Stanford University Computer Science Dept. Technical Report *Stan-CS-93-1495*, Nov. 1993

A Rewriting Technique for Using Delta Relations to Improve Condition Evaluation in Active Databases*

Elena Baralis
Jennifer Widom

Department of Computer Science
Stanford University
Stanford, CA 94305-2140

{baralis,widom}@cs.stanford.edu

Abstract

We give a method for improving the efficiency of condition evaluation during rule processing in active database systems. The method derives, from a rule condition, two improved conditions that can be used in place of the original condition when a previous value (true or false) of the original condition is known. The derived conditions are more efficient than the original condition because they replace references to entire database relations by references to *delta relations*, which typically are much smaller. Delta relations are accessible to rule conditions in almost all current active database systems, making this optimization broadly applicable. We specify an implementation of our rewriting method based on attribute grammars.

1 Introduction

Active database systems allow users to specify *event-condition-action* rules that are processed automatically by the database system in response to data manipulation by users and applications. A rule's *event* specifies what causes the rule to become triggered; typical triggering events are data modification or data retrieval. A rule's *condition* is a further qualification of a triggered rule, usually expressed as a predicate or query over the database. A rule's *action* is performed when the rule is triggered and its condition is true; actions usually are sequences of arbitrary database commands. Most current active database systems (both research prototypes and commercial systems) use this rule paradigm; see [HW93].

Rule processing in active database systems usually consists of an iterative cycle in which: (1) a triggered rule is selected; (2) the rule's condition is evaluated; (3) if the condition is true the rule's action is executed. In this paper we give a method for optimizing step (2) in this cycle for active databases that use the relational model.[1] Our method is based on the following two assumptions:

*This work was partially performed while the authors were at the IBM Almaden Research Center, San Jose, CA. At Stanford this work was supported by equipment grants from Digital Equipment Corporation and IBM Corporation. Baralis' permanent address is: Dipartimento di Automatica e Informatica, Politecnico di Torino, Corso Duca degli Abruzzi, 24-10129 Torino, Italy.

[1]Adapting the method to active OODB's is planned for further work; see Section 6.

1. During the rule processing cycle, sometimes the rule processor can be aware of the previous value of a rule's condition, either because the rule was evaluated earlier during rule processing, or because the rule is enforcing an integrity constraint that was consistent at the beginning of the transaction.

2. The rule language permits references to *delta relations*, which contain the data that was modified in a relation since the last time the rule was evaluated or since the beginning of the transaction.[2]

Although these assumptions may appear strong, we note that many of the prominent relational active database systems do satisfy the assumptions, including *A-RDL* [SKdM92], *Ariel* [Han92], *Heraclitus* [GHJ+93], and *Starburst* [WCL91]. Furthermore, for those active database systems that do not exactly satisfy this paradigm, e.g. *Postgres* [SJGP90], relatively straightforward modifications of our method should be applicable.

Given a rule r with condition C, our method statically derives from C two improved conditions, $PT(C)$ (*Previously True* condition) and $PF(C)$ (*Previously False* condition). When rule r is selected at run time, if it is known that r's condition was previously true, then $PT(C)$ is evaluated instead of C; similarly, if it is known that r's condition was previously false, then $PF(C)$ is evaluated instead of C. $PT(C)$ and $PF(C)$ reference delta relations where C references entire relations, so $PT(C)$ and $PF(C)$ are likely to be much more efficient to evaluate than C.

For generality (and for ease in proving correctness) we present our method using a powerful rule condition language based on relational algebra. The adaptation of our method for a condition language based on SQL or Quel is straightforward. We also specify our method as an attribute grammar; this allows a direct implementation of the method using a compiler-generator such as YACC [Joh75].

1.1 Previous Related Work

There is a clear connection between our work and the well-studied problem of *incremental evaluation*, especially as addressed in [QW91,RCBB89]. [RCBB89] proposes an incremental optimization technique for rule conditions with similar goals to ours. However, rule conditions are restricted to Select-Project-Join (SPJ) expressions, and all relations are required to have user-accessible tuple identifiers. In contrast, our rule conditions are more general than SPJ expressions (see Section 3), and user-level tuple identifiers are not required. [QW91] proposes a set of transformations that compute incremental changes to arbitrary relational expressions. Although the methods in [QW91] do apply to (a small variation on) the problem we consider, the application itself is not direct and is rather difficult to understand. Furthermore, the method in [QW91] sometimes determines that there have been additions to and deletions from an expression when in fact the net effect of the

[2] We assume that delta relations can be accessed efficiently, since in most systems delta relations are stored or indexed in main memory. We also assume that delta relations typically are much smaller than the corresponding entire relations.

changes is null; our method detects that there has been no change. In both [QW91] and [RCBB89], update modifications are modeled as deletions followed by insertions. We handle update modifications directly, which in some cases produces more efficient incremental expressions than those in [QW91,RCBB89], particularly when only certain attributes are updated (see the γ transformation in Section 4). Note also that our attribute grammar specification is a unique approach which leads directly to an implementation.

Some active database systems use methods based on *Rete* or *TREAT* networks [WH92] for efficient condition evaluation, e.g. [FRS93,Han92]. Unfortunately, these methods apply only to rule languages where references to a relation R implicitly reference delta relations for R, and to rule conditions that are restricted to SPJ expressions. We consider more general conditions, and we determine those scenarios in which R can be replaced by its delta relations. Commercial active rule languages appear to be following our model, so our techniques should be applicable in very practical settings.

Finally, note that in [CW90] we suggest techniques similar to those we present here, but only a very restricted case is described. In this paper we elaborate the suggestions of [CW90] in a general context.

1.2 Outline of the Paper

In Section 2 we give a more rigorous description of active database rule processing and we formalize the notion of delta relations. In Section 3 we define our condition specification language and we provide some examples. Section 4 is the core technical section: it contains our method for rule condition rewriting, several examples of the method, and a proof of its correctness. Section 5 specifies an implementation of the method using an attribute grammar. Finally, Section 6 concludes and proposes improvements and extensions to our technique.

2 Rule Processing and Delta Relations

Consider an active database system in which a set of event-condition-action rules are defined as described in Section 1. Suppose further that a set of user or application modifications are performed on the database, then rule processing is invoked. The pseudo-code in Figure 1 describes the general behavior of the system. Note that issues such as the "granularity" of rule processing (i.e. when rule processing is invoked relative to triggering events) and the method for selecting among multiple triggered rules do not affect our method. Furthermore, our method also applies to rule languages in which triggering events are implicit rather than explicit, e.g. [Han92,SKdM92].

When a rule's condition is evaluated and its action is executed, this occurs with respect to a database *transition*, i.e. the changes that have occurred since some previous database state. We consider a semantics in which each rule uses the transition since that rule was last selected, or since the original state (state `S` in Figure 1) if the rule has not yet been selected during rule processing. While this is the semantics taken by many active database systems, systems with slightly different

```
S = initial DB state
user or application modifications
S' = new DB state
repeat until no rules are triggered:
       select a triggered rule r
       evaluate r's condition based on S' and delta relations
       if true, execute r's action based on S' and delta relations
       S' = new DB state
```

Figure 1: Active database system behavior

semantics may require corresponding modifications to our method.

Delta relations encapsulate the changes that have occurred during a rule's transition, and they may be referenced in a rule's condition and action. For each relation R we assume four delta relations:

- *inserted(R)* contains the tuples inserted into R during the transition.
- *deleted(R)* contains the tuples deleted from R during the transition.
- *old-updated(R)* contains the pre-transition values of the tuples modified in R during the transition.
- *new-updated(R)* contains the current (i.e. new) values of the tuples modified in R during the transition.

In addition, *new-updated* and *old-updated* may be restricted to sets of attributes. Let $A_1, \ldots, A_n$ be attributes of relation R.

- *old-updated(R,*$\{A_1, \ldots, A_n\}$*)* contains the pre-transition values of the tuples in R for which at least one of $A_1, \ldots, A_n$ was modified during the transition.
- *new-updated(R,*$\{A_1, \ldots, A_n\}$*)* contains the current values of the tuples in R for which at least one of $A_1, \ldots, A_n$ was modified during the transition.

Typically, delta relations reflect the *net effect* of database modifications; that is, they contain only the net result of successive actions over the same tuple. This concept of net effect is used widely, e.g. [Han92,SKdM92,WF90], so we assume it here. Note that one implication of using net effects is that *inserted(R)*, *deleted(R)*, *old-updated(R)*, and *new-updated(R)* are disjoint.

We introduce four abbreviations that are used throughout the remainder of the paper:

(1) $\Delta^+(R) =$ *inserted(R)* $\cup$ *new-updated(R)*

(2) $\Delta^-(R) =$ *deleted(R)* $\cup$ *old-updated(R)*

(3) $\Delta^+(R, \{A_1, \ldots, A_n\}) =$ *inserted(R)* $\cup$ *new-updated(R,*$\{A_1, \ldots, A_n\}$*)*

(4) $\Delta^-(R, \{A_1, \ldots, A_n\}) =$ *deleted(R)* $\cup$ *old-updated(R,*$\{A_1, \ldots, A_n\}$*)*

Note that in (3) and (4), if the attribute list is empty, then Δ^+ and Δ^- degenerate to *inserted(R)* and *deleted(R)*. We informally refer to (both versions of) Δ^+ and Δ^- as *incremental* and *decremental changes,* respectively.

