

Partial Materialized Views

Gang Luo

IBM T.J. Watson Research Center

luog@us.ibm.com

Abstract

Early access to partial query results is highly desirable during exploration of massive data sets. However, it is challenging to provide transactionally consistent, immediate partial results without significantly increasing queries' execution time. To address this problem, this paper proposes a partial materialized view (PMV) method to cache some of the most frequently accessed results rather than all the possible results. Compared to traditional materialized views, the proposed PMVs do not require maintenance during insertion into base relations, and have much smaller storage and maintenance overhead. Upon the arrival of a query, the RDBMS first searches the PMV and returns to the user the cached partial results. Since a large portion of the PMV is cached in memory, this usually finishes within a millisecond. Then the RDBMS continues to execute the query to find the remaining results. The efficiency of our PMV method is evaluated through a simulation study, a theoretical analysis, and an initial implementation in PostgreSQL.

1. Introduction

Large data sets are common in practice, and the sizes of these data sets are becoming larger and larger. As a result, the capability of efficiently exploring massive data sets is urgently needed [19]. It has been widely recognized that early access to partial query results can provide the following benefits and greatly facilitate the exploration of massive data sets:

Benefit 1: The RDBMS becomes more user-friendly.

Benefit 2: Early termination of those queries with unsatisfactory partial results (e.g., if users would like to refine them) can greatly reduce the load on the RDBMS and significantly speed up the exploration process.

In practice, it is important to provide transactionally consistent, immediate partial results without significantly increasing queries' execution time, while statistical guarantee for the partial results is often not necessary.

For example, consider a retailer's customer service call center. When a customer calls in, the call center operator can offer him on-sale items that are of his interest. The operator first obtains all the items I_p that the customer recently purchased and then performs a query Q on two relations. From the first relation, Q retrieves all the items I_r that are related to at least one of the items in I_p . From the second relation R_{sale} , Q finds all the items in I_r that are currently on sale with a discount of at least $p\%$, where p is determined based on the loyalty of the customer. The

operator only needs to see partial results of Q in order to start making offers to the customer and no statistical guarantee is needed for these partial results. Nevertheless, these partial results have to be obtained quickly (before the customer hangs up). In commercial databases, a common practice is to use a separate R_{sale} for each store or each department. Consequently, many query templates are needed to support this application.

Database researchers have spent much effort on investigating techniques for providing partial query results. However, none of the existing techniques is completely satisfactory. These techniques fall into three categories:

- (1) Use non-blocking query processing to generate output tuples continuously [7, 20, 19].
- (2) Use special optimization techniques to find the first or top- k output tuples quickly [5, 8, 12, 22].
- (3) Use asynchronously updated replicas to provide output tuples quickly [3, 16].


Non-blocking query processing often increases a query's execution time significantly. During exploration of massive data sets, knowing beforehand whether or not the user wants to see all the query results is often infeasible. Hence, it is difficult to decide in advance whether traditional (blocking) query processing should be used to optimize the query's execution time, or non-blocking query processing should be used to generate output tuples continuously.

The optimization techniques for quickly finding the first or top- k output tuples are based on traditional (blocking) query processing and often require expensive I/Os. Hence, it can take much time (e.g., a few minutes) to generate the first or top- k output tuples. Moreover, in order to use these optimization techniques, the user needs to specify k , which can be difficult to know beforehand. If the user is not satisfied with the first or top- k output tuples and would like to see all the results, he has to re-execute the query. This re-execution wastes system resources and is slower than requiring all the results at the first time.

In general, there is a delay before the data in the master copy is transferred to an asynchronously updated replica. Thus, the query results provided by the replica can be transactionally inconsistent with the data in the master copy (e.g., a tuple is deleted from the master copy but still exists in the replica). This is unacceptable to many applications.

In this paper, we propose a *partial materialized view* (PMV) method that can provide immediate partial results without increasing queries' execution time much. These partial results are transactionally consistent and suitable for those applications that do not require statistical guarantee. Our idea is to reuse previous "hot" results. More

specifically, from previous queries' execution, some of the most frequently accessed results are remembered in the so-called PMVs. When a new query Q comes, the corresponding PMV is first searched and the found partial results are returned to the user. This often finishes within a millisecond, because a large portion of PMV is cached in memory. Then Q is executed to find the remaining results.



Compared to traditional materialized views (MVs) that store all possible results, our PMVs only store some of the most frequently accessed results and have smaller sizes. This saves most of the storage and maintenance overhead of traditional MVs while many queries can still have early access to partial results. Since PMVs are not used to provide all the query results, no maintenance of PMV is needed during insertion into base relations. To ensure that a large portion of PMVs is cached in memory and thus the return of partial results is quick, the size of each PMV has an upper bound. To increase our chance of using a PMV to provide partial results with only a limited storage, we continuously update the content in the PMV to adapt to the current query pattern, and restrict the maximum number of tuples that can be stored in the PMV for any single, so-called basic condition part. Whenever possible, both the maintenance and the update of PMVs are coupled with query execution for free. We investigate the performance of the PMV method with a simulation study, a theoretical analysis, and an initial implementation in PostgreSQL. Our results show that PMVs have minor overhead and can often provide partial results almost instantly. Also, the RDBMS can afford storing many PMVs.

The rest of the paper is organized as follows. Section 2 discusses the limitations of traditional MV method. Section 3 presents the details of the proposed PMV method. Section 4 investigates the performance of the PMV method. Finally, we discuss related work in Section 5 and conclude in Section 6.

2. Limitations of Materialized Views

A traditional method of speeding up query execution is to use MVs [17]. In this section, we first describe the queries that will be considered by the PMV method and then discuss the limitations of traditional MV method.

2.1 Query Specification

We consider one type of queries that are frequently encountered in practice (e.g., in form-based applications) – queries coming from templates of the following form:

q_i : `select L_s from R_1, R_2, \dots, R_n where C_{join} and C_{select}`

Here, L_s is the select list. C_{join} includes both the join condition among the $n \geq 1$ relations R_1, R_2, \dots , and R_n , and the selection conditions on a single relation that have no

parameters (e.g., $R_1.b=100$). $C_{select} = \wedge_{i=1}^m C_i$, where m is a number. Each C_i ($1 \leq i \leq m$) is a selection condition on a single relation R_{h_i} ($1 \leq h_i \leq n$). C_i takes one of the following two disjunctive forms, which accept one or more parameters:

Equality form: $\vee_{r=1}^{u_i} (R_{h_i}.a_{k_i} = v_{i,r})$.

Interval form: $\vee_{r=1}^{u_i} (v_{i,r} < R_{h_i}.a_{k_i} < w_{i,r})$. The intervals $(v_{i,1}, w_{i,1}), (v_{i,2}, w_{i,2}), \dots$, and (v_{i,u_i}, w_{i,u_i}) are disjoint from each other.

