Range Reporting for Moving Points on a Grid

Marek Karpinski[1], J. Ian Munro[2], and Yakov Nekrich[1]

[1] Department of Computer Science
University of Bonn
{marek,yasha}@cs.uni-bonn.de

[2] Cheriton School of Computer Science
University of Waterloo
imunro@uwaterloo.ca

Abstract. In this paper we describe a new data structure that supports orthogonal range reporting queries on a set of points that move along linear trajectories on a $U \times U$ grid. The assumption that points lie on a $U \times U$ grid enables us to significantly decrease the query time in comparison to the standard kinetic model. Our data structure answers queries in $O(\sqrt{\log U / \log \log U} + k)$ time, where k denotes the number of points in the answer. The above improves over the $\Omega(\log n)$ lower bound that is valid in the infinite-precision kinetic model. The methods used in this paper could be also of independent interest.

1 Introduction

Data structures for querying moving objects were extensively investigated in computational geometry and database communities. The orthogonal range reporting problem, i.e. the problem of storing a set of points S in a data structure so that all points in a query rectangle Q can be reported, was also extensively studied for the case of moving points. In this paper we describe a data structure that supports range reporting queries for a set of moving points on a $U \times U$ grid, i.e., when all point coordinates are positive integers bounded by a parameter U.

Previous and Related Results. The kinetic data structure framework proposed by Basch *et al.* [7] is the standard model for studying moving objects in computational geometry. The main idea of their approach is to update the data structure for a set S of continuously moving objects only at certain moments of time: updates are performed only when certain *events* changing the relevant combinatorial structure of the set S occur. For instance, the data structure may be updated when the order of projections of points on the x-axis changes or the closest pair of points in S changes; see e.g., [7, 15] for a more detailed description.

The kinetic variant of the range tree data structure was presented by Basch, Guibas, and Zhang [8]; their data structure uses $O(n \log^{d-1} n)$ space, answers d-dimensional queries in $O(\log^d n + k)$ time, and can be updated after each event in $O(\log^d n)$ time; henceforth k denotes the number of points in the answer. The two-dimensional data structure of Agarwal, Arge, and Erickson [2] supports range reporting queries in $O(\log n + k)$ time and uses $O(n \log n / \log \log n)$ space; the cost of updating their data structure after each event is $O(\log^2 n)$. As follows from standard information-theoretic arguments, the $O(\log n)$ query time is optimal in the infinite-precision kinetic model. Linear space kinetic data structures were considered by Agarwal, Gao, and Guibas [3] and Abam, de Berg, and Speckmann [1]. However these data structures have significantly higher query times: the fastest linear space construction [1] answers d-dimensional queries in $O(n^{1-1/d} + k)$ time.

A number of geometric problems can be solved more efficiently when points lie on a grid, i.e., when coordinates of points are integers[1] bounded by a parameter U. In the case of range reporting, significant speed-up can be achieved if the set of points S does not change. There are static data structures that support orthogonal range reporting queries in $O(\log \log U + k)$ time [17, 4]. On the other hand, if points can be inserted into or deleted from S, then any data structure that supports updates in $\log^{O(1)} n$ time needs $\Omega(\log n / \log \log n + k)$ time to answer a two-dimensional range reporting query [5]. This bound is also valid in the case when all points belong to a $U \times U$ grid.

Our Result. In this paper we consider the situation when coordinates of moving points belong to a $U \times U$ grid. Our data structure supports orthogonal range reporting queries in $O(\sqrt{\log U / \log \log U} + k)$ time. This result is valid in the standard kinetic model with additional conditions that all points move with fixed velocities along linear trajectories and all changes in the trajectories are known in advance. Queries can be answered at any time t, where t is a positive integer bounded by $U^{O(1)}$. Updates are performed only when x- or y-coordinates of any two points in S swap their relative positions, and each update takes poly-logarithmic time. The total number of events after which the data structure must be updated is $O(n^2)$. For instance, for $U = n^{O(1)}$ our data structure answers queries in $O(\sqrt{\log n / \log \log n} + k)$ time. Our result also demonstrates that the lower bound for dynamic range reporting queries can be surpassed in the case when the set S consists of linearly moving points. Our data structure uses $O(n \log^2 n)$ space and supports updates in $O(\log^3 n)$ time, but space usage and update cost can be reduced if only special cases of reporting queries must be supported. We describe a $O(n)$ space data structure that supports updates in $O(\log n)$ time and dominance queries in $O(\sqrt{\log U / \log \log U} + k)$ time. We also describe a $O(n \log n)$ space data structure that supports updates in $O(\log^2 n)$ time and three-sided[2] queries in $O(\sqrt{\log U / \log \log U} + k)$ time.

2 Overview

In section 3 we show that we can find the predecessor point of any $v \in U$ in the set S (with respect to x- or y-coordinates) in $O(\sqrt{\log U / \log \log U})$ time by answering a point location query among a set of segments. In fact, identifying the predecessor of a point q is the bottleneck of our query answering procedure.

In section 4 we describe the data structure that reports all points $p \in S$ that dominate the query point q, i.e. all points p such that $p.x \geq q.x$ and $p.y \geq q.y$; henceforth $p.x$ and $p.y$ denote the x- and y-coordinates of a point p. The query time of our data structure is $O(\sqrt{\log U / \log \log U} + k)$. The data structure is based on the modification of the d-approximate boundary [19] for the kinetic framework. The d-approximate boundary [19] enables us to obtain an estimation for the number of points in S that dominate an arbitrary point q. If a point q dominates a point on a d-approximate boundary $\mathcal{M}$, then q is dominated by at most $2d$ points of S; if q is dominated by a point on $\mathcal{M}$, then q is dominated by at least d points of S. In section 4 we show that a variant of a d-approximate boundary can be maintained under kinetic events. If a query point q is dominated by $k \leq \log n$ points of S, we can reduce the dominance query on S to a dominance query on a set that contains $O(\log n)$ points using a d-approximate boundary for $d = \log n$; see section 4. Otherwise, if $k > \log n$, we can answer a query in $O(\log n + k) = O(k)$ time using a standard kinetic data structure [2].

---

[1] For simplicity, we assume that all points have positive coordinates.

[2] The query range of a dominance query is a product of two half-open intervals. The query range of a three-sided query is a product of a closed interval and a half-open interval.

