# Reverse Furthest Neighbors in Spatial Databases 

Bin Yao, Feifei Li, Piyush Kumar<br>Computer Science Department, Florida State University<br>Tallahassee, Florida, USA<br>\{yao, lifeifei, piyush\}@cs.fsu.edu


#### Abstract

Given a set of points $P$ and a query point $q$, the reverse furthest neighbor (RFN) query fetches the set of points $p \in P$ such that $q$ is their furthest neighbor among all points in $P \cup\{q\}$. This is the monochromatic RFN (MRFN) query. Another interesting version of RFN query is the bichromatic reverse furthest neighbor (BRFN) query. Given a set of points $P$, a query set $Q$ and a query point $q \in Q$, a BRFN query fetches the set of points $p \in P$ such that $q$ is the furthest neighbor of $p$ among all points in $Q$. The RFN query has many interesting applications in spatial databases and beyond. For instance, given a large residential database (as $P$ ) and a set of potential sites (as $Q$ ) for building a chemical plant complex, the construction site should be selected as the one that has the maximum number of reverse furthest neighbors. This is an instance of the BRFN query. This paper presents the challenges associated with such queries and proposes efficient, R-tree based algorithms for both monochromatic and bichromatic versions of the RFN queries. We analyze properties of the RFN query that differentiate it from the widely studied reverse nearest neighbor queries and enable the design of novel algorithms. Our approach takes advantage of the furthest Voronoi diagrams as well as the convex hulls of either the data set $P$ (in the MRFN case) or the query set $Q$ (in the Brfn case). For the Brfn queries, we also extend the analysis to the situation when $Q$ is large in size and becomes disk-resident. Experiments on both synthetic and real data sets confirm the efficiency and scalability of proposed algorithms over the bruteforce search based approach.


## I. Introduction

Spatial databases have offered a large number of applications in last decade that shape the horizon of computing services from people's daily life to scientific research. For example, people rely on online map services to plan their trips; the deployment and query processing in large sensor networks [15] often require the design of location-aware algorithms. Driven by these increasing number of applications, efficient processing of important and novel query types in spatial databases has always been a focal point.

In this work, we identify a query type that has wide applications and requires novel algorithms for processing. For a large data set $P$ and any random query point $q$, we are interested in retrieving the set of points in $P$ that take $q$ as their furthest neighbors comparing to all points in $P$, i.e., collecting $q$ 's reverse further neighbors (RFN). This problem is referred as the monochromatic reverse furthest neighbor (MRFN) queries. It naturally has a bichromatic version as well (BRFN). Specifically, the query contains a set of query points $Q$ and one point $q \in Q$. The goal in this case is to find the set of points $p \in P$ so that they all take $q$ as their furthest neighbors compared to all points in $Q$.


Fig. 1. An Mrfn query example.
Consider the data set $P$ in Figure 1, suppose the query point $q$ is $p_{1}$, then its RFN includes $\left\{p_{6}, p_{7}, p_{8}\right\}$. Using $p_{7}$ as an example, visually its potential furthest neighbor is among $p_{1}, p_{2}$ and $p_{3}$. The dotted lines indicate the distances of these candidates to $p_{7}$ and among them $p_{1}$ is the furthest. Hence, $p_{7}$ is included in the RFN of $p_{1}$. On the other hand, $p_{5}$ is not included as $p_{5}$ 's furthest neighbor becomes $p_{7}$. This is an instance of the MRFN query. Note that the query point $q$ could be different from any existing point in $P$.


Fig. 2. An Brfn query example.
The BRFN query takes a set of query points as the input. With the same data set $P$, Figure 2 demonstrates the BRFN query. The BRFN of $q_{3}$ w.r.t $Q$ in Figure 2 is $\left\{p_{1}, p_{2}, p_{4}\right\}$. Using $p_{2}$ as an illustration, its distance to $q_{1}, q_{2}, q_{3}$ are denoted by the dotted lines and $q_{3}$ is the furthest. Similarly, $q_{1}$ 's BRFN is $\left\{p_{3}, p_{5}, p_{6}, p_{7}, p_{8}\right\} . q_{2}$ 's BRFN in this case is empty. A careful examination will confirm that no point in $P$ takes $q_{2}$ as its furthest neighbor from points in $Q$.

The motivation to study the RFN queries is largely inspired by an important query type that has been extensively studied recently, namely, the reverse nearest neighbor queries (RNN) [29], [20], [34], [5], [1], [35], [28], [26], [23], [33], [22], [27]. Intuitively, an RNN query finds the set of points taking the query point as their nearest neighbors and it also exists in both the monochromatic and bichromatic versions. Many applications that are behind the studies of the RNN queries naturally have the corresponding "furthest" versions. Consider next two examples for the MRFN and BRFN queries:

Example 1 For a large collection of points of interest in a region, every point would like to learn the set of sites that
take itself as their furthest neighbors compared to other points of interest. This has an implication that visitors to those sites (i.e., its reverse furthest neighbors) are highly unlikely to visit this point. Ideally, it should put more efforts in advertising itself to those sites.

Example 2 Given a database of residential sites as $P$ and a collection of potential locations $Q$ for building a chemical plant, due to the hazardous materials that may be produced by such a plant, we should select the location in $Q$ that is further away (compared to other choices in $Q$ ) from as many residential sites as possible, i.e., the point $q \in Q$ that has the largest number of reverse furthest neighbors in $P$.

The above example does not limit itself to spatial data sets, as long as there is a similarity metric between any two objects, the BRFN queries are applicable.

Example 3 Let $P$ be a large set of customers and $Q$ be a set of business competitors offering similar products in the market, suppose for each customer $p$ and each competitor $q$ there is a distance measure reflecting the rating of $p$ to $q$ 's product. The smaller value indicates a higher preference. For any competitor in $Q$, an interesting query is to discover the customers that dislike his product the most among all competing products in the market. This company could then carry out specialized analysis on this group of customers to identify potential drawbacks in its marketing strategy.

In the last example, an interesting observation is that $Q$ could be large and becomes disk-resident data set as well. As a result, the algorithm for the BRFN problem needs to be extended to the case where both $P$ and $Q$ are in external memory.

To the best of our knowledge, both the MrFn and Brfn have not been studied in the literature. The brute-force search algorithms for these problems are obviously too expensive to be of any practical use. It is worth mentioning that the BRFN problem has been briefly examined from a theoretical point of view recently [10]. However, large scale spatial databases are calling for a practical, efficient algorithm for this problem. More importantly, by taking the furthest neighbors, the geometric nature of the RFN problems has been changed from the RnN problems. Hence, we need to design new algorithms to process the RFN queries more efficiently by taking the new geometric perspectives into account.

Our Contributions. This work presents efficient algorithms for the MRFN and BRFN problems. We identify the important insights for the RFN problems based on the convex hulls of either $P$ or $Q$ and extend the basic idea to work with dynamic data sets as well as disk resident query groups. The paper is organized as follows: 1) We formulate the problem of the reverse furthest neighbors in Section II and survey the related work in Section III; 2) We propose effective pruning techniques based on several geometric insights to obtain two efficient, R-tree based exact algorithms for the MRFN problem
in Section IV; In particular, we design R-tree based algorithm to maintain $P$ 's convex hull dynamically so that our query algorithms work efficiently with dynamic data sets; 3) We present a practical and efficient R-tree based algorithm for the Brfn problem in Section V; 4) We discuss the generalization to handle query groups with large size for the BRFN problem in Section V-A; 5) We report a comprehensive experimental study with both synthetic and real data sets in Section VI.

## II. Problem Formulation

Let $P$ denote the spatial database in $d$-dimensional $E u$ clidean space. Our techniques could be easily generalized for any metric space where the distance between points satisfies the triangle inequality. For simplicity, we concentrate on the Euclidean distance where the distance between any two points $p$ and $q$ is denoted by $\|p-q\|$. The furthest neighbor of any point $p$ w.r.t a set of points $P$ is simply defined as:

Definition 1 The furthest neighbor of $p$ to a data set $P$ is defined as $\operatorname{fn}(p, P)=p^{*}$ s.t. $p^{*} \in P$, for $\forall p^{\prime} \in P$ and $p^{\prime} \neq p^{*}$, $\left\|p^{*}-p\right\| \geq\left\|p^{\prime}-p\right\|$. Ties are broken arbitrarily.

The monochromatic RFN query is formally defined as:
Definition 2 The Mrfn of $q$ w.r.t the data set $P$ is a set of points from $P$ that take $q$ as their furthest neighbors comparing to all points in $P$, i.e., MrFN $(q, P)=\{p \mid p \in$ $P, \operatorname{fn}(p, P \bigcup\{q\})=q$.

The bichromatic RFN query takes additionally a set of query points $Q$ as the input, and is formally defined in the follows.

Definition 3 The Brfn of $q \in Q$ w.r.t the data set $P$ and the query set $Q$ is a set of points from $P$ that take $q$ as their furthest neighbors comparing to all other points in $Q$, i.e., $\operatorname{BRFN}(q, Q, P)=\{p \mid p \in P, \operatorname{fn}(p, Q)=q\}$.

Assume that $|P|=n,|Q|=m$ and a page size of $B$, the brute-force search based approach for the MRFN (BRFN) problem takes $\mathcal{O}\left(n^{2}\right)(\mathcal{O}(m n))$ time complexity with $\mathcal{O}\left(n^{2} / B\right)$ $(\mathcal{O}(n / B)$ if $Q$ fits in memory or otherwise $\mathcal{O}(m n / B))$ I/Os.

