

Optimal Planar Orthogonal Skyline Counting Queries^{*}

Gerth Stølting Brodal and Kasper Green Larsen

MADALGO^{**}, Department of Computer Science, Aarhus University,
{gerth,larsen}@cs.au.dk

Abstract. The skyline of a set of points in the plane is the subset of maximal points, where a point (x, y) is maximal if no other point (x', y') satisfies $x' \geq x$ and $y' \geq y$. We consider the problem of preprocessing a set P of n points into a space efficient static data structure supporting orthogonal skyline counting queries, i.e. given a query rectangle R to report the size of the skyline of $P \cap R$. We present a data structure for storing n points with integer coordinates having query time $O(\lg n / \lg \lg n)$ and space usage $O(n)$ words. The model of computation is a unit cost RAM with logarithmic word size. We prove that these bounds are the best possible by presenting a matching lower bound in the cell probe model with logarithmic word size: Space usage $n \lg^{O(1)} n$ implies worst case query time $\Omega(\lg n / \lg \lg n)$.

1 Introduction

In this paper we consider orthogonal range skyline queries for a set of points in the plane. A point $(x, y) \in \mathbb{R}^2$ *dominates* a point (x', y') if and only if $x' \leq x$ and $y' \leq y$. For a set of points P , a point $p \in P$ is *maximal* if no other point in P dominates p , and the *skyline* of P , $\text{Skyline}(P)$, is the subset of maximal points in P .

We consider the problem of preprocessing a set P of n points in the plane with integer coordinates into a data structure to support *orthogonal range skyline counting* queries: Given an axis-aligned query rectangle $R = [x_1, x_2] \times [y_1, y_2]$ to report the size of the skyline of the subset of the points from P contained in R , i.e. report $|\text{Skyline}(P \cap R)|$. The main results of this paper are matching upper and lower bounds for data structures supporting such queries, thus completely settling the problem. Our model of computation is the standard unit cost RAM with logarithmic word size.

Previous Work. Orthogonal range searching is one of the most fundamental and well-studied topics in computational geometry, see e.g. [1] for an extensive list of previous results. For orthogonal range queries in the plane, with integer coordinates in $[n] \times [n] = \{0, \dots, n-1\} \times \{0, \dots, n-1\}$, the main results are

^{*} The full version of this paper is available at arxiv.org/abs/1304.7959.

^{**} Center for Massive Data Algorithmics, a Center of the Danish National Research Foundation, grant DNRF84.

Table 1. Previous and new results for skyline counting queries.

Space (words)	Query time	Reference
$n \frac{\lg^2 n}{\lg \lg n}$	$\frac{\lg^{3/2} n}{\lg \lg n}$	[6]
$n \lg n$	$\lg n$	[7]
$n \frac{\lg^3 n}{\lg \lg n}$	$\frac{\lg n}{\lg \lg n}$	[8]
n	$\frac{\lg n}{\lg \lg n}$	New

Table 2. Previous and new results for skyline reporting queries.

Space (words)	Query time	Reference
$n \lg n$	$\lg^2 n + k$	[9] (dynamic)
$n \lg n$	$\lg n + k$	[10, 7]
$n \frac{\lg n}{\lg \lg n}$	$\frac{\lg n}{\lg \lg n} + k$	[6]
$n \lg^\varepsilon n$	$(k + 1) \lg \lg n$	[11]
$n \lg^\varepsilon n$	$\frac{\lg n}{\lg \lg n} + k$	New
$n \lg \lg n$	$(k + 1)(\lg \lg n)^2$	[11]
$n \lg \lg n$	$\frac{\lg n}{\lg \lg n} + k \lg \lg n$	New
n	$(k + 1) \lg^\varepsilon n$	[11]

the following: For the orthogonal range *counting* problem, i.e. queries report the total number of input points inside a query rectangle, optimal $O(\lg n / \lg \lg n)$ query time using $O(n)$ space was achieved in [2]. Optimality was shown in [3], where it was proved that space $n \lg^{O(1)} n$ implies query time $\Omega(\lg n / \lg \lg n)$ for range counting queries.

For range *reporting* queries it is known that space $n \lg^{O(1)} n$ implies query time $\Omega(\lg \lg n + k)$, where k is the number of points reported within the query range [4]. The best upper bounds known for range reporting are: Optimal space $O(n)$ and query time $O((k + 1) \lg^\varepsilon n)$ [1], and optimal query time $O(\lg \lg n + k)$ with space $O(n \lg^\varepsilon n)$ [5]. In both cases $\varepsilon > 0$ is an arbitrarily small constant.

Orthogonal range skyline counting queries were first consider in [6], where a data structure was presented with space usage $O(n \lg^2 n / \lg \lg n)$ and query time $O(\lg^{3/2} n / \lg \lg n)$. This was subsequently improved to $O(n \lg n)$ space and $O(\lg n)$ query time [7]. Finally, a data structure achieving an even faster query time of $O(\lg n / \lg \lg n)$ was presented, however the space usage of that solution was a prohibitive $O(n \lg^3 n / \lg \lg n)$ [8]. Thus to date, no linear space solution exists with a non-trivial query time. Also, from a lower bound perspective, it is not known whether the problem is easier or harder than the standard range counting problem.

For orthogonal skyline reporting queries, the best bound is $O(n \lg n / \lg \lg n)$ space with query time $O(\lg n / \lg \lg n + k)$ [6], where k is the size of the reported skyline. Note that an $\Omega(\lg \lg n)$ search term is needed for skyline range reporting since the $\Omega(\lg \lg n)$ lower bound for standard range reporting was proved even for the case of determining whether the query rectangle is empty [4].

In [11] solutions for the sorted range reporting problem were presented, i.e. the problem of reporting the k leftmost points within a query rectangle in sorted order of increasing x -coordinate. With space $O(n)$, $O(n \lg \lg n)$ and $O(n \lg^\varepsilon n)$, respectively, query times $O((k + 1) \lg^\varepsilon n)$, $O((k + 1)(\lg \lg n)^2)$, and $O(k + \lg \lg n)$ were achieved, respectively. The structures of [11] support finding the rightmost (skyline) point in a query range ($k = 1$). By recursing on the rectangle above the reported point one immediately get the bounds for skyline reporting listed in Table 2, where only the linear space solution achieves query times matching those of general orthogonal range reporting.

Our Results. In Section 3 we present a linear space data structure supporting orthogonal range skyline counting queries in $O(\lg n / \lg \lg n)$ time, thus for the first time achieving linear space and improving over all previous tradeoffs. In Section 2 we show that this is the best possible by proving a matching lower bound. More specifically, we prove a lower bound stating that the query time t must satisfy $t = \Omega(\lg n / \lg(Sw/n))$. Here $S \geq n$ is the space usage in number of words and $w = \Omega(\lg n)$ is the word size in bits. For $w = \lg^{O(1)} n$ and $S = n \lg^{O(1)} n$, this bound becomes $t = \Omega(\lg n / \lg \lg n)$. The lower bound is proved in the cell probe model of Yao [12], which is more powerful than the unit cost RAM and hence the lower bound also applies to RAM data structures.

As a side result, we can also modify our counting data structure to support reporting queries. The details are in the full version of the paper. Our reporting data structure has query time $O(\lg n / \lg \lg n + k)$ and space usage $O(n \lg^\varepsilon n)$. The best previous reporting structure with a linear term in k has $O(\lg n / \lg \lg n + k)$ query time but $O(n \lg n / \lg \lg n)$ space [6]. The reporting structure can also be modified to achieve $O(\lg n / \lg \lg n + k \lg \lg n)$ query time and $O(n \lg \lg n)$ space. See Table 2 for a comparison to previous results.

Our lower bound follows from a reduction of reachability in butterfly graphs to two-sided skyline counting queries, extending reductions by Pătraşcu [13] for two-dimensional rectangle stabbing and range counting queries. Our upper bounds are achieved by constructing a balanced search tree of degree $\Theta(\lg^\varepsilon n)$ over the points sorted by x -coordinate. At each internal node we store several space efficient rank-select data structures storing the points in the subtrees sorted by rank-reduced y -coordinates. Using a constant number of global tables, queries only need to spend $O(1)$ time at each level of the tree.

