Paper Number 1138

Physical Database Design for Data Warehouses *

Wilburt Juan Labio[†] Dallan Quass[‡] Brad Adelberg[§]

Abstract

Data warehouses collect copies of information from remote sources into a single database. Since the remote data is cached at the warehouse, it appears as local relations to the users of the warehouse. To improve query response time, the warehouse administrator (WHA) will often materialize views defined on the local relations to support common or complicated queries. Unfortunately, the requirement to keep the views consistent with the local relations creates additional overhead when the remote sources change. The warehouse is often kept only loosely consistent with the sources: it is periodically refreshed with changes sent from the source. When this happens, the warehouse is taken off-line until the local relations and materialized views can be updated. Clearly, the users would prefer as little down time as possible. Often the down time can be reduced by *adding* carefully selected materialized views or indexes to the physical schema. This paper studies how to select the sets of supporting views and of indexes to materialize to minimize the down time. We call this the view index selection (VIS) problem. We present an A* search based solution to the problem as well as rules of thumb. We also perform additional experiments to understand the space-time tradeoff as it applies to data warehouses.

Keywords: data warehouses, materialized views, view maintenance, index selection, and physical database design, A*.

1 Introduction

Data warehouses collect information from many sources into a single database. This allows users to pose queries within a single environment and without concern for schema integration. Figure 1 shows a typical warehousing system. Relations R_{src}, S_{src}, and T_{src}, referred to as *source relations*, from sources 1, 2, and 3 respectively, are replicated at the warehouse as R, S, and T in order to answer user queries posed at the warehouse such as $R \bowtie S \bowtie T$. We refer to the replicated relations R, S, and T as *warehouse relations*. Consistency between the source relations and the warehouse relations is usually only loosely maintained: Changes to the source relations are queued and periodically shipped and applied to the warehouse relations. We call these changes *deltas*.

Queries posed at a data warehouse are often complex—involving joins of multiple relations as well as aggregation. Due to the complexity of these queries, *views* are usually defined — a view is a derived relation expressed in terms of the warehouse relations. Because the views are defined

*This work was supported by Hewlett Packard, by Philips, and by Rome Laboratories under Air Force Contract F30602-94-C-0237. This work was also supported by an equipment grant from Digital Equipment Corporation.

[†]Stanford University Department of Computer Science. e-mail: wilburt@cs.stanford.edu

[‡]Stanford University Department of Computer Science. e-mail: quass@cs.stanford.edu

[§]Stanford University Department of Computer Science. e-mail: adelberg@cs.stanford.edu

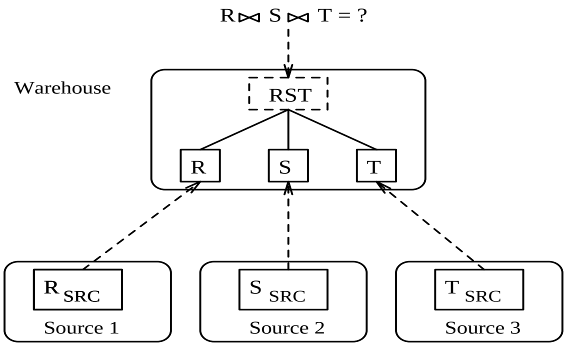

Figure 1: Warehouse with primary view.

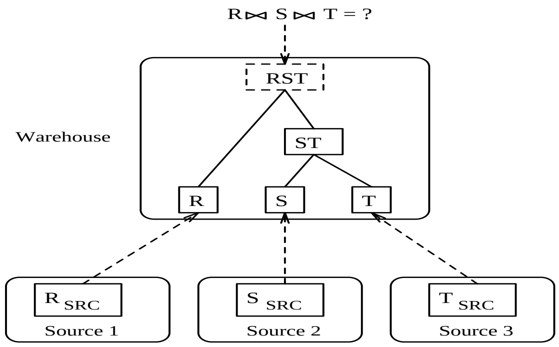

Figure 2: Warehouse with supporting view.

in terms of the warehouse relations, we refer to the warehouse relations also as *base relations*. For example, referring again to Figure 1, RST represents a view that is the expression $R \bowtie S \bowtie T$. Warehouses can store large amounts of data, and so in order to improve the performance of queries written in terms of the views, the views are often *materialized* by storing the result of the view at the warehouse. Unmaterialized views are called *virtual* views. Queries written in terms of materialized views can be significantly faster than queries written in terms of virtual views because the view tuples are stored rather than having to be recomputed.

Since materialized views are computed once and then stored, they become inconsistent as the deltas from the sources are applied to the base relations. In order to make a materialized view consistent again with the base relations from which it is derived, the view may either be recomputed from scratch, or *incrementally maintained* [GMS93] by calculating just the effects of the deltas on the view. These effects are captured in *view maintenance expressions* [GL95]. Each type of change (insertion, deletion, or update) requires a different expression. For example, if view RST in Figure 1 is materialized, the maintenance expression calculating the tuples to insert into RST due to insertions into R is $\triangle R \bowtie S \bowtie T$, where $\triangle R$ denotes the insertions into R.

Since the sizes of the views at a warehouse are usually so large and the changes small in comparison, it is often much cheaper to incrementally maintain the view than to recompute it from scratch. Incrementally maintaining a number of materialized views at a warehouse, even though cheaper than recomputing the views from scratch, may still involve a significant processing effort. To avoid impacting clients querying the warehouse views, view maintenance is usually performed at night during which time the warehouse is made unavailable for answering queries. A major concern for warehouses using this approach is that the views be maintained in time for the warehouse to be available for querying again the next morning. An important problem for data warehousing is thus: Given a set of materialized views that need to be maintained due to a set of deltas shipped from the data sources, how is it possible to reduce the total maintenance time?

Our approach to the problem of minimizing the time spent maintaining a set of views may seem counter-intuitive at first: add additional views and/or indexes. However, this is analogous to building indexes in traditional RDBMS's. For example, having an index on the key of a relation can vastly decrease the total time spent locating a particular tuple even though the index must be maintained as well. In this paper we will approximate maintenance time as the number of I/O's required and then endeavor to minimize the number of I/O's performed. We start with the number of I/O's required for maintaining the materialized views and the base relations at the warehouse. We then add a set of additional views and indexes that themselves must be maintained, but whose *benefit* (reduction in I/O's) outweighs the *cost* (increase of I/O's) of maintaining them. As an example, let us return to Figure 1. Suppose that in addition to materializing the *primary view*, RST, another view, ST, is also materialized. This situation is shown in Figure 2. By materializing view ST, the total cost of maintaining both RST and ST can be less than the cost of maintaining RST alone. For example, suppose that there are insertions to R but no changes whatsoever to S and T. To propagate the insertions to R onto RST, we must evaluate the maintenance expression that calculates the tuples to insert into RST due to insertions into R, which is $\triangle R \bowtie S \bowtie T$. With ST materialized, it is almost certain that this expression can be evaluated more efficiently as $\triangle R \bowtie ST$, joining the insertions to R with ST, instead of with S and T individually. Even if there are changes to S and T, the benefit of materializing ST may still outweigh the extra cost involved in maintaining it. Since the view ST is materialized to assist in the maintenance of the primary view RST, we call the view ST a *supporting view*.

In addition to materializing supporting views, it may also be beneficial to materialize indexes. Indexes may be built on the base relations, primary views, and on the supporting views. The general problem, then, is to choose a set of supporting views and a set of indexes to materialize such that the total maintenance cost for the warehouse is minimized. We call this the *View-Index Selection (VIS)* problem and it is the focus of this paper.

Below we list the primary contributions of this paper.

- We propose and implement an optimal algorithm based on A* that prunes as much as 99% of the possible supporting view and index sets to solve the VIS problem.
- Through both cost/benefit analysis and experimentation, we develop a number of rules of thumb that can help a warehouse administrator (WHA) find a reasonable set of supporting views and indexes to materialize in order to reduce the total maintenance cost.
- We compare the benefit of materializing supporting views as opposed to indexes, and discuss which should be chosen when the total storage space at the warehouse is constrained.
- We perform experiments to determine how sensitive the choice of supporting view and index sets are to the input parameters of the optimizer.

The rest of the paper proceeds as follows. Section 2 describes the VIS problem in detail. Section 3 presents the scope of our results and our approach to view maintenance. We describe our A*-based algorithm in Section 4. Section 5 develops rules of thumb for choosing a set of supporting views and indexes to materialize. We justify our rules both by a cost model analysis as well as by extensive experimentation using our A*-based algorithm. In Section 6, we report on additional experiments such as comparing the importance of indexes and supporting views when space is constrained. Next, in Section 7, we discuss how this paper relates to previous work in the area. Finally, we present our conclusions in Section 8.

2 General Problem

Having introduced the VIS problem, in this section we describe it fully and present an exhaustive search algorithm to solve it. We also show the worst case complexity of the VIS problem. Lastly, we present an example schema to illustrate the concepts introduced.

2.1 The Optimization Problem

An optimal algorithm must minimize the total cost of maintaining the warehouse. The total cost that we attempt to minimize is the sum of the costs of: (1) applying the deltas to the base relations, (2) evaluating the maintenance expressions for the materialized views, and (3) modifying affected indexes. The cost of maintaining one view differs depending upon what other views are available. It is therefore incorrect to calculate the cost of maintaining the original view and each of the additional views in isolation. Moreover, in order to derive the total cost it is necessary to consider the view selection and index selection together. If view selection is performed separately from index selection, it is not hard to concoct cases wherein a supporting view V is considered to be too expensive to maintain without indexes, but where V is actually part of the optimal solution since it may become feasible to maintain when the proper indexes are built.

To find the optimal solution, then, we must solve the optimization problem globally. One approach, proposed in Ross et al. [RSS96] (although this work does not consider indices), is to exhaustively search the solution space. Although exhaustive search is impractical for large problems, it illustrates the complexity of the problem and provides a basis of comparison for other solutions. The exhaustive algorithm works as follows (each stage is described below):

```
for each possible subset of supporting views
  for each possible subset of indexes on the views and base relations
    compute total cost of the materialized views and indexes and keep
    track of the supporting views and indexes that obtain the minimum cost
```

2.1.1 Choosing the views

In the first step we consider all possible subsets of the set of candidate views $\mathcal{C}$. As proposed in [RSS96], we consider as candidate views all distinct nodes that appear in a query plan for the primary view. Since the primary view is already materialized, it is not included in the candidate view set. For example, given a view $V = R \bowtie S \bowtie T$, $\mathcal{C} = \{RS, RT, ST\}$. In general, for a view joining n relations there are roughly $\mathcal{O}(2^n)$ different nodes that appear in some query plan for the view, one joining each possible subset of the base relations. Thus, to consider all possible subsets of $\mathcal{C}$, we need to evaluate roughly $\mathcal{O}(2^{2^n})$ different view states.