A data structure that supports dominance queries can be transformed into a data structure that supports arbitrary orthogonal range reporting queries [11, 18] by dividing the set S into subsets S_i and constructing dominance data structures for each S_i as described in section 5. However, we may have to delete a point p from one subset S_i and insert it into a subset S_j after a kinetic event. Unfortunately, the construction of [19] is static. It is not clear how (and whether) to modify the d-approximate boundary, so that insertions and deletions are supported. However, in our case the deleted (inserted) point always has the maximal or minimal x- or y-coordinate among all points in S_i (S_j). We will describe in section 5 how our dominance data structure can be modified to support these special update operations without increasing the query time. Our technique is similar to the logarithmic method and can be of independent interest. Thus we obtain the data structure for general orthogonal range reporting queries.

3 One-Dimensional Searching

Let $S_x(t)$ and $S_y(t)$ denote the sets of x- and y-coordinates of all points at time t. In this section we will describe how we can identify the predecessor of any q_x in $S_x(t)$ (resp. of q_y in $S_y(t)$) at current time t in $O(\sqrt{\log U / \log \log U})$ time using a linear space data structure.

Let $x_i(t) = a_i t + b_i$ be the equation that describes the x-coordinate of the point $p_i \in S$ at time t. The trajectory of the point in (t, x) plane (t-axis is horizontal) is a sequence of segments. Since we assume that all changes of point trajectories are known in advance, endpoints of all segments are known in advance. Two points swap ranks of their x-coordinates at time t if and only if their segments intersect at time t. We can find intersection points of all segments using the standard sweepline algorithm [10] in $O((n + f) \log n)$ time, where f is the number of segment intersections. We start the sweepline at $t = 0$ and move it to the right until n intersection points are identified or the last intersection point is found. These intersection points and the corresponding segments induce a subdivision of the (x, t) plane of size $O(n)$. We can construct the data structure for planar point location [12] that supports queries in $O(\sqrt{\log U / \log \log U})$ time. Let t_{max} be the largest t-coordinate of the already processed intersection point. For $t < t_{max}$, we can find the predecessor of any x by locating the segment lying immediately below the point (t, x). When $t = t_{max}$, we continue the sweepline algorithm and find the next n segment intersection points. The algorithm described in [10] finds f next segment intersection points $O(f \log n)$ time. Since the point location data structure [12] for a subdivision of size f can be constructed in $O(f)$ time, an amortized cost of processing a kinetic event is $O(\log n)$. We can de-amortize the update cost using standard techniques.

4 Dominance Queries

In this section we describe the data structure that reports all points from S that dominate the query point q, i.e. all points in the region $[q.x, +\infty) \times [q.y, +\infty)$. Our data structure is based on maintaining the d-approximate boundary for a set S. The notion of a d-approximate boundary is introduced in [19]; in this paper we change the definition and describe a kinetic version of this construction.

For a horizontal segment s, we denote by $\mathtt{start}(s)$ and $\mathtt{end}(s)$ x-coordinates of the left and the right endpoint of s; we denote by $y(s)$ the y-coordinate of all points on s. We will say that a segment s covers a point p if the x-coordinate of p belongs to $[\mathtt{start}(s), \mathtt{end}(s)]$. In this paper

Fig. 1. An example of a d-approximate boundary for $d = 6$.

we define a d-approximate boundary as a polyline $\mathcal{M}$ that consists of alternating horizontal and vertical segments, divides the plane into two parts, and satisfies the following properties:

Invariant 1 *Let s and r be two consecutive horizontal segments. Then $|\{p \in S \mid \texttt{start}(s) \leq p.x \leq \texttt{end}(r)\}| > d/2$.*

Invariant 2 *When a new segment s is inserted, the left endpoint of s is dominated by at most $3d/2$ points of S. The number of points in S that dominate the left endpoint of a segment $s \in \mathcal{M}$ does not exceed $2d$.*

Invariant 3 *When a new segment s is inserted, the right endpoint of s is dominated by at least d points. The number of points in S that dominate the right endpoint of a segment $s \in \mathcal{M}$ remains constant.*

Using Invariants 1-3, we can prove the following Lemma.

Lemma 1. *Every point on a d-approximate boundary $\mathcal{M}$ is dominated by at least d points and at most $2d$ points of S. There are $O(n/d)$ horizontal segments in $\mathcal{M}$.*

Proof: If a point $p \in \mathcal{M}$ is dominated by k points from S, then the left endpoint of some segment s is dominated by at least k points and the right endpoint of some segment r is dominated by at most k points. Hence, it follows from Invariants 2 and 3 that $d \leq k < 2d$. By Invariant 1, for two consecutive segments r and s there are more than $d/2$ points $p \in S$, such that $p.x$ belongs to the interval $[\texttt{start}(r), \texttt{end}(s)]$. For each point p, $p.x$ belongs to at most two such intervals; hence, the total number of segments is less than $8n/d$. $\qquad\qquad\square$

An example of a d-approximate boundary is shown on Fig 1. We will show below how the concept of a d-approximate boundary can be used to support dominance queries in $O(\sqrt{\log U / \log \log U} + k)$ time. Later in this section we will show how Invariants 1 -3 can be maintained.

Kinetic Boundary. We will use a kinetic variant of the d-approximate boundary, i.e. segments of the boundary move together with points of S. For every horizontal segment s in a boundary $\mathcal{M}$, let $\texttt{l}(s)$ denote the point with the largest y-coordinate such that $\texttt{l}(s).y < y(s)$ and let $\texttt{u}(s)$ denote the point with the smallest y-coordinate such that $\texttt{u}(s).y > y(s)$. Let $\texttt{first}(s)$ denote the point

4

with the smallest x-coordinate such that $\mathtt{first}(s).x > \mathtt{start}(s)$; let $\mathtt{last}(s)$ denote the point with the largest x-coordinate such that $\mathtt{last}(s).x < \mathtt{end}(s)$. We assume that $\mathtt{start}(s) = \mathtt{first}(s) - \frac{1}{2}$ and $y(s) = \mathtt{u}(s) - \frac{1}{2}$. That is, the left end and the y-coordinate of a segment change when $\mathtt{u}(s)$ and $\mathtt{first}(s)$ move. The right end of the previous segment and the x-coordinate of the connecting vertical segment change accordingly.