Preliminaries. If the coordinates of the input and query points are arbitrary integers fitting into a machine word, then we can map the coordinates to the range $[n]$ by using the RAM dictionary from [14], which support predecessor queries on the lexicographical orderings of the points in time $O(\sqrt{\lg n / \lg \lg n})$ using $O(n)$ space. This is less than the $O(\lg n / \lg \lg n)$ query time we are aiming for. Our solution makes extensive use of the below results from succinct data structures.

Lemma 1 ([15]). *A vector $X[1..s]$ of s zero-one values, with t values equal to one, can be stored in a data structure of size $O(t(1 + \lg s/t))$ bits supporting rank and select in $O(1)$ time, where $\text{rank}(i)$ returns the number of ones in $X[1..i]$, provided $X[i] = 1$, and $\text{select}(i)$ returns the position of the i 'th one in X .*

Lemma 2 ([16]). *Let $X[1..s]$ be a vector of s non-negative integers with total sum t . There exists a data structure of size $O(s \lg(2 + t/s))$ bits, supporting the lookup of $X[i]$ and the prefix sum $\sum_{j=1}^i X[j]$ in $O(1)$ time, for $i = 1, \dots, s$.*

Lemma 3 ([17, 18]). *Let $X[1..s]$ be a vector of integers. There exists a data structure of size $O(s)$ bits supporting range-maximum-queries in $O(1)$ time, i.e. given i and j , $1 \leq i \leq j \leq s$, reports the index k , $i \leq k \leq j$, such that $X[k] = \max(X[i..j])$. Queries only access this data structure, i.e. the vector X is not stored.*

2 Lower Bound

That an orthogonal range skyline counting data structure requires $\Omega(n \lg n)$ bits space, follows immediately since each of the $n!$ different input point sets of size n , where points have distinct x - and y -coordinates from $[n]$, can be reconstructed using query rectangles considering each possible point in $[n]^2$ independently, i.e. the space usage is at least $\lceil \lg_2(n!) \rceil = \Omega(n \lg n)$ bits.

In the remainder of this section, we prove that any data structure using $S \geq n$ words of space must have query time $t = \Omega(\lg n / \lg(Sw/n))$, where $w = \Omega(\lg n)$ denotes the word size in bits. In particular for $w = \lg^{O(1)} n$, this implies that any data structure using $n \lg^{O(1)} n$ space must have query time $t = \Omega(\lg n / \lg \lg n)$, showing that our data structure from Section 3 is optimal. Our lower bound holds even for data structures only supporting skyline counting queries inside 2-sided rectangles, i.e. query rectangles of the form $(-\infty, x] \times (-\infty, y]$. The lower bound is proved in the cell probe model of Yao [12] with word size $w = \Omega(\lg n)$. Since we derive our lower bound by reduction, we will not spend time on introducing the cell probe model, but merely note that lower bounds proved in this model applies to data structures developed in the unit cost RAM model. See e.g. [3] for a brief description of the cell probe model.

Reachability in the Butterfly Graph. We prove our lower bound by reduction from the problem known as *reachability oracles in the butterfly graph* [13]. A butterfly graph of degree B and depth d is a directed graph with $d + 1$ layers, each having B^d nodes ordered from left to right. The nodes at level 0 are the *sources* and the nodes at level d are the *sinks*. Each node, except the sinks, has out-degree B , and each node, except the sources, has in-degree B .

If we number the nodes at each level with $0, \dots, B^d - 1$ from left to right and interpret each index $i \in [B^d]$ as a vector $v(i) = v(i)[d-1] \dots v(i)[0] \in [B]^d$ (just write i in base B), then the node at index i at layer $k \in [d]$ has an out-going edge to each node j at layer $k+1$ for which $v(j)$ and $v(i)$ differ only in the k 'th coordinate. Here the 0'th coordinate is the coordinate corresponding to the least significant digit when thinking of $v(i)$ and $v(j)$ as numbers written in base B . Observe that there is precisely one directed path between each source-sink pair. For the s 'th source and the t 'th sink, this path corresponds to “morphing” one digit of $v(s)$ into the corresponding digit in $v(t)$ for each layer traversed in the butterfly graph.

The input to the problem of reachability oracles in the butterfly graph, with degree B and depth d , is a subset of the edges of the butterfly graph, i.e. we are given a subgraph G of the butterfly as input. A query is specified by a source-sink pair (s, t) and the goal is to return whether there exists a directed path from the given source s to the given sink t in G . Pătraşcu proved the following:

Theorem 1 (Pătraşcu [13], Section 5). *Any cell probe data structure answering reachability queries in subgraphs of the butterfly graph with degree B and depth d , having space usage S words of w bits, must have query time $t = \Omega(d)$, provided $B = \Omega(w^2)$ and $\lg B = \Omega(\lg Sd/N)$. Here N denotes the number of non-sink nodes in the butterfly graph.*

We derive our lower bound by showing that any cell probe data structure for skyline range counting can be used to answer reachability queries in subgraphs of the butterfly graph for any degree B and depth d .

Edges to 2-d Rectangles. Consider the butterfly graph with degree B and depth d . The first step of our reduction is inspired by the reduction Pătraşcu used to obtain a lower bound for 2-d rectangle stabbing: Consider an edge of the butterfly graph, leaving the i 'th node at layer $k \in [d]$ and entering the j 'th node in layer $k + 1$. We denote this edge $e_k(i, j)$. The source-sink pairs (s, t) that are connected through $e_k(i, j)$ are those for which:

1. The source has an index s satisfying $v(s)[h] = v(i)[h]$ for $h \geq k$, i.e. s and i agree on the $d - k$ most significant digits when written in base B .
2. The sink has an index t satisfying $v(t)[h] = v(j)[h]$ for $h \leq k + 1$, i.e. t and j agree on the $k + 1$ least significant digits when written in base B .

We now map each edge $e_k(i, j)$ of the butterfly graph to a rectangle in 2-d. For the edge $e_k(i, j)$, we create the rectangle $r_k(i, j) = [x_1, x_2] \times [y_1, y_2]$ where:

- $x_1 = v(i)[d - 1]v(i)[d - 2] \cdots v(i)[k]0 \cdots 0$ in base B ,
- $x_2 = v(i)[d - 1]v(i)[d - 2] \cdots v(i)[k](B - 1) \cdots (B - 1)$ in base B ,
- $y_1 = v(j)[0]v(j)[1] \cdots v(j)[k + 1]0 \cdots 0$ in base B , and
- $y_2 = v(j)[0]v(j)[1] \cdots v(j)[k + 1](B - 1) \cdots (B - 1)$ in base B .

The crucial observation is that for a source-sink pair, where the source is the s 'th source and the sink is the t 'th sink, the edges on the path from the source to the sink in the butterfly graph are precisely those edges $e_k(i, j)$ for which the corresponding rectangle $r_k(i, j)$ contains the point $(s, \text{rev}_B(t))$, where $\text{rev}_B(t)$ is the number obtained by writing t in base B and then reversing the digits.

We now collect the set of rectangles R , containing each rectangle $r_k(i, j)$ corresponding to an edge of the butterfly graph. Given an input subgraph G , we *mark* all rectangles $r_k(i, j) \in R$ for which the corresponding edge $e_k(i, j)$ is also in G . It follows that there is a directed path from the s 'th source to the t 'th sink in the subgraph G if and only if $(s, \text{rev}_B(t))$ is not contained in any *unmarked* rectangle in R .

Our goal is now to transform marked and unmarked rectangles to points, such that we can use a skyline counting data structure to determine whether a given point $(s, \text{rev}_B(t))$ is contained in an unmarked rectangle. Note that our reduction only works for the rectangle set R obtained from the butterfly graph, and not for any set of rectangles, i.e. we could not have reduced from the general problem of 2-d rectangle stabbing.

2-d Rectangles to Points. To avoid tedious details, we from this point on allow the input to skyline queries to have multiple points with the same x - or y -coordinate (though not two points with both coordinates identical). This assumption can easily be removed, but it would only distract the reader from the main ideas of our reduction. We still use the definition that a point (x, y) dominates a point (x', y') if and only if $x' \leq x$ and $y' \leq y$.