Sometimes our improved conditions require access to the "old" value of a relation, i.e. the relation's pre-transition value. We denote the old value of a relation R as R^O. While some active database rule languages provide a feature for accessing R^O directly, others do not. However, R^O always can be derived from the current value of R and R's delta relations, based on the equivalence:

$$R^O = (R - \Delta^+(R)) \cup \Delta^-(R)$$

3 Condition Language

In active database rule languages, conditions are sometimes expressed as predicates and sometimes as queries, where in the latter case the condition is true iff the query produces a non-empty result. It can easily be shown that the two representations are equivalent [HW93]; we represent conditions as predicates. The grammar of our condition language is given in Figure 2. The language is powerful enough to describe any condition expressible in relational algebra or calculus extended with aggregate functions, with the exception of duplicates and ordering conditions.

In the grammar's productions, terminal symbol `R` stands for a relation name and `R.A` for an attribute of relation `R`. The meaning of the language is mostly self-explanatory. Condition $\exists(Rexp)$ is satisfied iff relational expression $Rexp$ produces one or more tuples, while condition $\neg\exists(Rexp)$ is satisfied iff $Rexp$ produces no tuples. We assume a set semantics, i.e. no duplicates. Note the following points:

- Joins are expressed using cross-product (production 8) and selection (production 10).
- A selection that is a boolean formula of comparisons can be expressed in our language by using the following equivalences:

$$\sigma_{c_1 \wedge c_2} Rexp = \sigma_{c_1}(\sigma_{c_2} Rexp)$$
$$\sigma_{c_1 \vee c_2} Rexp = \sigma_{c_1} Rexp \cup \sigma_{c_2} Rexp$$

- Negation of selection predicates can be expressed by repeatedly applying DeMorgan's laws and then negating the innermost comparisons (i.e. $\leq$ becomes $>$, $=$ becomes $\neq$, etc.).
- Although for simplicity we have omitted this from our grammar, *Terms* may be arithmetic expressions over attributes and constants without affecting our method.

1.	$Cond$	$::=$	$\exists(Rexp)$
2.		$\mid$	$\neg\exists(Rexp)$
3.		$\mid$	$Cond_1 \wedge Cond_2$
4.		$\mid$	$Cond_1 \vee Cond_2$
5.		$\mid$	$(Cond)$
6.	$Rexp$	$::=$	`R`
7.		$\mid$	$Rexp_1 \cup Rexp_2$
8.		$\mid$	$Rexp_1 \times Rexp_2$
9.		$\mid$	$Rexp - SimpleRexp$
10.		$\mid$	$\sigma_{Compare} Rexp$
11.		$\mid$	$\pi_{AList} Rexp$
12.		$\mid$	$Aggr(Attr_1[, Attr_2])Rexp$
13.		$\mid$	$(Rexp)$
14.	$SimpleRexp$	$::=$	`R`
15.		$\mid$	$SimpleRexp_1 \cup SimpleRexp_2$
16.		$\mid$	$SimpleRexp_1 \times SimpleRexp_2$
17.		$\mid$	$\sigma_{Compare} SimpleRexp$
18.		$\mid$	$\pi_{AList} SimpleRexp$
19.		$\mid$	$(SimpleRexp)$
20.	$Compare$	$::=$	$Term_1 \ Op \ Term_2$
21.	$Term$	$::=$	$Attr$
22.		$\mid$	`Const`
23.	Op	$::=$	$> \mid < \mid \leq \mid \geq \mid = \mid \neq$
24.	$Aggr$	$::=$	`sum` $\mid$ `avg` $\mid$ `min` $\mid$ `max` $\mid$ `count`
25.	$AList$	$::=$	$Attr_1, \ldots, Attr_n$
26.	$Attr$	$::=$	`R.A`

Figure 2: Condition language syntax

- The *Aggr* operation (production 12) is for handling the aggregate functions `sum`, `avg`, `min`, etc. It extends a given relational expression with a new attribute containing the computed value of the aggregate function. The function is computed over the attribute $Attr_1$ and grouping may optionally be performed by specifying the $Attr_2$ attribute; see [CG85] for a detailed description of this construct.

- "Simple relational expressions" (*SimpleRexp*) are introduced to restrict the expressions that may appear as the second operand of a difference (see production 9). The need for this restriction is explained in Section 4.

- Most active database rule languages permit explicit references to delta relations in rule conditions. Therefore, technically we should include *Rexp* productions for *inserted(R)*, *deleted(R)*, etc. in our grammar. However, since there is no need to attempt improvement for these references, for clarity and simplicity we omit them from the grammar; adding them is trivial.

3.1 Examples

We give four examples of conditions expressed in our language. The first two examples use the following two relations:

```
EMPLOYEE(name, deptno, salary)
DEPARTMENT(deptno, name, location, budget)
```

which we abbreviate as `E` and `D` respectively.

Example 3.1: Informally: All USA employees earn at least 1/100 of their department's budget. Equivalently, no USA employee earns less than 1/100 of his department's budget. In our condition language:

$$\neg\exists(\sigma_{\texttt{E.salary<0.01*D.budget}}\sigma_{\texttt{E.deptno=D.deptno}}(\texttt{E} \times \sigma_{\texttt{D.location='USA'}}\texttt{D}))$$

Example 3.2: Informally: Some employee earns less than 10 or some department has a budget less than 100. In our condition language:

$$\exists(\sigma_{\texttt{E.salary<10}}\texttt{E}) \vee \exists(\sigma_{\texttt{D.budget<100}}\texttt{D})$$

The next two examples are adapted from a case study in [CW90] involving an electrical power distribution network; these examples use the following two relations:

```
WIRE(wire-id,from,to,voltage)
TUBE(tube-id,from,to,protected)
```

which we abbreviate as `W` and `T`, respectively.

Example 3.3: Informally: Some unprotected tube contains a high voltage ($> 5k$) wire. In our condition language:

$$\exists(\sigma_{\texttt{W.<from,to>=T.<from,to>}}(\sigma_{\texttt{W.voltage>5k}}\texttt{W} \times \sigma_{\texttt{T.protected=false}}\texttt{T}))$$

where we use $\sigma_{\texttt{W.<from,to>=T.<from,to>}}$ as an abbreviation for $\sigma_{\texttt{W.from=T.from}}\sigma_{\texttt{W.to=T.to}}$.

Example 3.4: Informally: Some tube contains no wires. More precisely, there is some tube such that no wire has the same `from` and `to` attributes. In our condition language:

$$\exists(\texttt{T} - \pi_{\texttt{schema(T)}}(\sigma_{\texttt{T.<from,to>=W.<from,to>}}(\texttt{T} \times \texttt{W})))$$

where we use `schema(T)` as an abbreviation for a list of all the attributes of `T`.

4 Derivation of Improved Conditions

Let C be a condition expressed using the language of Section 3. Suppose we want to evaluate C with respect to a database state S' and the transition from some previous database state S. Further suppose that the result of C in state S is known, i.e. C was either true or false in S.[3] In this case we use one of two improved conditions in place of C:

- $PF(C)$ (PF for "Previously False") is chosen when the outcome of the previous evaluation of C was false.
- $PT(C)$ (PT for "Previously True") is chosen when the outcome of the previous evaluation of C was true.

In general, $PF(C)$ only provides a useful improvement for conditions with existential quantifications ($\exists$) and disjunction, while $PT(C)$ only provides a useful improvement for conditions with negative existential quantification ($\neg\exists$) and conjunction. Intuitively, this is for the following reasons. When the condition was previously false, an existential quantification tells us that no data satisfied the relational expression. Hence we can check if data now satisfies the relational expression by checking changed data only. However, suppose we have a negative existential quantification. Then, since the condition was previously false, some data did previously satisfy the relational expression. In this case it is impossible to tell, by examining changed data only, whether the relational expression is now empty. If a previously false condition contains disjuncts, then we know that all disjuncts were false and we can improve each one based on that knowledge. However, if the condition contains conjuncts, then we don't know which conjuncts were previously false, and improvement is impossible. The same argument, in converse, holds for previously true conditions.[4]

In $PF(C)$ and $PT(C)$, each reference to a relation R is replaced by one of its corresponding incremental or decremental changes, $\Delta^+(R)$ or $\Delta^-(R)$, whenever possible. We define $PF(C)$ and

[3] Recall from Figure 1 that this will be the case whenever C's rule has been selected previously or the value of C was known in the initial state before user modifications.

[4] Certainly we might do somewhat better here, e.g. handle some special cases, or keep track of which conjuncts/disjuncts were true/false. We plan to investigate this as future work.