Different queries from the same template can have different u_i 's ($1 \leq i \leq m$). In Section 3.6, we will show how to extend our techniques to handle other forms of queries (e.g., aggregate queries, nested queries).

In the interval form case, $R_{h_i}.a_{k_i}$ can be a non-numerical (e.g., string) attribute. Also, $v_{i,r}$ ($1 \leq r \leq u_i$) can be $-\infty$ while $w_{i,r}$ can be $+\infty$, and “ $<$ ” can be replaced by “ \leq ”. In other words, the intervals can be either bounded or unbounded, open or closed. For ease of presentation, in the remainder of this paper, we always write an interval as an open bounded one, with the understanding that it can be closed and/or unbounded if necessary.

```
select R.a, S.e from R, S
where R.c=S.d and (R.f=f1 or R.f=f2 or ... or R.f=fh)
and (S.g=g1 or S.g=g2 or ... or S.g=gk);
```

Figure 1. An example query template E_{qt} .

For a large subset of queries of the q_t form, traditional query processing cannot produce output tuples quickly and continuously. These queries are our focus. Their query plans can be either fully pipelined or not fully pipelined. To illustrate the former case, let us consider the template E_{qt} in Figure 1. Suppose that an index exists on each selection/join attribute. The query plan fetches tuples from R using the index on $R.f$. For each retrieved tuple t_R , the index on $S.d$ is used to search S for matching tuples. If the selectivity of $S.g$ is low, the index on $S.d$ needs to be searched many times before the first query result tuple is obtained. This can take a few seconds in a lightly loaded RDBMS, and a few minutes in a heavily loaded RDBMS.

2.2 Limitations of Large Materialized Views

```
create materialized view VM as
select R.a, S.e, R.f, S.g from R, S where R.c=S.d;
```

Figure 2. An example large materialized view.

Existing techniques for automatically selecting MVs from query traces are based on “merging,” where the definition of each suggested MV is based on the common part of some of the queries [2, 33]. For example, for the template E_{qt} in Figure 1, existing automatic MV selection tools may suggest a materialized view V_M as shown in Figure 2. (The search procedure in V_M needs attributes $R.f$ and $S.g$.) As V_M needs to keep all the possible results for

queries from E_{qt} , V_M is fairly large. In general, due to the extreme storage and maintenance overhead of MVs [17], the RDBMS cannot keep a MV for each frequently used query template.

2.3 Limitations of Small Materialized Views

For the template E_{qt} in Figure 1, instead of using the big V_M in Figure 2 for all possible (f_i, g_j) pairs, one might wonder whether we could create multiple small MVs, one for each “hot” (f_i, g_j) pair, and use them to speed up query processing. These small MVs have the following advantages. First, a hot (f_i, g_j) pair appears frequently in queries from E_{qt} . Therefore, the RDBMS can use a small MV that is built for a hot (f_i, g_j) pair to partially answer a lot of queries from E_{qt} . Second, the combined size of these small MVs is a small percentage of that of the big V_M . Thus, the combined storage and maintenance overhead of these small MVs is smaller than that of V_M . Also, compared to V_M , these small MVs can be accessed more quickly, as they are more likely to be cached in memory.


However, existing techniques for answering queries using MVs [9, 14, 18] focus on shortening queries’ execution time. For many queries from the template E_{qt} , these small MVs cannot help much for that purpose. This is because typically, a query from E_{qt} contains both several hot (f_i, g_j) pairs and several cold (f_i, g_j) pairs. During the process of obtaining the results corresponding to the cold (f_i, g_j) pairs, the results corresponding to the hot (f_i, g_j) pairs can be computed inexpensively without using these small MVs.

For example, suppose that $(R.f=1, S.g=2)$ is the only hot (f_i, g_j) pair. We create a small materialized view V_{SM} from $(R.f=1, S.g=2)$ as follows:

create materialized view V_{SM} as select $R.a, S.e$
from R, S where $R.c=S.d$ and $R.f=1$ and $S.g=2$;

Consider the following query that comes from the template E_{qt} in Figure 1:

select $R.a, S.e$ from R, S
where $R.c=S.d$ and $(R.f=1$ or $R.f=3)$
and $(S.g=2$ or $S.g=4)$;



In order to obtain the results corresponding to the cold pair $(R.f=1, S.g=4)$, tuple(s) t_R of R where $R.f=1$ are fetched. Similarly, to obtain the results corresponding to $(R.f=3, S.g=2)$, tuple(s) t_S of S where $S.g=2$ are retrieved. Then it is not expensive to compute the (possibly in-memory) join between t_R and t_S and there is no need to use V_{SM} .

3. The Partial Materialized View Method

In this section, we present our PMV method for providing partial query results, which can overcome the limitations of traditional MV method. We first describe the main ideas. Then we go into the details of the method.

All discussions in Section 2.3 about small MVs are from the viewpoint of minimizing queries’ execution time. The main goal of our PMV method is to minimize the time of generating partial results. In this case, these small MVs for the hot (f_i, g_j) pairs become useful, as they can quickly

provide partial results to a large number of queries from the template E_{qt} .

For example, consider a query Q from the template E_{qt} in Figure 1. Q contains both several hot (f_i, g_j) pairs and several cold (f_i, g_j) pairs. The RDBMS answers Q in the following way:

Step 1: These small MVs are used to quickly obtain the partial results corresponding to the hot (f_i, g_j) pairs. These partial results are returned to the user and recorded in a temporary in-memory data structure D_S .

Step 2: Q is executed to obtain all the results. For each result tuple t , we check whether $t \in D_S$. If so (i.e., the user has already obtained t at Step 1), t is removed from D_S and not returned to the user. Otherwise if $t \notin D_S$, the RDBMS knows that t corresponds to some cold (f_i, g_j) pair and returns t to the user. In this way, each result tuple is returned to the user exactly once. (Query results can contain duplicate tuples. In the case that $t \in D_S$, if t is not removed from D_S and later another tuple $t'=t$ comes, the user can miss some result tuples.)

The above method will slightly increase query Q ’s execution time, as neither Step 1 nor the checking at Step 2 is needed in traditional query processing. However, this extra overhead is minor compared to the two benefits (user-friendliness, load reduction) of providing partial results that are mentioned in the introduction. Thus, for those applications of exploring massive data sets, it is worth to make this tradeoff. For the purpose of easy management, all the small MVs are combined into a single so-called PMV. This becomes our PMV method. More details of our method are described in the following subsections.

3.1 Definitions

We first introduce some definitions.

Partial materialized view. Consider a MV definition V_M . V_M may or may not exist in the RDBMS. Any subset V_{PM} of V_M is a *partial materialized view*. V_M is the *containing materialized view* of V_{PM} . The base relations of V_M are also called the base relations of V_{PM} . (Both MVs and PMVs are treated as multi-sets and thus can contain duplicate tuples.)