The next step of the reduction is to map the rectangles R to a set of points. For this, we first transform the coordinates slightly: For every rectangle $r_k(i, j) \in R$, having coordinates $[x_1, x_2] \times [y_1, y_2]$, we modify each of the coordinates in the following way: $x_1 \leftarrow dx_1 + (d - 1 - k)$, $x_2 \leftarrow dx_2 + d - 1$, $y_1 \leftarrow dy_1 + k$, and $y_2 \leftarrow dy_2 + d - 1$. The multiplication with d essentially corresponds to expanding each point with integer coordinates to a $d \times d$ grid of points. The purpose of adding k to y_1 and $(d - 1 - k)$ to x_1 is to ensure that, if two rectangles share a lower-left corner (only possible for two rectangles $r_k(i, j)$ and $r_{k'}(i', j')$ where $k \neq k'$), then those corners do not dominate each other in the transformed set of rectangles. We will see later that the particular placement of the points based on k also plays a key role. We use $\pi : [B^d]^4 \rightarrow [dB^d]^4$ to denote the above map. With this notation, the transformed set of rectangles is denoted $\pi(R)$ and each rectangle $r_k(i, j) \in R$ is mapped to $\pi(r_k(i, j)) \in \pi(R)$.

We now create the set of points P' containing the set of lower-left corner points for all rectangles $\pi(r_k(i, j)) \in \pi(R)$, i.e. for each $\pi(r_k(i, j)) = [x_1, x_2] \times [y_1, y_2]$, we add the point (x_1, y_1) to P' . The set P' has the following crucial property:

Lemma 4. *Let (x, y) be a point with coordinates in $[B^d] \times [B^d]$. Then for the two-sided query rectangle $Q = (-\infty, dx + d - 1] \times (-\infty, dy + d - 1]$, it holds that $\text{Skyline}(Q \cap P')$ contains precisely the points in P' corresponding to the lower-left corners of the rectangles $\pi(r_k(i, j)) \in \pi(R)$ for which $r_k(i, j)$ contains (x, y) .*

Proof. First let $p = (x_1, y_1) \in P'$ be the lower-left corner of a rectangle $\pi(r_k(i, j))$ such that $r_k(i, j)$ contains the point (x, y) . We want to show that $p \in \text{Skyline}(Q \cap P')$. Since $r_k(i, j)$ contains the point (x, y) , we have $x \geq \lfloor x_1/d \rfloor$ and $y \geq \lfloor y_1/d \rfloor$. From this, we get $dx + d - 1 \geq d\lfloor x_1/d \rfloor + (d - 1 - k) = x_1$ and $dy + d - 1 \geq d\lfloor y_1/d \rfloor + k = y_1$, i.e. p is inside Q . Since (x, y) is inside $r_k(i, j)$, we also have that $(dx + d - 1, dy + d - 1)$ is dominated by the upper-right corner of $\pi(r_k(i, j))$, i.e. $(dx + d - 1, dy + d - 1)$ is inside $\pi(r_k(i, j))$.

What remains to be shown is that no other point in $Q \cap P'$ dominates p . For this, assume for contradiction that some point $p' = (x'_1, y'_1) \in P'$ is both in Q and also dominates p . First, since p' is dominated by $(dx + d - 1, dy + d - 1)$ and also dominates p , we know that p' must be inside $\pi(r_k(i, j))$. Now let $\pi(r_{k'}(i', j')) \neq \pi(r_k(i, j))$ be the rectangle in $\pi(R)$ from which p' was generated, i.e. p' is the lower-left corner of $\pi(r_{k'}(i', j'))$. We have three cases:

1. First, if $k' = k$ we immediately get a contradiction since the rectangles $\pi(R)_k = \{\pi(r_{k'}(i', j')) \in \pi(R) \mid k' = k\}$ are pairwise disjoint and hence p' could not have been inside $\pi(r_k(i, j))$.
2. If $k' < k$, we know that $\pi(r_{k'}(i', j'))$ is shorter in x -direction and longer in y -direction than $\pi(r_k(i, j))$. From our transformation, we know that $(y_1 \bmod d) = k$ and $(y'_1 \bmod d) = k' < k$. Thus since p' dominates p , we must have $\lfloor y'_1/d \rfloor > \lfloor y_1/d \rfloor$. But these two values are precisely the y -coordinates of the lower-left corners of $r_k(i, j)$ and $r_{k'}(i', j')$. By definition, we get:

$$v(j')[0]v(j')[1] \cdots v(j')[k' + 1]0 \cdots 0 > v(j)[0]v(j)[1] \cdots v(j)[k + 1]0 \cdots 0.$$

Since $k' < k$, this furthermore gives us

$$v(j')[0]v(j')[1] \cdots v(j')[k'+1] > v(j)[0]v(j)[1] \cdots v(j)[k'+1].$$

From this it follows that

$$v(j')[0] \cdots v(j')[k'+1]0 \cdots 0 > v(j)[0] \cdots v(j)[k+1](B-1) \cdots (B-1),$$

i.e. the lower-left corner of $r_{k'}(i', j')$ is outside $r_k(i, j)$, which also implies that the lower-left corner of $\pi(r_{k'}(i', j'))$ is outside $\pi(r_k(i, j))$. That is, p' is outside $\pi(r_k(i, j))$, which gives the contradiction.

3. The case for $k' > k$ is symmetric to the case $k' < k$, just using the x -coordinates instead of the y -coordinates to derive the contradiction.

The last step of the proof is to show that no point $p = (x_1, y_1) \in P'$ can be in $\text{Skyline}(Q \cap P')$ but at the same time correspond to the lower-left corner of a rectangle $\pi(r_k(i, j))$ where $r_k(i, j)$ does not contain the point (x, y) . First observe that $(dx + d - 1, dy + d - 1)$ is contained in precisely one rectangle $\pi(r_{k'}(i', j'))$ for each value of $k' \in [d]$. Now let $\pi(r_k(i', j')) \neq \pi(r_k(i, j))$ be the rectangle containing $(dx + d - 1, dy + d - 1)$ amongst the rectangles $\pi(R)_k$. The lower-left corner of this rectangle is dominated by $(dx + d - 1, dy + d - 1)$ but also dominates p , hence p is not in $\text{Skyline}(Q \cap P')$. \square

Handling Marked and Unmarked Rectangles. The above steps are all independent of the concrete input subgraph G . As discussed, we need a way to determine whether a query point is contained in an unmarked rectangle or not. This step is now very simple in light of Lemma 4: First, multiply all coordinates of points in P' by 2. This corresponds to expanding each point with integer coordinates into a 2×2 grid. Now for every point $p \in P'$, if the rectangle $\pi(r_k(i, j))$ from which p was generated is marked, then we add 1 to both the x - and y -coordinate of p , i.e. we move p to the upper-right corner of the 2×2 grid in which it is placed. If $\pi(r_k(i, j))$ is unmarked, we replace it by two points, one where we add 1 to the x -coordinate, and one where we add 1 to the y -coordinate. We denote the resulting set of points $P(G)$. It follows immediately that:

Corollary 1. *Let G be a subgraph of the butterfly graph with degree B and depth d . Also, let (x, y) be a point with coordinates in $[B^d] \times [B^d]$. Then for the two-sided query rectangle $Q = (-\infty, 2d(x+1)-1] \times (-\infty, 2d(y+1)-1]$, it holds that $\text{Skyline}(Q \cap P(G))$ contains precisely one point from $P(G)$ for every marked rectangle in R that contains (x, y) , two points from $P(G)$ for every unmarked rectangle in R that contains (x, y) , and no other points, i.e. $|\text{Skyline}(Q \cap P(G))| - d$ equals the number of unmarked rectangles in R which contains (x, y) .*

Corollary 2. *Let G be a subgraph of the butterfly graph with degree B and depth d . Let s be the index of a source and t the index of a sink. Then the s 'th source can reach the t 'th sink in G if and only if $|\text{Skyline}(Q \cap P(G))| = d$ for the two-sided query rectangle $Q = (-\infty, 2d(s+1)-1] \times (-\infty, 2d(\text{rev}_B(t)+1)-1]$.*

Deriving the Lower Bound. The lower bound can be derived from Corollary 2 and Theorem 1 as follows. First note that the set R contains NB rectangles, since each rectangle corresponds to an edge of the butterfly graph and each of the N non-sink nodes of the butterfly graph has B outgoing edges. Each of these rectangles gives one or two points in $P(G)$. Letting n denote $|P(G)|$, we have $NB \leq n \leq 2NB$. From $N = d \cdot B^d \leq n$ we get $d \leq \lg n$ and $d = \Theta(\lg_B N)$.