C	$PF(C)$	$PT(C)$
$\exists(Rexp)$	$\exists(\vartheta(Rexp))$	$\exists(Rexp)$
$\neg\exists(Rexp)$	$\neg\exists(Rexp)$	$\neg\exists(\vartheta(Rexp))$
$Cond_1 \wedge Cond_2$	$Cond_1 \wedge Cond_2$	$PT(Cond_1) \wedge PT(Cond_2)$
$Cond_1 \vee Cond_2$	$PF(Cond_1) \vee PF(Cond_2)$	$Cond_1 \vee Cond_2$
$(Cond)$	$(PF(Cond))$	$(PT(Cond))$

Table 1: The PF and PT improved conditions

$Rexp$	$\vartheta(Rexp)$
R	$\Delta^+(\mathtt{R})$
$Rexp_1 \cup Rexp_2$	$\vartheta(Rexp_1) \cup \vartheta(Rexp_2)$
$Rexp_1 \times Rexp_2$	$(\vartheta(Rexp_1) \times Rexp_2) \cup (Rexp_1 \times \vartheta(Rexp_2))$
$Rexp_1 - SimpleRexp_2$	$(\vartheta(Rexp_1) - SimpleRexp_2) \cup (Rexp_1 \cap \vartheta'(SimpleRexp_2))$
$\sigma_{Compare} Rexp$	$\sigma_{Compare}(\vartheta(Rexp))$
$\pi_{AList} Rexp$	$\pi_{AList}(\vartheta(Rexp))$
$Aggr(Attr_1[, Attr_2])Rexp$	$Aggr(Attr_1[, Attr_2])Rexp$
$(Rexp)$	$(\vartheta(Rexp))$

Table 2: The ϑ transformation for relational expressions

$PT(C)$ by means of transformation rules based on the structure of C according to the grammar of Figure 2. The rules are given in Tables 1, 2 and 3; the rules are applied inductively to derive $PF(C)$ and $PT(C)$ for an arbitrarily complex condition. The correctness of these rules is shown in Section 4.2. We now provide more intuition for the rules, including an explanation for certain details of Tables 1, 2 and 3.

In Table 1, the rules for $PF(C)$ and $PT(C)$ use a transformation ϑ, which is applied to relational expressions. Table 2 contains the transformation rules for ϑ applied to an arbitrary relational expression $Rexp$, while Table 3 contains an additional transformation ϑ' applied to an arbitrary simple relational expression $SimpleRexp$. (Recall that simple relational expressions are those relational expressions that can appear as the second operand of a difference. Hence, ϑ' is applied in the fourth line of Table 2.) Intuitively, ϑ applied to a relational expression $Rexp$ produces an improved expression $Rexp'$ that computes the incremental changes to $Rexp$. When ϑ is applied to an expression with the difference operation, decremental changes to the second operand cause incremental changes to the entire expression. Hence, ϑ' is a "negated" version of ϑ that computes these decremental changes. Decremental changes can be computed only on *monotonic* relational expressions. This explains our introduction of simple relational expressions—simple relational expressions exclude difference and aggregate operators, which are non-monotonic.

$SimpleRexp$	$\vartheta'(SimpleRexp)$
R	Δ^-(R)
$SR_1 \cup SR_2$	$(\vartheta'(SR_1) - SR_2) \cup (\vartheta'(SR_2) - SR_1)$
$SR_1 \times SR_2$	$(\vartheta'(SR_1) \times SR_2^O) \cup (SR_1^O \times \vartheta'(SR_2))$
$\sigma_{Compare} SR$	$\sigma_{Compare}(\vartheta'(SR))$
$\pi_{AList} SR$	$\pi_{AList}(\vartheta'(SR)) - \pi_{AList} SR$
(SR)	$(\vartheta'(SR))$

Table 3: The ϑ' transformation for simple relational expressions

Observe the following points:

- For convenience, we use $\cap$ in ϑ applied to $Rexp_1 - SimpleRexp_2$. Expression $Rexp_1 \cap Rexp_2$ is equivalent to $Rexp_1 \times Rexp_2$ with a selection condition equating all corresponding attributes and appropriate projection.

- The computation of an aggregate function over a relational expression always requires the entire relational expression result (not just an incremental portion), thus aggregate expressions cannot be improved in the general case. However, conditions containing aggregate function expressions as operands still may be improved in their other operands.

- Note that in ϑ, the treatment of projection and union are substantially simpler than in general incremental query evaluation such as [QW91].

- ϑ' on cartesian products refers to "old" relational expression SR^O. Here we are using SR^O as an abbreviation for SR with all relation names R replaced by R^O, denoting the old value of R. (Recall that if old relations are not directly accessible, they can be derived from delta relations as described in Section 2.)

- For ϑ' applied to projections, if the attribute list $AList$ contains a key for the expression SR (as is often the case), then our formula can be simplified to $\pi_{AList}(\vartheta'(SR))$.

Note that grammar productions 14 – 20 from Figure 2 are not affected by the ϑ transformation, so they are not included in our tables.

While we expect the sizes of $\Delta^+(R)$ and $\Delta^-(R)$ in $PF(C)$ and $PT(C)$ to be small (much smaller than R), we can further reduce the sizes of $\Delta^+(R)$ and $\Delta^-(R)$ by ignoring all updates to attributes that do not influence the outcome of the condition. Let R' denote a specific reference to relation R in a condition C. We define the *relevant attribute set* $\rho(R', C)$ as the set of all attributes in R whose updates can affect the outcome of condition C with respect to reference R'. Let $Rexp(R')$ be any relational expression (or simple relational expression) in C containing the reference R'. For each attribute A_i of R, $A_i \in \rho(R', C)$ iff:

1. A_i appears in a selection comparison over $Rexp(R')$ (grammar productions 10, 17), or

2. A_i is used for aggregate computation or grouping on $Rexp(R')$ ($Attr_1$ or $Attr_2$ in grammar production 12), or

3. there is a union or difference applied to $Rexp(R')$ (grammar productions 7, 9, 15).

Given an original condition C and improved condition C' (either $PF(C)$ or $PT(C)$), we further improve C' as $\gamma(C')$, where γ replaces every reference $\Delta^+(R')$ or $\Delta^-(R')$ in C' by $\Delta^+(R', \rho(R', C))$ or $\Delta^-(R', \rho(R', C))$. This additional improvement has not been suggested previously (to the best of our knowledge); it can be very effective in practice, as shown by the examples in the next section.

4.1 Examples

We show the improved conditions for the four examples introduced in Section 3.1.

Example 4.1: All USA employees earn at least 1/100 of their department's budget. In our condition language:

$$C = \neg\exists(\sigma_{\texttt{E.salary}<0.01*\texttt{D.budget}}\sigma_{\texttt{E.deptno}=\texttt{D.deptno}}(\texttt{E} \times \sigma_{\texttt{D.location}='\texttt{USA}'}\texttt{D}))$$

C contains a negative existential quantification so, as explained in Section 4, only the previously true case is improved effectively. The improved conditions are:

$PF(C) = C$
$PT(C) = \neg\exists(\sigma_{\texttt{E.salary}<0.01*\texttt{D.budget}}\sigma_{dno}((\mathit{incr\text{-}E} \times \sigma_{USA}\texttt{D}) \cup (\texttt{E} \times \sigma_{USA}\mathit{incr\text{-}D})))$

where

dno is $\texttt{E.deptno} = \texttt{D.deptno}$
USA is $\texttt{D.location} =' \texttt{USA}'$
$\mathit{incr\text{-}E}$ is $\Delta^+(\texttt{E}, \{\texttt{deptno}, \texttt{salary}\})$
$\mathit{incr\text{-}D}$ is $\Delta^+(\texttt{D}, \{\texttt{deptno}, \texttt{location}, \texttt{budget}\})$

Example 4.2: Some employee earns less than 10 or some department has a budget less than 100. In our condition language:

$$C = \exists(\sigma_{\texttt{E.salary}<10}\texttt{E}) \vee \exists(\sigma_{\texttt{D.budget}<100}\texttt{D})$$

C contains a disjunction so, as explained in Section 4, only the previously false case is improved effectively. The improved conditions are:

$PF(C) = \exists(\sigma_{\texttt{E.salary}<10}\Delta^+(\texttt{E}, \{\texttt{salary}\})) \vee \exists(\sigma_{\texttt{D.budget}<100}\Delta^+(\texttt{D}, \{\texttt{budget}\}))$
$PT(C) = C$

Example 4.3: Some unprotected tube contains a high voltage ($> 5k$) wire. In our condition language:

$$C = \exists(\sigma_{\texttt{W.<from,to>=T.<from,to>}}(\sigma_{\texttt{W.voltage>5k}}\texttt{W} \times \sigma_{\texttt{T.protected=false}}\texttt{T}))$$

The improved conditions are:

$$PF(C) = \exists(\sigma_{join}(\sigma_{volt}\,incr\text{-}W \times \sigma_{unprot}\texttt{T}) \cup (\sigma_{volt}\texttt{W} \times \sigma_{unprot}\,incr\text{-}T)))$$
$$PT(C) = C$$

where

join is $\texttt{W.<from,to>} = \texttt{T.<from,to>}$
volt is $\texttt{W.voltage} > \texttt{5k}$
unprot is $\texttt{T.protected} = \texttt{false}$
incr-W is $\Delta^+(\texttt{W}, \{\texttt{voltage}, \texttt{from}, \texttt{to}\})$
incr-T is $\Delta^+(\texttt{T}, \{\texttt{from}, \texttt{to}, \texttt{protected}\})$

Note that this same example could not be improved using the method presented in [CW90].