2.1.2 Choosing the indexes

Now we must consider all possible subsets of the set of candidate indexes, $\mathcal{I}$. Candidate indexes, as defined in Finkelstein et. al. [FST88], are indexes on the following types of attributes:

- attributes with selection or join predicates on them. [1]
- key attributes for base relations where changes to the base relation include deletions or updates. When views are materialized on the base relations, key attributes of any base relation appearing in the view also qualify.
- attributes in GROUP BY or ORDER BY clauses.

Additional attributes can be candidates depending on the query optimizer being used. The reader is refered to [FST88] for more detail.

Since each materialized view will usually have candidate indexes, $\mathcal{I}$ must be recomputed at the beginning of every inner loop. The cardinality of $\mathcal{I}$ for a particular view state is roughly proportional to the number of materialized views and base relations in that state. Further, a particular view state contains between n and $\mathcal{O}(2^n)$ materialized views and base relations, so there can be as many as $\mathcal{O}(2^n)$ candidate indexes to consider. Since we must evaluate possible subsets of candidate indexes, the number of possible index states for a view state can be up to $\mathcal{O}(2^{2^n})$. (See Section 7 for an explanation of why standard approaches for index selection are not appropriate.) If we ignore global space constraints, we can do a little better than this since candidate indices on columns not involved in join predicates can be chosen greedily.

2.1.3 Computing the total update cost

Once a particular view and index state are chosen, obtaining the total cost is a query optimization problem in itself since it involves finding the most efficient query plan for each of the view maintenance expressions. Thus, the VIS problem for a single primary view joining n base relations contains roughly $\mathcal{O}(2^{2^n})$ query optimization problems in the most general case.

The query optimization itself is complicated by the presence of materialized views since the optimizer must also determine if it can use another materialized view in the query plan evaluating a maintenance expression. For example, given a view $V = R \bowtie S \bowtie T$, insertions to R are propagated onto V by the maintenance expression $\triangle R \bowtie S \bowtie T$. Suppose the view $ST = S \bowtie T$ is also materialized. The query optimization algorithm must consider the possibility of evaluating $\triangle R \bowtie S \bowtie T$ as $\triangle R \bowtie ST$ in finding the best query plan. This problem is known as "answering queries using views" [LMSS95].

To complicate matters, one batch of changes can generate multiple maintenance expressions that need to be evaluated. This happens due to different types of changes to the base relations. The maintenance expressions can be optimized as a group because of possible common subexpressions [RSS96]. This problem is known as the "multiple-query optimization" problem [Sel88].

2.2 Example

Consider the following base relations and view.

[1] In addition, the system must be able to use an index to process the predicate. This usually implies that the predicate is a simple comparison (except for $\neq$) or range operator and that the other operand is a constant or a column from a different table.

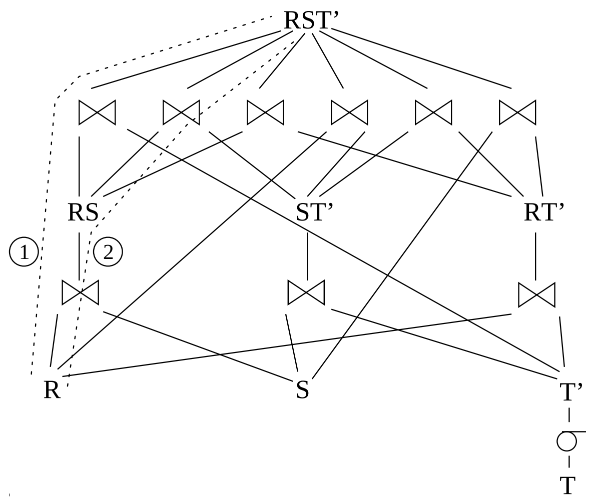

Figure 3: Example Schema.

```
R(R0,R1),  S(S0,S1),  T(T0,T1)

create view V(R0,R1,S0,S1,T0,T1) as
select *
from   R, S, T
where  R.R1 = S.S1 and S.S0 = T.T0 and T.T1 <= 10
```

Figure 3 shows an *expression dag* [RSS96] that includes all the nodes that could appear in a query plan for V, assuming the selection on $T.T1$ is pushed down. The view T' is the result of applying the selection condition to T. Under each view is the set of operations (join or select) that could be used to derive the view. For example, the view RST could be derived as the result of $R \bowtie S$ joined with T', or the result of $R \bowtie S$ joined with the result of $S \bowtie T'$, and so on. Each of the intermediate results could be materialized as a supporting view. Following the definition in Section 2.1.1, the set of candidate supporting views, $\mathcal{C}$, is $\{RS, ST', RT', T'\}$. Assuming V is materialized at a data warehouse (as well as the base relations), any possible subset of $\mathcal{C}$ might also be materialized as supporting views at the warehouse in order to minimize the total maintenance cost. In addition, indexes on V, the base relations, and the supporting views need to be considered.

It is useful to think of the expression dag in Figure 3 when considering the different *update paths* [RSS96] changes to base relations can take as they are propagated to the view. An update path corresponds to a specific query plan for evaluating a view maintenance expression. For example, the maintenance expression for propagating insertions to R onto V is to insert the result of $\triangle R \bowtie S \bowtie T'$ into V. The graph depicts seven possible update paths for this expression, two of which are shown in Figure 3: (1) $(\triangle R \bowtie S) \bowtie T'$, (2) $\triangle(R \bowtie S) \bowtie (S \bowtie T')$. Notice that the choice of update path can affect which indexes get materialized. If update path (1) is chosen, an index may be built on the join attribute of T' to help compute the maintenance expression. If path (2) is chosen however and view ST' is materialized, an index may be built on the join attribute of ST'. The implication of using different indexes depending upon which update path is chosen is discussed in Section 7.

Changes to base relations need to be propagated both to the primary view as well as to the supporting views that have been materialized. When propagating changes to several base relations onto several materialized views there are opportunities for multiple-query optimization. Results of maintenance expressions for one view can be reused when evaluating maintenance expressions for another view. For example, suppose view $RS = R \bowtie S$ is materialized. The result of propagating insertions to R onto RS, $\triangle R \bowtie S$, can be reused when propagating insertions to R onto V, $\triangle R \bowtie S \bowtie T'$, so that only the join with T' need be performed. In addition, common subexpressions can be detected between several maintenance expressions.

3 Problem Studied

As discussed in the previous section, the problem of choosing an optimal set of supporting views and indexes to materialize in order to minimize the maintenance cost of a view is very complex. While the algorithm we present is quite general, in order to perform cost model analysis and to simplify our coding effort, we have made simplifying assumptions to allow us to study the most important effects without getting lost in irrelevant detail. The resulting problem is still doubly-exponential and we feel that the insights we have gained from this study can lead to more general solutions. Furthermore, our assumptions are very similar to those made previously in the literature.

3.1 Database Model

[[[We should present these assumptions in a more coherent manner.]]] <==

We limit our consideration to maintaining a single select-join (SJ) view. Any combination of selections and joins in a view definition can be represented in this form. We assume for the simplicity of the cost model that the view does not involve self-joins and that all base relations have keys. In addition, we assume that all indexes are stored as B+-trees, that indexes are built on single attributes only, and indexes are built on relations and views stored as heaps. We consider the two most common physical join operators: nested-block joins and index joins.

We assume that the base relations from the source are replicated at the warehouse. In addition, we assume that *local selection conditions* (involving attributes of a single base relation) are always "pushed down" onto the base relations. When considering what additional data structures to materialize, we restrict ourselves to data structures that are themselves easily maintainable through SQL update statements. To this end we consider materializing supporting views that are *subviews* of the primary view. That is, they join a subset of the base relations in the primary view. We also consider building indexes on attributes in the base relations, primary view, and supporting views that are involved in selection and join conditions (involving attributes from two base relations).

A materialized supporting view V could thus be the result of applying selection conditions to a base relation (*e.g.*, $V = \sigma R$), or it could be the result of a join between two or more base relations, each having selection conditions pushed down (*e.g.*, $V = \sigma R \bowtie \sigma S$).

3.2 Change Propagation Model

We consider three types of deltas: insertions, deletions, and updates. We distinguish between two types of updates: Updates that alter the values of key attributes or attributes involved in selection conditions are called *exposed updates*; all other updates are called *protected updates*. Exposed updates can result in tuples being deleted from or inserted into the view. For this reason, we propagate exposed updates as deletions followed by insertions. Henceforth, all references to 'updates' should be interpreted to mean 'protected update'. Protected updates could also be propagated as deletions followed by insertions, but they can be applied directly to the view since they only change attribute values of tuples in the view, and never insert or remove tuples from the view.

We assume for the purposes of determining the cost of maintaining a view that each type of change to each base relation is propagated to the view and relevant supporting views separately. Therefore, the cost of maintaining a view or supporting view $V = R \bowtie S \bowtie T$ is the sum of the costs of propagating (onto V) each type of change to each of the base relations involved in V. For example:

- **Insertions** The cost of propagating insertions to R onto V is the cost of evaluating $\triangle R \bowtie S \bowtie T$, inserting the result into V, and updating the indexes of V. When propagating insertions it is often possible to reuse the results of evaluating insertions for one view in evaluating insertions for another. For example, if V is a supporting view of $V' = R \bowtie S \bowtie T \bowtie U$, then we allow the result of insertions to R onto V ($\triangle R \bowtie S \bowtie T$) to be reused when propagating insertions to R onto V' ($\triangle R \bowtie S \bowtie T \bowtie U$). In this respect we consider a limited form of multiple-query optimization.

- **Deletions** The cost of propagating deletions to R ($\bigtriangledown R$) onto V is the cost of evaluating $V \rtimes \bigtriangledown R$ (we use $\rtimes$ for semijoin), removing those tuples from V, and updating the indexes of V.

- **Updates** The cost of propagating updates to R (μR) onto V is the cost of evaluating $V \rtimes \mu R$ and updating those tuples in V. Note that because we only allow protected updates, we do not have to update the indexes of V since we only build indexes on attributes involved in selection conditions or keys and these attributes cannot be modified by a protected update.

4 Optimal Solution using A* algorithm

In this section we describe an optimal algorithm to solve the VIS problem and then show through experimental results that it vastly reduces the number of candidate solutions that must be considered.

4.1 Algorithm description

The A* algorithm is an improvement over exhaustive search because it attempts to prune the parts of the search space that cannot contain the optimal solution. In this section we describe how we have used the A* algorithm to solve the VIS problem. For further details on the A* algorithm itself, the reader is referred to [Nil71].