Other Notations. We summarize some of our notations here. The convex hull of a point set is defined as the smallest convex polygon containing all points. Intuitively, the convex hull is obtained by spanning an elastic band around the point set. The points touched by the elastic band become the vertices of the convex hull. We use $\mathcal{C}_{P}$ to denote both the ordered set of vertices and the convex polygon (the area enclosed by the elastic band) for the convex hull of $P$. Its meaning is clear from the context. An ordered list of points $p_{1} p_{2} \cdots p_{t} p_{1}$ represents a convex polygon defined by line segments $\left\{p_{1} p_{2}\right.$, $\left.p_{2} p_{3}, \ldots, p_{t} p_{1}\right\}$ with $t$ vertices. The concepts of convex hull and convex polygon work with any dimension $d$. We also used the furthest Voronoi diagram for a set of points $(\mathcal{F D}(P))$ and the furthest Voronoi cell for a single point w.r.t a set of points $(\mathrm{fvc}(p, P))$. The details will be defined when the related discussion emerges.

| Symbol | Description |
| :--- | :--- |
| $\\|p-q\\|$ | Euclidean distance between $p$ and $q$ |
| $\|\cdot\|$ | Size of a set |
| $(\subset) \subseteq$ | Both (strict) set and geometric containment |
| $\mathcal{C}_{P}\left(\mathcal{C}_{Q}\right)$ | Ordered set of vertices of $P$ 's $(Q$ 's) convex hull, |
|  | also is the convex polygon of $P$ 's $(Q$ 's) convex hull |
| $\operatorname{fn}(p, P)$ | The furthest neighbor of $p$ in $P$ |
| fvc $(p, P)$ | The furthest Voronoi cell of $p$ w.r.t $P$ |
| $\mathcal{F} \mathcal{D}(P), \mathcal{F D}(Q)$ | The furthest Voronoi diagram of $P(Q)$ |
| $n, m$ | $\|P\|$ and $\|Q\|$ respectively |
| $p_{1} p_{2}$ | A line segment between $p_{1}$ and $p_{2}$ |
| $p_{1} p_{2} \cdots p_{t} p_{1}$ | The convex polygon by $\left\{p_{1} p_{2}, \ldots, p_{t} p_{1}\right\}$ |

TABLE I
Notation USED.
When $A \subseteq B$ is used with geometric objects $A$ and $B$, it implies that all points from the area enclosed by $A$ are part of the points from the area enclosed by $B$. When $\subset$ is used in this context, it further requires all points from the area enclosed by $A$ are inside (while not touching) the boundary edges of $B$. Table I provides a quick reference to the main notations.

In this work we focuses on the two dimensional space. The main ideas developed could be generalized to higher dimensions and the details of such generalization is an interesting problem we will look at.

## III. Background and Related Work

R -tree [17] and its variants ( $\mathrm{R}^{*}$-tree as the representative [4]) have been the most widely deployed indexing structure for the spatial database, or data in multi-dimensions in general. Intuitively, R-tree is an extension of the $\mathrm{B}^{+}$-tree to higher dimensions. Points are grouped into minimum bounding rectangles (MBRs) which are recursively grouped into MBRs in higher levels of the tree. The grouping is based on data locality and bounded by the page size. An example of the R-tree is illustrated in Figure 3. R-tree has facilitated processing many important query types for spatial databases and data in multidimensional space. Two important classes of queries that are related to this work and have been extensively studied in the literature are nearest neighbor ( NN ) queries and range queries.

NN search has been thoroughly studied for the Euclidean space [24], [30], [19], [32], [21], [6], [14], [25], [18], [12]. In the Euclidean space, R-tree demonstrates efficient algorithms using either the depth-first [25] or the best-first [18] approach. The main idea behind these algorithms is to utilize branch and bound pruning techniques based on the relative distances between a query point $q$ to a given MBR $N$. Such distances include the mindist, the minmaxdist and the maxdist. The mindist measures the minimum possible distance for a point $q$ to any point in an MBR $N$; the minmaxdist measures the lower bound on the maximum distance for a point $q$ to any point in an MBR $N$; and finally, the maxdist simply measures the maximum possible distance between $q$ and any point in an MBR $N$. These distances are easy to compute arithmetically given $q$ and $N$. An example for these distances has been provided in Figure 3. The principle for utilizing them to prune the search space for NN search in R-tree is straightforward, e.g., when an MBR $N_{a}$ 's mindist to $q$ is larger than another MBR $N_{b}$ 's minmaxdist, we can safely prune $N_{a}$ entirely.


Fig. 3. The R-tree.

Another classical query type is the range query where the goal is to return all points that are fully contained by the query range (often as a rectangle). R-tree yields good performance for answering range queries [8]. The basic idea is to search all MBRs that intersect with the query range.

An interesting query type that has close relationship with NN search was defined in [22], in which the goal is to find the set of points from $P$ that take the query point $q$ as their nearest neighbors among all points in the data set $P$. This is the monochromatic reverse nearest neighbor query (monochromatic RNN). Due to its wide applications, RnN queries have received considerable attention since its appearance. In the monochromatic version [28], [26], [33], [22], [29], [27], the state of the art is the TPL method from [29]. TPL recursively filters the data by finding perpendicular bisectors between the query point and its nearest object and taking the half plane that is closer to the query point.

The bichromatic RNN also finds many applications [20], [1], [35], [26], [22], [27]. In this case, the query takes a set of query points $Q$ and a query point $q \in Q$. The set of points returned from $P$ all take $q$ as their nearest neighbors w.r.t other points in $Q$. The basic idea here is to use the Voronoi diagram and find the region that corresponds to the query point.

Most of the work for the RNN queries focused on the euclidean space. Many interesting variants have been studied. The RnN problem in graphs and road network was studied in [35]. Generalization to any metric space appeared in [1]. Continuous RnN was explored by [20], [31]. The RnN for moving objects was studied in [5]. Reverse $k N N$ search was examined by [29], [28]. Finally, the RnN for Ad-Hoc subspaces was solved by [34].

Our work explores the unique geometric property of the reverse furthest neighbors and proposes efficient, novel pruning techniques that are suitable in this case. In particular, we utilize the convex hull of either the data set in the case of MRFN or the query set in the case of BRFN for pruning. Hence, it is related with the RNN problem but based on significant different insights.

## IV. Monochromatic Reverse Furthest Neighbors

We search for efficient, R-tree based algorithms for the Mrfn queries in this section. Our approaches are inspired by the furthest Voronoi diagram and the convex hull.

## A. PFC: the Progressive Furthest Cell algorithm

For a set of points $P$, the furthest Voronoi diagram of $P$, denoted as $\mathcal{F} \mathcal{D}(P)$, is similar to the well-known Voronoi

(a) Furthest Voronoi diagram.

(b) Computing fvc $\left(p_{1}, \mathrm{P}\right)$.

Fig. 4. $\mathcal{F} \mathcal{D}(P)$ and its derivation.
diagram [3] of $P$ except that the space is partitioned with the furthest Voronoi cells instead of the nearest Voronoi cells.

Definition 4 For a space S , a point $p$ and a set of points $P(p \in P$ and $P \subseteq \mathcal{S})$, a convex polygon $\operatorname{fvc}(p, P) \subseteq \mathcal{S}$ is the furthest Voronoi cell of $p$ iff: $\forall$ point $p^{\prime} \subseteq \operatorname{fvc}(p, P)$, $\operatorname{fn}\left(p^{\prime}, P\right)=p$, i.e., any point inside $\operatorname{fvc}(p, P)$ will take $p$ as its furthest neighbor in $P . p^{\prime}$ is not necessarily a point in $P$. The furthest Voronoi diagram of $P(\mathcal{F} \mathcal{D}(P))$ is simply the collection of furthest Voronoi cells for all points in $P$.

An example of the furthest Voronoi diagram on an input space $\mathcal{S}:\left(x_{\ell}, y_{\ell}\right) \times\left(x_{h}, y_{h}\right)$ is illustrated in Figure 4(a) where $\operatorname{fvc}\left(p_{i}, P\right)$ is marked as region $\operatorname{fvc}\left(p_{i}\right)$. Computing the $\mathcal{F} \mathcal{D}(P)$ can be done by modifying the algorithm for the Voronoi diagram. To find the $\operatorname{fvc}\left(p_{i}, P\right)$, for each point $p_{j} \in P$ and $j \neq i$, the bisector line of the line segment $p_{i} p_{j}$ separates the space into two polygons and denote the one that does not contain $p_{i}$ as poly ${ }_{i j}^{j}$. Then $\operatorname{fvc}\left(p_{i}, P\right)$ is simply the intersection of poly ${ }_{i j}^{j}$ 's for all $j$ 's in $P$. The fact that poly $_{i j}^{j}$ is the polygon that is further away from $p_{i}$ among the two polygons (covering the whole space) partitioned by the bisector line of $p_{i} p_{j}$ straightforwardly yields the correctness of this algorithm. Figure 4(b) illustrates this idea for computing fvc $\left(p_{1}, P\right)$. One first finds poly ${ }_{12}^{2}$ with point $p_{2}$, followed by identifying poly ${ }_{13}^{3}$. The $\operatorname{fvc}\left(p_{1}, P\right)$ is simply the poly ${ }_{12}^{2} \bigcap$ poly $_{13}^{3}$.