Example 4.4: Some tube contains no wires. In our condition language:

$$C = \exists(\texttt{T} - \pi_{\texttt{schema(T)}}(\sigma_{\texttt{T.<from,to>=W.<from,to>}}(\texttt{T} \times \texttt{W})))$$

Since this example has a difference operator, transformation ϑ' is used as well as ϑ. Recall that $\texttt{W}^{\texttt{O}}$ and $\texttt{T}^{\texttt{O}}$ refer to the pre-transition (old) states of W and T, respectively. Note also that here delta relations cannot be restricted to updates on specific attributes (because of the difference operator), and that we are applying the simplified formula for ϑ' applied to projections since the projected attributes form a key. The improved conditions are:

$$PF(C) = \exists((\Delta^+(\texttt{T}) - \pi_t(\sigma_{join}(\texttt{T} \times \texttt{W}))) \cup (\texttt{T} \cap \pi_t(\sigma_{join}((\Delta^-(\texttt{T}) \times \texttt{W}^{\texttt{O}}) \cup (\texttt{T}^{\texttt{O}} \times \Delta^-(\texttt{W}))))))$$
$$PT(C) = C$$

where

t is $\texttt{schema(T)}$
join is $\texttt{T.<from,to>} = \texttt{W.<from,to>}$

As in Example 4.3, this condition could not be improved using the method in [CW90].

4.2 Correctness of the transformations

In the following, we assume that S and S' are two database states such that S occurs before S', and $\Delta^+(R)$ and $\Delta^-(R)$ contain the changes to any relation R between S and S', as defined in Section 2. Any reference to a relation R implicitly refers to the current state, i.e. R in state S', unless R^O is explicitly specified, in which case it refers to R in state S. Recall that sometimes we use the notation $Rexp^O$ to denote $Rexp$ with all relation references R replaced by R^O.

The proof of correctness of the $PF(C)$ and $PT(C)$ transformations is given by Theorem 1. However, before proving the Theorem we must prove a number of Lemmas. Note that, due to their length, most proofs appear in an Appendix rather than in the body of the paper. The first Lemma states that the decremental change to a relational expression is a subset of the relational expression in the "old" state (state S).

Lemma 1: If SR is a simple relational expression, then $\vartheta'(SR) \subseteq SR^O$.

Proof: See Appendix.

The next Lemma proves the correctness of the ϑ' transformation applied to simple relational expressions. It states that the ϑ' transformation yields an expression that produces all data in the old state but not in the current state.

Lemma 2: If SR is a simple relational expression, then $\vartheta'(SR)$ produces the data in $SR^O - SR$.

Proof: See Appendix.

The following Corollary is used in the proof of Lemma 4 below. It states that, for a given expression, the data in the old state is a subset of the data in the current state together with what has been deleted.

Corollary 1: If SR is a simple relational expression, then $SR^O \subseteq SR \cup \vartheta'(SR)$.

Proof: See Appendix.

Lemma 3 states that data added to a relational expression as defined by ϑ is a subset of the data in the relational expression; Lemma 4 states that the data in a relational expression is a subset of the data in the old state together with the data added.

Lemma 3: If Re is a relational expression, then $\vartheta(Re) \subseteq Re$.

Proof: See Appendix.

Lemma 4: If Re is a relational expression, then $Re \subseteq Re^O \cup \vartheta(Re)$. (Equivalently, if Re is a relational expression, then $Re - \vartheta(Re) \subseteq Re^O - \vartheta(Re)$.)

Proof: See Appendix.

The final Lemma proves the correctness of the ϑ transformation applied to relational expressions, as applied in our context. This Lemma states the equivalence of $\vartheta(Re)$ and Re when it is known that the evaluation of Re in state S produced no data. This Lemma is the cornerstone needed to prove our correctness Theorem.

Lemma 5: Let Re be a relational expression that produced no data in state S. Then $Re = \vartheta(Re)$ in state S'.

Proof: We show that $\vartheta(Re) \subseteq Re$ and $\vartheta(Re) \supseteq Re$, which proves the equivalence. Lemma 3 proves that $\vartheta(Re) \subseteq Re$. We know from Lemma 4 that $Re^O \cup \vartheta(Re) \supseteq Re$. By hypothesis $Re^O = \emptyset$, thus $\vartheta(Re) \supseteq Re$. $\square$

There is one more Corollary needed before the main Theorem:

Corollary 2: Let Re be a relational expression that produced no data in state S. Then $\vartheta(Re)$ is empty in state S' iff Re is empty in state S'.

Proof: Follows directly from Lemma 5, since $Re = \vartheta(Re)$ in state S'. $\square$

The following Theorem proves the main result of the paper, the correctness of the improved conditions $PF(C)$ and $PT(C)$.

Theorem 1: Let C denote any condition specified using the language of Figure 2. Then:
(a) If C was false in state S, C is true in state S' iff $PF(C)$ is true in state S'.
(b) If C was true in state S, C is true in state S' iff $PT(C)$ is true in state S'.

Proof: We prove our result using the ϑ transformation only. The correctness of the γ transformation is self-evident. The proof proceeds by structural induction on condition C.

1. $C = \exists(\mathit{Rexp})$
   (a) If C was false in S, then Rexp produced no data in S. By Corollary 2, $\vartheta(\mathit{Rexp})$ produces data in state S' iff Rexp does. Hence, $\exists(\mathit{Rexp})$ is true in S' iff $\exists(\vartheta(\mathit{Rexp}))$ is.
   (b) Holds trivially.

   $C = \neg\exists(\mathit{Rexp})$
   (a) Holds trivially.
   (b) If C was true in S, then Rexp produced no data in S. By Corollary 2, $\vartheta(\mathit{Rexp})$ produces data in state S' iff Rexp does. Hence, $\neg\exists(\mathit{Rexp})$ is true in S' iff $\neg\exists(\vartheta(\mathit{Rexp}))$ is.

2. $C = C_1 \wedge C_2$
   (a) Holds trivially.
   (b) Since C was true in S, both C_1 and C_2 were true in S. By the induction hypothesis, C_1 is true in state S' iff $PT(C_1)$ is true in state S'; similarly for C_2. Therefore, $C_1 \wedge C_2$ is true in state S' iff $PT(C_1) \wedge PT(C_2)$ is.

3. $C = C_1 \vee C_2$
   (a) Since C was false in S, both C_1 and C_2 were false in S. By the induction hypothesis, C_1 is true in state S' iff $PF(C_1)$ is true in state S'; similarly for C_2. Therefore, $C_1 \vee C_2$ is true in state S' iff $PF(C_1) \vee PF(C_2)$ is.
   (b) Holds trivially.

4. $C = (C)$. Holds trivially for both (a) and (b). $\square$

5 Attribute Grammar Implementation

Now we address the issue of implementing our approach. Suppose a new rule r with condition C is added to the database (or an existing rule's condition is to be improved). The following steps are followed:

1. C is translated into our condition language.
2. The improved conditions $PF(C)$ and $PT(C)$ are generated.
3. $PF(C)$ and $PT(C)$ are translated back to the rule's condition language.[5]
4. At run-time, whenever possible, the evaluation of C is replaced by the evaluation of $PF(C)$ or $PT(C)$.

Note that steps 1–3 are *static*: they are performed only once, when rule r is first defined. At run-time, the system evaluates an improved condition (step 4) just as it would have evaluated the original condition.

We consider in detail the implementation of step 2. The ϑ and γ transformations can be computed during the parsing of a condition to yield the improved conditions. Thus it is a natural choice to implement the derivation of improved conditions by means of an *attribute grammar*. In an attribute grammar, each symbol is allowed to have a fixed number of associated values, called attributes, and each grammar production has a set of attribute evaluation rules. Attributes can be used to pass information up a syntax tree: these are called *synthesized attributes*, and the evaluation rules associated with each production describe how the attributes' left-hand-side (LHS) values are computed from their right-hand-side (RHS) values. If instead the information flows down the syntax tree, these are called *inherited attributes*, and their evaluation rules describe how RHS attribute values are computed as a function of LHS values. For a more detailed description of attribute grammars refer to [ASU86].

Our attribute grammar computes the ϑ and γ transformations at the same time. At the end of the parsing process, the grammar produces the improved conditions $PF(C)$ and $PT(C)$ as attributes. The grammar uses one inherited attribute (γ) and eight synthesized attributes (ϑ, ϑ', PF, PT, A, C, E, EO), as follows:

- The ϑ and ϑ' synthesized attributes implement the ϑ and ϑ' transformations. At any time during the parsing process, Rexp.ϑ for a relational expression *Rexp* contains ϑ applied to *Rexp*. ϑ' is similarly defined for simple relational expressions.
- The PF and PT synthesized attributes are used to build the complete improved condition, using as building blocks the improved conditions provided by the ϑ attribute. At the end of the parsing process, these attributes contain the improved conditions $PF(C)$ and $PT(C)$, respectively.

