Efficient Processing of Top-k Dominating Queries on Multi-Dimensional Data*

Man Lung Yiu
Department of Computer Science
Aalborg University
DK-9220 Aalborg, Denmark
mly@cs.aau.dk

Nikos Mamoulis
Department of Computer Science
University of Hong Kong
Pokfulam Road, Hong Kong
nikos@cs.hku.hk

ABSTRACT

The top-k dominating query returns k data objects which dominate the highest number of objects in a dataset. This query is an important tool for decision support since it provides data analysts an intuitive way for finding significant objects. In addition, it combines the advantages of top-k and skyline queries without sharing their disadvantages: (i) the output size can be controlled, (ii) no ranking functions need to be specified by users, and (iii) the result is independent of the scales at different dimensions. Despite their importance, top-k dominating queries have not received adequate attention from the research community. In this paper, we design specialized algorithms that apply on indexed multi-dimensional data and fully exploit the characteristics of the problem. Experiments on synthetic datasets demonstrate that our algorithms significantly outperform a previous skyline-based approach, while our results on real datasets show the meaningfulness of top-k dominating queries.

1 Introduction

Consider a dataset $\mathcal{D}$ of points in a d-dimensional space $\mathcal{R}^d$. Given a (monotone) ranking function $F : \mathcal{R}^d \rightarrow \mathcal{R}$, a *top-$k$ query* [14, 9] returns k points with the smallest F value. For example, Figure 1 shows a set of hotels modeled by points in the 2D space, where the dimensions correspond to (preference) attribute values; traveling time to a conference venue and room price. For the ranking function $F = x + y$, the top-2 hotels are p_4 and p_6. An obvious advantage of the top-k query is that the user is able to control the number of results (through the parameter k). On the other hand, it might not always be easy for the user to specify an appropriate ranking function. In addition, there is no straightforward way for a data analyst to identify the most important objects using top-k queries, since different functions may infer different rankings.

Besides, a *skyline query* [2] retrieves all points which are not dominated by any other point. Assuming that smaller values are preferable to larger at all dimensions, a point p *dominates* another point p' (i.e., $p \succ p'$) when

$$(\exists\, i \in [1, d],\ p[i] < p'[i]) \wedge (\forall\, i \in [1, d],\ p[i] \leq p'[i]) \qquad (1)$$

where $p[i]$ denotes the coordinate of p in the i-th dimension. Continuing with the example in Figure 1, the skyline query returns points p_1, p_4, p_6, and p_7. [2] showed that the skyline contains the top-1 result for any monotone ranking function; therefore, it can be used by decision makers to identify potentially important objects to some database users. A key advantage of the skyline query is that it does not require the use of a specific ranking function; its results only depend on the intrinsic characteristics of the data. Furthermore, the skyline is not affected by potentially different scales at different dimensions (monetary unit or time unit in the example of Figure 1); only the order of the dimensional projections of the objects is important. On the other hand, the size of the skyline cannot be controlled by the user and it can be as large as the data size in the worst case. As a result, the user may be overwhelmed as she may have to examine numerous skyline points manually in order to identify the ones that will eventually be regarded as important.

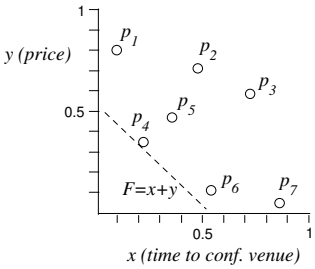

Figure 1: Features of hotels

From an analyst's point of view, an intuitive score function for modeling the importance of a point $p \in \mathcal{D}$ could be:

$$\mu(p) = |\ \{\ p' \in \mathcal{D} \mid p \succ p'\ \}\ | \qquad (2)$$

In words, the *score* $\mu(p)$ is the number of points dominated by point p. The following property holds for μ:

$$\forall\, p, p' \in \mathcal{D},\ p \succ p' \Rightarrow \mu(p) > \mu(p') \qquad (3)$$

Therefore, we can define a natural ordering of the points in the database, based on the μ function. Accordingly, the *top-k dominating* query returns k points in $\mathcal{D}$ with the highest score. For example, the top-2 dominating query on the data of Figure 1 retrieves p_4 (with $\mu(p_4) = 3$) and p_5 (with $\mu(p_5) = 2$). This result may indicate to an analyst the most popular hotels to the conference participants (considering price and traveling time as selection factors). Normally, a participant will try to book at p_4 and, if this hotel is fully-booked, try the next one (p_5). From this example, we can already see that a top-k dominating query is a powerful decision support tool, since it identifies the most significant objects in an intuitive way. From a practical perspective, top-k dominating queries combine the advantages of top-k queries and skyline

*Research supported by grant HKU 7160/05E from Hong Kong RGC.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.
VLDB '07, September 23-28, 2007, Vienna, Austria.
Copyright 2007 VLDB Endowment, ACM 978-1-59593-649-3/07/09.

queries without sharing their disadvantages. The number of results can be controlled without specifying any ranking function. In addition, data normalization is not required; the results are not affected by different scales or data distributions at different dimensions.

We are the first to recognize the importance of top-k dominating query as a data analysis tool and its advantages over top-k and skyline queries — Papadias et al. [23] did not explore such advantages although they introduced top-k dominating query as an extension of skyline query. In this paper, we identify the importance and practicability of the query and define some of its potential extensions. A simple evaluation method for top-k dominating queries, based on skyline computation, was proposed in [23]. The basic idea is to compute the skyline, find the top-1 object o in it (note that the top-1 point must belong to the skyline), remove o from $\mathcal{D}$ and iteratively apply the same procedure, until k results have been output. This skyline-based approach may perform many unnecessary score countings, since the skyline could be much larger than k. In addition, we note that the R-tree (used in the solution of [23]) may not be the most appropriate index for this query; since computing $\mu(p)$ is in fact an *aggregate* query, we can replace the R-tree by an *aggregate R-tree* (aR-tree) [17, 22].

Motivated by these observations, we propose specialized algorithms that operate on aR-trees. Our technical contributions include (i) a batch counting technique for computing scores of multiple points simultaneously, (ii) a counting-guided search algorithm for processing top-k dominating queries, and (iii) a priority-based tree traversal algorithm that retrieves query results by examining each tree node at most once. We enhance the performance of (ii) with *lightweight counting*, which derives relatively tight upper bound scores for non-leaf tree entries at low I/O cost. Furthermore, to our surprise, the intuitive *best-first* traversal order [13, 23] turns out not to be the most efficient for (iii) because of potential partial dominance relationships between visited entries. Thus, we perform a careful analysis on (iii) and propose a *novel, efficient tree traversal order* for it. Extensive experiments show that our methods significantly outperform the skyline-based approach. Finally, we define two interesting query variants; *aggregate* top-k dominating queries and *bichromatic* top-k dominating queries and show how our methods can be extended to process them.

The rest of the paper is organized as follows. Section 2 reviews the related work. Section 3 discusses the properties of top-k dominating search and proposes optimizations for the existing solution in [23]. We then propose eager/lazy approaches for evaluating top-k dominating queries. Section 4 presents an eager approach that guides the search by deriving tight score bounds for encountered non-leaf tree entries immediately. Section 5 develops an alternative, lazy approach that defers score computation of visited entries and gradually refines their score bounds when more tree nodes are accessed. Section 6 introduces extensions of top-k dominating queries and discusses their evaluation. In Section 7, experiments are conducted on both real and synthetic datasets to demonstrate that the proposed algorithms are efficient and also top-k dominating queries return meaningful results to users. Section 8 discusses alternative approaches for top-k dominating queries and query processing on non-indexed data. Finally, Section 9 concludes the paper.

2 Related Work

Top-k dominating queries include a counting component which is a case of multi-dimensional aggregation; in this section, we review related work on spatial aggregation processing. In addition, as the dominance relationship is relevant to skyline queries, we survey existing methods for computing skylines.

2.1 Spatial Aggregation Processing

R-trees [12] have been extensively used as access methods for multi-dimensional data and for processing spatial queries, e.g., range queries, nearest neighbors [13], and skyline queries [23]. The aggregate R-tree (aR-tree) [17, 22] augments to each non-leaf entry of the R-tree an aggregate measure of all data points in the subtree pointed by it. It has been used to speed up the evaluation of spatial aggregate queries, where measures (e.g., number of buildings) in a spatial region (e.g., a district) are aggregated.

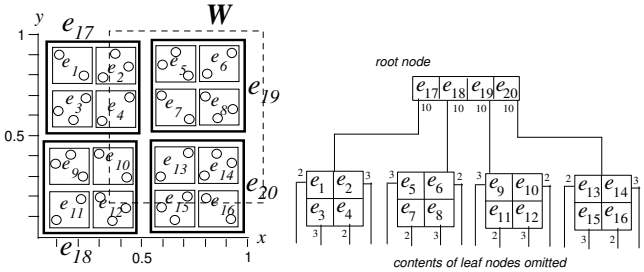

(a) a set of points (b) a COUNT aR-tree

Figure 2: aR-tree example

Figure 2a shows a set of points in the 2D space, indexed by the COUNT aR-tree in Figure 2b. Each non-leaf entry stores the COUNT of data points in its subtree. For instance, in Figure 2b, entry e_{17} has a count 10, meaning that the subtree of e_{17} contains 10 points. Suppose that a user asks for the number of points intersecting the region W, shown in Figure 2a. To process the query, we first examine entries in the root node of the tree. Entries that do not intersect W are pruned because their subtree cannot contain any points in W. If an entry is spatially covered by W (e.g., entry e_{19}), its count (i.e., 10) is added to the answer without accessing the corresponding subtree. Finally, if a non-leaf entry intersects W but it is not contained in W (e.g., e_{17}), search is recursively applied to the child node pointed by the entry, since the corresponding subtree may contain points inside or outside W. Note that the counts augmented in the entries effectively reduce the number of accessed nodes. To evaluate the above example query, only 10 nodes in the COUNT aR-tree are accessed but 17 nodes in an R-tree with the same node capacity would be visited.

