

Copyright © 1986, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

AN ANALYSIS OF RULE INDEXING IMPLEMENTATIONS
IN DATA BASE SYSTEMS

by

M. Stonebraker, T. Sellis and E. Hanson

Memorandum No. UCB/ERL M86/6

16 January 1986

Cover 3018

AN ANALYSIS OF RULE INDEXING IMPLEMENTATIONS
IN DATA BASE SYSTEMS

by

M. Stonebraker, T. Sellis and E. Hanson

Memorandum No. UCB/ERL M86/6

16 January 1986

ELECTRONICS RESEARCH LABORATORY
College of Engineering
University of California, Berkeley
94720

File 100

AN ANALYSIS OF RULE INDEXING IMPLEMENTATIONS IN DATA BASE SYSTEMS

M. Stonebraker, T. Sellis and E. Hanson

*Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA 94720*

Abstract

In this paper we discuss several alternate implementation schemes for rule indexing in a data base system. Two of the proposals have much in common with predicate locking, while two others resemble versions of physical locking. A performance analysis is conducted based on an abstract model of the rule indexing problem.

1. INTRODUCTION

There has been considerable discussion of the relative merits of predicate locking [ESWA76] and physical locking [GRAY78] to support concurrency control in relational data base systems. Physical locking has been implemented in most commercial relational systems with which we are familiar, and appears to be the tool of choice.

In this paper we show that similar considerations arise when a relational DBMS is extended to support rule processing for expert system applications. In this new context, tactics similar to both predicate locking and physical locking can be applied, and it is necessary to re-examine the best choice in light of the changed circumstances.

In Section 2 we indicate three different environments in which rule processing must be accomplished. Each will be seen to require very similar functionality. Then we turn in Section 3 to a discussion of ways to support these environments by using tactics similar to physical locking and predicate locking. Section 4 indicates an abstract model with which we can compare the performance of the different approaches, and an analysis of the cost of each scheme based upon the model. Section 5 then discusses predicted performance of the various implementations in a collection of different situations. Then, in Section 6 we discuss a collection of more sophisticated algorithms which may offer superior performance compared to the ones analyzed, and in Section 7 we draw some conclusions.

2. RULE PROCESSING ENVIRONMENTS

2.1. Triggers

The first environment of interest is support for active data bases which include triggers. One possible syntax was presented in a previous paper [STON85a], and others have been designed (e.g. [ESWA75]). We will use the following schema for the standard EMP relation for examples in this paper:

relation: EMP(name, age, salary, status)
storage structure B⁺-tree indexes on name, salary

A trigger to set Mike's salary equal to Sam's salary whenever Sam is given a raise would be expressed as follows:

```
range of E is EMP
replace ALWAYS EMP (salary = E.salary)
where EMP.name = "Mike"
and E.name = "Sam"
```

The required semantics are that this command should logically appear to run indefinitely. More realistically, whenever a command such as

```
replace EMP (salary = 1000) where EMP.name = "Sam"
```

is processed the trigger should be awakened to update Mike's salary.

More precisely, the system must store a (perhaps very large) collection of triggers:

```
T1: replace ALWAYS relname-1 (Target-list-1) where PREDICATE-1
```

.

```
Tn: replace ALWAYS relname-n (Target-list-n) where PREDICATE-n
```

When a user update (i.e. a replace or an append) is processed, e.g.

```
update-command relname (Target-list) where QUAL
```

the system must find all triggers T_i for which there exists a tuple *t* modified or inserted by the update such that:

```
t satisfies PREDICATE-i
and
t satisfies QUAL
and
```

Target-list contains an attribute which appears in Target-list-i or PREDICATE-i

Notice, that the only penalty for finding "false drops" (i.e. T_i for which the above condition is not true) is that time is wasted processing a trigger which does not actually do anything.

2.2. Inference

One possible approach to inferring data that is not present in the data base was also presented in [STON85a]. In that proposal columns of a relation are either real or virtual. Real fields are filled with normal data while virtual fields are inferred from rules associated with the relation. One possible rule would be:

```
replace DEMAND EMP (salary = E.salary)
where EMP.name = "Mike"
and E.name = "Sam"
```

This rule states that the salary of Mike should be the same as the salary of Sam. A second rule might require Sam's salary to be 1000 dollars as follows:

```
replace DEMAND EMP (salary = 1000)
where EMP.name = "Sam"
```

One wishes an environment whereby a user can ask queries of the form:

```
retrieve (EMP.salary) where EMP.name = "Mike"
```

and have the system infer the correct answer of 1000 from the above rules.