The algorithm takes as input the set of all possible views and indexes to materialize, $\mathcal{M}$. $\mathcal{M}$ does not include the base relations ($\mathcal{B}$) nor the primary view V but includes indexes that can be defined on them. (V and the base relations are constrained to be materialized.) The goal of the algorithm is to choose a subset $\mathcal{M}'$ of $\mathcal{M}$ to materialize such that the total cost, $\mathcal{C}$, is minimized. The total cost given a particular subset of views and indexes $\mathcal{M}'$ can be expressed as

$$C(\mathcal{M}') = \sum_{m \in (\mathcal{M}' \,\cup\, \mathcal{B} \,\cup\, \{V\})} \mathit{maint\_cost}(m, \mathcal{M}')$$

Function $\mathit{maint\_cost}(m, \mathcal{M}')$ returns the cost of propagating all changes to view or index m assuming only the views and indexes in $\mathcal{M}'$ are materialized (in addition to B and V).

Instead of directly searching the power set of $\mathcal{M}$, we set up the A* search to build the solution incrementally. It begins with an empty materialization set ($\mathcal{M}' = \phi$) and then considers adding single views or indexes. The algorithm terminates when a solution is found that has considered every view and index and is guaranteed to have the minimum total cost. We will call the intermediate steps reached in the algorithm *partial states*. Each partial state is described by the tuple $(\mathcal{M}_C, \mathcal{M}')$ where $\mathcal{M}_C$ is the set of features from $\mathcal{M}$ that have been considered and $\mathcal{M}'$ is the set of features from $\mathcal{M}_C$ that have been chosen to be materialized. For convenience, we will refer to set of unconsidered features, $\mathcal{M}_{\mathcal{U}}$, which is $\mathcal{M} - \mathcal{M}_C$.

Presented with a set of partial states from which to incrementally search, A* attempts to choose the most promising. It does so by *estimating* the cost of the best solution $\mathcal{M}' \cup \mathcal{M}'_{\mathcal{U}}$ ($\mathcal{M}'_{\mathcal{U}}$ is an estimate of the unconsidered features that would be chosen) that can be achieved from each partial state.

The *exact* cost of the best solution given a partial state can be decomposed as

$$\mathcal{C} = g + h$$

where g is the maintenance cost for the features chosen so far ($\mathcal{M}'$) and h is the maintenance cost for the features in $\mathcal{M}'_{\mathcal{U}}$. In general, g also needs $\mathcal{M}'_{\mathcal{U}}$ for its computation; that is, it is necessary to know which unconsidered features will be chosen in order to compute the maintenance cost of features in $\mathcal{M}'$. Fortunately, we can compute g using only $\mathcal{M}'$ so long as we impose a partial ordering on the features in $\mathcal{M}$ so that we only consider a feature when a decision has been made on every feature that affects its maintenance cost. Formally, a partial order $\prec$ is imposed upon $\mathcal{M}$ such that if a feature m_1 can be used in a query plan for propagating insertions to view m_2, then $m_1 \prec m_2$. Also, for an index m_1 on a view m_2, $m_2 \prec m_1$.

The exact formula for h is

$$\min_{\mathcal{M}'_{\mathcal{U}} \subseteq \mathcal{M}_{\mathcal{U}}} \Big(\sum_{m \in \mathcal{M}'_{\mathcal{U}}} \mathit{maint\_cost}(m, \mathcal{M}' \cup \mathcal{M}'_{\mathcal{U}}) \Big)$$

Unfortunately, this formula requires an exhaustive search to find the $\mathcal{M}'_{\mathcal{U}}$ that minimizes the equation.

Instead of performing this exhaustive search, we calculate a lower bound on h denoted $\hat{h}$. Using the lower bound, the A* algorithm can prune some of the partial states while still guaranteeing an

Input: $\mathcal{M}, \prec$
Output: Optimal $\mathcal{M}'$

Let state set $S = \{s\}$, where s is a partial state having
  $\mathcal{M}_\mathcal{C}(s) = \mathcal{M}'(s) = \phi$, and $\mathcal{M}_\mathcal{U}(s) = \mathcal{M}$ (base relations and V are materialized)
Loop
  Select the partial state $s \in S$ with the minimum value of $\hat{C}$
  If $\mathcal{M}_\mathcal{C}(s) \equiv \mathcal{M}$, return $\mathcal{M}'(s)$
  Let $S = S - \{s\}$
  For each view or index $m \in \mathcal{M}_\mathcal{U}(s)$ such that for all $m' \prec m$: $m' \in \mathcal{M}_\mathcal{C}(s)$
    Construct partial state s' such that
      $\mathcal{M}_\mathcal{C}(s') = \mathcal{M}_\mathcal{C}(s) \cup \{m\},\ \mathcal{M}_\mathcal{U}(s') = \mathcal{M}_\mathcal{U}(s) - \{m\},\ \mathcal{M}'(s') = \mathcal{M}_\mathcal{C}(s) \cup \{m\}$
    Construct partial state s'' such that
      $\mathcal{M}_\mathcal{C}(s'') = \mathcal{M}_\mathcal{C}(s) \cup \{m\},\ \mathcal{M}_\mathcal{U}(s'') = \mathcal{M}_\mathcal{U}(s) - \{m\},\ \mathcal{M}'(s'') = \mathcal{M}_\mathcal{C}(s)$
    Let $S = S \cup \{s'\} \cup \{s''\}$
  Endfor
Endloop

Table 1: A* Algorithm

optimal solution. Using $\hat{h}$, for any partial state we can compute a lower bound on $\mathcal{C}$ as

$$\hat{\mathcal{C}} = g + \hat{h}$$

Note that if $\mathcal{M}_\mathcal{C} \equiv \mathcal{M}$ then $\hat{\mathcal{C}} = \mathcal{C}$. We will develop an expression for $\hat{h}$ below but first we present the A* algorithm for the VIS problem.

The algorithm appears in Table 1. The state set S contains all active partial states. It initially contains only the partial state where none of the views and indexes have been considered. Each time through the loop the algorithm selects the partial state with the minimum lower bound on the cost. If the selected state has $\mathcal{M}_\mathcal{C} \equiv \mathcal{M}$, it is guaranteed to be the optimal choice. If the selected state is not a complete state, it is removed from the set of active states and for each view or index that can be added to the set of considered views and indexes without violating the partial order, two states are added to the set of active states: one with the view or index added to the chosen set ($\mathcal{M}'$), and one without.

The formula for $\hat{h}$ computes the cost of maintaining views and indexes in $\mathcal{M}_\mathcal{U}$ minus the upper bound of their benefit toward maintaining other views (including V).

$$\hat{h} = \sum_{m \in \mathcal{M}_\mathcal{U}} (h\_maint\_cost(m, \mathcal{M}') - max\_benefit(m, \mathcal{M}'))$$

We guarantee that any overestimation of the actual maintenance cost of m is more than compensated for by the overestimation of the benefit. Note that our function $\hat{h}$, although it achieves considerable pruning, can be improved. Deriving a tighter lower bound for h that can be computed efficiently is a subject for future research.

The function *h_maint_cost*$(m, \mathcal{M}')$ differs depending on whether m is a view or an index. If m is an index, the function returns the cost of maintaining m for all insertions and deletions that will be propagated to the view that m is on. (The details of our cost model are found in Appendix A.) If m is a view, the function returns the cost of propagating onto m insertions to each of the base relations referenced in m, plus the cost of propagating onto m deletions and updates to each of the base relations referenced in m assuming an index exists in m for the key attribute of each base relation. Note that when m is a view, we might overestimate the cost for propagating insertions since we are assuming that all other views in $\mathcal{M}_{\mathcal{U}}$ are not materialized (this is compensated for in *max_benefit*).

The function *max_benefit*$(m, \mathcal{M}')$ also differs depending on whether m is a view or an index. First we consider the case where m is an index.

1. If m is an index on a view v for the key attribute of a base relation R that is referenced in v, the function returns the cost of propagating deletions and updates from R to v without m minus the cost of propagating deletions and updates from R to v with m.

2. If m is an index on a view v for a join attribute that joins v to some relation R not referenced in v, the function sums for each view $v' \in \mathcal{M}_{\mathcal{U}}$ that includes R as well as all the relations in v and for every relation S in v' but not in v, the cost of scanning v (the maximum savings due to an index join using m when propagating insertions from s onto v').

3. If m is an index for both a key and a join attribute, the two benefits described are added.

Next we consider the case where m is a view. Intuitively, the maximum benefit of m is the cost of materializing m when propagating insertions to views for which m is a subview. The *max_benefit* function sums for each view $v' \in \mathcal{M}_{\mathcal{U}}$ that includes all the relations in m and for every relation S in v' but not in m, the cost of materializing m given the views and indexes in $\mathcal{M}_{\mathcal{C}}$.

4.2 Experimental results

To test the A* algorithm described in the previous section, we coded both it and the exhaustive algorithm described in Section 2. We then ran both algorithms on a variety of sample schemas. A summary of the results are presented in Table 2. Clearly, the A* algorithm is performing very well, pruning the vast majority of the search space. As the problems gets larger, due to more views or selection predicates, its relative performance increases as well. While it may still be possible to derive a tighter lower bound on h, even the algorithm as presented is a vast improvement over previously proposed algorithms.

5 Rules of Thumb

The A* algorithm presented in the last section yields optimal solutions but is often impractical except for small views. Fortunately, finding an optimal solution is not critical since there are

# of relations	# of selections	# of states visited exhaustive	A*	% pruned
2	0	32	11	67.7
2	1	192	21	89.1
2	2	960	28	97.1
2	4	960	29	97.0
3	1	2115072	17735	99.2
3	2	10575360	22809	99.8

Table 2: Comparison of A* and exhausive algorithms.

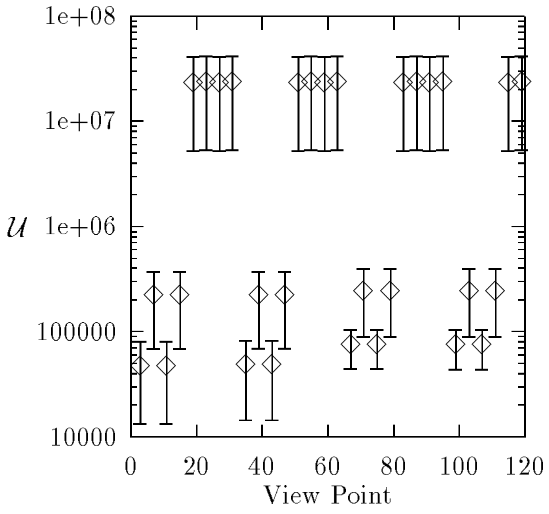

Figure 4: A Sample Solution Space

often many solutions that are close to optimal. In Figure 4, each point on the x-axis represents a particular view set. The range of values that a point can take (indicated by a vertical bar) is determined by the index sets that give the worst and the best update costs. This figure emphasizes two things: (1) there are several view sets that are close to optimal, and (2) index selection is very important even after picking a good view set. What is required, then, is to avoid poor view sets and then to pick a good index set.