2.2 Skyline Computation

Börzsönyi et al. [2] were the first to propose efficient external memory algorithms for processing skyline queries. The BNL (block-nested-loop) algorithm scans the dataset while employing a bounded buffer for tracking the points that cannot be dominated by other points in the buffer. A point is reported as a result if it cannot be dominated by any other point in the dataset. On the other hand, the DC (divide-and-conquer) algorithm recursively partitions the dataset until each partition is small enough to fit in memory. After the local skyline in each partition is computed, they are merged to form the global skyline. The BNL algorithm was later improved to SFS (sort-filter-skyline) [8] and LESS (linear elimination sort for skyline) [11] in order to optimize the average-case running time.

The above algorithms are generic and applicable for non-indexed data. On the other hand, [25, 16, 23] exploit data indexes to accelerate skyline computation. The state-of-the-art algorithm is the BBS (branch-and-bound skyline) algorithm [23], which is shown to be I/O optimal for computing skylines on datasets indexed by R-trees.

Recently, the research focus has been shifted to the study of queries based on variants of the dominance relationship. [20] propose a data cube structure for speeding up the evaluation of queries that analyze the dominance relationship of points in the dataset.

However, incremental maintenance of the data cube over updates has not been addressed in [20]. Clearly, it is prohibitively expensive to recompute the data cube from scratch for dynamic datasets with frequent updates. [6] identify the problem of computing *top-k frequent skyline* points, where the frequency of a point is defined by the number of dimensional subspaces. [5] study the *k-dominant skyline* query, which is based on the k-dominance relationship. A point p is said to k-dominate another point p' if p dominates p' in at least one k-dimensional subspace. The k-dominant skyline contains the points that are not k-dominated by any other point. When k decreases, the size of the k-dominant skyline also decreases. Observe that [20, 6, 5] cannot be directly applied to evaluate top-k dominating queries studied in this paper.

Finally, [28, 24] study the efficient computation of skylines for every subspace; [26] propose a technique for retrieving the skyline for a given subspace; [1, 15] investigate skyline computation over distributed data; [10, 7] develop techniques for estimating the skyline cardinality; [21] study continuous maintenance of the skyline over a data stream; and [4] address skyline computation over datasets with partially-ordered attributes.

3 Preliminary

In this section, we discuss some fundamental properties of top-k dominating search, assuming that the data have been indexed by an aR-tree. In addition, we propose an optimized version for the existing top-k dominating algorithm [23] that operates on aR-tees.

3.1 Score Bounding Functions

Before presenting our top-k dominating algorithms, we first introduce some notation that will be used in this paper. For an aR-tree entry e (i.e., a minimum bounding box) whose projection on the i-th dimension is the interval $[e[i]^-, e[i]^+]$, we denote its lower corner e^- and upper corner e^+ by

$$e^- = (e[1]^-, e[2]^-, \cdots, e[d]^-)$$

$$e^+ = (e[1]^+, e[2]^+, \cdots, e[d]^+)$$

Observe that both e^- and e^+ do not correspond to actual data points but they allow us to express dominance relationships among points and minimum bounding boxes conveniently. As Figure 3 illustrates, there are three cases for a point to dominate a non-leaf entry. Since $p_1 \succ e_1^-$ (i.e., full dominance), p_1 must also dominate *all* data points indexed under e_1. On the other hand, point p_2 dominates e_1^+ but not e_1^- (i.e., partial dominance), thus p_2 dominates some, but not all data points in e_1. Finally, as $p_3 \not\succ e_1^+$ (i.e., no dominance), p_3 cannot dominate any point in e_1. Similarly, the cases for an entry to dominate another entry are: (i) full dominance (e.g., $e_1^+ \succ e_3^-$), (ii) partial dominance (e.g., $e_1^- \succ e_4^+ \wedge e_1^+ \not\succ e_4^-$), (iii) no dominance (e.g., $e_1^- \not\succ e_2^+$).

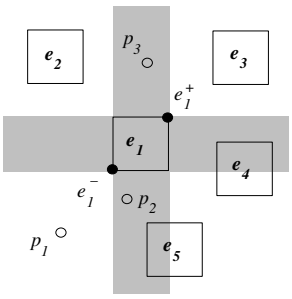

Figure 3: Dominance relationship among aR-tree entries

Given a tree entry e, whose sub-tree has not been visited, $\mu(e^+)$ and $\mu(e^-)$ correspond to the *tightmost* lower and upper score bounds respectively, for any point indexed under e. As we will show later, $\mu(e^+)$ and $\mu(e^-)$ can be computed by a search procedure that accesses only aR-tree nodes that intersect e along at least one dimension. These bounds help pruning the search space and defining a good order for visiting aR-tree nodes. Later in Sections 4 and 5, we replace the tight bounds $\mu(e^+)$ and $\mu(e^-)$ with loose lower and upper bounds for them ($\mu^l(e)$ and $\mu^u(e)$, respectively). Bounds $\mu^l(e)$ and $\mu^u(e)$ are cheaper to compute and can be progressively refined during search, therefore trading-off between computation cost and bound tightness. The computation and use of score bounds in practice will be further elaborated there.

3.2 Optimizing the Skyline-Based Approach

Papadias et al. [23] proposed a Skyline-Based Top-k Dominating Algorithm (STD) for top-k dominating queries, on data indexed by an R-tree. They noted that the skyline is guaranteed to contain the top-1 dominating point, since a non-skyline point has lower score than at least one skyline point that dominates it (see Equation 3). Thus, STD retrieves the skyline points, computes their μ scores and outputs the point p with the highest score. It then removes p from the dataset, incrementally finds the skyline of the remaining points, and repeats the same process.

Consider for example a top-2 dominating query on the dataset shown in Figure 4. STD first retrieves the skyline points p_1, p_2, and p_3 (using the BBS skyline algorithm of [23]). For each skyline point, a range query is issued to count the number of points it dominates. After that, we have $\mu(p_1) = 1$, $\mu(p_2) = 4$, and $\mu(p_3) = 1$. Hence, p_2 is reported as the top-1 result. We now restrict the region of searching for the next result. First, Equation 3 suggests that the region dominated by the remaining skyline points (i.e., p_1 and p_3) needs not be examined. Second, the region dominated by p_2 (i.e., the previous result) may contain some points which are not dominated by the remaining skyline points p_1 and p_3. It suffices to retrieve the skyline points (i.e., p_4 and p_5) in the constrained (gray) region M shown in Figure 4. After counting their scores using the tree, we have $\mu(p_4) = 2$ and $\mu(p_5) = 1$. Finally, we compare them with the scores of retrieved points (i.e., p_1 and p_3) and report p_4 as the next result.

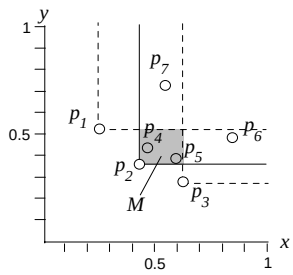

Figure 4: Constrained skyline

In this section, we present two optimizations that greatly reduce the I/O cost of the above solution by exploiting aR-trees. Our first optimization is called *batch counting*. Instead of iteratively applying separate range queries to compute the scores of the skyline points, we perform them in batch. Algorithm 1 shows the pseudocode of this recursive batch counting procedure. It takes two parameters: the current aR-tree node Z and the set of points V, whose μ scores are to be counted. Initially, Z is set to the root node of the tree and $\mu(p)$ is set to 0 for each $p \in V$. Let e be the current entry in Z to be examined. As illustrated in Section 3.1, if e is a non-leaf entry and there exists some point $p \in V$ such that $p \succ e^+ \wedge p \not\succ e^-$, then p may dominate some (but not guaranteed to dominate all) points indexed under e. Thus, we cannot imme-

diately decide the number of points in e dominated by p. In this case, we have to invoke the algorithm recursively on the child node pointed by e. Otherwise, for each point $p \in V$, its score is incremented by $\texttt{COUNT}(e)$ when it dominates e^-. BatchCount correctly computes the μ score for all $p \in V$, at a single tree traversal.

Algorithm 1 Batch Counting

```
    algorithm BatchCount(Node Z, Point set V)
    for all entries e ∈ Z do
        if Z is non-leaf and ∃p ∈ V, p ≻ e⁺ ∧ p ≯ e⁻ then
            read the child node Z′ pointed by e;
            BatchCount(Z′, V);
        else
            for all points p ∈ V do
                if p ≻ e⁻ then
                    μ(p):=μ(p)+COUNT(e);
```

Algorithm 2 is a pseudo-code of the Iterative Top-k Dominating Algorithm (ITD), which optimizes the STD algorithm of [23]. Like STD, ITD computes the top-k dominating points iteratively. In the first iteration, ITD computes in V' the skyline of the whole dataset, while in subsequent iterations, the computation is *constrained* to a region M. M is the region dominated by the reported point q in the previous iteration, but not any point in the set V of retrieved points in past iterations. At each loop, Lines 6–8 compute the scores for the points in V' in batches of B points each ($B \leq |V'|$). By default, the value of B is set to the number of points that can fit into a memory page. Our second optimization is that we sort the points in V' by a space-filling curve (Hilbert ordering) [3] before applying batch counting, in order to increase the compactness of the MBR of a batch. After merging the constrained skyline with the global one, the object q with the highest μ score is reported as the next dominating object, removed from V and used to compute the constrained skyline at the next iteration. The algorithm terminates after k objects have been reported.

For instance, in Figure 4, q corresponds to point $(0, 0)$ and $V = \varnothing$ in the first loop, thus M corresponds to the whole space and the whole skyline $\{p_1, p_2, p_3\}$ is stored in V', the points there are sorted and split in batches and their μ scores are counted using the BatchCount algorithm. In the beginning of the second loop, $q = p_2$, $V = \{p_1, p_3\}$, and M is the gray region in the figure. V' now becomes $\{p_4, p_5\}$ and the corresponding scores are batch-counted. The next point is then reported (e.g., p_4) and the algorithm continues as long as more results are required.