An important property for the $\mathcal{F} \mathcal{D}(P)$ is that any point from the furthest Voronoi cell of $p_{i}$ takes $p_{i}$ as its furthest neighbor among all points in $P$. Formally,

Lemma $1 \forall p^{\prime} \subseteq f v c(p, P)$, $f n\left(p^{\prime}, P\right)=p$. In addition, for a random point $p^{\prime}$ in space $\mathcal{S}$, if $f n\left(p^{\prime}, P\right)=p$, then $p^{\prime} \subseteq$ $f v c(p, P)$.

Proof: It is immediate by the construction of $\mathcal{F} \mathcal{D}(P)$.
By Lemma 1 , for a query point $q$, if we can compute the fvc $(q, P \bigcup\{q\})$, then the RFN of $q$ is simply those points in $P$ that are contained by the $\operatorname{fvc}(q, P \bigcup\{q\})$. The challenge is how to compute $\operatorname{fvc}(q, P \bigcup\{q\})$ efficiently when $P$ is large and stored on the disk. Obviously, a very expensive approach is to linear scan points in $P$ and apply the algorithm above, then perform a range query with the discovered furthest Voronoi cell of $q$. Next, we present a progressive version of this idea, using the help of the R-tree, where the $\operatorname{fvc}(q, P \bigcup\{q\})$ is computed incrementally and the points in $P$ that are contained


Fig. 5. PFC algorithm.
by the $\operatorname{fvc}(q, P \bigcup\{q\})$ are also obtained progressively in one pass.

The progressive algorithm. The detail of the progressive furthest cell (PFC) algorithm is listed in Algorithm 1. Given the query point $q$ and the R-tree index for the data set $P$, we maintain a priority queue $\mathcal{L}$ for storing MBRs and points. Objects in $\mathcal{L}$ are sorted in the decreasing order by their minmaxdist to $q$. The intuition is that the points that are further away from $q$ have higher chances being the RFN of $q$, hence we should consider these candidates with higher priority. Meanwhile, $\operatorname{fvc}(q, P \bigcup\{q\})$ is initialized to the whole space $\mathcal{S}$ and we also maintain two vectors. $V_{c}$ stores all points that are potentially the RFN of $q$ and $V_{p}$ stores all objects that may disqualify false positive candidates from $V_{c}$ in the postprocessing step. In the sequel, when the context is clear we simply use fvc $(q)$ to denote $q$ 's up-to-date furthest Voronoi cell.

At each iteration, the head entry $e$ of $\mathcal{L}$ is popped. If $e$ is a point, we find the bisector line for the line segment $q e$. As shown in Figure 4(b), the bisector line cuts the space $\mathcal{S}$ into two polygons and we identify the one that is further away from $q$, i.e., poly $_{q e}^{e}$. The $\operatorname{fvc}(q)$ is updated as the intersection of the current $\operatorname{fvc}(q)$ with poly ${ }_{q e}^{e}$. Next, we determine whether $e$ is a potential candidate. If $e \subseteq \operatorname{fvc}(q)$, then $e$ is pushed into $V_{c}$. Otherwise, it is discarded.

If $e$ is an MBR, we take the current furthest Voronoi cell of $q$ and calculate $e \bigcap \operatorname{fvc}(q)$. If this is empty we push $e$ into the pruning vector $V_{p}$, otherwise we retrieve each child of $e$ and insert into $\mathcal{L}$ the ones that do intersect with $\operatorname{fvc}(q)$. The rest are pushed into $V_{p}$.

There are two insights for the PFC algorithm. The $\mathrm{fvc}(q)$ is updated progressively and it always shrinks after each update. Another critical observation (following the first observation) is that if an MBR or a point does not intersect with the current $\operatorname{fvc}(q)$, it has no chance to be contained by the final $\operatorname{fvc}(q)$. Hence, we can safely claim that an entry $e$ is not a potential candidate as long as its intersection with the current $\operatorname{fvc}(q)$ is empty. However, $e$ cannot be disregarded if it is an MBR. The reason is that $e$ may contain points that could update (shrink) $\operatorname{fvc}(q)$ and disqualify those potential candidates identified in previous steps. As a result, a post-processing is required when $\mathcal{L}$ becomes empty, where entries in $V_{p}$ are checked in sequence and points contained by those entries will be used to update $\mathrm{fvc}(q)$. Finally, all points in $V_{c}$ are filtered once more using the final $\operatorname{fvc}(q)$.

```
Algorithm 1: PFC(Query \(q\); R-tree \(T\) )
    Initialize two empty vectors \(V_{C}\) and \(V_{p}\);
    Initialize the priority queue \(\mathcal{L}\) with \(T\) 's root node;
    \(\mathcal{L}\) orders entries in decreasing order of the minmaxdist;
    Initialize \(\operatorname{fvc}(q)=\mathcal{S}\);
    while \(\mathcal{L}\) is not empty do
        Pop the head entry \(e\) of \(\mathcal{L}\);
        if \(e\) is a point then
            Identify poly \({ }_{q e}^{e}\) and set \(\mathrm{fvc}(q)=\mathrm{fvc}(q) \bigcap\) poly \(_{q e}^{e}\);
            if \(f v c(q)=\emptyset\) then return;
            if \(e \subseteq f v c(q)\) then Push \(e\) into \(V_{c}\);
        else
            if \(e \bigcap f v c(q)\) is \(\emptyset\) then push \(e\) to \(V_{p}\);
            else
                Let \(u_{1}, \ldots, u_{f}\) be children MBRs of node \(e\);
                for \(i=1, \ldots, f\) do
                        if \(u_{i} \cap f v c(q) \neq \emptyset\) then
                        Insert the child node \(u_{i}\) into \(\mathcal{L}\);
                    else Insert the child node \(u_{i}\) into \(V_{p}\);
    Update fvc \((q)\) using points contained by entries in \(V_{p}\);
    Filter points in \(V_{c}\) using fvc \((q)\);
    Output \(V_{c}\); return;
```

Consider the example in Figure 5, assuming $p_{2}$ is first processed and it updates $\mathrm{fvc}(q)$ as shown in Figure 5(a). Since $p_{2}$ is contained by $\mathrm{fvc}(q)$, it is added to $V_{c}$. Next, the MBR $e$ is processed, obviously $e$ could not offer any possible RFN of $q$ as $e \bigcap \mathrm{fvc}(q)=\emptyset$. However, $e$ may contain points that update $\mathrm{fvc}(q)$ s.t. existing candidates are disqualified. Figure $5(\mathrm{~b}))$ shows that $p_{3}$ from $e$ could update $\operatorname{fvc}(q)$ and prune $p_{2}$. Hence, we need to add $e$ to $V_{p}$ for post-processing step.

PFC algorithm could terminate as soon as the fvc $(q)$ becomes empty (Line 9 in Algorithm 1). In practice, many points due to their geometric locations will not have a furthest Voronoi cell in the space (this phenomena will be explained in details in the next Section). By searching the space in the decreasing minmaxdist fashion, for these points we could quickly shrink their furthest Voronoi cells to empty, leading to efficient early termination. In addition, when the query point does have a furthest Voronoi cell, the PFC algorithm calculates the RFN in one pass.

## B. CHFC: the Convex Hull Furthest Cell algorithm

The PFC algorithm may scan a large number of points. Handling the false positives from the list of potential candidates in the post-processing step is expensive. An important lemma introduced next will significantly reduce the cost of searching for the RFN.

The convex hull $\mathcal{C}_{P}$ for a set of points $P$ is the smallest convex polygon defined by points in $P$ that fully contains $P$. We represent $\mathcal{C}_{P}$ as an ordered set of points that are vertices for the convex hull of $P$. For example, in Figure 6(a) $\mathcal{C}_{P}=$ $\left\{p_{1}, p_{2}, p_{3}, p_{4}, p_{5}\right\}$. We abuse the notation of $\mathcal{C}_{P}$ a little bit


Fig. 6. Furthest point to $p$ : Proof of Lemma 2
whenever the context is clear. It also represents the constrained area corresponding to $P$ 's convex hull. So $\mathcal{C}_{P}$ also denotes the shaded polygon $p_{1} p_{2} p_{3} p_{4} p_{5} p_{1}$ in Figure 6(a). Given a point $q$, we conjecture that the furthest point from $P$ to $q$ must be found in $\mathcal{C}_{P}$. This is indeed the case.

Lemma 2 Given $P$ and its convex hull $\mathcal{C}_{P}$, for a point $q$, let $p^{*}=f n(q, P)$, then $p^{*} \in \mathcal{C}_{P}$.

Proof: Suppose $\mathcal{C}_{P}=\left\{p_{1}, \ldots, p_{s}\right\}$, then the boundary edges for $\mathcal{C}_{P}$ are $E\left(\mathcal{C}_{P}\right)=\left\{p_{1} p_{2}, \ldots, p_{s} p_{1}\right\}$. First, we argue that $p^{*}$ must locate on one of the boundary edges. Assume this is not correct, as shown in Figure 6(a), then $p^{*} \subset \mathcal{C}_{P}$. For both $q \subseteq \mathcal{C}_{P}$ (e.g. $q_{1}$ ) and $q \nsubseteq \mathcal{C}_{P}$ (e.g. $q_{2}$ ), the line that extends $q p^{*}$ from $p^{*}$ must intersect with an edge from $E\left(\mathcal{C}_{P}\right)$ as $\mathcal{C}_{P}$ is a convex. In our example, it intersect with $p_{4} p_{5}$ on a point $p_{e}$, then obviously $q p_{e} \geq q p^{*}$. Next, we can show that the furthest point $p^{*}$ to $q$ from a line segment $p_{\ell} p_{r}$ must be either $p_{\ell}$ or $p_{r}$. With the help of the perpendicular line from $q$ to $p_{\ell} p_{r}$, as illustrated in Figure 6(b), this could be argued using proofs by contradiction again. This concludes that $q^{*}$ must belong to $\mathcal{C}_{P}$.