Given n , $w \geq \lg n$, and $S \geq n$, we now derive a lower bound on the query time. Setting $B = \frac{S}{n}w^2$ we have $B = \Omega(w^2)$ and $\lg B = \Omega(\lg \frac{Sd}{N})$ (as required by Theorem 1), where the last bound follows from $\lg \frac{Sd}{N} \leq \lg \frac{S \cdot \lg n}{n/2B} \leq \lg(2B \frac{S \cdot w}{n}) \leq \lg(2B^2) = O(\lg B)$. Furthermore we have $\lg \frac{Sw}{n} = \frac{1}{2} \lg(\frac{Sw}{n})^2 \geq \frac{1}{2} \lg(\frac{S}{n}w^2) = \frac{1}{2} \lg B$. From Theorem 1 we can now bound the time for a skyline counting query by $t = \Omega(d) = \Omega(\lg_B N) = \Omega(\lg n / \lg B) = \Omega(\lg n / \lg(Sw/n))$.

3 Skyline Counting Data Structure

In this section we describe a data structure using $O(n)$ space supporting orthogonal skyline counting queries in $O(\lg n / \lg \lg n)$ time. We describe the basic idea of how to support queries, and present the details of the stored data structure. The details of the query can be found in the full version of the paper.

The basic idea is to store the n points in left-to-right x -order at the leaves of a balanced tree T of degree $\Theta(\log^\varepsilon n)$, i.e. height $O(\log n / \log \log n)$, and for each internal node v have a list L_v of the points in the subtree rooted in v in sorted y -order. The *slab* of v is the narrowest infinite vertical band containing L_v . To obtain the overall linear space bound, L_v will not be stored explicitly but implicitly and rank-reduced using rank-select data structures, where navigation is performed using fractional cascading on rank-select data structures (details below). A 4-sided query R decomposes into 2-sided subqueries at $O(\log n / \log \log n)$ nodes (in Figure 1, R is decomposed into subqueries R_1 - R_5 , white points are nodes on the skyline within R , double circled points are the top-most points within each R_i). For skyline queries (both counting and reporting) it is important to consider the subqueries right-to-left, and the lower y -value for the subquery in R_i is raised to the maximal y -value of a point in the subqueries to the right. Since the tree T has non-constant degree, we need space efficient solutions for multislab queries at each node v . We partition L_v into blocks of size $O(\log^{2\varepsilon} n)$, and a query R_i decomposes into five subqueries (1-5), see Figure 2: (1) and (3) are on small subsets of points within a single block and can be answered by tabulation (given the *signature* of the block); (2) is a block aligned multislab query; (4) and (5) are for single slabs (at the children of v). For (2,4,5) the skyline size between points i and j (numbered bottom-up) can be computed as one plus the difference between the size of the skyline from 1 to j and 1 to k , where k is the rightmost point between i and j (see Figure 3, white and black circles and crosses are all points, crosses indicate the skyline from i to j , white circles from 1 to k , and white circles together with crosses from 1 to j). Finally, the skyline size from 1 to i can be computed from a prefix sum, if we for point i store the number of points in the skyline from 1 to $i - 1$ dominated by i (see

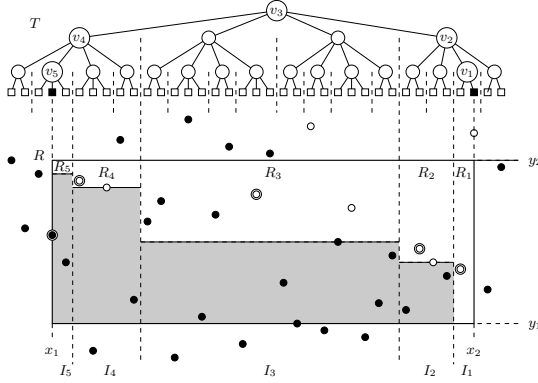


Fig. 1. The base tree T with $\Delta = 4$.

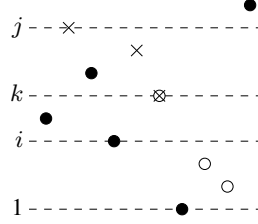


Fig. 3. $\text{SkyCount}_v(i, j)$.

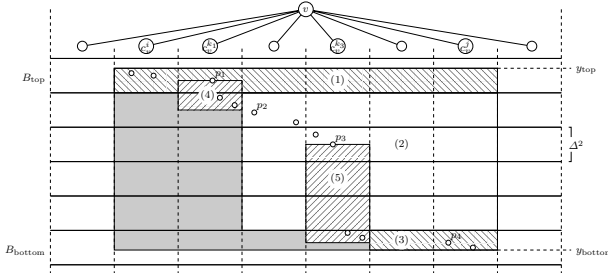


Fig. 2. Skyline queries for multislabs.

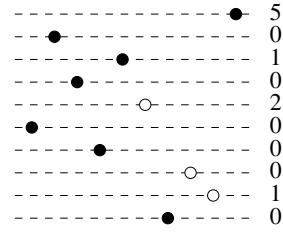


Fig. 4. $|\text{Skyline}(L_v[1..i])|$.

Figure 4, the skyline between 1 and 6 consists of the three white nodes, and the size is $6 - (2 + 0 + 0 + 0 + 1 + 0) = 3$.

The details of the construction are as follows. We let $\Delta = \max\{2, \lceil \lg^\varepsilon n \rceil\}$ be a parameter of our construction, where $0 < \varepsilon < 1/3$ is a constant. We build a balanced *base tree* T over the set of points P , where the leaves from left-to-right store the points in P in sorted order w.r.t. x -coordinate. Each internal node of T has degree at most Δ and T has height $\lceil \lg_\Delta n \rceil + 1$. (See Figure 1)

For each internal node v of T we store a set of data structures. Before describing these we need to introduce some notation. The subtree of T rooted at a node v is denoted T_v , and the set of points stored at the leaves of T_v is denoted P_v . We let $n_v = |P_v|$ and $L_v[1..n_v]$ be the list of the points in P_v sorted in increasing y -order. We let $I_v = [\ell_v, r_v]$ denote the x -interval defined by the x -coordinates of the points stored at the leaves of T_v , and denote $I_v \times [n]$ the *slab* spanned by v . The degree of v is denoted d_v , the children of v are from left-to-right denoted $c_v^1, \dots, c_v^{d_v}$, and the parent of node v is denoted p_v . A list L_v is partitioned into a sequence of blocks $B_v[1.. \lceil n_v/\Delta^2 \rceil]$ of size Δ^2 , such that $B_v[i] = L_v[(i-1)\Delta^2 + 1.. \min\{n_v, i\Delta^2\}]$. The *signature* $\sigma_v[i]$ of a block $B_v[i]$ is a list of pairs: For each point p from $B_v[i]$ in increasing y -order we construct a pair (j, r) , where j is the index of the child c_v^j of v storing p and r is the rank of p 's x -coordinate among all points in $B_v[i]$ stored at the

same child c_v^j as p . The total number of bits required for a signature is at most $\Delta^2(\lg \Delta + \lg \Delta^2) = O(\lg^{2\varepsilon} n \cdot \lg \lg n)$.

To achieve overall $O(n)$ space we need to encode succinctly sufficient information for performing queries. In particular we will *not* store the points in L_v explicitly at the node v , but only partial information about the points relative position will be stored.

Queries on a block $B_v[i]$ are handled using table lookups in global tables using the block signature $\sigma_v[i]$. We have tables for the below block queries, where we assume σ is the signature of a block storing points p_1, \dots, p_{Δ^2} distributed in Δ child slabs.

- Below(σ, t, i) Returns the number of points from p_1, \dots, p_t contained in slab i .
- Rightmost(σ, b, t, i, j) Returns k , where p_k is the rightmost point among p_b, \dots, p_t contained in slabs $[i, j]$. If no such point exists, -1 is returned.
- Topmost(σ, b, t, i, j) Returns k , where p_k is the topmost point among p_b, \dots, p_t contained in slabs $[i, j]$. If no such point exists, -1 is returned.
- SkyCount(σ, b, t, i, j) Returns the size of the skyline for the subset of the points p_b, \dots, p_t contained in slabs $[i, j]$.