Answering Queries. Our data structure is based on a d-approximate boundary $\mathcal{M}$ of S for $d = \log n$. For each segment $s \in \mathcal{M}$ we maintain the set $Dom(s)$ of all points that dominate the left endpoint of s. Obviously, the set $Dom(s)$ changes only when events concerning $\mathtt{first}(s)$ or $\mathtt{u}(s)$ take place.

All points of $Dom(s)$ are stored in a data structure D_s that supports dominance queries in $O(\log d + k) = O(\log \log n + k)$ time. We can use the data structure of [2], so that the space usage is $O(d)$ and updates after events are supported in $O(\log \log n)$ time. It is possible to modify the data structure of [2], so that points can be inserted into $Dom(s)$ or deleted from $Dom(s)$ in $O(\log \log n)$ time. All points of S are also stored in a kinetic data structure G that uses $O(n)$ space and supports dominance queries in $O(\log n + k)$ time and updates after kinetic events in $O(\log n)$ time. Again, we can use the result of [2] to implement G. Finally, we must be able to identify for each point $p \in S$ the segment $s \in \mathcal{M}$ that covers p. Using the dynamic union-split-find data structure of [13] or the van Emde Boas data structure [14], we can find the segment that covers any $p \in S$ in $O(\log \log n)$ time. When a new segment is inserted into or deleted from $\mathcal{M}$, the data structure is updated in $O(\log \log n)$ time.

Given a query point q, we identify the point p with the largest x-coordinate such that $p.x \le q.x$; this can be done in $O(\sqrt{\log U / \log \log U})$ time using the construction described in section 3. The point q dominates a point on $\mathcal{M}$ if and only if q dominates the left endpoint of the segment s that covers p or the left endpoint of the segment h that follows s. Suppose that q dominates a point on $\mathcal{M}$. Then q is dominated by at most $2d$ points of S. Let v be the left endpoint of a segment g, such that v is dominated by q. Each point $p \in S$ that dominates q also dominates v; hence, all points that dominate q belong to $Dom(g)$. We can use the data structure D_g and report all points that dominate q in $O(\log \log n + k)$ time. Suppose that q does not dominate any point on $\mathcal{M}$. Then there is at least one point on $\mathcal{M}$ that dominates q. Hence, there are $k \ge d = \log n$ points of S that dominate q. Using data structure G, we can report all points that dominate q in $O(\log n + k) = O(k)$ time.

Maintaining the d-Approximate Boundary. It remains to show how to maintain Invariants 1-3 after operations x-move and y-move. Henceforth, we will use the following notation. Suppose that $p.x < q.x$ before some kinetic event and $p.x > q.x$ after this event. Then, we say that p is x-moved behind q (q is x-moved before p). Suppose that $p.y < q.y$ before some kinetic event and $p.y > q.y$ after this event. Then, we say that p is y-moved above q (q is y-moved below p). Each kinetic event can be represented as a combination of at most one x-move and at most one y-move.

First, we consider the Invariant 2. Suppose that the left endpoint of an interval s is dominated by $2d$ points of S. Let h be the segment that precedes s, i.e., $\mathtt{end}(h) = \mathtt{start}(s)$. Let v be the vertical segment that connects h and s. We look for a point p with $p.x = \mathtt{start}(s)$ and $y(s) < p.y < y(h)$ such that p is dominated by $3d/2$ points. If such a point p exists, we set $p_l = p$ and search for a point p_r with $\mathtt{start}(s) < p_r.x < \mathtt{end}(s)$ and $p_r.y = p_l.y$ such that p_r is dominated by d points of S. If there is no such point, i.e., if the point $(\mathtt{end}(s), p_l.y)$ is dominated by at least d points, we replace s with a segment s' such that the left endpoint of s' is p_l and the right endpoint of s' is the point $(\mathtt{end}(s), p_l.y)$; see Fig 2a. In other words we change the y-coordinate of s to $p.y$. The new

Fig. 2. Updating the d-approximate boundary when the Invariant 2 is violated. For simplicity point of S are not shown. **(a)** There exists a point p_l but there is no p_r. **(b)** There exist both p_l and p_r. **(c)** There are neither p_l nor p_r.

set $Dom(s)$ contains $3d/2$ points and can be constructed in $O(d)$ time. If there is a point p_r with $\texttt{start}(s) < p_r.x < \texttt{end}(s)$ and $p_r.y = p_l.y$ that is dominated by d points of S, then we replace s with two new segments s' and s''. The left and right endpoints of s' are p'_l and p'_r respectively. The left endpoint of s'' is the point p'' such that $p''.x = p_r.x$ and $p''.y = y(s)$. See Fig 2b. The set $Dom(s')$ contains $3d/2$ points. There are at most $d/2$ points q such that $q.x > \texttt{start}(s)$ and $y(s) < q.y \le p.y$; hence, there are at most $d/2$ points q such that $q.x \ge p_r.x$ and $y(s) < q.y \le p.y$. Therefore, since p_r is dominated by d points of S, p'' is dominated by at most $3d/2$ points of S and $Dom(s'')$ contains at most $3d/2$ points. Since p'' is dominated by the right endpoint of the segment s, $Dom(s'')$ contains at least d points. We can construct $Dom(s')$ and $Dom(s'')$ and data structures $D_{s'}$ and $D_{s''}$ in $O(d)$ time.

If the right endpoint of h is dominated by at least $3d/2$ points, we shift the vertical segment v in $+x$ direction, so that the right endpoint of h is dominated by d points from S or the segment s is removed. That is, we identify the point r with $r.y = y(h)$ such that either $r.x = \texttt{end}(s)$ and r is dominated by at least d points of S or $r.x < \texttt{end}(s)$ and r is dominated by d points of S. If $r.x < \texttt{end}(s)$, we set $\texttt{end}(h) = \texttt{start}(s) = r.x$ and update $Dom(s)$, D_s accordingly ($O(d)$ points are removed from $Dom(s)$ and D_s). See Fig 2c. The new left endpoint of s is dominated by at most $3d/2$ points. If $r.x = \texttt{end}(s)$, we remove the segment s with D_s and $Dom(s)$.

The update procedure removes at most one segment and inserts at most two new segments that satisfy the Invariant 2. Hence, the update procedure takes $O(d) = O(\log n)$ time.