Algorithm 2 Iterative Top-k Dominating Algorithm (ITD)

```
    algorithm ITD(Tree R, Integer k)
 1: V:=∅; q:=origin point;
 2: for i := 1 to k do
 3:     M:=region dominated by q but by no point in V;
 4:     V′:=skyline points in M;
 5:     sort the points in V′ by Hilbert ordering;
 6:     for all batches V_c of (B) points in V′ do
 7:         initialize all scores of points in V_c to 0;
 8:         BatchCount(R.root,V_c);
 9:     V:=V ∪ V′;
10:     q:=the point with maximum score in V;
11:     remove q from V;
12:     report q as the i-th result;
```

4 Counting-Guided Search

The skyline-based solution becomes inefficient for datasets with large skylines as μ scores of many points are computed. In addition, not all skyline points have large μ scores. Motivated by these observations, we study algorithms that solve the problem directly, without depending on skyline computations. This section presents an *eager* approach for the evaluation of top-k dominating queries, which traverses the aR-tree and computes tight upper score bounds for encountered non-leaf tree entries immediately; these bounds determine the visiting order for the tree nodes. We discuss the basic algorithm, develop optimizations for it, and investigate by an analytical study the improvements of these optimizations.

4.1 The Basic Algorithm

Recall from Section 3.1 that the score of any point p indexed under an entry e is upper-bounded by $\mu(e^-)$. Based on this observation, we can design a method that traverses aR-tree nodes in descending order of their (upper bound) scores. The rationale is that points with high scores can be retrieved early and accesses to aR-tree nodes that do not contribute to the result can be avoided.

Algorithm 3 shows the pseudo code of the Simple Counting-Guided Algorithm (SCG), which directs search by counting upper bound scores of examined non-leaf entries. A max-heap H is employed for organizing the entries to be visited in descending order of their scores. W is a min-heap for managing the top-k dominating points as the algorithm progresses, while γ is the k-th score in W (used for pruning). First, the upper bound scores $\mu(e^-)$ of the aR-tree root entries are computed in batch (using the BatchCount algorithm) and these are inserted into the max-heap H. While the score $\mu(e^-)$ of H's top entry e is higher than γ (implying that points with scores higher than γ may be indexed under e), the top entry is deheaped, and the node Z pointed by e is visited. If Z is a non-leaf node, its entries are enheaped, after BatchCount is called to compute their upper score bounds. If Z is a leaf node, the scores of the points in it are computed in batch and the top-k set W (also γ) is updated, if applicable.

Algorithm 3 Simple Counting Guided Algorithm (SCG)

```
    algorithm SCG(Tree R, Integer k)
    H:=new max-heap; W:=new min-heap;
    γ:=0;                                  ▷ the k-th highest score found so far
    BatchCount(R.root,{e⁻ | e ∈ R.root});
    for all entries e ∈ R.root do
        enheap(H, ⟨e, μ(e⁻)⟩);
    while |H| > 0 and H's top entry's score > γ do
        e:=deheap(H);
        read the child node Z pointed by e;
        if Z is non-leaf then
            BatchCount(R.root,{e_c⁻ | e_c ∈ Z});
            for all entries e_c ∈ Z do
                enheap(H, ⟨e_c, μ(e_c⁻)⟩);
        else                                ▷ Z is a leaf
            BatchCount(R.root,{p | p ∈ Z});
            update W and γ, using ⟨p, μ(p)⟩, ∀p ∈ Z
    report W as the result;
```

As an example, consider the top-1 dominating query on the set of points in Figure 5. There are 3 leaf nodes and their corresponding entries in the root node are e_1, e_2, and e_3. First, upper bound scores for the root entries (i.e., $\mu(e_1^-) = 3$, $\mu(e_2^-) = 7$, $\mu(e_3^-) = 3$) are computed by the batch counting algorithm, which incurs 3 node accesses (i.e., the root node and leaf nodes pointed by e_1 and e_3). Since e_2 has the highest upper bound score, the leaf node pointed by e_2 will be accessed next. Scores of entries in e_2 are computed in batch and we obtain $\mu(p_1) = 5$, $\mu(p_2) = 1$, $\mu(p_3) = 2$. Since p_1 is a point and $\mu(p_1)$ is higher than the scores of remaining entries (p_2, p_3, e_1, e_3), p_1 is guaranteed to be the top-1 result.

4.2 Optimizations

Now, we discuss three optimizations that can greatly reduce the cost of the basic SCG. First, we utilize encountered data points to strengthen the pruning power of the algorithm. Next, we apply a

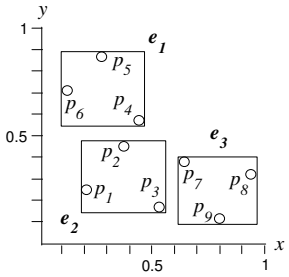

Figure 5: Computing upper bound scores

lazy counting method that delays the counting for points, in order to form better groups for batch counting. Finally, we develop a lightweight technique for deriving upper score bounds of non-leaf entries at low cost.

The pruner set. SCG visits nodes and counts the scores of points and entries, based only on the condition that the upper bound score of their parent entry is greater than γ. However, we observe that points which have been counted, but have scores at most γ can also be used to prune early other entries or points, which are dominated by them.[1] Thus, we maintain a pruner set F, which contains points that (i) have been counted exactly (i.e., at Line 15), (ii) have scores at most γ, and (iii) are not dominated by any other point in F. The third condition ensures that only minimal information is kept in F.[2] We perform the following changes to SCG in order to use F. First, after deheaping an entry e (Line 7), we check whether there exists a point $p \in F$, such that $p \succ e^-$. If yes, then e is pruned and the algorithm goes back to Line 6. Second, before applying BatchCount at Lines 10 and 14, we eliminate any entries or points that are dominated by a point in F.

Lazy counting. The performance of SCG is negatively affected by executions of BatchCount for a small number of points. A batch may have few points if many points in a leaf node are pruned with the help of F. In order to avoid this problem, we employ a *lazy counting* technique, which works as follows. When a leaf node is visited (Line 13), instead of directly performing batch counting for the points p, those that are not pruned by F are inserted into a set L, with their upper bound score $\mu(e^-)$ from the parent entry. If, after an insertion, the size of L exceeds B (the size of a batch), then BatchCount is executed for the contents of L, and all W, γ, F are updated. Just before reporting the final result set (Line 16), batch counting is performed for potential results $p \in L$ not dominated by any point in F and with upper bound score greater than γ. We found that the combined effect of the pruner set and lazy counting lead to 30% I/O cost reduction of SCG, in practice.

Lightweight upper bound computation. As mentioned in Section 3.1, the tight upper score bound $\mu(e^-)$ can be replaced by a looser, cheaper to compute, bound $\mu^u(e)$. We propose an optimized version of SCG, called Lightweight Counting Guided Algorithm (LCG). Line 10 of SCG (Algorithm 3) is replaced by a call to LightBatchCount, which is a variation of BatchCount. In specific, when bounds for a set V of *non-leaf* entries are counted, the algorithm avoids expensive accesses at aR-tree leaf nodes, but uses entries at non-leaf nodes to derive looser bounds.

LightBatchCount is identical to Algorithm 1, except that the recursion of Line 2 is applied when Z is at least two levels above leaf nodes and there is a point in V that partially dominates e; thus, the else statement at Line 5 now refers to nodes one level above the leaves. In addition, the condition at Line 7 is replaced by $p \succ e^+$; i.e., COUNT(e) is added to $\mu^u(p)$, even if p partially dominates entry e.

As an example, consider the three root entries of Figure 5. We can compute loose upper score bounds for $V = \{e_1^-, e_2^-, e_3^-\}$, without accessing the leaf nodes. Since, e_2^- fully dominates e_2 and partially dominates e_1, e_3, we get $\mu^u(e_2) = 9$. Similarly, we get $\mu^u(e_1) = 3$ and $\mu^u(e_3) = 3$. Although these bounds are looser than the respective tight ones, they still provide a good order of visiting the entries and they can be used for pruning and checking for termination. In Section 7, we demonstrate the significant computation savings by this lightweight counting (of $\mu^u(e)$) over exact counting (of $\mu(e^-)$) and show that it affects very little the pruning power of the algorithm. Next, we investigate its effectiveness by a theoretical analysis.

4.3 Analytical Study

Consider a dataset $\mathcal{D}$ with N points, indexed by an aR-tree whose nodes have an average fanout f. Our analysis is based on the assumption that the data points are uniformly and independently distributed in the domain space $[0, 1]^d$, where d is the dimensionality. Then, the tree height h and the number of nodes n_i at level i (let the leaf level be 0) can be estimated by $h = 1 + \lceil \log_f(N/f) \rceil$ and $n_i = N/f^{i+1}$. Besides, the extent (i.e., length of any 1D projection) λ_i of a node at the i-th level can be approximated by $\lambda_i = (1/n_i)^{1/d}$ [27].

We now discuss the trade-off of lightweight counting over exact counting for a non-leaf entry e. Recall that the *exact* upper bound score $\mu(e^-)$ is counted as the number of points dominated by its lower corner e^-. On the other hand, lightweight counting obtains $\mu^u(e)$; an upper bound of $\mu(e^-)$. For a given e^-, Figure 6 shows that the space can be divided into three regions, with respect to nodes at level i. The gray region M_2 corresponds to the maximal region, covering nodes (at level i) that are *partially* dominated by e^-. While computing $\mu(e^-)$, only the entries which are *completely inside* M_2 need to be further examined (e.g., e_A). Other entries are pruned after either disregarding their aggregate values (e.g., e_B, which intersects M_1), or adding these values to $\mu(e^-)$ (e.g., e_C, which intersects M_3).