The arguments to each of the above lookups consists of at most $|\sigma| + 2 \lg \Delta^2 + 2 \lg \Delta = |\sigma| + O(\lg \lg n) = O(\lg^{2\varepsilon} n \cdot \lg \lg n)$ bits and the answer is $\lg(\Delta + 1) = O(\lg \lg n)$ bits, i.e. each query can be answered in $O(1)$ time using a table of size $O(2^{\lg^{2\varepsilon} n \cdot \lg \lg n} \cdot \lg \lg n) = o(n)$ bits, since $\varepsilon < 1/3$.

For each internal node v of T we store the following data structures, each having $O(1)$ access time.

- $C_v(i)$ Compact array that for each i , where $1 \leq i \leq n_v$, stores the index of the child of v storing $L_v[i]$, i.e. $1 \leq C_v(i) \leq \Delta$. Space usage $O(n_v \lg \Delta)$ bits.
- $\pi_v(i)$ For each i , $1 \leq i \leq n_v$, stores the index of $L_v[i]$ in L_{p_v} , i.e. $L_{p_v}[\pi_v(i)] = L_v[i]$. This can be supported by constructing the select data structure of Lemma 1 on the bit-vector X , where $X[i] = 1$ if and only if $L_{p_v}[i]$ is in L_v . A query to $\pi_v(i)$ simply becomes a select(i) query. Space usage $O(n_v \lg(n_{p_v}/n_v)) = O(n_v \lg \Delta)$ bits.
- $\sigma_v(i)$ Array of signatures for the blocks $B_v[1..[n_v/\Delta^2]]$. Space usage $O(n_v/\Delta^2 \cdot \Delta^2 \cdot \lg \Delta) = O(n_v \lg \Delta)$ bits.
- Pred $_v(t, i)$ / Succ $_v(t, i)$ Supports finding the predecessor/successor of $L_v[t]$ in the i 'th child list $L_{c_v^i}$. Returns $\max\{k \mid 1 \leq k \leq n_{c_v^i} \wedge \pi_{c_v^i}[k] \leq t\}$ and $\min\{k \mid 1 \leq k \leq n_{c_v^i} \wedge \pi_{c_v^i}[k] \geq t\}$, respectively. For each child index i , we construct an array X^i of size $\lceil n/\Delta^2 \rceil$, such that $X^i[b]$ is the number of points in block $B_v[b]$ that are stored in the i 'th child slab. The prefix sums of each X^i are stored using the data structure of Lemma 2 using space $O((n_v/\Delta^2) \lg(\Delta^2))$ bits. The total space for all Δ children of v becomes $O(\Delta \cdot n_v/\Delta^2 \cdot \lg \Delta) = O(n_v)$ bits. The result of a Pred $_v(t, i)$ query is $\sum_{j=1}^{\lceil t/\Delta^2 \rceil - 1} X^i[j] + \text{Below}(\sigma_v(\lceil t/\Delta^2 \rceil), 1 + (t - 1 \bmod \Delta^2), i)$, where the first term can be computed in $O(1)$ time by Lemma 2 and the second term is a constant time global table lookup. The result of Succ $_v(t, i) = \text{Pred}_v(t, i)$ if $C_v[t] = i$, otherwise Succ $_v(t, i) = \text{Pred}_v(t, i) + 1$.

- $\text{Rightmost}_v(i, j)$ Returns the index k , where $i \leq k \leq j$, such that $L_v[k]$ has the maximum x -value among $L_v[i..j]$. Using Lemma 3 on the array of the x -coordinates of the points in L_v we achieve $O(1)$ time queries and space usage $O(n_v)$ bits.
- $\text{SkyCount}_v(i)$ Returns $|\text{Skyline}(L_v[1..i])|$. Construct an array X , where $X[i]$ is the number of points in $\text{Skyline}(L_v[1..i-1])$ dominated by $L_v[i]$. (See Figure 4) We can now compute $|\text{Skyline}(L_v[1..i])|$ as $i - \sum_{j=1}^i X[j]$. Using Lemma 2 the query time becomes $O(1)$ and the space usage $O(n_v)$ bits, since $\sum_{j=1}^{n_v} X[j] \leq n_v - 1$.
- $\text{SkyCount}_v(i, j)$ Returns $|\text{Skyline}(L_v[i..j])|$, computable by the following expression: $\text{SkyCount}_v(j) - \text{SkyCount}_v(\text{Rightmost}_v(i, j)) + 1$ (see Figure 3).

Finally, we store for each node v and slab interval $[i, j]$ the following data structures.

- $\text{Rightmost}_{v,i,j}(b, t)$ Returns k , where $L_v[k]$ is the rightmost point among the points in blocks $B_v[b..t]$ contained in slabs $[i, j]$. If no such point exists, -1 is returned. Can be solved by applying Lemma 3 to the array X , where $X[s]$ is the x -coordinate of the rightmost point in $B_v[s]$ contained in slabs $[i, j]$. A query first finds the block ℓ containing the rightmost point using this data structure, and then returns $(\ell - 1)\Delta^2 + \text{Rightmost}(\sigma_v[\ell], 1, \Delta^2, i, j)$. Space usage $O(n_v/\Delta^2)$ bits.
- $\text{Topmost}_{v,i,j}(b, t)$ Returns k , where $L_v[k]$ is the topmost point among the points in blocks $B_v[b..t]$ contained in slabs $[i, j]$. If no such point exists, -1 is returned. Can be solved by first using Lemma 3 on the array X , where $X[s] = s$ if there exists a point in $B_v[s]$ contained in slabs $[i, j]$. Otherwise $X[s] = 0$. Let ℓ be the block found using Lemma 3. Return the result of $(\ell - 1)\Delta^2 + \text{Topmost}(\sigma_v[\ell], 1, \Delta^2, i, j)$. Space usage $O(n_v/\Delta^2)$ bits.
- $\text{SkyCount}_{v,i,j}(b, t)$ Returns the size of the skyline for the subset of points in blocks $B_v[b..t]$ contained in slabs $[i, j]$. Can be supported by two applications of Lemma 2 on two arrays X and Y : Let $X[s] = \text{SkyCount}(\sigma_v[s], 1, \Delta^2, i, j)$, i.e. the size of the skyline of the points in block $B_v[s]$ contained in slabs $[i, j]$. Let $B_{v,i,j}[s]$ denote the points in $B_v[s]$ contained in slabs $[i, j]$. Let $Y[s] = |\text{Skyline}(B_{v,i,j}[1..s-1]) \setminus \text{Skyline}(B_{v,i,j}[1..s])|$, i.e. the number of points on $\text{Skyline}(B_{v,i,j}[1..s-1])$ dominated by points in $B_{v,i,j}[s]$. Space usage for X and Y is $O(n_v/\Delta^2 \cdot \lg \Delta^2)$ bits. We can compute $\text{SkyCount}_{v,i,j}(b, t) = \sum_{s=k}^t X[s] - \sum_{s=k+1}^t Y[s]$, where $k = \lceil \text{Rightmost}_{v,i,j}(b, t)/\Delta^2 \rceil$.

The total space of our data structure, in addition to the $o(n)$ bits for our global tables, can be bounded as follows. The total space for all $O(\Delta^2)$ multislab data structures for a node v is $O(\Delta^2 \cdot n_v/\Delta^2 \cdot \lg \Delta)$ bits. The total space for all data structures at a node v becomes $O(n_v \lg \Delta)$ bits. Since the sum of all n_v for a level of T is at most n , the total space for all nodes at a level of T is $O(n \lg \Delta)$ bits. Since T has height $O(\lg_\Delta n)$, the total space usage becomes $O(n \lg \Delta \cdot \lg_\Delta n) = O(n \lg n)$ bits, i.e. $O(n)$ words. The data structure can be constructed bottom-up in $O(n \log n)$ time.