For the materialized view V_M in Figure 2, we show an example partial materialized view V_{PM} in Figure 3.

relation R	relation S	materialized view V_M	partial materialized view V_{PM}																																																
<table border="1"><tr><td>a</td><td>c</td><td>f</td></tr><tr><td>1</td><td>4</td><td>1</td></tr><tr><td>1</td><td>5</td><td>1</td></tr><tr><td>7</td><td>6</td><td>3</td></tr></table>	a	c	f	1	4	1	1	5	1	7	6	3	<table border="1"><tr><td>d</td><td>e</td><td>g</td></tr><tr><td>4</td><td>2</td><td>7</td></tr><tr><td>5</td><td>2</td><td>7</td></tr><tr><td>6</td><td>8</td><td>9</td></tr></table>	d	e	g	4	2	7	5	2	7	6	8	9	<table border="1"><tr><td>a</td><td>e</td><td>f</td><td>g</td></tr><tr><td>1</td><td>2</td><td>1</td><td>7</td></tr><tr><td>1</td><td>2</td><td>1</td><td>7</td></tr><tr><td>7</td><td>8</td><td>3</td><td>9</td></tr></table>	a	e	f	g	1	2	1	7	1	2	1	7	7	8	3	9	<table border="1"><tr><td>a</td><td>e</td><td>f</td><td>g</td></tr><tr><td>1</td><td>2</td><td>1</td><td>7</td></tr></table>	a	e	f	g	1	2	1	7
a	c	f																																																	
1	4	1																																																	
1	5	1																																																	
7	6	3																																																	
d	e	g																																																	
4	2	7																																																	
5	2	7																																																	
6	8	9																																																	
a	e	f	g																																																
1	2	1	7																																																
1	2	1	7																																																
7	8	3	9																																																
a	e	f	g																																																
1	2	1	7																																																

Figure 3. An example partial materialized view.

Condition part. Consider the query template q_t in Section 2.1. A *condition part* is an m -tuple (d_1, d_2, \dots, d_m) , where for each i ($1 \leq i \leq m$):

- (1) If the selection condition C_i is of equality form, d_i is of the form $R_{h_i}.a_{k_i} = b_i$.

(2) If C_i is of interval form, d_i is of the form $b_i < R_{h_i}.a_{k_i} < c_i$.

A query result tuple t belongs to a condition part (d_1, d_2, \dots, d_m) if t satisfies all conditions d_i ($1 \leq i \leq m$). A condition part (d_1, d_2, \dots, d_m) is contained in another condition part $(d'_1, d'_2, \dots, d'_m)$ if whenever conditions d_i ($1 \leq i \leq m$) are true, conditions d'_i ($1 \leq i \leq m$) are also true.

For each selection condition C_i ($1 \leq i \leq m$) that is of interval form, let E_i denote the entire range of all possible intervals in C_i (e.g., $E_i = (-\infty, +\infty)$). We assume that the RDBMS knows multiple “dividing” values that can divide E_i into multiple non-overlapping “basic” intervals and these basic intervals fully cover E_i . Each basic interval is assigned a different id. The purpose of this division is discretization so that the problem becomes more tractable. The criterion for choosing dividing values is that the resulting basic intervals can be used to differentiate hot results from cold results.

In many form-based applications, for each selection condition C_i ($1 \leq i \leq m$) that is of interval form, the user is provided with both a list of from values and a list of to values. Each (from value, to value) pair chosen by the user forms an interval $(v_{i,r}, w_{i,r})$, where $1 \leq r \leq u_i$. In this case, these from values and to values can serve as dividing values. In other cases, we assume that either the person (e.g., DBA) who defines the PMV for the query template will specify the dividing values, or the continuous feature discretization technique [11] in machine learning can automatically learn dividing values from query traces.

Basic condition part. A condition part (d_1, d_2, \dots, d_m) is a *basic condition part*, if for each selection condition C_i ($1 \leq i \leq m$) that is of interval form, d_i is of the form $b_i < R_{h_i}.a_{k_i} < c_i$, where (b_i, c_i) is a basic interval.

A basic condition part (d_1, d_2, \dots, d_m) is stored in the following way:

- (1) If d_i is of the form $R_{h_i}.a_{k_i} = b_i$, value b_i is stored.
- (2) If d_i is of the form $b_i < R_{h_i}.a_{k_i} < c_i$, where (b_i, c_i) is a basic interval, the id of (b_i, c_i) is stored.

3.2 Organization of Partial Materialized Views

Consider a frequently used query template q_t (see Section 2.1). Suppose that the RDBMS cannot afford to keep a materialized view $V_M = (\text{select } L_s' \text{ from } R_1, R_2, \dots, R_n \text{ where } C_{\text{join}})$. Here, L_s' is the expanded select list that includes all the attributes in both C_{select} and the original select list L_s . (The search procedure in V_M needs the attributes in C_{select} .)

We build a partial materialized view V_{PM} for q_t :
 create partial materialized view V_{PM} as subset of
 select L_s' from R_1, R_2, \dots, R_n

where C_{join} with selection condition template C_{select} .

V_M is the containing MV of V_{PM} . All the tuples in V_{PM} satisfy the condition C_{join} .

The person who defines V_{PM} specifies an upper bound U_B for the size of V_{PM} . This U_B is used to constrain the storage and maintenance overhead of V_{PM} , and ensure that a significant portion of V_{PM} is cached in memory so that V_{PM} can be accessed quickly. Initially, V_{PM} is empty. Our goal is to use V_{PM} to provide immediate partial results to as many queries from the template q_t as possible.

In the template q_t , the original select list L_s is replaced with the expanded select list L_s' . This is to let all the attributes in C_{select} appear in query result tuples. As will be shown later, some result tuples are stored in V_{PM} . The attributes in C_{select} are needed to find partial results in V_{PM} . When the RDBMS obtains a query result tuple, it only returns the attributes in L_s to the user. Hence, the user still receives the same answer, as if V_{PM} did not exist and L_s in q_t had not been replaced by L_s' .

Each tuple of V_{PM} is composed of two parts: the “conceptual” basic condition part $bcp = (d_1, d_2, \dots, d_m)$, and attributes ats . ats is a query result tuple that includes all the attributes in the expanded select list L_s' , and belongs to bcp . bcp is “conceptual” in the sense that it is not actually stored in the tuple. Whenever needed, bcp is recovered from ats . We build an index I on bcp . If $m > 1$, I is a multi-attribute index. For example, for the template E_{q_t} in Figure 1, Figure 4 shows the corresponding PMV.

index I on bcp	$bcp=(f, g)$	a	e	f	g

Figure 4. Data structure of a partial materialized view V_{PM} .