Using Lemma 2 in a reverse angle, we can derive that only points in $\mathcal{C}_{P}$ will have reverse furthest neighbors w.r.t the data set $P$. In other words:

Lemma 3 For a data set $P$ and a point $p \in P$, $f v c(p, P)=\emptyset$ if $p \notin \mathcal{C}_{P}$. Furthermore, any point $p^{\prime} \subset \mathcal{C}_{P}$ (here $\mathcal{C}_{p}$ denotes the area enclosed by the convex hull), fvc $\left(p^{\prime}, P \bigcup\left\{p^{\prime}\right\}\right)=\emptyset$.

Proof: The first claim is a special case of the second claim, as all points from $P-\mathcal{C}_{P}$ are strictly (not on the edges) contained by the polygon area defined by $\mathcal{C}_{P}$. Hence, consider a random point $p^{\prime} \subset \mathcal{C}_{P}$, if $p^{\prime} \notin P$, update $P$ to $P^{\prime}=P \bigcup\left\{p^{\prime}\right\}$. By the property of the convex hull, it is straightforward to show that $\mathcal{C}_{P^{\prime}}=\mathcal{C}_{P}$ as $p^{\prime}$ is strictly inside $\mathcal{C}_{P}$ (see Lemma 4). This indicates that $p^{\prime} \notin \mathcal{C}_{P^{\prime}}$. Now, for any random query point $q$ in the space $\mathcal{S}$, by Lemma 2, $\operatorname{fn}\left(q, P^{\prime}\right) \neq p^{\prime}$ as $p^{\prime} \notin \mathcal{C}_{P}^{\prime}$. This indicates that there is not even a single point in the space will take $p^{\prime}$ as its furthest neighbor w.r.t $P^{\prime}$. By the definition of the furthest Voronoi diagram, $\operatorname{fvc}\left(p^{\prime}, P^{\prime}\right)=\emptyset$.

Combining this result with Lemma 1, another way to interpret Lemma 3 is that only the points from $\mathcal{C}_{P}$ will have the reverse furthest neighbors for a data set $P$. Given a query point $q$, we could view it as an insertion to $P$ and obtain a new data set $P^{*}=P \bigcup\{q\}$. It is critical to decide whether


Fig. 7. CHFC algorithm.
$q$ belongs to the set of vertices defining the convex hull of $P^{*}$. A good news is that $\mathcal{C}_{P^{*}}$ could be computed efficiently from $\mathcal{C}_{P}$ and $q$ alone, without looking at the rest of points in $P-\mathcal{C}_{P}$.

Lemma 4 For a data set $P$ and its convex hull $\mathcal{C}_{P}$, after adding a point $q$ to $P$, if $q$ is strictly contained by $\mathcal{C}_{P}\left(q \subset \mathcal{C}_{P}\right)$, then $\mathcal{C}_{P \cup\{q\}}=\mathcal{C}_{P}$; otherwise, $\mathcal{C}_{P \cup\{q\}}=\mathcal{C}_{\mathcal{C}_{P} \cup\{q\}}$.

Proof: This straightforwardly follows from the definition of the convex hull. Figure 7(a) gives its intuition.

Finally, it remains as a problem how to compute $q$ 's furthest Voronoi cell efficiently when $q$ is indeed a vertex for the convex hull of $P \bigcup\{q\}$. An important result is stated next that $\operatorname{fvc}(q, P)$ could be computed using only $\mathcal{C}_{P}$.

Lemma 5 For a data set $P$ and its convex hull $\mathcal{C}_{P}$, for a point $p \in \mathcal{C}_{P}, f v c(p, P)=f v c\left(p, \mathcal{C}_{P}\right)$.

Proof: We use proof by contradiction. Suppose this claim is not correct. Then there must be a point $p^{\prime} \in P-\mathcal{C}_{P}$ s.t. its furthest Voronoi cell could shrink the area enclosed by $\operatorname{fvc}\left(p, \mathcal{C}_{P}\right)$. However, by Lemma $3, \operatorname{fvc}\left(p^{\prime}, P=\emptyset\right)$. This completes the proof.

1) The Main Algorithm: Lemma 3, 4 and 5 (together with Lemma 1) immediately yield an efficient algorithm for the MRFN problem. Assuming for now that we have obtained $\mathcal{C}_{P}$ (details will be discussed soon), given a query point $q$, we can quickly return empty if $q \subset \mathcal{C}_{P}$. Otherwise, we set $P^{*}=$ $P \bigcup\{q\}$ and compute $\mathcal{C}_{P^{*}}$ using only $\mathcal{C}_{P}$ and $q$. Next, the furthest Voronoi cell of $q$ in $P^{*}$ is calculated using only $\mathcal{C}_{P^{*}}$. Finally, a range query using $\operatorname{fvc}\left(q, P^{*}\right)$ is executed on the Rtree of $P$ to retrieve the RFN of $q$. This is referred as the convex hull furthest cell (CHFC) algorithm.

Figure 7(a) demonstrates CHFC's idea. If the query point is $q_{1}$, since $q_{1} \subset \mathcal{C}_{P}$, CHFC algorithm will immediately return empty for $q_{1}$ 's RFN. Lemma 4 guarantees that $\mathcal{C}_{P \cup\left\{q_{1}\right\}}=\mathcal{C}_{P}$. Then, by Lemma $3 \operatorname{fvc}\left(q_{1}, P \bigcup\left\{q_{1}\right\}\right)=\emptyset$. Hence, by Lemma 1 , there will have no points taking $q_{1}$ as their furthest neighbors among all points in $P$. On the other hand, if the query point is $q_{2}$ where $q_{2} \not \subset \mathcal{C}_{P}$, CHFC algorithm will obtain $P \bigcup\left\{q_{2}\right\}$ 's convex hull by applying Lemma 4 . Essentially, only points from $\mathcal{C}_{P}$ and $q_{2}$ (i.e., $\left\{p_{1}, p_{2}, p_{3}, p_{4}, p_{5}, q_{2}\right\}$ ) will be used to compute the convex hull of $P \bigcup\left\{q_{2}\right\}$. The updated convex hull is simply $p_{1} p_{2} q_{2} p_{3} p_{4} p_{5} p_{1}$. Next, since $q_{2}$ is one of the points in this convex hull, we need to compute its furthest Voronoi cell and Lemma 5 achieves this using only points

```
Algorithm 2: CHFC(Query \(q\); R-tree \(T\) )
1 Compute \(\mathcal{C}_{P}\) with \(T\) using either the distance-priority or
    the depth-first algorithm [9];
    if \(q \subset \mathcal{C}_{P}\) then return \(\emptyset\);
    else
        Compute \(\mathcal{C}_{P^{*}}\) using \(\mathcal{C}_{P} \bigcup\{q\} ;\)
        Set \(\operatorname{fvc}\left(q, P^{*}\right)\) equal to \(\operatorname{fvc}\left(q, \mathcal{C}_{P^{*}}\right)\);
        Execute a range query using \(\operatorname{fvc}\left(q, P^{*}\right)\) on \(T\);
```

from $\mathcal{C}_{P \cup\left\{q_{2}\right\}}$. Finally, $q_{2}$ 's RFN includes all points in the area covered by its furthest Voronoi cell as Lemma 1 states. In this case, its RFN $p_{9}$ will be successfully retrieved.

The efficiency of the CHFC algorithm is achieved in two folds. First of all, all query points that are enclosed by $\mathcal{C}_{P}$ are extremely fast to deal with. Secondly, for the rest of query points, computing their furthest Voronoi cells becomes much more efficient by taking into account only the points in $\mathcal{C}_{P}$ (instead of doing this via $P$ ). In practice, $\left|\mathcal{C}_{P}\right| \ll|P|$. Of course, this reduction in size may not always be the case. It is fairly easy to construct examples where $P=\mathcal{C}_{P}$, e.g., $P$ only contains points on edges of a rectangle. Nevertheless, for most real life data sets we expect that $\left|\mathcal{C}_{P}\right| \ll|P|$ (or at least $\left.\left|\mathcal{C}_{P}\right|<|P|\right)$ holds and we could store $\mathcal{C}_{P}$ in main memory. This fact is verified in our experiments. For the case when this is violated, i.e., $\mathcal{C}_{P}$ is still too large to be stored in main memory, one can use the PFC algorithm or use the approximate technique we will discuss in Section V-A.
2) Computing P's convex hull: There are I/O efficient algorithms for computing the convex hulls of disk-based data sets. Specifically, convex hulls in two dimension can be computed in external memory with the help of sorting in external memory [16]. This means that we can find the convex hull of the disk-based data set $P(|P|=n)$ with $\mathcal{O}\left(\frac{n}{B} \log _{M} \frac{n}{M B}\right)$ I/Os for a main memory buffer with $M$ pages, assuming the page size is $B$. From a practical point of view and to deal with data in higher dimensions, one would like to compute the convex hull of $P$ using a R-tree. Fortunately, this problem has been studied in the literature and one can apply either the distance-priority or the depth-first algorithm from [9]. The complete CHFC algorithm is presented in Algorithm 2.