References

1. Chan, T.M., Larsen, K.G., Pătraşcu, M.: Orthogonal range searching on the RAM, revisited. In: 27th ACM Symposium on Computational Geometry, ACM (2011) 1–10
2. JáJá, J., Mortensen, C.W., Shi, Q.: Space-efficient and fast algorithms for multidimensional dominance reporting and counting. In: 15th International Symposium Algorithms and Computation. Volume 3341 of LNCS., Springer (2004) 558–568
3. Pătraşcu, M.: Lower bounds for 2-dimensional range counting. In: 39th Annual ACM Symposium on Theory of Computing, ACM (2007) 40–46
4. Pătraşcu, M., Thorup, M.: Time-space trade-offs for predecessor search. In: 38th Annual ACM Symposium on Theory of Computing, ACM (2006) 232–240
5. Alstrup, S., Brodal, G.S., Rauhe, T.: New data structures for orthogonal range searching. In: 41st Annual Symposium on Foundations of Computer Science, IEEE Computer Society (2000) 198–207
6. Das, A.S., Gupta, P., Kalavagattu, A.K., Agarwal, J., Srinathan, K., Kothapalli, K.: Range aggregate maximal points in the plane. In: 6th International Workshop Algorithms and Computation. Volume 7157 of LNCS., Springer (2012) 52–63
7. Kalavagattu, A.K., Agarwal, J., Das, A.S., Kothapalli, K.: Counting range maxima points in plane. In: International Workshop On Combinatorial Algorithms. Volume 7643 of LNCS., Springer (2012) 263–273
8. Das, A.S., Gupta, P., Srinathan, K.: Counting maximal points in a query orthogonal rectangle. In: 7th International Workshop on Algorithms and Computation. Volume 7748 of LNCS., Springer (2013) 65–76
9. Brodal, G.S., Tsakalidis, K.: Dynamic planar range maxima queries. In: 38th International Colloquium on Automata, Languages, and Programming. Volume 6755 of LNCS. Springer Verlag, Berlin (2011) 256–267
10. Kalavagattu, A.K., Das, A.S., Kothapalli, K., Srinathan, K.: On finding skyline points for range queries in plane. In: 23rd Annual Canadian Conference on Computational Geometry. (2011)
11. Nekrich, Y., Navarro, G.: Sorted range reporting. In: 13th Scandinavian Symposium and Workshops Algorithm Theory. Volume 7357 of LNCS., Springer (2012) 271–282
12. Yao, A.C.C.: Should tables be sorted? *Journal of the ACM* **28**(3) (1981) 615–628
13. Pătraşcu, M.: Unifying the landscape of cell-probe lower bounds. *SIAM Journal on Computing* **40**(3) (2011) 827–847
14. Beame, P., Fich, F.E.: Optimal bounds for the predecessor problem. In: 31st Annual ACM Symposium on Theory of Computing, ACM (1999) 295–304
15. Raman, R., Raman, V., Rao, S.S.: Succinct indexable dictionaries with applications to encoding k -ary trees and multisets. In: 13th Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM (2002) 233–242
16. Raman, R., Raman, V., Satti, S.R.: Succinct indexable dictionaries with applications to encoding k -ary trees, prefix sums and multisets. *ACM Transactions on Algorithms* **3**(4) (2007)
17. Sadakane, K.: Succinct data structures for flexible text retrieval systems. *Journal of Discrete Algorithms* **5**(1) (2007) 12–22
18. Fischer, J.: Optimal succinctness for range minimum queries. In: 9th Latin American Symposium on Theoretical Informatics. Volume 6034 of LNCS., Springer (2010) 158–169

The following appendix is not contained in the version of the paper appearing in the proceedings of the 14th Scandinavian Symposium and Workshops on Algorithm Theory.

A Skyline Range Counting Queries

To answer a skyline counting query $R = [x_1, x_2] \times [y_1, y_2]$, we identify the nodes on the paths in T from the two leaves storing x_1 and x_2 up to the lowest common ancestor of the two leaves. Let v_1, \dots, v_m be the set of these nodes in a right-to-left traversal in T (see Figure 1). The horizontal span of the query, $[x_1, x_2]$, is the concatenation of the span of at most one multislab I_1, \dots, I_m from each of v_1, \dots, v_m . For each such multislab I_ℓ we form a new subquery $R_\ell = I_\ell \times [z_\ell, y_2]$, completely spanning the multislab in the horizontal direction and vertically has a range $[z_\ell, y_2]$, where $z_1 = y_1$ and $z_\ell = \max\{z_{\ell-1}, y_\ell^{\max} + 1\}$, for $\ell = 2$ to m and y_ℓ^{\max} is the maximal y -coordinate of a point in $I_{\ell-1} \times [1, y_2]$. By definition of the R_ℓ queries, the skyline of the points contained within R is exactly the union of the skylines for each of the R_ℓ subqueries (see Figure 1), since the points in R_ℓ cannot be dominated by other points that are both in R and to the right of I_ℓ .

To navigate in T we need to find the index of the successor of y_1 and the predecessor y_2 in each of the L_{v_ℓ} lists. We start with y_1 and y_2 being the indexes at the root, and then use the $\text{Succ}_v/\text{Pred}_v$ structures at the nodes to find the successor of y_1 and predecessor of y_2 at all the nodes on the two paths from the root to x_1 and x_2 . To find the topmost point below y_2 in a multislab we use $\text{Topmost}_{v,i,j}$. To navigate y^{\max} values up and down between the levels of T we use $\pi_v(y^{\max})$ to move upwards and $\text{Succ}_v(y^{\max}, j)$ to move downwards to a slab j . These navigations can be performed in $O(1)$ time per node on the paths, i.e. total time $O(\lg_\Delta n)$.

What remains is to compute in $O(1)$ time the size the skyline within a query range R_ℓ . In the following we consider a query range that horizontally spans the child slabs $[i, j]$ of a node v , and vertically spans the indexes $[y_{\text{bottom}}, y_{\text{top}}]$ in L_v .

If the query range is within a single block of L_v (i.e. $\lceil y_{\text{bottom}}/\Delta^2 \rceil = \lceil y_{\text{top}}/\Delta^2 \rceil$), we compute the skyline size as

$$\text{SkyCount}(\sigma_v(\lceil y_{\text{top}}/\Delta^2 \rceil), 1 + (y_{\text{bottom}} - 1 \bmod \Delta^2), 1 + (y_{\text{top}} - 1 \bmod \Delta^2), i, j) .$$

Otherwise we decompose the skyline counting query into five subranges (1)-(5), see Figure 2. We first compute the y -coordinate of the rightmost point p_1 in the top block B_{top} of the query range using

$$p_1.y = \text{Rightmost}(\sigma_v(\lceil y_{\text{top}}/\Delta^2 \rceil), 1, 1 + (y_{\text{top}} - 1 \bmod \Delta^2), i, j) + \Delta^2 \lceil y_{\text{top}}/\Delta^2 - 1 \rceil ,$$

and compute the size of the skyline of the intersection of B_{top} and the query region by:

$$\text{SkyCount}(\sigma_v(\lceil y_{\text{top}}/\Delta^2 \rceil), 1, 1 + (y_{\text{top}} - 1 \bmod \Delta^2), i, j) . \quad (1)$$

Let k_1 be the slab containing p_1 , computable as $k_1 = C_v(p_1.y)$. If no point is found in block B_{top} , then $k_1 = i - 1$.

Next we compute the y -coordinate of the topmost point p_2 in the multislab query range spanning slabs $[k_1 + 1, j]$ and all blocks between B_{bottom} and B_{top} .

$$p_2.y = \text{Topmost}_{v, k_1+1, j}(\lceil y_{\text{bottom}}/\Delta^2 \rceil + 1, \lceil y_{\text{top}}/\Delta^2 \rceil - 1) .$$

In the same subrange we find the y -coordinate of the rightmost point p_3 using

$$p_3.y = \text{Rightmost}_{v, k_1+1, j}(\lceil y_{\text{bottom}}/\Delta^2 \rceil + 1, \lceil y_{\text{top}}/\Delta^2 \rceil - 1) .$$

Finally, the number of points on the skyline between p_2 and p_3 (including p_2 and p_3) is computed by

$$\text{SkyCount}_{v, k_1+1, j}(\lceil y_{\text{bottom}}/\Delta^2 \rceil + 1, \lceil y_{\text{top}}/\Delta^2 \rceil - 1) . \quad (2)$$

The slab containing the point p_3 is $k_3 = C_v(p_3.y)$. We compute the number of points on the skyline to the right of p_3 in block B_{bottom} by

$$\text{SkyCount}(\sigma_v(\lceil y_{\text{bottom}}/\Delta^2 \rceil), 1 + (y_{\text{bottom}} - 1 \bmod \Delta^2), \Delta^2, k_3 + 1, j) , \quad (3)$$

and the y -coordinate of the topmost point p_4 in block B_{bottom} contained in slabs $[k_3 + 1, j]$ by

$$p_4.y = \text{Topmost}(\sigma_v(\lceil y_{\text{bottom}}/\Delta^2 \rceil), 1 + (y_{\text{bottom}} - 1 \bmod \Delta^2), \Delta^2, k_3 + 1, j) \\ + \Delta^2 \lceil y_{\text{bottom}}/\Delta^2 - 1 \rceil .$$

The remaining points to be counted are the skyline points in slab k_1 between p_1 and p_2 , and in slab k_3 between p_3 and the point p_4 in block B_{bottom} . These values can be computed by

$$\text{SkyCount}_{c^{k_1}}(\text{Succ}_v(p_2.y + 1, k_1), \text{Pred}_v(p_1.y, k_1)) - 1 \quad (4)$$

$$\text{SkyCount}_{c^{k_3}}(\text{Succ}_v(p_4.y, k_3), \text{Pred}_v(p_3.y, k_3)) - 1 , \quad (5)$$

where we subtract one in both expressions, to avoid double counting p_1 and p_3 .