The virtues of this approach relative to utilizing the view mechanism to perform inference (e.g. [ULLM85, IOAN84]) have been presented in [STON85a]. In brief, the view mechanism is appropriate when a small collection of rules is present and most of them are applicable. The canonical example is the construction of the ANCESTOR view from a base PARENT relation. On the other hand, if there are many rules which define a particular view and most are not relevant to a given query, then our proposal excels. Moreover, our proposal supports conflicting rules. For example, a third rule might state that all managers must be paid 2000 dollars, e.g:

replace DEMAND EMP (salary = 2000) where EMP.status = "mgr"

If Mike is promoted from a worker to a manager, then his salary should be changed from 1000 to 2000. This requires a mechanism whereby multiple rules can apply in a given situation and one is designated at higher priority to be used in preference to the others. Such priority situations are common in AI applications and are straightforwardly supported in the scheme of [STON85a] but are much more difficult in rule systems which utilize the view processing system.

This environment requires storing a (perhaps very large) collection of DEMAND commands:

D₁: replace DEMAND (Target-list-1) where PREDICATE-1

D_n: replace DEMAND (Target-list-n) where PREDICATE-n

Then, the system must process a user command:

retrleve (Target-list) where QUAL

by finding all D_i for which there exists a tuple *t* such that:

t satisfies PREDICATE-i

and

t satisfied QUAL

and

Target-list intersect Target-list-i NON-EMPTY

For qualifying D_i, one must run an algorithm similar to query modification [STON75] to convert the original user retrieval into a new one which is then processed. The details of this algorithm appear in [STON85a]. Again, there is no semantic problem with "false drops"; they simply generate extra overhead.

2.3. Precomputed Answers to Commands

A third environment with similar requirements is one that allows queries in the query language to be data base objects [STON84]. For example, one could declare the salary field of the EMP relation to be a query language command. The value for this field for the employee named Mike would be:

range of E is EMP

retrleve (salary = E.salary) where E.name = "Sam"

Of course, Sam's salary would be the query returning a constant:

retrleve (salary = 1000)

The difference between this model and the previous one is that queries are bound to specific data tuples, and the value of the field is found by executing a query. For example, if all the employees in the shoe department share the same salary, then the query specification must be repeated for each one. On the other hand, the previous environment did not bind a query to a specific tuple, and a single DEMAND command could apply to all employees in the shoe department.

In this situation a desirable optimization strategy is to optionally cache the answer to stored queries [STON85b]. The next time one accesses the object, the data is already precomputed and

need not be rematerialized. However, the system must invalidate the precomputed object (e.g. the salary of Mike) if a subobject from which it is composed (e.g. the salary of Sam) is updated.

More precisely, this environment contains a collection of queries that have been precomputed, e.g:

Q_1 : retrieve (Target-list-1) where PREDICATE-1

Q_n : retrieve (Target-list-n) where PREDICATE-n

If an update command (a replace or append) is processed, e.g:

update-command (Target-list) where QUAL

the system must ascertain which retrieve commands must be invalidated. The required test is to find all Q_i for which there exists a tuple, t such that:

t satisfies PREDICATE- i

and

t satisfies QUAL

and

Target-list contains an attribute which appears in Target-list- i or PREDICATE- i

Notice that this test is identical to the trigger test.

3. SUPPORT FOR RULE PROCESSING

The first option for supporting rule processing would be to construct a theorem prover which would find all PREDICATE- i for which:

(QUAL AND PREDICATE- i) NOT EMPTY

Analysis of such a theorem prover is beyond the scope of this paper, and we restrict our attention to more conventional tactics. Also inference requires tuple-by-tuple processing on retrieves and bulk processing is not applicable. Hence, in the remainder of this paper we assume that tuples satisfying QUAL are determined in the course of normal query processing. For each such tuple, t the task is then the following:

Find all rules, D_i , T_i or Q_i depending on the environment, such that t satisfies PREDICATE- i .

Also, depending on the type of rules being used, there may be additional checking concerning presence of attributes.