Our goal is to use V_{PM} to provide immediate partial results to as many queries as possible. Hence, it is preferable to have a large number of basic condition parts stored in V_{PM} . In general, many query result tuples can belong to a single basic condition part, and it is not desirable to flood V_{PM} with all these tuples. Therefore, the person who defines V_{PM} specifies a constant F . For a basic condition part bcp , the RDBMS stores at most F result tuples (rather than all the possible result tuples) that belong to bcp in V_{PM} . This is different from the case of traditional MVs, where a materialized view V_M stores all the result tuples that satisfy the definition of V_M . Given the storage limit U_B of V_{PM} , for a query Q , this F makes a tradeoff between (a) the probability that V_{PM} can provide some partial results to Q , and (b) in the case that V_{PM} contains some partial results of Q , the number of partial result tuples that V_{PM} can provide to Q .

Let L denote the number of basic condition parts in V_{PM} . A_t denotes the average size of the tuples in V_{PM} . We have $U_B \leq L \times F \times A_t$. If $L=10K$, $F=2$, and $A_t=50B$, then the size of V_{PM} is no more than 1MB and thus the memory can hold many PMVs. As will be shown in Section 4.1, $L=10K$ can lead to a hit probability of 95%.

The design principles of our algorithm are as follows. The storage budget U_B is limited. Hence, V_{PM} should store

hot basic condition parts. (A hot basic condition part appears in a large number of queries.) This is to maximize the chance that V_{PM} can provide partial results to a query.

The query pattern can change from time to time. That is, the basic condition parts that are hot can keep changing. We want to automatically keep track of this change and update V_{PM} accordingly. Hence, all the basic condition parts in V_{PM} are managed by the CLOCK algorithm [29]: when V_{PM} is full, the RDBMS replaces the basic condition parts in V_{PM} that are no longer hot with the currently hot basic condition parts.

V_{PM} is initially empty. Before V_{PM} becomes full, content is filled into V_{PM} . When V_{PM} becomes full, the content in V_{PM} is updated as query pattern changes. Both the fill in process and the update process of V_{PM} should be as efficient as possible. Therefore, in the case that there is no change to the base relations of V_{PM} , the RDBMS only fills content into V_{PM} (if V_{PM} is not full) or changes the content of V_{PM} (if V_{PM} is full) for free when it obtains result tuples from query execution. There is no separate process for examining the base relations of V_{PM} .

Similarly, in the case that the base relations of V_{PM} get changed, the maintenance of V_{PM} should be as efficient as possible. Hence, whenever possible, the RDBMS couples the maintenance of V_{PM} with the execution of subsequent queries for free. Lastly, the use of V_{PM} needs to have minor influence on queries' execution time.

3.3 Handling Queries

When a query Q comes, the RDBMS performs the following operations:

Operation O_1 : The C_{select} of Q is broken into one or more non-overlapping condition parts. Each condition part is either a basic condition part itself or contained in a basic condition part.

Operation O_2 : For each generated condition part, the RDBMS checks whether there is a corresponding entry in V_{PM} . If so, the related tuples in V_{PM} are returned to the user as partial results. In this way, the RDBMS finds all the result tuples of Q that are in V_{PM} .

Operation O_3 : Q is executed to obtain all the result tuples. For those tuples that the user does not receive in Operation O_2 , the RDBMS returns them to the user now. Also, the content in V_{PM} is updated to reflect the observed change in the hot basic condition parts.

Operation O_1 : $C_{select} \Rightarrow$ Condition Parts

$C_{select} = \wedge_{i=1}^m C_i$. For each i ($1 \leq i \leq m$), there are two possible cases:

(1) C_i is of equality form $\bigvee_{r=1}^{u_i} (R_{h_i}.a_{k_i} = v_{i,r})$. Let set

$$S_i = \{R_{h_i}.a_{k_i} = v_{i,r} \mid 1 \leq r \leq u_i\}.$$

(2) C_i is of interval form $\bigvee_{r=1}^{u_i} (v_{i,r} < R_{h_i}.a_{k_i} < w_{i,r})$. For each r ($1 \leq r \leq u_i$), the RDBMS finds all the basic

intervals $J_{i,r}$ that overlap with the interval $(v_{i,r}, w_{i,r})$.

Let set $S_i = \bigcup_{r=1}^{u_i} \{R_{h_i}.a_{k_i} \in (v_{i,r}, w_{i,r}) \cap I_b \mid I_b \in J_{i,r}\}$.

C_{select} is broken into a number ($h \geq 1$) of "non-overlapping" condition parts $\prod_{i=1}^m S_i = \{cp_j \mid 1 \leq j \leq h\}$. For each condition

part cp_j ($1 \leq j \leq h$), there are two possible cases:

(1) cp_j is a basic condition part bcp_j itself.

(2) cp_j is contained in a basic condition part bcp_j .

In either case, bcp_j is called the *containing basic condition part* of cp_j .

Suppose that in the template E_{qt} in Figure 1, the selection condition on $S.g$ is of interval form rather than of equality form. Figure 5 shows an example of breaking the C_{select} of a query from E_{qt} into condition parts. The outer rectangle represents the entire query space, which is partitioned into non-overlapping basic condition parts as shown by the dashed lines. The gray rectangle represents the query. The C_{select} of this query is broken into nine condition parts. Each condition part is represented by the intersection of the gray rectangle and a dashed rectangle that is filled with either upward or downward diagonals.




Figure 5. An example of breaking the C_{select} of a query from E_{qt} into condition parts.

Operation O_2 : Returning Partial Results

A temporary in-memory data structure D_S is kept. For each condition part cp_j ($1 \leq j \leq h$) generated in Operation O_1 , a counter c_j is kept for its containing basic condition part bcp_j . Initially, D_S is empty and $c_j=0$ ($1 \leq j \leq h$). For each cp_j ($1 \leq j \leq h$), the index I on bcp_j is used to check whether cp_j 's containing basic condition part bcp_j exists in V_{PM} . There are two possible cases:

(1) bcp_j exists in V_{PM} . c_j is set to be the number of tuples in V_{PM} that belong to bcp_j . For each tuple t in V_{PM} that belongs to bcp_j , the RDBMS checks whether t belongs to cp_j . This is equivalent to checking whether t satisfies the C_{select} of query Q . If cp_j is a basic condition part itself, t must belong to cp_j . In contrast, if cp_j is contained in a basic condition part, t may or may not belong to cp_j . All the tuples in V_{PM} satisfy the condition C_{join} . Hence, if t satisfies C_{select} , t is returned to the user as a partial result, and recorded in D_S .

(2) bcp_j does not exist in V_{PM} . Nothing is done in this case.