In this section we propose rules of thumb that can help guide a WHA in choosing a reasonable set of supporting views and indexes without resorting to the full algorithm. The underlying theme of these rules of thumb is materialize a supporting view or an index if its benefit (reduction in I/O cost) is greater than its cost (increase in I/O cost). These rules of thumb function similarly to the rule "join small relations first" in query optimization. These are not hard and fast rules: many factors come into play and some rules tend to work against others. But we have found that the rules apply in general. Even when the rules do not apply, the cost-benefit analysis introduced in explaining each rule can help the WHA decide what to materialize in a specific situation.

We justify each rule of thumb through analysis and also through experimentation. The formulas we use in the analysis are rough approximations of the actual benefits and costs. However, a significantly more detailed and accurate cost model was used in the VIS optimizer (See Appendix A) which was used in the experiments. Since the rules of thumb are supported by the results of the VIS optimizer, it seems that the approximations used in these simpler formulas are reasonable.

The supporting experiments are performed for a view composed of only three base relations to keep the problem tractable. We expect that for a view joining more than three relations, the differences in the graphs would be more pronounced because there would be more opportunities to apply the rule. Due to space constraints, we only include the results of one representative experiment per rule although many more were conducted.

Notation	Description
$\mathcal{U}$	the total maintenance cost for the primary view, supporting views, and indexes
V	primary view, supporting view, or base relation
$\mathcal{R}(V)$	Base relations involved in V
$\lvert\mathcal{R}(V)\rvert$	Number of base relations involved in V
$\overline{\mathcal{R}(V)}$	Base relations not involved in V
$\mathcal{E}(V)$	Elements (materialized supporting views or base relations) joined in V
$\lvert\mathcal{E}(V)\rvert$	Number of elements joined in V
P_m	Number of pages of memory for database buffer
$P(V)$	Number of pages in V
$T(V)$	Number of tuples in V
$I(V)$	Number of insertions to V
$D(V)$	Number of deletions from V
$U(V)$	Number of updates to V
$P(\mathcal{R}(V))$	Sum of the pages in all of the base relations involved in V
$I(\mathcal{R}(V))$	Sum of the number of insertions to all of the base relations involved in V
$D(\mathcal{R}(V))$	Sum of the number of deletions from all of the base relations involved in V
$U(\mathcal{R}(V))$	Sum of the number of updates to all of the base relations involved in V
$P(\mathcal{E}(V))$	Sum of the pages in all of the elements of V
$P(V, R.A)$	Number of pages in an index on V for attribute $R.A$
$S(V, C)$	Number of tuples in V passing the selection condition C (if C is a join condition then it is the number of tuples in V that join with a single tuple in the other relation)

Table 3: Notation Used in Rules of Thumb

5.1 Schema and Notation

The statistics given in Table 3 are used in evaluating the rules of thumb. Since the rules of thumb are very approximate, the WHA needs only rough approximations of the statistics. The sensitivity of the results to estimation errors is studied in Section 6.2. Two points about the table need to be made. First, we also define the functions $P(\overline{\mathcal{R}(V)})$, $I(\overline{\mathcal{R}(V)})$, $D(\overline{\mathcal{R}(V)})$, and $U(\overline{\mathcal{R}(V)})$ to have their expected meaning. For instance, $P(\overline{\mathcal{R}(V)})$ denotes the sum of the pages in all of the base relations that are in the primary view but not in the view V. Second, if for the definition of $\mathcal{E}(V)$ there is more than one possible set of materialized supporting views and base relations that can be joined to derive V, then we assume that a set having the fewest number of elements (base relations

Schema	Relations	$T(V)$	$I(V)$	$D(V)$	Select	Join
1	R(<u>R0</u>, R1)	90M	1%	0.01%		$S(R,\ R1 = SO) = 3$
	S(<u>S0</u>, S1)	30M	1%	0.01%		$S(S,\ S0 = R1) = 1$
						$S(S,\ S1 = T0) = 3$
	T(<u>T0</u>, T1)	10M	1%	0.01%	$S(T,\ T1 <= 10)$ =1M	$S(T,\ T0 = S1) = 1)$
2	R(<u>R0</u>, R1, R2)	20M	1%	0.01%		$S(R,\ R0 = S1) = 1$
	S(<u>S0</u>, S1, S2)	20M	1%	0.01%	$S(S,\ T2 <= 20)$ =2M	$S(S,\ S1 = R0) = 1$
						$S(S,\ S2 = T0) = 1$
	T(<u>T0</u>, T1, T2)	20M	1%	0.01%		$S(T,\ T0 = S2) = 1)$

Figure 5: View Schemas

or supporting views) is chosen. For example, suppose that view V is defined joining R, S, and T, then $\mathcal{E}(V) = \{R, S, T\}$ and $\mathcal{R}(V) = \{R, S, T\}$. If another view $V' = RS$ is then materialized, $\mathcal{E}(V) = \{V', T\}$ because this set has only two elements (but $\mathcal{R}(V) = \{R, S, T\}$ still holds).

As mentioned, the rules of thumb proposed in this section will be supported with experimental results. All of the tests were run with one of two schemas depicted in Figure 5. The "Relations" column shows the attributes in each relation with the key attribute underlined. The next column ($T(V)$, using the notation in Table 3), gives the cardinality of each relation. The $I(V)$ and $D(V)$ columns give the number of insertions and deletions, respectively, as a percentage of $T(V)$. The updates were set to 0. The next two columns show the selection and join conditions using notation in Table 3. Schema 1 is a *linear join*, $V = R \bowtie S \bowtie \sigma_{10\%} T$, where both joins are on foreign keys. The relative cardinalities are $T(R) = 3 * T(S) = 9 * T(T)$. Schema 2 is also a linear foreign key join, $V = R \bowtie \sigma_{10\%} S \bowtie T$, but all of its relations have the same cardinalities. Both schemas have selection conditions with a selectivity of 10% but the conditions are placed on different relations.

5.2 When to Materialize Supporting Views

We give several rules of thumb governing which supporting views to materialize. The rules of thumb are based upon formulas estimating the benefit and cost of materializing a supporting view assuming that updates are done in batches, which is common in a data warehousing environment. We list the rules of thumb first, then analyze them using the formulas, and graph the results of our supporting experiments.

5.2.1 Benefits and Costs of Supporting Views

Rule 5.1 (Materialize Selective Supporting Views) *Materialize a supporting view V when $P(V) \ll P(\mathcal{E}(V))$.* ⊙

Rule 5.2 (Materialize Supporting Views Having No Deletions or Updates) *Materialize a supporting view V when $D(\mathcal{E}(V)) + U(\mathcal{R}(V)) = 0$.* ⊙

We assume in this section that a supporting view V does not *overlap* with any other materialized supporting view V' but we allow $R(V)$ to be a subset of $R(V')$. That is, for every other materialized supporting view V', either $\mathcal{R}(V) \cap \mathcal{R}(V') = \phi$, $\mathcal{R}(V) \subset \mathcal{R}(V')$, or $\mathcal{R}(V') \subset \mathcal{R}(V)$. The rule of thumb governing when to materialize overlapping supporting views is presented in [LQA96].

Rule 5.1 can support materializing a view even when Rule 5.2 doesn't hold. One instance where Rule 5.1 is likely to hold is when $V = S \bowtie \sigma_C T$, the join is from a foreign key in S to the key of T, and the selectivity of the selection condition C is low. These conditions imply $P(V) \ll P(\mathcal{E}(V))$.

In order to justify our rules of thumb and give a more detailed analysis of when to materialize a supporting view, we give approximate formulas for calculating the benefit and cost of materializing a supporting view (denoted as $\mathit{Benefit}_v$ and Cost_v). In general, a supporting view V should be materialized when $\mathit{Benefit}_v(V) > \mathit{Cost}_v(V)$. The formula for the benefit of a supporting view is:

$$\mathit{Benefit}_v(V) \approx \begin{cases} (|\mathcal{E}(V)| - 1) * I(\overline{\mathcal{R}(V)}) & \text{if } V \text{ is indexed on the appropriate join attributes and} \\ & \quad (|\mathcal{E}(V)| - 1) * I(\overline{\mathcal{R}(V)}) < P(\mathcal{E}(V)) - P(V) \\ P(\mathcal{E}(V)) - P(V) & \text{otherwise} \end{cases}$$

A materialized supporting view is beneficial because it contains pre-joined relations. The benefit of a supporting view V to a particular query (i.e., view maintenance expression) is therefore the difference between the cost of performing all of the joins in the query and the cost of just the joins between $\overline{\mathcal{R}(V)}$ and V. The costs of the joins depend on what type of join is used. If index joins are cheaper, then the benefit of materializing V is proportional to $|\mathcal{E}(V)| - 1$ (approximate number of index joins) times the number of insertions. Otherwise, if nested-block joins are used we assume that the insertions are always the smaller relation and that they will always fit in memory, so that the benefit is roughly the sum of the number of pages in the elements of V (the cost of joining with each of the elements of V) minus the number of pages in V (the cost of joining with V).

The cost of materializing a supporting view V is approximately the cost of propagating deletions and updates to $\mathcal{R}(V)$ onto V plus the cost of maintaining the indexes on V. The formula is:

$$\mathit{Cost}_v(V) \approx \begin{cases} D(\mathcal{R}(V)) + U(\mathcal{R}(V)) + \mathit{Cost}_i(V, *) & \text{if } V \text{ is indexed on the keys of base relations in} \\ & \mathcal{R}(V) \text{ and } D(\mathcal{R}(V)) + U(\mathcal{R}(V)) < P(V) * |\mathcal{R}(V)| \\ P(V) * |\mathcal{R}(V)| + \mathit{Cost}_i(V, *) & \text{otherwise} \end{cases}$$

In the formula, $\mathit{Cost}_i(V, *)$ denotes the cost of maintaining all of the indexes built on V. We defer giving a formula for the cost of materializing an index until Section 5.3. If V is indexed on the keys of the base relations and the cost of index joins is less that of nested-block joins, then the cost of maintaining V is proportional to the number of deletions and updates to $\mathcal{R}(V)$, since each deletion and update results in tuple lookups through the index. Otherwise, if nested-block joins are used, the cost is proportional to $P(V)$ times the number of base relations in V, since we have to scan V to find the tuples deleted or updated due to the changes to each base relation.