The first approach we consider is to view each PREDICATE- i as setting a predicate lock in the data base. When the tuple t is accessed, one must ascertain which predicate locks cover it. When the number of predicates is large, it is best to construct a predicate index to avoid a sequential search of all predicates. Given such an index, to find the predicates covering a tuple, one can perform an index lookup. In Section 3.1 we present a predicate indexing scheme for rule processing that has points in common with predicate locking proposals.

Alternatively, one can view each predicate as setting physical locks. Since each PREDICATE- i is a valid query qualification, a database command corresponding to it can be executed by the query processing engine. When this command executes, a special marker can be set on each accessed tuple instead of a conventional read or write lock. Such a marker, called a "trigger-me lock" or "t-lock," indicates that PREDICATE- i *might* cover the tuple. These locks can be consulted when the collection of rules that cover a tuple, t , is needed. We present an algorithm based on t-locks in Section 3.2.

Then, in Sections 3.3 and 3.4 we discuss two other variants of the basic predicate locking and physical locking techniques. These alternatives may prove attractive in certain situations.

3.1. Predicate Indexing

The goal of this scheme is to build a data structure that will allow the rule base to be efficiently searched to determine the PREDICATE- i that cover a specific tuple. In this section we assume that all PREDICATE- i have the following characteristics:

- 1) each PREDICATE- i restricts a single relation and contains a single tuple variable
- 2) each PREDICATE- i is a conjunction of terms of the form:

$$\text{constant-1} \leq \text{attribute-}i \leq \text{constant-2}$$

for $1 \leq i \leq F$, where F is the number of attributes for the relation. Constant-1 and constant-2 may be scalars or positive or negative infinity.

Notice, that this formulation does not allow the presence of join clauses in predicates. Moreover, the indexing proposal in this section cannot easily be extended to capture join predicates.

For each relation that has rules defined, we propose to build a special kind of R-tree [GUTT84a] on the total of F attributes that may have a restrictive term in a predicate. We select R-trees instead of some other multidimensional structure like k-D-B-trees [ROBI81] because rules do not represent point data in a multidimensional space, but rather rectangular regions. For example,

$$20 < \text{EMP.age} < 30 \text{ and } 10K < \text{EMP.salary} < 30K$$

defines a rectangle in the two dimensional space with dimensions EMP.age and EMP.salary.

An R-tree is a tree structure used to index rectangles in a geometric environment [GUTT84b]. It is clear that the predicates indicated above can be considered rectangles in the multidimensional space formed by considering each of the F attributes as a dimension. An intermediate node in an R-tree is a sequence of pairs ($RECT, PTR$), where $RECT$ is the description of a rectangle and PTR a pointer. The pointer PTR is used to point to another node all of whose rectangles are completely contained in $RECT$. The leaf nodes have the same format, except that there are no pointers.

One major difference between R-trees and B-trees is that the rectangles, $RECT$ found in intermediate nodes of an R-tree are not disjoint. Therefore, when descending the tree searching for the predicates which cover a specific tuple, t , one may investigate more than one path in the R-tree. When the indexed rectangles are large, as may be the case in this application, considerable overlap of bounding rectangles may be observed, and many paths may require investigation.

In this application, an alternate variation of R-trees may prove attractive. Any rule predicate can be decomposed into a collection of non-overlapping sub-rectangles whose union is the original predicate. These sub-rectangles can be judiciously chosen so that no bounding rectangle in any index level of the R-tree need be enlarged. Moreover, when a leaf-level page overflows and the page must be split, the bounding rectangle can usually be partitioned into two non-overlapping rectangles and then the sub-rectangles on the page which intersect the new boundary can be split. However, in the worst case where there are n rules that all mutually intersect, a sub-rectangle of each rule will occur on some leaf page. If the capacity of this page is exceeded, a split will not lower its occupancy and the leaf page must be extended with overflow pages.

In this case, only a single path in the tree must be investigated to find the predicates covering t . The cost is a perhaps much larger number of rules to index. We will call the new tree type the " R^+ -tree." A detailed investigation of its characteristics is presented in a separate paper [SELL86]. In the performance comparison that follows in Section 4, we analyze only R^+ -trees because we can only find closed-form expressions for expected costs in this case. Which variation actually performs best depends on the composition of the predicates.

It should be clear that one must descend an R^+ -tree from the root to one leaf node every time one wishes to find the predicates that cover a specific tuple. However, one need do no special maintenance of the tree when a tuple is inserted into the database or modified.