Operation O_3 : Returning Remaining Result Tuples and Updating Partial Materialized View

Query Q is executed to obtain all the result tuples. For each such result tuple t , the data structure D_S is checked to see whether the user has already obtained t in Operation

O_2 . If $t \in D_S$, t is removed from D_S . If $t \notin D_S$, the RDBMS performs the following operations:

- (1) Return t to the user.
- (2) Find the containing basic condition part bcp_j ($1 \leq j \leq h$) that t belongs to. For each basic condition part bcp , at most F query result tuples that belong to bcp can be stored in V_{PM} . If the counter $c_j < F$, t is added into V_{PM} and c_j is incremented by 1. This can require purging some basic condition part (and the associated query result tuples) from V_{PM} if V_{PM} has already been full. This case of $c_j < F$ is possible, e.g., as V_{PM} is not maintained immediately during insertion into the base relations of V_{PM} (see Section 3.4). In the case that $c_j = 0$, a new basic condition part bcp_j is added into V_{PM} .

After all the result tuples have been processed, the data structure D_S must be empty. D_S is freed.

3.4 Maintaining Partial Materialized Views

When the base relations of V_{PM} get changed, V_{PM} is maintained in a different way from traditional MVs. This is because V_{PM} is only a subset of its containing materialized view V_M . V_{PM} is not used to provide all the query results. As long as V_{PM} does not provide incorrect partial results, there is no need to change V_{PM} immediately. Rather, the maintenance of V_{PM} is deferred to when the RDBMS obtains result tuples from the execution of future queries for free. This minimizes the influence of V_{PM} on transactions that change the base relations of V_{PM} .

Upon a change ΔR_i to a base relation R_i ($1 \leq i \leq n$) of V_{PM} , there are three possible cases:

- (1) The change is an insert. This insert may generate new query result tuples. However, existing tuples in V_{PM} are not affected by this insert. Hence, V_{PM} is not maintained immediately.
- (2) The change is a delete. The join between ΔR_i and the other base relations R_j ($1 \leq j \leq n, j \neq i$) of V_{PM} is computed. For each join result tuple t , the index I on bcp is used to check whether $t \in V_{PM}$. (t must exist in V_{PM} 's containing MV V_M . However, since $V_{PM} \subseteq V_M$, t may or may not exist in V_{PM} .) If $t \in V_{PM}$, t is removed from V_{PM} .
- (3) The change is an update. Recall that all the attributes in C_{select} appear in the expanded select list L_s' . If this update does not change the attributes of R_i that appear in either L_s' or the condition C_{join} , it will not affect the existing tuples in V_{PM} . Hence, there is no need to maintain V_{PM} . (Deletion influences all the attributes of R_i and thus does not have this optimization.) Otherwise we proceed in a way similar to that in the case of deletion.

Note that in the case of delete or update, computing the join between ΔR_i and the other base relations can be costly. In many cases, we can avoid this join computation by building indices on some attributes of V_{PM} . Due to space constraints, the details of this method are available in [25].

3.5 Using Better Cache Management Method

Consider a basic condition part bcp that exists in the partial materialized view V_{PM} . Tuples in V_{PM} often have either a large number of attributes or some long attributes (e.g., detailed description). As a result, the combined size of all the tuples in V_{PM} that belong to bcp is usually much larger than the size of bcp . If we treat bcp as the page id, and all the tuples in V_{PM} that belong to bcp as the page, then V_{PM} looks much like a buffer pool. Hence, instead of using the CLOCK algorithm, the RDBMS can use other better buffer pool management algorithms (e.g., 2Q [23]) to manage V_{PM} . This will increase the probability that V_{PM} can provide partial results to queries from the template q_t . The experimental section 4.1 gives a performance comparison between CLOCK and 2Q.

3.6 Discussions and Summary of Advantages

Like traditional MVs, the standard locking protocol is used on PMVs to ensure serializability. When a query Q reads a partial materialized view V_{PM} in Operation O_2 , Q puts an S lock on V_{PM} . Then between Operations O_2 and O_3 , no other transaction can change the correct (V_{PM}) read result of Q by updating some base relation, as that would require updating V_{PM} with the acquisition of an X lock on V_{PM} . Hence, Q would not have read anomaly.

With minor changes in our algorithm, PMVs can be used to handle queries with distinct clauses. In Operation O_2 , only distinct tuples in the partial results obtained from the PMV are returned to the user and stored in the data structure D_S . In Operation O_3 , all distinct result tuples are first obtained from query execution. Then only those tuples that are not in D_S are returned to the user.

The above discussion focuses on non-aggregate queries, which are common these days. For example, both the call center scenario in the introduction and deep analytical tasks in real-time data warehouses require detailed data. With minor changes in the user interface, PMVs can also be used to handle aggregate queries (e.g., group by) or queries with order by clauses. The details are available in [25].

In certain cases, with some extension, PMVs can be used to handle nested queries. For example, consider a two-level nested query. The subquery appears in the where clause of the main query after an EXISTS operator. Suppose that we can quickly obtain tuples from the main query but checking the EXISTS condition is time-consuming. In this case, a PMV can be used to quickly generate partial results of the subquery. Then for some tuples from the main query, the process of checking the EXISTS condition can be sped up. Consequently, we can rapidly produce some partial results for the entire query.

The partial materialized view V_{PM} has the following advantages:

- (1) V_{PM} has small storage and maintenance overhead.
- (2) V_{PM} can provide immediate partial results to a large number of queries from the template q_t .

- (3) A large portion of provided partial results are hot results – they are frequently accessed by other queries from q_i . This is desirable for those applications where users care more about hot results than cold results. (For applications that users want to see random partial results, this can be a disadvantage. However, as shown in [10], in general it is difficult to provide random partial results.)
- (4) V_{PM} has minor influence on queries’ execution time.

4. Performance Evaluation of Partial Materialized View

The performance of our PMV method has been evaluated from three perspectives: (1) the probability that a PMV can provide partial results to a query, (2) the influence of the PMV method on queries’ execution time, and (3) the maintenance overhead of a PMV when its base relations get changed.

4.1 Probability of Being Useful

We first perform a simulation study to show that in a large number of cases, PMVs can provide partial results to a query. Consider a read-only database. We focus on those queries that come from the same template q_i . Assume that a partial materialized view V_{PM} is built for q_i . In Operation O_1 , the C_{select} of each query is broken into the same number $h \geq 1$ of condition parts, where each condition part is a basic condition part itself. The entire query space contains 1M basic condition parts bcp_i ($1 \leq i \leq 1M$). For each basic condition part, the number of query result tuples that belong to it is greater than F . As a result, for each basic condition part that exists in V_{PM} , F query result tuples are stored in V_{PM} . For each basic condition part in the C_{select} of a query, the probability that it is bcp_i ($1 \leq i \leq 1M$) is e_i . All the e_i ’s ($1 \leq i \leq 1M$) follow a Zipfian distribution with parameter α . That is, $e_i \propto 1/i^\alpha$.

