



US 20080052268A1

(19) **United States**

(12) **Patent Application Publication**
KOUDAS et al.

(10) **Pub. No.: US 2008/0052268 A1**

(43) **Pub. Date: Feb. 28, 2008**

(54) **METHOD AND APPARATUS FOR
OPTIMIZING QUERIES UNDER
PARAMETRIC AGGREGATION
CONSTRAINTS**

Related U.S. Application Data

(62) Division of application No. 10/828,839, filed on Apr. 21, 2004.

(60) Provisional application No. 60/464,256, filed on Apr. 21, 2003.

(76) Inventors: **NIKOLAOS KOUDAS**, Springfield, NJ (US); **Divesh Srivastava**, Summit, NJ (US); **Sudipto Guha**, Philadelphia, PA (US); **Dimitrios Gunopulos**, Irvine, CA (US); **Michail Vlachos**, Elmsford, NY (US)

Publication Classification

(51) **Int. Cl.**
G06F 17/30 (2006.01)

(52) **U.S. Cl.** **707/2; 707/E17**

Correspondence Address:


AT&T CORP.
ROOM 2A207
ONE AT&T WAY
BEDMINSTER, NJ 07921 (US)

(57) **ABSTRACT**

The present invention relates to a method and apparatus for optimizing queries. The present invention discloses an efficient method for providing answers to queries under parametric aggregation constraints.

(21) Appl. No.: **11/927,100**

(22) Filed: **Oct. 29, 2007**






FIG. 1

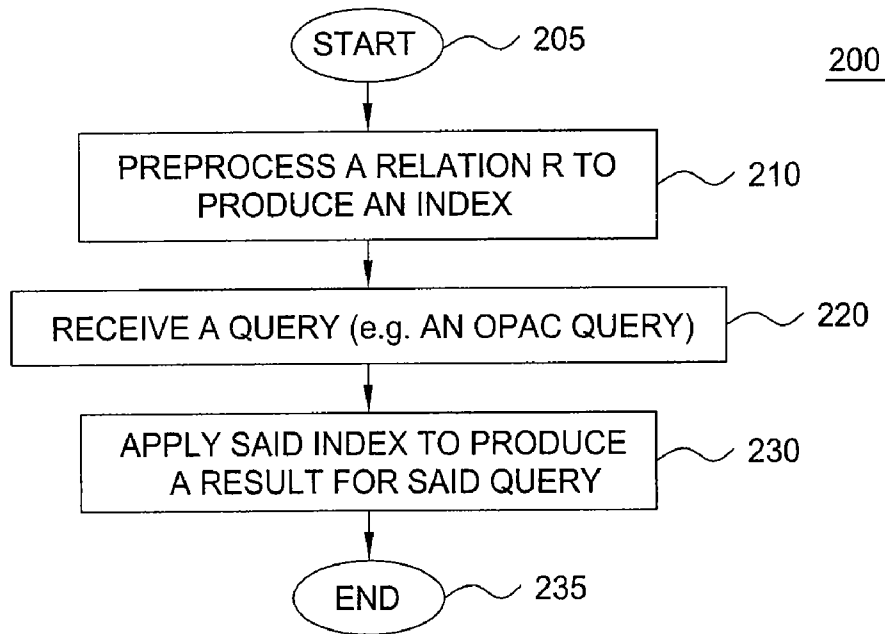


FIG. 2

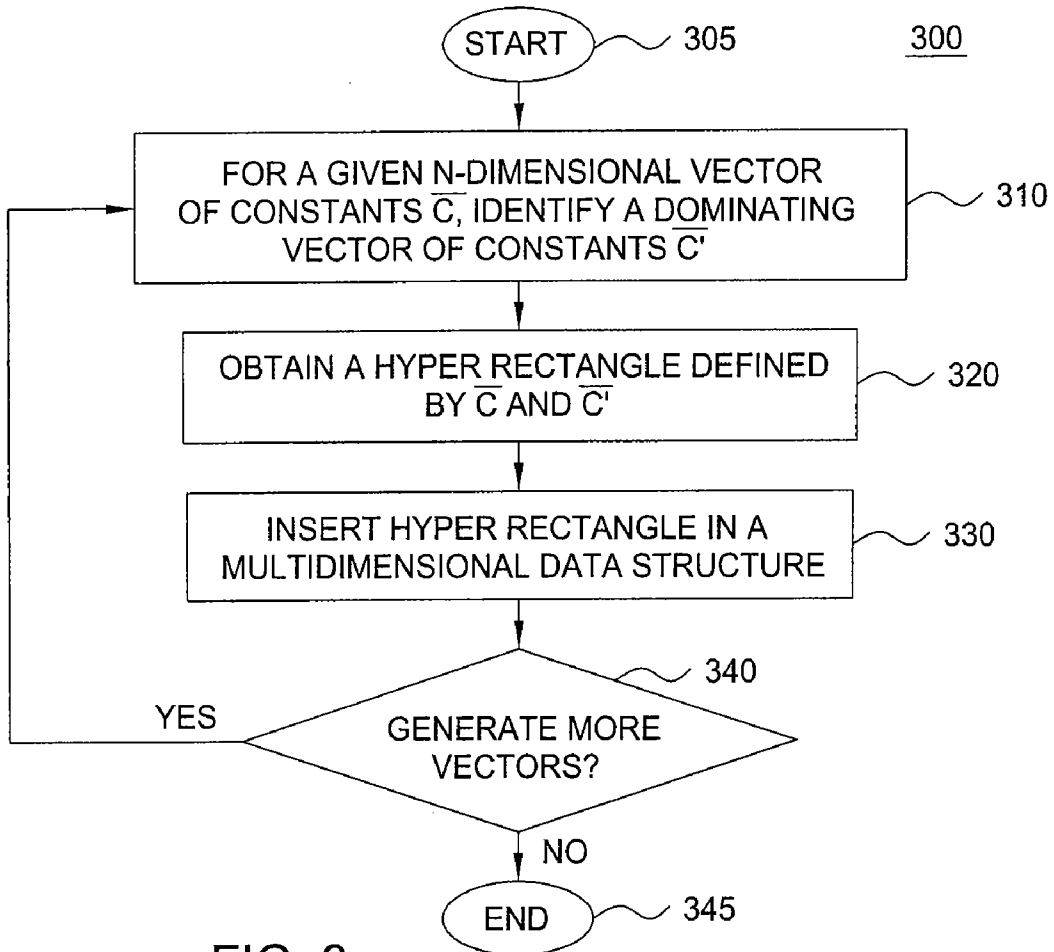


FIG. 3

Algorithm **GeneratePartitions**(ϵ, ϵ', D)