3.2. Basic Locking

In this scheme, a relation

RULES(id, rule-def)

will contain the rule base. The *id* field contains a unique identifier for the rule, and *rule-def* contains the definition of the rule, including its predicate.

For each rule defined, an access plan is constructed by the query optimizer. This plan is executed and each tuple it reads is marked with a t-lock containing an identifier for the current predicate. If a sequential scan of the relation is used, then all tuples in the relation will be marked. In this case we assume that conventional lock escalation will convert record locks to a relation lock. Otherwise, an index will be used for access and t-locks will be set on data records *and* on the key interval inspected in the index. Such index interval locks are required to correctly deal with insertion of new records, as explained momentarily. The exact form of the index locks may be specific to the type of index, and our analysis in Section 4 assumes that t-locks can be set between keys on the leaf level of the index. To simplify the analysis we also require that all predicates contain at least one indexed attribute, thus avoiding the use of any relation-level locks.

If a tuple t is inserted, then the collection of markers must be found for the new tuple. As a side effect of the insertion, values will be inserted into all indexes on the relation. If such a value is covered by a key-range lock, then a corresponding t-lock will be added to the data tuple containing the value.

To ascertain what collection of PREDICATE- i cover a tuple t , one first collects all the t-locks on t . Since these t-locks represent a superset of the predicates that actually match the tuple, relevant tuples in the RULES relation must be checked to determine whether t actually satisfies each one.

For example, the qualification:

EMP.salary = 1000 and EMP.age > 30

will set t-locks in the salary index and on all data records that it reads (i.e. those with salary = 1000). Not all of these will have qualifying ages. The reason such a cautious strategy must be adopted is that a non-indexed attribute may be modified so that a record matches a predicate it did not match before the change. For example, an employee may be aged from 30 to 31. Since there is no secondary index on age, the basic algorithm would have no way of discovering that it should now be marked, barring searching the salary index, a cost we wish to avoid if only age is updated. Because of this problem, t-locks must be set on all tuples that *potentially* satisfy a predicate based on the interval locks the predicate has set in the indexes.

This strategy is called basic locking because it sets t-locks on all objects for which a normal query would set read or write locks. Moreover, it requires no changes to conventional execution of access plans, so it can be properly called a locking mechanism. The advantage of this scheme is that it is closely coupled to normal query processing. New qualifications can be added using normal facilities, and locks for new tuples are found as byproducts of normal update processing.

In the next two sections we describe variations on basic predicate locking and basic physical locking.

3.3. Early Basic Locking

Notice that the above algorithm defers checking whether a given tuple actually satisfies a predicate until the tuple is accessed and the matching predicates are required. The basic locking algorithm does only a modest amount of work at the time a tuple is inserted or modified, leaving

the bulk of the overhead to the time the tuple is accessed. Hence, the algorithm can be classified as a *late* algorithm, in that it defers overhead whenever possible. If tuples are updated frequently, late algorithms should perform well. On the other hand, if a tuple is accessed often, the following *early* locking algorithm should prove advantageous.

In early basic locking, a collection of t-locks is constructed and saved on the tuple at the time it is inserted or modified just as in basic locking. However, an extra step is also performed, namely each corresponding predicate from the RULES relation is retrieved and checked against the tuple. Special *match markers* are stored on the tuple for all predicates that the tuple satisfies. In exchange for extra overhead at update time, this algorithm can deliver a list of qualifying rule-ids with no extra overhead at the time of access.

3.4. Early Predicate Indexing

An analogous early version of predicate indexing can also be constructed. In this situation, ascertaining which predicates cover a given tuple is not deferred to access time; rather it is done at insertion time. Hence, markers for the predicates that actually cover the tuple are stored on the tuple as in early basic locking. However, these markers are discovered by searching the predicate index rather than by consulting index interval locks in data indexes. Again, this scheme should be advantageous in "read mostly" environments.

4. PERFORMANCE ANALYSIS

The schemes discussed so far have dissimilar characteristics. Late predicate indexing requires all requests for the set of covering predicates to go through the index while early basic locking directly delivers the collection that qualify. When tuples are updated or inserted, there is no extra overhead paid by the late schemes while all early methods incur a substantial penalty. This section defines an abstract model of relevant operations so that performance of the various algorithms can be analyzed.