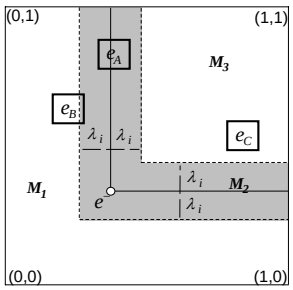

Figure 6: I/O cost of computing upper bound

Thus, the probability of accessing a (i-th level) node can be approximated by the area of M_2, assuming that tree nodes at the same level have no overlapping. To further simplify our analysis, suppose that all coordinates of e^- are of the same value v. Hence, the aR-tree node accesses required for computing the exact $\mu(e^-)$ can

[1] Suppose that a point p satisfies $\mu(p) \leq \gamma$. Applying Equation 3, if a point p' is dominated by p, then we have $\mu(p') < \gamma$.

[2] Note that F is the skyline of a specific data subset.

be expressed as[3]:

$$NA_E(e^-) = \sum_{i=0}^{h-1} n_i \cdot [(1 - v + \lambda_i)^d - (1 - v - \lambda_i)^d] \quad (4)$$

In the above equation, the quantity in the square brackets corresponds to the volume of M_2 (at level i) over the volume of the universe (this equals to 1), capturing thus the probability of a node at level i to be completely inside M_2. The node accesses of lightweight computation can also be captured by the above equation, except that no leaf nodes (i.e., at level 0) are accessed. As there are many more leaf nodes than non-leaf nodes, lightweight computation incurs significantly lower cost than exact computation.

Now, we compare the scores obtained by exact computation and lightweight computation. The exact score $\mu(e^-)$ is determined by the area dominated by e^-:

$$\mu(e^-) = N \cdot (1 - v)^d \quad (5)$$

In addition to the above points, lightweight computation counts also all points in M_2 for the leaf level into the upper bound score:

$$\mu^u(e) = N \cdot (1 - v + \lambda_0)^d \quad (6)$$

Summarizing, three factors N, v, and d affect the relative tightness of the lightweight score bound over the exact bound.

- When N is large, the leaf node extent λ_0 is small and thus the lightweight score is tight.
- If v is small, i.e., e^- is close to the origin and has high dominating power, then λ_0 becomes less significant in Equation 6 and the ratio of $\mu^u(e)$ to $\mu(e^-)$ is close to 1 (i.e., lightweight score becomes relatively tight).
- As d increases (decreases), λ_0 also increases (decreases) and the lightweight score gets looser (tighter).

In practice, during counting-guided search, entries close to the origin have higher probability to be accessed than other entries, since their parent entries have higher upper bounds and they are prioritized by search. As a result, we expect that the second case above will hold for most of the upper bound computations and lightweight computation will be effective.

5 Priority-Based Traversal

In this section, we present a *lazy* alternative to the counting-guided method. Instead of computing upper bounds of visited entries by explicit counting, we defer score computations for entries, but maintain lower and upper bounds for them as the tree is traversed. Score bounds for visited entries are gradually refined when more nodes are accessed, until the result is finalized with the help of them. For this method to be effective, the tree is traversed with a carefully-designed priority order aiming at minimizing I/O cost. We present the basic algorithm, analyze the issue of setting an appropriate order for visiting nodes, and discuss its implementation.

5.1 The Basic Algorithm

Recall that counting-guided search, presented in the previous section, may access some aR-tree nodes more than once due to the application of counting operations for the visited entries. For instance in Figure 5, the node pointed by e_1 may be accessed twice; once for counting the scores of points under e_2 and once for counting the scores of points under e_1. We now propose a top-k dominating algorithm which traverses each node at most once and has reduced I/O cost.

Algorithm 4 shows the pseudo-code of this Priority-Based Tree Traversal Algorithm (PBT). PBT browses the tree, while maintaining (loose) upper $\mu^u(e)$ and lower $\mu^l(e)$ score bounds for the entries e that have been seen so far. The nodes of the tree are visited based on a *priority order*. The issue of defining an appropriate ordering of node visits will be elaborated later. During traversal, PBT maintains a set S of visited aR-tree entries. An entry in S can either: (i) lead to a potential result, or (ii) be partially dominated by other entries in S that may end up in the result. W is a min-heap, employed for tracking the top-k points (in terms of their μ^l scores) found so far, whereas γ is the lowest score in W (used for pruning).

First, the root node is loaded, and its entries are inserted into S after upper score bounds have been derived from information in the root node. Then (Lines 8-18), while S contains non-leaf entries, the non-leaf entry e_z with the highest priority is removed from S, the corresponding tree node Z is visited and (i) the μ^u (μ^l) scores of existing entries in S (partially dominating e_z) are refined using the contents of Z, (ii) μ^u (μ^l) values for the contents of Z are computed and, in turn, inserted to S. Note that for operations (i) and (ii), only information from the current node and S is used; no additional accesses to the tree are required. Updates and computations of μ^u scores are performed incrementally with the information of e_z and entries in S that partially dominate e_z. W is updated with points/entries of higher μ^l than γ. Finally (Line 20), entries are pruned from S if (i) they cannot lead to points that may be included in W, and (ii) are not partially dominated by entries leading to points that can reach W.

Algorithm 4 Priority-Based Tree Traversal Algorithm (PBT)

algorithm PBT(Tree R, Integer k)
1: S:=new set; ▷ entry format in S: $\langle e, \mu^l(e), \mu^u(e)\rangle$
2: W:=new min-heap; ▷ k points with the highest μ^l
3: γ:=0; ▷ the k-th highest μ^l score found so far
4: **for all** $e_x \in R.root$ **do**
5: $\mu^l(e_x)$:=$\sum_{e \in R.root \wedge e_x^+ \succ e^-}$ COUNT(e);
6: $\mu^u(e_x)$:=$\sum_{e \in R.root \wedge e_x^- \succ e^+}$ COUNT(e);
7: insert e_x into S and update W;
8: **while** S contains non-leaf entries **do**
9: remove e_z: non-leaf entry of S with the highest *priority*;
10: read the child node Z pointed by e_z;
11: **for all** $e_y \in S$ such that $e_y^+ \not\succ e_z^- \wedge e_y^- \succ e_z^+$ **do**
12: $\mu^l(e_y)$:=$\mu^l(e_y) + \sum_{e \in Z \wedge e_y^+ \succ e^-}$ COUNT(e);
13: $\mu^u(e_y)$:=$\mu^l(e_y) + \sum_{e \in Z \wedge e_y^+ \not\succ e^- \wedge e_y^- \succ e^+}$ COUNT(e);
14: S_z:=$Z \cup \{e \in S \mid e_z^+ \not\succ e^- \wedge e_z^- \succ e^+\}$;
15: **for all** $e_x \in Z$ **do**
16: $\mu^l(e_x)$:=$\mu^l(e_x) + \sum_{e \in S_z \wedge e_x^+ \succ e^-}$ COUNT(e);
17: $\mu^u(e_x)$:=$\mu^l(e_x) + \sum_{e \in S_z \wedge e_x^+ \not\succ e^- \wedge e_x^- \succ e^+}$ COUNT(e);
18: insert all entries of Z into S;
19: update W (and γ) by $e' \in S$ whose score bounds changed;
20: remove entries e_m from S where $\mu^u(e_m) < \gamma$ and $\neg\exists e \in S,\ (\mu^u(e) \geq \gamma) \wedge (e^+ \not\succ e_m^- \wedge e^- \succ e_m^+)$;
21: **report** W as the result;

It is important to note that, at Line 21 of PBT, all non-leaf entries have been removed from the set S, and thus (result) points in W have their exact scores found.

To comprehend the functionality of PBT consider again the top-1 dominating query on the example of Figure 5. For the ease of discussion, we denote the score bounds of an entry e by the interval $\mu_\star(e)$=[$\mu^l(e), \mu^u(e)$]. Initially, PBT accesses the root node

[3]For simplicity, the equation does not consider the boundary effect (i.e., v is near the domain boundary). To capture the boundary effect, we need to bound the terms $(1 - v + \lambda_i)$ and $(1 - v - \lambda_i)$ within the range $[0, 1]$.

and its entries are inserted into S after their lower/upper bound scores are derived (see Lines 5–6); $\mu_\star(e_1)$=$[0,3]$, $\mu_\star(e_2)$=$[0,9]$, $\mu_\star(e_3)$=$[0,3]$. Assume for now, that visited nodes are prioritized (Lines 9-10) based on the upper bound scores $\mu^u(e)$ of entries $e \in S$. Entry e_2, of the highest score μ^u in S is removed and its child node Z is accessed. Since $e_1^- \not\succ e_2^+$ and $e_3^- \not\succ e_2^+$, the upper/lower score bounds of remaining entries $\{e_1, e_3\}$ in S will not be updated (the condition of Line 11 is not satisfied). The score bounds for the points p_1, p_2, and p_3 in Z are then computed; $\mu_\star(p_1)$=$[1,7]$, $\mu_\star(p_2)$=$[0,3]$, and $\mu_\star(p_3)$=$[0,3]$. These points are inserted into S, and W=$\{p_1\}$ with γ=$\mu^l(p_1)$=1. No entry or point in S can be pruned, since their upper bounds are all greater than γ. The next non-leaf entry to be removed from S is e_1 (the tie with e_3 is broken arbitrarily). The score bounds of the existing entries S=$\{e_3, p_1, p_2, p_3\}$ are in turn refined; $\mu_\star(e_3)$ remains $[0,3]$ (unaffected by e_1), whereas $\mu_\star(p_1)$=$[3,6]$, $\mu_\star(p_2)$=$[1,1]$, and $\mu_\star(p_3)$ =$[0,3]$. The scores of the points indexed by e_1 are computed; $\mu_\star(p_4)$=$[0,0]$, $\mu_\star(p_5)$=$[0,0]$, and $\mu_\star(p_6)$=$[1,1]$ and W is updated to p_1 with γ=$\mu^l(p_1)$=3. At this stage, all points, except from p_1, are pruned from S, since their μ^u scores are at most γ and they are not partially dominated by non-leaf entries that may contain potential results. Although no point from e_3 can have higher score than p_1, we still have to keep e_3, in order to compute the exact score of p_1 in the next round.

5.2 Traversal Orders in PBT

An intuitive method for prioritizing entries at Line 9 of PBT, hinted by the *upper bound principle* of [19] or the *best-first ordering* of [13, 23], is to pick the entry e_z with the highest upper bound score $\mu^u(e_z)$; such an order would visit the points that have high probability to be in the top-k dominating result early. We denote this instantiation of PBT by UBT (for Upper-bound Based Traversal).