Initialize:

Q: Queue of multidimensional constraint vectors

\mathcal{R} : R-tree

s, c, c' : constraint vectors

each coordinate of s is initially set to be equal to D and, and s is added to Q

- (1) while Q not empty
- (2) $\bar{c} = \text{headof}(Q)$
- (3) $(r, \bar{C}, p, S) = \text{LocateSolution}(\bar{c})$
- (4) if there is no rectangle r' in the R-tree \mathcal{R} that contains rectangle r and r not NULL
- (5) Insert (r, p, \bar{C}, S) to the R-tree \mathcal{R} by storing (r, p, \bar{C}) in a leaf index entry and maintaining a pointer to the set of tuple identifiers in the solution S on disk
- (6) CreateFront(Q,r)
- (7) endif
- (8) end-while

Algorithm **LocateSolution**(\bar{c})

Input: constant vector $\bar{c} = (c_1, \dots, c_n)$

Output: (r, \bar{C}, p, S)

- (1) $(p, S) = \text{MultiKnapsack}(\bar{c})$
- (2) if (S is NULL) return (NULL, NULL, 0, NULL)
- (3) for $i = 1$ to n
- (4) $c'_i = \frac{c_i}{1+\epsilon}$
- (5) $(p', S') = \text{MultiKnapsack}(\bar{c}')$
- (6) if (S' is NULL) return (NULL, NULL, 0, NULL)
- (7) if $((1 + \epsilon')p' > p)$
- (8) while $(p \geq \frac{p'}{1+\epsilon'})$
- (9) $\bar{c}_t = \bar{c}'; p_t = p'; S_t = S'$
- (10) for $i = 1$ to n
- (11) $c'_i = \frac{c'_i}{1+\epsilon'}$
- (12) $(p', S') = \text{MultiKnapsack}(\bar{c}')$
- (13) end-while
- (14) return (FormRect(\bar{c}_t, \bar{c}), \bar{c}_t, p_t, S_t)
- (15) else
- (16) return (FormRect(\bar{c}', \bar{c}), \bar{c}, p, S)

FIG. 4




FIG. 5





FIG. 6

FIG. 7



METHOD AND APPARATUS FOR OPTIMIZING QUERIES UNDER PARAMETRIC AGGREGATION CONSTRAINTS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a divisional of co-pending U.S. patent application Ser. No. 10/828,839, filed Apr. 21, 2004, which claims benefit of U.S. provisional patent application Ser. No. 60/464,256, filed Apr. 21, 2003, each of the above referenced applications is herein incorporated by reference.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] Embodiments of the present invention generally relate to a method for optimizing queries. More specifically, the present invention discloses an efficient method for providing answers to queries under parametric aggregation constraints.

[0004] 2. Description of the Related Art

[0005] In today's rapidly changing business landscape, corporations increasingly rely on databases to help organize, manage and monitor every aspect of their business. Databases are deployed at the core of important business operations, including Customer Relationship Management, Supply Chain Management, and Decision Support Systems. The increasing complexity of the ways in which businesses use databases creates an ongoing demand for sophisticated query capabilities.

[0006] Novel types of queries seek to enhance the way information is utilized, while ensuring that they can be easily realized in a relational database environment without the need for significant modifications to the underlying relational engine. However, as the size of databases continues to grow, coupled with the desire by users to formulate complex queries, traditional methods of responding to queries require a tremendous amount of computational cycles. This leads to delay and inefficiencies in the use of databases.

[0007] Therefore, there is a need for a method to efficiently provide answers to complex queries.

SUMMARY OF THE INVENTION

[0008] In one embodiment, the present invention generally relates to a method and apparatus for optimizing queries. The present invention discloses an efficient method for providing answers to queries under parametric aggregation constraints.

[0009] More specifically, the present invention discloses a method that preprocesses relations and constructs indices to efficiently provide answers to queries. e.g., optimization under parametric aggregation (OPAC) queries. The answers returned by the indices are approximate, not exact, (e.g., within an approximation factor) and provide guarantees for their accuracy. Moreover, the indices can be tuned easily to meet desired accuracy levels, providing a graceful tradeoff between answer accuracy and index space.

[0010] For example, a user may submit a query having aggregation constraints, where the user is willing to accept answers within a selectively defined approximation factor.

In doing so, the results can be generated quickly without having to solve a multi-attribute knapsack problem from scratch for every query having aggregation constraints.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] So that the manner in which the above recited features of the present invention can be understood in detail, a more particular description of the invention, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0012] FIG. 1 illustrates an example of an ϵ -Pareto set;

[0013] FIG. 2 illustrates a method of the present invention for efficiently providing a result to a query having aggregation constraints, e.g., an OPAC query;

[0014] FIG. 3 illustrates a method for preprocessing a relation R;

[0015] FIG. 4 illustrates the preprocessing method of the present invention in pseudo code format;

[0016] FIG. 5 illustrates the operation of GeneratePartitions of the present invention;

[0017] FIG. 6 presents an example of the partitioning generated by algorithm GeneratePartitions of the present invention; and

[0018] FIG. 7 illustrates a block diagram of the present query optimization method being implemented on a general purpose computer.

[0019] To facilitate understanding, identical reference numerals have been used, wherever possible, to designate identical elements that are common to the figures.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0020] In one embodiment, the present invention addresses a class of queries that is referred to as OPAC (optimization under parametric aggregation constraints) queries. Such queries aim to identify sets of database tuples that constitute solutions of a large class of optimization problems involving the database tuples. To illustrate this important class of queries, consider the following simple example.