Now we describe how the Invariant 1 can be maintained. Let h and s be two consecutive segments $(\texttt{end}(h) = \texttt{start}(s))$ and suppose that there are $d/2$ points that belong to $[\texttt{start}(h), \texttt{end}(s)]$. We replace h and s with one new segment g as follows. If the left endpoint of h is dominated by at least $3d/2$ points, then we set $\texttt{end}(h) = \texttt{end}(s)$ and remove the segment s, $Dom(s)$, and D_s. The point q with $q.y = y(h)$ and $q.x = \texttt{end}(s)$ is the new right endpoint of h. Since there are at most $d/2$ points $p \in S$ such that $\texttt{start}(s) \le p.x \le q.x$, q is dominated by at least d points of S. Hence, the new segment h satisfies Invariant 3. See Fig. 3a. If the point l with $l.x = \texttt{start}(h)$ and $l.y = y(s)$ is dominated by at most $3d/2$ points of S, we set $\texttt{start}(s) = \texttt{start}(h)$ and remove h, $Dom(h)$ and D_h. The set $Dom(s)$ and the data structure D_s are updated. See Fig. 3b. Since there are less than $d/2$ points $p \in S$, such that $p.y > y(s)$ and $\texttt{start}(h) \le p.x \le \texttt{end}(h)$, $O(d)$ new points are inserted into $Dom(s)$ and D_s. If the left endpoint of h is dominated by less than $3d/2$ points and the point l is dominated by more than $3d/2$ points, then we replace h and s with a new segment g. The left endpoint of g is the point m such that $m.x = \texttt{start}(h)$, $y(s) < m.y < y(h)$, and m is dominated by $3d/2$ points. The right endpoint of g is the point r with $r.x = \texttt{end}(s)$ and $r.y = m.y$. See Fig. 3c. The point r is dominated by at least d points of S. Hence, g satisfies Invariants 2 and 3.

Fig. 3. Updating the d-approximate boundary when the Invariant 1 is violated. For simplicity points of set S are not shown. **(a)** The left endpoint of h is dominated by $\geq 3d/2$ points. **(b)** The point l is dominated by at most $3d/2$ points. **(c)** The left endpoint of h is dominated by less than $3d/2$ points but l is dominated by more than $3d/2$ points.

Now we turn to Invariant 3. The number of points that dominate the right endpoint of a horizontal segment s changes: 1) if a point p with $p.x > \mathtt{end}(s)$ is y-moved above $\mathtt{u}(s)$ or below $\mathtt{u}(s)$; 2) if a point p with $p.y > y(s)$ is x-moved before $\mathtt{last}(s)$ or behind $\mathtt{last}(s)$.

First, we consider y-moves. Essentially, when a point p is y-moved, we shift the segment s in $+y$ or $-y$ direction so that the number of points that dominate the right endpoint of s remains unchanged. Suppose that a point $p = \mathtt{u}(s)$ with $p.x > \mathtt{end}(s)$ is y-moved below $\mathtt{l}(s)$; let y_n denote the y-coordinate of p.

(a) If $\mathtt{start}(s) < \mathtt{l}(s).x < \mathtt{end}(s)$, then we change the y-coordinate of s so that $y(s) = y_n - \frac{1}{2}$. The old point $\mathtt{l}(s)$ is added to $Dom(s)$ and D_s.

(b) If $\mathtt{l}(s).x < \mathtt{start}(s)$, then we also change the y-coordinate of s so that $y(s) = y_n - \frac{1}{2}$.

(c) If $\mathtt{l}(s).x > \mathtt{end}(s)$, then we change the y-coordinate of s so that $y(s) = \mathtt{l}(s).y - \frac{1}{2}$. We delete p from $Dom(s)$ and D_s and insert $\mathtt{l}(s)$ into $Dom(s)$ and D_s.

Suppose that a point p with $p.x > \mathtt{end}(s)$ is y-moved above $\mathtt{u}(s)$

(d) If $\mathtt{start}(s) < \mathtt{u}(s).x < \mathtt{end}(s)$, then we change the y-coordinate of s so that $y(s) = q.y - \frac{1}{2}$, where q is the point with the smallest y-coordinate such that $q.y > y_n$. The old point $\mathtt{u}(s)$ is removed from $Dom(s)$ and D_s. If $q = \mathtt{u}(h)$ for the segment h that precedes s, then the segment s and $Dom(s)$ are deleted.

(e) If $\mathtt{u}(s).x < \mathtt{start}(s)$, then we also change the y-coordinate of s so that $y(s) = q.y - \frac{1}{2}$, where q is the point with the smallest y-coordinate such that $q.y > y_n$. If $q = \mathtt{u}(h)$ for the segment h that precedes s, then the segment s and $Dom(s)$ are deleted.

(f) If $\mathtt{u}(s).x > \mathtt{end}(s)$, then we change the y-coordinate of s so that $y(s) = y_n - \frac{1}{2}$. We delete $\mathtt{u}(s)$ from $Dom(s)$ and D_s and insert p into $Dom(s)$ and D_s.

Observe that the number of points in D_s remains unchanged or increases by 1. See Fig 4.

We can handle the x-moves in a similar way. Let h be the horizontal segment that follows s in $\mathcal{M}$, i.e., $\mathtt{end}(s) = \mathtt{start}(h)$. When a point p is x-moved we change $\mathtt{end}(s)$ (and $\mathtt{start}(s)$) so that the number of points that dominate the right endpoint of s remains unchanged. Suppose that a point $p = \mathtt{last}(s)$ with $p.y > y(s)$ is x-moved behind $\mathtt{first}(h)$; let x_n denote the new x-coordinates of p and let q be the point with the smallest x-coordinate such that $q.x > x_n$.

(a) If $y(h) < \mathtt{first}(h).y < y(s)$, then we set $\mathtt{start}(h) = \mathtt{end}(s) = q.x - \frac{1}{2}$ and remove $\mathtt{first}(h)$ from $Dom(h)$ and D_h. If $q = \mathtt{first}(h')$ for the horizontal segment h' that follows h, then we delete the segment h and $Dom(h)$.