4.1. The Abstract Model

We assume that there are two relevant operations in the model. The first, called *update*, is to update a single field in a single tuple in a relation R which has F attributes. This single field is chosen at random from among the F candidates. The second, *access*, is to find all the predicates that cover a given tuple already existing in the data base. One reason for this model is that most database update operations can be modeled by single-attribute **replace** commands, i.e. update operations.

Moreover, the three applications presented in Section 2 are composites of these two operations. For example, the trigger environment requires tuples to be identified, the proposed data modification to be tentatively performed, and all predicates that cover these tuples to be found. This can be considered as a collection of pairs of basic operations in our model, each containing an update followed by an access. The other two example applications are similarly modeled.

The mix between updates and accesses is controlled by a parameter, P , which is the percentage of updates. The other parameters used in the analysis are presented in the following table.

Parameter	Description
C_1	The cost of evaluating a predicate for a given tuple
C_2	The cost of reading a page
B	The size of the page in bytes
W	The width of data records in the relation
N	The number of tuples in the relation
F	The number of fields in the relation
F_I	The number of fields in the relation with an index
S	The width of individual fields in the relation
4	The assumed width of pointers and t-locks
t	The number of rules
Q	The fraction of records matching a single term of a predicate (will vary with the model)
P	The fraction of update commands

The Model Parameters
Table 1

Moreover, we will analyze each of the following four models for the rule predicates:

Model 1: All predicates have a single clause restricting a single field z , i.e. they are of the form
 $\text{relation}.z = \text{value}$

Model 2: All rule predicates are of the form
 $\text{lvalue} \leq \text{relation}.z \leq \text{uvalue}$

Model 3: Each predicate has an equality restriction clause on all of the F attributes in the relation.
 $\text{relation}.z_1 = \text{value}_1 \text{ and } \dots \text{ and } \text{relation}.z_F = \text{value}_F$

Model 4: This model is the same as model 3, except that the individual clauses are range restrictions rather than exact-match terms.
 $\text{lvalue}_1 \leq \text{relation}.z_1 \leq \text{uvalue}_1 \text{ and } \dots \text{ lvalue}_F \leq \text{relation}.z_F \leq \text{uvalue}_F$

We now turn to the expected cost per operation for late predicate indexing and late basic locking for each of the four predicate models. The expected costs of the early versions of these implementations are straightforward to derive, and we omit them in the interest of brevity. In all cases, we assume that the general algorithm is utilized, and that special case features appropriate to a particular predicate class cannot be exploited. This corresponds to a realistic implementation since the composition of the rules is generally not known in advance. Also, we only count processing costs in excess of those required by any system to perform the retrieval and update operations.

4.2. Late Predicate Indexing

The predicate index must be built on all F attributes of the relation. Assuming 4 bytes for each pointer and W to be the width of a tuple, then each node can hold

$$p = \left\lfloor \frac{B}{2W+4} \right\rfloor$$

predicates. We will also assume for simplicity that all nodes are full. (Otherwise one can use some constant factor such as the one derived by Yao [YAO78]). In R^+ -trees the rules may be broken into more than one entry in models 2 and 4, and a reasonable approximation of the expected number of smaller pieces O that a rule must be broken into is

$$O = \begin{cases} 1 & \text{for models 1 and 3} \\ 2 & \text{for model 2} \\ \frac{F}{2^2} & \text{for model 4} \end{cases}$$

Given the above assumptions, the number of leaf pages, L in the index is:

$$L = \lceil \frac{t \cdot O}{p} \rceil$$

and the depth of the index is:

$$d = \lceil \log_p L \rceil$$

In this case the cost to find all predicates that overlap a specific tuple is, for models 1 and 3

$$cost = (d+1)(C_2 + C_1 \cdot p)$$

while for models 2 and 4

$$cost = (d+1)(C_2 + C_1 \cdot p) + (t \cdot Q - p)C_1 + \lceil \frac{t \cdot Q - p}{p} \rceil C_2 \quad \text{if } t \cdot Q > p$$

$$cost = (d+1)(C_2 + C_1 \cdot p) \quad \text{if } t \cdot Q \leq p$$

and the total cost per operation, TOTAL in the abstract model is:

$$TOTAL = (1-P)cost$$

4.3. Late Basic Locking

It is clear that an insert incurs zero extra overhead for this algorithm. The only cost occurs in finding covering predicates. The predicates corresponding to all t -locks must be accessed (at cost C_1 each) and then checked (at cost C_2) to find the ones that actually cover the tuple.