Nevertheless a closer look into PBT (Algorithm 4) reveals that the upper score bounds alone may not offer the best priority order for traversing the tree. Recall that the pruning operation (at Line 20) eliminates entries from S, saving significant I/O cost and leading to the early termination of the algorithm. The effectiveness of this pruning depends on the *lower* bounds of the best points (stored in W). Unless these bounds are tight enough, PBT will not terminate early and S will grow very large.

For example, consider the application of UBT to the tree of Figure 2. The first few nodes accessed are in the order: root node, e_{18}, e_{11}, e_9, e_{12}. Although e_{11} has the highest upper bound score, it *partially dominates* high-level entries (e.g., e_{17} and e_{20}), whose child nodes have not been accessed yet. As a result, the best-k score γ (i.e., the current lower bound score of e_{11}) is small, few entries can be pruned, and the algorithm does not terminate early.

Thus, the objective of search is not only to (i) examine the entries of large upper bounds early, which leads to early identification of candidate query results, but also (ii) eliminate partial dominance relationships between entries that appear in S, which facilitates the computation of tight lower bounds for these candidates. We now investigate the factors affecting the probability that one node partially dominates another and link them to the traversal order of PBT. Let a and b be two random nodes of the tree such that a is at level i and b is at level j. Using the same uniformity assumptions and notation as in Section 4.3, we can infer that the two nodes a and b not intersect along dimension t with probability[4]:

$$Pr(a[t] \cap b[t] = \varnothing) = 1 - (\lambda_i + \lambda_j)$$

a and b have a partial dominance relationship when they intersect along at least one dimension. The probability of being such is:

$$Pr(\bigvee_{t \in [1,d]} a[t] \cap b[t] \neq \varnothing) = 1 - (1 - (\lambda_i + \lambda_j))^d$$

The above probability is small when the sum $\lambda_i + \lambda_j$ is minimized (e.g., a and b are both at low levels).

The above analysis leads to the conclusion that in order to minimize the partially dominating entry pairs in S, we should prioritize the visited nodes based on their level at the tree. In addition, between entries at the highest level in S, we should choose the one with the highest upper bound, in order to find the points with high scores early. Accordingly, we propose an instantiation of PBT, called Cost-Based Traversal (CBT). CBT corresponds to Algorithm 4, such that, at Line 9, the non-leaf entry e_z with the highest level is removed from S and processed; if there are ties, the entry with the highest upper bound score is picked. In Section 7, we demonstrate the advantage of CBT over UBT in practice.

5.3 Implementation Details

A straightforward implementation of PBT may lead to very high computational cost. At each loop, the burden of the algorithm is the pruning step (Line 20 of Algorithm 4), which has worst-case cost quadratic to the size of S; entries are pruned from S if (i) their upper bound scores are below γ and (ii) they are not partially dominated by any other entry with upper bound score above γ. If an entry e_m satisfies (i), then a scan of S is required to check (ii).

In order to check for condition (ii) efficiently, we use a main-memory R-tree $I(S)$ to index the entries in S having upper bound score above γ. When the upper bound score of an entry drops below γ, it is removed from $I(S)$. When checking for pruning of e_m at Line 20 of PBT, we only need to examine the entries indexed by $I(S)$, as only these have upper bound scores above γ. In particular, we may not even have to traverse the whole index $I(S)$. For instance, if a non-leaf entry e' in $I(S)$ does not partially dominate e_m, then we need not check for the subtree of e'. As we verified experimentally, maintaining $I(S)$ enables the pruning step to be implemented efficiently. In addition to $I(S)$, we tried additional data structures for accelerating the operations of PBT (e.g., a priority queue for popping the next entry from S at Line 9), however, the maintenance cost of these data structures (as the upper bounds of entries in S change frequently at Lines 11-13) did not justify the performance gains by them.

6 Extensions

This section discusses interesting extensions to the basic form of top-k dominating queries we have studied so far. We note that the query types that are discussed here are original; to our knowledge they have not been mentioned or studied in the literature before.

6.1 Generic Aggregate Functions and Point Significance

We can generalize the top-k dominating query to include any aggregate function agg (i.e., instead of `COUNT`) and weights $w(p)$ of significance on points p (i.e., instead of all points having the same significance $w(p) = 1$). The generalized scoring function is defined as:

$$\mu_{agg}(p) = agg\ \{\ w(p') \mid p' \in \mathcal{D} \ \wedge\ p \succ p'\ \} \qquad (7)$$

It is not hard to see that our proposed techniques can be directly used for a generalized top-k dominating query, for distributive and monotone aggregate functions (like `SUM`, `MAX`, `MIN`) and weights of importance on the points. For this purpose, we can use an aggregate R-tree, where entries are augmented with the aggregate score of $w(p)$, for all points p under them.

[4]The current equation is simplified for readability. The probability equals 0 when $\lambda_i + \lambda_j > 1$.

Only slight modifications have to be made in our algorithms because the fundamental property of score dominance (in Equation 3) holds not only for COUNT (i.e., the default top-k dominating query), but also for SUM and MAX. The case for SUM can be directly solved by our algorithms. Regarding MAX, the counting operations (in ITD, LCG) and incremental refinement of score bounds (in PBT) need to be modified for MAX correspondingly. Interestingly, MAX provides us an opportunity to further optimize such counting operations and score refinements. As an example, Figure 7a shows the locations of the points with their weights in brackets. The points are indexed by a MAX aR-tree and the non-leaf entries e_2 and e_3 are augmented with the weights 0.9 and 0.7 respectively. Suppose that we need to compute $\mu_{max}(p_1)$, the score of p_1. We first access the child node of e_2 and update $\mu_{max}(p_1)$ to 0.9. Now, even though p_1 partially dominates e_3, we need not access the node of e_3 as it cannot further improve $\mu_{max}(p_1)$.

Note that query results for MIN can be obtained by evaluating a query for MAX. Specifically, assuming that the interval $[0, 1]$ is the domain of possible weights $w(p)$, our algorithms can be adapted as follows: (i) for each visited point (and entry), convert its weight $w(p)$ to $1 - w(p)$, (ii) evaluate the query for MAX to retrieve results, and (iii) at the end, transform each result value v to $1 - v$ for obtaining the final results.

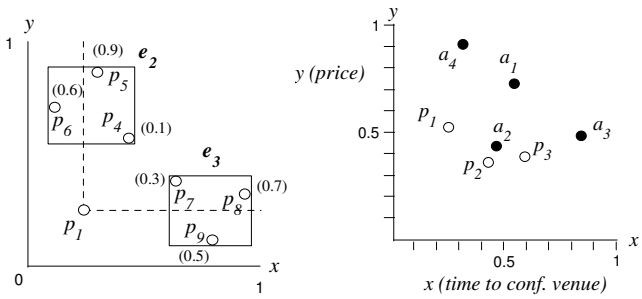

(a) Dominating MAX query (b) Bichromatic query

Figure 7: Variants of top-k dominating queries

6.2 Bichromatic Top-k Dominating Queries

Given a *provider* dataset $\mathcal{D}_P$ and a *consumer* dataset $\mathcal{D}_A$, the score of an object $p \in \mathcal{D}_P$ is defined as:

$$\mu_A(p) = |\ \{\ a \in \mathcal{D}_A \mid p \succ a\ \}\ | \tag{8}$$

A bichromatic top-k dominating query retrieves k data objects in $\mathcal{D}_P$ with the highest μ_A score. As an example of the applicability of this query, consider the points in Figure 7b, where $\mathcal{D}_P = \{p_1, p_2, p_3\}$ stores the feature values of different hotels (shown as white points) and $\mathcal{D}_A = \{a_1, a_2, a_3, a_4\}$ records the requirements for a hotel specified by different customers (shown as black points). For example, customer $a_1 = (0.55, 0.73)$ will only stay in a hotel whose x (time to the conference venue) and y (room price) values are at most 0.55 and 0.73 respectively. The bichromatic top-k dominating query could be used to find the most popular hotel; i.e., the one that fulfills the requirements of the largest number of customers. In this example, we have $\mu_A(p_1) = 2$, $\mu_A(p_2) = 3$, and $\mu_A(p_3) = 1$. Thus, the bichromatic top-1 point is p_2.

Algorithms ITD and LCG can be adapted for bichromatic queries with slight modifications. In particular, candidate points are accessed from the aR-tree on $\mathcal{D}_P$ while their scores are counted using the aR-tree on $\mathcal{D}_A$.

The extensions of PBT for bichromatic queries are more complex. Two sets S_P and S_A are employed for managing visited entries in $\mathcal{D}_P$ and $\mathcal{D}_A$ respectively, and initially they contain root entries of the corresponding tree. First, a non-leaf entry e_A (e.g., according to CBT order) is removed from S_A. After accessing the child node of e_A, its entries are inserted to S_A in order to refine score bounds of entries in S_P. Second, a non-leaf entry e_P (e.g., according to CBT order) is removed from S_P. After accessing the child node of e_P, its entries are inserted to S_P and their score bounds are refined by entries in S_A. Whenever score bounds of entries in S_P change, the result set W and the best-k score γ are updated. In addition, an entry $e_x \in S_P$ is pruned when its upper bound score $\mu_A^u(e_x)$ is below γ. On the other hand, an entry in S_A is pruned if it is not partially dominated by any entry in $e_x \in S_P$ with $\mu_A^u(e_x) \geq \gamma$. The above procedure repeats until S_A becomes empty and S_P contains the same objects as in W (i.e, all other entries in S_P have been eliminated).

7 Experimental Evaluation

In this section, we experimentally evaluate the performance of the proposed algorithms. All algorithms in Table 1 were implemented in C++ and experiments were run on a Pentium D 2.8GHz PC with 1GB of RAM. For fairness to the STD algorithm [23], it is implemented with the spatial aggregation technique (discussed in Section 2.1) for optimizing counting operations on aR-trees. In Section 7.1 we present an extensive experimental study for the efficiency of the algorithms with synthetically generated data. Section 7.2 studies the performance of the algorithms on real data and demonstrates the meaningfulness of top-k dominating points.