If $P$ is static, an obvious optimization for the CHFC algorithm is to pre-compute $\mathcal{C}_{P}$ and avoid the call to the distancepriority (or depth-first) algorithm at Line 1 completely. As argued above, we assume $\mathcal{C}_{P}$ can fit into main memory in this case, hence, any main memory convex hull algorithm could be applied in Line 4 . In fact, one could apply main memory dynamic convex hull algorithm [13] here to get an almost logarithm computation bound.
3) Dynamically Maintaining $\mathcal{C}_{P}$ : When $P$ is static, CHFC algorithm is extremely efficient after pre-computing $\mathcal{C}_{P}$ once. However, under dynamic insertion and deletion to $P$, CHFC has to compute $\mathcal{C}_{P}$ for every query and this could greatly degrade the query performance. If $P$ fits in main memory, recent theoretical study has confirmed that dynamically maintaining


Fig. 8. Dynamic update $\mathcal{C}_{P}$ using R-tree.
the convex hull of $P$ could be done almost in logarithm time [13]. The challenge is how to do this efficiently utilizing R-tree when $P$ is disk-resident.

We distinguish two cases, namely, point insertion and point deletion. Point insertion is easy to handle as suggested by the Lemma 4 , assuming $\mathcal{C}_{P}$ fits in memory. The main obstacle is when point is deleted as shown in Figure 7(b). Once again, the easy case is when the deleted point $p$ is inside the convex hull of $P\left(p \subset \mathcal{C}_{P}\right)$, e.g., $p_{7}$ in Figure 7(b). Obviously, in this situation $\mathcal{C}_{P}$ will not be affected by the deletion. On the other hand, when the deleted point $p$ is an existing vertex for the convex hull of $P\left(p \in \mathcal{C}_{P}\right)$, updating $\mathcal{C}_{P}$ is necessary. For example, deleting $p_{2}$ in Figure 7(b) will eventually lead to adding $p_{10}$ to $\mathcal{C}_{P}$.

When $p \in \mathcal{C}_{P}$ is deleted, we observe that most part of $\mathcal{C}_{P}$ is intact except for its left and right neighbors (denoted as $p_{\ell}$ and $p_{r}$ ). New vertices may be introduced in between $p_{\ell}$ and $p_{r}$. A straightforward solution is to retrieve all points, using R-tree, contained by the polygon $p_{\ell} p p_{r} p_{\ell}\left(p_{1} p_{2} p_{3} p_{1}\right.$ in Figure 7 (b)) and find the local convex hull for these points in main memory. But this could be potentially expensive (if $p_{\ell} p p_{r} p_{\ell}$ encloses large number of points).

A better approach is to adapt the qhull algorithm [7] to Rtree in this setting. In the qhull algorithm, it is shown that given two vertices $p_{\ell}$ and $p_{r}$ in the convex hull, a vertex (if it indeed exists) in the final convex hull that locates in between $p_{\ell}$ and $p_{r}$ could be located by finding the point that has the largest perpendicular distance to $p_{\ell} p_{r}$. This is true for both sides of $p_{\ell} p_{r}$. Recursively applying this step, all vertices in the final convex hull that locate in between $p_{\ell}$ and $p_{r}$ could be identified.

In our case, we only need to search for points on one side of the $p_{\ell} p_{r}$, i.e., the side that the deleted point $p$ locates at. An example demonstrating this idea is shown in Figure 8(a). When $p_{2}$ is deleted, we find $p_{\ell}=p_{1}$ and $p_{r}=p_{3}$ from $\mathcal{C}_{P}$. Searching along the side that $p_{2}$ locates at, the point with the largest perpendicular distance to $p_{1} p_{3}$ is $p_{7}$. It is added to $\mathcal{C}_{P}$. Next, we recursively apply the same step to $p_{1} p_{7}$ and $p_{7} p_{3}$, still only to the side that $p_{2}$ locates at. In this case, $p_{8}$ will be retrieved and one more recursion using $p_{1} p_{8}$ and $p_{8} p_{7}$ will terminate the search. The updated convex hull becomes $\mathcal{C}_{P}=\left\{p_{1}, p_{8}, p_{7}, p_{3}, p_{4}, p_{5}\right\}$. The only puzzle left is to doing these steps efficiently in R-tree.

Using an R-tree, our problem essentially reduce to the following. Given a line segment $p_{\ell} p_{r}$, find the point from the R-tree that has the largest perpendicular distance to $p_{\ell} p_{r}$ and it

```
Algorithm 3: \(\operatorname{MaxVP}\left(\right.\) Query \(p_{\ell}, p_{r}, p\); R-tree \(\left.T\right)\)
    Initialize the priority queue \(\mathcal{L}\) with \(T\) 's root node;
    \(\mathcal{L}\) orders entries in decreasing order of their maxVdist;
    \(\mathcal{L}\). minVdist keeps the max minVdist among all its entries;
    while \(\mathcal{L}\) is not empty do
        Pop the head entry \(e\) of \(\mathcal{L}\);
        if \(e\) is point then return \(e\);
        else
            for each child \(u_{i}\) of \(e\) do
                \(\operatorname{minVdist} u_{i}=\infty, \operatorname{maxVdist}{ }_{u_{i}}=-1\);
                for each corner point \(x\) of \(u_{i}\) do
                    if \(x\) and \(p\) are on the same side of \(p_{\ell} p_{r}\)
                    then
                            Get the perpendicular distance \(s_{x}\)
                from \(x\) to \(p_{\ell} p_{r}\);
                                    \(s_{x}\) updates minVdist \(_{u_{i}}, \operatorname{maxV} \operatorname{dist}_{u_{i}}\);
                if maxVdist \(_{u_{i}}>\mathcal{L}\).minVdist then
                            Insert \(u_{i}\) into \(\mathcal{L}\);
                            if minVdist \(_{u_{i}}>\mathcal{L}\). minVdist then
                    \(\mathcal{L} . \operatorname{minVdist}=\operatorname{minV}\) dist \(_{u_{i}} ;\)
```

must locate on one specified side of $p_{\ell} p_{r}$. This further reduces to the next problem, given a line segment $p_{\ell} p_{r}$, and an $\mathrm{MBR} e$, what is the maximum and minimum of the maximum possible perpendicular distances from a point in the MBR $e$ to the line segment $p_{\ell} p_{r}$ ?

Lemma 6 For an axis-parallel MBR e, both the maximum and minimum of the maximum possible perpendicular distances from a point in the $M B R$ e to any line segment $p_{\ell} p_{r}$ are bounded by the corner points of $e$.

Proof: W.l.o.g. assume that $p_{\ell} p_{r}$ and $e$ are in the positive quadrant of the plane. Let $\hat{c}$ be the unit vector perpendicular to $p_{\ell} p_{r}$ and point to the side of $p_{\ell} p_{r}$ we are interested in. We need to find $x^{*}=\arg \max _{x \in e} \hat{c}^{T} x$. Note that this is a linear program with four linear constraints, is feasible and bounded. Hence by the fundamental theorem of Linear Programming, one of the vertices of $e$ must yield the optimal solution [11]. The minimization problem can be handled similarly. A similar problem has also been studied by [9].

With Lemma 6, it is possible to migrate the aforementioned algorithm for the point deletion to an R-tree efficiently. We use minVdist and maxVdist to denote the minimum and maximum of the maximum possible perpendicular distance for a point from an MBR $e$ to a line segment $p_{\ell} p_{r}$. A special note in our setting is that minVdist and maxVdist are only bounded by the corner points of $e$ that locate on the same side of $p_{\ell} p_{r}$ as the deleted point $p$. Consider the example in Figure 8(b), suppose $p_{2} \in \mathcal{C}_{P}$ has been deleted. Its left and right neighbors in $\mathcal{C}_{P}$ are $p_{1}$ and $p_{3}$ respectively. First, we can safely prune all MBRs that locate completely on the other side of $p_{1} p_{3}$, e.g., $e_{3}$. Next, for the rest of MBRs, we calculate their minVdist

```
Algorithm 4: DynamicCH \(\left(\mathcal{C}_{P}\right.\), op, \(p\); R-tree \(\left.T\right)\)
    if \(o p\) is Insertion then
        if \(p \subset \mathcal{C}_{P}\) then return \(\mathcal{C}_{P}\);
        else return \(\mathcal{C}_{\mathcal{C}_{P} \cup\{p\}}\);
    else if \(o p\) is Deletion then
        if \(p \subset \mathcal{C}_{P}\) then return \(\mathcal{C}_{P}\);
        else
            Find \(p\) 's left (right) neighbor, \(p_{\ell}\left(p_{r}\right)\), in \(\mathcal{C}_{P}\);
            \(\mathcal{C}_{P}:\left\{\ldots, p_{\ell}, p, p_{r}, \ldots\right\} \longmapsto\left\{\ldots, p_{\ell}, p_{r}, \ldots\right\} ;\)
            Call QHullRtree \(\left(p_{\ell}, p_{r}, p, \mathcal{C}_{P}, T\right)\);
    /* qhull algorithm adapted to R-tree */
    QHullRtree ( \(p_{\ell}, p_{r}, p, \mathcal{C}_{P} ;\) R-tree \(T\) )
    Let \(p^{\prime}=\operatorname{MaxVP}\left(p_{\ell}, p_{r}, p, T\right)\);
    if \(p^{\prime}=\emptyset\) then return;
    else
        \(\mathcal{C}_{P}:\left\{\ldots, p_{\ell}, p_{r}, \ldots\right\} \longmapsto\left\{\ldots, p_{\ell}, p^{\prime}, p_{r}, \ldots\right\} ;\)
        QHullRtree \(\left(p_{\ell}, p^{\prime}, p, \mathcal{C}_{P}, T\right)\);
        QHullRtree ( \(p^{\prime}, p_{r}, p, \mathcal{C}_{P}, T\) );
```

and maxVdist using their corner points that locate on the same side of $p_{1} p_{3}$ as $p_{2}$. Note that we can safely prune an MBR if its maxVdist is smaller than the minVdist of some other MBRs, e.g., $e_{1}$ can be pruned in the presence of $e_{2}$ as the best point from $e_{1}$ cannot possibly beat at least one point from $e_{2}$ (in terms of the perpendicular distance to the query line segment). We exploit the R-tree nodes in a priority queue that orders entries by the decreasing maxVdist value. The search could terminate as soon as we find one point at the head of the queue or when the queue becomes empty. The Max-VerticalPoint (MaxVP) algorithm is listed in details in Algorithm 3.