Figure 2 illustrates the five partial counts computed. In the above we assumed that all queries ranges were non-empty. In case p_1 does not exist, then $k_1 = i - 1$ and (4) is not computed. If p_4 does not exist, then (5) stretches down to y_{bottom} . If p_2 and p_3 do not exist (p_2 and p_3 are the same point if (4) only contains one maximal point), then (2) and (5) are not computed, the leftmost slab of (3) is $k_1 + 1$, and (4) stretches down to $p_4.y + 1$.

To summarize, it follows that the skyline size for each multislab query R_ℓ can be computed in $O(1)$ time, and the total time for a skyline counting query becomes $O(\lg n / \lg \lg n)$.

B Skyline Range Reporting

In this section, we show how to extend our skyline range counting data structure from Section 3 to also support reporting. Given a query rectangle $R = [x_1, x_2] \times [y_1, y_2]$, we let v_1, \dots, v_m and I_1, \dots, I_m be defined as in Section A. The goal is to report the skyline for each of the subqueries $R_\ell = I_\ell \times [z_\ell, y_2]$ where $z_1 = y_1$ and $z_\ell = \max\{z_{\ell-1}, y_\ell^{\max} + 1\}$ for $\ell = 2$ to m and y_ℓ^{\max} is the maximal y -coordinate of a point in $I_1 \times [1, y_2]$. Using the approach from Section A we assume the z_ℓ 's have been computed as well as the index of the successor of y_1 and the index of the predecessor of z_ℓ in each of the L_{v_ℓ} lists. Recall the lists L_v are not stored explicitly.

To answer the query R_ℓ at a node $v = v_\ell$, let $[i, j]$ be the range of children of v that are spanned by R_ℓ in the horizontal direction and let y_{bottom} be the index of the successor of z_ℓ in L_v and y_{top} the index of the predecessor of y_2 in L_v . We first produce an output list Y_ℓ storing each point of $\text{Skyline}(R_\ell \cap P_v)$ as an index into L_v . The key observation for producing this list is that the skyline inside a query rectangle is the set of points produced by the following procedure: First report the rightmost point in the query range and then recurse on the query rectangle obtained by moving the bottom side of the query to just above the returned point.

We implement this strategy in the following. First, if R_ℓ is completely within one block (block $\lceil y_{\text{top}}/\Delta^2 \rceil$ of L_v), we answer it by first running

$$\text{Rightmost}(\sigma_v(\lceil y_{\text{top}}/\Delta^2 \rceil), 1 + (y_{\text{bottom}} - 1 \bmod \Delta^2), 1 + (y_{\text{top}} - 1 \bmod \Delta^2), i, j) .$$

Adding $\Delta^2 \lceil y_{\text{top}}/\Delta^2 - 1 \rceil$ to the returned value gives the index k into L_v of the rightmost point in the output. We add k to Y_ℓ and recurse on the query rectangle with y -range from $k + 1$ to y_{top} .

If the query range is not contained in one block, we use the decomposition into five queries that was introduced in Section A, see Figure 2. We define p_1, \dots, p_4 and (1), \dots , (5) as in Section A. The subquery (1) is answered as just described for the case of a query range completely within a block. The query (4) is answered using first the query $\text{Rightmost}_{c_v^{k_1}}(\text{Succ}_v(p_2.y + 1, k_1), \text{Pred}_v(p_1.y, k_1))$. Following that, we move the bottom of the query rectangle just above the returned point and recurse.

The query range (2) is answered by first repeatedly using $\text{Rightmost}_{v, k_1+1, j}$ to identify the blocks within the query range (2) containing points from $\text{Skyline}(R_\ell \cap P_v)$. Let q_1, \dots, q_t be the indexes into L_v of the rightmost points returned in each of these blocks, computable by

$$q_1 = \text{Rightmost}_{v, k_1+1, j}(\lceil y_{\text{bottom}}/\Delta^2 \rceil + 1, \lceil y_{\text{top}}/\Delta^2 \rceil - 1) ,$$

and for $r > 1$ (until no further point is found) by

$$q_r = \text{Rightmost}_{v, k_1+1, j}(\lceil q_{r-1}/\Delta^2 \rceil + 1, \lceil y_{\text{top}}/\Delta^2 \rceil - 1) .$$

Within each block $\lceil q_r/\Delta^2 \rceil$ we compute the additional points that should be reported within slabs $[k_1, j]$ from right-to-left, starting with

$$\text{Rightmost}(\sigma_v(\lceil q_r/\Delta^2 \rceil), 2 + (q_r - 1 \bmod \Delta^2), \Delta^2, k_1, j) ,$$

until no point is found or we find the first point f that should not be reported, i.e. f is dominated by q_{r+1} , which can be checked by the condition $\gamma < C_v(q_{r+1})$ or $\gamma = C_v(q_{r+1})$ and $q' = \text{Rightmost}_{v'}(f', q')$, where $\gamma = C_v(f)$, $v' = c_v^\gamma$ and $f' = \text{Pred}_v(f, \gamma)$, and $q' = \text{Pred}_v(q_{r+1}, \gamma)$.

The query (5) is answered by repeatedly using $\text{Rightmost}_{c_v^{k_3}}$ and finally we answer (3) using

$$\text{Rightmost}(\sigma_v(\lceil y_{\text{bottom}}/\Delta^2 \rceil), 1 + (y_{\text{bottom}} - 1 \bmod \Delta^2), \Delta^2, k_3 + 1, j)$$

and recursing above the returned point. It follows that the list Y_ℓ is produced in $O(1 + |\text{Skyline}(R_\ell \cap P_v)|) = O(1 + |\text{Skyline}(R_\ell \cap P)|)$ time. Summing over all lists Y_ℓ , we get a total time of $O(\lg n / \lg \lg n + k)$.

What remains is to map the indices in the lists Y_ℓ to the actual coordinates of the corresponding points. Using the π arrays, this can be done by repeatedly determining the position of $L_v[i]$ in L_{p_v} . Doing this for all $O(\lg n / \lg \lg n)$ levels of the tree allows one to deduce the global y -rank of the point corresponding to $L_v[i]$. Storing an additional $O(n)$ sized array mapping global y -ranks to the corresponding points gives a total running time of $O((1 + k) \lg n / \lg \lg n)$. To speed this up, we use the Ball-Inheritance structure of [1]. For completeness, we describe how this data structure is implemented in terms of the π arrays we have defined: Let $B \geq 2$ be a parameter. For every level j in the base tree T that is a multiple of B^i , for $i = 0, \dots, \lg_B \lg_\Delta n$, we let all nodes v at level j store the following array:

$\pi_v^{(B^i)}(j)$ For each $j, 1 \leq j \leq n_v$, stores the index of $L_v[j]$ in $L_{u(v)}$. Here $u(v)$ is the ancestor of v at the nearest level that is a multiple of B^{i+1} (excluding possibly the level storing v). This can be supported by constructing the select data structure of Lemma 1 on the bit-vector X , where $X[j] = 1$ if and only if $L_{u(v)}[j]$ is in L_v . A query $\pi_v^{(B^i)}(j)$ becomes $\text{select}(j)$. The space usage for $\pi_v^{(B^i)}$ becomes $O(n_v \lg(n_{u(v)}/n_v)) = O(n_v \lg(\Delta^{B^{i+1}})) = O(n_v B^{i+1} \lg \Delta)$ bits.