Name	Description
STD	Skyline-Based Top-k Dominating Algorithm [23]
ITD	Optimized version of STD (Sec. 3.2)
SCG	Simple Counting Guided Algorithm (Sec. 4)
LCG	Lightweight Counting Guided Algorithm (Sec. 4)
UBT	Upper-bound Based Traversal Algorithm (Sec. 5)
CBT	Cost-Based Traversal Algorithm (Sec. 5)

Table 1: Description of the algorithms

7.1 Experiments With Synthetic Data

Data generation and query parameter values. We produced three categories of synthetic datasets to model different scenarios, according to the methodology in [2]. UI contains datasets where point coordinates are random values *uniformly and independently generated* for different dimensions. CO contains datasets where point coordinates are *correlated*. In other words, for a point p, its i-th coordinate $p[i]$ is close to $p[j]$ in all other dimensions $j \neq i$. Finally, AC contains datasets where point coordinates are *anti-correlated*. In this case, points that are good in one dimension are bad in one or all other dimensions. Table 2 lists the range of parameter values and their default values (in bold type). Each dataset is indexed by an aR-tree with 4K bytes page size. We used an LRU memory buffer whose default size is set to 5% of the tree size.

Parameter	Values
Buffer size (%)	1, 2, **5**, 10, 20
Data size, N (million)	0.25, 0.5, **1**, 2, 4
Data dimensionality, d	2, **3**, 4, 5
Number of results, k	1, 4, **16**, 64, 256

Table 2: Range of parameter values

Lightweight counting optimization in Counting-Guided search. In the first experiment, we investigate the performance savings when using the lightweight counting heuristic in the counting-guided algorithm presented in Section 4. Using a default uniform dataset, for different locations of a non-leaf entry e^-, (after fixing all coordinates of e^- to the same value v), we compare (i) node accesses

of computing the exact $\mu(e^-)$ with that of computing a conservative upper bound $\mu^u(e)$ using the lightweight approach and (ii) the difference between these two bounds. Figure 8a shows the effect of v (i.e., location of e^-) on node accesses of these two computations. Clearly, the lightweight approach is much more efficient than the exact approach. Their cost difference can be two orders of magnitude when e^- is close to the origin. Figure 8b plots the effect of v on the value of upper bound score. Even though lightweight computation accesses much fewer nodes, it derives a score that tightly upper bounds the exact score ($\mu^u(e)$ is only 10% looser than $\mu(e^-)$). Summarizing, the lightweight approach is much more efficient than the exact approach while still deriving a reasonably tight upper bound score.

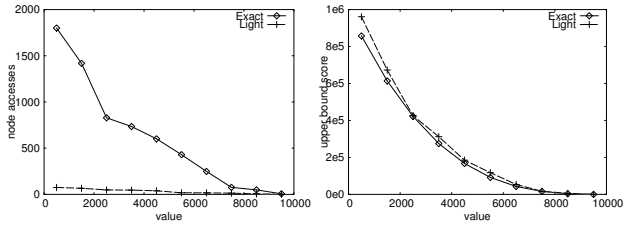

(a) Node accesses (b) Upper bound score

Figure 8: The effect of v, UI, N =1M, $d = 3$

Orderings in Priority-Based Traversal. In Section 5.2, we introduced two priority orders for selecting the next non-leaf entry to process at PBT: (i) UBT chooses the one with the highest upper bound score, and (ii) CBT, among those with the highest level, chooses the one with the highest upper bound score. Having theoretically justified the superiority of CBT over UBT (in Section 5.2), we now demonstrate this experimentally. For the default top-k dominating query on a UI dataset, we record statistics of the two algorithms during their execution. Figure 9a shows the value of γ (i.e., the best-k score) for both UBT and CBT as the number of loops executed. Note that in UBT/CBT, each loop (i.e., Lines 8–20 of Algorithm 4) causes one tree node access. Since γ rises faster in CBT than in UBT, CBT has higher pruning power and thus terminates earlier. Figure 9b plots the size of S (i.e., number of entries in memory) with respect to the number of loops. The size of S in CBT is much lower than that in UBT. Hence, CBT requires less CPU time than UBT on book-keeping the information of visited entries and negligible memory compared to the problem size. Both figures show that our carefully-designed priority order in CBT outperforms the intuitive priority order in UBT by a wide margin.

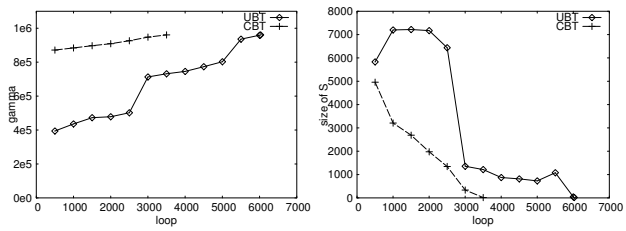

(a) Value of γ (b) Size of S

Figure 9: The effect of ordering priorities, UI, N =1M, $d = 3$

Comparison of all algorithms and variants thereof. We now compare all algorithms and their variants (STD, ITD, SCG, LCG, UBT, CBT) for the default query parameters on UI, CO, and AC datasets (Figure 10). In this and subsequent experiments, we compile the I/O and CPU costs of each algorithm, by charging 10ms I/O time per page fault, and show their I/O-CPU cost-breakdown. ITD performs much better than the baseline STD algorithm of [23] (even though STD operates on the aR-tree), due to the effectiveness of the batch counting and Hilbert ordering techniques for retrieved (constrained) skyline points. LCG and CBT significantly outperform ITD, as they need not compute the scores for the whole skyline, whose size grows huge for AC data. Note that the optimized version of counting-guided search (LCG) outperforms the simple version of the algorithm that computes exact upper bounds (SCG) by a wide margin. Similarly, for priority-based traversal, CBT outperforms UBT because of the reasons explained in the previous experiment. Observe that the best priority-traversal algorithm (CBT) has lower I/O cost than optimized counting-guided search (LCG), since CBT accesses each node at most once but LCG may access some nodes more than once during counting operations.

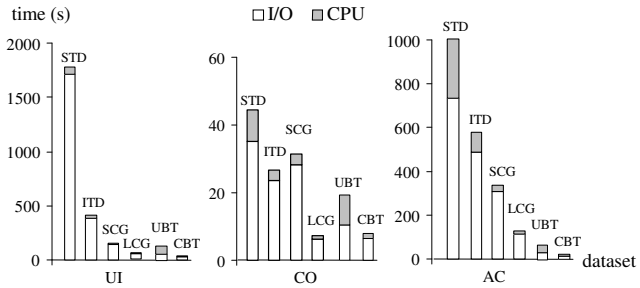

Figure 10: Query cost ($k = 16$, N =1M, $d = 3$)

In remaining experiments, we only compare the best algorithms from each gender (ITD, LCG, and CBT), for a wide range of query and system parameter values. First, we study the effect of the buffer on the performance of the algorithms. Figure 11 shows the cost of the algorithms as a function of buffer size (%). Observe that the costs of LCG and CBT with the smallest tested buffer (1% of the tree size) are still much lower than that of ITD with the largest buffer size (20%). Since CBT accesses each tree node at most once, its cost is independent of the buffer. Clearly, CBT outperforms its competitors for all tested buffer sizes. We note that the memory usage (for storing visited tree entries) of ITD, LCG, and CBT for UI data are 0.03%, 0.02%, 0.96% of the tree size, respectively, and are further reduced by 30% for CO data. For AC data the corresponding values are 2.72%, 0.11%, and 1.48%. Besides, their memory usage increases slowly with k and rises sublinearly with N. Even at $d = 5$, their memory usage is only two times of that at $d = 3$.

We also investigated the effect of k on the cost of the algorithms (see Figure 12). In some tested cases of Figure 12a, the cost of ITD is too high for the corresponding bar to fit in the diagram; in these cases the bar is marked with a "≈" sign and the actual cost is explicitly given. Observe that LCG and CBT outperform ITD in all cases. As k increases, ITD performs more constrained skyline queries, leading to more counting operations on retrieved points. CBT has lower cost than LCG for UI data because CBT accesses each tree node at most once. For CO data, counting operations in LCG become very efficient and thus LCG and CBT have similar costs. On the other hand, for AC data, there is a wide performance gap between LCG and CBT.

Figure 13 plots the cost of the algorithms as a function of the data dimensionality d. Again, ITD is inferior to its competitors for most of the cases. As d increases, the number of skyline points increases rapidly but the number of points examined by LCG/CBT increases at a slower rate. Again, CBT has lower cost than LCG for all cases. Figure 14 investigates the effect of the data size N on the cost of the algorithms. When N increases, the number of skyline points increases considerably and ITD performs much more batch counting operations than LCG. Also, the performance gap between