One might notice that we have not included the cost of propagating insertions in the formula above. To see why, consider a primary view $V^P = RST$ and a supporting view $V^S = ST$. The reason for ignoring the cost of propagating insertions is that the expression for propagating insertions to S onto V^S ($\triangle S \bowtie T$) is a subexpression of the expression for propagating insertions to S onto V^P ($R \bowtie \triangle S \bowtie T$). Therefore we can reuse the result and thus ignore the cost of propagating insertions to S onto V^S. A similar argument also holds for insertions to T. The only significant effect insertions have is in $\mathit{Cost}_i(V^S, *)$, as we will see in Section 5.3.

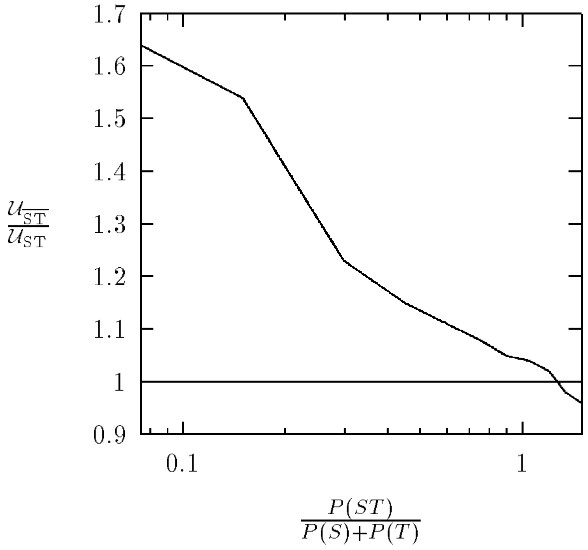

Figure 6: Support for Rule 5.1.

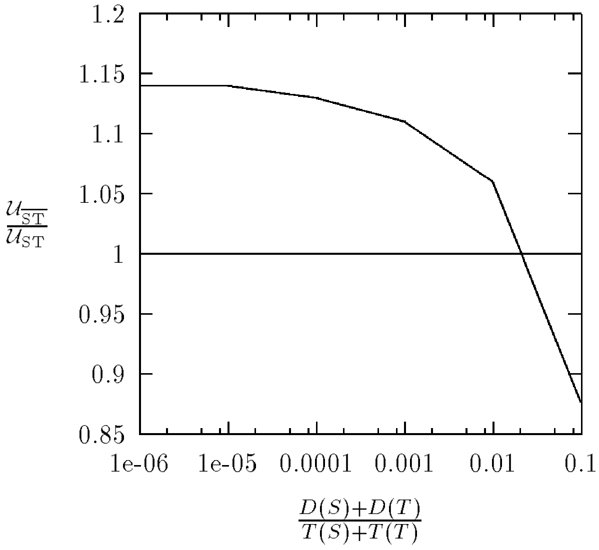

Figure 7: Support for Rule 5.2.

Rule 5.1 and Rule 5.2 are examples of rules that may work against or for each other. For instance if $D(\mathcal{E}(V)) + U(\mathcal{R}(V))$ is very high, it may not be beneficial to materialize V even if $P(V) \ll P(\mathcal{E}(V))$ holds. In the experiments, we set the parameters involved in one rule, while we varied the parameters in the other to show the effect.

Figure 6 shows the experimental support for Rule 5.1 (using Schema 1). In this experiment, we considered two view sets, one with the supporting view ST and one without. For both view sets, the whole index space was searched. The graph shows how the ratio of the total update cost without view ST over the total update cost with view ST materialized ($\mathcal{U}_{\overline{ST}}/\mathcal{U}_{ST}$) varies with $P(ST)/(P(S)+P(T))$. Therefore, it is beneficial to materialize ST when the line in the graph is above 1.0. As $P(ST)$ gets larger, it is less and less beneficial to materialize ST as predicted by Rule 5.1. Note that even when $P(ST) = P(S) + P(T)$, it is still beneficial to materialize ST for this schema because evaluating the maintenance expression $\triangle R \bowtie S \bowtie T$ is still more expensive than evaluating $\triangle R \bowtie ST$ (when ST is materialized). The reason is that there are more tuples in $\triangle R$ that match with S than with ST because of the selection condition on T.

In the next experiment (also on Schema 1), which supports Rule 5.2, the ratio $P(ST)/(P(S)+P(T))$ was set to 0.5. Figure 7 shows that as the deletion rates to S and T (as a fraction of $T(S)$ and $T(T)$) increase, it is less and less beneficial to materialize ST.

5.2.2 Things that Don't Matter

Rule 5.3 (Size Doesn't Matter) *In considering whether to materialize a supporting view V, the ratio of $P(V)$ to P_m doesn't matter.* $\odot$

The total number of pages in a supporting view relative to the number of pages of memory does not significantly impact the choice of whether to materialize the supporting view (unless of

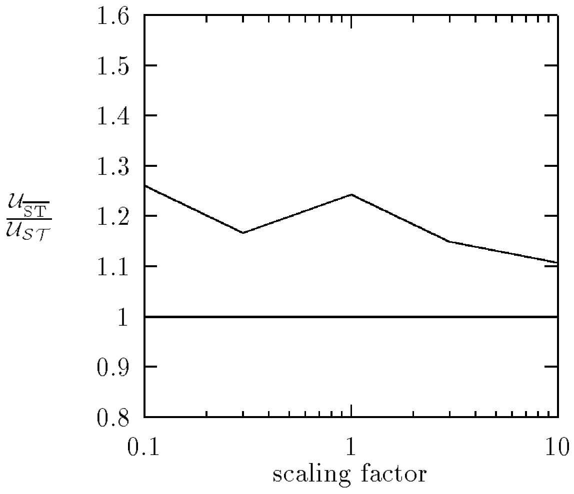

Figure 8: Support for Rule 5.3.

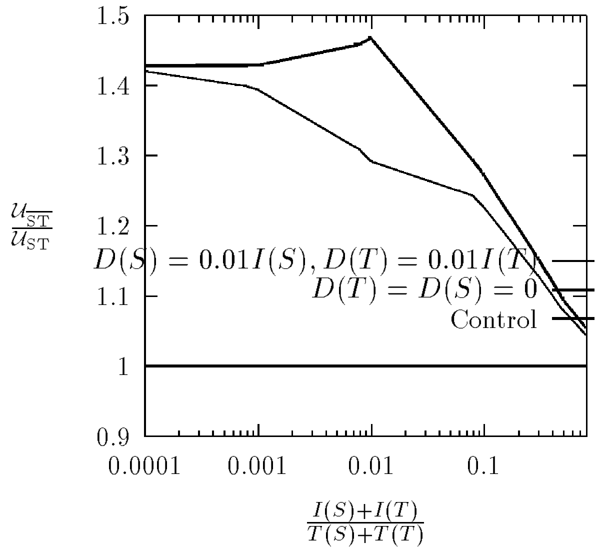

Figure 9: Support for Rule 5.4.

course the WHA is also trying to conserve space). Note that in the approximate formulas given for benefit and cost, P_m does not come into play. In our more detailed cost model (Appendix A), $P(V)$ relative to P_m has an impact primarily for index joins and index maintenance in which case small supporting views and indexes that fit entirely in memory have an advantage. But once a supporting view and its indexes grow beyond the size of memory then its size is not significant.

Figure 8 graphs the cost of maintaining two sets of supporting views for a primary view $V^P = RST$: one that includes a supporting view $V^S = ST$ and another where V^S is not materialized. We vary the sizes of all base relations as well as the number of changes to the base relations proportionately, while holding the number of pages of memory constant. Note that the size of V has little effect on the decision of whether to materialize it.

Rule 5.4 (Insertion Rate Doesn't Matter When No Deletions or Updates) *In considering whether to materialize a supporting view V, $I(\mathcal{R}(V))$ doesn't matter if $D(\mathcal{R}(V))+U(\mathcal{R}(V)) = 0$.* $\odot$

Let V be a supporting view that does not overlap with any other supporting view. The only time $I(\mathcal{R}(V))$ has a significant impact is in the cost of maintaining indexes on V. But without deletions or updates to $\mathcal{R}(V)$, if $I(\overline{\mathcal{R}(V)})$ is high enough that nested-block joins are performed between V and $\overline{\mathcal{R}(V)}$ instead of index joins, then indexes are not necessary on V and $I(\mathcal{R}(V))$ doesn't matter.

Figure 9 shows the results on an experiment performed on Schema 1. We considered two view sets, one with the supporting view ST and one without. We considered the ratio of the number of deletions and updates to the number of insertions to the two view sets as we varied the insertion rate. The deletion and update rates were set to 0 in one case, and 1/100th the insertion rate in another case. Figure 9 shows that insertion rate has little effect on the decision of whether to materialize ST if the deletion and update rates are set to 0. However, it does have an effect on whether to materialize ST if the deletion and update rates are non-zero.

5.3 When to Materialize Indexes

In this section we present rules of thumb for when to materialize indexes. The rules of thumb in this section generally apply to indexes on the primary view, supporting views, and base relations. Therefore, when we refer to a supporting view in this section we mean more generally a primary view, supporting view, or a base relation. Due to space constraints in the paper we are unable to give the results of our experimentation justifying the rules of thumb in this section. We limit ourselves to an analysis using our approximate formulas for benefit and cost.

Rule 5.5 (Build Indexes on Keys) *Build an index on a supporting view V for an attribute $R.A$ that is the key of base relation R involved in V if (1) $D(R)+U(R) > 0$, (2) $D(R)+U(R) \ll P(V)$, and (3) $I(\mathcal{R}(V)) + D(\mathcal{R}(V)) \ll P(V)$* $\odot$

Rule 5.6 (Build Indexes on Join Attributes—Sometimes) *Build an index on a supporting view V for an attribute $R.A$ that is involved in a join condition $R.A = S.B$ in the primary view when (1) $S \in \overline{\mathcal{R}(V)}$, (2) $I(\overline{\mathcal{R}(V)}) \ll P(V)$, and (3) $I(\mathcal{R}(V)) + D(\mathcal{R}(V)) \ll P(V)$* $\odot$

Rule 5.7 (Do Not Build Indexes on Local Selection Attributes) *Don't build an index on base relations R for an attribute $R.A$ involved in a selection condition C unless (1) indexes on R for attributes involved in join conditions have not been built, (2) a view $R' = \sigma_C R$ has not been materialized, (3) $S(R,C) \ll P(R)$, and (4) $I(R) + D(R) \ll P(R)$* $\odot$

Rule 5.8 (Build Indexes When the Index Fits In Memory) *Build an index on a supporting view V for an attribute $R.A$ if for any of the above Rules 5.5, 5.6, or 5.7, all but the final condition hold, and $P(V, R.A) < P_m$.* $\odot$

We now justify our rules by giving approximate formulas for the benefit and cost of materializing an index. Given an index on supporting view V for an attribute $R.A$, we give a single formula calculating cost but we give three formulas calculating benefit, one each for the following cases: (1) where the indexed attribute is a key of the underlying base relation ($Benefit_i^{key}(V, R.A)$), (2) where the indexed attribute is involved in a join condition ($Benefit_i^{jc}(V, R.A)$), and (3) where the indexed attribute is involved in a selection condition ($Benefit_i^{sc}(V, R.A)$).