[0021] Consider a large distributor of cables, who maintains a database relation R keeping track of the products in stock. Cable manufacturers ship their products in units, each having a specific weight and length. Assume that relation R has attributes uid (a unit identifier), manufacturer, weight, length and price, associated with each cable unit. A sample relation R is depicted in Table 1.

TABLE 1

Uid	Manufacturer	Weight	Length	Price
1	Optical Co.	30	40	50
2	Optical Co.	20	50	50
3	Optics Inc.	30	70	80

TABLE 1-continued

Uid	Manufacturer	Weight	Length	Price
4	Opticom Co.	20	20	10
5	Optics Inc.	20	20	20

[0022] Commonly, “queries” select cable units by imposing constraints on the total length and total weight of the units they are interested in, while optimizing on total price. Thus, the desired result is a set of tuples collectively meeting the imposed aggregate constraints and satisfying the objective function. Note that this is considerably different from selecting cable units (tuples) based on their individual attribute values.

[0023] For example, one query could request the set of cable units having the smallest total price, with total length no less than $L_c=90$ and total weight no less than $W_c=50$. A straight-forward solution to this query involves computing the total weight and length of each possible subset of cable units in R, identifying those that respect the constraints on length and weight, and returning the one with the lowest price. Clearly, such a brute force evaluation strategy is not desirable. In the example of Table 1, the answer set for this query would be {uid₂, uid₄, uid₅}, with a total price of 80.

[0024] A different query could seek to maximize the total price for a number of cable units requested, of total length no more than $L_c=90$ and of total weight no more than $W_c=50$. In this case, the answer set for this query would be {uid₁, uid₂} or {uid₃, uid₅} each with a total price of 100.

[0025] Finally, observe that L_c and W_c are parameters of these two OPAC queries, and different users may be interested in these queries, but with different values specified for each of these parameters.

[0026] Instances of OPAC queries are ubiquitous in a variety scenarios, including simple supplier-buyer scenarios (as illustrated by the present example), that use relational data stores. These queries can be easily generalized to more complex scenarios involving Business to Business interactions in an electronic marketplace. Any interaction with a database, requesting a set of tuples as an answer, specifying constraints over aggregates of attributes values, seeking to optimize aggregate functions on some measure attribute in the result set, is an instance of an OPAC query.

[0027] OPAC queries have a very natural mathematical interpretation. In particular, these queries represent instances of optimization problems with multiple constraints, involving the tuples and attributes of a database relation. Before describing the invention in greater detail, a series of mathematical definitions are presented below.

[0028] Let $R(A_1, \dots, A_n, P)$ be a relation, with attributes A_1, \dots, A_n, P . Without loss of generality assume that all attributes have the same domain. Denote by S a subset of the tuples of R and by $S_{A_i}, 1 \leq i \leq n$ and S_p , the (multiset of) values of attribute $A_i, 1 \leq i \leq n$ and P in S, respectively. Let $f_i, 1 \leq i \leq n$ and F denote aggregate functions (e.g., sum, max). The present invention considers atomic aggregation constraints of the form $f_i(S_{A_i}) \theta c_i$, where θ is an arithmetic comparison operator (e.g., \leq, \geq), and c_i is a constant, and complex aggregation constraints that are boolean combina-

tions of atomic aggregation constraints. The present invention refers to them collectively as aggregation constraints, denoted by ψ .

[0029] Definition 1: (General OPAC Query Problem): Given a relation $R(A_1, \dots, A_n, P)$, a general OPAC query Q specifies (i) a parametric aggregation constraint $\Psi_{\vec{v}}$, (ii) an aggregate function F, with optimization objective m (min or max), and (iii) a vector of constants \vec{c} . It returns a subset S of tuples from R as its result, such that (i) $\Psi_{\vec{v}=\vec{c}}(S_{A_1}, \dots, S_{A_n}) = \text{TRUE}$, and (ii) $(\forall S' \subseteq R, \Psi_{\vec{v}=\vec{c}}(S'_{A_1}, \dots, S'_{A_n}) = \text{TRUE}) \Rightarrow (F(S'_p) \leq_m F(S_p))$.

[0030] Intuitively, the result of a general OPAC query Q is a subset S of tuples of R that satisfy the parametric aggregation constraint $\Psi_{\vec{v}}$ (with the parameters \vec{v} instantiated to the vector of constants \vec{c}), such that its aggregate objective function is optimal (i.e., maximal under \leq_m) among all subsets of R that satisfy the (instantiated) parametric aggregation constraint.

[0031] It is evident that the result of a general OPAC query involves the solution of an optimization problem involving a (potentially) complex aggregation constraint on relation R. Depending on the specifics of the aggregate functions f_i, F , the nature of the aggregation constraint, and the optimization objective, different instances of the OPAC query problem arise.

[0032] The present invention considers the instance of the problem when the aggregate functions f_i, F return the sum of the values in their input multisets, the aggregation constraints are conjunctions of atomic aggregation constraints of the form $f_i(S_{A_i}) \leq c_i$, and the objective function seeks to maximize $F(S_p)$.

[0033] This formulation of an OPAC query gives rise to an optimization problem, namely the “multi-attribute knapsack” problem. Given this relationship between the specific form of the OPAC query and the multi-attribute knapsack problem, the present invention will refer to values of the function $F(S_p)$ as the “profit” for the set of tuples S. It is known that solving the knapsack problem, even in the simple instance involving a constraint on only one attribute (e.g.,

$$\sum_{x_i \in S_{A1}} x_i \leq c_1$$

and maximize

$$\sum_{x_j \in S_p} x_j$$

is NP-complete. However, this problem is solvable in pseudo-polynomial time with dynamic programming. For example, the pseudo-polynomial algorithm solving the knapsack problem in the single attribute case can serve as a basis for a solution of the multi-attribute problem as well. In particular, one could generate all solutions for one attribute, and pick the solution S that maximizes $F(S_p)$ among all solutions that satisfy the constraints on all attributes. The

form of the solution that is reported could vary; for example, the solution could be the set of tuple identifiers from R.

[0034] It is evident that every OPAC query Q determines an instance of a multi-attribute knapsack problem on relation R. Since the relation R can be very large, in the general case, solving the multi-attribute knapsack problem from scratch every time an OPAC query is posed is not at all pragmatic. Such an approach would be far from being interactive and, more importantly, it would be entirely DBMS agnostic, missing the opportunity to utilize the underlying DBMS infrastructure for query answering.