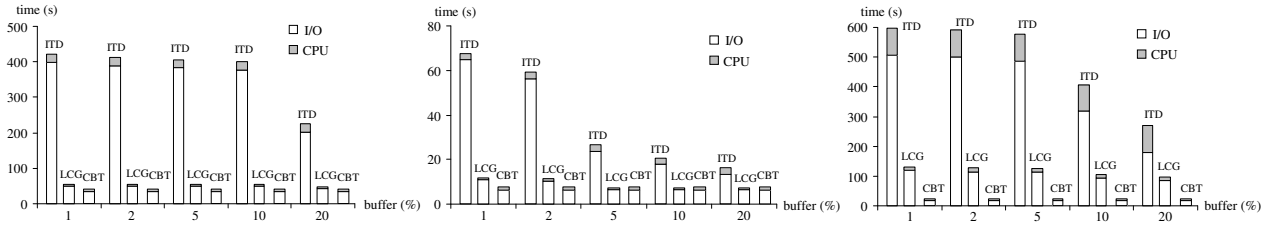

(a) UI (b) CO (c) AC

Figure 11: Cost vs. buffer size (%), $k = 16$, N =1M, $d = 3$

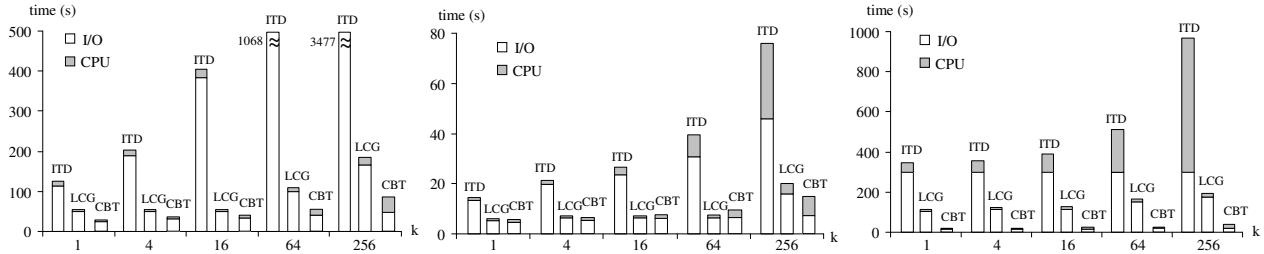

(a) UI (b) CO (c) AC

Figure 12: Cost vs. k, N =1M, $d = 3$

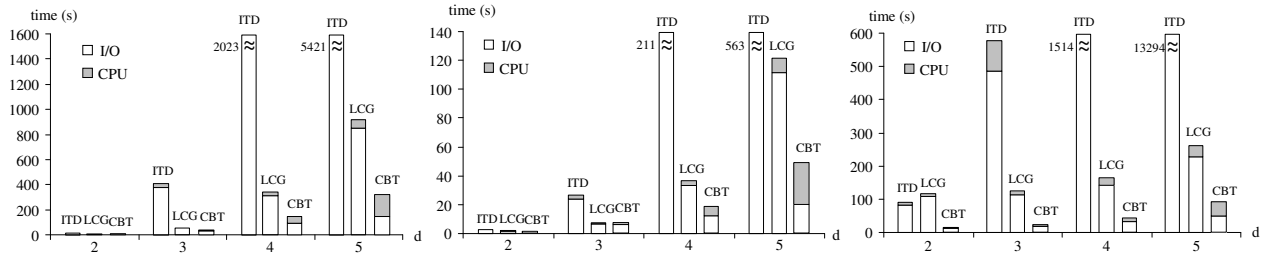

(a) UI (b) CO (c) AC

Figure 13: Cost vs. d, N =1M, $k = 16$

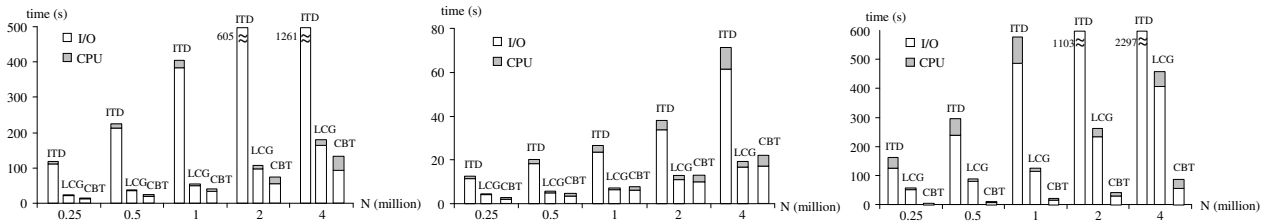

(a) UI (b) CO (c) AC

Figure 14: Cost vs. N, $d = 3$, $k = 16$

LCG and CBT widens.

Finally, we investigate the performance of the proposed algorithms for bichromatic top-k dominating queries. In this experiment, each *provider* dataset $\mathcal{D}_P$ (*consumer* dataset $\mathcal{D}_A$) contains 1 million points in a 3-dimensional space, k is set to 16 and the LRU memory buffer size is fixed to 5% of the sum of both tree sizes. Figure 15 illustrates the cost of the algorithms for different combinations of $\mathcal{D}_P$ and $\mathcal{D}_A$. For instance, the column UI/CO represents the combination that $\mathcal{D}_P$ is a UI dataset and $\mathcal{D}_A$ is a CO dataset. The least expensive case is CO/CO because few points are examined in $\mathcal{D}_P$ and the counting cost on $\mathcal{D}_A$ is low. On the other hand, the case AC/UI is the most expensive as many points need to be examined in $\mathcal{D}_P$ and the counting cost on $\mathcal{D}_A$ is also high. Observe that LCG and CBT outperform ITD in all cases. Except the cases CO/CO and CO/AC where LCG and CBT have similar costs, CBT outperforms LCG by a wide margin in the other 7 cases. In summary, for both monochromatic and bichromatic top-k dominating queries, CBT is the best algorithm, while in only few cases (for correlated datasets) its performance is similar to LCG.

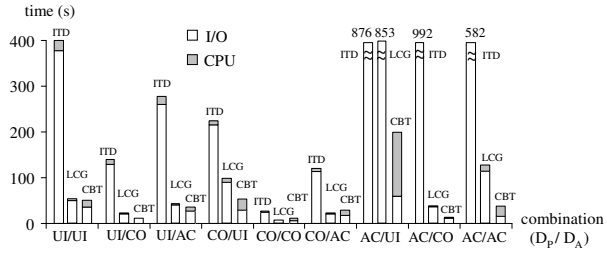

Figure 15: Bichromatic queries, $k = 16$, $N_P = N_A =$ 1M , $d = 3$

7.2 Experiments With Real Data

Datasets We experimented with three real multidimensional datasets: *FC*[5], *NBA*[6], and *BASEBALL*[7]. *FC* contains 581012 forest land cells (i.e., data objects), having four attributes: horizontal distance to hydrology (*hh*), vertical distance to hydrology (*vh*), horizontal distance to roadways (*hr*), and horizontal distance to fire points (*hf*). For *FC*, small values are preferable to large ones at all dimensions. *NBA* contains regular season statistics of 19112 NBA players (i.e., data objects). In order for the query to be meaningful, only few important attributes are selected for NBA players: games played (*gp*), points (*pts*), rebounds (*reb*), and assists (*ast*). *BASEBALL* consists of statistics of 36898 baseball pitchers (i.e., data objects). Similarly, few important attributes are chosen for baseball pitchers: wins (*w*), games (*g*), saves (*sv*), and strikeouts (*so*). In the last two datasets, large values are preferable for all dimensions and each player is uniquely identified by his/her name and year.

Performance Experiment Table 3 shows the cost of the algorithms on two largest datasets (*FC* and *BASEBALL*) for different values of k, by fixing the buffer size to 5% of the tree size. Observe that the cost of ITD becomes prohibitively expensive at high values of k. Clearly, CBT has the lowest cost and the performance gap between the algorithms widens as k increases.

	time (seconds)					
	FC			*BASEBALL*		
k	**ITD**	**LCG**	**CBT**	**ITD**	**LCG**	**CBT**
1	262.3	162.0	62.0	4.6	13.0	0.9
4	413.0	166.6	69.7	9.4	16.5	1.8
16	814.2	204.2	78.9	22.8	18.4	2.5
64	2772.7	282.2	99.4	69.7	22.8	3.5
256	9942.1	523.0	176.4	271.1	38.6	5.9

Table 3: Query cost vs. k, real datasets

Meaningfulness of top-k dominating query results Table 4 shows the dominating scores and the attribute values of the top-5 dominating players in the *NBA* and *BASEBALL* datasets. Readers familiar with these sports can easily verify that the returned results match the public view of super-star players. Although the ranking of objects by their μ-scores may not completely match with every personalized ranking suggested by individuals, a top-k dominating query at least enables them to discover some representative "top" objects without any specific domain knowledge. In addition, we note that some of the top-k results do not belong to the skyline. For example, the *NBA* player "Kevin Garnett / 2002" is the top-3 result, even though it is dominated by the top-1 result (i.e., not a skyline point). Similarly, the top-4 *BASEBALL* pitcher is dominated by the top-2. These players could not be identified by skyline queries.

Score	NBA Player / Year	gp	pts	reb	ast
18585	Wilt Chamberlain / 1967	82	1992	1952	702
18299	Billy Cunningham / 1972	84	2028	1012	530
18062	Kevin Garnett / 2002	82	1883	1102	495
18060	Julius Erving / 1974	84	2343	914	462
17991	Kareem Abdul-Jabbar / 1975	82	2275	1383	413

Score	BASEBALL Pitcher / Year	w	g	sv	so
34659	Ed Walsh / 1912	27	62	10	254
34378	Ed Walsh / 1908	40	66	6	269
34132	Dick Radatz / 1964	16	79	29	181
33603	Christy Mathewson / 1908	37	56	5	259
33426	Lefty Grove / 1930	28	50	9	209

Table 4: Top-5 dominating players

[5]Forest cover dataset, UCI KDD Archive. http://kdd.ics.uci.edu

[6]NBA Statistics v2.0. http://basketballreference.com/stats_download.htm

[7]The Baseball Archive v5.3. http://baseball1.com/statistics/

The next experiment compares the μ-score distribution of top-k dominating points and skyline points. For this, we first retrieve the skyline points and then compute top-k dominating points, by setting k to the number of skyline points (69 for *NBA* and 50 for *BASEBALL*). Figure 16 plots the scores of top-k dominating points and skyline points in the descending order. Clearly, top-k dominating points have much higher scores than skyline points, especially for the *BASEBALL* dataset. This indicates that the top-k dominating points may be more informative regarding the popularity of the players, when compared to skyline points. In addition, since the skyline sizes of both datasets cannot be controlled by users, the skyline results could appear too many or too few to the user. Finally, the skyline query does not provide the user an ordering of importance of the results, which could facilitate their post-processing.