[5]Note from Section 4 that our process is highly unlikely to yield constructs in $PF(C)$ and $PT(C)$ that are not available in the language used to specify the original rule condition.

$$
\begin{array}{llll}
1. & Cond & ::= & \exists(Rexp) \\
 & & & \mathsf{Rexp}.\gamma := \emptyset \\
 & & & \mathsf{Cond.PF} := \exists(\mathsf{Rexp}.\vartheta) \\
 & & & \mathsf{Cond.PT} := \exists(\mathsf{Rexp.E}) \\
 & & & \mathsf{Cond.C} := \exists(\mathsf{Rexp.E}) \\
2. & & | & \neg\exists(Rexp) \\
 & & & \mathsf{Rexp}.\gamma := \emptyset \\
 & & & \mathsf{Cond.PF} := \neg\exists(\mathsf{Rexp.E}) \\
 & & & \mathsf{Cond.PT} := \neg\exists(\mathsf{Rexp}.\vartheta) \\
 & & & \mathsf{Cond.C} := \neg\exists(\mathsf{Rexp.E}) \\
3. & & | & Cond_1 \wedge Cond_2 \\
 & & & \mathsf{Cond.PF} := \mathsf{Cond_1.C} \wedge \mathsf{Cond_2.C} \\
 & & & \mathsf{Cond.PT} := \mathsf{Cond_1.PT} \wedge \mathsf{Cond_2.PT} \\
 & & & \mathsf{Cond.C} := \mathsf{Cond_1.C} \wedge \mathsf{Cond_2.C} \\
4. & & | & Cond_1 \vee Cond_2 \\
 & & & \mathsf{Cond.PF} := \mathsf{Cond_1.PF} \vee \mathsf{Cond_2.PF} \\
 & & & \mathsf{Cond.PT} := \mathsf{Cond_1.C} \vee \mathsf{Cond_2.C} \\
 & & & \mathsf{Cond.C} := \mathsf{Cond_1.C} \vee \mathsf{Cond_2.C} \\
5. & & | & (Cond_1) \\
 & & & \mathsf{Cond.PF} := (\mathsf{Cond_1.PF}) \\
 & & & \mathsf{Cond.PT} := (\mathsf{Cond_1.PT}) \\
 & & & \mathsf{Cond.C} := (\mathsf{Cond_1.C})
\end{array}
$$

Figure 3: Attribute grammar for the *Cond* production

6. $Rexp$ ::= $\mathtt{R}$

$\mathsf{R}.\gamma := \pi_{\mathsf{schema(R)}}(\mathsf{Rexp}.\gamma)$
$\mathsf{Rexp}.\vartheta := \Delta^{+}(\mathtt{R}, \{\mathtt{R}.\gamma\})$
$\mathsf{Rexp.E} := \mathsf{R}$

7. | $Rexp_1 \cup Rexp_2$

$\mathsf{Rexp_1}.\gamma := \mathsf{All(Rexp_1)}$
$\mathsf{Rexp_2}.\gamma := \mathsf{All(Rexp_2)}$
$\mathsf{Rexp}.\vartheta := \mathsf{Rexp_1}.\vartheta \cup \mathsf{Rexp_2}.\vartheta$
$\mathsf{Rexp.E} := \mathsf{Rexp_1.E} \cup \mathsf{Rexp_2.E}$

8. | $Rexp_1 \times Rexp_2$

$\mathsf{Rexp_1}.\gamma := \mathsf{Rexp}.\gamma$
$\mathsf{Rexp_2}.\gamma := \mathsf{Rexp}.\gamma$
$\mathsf{Rexp}.\vartheta := (\mathsf{Rexp_1}.\vartheta \times \mathsf{Rexp_2.E}) \cup (\mathsf{Rexp_1.E} \times \mathsf{Rexp_2}.\vartheta)$
$\mathsf{Rexp.E} := \mathsf{Rexp_1.E} \times \mathsf{Rexp_2.E}$

9. | $Rexp_1 - SimpleRexp_2$

$\mathsf{Rexp_1}.\gamma := \mathsf{All(Rexp_1)}$
$\mathsf{SimpleRexp_2}.\gamma := \mathsf{All(SimpleRexp_2)}$
$\mathsf{Rexp}.\vartheta := (\mathsf{Rexp_1}.\vartheta - \mathsf{SimpleRexp_2.E}) \cup (\mathsf{Rexp_1.E} \cap \mathsf{SimpleRexp_2}.\vartheta')$
$\mathsf{Rexp.E} := \mathsf{Rexp_1.E} - \mathsf{SimpleRexp_2.E}$

10. | $\sigma_{Compare} Rexp_1$

$\mathsf{Rexp_1}.\gamma := \mathsf{Rexp}.\gamma \oplus \mathsf{Compare.A}$
$\mathsf{Rexp}.\vartheta := \sigma_{\mathsf{Compare}}(\mathsf{Rexp_1}.\vartheta)$
$\mathsf{Rexp.E} := \sigma_{\mathsf{Compare}}(\mathsf{Rexp_1.E})$

11. | $\pi_{AList} Rexp_1$

$\mathsf{Rexp_1}.\gamma := \mathsf{Rexp}.\gamma$
$\mathsf{Rexp}.\vartheta := \pi_{\mathsf{AList}}(\mathsf{Rexp_1}.\vartheta)$
$\mathsf{Rexp.E} := \pi_{\mathsf{AList}}(\mathsf{Rexp_1.E})$

12. | $Aggr(Attr_1[, Attr_2]) Rexp_1$

$\mathsf{Rexp_1}.\gamma := \mathsf{Rexp}.\gamma \oplus \mathsf{Attr_1.A} \oplus \mathsf{Attr_2.A}$
$\mathsf{Rexp}.\vartheta := \mathsf{Aggr(Attr_1, Attr_2)Rexp_1.E}$
$\mathsf{Rexp.E} := \mathsf{Aggr(Attr_1, Attr_2)Rexp_1.E}$

13. | $(Rexp_1)$

$\mathsf{Rexp_1}.\gamma := \mathsf{Rexp}.\gamma$
$\mathsf{Rexp}.\vartheta := (\mathsf{Rexp_1}.\vartheta)$
$\mathsf{Rexp.E} := (\mathsf{Rexp_1.E})$

Figure 4: Attribute grammar for the *Rexp* production

14. $SRexp$ ::= R

$\mathsf{R}.\gamma := \pi_{\mathsf{schema(R)}}(\mathsf{SRexp}.\gamma)$
$\mathsf{SRexp}.\vartheta' := \Delta^{-}(\mathtt{R}, \{\mathtt{R}.\gamma\})$
$\mathsf{SRexp.E} := \mathtt{R}$
$\mathsf{SRexp.EO} := \mathtt{R}^{0}$

15. | $SRexp_1 \cup SRexp_2$

$\mathsf{SRexp}_1.\gamma := \mathsf{All}(\mathsf{SRexp}_1)$
$\mathsf{SRexp}_2.\gamma := \mathsf{All}(\mathsf{SRexp}_2)$
$\mathsf{SRexp}.\vartheta' := (\mathsf{SRexp}_1.\vartheta' - \mathsf{SRexp}_2.\mathsf{E}) \cup (\mathsf{SRexp}_2.\vartheta' - \mathsf{SRexp}_1.\mathsf{E})$
$\mathsf{SRexp.E} := \mathsf{SRexp}_1.\mathsf{E} \cup \mathsf{SRexp}_2.\mathsf{E}$
$\mathsf{SRexp.EO} := \mathsf{SRexp}_1.\mathsf{EO} \cup \mathsf{SRexp}_2.\mathsf{EO}$

16. | $SRexp_1 \times SRexp_2$

$\mathsf{SRexp}_1.\gamma := \mathsf{SRexp}.\gamma$
$\mathsf{SRexp}_2.\gamma := \mathsf{SRexp}.\gamma$
$\mathsf{SRexp}.\vartheta' := (\mathsf{SRexp}_1.\vartheta' \times \mathsf{SRexp}_2.\mathsf{EO}) \cup (\mathsf{SRexp}_1.\mathsf{EO} \times \mathsf{SRexp}_2.\vartheta')$
$\mathsf{SRexp.E} := \mathsf{SRexp}_1.\mathsf{E} \times \mathsf{SRexp}_2.\mathsf{E}$
$\mathsf{SRexp.EO} := \mathsf{SRexp}_1.\mathsf{EO} \times \mathsf{SRexp}_2.\mathsf{EO}$

17. | $\sigma_{Compare} SRexp_1$

$\mathsf{SRexp}_1.\gamma := \mathsf{SRexp}.\gamma \oplus \mathsf{Compare.A}$
$\mathsf{SRexp}.\vartheta' := \sigma_{\mathsf{Compare}}(\mathsf{SRexp}_1.\vartheta')$
$\mathsf{SRexp.E} := \sigma_{\mathsf{Compare}}(\mathsf{SRexp}_1.\mathsf{E})$
$\mathsf{SRexp.EO} := \sigma_{\mathsf{Compare}}(\mathsf{SRexp}_1.\mathsf{EO})$