With these discussions, we could dynamically maintain $\mathcal{C}_{P}$ using the R-tree. This reduces the cost of the CHFC algorithm by avoiding the call on Line 1 in Algorithm 2. The DynamicCH algorithm in Algorithm 4 summarizes our idea presented above. We would like to point out that Algorithm 4 is presented for the two-dimensional case. For higher dimensions, certain generalization must be adapted. For example, instead of looking at one side for a segment from the convex hull, we need to examine one side of a plane.

## V. Bichromatic Reverse Furthest Neighbors

After resolving all the difficulties for the MrFN problem in Section IV, solving the BRFN problem becomes almost immediate. From the discussion in Section IV, all points in $P$ that are contained by $\operatorname{fvc}\left(q_{i}, Q\right)$ will have $q_{i}$ as their furthest neighbor. This immediately implies that Brfn $(q, Q, P)=$ $\{p \in P \wedge p \in \operatorname{fvc}(q, Q)\}$. Furthermore, Lemma 5 guarantees that $\operatorname{fvc}(q, Q)=\operatorname{fvc}\left(q, \mathcal{C}_{Q}\right)$. For example, in Figure 9 the reverse furthest neighbor of $q_{1}$ is $\left\{p_{4}\right\}$ which is contained in the $\operatorname{fvc}\left(q_{1}, Q\right)=\operatorname{fvc}\left(q_{1}, \mathcal{C}_{Q}\right)$. Lemma 3 indicates that only those points in $\mathcal{C}_{Q}$ will have reverse furthest neighbors. Hence, an efficient R-tree based query algorithm for finding the RFN of

```
Algorithm 5: BRFN(Query \(q, Q\); R-tree \(T\) )
    Compute the convex hull \(\mathcal{C}_{Q}\) of \(Q\);
    if \(q \subset \mathcal{C}_{Q}\) then return \(\emptyset ; / / q \notin \mathcal{C}_{Q}\)
    else
        Compute fvc \(\left(q, \mathcal{C}_{Q}\right)\);
        Execute a range query using \(\operatorname{fvc}\left(q, \mathcal{C}_{Q}\right)\) on \(T\);
```



Fig. 9. The BRFN algorithm.
$q$ is to only check those MBRs that intersect with $\operatorname{fvc}\left(q, \mathcal{C}_{Q}\right)$ when $q \in \mathcal{C}_{Q}$, e.g., in Figure 9 we need to access the MBR $e$ that intersects with $\operatorname{fvc}\left(q_{1}\right)$ to retrieve its reverse furthest neighbors. If $q=q_{4}$, we can return $\emptyset$ immediately. Algorithm 5 details this method.

## A. Disk-Resident Query Group

One limitation with our discussions so far is the problem of handling the query group with a massive size that do not fit in internal memory. As we have discussed, one could use the convex hull of $Q$ to reduce the number of points in the query group. There are I/O efficient algorithms for computing the convex hulls of disk-based data sets. In two dimension $\mathcal{C}_{Q}$ could be found with $\mathcal{O}\left(\frac{m}{B} \log _{M} \frac{m}{M B}\right)$ I/Os [16] for a main memory buffer with $M$ pages, assuming the page size is $B$ and $|Q|=m$. Alternatively, if $Q$ is indexed by R-tree, we could simply use the algorithm from [9]. For most cases, we expect that $\left|\mathcal{C}_{Q}\right| \ll|Q|$. However, one could easily construct special cases where $\left|\mathcal{C}_{Q}\right|=|Q|$, e.g., all points in $Q$ are vertices of a convex polygon. To handle such special instances, we propose to obtain an approximate convex hull of $Q$ using Dudley's approximation [36]. Dudley's construction generates an approximate convex hull of $Q$ (denote it as $\mathcal{A C}_{Q}$ ) with $\mathcal{O}\left(1 / \epsilon^{(d-1) / 2}\right)$ vertices with maximum Hausdorff distance of $\epsilon$ to the convex hull of $Q$. The Hausdorff distance measures how far two convex polygons $S_{1}$ and $S_{2}$ are from each other. Informally, the Hausdorff distance between $S_{1}$ and $S_{2}$, is the longest distance an adversary can force one to travel by choosing a point in one of the two sets, from where one then must travel to the other set. Formally, let $X$ and $Y$ be the vertices of $S_{1}$ and $S_{2}$ respectively, then:

$$
d_{H}\left(S_{1}, S_{2}\right)=\max \left(\sup _{x \in X} \inf _{y \in Y}\|x-y\|, \sup _{y \in Y} \inf _{x \in X}\|x-y\|\right)
$$

For fixed dimensions, this means that we can always get an approximate convex hull of $Q$ with constant number of vertices. Obviously, the smaller $\epsilon$ is, the more accurate the approximation is and the larger the size of the approximate
convex hull is (since Dudley's approximation is an inner approximation of the convex hull).

Roughly speaking, the edges of of $\mathcal{A C}_{Q}$ are within $\epsilon$ distances from edges of $\mathcal{C}_{Q}$. Clearly, there is a trade-off between the approximation quality and the size of $\mathcal{A C}_{Q}$. Henceforth, for a disk-based query group $Q$, we first compute its convex hull using the I/O efficient algorithm. If $\mathcal{C}_{Q}$ is still too large to fit in main memory, we replace $\mathcal{C}_{Q}$ with Dudley's approximation and specify the size of $\mathcal{A C}_{Q}$ of our choice.

The Dudley's approximation was proposed to work with main memory data sets [36]. For our purpose, we need to extend it to work with external data sets. To compute $\mathcal{A C}_{Q}$ in external memory, one can build an index that supports nearest neighbor search with logarithm IOs guarantees (such as the BBD tree[2]) on $\mathcal{O}\left(1 / \epsilon^{(d-1) / 2}\right)$ points scattered on the sphere containing the convex hull of $Q$ and then do a scan of the convex hull and compute the reverse nearest neighbor of each point read from $\mathcal{C}_{Q}$. This approach requires only linear I/Os on $\mathcal{C}_{Q}$ (excluding the I/Os to compute $\mathcal{C}_{Q}$ ). The following lemma summarizes this result. The details of this approach will appear in the full version.

Lemma 7 For a disk-based query group $Q$, one could always get a query group $Q^{\prime}$ with size $\mathcal{O}\left(1 / \epsilon^{(d-1) / 2}\right)$ using $\mathcal{O}\left(\frac{m}{B} \log _{M} \frac{m}{M B}\right)$ I/Os, where $m=|Q|, B$ is the page size and $M$ is the number of pages in the main memory, s.t. the Hausdorff distance between edges in $\mathcal{C}_{Q^{\prime}}$ and $\mathcal{C}_{Q}$ is at most $\epsilon$.

The introduction of Dudley's approximation only creates an additive error $\epsilon$ on the corresponding edges between $\mathcal{C}_{Q}$ and $\mathcal{A C}_{Q}$. Hence, for a query point $q$ and query group $Q$, if $\mathcal{C}_{Q}$ is too large to fit in main memory, we calculate $\mathcal{A C}_{Q}$. If $q$ is inside the convex polygon defined by $\mathcal{A C}_{Q}$ and away from any edge of this polygon by more than $\epsilon$, then it does not have reverse furthest neighbors. Otherwise, we find the nearest neighbor of $q$ in the vertices of $\mathcal{A C}_{Q}$, say $q^{\prime}$, and apply algorithm 5 using $q^{\prime}$ and $\mathcal{A C}_{Q}$. This algorithm is referred as the A-BRFN algorithm and our experiment confirms that it has high accuracy.

As a final note, the approximate convex hull idea could be also applied to the CHFC algorithm for the MRFN problem in Section IV-B when $\mathcal{C}_{P}$ is too large to fit in main memory.

## VI. EXPERIMENT

All proposed algorithms have been implemented into the widely used, disk-based spatial index library ${ }^{1}$, Standard geometric operations, e.g. convex polygon intersection, are provided by the CGAL library ${ }^{2}$. Finally, our I/O efficient approximate convex hull algorithm is developed based on the library from [36]. All experiments were executed on a Linux machine with an Intel 2 GHz cpu. For both R-tree and heapfiles the page size $B$ is set to 4 KB .

Data sets. The real data sets were obtained from the digital chart of the world server where points define the road

[^0]

Fig. 10. Different distribution types of synthetic data sets and query groups.