We compare the following two methods of managing all the basic condition parts in V_{PM} :

- (1) The **CLOCK** algorithm. V_{PM} is a queue with L entries that is managed by the CLOCK algorithm. Each entry can store one basic condition part bcp and F query result tuples that belong to bcp .
- (2) A simplified version of the **2Q** algorithm [23]. V_{PM} is composed of two queues: Am and Al . Am has N entries and is managed by the CLOCK algorithm. Each entry can store one basic condition part bcp and F query result tuples that belong to bcp . Al has $N'=50\% \times N$ entries and is a FIFO queue. Each entry stores one basic condition part. Upon the first time that a basic condition part bcp appears in the C_{select} of a query, bcp is put into Al . If during its stay in Al , bcp appears again in the C_{select} of another query, both bcp and F query result tuples that belong to bcp are moved to Am . Am is used to provide partial results to a query.

We assume that the storage requirement of a basic condition part is 4% of that of F query result tuples. Thus, given the same storage budget U_B of V_{PM} for both the CLOCK and the 2Q algorithms, we have $L=1.02 \times N$.

The purpose of the comparison between the CLOCK algorithm and the 2Q algorithm is to show that in a large number of cases, the simple CLOCK algorithm performs well. Also, CLOCK is not the best algorithm for managing all the basic condition parts in V_{PM} . In many cases, 2Q performs better than CLOCK. We leave it as an interesting area for future work to identify other algorithms that perform better than both CLOCK and 2Q.

We performed the following two experiments:

Number of bcps experiment. We fixed $N=20K$ and tested two cases:

- (i) $\alpha=1.07$. This is the high skew case. 10% of all the 1M basic condition parts get 90% of the chance of appearing in the C_{select} of a query.
- (ii) $\alpha=1.01$. This is the moderate skew case. 21% of all the 1M basic condition parts get 90% of the chance of appearing in the C_{select} of a query.

In either case, we varied h from 1 to 5. Recall that h is the number of basic condition parts in the C_{select} of a query.

PMV size experiment. We fixed $\alpha=1.07$ and $h=2$. We varied N from 10K to 30K. Recall that N determines the size of V_{PM} .

The hit probability is defined as the probability that V_{PM} can provide some partial results to a query Q . That is, if any of the h basic condition parts in the C_{select} of Q exists in V_{PM} , Q is “hit.” This definition is different from that in traditional caching [23], as our case is about “partial hit” while traditional caching is about “full hit.” In each test case, 1M queries were used to “warm up” V_{PM} . Then the hit probability was reported over the next 1M queries. (We also tested other numbers of “warm up” queries. The results were similar and thus omitted.)

For the number of bcps experiment, Figure 6 shows the hit probability results. The y-axis starts from 50%. h is the number of basic condition parts in the C_{select} of a query Q . If any basic condition part in the C_{select} of Q is “hit,” Q is “hit.” Hence, the hit probability approaches 100% quickly as h increases. The larger the α , the more queries focus on a few basic condition parts and thus the more likely these basic condition parts are cached in V_{PM} . Therefore, for a fixed algorithm (either CLOCK or 2Q) and a fixed h , the hit probability increases with α . For a fixed α and a fixed h , 2Q performs better than CLOCK, which is consistent with the results in [23].

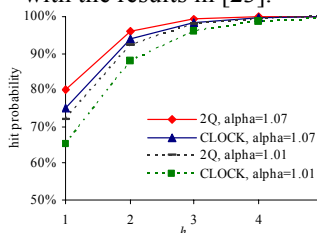


Figure 6. Hit probability (number of bcps experiment).

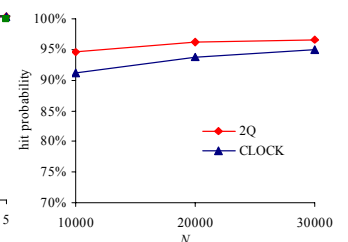


Figure 7. Hit probability (PMV size experiment).

Figure 7 shows the hit probability results from the *PMV* size experiment. The y-axis starts from 70%. The larger the N , the more basic condition parts and their corresponding query result tuples can be stored in V_{PM} , and thus the more likely V_{PM} can provide some partial results to a query. Therefore, the hit probability approaches 100% quickly as N increases. Again, for a fixed N , 2Q performs better than CLOCK.

4.2 Influence on Queries' Execution Time

In order to show that the PMV method has negligible influence on queries' execution time, we did a prototype implementation of our techniques in PostgreSQL Version 7.3.4 [28] for read-only database. Our measurements were performed with the PostgreSQL client application and server running on a computer with one 2.2GHz processor, 512MB main memory, one 40GB disk, and running the Microsoft Windows XP operating system. The default setting of PostgreSQL was used, where the buffer pool size is 1,000 pages. (We also tested larger buffer pool sizes. The results were similar and thus omitted.)

The relations used for the experiments followed the schema of the standard TPC-R Benchmark relations [32]:

customer (custkey, nationkey, ...),
orders (orderkey, custkey, orderdate, ...),
lineitem (orderkey, suppley, ...).

Table 1. Test data set.

	number of tuples	total size
customer	$0.15 \times s$ M	$23 \times s$ MB
orders	$1.5 \times s$ M	$114 \times s$ MB
lineitem	$6 \times s$ M	$755 \times s$ MB

s is the scale factor of the database. In our experiments, on average, each *customer* tuple matches ten *orders* tuples on the attribute *custkey*. Each *orders* tuple matches 4 *lineitem* tuples on the attribute *orderkey*.

We used the following two query templates:

Template T_1 : Find lineitems whose parts were provided by certain suppliers and sold on certain days.

```
select * from orders o, lineitem l
where o.orderkey=l.orderkey
and (o.orderdate= $d_1$  or ... or o.orderdate= $d_e$ )
and (l.suppley= $s_1$  or ... or l.suppley= $s_j$ );
```

Template T_2 : Find lineitems whose parts were provided by certain suppliers and sold to certain customers on certain days.

```
select * from orders o, lineitem l, customer c
where o.orderkey=l.orderkey and o.custkey=c.custkey
and (o.orderdate= $d_1$  or ... or o.orderdate= $d_e$ )
and (l.suppley= $s_1$  or ... or l.suppley= $s_j$ )
and (c.nationkey= $n_1$  or ... or c.nationkey= $n_g$ );
```

We built an index on each selection/join attribute. Before we ran queries, we ran the PostgreSQL statistics collection program on all the relations. For either template, due to the low selectivity of the selection attributes, the query plan is

not fully pipelined and thus traditional query execution cannot provide any result until it almost finishes.

For the template T_1 , each basic condition part is a 2-tuple (d_i, s_j) . For the template T_2 , each basic condition part is a 3-tuple (d_i, s_j, n_k) . We built two PMVs, one for T_1 and the other for T_2 . Either PMV contains 20K entries. Each entry can store one basic condition part bcp and F query result tuples that belong to bcp . (For each basic condition part, the number of query result tuples that belong to it is greater than F .)