[0035] In the present invention, the Pareto (or, dominating) set is the set of optimal solutions that are mutually incomparable, because improving one objective would lead to a decrease in another. In the present setting, the present invention considers a single optimization objective, but the present invention allows the user to dynamically specify the aggregation constraint parameters. Thus, the present invention can adapt the Pareto framework to the OPAC query problem.

[0036] Definition 2 (Pareto Set): The Pareto set P for an OPAC query defined on a relation R is the set of pairs $\{(\bar{c}, S)\}$ of all n-dimensional vectors $\bar{c}=(c_1, \dots, c_n)$ and associated solutions S, such that (a) there exists a solution $S \subseteq R$ with $f_i(S_{A_i})=c_i, 1 \leq i \leq n$, and (b) there is no other pair (\bar{c}', S') , such that $f_i(S'_{A_i})=c'_i, c'_i \leq c_i, 1 \leq i \leq n$ and $F(S'_P) > F(S_P)$.

[0037] This approach produces useful solutions where they are optimal both in terms of the parameters realizing them and the profit obtained. For any element (\bar{c}, S) of the Pareto set, there is no other solution with higher profit achieved by parameters at most as large in all dimensions as \bar{c} . Identifying such a set would be very informative as it contains valuable information about maximal profits.

[0038] The concept of Pareto points will be illustrated in an example with reference to Table 2 and FIG. 1. Specifically, given the relation S in Table 2, the Pareto points are the round points in FIG. 1.

TABLE 2

Relation S	a1	a2	Profit
t1	9	11	100
t2	11	9	100
t3	4	4	20

[0039] For example, (24, 24) is a Pareto point with profit 220, realized by the entire set of tuples. The vector (13,15) with profit 120 is also a Pareto point, the corresponding set of tuples being $\{t_1, t_3\}$, because no vector (c_1, c_2) with $c_1 \leq 13$ and $c_2 \leq 15$ has profit more than 120.

[0040] The notion of Pareto sets are defined for arbitrary classes of constraint problems and functions, not only for the multi-attribute knapsack. The constraint problems can be discrete and, in most such cases, the Pareto set can have exponentially many elements. This happens because linear programs can be posed in the Pareto framework, and the convex hull of the solution space for linear programs can have exponentially many (in the number of objects/variables) vertices.

[0041] The size of the Pareto set for an instance of the multi-attribute knapsack problem can be exponential in the

number of tuples, even if the number of attributes is a small constant. Consider, for example, the case where there is only one attribute, A_1 and a profit attribute P, and tuple i in relation R has the form $(2^i, 2^i)$. In this case, any subset of the tuples in R defines a unique cost and profit vector and no other set can achieve at least as small a cost and a higher profit. Therefore, all the subsets of tuples define dominating points.

[0042] To circumvent this problem, one can approximate the Pareto set. The ϵ -Pareto set, is a set of “almost” optimal solutions defined as: for every optimal solution c, the ϵ -Pareto set contains a solution that optimizes each of the optimization criteria within a fraction of ϵ (e.g., an approximation factor or an accuracy factor). It has been shown that the ϵ -Pareto set for a multi-objective knapsack problem can be computed efficiently and it is polynomial in size.

[0043] Given a relation R, functions f_i and F, and $\epsilon > 0$, the ϵ -Pareto set, P_ϵ is a set of solutions that almost dominate any other solution.

[0044] Definition 3 (ϵ -Pareto set): The ϵ -Pareto set for an OPAC query is a set of pairs $\{(\bar{c}, S)\}$ of n-dimensional vectors $\bar{c}=(c_1, \dots, c_n)$ and solutions S, such that, (a) there exists a solution $S \subseteq R$ with $f_i(S_{A_i}) \leq c_i, 1 \leq i \leq n$, and (b) there is no other pair (\bar{c}', S') , such that $f_i(S'_{A_i}) \leq c'_i, 1 \leq i \leq n, c'_i \leq (1+\epsilon) c_i, 1 \leq i \leq n$ and $F(S'_P) > (1+\epsilon) F(S_P)$.

[0045] The ϵ -Pareto set is illustrated as an example in FIG. 1. If ϵ is 0.25, the set of the rectangular points in FIG. 1 is an ϵ -Pareto set. For example, point (13,15) is in the ϵ -Pareto set because there is no vector with coordinates less than $1.25 * (13,15)$ that has profit more than $120 * 1.25$.

[0046] The ϵ -Pareto set is a very useful measure. Assuming that the size of this set is manageable, one could quickly generate an ϵ -Pareto set and utilize it for query answering. For example, one can show the following theorem:

[0047] Theorem 1: The size of the ϵ -Pareto set for an OPAC query instance defined on a relation R is polynomial in $|R|$ (the size of R) and $1/\epsilon$ but may be exponential in the number of attributes.

[0048] Proof: Assume that the n attributes of R are integers. Since f_i and F are polynomial functions (and more specifically sums), the domain of each of these functions cannot be more than $a^{|R|}$ for some constant $a > 1$. One can cover the space of $[1, a^{|R|}]$ with a set of geometrically increasing intervals with step $1+\epsilon$. To cover each domain one need

$$O\left(p\left(\frac{|R|}{\epsilon}\right)\right)$$

intervals (for some polynomial p). Taking the Cartesian product one can get a total of

$$O\left(p\left(\frac{|R|}{\epsilon}\right)^n\right)$$

hyper-rectangles. Clearly, taking one solution from the interior of each hyper-rectangle (if such a solution exists) results in an ϵ -Pareto set.

[0049] The present invention will provide the description of a technique suitable for efficiently answering OPAC queries over a database. In brief, the present method will provide an approximate solution but it will provide guarantees for an acceptable level of accuracy, which will expose useful tradeoffs that will be discussed below.

