T}'$ with an optimal K_j-shallow cutting on its input points in $\mathcal{T}$ by merging the optimal K_{j-1}-shallow cuttings of its two children nodes with the algorithm of Corollary 2. This takes $O(n)$ time per level of $\mathcal{T}$ and thus $O(n \log n)$ total time (since $K = \Omega(\log \log n)$) when $K_j < c_{\max} K_{j-1} \Rightarrow K_j < \frac{K}{c_{\max}^j}$. Thus $K_j = O(K)$ for all $j \in [1, h-1]$, where

$h = O(\log\log n)$ is the height of $\mathcal{T}'$, since even $K_{h-1} = O(\frac{K}{c_{\max}^{\log\log n}}) = O(\frac{K}{\log n})$. Hence all nodes in $\mathcal{T}$ are indeed equipped with optimal $O(K)$-shallow cuttings.

Assigning The Queries To The Equipped Cuttings. To complete the proof it suffices to assign the associated query points to the cells equipped in $\mathcal{T}$ in $O(n \log n)$ time. To do so we solve the *tree point location* problem on $\mathcal{T}$, the rectangular decompositions obtained from the planar projection of the shallow cuttings equipped in every node and all query points.

Problem 1. Given a complete binary tree where every node contains a planar rectangular decomposition and every root-to-leaf path is associated with a different planar query point, point-locate every query point on the decompositions of the left siblings of its associated path.

Lemma 3. *Offline tree point location on n query points and input decompositions of total size N can be solved deterministically in $O(N+n\log N)$ worst-case time and $O(n)$ space.*

Since the nodes of $\mathcal{T}$ are equipped with optimal $O(K)$-shallow cuttings the total size of all decompositions is $N := O(\frac{n}{K}\log n)$. Thus by Lemma 3 the total time to assign all queries is $O(\frac{n}{K}\log n + n\log\left(\frac{n}{K}\log n\right)) = O(n\log n)$ since $K = O(\log^{2+\epsilon} n)$.

Corollary 3. *Given n planar axis-aligned rectangles, the k enclosing pairs of rectangles can be reported deterministically in $O(n\log n+k)$ worst-case time and $O(n)$ space.*

Corollary 4. *Offline d-dimensional dominance reporting on n input points, n query points and k reported points can be solved deterministically in $O(n\log^{d-3} n+ k)$ worst-case time and $O(n)$ space for any constant dimension $d \geq 4$.*

6 Tree Point Location

Preliminaries. An *r-separator* of a graph of n vertices is a subset of $O(\sqrt{rn})$ vertices that divide the remaining vertices into $O(r)$ disjoint subgraphs (connected *components*) of $O(n/r)$ vertices each [14]. Given d-dimensional input points and query hyper-rectangles (*boxes*), the *offline d-dimensional orthogonal range reporting problem* reports the input points that are contained in every query box. The *offline planar rectangular point location problem* is defined as online planar orthogonal point location variant with the difference that all query points are given in advance and that the regions of the input decomposition are rectangles.

Lemma 4. *[?, Theorem 2.2] Offline r-separators on n planar input points can be computed deterministically in $O(n)$ worst-case time and space.*

Lemma 5. *[5, Lemma 4.2] Offline d-dimensional orthogonal range reporting on n input points, m query boxes and k reported points can be solved deterministically in $O(n \log_b^{d-1} n + b^{d-1} m \log_b^{d-1} n + k)$ worst-case time and $O(n)$ space, when parameter $b \geq 2$ and with sorted coordinates.*

Lemma 6. *[5, Lemma 4.1] Offline planar orthogonal point location on n query points and n disjoint orthogonal rectangles can be solved deterministically in $O(n)$ worst-case time and space, when $n \leq 2^{O(\sqrt{w})}$ and with sorted coordinates.*

Proof of Lemma 3. First we compute a *sparse* planar rectangular decomposition of the decomposition on every node of $\mathcal{T}$. In particular we transform the initial decomposition of size N_i in every node at level i into a planar graph in linear time and space and compute a $\left(\frac{N_i}{A}\right)$-*separator* of the graph for a parameter $A := 2^{\sqrt{w}}$, similarly to [6]. This produces a sparse decomposition of size $O(N_i/\sqrt{A})$ that divides the initial decomposition into $O(N_i/A)$ disjoint *components* of size $O(A)$ each that divide the plane to disjoint planar rectangles. By Lemma 4 this takes $O(N)$ time and space in total.