For the template T_1 , its combination factor is defined as $h=exf$. For the template T_2 , its combination factor is defined as $h=exfg$. In Operation O_1 , the C_{select} of each query from T_1/T_2 is broken into the same number (h) of condition parts, where each condition part is a basic condition part itself, and one of these h basic condition parts exists in the PMV. We performed three experiments. Each experiment was repeated a large number of times (a large number of runs). All the reported numbers are averaged over these runs.

Number of Tuples

In this experiment, we fixed $h=4$ and $s=1$. We varied F , the number of query result tuples stored in each entry of the




Figure 8. Overhead of our techniques (number of tuples experiment).

PMV, from 1 to 5. Figure 8 shows the impact of F on the overhead of our techniques. The overhead of our techniques increases with F . This is easy to understand, as for each entry of the PMV, F query result tuples stored there are checked. For a fixed F , the overhead of our techniques for the template T_2 is greater than that for the template T_1 . This is because T_2 is more complex than T_1 : T_2 joins three relations, while T_1 joins two relations. As a result, the basic condition parts generated from T_2 are more complex than those generated from T_1 . Also, the query result tuples of T_2 are longer than that of T_1 . Recall that in our PMV method, both basic condition parts and query result tuples are checked.

Combination Factor




Figure 9. Overhead of our techniques (combination factor experiment).

In this experiment, we fixed $F=3$ and $s=1$. We varied the combination factor h from 1 to 10. Figure 9 shows the impact of h on the overhead of our techniques. The larger the h , the more basic condition parts a query generated. Then

the RDBMS needs to spend more time on dealing with all these basic condition parts. As a result, the overhead of our techniques increases with h . Again, for a fixed h , the overhead of our techniques for the template T_2 is greater than that for the template T_1 .

Database Scale Factor

In this experiment, we fixed $h=4$ and $F=3$. We varied the database scale factor s from 0.5 to 2. The purpose of this experiment is to show that our techniques have negligible influence on queries' execution time.

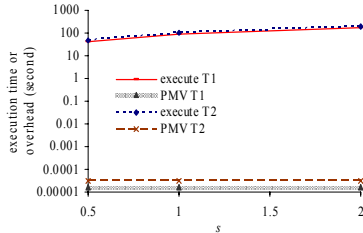


Figure 10. Query execution time vs. overhead of our techniques (database scale factor experiment).

represent the query execution time. The y-axis uses logarithmic scale.

Our techniques examine query result tuples rather than the data set. Also, our techniques mainly perform fast in-memory operations (recall that a significant portion of the PMV is cached in memory). Hence, compared to the query execution time, the overhead of our techniques is more than five orders of magnitude smaller. Since the cost of Operations O_1 and O_2 is less than the overhead of our techniques, the RDBMS can use the PMV to provide partial query results within a millisecond.

4.3 Maintenance Overhead

We use an analytical model to gain insight into the maintenance overhead of PMVs vs. MVs when their base relations get changed. A similar analytical model for MV maintenance has been validated in a commercial RDBMS (NCR Teradata) in [24]. Due to space constraints, the details of our model are available in [25]. (PostgreSQL currently does not support MVs. As a result, we were not able to compare the actual maintenance overhead of PMVs vs. MVs in PostgreSQL.)

Consider the template in Figure 1 and its corresponding partial materialized view V_{PM} . The materialized view V_M in Figure 2 is the containing MV of V_{PM} . We consider a single transaction T , where $p \times |\Delta R|$ tuples are inserted into R and $(1-p) \times |\Delta R|$ tuples are deleted from R . For both V_M and V_{PM} , the same updates must be performed on base relation R . Because of this, our model omits the cost of these updates and compares the maintenance cost of V_M and V_{PM} . The total workload TW for transaction T , which is defined as the total work done, is used as the cost metric.

Setting $|\Delta R|=1,000$, we present in Figures 11 and 12 the resulting performance of both the MV method and the PMV method. Figure 11 shows the total workload for

transaction T . The y-axis uses logarithmic scale. The maintenance of V_{PM} mainly performs cheap in-memory operations, while the maintenance of V_M requires a large number of expensive I/Os. Hence, for a fixed percentage of insertion p , maintaining V_{PM} is at least two orders of magnitude cheaper than maintaining V_M .

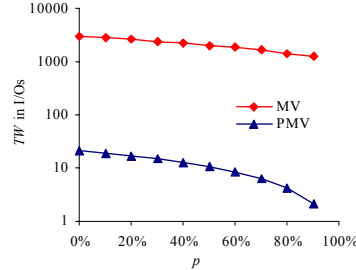


Figure 11. TW for transaction T .

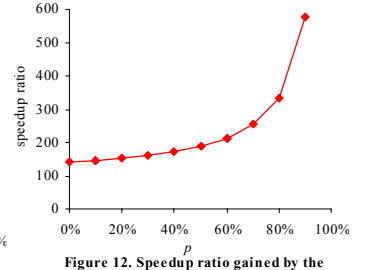


Figure 12. Speedup ratio gained by the partial materialized view method.

Inserting a tuple into V_M is less expensive than deleting a tuple from V_M . Also, there is no need to maintain V_{PM} in the presence of insertion into base relation R . As a result, both the maintenance overhead of V_M and the maintenance overhead of V_{PM} decrease as p increases. When $p=100\%$, the maintenance overhead of V_{PM} is 0. However, this cannot be shown in Figure 11, as the y-axis is on logarithmic scale.

Figure 12 shows the speedup ratio gained by maintaining V_{PM} over maintaining V_M . This speedup ratio increases with the percentage of insertion p , as there is no need to maintain V_{PM} during insertion into base relation R .

The above two-relation model can be easily extended to handle a (partial) MV defined on multiple base relations. However, this does not provide any insight not shown above and is thus omitted.

5. Related Work

To facilitate exploration of massive data sets, [20, 19] proposed using online aggregation to return approximate answers to the user immediately after a query is submitted to the RDBMS. Online aggregation focuses on aggregate queries. In contrast, our PMV method works for non-aggregate queries.

[1, 4] proposed building histograms “for free” by analyzing query results rather than checking the relation. In our case, if base relations do not change, the RDBMS both fills content into and updates PMVs “for free” by analyzing query results.

[30, 31] use partial indices to reduce index maintenance overhead. Upon an insertion into a relation R , the partial index I_p on R needs to be maintained immediately if this insertion satisfies the selection condition in the definition of I_p . In contrast, the PMV defined on R is not maintained immediately.

[26] proposed using sample MVs to support approximate query answering. A sample MV is a random sample of tuples in a MV. The maintenance of sample MVs is more expensive than that of PMVs, as randomness needs to be maintained in sample MVs. Also, since a sample MV does

not focus on hot query result tuples, the probability that it can provide partial results to a query is low. In a read-only environment, [15] proposed using icicles samples to support approximate query answering for key-foreign key join queries. In contrast, PMVs work in a general environment that allows updates.