[0050] FIG. 2 illustrates a method 200 of the present invention for efficiently providing a result to a query having aggregation constraints, e.g., an OPAC query. The method starts in step 205 and proceeds to step 210 where a preprocessing step is performed. Specifically, given a relation R (e.g., a database having a plurality of attributes or fields), the relation R is preprocessed to construct an index. In one embodiment, this step involves evaluating solutions to a multi-attribute knapsack problem on R, for only a select number of vectors of constants. A detailed description in the construction of the index is further provided below.

[0051] In step 220, a query having aggregation constraints is received by method 200. For example, the query can be an OPAC query.

[0052] In step 230, method 200 applies the index in response to the query having aggregation constraints to quickly provides an exact answer or an answer that is “guaranteed accurate” as suitably defined below. Method 200 ends in step 235.

[0053] In one embodiment, the present method is to preprocess relation R, constructing index structures enabling efficient answers to arbitrary OPAC queries. For a query Q, the present invention is intended to provide either the exact answer, or an answer that is “guaranteed accurate”, for suitably defined notions of accuracy. Moreover, the present construction will expose a tradeoff between accuracy and space, thereby providing the flexibility to fine tune the accuracy of the answers. The present invention quantifies the accuracy of answers to an OPAC query below

[0054] Definition 4 (ϵ, ϵ' -Accurate Answers): Let Q be an OPAC query specifying a vector of constants $\bar{c}=(c_1, \dots, c_n)$, having an answer S with profit P. For any $\epsilon, \epsilon'>0$, an ϵ, ϵ' -Accurate answer to Q, is a vector $\bar{c}'=(c'_1, \dots, c'_n)$ and an answer set S', such that $\forall_i, 1 \leq i \leq n, \bar{c}'_i \leq (1+\epsilon)c_i$ and $P'(1+\epsilon')>P$, where P' is the profit of an OPAC query specifying vector \bar{c}' of constants.

[0055] Assume that Q is a query specifying a vector \bar{c} of constants and that the answer to Q is a set $S \subseteq R$ with maximum profit P. An ϵ, ϵ' -accurate answer to Q is an answer set S' that is either the exact answer set S or it is an answer set corresponding to a query Q'. Query Q' specifies a vector of constants having values in each dimension less than or equal to $1+\epsilon$ of the corresponding values specified by Q. Moreover, the profit of Q' is strictly higher than a fraction of $1+\epsilon'$ of P. In the definition, without loss of generality, the present invention assumes the same ϵ fraction is used for all constant values. Different values for ϵ can be used for each of the values, if this is desirable, ϵ being defined as a vector in this case. In fact, the present invention does specify a different approximation factor, ϵ' , for the profit, to differentiate between the aggregate functions f and F.

[0056] As discussed in FIG. 2, the present invention preprocesses relation R, by constructing an index providing ϵ, ϵ' -accurate answers to OPAC queries. In one embodiment, the preprocessing will consist of solving the multi-attribute

knapsack problem exactly, for a select subset of the candidate query space of all possible OPAC queries. The present invention will then utilize these solutions towards providing ϵ, ϵ' -accurate answers to any candidate OPAC query on R.

[0057] This gives rise to the main problem addressed by the present invention. Given a relation R, an OPAC query without the vector of constants ϑ and ϵ, ϵ' , preprocess R constructing an index being able to efficiently provide ϵ, ϵ' -accurate answers to any OPAC query on R that provides the parameters (values) to the constant vector ϑ .

[0058] Consider the example of FIG. 1 again. Assume $\epsilon=\epsilon'=0.25$. Assume that the query (10,18) is given, that is, find a set of objects that satisfy these conditions and maximize the Profit. The set $\{t_1\}$ is an ϵ, ϵ' -accurate answer because it satisfies the constraints, and there is no other set that has higher profit even if the present invention relaxes the constraints by ϵ .

[0059] If the query was (10,10), the set $\{t_1\}$ is again an ϵ, ϵ' -accurate answer. Although the set does not satisfy the query constraints, it satisfies the relaxed constraints $((9, 11) \leq 1.25*(10,10))$, and has the highest profit among all solutions that satisfy these constraints.

[0060] FIG. 3 illustrates a method 300 for preprocessing a relation R. Specifically, for a relation $R(A_1, \dots, A_n, P)$, assume that the range of the Σ function applied on elements of each attribute A, has range $[O \dots D]$. Any candidate query Q specifies an n-dimensional vector of constants $\bar{c} \in [0 \dots D]^n$. The present invention will preprocess the space $[0 \dots D]^n$ of all vectors of constants that can be specified by a possible query, creating a number of partitions that aim to cover the space of all possible queries. The partitions will be constructed in a way such that, for all possible queries inside a partition, one can reason collectively about the properties and values of function F. Moreover, it will allow the present invention to derive an ϵ, ϵ' -accurate answer for any query falling inside a partition.

[0061] Method 300 starts in step 305 and proceeds to step 310. In step 310, given a n-dimensional vector of constants \bar{c} , method 300 identifies a dominating vector of constants \bar{c}' . Specifically, the property between n-dimensional vectors of constants are defined below.

[0062] Let $\bar{c}=(c_1, \dots, c_n), \bar{c}'=(c'_1, \dots, c'_n)$ be two n-dimensional vectors of constants. It is defined that \bar{c} is dominated by \bar{c}' , ($\bar{c} << \bar{c}'$) if $c_i \leq c'_i, 1 \leq i \leq n$.

[0063] The following observation is then made. Let Q, Q' be two queries on R, specifying vectors of constants \bar{c}, \bar{c}' , having result sets S, S' respectively. If $\bar{c} << \bar{c}'$ then

$$F(S_p) = \sum_{x_j \in S_p} x_j \leq F(S'_p) = \sum_{x_j \in S'_p} x_j.$$

[0064] Thus, if a vector of constants \bar{c} is dominated by a vector \bar{c}' , the profit one can achieve for \bar{c} is less than or equal to the profit one can achieve using the vector \bar{c}' . A consequence of the above observation is the following: Consider a sequence of queries, with vectors of constants, $\bar{c}_1 << \bar{c}_2 \dots << \bar{c}_m$. Observing the evolution of the values of F in each answer obtained starting from \bar{c}_m moving towards \bar{c}_1 , func-

tion F is monotonically non-increasing. The present technique will trace the evolution of function F along such sequences of dominated vectors. In order to be able to provide ϵ , ϵ' -accurate answers, one has to identify vectors of constants that cause the value of function F to change by an ϵ' fraction. At the same time, the coordinates of such vectors have to be related by ϵ as required by ϵ , ϵ' -accurate answers.