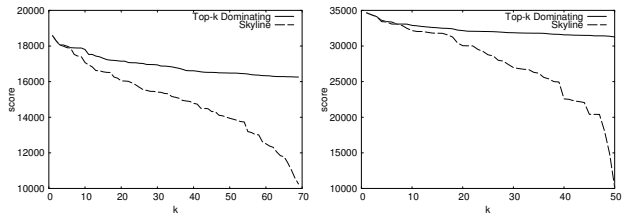

(a) NBA (b) BASEBALL

Figure 16: Score distribution in real datasets

8 Discussion

In this section, we discuss the limitations of alternative solutions for processing top-k dominating queries, based on dominating area computation or materialization. In addition, we study the evaluation of top-k dominating queries on non-indexed data.

8.1 Alternative Solutions

In order to gain a deeper understanding of the problem, we first explore straightforward solutions and discuss why they are either infeasible or inefficient. A simple approach for computing the top-k dominating points could be to find $K \gg k$ points with the highest *dominating area* first and then compute their scores in order to obtain the final result. The dominating area [15] of a point p is defined as $\Phi(p) = \prod_{i=1}^{d}(1 - p[i])$, assuming that the data domain is $[0, 1]^d$. This solution provides a good approximation of the actual results for uniform datasets. Nevertheless, if the data distribution is not uniform, the dominating area of a point p is not necessarily proportional to $\mu(p)$. Therefore, it is hard to set a bound K that guarantees no result loss. In addition, recall that the score $\mu(p)$ of a point p is not only decided by its own coordinates but also by other points in the dataset. Unlike measuring the dominating area, computing $\mu(p)$ (or even its close approximation) requires accesses to the data (i.e., incurs I/O cost).

Another approach is to materialize the scores for every point in the dataset such that top-k dominating queries can be processed very fast. However, the maintenance cost of the scores can be very high. Suppose that we have pre-computed the scores for all the points of Figure 4. If a point (e.g., p_7) is deleted from the dataset, the score of any point (e.g., p_1, p_2, p_4) dominating it must be decremented. Similarly, when a new point p' is inserted to the dataset, $\mu(p')$ needs to be computed and the score of any point dominating p' must be incremented. Since a single update may lead to score modifications for many points (which may be expensive to find), this approach is not appropriate for dynamic datasets.

8.2 Algorithms for Non-indexed Data

We now discuss how top-k dominating queries can be evaluated efficiently on non-indexed data. To ease our discussion, we assume that Y memory pages are available, the page capacity is B and the dataset contains N tuples.

The basic block-nested-loop join algorithm can be adapted to compute the scores of all data points and then return top-k results. However, this method requires $\frac{N}{B} \cdot (1 + \frac{N}{YB})$ page accesses and does not scale well for large datasets. A better approach would be to adapt the skyline-based solution in Section 3.2 for non-indexed data. This solution is composed of two main operations: (i) finding the skyline (or constrained skyline), and (ii) counting the scores of the retrieved points (in batch). The skyline operation can be implemented by LESS [11], the state-of-the-art external memory skyline algorithm on non-indexed data. The counting operation, implemented by scanning the dataset, can be performed in batches of every YB points (i.e., available memory). Although this approach only computes the scores for skyline points (and constrained skyline points), its performance deteriorates for datasets with large skyline. Another disadvantage is that the worst case I/O-cost of LESS is still quadratic to N, according to [11].

We assert that the best solution is to first bulk-load an aR-tree (e.g., using the algorithm of [18]) from the dataset and then compute top-k dominating points by our algorithms. Bulk-loading requires externally sorting the points at $\frac{N}{B} \cdot (2 + 2 \cdot \lceil \log_{Y-1} \frac{N}{YB} \rceil)$ disk page accesses, which scales well for large datasets. As we showed in Section 7, CBT and LCG (our best algorithms) are scalable.

9 Conclusion

In this paper, we studied the interesting and important problem of processing top-k dominating queries on indexed multi-dimensional data. Although the skyline-based algorithm in [23] is applicable to the problem, it suffers from poor performance, as it unnecessarily examines many skyline points. This motivated us to develop carefully-designed solutions that exploit the intrinsic properties of the problem for accelerating query evaluation. First, we proposed ITD, which integrates the algorithm of [23] with our optimization techniques (batch counting and Hilbert ordering). Next, we developed LCG, a top-k dominating algorithm that guides search by computing upper bound scores for non-leaf entries, and utilizes a lightweight (i.e., I/O-inexpensive) technique for computing upper bound scores. Then, we proposed I/O efficient algorithm CBT that accesses each node at most once. The effectiveness of our optimizations (lightweight counting technique in LCG and traversal order in CBT) was analyzed theoretically.

Our experimental study suggests that LCG and CBT are the best algorithms, typically being several times faster than ITD; a method that is already significantly faster than the naive skyline-based approach of [23]. LCG and CBT scale well with the buffer size, k, and the data size N. In addition, they scale better than ITD with the data dimensionality d. Nevertheless, the performance of LCG varies according to the data distribution. CBT outperforms LCG for uniform and anti-correlated data, while the two methods have similar cost for correlated data. As CBT has stable performance across different data distributions, its usage for top-k dominating queries is recommended. Our experiments on real datasets demonstrate that top-k dominating queries may deliver more useful results than skyline queries. Finally, for the first time in the literature, we defined and studied interesting variants of top-k dominating queries; queries with arbitrary aggregate functions (as opposed to `COUNT`), queries where points carry weights of importance (as opposed to all points having the same importance), and bichromatic top-k dominating queries (where dominance is counted on another dataset).

In the future, we plan to develop specialized algorithms for non-indexed data that rely on spatial hashing. Finally, we intend to devise cheap, approximate techniques that compute the top-k dominating set with some error guarantee.

10 References

[1] W.-T. Balke, U. Güntzer, and J. X. Zheng. Efficient Distributed Skylining for Web Information Systems. In *EDBT*, 2004.

[2] S. Börzsönyi, D. Kossmann, and K. Stocker. The Skyline Operator. In *ICDE*, 2001.

[3] A. R. Butz. Alternative Algorithm for Hilbert's Space-Filling Curve. *IEEE Trans. Comput.*, C-20(4):424–426, 1971.

[4] C.-Y. Chan, P.-K. Eng, and K.-L. Tan. Stratified Computation of Skylines with Partially-Ordered Domains. In *SIGMOD*, 2005.

[5] C.-Y. Chan, H. Jagadish, K.-L. Tan, A. Tung, and Z. Zhang. Finding k-Dominant Skylines in High Dimensional Space. In *SIGMOD*, 2006.

[6] C.-Y. Chan, H. Jagadish, K.-L. Tan, A. Tung, and Z. Zhang. On High Dimensional Skylines. In *EDBT*, 2006.

[7] S. Chaudhuri, N. Dalvi, and R. Kaushik. Robust Cardinality and Cost Estimation for Skyline Operator. In *ICDE*, 2006.

[8] J. Chomicki, P. Godfrey, J. Gryz, and D. Liang. Skyline with Presorting. In *ICDE*, 2003.

[9] R. Fagin, A. Lotem, and M. Naor. Optimal Aggregation Algorithms for Middleware. In *PODS*, 2001.

[10] P. Godfrey. Skyline Cardinality for Relational Processing. In *FoIKS*, 2004.

[11] P. Godfrey, R. Shipley, and J. Gryz. Maximal Vector Computation in Large Data Sets. In *VLDB*, 2005.

[12] A. Guttman. R-Trees: A Dynamic Index Structure for Spatial Searching. In *SIGMOD*, 1984.

[13] G. R. Hjaltason and H. Samet. Distance Browsing in Spatial Databases. *TODS*, 24(2):265–318, 1999.

[14] V. Hristidis, N. Koudas, and Y. Papakonstantinou. PREFER: A System for the Efficient Execution of Multiparametric Ranked Queries. In *SIGMOD*, 2001.

[15] Z. Huang, C. S. Jensen, H. Lu, and B. C. Ooi. Skyline Queries Against Mobile Lightweight Devices in MANETs. In *ICDE*, 2006.

[16] D. Kossmann, F. Ramsak, and S. Rost. Shooting Stars in the Sky: An Online Algorithm for Skyline Queries. In *VLDB*, 2002.

[17] I. Lazaridis and S. Mehrotra. Progressive Approximate Aggregate Queries with a Multi-Resolution Tree Structure. In *SIGMOD*, 2001.

[18] S. T. Leutenegger, J. M. Edgington, and M. A. Lopez. STR: A Simple and Efficient Algorithm for R-Tree Packing. In *ICDE*, 1997.

[19] C. Li, K. C.-C. Chang, and I. F. Ilyas. Supporting Ad-hoc Ranking Aggregates. In *SIGMOD*, 2006.

[20] C. Li, B. C. Ooi, A. Tung, and S. Wang. DADA: A Data Cube for Dominant Relationship Analysis. In *SIGMOD*, 2006.

[21] X. Lin, Y. Yuan, W. Wang, and H. Lu. Stabbing the Sky: Efficient Skyline Computation over Sliding Windows. In *ICDE*, 2005.

[22] D. Papadias, P. Kalnis, J. Zhang, and Y. Tao. Efficient OLAP Operations in Spatial Data Warehouses. In *SSTD*, 2001.

[23] D. Papadias, Y. Tao, G. Fu, and B. Seeger. Progressive Skyline Computation in Database Systems. *TODS*, 30(1):41–82, 2005.

[24] J. Pei, W. Jin, M. Ester, and Y. Tao. Catching the Best Views of Skyline: A Semantic Approach Based on Decisive Subspaces. In *VLDB*, 2005.

[25] K.-L. Tan, P.-K. Eng, and B. C. Ooi. Efficient Progressive Skyline Computation. In *VLDB*, 2001.

[26] Y. Tao, X. Xiao, and J. Pei. SUBSKY: Efficient Computation of Skylines in Subspaces. In *ICDE*, 2006.

[27] Y. Theodoridis and T. K. Sellis. A Model for the Prediction of R-tree Performance. In *PODS*, 1996.

[28] Y. Yuan, X. Lin, Q. Liu, W. Wang, J. X. Yu, and Q. Zhang. Efficient Computation of the Skyline Cube. In *VLDB*, 2005.