In general, an index should be materialized on supporting view V for an attribute $R.A$ if $Benefit_i(V, R.A) > Cost_i(V, R.A)$ where $Benefit_i(V, R.A)$ is $Benefit_i^{key}(V, R.A)$, $Benefit_i^{jc}(V, R.A)$, or $Benefit_i^{sc}(V, R.A)$ depending upon how $R.A$ is used in the view. If $R.A$ is involved in multiple cases, such as being both a key of an underlying base relation as well as involved in a join condition, then the benefit of materializing an index for $R.A$ is the sum of the individual benefits.

Due to space limitations we cannot provide an analysis of the following approximate formulas. We simply note that rules 5.5, 5.6, and 5.7 governing when to materialize each type of index follow closely from the corresponding benefit formulas, and that the cost of maintaining an index is limited by the size of the index if the index fits entirely in memory, justifying Rule 5.8.

$$\mathit{Benefit}_i^{key}(V, R.A) \approx \begin{cases} (P(V) - D(R)) + (P(V) - U(R)) & \text{if } D(R) < P(V) \text{ and } U(R) < P(V) \\ (P(V) - D(R)) & \text{if } D(R) < P(V) \text{ and } U(R) \geq P(V) \\ (P(V) - U(R)) & \text{if } D(R) \geq P(V) \text{ and } U(R) < P(V) \\ 0 & \text{otherwise} \end{cases}$$

$$\mathit{Benefit}_i^{jc}(V, R.A) \approx \begin{cases} P(V) - S(V,C) * I(\overline{\mathcal{R}(V)}) & \text{if (1) join condition } C \text{ is } R.A = S.B, \\ & \text{(2) } S \in \overline{\mathcal{R}(V)}, \text{ and} \\ & \text{(3) } S(V,C) * I(\overline{\mathcal{R}(V)}) < P(V) \\ 0 & \text{otherwise} \end{cases}$$

$$\mathit{Benefit}_i^{sc}(V, R.A) \approx \begin{cases} P(V) - S(V,C) & \text{if (1) } V \equiv R \text{ (} V \text{ is a base relation),} \\ & \text{(2) } C \text{ is a selection condition on } R.A, \\ & \text{(3) indexes on } V \text{ for attributes involved in join conditions} \\ & \text{ have not been built,} \\ & \text{(4) a view } V' = \sigma_C V \text{ has not been materialized, and} \\ & \text{(5) } S(V,C) < P(V) \\ 0 & \text{otherwise} \end{cases}$$

$$\mathit{Cost}_i(V, R.A) \approx \begin{cases} P(V, R.A) & \text{if } P(V, R.A) < P_m \\ I(\mathcal{R}(V)) + D(\mathcal{R}(V)) & \text{otherwise} \end{cases}$$

6 Results

In the previous section, experimental results were used to validate the analysis supporting the rules of thumb. In this section, we present results that were borne directly from experimentation. In particular, we attempt to answer the following questions:

- Are views or indexes better when space is constrained?
- How sensitive is the optimal solution to the WHA's estimates of system parameters?

Due to space constraints, we only present the results of only one representative experiment for each question although many more were performed. In addition, in the full version of the paper [LQA96], we also consider whether protected updates should be treated atomically or split into pairs of insertions and deletions. The experiments shown in this section were all run on Schema 1 (described in Section 5.1). Although this schema is composed of only 3 relations, we believe our results to be more general because we have explored a number of larger schema with heuristic search algorithms and the results so far support those reported here.

6.1 Are Views or Indexes Better When Space is Constrained?

In this paper, we have shown how to find the optimal set of supporting views and indexes to materialize in order to minimize the total maintenance time. Sometimes, however, the amount of additional storage required is prohibitive. In these cases, one may ask how much storage is required to attain most of the performance gains and which structures should be materialized. We consider these questions for Schema 1 under two different update loads. In both experiments, we gradually increase the available storage from that required to materialize the primary view (RST) to that required by the optimal solution for the unconstrained problem. For generality, we measure the additional space as a fraction of the space required to store the base relations. At each point we find the best solution that fits in the available storage. The cost of this solution relative to the non-constrained optimum is plotted on the y-axis.

The results of the experiments are shown in Figure 10(a) and (b). As the graphs indicate, the schemas evolve in discrete steps - only changing when enough storage becomes available to add a new index or materialized view. The number of steps in the progression is too large (52 in (a) and 25 in (b)) to show every schema change but the results are summarized in Figure 11. The numbers next to features indicate in what order they are added as storage increases. Using Figure 11(a) as an example, the experiment starts with only the base relations and primary view materialized – they are numbered 0. The next features to be added are indexes on the keys of the base relations present in the view RST, starting with $T0$ and then adding $S0$ and $R0$. Next, the selection node T' is materialized and an index built on its attribute $T'0$. The reason that it takes 52 steps to add all 10 numbered feature sets is that a new feature is often added at the expense of an older one. For instance, when the view T' is materialized, the index on $R0$ in RST is dropped until enough space is available to add it again. The graphs in Figure 10 are also annotated with the feature numbers to help indicate which features most impact the update performance.

The first important point to note from this experiment is that under both update loads, a large portion of the total update savings can be achieved with a reasonably small amount of additional storage. Note the large drop in I/Os for the high-update experiment that results from materializing view T' (feature 3) and then adding indexes on $T0$ and $S0$ again (they were dropped earlier to make space for T'). The next large drop occurs after enough space is found to materialize ST (feature 5). By the time point A (which corresponds to features 1,2 and 5) is reached, the update cost is within 5% of the optimal cost. This is encouraging for warehouses that have space constraints. It should be noted that even though the extra storage required for the views and indexes does not seem that large compared to the warehouse relation sizes ($\approx 25\%$), there will typically be many views defined over the same relations so the total storage required by views and indexes can be larger than that of warehouse relation when the warehouse is considered in its entirety.

It is interesting to see how the two images of Figure 11 are supported by our rules of rules of thumb. Because RST is such a large relation, and there are deletions (but relatively few) to warehouse relations R, S, and T, by Rule 5.5, indexes should be built on RST for the keys of each of the warehouse relations. Also, because of the selection condition on T, the materialized view T' is much smaller than T. Therefore, by Rule 5.1 view T' should be materialized. Finally, note that view ST is not materialized until near the end. Even though the number of pages in ST is less than

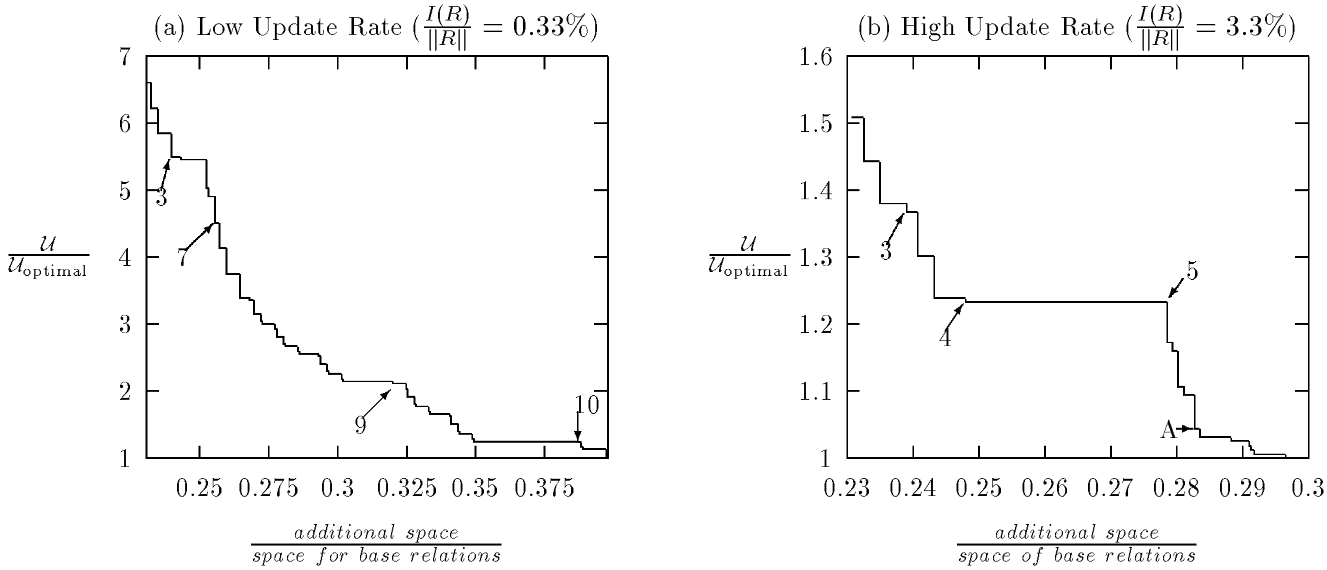

Figure 10: Effects of space on update cost.

the sum of the pages in S and T and should be materialized by Rule 5.1, ST is a relatively large structure to materialize in comparison to the indexes. Therefore, we find that the maintenance cost is minimized overall in this case by materializing several small beneficial structures (*i.e.*, indexes) than by materializing one large one (*i.e.*, view ST). It isn't until the most useful indexes have already been materialized that view ST is chosen for materialization.

6.2 Sensitivity Analysis

So far, this paper has focused on finding the optimal solution to the VIS problem. Just how well this solution works on the actual warehouse depends on how closely the input parameters, such as relation sizes and delta rates, match the real values of the system.[2] An important question for the WHA, then, is just how sensitive the optimizer is to the estimates of the input parameters. Clearly, one would hope that the optimal solution for the estimates is at least a good solution for systems with only slightly different parameters. In this section, we investigate just how badly optimal solutions decay at neighboring points. Due to space constraints, we consider only the estimate of insertion and deletion rates. (The reader is referred to [LQA96] for more experimental results.)

In this experiment, we varied the combined insertion and deletion rates to each base relation such that the ratio $\frac{I(R)+D(R)}{||R||} = \frac{I(S)+D(S)}{||S||} = \frac{I(T)+D(T)}{||T||}$ increased from 0.001 to 0.1 in five steps. At each step, we found the optimal solution and then plotted its performance over the entire range. The results, which are shown in Figure 12, suggest that except for a small region in the middle of the graph, the choice of optimum in not sensitive to the combined insertion-deletion rate. For instance, the optimal solution for an estimated ratio of 0.001 is still optimal even when the ratio grows to 0.01. The only area where the optimizer seems sensitive is in the range shown in the middle of the

[2]It also depends on how closely the VIS optimizer's cost model follows that of the dbms. This concept is discussed in [FST88].