Fig. 11. MRFN algorithms: cpu computation and IOs analysis.
networks from California (CA), San Francisco (SF) and USA (US). CA also contains large number of points of interest (e.g, restaurants, resorts). These data sets are available online ${ }^{3}$. To create a data set of larger size, we merge them into one data set (denoted as Map) after normalizing each data set into the space $L=(0,0) \times(100000,100000)$. Points in CA, SF and US have various degrees of skew distributions in the space. Map contains 476, 578 number of points. There are three kinds of synthetic data sets used in the experiment: uncorrelated uniformly generated (UN) points, correlated bivariate (CB) points and random-cluster distribution (R-Cluster) points in $L$. An example of the CB and R-cluster data sets is illustrated in Figure 10. In this work we concentrate on the two-dimensional space. How to generalize our algorithms to higher dimensions as well as the experiments for those are interesting extensions to this work.

Performance measurement. For all algorithms, we measured their performance using two metrics, namely, the CPU time and the number of IOs. We would like to highlight that the CPU time measure the pure computation cost of an algorithm. It is NOT the total execution time of the algorithm. The total execution time is simply the CPU time plus the IO cost. In most cases, the IO cost dominates the overall execution time. Finally, by default 1,000 queries were generated for each experiment and we report the average for one query.

## A. Algorithms for the MrFn problem

Experiment Setup. For the MRFN problem, we generated synthetic data sets whose size are equal to the size of Map data set. The query point is randomly selected from the space $L$. For brevity, only the results from the UN, R-Cluster and Map data sets are reported.

[^1]

Fig. 12. DynamicCH algorithms: deleting vertex of $\mathcal{C}_{P}$.

The PFC and CHFC algorithms. We first study the case when $P$ is static. In this setting, we can optimize the CHFC algorithm by pre-computing $P$ 's convex hull using R-tree and store $\mathcal{C}_{P}$ for query processing. Figure 11 compares the cost of these two algorithms against the brute-force search (BFS) based approach. To make the comparison fair, we optimize the BFS algorithm by terminating the search as early as possible. For a point $p \in P$, we first calculate its distance to $q$ and then calculate its distance to every other point $p^{\prime}$ in $P$. However, whenever there is a $p^{\prime}$ s.t. $\left|p-p^{\prime}\right|>|p-q|$, we know for sure that $p \notin \operatorname{MrFN}(q, P \bigcup\{q\})$ and the BFS algorithm continues to the next point in $P$. Clearly, Figure 11 indicates that both PFC and CHFC outperform the BFS algorithm in terms of both the CPU cost and the number of IOs. Especially, for the dominant cost of IOs, both algorithms are at least two orders of magnitude better.

Among the three data sets, BFS algorithm has almost constant computation cost as well as the number of IOs. For both the PFC and CHFC algorithms, R-Cluster data set results in higher costs. This is easily explained by the more scattered geometric locations of points in the R-Cluster case. When $|P|$ is static and $\mathcal{C}_{P}$ is known, it is easy to expect that CHFC outperforms PFC, as for a large number of query points CHFC could be extremely efficient. This is confirmed by Figure 11.

Dynamically maintaining $\mathcal{C}_{P}$. One advantage of the PFC algorithm is that it requires no additional effort to work with dynamic data set. On the other hand, when insertion and deletion of points happen in $P$, CHFC algorithm has to dynamically maintain $\mathcal{C}_{P}$ using Algorithm 4. Following our discussions, insertion of points is easy to handle and the worst case is the deletion of a point $p$ that happens to be one of the vertices in $\mathcal{C}_{P}$. Section IV-B. 3 discussed this problem in details and we executed an experiment to verify the efficiency of the DynamicCH algorithm in this worst case. For a data set $P$, we first computed its $\mathcal{C}_{P}$ using the R-tree based algorithm from [9]. Then, for each point $p$ in $\mathcal{C}_{p}$, we deleted $p$ and computed the new convex hull for $P-\{p\}$. Note that, for each such point $p$ we always started with the original $\mathcal{C}_{P}$ and $P$. The average cost of maintaining the convex hull after deleting one point in $\mathcal{C}_{P}$, in terms of both the cpu computation and number of IOs, of our DynamicCH algorithm is shown in Figure 12. They essentially reflect the cost of QHullRtree algorithm based on the MaxVP algorithm when we randomly


Fig. 13. Pruning with the convex hull of $Q$.
delete one point from $\mathcal{C}_{P}$. We repeated this experiment for data sets of different sizes as well as different distributions. Figure 12 clearly shows that our approach is highly effective, with cpu computation cost at around 0.2 milliseconds and IOs lower than 10 . More importantly, our algorithm scales nicely with the increase in size for the data set $P$. Uniform distributed data set generates higher maintenance cost as the chance of containing more points in the polygon area $p_{\ell} p p_{r} p_{\ell}$ is higher.

Given these results, we can safely claim that CHFC is extremely efficient in practice. Of course, there definitely exists data sets such that dynamically maintaining the convex hull could be expensive $\left(p_{\ell} p p_{r} p_{\ell}\right.$ contains large number of points), or $\mathcal{C}_{P}$ does not fit into main memory. We could use PFC algorithm in those cases.

## B. The BRFN algorithm

Experiment Setup. For the Brfn problem, the cost of the query depends on several critical factors. They include the location of $Q$, the range (its area $\mathcal{A}$ ) of $Q$ and how points are distributed within the range. Lastly, the size of $Q$ also plays a role. Given these observations, a random query is generated as follows. We specify $|Q|$ and its $\mathcal{A}$ (as a percentage of the area for the entire space $L$ ). Next a random location in $L$ is selected as the center of the query group. Three types of distributions discussed above were used to generate $|Q|$ number of points within the specified area. Examples of query groups with these distributions are given in Figure 10. In addition, we randomly select a query point $q \in Q$ after generating $Q$. We distinguish two cases: $q$ is randomly selected from $Q$ ( $R_{Q}$ strategy) or $q$ is randomly selected from $\mathcal{C}_{Q}\left(R_{\mathcal{C}_{Q}}\right.$ strategy).

Pruning power of $\mathcal{C}_{Q}$. Figure 13 shows that for various query distributions, $\mathcal{C}_{Q}$ could significantly reduce the size of the query group. We plot the averages together with the $5 \%-95 \%$ confidence interval. For $Q=1000,\left|\mathcal{C}_{Q}\right|$ is only about $2 \%$ of $|Q|$. Furthermore, the pruning power of $\mathcal{C}_{Q}$ increases (see Figure 13(b)) for larger $|Q|$ 's as $\left|\mathcal{C}_{Q}\right|$ grows at a much slower pace than $|Q|$ does. Lastly, Figure $13(a)$ indicates that $\frac{\left|\mathcal{C}_{Q}\right|}{|Q|}$ is roughly a constant over the query area and UN distributed query groups have larger $\left|\mathcal{C}_{Q}\right|$ over correlated and clustered groups. In most of our experiments, the results from the CB and R-Cluster query groups are quite similar. Hence, we only report the results from UN and CB query groups.

Algorithm BRFN. The BRFN algorithm (Algorithm 5) takes advantage of the convex hull of $Q$ and converts the BRFN


Fig. 14. BRFN algorithms: cpu computation and IOs analysis, Map data set.
problem into a simple range query with a convex polygon. Figure 14 confirms its superior performance over the bruteforce search (BFS) algorithm. The BFS algorithm in the BRFN case simply calculates the furthest neighbor among all points in $Q$ for each point $p \in P$. If it is equal to $q$, then $p$ is included in the Brfn $(q, Q, P)$. Note that in Figure 14 all the $y$-axes are plotted in the $\log$ scale. The BrFN algorithm is in general 2 to 4 orders of magnitude more efficient than the naive approach and the gap in performance is enlarging as $|Q|$ increases (Figure 14(c) and 14(d)). The cost of the BRFN algorithm does increase w.r.t the increase in the area of the query group (Figure 14(a) and 14(b)), but it is not very sensitive to such increases as shown in the log scale. For the BRFN algorithm, if $q$ is a random point from $Q-\mathcal{C}_{Q}$, obviously no range query is required on the R-tree and its query cost is tiny. Given the fact that majority of points in $Q$ will not be a member of $\mathcal{C}_{Q}$ (as shown by the convex hull pruning experiments from this section), the BRFN algorithm will be highly efficient for the case of $q \in Q$ as reflected by Figure 14. The $U N$ type of query groups reduces the cost of the BrFn algorithm compared to the $C B$ type in Figure 14, in terms of both cpu computation and IOs. The furthest cells of query points are distributed more evenly in the $U N$ query groups due to the uniformity. This indicates smaller furthest cells than those from the $C B$ query groups. Hence, BrFn is more efficient on the $U N$ type of query groups.

Finally, we observe an interesting pattern for the IOs cost of the BRFN algorithm for query groups of different sizes. It is in fact decreasing while $|Q|$ increases (Figure 14(d)). Recall that $\left|\mathcal{C}_{Q}\right|$ increases much slower than $|Q|$ does. Hence, for the $R_{Q}$ strategy, the chance that $q \in \mathcal{C}_{Q}$ is getting smaller. This will reduce the overall cost. On the other hand, for the $R_{\mathcal{C}_{Q}}$ strategy, as $\left|\mathcal{C}_{Q}\right|$ still contains more points when $|Q|$ increases, the furthest cell of an individual point $q \in \mathcal{C}_{Q}$ is smaller which leads to the reduction in the query cost.