Therefore, the total cost per operation in the abstract model is:


$$TOTAL = (1-P)t \cdot Q \cdot (C_1 + C_2)$$

5. PERFORMANCE RESULTS

In order to compare the four implementations we set the following parameters to constants as indicated:


$$\begin{aligned} C_1 &= 10 \\ C_2 &= 30 \\ B &= 2,000 \\ W &= 100 \\ N &= 1,000,000 \\ F &= 10 \\ S &= 10 \\ F_I &= 3 \end{aligned}$$

One can interpret C_1 and C_2 as times in msec; the other parameters are typical of current applications. It can be noted that models 1 and 3 will yield the same expected total cost per basic operation for any particular setting of the model parameters. Hence, they differ only in what values of Q are intuitively reasonable. In all cases, predicate indexing performance is sensitive to the total number of rules, t , and we set that to 10,000. Figure 1 plots expected cost per operation, TOTAL, for models 1 and 3 for $t \cdot Q = 1$ as P is varied from 0 to 1. The labels LPI, LBL, EPI and EBL in Figure 1 stand for late predicate indexing, late basic locking, early predicate indexing, and early basic locking respectively. The other figures are similarly labeled. For



Models 1 and 3 : Costs vs Update Probability with $t.Q = 1$.
Figure 1

models 2 and 4 we set $t.Q$ to be 15 to simulate each tuple being covered by 15 rules and performed the same analysis as above. The results for model 4 appear in Figure 2; those for model 2 were similar, and are omitted for brevity.




Model 4 : Costs vs Update Probability with $t.Q = 15$.
Figure 2

In our analysis, the cost of EBL is the same as that for BL except that $(1-P)$ is replaced by P . The same is true for EPI and LPI. This yields a performance crossover point between the early and late versions of the respective algorithms at $P = .5$. In reality, the early algorithms have slightly more overhead (to save match markers on tuples) than the late equivalents. Thus, the actual crossover point would likely occur when P was slightly less than .5 in an empirical study.

Notice that the variants of basic locking, LBL and EBL, dominate the predicate indexing schemes in both situations. However, the two versions of predicate indexing will offer superior performance when a tuple is covered by a sufficiently large number of rules. Basic locking must access and check a predicate for each t -lock in a tuple. This requires one random disk access per t -lock, whereas the qualifying predicates are clustered in the predicate index. This absence of clustering dooms basic locking to poor performance when the number of t locks becomes sufficiently large. On the other hand, basic locking excels when the expected number of rules which cover a tuple is low.


To ascertain when predicate indexing becomes attractive, we varied the expected number of rules that cover a tuple, $t \cdot Q$, for a constant update probability, P , of 0.15. The results for models 1 and 3 appear in Figure 3; those for models 2 and 4 were similar, and are thus not shown. Notice that predicate indexing becomes the preferred option when 17 rules cover each tuple. Figure 4 summarizes our results by giving the areas in the $(t \cdot Q, P)$ space where each of the implementations seems to work best.

In the case of triggers and invalidation of materialized objects, the update and access operations occur in pairs, and a value of $P=.5$ is thus expected. In a system using only these kinds of rules, a late scheme is probably preferred since it is slightly easier to implement. In an inference environment, the access operation occurs at tuple read time, so P is exactly the fraction of data base operations that are writes. In many environments, this fraction is substantially less than 0.5, so an early algorithm would normally be preferred.



Models 1, 2 and 3 : Costs vs Expected Number of Predicates
Covering a Single Tuple ($P = 0.15$)

Figure 3



Models 1 and 3: Regions of Best Performance for
Different Values of P and $t.Q$.

Figure 4

6. OTHER APPROACHES

It is clear from the preceding section that the best choice varies with the expected environment. Hence, it is feasible, although not very attractive, to implement more than one of the options and then choose an implementation based on the following:

1. the expected number of rules that will cover a given tuple
2. the update probability.

Moreover, there are considerations concerning the generality of predicates required. If one needs predicates that include joins, then the predicate indexing schemes do not work, and a locking scheme must be employed. One possible composite scheme that overcomes this deficiency is now presented. A multi-relation predicate can be decomposed into a query plan that includes selections, projections, and a specific join algorithm run on pairs of relations (e.g. [WONG76, SELI79]). One could insert any selection or projection subquery in the plan into the appropriate rule processing system (e.g. late basic locking). Then one could physically mark any pairs of tuples that satisfied the join clauses. If any subquery t-locks were broken during normal query processing, processing could then be done to maintain the predicate index and determine if any new pairs of tuples matched under the join predicate. It is possible that such a scheme may be advantageous. Other schemes to extend predicate indexing to join predicates are a subject for future research.