[0065] Let $[0 \dots D]^n$ be the domain of all possible vectors of constants and consider one of these vectors, $\bar{c} \in [0 \dots D]^n$. Let $S_{\bar{c}}$ be the solution to the query with vector of constants \bar{c} and $F(S_{\bar{c}})$ be the associated profit. The present invention will aim to identify the vector of constants \bar{c}' by manipulating the coordinates of vector \bar{c} by fractions of $1+\epsilon$, such that (a) $\bar{c}' \ll \bar{c}$, (b) $(1+\epsilon') F(S_{\bar{c}'}) > F(S_{\bar{c}})$, where $S_{\bar{c}'}$ the solution to the OPAC query with vector of constants \bar{c}' and (c) vector \bar{c}' is minimal.

[0066] Returning to FIG. 3, method 300 in step 320 obtains a hyper rectangle or partition defined by vectors \bar{c} and \bar{c}' . By definition, any query with vector of constants inside the hyper rectangle has \bar{c}' as an ϵ , ϵ' -answer.

[0067] In step 330, method 300 then inserts the hyper rectangle into a multidimensional data structure, e.g., an R-tree. Method 300 then proceeds to step 340.

[0068] In step 340, method queries whether additional vectors are to be generated. In other words, whether addition partitions are needed to cover the space of all possible queries. If the query is negatively answered, then method 300 ends in step 345. If the query is positively answered, then method 300 returns to step 310, where the process is repeated to generate additional partitions to cover the space of all possible queries.

[0069] FIG. 4 illustrates the preprocessing method 300 in pseudo code format. Given the domain of possible vectors of constants $[0 \dots D]^n$, it starts exploring the space by considering the vector corresponding to the upper right corner of the space. This is a vector \bar{c} that dominates all other vectors and consequently according to the observation made above, corresponds to an OPAC query having the maximum profit. This vector is inserted to a queue and the algorithm iterates while the queue is not empty. For each vector \bar{c} in the queue, the algorithm aims to construct an ϵ , ϵ' -accurate answer for it (and subsequently for each vector dominated by \bar{c}). The algorithm invokes function LocateSolutions with parameter \bar{c} (line (3) in Algorithm GeneratePartitions). This function returns a hyper rectangle r corresponding to a region of space of $[0 \dots D]^n$, a solution S , a vector C and a profit p . The semantics associated with this result is that S is the solution to a query specifying vector of constants C having profit p and forms an ϵ , ϵ' -accurate answer for each vector of constants inside r . The hyper rectangle is inserted in a multidimensional data structure, such as an R-tree, along with the associated profit, vector C and solution S . Along with the leaf index entry for r , the profit p and vector C are stored as well as a pointer, pointing to the solution (set of tuple identifiers) S on disk. Consequently, the index acts as a secondary structure pointing to ϵ , ϵ' -accurate answers on disk. Finally, (in line 6) function "GeneratePartitions" generates a set of n vectors of constants. This set is constructed in a way such that that no vector in Vec dominates another; namely $\nexists \bar{c}', \bar{c}'' \in Vec, \bar{c}' \ll \bar{c}''$. This set of vectors is constructed by calling function "CreateFront", which accepts as parameters the coordinates of the newly formed hyper rectangle r

and the queue \emptyset . Given a hyper-rectangle r with lower left corner $\bar{c}'=(c'_1, \dots, c'_n)$, and upper right corner $\bar{c}=(c_1 \dots c_n)$, CreateFront creates a set of n vectors that together dominate the entire space dominated by \bar{c} , excluding the space spanned by r . This is the set of vectors $\bar{c}_i=(c_1, \dots, c_{i-1}, c'_i, c_{i+1}, \dots, c_n)$, $1 \leq i \leq n$. Each one of these vectors is inserted in the queue Q , unless it is already in the queue.

[0070] Function "LocateSolution" accepts as a parameter a vector \bar{c} and identifies a hyper rectangle $r \subset [0 \dots D]^n$. Let r correspond to a hyper rectangle defined by vectors (\bar{c}', \bar{c}) (where \bar{c}' the lower left and \bar{c} the upper right corners); \bar{c} is the vector provided to LocateSolution at input and \bar{c}' a vector corresponding to a query having a profit no less than an $1+\epsilon'$ fraction of the profit of the query specifying vector of constants \bar{c} . Given a specific \bar{c} , the search for a \bar{c}' with the aforementioned properties is performed in function LocateSolution. Initially (line (1)) function MultiKnapsack (employing any pseudo polynomial algorithm for solving the multi attribute Knapsack problem) is called to determine a solution S (set of tuples) and the profit p to the query with vector of constants \bar{c} . Then, the vector \bar{c}' is formed by decreasing each coordinate of \bar{c} by an $1+\epsilon$ fraction and function MultiKnapsack is called again, this time with vector \bar{c}' , to determine the solution S' with profit p' , to a query with constant vector \bar{c}' . Two cases of may interest arise.

[0071] First, if the profit p' of the query with vector \bar{c}' is larger than an $1+\epsilon'$ fraction of the profit of the query with vector \bar{c} , then algorithm LocateSolutions attempts to minimize vector \bar{c}' by successively reducing its coordinates by a $1+\epsilon$ fraction, updating the solution S' and profit p' attainable (lines (5)-(12)). If it succeeds, the algorithm forms a hyper rectangle using the minimal vector \bar{c}' and \bar{c} , returning it along with vector \bar{c}' , the profit p' and solution S' (line (14)) of \bar{c}' .

[0072] Second, if, on the other hand, the profit of the query with vector \bar{c}' is smaller than an $1+\epsilon'$ fraction of the profit of the query with vector \bar{c} , then algorithm LocateSolutions does not attempt to reduce vector \bar{c}' further; it forms a hyper rectangle consisting of the (\bar{c}', \bar{c}) returning it along with the associated profit p vector \bar{c} and solution S of \bar{c} (line (16)).