(b) If $\mathtt{first}(h).y > y(s)$, then we set $\mathtt{start}(h) = \mathtt{end}(s) = x_n - \frac{1}{2}$; we also remove $\mathtt{first}(h)$ from

Fig. 4. The segment s is shifted so that Invariant 1 is maintained after a y-move. Figures **(a)**, **(b)**,**(c)**, **(d)**, **(e)**, and **(f)** correspond to cases (a), (b), (c), (d), (e), and (f) respectively.

$Dom(h)$ and D_h and insert p into $Dom(h)$ and D_h.

(c) If $\texttt{first}(h).y < y(h)$, then we set $\texttt{start}(h) = \texttt{end}(s) = q.x - \frac{1}{2}$. If $q = \texttt{first}(h')$ for the horizontal segment h' that follows h, then we delete the segment h and $Dom(h)$.

Suppose that a point p with $p.y > y(s)$ is x-moved before $\texttt{last}(s)$.

(d) If $y(h) < \texttt{last}(s).y < y(s)$, then we set $\texttt{start}(h) = \texttt{end}(s) = x_n - \frac{1}{2}$; we add $\texttt{last}(s)$ to $Dom(h)$ and D_h.

(e) If $\texttt{last}(s).y > y(s)$, then we set $\texttt{start}(h) = \texttt{end}(s) = \texttt{last}(s).x - \frac{1}{2}$; we also remove p from $Dom(h)$ and D_h and insert $\texttt{last}(s)$ into $Dom(h)$ and D_h.

(f) If $\texttt{last}(s).y < y(h)$, then we set $\texttt{start}(h) = \texttt{end}(s) = x_n - \frac{1}{2}$. We add $\texttt{last}(s)$ to $Dom(h)$ and D_h

Again the number of points in D_s remains unchanged or increases by 1.

We will show in the full version that we update $\mathcal{M}_i$ because Invariants 2 or 1 are violated at most once for $\Omega(d)$ events. Update procedures for maintaining Invariants 1 and 2 involve inserting and deleting a constant number of segments into $\mathcal{M}$ and the data structure D_s for every segment s contains $O(d)$ points. Hence, we must perform $O(1)$ amortized updates of data structures D_s after each kinetic event. Since every update of D_s takes $O(\log \log n)$ time, the amortized cost of updating after an event is $O(\log \log n)$. Since each data structure D_s can be constructed in $O(\log n)$ time, the worst-case update time is $O(\log n)$.

Construction of a d-Approximate Boundary. Now we show that d-approximate boundary can be constructed in $O(n)$ time if points are sorted by x- and y-coordinates. Since a d-approximate boundary consists of alternating horizontal and vertical segments, it suffices to determine the endpoints of horizontal segments. We can guarantee that the left endpoint of each segment is dominated by $3d/2$ points of S and the right endpoint of each segment is dominated by d points of S using the following algorithm. Lists L_x and L_y contain points of S sorted in descending order of their

x-coordinates and y-coordinates respectively. With every element of L_x we store a pointer to its position in L_y and vice versa. The pointer ptr_x (ptr_y) points to the first not yet processed element in L_x (L_y). We construct a sequence of horizontal and vertical segments so that the left endpoint of each segment is dominated by $3d/2$ points and the right endpoint of each segment is dominated by d points.

We assign ptr_y to the $(3d/2+1)$-th element of L_y and ptr_x to the first element of L_x; the point p_l with $p_l.y = \mathrm{ptr}_y.y - 1/2$ and $p.l.x = 0$ is the left endpoint of the first segment. Clearly, p_l is dominated by $3d/2$ points. (1) The right endpoint of the segment with left endpoint p_l can be found as follows. We traverse elements of L_x that follow ptr_x until $d/2+1$ points q such that $q.y > \mathrm{ptr}_y.y$ are visited and update ptr_x accordingly. The point p_r with $p_r.x = \mathrm{ptr}_x.x - 1/2$ and $p_r.y = p_l.y$ is the right endpoint of the currently constructed horizontal segment. (2) We identify the left endpoint of the next segment by traversing elements of L_y that follow ptr_y until $d/2$ points q such that $q.x > \mathrm{ptr}_x.x$ are visited or we reach the end of the list L_y. The pointer ptr_y is updated accordingly. If we reached the end of L_y, then all points are processed and the algorithm is completed. Otherwise we set $p_l.x = \mathrm{ptr}_x.x - 1/2$ and $p_l.y = \mathrm{ptr}_y.y - 1/2$, go to step (1) and determine the right endpoint of the next segment.

Theorem 1. *There exists a linear space data structure that supports dominance queries for a set of linearly moving points on $U \times U$ grid in $O(\sqrt{\log U / \log \log U} + k)$ time and updates after kinetic events in $O(\log n)$ worst-case time. The amortized cost of updates is $O(\log \log n)$. If trajectories of the points do not change, then the total number of kinetic events is bounded by $O(n^2)$.*

5 Orthogonal Range Reporting Queries

Three-sided range reporting queries and orthogonal range reporting queries can be reduced to dominance queries with help of standard techniques. However to apply these techniques in our scenario, we must modify the data structure of section 4, so that insertions and deletions are supported in some special cases. At the end of this section we demonstrate how our data structure with additional operations can be used to support arbitrary orthogonal range reporting queries.

Additional Update Operations. We first describe the data structure that supports insertions and deletions in two special cases: Let $x_{\min}$ and $x_{\max}$ be the smallest and the largest x-coordinates of points in S. The operation insert_x^+ inserts a point p with $x_{\max} < p.x$. The operation delete_x^+ deletes a point p with $p.x = x_{\max}$. Operations insert_x^- and delete_x^- insert and delete a point whose x-coordinate is smaller than x-coordinates of all other points in S. It is easy to augment our data structure so that insert_x^- and delete_x^- are supported: the inserted (deleted) point is either below a d-approximate boundary $\mathcal{M}$ or dominates only the leftmost horizontal segment of $\mathcal{M}$. Hence, each insert_x^- and delete_x^- affects the data structure D_s for at most one segment s. Essentially we can handle insert_x^- and delete_x^- in the same way as x-moves for the leftmost segment s. Maintaining the d-approximate boundary $\mathcal{M}$ after insert_x^+ and delete_x^+ is more involved: since the y-coordinate of a newly inserted (deleted) point can be larger than the y-coordinates of all (other) points in p, we may have to update data structures D_s for all segments $s \in \mathcal{M}$ after a single update operation. Below we describe how insert_x^+ and delete_x^+ can be supported.