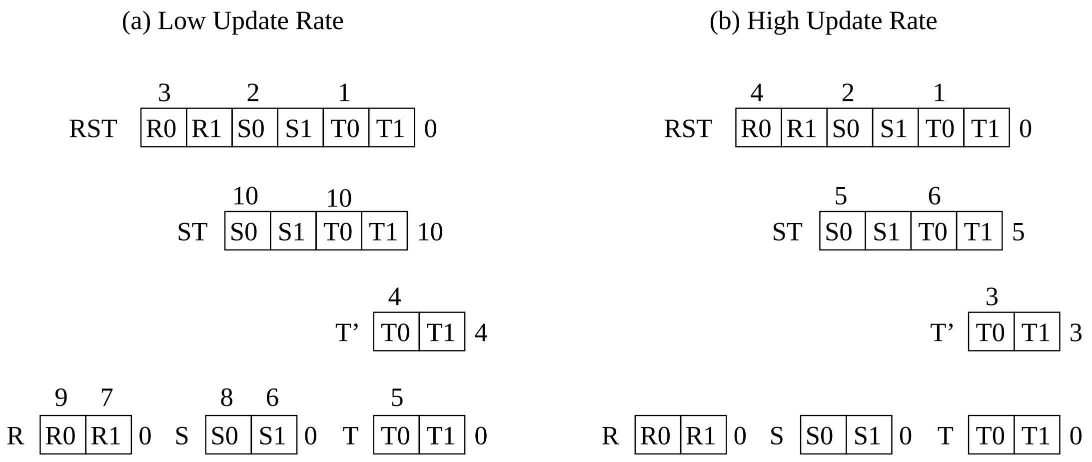

Figure 11: Evolution of the physical design.

graph where an order of magnitude error in estimation can lead to a three-fold performance hit or worse. This sensitive region corresponds to the point when the insertion-deletion rate to the base relations becomes large enough that it is no longer worthwhile to build indexes on their attributes.

This experiment is typical of many sensitivity analyses that we have performed. The optimal solutions perform well across a wide range of parameter values except for a few small regions that correspond to major schema changes. This is reassuring. One must be careful, however, in over-generalizing this result. It is likely that in schemas with more relations there will be more frequent shifts in the optimal schema. Whether these shifts will result in large differences in the maintenance cost is a subject for future research.

7 Related Work

Previous work related to this paper falls into two categories, depending on the context in which it was written: physical database design and rule condition maintenance.

7.1 Physical Database Design

Three costs must be balanced in physical database design for warehouses: (1) the cost of answering queries using warehouse relations and additional structures, (2) the cost of maintaining additional structures, and (3) the cost of secondary storage. We have assumed that the primary view is materialized, which minimizes the cost of (1), and focused on choosing supporting view and indices such that the cost of (2) is minimized. We have also considered how constraining cost (3) affects our results.

This problem was first studied by Roussopoloulos [Rou82]. The additional structures considered for materialization are view indices, rather than the views themselves, to save on storage. A view index is similar to a materialized view except that instead of storing the tuples in the view directly, each tuple in the view index consists of pointers to (or equivalently, tuple id's of) the tuples in the base relations that derive the view tuple. No other type of index are considered. (In our paper we

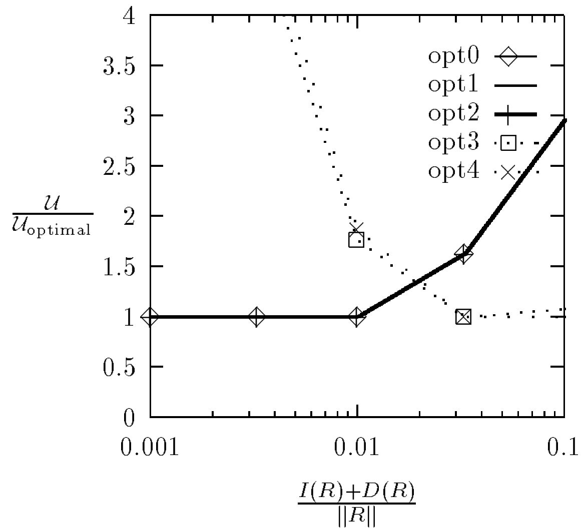

Figure 12: Sensitiviy of Optimal Solutions to Insert/Delete Rates.

choose to maintain the actual views since the cost of secondary storage is now much lower and no commercial database supports view indices.)

The Roussopoloulos paper presents an elegant algorithm based on A* and the approximate knapsack problem to find an optimal solution to the view selection problem. The algorithm, however, works because of two simplifying assumptions. First, it uses a very simple cost model for updating a view: the cost is proportional to the size of the view. But when views are incrementally maintained, the cost of maintenance is proportional not only to the size of the view but also to the sizes of the changes, the base relations, and subviews. We have shown in Section 2 that the cost of maintenance is a complex query optimization problem and cannot be estimated without knowing which subviews are materialized. Second, the Roussopoloulos algorithm does not consider index selection (other than view indices). We have shown in Section 6.1 that index selection has a significant impact on choosing which subviews to materialize, since the proper indexes can make a materialized subview less costly to maintain. Relaxing either of the above two assumptions invalidates the use of the Roussopoloulos algorithm. Still, this is a good first treatment of the subject and the author presents experimental results for the algorithm.

More recently, Ross et al. [RSS96] examines the same problem. They describe an exhaustive search algorithm to solve the view selection problem but without considering indexes. They also propose heuristics for pruning the space to search. We have extended their work by considering indexes, developing rules of thumb for choosing supporting views and indexes using cost model analysis, and presenting an improved optimal algorithm. We have also implemented our algorithm and used it generate experimental results that support the rules of thumb as well as answer questions such as whether to materialize indexes or views when space is constrained.

Other work has looked at the initial problem of choosing a set of primary views such that the cost of (1) is minimized, while ensuring that the costs of (2) and (3) are not too high. [SP89] considers this problem in the case of distributed views. [HRU96] has investigated this problem for the case of aggregate views. Tsatalos et al. [TSI94] consider materializing views in place of the base relations in order to improve query response time. Rozen et al. [RS91] look at this problem as adding a set of "features" to the database.

In particular, the index selection part of our VIS problem has been well-studied [FST88,CBC93] in the context of physical database design. Choosing indexes for materialized views is a straightforward extension. What is troublesome, however, is that the previous algorithms require the queries (and their frequencies) on each base relation as inputs. This information is used in pruning the search space of indexes to consider. In the VIS problem there are no user generated queries on the base relations or supporting views since they are all handled by the primary views: The only queries on base relations or supporting views are generated by maintenance expressions. Unfortunately, the set of generated queries depends on the update paths chosen for each type of delta. Recall from the example that if a view ST exists, the maintenance expression $\triangle R \bowtie S \bowtie T$ could be answered either from the base relations or as $\triangle R \bowtie ST$. The choice between the two update paths depends on whether there is an index on ST, which has not yet been determined. Thus one cannot determine the query set on each base relation and supporting view without knowing which indexes are present, which makes the algorithms proposed in previous work unusable here.

7.2 Rule Condition Maintenance

Previous work on active database and production systems also relates to the VIS problem we have described. Many authors have considered how to evaluate trigger conditions for rules. This can be considered a view maintenance problem where a rule is triggered whenever the view that satisfies its condition becomes non-empty. Wang and Hanson [WH92] study how the production system algorithms Rete [For82] and TREAT [Mir87] perform in a database environment. An extension to TREAT called A-TREAT is considered in [Han92]. Fabret et al. [FRS93] took an approach similar to ours by considering how to choose supporting views for the trigger condition view. Translated into our terminology, the rule of thumb they developed is essentially to materialize a supporting view when it is *self-maintainable*; *i.e.*, when it can be maintained for the expected changes to the base relations by referencing the changes and the view itself, but without referencing any base relations. For example, given a primary view $V^P = RST$ with only deletions (no insertions) expected to base relations S and T, then supporting view view $V^S = ST$ is a self-maintainable view. We have found through the results of our experimentation that for our environment almost the opposite is true. We have a rule of thumb that specifies to materialize a view when no deletions (insertions are fine) are expected to the base relations involved in the view.

Segev et al. [SF91, SZ91] consider a similar problem in expert systems. They also assume small deltas and ubiquitous indexes. They do not, however, consider maintaining subviews of the primary view, but instead describe *join pattern indexes*, which are specialized structures for maintaining materialized views. Join pattern indexes are an interesting approach, but require specialized algorithms to maintain. They cannot be maintained with SQL update statements, which is necessary for our approach because we want the WHA to be able to choose a set of supporting views and indexes and maintain them without having to write specialized code.

A major difference between all of these studies and this one is that they consider a rule environment where changes in the underlying data are proprogated immediately to the view. Hence, the size of the deltas sets are relatively small, which means that index joins will usually be much cheaper than nested-block joins. They therefore assume that indexes exist on all attributes involved in selection and join conditions. However, in the data warehousing environment studied here, a large number of changes are propagated at once, and the cost of maintaining the indexes often outweighs any benefit obtained by doing index joins, so it is not correct to assume that indexes exist on all attributes involved in selection and join conditions.

8 Conclusions

This paper considered the VIS problem, which is one aspect of choosing good physical designs for relational databases used as data warehouses. We described and implemented an optimal algorithm based on A* that vastly prunes the search space compared to previously proposed algorithms [RSS96]. Since even the A* algorithm is impractical for many real world problems, we developed rules of thumb for view and index selection. These rules were validated through both analysis and experimental results.

By running experiments with the optimal algorithm, we studied how space can be best used when it is constrained: whether for materializing indexes or supporting views. Our results indicate that building indices on key attributes in the primary view lead to solid maintenance cost savings with modest storage requirements.

In the future we plan to develop and compare a number of heuristics for pruning the exhaustive search space so that good solutions can be found through limited search.

Finally, we would like to remark that the VIS problem is critical in the eventual success of RDBMS as the supporting technology of data warehouses. Warehouses will be required to store increasingly more data and to answer ever more complicated queries. Tools to help WHAs devise physical designs that meet the customer's constraints on query response time, storage costs, and maintenance time will become a necessary component in an overall data warehousing strategy.

References

[CBC93] S. Choenni, H. Blanken, and T. Chang. On the selection of secondary indices in relational databases. *Data and Knowledge Engineering*, 11:207–33, 1993.

[For82] C. L. Forgy. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, 19:17–37, 1982.