[0073] Function MultiKnapsack with parameter \bar{c} is guaranteed to return a non empty solution, if a solution that satisfies the constraints exists. The Function will return a null solution if there is no subset of relation R satisfying the constraints imposed by vector \bar{c} . To formalize this notion the present invention defines the feasible region of relation R below.

[0074] (Feasible Region): Let $R(A_1, \dots, A_n, P)$ be a relation with tuples $(A_1 \dots A_n)$. Assume that the Σ function applied on elements of each attribute A_i and P has range $[0 \dots D]$ as well. The feasible region of R is the set of all vectors c dominated by vector (D, \dots, D) , that dominate at least one tuple of R .

[0075] Algorithm GeneratePartitions progressively reduces the values in each dimension of vectors from the queue and eventually vectors generated by CreateFront will be outside the feasible region. As soon as a vector falls outside the feasible region of R , function MultiKnapsack returns null and progressively the number of elements in the queue decreases. The following example illustrates the operation of the algorithm.

[0076] FIG. 5 illustrates an operation of GeneratePartitions of the present invention. Consider the relation in FIG. 5. There are only two tuples, t_1 and t_2 , and the range of the Σ function on both attributes is $[0,15]$. Assume that $\epsilon = \epsilon' = 0.25$.

[0077] The algorithm starts at (15,15) and finds that the solution at vector (15,15) is $\{t_1, t_2\}$ with Profit **200**. The next vector that is investigated by LocateSolution is $1/1.25(15, 15) = (12,12)$. The best solution at (12, 12) is either $\{t_1\}$, or $\{t_2\}$, both with Profit **100**. Since $100 * 1.25 < 200$, the algorithm does not extend this rectangle further. It adds the rectangle **520** $R_1 = [(15,15), (12,12)]$ in the R-tree, and associates with this rectangle the solution at the top corner, namely $\{t_1, t_2\}$, with Profit **200**.

[0078] CreateFront then adds the following 2 points in the queue: (15,12) and (12,15). At this point, the top of the queue is (15,12). LocateSolution finds that a solution for this point is either $\{t_1\}$ or $\{t_2\}$, both with Profit **100**. The algorithm finds the bottom corner which is $1/1.25(15,12) = (12,10)$ (rounding up the numbers to simplify the example) which still has a solution with profit **100**, so it extends this to $1/1.25(12,10) = (10,8)$ which still has the same solution ($\{t_1\}$), and then to (8, 7), which does not have any feasible solution. So at this point it backs up, and adds the rectangle **530** $R_2 = [(15,12), (10,8)]$ with solution $\{t_1\}$ and Profit **100** to the R-tree. CreateFront then adds the points (15,8) and (10,12) to the queue. Similarly, considering vector (12,15) leads to the insertion of rectangle **540** R_3 in the R-tree, and vector (15, 8) leads to the insertion of rectangle **550** R_4 . The next vector at the top of the queue is (10,12), which is contained in rectangle R_3 . So no new rectangle is created, but the vectors (10,10) and (8,12) are inserted in the queue by CreateFront. The operation of the algorithm for this example is illustrated in FIG. 5, where FIG. 6 presents an example of the partitioning generated by algorithm GeneratePartitions.

[0079] The outcome of algorithm GeneratePartitions, is a multi-dimensional index (e.g., an R-tree), providing access to a collection of hyper rectangles. For a query Q specifying a vector of constants \bar{c} , the present invention obtains an ϵ, ϵ' -accurate answer.

[0080] Specifically, the present invention use \bar{c} and search the multidimensional index seeking hyper rectangles containing \bar{c} and let Ans be the set of hyper rectangles identified. Each of these hyper rectangles obtained from the leaves of the index, has a profit and a vector associated with it and points to a set of tuple identifiers on disk. Any of these solutions is an ϵ, ϵ' answer and returning the vector, profit, and set of tuple identifiers associated with any hyper rectangle in Ans, suffices. In the case Ans is empty, then there is no feasible answer to Q in relation R and thus there is no possible ϵ, ϵ' -accurate answer.

[0081] Consider again the example of FIG. 5. Assume one is given query (14,14). This query vector falls in rectangle $R_1[(15,15), (12,12)]$. This query can be answered with the solution $\{t_1, t_2\}$, which is associated with R_1 , and which has Profit **200**. This is an ϵ, ϵ' -accurate answer. From the definition of the ϵ, ϵ' -accurate answer; the present invention has to return a solution which satisfies constraints $c'_i < (1+\epsilon)c_i$ and profit $P'(1+\epsilon') > P$. For the constraint vector (14,14), the best solution is either $\{t_2\}$ or $\{t_1\}$, with profit **100**. Since $15 \leq (1+\epsilon)14 = 1.25 * 14$, the solution $\{t_1, t_2\}$ satisfies the

relaxed constraints, and, since $200 * 1.25 > 100$, the profit constraint is satisfied as well. Intuitively what happens is that, the present invention does not give an exact answer; but the present invention provides an answer that is at least as good (and may in fact be much better) if the user is willing to relax the constraints by a factor of ϵ .

[0082] Let's assume a query (14, 9). This vector falls in rectangle R_2 . The solution that is returned is the one associated with rectangle $r_2: \{t_1\}$, with Profit **100**. In this case, $c'_i < 1.25 c_i$, and $P' = P$, so this is a ϵ, ϵ' -accurate answer as well.

[0083] It should be noted that the present method guarantees that every feasible member of the space of all possible constant vectors will be contained in at least one hyper rectangle. The algorithm will cover the space of candidate query vectors using hyper rectangles.

[0084] In particular: Let \bar{c} be the constant vector associated with an OPAC query Q . Assume there exists a subset of tuples that satisfy the constraints \bar{c} . Then algorithm GeneratePartitions, creates at least one hyper rectangle containing vector \bar{c} .

[0085] The fact that the GeneratePartitions method covers the entire feasible space can be demonstrated. The first vector dominates the entire space. Each iteration of the algorithm takes a vector from the queue, uses this vector to form the upper right corner of a new hyper rectangle, and adds a new set of vectors in the queue that together dominate the space that the original vector dominated with the exception of the space of the hyper-rectangle. Since the algorithm terminates when no vectors are in the queue, it follows that the entire feasible space is covered by hyper rectangles.