Our approach is similar to the logarithmic method [9, 16] that is used to transform static data structures into data structures that support insertions. We construct the data structure of section 4 augmented with insert_x^- and delete_x^- for sets $H_2, H_3, \ldots, H_m$. A point p x-overlaps with point q if

$p.x > q.x$. A point p x-overlaps with a set S if it x-overlaps with at least one point $q \in S$. Each set H_i satisfies the following conditions:

1. For $i = 2, \ldots, m - 1$, H_i contains between 2^{2i-1} and 2^{2i+2} points; H_m contains between 2^{2i-2} and $2^{2i+2} + 2^{2i}$ points
2. Each point of H_i x-overlaps at most 2^{2i-4} points in H_{i-1}
3. At most 2^{2i-3} points from H_i x-overlap with H_{i-1}

As follows from conditions 2 and 3, no element of H_{i+1} x-overlaps with H_{i-1}: The rightmost point in H_{i+1} x-overlaps at most 2^{2i-2} leftmost points in H_i by condition 2. Only 2^{2i-3} rightmost points in H_i can x-overlap with a point in H_{i-1}. Hence, any $q \in H_{i+1}$ that x-overlaps a point in H_{i-1} would x-overlap $|H_i| - 2^{2i-3} > 2^{2i-2}$ points in H_i, which contradicts condition 2.

For $j = 1, \ldots, m$, we maintain $\max_y(j) = \max\{p.y | p \in H_j\}$ and $\min_x(j) = \min\{p.x | p \in H_j\}$. All $\min_x(j)$ are stored in a kinetic binary search tree. For each j, $2 \leq j \leq m$, a point p_j such that $p_j.y = \max_y(j)$ and $p_j.x = j$ is stored in a kinetic data structure Y; since Y contains $O(\log n)$ points, Y supports dominance reporting queries in $O(\log \log n + k)$ time. The data structure D_j, $j = 2, \ldots, m$, contains all points from H_j and supports dominance queries, x-moves and y-moves as described in section 4. To speed-up update operations, we store only one data structure G for all points in $\cup_{j=2}^{m} H_j$. G is implemented as described in [2], supports dominance queries in $O(\log n + k)$ time and can be modified to support arbitrary updates as well as kinetic events in $O(\log n)$ time. We also store one data structure $\mathcal{V}$ that contains x-coordinates of all points $\cup_{j=2}^{m} H_j$ and enables us to search in the set of x-coordinates at any time d. That is, all D_j share one data structure G and one data structure $\mathcal{V}$. For each j we also store all points of H_j in a list L_j that contains all points from H_j in the descending order of their y-coordinates.

Given a query $Q = [a, +\infty) \times [b, +\infty)$, we can find in $O(\log \log n)$ time the smallest index j, such that at least one point in H_j has x-coordinate smaller than a. Then, as follows from conditions 2 and 3, x-coordinates of all points in $H_2 \cup \ldots \cup H_{j-1}$ are greater than a, and both H_j and H_{j+1} may contain points whose x-coordinates are greater than or equal to a. Sets H_f, $f > j + 1$, contain only points whose x-coordinates are smaller than a. Using Y, we can identify all H_f such that $f < j$ and H_f contains at least one point p with $p.y \geq b$. For every such f all points p such that $p \in H_f$ and $p.y \geq b$ can be reported by traversing the list L_f. Hence, reporting all points $p \in H_f$ such that $f < j$, $p.x \geq a$ and $p.y \geq b$ takes $O(\log \log n + k)$ time. We can report all points in H_j and H_{j+1} that belong to $Q = [a, +\infty) \times [b, +\infty)$ in $O(\sqrt{\log U / \log \log U} + k)$ time using data structures D_j and D_{j+1} respectively.

It remains to show how conditions 1-3 above can be maintained. Clearly, conditions 1-3 are influenced by x-moves and operations insert_x^-, delete_x^-, insert_x^+ and delete_x^+; y-moves cannot violate them. We say that a set S is x-split into sets S_1 and S_2 if $S_1 \cup S_2 = S$ and the x-coordinates of all points in S_1 are larger than the x-coordinates of all points in S_2. After an operation insert_x^+ or delete_x^+, we re-build the data structure for H_2 in $O(1)$ time. When the number of elements in a set H_j, $j < m$, becomes smaller than 2^{2j-1} or greater than 2^{2j+1}, we x-split the set $H_{j+1} \cup H_j$ into two new sets H_j' and H_{j+1}', so that H_j' contains 2^{2j} points and H_{j+1}' contains $|H_j| + |H_{j+1}| - 2^{2j}$ points. If the number of elements in H_m exceeds $2^{2m+2} + 2^{2m}$ we x-split H_j into sets H_m' and H_{m+1}' that contain 2^{2m} and 2^{2m+2} points respectively. Suppose that the number of points in H_m is smaller than 2^{2m-2}. If $|H_{m-1}| + |H_m| \leq 3 \cdot 2^{2m-2}$, we set $H_{m-1}' = H_{m-1} \cup H_m$ and decrement m by 1. If $|H_{m-1}| + |H_m| > 3 \cdot 2^{2m-2}$, we x-split $H_{m-1} \cup H_m$ into H_{m-1}' and H_m' that contain 2^{2m-2} and $|H_{m-1}| + |H_m| - 2^{2m-2}$ points respectively. We also take care that conditions 2 and 3 are maintained. If, as a result of x-moves, the number of points in some H_j that x-overlap H_{j-1} exceeds 2^{2j-3}, or

there is at least one point in H_j that x-overlaps more than 2^{2j-4} points in H_{j-1}, then we x-split $H_{j-1} \cup H_j$ into sets H'_{j-1} and H'_j that contain $|H_{j-1}|$ and $|H_j|$ elements respectively. Each set H_j is rebuilt at most once after a sequence of $\Theta(H_j)$ special insert or delete operations. Each H_j is also re-built at most once after a sequence of $\Theta(|H_j|)$ x-moves, i.e., after $\Theta(H_j)$ kinetic events. We can maintain the list of points in each H_j sorted by their y-coordinates; hence, the data structure D_j for a newly re-built set H_j can be constructed in $O(|H_j|)$ time. Therefore, the amortized cost of updates and kinetic events is $O(\log n)$. We can de-amortize update costs using the same techniques as in the logarithmic method [16].