[13] uses chunks to cache OLAP query results in the middle tier. [13] focuses on aggregate queries in a read-only environment, and imposes an order on each dimension if no implicit order exists. In contrast, our method works for non-aggregate queries in a general environment that allows updates, and does not impose non-natural orders on attribute values.

In a data stream environment, to speed up the processing of continuous multi-way windowed join queries, [6] proposed caching a subset of the join result tuples of some of the streams. If a key value v exists in the cache, all the join result tuples related to v must also exist in the cache. This requires maintaining the cache immediately upon arrival of new tuples from the streams. In contrast, upon insertion into base relations, PMVs are not maintained immediately.

In a distributed data integration environment, [21] and [27] define a PMV as a MV whose definition contains a subset of all the attributes and a where clause, respectively. Both PMV definitions are different from the one used in this paper.

6. Conclusion

We have presented a partial materialized view method that can provide transactionally consistent, immediate partial query results to the user without increasing queries' execution time much, by caching hot query results in PMVs. Our experiments with a simulation study, a theoretical analysis, and a prototype implementation in PostgreSQL show that PMVs have low storage and maintenance overhead. In a large number of cases, they can provide partial results almost instantly. Many PMVs can reside in the RDBMS simultaneously. Also, our method has negligible influence on queries' execution time. Hence, our PMV method can greatly facilitate the exploration of massive data sets. Besides providing partial query results, our techniques can be extended to address other problems, such as ranking query result tuples according to their popularity. The details are available in [25].

Acknowledgements

We would like to thank Jiuxing Liu, Jeffrey F. Naughton, Ying-li Tian, Michail Vlachos, Haijing Wang, Michael W. Watzke, and Hao Yang for helpful discussions.

References

[1] A. Aboulnaga, S. Chaudhuri. Self-tuning Histograms: Building Histograms without Looking at Data. SIGMOD Conf. 1999: 181-192.
[2] S. Agrawal, S. Chaudhuri, and V.R. Narasayya. Automated Selection of Materialized Views and Indexes in SQL Databases. VLDB 2000: 496-505.

[3] C. Bornhövd, M. Altnel, and S. Krishnamurthy et al. DBCache: Middle-tier Database Caching for Highly Scalable e-Business Architectures. SIGMOD Conf. 2003: 662.
[4] N. Bruno, S. Chaudhuri, and L. Gravano. STHoles: A Multidimensional Workload-Aware Histogram. SIGMOD Conf. 2001: 211-222.
[5] N. Bruno, S. Chaudhuri, and L. Gravano. Top-k Selection Queries over Relational Databases: Mapping Strategies and Performance Evaluation. TODS 27(2): 153-187, 2002.
[6] S. Babu, K. Munagala, and J. Widom et al. Adaptive Caching for Continuous Queries. ICDE 2005: 118-129.
[7] S. Chandrasekaran, O. Cooper, and A. Deshpande et al. TelegraphCQ: Continuous Dataflow Processing for an Uncertain World. CIDR 2003.
[8] M.J. Carey, D. Kossmann. On Saying "Enough Already!" in SQL. SIGMOD Conf. 1997: 219-230.
[9] S. Chaudhuri, R. Krishnamurthy, and S. Potamianos et al. Optimizing Queries with Materialized Views. ICDE 1995: 190-200.
[10] S. Chaudhuri, R. Motwani, and V.R. Narasayya. On Random Sampling over Joins. SIGMOD Conf. 1999: 263-274.
[11] J. Dougherty, R. Kohavi, and M. Sahami. Supervised and Unsupervised Discretization of Continuous Features. ICML 1995: 194-202.
[12] D. Donjerkovic, R. Ramakrishnan. Probabilistic Optimization of Top N Queries. VLDB 1999: 411-422.
[13] P. Deshpande, K. Ramasamy, and A. Shukla et al. Caching Multidimensional Queries Using Chunks. SIGMOD Conf. 1998: 259-270.
[14] J. Goldstein, P. Larson. Optimizing Queries Using Materialized Views: A Practical, Scalable Solution. SIGMOD Conf. 2001: 331-342.
[15] V. Ganti, M. Lee, and R. Ramakrishnan. ICICLES: Self-Tuning Samples for Approximate Query Answering. VLDB 2000: 176-187.
[16] H. Guo, P. Larson, and R. Ramakrishnan et al. Relaxed Currency and Consistency: How to Say "Good Enough" in SQL. SIGMOD Conf. 2004: 815-826.
[17] A. Gupta, I.S. Mumick. Materialized Views: Techniques, Implementations, and Applications. MIT Press, 1999.
[18] A.Y. Halevy. Answering Queries Using Views: A Survey. VLDB J. 10(4): 270-294, 2001.
[19] J.M. Hellerstein, P.J. Haas, and H. Wang. Online Aggregation. SIGMOD Conf. 1997: 171-182.
[20] P.J. Haas, J.M. Hellerstein. Ripple Joins for Online Aggregation. SIGMOD Conf. 1999: 287-298.
[21] R. Hull, G. Zhou. A Framework for Supporting Data Integration Using the Materialized and Virtual Approaches. SIGMOD Conf. 1996: 481-492.
[22] I.F. Ilyas, W.G. Aref, and A.K. Elmagarmid. Supporting Top-k Join Queries in Relational Databases. VLDB J. 13(3): 207-221, 2004.
[23] T. Johnson, D. Shasha. 2Q: A Low Overhead High Performance Buffer Management Replacement Algorithm. VLDB 1994: 439-450.
[24] G. Luo, J.F. Naughton, and C.J. Ellmann et al. A Comparison of Three Methods for Join View Maintenance in Parallel RDBMS. ICDE 2003: 177-188.
[25] G. Luo. Partial Materialized Views, full version. Available at http://www.cs.wisc.edu/~gangluo/partial_full.pdf, 2006.
[26] F. Olken, D. Rotem. Maintenance of Materialized Views of Sampling Queries. ICDE 1992: 632-641.
[27] Oracle Label Security. <http://www.oracle.com/technology/docs/dep/label/security/pdf/olsag.pdf>, 2000.
[28] PostgreSQL homepage, 2005. <http://www.postgresql.org>.
[29] A. Silberschatz, P. Galvin, and G. Gagne. Operating System Concepts, Sixth Edition. John Wiley, 2002.
[30] P. Seshadri, A.N. Swami. Generalized Partial Indexes. ICDE 1995: 420-427.
[31] M. Stonebraker. The Case for Partial Indexes. SIGMOD Record 18(4): 4-11, 1989.
[32] TPC Homepage. TPC-R benchmark, www.tpc.org.
[33] D.C. Zilio, J. Rao, and S. Lightstone et al. DB2 Design Advisor: Integrated Automatic Physical Database Design. VLDB 2004: 1087-1097.