18. | $\pi_{AList} SRexp_1$

$\mathsf{SRexp}_1.\gamma := \mathsf{SRexp}.\gamma$
$\mathsf{SRexp}.\vartheta' := \pi_{\mathsf{AList}}(\mathsf{SRexp}_1.\vartheta') - \pi_{\mathsf{AList}}(\mathsf{SRexp}_1.\mathsf{E})$
$\mathsf{SRexp.E} := \pi_{\mathsf{AList}}(\mathsf{SRexp}_1.\mathsf{E})$
$\mathsf{SRexp.EO} := \pi_{\mathsf{AList}}(\mathsf{SRexp}_1.\mathsf{EO})$

19. | $(SRexp_1)$

$\mathsf{SRexp}_1.\gamma := \mathsf{SRexp}.\gamma$
$\mathsf{SRexp}.\vartheta' := (\mathsf{SRexp}_1.\vartheta')$
$\mathsf{SRexp.E} := (\mathsf{SRexp}_1.\mathsf{E})$
$\mathsf{SRexp.EO} := (\mathsf{SRexp}_1.\mathsf{EO})$

Figure 5: Attribute grammar for the *SimpleRexp* production

20.	$Compare$	$::=$	$Term_1 \; Op \; Term_2$
			$\mathsf{Compare.A} := \mathsf{Term_1.A} \oplus \mathsf{Term_2.A}$
21.	$Term$	$::=$	$Attr$
			$\mathsf{Term.A} := \mathsf{Attr.A}$
22.		$\mid$	`Const`
			$\mathsf{Term.A} := \emptyset$
23.	Op	$::=$	$> \mid < \mid \leq \mid \geq \mid = \mid\ != $
24.	$Aggr$	$::=$	`sum` $\mid$ `avg` $\mid$ `min` $\mid$ `max` $\mid$ `count`
25.	$AList$	$::=$	$Attr_1, \ldots, Attr_n$
26.	$Attr$	$::=$	`R.a`
			$\mathsf{Attr.A} := \mathsf{R.a}$

Figure 6: Attribute grammar for the remaining productions

- The γ inherited attribute allows us to progressively build top-down the relevant attributes set ($\rho(R, C)$ from Section 4), by adding the attributes involved in all predicates and aggregate functions while descending the parse tree. When the relation terminal symbol is reached (`R`), the reduced delta relation is defined using the relevant attribute set contained in the γ attribute. Note that for γ we use $\mathsf{schema(R)}$ to denote all attributes of relation `R`, and $\mathsf{All(Rexp)}$ to denote all relation attributes in expression $Rexp$.

- The A, C, E, and EO synthesized attributes are needed to aid the propagation process. Attribute A allows us to extract from a predicate or aggregate definition the list of relevant attributes. Attributes C, E, and EO propagate up the parse tree the definition of the non-improved condition, and the current and old relational expressions, so that they are available each time they are required in the definition of $PF(C)$, $PT(C)$, ϑ, and ϑ'.

In Figures 3, 4, and 5, the attribute grammars for the $Cond$, $Rexp$, and $SimpleRexp$ productions are given, while the remaining productions are given in Figure 6. The renaming of tokens appearing in both the LHS and RHS productions, although not necessary, improves the readability of the grammar. The $\oplus$ operator performs list concatenation; it is used to build attribute lists. Note that not every production needs to compute every attribute, since different attributes have different scopes; for details see [ASU86].

We are currently building a prototype condition rewriter using the parser-generator tools *LEX* [Les75] and *YACC* [Joh75]. We hope to incorporate our algebraic condition rewriting facility into an active database rule system, most likely Starburst [WCL91]. This should be relatively straightforward: Delta relations as used in this paper are directly available in Starburst. The Starburst SQL-based condition language easily translates to and from our condition language. We can generate the improved conditions and store them with the original condition in the Starburst Rule Catalog [WCL91]. It is then sufficient to add to the run-time rule processor the logic for:

(a) remembering a previous outcome of condition evaluation for each rule, and (b) choosing an improved condition for evaluation in place of the original condition whenever possible. Based on the implementation architecture of the Starburst Rule System [WCL91], both of these tasks can be performed easily and efficiently.

6 Conclusions and Future Work

We have described a method for improving the condition evaluation phase of active database rule processing. Our method is based on rule conditions expressed in an extension of relational algebra; this provides both a logical formulation and a framework that can apply to multiple rule languages. We have proven the correctness of our approach, and we have specified an implementation based on an attribute grammar.

As future work we plan to:

- Improve our handling of aggregate functions
- Extend our condition language to more succinctly express certain constructs, so that we can improve our rewriting for these constructs (e.g. negative subqueries, which are now expressed as relational difference)
- Handle special cases where $PF(C)$ and $PT(C)$ can improve conditions currently not improved (recall Section 4)
- Consider similar methods in the context of deductive and object-oriented active database systems (e.g. we hope to use these methods in the IDEA project [CM93], where they should apply in a straightforward way)
- Implement our method in the Starburst Rule System and experiment with its practical effectiveness under a variety of database and rule processing loads

Finally, although we believe that our improved conditions will be much cheaper to evaluate in most cases, this does assume a reasonably sophisticated query optimizer (i.e. an optimizer that can exploit references to very small relations). In addition, there may be states that arise during rule processing in which it actually is cheaper to evaluate the original condition rather than the improved condition. We plan to develop a framework whereby we can characterize those states in which incremental conditions may not be cheaper (e.g., by using run-time statistics). We also hope to investigate query optimization strategies that handle our improved conditions in an "intelligent" way.

Appendix

We prove the four Lemmas and one Corollary whose proofs were omitted in Section 4.2. Recall that we assumed S and S' were two database states such that S occurred before S', and $\Delta^+(R)$

and $\Delta^-(R)$ contain the changes to any relation R between S and S'.

Lemma 1: If SR is a simple relational expression, then $\vartheta'(SR) \subseteq SR^O$.

Proof: The proof proceeds by induction on the structure of SR.

Base case. $SR = R$. By the definition of delta relations, $\Delta^-(R) \subseteq SR^O$, so $\vartheta'(R) \subseteq SR^O$.

Induction.

1. $SR = SR_1 \cup SR_2$. By definition $\vartheta'(SR) = (\vartheta'(SR_1) - SR_2) \cup (\vartheta'(SR_2) - SR_1)$. By the induction hypothesis, $\vartheta'(SR_1) \subseteq SR_1^O$ and $\vartheta'(SR_2) \subseteq SR_2^O$. Then $\vartheta'(SR) \subseteq (SR_1^O - SR_2) \cup (SR_2^O - SR_1) \subseteq (SR_1^O \cup SR_2^O)$.

2. $SR = SR_1 \times SR_2$. By definition $\vartheta'(SR) = (\vartheta'(SR_1) \times SR_2^O) \cup (SR_1^O \times \vartheta'(SR_2))$. By the induction hypothesis, $\vartheta'(SR_1) \subseteq SR_1^O$ and $\vartheta'(SR_2) \subseteq SR_2^O$. Then $\vartheta'(SR) \subseteq (SR_1^O \times SR_2^O) \cup (SR_1^O \times SR_2^O) = SR_1^O \times SR_2^O$.

3. $SR = \sigma_{Compare} SR_1$. By definition $\vartheta'(SR) = \sigma_{Compare}(\vartheta'(SR_1))$. By applying the induction hypothesis on SR_1, we obtain $\vartheta'(SR) \subseteq \sigma_{Compare}(SR_1^O) = SR^O$.

4. $SR = \pi_{AList}(SR_1)$. By definition $\vartheta'(SR) = \pi_{AList}(\vartheta'(SR_1)) - \pi_{AList}(SR_1)$. By applying the induction hypothesis on SR_1, $\vartheta'(SR) \subseteq \pi_{AList}(SR_1^O) - \pi_{AList}(SR_1) \subseteq \pi_{AList}(SR_1^O) = SR^O$.

□

Lemma 2: If SR is a simple relational expression, then $\vartheta'(SR)$ produces the data in $SR^O - SR$.

Proof: The proof proceeds by induction on the structure of SR.

Base case. $SR = R$. We must prove that $R^O - R = \vartheta'(R) = \Delta^-(R)$. This follows directly from the definition of $\Delta^-(R)$.

Induction.