Lemma 2. *There exists a $O(n)$ space data structure that supports dominance queries on $U \times U$ grid in $O(\sqrt{\log U / \log \log U} + k)$ time. Updates after kinetic events and operations* insert_x^+, delete_x^+, insert_x^- *and* delete_x^- *are supported in $O(\log n)$ time.*

Three-Sided Reporting Queries. Now we are ready to describe data structures that support three-sided reporting queries, i.e., the query range is a product of a closed interval and a half-open interval. We apply the standard method used in e.g. [19], [18] to augment data structure for dominance queries.

Let T be an arbitrary balanced tree with constant node degree on the set of x-coordinates of all points in S. We associate an interval $(a_l, b_l]$ with each leaf l of T, where a_l is the predecessor of the smallest value m_l stored in the node l, and b_l is the largest value stored in the node l. We associate an interval $(a_l, +\infty)$ with the rightmost leaf l. With each internal node v of T we associate an interval $\mathtt{int}(v) = \cup \mathtt{int}(v_i)$ for all children v_i of v. Let S_v be the set of points $p \in S$ such that $p.x \in \mathtt{int}(v)$. In every internal node v we store two data structures $\mathcal{L}_v$ and $\mathcal{R}_v$ that support dominance reporting queries open to the left and open to the right (i.e., queries $(-\infty, a] \times (-\infty, b]$ and $[a, +\infty) \times (-\infty, b]$) for the set S_v. In each node v we also store the list of points in S_v sorted by their y-coordinate.

Given a three-sided query with $Q = [a, b] \times (-\infty, c]$ we can find in time $O(\sqrt{\log U / \log \log U})$ the node v such that $[a, b] \subset \mathtt{int}(v)$, but $[a, b] \not\subset \mathtt{int}(v_i)$ for all children v_i of v. Suppose $\mathtt{int}(v_j) \subset [a, b]$ for $j = r, r, +1, \ldots, q$. Then x-coordinates of all points in children $v_r, \ldots, v_q$ of v belong to $[a, b]$. We can report all points p in $v_r, \ldots, v_q$ whose y-coordinate do not exceed c using sorted lists of points in $S_{v_r}, \ldots, S_{v_q}$. We also answer two dominance queries $Q_1 = [a, +\infty) \times (-\infty, c]$ and $Q_2 = (-\infty, b] \times (-\infty, c]$ with help of data structures $\mathcal{R}_{v_{r-1}}$ and $\mathcal{L}_{v_{q+1}}$ respectively.

After a kinetic event affecting two points p and q from the same set S_v, the data structures $\mathcal{L}_v$ and $\mathcal{R}_v$ are updated. After a kinetic event that affects points p and q that belong to two neighbor sets S_{v_i} and $S_{v_{i+1}}$ respectively, we swap p and q: p is removed from S_{v_i} and inserted into $S_{v_{i+1}}$, and q is removed from $S_{v_{i+1}}$ and inserted into S_{v_i}. In this case a constant number of operations insert_x^- and delete_x^- (resp. insert_x^+ and delete_x^+) is performed. Each point belongs to $O(\log n)$ sets S_v. Hence, the space usage is $O(n \log n)$ and an update after a kinetic event takes $O(\log^2 n)$ time.

Lemma 3. *There exists a $O(n \log n)$ space data structure that supports three-sided reporting queries for a set of linearly moving points on a $U \times U$ grid in $O(\sqrt{\log U / \log \log U} + k)$ time and updates after kinetic events in $O(\log^2 n)$ time. If trajectories of the points do not change, then the total number of kinetic events is $O(n^2)$.*

Orthogonal Range Reporting Queries. In a similar way to Lemma 2 we can extend the data structure to support update operations in two other special cases. Let $y_{\min}$ and $y_{\max}$ be the smallest and the largest y-coordinates of points in S. The operation insert_y^+ inserts a point p with $y_{\max} < p.y$. The operation delete_y^+ deletes a point p with $p.y = y_{\max}$. Operations insert_y^- and delete_y^- insert and

delete a point whose y-coordinate is smaller than y-coordinates of all other points in S. Again, it is easy to modify the data structure of Lemma 3 so that it supports insert_y^- and delete_y^- because these operations affect at most one segment of $\mathcal{M}$. We can support insert_y^+ and delete_y^+ using the same construction as in Lemma 2.

A point p y-overlaps with a point q if $p.y > q.y$. Analogously to Lemma 2, points are stored in sets $V_2, \ldots, V_m$ and we maintain the invariants:

1. For $i = 2, \ldots, m-1$, V_i contains between 2^{2i-1} and 2^{2i+2} points; V_m contains between 2^{2i-2} and $2^{2i+2} + 2^{2i}$ points
2. Each point of V_i y-overlaps at most 2^{2i-4} points from V_{i-1}
3. At most 2^{2i-3} points from V_i y-overlap with V_{i-1}

Elements of V_i are stored in an augmented data structure E_i of Lemma 2 so that kinetic events and operations insert_x^-, delete_x^-, insert_x^+, delete_x^+, insert_y^- and delete_y^- are supported. For $j = 1, \ldots, m$, we maintain $\max_x(j) = \max\{p.x | p \in V_j\}$ and $\min_y(j) = \min\{p.y | p \in V_j\}$. All $\min_y(j)$ are stored in a kinetic binary tree. For each j, we store a point p_j with $p_j.x = \max_x(j)$ and $p_j.y = j$ in a kinetic data structure X; X supports dominance reporting queries in $O(\log \log n + k)$ time and updates in $O(\log n)$ time. For each j we also store all points of V_j in a list L_j' that contains all points from V_j in the descending order of their x-coordinates.

Given a query $Q = [a, +\infty) \times [b, +\infty)$, we can find in $O(\log \log n)$ time the smallest index j, such that at least one point in V_j has y-coordinate smaller than b. According to conditions 2 and 3 above, V_j and V_{j+1} may contain points whose y-coordinate are greater than or equal to b. The y-coordinates of all points in $V_2 \cup \ldots \cup V_{j-1}$ are greater than b. The y-coordinates of all points in sets V_i, $i > j+1$, are smaller than b.