Also, if one restricts the predicates to model 1, then the predicate indexing scheme degenerates to a conventional B-tree which can be integrated with the secondary indexing mechanism. Hence, the implementation difficulty is eased in this special case.

Another improvement to basic locking is to organize the collection of rules so that the rules are clustered in a way similar to the predicate index described above. This will improve the performance of basic locking in environments with many rules covering the average tuple because it will cluster rules that must be checked together onto a smaller number of disk pages.

In addition, basic locking can be improved if additional complexity is tolerable. We will briefly describe another more sophisticated marking scheme we call "mark intersection" that may offer superior performance. The primary motivation behind the mark intersection algorithm is to reduce the number of unnecessary searches into the RULES relation (false drops) by carefully analyzing the t-locks on a tuple.

Consider the following general form of a predicate:

$P: p_1 \text{ and } \dots \text{ and } p_k$

Here, p_j is a restriction term such as " $R.A = 5$ " or " $50 \leq R.A \leq 100$ ", and $1 \leq k \leq F$. The attributes for which a predicate has a restriction term will be denoted by a_{i_j} , where $1 \leq j \leq k$, and i_1 through i_k are the indexes of the attributes with a restriction term. We say that a predicate *partially matches* a tuple if at least one predicate term p_j , $1 \leq j \leq k$, matches an attribute a_{i_j} of the tuple. For example, consider the following relation and predicates:

$R(A, B, C)$ – indexes on A and B, no index on C

contents:

$t: \langle A=5, B=10, C=14 \rangle$

$D_1: A=4 \text{ and } C=15$

$D_2: A=4 \text{ and } B=10$

$D_3: A=5 \text{ and } B=10$

Given this relation and set of predicates, we would say that D_2 partially matches t on attribute B, D_3 partially matches t on both A and B, and D_1 has no partial matches for t . In general, a predicate P matches a tuple if and only if the tuple partially matches P on all k terms of P . Attributes for which P does not have a term need not be considered. Thus, D_3 matches t , and D_1 and D_2 do not. Furthermore, if a predicate does not partially match a tuple on all terms, then it definitely does not match the tuple.

Given this background, the mark intersection algorithm can be described. First, the algorithm requires every predicate to have at least one term on an indexed attribute. Each rule sets locks in all indexes for which it has a restriction term, unlike the basic locking algorithms which set locks in only one index. Besides just locking the indexes, a list of rule-ids is stored on every indexed attribute of every tuple in the database. A rule-id is stored on a tuple attribute if and only if that rule partially matches the tuple on that attribute and the attribute has an index. Associated with each rule-id stored with an attribute is information giving the positions of the attributes for which the rule predicate has a restriction term. This information can be stored with each rule-id as a bit string $b_1 \dots b_F$ where b_i is 1 if the rule has a restriction term on attribute i , and 0 otherwise. In the example database, t will have the following lists of marks and associated bit strings on its attributes:

$t: \langle A=5; [\langle D_3, 110 \rangle], B=10; [\langle D_2, 110 \rangle, \langle D_3, 110 \rangle], C=15 \rangle$

Notice that attribute C of t has no list of rule-ids or bit strings since R.C has no index.

When a tuple is inserted into a relation, the lists of markings for the tuple attributes must be determined. This is easily accomplished since all indexes must be consulted, and the marks can be determined at low cost when the indexes are searched. Furthermore, if the tuple is modified with a replace command, the list of marks on an attribute changes only if that attribute was modified. This means that only indexes corresponding to updated attributes need be inspected and appropriate markings recomputed. Thus, no more index I/O is required than for normal query processing.

The set S of predicates that *might* match a tuple, t , is determined by finding the set of all $\langle \text{rule-id}, \text{bit-string} \rangle$ pairs on the tuple such that the following holds:

for every attribute of R that has an index, and for which the bit-string element b_i is a 1, an identical $\langle \text{rule-id}, \text{bit-string} \rangle$ pair occurs on the mark list for attribute i .