1. $SR = SR_1 \cup SR_2$.

$$\begin{array}{lll} \vartheta'(SR) & = (\vartheta'(SR_1) - SR_2) \cup (\vartheta'(SR_2) - SR_1) & \text{def. of } \vartheta' \\ & = ((SR_1^O - SR_1) - SR_2) \cup ((SR_2^O - SR_2) - SR_1) & \text{induction hyp.} \\ & = (SR_1^O - (SR_1 \cup SR_2)) \cup (SR_2^O - (SR_2 \cup SR_1)) & (A - B) - C = A - (B \cup C) \\ & = (SR_1^O \cup SR_2^O) - (SR_1 \cup SR_2) & (A - B) \cup (C - B) = (A \cup C) - B \\ & = SR^O - SR & \text{def. of SR} \end{array}$$

2. $SR = SR_1 \times SR_2$.

$$\begin{array}{lll} & SR^O - SR = & \\ = & SR^O - (SR^O \cap SR) & A - B = A - (A \cap B) \\ = & (SR_1^O \times SR_2^O) - ((SR_1^O \times SR_2^O) \cap (SR_1 \times SR_2)) & \text{def. of } SR \\ = & (SR_1^O \times SR_2^O) - ((SR_1^O \cap SR_1) \times (SR_2^O \cap SR_2)) & \text{schema (A) = schema (C) and} \\ & & \text{schema (B) = schema (D) thus} \\ & & (A \times B) \cap (C \times D) = (A \cap C) \times (B \cap D) \\ = & (SR_1^O \times SR_2^O) - ((SR_1^O - (SR_1^O - SR_1)) \times & A \cap B = A - (A - B) \\ & (SR_2^O - (SR_2^O - SR_2))) & \end{array}$$

$$
\begin{aligned}
&= (SR_1^O \times SR_2^O) - ((SR_1^O \times SR_2^O) - (SR_1^O \times && \text{"distribution" of } \times \\
&\quad (SR_2^O - SR_2)) - ((SR_1^O - SR_1) \times \\
&\quad (SR_2^O - (SR_2^O - SR_2)))) \\
&= (SR_1^O \times SR_2^O) \cap ((SR_1^O \times (SR_2^O - SR_2)) \cup && A-(A-B) = A \cap B \\
&\quad ((SR_1^O - SR_1) \times (SR_2^O - (SR_2^O - SR_2)))) \\
&= (SR_1^O \times SR_2^O) \cap ((SR_1^O \times \vartheta'(SR_2)) \cup && \text{induction hyp.} \\
&\quad (\vartheta'(SR_1) \times (SR_2^O - \vartheta'(SR_2)))) \\
&= (SR_1^O \times SR_2^O) \cap ((SR_1^O \times \vartheta'(SR_2)) \cup && \text{Lemma 1: } SR_1^O \supseteq \vartheta'(SR_1) \\
&\quad (\vartheta'(SR_1) \times \vartheta'(SR_2^O)) \cup (\vartheta'(SR_1) \times (SR_2^O - \vartheta'(SR_2))) \\
&= (SR_1^O \times SR_2^O) \cap \vartheta'(SR) && \text{def. of } \vartheta' \\
&= \vartheta'(SR) && \text{Lemma 1}
\end{aligned}
$$

3. $SR = \sigma_{Compare}(SR_1)$.

$$
\begin{aligned}
\vartheta'(SR) &= \sigma_{Compare}(\vartheta'(SR_1)) && \text{def. of } \vartheta' \\
&= \sigma_{Compare}(SR_1^O - SR_1) && \text{induction hyp.} \\
&= \sigma_{Compare}(SR_1^O) - \sigma_{Compare}(SR_1) && \text{distribution of } \sigma \text{ over diff.} \\
&= SR^O - SR && \text{def. of SR}
\end{aligned}
$$

4. $SR = \pi_{AList} SR_1$.

$$
\begin{aligned}
\vartheta'(SR) &= \pi_{AList}(\vartheta'(SR_1)) - \pi_{AList}(SR_1) && \text{def. of } \vartheta' \\
&= \pi_{AList}(SR_1^O - SR_1) - \pi_{AList}(SR_1) && \text{induction hyp.} \\
&= \pi_{AList}(SR^O) - \pi_{AList}(SR) && \pi_{AL}(A-B) - \pi_{AL}(B) = \pi_{AL}(A) - \pi_{AL}(B) \\
&= SR^O - SR && \text{def. of SR}
\end{aligned}
$$

5. $SR = (SR_1)$. Holds trivially.

□

Corollary 1: If SR is a simple relational expression, then $SR^O \subseteq SR \cup \vartheta'(SR)$.

Proof:

$$
\begin{aligned}
SR \cup \vartheta'(SR) &= SR \cup (SR^O - SR) && \text{Lemma 2 (def. of } \vartheta'(SR)) \\
&= SR \cup SR^O && A \cup (B - A) = A \cup B \\
&\supseteq SR^O
\end{aligned}
$$

□

Lemma 3: If Re is a relational expression, then $\vartheta(Re) \subseteq Re$.

Proof: The proof proceeds by induction on the structure of Re.

Base case. $Re = R$. $\vartheta(Re) = \Delta^+(R) \subseteq R$ by the definition of $\Delta^+(R)$.

Induction.

1. $Re = Re_1 \cup Re_2$. By the induction hypothesis, $\vartheta(Re_1) \subseteq Re_1$ and $\vartheta(Re_2) \subseteq Re_2$. Then $Re = Re_1 \cup Re_2 \supseteq \vartheta(Re_1) \cup \vartheta(Re_2) = \vartheta(Re_1 \cup Re_2) = \vartheta(Re)$.

2. $Re = Re_1 \times Re_2$. By the induction hypothesis, $\vartheta(Re_1) \subseteq Re_1$ and $\vartheta(Re_2) \subseteq Re_2$. Then $Re_1 \times Re_2 \supseteq \vartheta(Re_1) \times Re_2$ and $Re_1 \times Re_2 \supseteq Re_1 \times \vartheta(Re_2)$. Then $Re = Re_1 \times Re_2 \supseteq (\vartheta(Re_1) \times Re_2) \cup (Re_1 \times \vartheta(Re_2)) = \vartheta(Re_1 \times Re_2) = \vartheta(Re)$.

3. $Re = Re_1 - SRe_2$.

$$\begin{array}{lll}
 & \vartheta(Re) = & \\
= & (\vartheta(Re_1) - SRe_2) \cup (Re_1 \cap \vartheta'(SRe_2)) & \text{def. of } \vartheta \\
= & (\vartheta(Re_1) - SRe_2) \cup (Re_1 \cap (SRe_2^O - SRe_2)) & \text{Lemma 2} \\
= & (\vartheta(Re_1) - SRe_2) \cup ((Re_1 \cap SRe_2^O) - SRe_2) & A \cap (B - C) = (A \cap B) - C \\
= & (\vartheta(Re_1) \cup (Re_1 \cap SRe_2^O)) - SRe_2 & (A - B) \cup (C - B) = (A \cup C) - B \\
\subseteq & (\vartheta(Re_1) \cup Re_1) - SRe_2 & Re_1 \cap SRe_2^O \subseteq Re_1 \\
\subseteq & Re_1 - SRe_2 & \text{induction hyp. on } Re_1 \\
= & Re & \text{def. of } Re
\end{array}$$

4. $Re = \sigma_{Compare}(Re_1)$. By the induction hypothesis, $Re_1 \supseteq \vartheta(Re_1)$. $\sigma_{Compare}(Re_1) \supseteq \sigma_{Compare}(\vartheta(Re_1))$. Then $\vartheta(\sigma_{Compare}(Re_1)) = \sigma_{Compare}(\vartheta(Re_1)) \subseteq \sigma_{Compare}(Re_1) = Re$.

5. $Re = \pi_{AList}(Re_1)$. By the induction hypothesis, $Re_1 \supseteq \vartheta(Re_1)$. $\pi_{AList}(Re_1) \supseteq \pi_{AList}(\vartheta(Re_1))$. Then $\vartheta(\pi_{AList}(Re_1)) = \pi_{AList}(\vartheta(Re_1)) \subseteq \pi_{AList}(Re_1) = Re$.

6. $Re = Aggr(Attr_1, Attr_2)(Re_1)$. $\vartheta(Re) = Re$ so it holds trivially.

7. $Re = (Re_1)$. Trivial case.

$\square$

Lemma 4: If Re is a relational expression, then $Re \subseteq Re^O \cup \vartheta(Re)$. (Equivalently, if Re is a relational expression, then $Re - \vartheta(Re) \subseteq Re^O - \vartheta(Re)$.)

Proof: The proof proceeds by induction on the structure of Re.

Base case. $Re = R$. By the definition of delta relations, $R = (R^O \cup \Delta^+(R)) - \Delta^-(R)$. Therefore $R \subseteq R^O \cup \Delta^+(R)$.

Induction.

1. $Re = Re_1 \cup Re_2$. By the induction hypothesis, $Re_1 \subseteq \vartheta(Re_1) \cup Re_1^O$ and $Re_2 \subseteq \vartheta(Re_2) \cup Re_2^O$. Then $Re = Re_1 \cup Re_2 \subseteq (Re_1^O \cup \vartheta(Re_1)) \cup (Re_2^O \cup \vartheta(Re_2)) = (Re_1^O \cup Re_2^O) \cup (\vartheta(Re_1) \cup \vartheta(Re_2)) = Re^O \cup \vartheta(Re)$.