Given an index i into L_v , we can now recover $L_v[i]$ by using the π arrays to first jump B levels up, then B^2 levels up and so forth. The number of jumps becomes $O(\lg_B \lg_\Delta n)$ and hence we get a query time of $O(\lg n / \lg \lg n + k \lg_B \lg_\Delta n)$. The total space usage for all π arrays becomes

$$O\left(\sum_{i=1}^{\lg_B \lg_\Delta n} \frac{\lg_\Delta n}{B^i} \cdot B^{i+1} \lg \Delta\right) = O(n \lg n \cdot (B \lg_B \lg_\Delta n))$$

bits. Setting $B = \lg^\varepsilon n$ for an arbitrarily small constant $\varepsilon > 0$ gives a data structure with query time $O(\lg n / \lg \lg n + k)$ and space usage $O(n \lg^\varepsilon n)$ words. Setting $B = 2$ gives a data structure with query time $O(\lg n / \lg \lg n + k \lg \lg n)$ and space usage $O(n \lg \lg n)$ words as claimed.

C Butterfly Graph Example

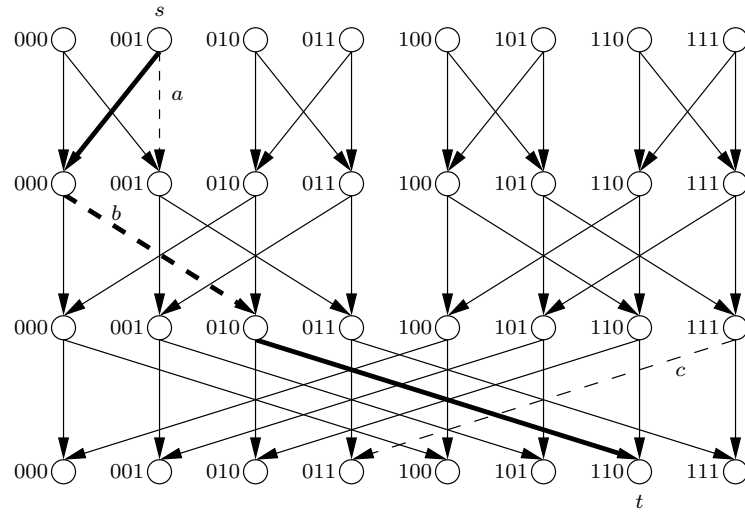


Fig. 5. A butterfly with degree $B = 2$ and depth $d = 3$. The path shown in **bold** is the unique path from the source $s = 001$ to the sink $t = 110$. A concrete input to the *reachability oracles in the butterfly graph* problem consists of a subset of the edges of the butterfly. An example input is obtained by deleting the dashed edges labelled a , b and c . For that input, there is no path from the source s to the sink t since the edge b is not part of the input.

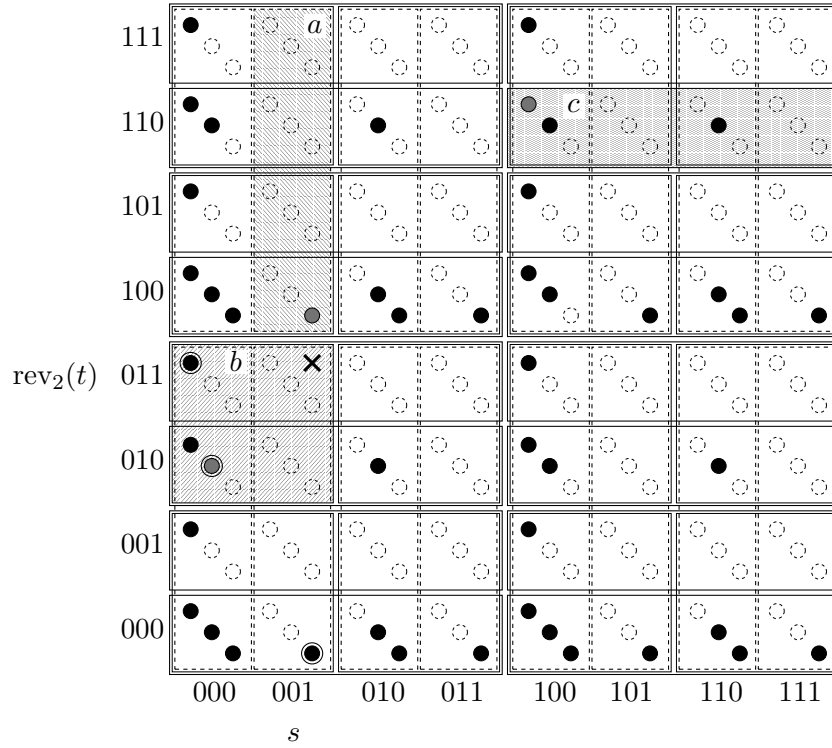


Fig. 6. The butterfly with degree $B = 2$ and depth $d = 3$ from Figure 5 translated to a set of rectangles. The dashed rectangles correspond to the edges a , b and c from Figure 5. Every grid point is replaced by up to d points placed on a diagonal. Each rectangle obtained from an edge of the butterfly graph produces one point on the diagonal corresponding to the rectangle's lower left corner. The points corresponding to the rectangles obtained from edges a , b and c are shown in gray. The query corresponding to the source $s = 001$ and the sink $t = 110$ in Figure 5 is translated to the two-sided skyline query rectangle with its upper right corner at the \times . The double circled points are the points on the skyline of the query range and these correspond exactly to the lower left corners of the rectangles containing the \times .

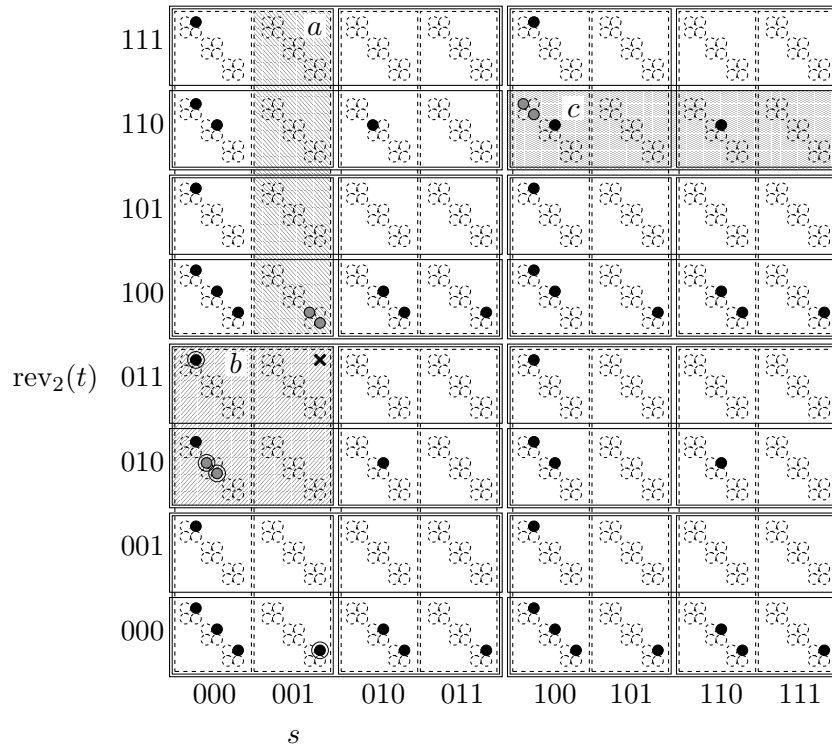


Fig. 7. Points corresponding to unmarked rectangles are replaced by two points. The example from Figure 5 and Figure 6 has three unmarked rectangles, corresponding to edges a , b and c of the butterfly graph. As shown, these rectangles become two (gray) input points and marked rectangles are represented by only one input point. The upper right corner of the two-sided query rectangle corresponding to the source $s = 001$ and sink $t = 110$ in the previous examples is shown as a \times . The double circled points are the points on the skyline of the query range. As can be seen, the unmarked rectangle corresponding to the edge labelled b contributes two points to the skyline of the query \times .