[FRS93] F. Fabret, M. Regnier, and E. Simon. An adaptive algorithm for incremental evaluation of production rules in database. In *Proceedings of International Conference on Very Large Data Bases*, pages 455–66, 1993.

[FST88] S. Finkelstein, M. Schkolnick, and P. Tiberio. Physical database design for relational databases. *ACM Transactions on Database Systems*, 13(1):91–128, 1988.

[GL95] T. Griffin and L. Libkin. Incremental maintenance of views with duplicates. In M. Carey and D. Schneider, editors, *Proceedings of ACM SIGMOD 1995 International Conference on Management of Data*, pages 328–339, San Jose, CA, May 23-25 1995.

[GMS93] A. Gupta, I. Mumick, and V. Subrahmanian. Maintaining views incrementally. In *Proceedings of ACM SIGMOD 1993 International Conference on Management of Data*, Washington, DC, May 26-28 1993.

[Han92] E. Hanson. Rule condition testing and action execution in Ariel. In *Proceedings of 1992 ACM SIGMOD*, pages 49–58, 1992.

[HRU96] V. Harinarayan, A. Rajaraman, and J.D. Ullman. Implementing data cubes efficiently. In *Proceedings of 1996 ACM SIGMOD*, 1996.

[LMSS95] A. Levy, A. Mendelzon, Y. Sagiv, and D. Srivastava. Answering queries using views. In *Proceedings of the Fourteenth Symposium on Principles of Database Systems (PODS)*, pages 95–104, San Jose, CA, May 22-24 1995.

[LQA96] W. Labio, D. Quass, and B. Adelberg. Physical database design for data warehouses - the vis problem. Technical report, Stanford University, 1996. Available by anonymous ftp from db.stanford.edu in /pub/labio/1996.

[Mir87] D. P. Miranker. Treat: A better match algorithm for ai production systems. In *Proceedings of AAII 87 Conference on Artificial Intelligence*, pages 42–47, August 1987.

[ML89] L.F. Mackert and G.M. Lohman. Index scans using a finite lru buffer: A validated i/o model. *ACM Transactions on Database Systems*, 14(3):401–24, 1989.

[Nil71] N. Nilsson. *Problem solving methods in artificial intelligence*. McGraw-Hill, 1971.

[Rou82] N. Roussopoulos. View indexing in relational databases. *ACM Transactions on Database Systems*, 7(2):258–90, 1982.

[RS91] S. Rozen and D. Shasha. A framework for automating physical database design. In *Proceedings of International Conference on Very Large Data Bases*, pages 401–11, 1991.

[RSS96] K.A. Ross, D. Srivastava, and S. Sudarshan. Materialized view maintenance and integrity constraint checking: Trading space for time. In *Proceedings of 1996 ACM SIGMOD*, 1996.

[Sel88] T. Sellis. Multiple-query optimization. *ACM Transactions on Database Systems*, 13(1):23–52, 1988.

[SF91] A. Segev and W. Fang. Optimal update policies for distribtued materialized views. *Management Science*, 17(7):851–70, 1991.

[SP89] Arie Segev and Jooseok Park. Updating distributed materialized views. *IEEE Transactions on Knowledge and Data Engineering*, 1(2):173–184, June 1989.

[ST85] M. Schkolnick and P. Tiberio. Estimating the cost of updates in a relational database. *ACM Transactions on Database Systems*, 10(2):163–79, 1985.

[SZ91] A. Segev and J.L. Zhao. Data management for large rule systems. In *Proceedings of International Conference on Very Large Data Bases*, pages 297–307, 1991.

[TSI94] O.G. Tsatalos, M.H. Solomon, and Y.E. Ioannidis. The gmap: A versatile tool for physical data independence. In *Proceedings of International Conference on Very Large Data Bases*, pages 367–78, 1994.

[WH92] Y. Wang and E. Hanson. A performance comparison of the rete and treat algorithms for testing database rule conditions. In *Proceedings of International Conference on Very Large Data Bases*, pages 88–97, 1992.

A Cost Model

In this section we give our formulas for deriving the overall cost of maintaining a set of views due to changes to the warehouse relations. The formulas are based upon cost models for queries and updates [ST85] appearing elsewhere. The formulas represent a fairly accurate and detailed cost model, upon which we based our implementation of an algorithm that used exhaustive search to find the optimal set of supporting views and indexes for a given primary view. The results of experiments using this algorithm were used to justify our rules of thumb in Section 5 and our results in Section 6.

The main formula given in this section is $Cost_v(\mathcal{V})$, which is the cost of maintaining a set of views $\mathcal{V}$. The other formulas are used to support $Cost_v(\mathcal{V})$. We redefine the approximate formulas given in Section 5 for the cost of maintaining a view or an index in this section to use our more detailed cost model. Note that we do not give formulas for benefit, but one can derive $Benefit_v(V) = Cost_v(\mathcal{V}) - Cost_v(\mathcal{V} \cup \{V\})$.

Table 3 lists additional statistical functions that are used in the cost formulas in this section, In addition to the notation of Table 3, we need to define $H(V, R.A)$ as the height of an index on V for attribute $R.A$. Note that much of the statistical information for views can be derived from statistical information for the warehouse relations and the selectivities of selection and join conditions.

Table 4 gives our formula for $Cost_v(\mathcal{V})$ and its supporting formulas. Note that $Eval(expr)$ is the traditional query optimization cost function. In the formulas we use $\triangle R$, $\bigtriangledown R$, and μR to represent the set of insertions, deletions, and updates to R respectively. We have implemented an exhaustive-search query optimizer that calculates $Eval(expr)$ by considering all possible query plans. It uses as the cost estimates for each operator in the tree the formulas appearing in Table 5. The optimizer evaluates the cost of each possible query plan and selects the plan with the minimum cost. In addition, the optimizer considers possibly using materialized views in the evaluation of the expression, and considers reusing results of other expressions (which have been saved in $\triangle V_R^{save}$ relations).

Two more formulas need to be explained:

$$yao(n, p, k) = \begin{cases} k & k < p/2 \\ (k+p)/3 & p/2 < k <= 2p \\ p & 2p < k \end{cases}$$

The *yao* function returns an estimate of the number of page read operations given that k out of n tuples are read from a relation spanning p pages. The yao function assumes that either the memory buffer is large enough to hold the entire relation, or that the tuple accesses have been sorted beforehand so that tuples from the same page will be requested one after the other. Since the assumption that a relation fits entirely in memory is unrealistic for a data warehouse and we assume that tuple accesses are not usually sorted beforehand, our formulas often make use of a function Y_{WAP} presented in [ML89] for estimating the number of page read operations given k

tuple fetches and a memory buffer of m pages.

$$Y_{WAP}(n, p, k, m) \begin{cases} min(k, p) & p <= m \\ k & p > m \text{ and } k <= m \\ m + (k - m)(p - m)/p & p > m \text{ and } k > m \end{cases}$$

[[[What is the parameter order they give?]]] <==

Name	Formula	Description
$Cost_v(\mathcal{V})$	$\sum_{V \in \mathcal{V}} Cost_v(V)$	Derive cost to maintain a set of views by summing cost to maintain each view.
$Cost_v(V)$	$\sum_{R \in \mathcal{R}(V)}(Prop_{ins}(R,V)$ $+Prop_{del}(R,V) + Prop_{upd}(R,V))$	Sum the cost of propagate changes to each relation into V
$Prop_{ins}(R,V)$	$Eval(\triangle R \bowtie R_2 \bowtie \ldots \bowtie R_k \rightarrow \triangle V_R)$, $+Apply_{ins}(\triangle V_R, V)$ $+Apply_{ins}(\triangle V_R, \triangle V_R^{save})$ $+ApplyIx(\triangle V_R, V)$	Evaluate effect on V of $\triangle R$, which we call $\triangle V_R$, where $\{R, R_2, \ldots, R_k\} = \mathcal{R}(V)$ Insert $\triangle V_R$ into V Save it for possible reuse as $\triangle V_R^{save}$ (small cost anyway) Update indexes on V
$Prop_{del}(R,V)$	$Eval(V \ltimes_{\text{key of R}} \nabla R \rightarrow \nabla V_R)$ $+Apply_{delupd}(\nabla V_R, V)$ $+ApplyIx(\nabla V_R, V)$	Evaluate effect on V of ∇R, which we call ∇V_R Delete ∇V_R from V Update indexes on V
$Prop_{upd}(R,V)$	$Eval(V \ltimes_{\text{key of R}} \mu R \rightarrow \mu V_R)$ $+Apply_{delupd}(\mu V_R, V)$	Evaluate effect on V of ∇R, which we call μV_R Update μV_R in V
$Apply_{ins}(R,V)$	$P(R)$	Append tuples in R to V
$Apply_{delupd}(R,V)$	$yao(T(V), P(V), T(R))$	Delete or update tuples of R in V ($R \subseteq V$). Exact locations of tuples of R in V are derived when R is derived. If index join is used to derive R instead of nested-block join, then use $Y_{WAP}(T(V), P(V), T(R), P_m)$ instead of $yao(T(V), P(V), T(R))$.
$ApplyIx(R,V)$	$\sum_{R.A \in \text{indexes on } V}$ $(Y_{WAP}(T(V), P(V, R.A), T(R) * (H(V, R.A) - 1)$ $+Y_{WAP}(T(V), P(V, R.A), T(R)))$	For each index on V, sum approximate number of index pages to read assuming root cached, plus approximate number of index pages to write (leaves only).

Table 4: Cost Formulas

Operator	Formula	Description
$Eval$(Nested-block Join $E_1 \bowtie E_2$)	$Eval(E_1) + \lceil P(E_1)/P_m \rceil * Eval(E_2)$	Assume try to fit as much of left-hand expression result in memory as possible, then evaluate right-hand expression.
$Eval$(Index Join $E join V$)	$Eval(E)$	Cost of evaluating the left hand expression
	$+Y_{WAP}(T(V), P(V, S.B), T(E) * X, P_m/2)$	Let $X = H(V, S.B) - 2 + \lceil P(V, S.B) * S(V, JC)/T(V) \rceil$. Let the join condition JC be on indexed attribute $S.B$ in V, then Y_{WAP} is the number of index pages to read, assuming buffer memory is split between index and relation
	$+Y_{WAP}(T(V), P(V), T(E) * S(V, JC), P_m/2)$	Number of relation pages to read
$Eval$(Relation Scan V)	$P(V)$	
$Eval$(Index Scan V)	$H(V, S.B) - 1 + \lceil P(V, S.B) * \frac{S(V,SC)}{T(V)} \rceil$	Let the selection condition SC be on indexed attribute $S.B$ of V, then this line computes the number of index pages to read.
	$+Y_{WAP}(T(V), P(V), S(V, SC), P_m)$	Number of relation page to read

Table 5: Query-Optimizer Cost Formulas