2. $Re = Re_1 \times Re_2$.

$$\begin{array}{lll}
 & Re^O \cup \vartheta(Re) = & \\
= & (Re_1^O \times Re_2^O) \cup (\vartheta(Re_1) \times Re_2) \cup (Re_1 \times \vartheta(Re_2)) & \text{def. of } Re \text{ and } \vartheta(Re) \\
\supseteq & ((Re_1^O - \vartheta(Re_1)) \times (Re_2^O - \vartheta(Re_2))) \cup & Re^O \supseteq Re^O - \vartheta(Re) \\
 & (\vartheta(Re_1) \times Re_2) \cup (Re_1 \times \vartheta(Re_2)) & \\
\supseteq & ((Re_1 - \vartheta(Re_1)) \times (Re_2 - \vartheta(Re_2))) \cup & \text{induction hyp.: } Re - \vartheta(Re) \subseteq Re^O - \vartheta(Re) \\
 & (\vartheta(Re_1) \times Re_2) \cup (Re_1 \times \vartheta(Re_2)) & \\
= & (Re_1 \times Re_2 - Re_1 \times \vartheta(Re_2) - \vartheta(Re_1) \times & \text{“distribution” of } \times \\
 & (Re_2 - \vartheta(Re_2))) \cup (\vartheta(Re_1) \times Re_2) \cup (Re_1 \times \vartheta(Re_2)) & \\
= & (Re_1 \times Re_2) \cup (\vartheta(Re_1) \times Re_2) \cup (Re_1 \times \vartheta(Re_2)) & \text{if } D \supseteq C \text{ (Lemma 3) then} \\
 & & (A - B - C) \cup D \cup B = A \cup D \cup B \\
= & Re_1 \times Re_2 & \text{Lemma 3} \\
= & Re & \text{def. of } Re
\end{array}$$

3. $Re = Re_1 - SRe_2$.

$$
\begin{array}{lll}
 & Re^O \cup \vartheta(Re) = & \\
= & (Re_1^O - SRe_2^O) \cup (\vartheta(Re_1) - SRe_2) \cup & \text{def. of } Re \text{ and } \vartheta(Re) \\
 & (Re_1 \cap \vartheta'(SRe_2)) & \\
\supseteq & ((Re_1^O - \vartheta(Re_1)) - SRe_2^O) \cup & Re^O \supseteq Re^O - \vartheta(Re) \\
 & (\vartheta(Re_1) - SRe_2) \cup (Re_1 \cap \vartheta'(SRe_2)) & \\
\supseteq & ((Re_1 - \vartheta(Re_1)) - SRe_2^O) \cup & \text{induction hyp.: } Re_1 - \vartheta(Re_1) \subseteq Re_1^O - \vartheta(Re_1) \\
 & (\vartheta(Re_1) - SRe_2) \cup (Re_1 \cap \vartheta'(SRe_2)) & \\
\supseteq & ((Re_1 - \vartheta(Re_1)) - (SRe_2 \cup \vartheta'(SRe_2))) \cup & \text{Corollary 1: } SRe_2 \cup \vartheta'(SRe_2) \supseteq SRe_2^O \\
 & (\vartheta(Re_1) - SRe_2) \cup (Re_1 \cap \vartheta'(SRe_2)) & \\
= & ((Re_1 - \vartheta(Re_1) - SRe_2) - \vartheta'(SRe_2)) \cup & A - (B \cup C) = (A - B) - C \\
 & (\vartheta(Re_1) - SRe_2) \cup (Re_1 \cap \vartheta'(SRe_2)) & \\
= & (((Re_1 - \vartheta'(SRe_2) - \vartheta(Re_1)) \cup & (A - C) \cup (B - C) = (A \cup B) - C \\
 & \vartheta(Re_1)) - SRe_2) \cup (Re_1 \cap \vartheta'(SRe_2)) & \\
= & (((Re_1 - \vartheta'(SRe_2)) \cup \vartheta(Re_1)) - SRe_2) \cup & (A - B) \cup B = A \cup B \\
 & (Re_1 \cap \vartheta'(SRe_2)) & \\
= & ((Re_1 - \vartheta'(SRe_2)) \cup \vartheta(Re_1) \cup & \text{Lemma 2: } SRe_2 \cap \vartheta'(SRe_2) = \emptyset \\
 & (Re_1 \cap \vartheta'(SRe_2))) - SRe_2 & \\
= & (Re_1 \cup \vartheta(Re_1)) - SRe_2 & (A - B) \cup (A \cap B) = A \\
= & Re_1 - SRe_2 & \text{Lemma 3} \\
= & Re & \text{def. of } Re
\end{array}
$$

4. $Re = \sigma_{Compare}(Re_1)$. By the induction hypothesis, $\sigma_{Compare}(Re_1) \subseteq \sigma_{Compare}(Re_1^O \cup \vartheta(Re_1)) = \sigma_{Compare}(Re_1^O) \cup \sigma_{Compare}(\vartheta(Re_1)) = Re^O \cup \vartheta(Re)$.

5. $Re = \pi_{AList}(Re_1)$. By the induction hypothesis, $\pi_{AList}(Re_1) \subseteq \pi_{AList}(Re_1^O \cup \vartheta(Re_1)) = \pi_{AList}(Re_1^O) \cup \pi_{AList}(\vartheta(Re_1)) = Re^O \cup \vartheta(Re)$.

6. $Re = Aggr(Attr_1, Attr_2)(Re_1)$. The subset relation $Aggr(Attr_1, Attr_2)(Re_1) \subseteq Aggr(Attr_1, Attr_2)(Re_1^O) \cup Aggr(Attr_1, Attr_2)(Re_1)$ holds trivially.

7. $Re = (Re_1)$. Trivial case.

$\square$

Acknowledgements

We thank Bobbie Cochrane for useful suggestions and Stefano Ceri for enabling the collaboration.

References

[ASU86] A.V. Aho, R. Sethi, and J.D. Ullman. *Compilers: principles, techniques, and tools.* Addison-Wesley, Reading, Massachusetts, 1986.

[CG85] S. Ceri and G. Gottlob. Translating SQL into relational algebra: Optimization, semantics, and equivalence of SQL queries. *IEEE Transactions on Software Engineering*, 11(4):324–345, April 1985.

[CM93] S. Ceri and R. Manthey. First specification of Chimera. Technical Report IDEA.DD.2P.004.02, Politecnico di Milano, Milan, Italy, May 1993.

[CW90] S. Ceri and J. Widom. Deriving production rules for constraint maintenance. In *Proceedings of the Sixteenth International Conference on Very Large Data Bases*, pages 566–577, Brisbane, Australia, August 1990.

[FRS93] F. Fabret, M. Regnier, and E. Simon. An adaptive algorithm for incremental evaluation of production rules in databases. In *Proceedings of the Nineteenth International Conference on Very Large Data Bases*, pages 455–466, Dublin, Ireland, August 1993.

[GHJ+93] S. Ghandeharizadeh, R. Hull, D. Jacobs, et al. On implementing a language for specifying active database execution models. In *Proceedings of the Nineteenth International Conference on Very Large Data Bases*, Dublin, Ireland, August 1993.

[Han92] E.N. Hanson. Rule condition testing and action execution in Ariel. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 49–58, San Diego, California, June 1992.

[HW93] E.N. Hanson and J. Widom. An overview of production rules in database systems. *The Knowledge Engineering Review*, 8(2):121–143, June 1993.

[Joh75] S.C. Johnson. YACC — yet another compiler compiler. Technical Report CSTR 32, Bell Laboratories, Murray Hill, New Jersey, 1975.

[Les75] M.E. Lesk. LEX — a lexical analyzer generator. Technical Report CSTR 39, Bell Laboratories, Murray Hill, New Jersey, 1975.

[QW91] X. Qian and G. Wiederhold. Incremental recomputation of active relational expressions. *IEEE Transactions on Knowledge and Data Engineering*, 3(3):337–341, September 1991.

[RCBB89] A. Rosenthal, S. Chakravarthy, B. Blaustein, and J. Blakeley. Situation monitoring for active databases. In *Proceedings of the Fifteenth International Conference on Very Large Data Bases*, pages 455–464, Amsterdam, The Netherlands, August 1989.

[SJGP90] M. Stonebraker, A. Jhingran, J. Goh, and S. Potamianos. On rules, procedures, caching and views in data base systems. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 281–290, Atlantic City, New Jersey, May 1990.

[SKdM92] E. Simon, J. Kiernan, and C. de Maindreville. Implementing high level active rules on top of a relational DBMS. In *Proceedings of the Eighteenth International Conference on Very Large Data Bases*, pages 315–326, Vancouver, British Columbia, August 1992.

[WCL91] J. Widom, R.J. Cochrane, and B.G. Lindsay. Implementing set-oriented production rules as an extension to Starburst. In *Proceedings of the Seventeenth International Conference on Very Large Data Bases*, pages 275–285, Barcelona, Spain, September 1991.

[WF90] J. Widom and S.J. Finkelstein. Set-oriented production rules in relational database systems. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 259–270, Atlantic City, New Jersey, May 1990.

[WH92] Y.-W. Wang and E.N. Hanson. A performance comparison of the Rete and TREAT algorithms for testing database rule conditions. In *Proceedings of the Eighth International Conference on Data Engineering*, Phoenix, Arizona, February 1992.