Then we assign the query points associated with the right sibling of every node of $\mathcal{T}$ to a component of its sparse decomposition. Specifically we follow the approach of Chan, Larsen, Pătraşcu [5, Section 4.2] and assign collectively all query points to all respectively associated components in $\mathbb{T}$ by solving a single offline 3D range reporting problem where the query points are considered as input points and the input components are transformed to a set of disjoint query boxes. By Lemma 5 this costs $O(n \log_b^2 N + b^2 (\frac{N}{\sqrt{A}}) \log_b^2 N + n \log N)$ time, which is $O(n \log N)$ when parameter $b := A^{1/2-\alpha}$ for a constant $\alpha > 0$. Finally we point locate every query point within each component. Since the size of each component of size $2^{O(\sqrt{w})}$ the offline point location algorithm of Lemma 6 completes this step in $O(N)$ time. The coordinates are presorted in linear time [5].

Acknowledgements. We thank Prof. Christos Makris from the University of Patras, Greece, for preliminary work on this problem.

References

1. Peyman Afshani. On dominance reporting in 3D. In Dan Halperin and Kurt Mehlhorn, editors, *Algorithms - ESA 2008*, volume 5193 of *Lecture Notes in Computer Science*, pages 41–51. Springer Berlin Heidelberg, 2008.
2. Peyman Afshani and Konstantinos Tsakalidis. Optimal deterministic shallow cuttings for 3D dominance ranges. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1389–1398, 2014.
3. Jon Louis Bentley and Derick Wood. An optimal worst case algorithm for reporting intersections of rectangles. *Computers, IEEE Transactions on*, C-29(7):571–577, July 1980.
4. Timothy M. Chan. Persistent predecessor search and orthogonal point location on the word RAM. *ACM Trans. Algorithms*, 9(3):22:1–22:22, June 2013.

5. Timothy M. Chan, Kasper Green Larsen, and Mihai Pătraşcu. Orthogonal range searching on the RAM, revisited. In *Proceedings of the Twenty-seventh Annual Symposium on Computational Geometry*, SoCG '11, pages 1–10, New York, NY, USA, 2011. ACM.
6. Timothy M. Chan and Mihai Pătraşcu. Transdichotomous results in computational geometry, I: Point location in sublogarithmic time. *SIAM Journal on Computing*, 39(2):703–729, 2009.
7. Bernard Chazelle and LeonidasJ. Guibas. Fractional cascading: I. A data structuring technique. *Algorithmica*, 1(1-4):133–162, 1986.
8. P. Emde Boas, R. Kaas, and E. Zijlstra. Design and implementation of an efficient priority queue. *Mathematical systems theory*, 10(1):99–127, 1976.
9. Michael L. Fredman and Dan E. Willard. Trans-dichotomous algorithms for minimum spanning trees and shortest paths. *J. Comput. Syst. Sci.*, 48(3):533–551, June 1994.
10. Prosenjit Gupta, Ravi Janardan, Michiel Smid, and Bhaskar Dasgupta. The rectangle enclosure and point-dominance problems revisited. *International Journal of Computational Geometry and Applications*, 07(05):437–455, 1997.
11. Yijie Han. Deterministic sorting in O(n log logn) time and linear space. *J. Algorithms*, 50(1):96–105, January 2004.
12. George Lagogiannis, Christos Makris, and Athanasios Tsakalidis. A new algorithm for rectangle enclosure reporting. *Information Processing Letters*, 72(56):177 – 182, 1999.
13. Der-Tsai Lee and Franco P. Preparata. An improved algorithm for the rectangle enclosure problem. *Journal of Algorithms*, 3(3):218 – 224, 1982.
14. R. Lipton and R. Tarjan. Applications of a planar separator theorem. *SIAM Journal on Computing*, 9(3):615–627, 1980.
15. Christos Makris and Konstantinos Tsakalidis. An improved algorithm for static 3D dominance reporting in the pointer machine. In Kun-Mao Chao, Tsan-sheng Hsu, and Der-Tsai Lee, editors, *Algorithms and Computation*, volume 7676 of *Lecture Notes in Computer Science*, pages 568–577. Springer Berlin Heidelberg, 2012.
16. Franco P. Preparata and Michael I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, 1985.
17. Vijay Vaishnavi and Derick Wood. Data structures for the rectangle containment and enclosure problems. *Computer Graphics and Image Processing*, 13(4):372 – 384, 1980.