[0086] The following result demonstrates that the answer to any OPAC query, Q , specifying a vector of constants \bar{c} , obtained from the index, is an ϵ, ϵ' -accurate answer.

[0087] Theorem: Specifically, Let Q be a query specifying vector \bar{c} as a constant vector. Let r be a hyper rectangle containing \bar{c} , generated by algorithm GeneratePartitions. The answer to Q returned from the index, consisting of a vector; a set of tuples, and a profit is an ϵ, ϵ' -accurate answer.

[0088] The above theorem can be demonstrated. Let p_2 be the profit of the lower left corner vector \bar{c}_2 , and p_1 be the profit of the upper right corner vector \bar{c}_1 of multidimensional rectangle r . Assume a vector \bar{c} located inside r .

[0089] If $p_2(1+\epsilon') \geq p_1$, then the lower left corner is an ϵ, ϵ' -accurate answer: since \bar{c}_2 is dominated by \bar{c}_1 , the optimal profit for \bar{c} is at most p_1 , and therefore at most an $1+\epsilon'$ fraction higher than p_2 . In this case algorithm GeneratePartitions stores in the index entry, along with r , a vector with coordinates equal to \bar{c}_2 and a profit p_2 and thus the answer returned is an ϵ, ϵ' -accurate answer.

[0090] If on the other hand $p_2(1+\epsilon') < p_1$, then the profit of the lower left corner may be more than a fraction of $1+\epsilon'$ smaller than the profit of the query. However, by the construction of the algorithm GeneratePartitions this can only happen if all the values of the vector \bar{c} are within an $1+\epsilon$ fraction of the values of the upper right vector \bar{c}_1 . In this case, \bar{c}_1 provides an ϵ, ϵ' -accurate answer since it gives a much better profit with just an ϵ relaxation of the constraints. Algorithm GeneratePartitions will associate the vector \bar{c}_1 and its profit with the index entry, along with r .

[0091] Restricting the present invention to monotone classes of aggregation functions, one can improve the computational aspects related to the construction of an ϵ -Pareto set. The following theorem shows that the total number of hyper-rectangles represented in the index is polynomial.

[0092] Theorem: The number of vectors that create new multidimensional rectangles at any step of the execution of algorithm GeneratePartitions, is polynomial to $1/\epsilon$.

[0093] The above theorem can be demonstrated. Assume that the range of the Σ function applied on the A_i attribute values of a non-empty subset of the tuples is $[1 \dots D]$ for each i . Here it is assumed that the attributes have non zero values; to deal with zero values that have to be added to the interval from zero to the smallest non zero value as one additional interval in the partition.

[0094] Then the present invention can partition the range in

$$o\left(\frac{\log D}{\log(1+\epsilon)}\right) = o\left(\frac{\log D}{\epsilon}\right)$$

intervals (for $0 < \epsilon < 1$), geometrically increasing with step $1+\epsilon$. Taking the Cartesian product of the n attributes, the present invention creates

$$o\left(\left(\frac{\log D}{\epsilon}\right)^n\right)$$

n dimensional points. Note that, by the construction of the algorithm, every vector inserted in the queue corresponds to one of these points. It follows that the number of hyper rectangles in the index is polynomial to $1/\epsilon$ and to the size of the range of the Σ function applied on the attribute values, and is exponential to the number of the attributes.

[0095] FIG. 7 illustrates a block diagram of the present query optimization method being implemented on a general purpose computer 700 or any other hardware equivalents. Specifically, the system can be employed to efficiently provide results in response to a query having aggregation constraints. In one embodiment, the system deploys a query optimization module 740 that executes various methods or algorithms as disclosed above relating to query optimization.

[0096] Thus, system 700 comprises a processor (CPU) 710, a memory 720, e.g., random access memory (RAM) and/or read only memory (ROM), the query optimization module 740, and various input/output devices 730, (e.g., storage devices, including but not limited to, a tape drive, a floppy drive, a hard disk drive or a compact disk drive, a receiver, a transmitter, a speaker, a display, a clock, an output port, a user input device (such as a keyboard, a keypad, a mouse, and the like, or a microphone for capturing speech commands).

[0097] It should be understood that the query optimization module 740 can be implemented as one or more physical devices that are coupled to the CPU 710 through a communication channel. Alternatively, the query optimization module 740 can be represented by one or more software applications (or even a combination of software and hardware, e.g., using application specific integrated circuits (ASIC)), where the software is loaded from a storage medium, (e.g., a magnetic or optical drive or diskette) and operated by the CPU in the memory 720 of the computer. As such, the query optimization module 740 (including associated data structures) of the present invention can be stored on a computer readable medium, e.g., RAM memory, magnetic or optical drive or diskette and the like.

[0098] While the foregoing is directed to embodiments of the present invention, other and further embodiments of the invention may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

1. A method for generating an index for use with query having aggregation constraints, comprising:

identifying a dominating vector of constants, \bar{c}' for a given n -dimensional vector of constants \bar{c} ;

obtaining a first partition defined by said vector \bar{c} and said vector \bar{c}' ; and

inserting said first partition into a multidimensional data structure.

2. The method of claim 1, further comprising

excluding said first partition; and

wherein said steps are iteratively repeated until an entire feasible space dominated by said vector of constants \bar{c} has been partitioned.

3. An apparatus for generating an index for use with query having aggregation constraints, comprising:

means for identifying a dominating vector of constants, \bar{c}' for a given n -dimensional vector of constants \bar{c} ;

means for obtaining a first partition defined by said vector \bar{c} and said vector \bar{c}' ; and

means for inserting said first partition into a multidimensional data structure.

4. A computer-readable medium having stored thereon a plurality of instructions, the plurality of instructions including instructions which, when executed by a processor, cause the processor to perform the steps comprising of:

identifying a dominating vector of constants, \bar{c}' for a given n -dimensional vector of constants \bar{c} ;

obtaining a first partition defined by said vector \bar{c} and said vector \bar{c}' ; and

inserting said first partition into a multidimensional data structure.

* * * * *