Using X, we can identify all V_f such that $f < j$ and V_f contains at least one point p with $p.x \geq a$. For every such f, all points p such that $p \in H_f$ and $p.x \geq a$ can be reported by traversing the list L_f'. Hence, reporting all points $p \in V_f$ such that $f < j$, $p.x \geq a$ and $p.y \geq b$ takes $O(\log \log n + k)$ time. We can report all points in H_j and H_{j+1} that belong to $Q = [a, +\infty) \times [b, +\infty)$ in $O(\log \log U + k)$ time using data structures E_j and E_{j+1} respectively.

When a point p is inserted with an operation insert_x^+ or insert_x^-, we identify the data structure E_i, such that $p.y > \min_y(i)$ but $p.y < \min_y(j)$ for all $j < i$, and insert p into E_i as described in Lemma 2. When a point p is deleted with operations delete_x^+ or delete_x^-, we delete p from the data structure E_i such that $p \in H_i$. When a point is inserted or deleted with operations insert_y^- or delete_y^-, we insert or delete this point into the set V_m. If a point is inserted or deleted with operations insert_y^+ or delete_y^+, we rebuild the data structure E_2. Invariants 1-3 for sets V_i can be maintained under insert_y^+, delete_y^+, and kinetic events with the same method that was used in Lemma 2 to maintain sets H_i.

Hence, we obtain

Lemma 4. *There exists a data structure that supports dominance queries on $U \times U$ grid in $O(\sqrt{\log U / \log \log U} + k)$ time. Kinetic events and operations* insert_x^+, delete_x^+, insert_x^-, delete_x^-, insert_y^+, delete_y^+, insert_y^-, *and* delete_y^- *are supported in* $O(\log n)$ *time.*

We can obtain a $O(n \log n)$ space data structure for three-sided queries that supports operations insert_y^+, delete_y^+, insert_y^-, and delete_y^- as well as kinetic events in $O(\log^2 n)$ time in the same way as in the proof of Lemma 3. Using the same technique once again, we obtain the result for general two-dimensional range reporting queries stated in Theorem 2.

Theorem 2. *There exists a $O(n \log^2 n)$ space data structure that supports orthogonal range reporting queries on $U \times U$ grid in $O(\sqrt{\log U / \log \log U} + k)$ time and updates after kinetic events in $O(\log^3 n)$ time. If trajectories of the points do not change, then the total number of kinetic events is $O(n^2)$.*

6 Conclusion and Open Problems

In this paper we describe a data structure for a set of moving points that answers orthogonal range reporting queries in $O(\sqrt{\log U / \log \log U} + k)$ time. The query time is dominated by the time needed to answer a point location query; the rest of the query procedure takes $O(\log \log n + k)$ time. Thus a better algorithm for the point location problem would lead to a better query time of our data structure. Proving any $\Omega(\log \log U)$ lower bound for our problem, i.e., proving that kinetic range reporting is slower than static range reporting, would be very interesting.

While kinetic data structures usually support arbitrary changes of trajectories and only require that points follow constant-degree algebraic trajectories, we assume that points move linearly and all changes of trajectories are known in advance. An interesting open question is whether we can construct a kinetic data structure that achieves $o(\log n / \log \log n)$ query time without these additional assumptions.

References

1. M. A. Abam, M. de Berg, B. Speckmann *Kinetic kd-Trees and Longest-Side kd-Trees*, SIAM J. Comput. 39(4), 1219-1232, 2009.
2. P. Agarwal, L. Arge, J. Erickson, *Indexing Moving Points*, J. Comput. System Sci. 66, 207-243, 2003.
3. P. Agarwal, J. Gao, L. Guibas, *Kinetic Medians and kd-Trees*, Proc. ESA 2002, 5-16.
4. S. Alstrup, G. S. Brodal, T. Rauhe, *New Data Structures for Orthogonal Range Searching*, Proc. FOCS 2000, 198-207.
5. S. Alstrup, T. Husfeldt, T. Rauhe, *Marked Ancestor Problems*, Proc. FOCS 1998, 534-544.
6. L. Arge, J. S. Vitter, "Optimal External Memory Interval Management", SIAM J. on Computing 32(6), 1488-1508, 2003.
7. J. Basch, L. Guibas, J. Hershberger, *Data Structures for Mobile Data*, J. Algorithms 31, 1-28, 1999.
8. J. Basch, L. Guibas, L. Zhang, *Proximity Problems on Moving Points*, Proc. 13th ACM Symposium on Computational Geometry 1997, 344-351.
9. J. L. Bentley, *Decomposable Searching Problems*, Information Processing Letters 8(5), 244-251 (1979).
10. J. L. Bentley T. Ottmann, *Algorithms for Reporting and Counting Geometric Intersections*, IEEE Trans. Comput. 28, 643-647, 1979.
11. B. Chazelle, L. J. Guibas, *Fractional Cascading: I. A Data Structuring Technique*, Algorithmica 1(2), 133-162, 1986.
12. T. M. Chan, M. Patrascu, *Transdichotomous Results in Computational Geometry, I: Point Location in Sublogarithmic Time* , SIAM J. Comput. 39(2), 703-729, 2009.
13. P. F. Dietz, R. Raman, *Persistence, Amortization and Randomization*, Proc. SODA 1991, 78-88.
14. P. van Emde Boas, *Preserving Order in a Forest in Less Than Logarithmic Time and Linear Space*, Inf. Process. Lett. 6(3): 80-82 (1977).
15. L. J. Guibas, *Kinetic Data Structures: a State of the Art Report*, Proc. 3rd Workshop on the Algorithmic Foundations of Robotics 1998, 191-209.
16. M. H. Overmars, *Design of Dynamic Data Structures*, Springer-Verlag New York, Inc., Secaucus, NJ, 1987.
17. M. H. Overmars, *Efficient Data Structures for Range Searching on a Grid*, J. Algorithms 9(2), 254-275 (1988).
18. S. Subramanian, S. Ramaswamy, *The P-range Tree: A New Data Structure for Range Searching in Secondary Memory*, Proc. SODA 1995, 378-387.
19. D. E. Vengroff, J. S. Vitter, *Efficient 3-D Range Searching in External Memory*, Proc. STOC 1996, 192-201.