## C. Query groups of large size

Experiment Setup. Concentrating on the BrFn problem, we use the same setup as in Section VI-B, but with larger query groups. The algorithm proposed in Section V-A for the Brfn query has much lower cost than the brute-force search and very good approximation qualities. For brevity those results were omitted and we focused on reporting its query efficiency.
Approximate convex hull based algorithm. There are cases


Fig. 15. Approximate convex hulls: $\mathcal{A}=3 \%$, Map data set.
where not only $|Q|$ is large, but also $\left|\mathcal{C}_{Q}\right|$. These cases rarely happen in practice. But it is not entirely impossible. Efforts have been devoted in Section V-A to deal with such scenarios when they do arise. The idea is to develop IO efficient approximate convex hulls $\left(\mathcal{A C}_{Q}\right)$. Figure 15 reports the performance of such an approach. The algorithm we have developed could be applied on the $\mathcal{A C}_{Q}$ instead of the $\mathcal{C}_{Q}$ and we denote it as A-BrFN. Our algorithm outperforms the basic brute-force search by orders of magnitude. By treating $Q$ as external memory data sets, computing its approximate convex hull incurs a cost that is similar to the sorting in external memory. Hence, we see an increase in terms of both cpu computation and IOs for our algorithm when $|Q|$ increases.

From the discussion in Section V-A, the error introduced by the approximation is independent of $|Q|$ and is only determined by $\left|\mathcal{A C}_{Q}\right|$. When $\left|\mathcal{A C}_{Q}\right|=20, \epsilon$ is 0.01 in two dimension. Hence, our algorithm achieves excellent approximation qualities.

## D. Scalability of various algorithms

Experiment Setup. Finally, we investigate the scalability of our algorithms w.r.t the size of the data set $P$. We use the same setup as in Section VI-A and VI-B, but with R-Cluster data sets of different sizes as $P$.

Scalability. For the BRFN problem, we fix the query area as $\mathcal{A}=3 \%$ and $|Q|=1000$. The $C B$ query type is used. Figure 16 shows the scalability of various algorithms. More specifically, Figure 16(a) and 16(b) show the results for the MRFN queries. Figure 16(c) and 16(d) show the results for the Brfn queries. Note that in both Figure 16(b) and 16(d), the number of IOs were plotted in the log scale. Compared to the brute-force search, all of our algorithms scale nicely with the increase in size of $|P|$. When $|P|=2,000,000$, BFS takes


Fig. 16. Scalability: cpu computation and IOs, R-Cluster data set.
more than 10, 000 IOs whereas our algorithms take from only a few to less than 100 IOs.

## VII. CONCLUSION

This paper studies the reverse furthest neighbor queries that has many real life applications. Our work solves the RFN queries in both monochromatic and bichromatic versions. We present R-tree based, efficient algorithms for both MrFN and BrFN problems with excellent pruning capability. All of our algorithms allow dynamic updates to the data sets. Furthermore, it has been adapted to work with disk-resident query groups in the BRFN case. Future work includes generalizing our algorithms to higher dimensions, dealing with moving points and continuous queries, and answering RFN queries in a road-network or any Ad-Hoc subspaces.

## VIII. Acknowledgment

We would like to thank the anonymous reviewers for their valuable comments that improve this work. Bin Yao is supported in part by CRC Planning Grant from FSU. Piyush Kumar is supported in part by NSF grant CCF-0643593.

## REFERENCES

[1] E. Achtert, C. Böhm, P. Kroger, P. Kunath, A. Pryakhin, and M. Renz, "Efficient reverse k-nearest neighbor search in arbitrary metric spaces," in SIGMOD, 2006.
[2] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu, "An optimal algorithm for approximate nearest neighbor searching in fixed dimensions," Journal of ACM, vol. 45, no. 6, pp. 891-923, 1998.
[3] F. Aurenhammer, "Voronoi diagrams - a survey of a fundamental geometric data structure," ACM Computing Survey, vol. 23, no. 3, pp. 345-405, 1991.
[4] N. Beckmann, H. P. Kriegel, R. Schneider, and B. Seeger, "The R*tree: an efficient and robust access method for points and rectangles," in SIGMOD, 1990.
[5] R. Benetis, C. S. Jensen, G. Karciauskas, and S. Saltenis, "Nearest and reverse nearest neighbor queries for moving objects," The VLDB Journal, vol. 15, no. 3, pp. 229-249, 2006.
[6] S. Berchtold, C. Böhm, D. A. Keim, and H. P. Kriegel, "A cost model for nearest neighbor search in high-dimensional data space," in $P O D S$, 1997.
[7] M. Berg, M. Kreveld, M. Overmars, and O. Schwarzkopf, Computational geometry: algorithms and applications. Springer, 1997.
[8] C. Böhm, "A cost model for query processing in high dimensional data spaces," ACM Transaction on Database Systems, vol. 25, no. 2, pp. 129-178, 2000.
[9] C. Böhm and H.-P. Kriegel, "Determining the convex hull in large multidimensional databases," in DaWaK, 2001.
[10] S. Cabello, J. M. Diaz-Banez, S. Langerman, C. Seara, and I. Ventura, "Reverse facility location problems," European Journal of Operational Research, vol. to appear, 2008.
[11] Chvatal and Vasek, Linear Programming, ser. A Series of Books in the Mathematical Sciences. New York: W. H. Freeman and Company, 1983.
[12] B. Cui, B. C. Ooi, J. Su, and K.-L. Tan, "Contorting high dimensional data for efficient main memory kNN processing," in SIGMOD, 2003.
[13] E. D. Demaine and M. Patrascu, "Tight bounds for dynamic convex hull queries (again)," in $S o C G, 2007$.
[14] K. Deng, X. Zhou, H. T. Shen, K. Xu, and X. Lin, "Surface k-NN query processing," in ICDE, 2006.
[15] S. P. Fekete and A. Kroller, "Geometry-based reasoning for a large sensor network," in SoCG, 2006.
[16] M. T. Goodrich, J.-J. Tsay, D. E. Vengroff, and J. S. Vitter, "Externalmemory computational geometry," in FOCS, 1993.
[17] A. Guttman, "R-trees: a dynamic index structure for spatial searching," in SIGMOD, 1984.
[18] G. R. Hjaltason and H. Samet, "Distance browsing in spatial databases," ACM Trans. Database Syst., vol. 24, no. 2, 1999.
[19] H. V. Jagadish, B. C. Ooi, K.-L. Tan, C. Yu, and R. Zhang, "iDistance: An adaptive $\mathrm{B}^{+}$-tree based indexing method for nearest neighbor search," ACM Trans. Database Syst., vol. 30, no. 2, pp. 364-397, 2005.
[20] J. M. Kang, M. F. Mokbel, S. Shekhar, T. Xia, and D. Zhang, "Continuous evaluation of monochromatic and bichromatic reverse nearest neighbors," in ICDE, 2007.
[21] M. R. Kolahdouzan and C. Shahabi, "Voronoi-based k nearest neighbor search for spatial network databases," in VLDB, 2004.
[22] F. Korn and S. Muthukrishnan, "Influence sets based on reverse nearest neighbor queries," in SIGMOD, 2000.
[23] F. Korn, S. Muthukrishnan, and D. Srivastava, "Reverse nearest neighbor aggregates over data streams," in $V L D B, 2002$.
[24] K. Mouratidis, M. L. Yiu, D. Papadias, and N. Mamoulis, "Continuous nearest neighbor monitoring in road networks," in $V L D B, 2006$.
[25] N. Roussopoulos, S. Kelley, and F. Vincent, "Nearest neighbor queries," in SIGMOD, 1995.
[26] A. Singh, H. Ferhatosmanoglu, and A. Tosun, "High dimensional reverse nearest neighbor queries," in CIKM, 2003.
[27] I. Stanoi, M. Riedewald, D. Agrawal, and A. E. Abbadi, "Discovery of influence sets in frequently updated databases," in $V L D B, 2001$.
[28] Y. Tao, D. Papadias, and X. Lian, "Reverse kNN search in arbitrary dimensionality," in $V L D B, 2004$.
[29] Y. Tao, D. Papadias, X. Lian, and X. Xiao, "Multidimensional reverse kNN search," The VLDB Journal, vol. 16, no. 3, pp. 293-316, 2007.
[30] Y. Tao, D. Papadias, and Q. Shen, "Continuous nearest neighbor search," in $V L D B, 2002$.
[31] T. Xia and D. Zhang, "Continuous reverse nearest neighbor," in ICDE, 2006.
[32] X. Xiong, M. F. Mokbel, and W. G. Aref, "SEA-CNN: scalable processing of continuous k-nearest neighbor queries in spatio-temporal databases," in ICDE, 2005.
[33] C. Yang and K.-I. Lin, "An index structure for efficient reverse nearest neighbor queries," in $I C D E, 2001$.
[34] M. L. Yiu and N. Mamoulis, "Reverse nearest neighbors search in ad hoc subspaces," IEEE Transactions on Knowledge and Data Engineering, vol. 19, no. 3, pp. 412-426, 2007.
[35] M. L. Yiu, D. Papadias, N. Mamoulis, and Y. Tao, "Reverse nearest neighbors in large graphs," IEEE Transactions on Knowledge and Data Engineering, vol. 18, no. 4, pp. 540-553, 2006.
[36] H. Yu, P. K. Agarwal, R. Poreddy, and K. R. Varadarajan, "Practical methods for shape fitting and kinetic data structures using core sets," in SoCG, 2004.


[^0]:    ${ }^{1}$ www.research.att.com/~marioh/spatialindex/index.html
    ${ }^{2}$ www.cgal.org

[^1]:    ${ }^{3}$ www.cs.fsu.edu/~lifeifei/SpatialDataset.htm