For example, examining t as shown above, it is not possible to rule out a match on t for D_3 since b_1 and b_2 are 1 and $\langle D_3, 110 \rangle$ is on the list for A and B. However it is possible to rule out a match on D_2 since $b_1=1$ for D_2 , but $\langle D_2, 110 \rangle$ is not on the mark list for attribute A. The collection of predicates that cannot be ruled out must be checked just as in the basic locking algorithm.

This scheme may offer good performance because it reduces the number of predicates that must be accessed and checked compared to the basic locking scheme. An analysis of the cost of this scheme is the subject of a separate paper [HANS86].

7. CONCLUSIONS

We have presented alternate implementations for rule indexing in a data base system. Our first proposal resembles predicate locking and uses a variation of R-trees to index the rule set. The other approach resembles physical locking. We analyze both early and late versions of each proposal.

Our performance analysis results show that it is not possible to choose one implementation to support efficiently every rule-based environment. Physical locking is most promising in simpler environments where the expected number of rules that cover a specific tuple is low. On the other hand, predicate indexing dominates at higher numbers of rules. The early versions perform better at update probabilities under 0.5, while the late alternatives excel at higher probabilities. Moreover, we proposed extensions to predicate indexing and basic locking that attempt to overcome their disadvantages. Analysis of these schemes and investigation of other extensions are a topic of future research.

8. REFERENCES

- [ESWA75] Eswaran, K., *A General Purpose Trigger Subsystem and Its Inclusion in a Relational Data Base System*, IBM Research, Report No RJ 1833, San Jose, CA, July 1976.
- [ESWA76] Eswaran, K.P., J.N. Gray, R.A. Lorie, I.L. Traiger, *The Notions of Consistency and Predicate Locks in a Database System*, Communications of the ACM, (19), 11, 1976.
- [GRAY78] Gray, J.N., *Notes on Data Base Operating Systems*, IBM Research, Report No RJ 2254, San Jose, CA, August 1978.
- [GUTT84a] Guttman, A., *R-Trees: A Dynamic Index Structure for Spatial Searching*, Proceedings of the 1984 ACM-SIGMOD International Conference on Management of Data, Boston, MA, June 1984.
- [GUTT84b] Guttman, A., *New Features for Relational Database Systems to Support CAD Applications*, PhD Thesis, University of California, Berkeley, June 1984.
- [HANS86] Hanson, E. et. al., *Lock-Based Algorithms for Rule Processing in Databases*, (in preparation).
- [IOAN84] Ioannidis, Y., et al., *Enhancing INGRES with Deductive Power*, Position Paper, Proceedings of the 1st International Workshop on Expert Data Base Systems, Kiawah Isl., SC, October 1984.
- [ROBI81] Robinson, J. T., *The K-D-B tree: A Search Structure for Large Multidimensional Dynamic Indexes*, Proceedings of the 1981 ACM-SIGMOD International Conference on Management of Data, April 1981.

- [SELI79] Selinger, P. et. al., *Access Path Selection in a Relational Database Management System*, Proceedings of the 1979 ACM-SIGMOD International Conference on Management of Data, Boston, MA, June 1979.
- [SELL86] Sellis, T. et. al., *An Analysis of R⁺-Trees*, (in preparation).
- [STON75] Stonebraker, M., *Implementation of Integrity Constraints and Views by Query Modification*, Proceedings of the 1975 ACM-SIGMOD International Conference on Management of Data, San Jose, CA, June 1975.
- [STON84] Stonebraker, M. et. al., *QUEL as a Data Type*, Proceedings of the 1984 ACM-SIGMOD International Conference on Management of Data, Boston, MA, June 1984.
- [STON85a] Stonebraker, M., *Triggers and Inference in Data Base Systems*, Proceedings of the Islamorada Expert Data Base Conference, February 1985 (to appear in Springer-Verlag book).
- [STON85b] Stonebraker, M. et. al., *Extending a Data Base System with Procedures*, UC Berkeley, Memo No. UCB/ERL/M85/59, Berkeley, July 1985.
- [ULLM85] Ullman, J., *Implementation of Logical Query Languages for Data Bases*, Proceedings of the 1985 ACM-SIGMOD International Conference on Management of Data, Austin, TX, May 1985.
- [WONG76] Wong, E., Youssefi K., *Decomposition: A Strategy for Query Processing*, ACM Transactions on Database Systems, (1), 3, September 1976.
- [YAO78] Yao, A. C., *On Random 2-3 Trees*, Acta Informatica, (9), 2, 1978.